Psychophysical Dissociation of Phase and Position Cues for Visual Motion Perception

Alexandra Hibble

A thesis submitted for the degree of

Doctor of Philosophy



St John's College
Department of Experimental Psychology
University of Oxford
Trinity 2022

Abstract

Visual motion perception is essential for viewing and interacting with our environment. The motion of an object in the visual field can be described by the changes in the position of a feature over time ('position cues' in the spatial domain), or by changes in the phase of the spatial frequency components in the Fourier domain ('phase cues'). This thesis investigates the relative extent to which people rely on these cues to infer movement in the visual scene.

I showed that apparent motion discrimination with phase cues and position cues can be psychophysically dissociated by temporal interval (Chapter 2), and spatial frequency (Chapter 3), and that displacement thresholds for position cues increase more rapidly with eccentricity than do phase cues (Chapter 3). When phase and position cues are conflicted, I demonstrated that these spatial frequency dissociations can predict reversals in reported motion direction, both with changing spatial frequency, and between foveal and peripheral viewing (Chapters 4 and 5). The reliance on position cues to infer motion correlates with the ability to discriminate shapes, such that when shape discrimination is impaired motion discrimination is reliant on phase cues, in both normal observers and in the scotoma of V1 lesioned patients (Chapter 6). This shows that the visual system treats phase and position cues as separate sources of information which could be analysed separately.

These psychophysical dissociations may reflect the properties of different detectors in the visual field. Correlations of models of cell density distributions with the displacement thresholds across eccentricity suggest that the pα and pβ retinal ganglion cells form the physiological substrates for the input to these phase and position systems respectively (Chapter 7).

This thesis presents a psychophysical dissociation between phase and position cues, and suggests that these cues are initially processed by separate systems, with distinct physiological substrates.

Acknowledgements

I must thank my supervisor Paul Azzopardi, for your patience, and willingness to exchange late night emails – hopefully for the last time! Thanks also to my secondary supervisor Hannah Smithson, for such a warm welcome to the Oxford Perception Lab, and to the OPL members past and present, especially Sarah Prosciutto for all the late nights. I am also grateful to my examiners, Professor Timothy Ledgeway and Professor Brian Rogers, for such an enjoyable viva.

To all of my friends at St. John's, past and present, this thesis would not be where it is without your love and support. In particular to Anna and Natalia, for being the three musketeers throughout this final stretch; to the lovely ladies of Juxon street: Rowena, and my Jars behind Bars, Laura and Catriona; to Ollie, for being fantastisch through it all; and to my home friends, Rhianna and Soph: a huge thank you and I love you all.

Most of all, to my parents. Thank you for all these years of unwavering support. I love you forever, and I hope you can emotionally recover from this.

This work was supported by the generous funding from the St. John's Lamb and Flag scholarship, the Clarendon Fund, and the Department of Experimental Psychology.

Table of Contents

Abstract 2

Acknowledgements 3

Table of Contents 4

Chapter 1: Introduction 6

1.1 Cues for motion perception 6
1.2 Are both phase and position cues used to discriminate motion? 10
1.3 Are phase and position cues processed by different systems? 15
1.4 Models for the analysis of phase and position cues 24
1.5 The organisation of the visual system for motion perception 37
1.6 Conclusion 47
1.7 Thesis Overview 47

Chapter 2: Spatial and temporal frequency dissociations of phase and position cues 53

2.1 Introduction 53
2.2 Temporal Integration Dissociation of Phase and Position Systems 54
2.3 Spatial Frequency Dissociation of Phase and Position Systems 71
2.4 Motion Energy Analysis 97
2.5 Conclusion 113
2.6 Appendix 115

Chapter 3: Psychophysical dissociation of phase and position cues 130

3.1 Introduction 130
3.2 Method 136
3.3 Results 139
3.4 Discussion 146
3.5 Appendix 152

Chapter 4: Perceived motion in Gabor based conflicting cue stimuli 156

4.1 Introduction 156
4.2 Opposing motion Gabors presented across the visual field 158
4.3 Differentiating between the Spatial Scaling and Feature Blur hypotheses 169
4.4 Discussion 187

Chapter 5: Perceived motion in reverse phi stimuli 191

5.1 Introduction 191
5.2 Reverse phi grating at fovea 193
5.3 Reverse phi square stimulus across eccentricity 200
5.4 Discussion 213

Chapter 6: Residual motion abilities following lesions of striate cortex 216

6.1 Introduction 216
6.2 Method 221
6.3 Experimental Methods and Results 224
6.4 Discussion 243

Chapter 7: Correlations between psychophysics and physiology 251

7.1 Introduction 251
7.2 Method 257
7.3 Results 268
7.4 Discussion 271

Chapter 8: Discussion 277

8.1 Overview of results 278
8.2 Main contributions 288
8.3 Future Directions 292
8.4 Conclusion 298

Reference List 300

Chapter 1: Introduction

The ability to perceive motion is a fundamental property of the human visual system enabling observers to perceive, navigate and interact with the environment. A Newtonian definition of motion is the change in position of a feature with time, which can also be expressed in the frequency domain, as a change in phase of the Fourier components of the image. For the majority of moving objects, describing the change in spatial position of the image features, and the change in phase of the Fourier components, are mathematically equivalent descriptions. However, laboratory stimuli can be constructed where these two descriptions of motion predict different directions of motion, or stimuli where only one of these descriptions corresponds to a coherent motion direction.

Psychophysical experiments have shown that human participants can detect and discriminate motion when the change in spatial position of features do not give a coherent cue for direction, but when the changes in the phase of the Fourier components predict coherent motion (Van Santen & Sperling, 1984). Participants are also able to discriminate motion in stimuli that are carefully constructed to contain no changes in phase, where motion is produced solely through changes in feature position (Lu & Sperling, 1995). This suggests that humans are able to discriminate motion using both descriptions of motion: changes in phase and changes in the spatial position of features. This thesis investigates whether these phase and position cues are processed by separate systems using psychophysical dissociations, and whether these imply that the different descriptions of motion are processed by different mechanisms in the visual system.

In this chapter I will lay out evidence to suggest that humans are able to discriminate motion using phase and position cues independently, and discuss the evidence as to whether these cues are processed by two separate systems. I will introduce some of the models proposed as mechanisms to extract phase and position cues, and give an overview of the physiology of the visual motion system in the human brain.

1.1 Cues for motion perception

A Newtonian description of motion refers to a change in the position of a feature with time. However, motion can also be described by the changes in phase of the spatial frequency components in the Fourier domain. Both the changes in position and the changes in phase are sources of information, or cues, that could be used by the visual system to infer the motion of an object in the visual scene. These two descriptions of motion are referred to throughout this thesis as ‘position’ and ‘phase’ cues, respectively.

1.1.1 Position cues

The change in the position of a feature with time is described in this thesis as the 'position cue'. A visual feature is a local image characteristic that encodes relevant information for identifying objects and surfaces within the visual input, where features are 'distinctive [and] relevant' (Weinmann, 2013). Information is not distributed equally across a visual scene: the change in luminance at two neighbouring locations that denotes an edge is more informative than identical luminance at two neighbouring locations, so a feature represents a spatially restricted area of useful information. This means that features are inextricably linked to their position in the visual input, hence the terminology 'position' cue.

Features in an image could be constructed from a variety of image properties, such as changes in luminance, contrast and texture or colour. In most visual stimuli, there are multiple image characteristics that could be used to define the 'features' in the image. For example, in a sine wave the features are recognisable as the bars of the grating, yet the position of these bars could be defined through the zero-crossings, where zero-crossings refer to the location where the function (here the sine wave) crosses zero, the locations of minimum or maximum luminance, and the extended regions of lower and higher luminance. Research in the field of computer vision has also operationalised features as basic visual elements, with Marr's primal sketch produced from zero-crossings for an image filtered at several spatial scales (Marr, 1980), and more recent work in the field specifying edges, corners and regions (or 'blob') as the 'visual features' (Li *et al.*, 2015).

More complex constructions of image features were demonstrated by Lu & Sperling who produced gratings where features are produced from luminance and contrast modulations, but also depth, motion-motion and flicker modulations (Lu & Sperling, 1995). This family of gratings demonstrates that there is not a single image property that universally denotes a 'feature': in the depth modulated gratings there was no correlation in the luminance of the pixels between frames, only as a space-varying correlation between the right and left eye images, and in the motion modulated gratings the 'bars' of the grating were produced by dots moving in opposing directions, rather than local changes in luminance or contrast.

1.1.2 Phase cues

Any visual image can be described as a set of Fourier components: a set of sine waves with specified spatial frequencies, amplitudes, orientations, and phases. An alternative description of motion in the visual input is by the change in phase of its Fourier components.

1.1.3 Congruent motion cues

Typically changes in phase (phase cues) are also explicit in the spatial domain, producing features that change position, and that changes in feature position (position cues) can typically be described

by changes in phase: these are two different ways to describe the same information. This is exemplified in a drifting sine wave grating, which contains a single spatial frequency, and can be described in three dimensions: the x and y extent in space, and in time (t) (see Figure 1.1). The change in position of the features in the grating is visible in the displacement of the regions of higher and lower luminance (Figure 1.1, top right), and could be described using various features within the grating. The changing phase of the sine wave grating could be described by the oriented slope of luminance when plotted in x-t coordinates (Figure 1.1, bottom left), or made explicit when two-dimensional Fourier analysis is applied to the x-t representation of the moving grating. The resultant amplitude spectrum would consist of a pair of pulses describing the spatiotemporal frequency of the grating (as Fourier transforms are mirror symmetric, this represents the single frequency present in the stimulus), and the quadrants of the amplitude spectrum that these pulses are displayed in would describe the motion direction of the grating (see Figure 1.1, bottom right).

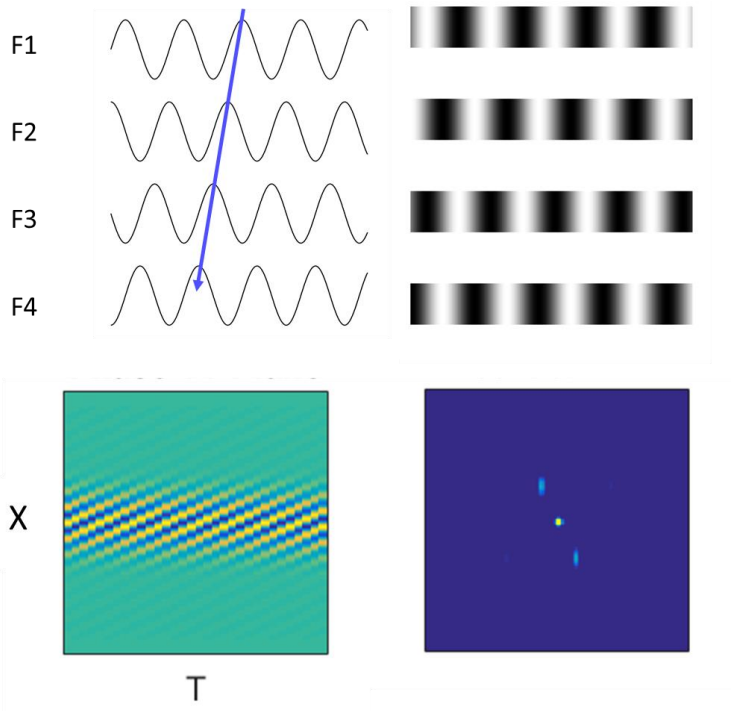


Figure 1.1. Schematic describing a sine wave grating displaced leftwards by 90° on each frame, its XT projection and centred Fast Fourier Transform (FFT). Top Row: the cross-section of 4 frames of a sine wave grating displaced by 90° on each frame (left), and the 2D appearance of the grating to an observer, in X-Y coordinates. Bottom Left: the X-T transformation of 32 frames of the sine wave grating, where the diagonal stripes depict the change in position of the peaks and troughs of the grating over successive frames of the stimulus. Bottom right: the resultant amplitude spectrum of the Fourier transform of this X-T input, where the pale blue peaks represent the frequency of the sine wave. The distance of the peaks from the origin represents the spatial frequency of the input, with increasing distance representing higher spatial frequencies, the brightness of the peaks the amplitude (power), and the angle of the signals from this origin refer to the orientation of the frequencies in the XT input. This is a centred FFT, where the quadrants have been rearranged to display the lowest frequencies at the centre of the image, and progressively higher SFs towards the outside. This rearrangement means that leftward motion produces a centred FFT with frequencies on the rightward diagonal and vice versa. Motion energy is calculated by comparing the amplitude of the frequencies in each quadrant in diagonal pairs (ME = (Q1+Q3) – (Q2 + Q4)): for the leftward displaced sine wave the positive and negative peaks representing the frequency of the sine wave exist in quadrants 2 and 4, producing an amplitude spectrum with a negative motion energy, consistent with rightward motion (see Chapter 2.4 for a more thorough explanation).

1.1.4 Conflicting motion cues

For the sine wave grating displayed in Figure 1.1 the displacement of features and the phase of the Fourier components both signal motion in the same direction. However, this is not the case for all stimuli, such as in the case of the Missing Fundamental Grating. The Missing Fundamental Grating (MFG: Adelson, 1982; Adelson & Bergen, 1985) is produced by subtracting the fundamental frequency from a square wave of matched frequency and amplitude. Square waves are formed from the sum of the fundamental frequency and all of its odd harmonics, with the harmonics having n times the frequency and 1/n the amplitude of the fundamental frequency. Subtracting the fundamental frequency from a square wave grating produces a grating that retains the periodicity of the fundamental frequency, albeit with a scalloped shape profile, but the lowest harmonic with the highest amplitude is now the third harmonic (See Figure 1.2, top row). A forward displacement of the MFG by ¼ of the wavelength of the fundamental frequency, equivalent to a 90^{o} phase shift of the fundamental frequency, displaces the features forward. However, as the third harmonic has three times the frequency of the fundamental frequency, this produces a phase shift of ¾ in the third harmonic (otherwise expressed as 270^{o}), which is equivalent to a backwards 90^{o} phase displacement. The other odd harmonics in the grating are similarly aliased in forwards (if in the form 4n+1) or backwards (if in the form 4n-1) directions. The summed amplitude of the 4n-1 harmonics is greater than the summed amplitude of the 4n+1 harmonics, therefore the motion energy of the grating is in the opposite direction to the displacement of the features in the grating. This can be seen in the Fourier spectrum in Figure 1.2, whereby the diagonally oriented pulses describing the motion energy of the 90^{o} displaced MFG are in the opposite direction to the displacement of the features in the grating, with the two brightest diagonal pulses describing the motion energy of the 3rd harmonic. As Fourier components are infinite in extent, this 270^{o} phase shift is aliased as a 90^{o} phase shift. This mathematical property - the third harmonic displaced by 90^{o} of the fundamental frequency resulting in an 270^{o} phase shift equivalent to a backwards 90^{o} displacement – means that displacing the MFG by 90^{o} of the fundamental frequency produces a grating where the features (position cue) are displaced forwards, and the grating phase (phase cue) signals motion in the reverse direction. Thus the MFG is an example of a stimulus where the phase and position cues signal motion in opposing directions.

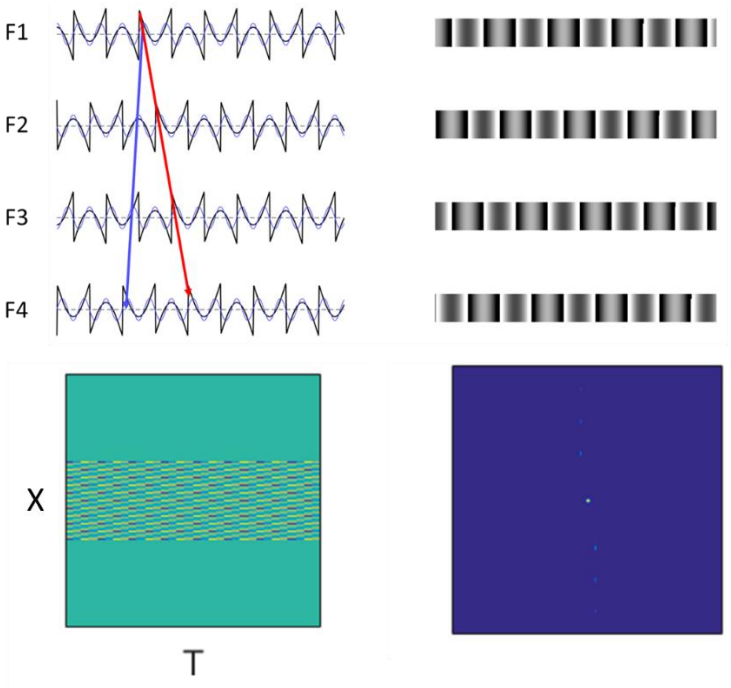


Figure 1.2. Schematic displaying a Missing Fundamental Grating (MFG), displaced by 90° on each frame (top row), and its appearance in XT coordinates and the 2D Fourier spectrum of the XT input (bottom row). Top Left: the cross-section of 4 frames of the MFG displaced by 90° on each frame with blue arrows indicating the direction of the aliased third harmonic, and the red arrows indicating the direction of displacement of the features. Top Right: the 2D appearance of the grating to an observer, in X-Y coordinates. Bottom Left: the X-T transformation of 32 frames of the MFG, and bottom right: the centred FFT, the resultant amplitude spectrum of the Fourier transform of this X-T input. The centred FFT displays the zero frequency component of the stationary Hanning window (the yellow centre peak), and positive and negative peaks along the right diagonal. Each peak refers to the odd harmonics in the MFG, with the f3 harmonic closest to the centre and with the highest amplitude of the harmonics, and each further peak along the diagonals corresponding to the higher odd harmonics, with lower amplitudes. This centred FFT has been cropped, with only the amplitudes of the f3, f5 and f7 harmonics visible in this display. Note that due to the rearrangement of the quadrants in the FFT, motion is represented along the opposite diagonal: the presence of peaks along the rightward diagonal in the centred FFT demonstrates that the changing phase of the Fourier components in the stimulus are in the leftward direction. Therefore whilst the MFG is displaced to the right, the Fourier components of the MFG (determined by the changing phase of the odd harmonics) signal motion in both leftwards (3rd,7th,11th etc. harmonics) and rightwards (5th, 9th, 13th etc.) harmonics, with the lowest frequency third harmonic dominating reported motion direction.

1.2 Are both phase and position cues used to discriminate motion?

To determine whether participants are able to independently discriminate motion using phase or position cues, stimuli can be constructed that minimise, remove, or render cues incoherent, producing stimuli that contain useful directional information in solely the phase or position cues. These stimuli can be described as ‘pure’ stimuli. The ability to discriminate motion using these stimuli would provide evidence that human participants are sensitive to both phase and position cues.

1.2.1 Pure position stimuli

One method to produce stimuli which contain position but not phase cues, is to attempt to remove phase cues entirely. Illusory figures can be produced by aligning pacmen, circles with missing right-

angled segments, to form an illusory square figure (Kanisza, 1955, Figure 1.3, left). Whilst the square figure appears to be clearly separated from the background, and may appear to be a different luminance or at a different depth to its surroundings (Watanabe & Oyama, 1988), there are no physical contours separating the figure from the background. Kanizsa figures can be used to construct motion stimuli by forming rows of the aligned pacmen, and having them rotate so as to move the illusory figure along the row of aligned pacmen (Figure 1.3, right). If the pacmen which aren't involved in producing the square also rotate at random in different directions on each frame, this will prevent any coherent local motion signals, and match the noise from the pacmen rotating to produce the aligned illusory contour (Seifert & Hock, 2014). This produces an illusory square displaced horizontally, but without any coherent phase cues. However, one complication in the use of the displaced Kanizsa figures is that whilst by definition there are no coherent phase cues for discriminating motion direction available in the visual input, the visual system responds to these illusory contours in a similar way to real contours. This can be shown psychophysically, whereby participants exhibit similar changes in threshold facilitation and elevation along the contour (Dresp & Bonnet, 1991), and adaptation can occur between real and illusory contours, and vice versa (tilt aftereffect; Berkley *et al.*, 1994). The treatment of illusory contours as real contours defined by luminance could be caused by lateral inhibition, possibly even at the level of the retinal ganglion cells (Coren, 1991). The percept of illusory contours is abolished when the stimuli are presented chromatically at iso-luminance (Cavanagh, 1987), strongly suggesting that the although the illusory contour by definition is iso-luminant (since there is no contour present), apparent luminance differences are being introduced using lateral inhibition generated from the aligning elements. If so, it's possible that systems involved in processing phase cues may respond to these illusory contours as luminance defined contours, and that this will introduce a phase cue into the visual input, rendering these stimuli inappropriate as 'pure' position cue stimuli.

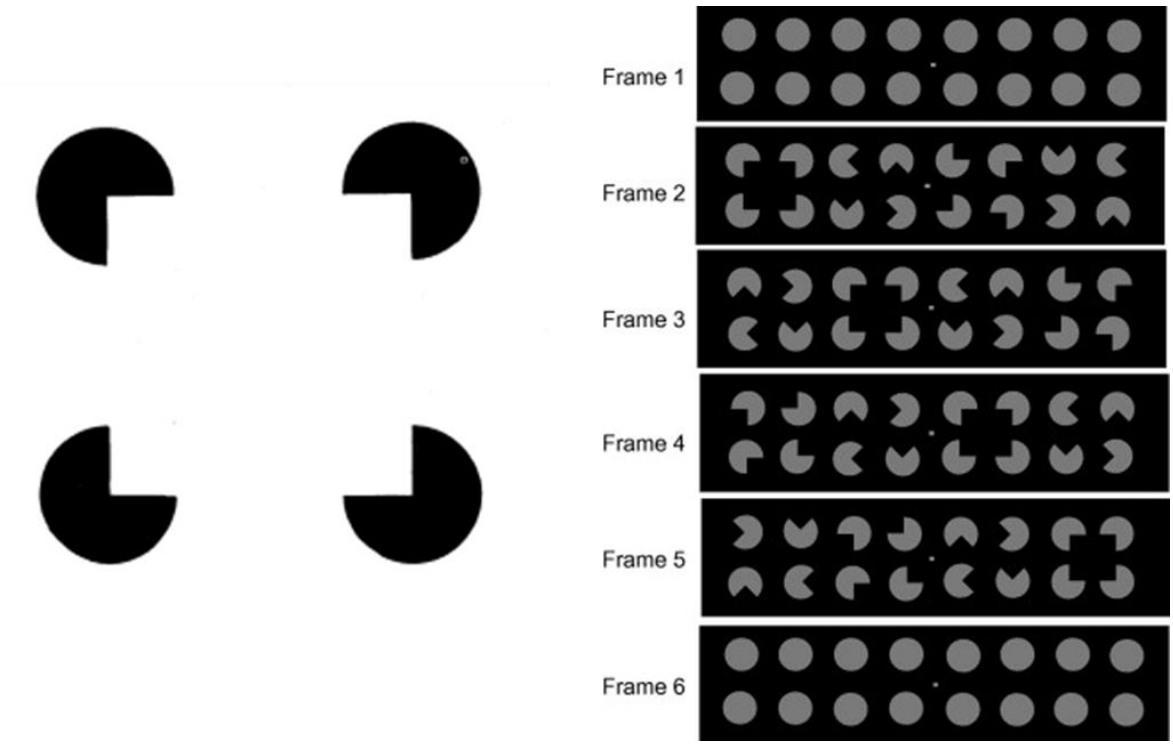


Figure 1.3. Left: From Watanabe & Oyama (1988), Kanizsa square, with a central illusory square that appears brighter than the background and outlined by illusory contours. Right: From Seifert & Hock (2014), displaying the six frames of a stimulus that uses illusory figures, to produce motion without coherent phase cues. 'Pacmen' shapes with missing corners are aligned to produce a strong percept of a figure where the edges of the figure are illusory and created without luminance borders. On each frame the pacmen rotate, producing an illusory square that appears to be displaced horizontally. As the borders of the square are illusory, the phase cues in the stimulus are confined to the rotation of the pacmen, which does not contain directional information.

A different, but more prevalent approach, has been to construct motion stimuli where the visual characteristics defining the features change on each frame. The change in visual information available on each frame can eliminate the correspondence in the phases of the Fourier components of the stimuli, or render them incoherent and not a viable cue for motion discrimination.

Baro & Levinson interleaved pairs of counterphase sine wave gratings and counterphase Missing Fundamental Gratings, both constructed using the same fundamental frequency (see Figure 1.4, Baro & Levinson, 1988). The sine wave contains only the fundamental frequency, whereas the Missing Fundamental Grating contains features with the same periodicity as the sine wave, due to the sum of the odd harmonics of the fundamental frequency, without containing the fundamental frequency. This means that there are no Fourier components in common between the two gratings, and so no possible phase cues. On each frame the grating, either a sine wave or Missing Fundamental, was displaced by a distance equivalent to ¼ of the fundamental frequency. Motion direction can be discriminated in the direction of the feature displacement, despite the different grating profiles on each frame producing substantially different appearances of the features – a scalloped bar on the MFG frames compared to the smooth increase and decrease in the sine wave

grating. This could be achieved by using the displacement of the zero-crossings as the 'feature', which share a direction and are displaced in a consistent 90^{o} between frames although with a different gradient, or could be produced by some mapping between the bars of the MFG and sine waves by extracting a general concept of feature and location for the 'bars' on each frame, and tracking the displacement of the 'feature' between frames without depending on the low-level image characteristics that produce the feature. The ability of participants to discriminate motion direction in this stimulus demonstrates sensitivity to position cues in the absence of phase cues.

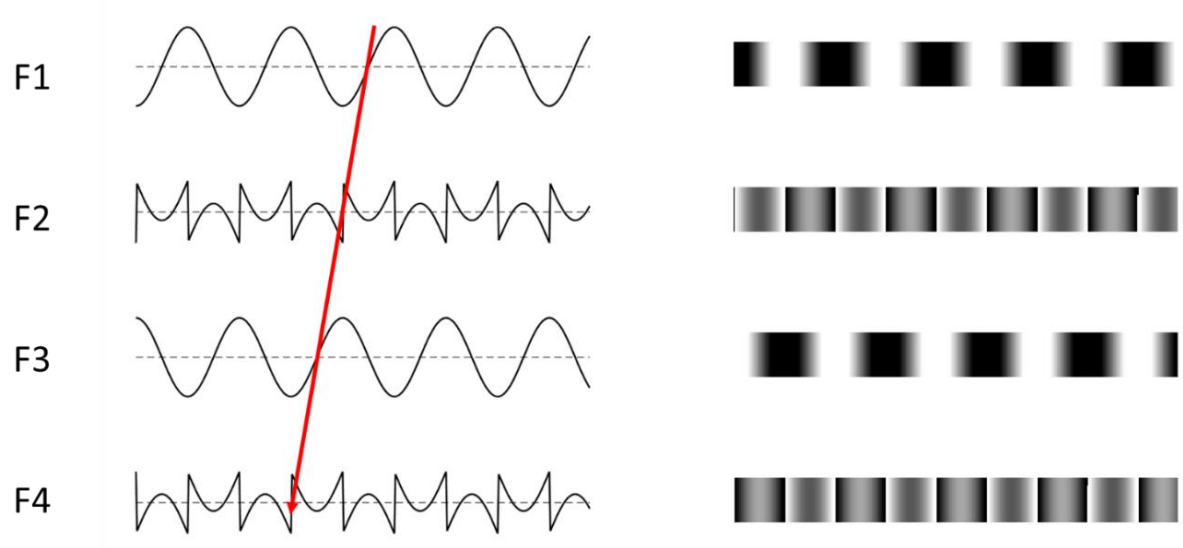


Figure 1.4. Figure displaying 4 frames of the MF1F, designed by Baro & Levinson (1988). The left hand panel shows the cross-section of each frame, with the red arrow demonstrating the coherent displacement of the zero-crossings between each frame (a possible feature used to achieve motion discrimination), and the right hand panel displaying how the stimulus appears to the participant on each frame.

A more complex version of this idea – interleaving gratings of different identities – was presented in Lu & Sperling's demonstration of 'Attention generated motion perception', (Lu & Sperling, 1995). Participants were shown a drifting grating, which on alternative frames was a depth modulated or texture modulated grating (Figure 1.5). The motion direction of the stimulus is ambiguous, in that there is no explicit match between textural elements and depth, and perceived motion direction is dependent on which image characteristics are matched between each frame e.g. coarse texture regions matched with protruding regions, rather than receding regions. The successful motion discrimination performed by participants demonstrates an ability to track feature displacement despite a lack of corresponding image characteristics, and suggests that the visual system is extracting feature location and feature salience.

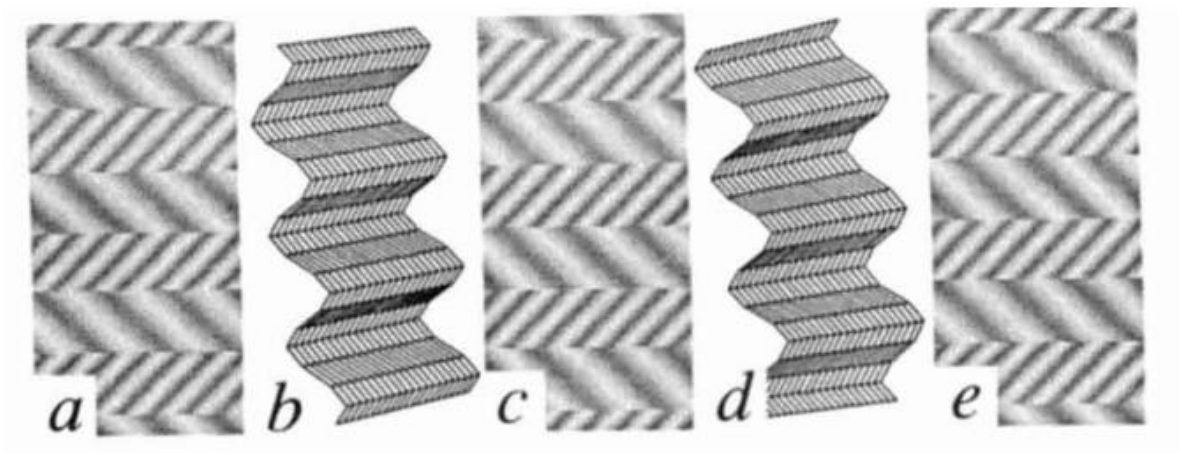


Figure 1.5. From Lu & Sperling (1995b): Schematic of each of the first five frames (a to e) of an ambiguous grating stimulus, produced from interleaving frames of a sinusoidal grating modulated in depth, or in the textural elements. When instructed in an explicit mapping between frames, participants were able to consistently discriminate the direction of motion, demonstrating an ability for top-down attention to produce a salience map from which motion discrimination can be achieved: motion from the displacement of salient features, without reliance on corresponding image characteristics between frames.

1.2.2 Pure phase stimuli

Methods employed to produce pure phase stimuli focus on rendering the position cues incoherent, or by finding parameters under which the feature or position in congruent stimuli cannot be detected. Examples of these stimuli, and evidence that human participants can discriminate motion using phase cues, can be seen in the 'real' world, and not just in psychophysics laboratories: the direction and velocity of motion in the murmuration of starlings, wind across a field of long grass, or the flow of water along a river, can be easily perceived without the change of position of features with time. These examples can be considered as lacking coherent motion between features, because the features are spatially unresolvable, or as the motion of individual elements has a high lack of correspondence between features with time, rendering the tracking of features an unreliable cue to motion direction.

Stimuli containing phase cues but incoherent position cues can be demonstrated in psychophysics laboratories, such as the 'pedestal' gratings, introduced by Van Santen & Sperling (Van Santen & Sperling, 1984). These are formed from the sum of a stationary low contrast sine wave and a matched frequency drifting sine wave with half the amplitude of the stationary sine wave (Figure 1.6). The phase cue is produced by the change in the phase of the drifting sine wave, whilst the position cue in the grating is determined by the luminance profile of the sum of the gratings. As the stationary sine wave has twice the amplitude of the drifting sine wave, the peaks and troughs of the grating will appear to oscillate back and forth between frames around the location of the features of the stationary sine wave. Therefore, this change in location of the features of the image between frames, whether carried by the maxima and minima of the luminance or by the zero crossing of the grating, does not provide a coherent cue for motion perception. After training to 'attend to the

wind' participants were able to discriminate the motion direction of the drifting component, and to ignore the incoherent position cue. The pedestal gratings are an example of 'objectless' motion perception, where the motion of the phase cue can be discriminated, without being attached to a displaced feature.

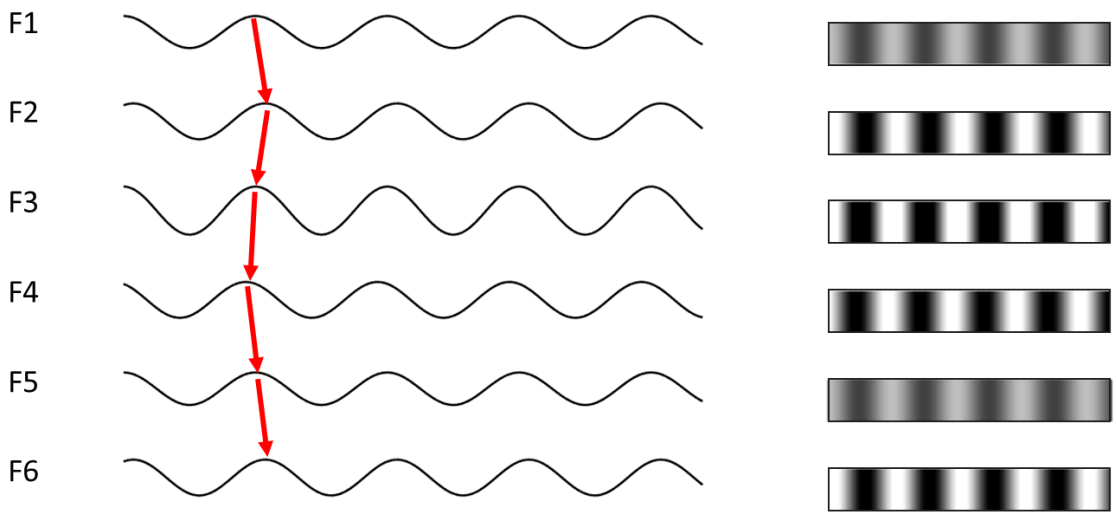


Figure 1.6. Figure displaying 6 frames of the pedestal grating (Van Santen & Sperling, 1984), formed from a drifting sine wave grating summed with a stationary sine wave grating, with the same spatial frequency and twice the amplitude. This can be represented as sin(x+t) + 2 sin(x), where x is the spatial frequency of the grating and t is the displacement in time. Left: the cross-sections of the pedestal grating, and right, as it appears to the observer. The red arrows on the cross-section trace the location in space of one of the grating peaks, demonstrating that there is no consistent spatial displacement of the peaks, despite the coherent phase cue across frames.

1.3 Are phase and position cues processed by different systems?

Differences in the sensitivity to phase and position cues across spatial and temporal frequency ranges, contrast etc. can be used to investigate whether these cues are processed by different systems. Being able to describe motion in terms of its phase or position cues, and show that participants can discriminate motion direction selectively using either cue is not sufficient evidence to conclude that there are two separate systems involved in processing motion: it could be that there is a single system able to extract motion using transformations of the visual input that render the distinction between cues meaningless. Equally, differences in reported motion could be the result of the image characteristics defining the stimuli. More compelling evidence to suggest that there are different systems responding independently to phase and position cues would be produced from comparing psychophysical dissociations between stimuli with similar image characteristics. The pattern of psychophysical dissociations could be used to produce hypotheses about the operations of the putative underlying mechanisms, and suggest potential physiological substrates.

Carefully constructed stimuli can be used to investigate psychophysical dissociations between phase and position cues. There are three approaches that could be taken to explore psychophysical dissociations between phase and position cues. Firstly, the parameter spaces for pure phase and position cues could be explored to see whether there are systematic differences. However, previously published constructions of pure phase and position cues vary substantially in their spatial and temporal structures, so it is difficult to find stimuli sufficiently similar that differences in participants' responses can be attributed to the phase or position system, rather than the construction of the stimuli. Secondly, the parameter space could be explored using stimuli where the visual cues are available to both phase and position systems and find a parameter range where one cue cannot be detected, such as phi motion. This can additionally be used to distinguish between the description of motion in the visual input, and the parameter space in which the detectors for that system can utilise these motion cues. Finally, conflicting cue stimuli could be used, where the reported motion direction tells us which cue has the greater relative strength at that presentation.

1.3.1 Congruent phase and position cues

The majority of visual stimuli depicting motion, as with real-world moving objects, contain congruent phase and position cues, whereby the changes in phase of the Fourier components are also explicit in the spatial domain, and the changes in feature position can also be described by changes in phase of the Fourier components. However, just because the information in the visual input can be described by both phase and position cues, does not mean that these cues are equally detectable: which cues are utilised by the motion systems is a function of the cue and the properties of the detectors.

If there are regions of parameter space where it can be demonstrated that one cue is unavailable, but that motion discrimination remains high, this would be evidence to suggest that the two cues are processed by different systems. This strategy could potentially be accomplished looking for parameters where motion discrimination is occurring, but the position cue is potentially unavailable. This could be achieved in stimuli where the change in position of the features is spatially unresolvable, but motion discrimination is occurring.

This can be seen in demonstrations of apparent motion, where the percept of motion can be produced between two spatially separated locations, by varying temporal onset and offset of stimuli at these locations. Foster *et al.* (1989) constructed an apparent motion stimulus using two separated pixels, finding the minimum pixel separation that would reliably elicit a motion percept at 5 peripheral viewing locations between 5 and 25 degrees eccentricity. Foster *et al.* (1989) measured the minimum pixel separation that participants could reliably report spatial separation between the two locations. At all eccentricities the thresholds for motion perception (phase cue) were

consistently lower (smaller separation required) than the thresholds for judging spatial separation. This demonstrates that motion perception can occur even when the spatial separation between two locations cannot be distinguished. Furthermore, depending on the spatial separation and the temporal interval, participants report seeing either alternating onset at the different locations, one figure moving between the locations (beta motion), or the sensation of motion from a shadowy occluding figure, moving between the locations with continuous onset (phi motion), (Exner, 1875; Wertheimer, 1912; Steinman *et al.*, 2000). Phi motion can be considered objectless motion, as the sensation of motion is not tied to an object. The sensation of 'objectless motion' without a visible feature changing location at high temporal frequencies suggests that at high temporal frequencies only the phase cue is available to motion detectors. This implies that the upper temporal frequency limit for motion discrimination for position cues is lower than that for phase cues. Apparent motion stimuli therefore can be used to define regions of parameter space where the spatial and temporal parameters can be constructed so that motion discrimination is occurring, but the position is spatially unresolvable and features are below the detection threshold. A similar effect can be seen in low contrast drifting sine waves, where motion discrimination occurs despite features of the grating being below detection threshold (Kelly, 1979).

However, this approach relies on the assumption that the mechanism for extracting motion from position cues in the image is reliant on explicitly detecting features and their positions. Some proposed models for the extraction of position cues do not rely on the explicit detection of features (e.g. Filter-rectify-filter models, discussed in Chapter 1.4.2). These experiments argue that the discrimination threshold for phase cues for these stimuli, at these parameters, is lower than the detection threshold for spatial separation and feature identity: these factors may or may not constrain models extracting the position cue. This reveals a critical flaw in this experimental strategy whereby in congruent motion stimuli it is difficult to say whether motion discrimination is occurring due to phase or position cues, and our interpretation of the data relies on underlying assumptions about the mechanisms used to extract motion information from the visual scene.

1.3.3 Conflicting phase and position cues

Visual stimuli can contain conflicting phase and position cues, where the two cues signal motion in different directions. Classic examples of conflicting cue stimuli are the Missing Fundamental Grating (Adelson & Bergen, 1985, Figure 1.2), and reverse phi stimuli (Anstis & Rogers, 1975, Chubb & Sperling, 1989, Figure 1.7). Participants are able to detect and discriminate motion across these different stimuli, reporting motion in either the phase or position cue direction depending on the parameters of the stimuli (Reverse Phi: Anstis, 1970; Chubb & Sperling, 1989; Gorea, 1995; Edwards & Nishida, 2004; MFG: Georgeson & Shackleton, 1989; Georgeson & Harris, 1990). This

demonstrates that participants are able to process both phase and position cues, but also supports a psychophysical dissociation between two systems processing phase and position cues separately.

One stimulus class often used to demonstrate conflict between the phase and position systems, and to investigate different ranges of spatiotemporal sensitivity of these systems, are 'reverse phi' stimuli. Reverse phi stimuli are constructed from the displacement of a feature with a simultaneous reversal in contrast (Anstis, 1970), such as a bar which is displaced horizontally across an image, reversing contrast on each displacement (Chubb & Sperling, 1989, Figure 1.7). This can be described in terms of the change of orientation in the zero-crossings of the image, or in terms of phase and position cues. This reversal of stimulus contrast reverses the orientation of the luminance in the space-time plot, producing a stimulus where the change in phase of the Fourier components are in the opposite direction to the direction of the displacement of the image features.

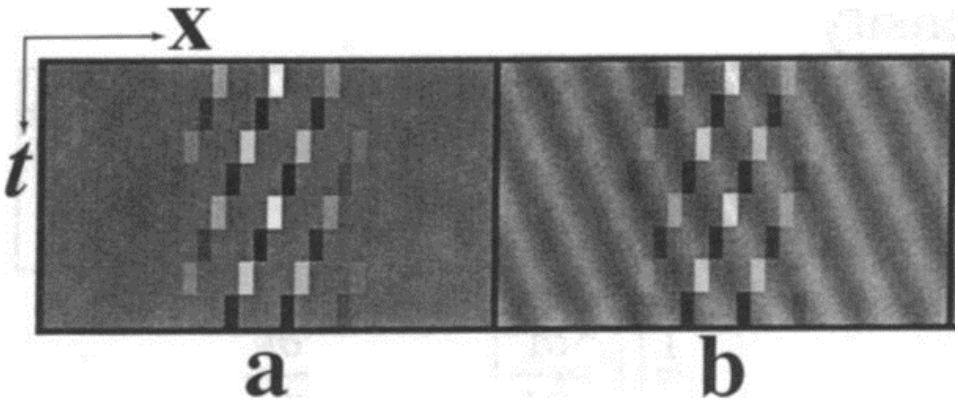


Figure 1.7. Figure from Chubb & Sperling (1989). Left: Box a displays the 8 frames of the contrast reversing grating, where each row in the image corresponds to a frame displayed in the experiment, producing an x-t schematic of the experiment. Right: In box b Chubb & Sperling overlay a sinusoid grating which indicates the predicted direction of the phase cue in the image. The phase cue indicates motion to the right, whereas the reversing contrast bars, and therefore the position cue, are displaced to the left.

The effect of stimulus displacement on determining reported motion direction in reverse phi stimuli was demonstrated by Anstis and Rogers (1975), where discs presented with less than 10 minutes of arc displacement were reported in the phase cue direction, and above 10 minutes of arc, in the position cue direction. Chubb & Sperling (1989) also investigated the effect of displacement size on reported cue direction, using viewing distances of 2 and 8 metres to change the spatial frequency on the retina of their reverse phi grating. When the grating was displayed at a spatial frequency of 4.17 cycles per degree motion was reported in the position cue direction, whereas at a spatial frequency of 16.67 cycles per degree motion was reported in the phase cue direction. Also using a reverse phi grating, Gorea (1995) showed an interaction between the spatial and temporal frequency and reported motion direction, where motion was reported in the position cue direction for stimuli that were presented at less than 3.75Hz and 1 cycle per degree, but consistently reported in the phase

cue direction at greater temporal and spatial frequencies. Edwards and Nishida (2004) found a similar effect of temporal frequency, varying the velocity of a reverse phi random dot display. At low speeds (6 degrees per second) motion was reported in the position cue direction, and at high speeds (27 degrees per second), in the phase cue direction. Finally, Wehrhahn (2006) investigated reported motion direction in an apparent motion stimulus consisting of a pair of opposite contrast bars, presented at varying spatial and temporal offsets. These are reverse phi stimuli which don't overlap in spatial extent, unlike the others discussed here, which could suggest that the reverse phi phenomenon – reporting motion in the phase cue direction - is driven by falling within the appropriate spatio-temporal receptive fields rather than purely in spatial coordinates. Below an offset of 6 minutes of arc stimuli were reported in the phase cue direction, and at separations of 12 minutes of arc and above, motion was reported in the position cue direction. Across these various constructions of reverse phi stimuli, motion is reported in the phase cue direction at high spatial and temporal frequencies, and in the position cue direction at low spatial and low temporal frequencies. This difference in spatiotemporal sensitivity for the phase and position cues could indicate that the two cue types are processed by two separate systems.

In experiments using reverse phi stimuli of different constructions motion is reported in the position cue direction at low temporal frequencies and in the phase cue direction at high temporal frequencies (Edwards and Nishida, 2004, Gorea,1994, Wehrhahn, 2006). A similar finding has been reported in plaids. When the sum of two sinusoid gratings oriented at ±84° from vertical is displaced by $^{3}/_{8}$ of its wavelength, the luminance (phase cue) and contrast envelope (position cue) are displaced in opposing directions. When presented at temporal frequencies below 3Hz participants reported motion in the position cue direction, but above 3Hz in the phase cue direction (Derrington & Ukkonen, 1999). This is consistent with the rapid decline in sensitivity at high temporal frequencies for gratings defined by position cues, but not for phase cues (Lu & Sperling, 1995), and the presence of objectless motion (phi) in apparent motion stimuli presented at high temporal frequencies. These findings suggest that access to the position cue declines at high temporal frequencies, suggesting that systems processing phase and position cues have detectors operating across different temporal frequency ranges, and that discriminating between phase and position systems may be easier when using stimuli presented at higher temporal frequencies.

Research into spatial frequency using reverse phi gratings often confounds the spatial frequency of the features, and the magnitude of the displacement, as only one of these dimensions is explored. Smith & Ledgeway (2001) investigated whether there is a bias towards phase and position cues at different spatial frequencies, and separately manipulated both the spatial frequencies contained in the stimuli and the magnitude of displacement. They constructed a set of RDKs with varying element

sizes, and applied no filter, or high pass filters at 1, 2 or 4 cycles per degree (Figure 1.8). The impact of these filters was to change the lowest spatial frequency present in the stimuli, without changing the position of the features – in the same way that the MFG retains the periodicity of its lowest spatial frequency when this is removed. This can be seen in Figure 1.8, where the outline (zero-crossing) of the large elements is consistent between the filtered and unfiltered conditions, despite the loss of low spatial frequencies. Both standard (same polarity), and reverse phi (reversed polarity) conditions were constructed, displacement between frames was varied between 0 and 120arcmin, and participants were asked to report the motion direction of the RDK.

If the position cue is used to discriminate motion direction in these stimuli then motion should be reported in the direction of feature displacement, regardless of the polarity of the frames, up to a limiting displacement that is too large for the features to be matched between. If the phase cue is being used then reported motion direction should reverse with reversals of contrast, and performance will be determined by the spatial frequency components, and not the element sizes, in the image. Reported motion direction should be determined by the lowest spatial frequencies present in the image, and should exhibit aliasing with the magnitude of displacement. In the high pass filtered stimuli there should therefore be no difference in reported motion between stimuli with different element sizes, as the lowest spatial frequency in the image is determined by the high pass filtering.

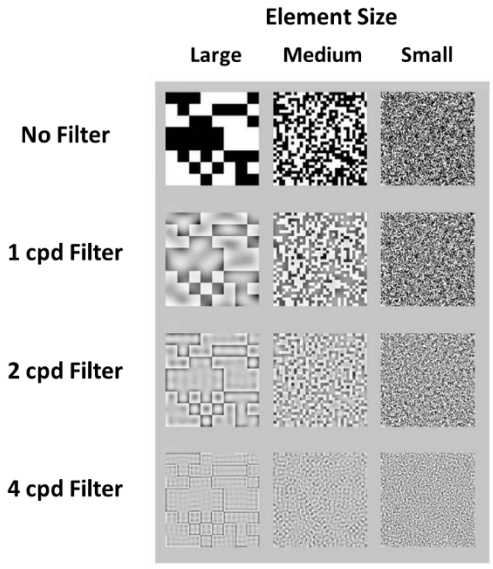


Figure 1.8. Figure from Smith & Ledgeway (2001) displaying example stimuli from their experiment: RDKs with different sized elements, and with a high pass filter applied to some stimuli to remove spatial frequencies below 1, 2 or 4 cycles per degree. These frames were displaced, by a magnitude between 0 and 120 arcmin, to produce an apparent motion stimulus. Smith & Ledgeway (2001) also included a reverse phi condition, where the second frame reversed in polarity, simultaneously with the displacement.

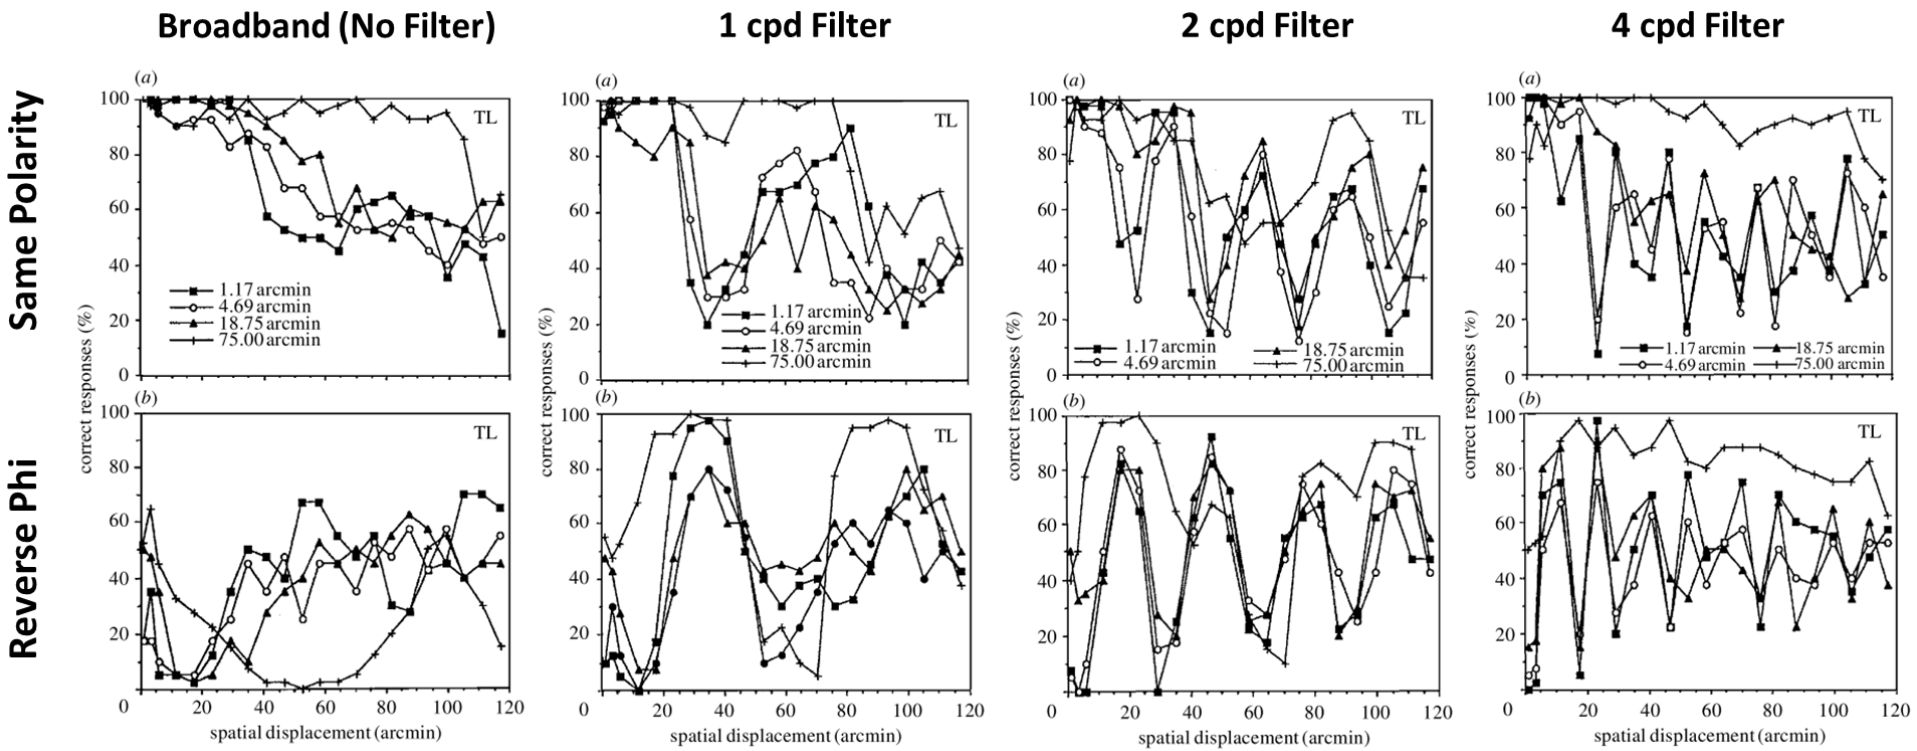


Figure 1.9. From Smith & Ledgeway (2001), displaying the results of participant TL, for each of the filtered conditions from broadband (far left) to high-pass filtered (4 cpd, far right column). The standard condition (same polarity) is represented on the top row, and the figures on the bottom row come from the reverse phi condition. The four element sizes are represented on each plot, with a different symbol for each (legend on figure). Displacement magnitude between frames 1 and 2 of the stimuli is plotted on the x axis, from 0 to 120 arcmin, and motion discrimination is plotted on the y axis, where above 50% motion is reported in the displacement direction, and below 50% is reported in the direction opposite to the displacement direction. Across the majority of stimuli, motion is reported in the phase cue direction, seen by the reversed reported motion direction in the reverse phi condition (bottom row), and by the change in spatial aliasing with displacement as the lowest spatial frequency in the image is increased by high pass spatial filtering. However, for stimuli filtered at 4 cpd, the largest element size (75 arcmin) is reported in the position cue direction, as demonstrated by the lack of reversal of reported motion in the reverse phi condition, and the lack of spatial aliasing with different displacement magnitudes. Motion is reported in the displacement direction, suggesting that it is produced from a position cue that is dependent on the features in the image – the large element size – rather than the spatial frequency components of the image.

Reported motion in the broadband stimulus was predicted by the phase cue, as demonstrated by the reversal of reported motion direction in the reverse phi condition (Figure 1.9). Further evidence that reported motion direction was reliant on the phase cue comes from the spatial aliasing visible in 1, 2 and 4 cpd high pass filtered stimuli, whereas the lowest spatial frequency in the image increases, from 1 to 4 cpd, the displacement needed to produce aliasing in the reversed direction decreases, producing more cycles of reversals within the same range of magnitudes of displacement (Across the columns, Figure 1.9). The cyclical relationship produced by aliasing is further evidence that the phase cue is determining reported motion direction, as the reported direction for the standard and reverse phi stimuli were in counterphase, demonstrating that the reversal of contrast was acting as an 180^{o} phase displacement of the stimulus (Smith & Ledgeway, 2001).

However, for the RDKs high pass filtered at 4 cpd this cyclical relationship was not seen for the stimulus with the largest element size (75 arcmin), and there was no reversal of reported direction in the reverse phi condition (Figure 1.9, rightmost column). These two factors are consistent with motion in this stimulus being reported using the position cue, suggesting that despite the spatial filtering, the position system is still active in this spatial frequency range provided the stimulus element size is sufficiently large. The lowest spatial frequencies in this stimulus are the same as in the stimuli of the smaller element sizes. However, the spatial structure – size of the elements – has not been removed by the spatial frequency filtering. This suggests that the position system is using the spatial structure – the features of the image – to determine perceived motion direction.

A final psychophysical dissociation between phase and position cues comes from comparing discrimination and thresholds for stimuli presented monocularly and interocularly. Using monocular and interocular displays, Lu & Sperling (1995a) demonstrated that the system processing motion from phase cues was a predominantly monocular processes. Motion discrimination for luminance and contrast modulated gratings can be achieved in interocular displays, where alternate frames of a grating displaced by 90^{o} are presented to each eye (Shadlen & Carney, 1986). This 90^{o} displacement means that the left and right eyes are receiving gratings in counterphase, and to discriminate the direction of motion, must integrate between the left and right eye images. Adding a stationary pedestal to these gratings does not interfere with the phase cue, as it produces no directional Fourier components as it has no phase changes, but will disrupt the position cues, as the addition of the pedestal renders the displacement of the features across frames incoherent (Van Santen & Sperling, 1984). When pedestaled luminance and contrast modulated gratings are displayed interocularly, participants can no longer discriminate motion direction (Lu & Sperling, 1995), but motion discrimination is possible when the gratings are presented monocularly. This demonstrates that the non-pedestalled motion discrimination was carried by position cues, and not phase cues.

For motion-modulated gratings that contain no coherent phase cues, where motion discrimination is reliant on the position cue, there is no substantial threshold change between non-pedestalled monocular and interocular viewing, compared to the 10-fold increase for luminance-modulated and 3-fold increase for contrast modulated gratings (Lu & Sperling, 1995). This suggests that the mechanism processing position cues is indifferent to whether stimuli are presented monocularly or interocularly, whilst phase cues cannot be processed interocularly. This has two key implications: firstly, it suggests that phase and position cues are processed by two separate mechanisms. Secondly, this may give some insights into the potential physiological substrates for phase and position systems, where the site of binocular combination is thought to be V1 (DeAngelis *et al.*, 1991; Ohzawa *et al.*, 1990). Looking for substrates with and without binocular combination may be an avenue to explore in searching for physiological substrates for the phase and position systems.

Across these experiments using conflicting cue stimuli there is a clear bias where motion is reported in the phase cue direction at high spatial and temporal frequencies, and in the position cue direction at low spatial and low temporal frequencies – consistent with the findings in congruent stimuli. The reported direction of motion in these conflicting cue stimuli can therefore demonstrate which cue is dominating the motion percept and therefore has a greater relative strength. This method could be used to map the spatial and temporal parameters across which the sensitivity to phase and position cues differs.

The detectability of a motion cue is the function of the strength of the cue and the sensitivity of the detector. Conflicting cue stimuli contain different phase and position cues e.g. In the MFG the position cue is defined by the changing position of the features in the grating, whilst the phase cue is defined by the aliased odd harmonics. To investigate whether phase and position cues are processed by separate mechanisms, the confound of image characteristics must be controlled for. This could be in stimulus design – producing phase and position cues with more similar image characteristics – or by comparing across sufficient stimulus classes to determine that psychophysical dissociations are the result of underlying separate mechanisms extracting phase and position cues, rather than the image characteristics of individual stimuli.

1.3.4 Conclusion

Psychophysical dissociations between phase and position cues are essential in constructing models of motion perception, both in forming hypotheses about underlying mechanisms, and to investigate the possible physiological substrates. Experiments using both congruent and conflicting cues demonstrate consistent psychophysical dissociations. These could reflect the behaviour of two separate systems processing phase and position cues, with detectors operating across different temporal frequency ranges, with sensitivity to position cues declining at relatively lower temporal

frequencies than phase cues. Research into the relative sensitivity of the systems processing phase and position cues across spatial frequency suggests that whilst there may be more substantial overlap (Smith & Ledgeway, 2001), the phase system may be sensitive to higher spatial frequencies than the position system, as motion can be discriminated from phase cues when spatial separation is unresolvable and hence position cues are unavailable (Foster *et al.*, 1989). Both these differences in spatial and temporal sensitivity have implications for the arrangements of the initial detectors or filters in models produced to analyse motion.

1.4 Models for the analysis of phase and position cues

Producing plausible models for analysing phase and position cues provides potential descriptions of the mechanisms needed to take place in the visual system, where the differences between models can be used to discuss whether a single mechanism is plausible. The complexity of the transformations required can also be used to discuss possible physiological substrates.

1.4.1 Models for extracting motion from phase cues

Phase cues in a stimulus can be described in spatiotemporal parameters, where the velocity of stimuli can be described in intensity (I) across x, y and time coordinates: I(x,y,t), where the steeper the gradient in t, the higher the velocity of the stimulus. Similarly, the Fourier description of a moving stimulus is also oriented, where the changing phase of a stimulus moving at a constant velocity is described as motion energy in frequency space, that is an oriented plane through the origin (see Figures 1.1 and 1.2). Using these representations of motion, motion models have been proposed that extract and interpret the orientation of stimuli in space-time or in spatio-temporal frequency space.

A popular class of models for the analysis of phase cues by the visual system suggests that the orientated slope of luminance in space-time could be extracted by a motion detector formed from a pair of detectors that are oriented in this space-time dimension (Reichardt, 1957; Adelson & Bergen, 1985; Van Santen & Sperling, 1985). The Reichardt detector model (Figure 1.9, left) receives a point light input, a single luminance value, from two spatially separate locations. These are passed through a temporal filter which applies a linear, time-invariant filter with attenuation and phase shift. This signal is then multiplied with the original and opposite signal (e.g. left hand original with transformed right hand signal), which introduces non-linearity into the model, and sent through a time averaging unit. The outputs of the left and right sides of the model are summed, and the sign of the output determines calculated motion direction. This model was extended in the 'Elaborated Reichardt detector' model (Figure 1.10, right), where Van Santen & Sperling (1985) propose that the model input is now a pair of spatial filters arranged in quadrature, rather than a single point-light input. This is more physiologically plausible given the spatial extent of receptive fields in the visual

system, and logistically plausible, given the number of receptive fields that would be needed to tile the retina if the input was from a single photoreceptor. The use of a spatial filters, arranged in quadrature pairs, means that the Elaborated Reichardt model extracts motion energy from the spatio-temporal domain.

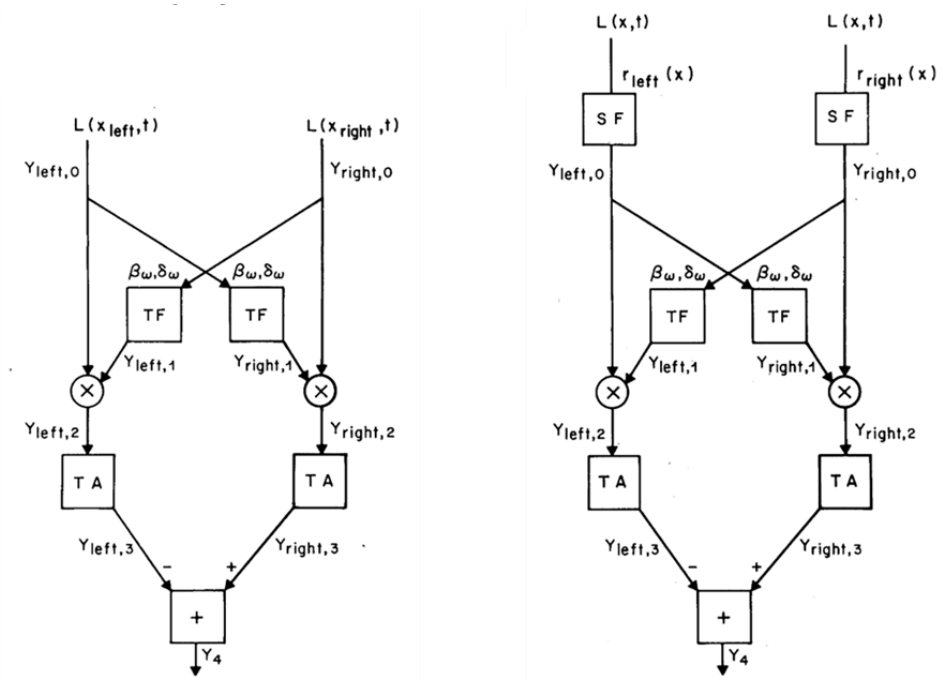


Figure 1.10 Figures from Van Santen & Sperling (1984). Left: Reichardt detector model, where TF refers to the temporal filter, the 'X' to the unit that multiplies the original and transformed inputs, and TA to the time averaging unit. The outputs of the left and right sides of the model are summed (+), and the sign of the output determines calculated motion direction. Right: Elaborated Reichardt detector model, where Van Santen & Sperling propose that the model input is now the stimulus passed through a spatial filter, rather than a single point-light input.

The Elaborated Reichardt detector model is affected by the aperture problem: motion can only be extracted perpendicular to the local contour, and for these models, the aperture problem also limits the ability of these models to extract velocity . This means that these first stage motion detectors will only produce a signed motion signal with 180^{o} specificity. This necessitates a second stage, where multiple first-stage directionally selective units are combined to produce an accurate direction estimate, and with overlapping receptive fields to extract velocity from the spatiotemporal luminance input (Heeger, 1987). Therefore to produce accurate velocity estimates, the model requires units with a range of spatio-temporal sensitivities. This could be achieved by having various ranges of the separation in space-time of the spatial filters in the Elaborated Reichardt detector model. These potentially correspond with the width of centre-surround receptive fields in the early stages of the visual system. Overlapping receptive fields are mandated by this, but reduce the ability of a system to precisely locate the position of a feature: this may be a potential physiological difference that may underpin the initial stages of models processing phase and position cues, reflected in models of motion cue analysis. These motion energy models would predict the motion

response to phase cues without coherent features (pedestal gratings; Van Santen & Sperling, 1984), and the reversal in reported motion direction in reverse phi stimuli (Anstis, 1970).

An alternative class of models for extracting motion produced from changes in phase are 'gradient models', which use partial derivatives of the oriented luminance spatio-temporal plane to directly estimate velocity. This eliminates the need for a second stage comparison of the outputs of first-stage detectors, but multiple derivatives have to be computed on the same input. This is needed to cope with the brightness constraint, and guard against the aperture problem. The brightness constraint describes the assumption that the intensity of points within the visual input are conserved during motion introduces unknowns into the motion constraint equation (Johnston *et al.*, 1992), and can introduce zero value denominators into a partial derivative. When denominators tend towards zero, this implies an infinite (impossible) velocity. It is unlikely that a zero denominator would be present in both a first and second stage partial derivative (Poston & Stewart, 1978), therefore reducing the chance that the model will produce impossible velocity estimates. The extended model of Johnston *et al.* (1992) is potentially biologically plausible, where the convolution of the visual input with separable space-time filter kernels could produce spatial and temporal partial derivatives. Johnston *et al.* (1992) model this zero-order spatio-temporal filter with a Gaussian profile in the space domain and a Lorentzian profile in the temporal domain, similar to descriptions of receptive fields in V1.

Johnston *et al.* (1992) examined whether these gradient models could account for a range of motion phenomena. Testing their model across various visual motion stimuli, they calculated partial derivatives for the intensity in the space and time dimensions, and used these to produce constraint models which are solved by the combination of multiple local units. They extended the model further, constraining it to minimise the change in illumination relative to the local mean. The model could accurately predict the direction of drifting gratings, and accounted for the effects of spatial frequency on perceived speed (Smith & Edgar, 1990). The extended Johnston *et al.* (1992) model contains a rectification within the image processing, unlike the motion energy models based on the Reichardt detector. Interestingly, when the model was presented with the reverse phi grating of Chubb & Sperling (1989) at different scales, the model reported motion in the reversed direction at the far viewing (high spatial frequencies) and in the veridical, forward direction at the near viewing condition (lower spatial frequencies), consistent with the human psychophysical data. This is in direct opposition to Chubb & Sperling's interpretation of the data, where they argue that the psychophysical data provides evidence that humans perceive motion using both a motion energy (phase cue) and separate full-wave rectification system (Chubb & Sperling, 1989). The extended Johnston *et al.* (1992) gradient model cannot account for all motion stimuli however, where it did

not produce a coherent motion signal for the drifting temporal frequency modulation grating stimulus of Chubb & Sperling (1988), despite high accuracy in human participants. Furthermore, even with the extended rectification stage, the Johnston *et al.*, (1992) model would not produce coherent motion signals when applied to proposed pure position cue stimuli that rely on salience-driven decisions (Lu & Sperling, 1995).

Both the classes of motion detectors reliant on linear filters, exemplified by the Elaborated Reichardt detector model, and the motion detectors using gradient schemes, such as the extended model of Johnston *et al.* (1992), are arguably biologically plausible explanations for the initial discrimination of motion direction from phase cues. These both suggest that the first stage could be carried out by centre-surround opponent retinal ganglion cells, and the second stage by simple cells in V1.

Both motion energy and motion gradient models rely on low level motion detectors, formed from cells at the early stages of the visual system. Evidence for the existence of motion detectors comes from motion after effects (MAE), where following sustained viewing of a moving surface, a stationary surface will appear to drift in the opposite direction (Wohlgemuth, 1911). This suggests that a detector tuned to that direction and velocity has been habituated by the sustained viewing of the motion input. Aristotle was potentially the first to describe the MAE in 300 BC: "when persons turn away from looking at objects in motion, e.g. rivers, and especially those which flow very rapidly, they find that the visual stimulations still present themselves, for the things really at rest are then seen moving" (reported in Ross 1931), with similar effects reported by Addams (1834) after viewing the Falls of Foyers in Loch Ness. Critically, these motion after effects are not dependent on features detection within the visual array as they have been classically demonstrated with flowing bodies of water, which we can consider a feature-less stimulus, if presented at a scale when no features are detectable (cf. Kelly, 1979). Furthermore, when viewing the stationary surface following adaptation there is the strong sensation of motion in the opposite direction to the adapting surface, but any features on the stationary surface remain stationary. This suggests that the MAE is reliant on phase cues within the visual display.

Motion after effects also occur following viewing of conflicting cue stimuli, with the direction of the MAE determined by the phase cue, and not the position cue. Using two slide projectors Anstis & Rogers (1975) produced a reverse phi disc stimulus, where the positive image and slightly displaced negative image of a disc were projected separately, with a slow dissolve between one projector decreasing and the other increasing in intensity. The displacement produced the change in position between the two discs. Depending on the size of the displacement between the two discs participants reported motion either in the direction of the displacement, reporting that the disc

appeared to have 5 or 6 intermediate positions between the start and end position, or reported motion in the opposite direction to the displacement, the reverse phi effect. The use of the two slide projectors meant that Anstis & Rogers could adapt the speed of the display. They reported that when using a slow fade, lasting 30 seconds, a strong motion after effect was produced in the reverse phi direction. This suggest that MAE is indeed carried by the phase cues and not position cues in the visual array. This provides psychophysical support for the presence of motion detector units responsive to phase cues, which could form the initial stages of the Extended Reichardt and gradient models.

1.4.2 Models for extracting motion from position cues

The change in position of a feature with time could be solved either by a model that transforms the visual input to render salient features as explicit, and then performs operations similar to those discussed in the previous section, or by a model that detects features, and can correlate feature location across time.

This first class of models can be grouped as 'rectification models', where the visual input is transformed, before a further, often motion energy type analysis, is applied. Many visual stimuli contain information that is available to motion energy or gradient models without rectification, such as a drifting sine wave, where the change in luminance across the visual image is produced by the changing phase. However, more complex visual stimuli exist in which participants can discriminate motion direction clearly despite the lack of coherent luminance information, such as contrast modulated gratings with dynamic spatially broadband noise (Smith & Ledgeway, 1998), necessitating a second model for motion processing.

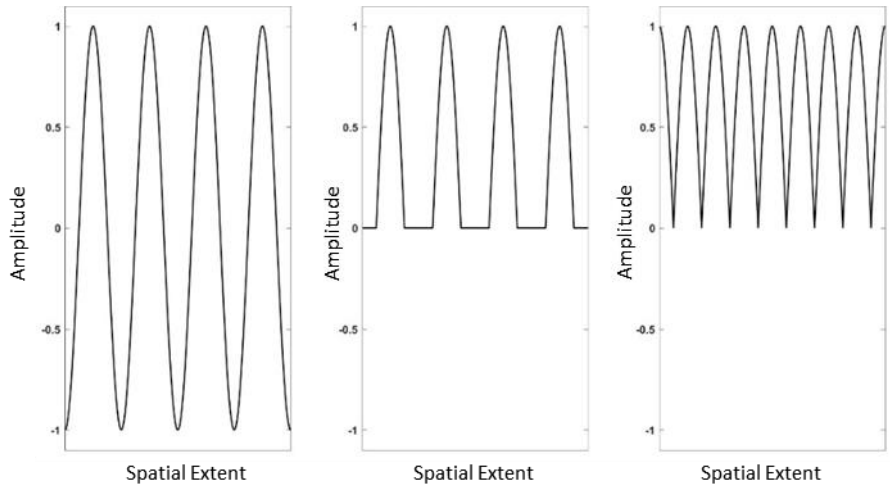


Figure 1.11. Figure plotting a full sine wave (left panel), and demonstrating two standard rectification operations: the positive half wave rectification of a sine wave (central panel), where all negative values are removed, and the full wave rectification of a sine wave (right panel), where all negative values are converted into positive values.

Different models have been proposed, often following the 'filter-rectify-filter' model. In these models the initial input is filtered such as through a bandpass spatial filter (Ledgeway & Hess, 2000), an orientation selective spatial filter (Wilson *et al.*, 1992), or a temporal filter (Chubb & Sperling, 1989), and a subsequent rectification operation occurs. The rectification operation varies in different models, but is most commonly through half or full wave rectification (see Figure 1.11, Solomon & Sperling, 1994), or by squaring the filtered input (Wilson *et al.*, 1992; Chubb & Sperling, 1991). This rectification transforms the image, introducing new spatial frequency components. Finally, the rectified output is passed through a second filter stage, which differs from the first, and could be a low-pass spatial filter (Ledgeway & Hess, 2000), or a filter stage that selects for different orientations and lower spatial frequencies (Wilson *et al.*, 1992), again introducing spatial frequencies into the model for a later stage of motion processing. Many models then use this filtered output as input to a standard motion energy analysis stage (Figure 1.12, Wilson *et al.*, 1992, Chubb & Sperling, 1991).

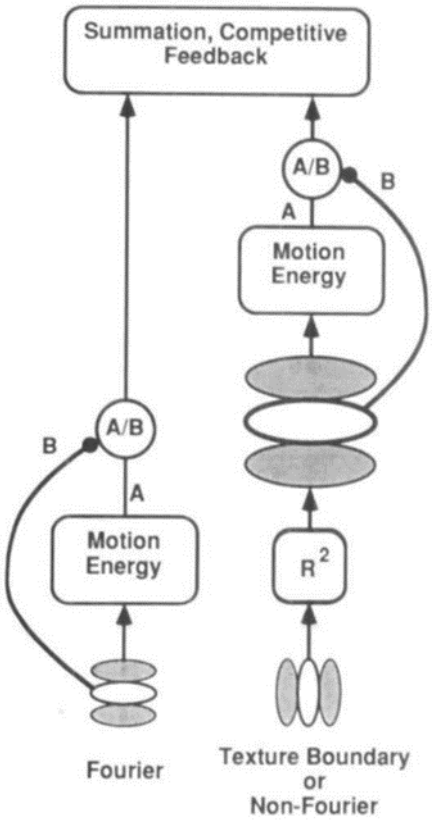


Figure 1.12. From Wilson *et al.*, (1992). Left-hand diagram: The 'Fourier' pathway is equivalent to the motion energy model of the elaborated Reichardt detector, but preceded by orientation-selective filtering. The oriented filtering stage also produces a feed- forward gain-control signal (arrow labelled B) which is applied to the output of the motion energy analysis. Right-hand diagram: The 'Non-Fourier' pathway is the pathway proposed to extract 'texture boundaries'. The model starts with the same oriented filtering, then the rectification stage where the filtered input is squared (R2 box), before being passed to a second filtering stage, with a different orientation and lower spatial frequency. This filtered output then passes through the same motion energy and gain control stages as the Fourier pathway. Wilson *et al.* (1992)'s model also proposes how the two pathways are combined: in V5/MT , where the response from both pathways are weighted by a cosine function, and compared using competitive inhibition to produce a proposed direction signal.

The operations in these rectification models are potentially physiologically plausible, where spatial filtering could be achieved through the various widths of the centre-surround receptive fields, like those of retinal ganglion cells, and oriented spatial filtering, as specified by Wilson *et al.* (1992), could be achieved through elongated neighbouring excitatory and inhibitory subregions, such as in the simple cells of V1 (Hubel & Wiesel, 1962, 1968). Likewise, some rectification operations are also physiologically plausible, where full wave rectification could occur through the sum of responses of an ON-centre and OFF-centre unit with matching spatial properties, whereas positive half wave rectification could occur through taking solely the outputs of ON-centre units, and negative half-wave rectification through the outputs of OFF-centre units.

Simulations using plaids (Wilson *et al.*, 1992), texture quilts (Chubb & Sperling, 1991) and Gabor micro-patterns (Ledgeway & Hess, 2000) have shown that these filter-rectify-filter models can accurately extract motion from contrast modulation in the image that would be 'invisible' to a motion energy detector. However, the initial filtering stage of these models will remove phase cues carried by sinusoidal variations in luminance that fall outside of the filter. These models often contain low-pass filters such as the second set of filters in the Wilson *et al.*, (1992) model. If the visual system is employing a filtering stage in the system used to process position cues, this could explain the spatial frequency bias seen in conflicting cue stimuli, with conflicting cue stimuli preferentially reported in the phase cue direction at high spatial frequencies and the position cue direction at low spatial frequencies (seen in reverse phi stimuli; Chubb & Sperling, 1989, Gorea, 1994). This effect of filtering also means that these models are not able to completely model the gamut of human motion perception for both phase and texture-based position cues. Since motion energy based models cannot extract position cues from the texture-based stimuli (e.g. texture quilts, Gabor micropatterns), this suggests that filter-rectify-filter models cannot be the sole motion processing mechanism. Wilson *et al.* (1992) suggests that the visual system contains both a standard motion energy analysis and a filter-rectify-filter analysis (see Figure 1.12).

Several pieces of evidence support the idea of a different system for extracting motion produced from position cues, using the explicit change in position of the image features. Firstly, the evidence that human participants are capable of coherently discriminating motion from complex stimuli that would be invisible to filter-rectify-filter models, such as the interleaved gratings of Lu & Sperling (1995b) where motion is determined by attributing salience between features with different identities. Secondly, the experiments of Solomon & Sperling (2004) who assessed whether position cues were reliant on full or half-wave rectification. Solomon & Sperling (2004) produced two novel forms of micro-balanced stimuli: 'half wave' stimuli where position cue information is preserved following a thresholded half wave rectification, but lost when a full wave rectification is applied, and

'full wave' stimuli, where the position cue information is lost following a half wave rectification, but preserved following a full wave rectification. Participants were able to coherently discriminate motion direction in both of these stimuli, demonstrating that whatever mechanisms are extracting position cues from the stimuli, that they are not reliant on either mechanism to the exclusion of the other. Therefore the mechanism extracting position cues may use both full and half-wave rectifications of the visual input, or a further mechanism that produces the same directional response as a full or half wave rectification, but uses a different method.

Gabor micropatterns can also be used to demonstrate that motion discrimination can be produced in the absence of visual information that could be extracted by motion energy models, or that could be made explicit by filter-rectify-filter transformations, indicating the presence of a mechanism performing explicit tracking of the change in position of features. The displacement of a Gabor stimulus could be discriminated by the change in position of the Gabor or the change in the phase of the Fourier components (a symmetrical band of spatial frequencies, centred at the frequency of the sinusoid carrier). Boulton & Baker (1993) demonstrated that when multiple Gabors were presented in a high density 'Gabor micro-pattern' motion was discriminated using the phase of the Gabor, as depending on the displacement of the Gabor, motion was reported veridically, in the direction of displacement, or aliased backwards. However, at low densities, when a single Gabor was displayed and displaced, no aliasing occurred, and motion was reported veridically at all displacements. At low densities the Gabor is a more salient stimulus (or alternatively, at high densities, the Gabor micro-pattern is vulnerable to the correspondence problem). If the change between the use of the phase and position cues for motion discrimination is due to the increased saliency of the single Gabor, this offers support for a models where position cues are detected using their salience, whereas filter-rectify-filter models do not rely on explicitly detecting features or their locations. The veridical motion discrimination for low density Gabor micro-patterns could be achieved either via filter-rectify-filter models, such as through a full-wave rectification of the input, or through a position model where the changing positions of the Gabor was tracked across frames. Ledgeway & Hess (2000) multiplied the original Boulton & Baker (1993) Gabor micro-pattern with a carrier, and randomised the internal structure of the Gabors: removing the phase cues, and low level visual information that would be used by a filter-rectify-filter model. Motion discrimination for large displacements was preserved, demonstrating that motion discrimination was not solely reliant on a filter-rectify-filter model, and that motion discrimination could be achieved by a position cue tracking model.

Finally, two psychophysical studies have found results where the response of the position system is better predicted by the displacement of features, rather than motion energy analysis applied to a

filtered and rectified input. Seiffert & Cavanagh (1998) looked at the minimum thresholds of displacement for discriminating motion in gratings displayed rotating around an annulus at different temporal frequencies. For luminance defined gratings (containing phase cues), motion thresholds were determined by a minimum velocity, as expected for a system processing motion using spatio-temporal filters. However, for gratings that were texture-based or stereo-based (containing position, but not phase cues), motion thresholds were approximately constant across temporal frequency, and determined by minimum displacement (Seiffert & Cavanagh, 1998). This suggests that the texture and stereo-based gratings were not processed by a filter-rectify-filter model, which would predict a scaling of thresholds with temporal frequency, especially if using motion energy processes post-rectification, as suggested by Wilson *et al.* (1992). Instead these results are best explained by a mechanism tracking the change in position of the features in the grating, insensitive to the length of temporal interval elapsing between the different positions that the features take. Smith & Ledgeway's (2001) high-pass filtered RDK stimuli can also be used to discuss whether a system processing motion through position cues is better characterised by a filter-rectify-filter model, or by a model tracking the displacement of stimulus features. Smith & Ledgeway (2001) demonstrated that motion could be reported either in the phase cue or position cue direction. RDK stimuli with small elements showed aliasing of the lowest spatial frequency with the displacement magnitude, and reversals of reported motion direction when polarity was reversed between frames: both results demonstrating that motion direction was discriminated using the phase cue. However, for high pass filtered RDK stimuli with larger elements (75 arcmin), motion was reported in the displacement direction, and showed no aliasing with displacement magnitude, and no reversal of reported motion with reverse phi. A filter-rectify-filter model using a full wave rectification would predict that motion would be reported in the displacement direction, but does not explain the shift in reported motion direction between small and large elements since the use of the high pass frequency filter dissociates the spatial structure of the images from the spatial frequency components. A position system based on item salience would predict this shift if the size of the smaller elements is below a critical threshold, which may be relatively high for these stimuli, given the correspondence problem in RDKs.

An alternative model for processing position cues can be taken from the field of visual search, with 'saliency maps'. It is thought that visual search can be performed by a bottom-up process, driven by the salience of features in the visual scene, or a top-down process, where the 'spotlight of attention' sweeps the scene to find stimuli meeting a specific criteria. The bottom-up process is reliant on visual features being assigned a salience value. Koch & Ullman (1985) suggest that this could be produced using parallel processing of the visual scene by various feature maps e.g. colour,

orientation, intensity, which are then combined into a topographic saliency map, and the most salient features determined by a winner takes all network (Koch & Ullman, 1985). In a model of motion perception, position cues could be tracked, not by their physical properties, but by their value and location in a salience map. This idea is similar to the long-range motion mechanism of Braddick (1980), where the location of features are mapped between frames, with Cavanagh (1992) specifying an explicit role for attention in this process. Lu *et al.* (1999) have demonstrated salience as a critical component of perceiving motion from position cues, using isoluminant chromatic gratings. Participants can discriminate motion in a isoluminant chromatic grating, which by definition has no phase cues due to the isoluminance. However, once the chromatic bars are set to be perceptually equally salience, motion discrimination is abolished, demonstrating that the position cue is reliant on salience.

One implementation by which features could be extracted from the visual input comes from early work in computer vision, pioneered by Marr, where Marr & Ullman (1981), building upon their models for edge detection using zero-crossing segments (Marr *et al.*, 1979), developed a model that could rapidly extract a wide range of velocities, even at very small displacements (Figure 1.13). Using Gaussian masks of different widths, the contours in the spatiotemporal luminance planes (the regions of sharp intensity changes), are highlighted. For the extraction of motion, these zero-crossings can be calculated separately on the distribution of intensity (or luminance) in the space (S) and time (T) domains. Marr & Ullman (1981) suggest that a directionally selective unit can be produced through the combination of two S units (derivative in space) and one T unit (derivative in time) (Figure 1.13). The S units will be positioned overlapping in space with small separation, with one ON-centre, activated by zero-crossings, and one OFF-centre, inhibited by zero-crossings, where the spatial arrangement of the S units - overlapping receptive fields with a small separation – could be used as a criterion for assessing potential physiological substrates for a model. As Koch & Ullman (1985) suggest, a similar process could occur across 8 spatial scales, for luminance, colour and orientation, to produce a topographic salience map. A computational implementation of a refined version of Koch & Ullman's model (Koch & Ullman, 1985) was able to complete a complex visual search task, demonstrating the model's ability to extract salience in a static scene (Itti & Koch, 2000): Testing more recent implementations of this model on complex motion stimuli would be an informative next step in determining whether the salience map hypothesis is a plausible mechanism for the motion discrimination of position cues.

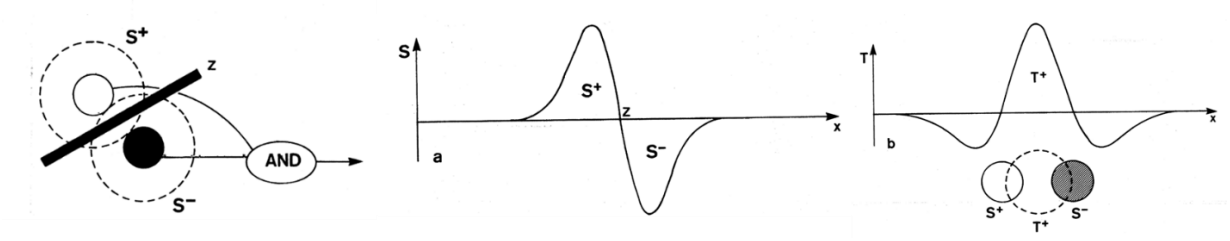


Figure 1.13. From Marr & Ullman (1981). Representations of the positive and negative S units, arranged either side of a zero-crossing 'Z' (left and central). Right hand panel: proposed organisation of the two S units with the T unit, to produce a directionally selective unit that produces a signed direction response, with a specificity of 180°. To detect stationary zero-crossings, the model requires a second, negative T unit.

Salience is not a property of a feature alone, but is a property of the relationship of features and the surrounding that they are displayed in: this necessitates a relatively large receptive field, to determine which features define salience, and sensitivity to different parameters that might define salience. One plausible site for such a salience map is V4 (Mazer & Gallant, 2003), where cells are sensitive to orientation and colour (Schein & Desimone, 1990), and receptive fields are larger than those in V1 and V2 (Smith *et al.*, 2001). The salience map hypothesis suggests that the ability to discriminate motion from position cues is produced as a by-product of the transformation of the visual scene into salience map, and not as the primary function. Discriminating motion direction in the salience map hypothesis may be the result of updating the location of salient features. The temporal limit on motion discrimination would therefore be imposed by memory, or maintenance of the current salience map, and not the extent of a receptive field oriented in time. Research using blank inter-frame intervals has demonstrated a stark difference in the temporal integration period for phase and position cues, where MFG with blank inter-frame intervals are reported in the position cue direction even with blank IFIs of 320ms, (Georgeson & Harris, 1990), whereas motion is reported in the phase cue direction for IFIs of 20ms and below. A similar distinction occurs in Ternus displays, where at short IFIs the appearance of motion as individual elements is governed by low-level processes, likely using the phase cues, whereas at long IFIs the appearance of motion as coherent groups is governed by high-level processes, feature-tracking (Braddick, 1980).

These two solutions for processing motion from position cues make separate claims about the separability of the phase and position systems. The 'rectification' hypothesis suggests that phase and position cues are processed by the same system, with a front end rectification stage for position cues, whilst the 'salience map' hypothesis suggests that the phase and position systems are entirely separate. The production of salience maps may use similar visual transformations as the initial stages of the rectification models: the difference between these models is whether features are

subsequently processed by a motion specific system, or if motion is tracked between salience maps using updated locations with time.

The rectification model and salience map model both have some plausible physiological properties. Some of the transformations suggested by rectification models are physiologically plausible, where spatial and orientation filtering could be achieved through the centre-surround receptive fields of retinal ganglion cells, and full wave rectification could occur through the sum of responses of an ON-centre and OFF-centre unit with matching spatial properties, whereas half wave rectification could occur through taking solely the outputs of ON-centre units. Likewise, a salience map could be plausibly represented in V4, with neurons sensitive to both colour and orientation (Schein & Desimone, 1990).

It is also possible there are both systems present for processing position cues, an idea endorsed by Lu & Sperling in their 'three-systems architecture for motion perception: a motion energy system to process phase cues, and then two systems able to process stimuli 'invisible' to a motion energy model - a filter-rectify-filter model, and a third system where attention based feature tracking occurs (Lu & Sperling, 1995a; Lu & Sperling, 2001). There is experimental support for this idea, whereby contrast modulated-gratings fail a pedestal test when presented at low contrasts, but are unaffected by the inclusion of a pedestal at high contrasts (Derrington & Ukkonen, 2000). This suggests that low contrast gratings are processed by a position system reliant on tracking and identifying features (which the pedestal test renders an incoherent cue for motion direction, Lu & Sperling, 1995a), whilst the high contrast gratings may be processed by a filter-rectify-filter model. Similar support for position cues being processed by either a rectified motion energy model, or by feature tracking, comes from Smith (1994). Smith (1994) produced a contrast modulated version of a conflicting cue grating (the 3f4f, Hammett *et al.*, 1993), and applied both inter-frame intervals which disrupt motion energy processing and masking which disrupts the appearance of features, and so disrupts feature tracking. Smith demonstrated that motion was reported in the phase cue direction when masking was applied, indicating that motion energy analysis following rectification was occurring. When IFIs were introduced to the unmasked stimulus, motion was reported in the position cue direction, demonstrating that a feature tracking mechanism was also able to discriminate motion for this stimulus.

1.4.3 Conclusion

These models of potential motion processing mechanisms aim to produce parsimonious explanations for human motion perception. The models presented here demonstrate different methods for extracting phase and position cues from the visual scene, demonstrating how the same image characteristics could be processed by different computations. These provide plausible

mechanisms that could support separate systems for extracting and processing motion from the different descriptions of motion.

Several simple biologically plausible mechanisms have been proposed for processing phase cues, and obtained consensus across several mathematically similar models. There are two distinct models for the processing of position cues: one reliant on transformation of the visual input and then the use of the phase processing system, and one suggesting that position cues are processed using salience maps of the visual input. Both models also have plausible physiological substrates for some aspects, such as half and full wave rectification through activity in ON-centre and OFF-centre retinal ganglion cells in the rectification model, and the representation of salience in V4 for the salience model.

Psychophysical data from 'pure' cue experiments are best explained by two separate models for processing phase and position cues, with the production of stimuli that are 'invisible' to either phase models, or position models, but where motion can be successfully discriminated e.g. pedestal gratings that contain no coherent position cues (Van Santen & Sperling, 1984), and drift-balanced and micro-balanced stimuli that contain no coherent phase cues (Chubb & Sperling, 1988). Likewise, the difference in interocular processing for phase and position cues (Lu & Sperling, 1995) supports the initial separation of motion mechanisms for the processing of phase and position cues, reflected in the different initial stages of the models for phase and position cues.

Further evidence suggesting that phase and position cues are processed by separate systems, comes from psychophysical experiments investigating whether participants can integrate between phase and position cues. This relies on the assumption that if the phase and position cues are processed by the same system, then observers should be able to integrate motion between the phase and position cues. Ledgeway & Smith (1994) demonstrated that motion can be perceived across interleaved stimuli containing different position cues (different contrast modulated gratings), but not between interleaved stimuli containing phase and position cues (luminance and contrast modulated gratings). This failure to integrate between phase and position systems suggests that separate mechanisms process phase and position cues.

The models described in this section describe how different transformations of information in the visual array could be used to separately process phase and position cues. They provide a theoretical description of how separate motion processing mechanisms could co-exist in the visual system, where the presence of multiple motion processing mechanisms might explain the patterns of psychophysical dissociation observed for stimuli containing conflicting or single phase and position cues. To discriminate between the possible mechanisms producing motion discrimination in the human visual system, the best evidence comes from looking stimuli that produce different

predictions of motion in the phase and position descriptions: reversals of perceived motion in reverse phi, or the sensation of motion from paired stationary luminance modulations (Flynn & Shapiro, 2018) strongly endorse a mechanism that uses the spatio-temporal luminance gradient, and the perception of motion in stimuli without coherent phase cues (Solomon & Sperling, 2004; Lu & Sperling, 1995b) suggest a mechanism using position without spatio-temporal luminance gradients.

These models are divided into different categories based on the mechanisms applied to the visual input, and not the description of the stimuli themselves. For example, whilst the Elaborated Reichardt detector is reported in Chapter 1.4.1. as a 'model for phase cues', it will respond to equally well to luminance signals produced by changes in position as changes in phase: it is not discriminating between the two descriptions of motion. This is the same for all models presented in this chapter – when applied to stimuli where the mathematical description of the change in phase of the Fourier components is equivalent to the change in spatial position of the features in an image, such as for a drifting sine wave grating, mechanisms designed to extract 'phase' or 'position' cues would produce the same response to these stimuli. Therefore manipulations of phase and position cues cannot be said to solely activate mechanisms designed to process phase or position cues, as for the majority of stimuli both mechanisms will be simultaneously active.

1.5 The organisation of the visual system for motion perception

Psychophysical data demonstrate that humans are able to discriminate motion using both phase and position cues, with psychophysical dissociations in sensitivity across spatial and temporal parameters suggesting that these cues may be processed by separate mechanisms, or at least initially processed using different motion detectors. Whilst models of motion mechanisms can be produced to provide predictions of motion direction that align with the perceptual experience of observers, this is not evidence that these mechanisms exist in the visual system, or that the visual system separately processes motion information aligned along the description of motion in terms of phase and position cues.

An alternative method for investigating how the visual system might process visual motion is to look directly at the physiology of the visual system and investigate what mechanisms are plausible given the response and functions of cells. This method could also be applied to known visual pathways, where activity in pathways could be correlated with perceptual experience, and similarly, changes in perceptual experience when pathways are selectively lesioned could be used to produce more causal relationships between pathways and possible mechanisms.

1.5.1 Properties of cells in the early visual system

If there are separate motion systems that process motion using either phase or position cues, then early receptors in the visual system should be optimised for the detection of these cues.

An obvious flaw in the proposal that the visual system extracts motion information through the changes in phase of the Fourier components, is that it is a physiologically unrealistic proposition. This is outlined in Howards & Rogers (1985), where they describe the four necessary components of the detectors needed for a system to perform Fourier analysis: narrow spatial-frequency bandwidth, with infinite size, spatially homogenous and able to simultaneously encode both amplitude and phase. The visual system meets none of these criteria. Receptive fields in the visual system respond to information located within bounded regions of space, and so detectors must have a finite size. Similarly, the visual system does not have spatial frequency detectors with the necessary narrow bandwidth, as shown by Blakemore & Campbell (1969), where adaptation to a sinusoid grating caused elevation in the thresholds of higher and lower spatial frequencies. The bandwidth of these spatial frequency detectors is furthermore not uniform, estimated at approximately 2.2 octaves for the lowest spatial-frequency channel and about 1.3 octaves for the highest spatial-frequency channel (Wilson, 1991), meaning that the precision of a putative Fourier analysis would vary with spatial frequency. Finally, the increasing size of receptive fields with periphery (Perry & Cowey, 1984) demonstrates the lack of spatial homogeneity of the visual system.

Similarly, the visual system is also not optimised to detect position in isolation: a perfect position detector would have a large spatial frequency bandwidth, and very small receptive fields. The insufficiently narrow spatial tuning of the visual system for Fourier analysis, is still too broad for an optimal position detector. Gabor (1946) noted this issue, looking at temporal frequency and time in acoustic signals, which is analogous to spatial frequency and position, and reported that a compromise which would reduce uncertainty equally for both dimensions was the ‘Gabor’, named in his honour: a pair of detectors with a Gaussian profile, arranged in quadrature (a phase shift of 90^{o}).

The components to form these Gabor detectors have been found physiologically, with receptive fields for adjacent simple cortical cells in the cat having with phase shifts in quadrature (Pollen and Ronner, 1981), and similarly, that 2D Gabors could describe the spatial properties of the simple cortical cells in a cat, (Jones & Palmer, 1987). Gabor detectors, or spatial filters arranged in quadrature, are specified by some motion models as the model input, such as in the Elaborated Reichardt detector (Van Santen & Sperling, 1985). This convergence between the mathematical compromise to detect phase and position, the observed physiology, and the utilisation of Gabors in descriptions of motion mechanisms suggest that Gabor functions may pose a viable and likely model for the initial inputs to motion processing mechanisms.

The possible use of Gabors for the initial inputs to motion processing mechanisms suggests that, at least in the initial stages, motion is not extracted separately from phase and position cues. This would be consistent with both motion energy models and some of the filter-rectify-filter models discussed in Chapter 1.4. If Gabors are used as the initial input for motion processing models, then this would suggest that spatial frequency dissociations between stimuli are introduced by later filter and rectification stages, and the pattern of spatial frequency dissociations could be used to further investigate which rectifications are being used. However, it is also plausible that there are separate motion systems receiving the same initial input, and that it is later transformations that make either the phase or the position cues more explicit using different comparisons between detectors.

1.5.2 The organisation of motion processing pathways

In the primate brain, the posterior bank of the superior temporal sulcus, referred to as MT (Allman & Kaas, 1971), or V5 (Zeki, 1980), is the cortical region labelled the 'motion area', due to its high density of directionally selective cells (Zeki, 1978, Zeki, 1980), and the impairment to motion perception abilities following lesions of the area (macaques; Newsome & Paré, 1988; humans; Baker *et al.*, 1991). Reported motion direction correlates with activity in this region (Britten *et al.*, 1992), even for illusory motion stimuli (Zeki *et al.*, 1993), and direct stimulation of cells in V5/MT can be used to bias reported motion direction (Salzman *et al.*, 1992), demonstrating the region's critical involvement with the perception of motion. A critical role played by V5/MT is to integrate the conflicting motion signals in the visual array. In complex stimuli produced from multiple sine wave components the displacement of the stimulus will result in each component having a different change in phase, potentially signalling motion in different directions, yet for most stimuli observers perceive one coherent motion direction. Likewise, stimuli containing phase and position cues signalling motion in conflicting directions are also perceived as having a single coherent direction, bar some cases of motion transparency (Hammet *et al.*, 1993), suggesting that there is integration between these conflicting cues. This may occur in V5/MT, where activity correlates with the perceived direction rather than the direction of the individual components, by solving the intersection of constraints of the inputs to V1 (Movshon *et al.*, 1985, Simoncelli & Heeger, 1998).

Activity in V5/MT is produced from two projections: a 'cortical' route, where activity is carried along retinal projections to the lateral geniculate nucleus (LGN), from the LGN to striate cortex, and from striate cortex to V5/MT , and a 'midbrain' route, where activity is carried along retinal projections to the superior colliculus, from the superior colliculus to the pulvinar, and from the pulvinar to V5/MT (Figure 1.14). Both the cortical route and midbrain show activity in response to viewing motion stimuli (V1: McKeefry *et al.*, 1997; LGN & Superior colliculus: Schneider & Kastner, 2005). Event-related potentials also suggest that both the cortical and midbrain pathway are active in human

observers viewing motion stimuli. These studies report two peaks of activity in V5/MT following motion onset, and the first of these precedes activity in V1, and is thought to reflect the activity of the midbrain pathway (Buchner *et al.*, 1997; Stevens *et al.*, 2007; Ffytche *et al.*, 1995). The amplitude of this non-V1 peak is smaller than the second peak following V1 activity (Stevens *et al.*, 2007), which is consistent with the relative size of the cortical and midbrain pathways. One caveat in these imaging studies is that activity per se doesn't represent a sensitivity for motion, it may be a response to the temporal frequencies in the stimuli. This is acknowledged for imaging of V1, where activity seems to be related more to the local and spatial properties of the stimulus rather than to the motion processing in V1 (Smith *et al.*, 1998; McKeefry *et al.*, 1997). However, this activity does support an interpretation of V5/MT as the site of visual motion perception, receiving projections carrying local motion signals that will be integrated into the coherent percept in V5/MT (Movshon & Newsome, 1996).

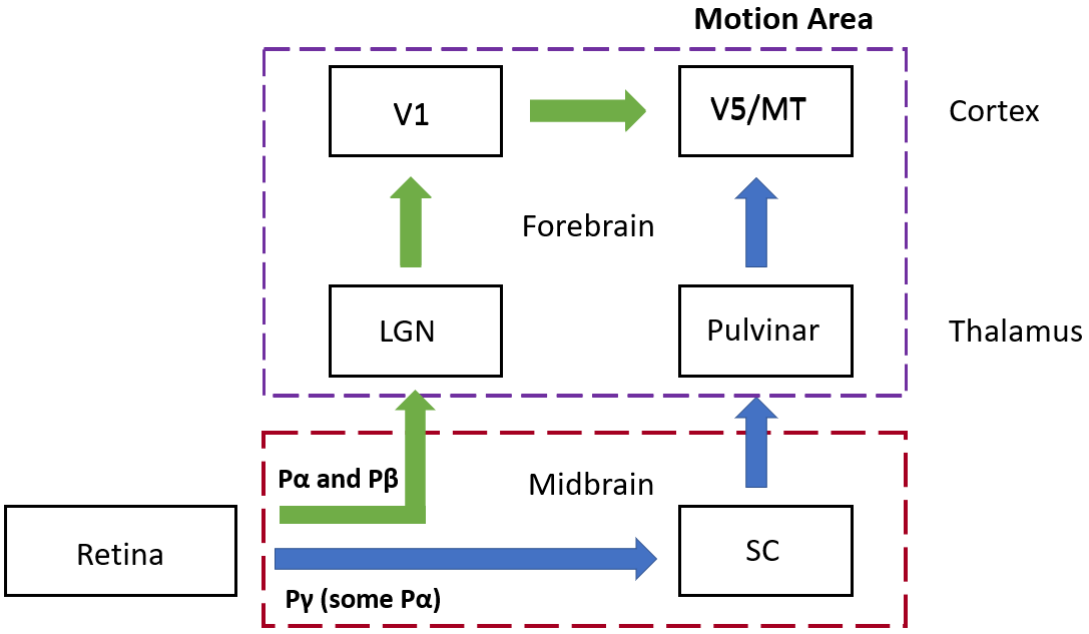


Figure 1.14. A schematic of the major projections to V5/MT : the cortical projection of pα and pβ RGCs to the LGN, through V1, and to V5/MT , and the midbrain projection of pγ and some pα RGCs to the superior colliculus, which projects to the pulvinar and to V5/MT .

The involvement of the cortical and midbrain pathways in V5/MT activity and accompanying motion perception abilities has been demonstrated by lesion studies. Recordings of V5/MT activity have found weakened but present activity following lesions of V1 (Rodman *et al.* 1990; Girard *et al.*, 1992). The reduction in this activity demonstrates the loss of the cortical pathway, with projects from the retina to V5/MT via V1, and the presence of the residual activity demonstrates the activity of the midbrain which is directionally selective (Rosa *et al.*, 2000), and sufficient for reduced but present motion perception abilities in macaques (Rodman *et al.* 1990; Mohler & Wurtz, 1977). The

midbrain pathway is demonstrated as the cause of this residual activity whereby following V1 lesions the midbrain pathway is still active in response to motion stimuli, with activity seen in the pulvinar (macaques; Warner *et al.*, 2015; humans; Barleben *et al.*, 2014), and in the superior colliculus (Tran *et al.*, 2019). Finally, lesioning the midbrain pathway in addition to V1 abolishes V5/MT activity along with the residual visual abilities (Rodman *et al.* 1990; Mohler & Wurtz, 1977; Gross, 1991; Kinoshita *et al.*, 2019).

These two pathways receive inputs from different classes of retinal ganglion cells with vastly different sized populations (Figure 1.14). In the primate retina there are approximately 20 classes of retinal ganglion cell (RGC) (Dacey, 2000; Yan *et al.*, 2020; Dacey, 2004). The primate retina mostly consists of pβ (midget) RCGs which make up ~70% of the total RGC population, with pα RGCs and pγ RGCs each accounting for a further 10% (Perry & Cowey, 1984). Since these estimates, the small bistratified RGCs have also been identified in the retina (Dacey & Lee, 1994), which make up a small but significant fraction of the total RGCs. The cortical pathway, which refers to the projection of visual signals from the retina via LGN and V1 to V5/MT , receives input from all of the pβ RGCs and the majority of pα RGCs and the small bistratified RGCs (Perry & Cowey, 1984; Dacey & Lee, 1994). The midbrain pathway, where visual information is projected to V5/MT via the superior colliculus and pulvinar, receives input from the pγ RGCs and a small proportion of the pα RGCs (Perry *et al.*, 1984; Crook *et al.*, 2008). The midbrain and cortical pathway both contain pα RGCs but these may subserve different visual functions as a consequence of the different areas of the brain that they project to.

The functional properties of these two pathways can be inspected to determine whether there are any correlates with the dissociation of phase and position systems, and these can be compared in terms of the different brain regions of the pathways, or by the difference in the initial input of RGC classes.

1.5.3 Cortical pathway

The 'cortical pathway' refers to the projections of pα and pβ RGCs to LGN, to V1 and finally to V5/MT. Along the cortical pathway, the projections from the different classes of RGC are kept separate. The pα RGCs synapse in layers 1 and 2 of the LGN, referred to as the magnocellular layers and comprising ~10% of the LGN neurons, and the pβ RGCS synapse in layers 3 to 6 of the LGN, referred to as the parvocellular layers, comprising of ~80% of the LGN neurons (Perry & Cowey, 1984). The small bistratified RGCs (Dacey & Lee, 1994) project to the koniocellular layers of the LGN, which are 'intralaminar' layers sandwiched between the layers of the LGN structure (Roy *et al.*, 2009), and comprise ~8% of LGN neurons (Dacey, 2000). These projections remain separate in the striate cortex, where the parvocellular layers of the LGN project to layer 4Cβ of V1, and the

magnocellular layers to layer 4Cα of V1 (Chatterjee & Callaway, 2003). Due to this separation, in this thesis I will refer to the pα RGC projection to the magnocellular layers of LGN as the 'magnocellular pathway' and the pβ RGC projection to the parvocellular layers of LGN as the 'parvocellular pathway'.

Projections from V1 to extrastriate cortex are less completely segregated: there is substantial interaction between the two pathways in V1 (Van Essen *et al.*, 1992; Malpeli *et al.*, 1981), but the two pathways retain some functional segregation. This can be demonstrated using selective lesions to the parvocellular and magnocellular layers of the LGN, and examining the change in extrastriate activity, where the lesion of layers of the LGN can be considered a proxy for lesioning the parvocellular and magnocellular layers of V1 (Maunsell *et al.*, 1990; Ferrera *et al.*, 1992, 1994). Lesions to the magnocellular layers of the LGN greatly reduced, but did not completely abolish, activity in V5/MT , whilst lesions to the parvocellular pathway had only a small effect on V5/MT activity (Maunsell *et al.*, 1990). Consistent with this result, Nassi & Callaway (2006) examined projections to V5/MT /MT using retrograde tracer, and reported that whilst the majority of the input to V5/MT was formed from the direct projection from the magnocellular layers of V1, there was an indirect projection to V5/MT from the parvocellular layers. These results suggest that activity in V5/MT , the cortical motion areas (Zeki, 1980), is predominantly driven by the magnocellular layers of V1.

V5/MT is not the only area in extrastriate cortex to receive a significant magnocellular projection. Area V4 in the cortex is reportedly specialised for processing chromatic information (Zeki, 1973), but also contains neurons sensitive to orientation (Zeki, 1983), and cells that respond to both colour and orientation (Schein & Desimone, 1990). Lesions to the parvocellular layers of the LGN did reduce the magnitude of activity in V4, but a similar reduction in activity was found following lesions of the magnocellular LGN pathway (Ferrera *et al.*, 1992; Ferrera *et al.*, 1994), demonstrating that input to V4 is less decisively separated between the magnocellular and parvocellular layers. There are reciprocal projections between V4 and V5/MT (Maunsell & Van Essen, 1983; Ungerleider & Desimone, 1986), suggesting that whilst characterised as the cortical area for colour, there may be some involvement with motion perception, such as motion-induced position shift (De Valois & De Valois, 1991; Ramachandran & Anstis, 1990). The projections to V4 from both parvocellular and magnocellular layers of the LGN (Ferrera *et al.*, 1992; Ferrera *et al.*, 1994), suggest that post-V1 the separation of the parvocellular and magnocellular pathways is less distinct, and demonstrates the involvement of both pathways in the many cortico-cortical pathways responsible for the complexities of the human visual experience (Zeki & Shipp, 1988; Salin & Bullier, 1995).

One method to determine differences in function between the cortical and midbrain pathways is to demonstrate a double-dissociation, separately lesioning each pathway and observing the changes in visual perception. However, a critical complication in this approach is that lesions of the midbrain pathway produce impairments in eye movements (macaques: Schiller *et al.*, 1987; humans: Pierrot-Deseilligny *et al.*, 1991), and in spatial attention (Robinson & McClurkin, 1989; Krauzlis *et al.*, 2013), producing limiting confounds for any studies aiming to investigate the activity of the cortical pathway without simultaneous activity of the midbrain pathway. Given the low contrast sensitivity and poor acuity of residual vision mediated by the midbrain following lesions of V1 (Barbur *et al.*, 1994), studies investigating the activity of the cortical pathways often use stimuli that are unlikely to produce significant activity in the midbrain: with such stimuli, we can consider the midbrain pathway functionally lesioned.

Investigations into the perceptual correlates of the cortical pathway have focused on the difference between the magnocellular and parvocellular projections within the cortical pathway. The functional properties of these pathways have been examined by producing selective lesions of the parvocellular and magnocellular layers in LGN and observing the consequent deficits in visual abilities of macaques. Lesions to the parvocellular layers significantly impaired the ability on aspects critical to feature detection. Discrimination of colour and texture was significantly impaired (Merigan & Eskin, 1986; Schiller *et al.*, 1990), and there was a 3-4 fold reduction in visual contrast sensitivity across all static spatial frequencies tested (Merigan *et al.*, 1991a). However, parvocellular lesions did not affect sensitivity to temporal motion, with macaques having “normal sensitivity for stimuli of high temporal, low spatial frequency” (Merigan & Eskin, 1986).

Conversely, lesioning the magnocellular layers of the LGN had little effect on visual spatial acuity, but did dramatically impair temporal acuity, measured by critical-flicker-fusion thresholds (Schiller *et al.*, 1990), and the discrimination of low contrast stimuli, which were impaired significantly more at higher temporal frequencies (Merigan *et al.*, 1991b; Schiller *et al.*, 1990). Taken together, these results suggest that the parvocellular layers of the LGN have high spatial acuity which declines with increasing temporal frequencies, whilst the magnocellular layers of the LGN are particularly sensitive for high temporal frequencies.

These selective lesions studies demonstrate that the magnocellular layers are a likely substrate for motion perception, as are the parvocellular layers for detecting image features. However, whilst these lesion studies highlight the role of the parvocellular layers in the detection of static features, these are the stimulus elements essential for discriminating motion from position cues when phase cue information is unavailable, such as in the high spatial frequency reliant translating Kanizsa

squares stimulus proposed by Seifert & Hock (Seifert & Hock, 2014). This is tentative evidence suggesting that the parvocellular layers of the LGN may operate as the input for the position system. The lack of broadband temporal sensitivity is only critical for models of position cue processing that are based in spatiotemporal coordinates – such as the filter-rectify-filter models – and not for models that are not dependent on the extent of time between changes in position, such as the salience map hypothesis.

1.5.4 Midbrain pathway

The midbrain pathway consists of approximately one tenth of the number of RGC projections, with 7-10% of all retinal ganglion cells projecting to the superior colliculus (Perry & Cowey, 1984), with the vast majority of all other RGC axons projecting to V1, via the LGN. As shown in Figure 1.13 this projection is formed mostly from pγ and pε classes of retinal ganglion cells (Perry & Cowey, 1984; Rodieck & Watanabe, 1993), but also from a smaller proportion of the much more numerous pα, or 'parasol' cells (Leventhal *et al.*, 1981; Schiller & Malpeli, 1977; Crook *et al.*, 2008). These RGCs project to the superior colliculus (Perry *et al.*, 1984), where they synapse with cells projecting to the pulvinar (Benevento & Standage, 1983), and project from pulvinar to V5/MT (Standage & Benevento, 1983; Berman & Wurtz, 2010).

Studies of motion perception in the midbrain pathway have primarily been conducted in patients with lesions of V1. Residual vision following lesions of V1 is qualitatively different from vision mediated by both the cortical and midbrain pathways, where patients report that vision in the scotoma is 'vague and shadowy… neither shape nor colour can be attributed to it' (Riddoch, 1917). Patients with V1 lesion are impaired across all detection and discrimination tasks at low contrast, as demonstrated by the substantially lower spatial contrast sensitivity function in patients' scotoma when compared to their sighted field (Sahraie *et al.*, 2008**,** Figure 1.15, central panel). Patients' ability to detect and discriminate motion is far better preserved than their sensitivity to colour and form. Across patients Riddoch reported a consistent bias where patients were able to detect stimuli when moving that they could not detect when static (Riddoch, 1917). Sahraie *et al.* (2008) has replicated this result more recently demonstrating a significant increase in the detection of a Gabor with spatial frequency of 1cpd when displayed at 10 Hz, compared to detection when static (Figure 1.15, right panel). This could be explained by mappings of patients' direction discrimination across a range of temporal frequencies and contrasts, which show that patients are especially impaired at low temporal frequency, below 5 degrees per second, and at low contrasts, below 50%, (see Figure 1.16, RDKs: Giaschi *et al.*, 2003, bars: Barbur *et al.*, 1994). Patients also have a bias for stimuli at lower spatial frequencies, with studies employing gratings and Gabors classically of 1 cpd, but also as low as 0.3 cpd (Morland *et al.*, 1999). As shown by Barbur *et al.* (1993) in patient GY, the detection of

stimuli rapidly declines for stimuli presented at spatial frequencies greater than 1 cpd (Figure 1.16, left panel).

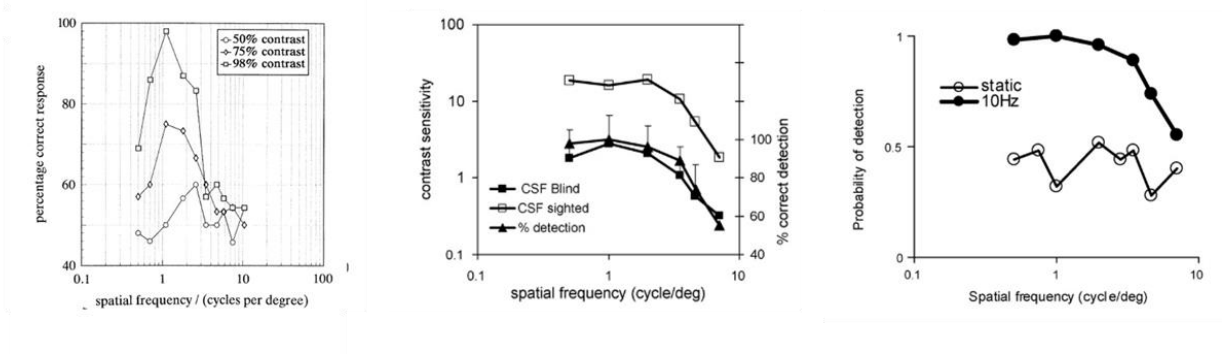


Figure 1.15. Left: From Barbur *et al.* (1993): Percent correct for patient GY, across spatial frequencies measured at 50, 75 and 98% contrast. Performance demonstrates a narrow range of sensitivity, centred around 1 cpd. Centre and Right: Figures from Sahraie *et al.* (2008) demonstrating the spatial and temporal frequency sensitivity of 8 V1 lesioned patients. The central panel displays the comparison between the sighted and blind field of patients, using 10° Gabors displayed across a range of contrasts and spatial frequencies: demonstrating the substantially lower contrast sensitivity in the blind field (scotoma). The rightmost panel demonstrates how patients' motion detection in the blind field dramatically increased with when the stimulus was temporally modulated instead of static.

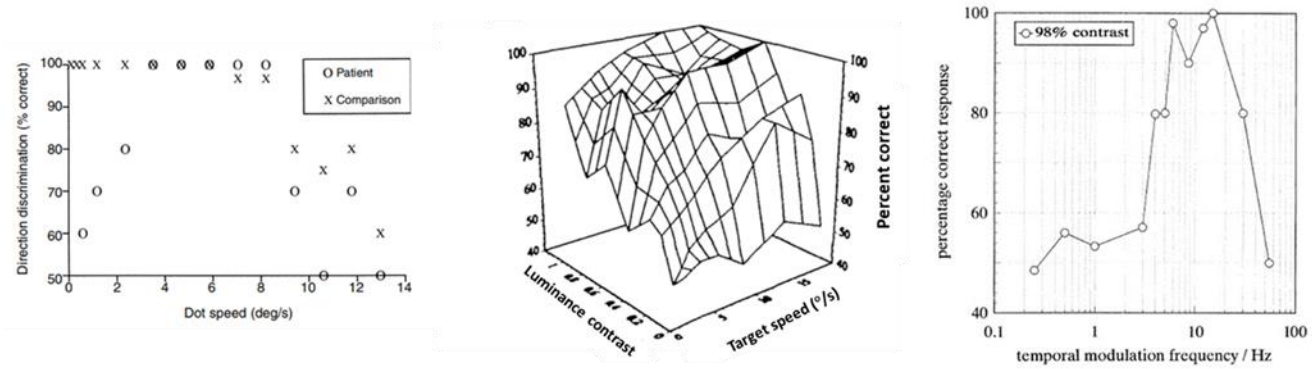


Figure 1.16. Left: From Giaschi *et al.* (2003). Comparison between motion discrimination of a moving dot pattern between a patient with a V1 lesion and three age matched control participants. The patient has substantially impaired abilities for displays presented at speeds below 3.5 degrees per second when compared to the matched controls. Central: Barbur *et al.* (1993). Direction discrimination of a moving bar by patient GY, as a function of contrast and target speed. Direction discrimination is high for high contrast fast moving bars, but impaired at low contrasts and speeds. Right: From Barbur *et al.* (1994), Gabor where the sinusoid carrier had a spatial frequency of 1.1 cpd.

This low spatial frequency bias for midbrain-mediated motion perception could reflect the spatial frequency sensitivity of the RGCs in the midbrain pathway, or could reflect the reduction in overall spatial resolution of the pathway due to the substantially fewer total number of RGCs compared to the cortical pathway. In reference to possible dissociations between phase and position systems, Azzopardi & Hock (2011) reported that patient GY was impaired at discriminating shape, and on a conflicting cue stimulus reported motion in the phase cue direction. This suggests that the midbrain pathway can process phase cues for motion, and replicates the finding that the position system can

only process motion from position cues when the features of the stimuli are discriminable (Smith & Ledgeway, 1994; Cavanagh, 1990). This may suggest that the midbrain pathway is unable to process motion from position cues, or that the stimuli used were insufficiently discriminable, and that with more visible features the midbrain pathway could potentially process position cues for motion.

In normal observers with intact cortical and midbrain pathways, Pizlo *et al.* (Pizlo *et al.*, 2008) conducted an fMRI study to observe activity in the midbrain and cortical pathway when observers viewed an apparent motion and a 'phi' stimulus. Pizlo *et al.* employed two circular stimuli: one composed of eight dots rotating through nine positions, and one composed of a single dot moving around the same nine positions. When displayed at 220 milliseconds per revolution, both the eight dot and single dot conditions were perceived by participants as beta motion, smooth motion of the dots around the circle, despite only nine locations of the circle being occupied. However, for the eight dot condition, when presented at 40 milliseconds per revolution, phi motion was experienced by participants: a shadow, the same luminance as the background, rapidly travelling around the circle obscuring the perceptually stationary filled 9 locations. For the single dot condition displayed at the same speed of 40 milliseconds per revolution, participants experienced beta motion. The 'phi' stimulus is an example of objectless motion, a stimulus containing both phase and position cues, but where the detection of the position cue is apparently below threshold (Wertheimer, 1912; Steinman *et al.*, 2000). The apparent motion stimulus has a similar construction, but with detectable phase and position cues. Pizlo *et al.* (Pizlo *et al.*, 2008) reported significant activity in the midbrain pathway for both stimuli, but favouring the phi stimulus. This suggests that the midbrain pathway is responsive to motion stimuli which don't contain position cues, and thus could be considered as a possible substrate for the phase system.

These results provide preliminary evidence that the midbrain pathway may be able to process phase cues (Azzopardi & Hock, 2011; Pizlo *et al.*, 2008). Selective lesions of the magnocellular and parvocellular LGN layers in macaques produce a pattern of dissociations in spatial and temporal frequency that suggests that the magnocellular pathway may be better suited for processing phase cues, and the parvocellular pathway for processing position cues. The midbrain pathway is composed of RGC projections from the pα and pγ class, with no projections from the pβ RGC class. If we assume that the pα RGCs are producing pathways with a similar spatiotemporal sensitivity in the cortical and midbrain pathways, then the response of the midbrain pathway to phase cues (Pizlo *et al.*, 2008), and the preference for phase over position cues in patient GY with lesioned V1 (Azzopardi & Hock, 2011), supports a tentative interpretation of the pα RGCs as an input to the phase system, across the midbrain and cortical pathways, and the pβ RGCs as an input to the position system, in the cortical pathway only.

1.6 Conclusion

Motion in the visual input can be described in terms of phase and position cues. Human participants are able to perceive and discriminate motion using either cues, with psychophysical dissociations reported when stimuli are displayed with different spatial and temporal parameters. These psychophysical dissociations may be the result of differences in the systems and mechanisms used to produce coherent motion percepts from these two classes of motion cue. Models of motion perception designed to extract motion from the visual input contain different proposed transformations and mechanisms of the visual input for phase and position cues. Whilst the majority of the visual input contains congruent phase and position cues, often from the same visual features, there are stimuli 'invisible' to all models suggested. This supports the hypothesis that the visual system contains two separate systems for processing phase and position cues.

Another approach to determine whether human motion perception is subserved by two separate motion processing systems is to look at the structure and function of the human visual system, to investigate whether there are plausible candidates for separate systems for processing phase and position cues. These data suggest that the physiological substrates for the systems processing phase and position cues may not be a dissociation between the cortical and midbrain pathways, but instead a dissociation between the classes of retinal ganglion cell forming the initial inputs for two separate systems. The cortical pathway contains both pα and pβ RGCs, projecting to the magnocellular and parvocellular layers of the LGN and V1, whilst the midbrain receives projections from pα and pγ RGCs. If the pα and pβ RGCs form the substrates for the phase and position systems respectively, this could explain the overall loss of motion sensitivity following lesions to V1 but explain the possible residual motion discrimination for phase cues (Azzopardi & Hock, 2011).

1.7 Thesis Overview

In this thesis, I propose that the human visual system has two distinct systems that process the phase and position cues for motion, and that these two systems have different physiological substrates. To investigate this claim, I explored whether the psychophysical dissociation between the spatial frequency sensitivity of the phase and position systems can explain qualitative differences between motion perception at the fovea and periphery, and whether these dissociations can be correlated with physiological substrates.

I collected motion discrimination thresholds at fovea and across periphery to investigate differences in the spatial frequency sensitivity to phase and position cues across the visual field. I then investigated reported motion direction in two forms of conflicting cue stimuli to determine whether eccentricity driven reversals in reported motion could be explained by differences in the relative

sensitivity of the phase and position systems to the spatial frequencies of the stimuli. To investigate whether these psychophysical dissociations can be correlated with physiological substrates, I investigated sensitivity to phase and position cues in a patient with a lesion to V1. Finally, I compared simulated cell density distributions of different classes of retinal ganglion cells and cells in LGN, with the displacement thresholds for phase and position cues gathered across the retina in Chapter 3.

Psychophysical dissociation

In the first half of the thesis I investigated motion perception for phase and position cues, to explore whether patterns of reported motion support a psychophysical dissociation of the separate systems processing phase and position cues.

Chapter 2: Exploring stimuli for spatial and temporal frequency dissociation of phase and position cues

Researchers investigating phase and position cues for motion perception have proposed that the ability to discriminate motion from phase or position cues is different across the spatiotemporal domain: when viewing congruent stimuli motion discrimination is reliant on phase, and not position cues, at the highest spatial and temporal frequencies (Foster *et al.*, 1989; Wertheimer, 1912). Similarly, in when viewing conflicting cue stimuli, participants report motion direction in the position cue direction at low spatial and temporal frequencies, and in the phase cue direction at high spatial and temporal frequencies (Gorea, 1994). These experiments have used a wide variety of stimuli, and it is possible that their results reflect differences in stimulus construction, rather than the underlying sensitivity of the system processing these cues.

In Chapter 2 I developed and tested a novel grating stimulus designed to isolate phase cues to conduct preliminary investigations into differences in temporal integration between the phase and position systems. I investigated the potential psychophysical dissociation of the phase and position systems, using four Gabor stimuli. I tested stimuli where coherent motion was produced through the a phase shift of the carrier or the displacement of the aperture, potentially mapping to phase or position cues, and 'conflicting' cue stimuli in which the phase change of the carrier and the displacement of the aperture signal motion in conflicting directions. The aim of the single cue stimuli was to investigate the response of the phase or position system in isolation, whereas using conflicting cue stimuli enabled me to compare the relative strength of the phase and position systems in response to the stimuli presented. I used a three-dimensional Fourier analysis (Shapiro *et*

al., 2010), to describe the phase cues in the stimuli, and discussed the application of Fourier analysis on rectified images of the stimuli for describing position cues.

Chapter 3: Spatial frequency sensitivity across the visual field

Research using apparent motion stimuli has suggested that at high spatial frequencies, participants are able to discriminate motion using phase cues between two spatially unresolvable locations e.g. when the position cues are below the detection threshold (Foster *et al.*, 1989). This could suggest that there are two distinct systems processing phase and position cues, with different spatio-temporal arrangements of the initial detectors.

In Chapter 3 I collected displacement thresholds for motion discrimination using Gabors where motion was produced from the phase shift of the carrier or the displacement of the aperture, at fovea and up to 15 degrees eccentricity. The phase shift of the carrier produces a congruent position cue, due to the corresponding displacement of the carrier features, and the limitations of this stimulus are discussed. These stimuli were chosen as to investigate the dissociation between the phase and position systems in the spatial frequency domain across the visual field. I compared the gradient of displacement thresholds across eccentricity to discuss whether differences in displacement thresholds for these carrier and aperture manipulations support an interpretation of two separate motion processing systems. I also compared these gradients with published cortical magnification gradients to make preliminary correlations with physiological substrates.

Chapter 4: Perceived motion in Gabor based conflicting cue stimuli

In Chapter 3 I demonstrated that the difference in displacement thresholds for Gabors where motion is produced either through a phase shift of the carrier or the displacement of the aperture increases with periphery. This could explain the difference in the reported motion direction for conflicting cue stimuli which, when viewed foveally, are more likely to be reported in the position cue direction, and when viewed peripherally, more likely to be reported in the phase cue direction (Chubb & Sperling, 1989). For Gabor stimuli containing conflicting cues, this change in reported motion cannot be accounted for by scaling the spatial frequency of the stimulus (Gurnsey & Biard, 2012). This could be explained by a dissociation in spatial frequency sensitivity functions for the phase and position systems, which scale differently with eccentricity.

An alternative explanation for the change in perceived motion between the fovea and periphery in conflicting cue stimuli was proposed by Shapiro *et al.* (Shapiro *et al.*, 2011), who suggested that the carrier and aperture cues are processed separately at the fovea and integrated in the periphery. In Chapter 4 I compared whether reported motion direction in Gabor based conflicting cue stimuli is

better accounted for by Shapiro's explanation, or by my hypothesis that a dissociation in spatial frequency sensitivity functions for the phase and position systems which scale differently with eccentricity.

Chapter 5: Perceived motion in reverse phi stimuli

Reverse phi stimuli are produced from the simultaneous displacement and contrast reversal of a feature. The reversal in contrast reverses the direction of the phase cues in the stimulus, producing a conflicting cue stimulus where the phase and position cues are derived from the same feature and signal motion in opposing directions.

In Chapters 3 and 4, I proposed that the reversals in perceived motion for conflicting cue stimuli presented at different spatial frequencies or at different locations in the visual field, could be explained by differences in the spatial frequency sensitivity of the phase and position systems. This hypothesis was developed using single and conflicting cue Gabor stimuli. To further validate this hypothesis, I explored whether it could explain the reported motion in the reverse phi stimulus, presented in fovea and periphery.

Physiological dissociation

In the second half of the thesis, I investigated potential physiological substrates for the phase and position motion systems.

Chapter 6: Residual motion abilities following lesions of striate cortex

Motion perception in the primate brain appears to be mediated through two major pathways: the midbrain route, where signals from the retina reach V5/MT via the superior colliculus and the pulvinar, and a second, larger pathway, where signals from the retina reach V5/MT via the lateral geniculate nucleus and striate cortex. Following lesions to striate cortex, visual capabilities are much reduced. Patients have some residual motion perception but have a substantial loss in spatial acuity and the ability to localise and discriminate shape (Riddoch, 1917). As shown in apparent motion studies, and low contrast sine waves (Kelly *et al.*, 1979), once access to position cues is reduced motion perception is reliant on phase cues in the image. The residual motion discrimination capabilities shown in V1 lesioned patients use stimuli containing congruent phase and position cues (bars; Azzopardi & Cowey, 2001, circles; Barbur *et al.*, 1993, sine wave gratings; Morland *et al.* 2004) could be produced from phase cues, and be obscuring the loss of sensitivity to position cues.

One case study testing this hypothesis has reported that when viewing a reverse phi stimulus, patient GY reported motion in the phase and not position cue direction in his blind field, whilst

reporting the same stimulus in the position cue direction in an eccentricity matched location in his sighted field. This is preliminary evidence that lesioned patients may have access to phase but not position cues. These pieces of evidence suggests that residual motion may be mediated solely by the phase system, which would suggest that the midbrain pathway may be a physiological substrate for the phase, but not the position system.

In this chapter I explored the residual motion capabilities of a patient with a striate cortex lesion across a series of congruent and conflicting motion cue stimuli.

Chapter 7: Correlations between psychophysics and physiology

The phase and position systems can be psychophysically dissociated in terms of spatial frequency, as discussed in this thesis, but also in terms of sensitivity to temporal frequency (Lu & Sperling, 1995) and chromatic information (Lu *et al.*, 1999). This dissociation correlates with the functional properties of the parvocellular and magnocellular layers of the lateral geniculate nucleus, revealed through lesion studies (Merigan & Eskin, 1986, Schiller *et al.*, 1990). These receive input from separate classes of retinal ganglion cells (RGCs).

These retinal ganglion cells can be considered as sampling units for visual systems, where the density and arrangement of these sampling units would specify the limit of spatial resolution across the retina. This approach has been applied to research into human acuity for static stimuli, where the change in density of pβ RGCs across the retina correlates well with the decrease in acuity for static stimuli presented at different locations on the retina (Watson, 2014). If different classes of retinal ganglion cells, with their distinct density and arrangement, formed the separate inputs to the phase and position systems, then this could explain differences in displacement thresholds across the visual system, and be used to form testable hypotheses about which classes of retinal ganglion cell form the input to which motion systems.

I proposed that the phase system is produced from input from the pα RGCs and the position system from pβ RGCs. These classes of RGC have different distributions of cell density across the retina, which would produce different resolution limits across the retina for the two systems. The arrangement of the receptive fields also differs, with the overlap of receptive fields in pα RGCs estimated as up to 3.5 times higher than the overlap seen in pβ RGCs (Perry & Cowey, 1984). In this chapter I produce models of retinal ganglion cell density for pα and pβ RGCs, for their density, and models that include both density and estimated receptive field overlap. I also produce models for the density distributions of cells in the magnocellular and parvocellular layers of the LGN. I

correlated the displacement thresholds collected for Gabor stimuli in Chapter 3 with these models, to propose potential physiological substrates for the phase and position systems.

Chapter 2: Spatial and temporal frequency dissociations of phase and position cues

2.1 Introduction

The motion of a stimulus can be described by changes in the phase of the Fourier components of the image, 'phase cues', and changes in the position of salient features, 'position cues'. For the majority of motion stimuli, these two descriptions are mathematically equivalent. The motion of a drifting sine wave could be described by the changing phase of the fundamental frequency, or by the change in spatial position of the features within the sine wave, such as the peaks, troughs, and zero-crossings. Both descriptions of the stimulus would predict motion in the same direction. However, laboratory generated stimuli can be designed whereby the two descriptions of the stimuli would produce predictions of motion in different directions, such as Missing Fundamental Gratings (Figure 1.2) and reverse phi stimuli (Figure 1.7). These are referred to as 'conflicting cue' stimuli.

The direction that motion is reported in for conflicting cue stimuli depends on the parameters that the stimuli are presented at. For both Missing Fundamental Gratings and reverse phi stimuli, motion is reported in the position cue direction at low temporal frequencies, and in the phase cue direction at high temporal frequencies (MFG: Georgeson & Harris, 1990; Reverse phi RDKs: Edwards & Nishida, 2004). A similar dissociation in reported motion direction for conflicting cue stimuli is observed with differences in spatial frequency presentations, where motion is reported in the position cue direction for stimuli presented at low spatial frequencies, and in the phase cue direction at high spatial frequencies (Chubb & Sperling, 1989). This is consistent with evidence that the motion discrimination of phase cues is possible at high spatial frequencies, where the spatial separation of the locations producing the motion stimulus is not discriminable, suggesting that position cues for motion would be impaired at these spatial frequencies (Foster *et al.*, 1989).

These dissociations in sensitivity to spatial and temporal frequency between phase and position cues suggest that the phase and position cues may be discriminated by separate motion processing systems. To support this interpretation, it is necessary to investigate these spatial and temporal frequency dissociations, using experiments containing stimuli which are matched as closely as possible on the construction of the stimulus, but contain separate phase and position cues. Psychophysical dissociations presented in these stimuli could then be more confidently attributed to the operation of two distinct systems processing phase and position cues.

Whilst we can describe motion stimuli in terms of the change in phase of Fourier components and spatial position, and can observe that participants can report motion is either of these predicted directions, is not evidence that the visual system contains mechanisms that utilise these descriptions

to extract predicted motion direction. This chapter is exploring whether we can produce stimuli where these descriptions of motion can be dissociated, and if so, if there are consistent patterns in reported motion direction that would suggest that these stimuli are processed by more than one system or mechanism. In this chapter I assess a temporal dissociation between phase and position cues, looking, not at temporal frequency, but at the temporal integration dynamics between motion discrimination for phase and position cues (Chapter 2.2). I investigate the possible spatial frequency dissociation between phase and position systems by investigating reported motion direction in conflicting cue Gabors, presented peripherally (Chapter 2.3). Finally, I apply Fourier analysis to the stimuli presented in this chapter and discuss the application to models of motion perception (Chapter 2.4).

2.2 Temporal Integration Dissociation of Phase and Position Systems

Grating stimuli are popular in psychophysics as they can be varied easily in terms of spatial frequency (cycles per degree of visual angle), and temporal frequency (cycles per second), and motion can described in terms of the change in phase of their Fourier components. In a grating stimulus the position cue is produced from the features of the image, constructed from the sum of all harmonic components. The phase cue in a grating is determined by the most detectable component, most often the lowest frequency component (Brown & He, 2000).

Conflicting cue stimuli contain both phase and position motion cues signalling motion in differing directions. The Missing Fundamental Grating (MFG) is an example of a grating where the phase of the lowest harmonic and the features defined by luminance contrast signal motion in opposing directions when the grating is displaced by 90° of the grating period on each frame. The MFG (Adelson, 1982; Adelson & Bergen, 1985) is produced by subtracting the fundamental frequency from a square wave. As the harmonic components are all presented in phase they sum to produce features which retain the periodicity of the fundamental on each frame, but now the lowest harmonic, with the highest amplitude, is the third harmonic (See Figure 2.1). A forward displacement of the grating by the equivalent of 90^{o} of the (absent) fundamental frequency displaces the features forward but produces a phase shift of 270^{o} in the third harmonic, which is equivalent to a backwards 90^{o} shift. This results in a backwards stimulus motion being cued in the forwards direction by the position cue and in the reverse direction by the phase cue. The perceived motion in the MFG is highly dependent on the contrast and temporal frequency of the stimulus, where high temporal frequencies and low contrast bias perceived motion in the direction cued by the change in phase. Presentations at low temporal frequency or high contrast bias in the direction cued by the change in position (Georgeson & Harris,1990; Georgeson & Shackleton, 1989).

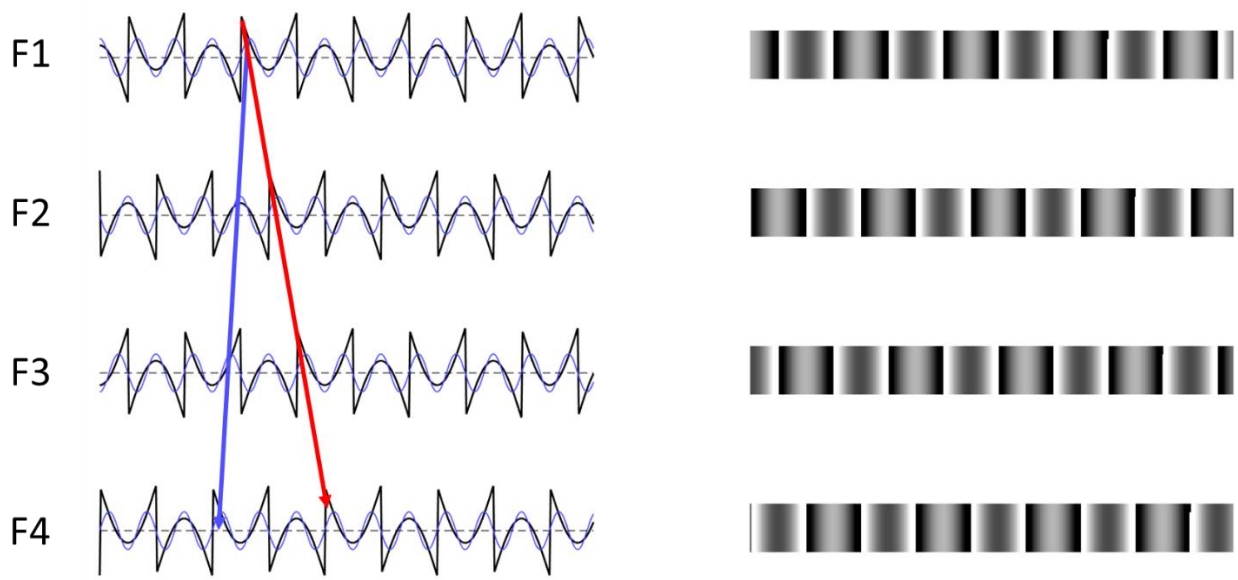


Figure 2.1. Figure displaying the cross-section (black line, left hand panel) and appearance (right hand panel) of the Missing Fundamental Grating (Adelson, 1982). I have displayed the third harmonic (blue line) on the cross-section, to demonstrate the backward aliasing in the third harmonic (blue arrow), despite the forward displacement of the features (red arrow).

A similar grating was described by Hammett *et al.* (Hammett *et al.*, 1993). This was produced from the sum of two equal amplitude sine waves with frequencies in the ratio 3:4, where displacing the grating by 90° (of the fundamental frequency that would have produced these stimuli as its third and fourth harmonics), similarly displaces the features forward, but the phase of the lowest frequency component is displaced backwards (Figure 2.2). A contrast modulated version of the 3f4f was demonstrated by Smith (1994), who multiplied the 3f4f grating with two-dimensional (2D) static noise (Smith, 1994). This stimulus still has opposing phase and position cue directions – if the phase cue is made explicit by some form of rectification. The stimulus was reported in the phase cue direction, demonstrating that the contrast modulated grating could be processed by a motion energy mechanism, following rectification. Smith (1994) demonstrated that motion was indeed being processed using the phase cue by adding a mask to the stimulus. Stationary masks do not produce a change in phase of the Fourier components of an image, so have no impact on the phase cue. Stationary masks do interfere with position cues, since the additional spatial frequencies can sum with other present spatial frequencies and obscure the form and location of features – much as the 4f harmonic is doing in the original 3f4f stimulus. The addition of a mask had no large effect, demonstrating that motion was being reported using the phase cue, following some form of rectification. Smith (1994) demonstrated that the position cues available in the contrast modulated 3f4f were also available to motion detectors, using inter-frame intervals, which disrupt processing of motion energy models by extending the stimuli outside of the spatio-temporal receptive fields. When inter-frame intervals were applied motion was reported in the displacement direction, the

position cue. This was convincingly demonstrated by the further addition of a mask: for the masked stimulus with IFI, motion discrimination was markedly impaired, as both the phase and position cue were disrupted. These results suggest that the contrast-modulated 3f4f, like the 3f4f of Hammet *et al.*, (1993), has conflicting phase and position cues.

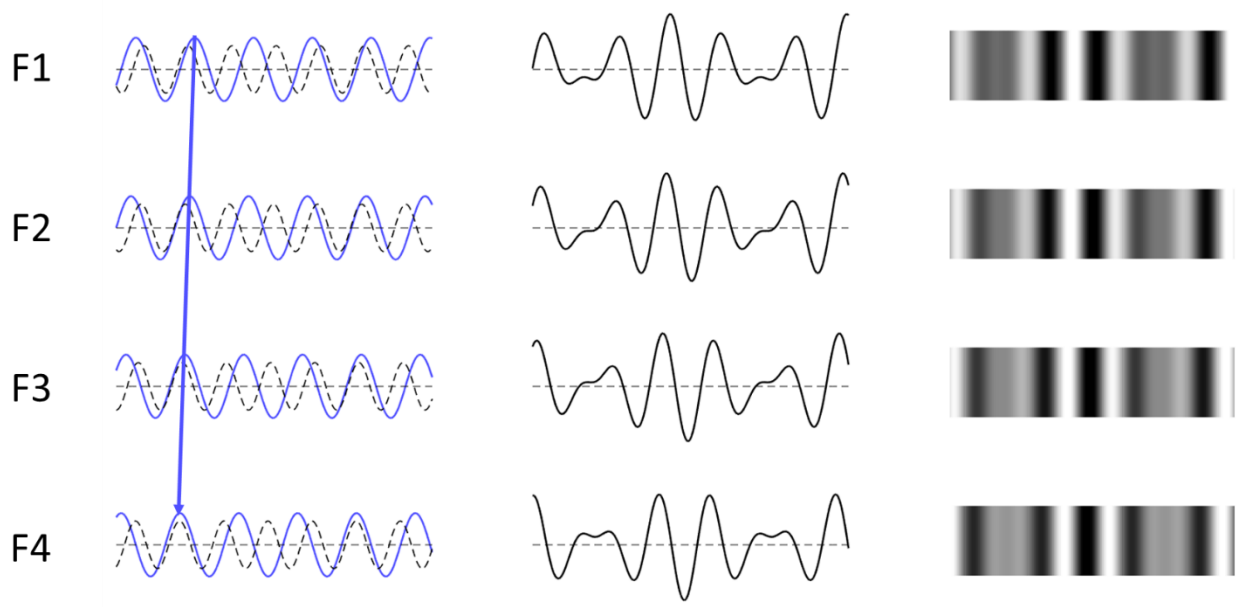


Figure 2.2. Figure displaying 4 frames of the construction (leftmost panel), the cross-sections (central panel) and appearance (righthand pane) of the 3f4f stimulus of Hammet *et al.* (1993). The leftmost column separately plots the third and fourth harmonics, with the arrow displaying the backwards direction of the phase cue, the 270° phase shift of the third harmonic, aliased as a backwards 90° displacement.

A 'pure' phase grating should contain a change of phase of the Fourier components of the grating, but without a corresponding change in position of the salient features. 'Pedestal' gratings, introduced by Van Santen & Sperling, sum a stationary low contrast sine wave and a matched frequency drifting sine wave with half the amplitude of the stationary sine wave (Van Santen & Sperling, 1984). The phase cue is carried by the drifting sine wave coherently across frames provided that it is sufficiently above participants' detection threshold. The position cue in the grating is determined by the amplitude of the sum of the gratings. As the stationary sine wave contributes twice as much amplitude as the drifting sine wave, the peaks and troughs of the visible grating will appear to oscillate around the location of the stationary amplitude. This produces an incoherent direction of motion for the position cue. After training to 'attend to the wind', participants are able to report the motion direction of the drifting component, and discard the incoherent position cue, producing a 'pure' first order motion stimulus. The pedestal gratings are an example of 'objectless' motion perception, where the motion of the phase cue can be discriminated, without being attached to a displaced feature.

A grating stimulus in which motion is solely determined by the change in location of the features is more difficult to construct because the features in a compound grating are the result of the sums of the harmonic components of the grating. Any change in the position of the grating features will be described by a change in phase of these harmonic components, producing a phase cue. However, one method of removing a correspondence in the phase of the Fourier components between frames is to render phase cues incoherent, by changing the identity of the stimulus on each frame. Any changes in phase of Fourier components will be at random, and across frames not informative as to the direction of motion. Cavanagh *et al.* demonstrated that an apparent motion stimulus which changed identity between frames, such as by texture, colour and depth, still produced a strong percept of motion in that the salient feature on each frame changed location, regardless of the properties defining the figure as salient (Cavanagh *et al.*, 1989). Grating based stimuli with a duty cycle of ½ do not undergo the same figure-ground segmentation, since there is no delineation between figure and ground. However, if features can be mapped between gratings, but formed of different harmonic components, then there should be no corresponding phase cue between the different gratings. This approach was taken by Baro & Levinson: pairs of MFGs were interleaved with a sine wave grating, produced using the same fundamental frequency, which has the same periodicity as the MFG (Baro & Levinson, 1988). I shall refer to this as the 'MF1F stimulus' (see Figure 2.3). Using a 90^{o} phase displacement ensures that the pairs of frames are in counterphase, so any coherent motion must be produced from tracking between frames. Baro & Levinson (1988) found that for two of four participants this stimulus can produce a coherent perception of motion, if using a high contrast grating with a sufficiently slow velocity. For the remaining two participants, one showed some motion discrimination abilities at the lowest velocities, and the fourth participant was at chance performance across all velocities. This inconsistency in performance could potentially be improved by further training on the stimulus, with explicit instructions about which features to match across frames, similar to the instructions used in complex feature mapping stimuli (Lu & Sperling, 1995; Lu *et al.*, 1999). For the stimuli presented in this chapter I found that the subjective motion strength increased with rectangular apertures orthogonal to the orientation of the grating, with similar effects also found for direction discrimination in plaids and the barberpole illusion (Mussap & Crassini, 1993). This effect may operate only for the phase cue, but could potentially improve performance in the MF1F stimulus. Matching between the features of the sine wave and MFG could be encouraged by asking participants to match the subjective contrast between the two gratings prior to the experiment, thus reducing the differences in feature appearance, without changing feature identity.

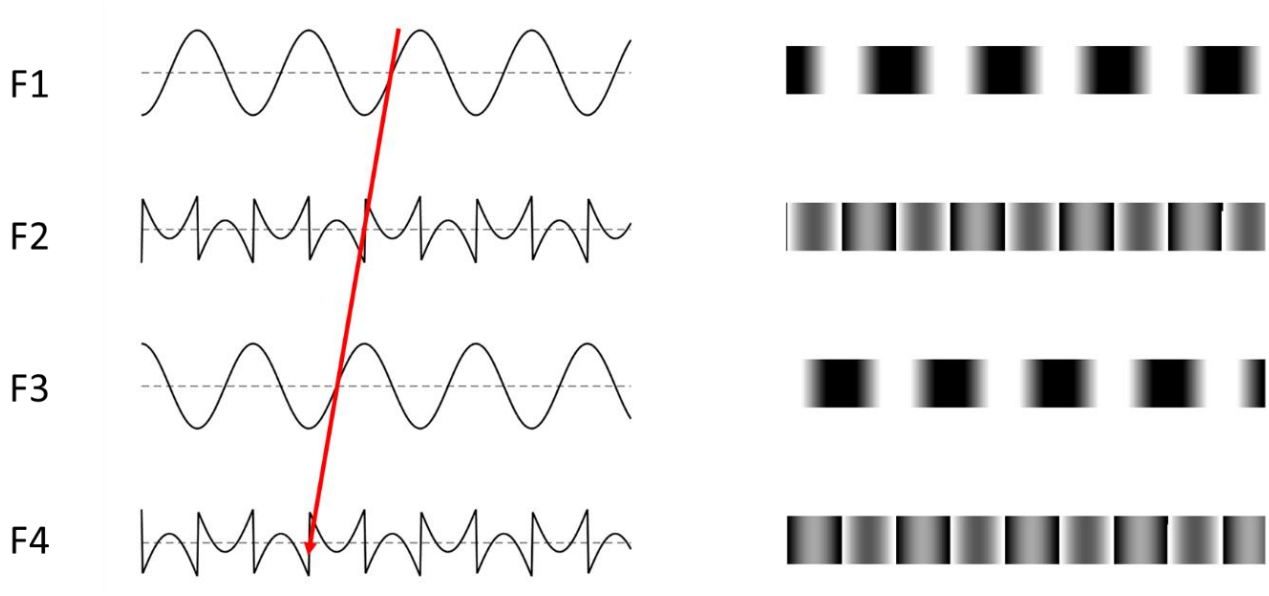


Figure 2.3 The 'MF1F', designed by Baro & Levinson, (1988). Each row corresponds to one frame of the stimulus. On each consecutive frame, the grating has been displaced by 90° of the fundamental frequency. Left; Profile plots for each frame of the MF1F stimulus; interleaved sine wave gratings and Missing Fundamental Gratings. The red arrow marks the leftward displacement of the zero-crossings, which could be used as a cue to track the displacement of the grating features. Right; Image of the MF1F stimulus presented on each frame.

I devised a matching stimulus to Baro & Levinson's (1988) 'MF1F' grating, where I interleaved frames of the MFG, but with frames of the third harmonic to match the phase cue rather than the features, the 'MF3F' (Figure 2.4). The third harmonic is three times the periodicity of the features, and when displaced by ¼ of a cycle, the peaks and troughs fall so that they are equidistant from the peaks and troughs of the MFG. If the grating is displaced by 90° of the fundamental frequency on each frame, with the change of grating identity on consecutive frames, then this displacement between frames should be perceived using the third harmonic, and so act as a 'pure' phase cue stimulus. As in the MF1F, the MFG and the third harmonic are in counterphase with their paired frames, so there are no motion cues between pairs of frames.

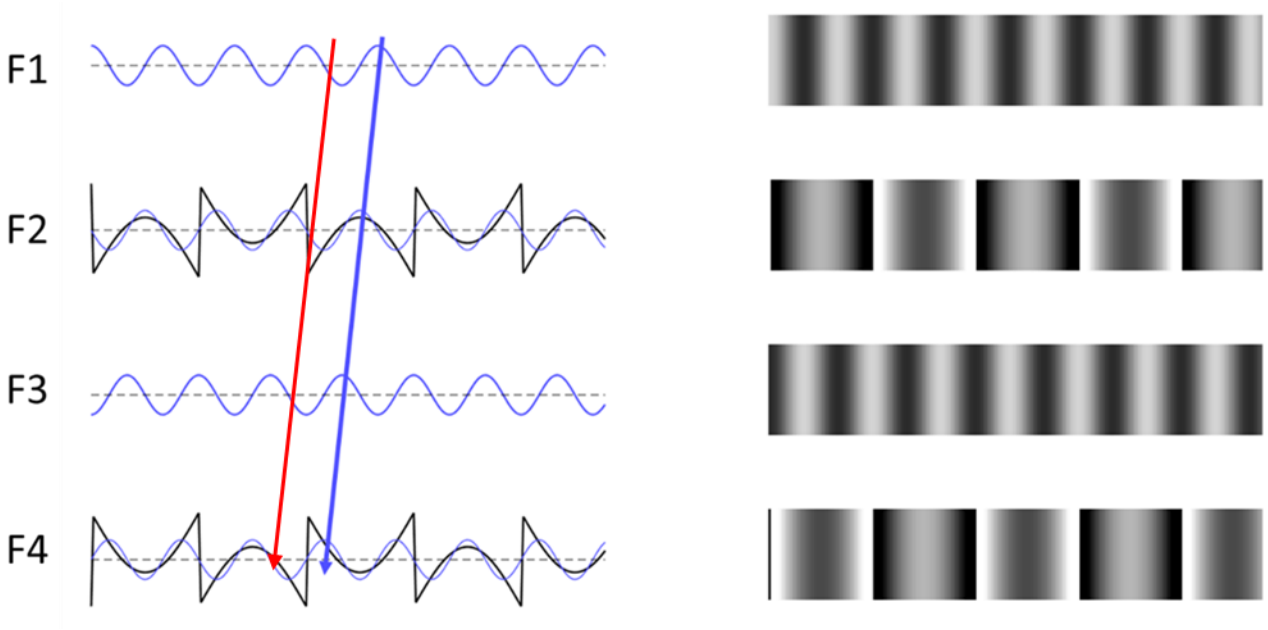


Figure 2.4. The 'MF3F', my proposed novel 'pure' first order stimulus. Each row corresponds to one frame of the stimulus. On each consecutive frame, the grating has been displaced by 90° of the fundamental frequency. Left; Profile plots for each frame of the MF3F stimulus; interleaved antiphase third harmonic and missing fundamental grating, with the third harmonic traced on the missing fundamental frames in blue, to show the quarter cycle phase shifts between each frame. The blue arrow indicates the backward aliasing in the third harmonic, and the red arrow indicates the position cue confound, the changes in zero-crossing location that could be tracked between frames. Right; Image of the MF3F stimulus presented on each frame.

To be classified as a 'pure' phase stimulus, the MF3F grating needs to produce a coherent and consistent motion percept, and carry only phase cues with no trackable position changes of features across or between frames. My first experiment assessed whether participants could consistently report motion in the MF3F stimulus, and whether this was in the direction predicted by the phase cue. The presence of phase artefacts were assessed in Chapter 2.4.2, where Fourier analysis was applied to the stimuli.

My second experiment compared motion discrimination across a range of durations of blank interframe intervals (IFI), for a novel grating based stimulus containing solely phase cues, the MF3F. The Missing Fundamental Grating (MFG), contains a matched phase cue, and a position cue signalled motion in the opposite direction. Inserting blank inter-frame-intervals (IFIs) is thought to abolish the phase cues in a stimulus, where the system processing phase cues is highly sensitive to the temporal information in the stimulus. Blank IFIs bias motion towards the position cue direction in micro-pattern Gabor arrays (Baker & Braddick, 1985; Bex & Baker, 1999), in the contrast-modulated 3f4f (Smith, 1994), and in the Missing Fundamental Grating, where IFIs above 50 ms produce a reversal of perceived motion from the phase direction to the position cue direction (Georgeson & Harris, 1990). Motion discrimination remains high in MFG across increasing durations of the blank IFI,

suggesting that the discrimination of the position cue is not impaired by the insertion of blank IFIs. This indicates that motion in the phase and position mechanisms may operate differently with time, where the phase system may integrate luminance across space and time dimensions, whilst the position system seems invariant to the increase of blank IFIs, suggesting that the position mechanism is not integrating feature locations across time.

Method

In this section I describe the methods for two experiments conducted using the MF3F grating: Experiment 1, in which I compared reported motion direction for the MF3F grating and a third harmonic grating presented at 5 different spatial frequencies, and Experiment 2, in which I compared reported motion direction in the MF3F grating and an MFG grating, presented with 5 different durations of blank inter-frame intervals. All participants participated in Experiment 1 first, followed by Experiment 2.

Prior to the two experiments all participants completed a set of training trials. In these, the MF3F grating stimulus was displayed at 100% contrast with a fundamental frequency of 1 cpd (thus third harmonic at 3 cpd), and trials were labelled with the direction of motion. The training trials ended when the participant reported that they had 'high confidence' on 3 consecutive trials; no participant required more than 20 trials. The two experiments were run in the same two-hour session, with a 15-minute break between Experiment 1 and Experiment 2.

Ethics

This study was approved by the Medical Sciences Inter-divisional Research Ethics Committee (IDREC) at the University of Oxford in accordance with the procedures laid down by the University for Ethical Approval of all research involving human participants, under reference R72149/RE001.

Participants

Four participants (one female), aged 24, 26, 28 and 64, with normal or corrected-to-normal vision.

Apparatus

Stimuli were designed and presented in MATLAB 2019, and displayed on a Dell XPS 15 LCD screen laptop (1920 x 1200 pixels), with NVIDIA® GeForce RTX graphics card and a maximum luminance of 318 cd/m^2. These experiments were run remotely, due to the COVID-19 pandemic preventing access to laboratory testing conditions. Perceptual gamma correction of the screen was set for each participant using the achromatic half tone matching procedure detailed in Xiao *et al.* (Xiao *et al.*, 2011). Participants were tested under daylight conditions, using binocular vision. Stimuli were

presented centrally on the screen and participants were asked to fixate throughout the task using a small white fixation cross (0.15 degrees of visual angle) presented centrally between trials. Participants were seated 180 cm from the screen and input responses using a wireless Logitech keyboard.

Stimuli

Stimuli were viewed through a rectangular vertical aperture set to the width of one cycle of the stimulus, and a height of five cycles of the fundamental frequency stimulus. This produces stimuli of 1 by 5 degrees of visual angle for all stimuli in Experiment 2, where stimuli all had a fundamental frequency of 1 cpd, and stimuli ranging from 0.125 cpd by 0.625 cpd (fundamental frequency of 8 cpd) to 2 by 10 degrees (fundamental frequency of 0.5 cpd), in Experiment 1.

Apparent motion was produced by displacing the grating by a distance corresponding to a 90^{o} shift of the fundamental frequency each frame. Each presentation consisted of 8 frames, each displayed for 100 ms (equivalent to 2.5Hz if no IFI included). The initial phase was randomly assigned for each trial, and for the MF3F stimulus it was randomly determined whether the grating would start with a missing fundamental or third harmonic frame. The gratings were horizontal, with the bars of the grating orthogonal to the direction of motion. The screen background was matched to the mean luminance of the stimulus, as was frame luminance of the IFIs in Experiment 2. As the screen was gamma corrected perceptually for each participant, mean luminance will have varied slightly for each participant.

Experiment 1 stimuli. There were two stimuli in this experiment: the 'MF3F' grating, which consists of interleaved frames of the missing fundamental and third harmonic gratings, and the third harmonic. These were displayed at five spatial frequencies (fundamental frequency of 0.5, 1, 2, 4, and 8 cpd), with the third harmonic equivalent to a sine wave of 3 times these spatial frequencies. The stimuli were displayed at matched amplitudes of the fundamental frequency, which produced different contrasts in the stimuli, with Missing Fundamental Gratings therefore displayed at a contrast of pi/4 (78.54%) and the third harmonics displayed at 1/3 (33.33%).

Experiment 2 stimuli. There were two stimuli in these experiments: the 'MF3F' grating, which consists of interleaved frames of missing fundamental and third harmonic gratings, and the missing fundamental grating. These were displayed with a fundamental frequency of 1 cpd, chosen as this was consistently the highest performance on the MF3F Spatial experiment. Between each frame of the stimuli, a blank frame was inserted, matched to the mean luminance of the stimulus. I

used 5 durations of IFI, chosen for comparability with Georgeson & Harris (Georgeson & Harris, 1990); 0 ms, 20 ms, 40 ms, 80 ms and 160 ms. The stimuli were displayed at 10% contrast of the fundamental frequency.

Procedure

Both experiments entailed a one-interval direction discrimination task. Following each trial participants used the arrow keys of the keyboard to report whether stimuli moved upwards or downwards. Each experiment contained two stimuli conditions, presented across five levels of the dependent variable (IFI or spatial frequency). There were 40 trials for each stimulus condition, randomly ordered for stimulus type and condition, producing a total of 400 trials per experiment. Each experiment took participants an hour to complete.

Results

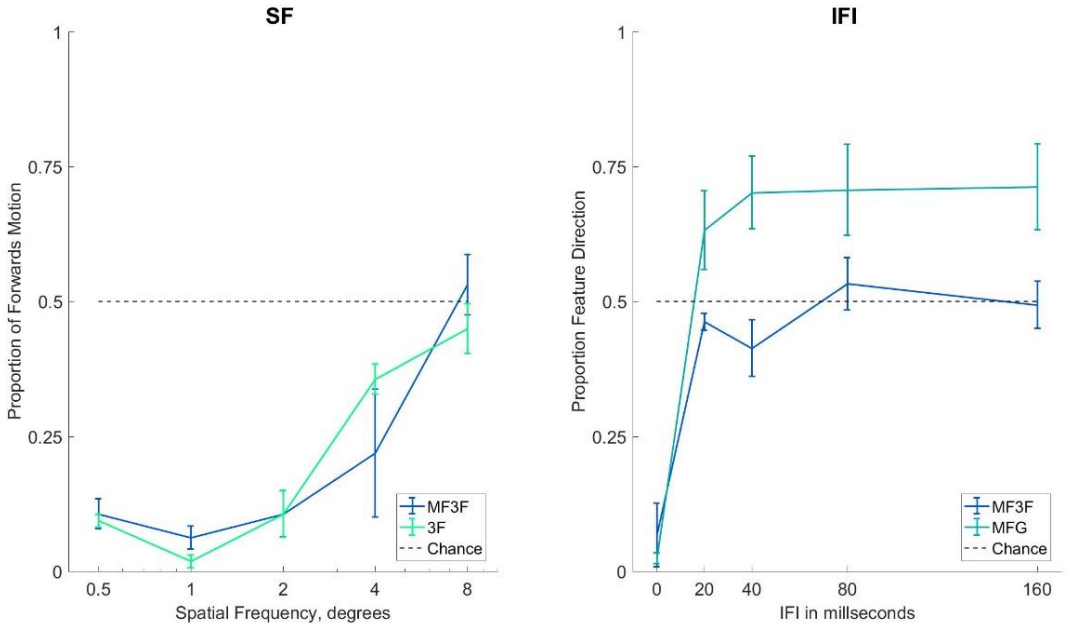


Figure 2.5 Figures showing proportion of trials reported in the displacement direction, for the MF3F and 3F stimuli across 5 spatial frequencies (left), and the MFG and MF3F stimuli at 5 lengths of inter-frame intervals (right). The dashed lines represented the best-fit non-linear regressions, and the dashed black line at 0.5 illustrates the chance performance level. Data were averaged across the four observers in each experiment, and plotted with standard error of the mean bars. In both figures performance above chance reflects motion reported in the direction of the displacement of the fundamental frequency, which is not included in any stimulus but serves as a reference for the stimulus construction. Spatial frequency is plotted in terms of the fundamental frequency, so the actual frequency of the third harmonic is three times labelled values.

To analyse the data in both experiments I operationalised reported motion direction as the proportion of trials reported in the direction of the displacement of the fundamental frequency, if it had been included in the stimulus. I referred to this as proportion correct, where one would expect the proportion correct to be below chance if using the third harmonic to discriminate motion direction, since the third harmonic is being displaced by 270^{o} on each frame, which is aliased as a backwards 90^{o} phase displacement. The results from Experiments 1 and 2 are displayed in Figure 2.5, where mean proportion correct across participants is plotted against spatial frequency (Experiment 1) or length of blank inter-frame-interval (Experiment 2), with error bars displaying the standard error of the mean.

In Experiment 1 proportion correct is below chance for both the 3F and MF3F conditions for the lowest 4 spatial frequencies, and at chance for the highest spatial frequency that the stimulus was displayed at: 8 cpd of the fundamental frequency, equivalent to 24 cpd for the third harmonic. In Experiment 2 proportion correct is below chance at the 0 ms IFI for both the MF3F and the MFG, but appears to be at chance for all non-zero IFI lengths for the MF3F, and above chance for the MFG. The figures for individual participant performance are displayed in the Appendix for this chapter, in Chapter 2.6.1.

To analyse the effect of varying spatial frequency and inter-frame interval length on reported motion direction in the MF3F grating and comparison grating conditions, I used non-linear mixed effects modelling, using the nlme package in R (Pinheiro *et al.*, 2013), and following the procedure outlined in Appendix 2.6.3. The nlme package fits a specified equation to the data, producing a set of parameters and coefficients for the overall dataset, and at each fixed effect level. The models fitted in this chapter describe the relationship between spatial frequency or inter-frame interval and performance at each condition.

Effect of Spatial Frequency on reported motion direction

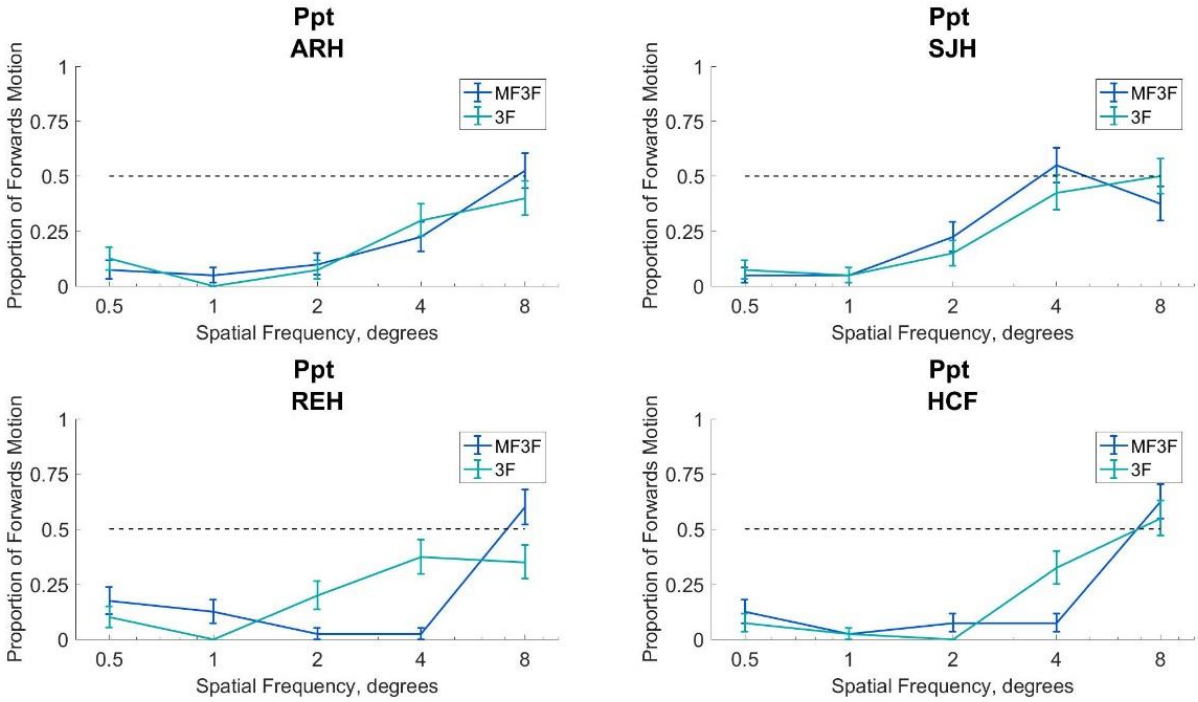


Figure 2.6. Figure displaying the data for each of the 4 participants, plotting the proportion of trials reported in the displacement direction, for the MF3F and 3F stimuli across 5 spatial frequencies. The dashed black line at 0.5 illustrates the chance performance level, where performance above chance reflects motion reported in the direction of the displacement of the fundamental frequency, which is not included in any stimulus but serves as a reference for the stimulus construction. Spatial frequency is plotted in terms of the fundamental frequency, so the actual frequency of the third harmonic is three times labelled values. Error bars represent the standard error of the mean across the 40 trials displayed per stimulus.

The data from Experiment 1 (see Figure 2.5 Left for data averaged across the four participants, and Figure 2.6 for each of the four participants individually), are best described by a quadratic fitting proportion correct by spatial frequency (SF), in logarithm of base ten units;

$$\text{proportion correct} \sim a*\log10(SF)^2 + c$$

I included both stimulus condition and participant identity as fixed effects in the model, across which the two fixed parameters could vary. As can be seen in Table 2.1 both parameters contribute significantly to the model, indicating its appropriate fit to the data. There was negligible variation across the fixed parameters at either the stimulus condition or the participant level, and the coefficients for both parameters when examined for the MF3F data and the MFG data were identical for the first 5 decimal places.

Table 2.1 Table reporting the parameter values for the model fitted to data in Experiment 1.

Parameter	Coefficient	Std. dev	Coefficient (MF3F data)	Coefficient (3F data)	Variation across condition level	Variation across ppt. level	df	t value	*p* value
a	0.551	0.052	0.551	0.551	<0.001	<0.001	31	10.516	<0.001
c	0.055	0.021	0.055	0.055	<0.001	<0.001	31	2.611	0.014

To investigate whether there was a significant difference between the MF3F and 3F stimulus conditions I compared my model, where both condition and participant identity were included as fixed effects, to a reduced model, where only participant identity was included as a fixed effect. There was no significant difference in goodness of fit between the two models (Log Likelihood Ratio = .000, p = 1.000), indicating no significant difference in performance between the two stimulus conditions.

Effect of varying IFI on perceived direction of motion

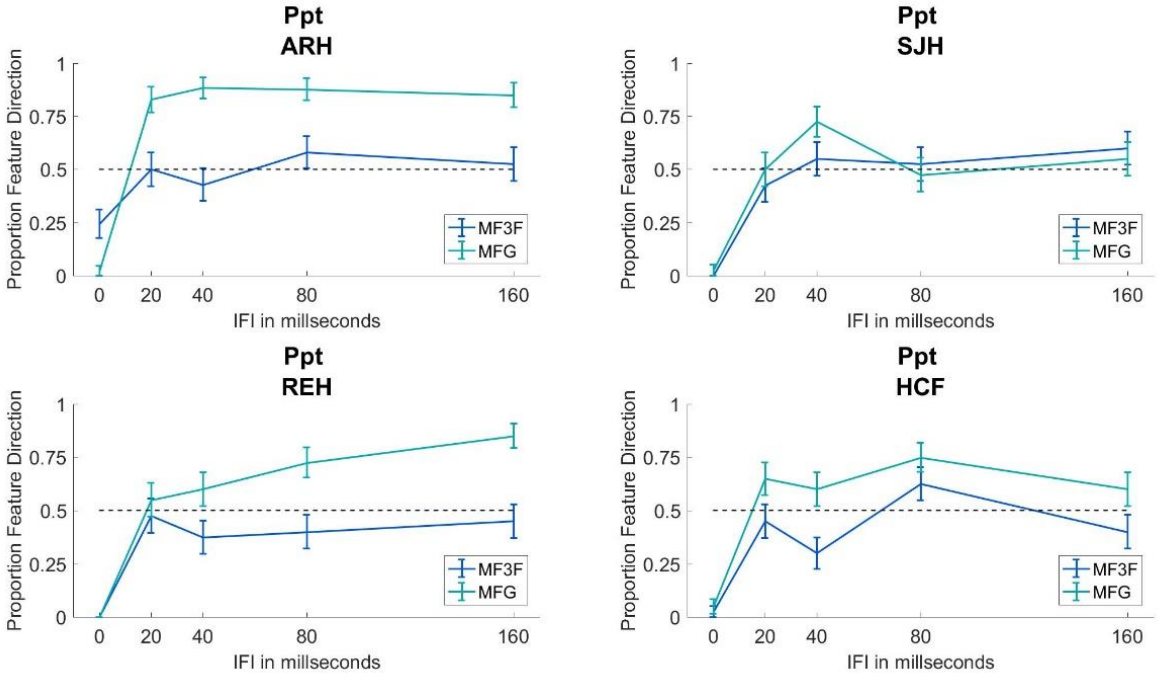


Figure 2.7. Figure displaying the data for each of the 4 participants, plotting the proportion of trials reported in the displacement direction, for the MF3F and 3F stimuli across 5 inter-frame interval lengths. The dashed black line at 0.5 illustrates the chance performance level, where performance above chance reflects motion reported in the direction of the displacement of the fundamental frequency, which is not included in any stimulus but serves as a reference for the stimulus construction. Error bars represent the standard error of the mean across the 40 trials displayed per stimulus.

The data from Experiment 2 (see Figure 2.5 Right for data averaged across participants, and Figure 2.7 for individual participant data), were best fitted by a 3 parameter asymptotic regression model, where the a parameter controls the scale of the curve, the b parameter the steepness of the curve, and the c parameter determines the asymptote on the y axis, corresponding to the maximum proportion correct; proportion correct ~ a*(exp(-b * IFI)) + c),

I included both stimulus condition and participant identity as fixed effects in my model, over which the three fixed parameters could vary. As can be seen in Table 2.2 all three parameters contributed significantly to the model, indicating its appropriate fit to the data. There was variation across both the a and c (asymptote) parameters at the condition and participant level, and negligible variation across the b parameter.

Table 2.2. Table reporting the parameter values for the model fitted to data in Experiment 2.

Parameter	Coefficient	Std. dev.	Coefficient (MF3F data)	Coefficient (MFG data)	Variation across condition level	Variation across ppt. level	df	t value	*p* value
a	-0.552	0.094	-0.446	-0.657	0.116	<0.001	30	5.889	< 0.001
b	0.114	0.035	0.114	0.114	<0.001	<0.001	30	3.294	0.003
c	0.598	0.078	0.508	0.687	0.097	0.069	30	7.626	< 0.001

The coefficients for the a and c parameter differ between the MFG and MF3F data. To investigate whether this demonstrated a significant difference between the MF3F and MFG stimulus conditions, I compared the full model, where both stimulus condition and participant identity were included as fixed effects, to a reduced model, where only participant identity was included as a fixed effect. There was a significant decrease in goodness of fit when condition was not included in the model (Log Likelihood Ratio = 16.581, $p < .001$), demonstrating a significant difference in performance between the two conditions.

To determine whether performance in the MF3F and MFG conditions reached asymptote at chance, I extracted the asymptote from the model for all participants for the MF3F and MFG data and ran a one sample t test between the intercept values and 0.5. There was no significant difference in the MF3F condition ($p = .896$), demonstrating that performance asymptotes at chance for the MF3F

condition at long IFIs. However, for the MFG condition there was a significant difference between participants' performance asymptote and chance, ($t(3) = 5.387$, $p = .013$, $M = 0.698$), with a confidence interval that was significantly higher than chance (95% CIS [0.581 0.831]), indicating that in the MFG condition, motion at high IFIs was reported in the position cue direction.

Discussion

In Experiment 1 I investigated how reported motion direction in the MF3F and third harmonic gratings varied by spatial frequency. Participants could successfully discriminate motion direction in the MF3F stimulus, across a range of spatial frequencies. There was no significant difference between proportion correct between the MF3F and the 3F, where performance was at chance by 8 cpd, corresponding to 24 cpd of the third harmonic. This suggests that when presented at maximum contrast, there was no disruption to performance caused by the switching between the MFG and third harmonic frames in the MF3F stimulus.

In Experiment 2, I compared motion discrimination across a range of durations of blank interframe intervals (IFI) for a novel grating based stimulus containing solely phase cues, the MF3F, with the Missing Fundamental Grating (MFG), which contained the same phase cues, and an additional position cue signalling motion in the opposing direction. I demonstrated that sensitivity of the phase system declined rapidly with increasing IFI length, shown by a loss of motion discrimination in the pure phase cue grating, and in the reversal of reported motion direction in the MFG, from the phase cue direction when there was no IFI included, to the position cue direction when IFIs were included. When viewing the MFG motion is reported in the phase cue direction when no IFI is included between frames of the stimulus, and in the position cue direction when IFIs are included in the stimulus. Once the IFI reaches a critical duration, motion is reported in the position cue direction, but the increasing length of the IFI has no further effect on motion discrimination performance. This is consistent with the results of Georgeson & Harris (1990) who also reported motion reversals in the MFG with the introduction of blank inter-frame intervals. The difference between the phase and position cues in motion discrimination across IFI duration suggests that different mechanisms are responsible for processing phase and position cues.

I suggest that the results indicate a difference in how motion cues are integrated across time, with the results in the phase cue condition consistent with motion from phase cues being extracted by detectors oriented in space-time. One suggested implementation is through detectors formed from a quadrature pair of spatiotemporal filters (Reichardt, 1957; Adelson & Bergen, 1985; Van Santen & Sperling, 1985), responding to oriented luminance in space-time. The insertion of short blank IFIs

would reduce the activity of the motion detector since it would be integrated with the signed motion signals from the 'on' frames of the stimulus. Longer IFIs would abolish the motion response completely, because the separation in time between the 'on' frames of the stimulus exceeds the temporal integration window of the detectors.

The high motion discrimination across long durations of blank IFIs in MFG grating, once it is reported in the position cue direction, suggests that the position system is not integrating the features of the MFG with the blank IFIS. These results could be explained by a system, where the information about features and feature position is updated when feature position changes, rather than the continuous integration of luminance information in time at each location, which I suggest is occurring in the phase cue condition.

This experimental procedure could be used to estimate the temporal integration properties of the phase and position systems, where motion discrimination or detection could be measured across a range of IFI lengths presented at a range of contrasts and spatial frequencies. Replicating the current experiment at a higher contrast will likely extend the temporal window for the phase system, as the higher contrast of the phase cue will require a longer duration of blank IFI for the integrated contrast to produce a null motion output. The MF1F of Baro & Levinson (1988) could be included to map the temporal dynamics of the position system, rather than relying on the reported motion direction in the MFG. With a wider window of temporal integration, and the inclusion of the MF1F stimulus, the dynamics of how temporal integration varies between the phase and position systems can be assessed. This preliminary evidence for the difference in temporal integration suggests that different systems are processing the phase and position cues, using different mechanisms.

A critical flaw in this stimulus is that it does not completely remove the available motion cues from the change in spatial position of features between frames, as it was initially designed to do so. Whilst the change in identity between the MFG and the third harmonic abolishes the mapping of features with the same overall periodicity, there are still features present in the stimulus that could be tracked across frames. This can be seen in Figure 2.4, where the areas of minimum luminance on the MFG can be tracked to troughs in the 3F, and likewise areas of maximum luminance on the MFG to peaks in the 3F frames, or mapped across the zero-crossings between frames. Critically, the direction of motion signalled by this residual position cue is in the same direction as the phase cue. This means that it is impossible to validate which cue is being used to discriminate motion just from reported direction alone. One method of verifying whether motion discrimination in the MF3F is produced from the phase cue of the third harmonic, or by an undefined position cue, would be to compare contrast sensitivity functions for the MF3F and a third harmonic grating. If motion

discrimination relies on the third harmonic, then similar contrast sensitivity functions should be produced whether the third harmonic is presented in isolation, or whether interleaved with the frames of the MFG.

One possibility is that the position cue is below the threshold for discrimination for participants, so although the cue is present in the stimulus, it is not used by participants in discriminating direction. The change in reported direction in the MFG suggests that participants are able to use position cues to maintain motion discrimination even at long durations of ISI. Therefore the lack of motion discrimination in the MF3F at longer durations of ISI could be used to argue that the position cue is not being used to discriminate motion. However, this would not be useful evidence in determining whether motion discrimination at 0 ISI is relying on the phase or the position cue. The periodicity of features in the MFG is $^1/_3$ of the periodicity of the remaining position cues in the MF3F, as the features in the MFG are constrained by the fundamental frequency whilst the features in the MF3F are constrained by the periodicity of the third harmonic. One method to directly test whether the position cue in the MF3F is available to participants would be to test motion discrimination in the MF1F, matching the spatial frequency and amplitude of the third harmonic in the MF3F to the missing fundamental frequency of the MFG. If motion discrimination is still below chance, then this would be evidence to suggest that the residual position cues in the MF3F are not being used to discriminate motion, and instead that motion is determined by the phase cue, as was the original intent in designing the stimulus.

A further potential limitation of these stimuli is that the MFG changes appearance at low contrasts. In static frames of the MFG when presented at detection threshold, MFG presented below 1 cpd appears to have a square wave profile, but MFG presented at higher spatial frequencies appear as the third harmonic (Campbell *et al.*, 1978). For apparent motion renderings of the MFG, Georgeson & Harris (Georgeson & Harris, 1990) reported that, for MFG at 1 cpd at low contrasts (4%) and high temporal frequencies, the stimulus has a 'third harmonic like' appearance, but with increasing contrast and decreasing temporal frequency, the 'square wave' appearance is restored. This produces a potential confound in the MF3F, where, if the MFG is perceived as a third harmonic, and not in its lower periodicity scalloped square wave profile, this introduces matching features between the MFG and third harmonic frames, introducing position cue into my MF3F pure phase cue stimulus. The reported direction of motion would not change, hence this confound would only be detected by asking participants to report the subjective appearance of the MF3F, or by testing the MF3F at longer inter-frame intervals (IFIs), where performance would not be abolished if a position confound was occurring. At 12% contrast and above, all MFG were reported as appearing 'square wave', regardless of whether they were reported in the phase or position cue direction (Georgeson

& Harris, 1990) . This indicates that this potential confound can be avoided by presenting MF1F and MF3F at moderate contrast levels, and that this appearance confound may be dissociable from the cues used to determine motion direction. For these reasons, and the presence of the confounding position cue in the MF3F, whilst the MF1F and MF3F are on the surface a compelling pair of gratings for testing psychophysical dissociations of phase and position cues, I use alternative stimuli throughout the rest of the thesis.

Replicating the experiment with a different stimulus set, that doesn't contain this artefact or the contrast issues inherent in using the MFG, would be a vital step in further investigating temporal integration differences between motion cues. Van Santen & Sperling (1984) demonstrated that the addition of a pedestal of matched frequency to a drifting sinusoid grating could be used to obscure the change in position of grating features, without disrupting the change in phase of the Fourier components. A pedestal can't be added to the MF3F, as it relies on matched pairs of interleaved gratings: displaced by 90^{o} on each frame, to produce pairs of gratings in counterphase, where stimulus features cannot be tracked between frames, or across the pairs of frames: the pedestal relies on the grating being a drifting grating, with displacements between frames of less than 90^{o}. However, the same experimental design could be used, but using a pedestalled Gabor with a drifting carrier. The addition of a pedestal would produce a Gabor stimulus that contains motion cues solely in the change in phase of the Fourier components, and without the artefact of the changing spatial position of grating features.

This pedestalled multi-frame Gabor with a drifting carrier could be used to investigate the effects of temporal integration, and whether these vary when the motion in a stimulus is produced from changes in Fourier components, and changes in spatial position. Motion discrimination across increasing inter-frame intervals could be compared between a pedestalled and a non-pedestalled drifting Gabor, where both contain the same change in phase of Fourier components, but the non-pedestalled Gabor also contains the congruent change in spatial position of the grating features. If motion discrimination in the pedestalled Gabor declines rapidly with increasing IFIs, but does not in the non-pedestalled Gabor, then this would suggest that motion discrimination from changes in phase of the Fourier components and from changes in spatial position have different temporal integration functions. This would be suggestive that these cues can be processed by different mechanisms, which could suggest two separate systems for motion processing.

2.3 Spatial Frequency Dissociation of Phase and Position Systems

A visual stimulus popular in psychophysics is the Gabor. These are used in this thesis in an attempt to produce stimuli where the descriptions of motion in terms of changes in phase of the Fourier components, and changes in the spatial position, can be somewhat dissociated.

A Gabor is the product of a sinusoid wave multiplied by a 1D Gaussian function (see Figure 2.8 top row), and a Gabor patch is the product of a sine wave grating multiplied by a 2D Gaussian function (see Figure 2.8 bottom row). The cross-section of a Gabor patch perpendicular to the orientation of the carrier will resemble a 1D Gabor (see Figure 2.8 top row right), and a cross-section taken parallel to the orientation will have the profile of the Gaussian window (Figure 2.8 top row central). The weighting of the sine wave grating with the 2D Gaussian window produces a symmetric band of frequencies, centred on the spatial frequency of the sinusoid. This can be understood in terms of the Fourier transforms of both components: the Fourier transform of a Gaussian function is another Gaussian centred at 0, whereas the Fourier transform of a sine wave is a single frequency. The width of the Fourier transform of the Gaussian function on the frequency axis is inversely proportional to the width of the original Gaussian function, therefore Gaussian windows with larger standard deviations will produce Gabors with a narrower band of spatial frequencies. Likewise the single orientation of the sine wave grating, when weighted by the Gaussian window, produces a Gabor with a band of orientations centred on the original orientation of the sine wave grating. The Gabor patch can be thought of as 'knowing what is where', with the frequency of the sine wave grating providing the 'what', and the spatial coordinates of the Gaussian window provide the 'where'. This quality has made them very popular in communication analysis, image analysis, and also in visual neuroscience as it is thought that the properties of the receptive fields of simple cells in V1 have a Gabor like response (Daugman, 1985). For vision research the Gabor is a popular stimulus as the spatial location and frequency of the stimulus can be well defined within a specific band of frequencies and orientations, given the original frequency and orientation of the sine wave grating, and the standard deviation of the Gaussian window. The Gaussian window produces a smooth transition from peak contrast, to zero contrast without introducing errant high spatial frequencies in the construction of the aperture, unlike circular or rectangular apertures.

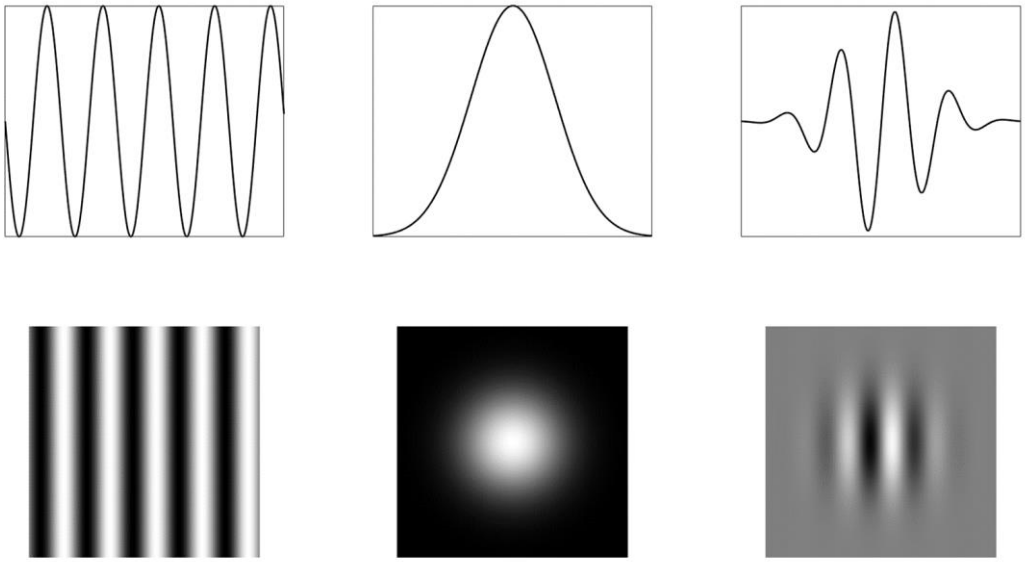


Figure 2.8. Schematic displaying a 1D Gabor (top row), and 2D Gabor patch (bottom row). Top row left: the profile of a sine wave, top row central: the profile of a Gaussian function, top row right: the profile of the 1D Gabor function that results from weighting the sine wave by the 1D Gaussian function. Bottom row left: image of a 2D sine wave grating, bottom row central: image of a 2D Gaussian window, bottom row right: image of the 2D Gabor, produced from weighting the sine wave grating by the 2D Gaussian window.

The separable frequency and position components of the Gabor patch make it a useful stimulus for investigating phase and position cues, where phase cues are produced from changing the phase of the sinusoid grating, and position cues by displacing the position of the Gaussian window. When we consider the effect of applying the Gaussian window in terms of the effect on the Fourier components of the carrier – modifying the bandwidth of frequencies present in the Fourier transform of the carrier - this explains why changing the position of the Gaussian aperture can be considered a position cue, without producing a corresponding phase cue. However, changing the phase of the carrier will produce simultaneous position cues, as the peaks, troughs and zero-crossings of the Gabor will change position. This introduces a limitation into the use of Gabors in investigating dissociations of phase and position cues: it is impossible to produce a Gabor that solely contains a phase cue, as the change in phase of the carrier by definition produces changes in position of the features within the Gabor. To make this distinction explicit, I will refer to the changes in phase of the carrier as 'carrier' cues, and changes in the displacement of the aperture as 'aperture' cues.

Conflicting cue Gabors are produced when the direction of the phase displacement of the carrier and Gaussian window displacement are in different directions. Due to the aperture problem the motion of the phase cue will only be perceived in directions orthogonal to the orientation of the sine wave grating. The position cue is not limited to being perceived in perpendicular directions only, as it is

determined by the spatial location of the Gaussian window, where displacements in any direction are equally meaningful.

Perceived motion direction in conflicting cue Gabors could be reported in the direction of the carrier displacement, the direction of the Gaussian window displacement, or in some cases, in a trajectory intermediate to these. This has been shown in Gabor stimuli where the displacement of the Gaussian envelope is orthogonal to the phase displacement of the sine wave. When Gabor stimuli are displaced, the perceived motion direction of these orthogonal Gabor stimuli is highly dependent on the retinal location of the stimulus. When viewed foveally both single Gabors and displays containing multiple Gabors with orthogonally oriented cues are perceived in the position cue direction, and when viewed in the periphery, single and group Gabor displays are perceived as having a diagonal trajectory intermediate to the direction of the phase and position cues (Tse & Hsieh, 2006; Shapiro *et al.*, 2011). Perceived motion in conflicting cue Gabors is also affected by temporal parameters, where in dense displays of multiple Gabors with opposing cue directions presented close to fovea, Boulton & Baker reported that at short stimulus-onset asynchronies (SOA), motion is reported in the phase cue direction, and at long SOAs in the position cue direction (Boulton & Baker, 1993).

The purpose of this experiment was to investigate whether coherent motion can be perceived in the carrier and aperture cue direction in two-frame orthogonal and opposing arrangements of conflicting cue Gabor stimuli, and whether there consistent biases towards the aperture or carrier cue direction when the spatial frequency of the carrier and size of the Gabor are manipulated. I also included Gabor stimuli containing solely a displacement in Gaussian aperture, or phase displacement of the carrier, to examine the motion discrimination of the carrier and aperture cues when presented in isolation. This produced a stimulus set that can be relatively confined to a spatial location, due to the use of just two frames, and so able to test the sensitivity to changes in carrier phase and aperture displacement both with spatial frequency, and at different locations on the retina. I used a single Gabor stimulus rather than an array of Gabors, as the manipulation of spatial frequency, with the fixed relationship to the standard deviation of the Gaussian aperture, meant that low spatial frequency presentations of the Gabors would rapidly increase in size: using an array of Gabors would drastically reduce the localisation of the Gabors to the eccentric testing location. Using multiple-frame Gabor stimuli may also have biased reported motion towards the aperture displacement direction, and so using a two-frame Gabor may be a better choice to investigate the effects of spatial frequency on reported motion as the relative strength of the aperture cue and carrier cue will be more balanced. Using a flat-topped Gaussian window with sufficiently large

standard deviation to display multiple cycles of the carrier should also increase the relative strength of the carrier cue.

Method

Ethics

This study was approved by the Medical Sciences Inter-divisional Research Ethics Committee (IDREC) at the University of Oxford in accordance with the procedures laid down by the University for Ethical Approval of all research involving human participants, under reference R73572/RE001.

Participants

Six participants, (1 male) aged 20-29. All participants had normal or corrected to normal vision, and all participants were naïve to the hypothesis.

Apparatus

The experiment was conducted in a dimly lit room. Stimuli were presented on a CRT monitor (Mitsubishi Diamond Pro 2070, 1600 × 1200 pixels), with maximum luminance of 81.32 cd/m^2 and gamma corrected using the ColorCAL MK II colorimeter (Cambridge Research Systems). The stimuli were rendered and presented using the MATLAB toolbox 'PsychToolbox' toolbox, in MATLAB R2014a. Participants were seated 125 cm from the CRT screen, and this viewing distance was maintained with a chin rest, adjusted for each participant's height. Viewing was monocular, with each participant's right eye covered by an opaque eyepatch.

Stimuli

Stimuli consisted of two-frame apparent motion windowed sinusoidal gratings, with each trial consisting of two frames. In this experiment, stimuli were not weighted by a Gaussian aperture, but with an aperture similar in profile to a Gaussian window, with the central region of the aperture passing gratings at full contrast. This aperture, a 'flat-topped Gaussian' consists of a central circular aperture with a radius of 1 cycle, which passed 100% of the sinusoid grating, and an annulus where contrast falls off following a Gaussian profile, that spanned 2.5 cycles of the grating. This produces stimuli with effectively 3 visible cycles, and a gradual taper to the background luminance, as shown in Figure 2.9. The aperture's width was also a fixed proportion of the spatial frequency of the grating, so that for all spatial frequencies in the experiment, an identical number of cycles were visible.

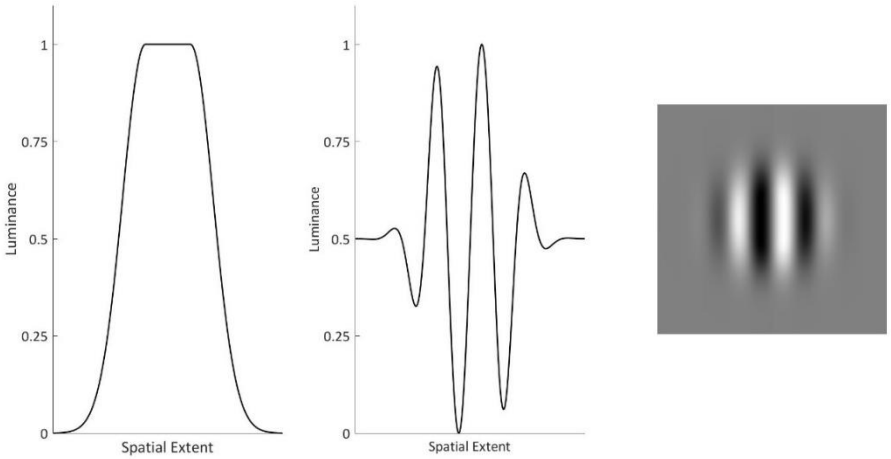


Figure 2.9. Schematic displaying the structure of the 'flat-topped Gaussian' used in Chapter 2.3, with the cross-section of the flat-topped aperture in the leftmost panel, the central panel displaying the cross-section of the resulting Gabor stimulus when a sinusoid grating is weighted by the 2D flat topped Gaussian function, and the rightmost panel displaying the resulting 'Flat- topped Gabor' as seen by the participants. The central region, the 'flat top', is equal to one cycle of the sinusoid grating.

I used a 'flat-topped Gaussian' as the aperture in this experiment to reduce the interaction between the phase of the sinusoid grating and the contrast of the Gabor. The contrast of the Gabor is the result of the contrast of the sinusoid grating and the contrast of the Gaussian envelope. With a Gabor produced from a narrow Gaussian aperture the phase of the sinusoid grating affects the resulting Gabor contrast, due to which sections of the sinusoid grating fall under the peak of the Gaussian aperture. If the sinusoid grating is presented at 180^{o}, the resulting contrast will be at the theoretical maximum for that combination of sinusoid and Gaussian aperture, as the peak of the sinusoid will align with the peak of the Gaussian window. However, if the sinusoid is presented at other phases the resultant Gabor will have a lower contrast, with the lowest contrasts possible when the Gabor is produced from a sinusoid with a phase of 90^{o} or 270^{o}. Equally, if the Gaussian envelope changes position then this will also change the contrast of the Gabor, as the alignment of the Gaussian envelope and the peaks and troughs of the sinusoid grating will change. This will produce a contrast cue, and hence a phase cue, from a change in position of the Gaussian envelope. Therefore for Gabor stimuli with narrow Gaussian apertures a 'pure' position cue will have a confounding phase cue. This confound can be reduced by having wider Gaussian apertures, where more cycles of the sinusoid will align with the peak of the Gaussian aperture, or be avoided more completely by having a central region of the aperture that extends over one cycle or more, at maximum contrast. This interaction between the size of Gabors and their phase has been shown psychophysically (Foley *et al.*, 2007), further demonstrating the need to either increase the size of the Gabors, which reduces their use in targeting specific regions of the retina, or change the profile of the window used

to weight the sinusoid grating. The flat-top Gaussian used as the aperture in this experiment avoids this confound by merit of its central maximum contrast region.

Gratings were displayed at one of 9 spatial frequencies, distributed logarithmically between 0.3 cpd and 12 cpd. The phase of the first frame was randomly assigned on each trial, and stimuli were presented at 100% contrast. The Gabors were presented at one of 4 orientations; 0°, 90°, 180° or 270°, and had the same orientation on both the first and second frames of each trial.

There are four conditions in the experiment, where changes in the position of the aperture and the phase of the grating are combined to produce a two-frame apparent motion stimulus. There are two single cue conditions: the 'Carrier' condition, where the aperture is stationary, but there is a quarter cycle phase shift on the second frame, and the 'Aperture' condition, where the phase of the grating remains the same on the two frames, but the aperture is displaced on the second frame, parallel to the orientation of the grating. The displacement is by one quarter of the wavelength of the fundamental frequency, equivalent in distance to a quarter-cycle phase shift. There are two double cue conditions, which contain both a phase change of the grating and a position change of the aperture. The 'Opposing Motion' condition has the aperture displacement and phase shift in opposing directions, e.g. a vertical grating with a phase shift to the left and aperture displacement to the right. The 'Orthogonal Motion' condition has the aperture displacement and phase shift in orthogonal directions, e.g. a vertical grating with a phase shift to the left and aperture displacement downwards. These are depicted in Figure 2.10.

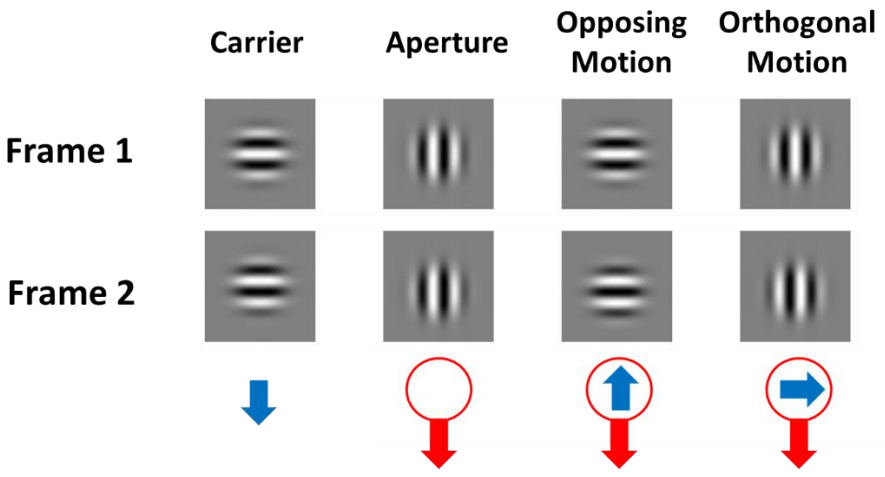


Figure 2.10 Schematic of the four conditions included in the experiment. The blue arrow refers to the direction of the phase displacement, and the red circle to the direction of the envelope displacement. From left to right the stimuli are the Carrier, Aperture, Opposing and Orthogonal Motion stimuli, with the top row representing the first frame and the bottom row the second, in a trial that would be classified as 'downward' motion, as in conflicting cue trials the position cue direction was used as the default correct direction.

Procedure

The experiment used a method of constant stimuli procedure, with stimuli presented in a randomised order. Each trial consisted of two frames of the stimulus, where the displacement in the phase or position cue between the first and second frame produced the apparent motion cue. Each frame was presented for 250 ms, and the two stimulus frames were followed by a blank screen matched to the background luminance of the stimulus frames. The blank screen was displayed while participants reported motion direction using the 4 arrow keys on the keyboard, and then rated their confidence in their response, using the 1 to 4 keys, where 1 corresponded to 'very low confidence' and 4 to 'very high confidence'. Once participants had reported perceived motion direction and their confidence rating, they initiated the next trial by pressing the spacebar.

There were 60 trials for each of the nine spatial frequencies and each of the four stimulus conditions, 2160 trials in total. The experiment was run over two sessions, with trials arranged into 14 blocks of 155 trials (with the final block containing only 145 trials). Participants were encouraged to have a five-minute break between each block of trials. Each session took two hours, with the total experiment length of up to four hours. Sessions took place within the same week, but not all within a 24-hour period.

Fixation was encouraged by means of a 0.25° by 0.25° white cross, presented on the left side of the screen, where the stimulus was presented at the centre of the screen, so projected onto the retina at ten degrees of eccentricity. Eye tracking to monitor fixation was not possible with the COVID-19 restrictions in place during this experiment. Participants could abort trials in which they felt their fixation was poor, and this trial would be randomly shuffled into the remaining trials. Each participant aborted fewer than 5 trials per experimental session.

Results

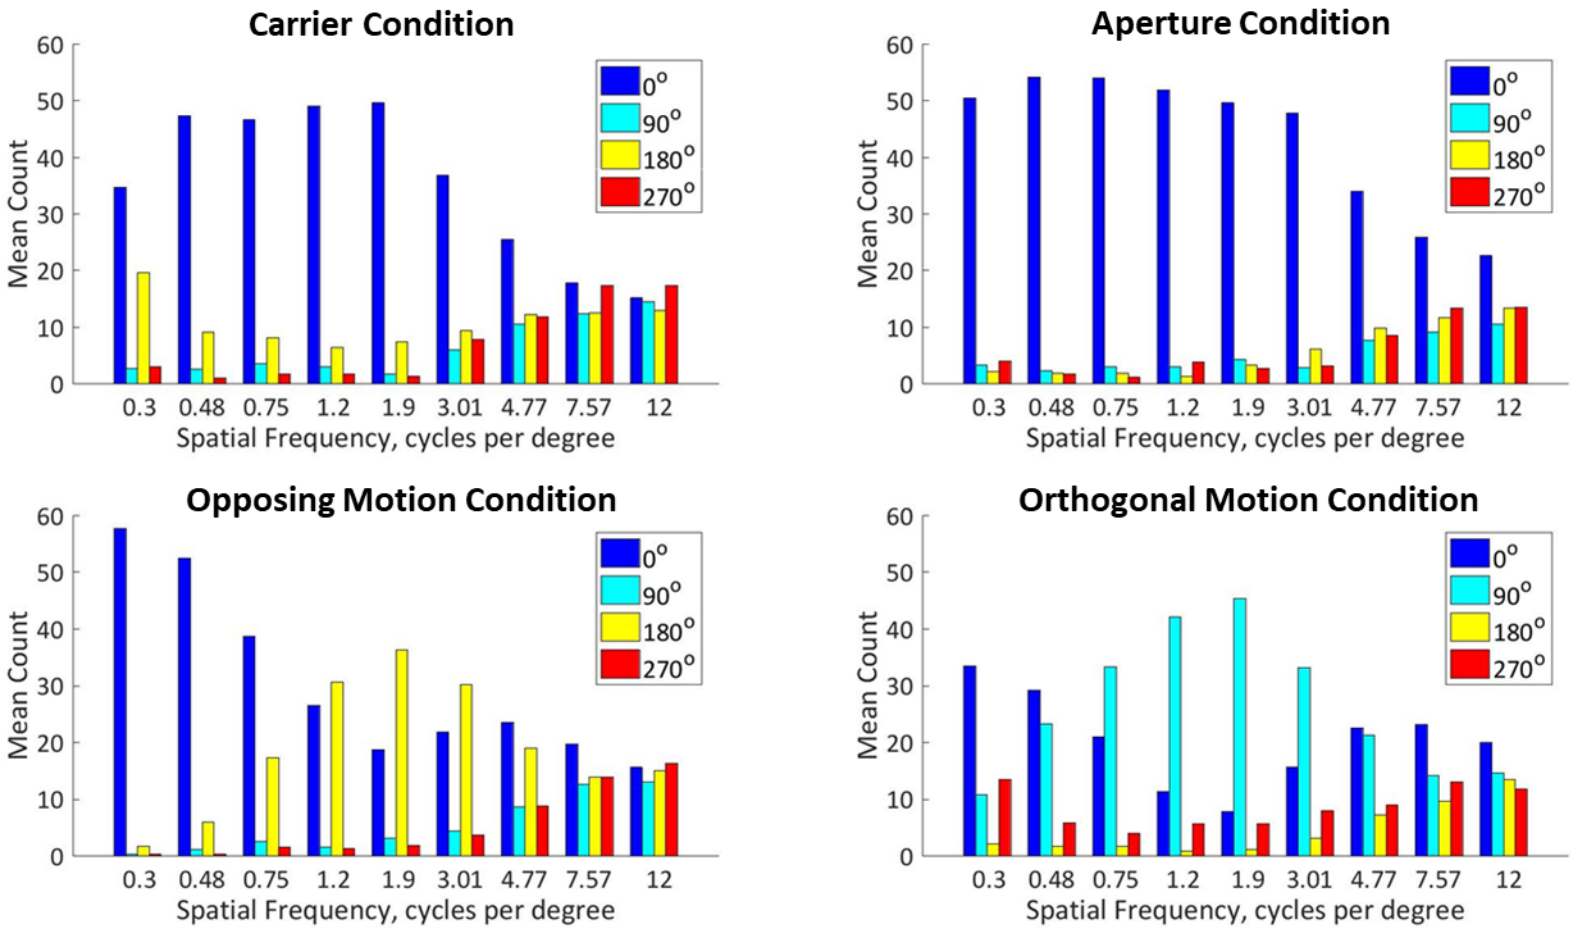


Figure 2.11 Bar graphs displaying the mean motion trajectory across participants, for each of the 9 spatial frequencies and for each of the four Gabor conditions. The correct motion trajectory is 0° for the Carrier and Aperture conditions (top row), as these contain a single motion cue. For the Opposing Motion condition (bottom left panel) the position cue trajectory is 0° and the phase cue trajectory is at 180°, and for the Orthogonal Motion condition (bottom right panel) the position cue trajectory is at 0°, and the phase cue trajectory is at 90°.

Data were collected from six participants, in four stimulus conditions, displayed at 9 spatial frequencies, with 60 trials per spatial frequency for each stimulus condition. The Carrier and Aperture motion conditions were stimuli that contained a single cue direction; either the phase of the carrier is shifted by 90^{o} (carrier cue) or the aperture is displaced by one-quarter of the wavelength of the fundamental frequency of the carrier (aperture cue). The Opposing and Orthogonal Motion conditions contained both a phase shift of the carrier and a displacement of the aperture position, in directions that were 90° apart in the Orthogonal Motion condition, and 180° apart in the Opposing Motion condition. Data were summarised as the number of trials reported in each direction, where at each SF participants reported motion in one of the four possible cardinal directions, with 60 trials per SF (see Figure 2.11 for the data aggregated across the group).

The four stimulus conditions can be grouped into 'single' (Carrier and Aperture conditions) and 'conflicting' (Opposing and Orthogonal Motion conditions) cue conditions. In the single cue conditions there is only one 'correct' reported motion direction, whereas in the conflicting cue conditions there are two valid directions that motion could be reported in. For this reason, some analyses in this section will be run separately for the single and conflicting cue conditions.

The reported motion direction was re-coded as motion trajectory using the difference in the reported motion direction (0^{o}, 90^{o}, 180^{o} and 270^{o}) relative to the cued direction. In the single cue conditions, this means that the correct reported motion direction is at 0^{o}. In the conflicting cue conditions I re-coded reported motion direction against the aperture cue direction, meaning that for the Opposing Motion condition the carrier cue trajectory is at 180^{o}, and for the Orthogonal Motion condition, the carrier cue trajectory is at 90^{o}.

Figure 2.11 displays the histograms of the mean number of trials reported in each direction by the 6 participants, at each of the 9 spatial frequencies that stimuli were displayed at, for each of the four stimulus conditions. For the Carrier and Aperture conditions (top row) the majority of trials are reported in the correct direction, at 0^{o}, especially at the low and mid-range spatial frequencies. For the Carrier condition, performance nears chance at the highest spatial frequencies, whilst performance remains higher than chance in the Aperture condition at the highest spatial frequencies. In the Opposing Motion condition (bottom left panel) trials are reported predominantly in the aperture cue direction (0^{o}) at low spatial frequencies and in the carrier cue direction (180^{o}) at mid-range spatial frequencies. Similarly, for the Orthogonal Motion condition (bottom right panel), trials are reported predominantly in the aperture cue direction (0^{o}) at low spatial frequencies, and in

the carrier cue direction (90°) at mid-range spatial frequencies. These histograms are presented for each participant in the appendix of Chapter 2, in Chapter 2.6.1.1.

Exploratory analysis of reported motion

I ran an initial exploratory analysis of the reported trajectories as group data, using a chi square test of independence, to compare the distribution of trajectories between the Carrier and Aperture conditions, and between the Opposing and Orthogonal Motion conditions. In these analyses participant identity was not treated as a fixed effect, to simplify this initial analysis.

Comparison of trajectories in the Carrier and Aperture conditions

The distribution of trajectories in the Carrier and Aperture conditions were compared using a 2 (condition) x 9 (spatial frequency) chi-square test of independence. . There was a significant overall difference between the Carrier and Aperture conditions in the distribution of trajectories across the four possible directions, $x^2(3) = 137.391$ $p < .001$, Cramer's V =.146 and this was significant at all 9 spatial frequencies, $p \leq .002$. This is displayed in Table 2.3.

I ran a second chi squared test of independence on the trajectory data aggregated across spatial frequency to compare the distribution of trajectories in the 4 directions, between the Carrier and Aperture conditions. There was a significant difference between the number of trajectories in the 0° direction ($x^2(1) = 114.00$, $p < .001$), with significantly more trials reported in the 0° direction in the Aperture condition. There was also a significant difference in the 180° relative direction, $x^2(1) = 99.50$, $p < .001$, with significantly more trials reported in the 180° direction in the Carrier condition.

I explored both of these results further, at each spatial frequency, to compare the distribution of the four trajectories at each of the 9 spatial frequencies, for the Carrier and Aperture conditions, by running a 4 (direction) x 9 (spatial frequency) chi squared test of independence using Bonferroni corrected significance level (0.0014) of the adjusted standardised residuals. These can be seen in Table 2.12 in the appendix (Chapter 2.6.1.2).There were significantly more trials reported in the 0° direction in the Aperture condition at 7 of the 9 spatial frequencies ($p \leq .003$), but no significant difference at 1.2 cpd and 1.9 cpd ($p \geq 0.0833$). The greater number of trials reported in the 180° direction in the Carrier condition was significant for the lowest four spatial frequencies, ($p < .001$).

Table 2.3 Table displaying the results of an exploratory chi square test of independence to investigate whether there was a significant difference in the distribution of trajectories between the Carrier and Aperture conditions, with the analysis ran on data aggregated across the six participants.

Spatial Frequency, cycles per degree	Pearson Chi-Square	df	*p* value	Cramer's V
0.30	103.123	3	<0.001	0.378
0.48	33.128	3	<0.001	0.215
0.75	28.032	3	<0.001	0.197
1.20	25.164	3	<0.001	0.187
1.90	18.778	3	<0.001	0.161
3.01	31.147	3	<0.001	0.208
4.77	14.701	3	0.002	0.143
7.57	14.895	3	0.002	0.144
12.00	15.645	3	0.001	0.147
Total	137.391	3	<0.001	0.146

Comparison of re-coded trajectories in the Opposing and Orthogonal Motion conditions

I compared reported direction between the Opposing and Orthogonal Motion conditions, using re-coded responses, so that the carrier and aperture cue directions were compared. Instead of comparing motion trajectory in the form 0°, 90°, 180° or 270°, I coded directions as in the 'Aperture cue', 'Carrier cue', 'Error 1', 'Error 2' directions, with correspondences displayed in Table 2.4.

Table 2.4. Table outlining the re-coded trajectories for the Opposing and Orthogonal Motion conditions, taking the 0°, 90°, 180° and 270° directions into 'Aperture, 'Carrier', 'Error 1' and 'Error 2' directions.

	Aperture	Carrier	Error 1	Error 2
Opposing	0	180	270	90
Orthogonal	0	90	270	180

I ran a 2 (Opposing vs Orthogonal Motion) x 9 (SF) chi squared test, on data aggregated across the six participants. There was a significant difference in the distribution of reported directions between the Opposing and Orthogonal Motion conditions, $x^2(3) = 220.235$, $p < .001$, Cramer's V = 0.184, and this was significant for the lowest 6 of the 9 spatial frequencies, $p \leq .003$. These are reported in Table 2.5.

Table 2.5. Table displaying the results of an exploratory chi square test of independence to investigate whether there was a significant difference in the distribution of trajectories between the Opposing and Orthogonal conditions, with the analysis ran on data aggregated across the six participants.

Spatial Frequency, cpd	Pearson Chi-Square	df	*p* value	Cramer's V
0.30	162.030	3	<0.001	0.474
0.48	125.455	3	<0.001	0.417
0.75	63.831	3	<0.001	0.298
1.20	62.602	3	<0.001	0.295
1.90	37.658	3	<0.001	0.229
3.01	13.697	3	0.003	0.138
4.77	2.020	3	0.568	0.053
7.57	6.198	3	0.102	0.093
12.00	5.125	3	0.163	0.084
Total	220.235	3	<0.001	0.184

I ran a 2 (Opposing vs Orthogonal Motion) x 4 (directions) chi squared test, on data aggregated across the six participants. There were significantly more trials reported in the aperture cue direction in the Opposing than the Orthogonal Motion condition, $x^2(1) = 184.96$, $p < .001$, and also significantly more trials reported in the carrier cue direction in the Orthogonal than the Opposing Motion condition, $x^2(1) = 110.25$, $p < .001$. Finally, there were significantly more trials reported in

the Error 1 direction in the Orthogonal than the Opposing Motion condition, $x^2(1) = 46.24$, $p < .001$. There was no significant difference in the proportion of trials reported in the Error 2 direction, $p = 0.057$.

I explored these results further, at each spatial frequency, to compare the distribution of the four directions at each of the 9 spatial frequencies, for the Opposing and Orthogonal Motion conditions, by running a 2 (condition) x 4 (direction) x 9 (spatial frequency) chi squared test of independence using Bonferroni corrected significance level (0.0014) of the adjusted standardised residuals. These can be seen in Table 2.13 in the appendix (Chapter 2.6.1.2).There were significantly more trials reported in the aperture cue direction in the Opposing Motion condition than the Orthogonal motion condition for the lowest 5 spatial frequencies (p <.001). There were significantly more trials reported in the carrier cue direction for the Orthogonal Motion condition than the Opposing Motion condition for the lowest 5 spatial frequencies, (p < .001*).* There were significantly more trials reported in the Error 1 direction in the Orthogonal Motion condition than the Opposing Motion condition for 0.3 cpd, 0.48 cpd and 1.2 cpd (p < .001).

Analysis of reported motion

The chi squared analyses suggest that there may be a significant difference in the accuracy of reported motion directions between the Carrier and Aperture conditions, with consistently more trials reported in the correct 0^o direction in the Aperture condition than the Carrier condition. The chi squared analysis on the re-coded directions reported in the conflicting cue conditions suggest that more trials are reported in the aperture cue direction in the Opposing Motion condition, and in the carrier cue direction in the Orthogonal Motion condition.

To investigate whether these differences were significant across participants, and to better describe these effects by spatial frequency, I transformed the data into binary arrays. For the Carrier and Aperture conditions I grouped data into 'correct' and 'incorrect' trials, where trials with a 0^o were correct. For the Opposing and Orthogonal Motion conditions I performed two analyses; one looking at proportion correct, where correct refers to trials reported in the carrier or aperture cue direction, and incorrect in the other two non-cue directions, and a second analysis comparing the number of trials reported in the carrier and aperture cue directions. In Figure 2.12 I have plotted these binary arrays for the Carrier and Aperture conditions (left panel), where proportion correct refers to the mean proportion of trials reported in the 0^o trajectory, and also plotted the binary arrays for the Opposing and Orthogonal motion condition (right panel), where proportion correct refers to the mean number of trials reported in either the carrier or aperture cue direction. The chance level

differs between the two panels, since the single cue conditions have one correct direction out of four (25% chance), and the conflicting cue directions have two correct directions out of the four (50% chance). This limits the comparisons that can be carried out between the single and conflicting cue conditions.

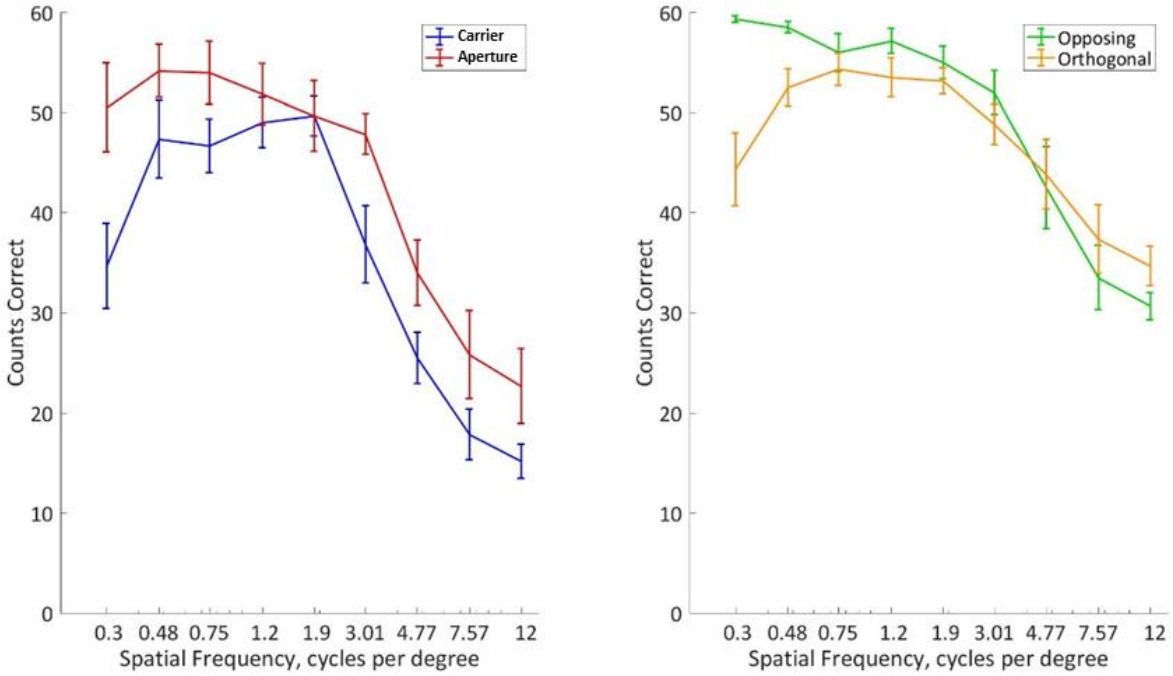


Figure 2.12 Figure displaying the mean number of trials (counts) for the Carrier and Aperture conditions which contain a single cue and so have one correct direction out of four (left panel) and the Opposing and Orthogonal Motion conflicting cue conditions, where counts correct refers to trials reported in the carrier or aperture cue direction (right panel). This produces two valid motion cue directions out of a possible four directions for the Opposing and Orthogonal Motion conditions, hence the asymptote at 30 trials out of a possible 60, whilst the Carrier and Aperture conditions will theoretically asymptote at 15 trials.

Analysis of single cue conditions

Comparing proportion correct in the Carrier and Aperture conditions

The proportion correct data for the Carrier and Aperture conditions (see Figure 2.12) was produced by taking the number of trials each participant reported in the 0^o direction as a proportion of the total number of trials. These data were fit best by a third order polynomial;

$$\text{proportion correct} \sim a^* \log10\,(SF)^3 + b^* \log10\,(SF)^2 + c^* \log10(SF) + d$$

The parameters and coefficients are reported in Table 2.6 with variation around the a and b parameters at the condition level, and around the b, c and d parameters at the participant level.

Table 2.6. Table reporting the parameters and coefficients for the model fitting the proportion correct data for the Carrier and Aperture conditions.

Parameter	Coefficient	Std. error	Coefficients for the Carrier condition	Coefficients for the Aperture condition	Std. dev. Across condition level	Std. dev. Across ppt. level	df	t value	*p* value
a	0.369	0.132	0.502	0.237	0.151	<0.001	93	2.789	0.006
b	-0.749	0.161	-0.926	-0.573	0.195	0.109	93	-4.664	<0.001
c	-0.164	0.041	-0.164	-0.164	<0.001	0.098	93	-3.972	0.001
d	0.874	0.033	0.874	0.874	<0.001	0.102	93	26.484	<0.001

To determine whether there was a significant difference between stimulus conditions, I ran the full model including condition and participant as fixed effects and a reduced model where only the participant was included as a fixed effect. There was a significantly better fit to the data when condition was included as a fixed effect (Log Likelihood Ratio = 36.567, $p < .001$), demonstrating that the two conditions are fitted by different parameters, and are significantly different. Performance for the Aperture condition was higher at all spatial frequencies, and particularly at the lowest spatial frequencies, where the Carrier condition had a steep decline in performance, and the Aperture condition only a shallow decline (see Figure 2.12, left).

Analysis of conflicting cue conditions

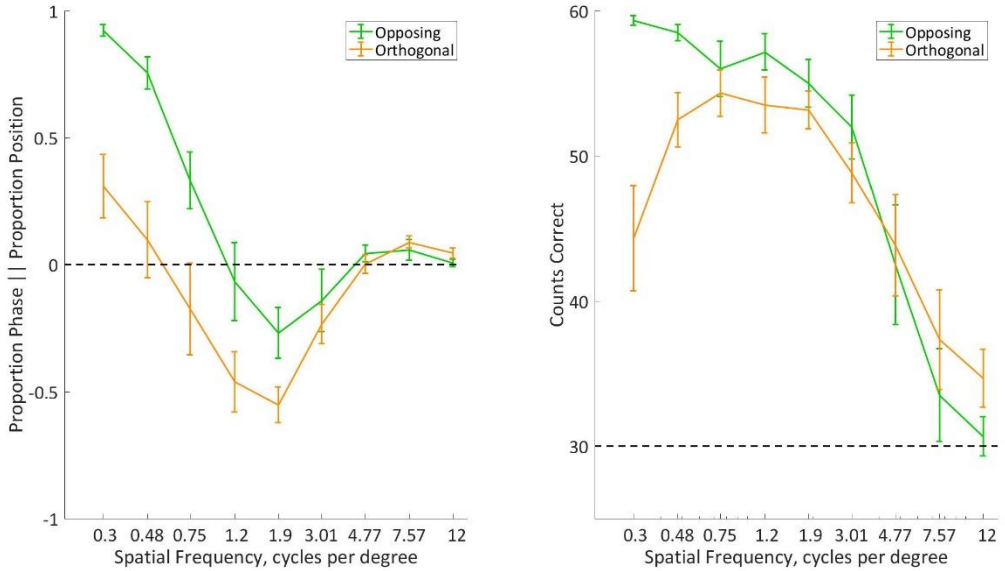


Figure 2.13. Two methods of representing the reported trajectories in the Opposing and Orthogonal Motion conditions on a two dimensional figure. The left hand panel displays the proportion of the trials reported in the carrier or aperture cue direction, where proportions were calculated for each participant, and mean proportion plotted here with error bars representing the standard error of the mean. A proportion > 0.5 indicates that motion was predominantly reported in the aperture cue direction, and a proportion < 0.5 indicates that motion was predominantly reported in the carrier cue direction. The right hand panel plots the mean number of trials reported in either the carrier or aperture cue direction, aggregated together as 'correct' responses, where the incorrect responses would be reported in the 90° or 270° direction for the Opposing Motion condition, and in the 180° or 270° direction for the Orthogonal Motion condition.

I ran three separate analyses to investigate reported motion direction in the conflicting cue conditions. I first verified that there was a significant difference in the proportion of the number of trials reported in either the carrier or the aperture cue direction by spatial frequency, separately within the Opposing and Orthogonal Motion conditions (Figure 2.13, left hand panel). I then compared whether there was a significant difference between the Opposing and Orthogonal motion conditions in the overall proportion correct – trials reported in either the carrier or the aperture cue direction (Figure 2.13, right hand panel). Finally, I analysed whether there was a significant difference in the proportion of trials reported in the carrier or aperture cue direction, by comparing the proportion of trials reported in these directions, comparing between the Opposing and Orthogonal Motion conditions (Figure 2.13, left hand panel).

Comparing the proportion of trials reported in the carrier and aperture cue direction, separately within the Opposing and Orthogonal Motion conditions

To determine in which direction motion was predominantly reported, I analysed the number of trials reported in the carrier and aperture cue directions, separately for the Opposing and Orthogonal Motion conditions, across the 9 spatial frequencies. I calculated the proportion of trials reported in the carrier and aperture cue directions, so that a proportion greater than 0.5 indicates a majority of trials reported in the aperture cue direction, below 0.5 indicates a majority of trials reported in the carrier cue direction, and a proportion of 0.5 indicates equal proportions of trials reported in the carrier and aperture cue directions (see Figure 2.13, left hand panel). I analysed the proportion of trials reported in the carrier and aperture cue directions separately for the Opposing and Orthogonal Motion conditions, using a series of Bonferroni corrected one sample t tests comparing the proportion of trials to 0.5, across all six participants.

The results are reported in Table 2.7. In the Opposing Motion condition there is a significant difference in the proportion of trials reported in the carrier and aperture cue directions. At 0.30 cpd, 0.47 cpd and 0.75 cpd there are significantly more trials reported in the aperture cue direction, and at 1.90 cpd there are significantly more trials reported in the carrier cue direction. In the Orthogonal motion condition there are significantly more trials reported in the aperture cue direction at 0.30 cpd, 7.57 cpd and 12.00 cpd, and significantly more trials reported in the carrier cue direction at 1.20 cpd, 1.90 cpd and 3.01 cpd.

Table 2.7. Table displaying the results of the Bonferroni corrected one sample t tests, run at each spatial frequency on the proportion of trials reported in the carrier or aperture cue direction. Analyses were run separately for the Opposing and Orthogonal motion conditions, and presented in adjoined tables here for ease of comparison.

Spatial Frequency, cpd	Opposing Motion				Orthogonal Motion			
	Mean	SE	t value	*p* value	Mean	SE	t value	*p* value
0.30	0.972	0.010	45.544	<.001*	0.940	0.058	4.130	0.009*
0.47	0.898	0.322	12.351	<.001*	0.547	0.090	0.521	0.625
0.75	0.692	0.057	3.376	0.020*	0.379	0.099	-1.221	0.248
1.20	0.464	0.078	-0.469	0.659	0.212	0.064	-4.495	0.006*
1.90	0.338	0.061	-2.652	0.045*	0.145	0.044	-7.984	<0.001*
3.01	0.433	0.749	-0.891	0.414	0.318	0.062	-2.937	0.032*
4.77	0.563	0.032	1.959	0.108	0.523	0.028	0.826	0.447
7.567	0.584	0.040	2.100	0.090	0.630	0.046	2.850	0.036*
12.00	0.510	0.030	0.320	0.762	0.585	0.031	2.836	0.036*

Comparing proportion correct in the Opposing and Orthogonal Motion conditions

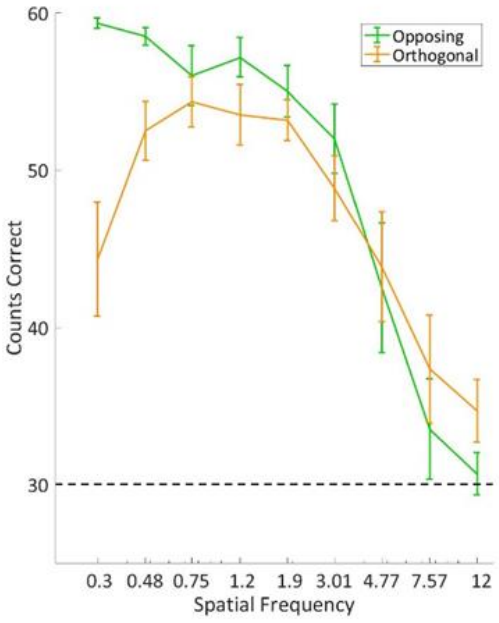


Figure 2.14. Figure displaying the proportion correct – the proportion of total trials reported in either the carrier or the aperture cue direction – for the Opposing and Orthogonal Motion conditions. Proportion correct was calculated for each participant, and the error bars represent the standard error of the mean.

I also calculated proportion correct for the Opposing and Orthogonal Motion conditions, by summing the trials reported in the carrier and aperture cue directions – 0° and 180° for the Opposing Motion condition and 0° and 90° for the Orthogonal Motion condition– and calculating these as a proportion of the total 60 trials per stimulus condition (See Figure 2.14). Performance for the Opposing Motion condition was higher at the lowest 5 spatial frequencies, and lower at the two highest spatial frequencies, with the largest difference between conditions at the lowest spatial frequencies where there was a steep decline in the Orthogonal Motion performance, but the Opposing Motion performance had plateaued. These data were best described by a third order polynomial;

$$\text{proportion correct} \sim a^* \log10\,(SF)^3 + b^* \log10\,(SF)^2 + c^* \log10(SF) + d$$

The parameters and coefficients of the model are reported in Table 2.8. Whilst the c parameter did not contribute significantly to the fixed parameters of the model, this model was a significantly better fit to the data than a reduced third order polynomial without the c parameter (Log Likelihood Ratio = 29.1589 , p <.0001). There is very little difference in the a and b parameters between the model fit for the Opposing and Orthogonal Motion condition data, but the c parameter is lower, and the d parameter higher in the Opposing Motion condition, and these parameters also show variation across the participant level. The d parameter is the intercept of the model, and indicates that there is a greater overall performance in the Opposing Motion condition, as displayed in Figure 2.14.

Table 2.8. Table reporting the parameters and coefficients for the third order polynomial model fitted to the proportion correct data in the Opposing and Orthogonal motion conditions.

Parameter	Coefficient	Std. error	Coefficients for the Opposing Motion condition	Coefficients for the Orthogonal Motion condition	Std. dev. across the condition level	Std. dev. across the ppt. level	df	t value	*p* value
a	0.127	0.063	0.127	0.127	< 0.001	< 0.001	93	2.019	0.046
b	-0.404	0.059	-0.404	-0.404	< 0.001	< 0.001	93	-6.860	< 0.001
c	-0.095	0.059	-0.155	-0.034	0.068	0.071	93	-1.614	0.110

d	0.938	0.031	0.964	0.912	0.033	0.056	93	30.433	< 0.001

To determine whether there was a significant difference between stimulus conditions, I ran the full model including stimulus condition and participant as fixed effects, and a reduced model where only participant was included as a fixed effect. There was a significantly better fit to the data when condition was included as a fixed effect (Log Likelihood Ratio = 22.627, $p < .001$), demonstrating a significant difference in proportion correct between the two conditions

Comparing the proportion of trials reported in the carrier and aperture cue direction between the Opposing and Orthogonal Motion conditions

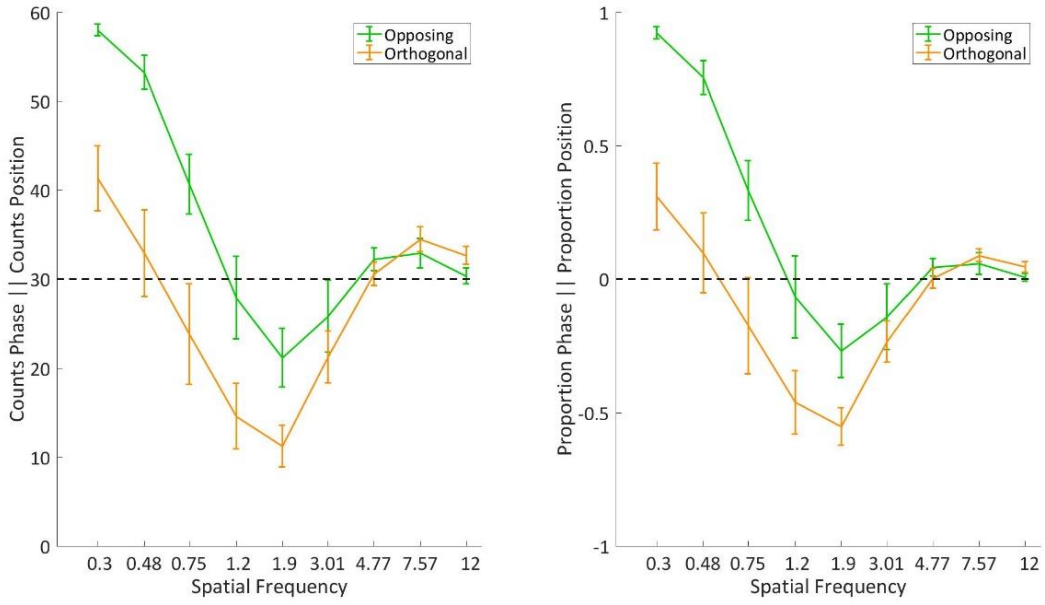


Figure 2.15. Figure displaying the proportion of trials reported in the carrier and aperture cue direction for the Opposing and Orthogonal Motion conditions. In the left panel, this was calculated using the proportions of the total number of trials, and in the right panel, these values were scaled by the proportion of trials reported in either the carrier or aperture cue direction, so only correct trials were included. For my analysis I used the right hand panel; data that only looked at the trials reported in a valid motion cue direction.

The reported directions in the Opposing and Orthogonal Motion conditions were used to calculate the proportion of trials reported in the aperture cue direction, as a proportion of the total number of trials reported in either the carrier or aperture cue direction. In this method of opposing the number of trials reported in the carrier or aperture cue direction, errors were not included (trials reported in

the 90° and 270° direction in the Opposing Motion condition and in 180° and 270° in the Orthogonal Motion condition). This is displayed in the left hand panel of Figure 2.15. An alternative method for calculating the relative proportion of responses in the carrier or aperture cue direction is to calculate the relative performance of the two conditions and scale by the proportion of trials reported in either of these correct) directions. This scaled data are displayed in the right hand panel of Figure 2.15. This can be represented as;

$$(\text{aperture} - \text{carrier}) / ((\text{aperture} + \text{carrier})/60)$$

This produces very similar results to the unscaled data, but has a 'proportion correct' closer to chance across all spatial frequencies. This suggests that including or excluding the error trials has little effect on the relationship of reported motion direction by spatial frequency – that the switch between using the aperture cue at low spatial frequencies and carrier cue at mid-range spatial frequencies is driven by the relative strength of the two cues, and not by a dramatic decline in performance.

As can be seen in Figure 2.15, trials for the Opposing Motion condition were reported predominantly in the aperture cue direction for the lowest three spatial frequencies, and in near- equal proportions in the carrier and aperture cue directions for the all spatial frequencies above 1 cpd. In contrast, motion in the Orthogonal Motion condition was reported in predominantly the aperture cue direction only at the lowest spatial frequency, and in the carrier cue direction for spatial frequencies from 0.75 cpd to 3.01 cpd. Performance was very similar between the two conditions at the three highest spatial frequencies. I analysed the unscaled data (right hand panel of Figure 2.15), fitting the 'proportion correct' with a 4^{th} order polynomial, which captures the inflection point at low spatial frequency and two turning points at mid and spatial frequency;

$$\text{proportion correct} \sim a*\log10(SF)^4 + b*\log10(SF)^3 + c*\log10(SF)^2 + d*\log10(SF) + e$$

The parameters and coefficients are displayed in Table 2.9, demonstrating that the model had a significant contribution from each of the five parameters, suggesting that it was an appropriate choice of equation. There was little variation across the a and b parameters at either the condition or participant level. However, the coefficients for the c and d parameters were lower in and the coefficient for the e parameter was higher for the Opposing Motion condition than the Orthogonal motion condition.

Table 2.9. Table reporting the parameters and coefficients for the fourth order polynomial model fitted to the proportion of trials reported in the carrier and aperture cue directions in the Opposing and Orthogonal motion conditions.

Parameter	Coefficients	Std. error	Opposing Motion condition	Orthogonal Motion condition	Std. deviation across the condition level	Std. deviation across the ppt. level	df	t value	*p* value
a	-1.699	0.220	-1.699	-1.699	< 0.001	< 0.001	92	-7.715	<.001*
b	1.479	0.261	1.479	1.479	< 0.001	< 0.001	92	5.667	<.001*
c	1.295	0.111	1.266	1.324	0.052	0.149	92	11.686	<.001*
d	-0.849	0.097	-0.925	-0.773	0.086	0.069	92	-8.788	<.001*
e	0.394	0.091	0.495	0.292	0.112	0.125	92	4.347	<.001*

From these coefficients we can extract information about the fitted model, such as its turning points or local maxima and minima. In these experiments I was most interested in the local minima, as this refers to the spatial frequency and proportion value corresponding to the point where the highest proportion of trials were reported in the carrier cue direction. The local minima was at 1.71 cpd and 0.198 'proportion correct' for the Opposing Motion condition, and at 1.93 cpd and 0.357 'proportion correct' for Orthogonal Motion condition. This demonstrates that there is a greater proportion of trials reported in the carrier cue direction in the Orthogonal Motion condition than the Opposing Motion condition, and that the highest proportion of trials reported in the carrier cue direction were reported at a lower spatial frequency in the Opposing than the Orthogonal Motion condition. This, along with the differences in the c, d and e parameter coefficients reported in Table 2.9, shows that the two curves aren't simple translations along the x or y axis, but that there is a more complex relationship between trials reported in carrier and aperture cue directions with spatial frequency, between the two stimulus conditions.

To determine whether there was a significant difference between conditions, I compared the goodness of fit between the full model which included condition and participant as fixed effects, and a reduced model where only participant was included as a fixed effect. There was a significantly better fit to the data when condition was included as a fixed effect (Log Likelihood Ratio = 53.881,

$p < .001$) demonstrating that the two conditions are fitted by different parameters, and are significantly different.

Discussion

Reported motion direction, re-coded as motion trajectory, for all four Gabor stimulus conditions showed a similar effect of spatial frequency on performance, where the proportion of trials in a valid motion cue direction follows a quadratic relationship with spatial frequency, with peak performance at low spatial frequencies and performance declining to chance level with increasing spatial frequency (Figure 2.12). In the conflicting cue conditions, the dominant direction of reported motion, where motion was reported in the aperture cue direction at low spatial frequencies, predominantly in the carrier cue direction at mid-range spatial frequencies, and at chance at the highest spatial frequencies (Figure 2.13, left hand panel), can also be compared. This relationship was observed across all participants (Figure 2.25-2.31), with all paticipants reporting motion in the carrier cue direction more than the aperture cue direction at at least one spatial frequency. However, there was individual variation in the bias towards the carrier cue direction. Participant MXG showed the strongest motion reversal for both the Opposing and Orthogonal Motion stimuli, whereas for participants LMJ and SJH reported motion direction was more evenly split between the carrier and aperture cue directions. There is also individual variation in the difference between motion reversals in the Opposing and Orthogonal conditions, with some participants showing substantially more bias towards the carrier cue in the Opposing than the Orthogonal Motion condition (participants ARD and ARH), and some participants showing only a slight increase in this bias for the Opposing Motion condition (participant SJH). These differences appear to be driven

The reported direction in the conflicting cue conditions indicates which motion cue has a greater relative strength at that eccentricity. Both Opposing and Orthogonal Motion conditions showed a consistent bias towards the carrier cue at mid-range spatial frequencies, between 1.2 cpd and 3.01 cpd in the Opposing Motion condition, and between 0.75 cpd and 3.01 cpd in the Orthogonal Motion condition. The reported motion in the Carrier and Aperture stimulus conditions does not follow the same trend in results; performance in the Aperture condition was higher than performance in the Carrier condition at all spatial frequencies, and significantly higher at 0.3 cpd and 3.01 cpd. This suggests that proportion correct on single cue conditions is not a good predictor of motion trajectory on conflicting cue conditions. This could be a consequence of the simplistic measurement of single cue performance; assessing contrast thresholds for the single cues may be a better descriptor of the strength of the carrier and aperture cues at each spatial frequency, and produce a better predictor for conflicting cue performance.

In the two conflicting cue conditions, participants consistently reported motion in either the carrier and aperture cue directions, except at the highest spatial frequencies, as demonstrated with group level chi squared analysis, and the non-linear mixed effects models on proportion correct and the proportion of trials reported in the carrier and aperture cue directions. This indicates that performance at mid-range spatial frequencies is a result of participants choosing either the carrier or aperture cue in near-equal proportions, and not incoherent motion produced through equally balanced cue strengths. The similarity in results between the Opposing and Orthogonal Motion conditions, despite the different arrangement of the carrier and aperture cues within the stimuli, supports this interpretation of results. Reported motion is not just the result of errors distributed along a single motion axis, but a reversal in reported trajectory. These reversals in motion trajectory are consistent with spatial frequency driven reversals in reverse phi stimuli (Chubb & Sperling, 1989). Chubb & Sperling also demonstrated motion reported in the position cue direction at low spatial frequencies, and in the phase cue direction at high spatial frequencies. This similarity between the results in Gabor experiments using carrier and aperture cues and results in experiments using reverse phi stimuli with phase and position cues suggests that there may be differences in the spatial frequency sensitivity of the phase and position motion systems, and that this may be a fruitful avenue for exploration.

There is a significant difference in the distribution of trials reported in the carrier and aperture cue direction in the Opposing and Orthogonal Motion conditions. In the Opposing Motion condition there were more trials reported in the aperture cue direction, and in the Orthogonal Motion condition, more trials reported in the carrier cue direction. At the lowest spatial frequencies, this difference could be partially explained by differences in the distribution of errors in carrier and aperture cues, which affect the Opposing and Orthogonal Motion conditions differently due to the different arrangement of the carrier and aperture cues. Errors in the Carrier condition at low spatial frequencies were predominantly in the opposite direction to the carrier cue direction (180°). This is probably due to the aperture effect, where motion in phase shifted grating stimuli viewed through an aperture was perceived orthogonal to the grating orientation. In comparison, the displacement of an aperture over a grating does not have any similar constraints in perceived direction, so there were no directional biases in errors associated with the aperture cue. These carrier-cue associated errors can be seen in the 270° direction at low spatial frequencies in the Orthogonal Motion condition. However, in the Opposing Motion condition errors more likely to occur along the vertical axis from the carrier cue, and so errors will fall into the 0° and 180° directions, and be classified as reported cue trials. The carrier cue associated errors could explain this difference in performance

between the Opposing and Orthogonal Motion conditions at low spatial frequencies, but not mid-range spatial frequencies. These errors will likely vary between participants, as can be seen in the high proportion of errors in the Carrier condition at 0.3 cpd, where participants ARH, CJH and LMJ (Figures 2.26-2.28) all reported a substantial number of trials in the opposite direction to the phase displacement, whilst this error was less apparent in the other three participants.

The difference in the distribution of trials between the carrier and aperture cue directions could also be affected by the orientation of the grating relative to the aperture cue. The orientation of the grating in relationship to the aperture cue direction is horizontal in the Opposing Motion condition, and vertical in the Orthogonal Motion condition. These cannot be adjusted, since the carrier cue can only be perceived orthogonal to the grating orientation, but could be contributing to a difference in strength of the aperture cue of the stimulus. This could be explored by comparing performance, or contrast sensitivity, to aperture cues in Gabors of various orientations. This suggests that the orientation of the carrier and aperture cues in these stimuli has a significant effect on motion trajectory, but that these effects – differences in error distribution due to the orientation of the cues, and potential differences in the strength of the aperture cue – are unavoidable in Gabor stimuli with these cue orientations.

Asking participants to report the perceived motion direction of conflicting cues revealed which cue was perceived as having a greater relative strength. Longer frame displays of Gabor stimuli with an orthogonal orientation of carrier and aperture cues have reported motion purely in the aperture cue direction at fovea, but when viewed peripherally reported motion in a diagonal trajectory, intermediate between the carrier and aperture cue directions (Shapiro *et al.*, 2010; Cavanagh & Tse, 2019). In my experimental procedure, it is unclear whether participants were perceiving trials in the Orthogonal Motion condition as having a diagonal trajectory of intermediate angle between the aperture cue (0^o) and carrier cue (90°) directions. Participants could be choosing an angular criterion and using this to assign reported motion into the carrier or aperture cue directions, or the two-frame display may cover an insufficient distance for motion integration to occur. To differentiate between these two possibilities more response directions are required. Similarly, perceived motion in the Opposing Motion condition may be produced from an integrated percept of the carrier and aperture cues, but this would result in the stimulus being perceived to move slower, as suggested by Hughes (Hughes, 2018), or covering a shorter displacement distance. Neither of these percepts would be revealed from reported direction. This current methodology does demonstrate the use of conflicting cues in two-frame single Gabors.

The simplicity of the Gabor stimulus construction and flexibility in stimulus parameters makes them an ideal stimulus to investigate the spatial and temporal profiles of underlying physiological systems that may separately process the phase and position cues. The carrier and aperture cues in these Gabor stimuli can be varied separately in spatial frequency, by changing the amount of carrier or aperture displacement. In this experiment I matched the displacement of the carrier and aperture cues which prevents relative motion between frames of the Opposing Motion stimulus. At the peripheral location tested, this matched displacement produced a reversal of predominant motion from aperture to carrier cue in the Opposing and Orthogonal Motion conditions. Potentially decreasing the displacement in the aperture cue would have produced a greater number of trials reported in the carrier cue direction. In Gabor stimuli presented over multiple frames, the temporal frequency of the grating and window displacements can be varied separately (Heller *et al.*, 2021). The contrast of the carrier and aperture cues cannot be varied separately, when presented in the conflicting cue arrangement, but unlike other conflicting cue stimuli, such as the MFG and reverse phi stimuli, the carrier and aperture cues can be presented separately, and contrast sensitivity functions mapped for the carrier and aperture cues.

This experiment successfully demonstrated that two-frame displays of individual Gabors can capture consistent reports of motion in the carrier and aperture cue directions, in both single and conflicting cue conditions. Longer frame Gabor experiments suggest an integration between the carrier and aperture cue into the percept of one object smoothly drifting on diagonal trajectory (Tse *et al.*, 2006). Researchers have varied the temporal frequency of the grating and aperture separately (Heller *et al.*, 2021) to produce models of cue integration. Including the single cue Gabor conditions, and asking participants to report direction in addition to distance, could further improve these models of integration. Most natural stimuli will contain both phase and position cues, which may signal motion in congruent or conflicting directions, and with different magnitudes. The response to these stimuli is unlikely to be a linear sum of the response to the individual components, so separable cues which can be assessed independently and in combination are required.

The Gabor stimuli described in this section provide a set of stimuli where responses can be measured for displacements of the carrier and Gaussian aperture independently, or when these cues are pitted against each other. However, the manipulations of the carrier phase change both the phase of the Fourier components of the stimulus, and changes in the spatial position of features within the Gabor, even when the Gaussian aperture is stationary. Gabor stimuli which contain changes in phase would therefore produce the same predicted direction of motion for mechanisms that extract motion from changes in the Fourier components of stimuli, and for mechanisms using

the changing location of features within stimuli. Gabor stimuli in which motion is produced through a displacement of the Gaussian window potentially contain no changes in the phase of the Fourier components of the image, depending on how stimuli are constructed. In the 'Aperture' condition, the carrier is oriented in parallel with the displacement of the aperture and the same phase used on every frame, meaning that there are no changes in phase of the Fourier components between frames.

2.4 Motion Energy Analysis

In this chapter, I have presented grating and Gabor stimuli that attempted to produce motion stimuli with separable descriptions in the change in phase of Fourier components and the change in spatial location. The changes in the phase of the Fourier components of stimuli can be verified using Fourier analysis. In this section, I have applied the motion energy analysis of Shapiro *et al.* (Shapiro *et al.*, 2010) to my Gabor and grating stimuli. This analysis applies three-dimensional Fourier analysis to the three-dimensional representation of my stimulus, where time is the third dimension. This analysis can be applied to both the phase and position cues in the stimuli, where Fourier analysis run on the untransformed stimulus input will calculate motion energy using the changes in phase of the Fourier components of the image, i.e. the phase cues. Transforming the stimulus input by full wave rectification takes differences in contrast that would have been 'invisible' to the Fourier analysis, and renders them as luminance values; this can then be analysed by the same three-dimensional Fourier analysis, to estimate the predicted direction of position cues in my stimuli. However, whilst the change in location of a gaussian aperture can be revealed through full wave rectification, the change in position of features within a grating stimulus are not apparent, restricting the use of the nME analysis in determining whether stimuli contain position cues.

2.4.1 nME Analysis of Gabor stimuli

Method

In the Gabor experiment, I used a two-frame Gabor stimulus to limit the change in location between stimuli that contained a change in position of the Gaussian aperture, and those that did not. This analysis cannot be run accurately on two-frame stimuli; for the purposes of this analysis, I used a 32-frame stimulus.

The analysis began by constructing a three dimensional representation of the stimuli. For all grating stimuli and Gabors, I used 1 cpd as the frequency for the fundamental, with an initial phase of 180^{o}. Gabor stimuli are windowed by a Gaussian with a diameter of 1.5 cycles, and the grating stimuli by a

rectangular aperture of 1° by 5°. To produce the three dimensional representation, I produced 32 frames of the stimulus, by applying the relevant phase and position shifts specified by each stimulus design. For example, for the orthogonal Gabor, displacing the phase of the horizontal grating by 90° on each frame and displacing the position of the Gaussian aperture rightwards, by the equivalent distance. To build the 3D representation of each stimulus I replicated each frame 60 times, to produce a cube with equal length sides, and stored these frames sequentially in a three dimensional matrix. I also included blank frames at the start and end of the time axis, to remove onset and offset artefacts, and padded the x and y dimensions similarly, maintaining the equal dimensions of the cube. This is the 'first order cube', with the X and Y dimensions as the facing planes of the matrix, and the third order dimension, T, representing time (Figure 2.16).

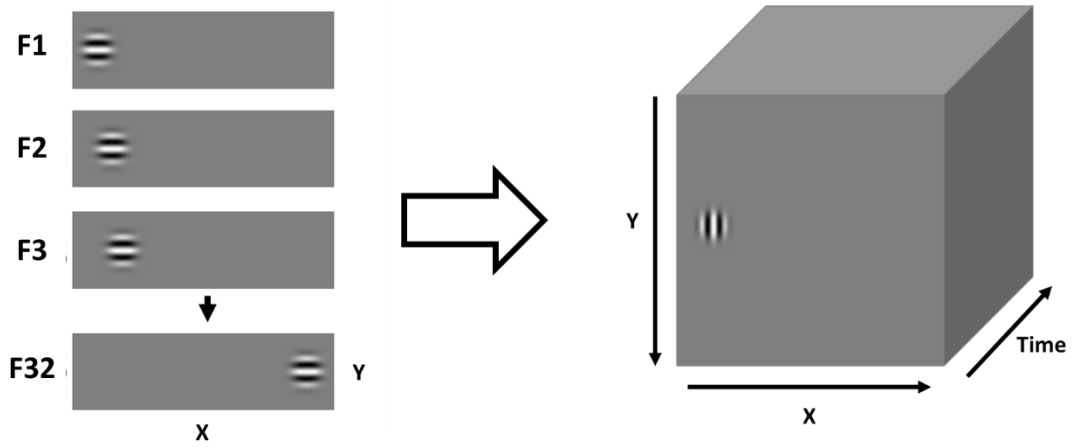


Figure 2.16. Schematic demonstrating the construction of the 3D representation of stimuli analysed in this section; 32 frames of the stimulus are produced. These frames are the XY projections of the stimulus as seen by the participants, with the excess mid-luminance background extensive in Y trimmed here for parsimony. As the frame number increases the Gaussian envelope is displaced rightwards, and the phase of the grating is displaced downwards. To produce a more continuous stimulus, the 32 frames of the Gabor stimulus are each replicated 60 times and then ordered along the third dimension to produce an equal length third side, the time dimension.

The next stage in the analysis was to produce projections onto the XT and YT faces, for the first and second order versions of the stimulus cube. To produce the second order stimulus cube, I took the first order cube and removed the zero-frequency component (the mean luminance), and then full wave rectify all values (taking the absolute value). The XT and YT projections are mathematically equivalent to taking the mean along the x-dimension to produce values for the YT plane, and along the y dimension to produce values in the XT plane. These projections may be envisioned as looking down on the cube through the top face (XT plane) or looking horizontally through the cube, onto its left face (YT plane). These XT and YT projections are images of the stimulus motion in time, just looking at the x or y directions. In Figure 2.17, the XT and YT projections for the orthogonal stimulus can be seen.

In this orthogonal stimulus, the Gaussian window of the Gabor is displaced horizontally rightwards and the phase of the carrier displaced orthogonally in the vertical direction, downwards. When looking at the first order planes there should be no change in the angle of spatiotemporal luminance in the XT plane, although there should be an overall change in location, and there should be an angled gradient of luminance in the YT plane, representing the change in phase in the vertical direction.

These can be seen in the first order XT and YT planes below (Figure 2.17, top row). The second order XT and YT planes are taken from the full wave rectified version of the stimulus cube. This transformation is designed to segment input into features and background, so one would expect the overall location of the Gabor to be dominating the XT projection, with a downward slope in the XT plane, to represent the change in x-coordinate of the Gabor with time, but no coherent direction in the YT plane, as the overall location of the Gabor did not change in the Y dimension (Figure 2.17, bottom row).

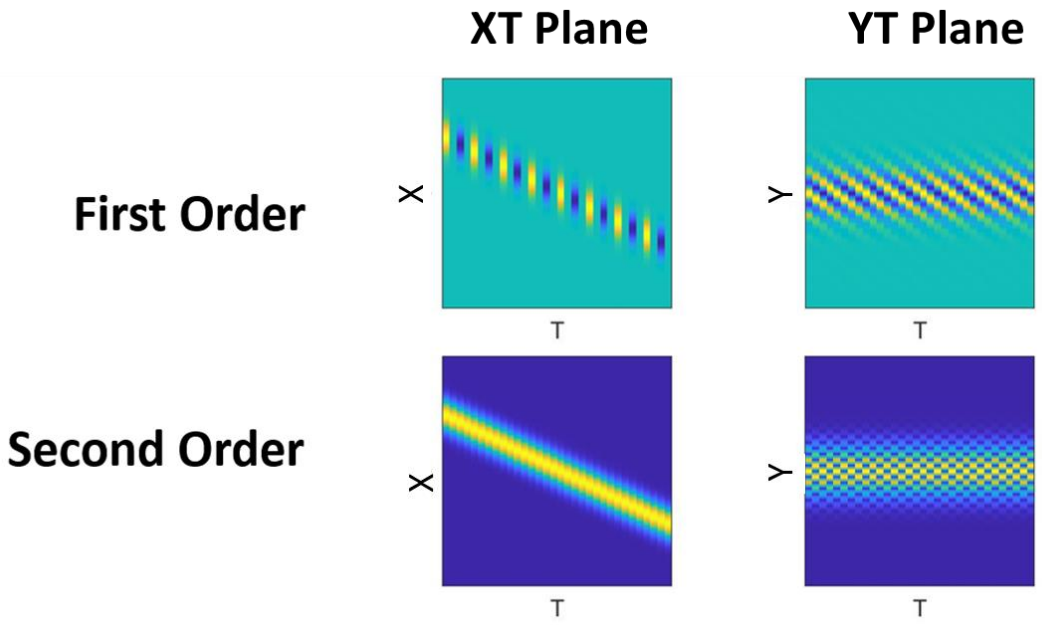


Figure 2.17. Figure displaying the XT and YT projections, or 'planes', from the first and second order transformations of the stimulus cube for the Orthogonal Motion stimulus. The angle of luminance in the projections describes the change of in luminance with time, where diagonally oriented luminance indicates coherent motion direction, and vertical or horizontal luminance represents stationary components. The brighter the pixel, the higher the luminance, with turquoise representing the mean luminance of the stimulus.

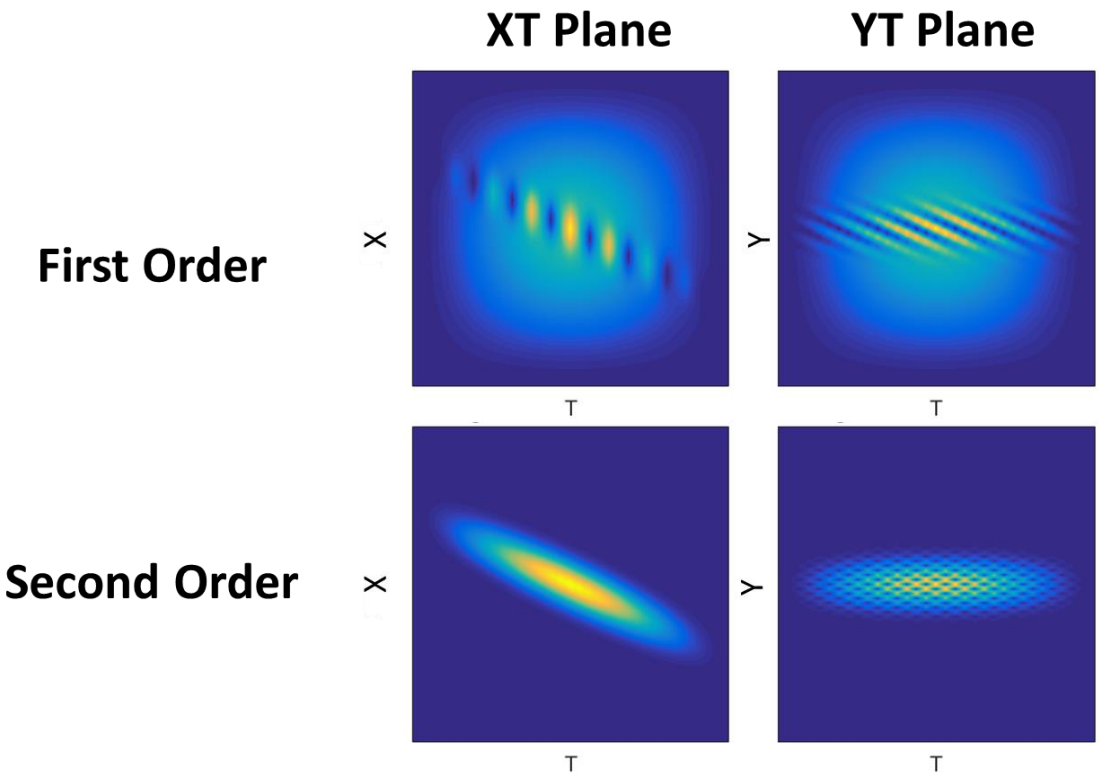


Figure 2.18. Figure displaying the result of Hanning windows applied to the XT and YT planes of the first and second order transformations of the stimulus cube for the Orthogonal Motion stimulus. The Hanning window smooths the onset and offset of the 2D XT and YT planes, which removes the edge artefacts that would otherwise appear in the Fourier analysis.

I ran a 2D fast Fourier transform on the first and second order XT and YT planes. The output of the Fourier transform is in the amplitudes of the detected frequencies in the image. The quadrants of the output are arranged so that the centre of the square refers to the lowest spatial frequencies, and the angle of the signals from this origin refer to the orientation of the frequencies in the XT or YT input (First order; Figure 2.19, Second order; Figure 2.20). This rearrangement of the quadrants means that rightward motion is represented by a leftward diagonal, and vice versa in the XT, and that in the YT projection, the positive nME value associated with downward motion is seen in the quadrants along the left diagonal, and upward motion would be represented by amplitudes in the quadrants along the right diagonal. The brighter the pixel, the greater the amplitude of the detected frequency.

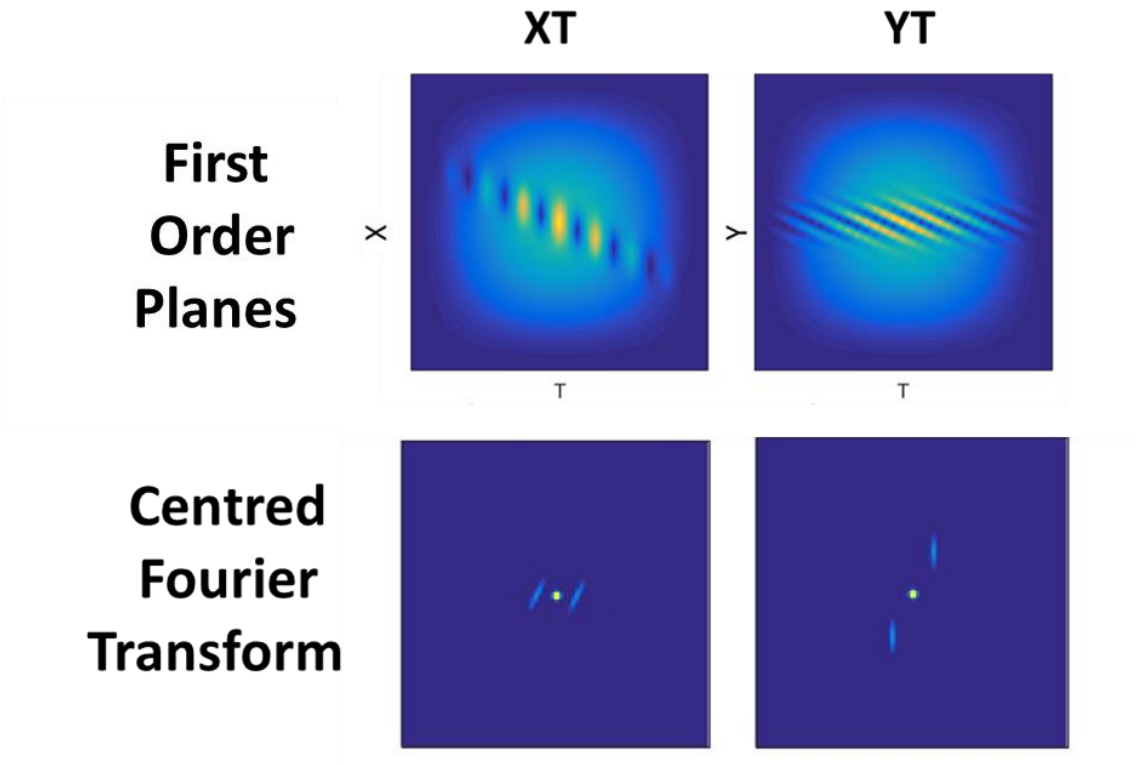


Figure 2.19. Figure displaying the centred fast Fourier transformations of the first order XT (top left) and YT (top right) planes. The Fourier analysis of the first order stimulus produces an nME value of -0.123 in the XT plane and 0.276 in the YT plane, suggesting that a first order system will perceive the stimulus as moving downwards.

The centred fast Fourier transforms can be used to calculate whether there is a substantial directional signal in a stimulus. If there is a significant imbalance in the amplitude, or 'motion energy', in the opposing quadrants of the Fourier transform, then this indicates that there may be a coherent directional signal. One way of representing this is by using polar plots (Figure 2.21 and Figure 2.22, bottom row), where the motion energy is plotted for each centred 15 degree segment of the centred Fourier transform. This is a method of plotting the central Fourier transforms that can make more explicit any bias in the orientation in the motion energy signal.

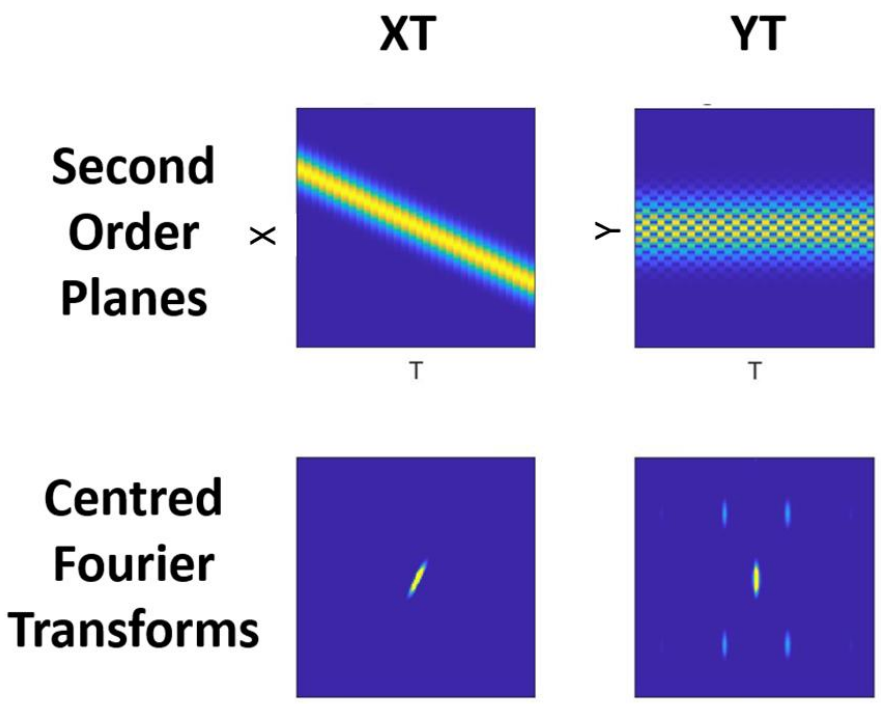


Figure 2.20. Figure displaying the centred fast Fourier transformations of the second order XT (top left) and YT (top right) planes. The quadrants in the Fourier transform of the YT planes have similar overall motion energy (nME = -0.018),whereas the quadrants of the XT plane have a clear imbalance between quadrants 1 and 3, and 2 and 4, resulting in a high nME value of 0.628. This indicates that a second order system would report this stimulus as moving rightwards, as seen by the leftward oriented streak in the XT centred Fourier Transform.

From the Fourier output, I calculated motion energy (ME) by summing the power in each quadrant of the Fourier output, where quadrants are labelled from the bottom left clockwise, due to my rearrangement for display. These quadrants are then opposed, to compare the summed motion energy in quadrants 1 and 3 with the summed motion energy in quadrants 2 and 4; assessing if there is an asymmetric distribution of motion energy across the vertical axis. A normalised measure of motion energy (nME) is calculated by dividing the ME value by the sum of all quadrants, (Equation 2), to account for contrast-gain control in motion systems (Lu & Sperling, 1996).

$$ME = (Q1+Q3) - (Q2 + Q4) \quad (1)$$

$$nME = ((Q1+Q3) - (Q2 + Q4)) / (Q1+Q2+Q3+Q4) \quad (2)$$

The nME values can be used to predict the perceived motion of these stimuli when viewed by a first or second order system. The sign of the ME gives the predicted direction, where a positive ME in the

XT plane refers to rightward motion and a negative ME to leftward motion, and in the YT plane, upward and downward motion. The magnitude of the nME value refers to the strength of that motion signal, and, by comparing the normalised motion energy values of the XT and YT planes, I can predict which direction will dominate the percept when viewed by a first order or second order system. For the Orthogonal Motion stimulus that I have used to illustrate this analysis procedure, the Fourier transform of the first and second order stimulus cube reveals conflicting dominant motion energy signals. The Fourier analysis of the first order stimulus produces an nME value of -0.123 in the XT plane and 0.276 in the YT plane, suggesting that a first order system will perceive the stimulus as moving downwards (see Figure 2.21). However, the Fourier analysis of the second order stimulus produces an nME value of 0.628 in the XT plane and -0.018 in the YT plane, suggesting that a second order system will perceive the stimulus as moving rightwards (see Figure 2.22). Therefore this Fourier analysis procedure can correctly extract the conflicting phase and position cues in the Orthogonal Motion stimulus, and can be applied to other motion stimuli as an analytical method to detect the presence of phase and position motion cues.

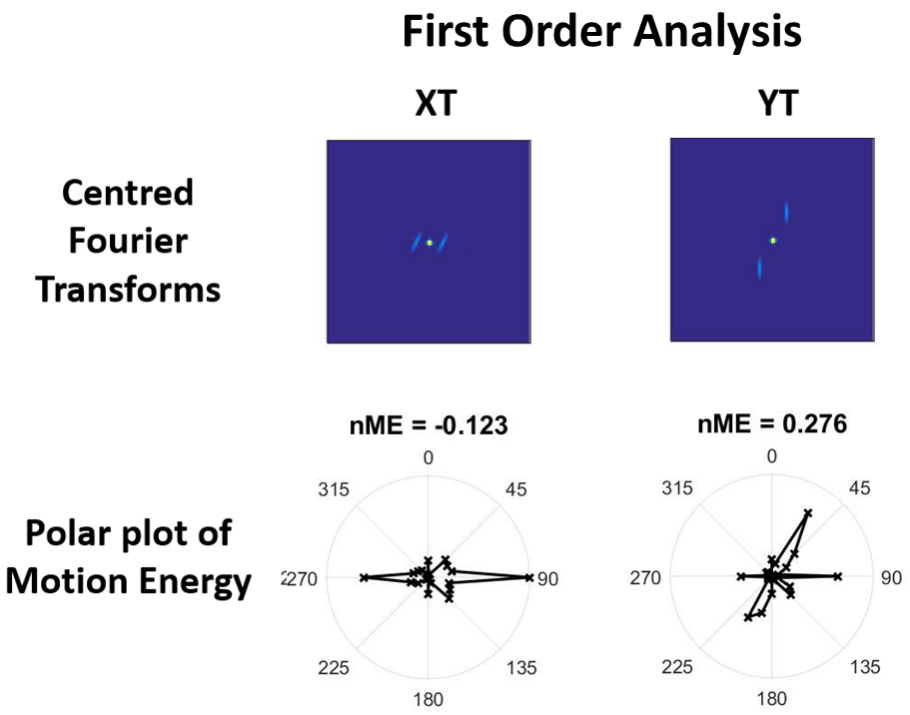


Figure 2.21. Top Row; Figure displaying the centred Fourier transforms of the XT and YT planes of the first order transformation of the stimulus cube for the Orthogonal Motion stimulus. Bottom Row; Polar plot representation of the motion energy values in the centred Fourier transforms, where the motion energy is plotted for each centred 15 degree segment of the centred fast Fourier transform. The nME value is larger for the analysis of the YT plane, suggesting that if viewed by a system that uses phase cues to discriminate motion, that the stimulus would be viewed as moving downwards.

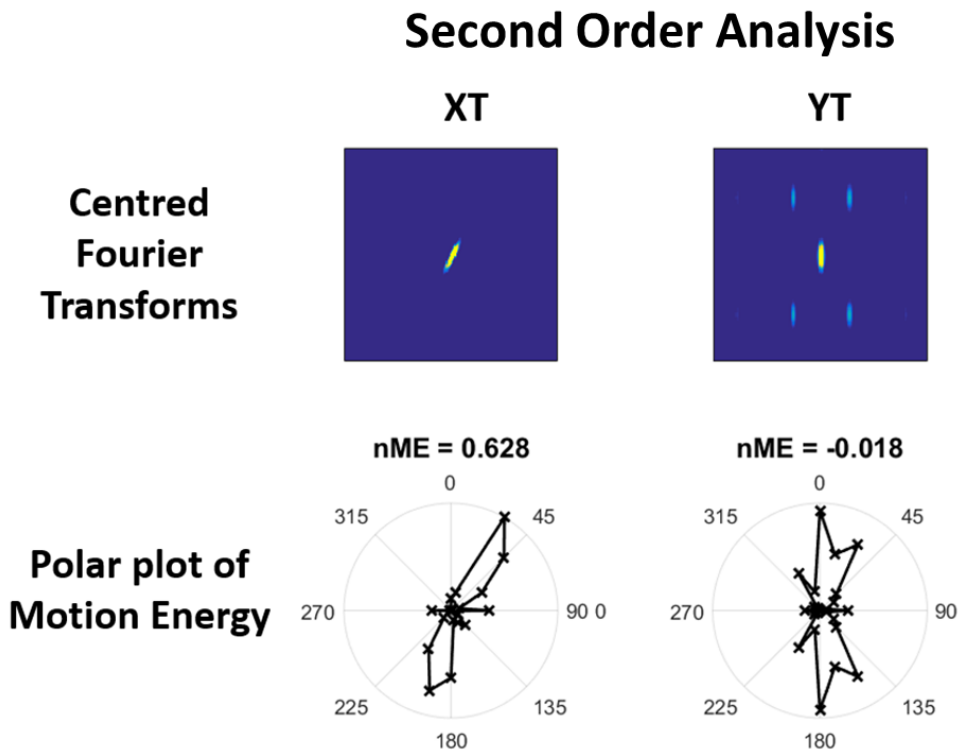


Figure 2.22. Top Row; Figure displaying the centred Fourier transforms of the XT and YT planes of the second order transformation of the stimulus cube for the Orthogonal Motion stimulus. Bottom Row; Polar plot representation of the motion energy values in the centred Fourier transforms, where the motion energy is plotted for each centred 15 degree segment of the centred fast Fourier transform. The nME value is larger for the analysis of the XT plane, suggesting that if viewed by a system that uses position cues to discriminate motion, that the stimulus would be viewed as moving rightwards.

Results

I ran this analysis for all 4 Gabor stimuli included in Chapter 2.3, and a further Gabor stimulus from Chapter 3, the 'Aperture + Rotation' stimulus, which was constructed from a drifting Gaussian window with simultaneous rotation of the sinusoid grating (See Figure 2.23, fifth column). The aim of this analysis was to verify the direction of the Fourier components within these Gabor stimuli, which are analysed in the first order analysis. The second order analysis applies a full wave rectification to the 3D representation of the stimulus before running the 3D Fourier analysis, with this transformation making explicit the change in position of the Gaussian aperture. The first order analysis therefore assess whether stimuli contain 'phase' cues, and the second order analysis assesses the presence of position cues. However, the second order analysis does not reveal all forms of position cue, such as the change in the position of carrier features produced by a phase shift of the carrier, and so an absence of a signal in the second order analysis is not sufficient to conclude that a stimulus contains no position cues.

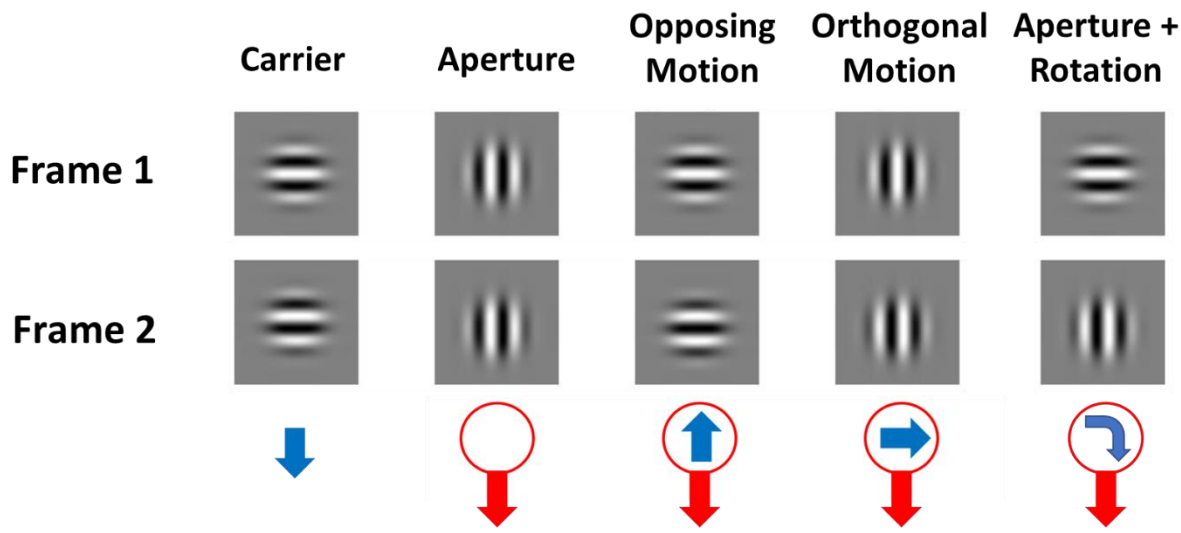


Figure 2.23. Figure displaying the two frames of downward motion for the five stimuli included in the nME analysis of this section.

The results of the Fourier analysis for the XT and YT planes from the first and second order transformations of the five stimuli are reported below in Table 2.10, and the figures detailing the Fourier analysis of the XT and YT projections of the first and second order transformations of the stimulus cube are displayed in the appendix, Chapter 2.6.3. There are small artefacts that cause small motion energy values, caused by spurious pixel correlations between frames, but the motion signals are substantially larger, and align with the design of the stimuli.

The Carrier stimulus has the largest nME value in the first order YT plane, consistent with the phase shift of the carrier, and negligible values in the first order XT plane. There were negligible nME values in either of the second order planes, consistent with the stationary nature of the aperture. Note that although this first order analysis reveals the phase cue, the lack of nME values in the second order analysis is not evidence that there are no position cues in the image, just that there are not position cues that can be detected by a Fourier analysis on a full wave rectification: the change in location of the carrier following the phase shift is a valid position cue.

The Aperture and Aperture + Rotation stimuli were designed to contain position cues, with no corresponding changes in the phase of the Fourier components. Both stimuli have negligible nME values in the first order analysis, confirming the absence of phase cues. They have negligible nME values in the XT second order plane, but have large nME values in the second order YT plane, consistent with the stimulus design – a downward displacement of the aperture.

The two conflicting cue Gabor stimuli were designed to contain phase and position cues in different directions. The Orthogonal motion condition has a substantial nME value in the first order XT plane,

and negligible nME value in the first order YT plane, indicating the presence of the horizontal phase cue produced by the phase shift of the carrier, and similarly the Opposing Motion stimuli has a substantial nME values in the first order YT plane, and negligible nME value in the first order XT plane, due to the phase shift of the vertically oriented carrier. Both the Opposing and Orthogonal motion stimuli have substantial nME values in the second order YT planes, and negligible nME values in the second order YT planes, consistent with the stimulus construction of a downwards Gaussian envelope displacement.

Table 2.10 Displaying nME values for the five analysed Gabor conditions, with the largest nME values in bold.

Stimulus	First Order		Second Order		1st Order Direction	2nd Order Direction
	XT	YT	XT	YT		
Carrier	-0.0798	**0.2761**	-0.0355	-0.0177	Down	-
Aperture	-0.0779	0.0190	-0.0177	**0.6284**	-	Down
Opposing Motion	-0.0796	**-0.4256**	-0.0355	**0.3127**	Up	Down
Orthogonal Motion	**0.2761**	-0.1233	-0.0177	**0.6284**	Left	Down
Aperture+Rotation	-0.0775	-0.0724	-0.0177	**0.3133**	-	Down

Discussion

The nME values produced for the first order analyses of my Gabor stimuli confirm that I have successfully produced stimuli containing phase cues (Carrier, Opposing and Orthogonal motion stimuli), and stimuli that do not contain phase cues (Aperture and Aperture+Rotation stimuli). The first order Fourier analysis of the multi-frame versions of these pure and conflicting cue stimuli support their use in psychophysical experiments investigating motion discrimination of phase cues, and, more importantly, demonstrates that the Aperture and Aperture+Rotation stimulus conditions can explicitly labelled as 'pure' position cue stimuli with an absence of a coherent phase cue.

The second order analysis demonstrates a substantial nME value reflecting the displacement of the aperture in the Aperture, Opposing Motion, Orthogonal Motion and the Aperture+Rotation stimuli. The presence of substantial nME signals in the first and second order analyses of the conflicting cue stimuli is consistent with the behavioural data in Chapter 2.3. The Opposing Motion stimulus has a substantial upward first order signal (nME = -0.4256), and downward second order signal (nME =

0.3127), reflecting the two predominant reported motion directions when this stimulus was viewed peripherally (Figure 2.11). Similarly, the Orthogonal Motion stimulus has a substantial rightward first order signal (nME = 0.2761), and downward second order signal (nME = 0.6284), also the two predominant reported motion directions when this stimulus was viewed peripherally (Figure 2.11).

The second order analysis was included as it can identify the overall position changes of the gaussian aperture, and is one mechanism by which a filter-rectify-filter model, such as Wilson et al. (1992), might operate: the same mechanism to extract both phase and position cues, but with different initial image processing to render position cues that are 'invisible' to first order analyses. However, the second order analysis is only one method to reveal position cues, and does not indicate the presence of the position cues caused by the phase shift of the carrier in the Carrier, Opposing and Orthogonal motion stimuli. Therefore the absence of an nME value in the second order analysis of the Carrier stimulus is not taken to be evidence of the absence of a position cue in the stimulus.

The proof-of-concept shown by the second order analysis in revealing the position cue caused by the aperture displacement could be used to extend this model: we know that there are position cues caused by the displacement of features in the carrier following a phase shift, therefore a second-order analysis with an image rectification that makes these position cues explicit could then be used to argue for possible rectification methods within the visual system.

2.4.2 nME analysis of Complex Gratings

The nME analysis was also applied to the novel grating stimuli presented in Chapter 2.2, the MF3F, to verify the presence of a coherent phase cue between the interleaved frames. The nME analysis was also applied to the third harmonic presented in isolation, as this should have the same change in phase of the Fourier components as the MF3F. The nME analysis was also applied to the MF1F of Baro & Levinson (1988) to verify the proposed absence of a change in phase of the Fourier components, as this would validate the stimulus as a 'pure' position cue stimulus. The MFG was also included in this analysis as a further comparison, as frames of the MFG are included in both the MF1F and the MF3F, and the sine wave with matched fundamental frequency.

Method

To analyse the grating stimuli presented in Chapter 2.2, I followed the same procedure for the construction of the 3D representation of the stimuli (x,y,time), and the Fourier analysis, as described in section 2.4.1 for the Gabor stimuli. I again used a 32 frame stimulus for the analysis, although only 8 frames were used in the experiments presented in this Chapter 2.2.

I produced the three dimensional stimulus cubes for the two potential 'pure' grating stimuli, the MF1F and MF3F, and three comparison stimuli, the MFG, and the fundamental frequency and third harmonic used for the construction of these stimuli. All stimuli were rendered at 1 cpd of the fundamental frequency and at the maximum contrast of the fundamental frequency, and weighted by the same vertical 5° by 1° rectangular aperture that was used to display the stimuli in the two experiments in Chapter 2.2. The YT planes of the first order rendering can be seen in Figure 2.24, and were weighted with a 2D Hanning window prior to the Fourier analysis.

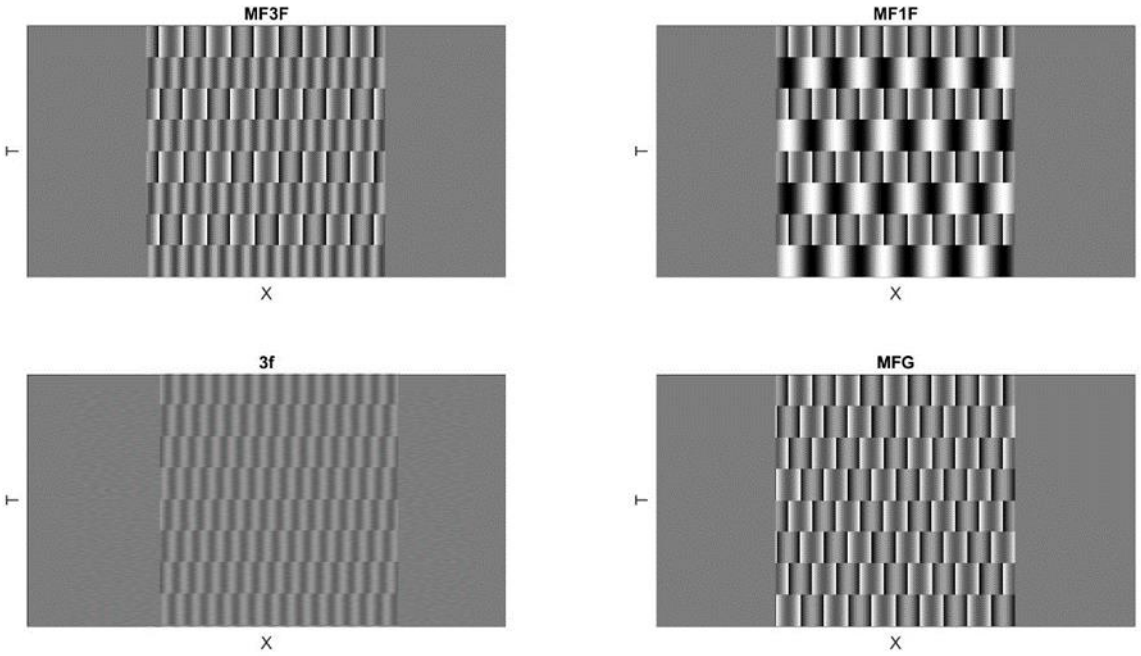


Figure 2.24. Figure displaying the frames from the YT planes, which are the input to the Fourier analysis. Position cues will be produced from the change in feature position produced from the phase shift of the grating in the MFG and 3f, and can be seen in the diagonal trajectories in the YT frames. In the MF1F and MF3F the position cues from the change in locations of peaks and troughs of the gratings, or zero-crossings, can be tracked along the diagonals in the YT planes. These position cues are not be detected by the second order analysis, which relies on figure-ground segmentation to extract position cues.

Results

The nME values for the Fourier analysis of the XT and YT planes of the first and second order stimulus cubes for the five stimuli are reported in Table 2.11. All gratings were oriented horizontally and displaced in the vertical plane, by a 90^{o} phase shift of the fundamental frequency. Expected motion is therefore only in the YT plane, with the direction of the fundamental frequency producing an expected positive nME value. The nME values of the XT reflects noise from spurious correlations of pixels between frames. The second order analysis is included to demonstrate the inability of the current analysis to detect position cues, which are produced from the phase displacements of

carriers. A 'true' position cue detector would be able to extract the position cues in the sine wave, 3F, MFG, and MF1F, and verify whether there are coherent displacements of features between interleaved frames in the MF3F, as Figure 2.24 would suggest.

There is a substantial downward signal in the YT planes of the first order rendering of the MF3F (nME= -0.1920), revealing that there are coherent changes in phase of the Fourier components between the interleaved frames of the grating. The negative value of the nME is in the expected direction, matched to the direction of the third harmonic, as the displacement of the frames by 90^{o} of the fundamental frequency from which both gratings are constructed will be aliased backwards in the third harmonic. This suggests that, as intended, there was a coherent phase cue in the MF3F stimulus.

There was no significant nME value in the MF1F stimulus, where the nME value for the YT plane produced falls within the bounds of the signals produced in the horizontal nME plane, which correspond to incoherent noise in the stimulus (nME range -0.0332 to -0.0418). This validates the use of the MF1F as a 'pure' position cue stimulus.

The sine wave grating, MFG and third harmonic all had the expected substantial nME value in the first order YT analysis (sine wave grating = 0.3350; MFG= -0.3826; third harmonic= -0.3953). The nME value for both the 3F and the MFG is negative, demonstrating that the third harmonic (and other 4n-1 harmonics) are determining the phase cues in the MFG, through the aliasing of the third harmonic due to the phase shift of the grating equivalent to a 90^{o} phase displacement of the fundamental frequency.

Table 2.11. Table displaying the nME values produced from the Fourier analysis of the XT and YT planes of first and second order renderings of 5 stimuli; the MF1F and MF3F, the conflicting cue MFG, and the fundamental frequency and third harmonic involved in the construction of these stimuli.

Stimulus	1st order XT	1st order YT	2nd order XT	2nd order YT
Sine wave	-0.0418	**0.3350**	-0.0067	-0.0177
Third harmonic	-0.0413	**-0.3953**	-0.0067	-0.0172
MFG	-0.0394	**-0.3826**	-0.0066	-0.0079
MF3F	-0.0332	**-0.1920**	-0.0063	-0.0092
MF1F	-0.0400	-0.0383	-0.0047	-0.0109

Discussion

The nME analysis of these gratings validates the presence of a coherent phase cue in the MF3F, with the direction of the nME consistent with this phase cue being produced by the third harmonic. The magnitude of the nME YT value is lower in the MF3F than the third harmonic and MFG gratings, likely due to the additional noise caused by the change in grating identity on each frame reducing the value of the motion energy once normalised using equation 2 (Section 2.4.1, Methods). However, to use the MF3F as a 'pure' phase cue, the absence of position cues must be verified. The second order analysis was not expected to produce coherent signals for any of the grating stimuli tested, as the full wave rectification in the analysis removes the directional signal produced from the quarter cycle phase shifts in the grating. Zero-crossings in the stimulus can be matched across frames and would offer a coherent cue for direction, and in the same direction as the phase cue. This means that in the current experimental results (Chapter 2.2), it is impossible to differentiate whether participants were using the phase cue, or some position cue such as the displaced zero-crossings, or mapping between areas of high and low luminance between frames.

The nME analysis found no coherent motion energy, or phase cue, in the MF1F, consistent with the stimulus design (Baro & Levinson, 1988). The motion discrimination results in Baro & Levinson suggest that for at least 2 of the 4 participants, that there is an available cue for motion in the image; Baro & Levinson suggest that participants are tracking features between the interleaved frames of the MFG and the fundamental frequency (Baro & Levinson, 1988). This validates the use of the MF1F as a grating based 'pure' position cue stimulus, and also suggests that the MF1F would be a useful stimulus to test future image analysis algorithms to see whether these can extract the position cue from the stimulus.

The position cue in these grating stimuli is not extracted in the second-order analysis, which means that the nME method cannot be used to assess whether or not position cues are present. This limits the use of the MF3F as a potential motion stimulus.

2.4.3 Discussion

This motion energy analysis can be used to verify the presence or absence of phase cues in motion stimuli, as it uses three-dimensional Fourier analysis to assess the change in phase of the Fourier components. The nME analysis in this chapter has validated the presence of phase cues in the Carrier and conflicting cue Gabors (Chapter 2.3), and the novel MF3F stimulus (Chapter 2.2).

Crucially, it has also validated the absence of phase cues in the Aperture and Aperture + Rotation Gabor stimuli and the MF1F of Baro & Levinson (1988). These stimuli could therefore be used as 'pure' position cue stimuli, the Gabor in two-frame displays that limit the change in retinal position of the stimuli, and the MF1F grating in longer frame displays where temporal frequency could be varied. The nME analysis is unable to validate the absence of position cues in the image, as the second order analysis is just one method by which a system could extract motion cues, by performing a full-wave rectification of the stimulus and then applying the same three-dimensional Fourier analysis. Thus the nME cannot be used to draw conclusions about the presence or absence of position cues in motion stimuli.

The nME analysis is a method to validate stimulus construction, and it not intended to be a description of mechanisms in the visual system: the use of Fourier analysis is not a suggestion that the visual system is using Fourier analysis. This also means that the nME values are limited in their utility, as they are useful only to indicate the direction of motion of phase cues in the image, and the magnitude of the nME value is useful only to differentiate between nME values that are random noise, and nME values that indicate a significant motion cue. To render nME values as more useful predictors of motion strength, one method would be to adjust the amplitude of the detected spatial and temporal frequencies in the model output by the human contrast sensitivity function , or to apply an initial filter to the input that filters images using measured sensitivity to both spatial and temporal frequencies.

The nME analysis is limited in its application to apparent motion stimuli with few frames, as with fewer frames the increased noise caused by spurious luminance correlations between pixels will obscure the 'signal' of the changing phase of the Fourier components of the image. In this Chapter the grating stimuli were presented in eight-frame apparent motion displays (Chapter 2.2), and the Gabor stimuli as two-frame apparent motion displays (Chapter 2.3), whilst the analysis was run on 32-frame renderings of the stimuli. Whilst these stimuli were constructed to produce (or not produce) Fourier components and this analysis was used to verify their presence, on more complex stimulus designs this may restrict the use of this analysis.

This method is currently restricted to detecting position cues using Fourier analysis following an image rectification, in this instance the full wave rectification as used in Shapiro *et al*. (2010). This models the 'second order system' of Lu & Sperling's three systems architecture (Lu & Sperling, 1995). The full-wave rectification of the image means that some position cues, previously 'invisible' to the Fourier analysis now have associated phase shifts that can be extracted through Fourier

analysis. This full-wave rectification makes 'visible' features that have clear figure-ground segmentation, such as a Gabor, where the full-wave rectification makes the Gaussian envelope an explicit feature where the change in position is then a change in phase once full wave rectified. However, continuous gratings with 90° displacement on each frame have no coherent motion signal when full-wave rectification is applied, since the 'lobes' of the grating could be mapped both forwards and backwards on each consecutive frame. Both full and half-wave rectification are biologically plausible mechanisms, through the output of responses of ON-centre and OFF-centre units with matching spatial properties, such as centre-surround retinal ganglion cells or simple cells in V1. The nME analysis could easily be adjusted to include different rectification algorithms, such as squaring the input (Wilson *et al.*, 1992; Chubb & Sperling, 1991), whereby the choice of rectification could be chosen to produce an nME analysis that more closely models hypothesised mechanisms in the visual system.

A further extension to the nME model would be to include additional filtering stages, prior to the image rectification, such as an orientation selective spatial filter (Wilson *et al.*, 1992), or a temporal filter (Chubb & Sperling, 1989). This stage could also be used to implement the human contrast sensitivity function. A second filtering stage is also included in many models of motion processing, and could be used to further extend the nME. Suggestions include a low-pass spatial filter (Ledgeway & Hess, 2000), or a filter stage that selects for different orientations and lower spatial frequencies (Wilson *et al.*, 1992). Both of these filters introduce additional spatial frequencies into the transformed image input, which could be relevant for interpreting the bias towards phase and position cues. With sufficient improvements to the nME analysis, models could then be compared to determine which are a better model of human vision, to investigate different physiologically plausible mechanisms for motion perception. Applying different filters could also make explicit whether the transformations of images are introducing or eliminating essential stimulus characteristics. The drifting Kanizsa squares of Seiffert & Hock (2014) can be described as a 'pure' position cue stimulus, since the lack luminance boundaries in the Kanizsa squares, and random rotation of the pacmen producing the Kanizsa squares, mean that there are no changes in phase of Fourier components that would produce a coherent motion direction signal. However, if lateral inhibition is causing these illusory contours to be processed as 'real' contours, as psychophysical data suggests (Dresp & Bonnet, 1991; Berkley *et al.*, 1994), then the stimulus is only a 'pure' position stimulus in abstract terms, and not after initial image transformations of the visual system. An extended nME model could make these artefacts explicit.

Models based on image transformation may be unable to fully explain all cases of motion discrimination, especially for stimuli requiring explicit instruction. Lu & Sperling (1999) presented interleaved frames of grating stimuli produced from variations across a variety of image properties, such as texture, depth and orientation. Across the frames, there was no stimulus-driven variation in salience, and participants were explicitly instructed about which features to track motion displacement across. Thus, salience was applied top-down, based on arbitrary mappings, and not due to visual salience of the features. The ability of participants to perform motion discrimination across these gratings, albeit after several hundred trials of training, is evidence that motion perception can occur for stimuli solely based on assigned salience. This cannot be accounted for in models of motion perception using image transformation. However, this class of top-down salience driven motion perception could be accounted for by the salience map model of motion perception. These models could then be applied to stimuli that contain solely position cues but do not require explicit instruction to select which features to attend to, such as the MF1F of Baro & Levinson (1988), or the micro-balanced stimuli of Solomon & Sperling (1994). If the salience model can successfully extract motion from these constructions of 'pure' position cue stimuli, where other extended-nME models may fail, then this would convincing evidence in favour of the salience map hypothesis.

2.5 Conclusion

The aim of this chapter was to identify psychophysical stimuli containing phase and position cues, that could be used to explore psychophysical dissociations. If the stimuli are sufficiently well matched for image characteristics, any psychophysical dissociations could be used to discuss whether the data support the existence of multiple motion processing systems, and these psychophysical dissociations used to form hypotheses about the mechanisms within these systems.

Grating based stimuli and Gabor stimuli were explored psychophysically, and the nME analysis applied to validate the phase cues within these stimuli. The MF1F grating (Baro & Levinson, 1988) and the Aperture and Aperture+Rotation Gabors were shown to contain no phase cues, and thus motion is described purely by the change in position of salient features. In contrast, this chapter did not successfully present stimuli that contained phase cues without a corresponding position cue, with both the MF3F and the stationary Gabor with phase-shifted carrier containing a change in position of features that could be tracked between frames.

In this chapter, I have demonstrated a possible dissociation in the temporal integration between motion discriminated in the direction predicted by the change in phase of the Fourier components, and the change in position of the stimulus features, using the MFG and a novel stimulus, the MF3F. These results were limited by the presence of congruent position cue in the MF3F, where there were changes in the spatial position of features between frames that may have been used by participants to perform the motion discrimination. These results may suggest that the visual system is able to perform motion discrimination using different cues, at different time scales.

I have also demonstrated a bias in reported motion direction with spatial frequency in conflicting cue Gabor stimuli, where motion is reported in the direction of the displacement of the Gaussian aperture at low spatial frequencies, and in the direction of the change in phase of the carrier at high spatial frequencies. However, the change in phase of the Fourier components in the Gabors result in a corresponding change in position of features. Therefore the 'carrier condition' Gabor cannot be considered a 'pure' stimulus where motion is solely described through changes in the phase of the Fourier components, and the two conflicting conditions (Opposing and Orthogonal motion), also contain this confound.

The Opposing and Orthogonal Gabors are reported in the direction predicted by the change in spatial position of the features at low spatial frequencies, and in the change in phase of the Fourier components at high spatial frequencies. This is the same relationship as seen in reverse phi stimuli (gratings: Chubb & Sperling, 1989; Gorea, 1994), where motion is predicted in opposing directions by the change in the phase of the Fourier components and the change in spatial position of the features, but without the confound seen in Gabor stimuli. This similarity in results could suggest that whilst the confound is present in the Gabor stimuli, participants may be relying on the change in phase of the Fourier components at high spatial frequencies, and not using the cue of the changing location of the carrier features. The change in reported direction for the conflicting cue Gabor stimuli could be explained by the presence of two motion systems, one processing phase and the other position cues, that have distinct spatial frequency sensitivity functions.

Finally, I have used the three-dimensional Fourier analysis of Shapiro *et al.* (Shapiro *et al.*, 2010) to verify the presence and absence of phase cues in the stimuli discussed in this chapter. I suggest applications for extending this analysis to analyse position cues, and it's potential for producing models of the potential mechanisms used to compute motion direction from phase and position cues in the visual system.

2.6 Appendix

2.6.1 Appendix for Section 2.3

2.6.1.1. Histograms for individual participant data

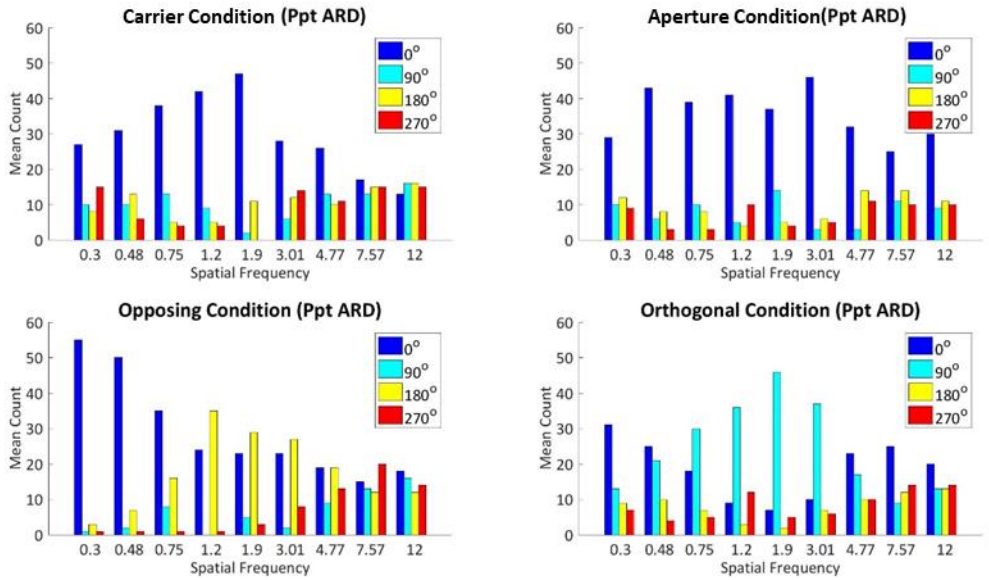


Figure 2.25 Reported motion trajectories for each of the 4 stimulus conditions, across all spatial frequencies tested, for participant ARD.

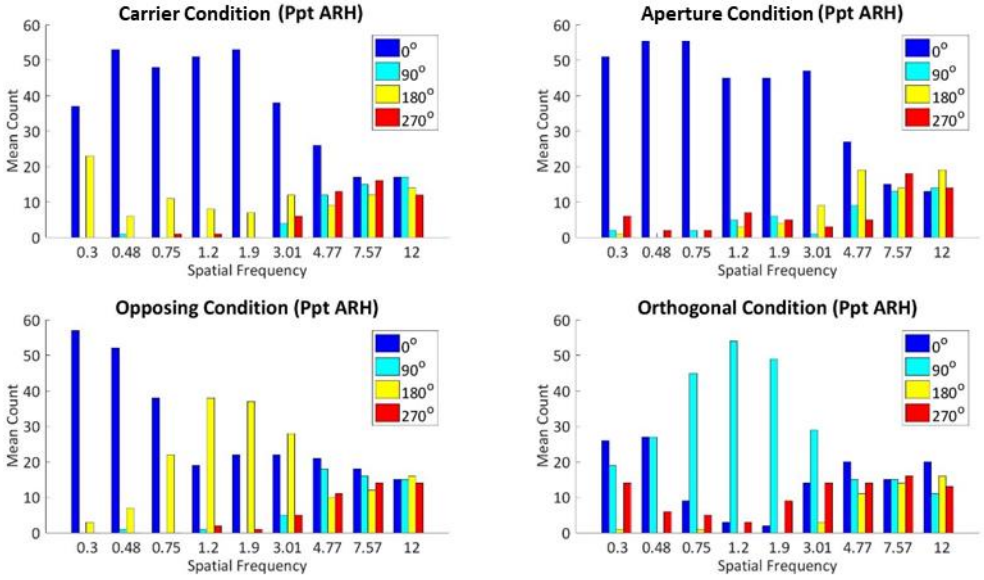


Figure 2.26. Reported motion trajectories for each of the 4 stimulus conditions, across all spatial frequencies tested, for participant ARH.

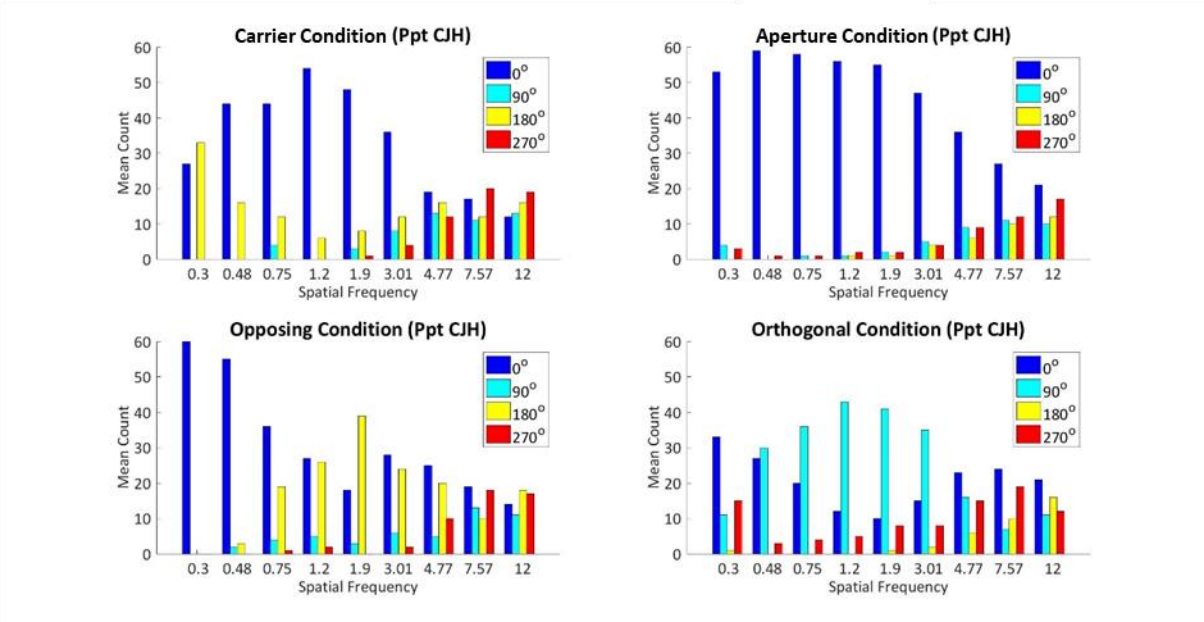


Figure 2.27. Reported motion trajectories for each of the 4 stimulus conditions, across all spatial frequencies tested, for participant CJH.

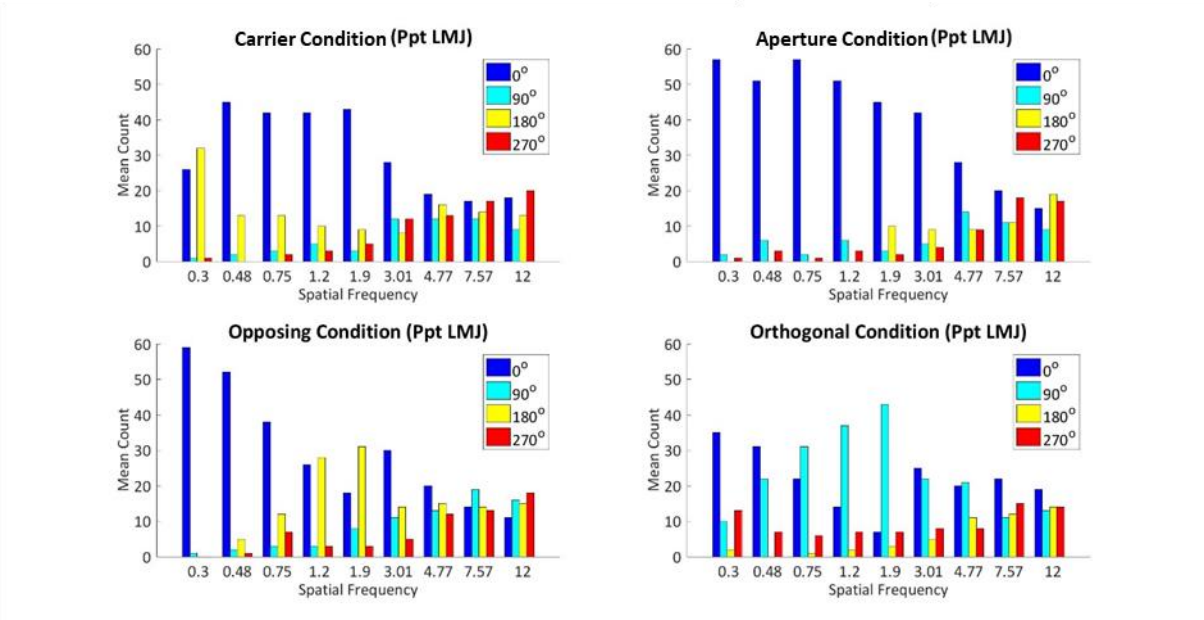


Figure 2.28. Reported motion trajectories for each of the 4 stimulus conditions, across all spatial frequencies tested, for participant LMJ.

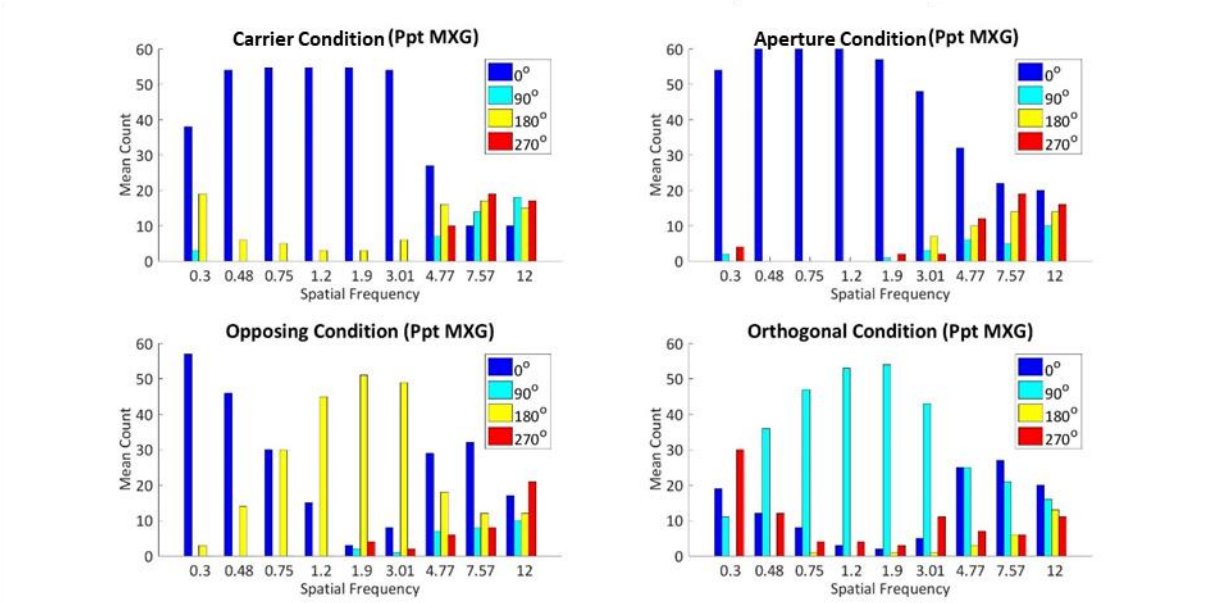


Figure 2.29. Reported motion trajectories for each of the 4 stimulus conditions, across all spatial frequencies tested, for participant MXG.

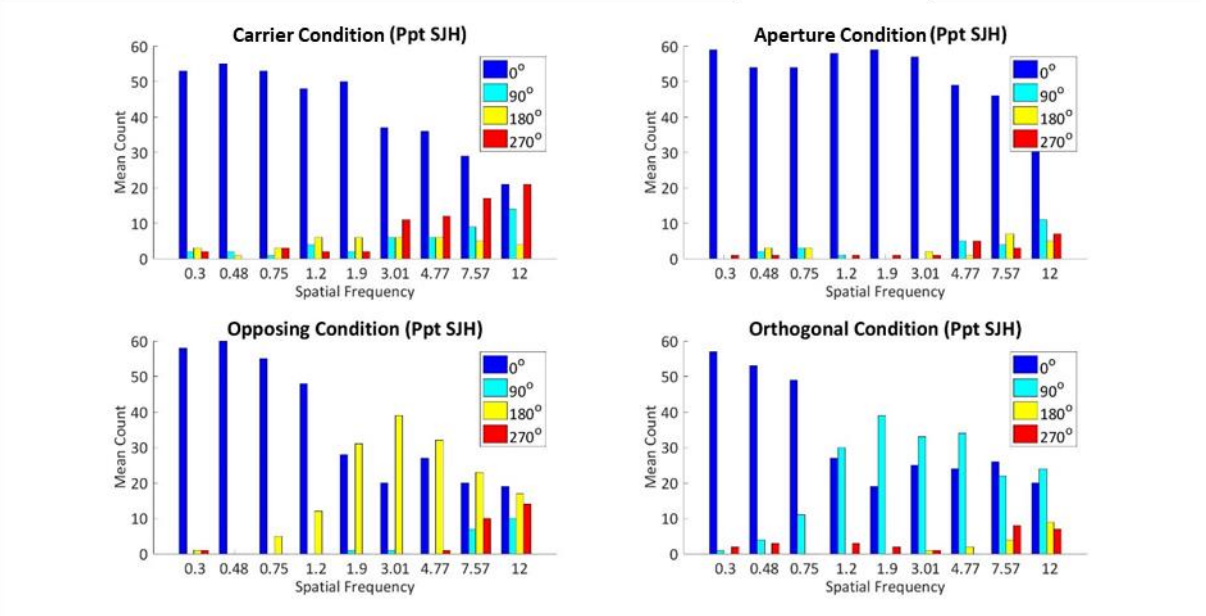


Figure 2.30. Reported motion trajectories for each of the 4 stimulus conditions, across all spatial frequencies tested, for participant SJH.

2.6.1.2 Comparison of trajectories in the Carrier and Aperture Motion conditions

There was a significant difference in the distribution of the trajectories across the four directions when compared between the Carrier and Aperture conditions, at each of the 9 spatial frequencies that stimuli were displayed at.

I explored this result further to compare the distribution of the four trajectories at each of the 9 spatial frequencies, for the Carrier and Aperture conditions, by running a 4 (direction) x 9 (spatial frequency) chi squared test of independence using Bonferroni corrected significance level (0.0014) of the adjusted standardised residuals. These can be seen in Table 2.12 where a negative adjusted residual indicates more trials reported in that direction in the Aperture condition, and a positive adjusted residual indicates more trials reported in that direction in the Carrier condition. There were significantly more trials reported in the 0^o direction in the Aper condition at 7 of the 9 spatial frequencies ($p \leq .003$), but no significant difference at 1.2 cpd and 1.9 cpd ($p > 0.083$). The greater number of trials reported in the 180^o direction in the Carrier condition was significant for the lowest four spatial frequencies, ($p < .001$).

Table 2.12 Table displaying the 4 (direction) x 9 (spatial frequency) chi squared test of independence, comparing the distribution of trajectories between the Carrier and Aperture conditions.

Spatial Frequency, cpd	Trajectory, in degrees	Adjusted residual (z score)	Chi Square Value	*p* value
.30	0	-7.80	60.80	<.0001*
.30	90	-.70	.50	.4795
.30	180	10.10	102.90	<.0001*
.30	270	-1.00	.90	.3428
.48	0	-4.20	17.90	<.0001*
.48	90	.20	.00	1.000
.48	180	5.70	32.30	<.0001*
.48	270	-1.00	1.00	.3173
.75	0	-4.50	19.90	<.0001*
.75	90	.50	.20	.6547
.75	180	5.10	26.30	<.0001*
.75	270	.70	.50	.4795

1.20	0	-1.70	3.00	.0833
1.20	90	.00	.00	1.000
1.20	180	4.60	20.90	<.0001*
1.20	270	-2.30	5.40	.0201
1.90	0	.00	.00	1.000
1.90	90	-2.70	7.50	.0062
1.90	180	3.10	9.90	.0017
1.90	270	-1.70	2.80	.0943
3.01	0	-5.40	29.10	<.0001*
3.01	90	2.70	7.40	.0065
3.01	180	2.10	4.50	.0339
3.01	270	3.60	13.10	.0003*
4.77	0	-3.80	14.50	.0001*
4.77	90	1.80	3.10	.0783
4.77	180	1.30	1.80	.1797
4.77	270	2.00	3.90	.0483
7.57	0	-3.70	13.80	.0002*
7.57	90	1.80	3.40	.0652
7.57	180	.50	.20	.6547
7.57	270	2.10	4.20	.0404
12.00	0	-3.60	13.00	.0003*
12.00	90	2.20	4.90	.0269
12.00	180	-.20	.00	1.0000
12.00	270	2.00	3.80	.0513
Total	0	-10.70	114.00	<.0001*
Total	90	2.70	7.10	.0077
Total	180	10.00	99.50	<.0001*
Total	270	2.70	7.30	.0069

2.6.1.3 Comparison of trajectories in the Opposing and Orthogonal Motion conditions

I compared the distribution of the re-coded four directions at each of the 9 spatial frequencies, for the Opposing and Orthogonal Motion conditions, by running a 2 (condition) x 4 (direction) x 9 (spatial frequency) chi squared test of independence using Bonferroni corrected significance level (0.0014) of the adjusted standardised residuals. These are reported in Table 2.13. There were significantly more trials reported in the position cue direction in the Opposing Motion condition than the Orthogonal motion condition for the lowest 5 spatial frequencies (p <.001). There were significantly more trials reported in the phase cue direction for the Orthogonal Motion condition than the Opposing Motion condition for the lowest 5 spatial frequencies, (p < .001*).* There were significantly more trials reported in the Error 1 direction in the Orthogonal Motion condition than the Opposing Motion condition for 0.3 cpd, 0.48 cpd and 1.2 cpd (p < .001).

Table 2.13 Table displaying the 4 (direction) x 9 (spatial frequency) chi squared test of independence, comparing the distribution of re-coded trajectories between the Opposing and Orthogonal Motion conditions.

Spatial Frequency, cpd	Re-Coded Direction	Adjusted residual (z score)	Chi Square Value	*p* value
.30	Position	12.60	158.76	<.0001*
.30	Phase	-6.70	44.89	<.0001*
.30	Error 1	-9.20	84.64	<.0001*
.30	Error 2	-2.90	8.41	.0037
.48	Position	11.20	125.44	<.0001*
.48	Phase	-9.00	81.00	<.0001*
.48	Error 1	-4.50	20.25	<.0001*
.48	Error 2	-2.30	5.29	.0214
.75	Position	7.90	62.41	<.0001*
.75	Phase	-7.20	51.84	<.0001*
.75	Error 1	-1.50	2.25	.1336
.75	Error 2	-.20	.04	.8415
1.20	Position	7.30	53.29	<.0001*
1.20	Phase	-5.30	28.09	<.0001*

1.20	Error 1	-3.90	15.21	.0001*
1.20	Error 2	.80	.64	.4237
1.90	Position	5.80	33.64	<.0001*
1.90	Phase	-4.30	18.49	<.0001*
1.90	Error 1	-2.10	4.41	.0357
1.90	Error 2	1.00	1.00	.3173
3.01	Position	3.00	9.00	.0027
3.01	Phase	-1.30	1.69	.1936
3.01	Error 1	-2.70	7.29	.0069
3.01	Error 2	.50	.25	.6171
4.77	Position	.50	.25	.6171
4.77	Phase	-1.10	1.21	.2713
4.77	Error 1	-.20	.04	.8415
4.77	Error 2	1.10	1.21	.2713
7.57	Position	-1.60	2.56	.1096
7.57	Phase	-.20	.04	.8415
7.57	Error 1	-.20	.04	.8415
7.57	Error 2	2.30	5.29	.0214
12.00	Position	-2.10	4.41	.0357
12.00	Phase	.20	.04	.8415
12.00	Error 1	.60	.36	.5485
12.00	Error 2	1.50	2.25	.1336
Total	Position	13.60	184.96	<.0001*
Total	Phase	-10.50	110.25	<.0001*
Total	Error 1	-6.80	46.24	<.0001*
Total	Error 2	1.90	3.61	.0574

2.6.2 Appendix for Section 2.4.1

Displayed in this section are the figures illustrating the nME analysis described in Section 2.4.1, where I ran a Fourier analysis on 3 dimensional constructions of five stimuli; the four stimuli used in the experiment in Section 2.3, and Chapter 4, and the Aperture + Rotation stimulus, used in Chapter 3. The analysis was run on 32-frame stimulus displays, with horizontal motion; the experiments were run using 2-frame stimuli with vertical motion (where in conflicting cue stimuli the vertical motion refers to the position cue).

In the Carrier condition, the vertical sinusoid grating is weighted by the flat-topped 2D Gaussian function (the aperture). The aperture remains stationary, but the sinusoid grating is displaced by 90° on each frame.

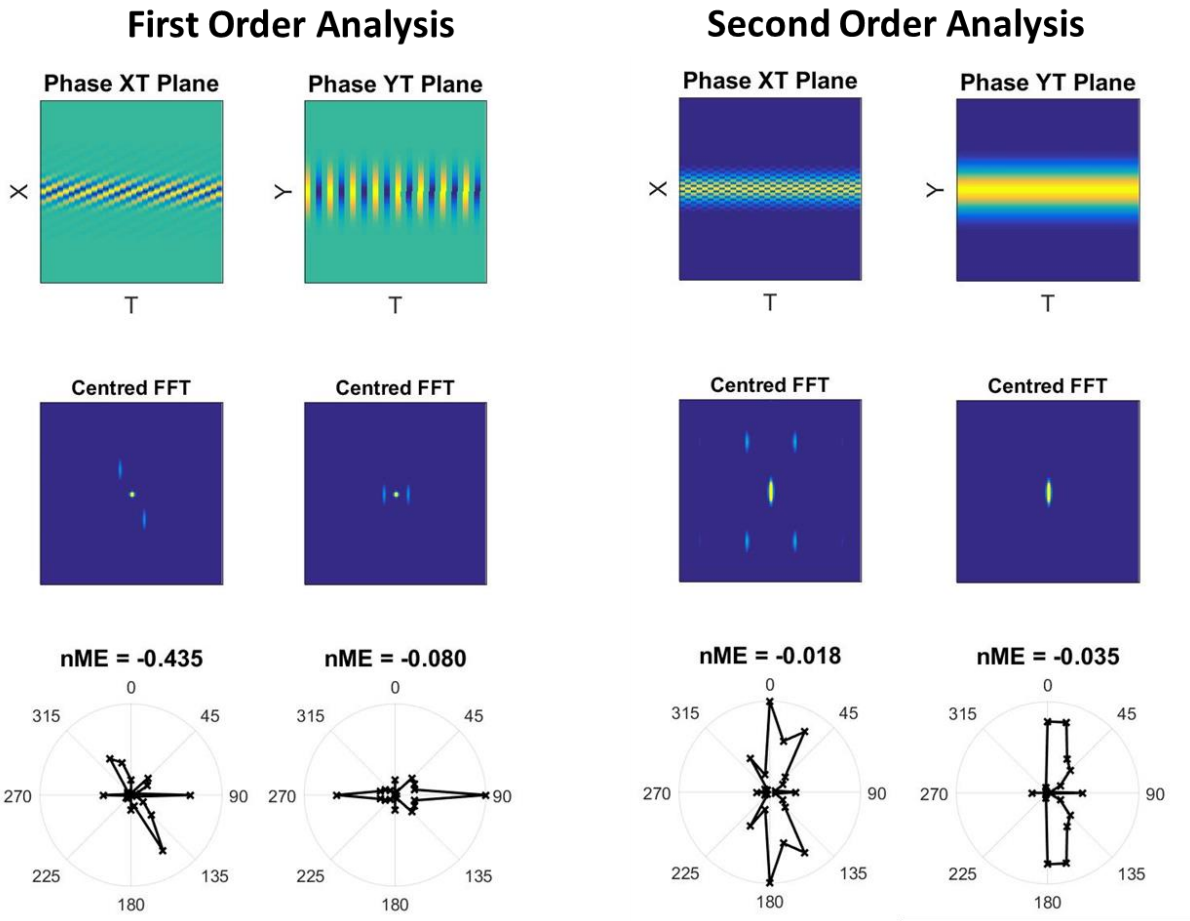


Figure 2.31 First and Second order analysis of the Carrier stimulus.

In the Aperture condition, the horizontal sinusoid grating is weighted by the 2D flat-topped Gaussian function (aperture). The sinusoid grating remains stationary on each frame, and the aperture is displaced horizontally on each frame, by a distance equivalent to one quarter of the grating wavelength.

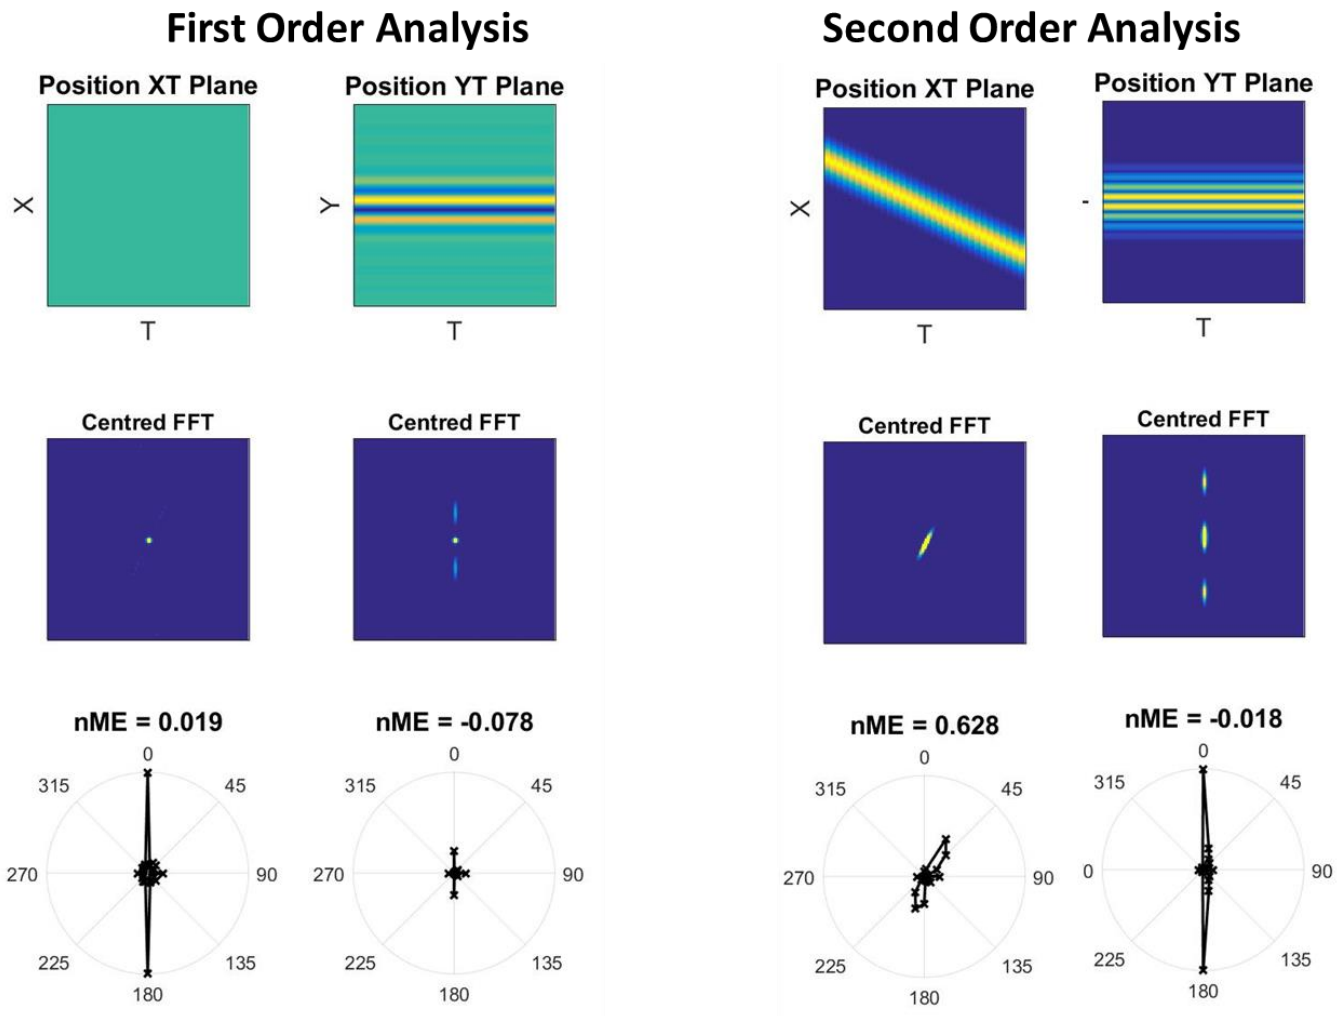


Figure 2.32. First and Second order analysis of the Aperture stimulus.

In the 'Opposing Motion' condition the vertical sinusoid grating is weighted by the 2D flat-topped Gaussian function (aperture). On each frame the aperture is displaced horizontally leftwards, and the grating is displaced by 90°.

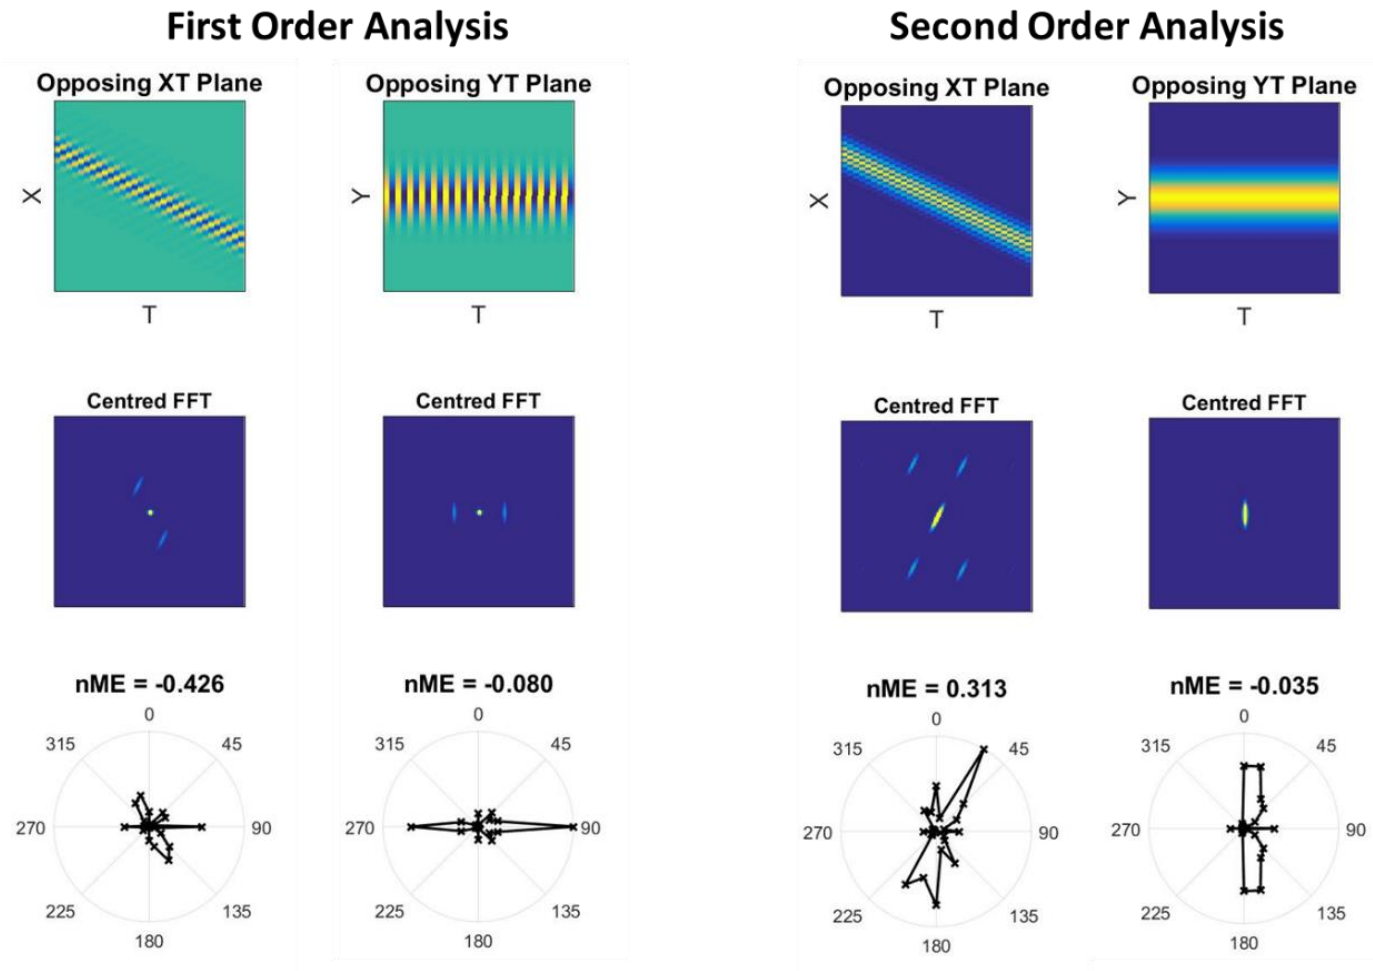


Figure 2.33 First and Second order analysis of the Opposing Motion stimulus.

In the 'Orthogonal Motion' condition the horizontal sinusoid grating is weighted by the flat-topped 2D Gaussian function (aperture). On each frame the aperture is displaced horizontally leftwards, and the grating is displaced by 90°.

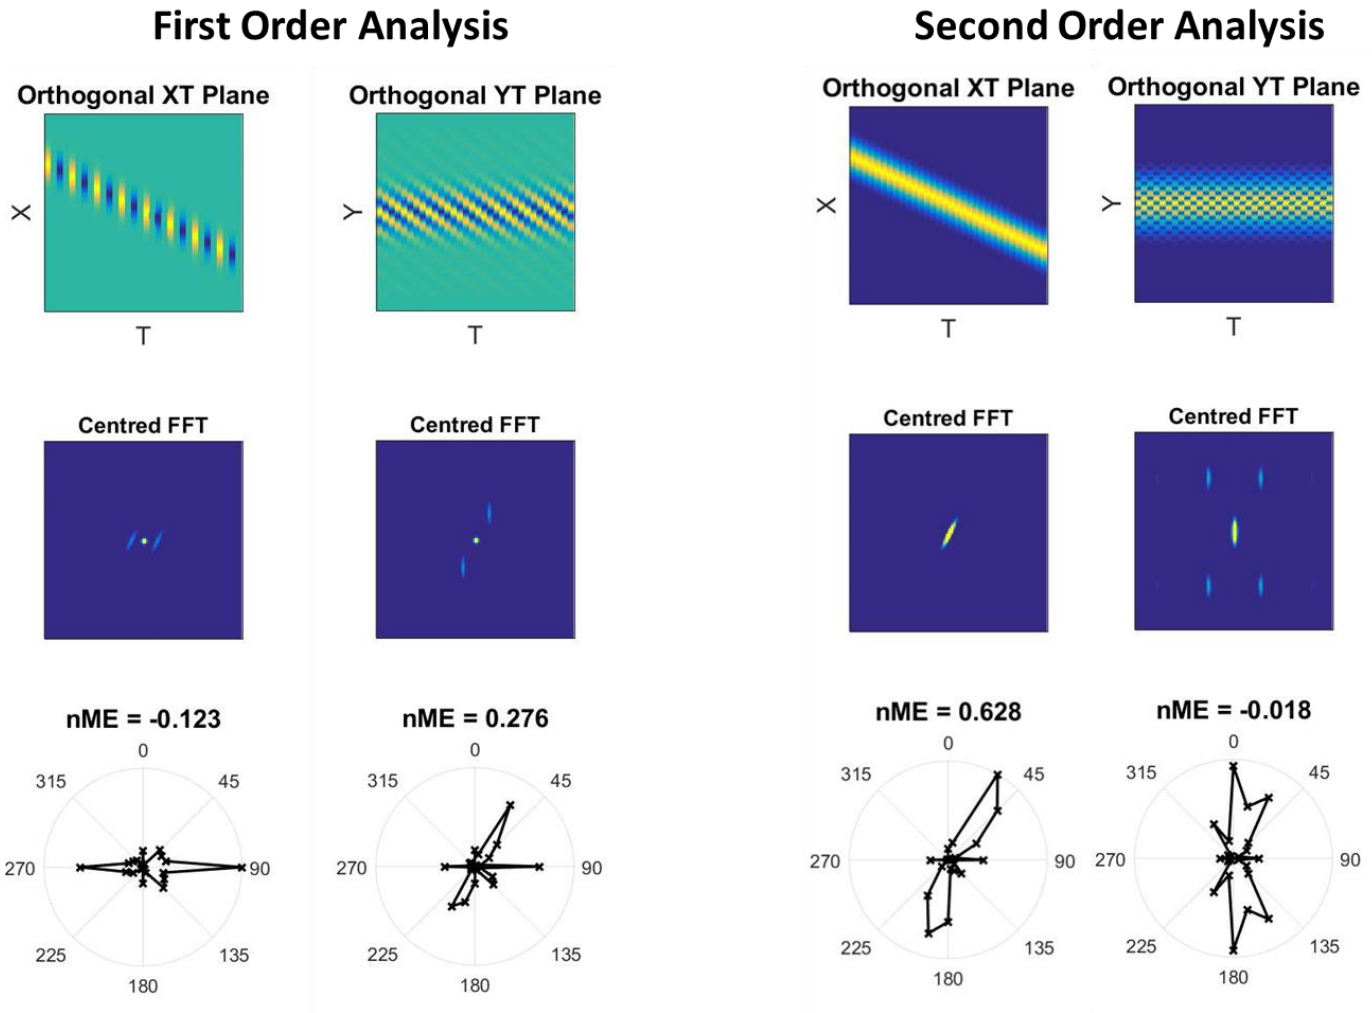


Figure 2.34. First and Second order analysis of the Orthogonal Motion stimulus.

In the Aperture+Rotation condition the horizontal sinusoid grating is weighted by the flat-topped 2D Gaussian function (aperture). On each frame the aperture is displaced horizontally leftwards, and the grating is rotated by 90°.

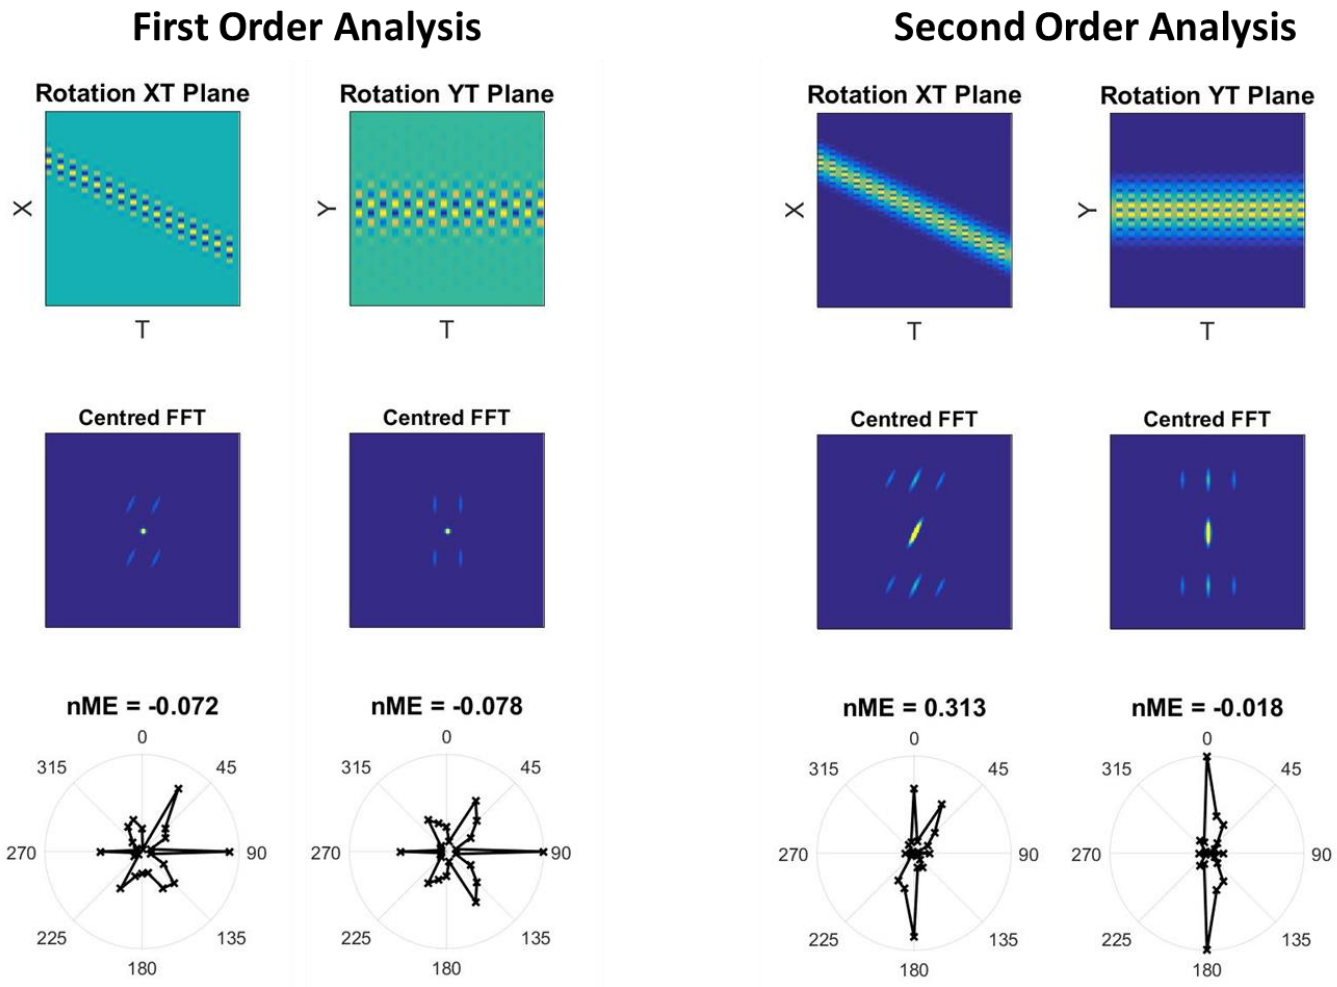


Figure 2.35. First and Second order analysis of the Aperture+Rotation stimulus.

2.6.3 Statistical analysis within the thesis

Choice of analysis method

Across the thesis a primary method of analysis that I employ is "non-linear mixed effects regression" using the nlme package in R (Pinheiro *et al.*, 2013). Non-linear mixed effects regression, or modelling, describes the relationship between a continuous independent and dependent variable, such as spatial frequency and direction discrimination performance. A linear relationship between the variables refers to equations where parameters are not raised to the power of x, whereas non-linear relationships between variables involve parameters raised to the power x. The nlme package fits the given expression to the data, using the Gauss–Newton algorithm to minimise the sum of squares, performing a non-linear version of the least squares fitting which is used in linear models. In this thesis I analyse data with various non-linear relationships, such as the four-parameter logistic model in Chapter 5 to describe reported direction of a reverse phi grating, and the three parameter asymptotic regression model in Chapter 2 to describe the asymptotic performance of the MF3F and MFG when blank IFIs are inserted, and the sum of Gaussian models to describe the frequency distribution of reported trajectories in Chapter 4. The nlme package is also suitable for use with linear models, such as the second and third order polynomials used to describe the relationship between direction discrimination and spatial frequency for phase and position cues in Chapter 3.

This method is appropriate for the data contained within this thesis as most experiments in the thesis the different experimental conditions can be described by the same equation, where a significant difference between the conditions would produce a significant difference in the coefficients of the parameters describing experimental condition. If two conditions within an experiment produce data with different forms, such as the frequency distributions of trajectories for the Opposing and Orthogonal Motion conditions in Chapter 4, then data can be analysed separately, or possibly rearranged into similar formats: I took both these approaches for different analyses of this data.

In this thesis an additional complication for the analysis is the use of repeated measures: in each experiment there are multiple stimulus conditions, such as different types of grating, and participants will take part in both conditions. To ensure the independence of residuals across the different levels of the model it is essential to correctly nest these properties within the model so that an error term can vary separately across the stimulus condition and for each participant.

The coefficients for the parameters in the equation are able to vary across various nested random effects, such as stimulus condition and participant identity, enabling modifications of the general model fit for each condition. The nlme package fits the given equation to the data, using the Gauss–

Newton algorithm to minimise the sum of squares, performing a non-linear version of the least squares fitting which is used in linear models. The use of an equation to describe the relationship between the variables enables us to describe more fully the difference between stimulus conditions, and depending on the equation chosen the parameters can be related to underlying processes.

Analysis procedure

The analysis procedure is consistent across all applications of the nlme package across my thesis.

Choosing the initial model

I begin with organising my data, if necessary transforming the data e.g. converting the number of trials reported in each interval into 'proportion correct' or calculating a frequency distribution of reported trajectories. I then plot the data to determine the overall relationship between variables, looking for the most parsimonious equation that contained parameters with explanatory power, such as asymptotes and thresholds in the four parameter logistic function fitted in Chapter 5.2.

Once I have determined which equation best describes the data I run an initial analysis to fit the equation to the data. I then verify that the equation does describe the data well, by using the nlme package to see whether all the equation parameters contribute significantly to fitted model. This stage can be used to refine the model, producing simpler models, or could reveal that the data are poorly fitted by this model, and a new base equation should be chosen.

Analysing the data

Once the initial model has been chosen, the data are then grouped by the fixed factors, such as experimental condition, eccentricity and participant identity. This grouping is critical as this allows random effects to vary across these levels separately, as well as the fixed effects, which are critical in determining whether there is a significant difference between conditions. Once grouped, the data can be fitted by the model.

The nlme package will produce a model with overall coefficients for each of the parameters, describing the overall relationship of the independent and dependent variables. These parameters have adjusted values at the levels of the fixed effects, which enables the model to describe the data more accurately. Each level also has a separate random error.

Assessing whether there is a significant difference between models

We can determine whether there is a significant difference between stimulus conditions by producing full and reduced models of the data, and comparing the goodness of fit between full and reduced models. The full model is the model fitted to the data when all of the fixed effects have

been labelled, and the reduced model refers to a model where some fixed effects, such as stimulus condition, have not been labelled. Reduced models therefore make no distinction between data points between the unlabelled conditions. The reduced model produced will have the same overall equation and similar coefficients at the overall level as the full model, but won't have the specified modifications to the parameters at the un-labelled level. If there is a significantly better goodness of fit of the model when stimulus condition is included as a fixed effect, then this demonstrates a significant difference between conditions. Goodness of fit can be assessed by calculating the difference in Akaike Information Criterion (AIC) or the Log Likelihood Ratio between full and reduced models.

The AIC is calculated using the models maximum likelihood estimation as the model of goodness of fit (the natural logarithm of the maximum likelihood multiplied by -2), and adds a penalty for the number of parameters (twice the number of parameters). This penalisation for the number of parameters means that the AIC can compare between models fitted to the same dataset, but with different forms, and that conceptually the AIC balances goodness of fit with model simplicity. When comparing the AIC value between full and reduced models, the lower AIC value indicates better fit, and a model can be considered significantly better if the AIC value is reduced by at least 2 (Anderson & Burnham, 2004).

The Log Likelihood Ratio can also be used to compare whether there is a significant difference between conditions in my thesis, as the full and reduced models fitted to my experimental data can be considered nested. If there is not a significant variance between conditions (the level removed in the reduced model) then the random effects across the parameters for that level should be small. The Log Likelihood ratio assess whether the difference between the full and reduced model is substantially larger expected from solely sampling error, or whether the difference reflects a significant difference between the conditions included in the full model. The Log Likelihood Ratio test is included in the nlme package, with effect size and significance values.

Chapter 3: Psychophysical dissociation of phase and position cues

3.1 Introduction

Motion in the visual array can be described in terms of the change in phase of the Fourier components or by the change in spatial position of image features. These two descriptions of image motion are mathematically equivalent. However, stimuli can be produced where these two descriptions produce conflicting predictions about motion direction, as in the Missing Fundamental Grating (MFG) and reverse phi stimuli. Stimuli can also be produced where one of these descriptions is rendered incoherent. The change in position of features in a drifting sine wave grating can be obscured by introducing a pedestal (Van Santen & Sperling, 1984), and complex stimuli can be produced where there is no coherent motion produced from the change in phase of Fourier components of the image, but the changing position of salient features can be tracked across frames to produce a coherent motion percept (Lu & Sperling, 1995b).

Models of human motion perception have proposed different mechanisms by which these different motion cues might be extracted, relying on different descriptions of the visual input. If different mechanism are involved in separately processing changes in phase and changes in position, these may require different neural circuitry, and therefore differences may be apparent under psychophysical investigation. A challenge in this approach is the difficulty of producing stimuli that selectively contain the cues of interest, or in producing stimuli where the image characteristics of different stimuli are sufficiently similar that differences in psychophysical performance can be attributed to potential underlying mechanisms, and not differences in the stimuli used.

Motion perceived from the change in spatial location of features (position cues) is dependent on localising and identifying the stimulus (Cavanagh, 1990), and can occur over large blank inter-frame intervals (Georgeson & Harris, 1990), suggesting that the motion cue is not dependent on the temporal properties of the input falling within a distinct temporal window, just that there is a change in location with time. In contrast, motion perceived from the change in phase of Fourier components (phase cues) can be successfully discriminated in stimuli where motion is unattached to an object (Van Santen & Sperling, 1984; Wertheim, 1912; Steinman *et al.* 2000), and without a change in position (Flynn & Shapiro, 2018).

Reported motion is biased towards the phase cue direction for stimuli presented at low contrasts in conflicting cue stimuli produced in the Missing Fundamental Grating (MFG), where the phase and position cue have different periodicities (Georgeson & Harris, 1990), and in reverse phi stimuli

(Edwards & Nishida, 2004). Temporal frequency also biases reported motion, towards the phase cue at high temporal frequencies and in the position cue direction at low temporal frequencies, seen in Missing Fundamental Gratings (Georgeson & Harris, 1990), and in orthogonal Gabor stimuli, where the temporal frequency can be manipulated for the sine wave grating and Gaussian aperture separately (Heller *et al.*, 2021). These results are consistent with Van Santen & Sperling's demonstration in pedestals that motion discrimination, carried by the phase cue, was high across increasing temporal frequencies, but the perception of the 'wobble', the incoherent motion of the position cue, was abolished as temporal frequency increased (Van Santen & Sperling, 1984). There is also compelling support for a spatial frequency dissociation between the phase and position systems, where reported motion direction in conflicting cue stimuli is biased towards the position cue at low spatial frequencies, and the phase cue at high spatial frequencies. This is reported for reverse phi bar and grating stimuli viewed foveally where spatial frequency was manipulated using viewing distance (Chubb & Sperling, 1989; Benton *et al.*, 1997), and in discs, where motion was reported in the phase cue direction at small displacements between frames, and in the position cue direction at large displacements (Anstis & Rogers, 1975). I have shown a similar effect in two-frame Gabor stimuli with conflicting cues viewed peripherally, where in Gabors with the phase and position cues arranged in opposition or orthogonally, motion was reported in the position cue direction at low spatial frequencies and in the phase cue direction at higher spatial frequencies (Chapter 2).

The most compelling evidence for a psychophysical dissociation between phase and position cues comes from changes in perceived motion for conflicting cue stimuli between fovea and periphery. In reverse phi stimuli, a stimulus which is reported in the position cue direction at fovea is reported in the phase cue direction when viewed peripherally (grating; Chubb & Sperling, 1989; Anstis, Cavanagh & Mather, 1985: RDK; Edwards & Nishida, 2004). Whilst a Gabor stimulus cannot be as easily divided into its phase and position cue components, as phase shifts of the carrier produce a corresponding change in spatial position of the carrier features, the aperture displacement can be considered a 'position cue', providing that the Gabor is designed so that either the carrier remains stationary or the phase is displaced along a different orientation to the aperture displacement. This is the case in Gabors where the carrier is oriented in parallel with the direction of aperture displacement, and so any phase shifts will produce an orthogonal phase cue.

When viewed foveally, an orthogonal Gabor is perceived as moving in the aperture displacement direction (position cue) but when viewed in the periphery, the perceived trajectory of motion is reported at an intermediate trajectory between the direction of the aperture displacement and phase shift of the carrier (Tse & Heich, 2006; Shapiro *et al.*, 2010; Gurnsey & Biard, 2012).

This bias towards position cues at fovea, and towards phase cues in periphery across these different stimuli could be caused by effective change in spatial frequency sensitivity that occurs between the fovea and periphery. If the spatial frequency sensitivity functions for the phase and position systems are both displaced to lower spatial frequencies with eccentricity, as may be seen in the response to spatial frequencies measured by static and drifting gratings (Virsu & Rovamo, 1979), then a low spatial frequency stimulus, which produced a stronger response in the position system at fovea, may now produce a stronger response in the phase system when viewed peripherally. This difference in the change in spatial frequency sensitivity of the phase and position systems with eccentricity could also explain the change in perceived direction in conflicting cue Gabor stimuli. Drifting orthogonal Gabors are seen solely in the position cue direction (aperture displacement) at the fovea, and in the periphery they are reported at trajectories intermediate to the aperture displacement and the direction of the phase shift of the carrier. This could be explained by an increasing relative strength of the phase system in the periphery. Consistent with this interpretation, the reported motion in these conflicting cue Gabor stimuli cannot be scaled by spatial frequency to produce the same reported motion trajectory (Gurnsey & Biard, 2012).

Thus, evidence from different constructions of conflicting cue stimuli suggests a difference in the spatial frequency sensitivity of the phase and position systems, where the phase system is sensitive to motion at higher spatial frequencies than in the position system. The difference in reported motion trajectory between fovea and periphery cannot be explained through the mechanism of a single system, as demonstrated by failures to scale the reported motion direction in Gabors with orthogonally oriented carrier and aperture cues. At the fovea, motion is reported in the aperture cue direction, across all spatial frequencies, whilst in periphery the reported trajectory of motion is a diagonal trajectory intermediate to the carrier and aperture cue directions, demonstrating integration between these cued directions (Gurnsey & Biard, 2012). Therefore, no matter the size of the stimulus presented at fovea, the difference in reported motion will be different from reported motion in the periphery. This integration between motion cues in the periphery could be explained by a lack of integration at the fovea, or the retained presence of the integration mechanism, but with such a small contribution from the carrier cue, that there is no effect on reported direction. Alternatively, these results could be explained by separate phase and position systems, with differences in spatial frequency sensitivity that change differently across the retina, resulting in an increased bias towards the carrier cue in the periphery as the two spatial frequency functions increasingly diverge.

Testing spatial frequency sensitivity

One method of assessing the spatial frequency function of a system is to measure the contrast sensitivity function, where the amount of contrast needed for detection or discrimination of a stimulus is calculated across a range of spatial frequencies. This method is excellent for sinusoid grating stimuli which contain a single sine wave, and has successfully demonstrated the changing spatial frequency sensitivity to drifting and static gratings across eccentricity (Virsu & Rovamo, 1979). However, this method is less appropriate for stimuli where the motion cue is not determined from a single frequency grating. In this experiment I used Gabor stimuli where the phase cue is produced from displacement of the carrier (alongside position cues in the corresponding feature displacement), but the position cue derives from the from displacement of the Gaussian envelope, calculated through the displacement of the Gabor centroid (Toet & Koenderink, 1988). Therefore, the contrast sensitivity approach may not be appropriate for comparing between the two cues, since the displacement of the aperture is not perceived using the sinusoid grating of the stimulus.

An alternative method for mapping spatial frequency sensitivity in more complex stimuli is to display suprathreshold apparent motion stimuli and measure performance in terms of proportion correct, across a range of displacements of the stimulus between frames. This method has been used for apparent motion RDKs (Baker & Braddick, 1985), and phi motion stimuli between two single dots at two locations (Foster *et al.*, 1989), but can also be applied to grating stimuli, for example, where Levi *et al.* compared displacement thresholds for grating acuity and unreferenced motion (Levi *et al.*, 1984). An advantage of this method is that a displacement threshold, the displacement corresponding to a certain performance level, can then be plotted against eccentricity, and used to compare between different psychophysical stimuli (Foster *et al.*, 1989), or against physiological estimates (Baker & Braddick, 1985). Displacement thresholds can also be compared for data collected between different experiments, as in Levi *et al.* (Levi *et al.*, 1984), where the displacement thresholds calculated for unreferenced motion across eccentricity were compared to published values on grating and vernier acuity (Westheimer *et al.*, 1979, 1982) to propose a psychophysical dissociation between systems processing motion and position cues. Displacement thresholds measured across eccentricity have also been compared to physiological substrates (Cowey & Rolls, 1974; Levi *et al.*, 1985; Strasburger *et al.*, 2011), an approach I explore later in the thesis.

Testing phase and position cues

Ideally, this hypothesis would be tested using stimuli where coherent motion discrimination is reliant solely on the phase or the position cues in an image. This ideal stimulus set would also contain

relatively similar image characteristics for the 'pure' phase and position cues, so that psychophysical dissociations could be attributed to the presence of phase or position cue, and thus used to draw inferences about possible separable mechanisms for analysing these cues, rather than differences in performance being the result of image characteristics. One possibility for a 'phase' cue stimulus would be to use the pedestal stimulus of Van Santen & Sperling (1984), which contains coherent motion through a phase cue, but the changing position of features on each frame is incoherent. Various stimulus manipulations could be used to produce a putative counterpart to this stimulus, such as removing the pedestal to restore coherent motion between features, and using a long inter-frame interval to potentially remove the phase cues. This approach relies on the assumption that long-interframe intervals would be successful in silencing phase cues, but introduces other issues in attempting to match the temporal characteristics of the two stimuli.

Two frame Gabors offer a reasonable compromise in this search for stimuli with similar image characteristics that can be manipulated to produce phase and position cues separably. Displacing the aperture of a Gabor produces a position cue, and potentially without corresponding phase cues depending on the stimulus design. Changing the phase of the carrier produces a phase cue, albeit with a corresponding position cue due to the change in position of the features of the carrier. The Gabor stimulus is relatively simple in the spatial frequency domain compared to other stimuli designed to isolate position cues, such as the MF1F of Baro & Levinson (1981), or interleaved gratings defined by depth, motion-modulated motion, texture etc. (Lu & Sperling, 1995b). Coherent motion in Gabor stimuli can be produced in two-frame stimuli, unlike pedestals which require a full cycle of frames (the reciprocal of the displacement on each frame) to ensure that the 'wobble' of the features on each frame can't be used as a position cue for motion discrimination. Finally, the aperture and carrier of the Gabor can be manipulated separately, or within the same stimulus to produce conflicting cue stimuli, such as the Opposing and Orthogonal Gabors in Chapter 2.3. This means motion discrimination of the carrier and aperture cues can be tested in isolation, and pitted against each other to directly compare the relative strength of the cues. Testing the carrier and aperture cues in isolation or in conflict provides different information to determine whether psychophysical dissociations support a hypothesis of separate systems for processing phase and position cues.

In this chapter I test the hypothesis of a dissociation between the spatial frequency sensitivity functions for the phase and position systems which increases with eccentricity, by measuring direction discrimination for a set of two-frame Gabor stimuli (Figure 3.1), at the fovea and 3 peripheral locations.

The first construction of the Gabor stimul is referred to as the 'Carrier condition', and consists of a stationary Gaussian aperture and a sinusoid grating, or 'carrier', where the phase of the carrier is shifted by 90^{o} between frames. The motion in this stimulus can be described by the change in phase of the Fourier components of the carrier, but can equally be described by a change in the spatial location of the features of carrier e.g. the displacement of the peaks, troughs and zero-crossings. The second construction of Gabor stimuli is the 'Aperture' condition, produced from the displacement of the Gaussian window and simultaneous rotation of the sinusoid carrier. The carrier rotation means that there is no coherent change in phase of the Fourier components between frames, and additionally disrupts local luminance gradients between pixels across frames. The motion cue for the Aperture condition therefore comes from the displacement of the spatial position of the Gaussian window, and not from features of the carrier. The displacement of the Gaussian window is by a quarter of the wavelength of the fundamental frequency of the carrier, and therefore equivalent in distance to the distance travelled by the peaks and troughs of the carrier in the Carrier condition, produced from the phase shift of the carrier. Therefore the distance moved by features in the Carrier condition is equal to the distance moved by the gaussian aperture in the Aperture condition.

A third condition was included in this experiment, and the data are analysed separately and presented in Chapter 4. This is the 'Opposing Motion' condition, produced from a displacement of the Gaussian window and phase shift of the sinusoid carrier in opposite directions, as introduced previously in Chapter 2.3 as the 'Opposing Motion' condition. Again, the displacement of the Gaussian window is set to the equivalent distance of a 90^{o} phase shift of the carrier frequency. The purpose of the Opposing Motion condition was to explore reported motion direction across spatial frequency and eccentricity in conflicting cue Gabor stimuli, whereas the single cue conditions presented in this chapter are used to characterise the spatial frequency sensitivity to carrier and aperture cues in isolation, and to compare the changes in displacement thresholds across eccentricity.

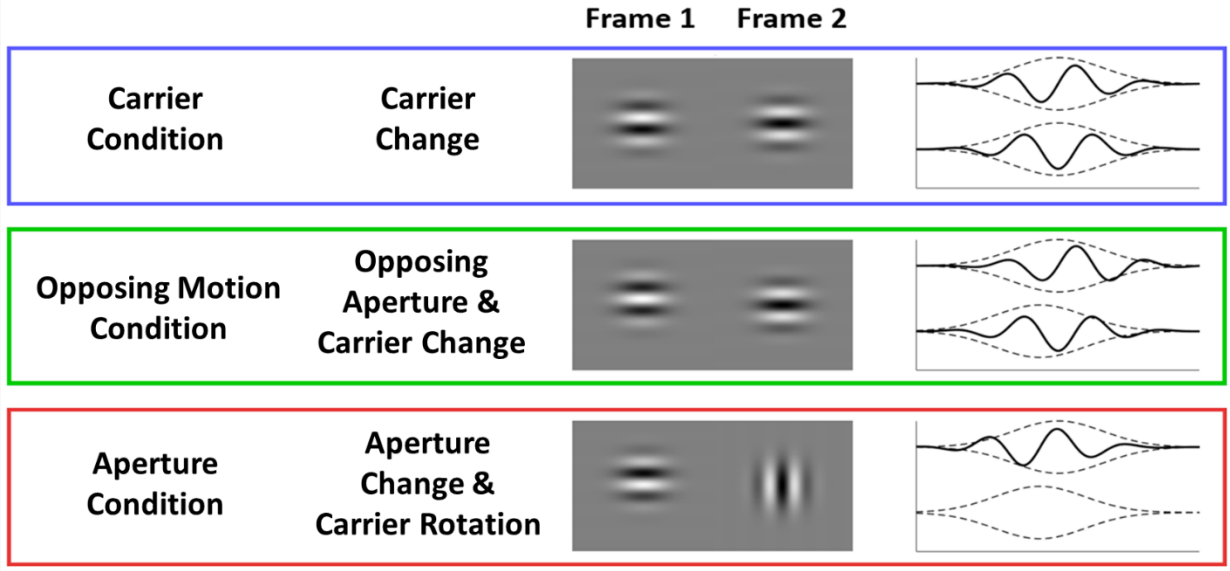


Figure 3.1. Schematic demonstrating the three experimental conditions. Column 1 is the condition name, column 2 describes the motion cues within the condition, column 3 depicts Frame 1 and Frame 2 of the stimulus, as they appeared to the participant, and column 4 displays the cross-sections taken horizontally across the grating, with the Gaussian window profile displayed as a dashed line. This pair of frames results in an apparent motion percept of the stimulus moving in the horizontal plane; leftward in the Carrier and Aperture conditions, and in the Opposing Motion condition either left, if the aperture cue dominates motion percept, or right, if the carrier cue dominates the motion percept. In this experiment stimuli moved in the vertical plane, as shown in column 3, whilst column 4 displays the stimuli cross-sections horizontally.

The aim of this experiment was to investigate whether there were differences in motion discrimination resulting from motion produced from a change in carrier phase (carrier condition), or a change in aperture location (aperture condition). Motion discrimination for all three classes of stimuli was measured across spatial frequency, and at four retinal locations, to see whether there was a difference in sensitivity both across spatial frequency. For the Carrier and Aperture conditions, displacement thresholds were interpolated for 75% proportion correct, and plotted against eccentricity.

3.2 Method

Ethics

This study was approved by the Medical Sciences Inter-divisional Research Ethics Committee (IDREC) at the University of Oxford in accordance with the procedures laid down by the University for Ethical Approval of all research involving human participants, under reference R51336/RE001.

Participants

Seven participants (three male) aged 21-57 years old, with normal or corrected-to-normal vision.

Apparatus

Stimuli were presented on a CRT monitor (Mitsubishi Diamond Pro 2070, 1600 × 1200 pixels), and gamma corrected using the ColorCAL MK II colorimeter (Cambridge Research Systems), with a maximum luminance of 81.33 cd/m^2. Viewing distance was set at 180 cm from the monitor for the conditions viewed foveally and at 5 degrees eccentricity and at 110 cm for conditions viewed at 10 and 15 degrees eccentricity, and maintained with a chin rest. Viewing was monocular, with participants' left eyes covered by an opaque eyepatch, and fixation was monitored using the EyeLink 1000.

Stimuli

Stimuli were constructed in MATLAB from sine wave gratings (carrier) weighted by a 2D Gaussian window, with standard deviation of 2.5 cycles of the grating wavelength. This sets the ratio of Gaussian envelope size to grating frequency, so that the same number of cycles are displayed for each Gabor, regardless of spatial frequency. The 'Carrier' condition was produced by phase-shifting the carrier by a quarter of a cycle, the 'Aperture' condition by rotating the carrier by 90 or 270 degrees, and displacing the envelope by distance equal to quarter of a cycle. The 'Opposing Motion' condition (results reported in Chapter 4) was produced by phase-shifting the carrier by a quarter of a cycle and. displacing the envelope by distance equal to quarter of a cycle. The initial phase of the carrier for each stimulus was randomly assigned on each trial, and stimuli were presented at maximum contrast. The spatial frequency of the Gabors ranged from 0.3° to 32°, with 7 spatial frequencies presented for the conditions viewed at fovea and 5 degrees eccentricity , and 6 spatial frequencies presented for the conditions viewed at 10 or 15 degrees eccentricity.

In this experiment Gabors were displayed in a two-frame apparent motion procedure which, as demonstrated in Chapter 2.3, produces a coherent motion percept when presented in the periphery for both phase and position cues, and minimises the difference in the change of location between the position and phase cue stimuli. I displayed each frame for 250 ms, as a balance between the high temporal frequency bias for carrier cues, and low temporal frequency bias for aperture cues seen in Gabor stimuli (Heller *et al.*, 2021).

The normalised Motion Energy (nME) analysis, described in Chapter 2.4, has been used to assess the Fourier components in the Carrier and Aperture conditions. There was a significant nME value for

vertical motion in the first order analysis for the Carrier condition, but no coherent nME value for any motion direction in the second order analysis, where the same Fourier analysis is applied to the stimulus, following a full-wave rectification. In the Aperture condition, the opposite pattern of results was found, where there was no coherent motion in the first order analysis, but following a full-wave rectification, there was a significant nME signal for vertical motion (Table 3.1). The nME analysis is designed to characterise the Fourier components within the stimuli, and not the 'position' cues, which could be defined by a vast range of image characteristics, and made explicit by a range of different image transformations. The nME used in this thesis uses a full-wave rectification to render the position cue of the Gaussian aperture explicit in the Aperture condition, but this transformation doesn't reveal the change in position of the carrier features in the Carrier condition. Therefore whilst the nME does show a double dissociation in the nME signals for the first and second order analysis, this can only be used to verify the lack of phase cues in the Aperture condition, and does not indicate a lack of position cues in the Carrier condition.

Table 3.1. Table displaying the nME values calculated in Chapter 2.4 for 32-frame renderings of the Gabor stimuli where motion is produced from a change in the Carrier, both the Carrier & Aperture in opposing directions, or Aperture, that are used in the experiment described in this chapter.

Stimulus Condition	First Order		Second Order	
	XT	YT	XT	YT
Carrier	-0.0798	**0.2761**	-0.0355	-0.0177
Opposing Motion	-0.0796	**-0.4256**	-0.0355	**0.3127**
Aperture	-0.0775	-0.0724	-0.0177	**0.3133**

Procedure

I used a two-alternative forced choice procedure on each trial, where the participant was asked to report 'in which interval did the stimulus move upwards?' using the keyboard. Each trial consisted of two intervals, each containing two frames of the apparent motion stimuli with counterbalanced presentation across the two intervals; AB, or BA, where AB corresponds to an upward apparent motion percept, and BA to a downward apparent motion percept. Each trial lasted 1000 ms (250 ms per stimulus frame, and 500 ms inter-trial blank screen), and was followed by a blank screen of mean luminance, where the participant reported the interval in which they had perceived upwards motion, and initiated the next trial using the spacebar.

At each eccentricity, participants viewed stimuli in three conditions, 'Carrier, 'Opposing Motion' and 'Aperture' (see Figure 3.1), at 7 spatial frequencies per condition at fovea and 5 degrees eccentricity, and at 6 spatial frequencies at 10 and 15 degrees eccentricity, with 45 trials per condition. All stimulus conditions and spatial frequencies were interleaved, producing 945 (or 810) trials per eccentricity.

Participants completed the experiment across 4 sessions, which each lasted between 90-120 minutes, with one eccentricity tested during each session. Participants started with foveal viewing, and the eccentricity of the viewing location was increased by 5 degrees each testing session.

Fixation was monitored using the Eyelink 1000 recording at 500 Hz, with an initial calibration session at the start of the experiment, and recalibrations during the experiment every 15 minutes. Participants were also able to abort trials where they felt that their fixation had drifted, and these trials were shuffled into the list of future trials.

3.3 Results

Removal of trials/Data pre-processing

The data consists of participants' responses to the three conditions (Carrier, Opposing Motion and Aperture), presented at the fovea and 5, 10 and 15 degrees eccentricity, with ~45 trials per condition. The data for individual participants is presented in the appendix of this chapter, and the Opposing Motion condition is analysed separately and reported in Chapter 4.2.

Fixation during the experiment was measured with the EyeLink 1000. I removed trials with poor fixation, defined as eyes drifting more than 1 degree of visual angle from the fixation cross, during the trial. This reduced the total number of trials per participant. I compared the number of trials for each participant across condition and eccentricity; spatial frequency within the analysis was not included because a different range of spatial frequencies were used across each eccentricity. The range of spatial frequencies was chosen to produce a similar range of performance at each eccentricity.

I ran a 3 (condition) x 4 (eccentricity) repeated measures ANOVA on the mean number of trials retained following removal for loss of fixation. There was no significant main effect of condition (p = .805), or of eccentricity (p = .419) and no significant interaction between them (p = .789). There was a significant difference between participants ($F(1,6) = 531.005$, $p < .001$, $\eta_p^2 = .989$).

Comparison of motion discrimination in Carrier and Aperture conditions

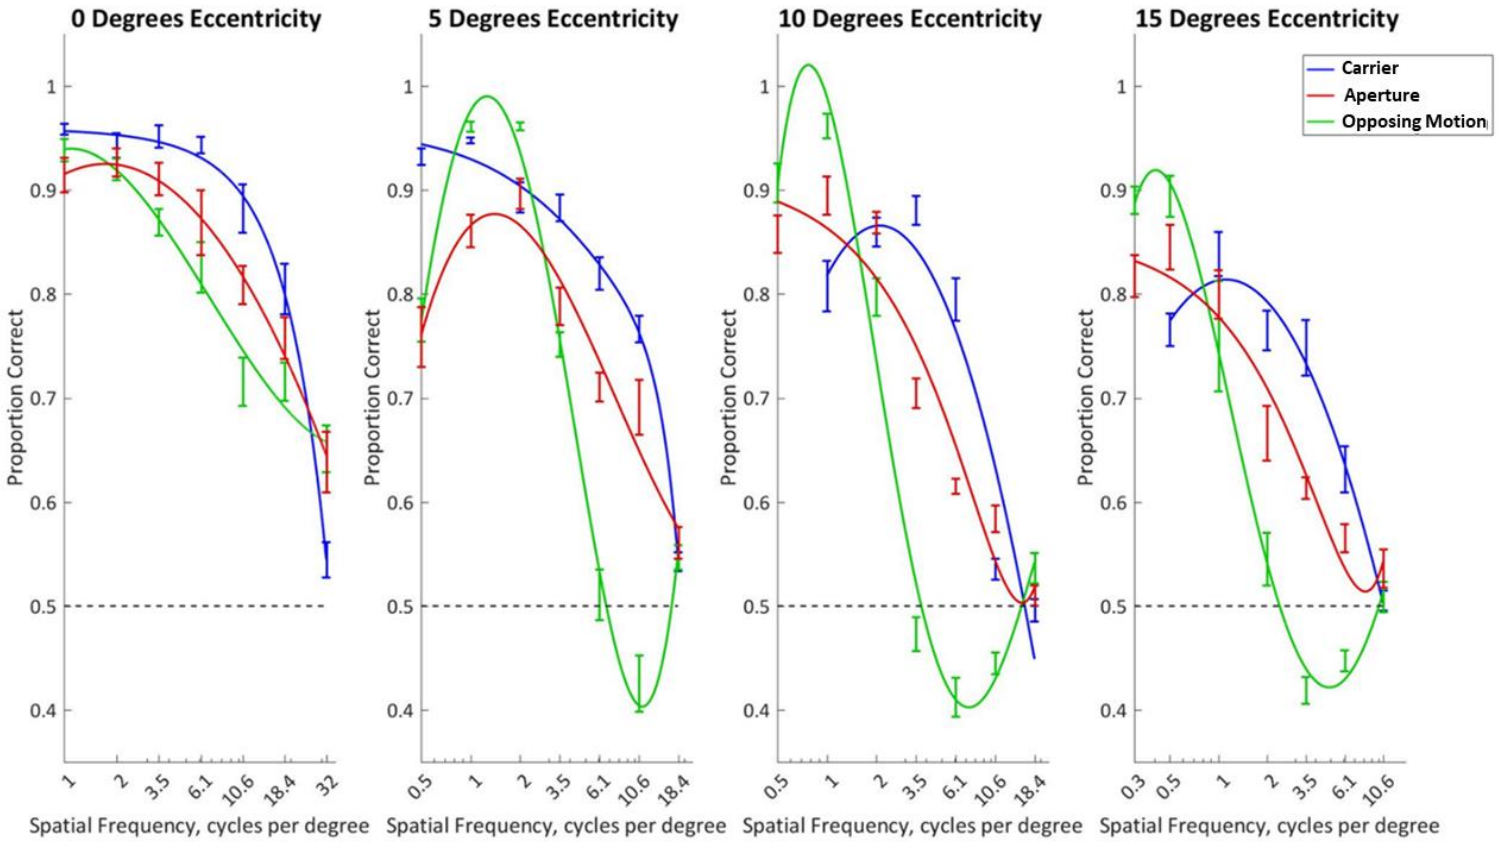


Figure 3.2. Figure plotting the mean proportion correct in the Carrier and Aperture conditions by spatial frequency, at the four eccentric locations tested. The blue curve plots the data collected in the Carrier condition, the red curve the data collected in the Aperture condition, and the green curve is the Opposing Motion data, which is reported in Chapter 4.2. The bars represent the standard error of the mean, displaying the range of participant responses.

The Carrier and Aperture conditions data were collected across a range of spatial frequencies, at 4 eccentricities; fovea, and 5, 10 and 15 degrees eccentricity. Motion discrimination performance was operationalised as 'proportion correct'; the proportion of trials where the correct interval was reported as containing upward motion. Proportion correct plotted against the spatial frequency of the stimuli was best described by a quadratic function at each eccentricity, when data was modelled including participant identity (data were displayed using second and third order polynomial curves of best fit in Figure 3.2), for both Carrier and Aperture conditions. To analyse the relationship between proportion correct and spatial frequency and eccentricity, for the two conditions, I fitted a quadratic model to the data, using the nlme package in R (Pinheiro *et al.*, 2013). This enabled the modelling of the entirety of the motion discrimination data, including eccentricity, condition and participant as nested levels of fixed effects in the quadratic description of the data. The model fits one overall function to the data set, and then each of the three parameters in the model (a,b,c) can vary across each level of eccentricity, condition and participant identity, to reduce the non-linear sum of squares error.

$$\text{Proportion correct} \sim a * \log10(SF)^2 + b * \log10(SF) + c$$

At each eccentricity a different but overlapping range of spatial frequencies were tested; there were 9 spatial frequencies included in the experiment, with each eccentricity sampling 6 or 7, with the same 5 spatial frequencies included for all eccentricities. Non-linear mixed effects models can fit data with missing values, so I was able to include data from all eccentricities in the same model.

The model was produced by fitting the data with a three parameter quadratic function, fitting proportion correct by the logarithm of base 10 of spatial frequency, using the nlme procedure described in Chapter 2.6.3.

The model converged for all participants, with parameters and their respective coefficients reported in Table 3.2. There was a significant contribution of both the a parameter and the c (intercept) parameter. There was no significant contribution of the b parameter to the model fit. However, there was a significant decrease in the goodness of fit of the model when the b parameter was removed (Log Likelihood Ratio = 45.576, $p < .001$), so I included the b parameter in the model despite the non-significant contribution. There was variation around the b and c parameters at the eccentricity, condition and participant nested levels of fixed effects, and variation around the a parameter solely at the condition level.

Table 3.2. Table displaying the parameters for the quadratic model fitting the proportion correct by spatial frequency, with stimulus condition and participant identity included as fixed effects.

Parameter	Coefficient	Std. dev.	Std. dev. across eccentricity	Std. dev. across stimulus condition	Std. dev. across participant	df	t value	*p* value
a	-0.240	0.020	0.000	0.012	0.000	306	-11.701	<.001*
b	0.013	0.057	0.103	0.035	0.037	306	0.234	0.815
c	0.876	0.021	0.027	0.033	0.065	306	41.905	<.001*

To assess whether there was a difference in proportion correct between the Carrier and Aperture conditions, I ran the model on the initial dataset, where data were grouped by eccentricity, condition and participant (full model), and on the same dataset, but with data grouped only by eccentricity and participant (reduced model). There was a significant decrease in goodness of fit when condition was removed as a fixed effect (Log Likelihood Ratio =15.034 , $p < .002$), (full model: AIC = -641.369, Log Likelihood = 333.685; reduced model: AIC = -632.335, Log likelihood = 326.167), demonstrating that the model produced a significantly different set of parameters to describe proportion correct by spatial frequency in the Carrier and Aperture conditions.

Similarly, to assess whether there was a significant effect of eccentricity on proportion correct, I compared the full model against a simplified model, where only stimulus condition and participant were included as fixed effects. There was a significant decrease in goodness of fit when eccentricity was removed as a grouping variable (Log Likelihood Ratio = 205.646, $p < .0001$), (full model: AIC = -641.369, Log Likelihood = 333.685; reduced model: AIC = -441.723, Log likelihood = 230.862), demonstrating that the model produced a significantly different set of parameters to describe proportion correct at each eccentricity.

Displacement thresholds

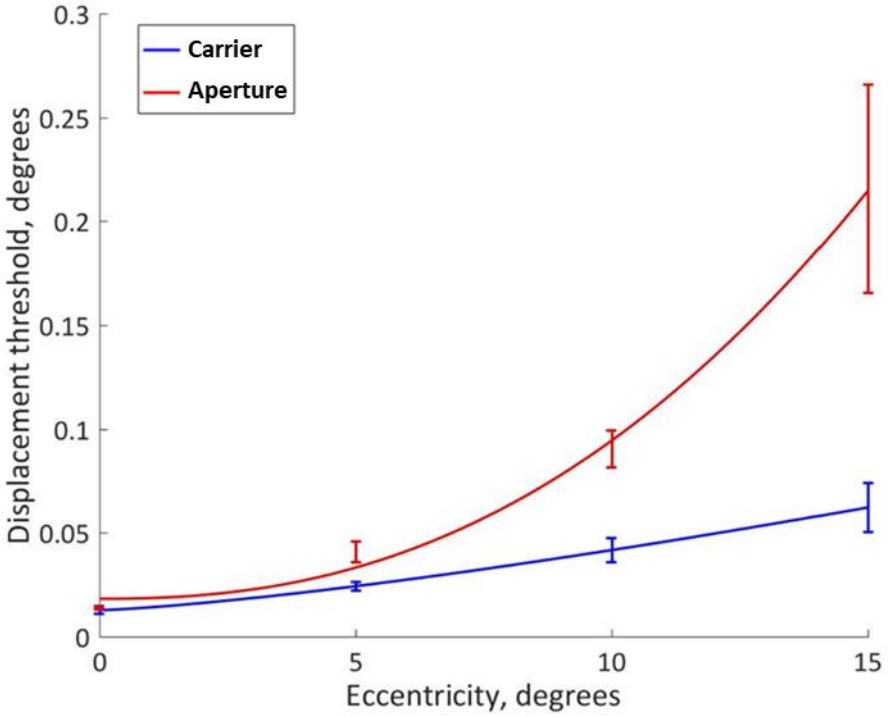


Figure 3.3 Figure plotting the displacement thresholds by eccentricity for the Carrier and Aperture conditions. Displacement thresholds correspond to the displacement of the Gabor at 75% proportion correct, where the displacement corresponds to a quarter of the wavelength of the Gabor stimulus. As before, the red curve refers to the Aperture condition and the blue curve to the Carrier condition.

For each participant, for the Carrier and Aperture conditions, and at each eccentricity, I fitted both the three parameter quadratic function described above, or a four term polynomial;

$$\text{Proportion correct} \sim a*\log10(SF)^3 + b*\log10(SF)^2 + c*\log10(SF) + d$$

which was able to fit the data better when participants had a string of chance proportion correct at the highest spatial frequencies. Using the function with the best fit to the data, I extracted the spatial frequencies that correspond to 75% proportion correct, for each participant at each eccentricity.

The spatial frequency value corresponding to 75% proportion correct can be transformed into a displacement threshold by taking a quarter of the reciprocal (1/4*SF), since each Gabor was displaced in phase or position by a quarter of the wavelength of that Gabor.

I compared displacement thresholds between the Carrier and Aperture conditions using the nlme package in R, where I fitted displacement thresholds against eccentricity using a 2 parameter quadratic;

Displacement Threshold ~ a*eccentricity^2 + c

and included stimulus condition and participant identity as fixed effects. There was no variation across the intercept parameter at the condition or participant level (< 10^{-7} *5.2), so I excluded the intercept from the fixed effects structure. This suggests that there is no significant difference in displacement thresholds at fovea. There is no loss of goodness of fit, as demonstrated by an ANOVA between the full 2 parameter model and the simplified 2 parameter model, (p = 1.000).

I then analysed whether there was a significant difference across conditions, by running my simplified 2 parameter model across the data grouped by condition and participant, and the data grouped solely by participant. There was a significant difference between the two models (Log Likelihood Ratio = 72.084, p < .001), with substantially higher AIC and lower log-likelihood values for the model without labelled conditions (full model: AIC = -208.141, Log Likelihood = 75.029; reduced model: AIC = -142.057, Log likelihood = 75.029). This demonstrates that the data are fit significantly better when allowed to vary in steepness (parameter a) by condition; there is a significant difference between condition, with a steeper rate of threshold increase in the Aperture than the Carrier condition.

The displacement threshold data were best described by a 2 parameter quadratic equation, where I included both stimulus condition and participant as fixed effects. There was no variation in the intercept between stimulus condition and participants, so model fits varied only in the a parameter. The model coefficients and parameters are reported in Table 3.4, where both the a and c parameter significantly contributed to the model, and there was a significant difference between conditions, as demonstrated by the reduction in goodness of fit when condition is removed as a fixed effect (Log Likelihood Ratio = 72.084, p < .001). This model demonstrated a significantly greater change in displacement thresholds with eccentricity in the Aperture condition, with the a parameter approximately 2.69 times greater than the a parameter in the Carrier condition.

Table 3.4. Table displaying the model parameters for the 2 parameter quadratic equation that describes the relationship between displacement threshold and eccentricity.

Parameter	Coefficient	Std. dev.	Std. dev. across stimulus condition	Std. dev. across participants	df	t value	*p* value
a	0.001	<0.001	<0.001	<0.001	41	2.322	0.025*
c	0.014	0.004	-	-	41	3.441	0.001*

E2 values

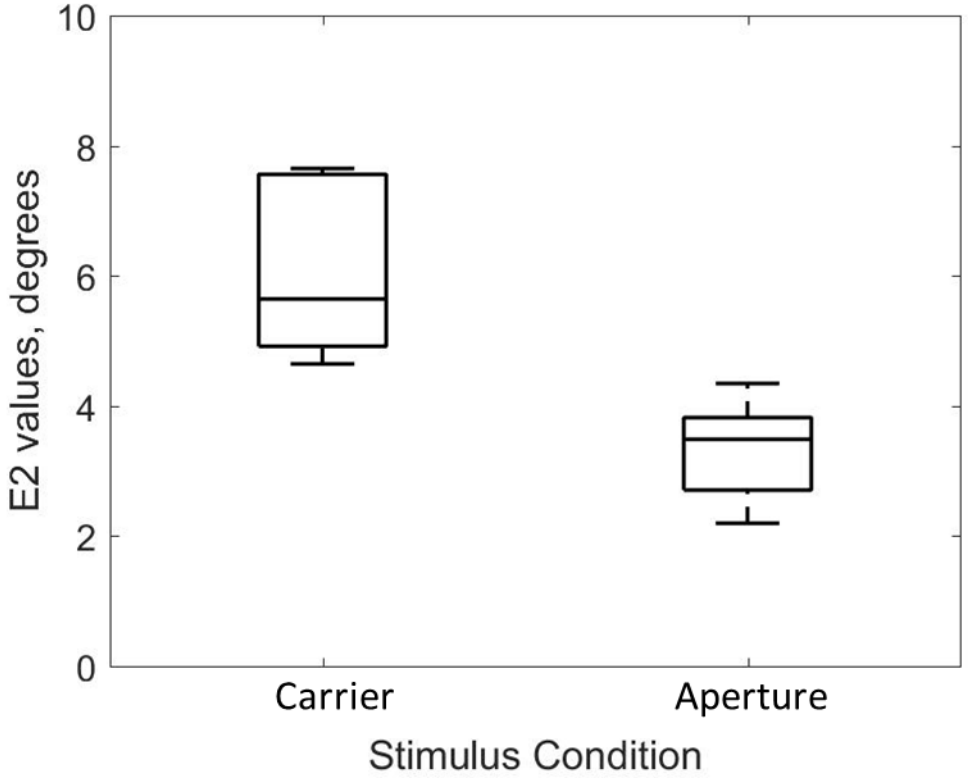


Figure 3.4. Boxplot displaying the range of E2 values across participants, for the Carrier and Aperture conditions, where the E2 values were higher in the Carrier (Mean = 5.755, standard error = 0.509) than the Aperture condition (Mean = 3.342, standard error = 0.323), indicating a shallower slope of displacement thresholds with eccentricity.

The 'E2 value' (Levi *et al.*, 1985), refers to the eccentricity at which the displacement threshold recorded at fovea has doubled in size. A low E2 value therefore demonstrates a rapid increase in displacement threshold with eccentricity, and a higher E2 value, a shallower increase in displacement threshold with eccentricity.

E2 values were extracted for each participant for the Carrier and Aperture conditions, from interpolating the spatial frequency corresponding to double the foveal displacement value. I interpolated to find the E2 value using the fitted quadratic functions, and calculated E2 values for Carrier and Aperture conditions for each participant (Figure 3.4).

I compared E2 values using a one way repeated measures ANOVA. There was a significant difference between E2 values between the two stimulus conditions, ($F(1,6) = 9.327$, $p = .022$, $\eta_p^2 = .609$), where the E2 value was significantly higher in the Carrier condition than the Aperture condition (Carrier; Mean = 5.755, standard error = 0.509: Aperture; Mean = 3.342, standard error = 0.323). There was a significant difference also between subjects effect, ($F(1,6) = 809.674$, $p < .001$, $\eta_p^2 = .993$), as the range of E2s varied between participants.

3.4 Discussion

Direction discrimination by spatial frequency and eccentricity

In this experiment I measured direction discrimination of the Carrier and Aperture conditions across the range of spatial frequencies, at four locations across the visual field. I hypothesised that there would be a difference in spatial sensitivity between the two conditions, with greater motion discrimination of the Carrier than the Aperture condition at high spatial frequencies, with the difference increasing when viewed peripherally. This would offer an explanation for why Gabors containing carrier and aperture displacements in orthogonal directions were reported in the aperture displacement direction when viewed foveally, and in an intermediate direction between the direction of the carrier and aperture displacements, when viewed peripherally.

There is a significant difference between proportion correct in the Carrier and Aperture conditions at all eccentricities. Proportion correct can be described by a quadratic function, with a higher proportion correct in the Carrier condition than the Aperture condition across the majority of spatial frequencies, particularly at the fovea and 5 degrees eccentricity. The quadratic curves describing proportion correct are displaced towards lower spatial frequencies in the Aperture condition with increasing eccentricity, than the quadratic curves describing proportion correct in the Carrier condition (Figure 3.2). Consistent with my hypothesis, the position of the curves describing proportion correct by spatial frequency along the spatial frequency axis are displaced to higher spatial frequencies for the Carrier condition, and this displacement increased with eccentricity. This can be seen in the individual participant figures in the appendix (Chapter 3.5). These results are therefore a consistent explanation for the effects of spatial frequency and eccentricity observed in conflicting cue Gabor stimuli, where motion direction is biased towards the aperture displacement at the fovea and low spatial frequencies, and towards the carrier displacement at higher spatial frequencies and in the periphery (Chapter 2.2; Shapiro *et al.,* 2010; Gurnsey & Biard, 2012). My results could also be used to explain conflicting cue performance on reverse phi stimuli, where motion is reported in the phase cue direction at higher spatial frequencies and with increasing eccentric location, (Chubb & Sperling, 1989).

Displacement thresholds

At each eccentricity I extracted the spatial frequency corresponding to 75% correct for the Carrier and Aperture conditions. The displacement in the Carrier and Aperture conditions for each stimulus was proportionate with the stimulus wavelength; the phase of the grating was displaced by 90^{o}, or the Gaussian aperture was displaced by ¼ of the stimulus wavelength. The displacement threshold for each participant was therefore ¼ of the wavelength of the threshold spatial frequency These displacement thresholds were plotted by eccentricity, and there was a significant and substantial

increase in the rate of change of the of displacement threshold by eccentricity for the Aperture condition compared to the Carrier condition (Figure 3.3).

If the logic of the cortical scaling literature is followed, the significant difference in the rate of increase of the displacement thresholds with eccentricity for the Carrier and Aperture conditions, is suggestive of a separate mechanism for the analysis of these separate stimuli, as originally intended as the experimental design. These results could reflect two potential mechanisms extracting motion using the different descriptions of the stimuli (the change in phase of the carrier, and the change in spatial position of the aperture), such as the Fourier and non-Fourier pathways described in Wilson *et al.* (1992) (Figure 1.12), a singular mechanism is using separate rectifications to make the different descriptions of the stimuli more explicit, or some further arrangement. However, this difference in displacement thresholds could be due to the difference in the image characteristics producing the displacement – the change in spatial position of the carrier features vs. the change in position of the Gaussian aperture. With the confound of the inclusion of both phase and position cues in the Carrier condition, it is more difficult to claim that these data reveal two separate underlying mechanisms for extracting phase and position cues.

The rate of increase of displacement thresholds can be compared between stimuli and different studies under the assumption that stimuli processed by the same psychophysical systems will have similar changes in acuity across the visual field, where acuity is indexed by displacement threshold (Weymouth, 1958). One measurement used to compare the rates of increase in displacement thresholds is the 'E2 value' (Levi *et al*., 1985), which refers to the eccentricity at which the displacement threshold recorded at fovea has doubled in size. A low E2 value therefore demonstrates a rapid increase in displacement threshold with eccentricity, and a higher E2 value, a shallower increase in displacement threshold with eccentricity.

There was a significant difference between my extracted E2 values between for Carrier and Aperture conditions, with a mean E2 value of 5.76 degrees for the Carrier condition, and a mean E2 value of 3.34 degrees in the Aperture condition, reflecting the substantially shallower increase in displacement thresholds with eccentricity in the Carrier condition. However, these E2 estimates are limited in their precision, as I tested solely at fovea and 5 degrees eccentricity and interpolated for each participant. Published E2 values for acuity in static gratings (Strasburger *et al*., 2011) are similar to my E2 values extracted from the displacement thresholds in the Aperture condition. Published E2 values for motion stimuli are substantially higher, reflecting a shallower increase in displacement thresholds with eccentricity for motion stimuli. Studies of unreferenced motion reported E2 values of 5.8 and 10.5 degree. (Levi *et al*., 1984), and 6.3, 7.9, 8.6 and 11.1 degrees (Whitaker *et al*., 1992),

and these are higher, but similar to my E2 values for the Carrier condition. This difference could be the result of the confound in the Carrier condition: the change in phase of the carrier produces a simultaneous position cue in the change in spatial position of the features. If the Carrier condition is perceived using both the position and phase cue, potentially depending on spatial frequency, we might expect the displacement thresholds to be biased towards the higher displacement thresholds of the Aperture condition, which contains solely a position cue. The potential bias in the displacement thresholds would then be reflected in a lower E2 value than for a 'pure' phase cue stimulus. This is a possible explanation for the E2 values in the Carrier condition are lower than the E2 values reported for other stimuli that potentially better isolate phase cues in their stimulus construction.

The difference between the E2 values reported in the literature for grating acuity and unreferenced motion mirrors the difference in the E2 values calculated for the Aperture and Carrier conditions, and supports my proposal that the phase and position cues are processed by separate psychophysical systems. These E2 values are a useful metric enabling comparison across published studies, which used different stimuli and methodologies. However, these comparisons are limited due to the low precision of my data, which is compounded by the different functions used to fit displacement thresholds with eccentricity, whereby a linear or a quadratic fit may produce very different estimates of E2 values when interpolating between data points.

I have proposed that the position system is responsible for processing the identity and location of visual stimuli, and that the ability to track salient features by changes in position is a by-product of a system that processes the location and identity of stimuli in the visual field. The similarity between the E2 values published for stationary grating acuity and the E2 values reported in my Aperture condition supports this interpretation. This finding is relevant because whilst it mirrors the accepted evidence that the decline for acuity is steeper than the decline for motion thresholds (Foster *et al*., 1989), this is the first data, to my knowledge, that demonstrates that motion extracted from position cues shows a similar decline with eccentricity as static acuity thresholds.

Limitations of the experiment

The choice of stimulus limits of the interpretation of the data in regards to possible motion systems. The impetus in using these Gabor stimuli was that the stimuli would be selectively processed by a system extracting motion from phase cues, and a system extracting motion from position cues, corresponding to the Carrier and Aperture stimulus conditions. However, as discussed previously, the change in phase of the sinusoid carrier can be described both in terms of the change in phase of

the Fourier components, but equally by the change in position of the carrier features. The features within the Gabor, such as the areas of higher and lower luminance and zero crossings will be displaced with changes of the grating phase, and thus could provide a cue for the position system. In Chapter 2.2 I replicated the results of Georgeson & Harris (Georgeson & Harris, 1990) and demonstrated that participants were able to discriminate the motion direction of the position cue in the missing fundamental grating (MFG), when presented with long blank inter-frame intervals, and could report it in the position cue direction. This is evidence of the position system's ability to track the displacement of periodic gratings.

In my Gabor experiments, this confound may not exert a large effect on displacement thresholds in the Carrier condition, especially when tested at higher eccentricities. This is due to the decline in motion discrimination in the Aperture condition at lower spatial frequencies rather than the Carrier condition. The range of spatial frequencies where the displacement threshold for the Carrier condition is occurring is in a region where the contribution from the position system is minimal. Support for the lack of position system involvement in determining Carrier condition displacement thresholds comes from the similarity in the rate of change of phase cue thresholds across eccentricity, with the rate of change of the displacement thresholds for unreferenced motion stimuli designed to further limit the involvement of the position system (Levi *et al*., 1984; Whittaker *et al*, 1992). However, this confound could affect the discrimination of motion direction in the Carrier condition at low spatial frequencies, where motion discrimination is high in the Aperture condition and so the position system may be able to process motion in the Carrier condition. This could explain why there is little difference in motion discrimination performance between the Carrier and Aperture conditions at low spatial frequencies, where on some trials, the position mechanism may have been responsible for perceived motion. This confound could be avoided by presenting the Gabor stimuli at higher temporal frequencies (Murray *et al*., 1983), disrupting the position system by presenting multiple Gabors in the display (Boulton & Baker, 1993), or using stimuli where the phase cue is not attached to a trackable feature, such as in a pedestal, (van Santen & Sperling, 1984), or using my novel stimulus the 'MF3F' (once validated through a pedestal masking experiment).

Replicating this experiment using grating based stimuli to produce 'pure' phase and position cues, will also impact the precision of the eccentric locations tested, as we can restrict the width or height of the gratings. These gratings could also be presented in narrow vertical apertures to test motion discrimination at multiple locations between fovea and 5 degrees eccentricity, where displacement thresholds at these locations would improve the precision of my estimated displacement threshold rates of increase with eccentricity for comparison with other psychophysical studies which sample the near-foveal region more densely. Ideally, the pedestal stimulus would be paired with a grating

based stimulus that contains position cues, but without a corresponding change in the phase of the Fourier components. The MF1F of Baro & Levinson (1981) is one possibility.

Aside from the potential confounding position cue in the Carrier condition, another stimulus property that may have affected displacement thresholds is the temporal frequency that stimuli were presented at. In this experiment I used two-frame presentations of the stimuli, to reduce the difference in the change in position on the retina of the Aperture condition, compared to the stationary Carrier condition. If the experiment was replicated with different grating-based stimuli, where the position cue was the change in grating features, rather than the displacement of an aperture, then temporal frequency could be varied and the experiment replicated at different temporal frequencies. Lu & Sperling (1995a) using depth and motion modulated gratings have suggested that motion discrimination performance in the position system declines with increasing temporal frequency (Lu & Sperling, 1995a). Replicating this experiment with grating based stimuli at higher temporal frequencies may produce an even greater disparity in the displacement thresholds between stimuli selectively containing phase and position cues. Manipulating both spatial and temporal frequency across eccentricity could be used to map the spatio-temporal surfaces for motion discrimination to better investigate whether different mechanisms are present to extract phase and position cues, or whether the results in this chapter are driven by differences in stimulus image characteristics.

A further theoretical issue with the approach in this chapter is that the descriptions of motion cues may not map to the cues used by motion processing mechanisms e.g. just because motion stimuli can be described by changes in phase of Fourier components, does not mean that there exists a mechanism in the human visual system performing Fourier analysis.

Conclusion

The significant difference between the spatial frequency sensitivity functions of the Carrier and Aperture conditions across eccentricity, and the significant difference in the increase in displacement threshold with eccentricity, are suggestive of a psychophysical dissociation between the displacement of the carrier and the aperture of the Gabor stimuli. This could be driven by the different image characteristics of these two stimuli, or reflect differences in the systems extracting motion from these stimuli.

The higher spatial sensitivity to the phase cues in comparison with the position cues, suggests that motion extracted from phase cues is not reliant on extracting feature information. This is seen in pedestal gratings, where the displacement of the features back and forth on each frame produces an incoherent motion direction , but participants are able to discriminate motion using the phase cue

(Van Santen & Sperling, 1984). Demonstration of apparent motion – where motion is produced between two spatially unresolvable locations (Exner, 1875) – suggest that motion discrimination is also not reliant on detecting the positions that features are presented at. A more recent demonstration by Foster *et al*. (1989) compared the discrimination of the spatial separation of two locations with the direction discrimination of apparent motion produced, by alternating whether these locations displayed a point of light with 100 ms onset asynchrony (Foster *et al*., 1989). Foster *et al*., (1989) replicated Exner's (1875) results, finding significantly lower displacement thresholds for motion discrimination than spatial discrimination, and reported a similar rate of change of these displacement thresholds with eccentricity as the rates of change with eccentricity for the displacement thresholds in the Carrier and Aperture conditions. The discrimination of motion between two spatially unresolvable locations is not, by itself, evidence that motion discrimination is reliant on the phase cue. Filter-rectify-filter models are proposed as a mechanism to extract motion from the position cues, without explicitly detecting position. However, for models of position cue processing that explicitly detect the position of features, such as feature correlation models, the limit of spatial resolution for motion detection will be the same limit as to detect spatial separation. In these models, motion would not be discriminated between locations below the limit of spatial resolvability. The similarity in the change in displacement thresholds across eccentricity for the phase cue in this experiment, and the fine-grained motion illusion in Foster *et al*. (1989), support an interpretation of phi motion as objectless motion, processed as a phase cue rather than a position cue. This could explain the phenomena reported by Walls (Walls, 1942), where motion stimuli displayed in the periphery has "a saliency and attention value quite out of proportion to the clarity with which they are actually discriminated".

These data provide a parsimonious explanation for the reversals of reported direction observed in conflicting cue stimuli, presented at different spatial frequencies, or different eccentricities due to the relative strength of the Carrier and Aperture conditions. These reversals in reported motion come from a handful of studies that sampled between two spatial frequencies or eccentricity: the experiment in this chapter systematically varied both spatial frequency and eccentricity, to demonstrate this explanation.

3.5 Appendix

The following figures depict proportion correct for each of the three stimulus conditions (data for the Opposing Motion condition is analysed and discussed in Chapter 4), plotted against spatial frequency, at the four tested eccentrities for each of the 7 participants. The Opposing Motion data was determined as 'correct' if the interval that motion was reported in was the interval in which the aperture cue moved upwards, so values below chance represent spatial frequencies at which the majority of trials were reported in the interval in which the carrier cue was upwards.

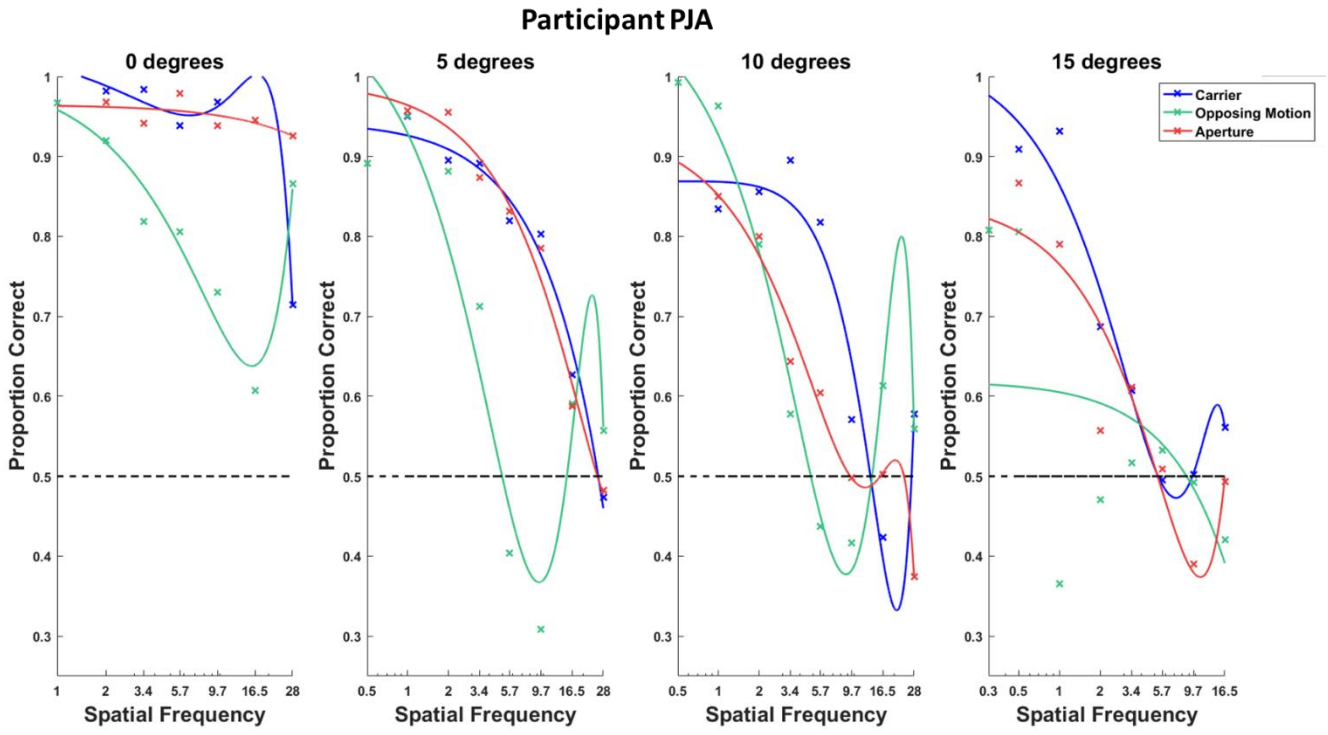


Figure 3.5. Figure plotting the proportion correct for participant PJA in the Carrier (blue curve), Opposing Motion (green curve) and Aperture (red curve) conditions by spatial frequency, at the four eccentric locations tested. The data (crosses) were fitted by two or three parameter polynomials, and displacement thresholds interpolated from these curves.

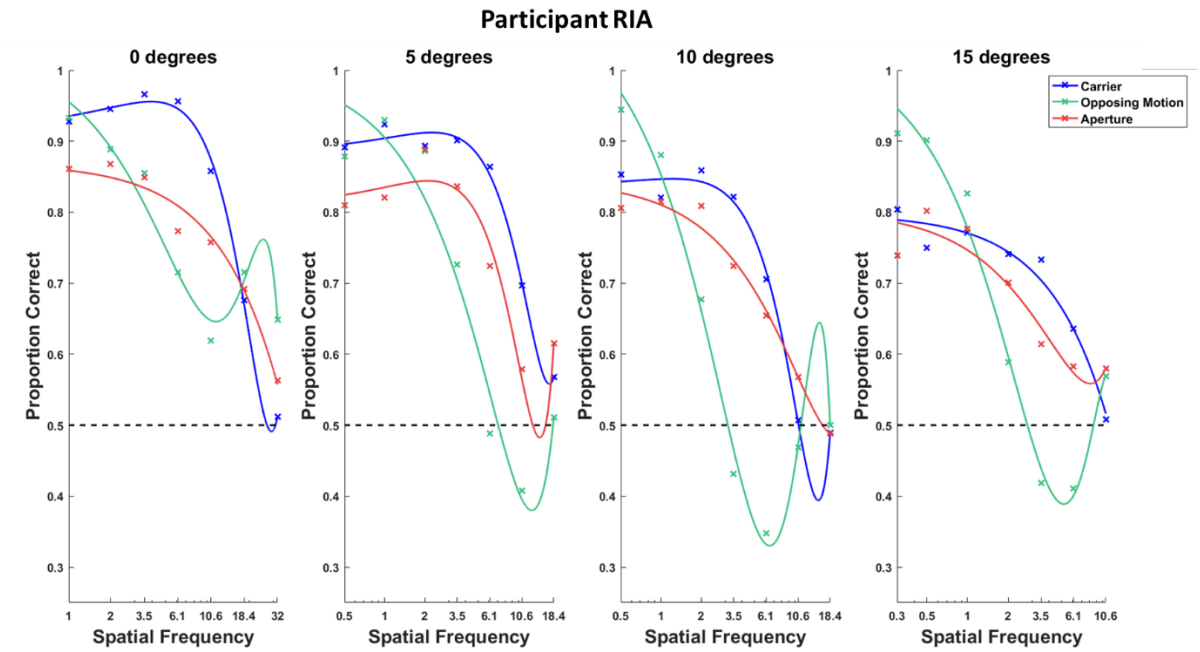


Figure 3.6. Figure plotting the proportion correct for participant RIA in the Carrier (blue curve), Opposing Motion (green curve) and Aperture (red curve) conditions by spatial frequency, at the four eccentric locations tested. The data (crosses) were fitted by two or three parameter polynomials, and displacement thresholds interpolated from these curves.

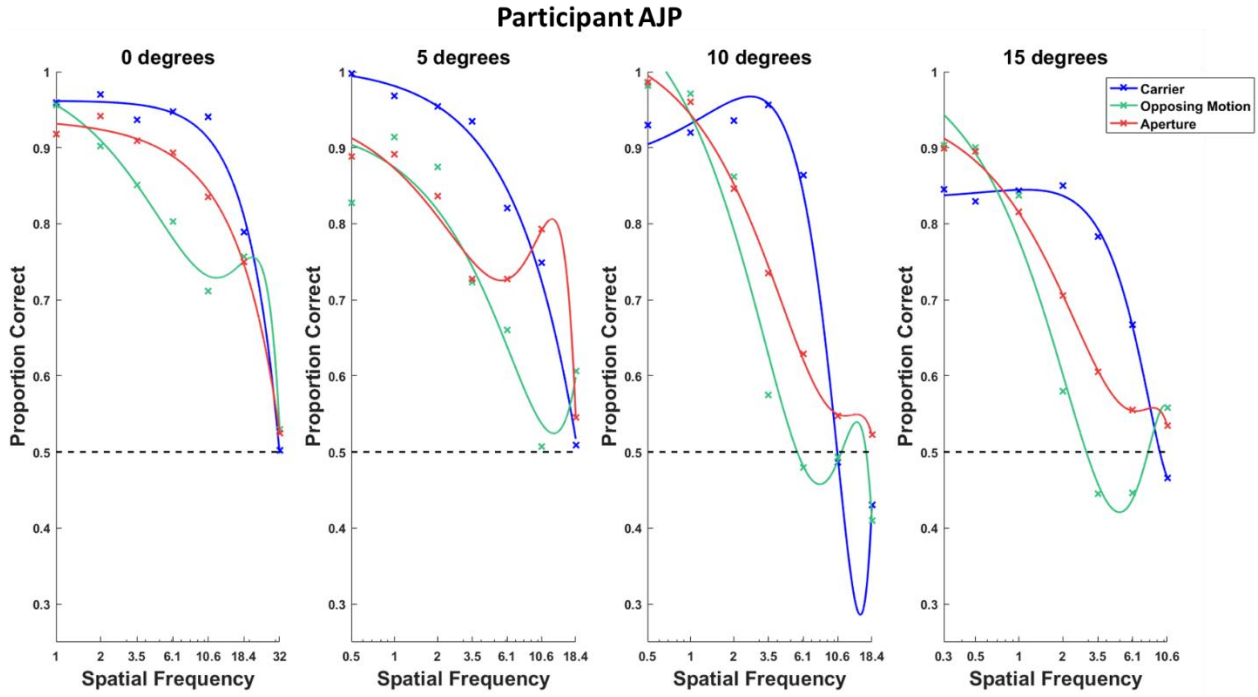


Figure 3.7. Figure plotting the proportion correct for participant AJP in the Carrier (blue curve), Opposing Motion (green curve) and Aperture (red curve) conditions by spatial frequency, at the four eccentric locations tested. The data (crosses) were fitted by two or three parameter polynomials, and displacement thresholds interpolated from these curves.

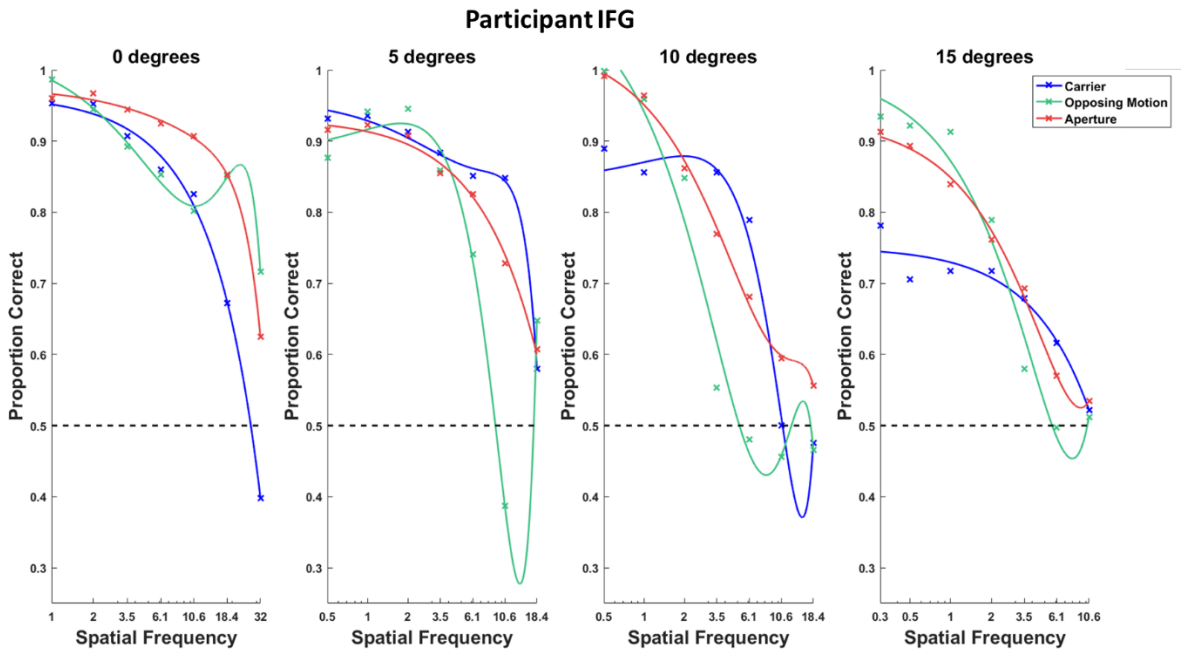


Figure 3.8. Figure plotting the proportion correct for participant IFG in the Carrier (blue curve), Opposing Motion (green curve) and Aperture (red curve) conditions by spatial frequency, at the four eccentric locations tested. The data (crosses) were fitted by two or three parameter polynomials, and displacement thresholds interpolated from these curves.

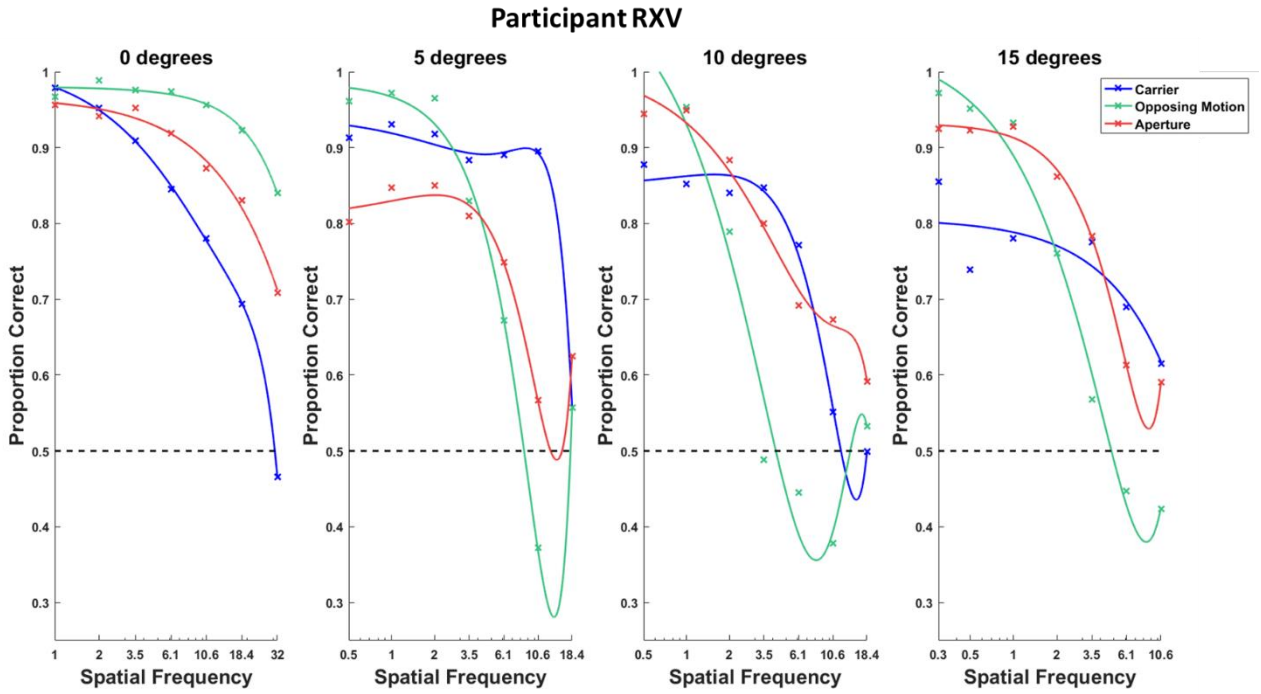


Figure 3.9. Figure plotting the proportion correct for participant RXV in the Carrier (blue curve), Opposing Motion (green curve) and Aperture (red curve) conditions by spatial frequency, at the four eccentric locations tested. The data (crosses) were fitted by two or three parameter polynomials, and displacement thresholds interpolated from these curves.

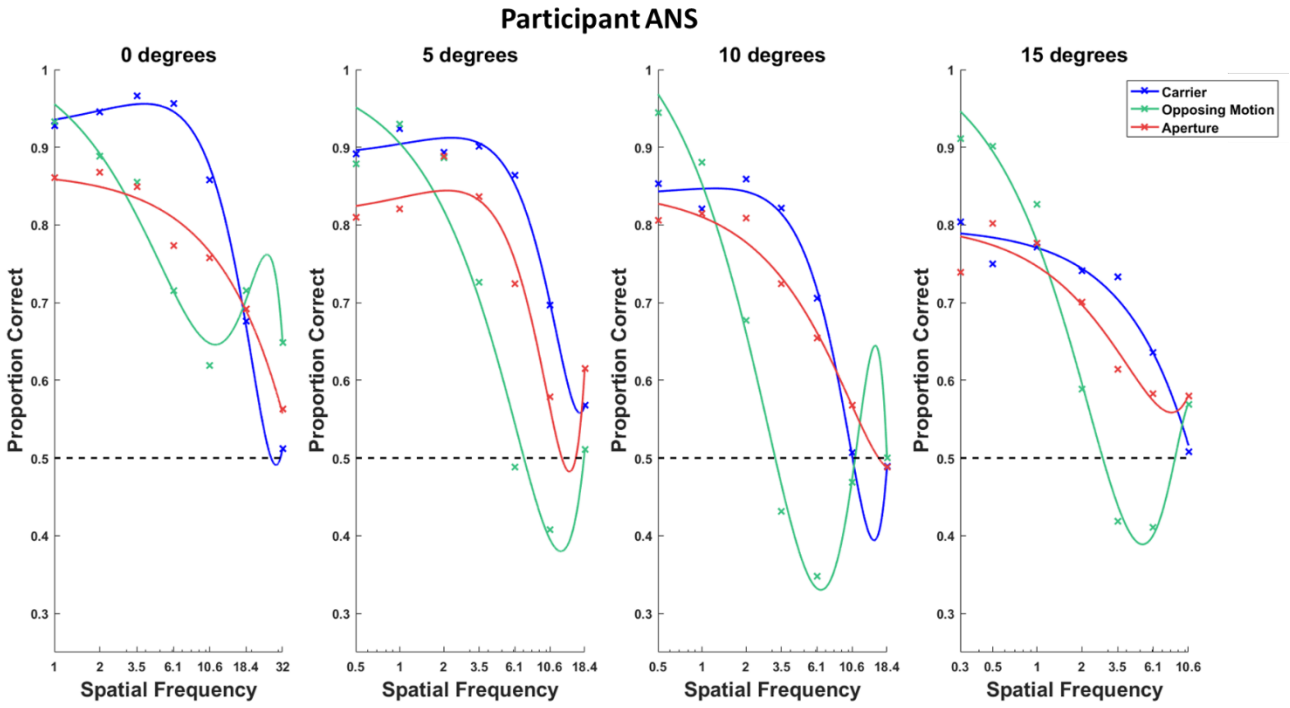


Figure 3.10. Figure plotting the proportion correct for participant ANS in the Carrier (blue curve), Opposing Motion (green curve) and Aperture (red curve) conditions by spatial frequency, at the four eccentric locations tested. The data (crosses) were fitted by two or three parameter polynomials, and displacement thresholds interpolated from these curves.

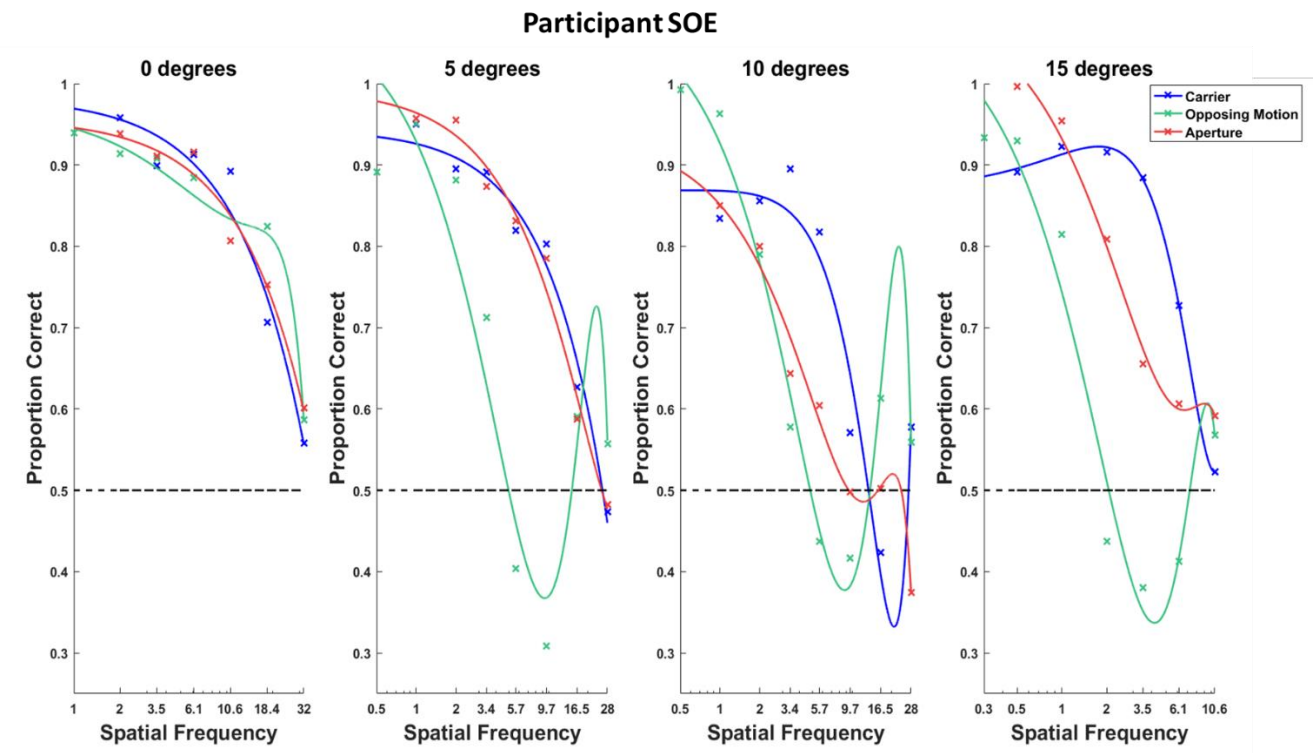


Figure 3.11. Figure plotting the proportion correct for participant SOE in the Carrier (blue curve), Opposing Motion (green curve) and Aperture (red curve) conditions by spatial frequency, at the four eccentric locations tested. The data (crosses) were fitted by two or three parameter polynomials, and displacement thresholds interpolated from these curves.

Chapter 4: Perceived motion in Gabor based conflicting cue stimuli

4.1 Introduction

In Chapter 3 I assessed the spatial frequency sensitivity to changes in the phase of Fourier components and changes in spatial position using Gabor stimuli, at the fovea and across periphery. I found that across all the retinal locations tested that participants' motion discrimination, measured by proportion correct, could be best described by quadratic functions, that were displaced to lower spatial frequencies with increasing eccentricity. Performance declined more rapidly for Gabors where motion was produced through the change in spatial position of the aperture (Aperture condition), than for the Gabors where motion was produced from the change in phase of the carrier (Carrier condition), which could be modelled by the quadratic function being more rapidly displaced towards lower spatial frequencies with increasing eccentricity. Furthermore, the change in displacement thresholds across eccentricity also varies between the Carrier and Aperture conditions, with a much more rapid increase in displacement thresholds with eccentricity for the Aperture condition. I suggested that these results demonstrate a psychophysical dissociation between motion extracted from the change in phase of the carrier and motion extracted from the change in spatial position of the aperture. This could be caused potentially be caused by the activity of two separate mechanisms, one that processes the phase cues for motion and one that processes position cues for motion. Equally, this could difference could be caused by the difference in stimulus characteristics and processed by the same mechanism, or reflect the activity of one mechanism processing both stimulus conditions, but applying different rectifications to the visual input.

I propose that the results described in Chapter 3 can be described as different spatial sensitivity functions, which are biased towards low spatial frequencies for displacements in Gabor aperture, and towards the higher spatial frequencies when motion is produced from the change in phase of the carrier. If this is the case, then reported motion in stimuli which contain both carrier and aperture displacements, but signal motion in conflicting directions, should have a predictable relationship with spatial frequency. I would predict a bias towards the aperture displacement direction at low spatial frequencies and towards the carrier phase shift at higher spatial frequencies, with this bias increasing with increasing eccentricity, reflecting the increasing difference in displacement thresholds observed in Chapter 3 (Figure 3.3). Assessing reported motion in conflicting cue stimuli can therefore be used to test whether this spatial scaling hypothesis accurately represents the organisation of the motion processing system.

Most research using Gabors has used multi-frame stimulus displays that will bias motion in the direction of the aperture displacement, as with multiple frames the aperture change will be displaced over an extended distance across the retina. I have shown in Chapter 2 that, using a two-frame Gabor presented at 10 degrees eccentricity, the perceived motion direction in conflicting cue Gabors is determined by spatial frequency. Similar results were reported by Gurnsey & Biard (Gurnsey & Biard, 2012), with a many-framed Gabor stimulus, where the bias towards the carrier cue was increased for smaller Gabors with higher spatial frequencies. This is consistent with my spatial scaling hypothesis.

For conflicting cue Gabor stimuli presented at or near the fovea, motion is overwhelmingly reported in the direction of the aperture displacement, as demonstrated by the many-framed drifting orthogonal Gabor (Tse, 2006; Shapiro *et al.*, 2010). As the viewing location becomes increasingly eccentric, motion in orthogonal Gabors is reported in a direction intermediate to the aperture displacement and the direction along which the phase of the carrier is shifted, with a greater bias towards the carrier shift at greater eccentricities (Shapiro *et al.,* 2010; Gurnsey & Biard , 2012). This does demonstrate an effect of eccentricity on perceived motion, but suggests that motion at the fovea is perceived solely in the position cue direction. Hughes *et al.* (Hughes *et al.*, 2015) reported that the perceived motion trajectory of a drifting oblique Gabor, which was binocularly tracked and thus displayed foveally, was slightly biased by the motion direction of the grating. Boulton & Baker (1993) showed that for Gabors arranged with the aperture displacement and carrier shift in opposing directions, when viewed at 2.5 degrees eccentricity motion was reported in the aperture displacement direction for single Gabors, but in the direction of the carrier shift when multiple Gabors were presented as a 'micro-patttern'. This demonstrates the impact of stimulus parameters, whereby increasing the number of Gabors present could have reduced the relative strength of the aperture displacement as there may have been an increase in a possible correspondence problem between Gabors, or increased the relative strength of the carrier shift as the change in phase of the Fourier components would be identical across all Gabors presented.

I suggest that the lack of motion perceived in the carrier shift direction for orthogonal Gabors presented at the fovea in Gabors is a function of the stimulus construction, where the use of many frames is increasing the strength of the aperture cue proportional to the carrier cue, and so dominating the motion percept. I predict that, with a shorter frame display of the conflicting cue Gabors, reported motion will show the spatial frequency bias predicted by the spatial scaling hypothesis when viewed foveally, for both opposing and orthogonal arrangements of the aperture displacement and change in phase of the carrier. This was assessed across two experiments, both

using two-frame Gabors, with equal distant displacements of the Gaussian aperture and phase shifts of the carrier, with these cues arranged in opposing and orthogonal directions.

4.2 Opposing motion Gabors presented across the visual field

To assess the role of both spatial frequency and eccentricity, in determining perceived motion in conflicting cue conditions, I measured reported motion for a two-frame conflicting cue Gabor, where a horizontal Gabor simultaneously underwent a phase displacement of the carrier and a displacement of the Gaussian aperture in the opposite direction. These data were collected in the experiment reported in Chapter 3, alongside the Carrier and Aperture stimuli conditions. These data are presented here separately as they were collected following a different theoretical rationale, and require separate analysis due to the presence of two valid motion directions in each Gabor, whilst the data reported in Chapter 3 had a single valid motion direction.

Method

In this section, I describe the results from the Opposing Motion condition, which contained a phase displacement of the carrier and a displacement of the Gaussian aperture arranged in opposing directions. These 'Opposing Motion' Gabors were displayed across seven spatial frequencies at four retinal locations, the fovea and 5, 10 and 15 degrees eccentricity. For further details on the methods, see Chapter 3.2.

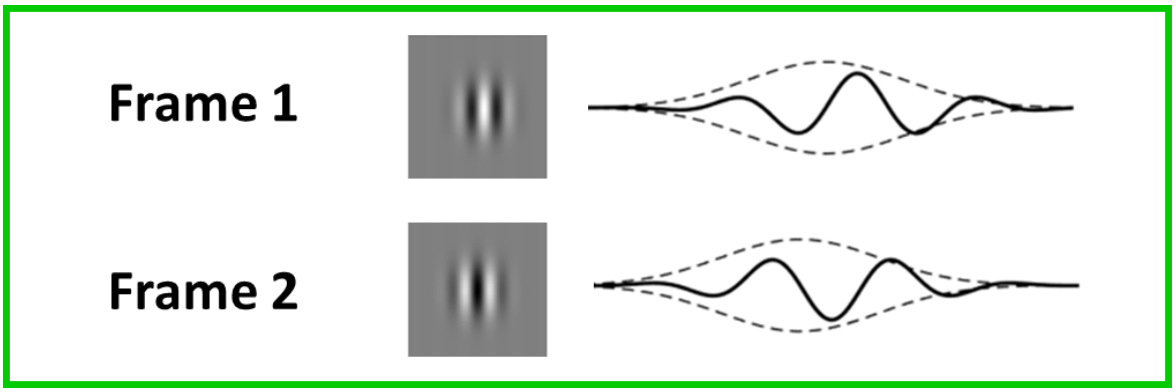


Figure 4.1. Schematic of the Opposing Motion condition. The grating is displaced by 90° rightwards (carrier cue), and the window (aperture cue) is displaced leftwards, by a distance equivalent to a quarter of the wavelength of the grating. The central column displays the Gabor as it appears to the participant, and the righthand column the cross-section of the stimulus, with the dashed line indicating the Gaussian envelope.

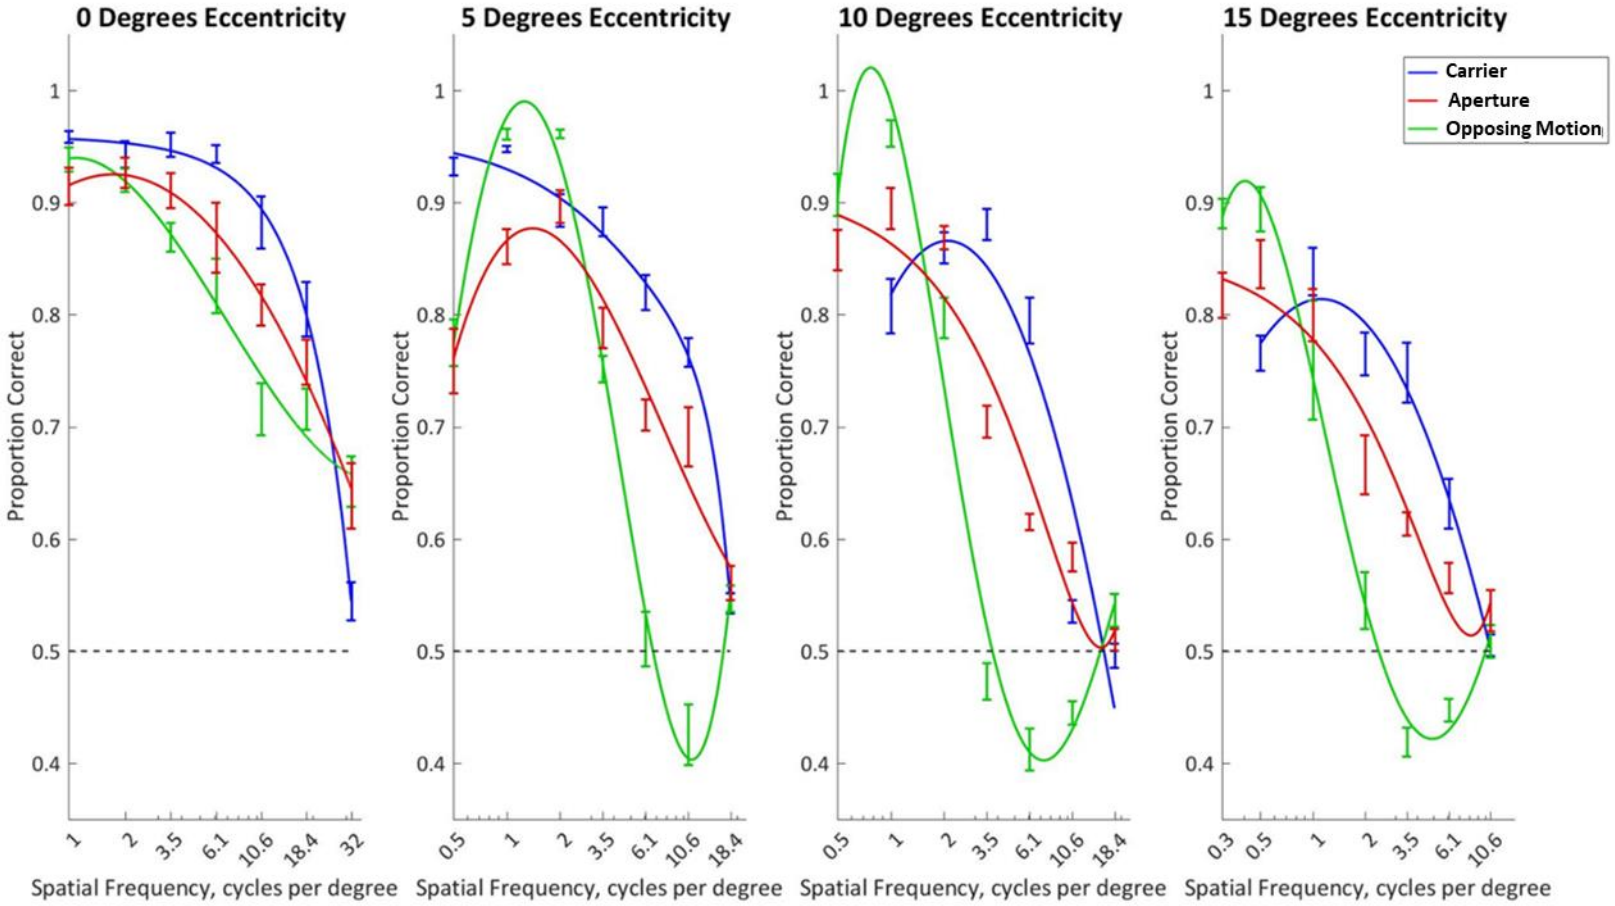


Figure 4.2. This figure displays the three stimulus conditions, at the four locations in the visual field at which participants viewed them. Proportion correct is used to compare the Opposing Motion condition data to the Carrier and Aperture conditions, presented in Chapter 3. For the Opposing Motion condition data 'correct' refers to trials reported in the aperture cue direction, so data points significantly above 0.5 are predominantly reported in the position cue direction, and data reported below 0.5 is predominantly reported in the carrier cue direction. The green curves refer to the Opposing Motion condition, the data described in this chapter, with the blue and curves referring to the Carrier condition and Aperture conditions data presented in Chapter 3.

In this chapter I used a 2afc procedure, where participants reported in which interval motion was perceived upwards. In the Opposing Motion condition, where both intervals include an upward and a downward cue, in either the phase shift of the carrier (carrier cue), or the displacement of the aperture (aperture cue), I coded the trials where participants reported the interval with the upward aperture displacement as 'correct'. The proportion of correct trials can be plotted against spatial frequency, as shown in Figure 4.2, and it is this 'proportion correct' that is used to analyse the use of the carrier and aperture displacements in this section.

The proportion correct responses made in the fovea and periphery have a different relationship with spatial frequency. At the fovea, participants' performance is best described by a quadratic function (Equation 1), where proportion correct peaks at the lowest spatial frequency and declines with increasing spatial frequency. In the three peripheral testing locations, proportion correct was better described by a third order polynomial (Equation 2), with the highest proportion correct at low spatial frequencies and turning point, with lowest proportion correct, at mid-high spatial frequencies. Due to the difference in the shape of the data I analysed foveal and peripheral data separately.

$$\text{Equation 1;}\quad \text{Proportion correct} \sim a*\log10(SF)^2 + b*\log10(SF) + c$$

$$\text{Equation 2;}\quad \text{Proportion correct} \sim a*\log10(SF)^3 + b*\log10(SF)^2 + c*\log10(SF) + d$$

Analysis of foveal data

I investigated the relationship between the proportion of trials reported in the aperture displacement direction at each spatial frequency at the fovea (Figure 4.2, lefthand panel). Using non-linear mixed effects modelling I fit the data, including participant as a fixed effect, and compared the goodness of fit for three linear models; a straight line, a quadratic function and the third order polynomial that best fit the peripheral data. The data were significantly better fit by a quadratic model (AIC = -82.945) than a straight line (AIC = -76.93224), (Log Likelihood Ratio = 10.017, p = .007), and there was no improvement in goodness of fit between the quadratic model and a third order polynomial (p = .547). The coefficients for the quadratic model are reported in Table 4.1, where despite the lack of significance of the a parameter, the model is a better fit than a one or three parameter model potentially reflecting the inclusion of participant as a fixed effect.

Table 4.1. Table displaying the model information for the model fitted to the foveal data.

Parameter	Coefficient	Std. dev	Std. dev across participants	df	t value	*p* value
a	-0.037	0.054	<0.001	40	-0.684	0.498
b	-0.149	0.086	0.034	40	-1.734	0.091
c	0.952	0.033	0.045	40	28.598	<0.001*

Analysis of peripheral data

The peripheral data sets could be included in one model, as non-linear mixed model analyses can handle data sets with overlapping but different ranges of x value (SF), providing the data can be modelled by the same underlying equation. As can be seen in Figure 4.2 the proportion correct data for the three peripheral conditions have similar relationships with spatial frequency , with a peak at lower spatial frequencies and a turning point at higher spatial frequencies.

I fit the peripheral data with a third order polynomial, to proportion correct by the logarithm of base 10 of spatial frequency, and included eccentricity as a categorical factor and a fixed effect variable, and participant identity as a second fixed effect;

$$\text{proportion correct} \sim a*\log10(SF)^3 + b*\log10(SF)^2 + c*\log10(SF) + d.$$

Whilst the c parameter has a smaller contribution to the model than the other parameters ($t = -1.729$, $p = .086$), the model fit was significantly impaired when the c parameter was removed, (Log Likelihood Ratio = 83.643, $p < .001$). The coefficients for this model are reported in Table 4.2, along with the standard deviation of each parameter across the two fixed effects of eccentricity and participant.

Table 4.2. Table displaying the model information for the model fitted to the peripheral data.

Parameter	Coefficient	Std. dev	Std. dev across eccentricity	Std. dev across participant	df	t value	*p* value
a	0.885	0.098	<0.001	<0.001	123	9.009	<.001*
b	-1.014	0.263	0.393	<0.001	123	-3.857	<.001*
c	-0.314	0.182	0.303	<0.001	123	-1.729	0.0863
d	0.898	0.051	0.081	<0.001	123	17.642	< .001*

To analyse the effect of eccentricity on performance, I produced a comparison reduced model effects that included only participant identity as a fixed effect (reduced model). I compared the fit to the data between the final (AIC = -162.431) and reduced (AIC = -130.573) models, and found a significant decrease in goodness of fit when eccentricity was removed from the model (Log Likelihood Ratio = 37.858 p <.001), demonstrating that a better fit to the data was produced when the parameters were allowed to vary by eccentricity. This indicates a significant difference in the relationship between proportion correct and spatial frequency across the three peripheral viewing locations.

Using the 2afc procedure means that incorrect trials could be produced by motion reported in the carrier phase shift direction (hence a downward position cue interval), and trials where motion is reported erroneously. This means that I can only conclude that the motion is reported in the direction of the phase shift of the carrier or the displacement of the aperture when mean performance is significantly above or below chance.

In the three peripheral conditions, the modelled proportion correct falls below chance for at least one of the spatial frequencies tested, for all participants. I ran a one sample t test at each eccentricity on the set of the minimum proportion correct for each participant at that eccentricity, with a Bonferroni correction for multiple comparisons (corrected alpha p = 0.01667).

There was no significant difference between mean minimum proportion correct and chance at 5 degrees eccentricity (p = .088), but there was a significant difference at 10 degrees eccentricity (mean proportion correct = 0.363, t(6) = -4.497, p = .004), and at 15 degrees eccentricity (mean proportion correct = 0.348, t(6) = -3.461, p = .0134).

Peaks and troughs

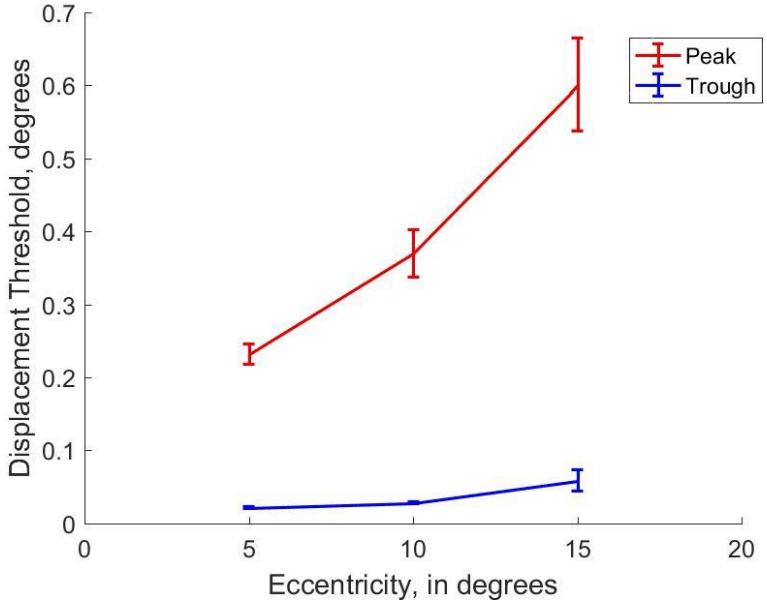


Figure 4.3. Figure displaying the displacement thresholds corresponding to the maxima and minima of participants' proportion correct for reported motion direction in the Opposing Motion condition, extracted from the fitted third order polynomials. The peaks of these curves correspond to the spatial frequency at which the most trials were reported in the aperture displacement direction, and the troughs, the most trials reported in the direction of the phase shift of the carrier. These values are plotted here in units of displacement, for each of the three peripheral viewing locations.

The effect of eccentricity may also be observed by comparing the locations of the peaks and troughs of the third order polynomials fitted to the proportion correct data, for the three peripheral conditions (Figure 4.4). The maxima and minima were extracted from third order polynomials fitted to each participants' data individually, at each eccentricity, displayed in Figure 4.3. I refer to the peaks and troughs as 'turning points', where 'turning point type' refers to whether the data point corresponds to a peak or a trough. To simplify comparisons with the displacement thresholds presented in Chapter 3 and Chapter 5 I analysed these turning points in units of degrees, corresponding to the displacement distance of the aperture (which is the same distance as covered by the change in spatial position of the carrier features when the carrier is phase shifted), rather than in units of cycles per degree.

I used a non-linear mixed effects modelling to compare different fits for the data, where I included both turning point type (peak or trough) and participant as fixed effects, and found that the data was best described by a single parameter linear fit;

turning point ~a*eccentricity.

The a parameter contributes significantly to the goodness of fit of the model, t(28) = 3.358, p = .050. The a parameter has a coefficient of .050 for the peaks condition, and 0.004 for the trough condition, with a standard deviation of 0.013, and across participants, a standard deviation of 0.003.

There is a significantly better goodness of fit of the model when the model includes both turning point type and participant identity as fixed effects, than when only participant identity is included as a fixed effect, (Log Likelihood Ratio =73.339 p <.001). This demonstrates a significant difference between the two conditions, and as the only parameter value in the model is the slope, I concluded that there is a significantly different slope between the turning point values with eccentricity for the peaks and troughs.

Discussion

I expected to see a reversal in proportion correct with spatial frequency at each eccentricity, with motion reported significantly above chance at the low spatial frequencies, corresponding to being reported in the aperture cue direction, and significantly below chance at the high spatial frequencies, corresponding to being reported in the carrier cue direction. This prediction was upheld in the three peripheral locations, but not when the Opposing Motion Gabor stimulus was viewed foveally.

Foveal viewing

At fovea, the same pattern is not seen in reversal of reported motion as is seen in the peripheral data. Motion discrimination is best described by a quadratic function, and proportion correct may be solely determined by the aperture cue. This result mirrors the reported motion direction in Shapiro *et al.*'s Break of the Curveball, where participants reported motion in the aperture cue direction when viewing an orthogonal Gabor stimulus foveally (Shapiro *et al.*, 2010).

At the fovea, the quadratic shape of proportion correct by spatial frequency in the direction of the aperture cue, could reflect no effect of the opposing carrier cue. This is the argument put forth by Shapiro *et al.* (2010), where participants rotated an arrow to report motion direction of an orthogonally oriented conflicting cue Gabor, and reported no significant bias in motion away from the direction of the aperture cue. Shapiro *et al.* (Shapiro *et al.*, 2010) propose that the foveal bias for the aperture cue is due to a lack of integration between the carrier and aperture cues at the fovea. This is not the case in the reverse phi stimuli viewed foveally, where Chubb & Sperling (Chubb & Sperling, 1989) report that perceived motion can be reversed by manipulating the viewing distance, and therefore the spatial frequency of the stimulus on the retina. The Opposing Motion Gabor

contains position cues in both cue directions: the displacement of the aperture, but also in the change in carrier phase producing corresponding changes in position of the carrier features. It is possible that the absence of reversal in the Opposing Motion Gabor is due to the presence of additional phase and position cues in the stimulus, whereas the reverse phi stimuli have a dissociation in the direction of phase and position cues.

Alternatively, the absence of motion reversal at the fovea could be that the relative strength of the aperture cue is far greater than the strength of the carrier cue: the competition between the two cues that drives motion reversal in the periphery is still occurring, but the relative strength of the carrier cue is so weak that there is no observable effect on reported motion direction. I had thought that this bias at the fovea towards the aperture cue direction would be reduced in a two-frame stimulus, as this dramatically reduces the change in location between the start and end frame for the stimulus. In Chapter 3 I compared motion discrimination across the same spatial frequency and eccentricity ranges for Gabors where motion was produced through a displacement in the aperture position, and Gabors where motion was produced through a phase shift of the carrier. The proportion correct for these two conditions was relatively similar at fovea, but is a crude measure of relative strength. One potential method to investigate the relative strength of these cues, in an attempt to interpret the bias in the Opposing Motion Gabor towards the aperture cue, would be to measure contrast sensitivity functions for the two single cue conditions.

Peripheral viewing

In the three peripheral viewing locations, reported motion showed the predicted reversals with spatial frequency. The proportion correct data for the Opposing Motion condition was modelled by a third order polynomial with two inflection points (see Figure 4.2). Proportion correct peaked at the low spatial frequencies, was below chance at the mid to high spatial frequencies, and tended to chance performance at the highest spatial frequencies. Furthermore, the inflection points where proportion correct reverses were consistent with the locations where the spatial frequency sensitivity functions of the Carrier and Aperture conditions can be seen to separate along the spatial frequency axis (Figure 4.2). This suggests that the reversals of perceived motion observed for the peripheral viewing locations are driven by the change in spatial sensitivity functions of the Carrier and Aperture conditions which I reported in Chapter 3.

From the fitted functions for each participant, I extracted the spatial frequency corresponding to the peaks and troughs of the polynomial. These can be conceptualised as the locations at which the

aperture cue was perceived as relatively strongest compared to the carrier cue, and likewise, the trough in the data as the spatial frequency corresponding to the carrier cue being perceived as the strongest compared to the aperture cue. These show a significantly different gradient of change with eccentricity, where the peak value, the performance in the aperture cue direction, increases much more with eccentricity than does the trough value, the maximum performance in the carrier cue direction (Figure 4.3). This is consistent with the displacement thresholds extracted for the carrier and aperture conditions reported in Chapter 3, where the two cues were presented independently, which showed the same trend – a greater increase in displacement thresholds with eccentricity in the aperture condition than the carrier condition (Figure 3.3). This supports the interpretation that the spatial sensitivity to displacements of the aperture is displaced to lower spatial frequencies at a more rapid rate than the spatial sensitivity to phase shifts of the carrier, of equivalent distance. The distance of the peaks and trough from chance performance is not symmetrical. This could indicate that the aperture cue is still exerting an effect on perceived direction, even when the majority of trials are reported in the carrier cue direction, or could reflect a regression to chance performance with increasing spatial frequency.

The shape of proportion correct results with spatial frequency is identical across the three peripheral conditions. The peak and trough of the proportion correct data can be explained by a difference in the relative strength of the carrier and aperture cues, but there are two possible explanations for the ascending arm following the trough, where performance at the highest spatial frequencies tested at each of the peripheral conditions is at chance. First, it could represent an increased use of the aperture cue at high spatial frequencies, suggesting a narrow range for carrier cue performance, or alternatively could reflect a return to chance performance, when spatial frequencies are high and acuity to the carrier and aperture cues are reduced. To differentiate between increased use of the carrier cue direction at high spatial frequencies, or performance showing a return to chance performance at the highest spatial frequencies, requires a methodology that collects reported motion direction, and so can differentiate between lower proportion correct, and trials reported in the phase or position cue direction.

Limitations of the methodology

The choice of the 2afc procedure limits our interpretation of the data. Since performance is so heavily biased in the aperture cue direction, this was taken as the default ‘correct’ direction. The use of the 2afc procedure means that we cannot differentiate between errors and trials reported in the carrier cue direction, and so cannot conclude that motion is reported in the carrier cue direction

unless motion is significantly below chance. This is especially an issue at the fovea, where motion is predominantly reported in the aperture cue direction, but the spatial scaling hypothesis would suggest at least some impact of the conflicting carrier displacement.

One method to negate this methodological issue is to compare proportion correct in the Opposing Motion condition to proportion correct in a two-frame Gabor stimulus, which contains solely a displacement of the aperture cue. While the Aperture condition included in this experiment (and reported in Chapter 3) includes a displacement of the aperture, it also includes a rotation of the carrier between frames, so there is a difference in Gabor orientation on each frame. The change in orientation adds an additional cue, that may decrease participant performance due to spurious pixel correlations between frames in horizontal directions. This could explain why performance in the Opposing Motion condition presented in this chapter is substantially higher than performance in the Aperture condition in the peripheral viewing locations (Figure 4.2), and therefore indicate that direct comparisons between the Opposing Motion and Aperture conditions is not an effect measure for identifying the strength of the aperture cue at fovea. Due to the rotation of the carrier in the Aperture condition, the orientation in the Aperture and Opposing Motion conditions will only match on one of the two frames presented in each interval. Rider *et al.* reported that the perceived velocity of Gabors was affected by grating orientation, with slower perceived velocity for Gabors with orthogonal rather than parallel orientation (Rider *et al.*, 2014). The normalised Motion Energy Analysis conducted in Chapter 2.4, and also displayed in Table 3.1, also suggests an effect of carrier orientation. There is a greater normalised Motion Energy (nME) value for the aperture cue in a conflicting cue stimulus with a carrier with parallel orientation (Orthogonal Motion condition, nME = 0.628, Position stimulus = 0.628), than an orthogonal orientation (Opposing Motion condition, nME = 0.313) to the displacement of the aperture. The effect of the grating orientation on the performance when the Gaussian window is displaced also prevents a direct comparison of proportion correct between the Aperture condition in Chapter 3 and the Opposing Motion condition. A solution to this limitation is to change the methodology and assess reported motion direction, rather than using the 2afc procedure and asking participants to report ‘the interval in which motion direction was perceived upwards’.

Conclusion

My results demonstrate that, when viewed in the periphery, Gabors with an aperture displacement and carrier phase shift in opposing directions (Opposing Motion condition) are reported in the aperture displacement direction at low spatial frequencies and in the carrier phase shift direction at

high spatial frequencies. When performance in the Opposing Motion condition is plotted alongside Gabors which contain motion solely through an aperture displacement or carrier phase shift (Aperture and Carrier conditions, Chapter 3) (Figure 4.1), the inflection point of the proportion correct in the Opposing Motion condition occurs where the performance in the Aperture condition declines, but the performance in the Carrier condition remains high. This suggests that the spatial frequency sensitivity of the Carrier and Aperture conditions is driving the reversals of perceived motion in the Opposing Motion Gabor stimulus, as suggested by my spatial scaling hypothesis. However, at the fovea, there is no reversal of motion in the Opposing Motion condition, despite the maintained difference in the spatial scaling functions of the Carrier and Aperture conditions, and the reversals of reported motion in the reverse phi stimuli at fovea (Chubb & Sperling, 1989). I suggest that this is due to the methodology used, where it is not possible to distinguish between trials reported in the carrier cue direction and errors, and the higher relative strength of the aperture cue rather than the carrier cue, which is visible in the asymmetry of the peaks and troughs in the peripheral reversals of motion in the Opposing Motion condition.

Determining whether there are reversals of motion, specifically motion reported in the carrier cue direction, at fovea, is important in determining whether the same mechanisms are involved in motion perception at fovea and periphery, and results reflect the imbalance between the relative strength of the carrier and aperture cues in the Gabor stimulus, or whether motion perception is qualitatively different between the fovea and periphery. If reported motion is driven by the differential sensitivity to the aperture and carrier cues, determined by the changing spatial sensitivity functions with eccentricity, these result could potentially be applied to understanding of conflicting cue stimuli with different constructions, and to develop wider hypotheses about motion processing abstracted from the individual image characteristics of different stimulus constructions.

Reported motion direction in reverse phi stimuli can be manipulated by changes in spatial frequency, where motion is reported in the phase cue direction at high spatial frequencies, and in the position cue direction at lower spatial frequencies (Chubb & Sperling, 1989, Benton et al., 1997). If the conflicting Gabor stimuli follow a similar pattern to the reverse phi stimuli, then this might suggest that motion direction in the Gabor stimuli is also driven by sensitivity to the phase and position cues in the image, and that motion in the carrier cue direction is determined by the phase cue of the phase shift of the carrier, and not by the change in spatial position of the features of the carrier. If so, then we would be better able to use Gabor stimuli as a stimulus to probe wider questions about responses to phase and position cues by the visual system, and investigate whether there are mechanisms is that visual system that do differentially process visual information using the Fourier and spatial position descriptions of the visual input.

To further investigate whether reported motion in Opposing Motion Gabors is determined by similar or different mechanisms at the fovea and near periphery I ran a further experiment, where I displayed two-frame conflicting cue Gabors across a similar range of spatial frequencies, at the fovea and at 10 degrees eccentricity, but changed the method for reporting direction, so that participants reported the direction of perceived motion by rotating an arrow, and thus could report direction in any degree (0° to 360°).

4.3 Differentiating between the Spatial Scaling and Feature Blur hypotheses

To fully support the spatial scaling hypothesis, it is essential to assess whether there is any effect on reported motion from the presence of the carrier cue at fovea, and whether there is a variation in reported motion direction with changing spatial frequency. I also investigated reported motion at 10 degrees eccentricity, to verify that the reversals of reported motion in the periphery accurately reflected a change in perceived motion direction and were not an increase in errors due to the presence of a conflicting cue.

An alternative explanation for the change in perceived motion between the fovea and periphery in conflicting cue stimuli was proposed by Shapiro *et al.* in the 'feature blur hypothesis', (Shapiro *et al.*, 2011), who suggested that the carrier and aperture cues were processed separately at the fovea, and integrated in the periphery. The lack of integration of the carrier and aperture cues in the fovea does not necessarily prove the absence of an integration mechanism at fovea; the cues may still be integrated, but the contribution from the carrier cue is too small to exert an effect on reported motion trajectory.

In this experiment I compared whether reported motion direction in Gabor based conflicting cue stimuli is better accounted for by Shapiro's explanation, or by my hypothesis; a dissociation in spatial frequency sensitivity functions for the carrier and aperture cues which scale differently with eccentricity. To enhance the ability to discriminate between my explanation and the 'feature blur hypothesis', I included the same stimulus as Shapiro *et al.*, a Gabor with an orthogonal relationship between the carrier and aperture cues, but maintained the use of two-frame apparent motion presentation to increase the likelihood of detecting the contribution of the carrier cue to reported motion trajectory. In this experiment participants reported motion direction by rotating a response arrow with degree precision, to differentiate between trials reported in the carrier and aperture cues direction, but to also account for trials reported at intermediate trajectories between these directions.

Method

Ethics

This study was approved by the Medical Sciences Inter-divisional Research Ethics Committee (IDREC) at the University of Oxford in accordance with the procedures laid down by the University for Ethical Approval of all research involving human participants, under reference R73572/RE001.

Participants

7 participants (3 male) aged 20-29. All participants had normal or corrected to normal vision, and all participants were naïve to the hypothesis.

Apparatus

The experiment was conducted in a dimly lit room. Stimuli were presented on a CRT monitor (Mitsubishi Diamond Pro 2070, 1600 × 1200 pixels), with a maximum luminance of 81.32 cd/m^2, and gamma corrected using the ColorCAL MK II colorimeter (Cambridge Research Systems). The stimuli were rendered and presented, using the MATLAB toolbox 'PsychToolbox' toolbox, in MATLAB R2014a. Viewing distance was set at 180 cm from the monitor and maintained with a chin rest. Viewing was monocular with each participant's right eye covered by an opaque eyepatch.

Stimuli

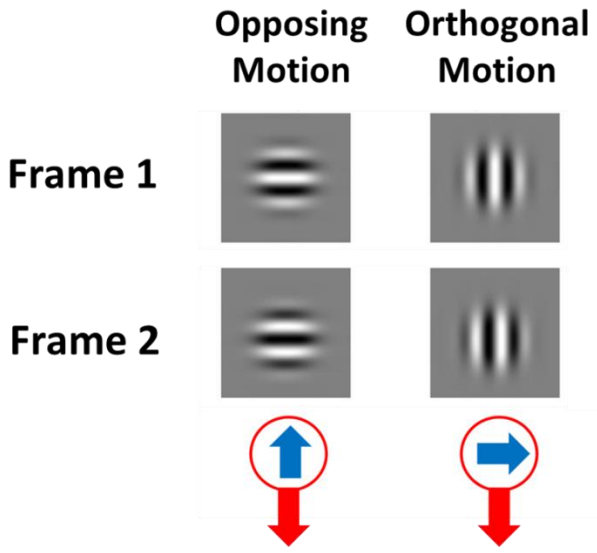


Figure 4.4 Schematic displaying the two stimuli used in this section; the Opposing Motion condition where the phase of the Gabor is simultaneously displaced in the opposite direction to the displacement of the Gaussian envelope of the Gabor, and the Orthogonal Motion stimulus, where the Gabor is oriented parallel with the direction of envelope displacement, and phase and envelope displacement occur simultaneously producing orthogonally oriented motion cues. The envelope displacement is equal in distance to a quarter of the wavelength of the sinusoidal carrier, so equal to the 90^o phase shift of the Gabor.

Stimuli were two-frame apparent motion windowed sinusoid gratings, with each trial consisting of two frames (see Figure 4.4). Stimuli were produced from sinusoidal gratings enveloped by an Gaussian aperture. This produced stimuli with effectively 3 visible cycles, and a gradual taper to the background luminance,

Gratings were presented at 7 spatial frequencies (0.81, 1.33, 2.20, 3.60, 5.92, 9.73 and 16.00 cycles per degree). Because the size of the aperture was set by multiples of the cycles of the current sinusoid grating, the size of the stimulus varied with spatial frequency, but still ensured that the same extent of each grating was visible through the aperture for the different spatial frequencies of the stimulus. Grating phase was randomly assigned on each trial, and stimuli were presented at 100% contrast. The gratings were presented at a randomly assigned orientation, between 0 and 360 degrees, and had the same orientation on both the first and second frames of each trial.

The experiment contained two stimulus conditions, both described previously in Chapter 2. The 'Opposing Motion' condition had the aperture displacement and phase shift in opposing directions, e.g. a vertical grating with a phase shift to the left and aperture displacement to the right. The 'Orthogonal Motion' condition had the aperture displacement and phase shift in orthogonal directions, e.g. a vertical grating with a phase shift to the left and aperture displacement downwards. The nME analysis of the two stimuli, reported in Chapter 2.4.1 and reproduced in Table 4.3, suggests that the orientation of the grating relative to the position cue displacement may affect the strength of the perceived position cue, which has approximately twice the nME value in the Orthogonal than the Opposing Motion condition. Psychophysically, the opposite behaviour has been seen with Rider *et al.* (Rider *et al*, 2014) reporting that drifting Gabors were perceived with a higher velocity when the orientation of the grating was parallel to the position displacement, despite the lack of significant difference in acuity in static horizontal and vertical gratings in normal observers (Meiusi *et al.*, 1993). The nME analysis also suggests that there may be a difference in the relative strength of the carrier cue caused by the different grating orientations, with a higher nME value for downward motion (nME = -0.4256) in the Opposing Motion condition than the nME value for horizontal motion in the Orthogonal condition (nME = 0.2761).

Table 4.3. Table reporting nME values of example Opposing and Orthogonal Motion stimuli, from Chapter 2.4.1.

Cue Arrangement	First Order		Second Order	
	XT	YT	XT	YT
Opposing	-0.0796	**-0.4256**	-0.0355	**0.3127**
Orthogonal	**0.2761**	-0.1233	-0.0177	**0.6284**

Procedure

Each trial consisted of two frames of the stimulus, each presented for 250ms, and followed by a screen matched to the background luminance of the stimulus frames with a centrally presented arrow. The arrow was displayed at maximum luminance (81.33 cd/m^2), which was 2 degrees of visual angle long, and the head spanned a width of 0.5 degrees of visual angle, while the width of the arrow outline was 0.15 degrees of visual angle. The initial orientation of the arrow was randomised on each trial.

Following each trial, participants used the centrally presented arrow to report the perceived motion direction of the stimulus. Participants rotated the arrow using the left and right arrow keys on the keyboard, where each button press corresponded to a change in arrow orientation of 1 degree, updated on each frame. Participants submitted the arrow orientation using the spacebar, and this initiated the next trial.

Fixation was encouraged by use of a 0.25 by 0.25 degrees of visual angle white cross, presented centrally for 150 ms before each trial. Fixation was not recorded in these experiments, due to the COVID-19 pandemic restrictions on laboratory testing conditions, where the repeated calibrations needed for recording fixation using the EyeLink 1000 would not have been permissible because of the necessary proximity between the experimenter and participant during calibration. Participants could abort trials on which they felt their fixation was poor, and those aborted trials would be randomly shuffled into the remaining trials; this happened fewer than 10 times for each participant.

There were 50 trials for each of the 7 spatial frequencies and each of the 6 stimulus conditions, 2,100 trials in total. The experiment was run across two sessions, with trials arranged into 14 blocks of 150 trials. Participants were encouraged to have a five-minute break between blocks. Each session took approximately two hours, with the total experiment length of four hours. Sessions took place within the same week, but not within a 24-hour period.

Results

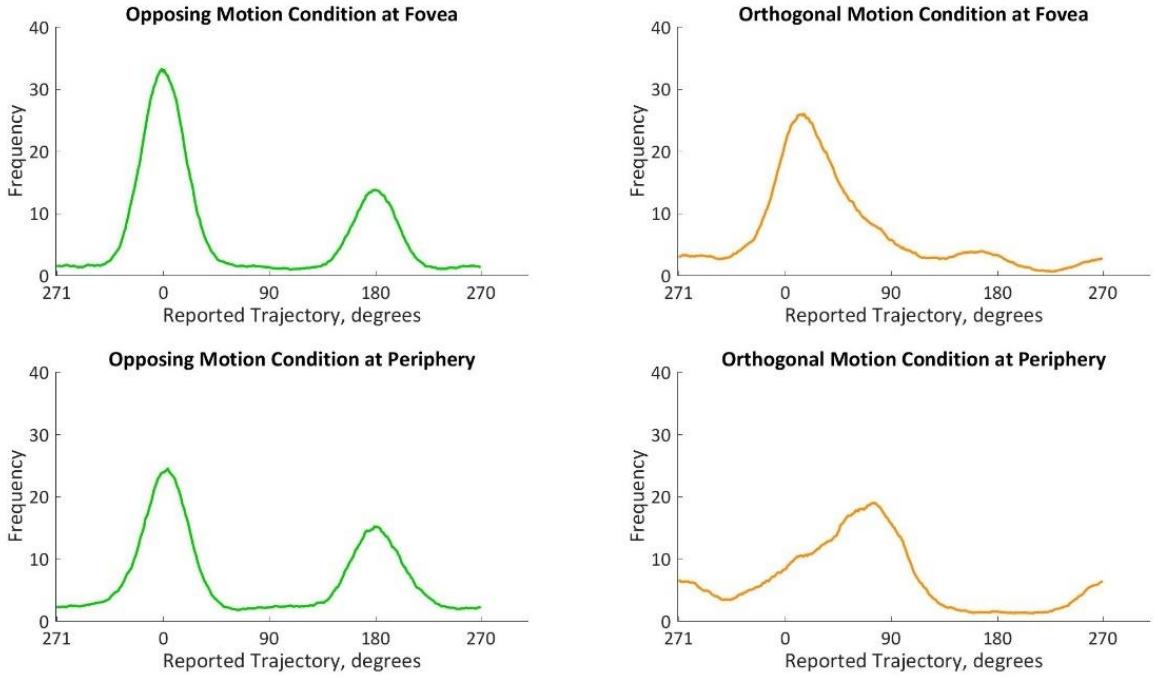


Figure 4.5. This figure displays the aggregated frequency distributions of reported motion trajectories across participants and spatial frequencies. The aperture cue direction is at 0° for both conditions, and the carrier cue direction is at 180° in the Opposing Motion condition and 90° in the Orthogonal Motion condition. The green curves refer to the Opposing Motion condition, and the orange curves to the Orthogonal Motion condition, with the top row reporting data collected foveally, and the bottom row reporting data collected peripherally.

Participants reported the perceived motion of stimuli using a rotating arrow. I have taken the difference between the position cue direction and the reported direction as our dependent variable, and refer to this as the reported 'trajectory'. I can analyse these trajectories in terms of the number of trajectories reported in a particular direction, the circular mean of the trajectories, or as a frequency distribution. To produce the frequency distribution of the trajectories, I calculated the number of times a participant had reported each angle in the 50 trials. This could then be plotted as a frequency distribution against the trajectory angle (see Figure 4.5). Although the data are circular, the frequency distributions can be plotted on a linear scale plotting from -90 degrees to +269 degrees. I chose to 'cut' the data at 270° rather than 180°, so that the Opposing Motion condition, which has 180° separation between the carrier and aperture cue trajectories, could be plotted as the sum of two Gaussians without interruption. This transformation of the frequency distribution of the circular trajectory data onto a linear scale enables the use of traditional linear statistics.

In the Orthogonal condition trials were presented with the conflicting carrier cue randomly allocated to either 90^{o} and 270^{o} from the direction of the aperture cue. For these analyses, I had replotted data so that the carrier cue was consistently plotted at 90°. The Opposing Motion condition required no such transformation. Across all analyses in this section, I excluded the data at the highest spatial frequency, 16 cpd at fovea and 5.92 cpd at 10 degrees eccentricity. The data at these two spatial

frequencies were too noisy, with peaks in the frequency distribution of trajectories emerging only when data were aggregated across participants.

I analysed the data by fitting and comparing different models to the frequency distribution of reported motion trajectories. If reported motion was solely reported in the aperture cue direction when viewing the Opposing or Orthogonal conditions, then I would expect the frequency distributions of participants' reported trajectories to be symmetrical around 0°, and described by a Gaussian distribution. A deviation from this model could suggest that participants were reporting some trials in the carrier cue direction.

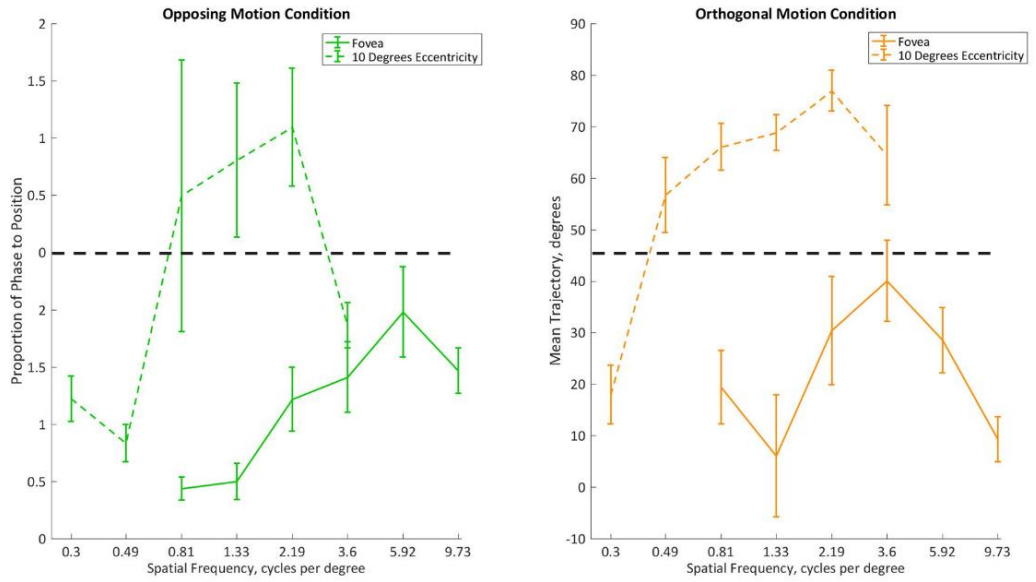


Figure 4.6. Left; Opposing Motion condition. This plot reports the relative proportion of trials reported in the carrier or aperture cue direction. These values were calculated from the ratio of the peak in the carrier cue direction (180°) to the peak in the aperture cue direction (0°), for each participants' data fitted by the nlme model. Values below 0 indicate a greater proportion of trials reported in the aperture cue direction, and values above 0 indicate that a greater proportion of trials were reported in the carrier cue direction. These proportions were calculated for each participant, and mean proportion is displayed here, with error bars representing the standard error of the mean. Right; Orthogonal Motion condition. This plot reports the location in degrees of the fitted Gaussian to participants' reported motion trajectory, in the nlme model for each participant. The aperture cue was at 0° and the carrier cue at 90°, where 45° represents the midpoint between the carrier and aperture cue directions. These graphs show that both Opposing and Orthogonal motion stimuli were reported in the aperture cue direction when viewed foveally (solid line), and in the carrier cue direction when viewed peripherally (dashed line), for all but the lowest spatial frequencies.

In the Opposing Motion condition I would expect trials reported in the carrier cue direction to be clustered symmetrically around 180°. To assess whether trials were reported in the carrier cue direction I compared the goodness of fit of the data when fitted by a single Gaussian model centred at the aperture cue (0°), with a sum of two Gaussians model, consisting of one Gaussian centred at

the aperture cue direction and one Gaussian centred at the carrier cue (180°) direction. The relative height of the Gaussians fitted at 0^o and 180^o for each participant, at each spatial frequency, for foveal and peripheral data, could be compared to map the changing distributions of trajectories . These data are displayed in the left hand panel of Figure 4.6, and whilst this analysis is not presented in this chapter, this is a useful illustration of the change from trials reported predominantly in the aperture cue direction when viewed foveally, to being reported predominantly in the carrier cue direction when viewed at 10 degrees eccentricity.

In the Orthogonal condition, the carrier cue was at the 90° direction relative to the aperture cue at 0^o. Visual inspection of the data suggested that the data were not bimodal, with a peak at the carrier and aperture cues, but rather that trajectories were reported as a weighted integration between the carrier and aperture cues. This can be modelled by a single Gaussian where the mean parameter is situated at any location between 0° and 90°. If there were a significant number of trials reported in the carrier cue, the data would be better fitted by the single Gaussian with a 'free' mean, where the position of the mean parameter would be greater than 0, and approached 90° with increasing reports in the carrier cue direction. The location of the Gaussian along the distribution of reported trajectories corresponds to the mean trajectory reported by participants. These mean locations can be compared for each of the spatial frequencies tested, and between the foveal and peripheral testing locations, as shown in Figure 4.6 (right hand panel). This analysis is not included in this chapter, for brevity, but Figure 4.6 is a useful illustration of the changing location of the fitted Gaussian.

The equations used in these analyses are listed below:

Sum of two Gaussians (double Gaussian) with one Gaussian fixed at 0° and the other at 180° (comparison for Opposing Motion);

$$\text{frequency} \sim a1*\exp(-((\text{degrees})/c1)^2) + a2*\exp(-((\text{degrees}-180)/c2)^2) + d$$

where a is the height of the Gaussian, b (if included) is the mean of the Gaussian and specifies the location of the peak along the degrees axis, c is the width of the Gaussian, and d refers to an additional intercept, which accounts for the noise in the frequency data.

Single Gaussian with a fixed mean at 0^o and included intercept (base model);

$$\text{frequency} \sim a*\exp((-\text{degrees}/c)^2) + d$$

Single Gaussian with a variable mean and included intercept (comparison for Orthogonal Motion);

$$\text{frequency} \sim a*\exp((-\text{degrees}-b\ /c)^2) + d$$

Foveal Results

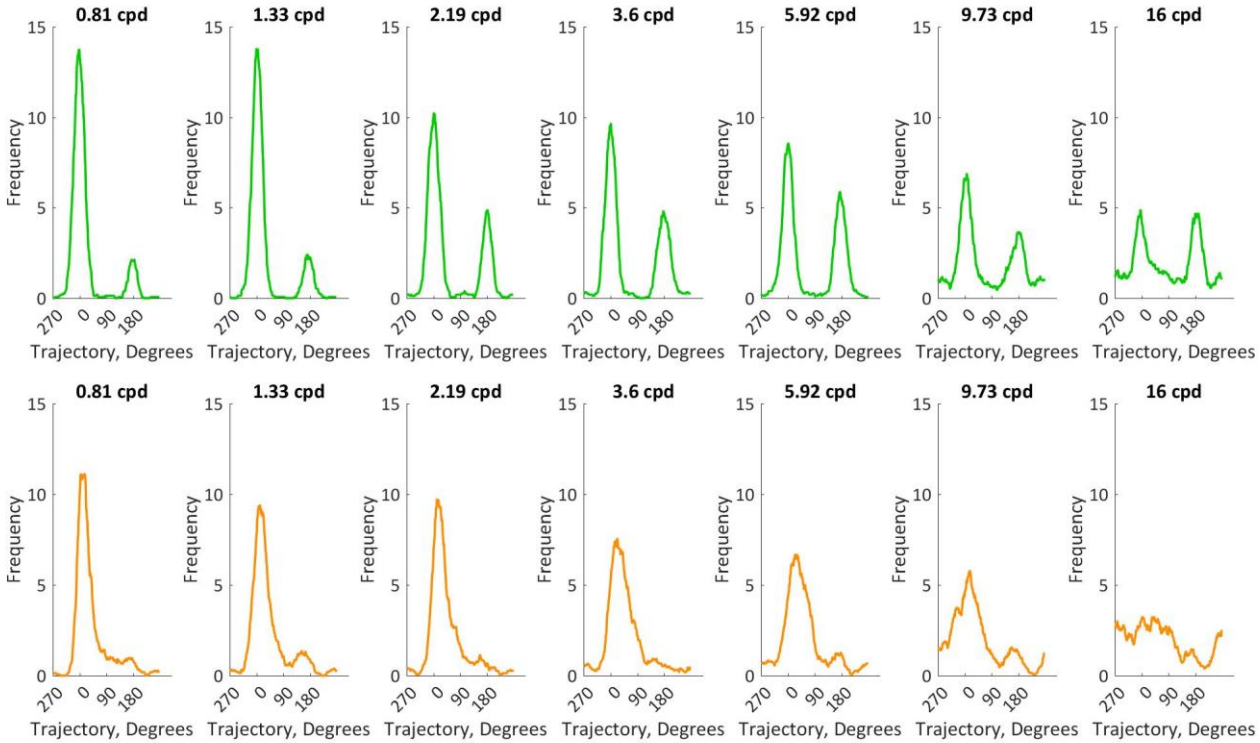


Figure 4.7. This figure displays the distribution of reported motion trajectories collected at the fovea across the 7 participants, for the Opposing Motion condition (green; top row), and the Orthogonal Motion condition (orange; bottom row), with each column displaying data for the spatial frequency that these Gabors were displayed at. I fitted the distribution of participants' reported motion trajectories with a smoothing window of 15 degrees for display here, but analysis was carried out on the un-smoothed distributions of reported motion trajectories. For both Opposing and Orthogonal conditions, motion is predominantly reported in the aperture cue direction, as shown by the peak at 0° for the frequency of trials reported at each direction.

Opposing Motion at Fovea

The trajectories reported when viewing the Opposing Motion condition at the fovea clustered around 0° and 180° (Figure 4.7, top row). Using the nlme package in R (Pinheiro *et al.*, 2013), I fitted the Opposing motion data by the single Gaussian model centred at 0°, and by the sum of two Gaussians model centred at 0° and 180°, and included spatial frequency and participant as fixed effects. I compared goodness of fit, using the 'anova' function in R, and found that there was a significantly better fit of the data by the double Gaussian model compared to the single Gaussian model, (Log Likelihood Ratio = 20313.71, p <.001).

The coefficients of the double Gaussian model are reported in Table 4.4. There was substantial variation in the a2 parameter, the height of the Gaussian fixed at the carrier cue trajectory (180°), which was approximately a third of the height of the a1 parameter. The width of the Gaussians at 0° and 180° (c1 and c2 parameters) were very similar, so it can be concluded that approximately three times more trials were reported in the aperture cue direction than the carrier cue direction.

Table 4.4. Table reporting the parameters of the sum of two Gaussians model, that best fit the frequency distribution of reported motion trajectories in the Opposing Motion condition when viewed foveally. Parameter a1 refers to the height of the Gaussian centred at 0°, and a2 the Gaussian centred at 180°. Parameter c1 refers to the width of the Gaussian centred at 0°, and c2 the width of the Gaussian centred at 180°. Parameter d refers to the intercept of the model, representing the noise in the frequency distribution data.

Parameter	Coefficient	Std. error	Std. dev. Across spatial frequency	Std. dev. Across participants	t value	df	*p* value
a1	0.731	0.080	0.174	0.247	9.076	15074	<.001*
c1	26.021	0.744	0.002	4.739	34.978	15074	<.001*
a2	0.262	0.039	0.064	0.188	6.718	15074	<.001*
c2	27.000	1.947	3.982	6.185	13.868	15074	<.001*
d	0.014	0.007	0.016	0.022	1.903	15074	.057

I assessed whether there was a significant effect of spatial frequency by fitting the double Gaussian model with both spatial frequency and participant as fixed effects, and comparing the goodness of fit to a reduced model, where only participant was included as a fixed effect. This found a significant reduction in goodness of fit, when spatial frequency was not included as a fixed effect (Log Likelihood Ratio = 24601.23, p < .001). This demonstrates a significant effect of spatial frequency,

where there is a significant difference in the proportion of trials reported in the carrier and aperture cue directions across the Opposing Motion Gabor stimuli presented at different spatial frequencies.

Because the widths of the Gaussians were matched between the carrier and aperture cue directions, the height of the Gaussians can be used to investigate the relative proportion of trials reported in the carrier and aperture cue directions. I replicated the comparison of the single and sum of two Gaussians model for the data separately for each spatial frequency, to investigate whether there was a significant number of trials reported in the carrier cue direction at each spatial frequency, and the relative proportion of trials reported in the carrier and aperture cue directions. These results are reported in Table 4.5, where a significant p value and positive Log Likelihood Ratio indicates a significantly better fit of the double than the single Gaussian model, for data at that spatial frequency.

There was a significantly better fit to the data using the sum of two Gaussians model at every spatial frequency, and the height of the Gaussian centred at the carrier cue (a2 parameter) increased with spatial frequency, demonstrating more trials reported in the carrier cue direction at higher spatial frequencies.

Table 4.5. Table describing the sum of Gaussians model fitted to the frequency distribution of reported direction trajectories to the Opposing Motion Gabor stimulus, fitted at each spatial frequency tested. The relative proportion of trials reported in the cue directions can be approximated by the heights of the Gaussians fitted to the trajectories reported in the aperture cue direction (a1) and the carrier cue direction (a2). The table also reports whether the trajectory data was significantly better fit by the sum of Gaussians model than the single Gaussian model.

Spatial frequency (cpd)	a1	a2	AIC	Log Likelihood	df	Log Likelihood Ratio	*p* value
0.81	0.977	0.153	-7380.245	3701.122	11	1323.132	<.0001*
1.33	0.976	0.162	-7830.580	3926.290	11	2043.799	<.0001*
2.19	0.730	0.335	-8568.158	4295.079	11	4767.215	<.0001*
3.60	0.671	0.328	-6002.51	3012.255	11	2526.05	<.0001*
5.92	0.602	0.398	-8758.691	4390.345	11	5228.395	<.0001*
9.57	0.555	0.269	-6025.175	3023.587	11	2514.76	<.0001*

Orthogonal Motion at Fovea

The trajectories reported when viewing the Orthogonal Motion condition at the fovea were reported between 0° and 90°, clustering closer to 0° (Figure 4.7, bottom row). For the Orthogonal Motion condition I fitted data with a single Gaussian model with a mean fixed at 0°, and data fitted with a single Gaussian model where the mean was not fixed. If there was a significant effect of the carrier cue on the trajectories in the Orthogonal condition, then there should be a significantly better fit of the data by the model with the free roaming mean parameter, and the mean of the Gaussian should be substantially higher than 0°.

I fitted both models using the non-linear mixed effects package in R, and included spatial frequency and participant as fixed effects. There was a significantly better fit to the data when the mean of the Gaussian was allowed to vary, (Log Likelihood Ratio = 15806.08 $p < .001$) compared to the fixed mean model.

The parameters of the model, with their respective coefficients, significance and standard deviation across fixed effects, are reported in Table 4.6. Each of the four parameters contributed significantly to the Gaussian model, and the mean was substantially higher than 0°, at 20.82° (SE = 4.25°), demonstrating a significant effect of the conflicting carrier cue, but with a mean trajectory reflecting a greater relative strength of the aperture than of the carrier cue as this value is less than 45°, which is the midway point between the aperture cue (0°) and carrier cue directions (90°).

Table 4.6 Table reporting the parameters of Gaussians model that best fit the frequency distribution of reported motion trajectories in the Orthogonal Motion condition when viewed foveally. Parameter a1 refers to the height of the Gaussian, parameter b1 to the mean of the Gaussian, parameter c1 refers to the width of the Gaussian and parameter d refers to the intercept of the model, representing the noise in the frequency distribution data.

Parameter	Coefficients	Std. error	Std. dev. Across Spatial Frequency	Std. dev. Across participants	t value	df	*p* value
a1	0.794	0.084	0.180	0.195	9.479	15075	<.001*
b1	20.821	4.248	8.885	10.810	4.902	15075	<.001*
c1	39.172	5.017	10.409	13.031	7.808	15075	<.001*
d1	0.006	0.002	<0.001	0.009	3.042	15075	<.001*

I assessed whether there was a significant effect of spatial frequency, by fitting the model with both spatial frequency and participant as fixed effects, and comparing the goodness of fit to a reduced model, where only participant was included as a fixed effect. This found a significant reduction in goodness of fit when spatial frequency was not included as a fixed effect (Log Likelihood Ratio = 15458.12, $p < .001$), demonstrating a significant difference in the proportion of trials reported in the carrier and aperture cue directions across spatial frequency.

I explored the effect of spatial frequency on reported direction trajectory further, by fitting the Gaussian model with the free mean to the data separately at each spatial frequency. I ran a Bonferroni corrected one sample t test on the Gaussian peak locations (mean, b1 parameter), to test whether the mean location was significantly different from 0^{o}, which would demonstrate a significant effect of the conflicting carrier cue on reported motion trajectory. The results of these individual t tests are reported in Table 4.7. Using the corrected alpha level of 0.0083, the fitted Gaussians peaks were significantly higher than 0° for the lowest five of the six spatial frequencies, demonstrating a significant effect of the carrier cue on reported motion trajectories in the Orthogonal Motion Gabor stimulus at fovea.

Table 4.7. Table displaying the results of Bonferroni corrected one sample t tests run on single Gaussian models fitted to reported direction trajectory frequency distributions for the Orthogonal Motion Gabor when viewed foveally, for each spatial frequency tested

Spatial frequency (cpd)	b1	Std.	t value	df	*p* value
0.81	12.787	4.470	6.396	6	0.003
1.33	14.098	5.535	5.696	6	0.005
2.19	21.333	5.098	9.357	6	0.001
3.60	31.921	11.413	6.254	6	0.003
5.92	34.751	15.910	4.884	6	0.008
9.73	10.358	13.966	1.659	6	0.173

Peripheral Results

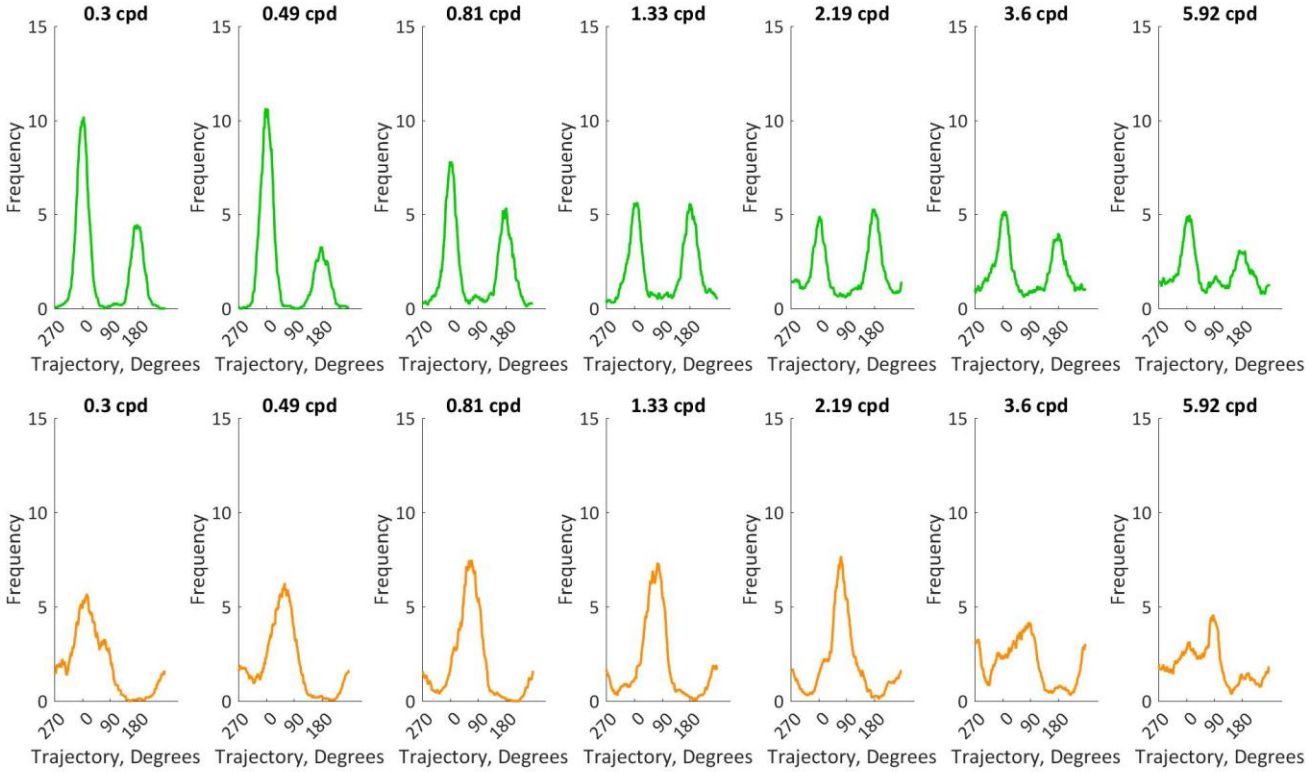


Figure 4.8. This figure displays the distribution of reported motion trajectories collected at 10 degrees eccentricity (peripheral condition) across the 7 participants for the Opposing Motion condition (green; top row), and the Orthogonal Motion condition (orange; bottom row), with each column displaying data for the spatial frequency that these Gabors were displayed at. I fitted the distribution of participants' reported motion trajectories with a smoothing window of 15 degrees for display here, but analysis was carried out on the un-smoothed distributions of reported motion trajectories. Motion was reported in both carrier and aperture cue directions, as indicated by the peaks at 0° and 180° in the Opposing Motion condition, and the peaks at 0° and 90° in the Orthogonal Motion condition.

Opposing Motion condition at 10 degrees eccentricity

The trajectories reported when viewing the Opposing Motion condition at 10 degrees eccentricity clustered around 0° and 180° (Figure 4.8, top row). I compared the goodness of fit of the reported trajectories in the Opposing Motion condition at 10 degrees eccentricity for the data fitted with a double Gaussian with means fixed at 0° and 90°, and with a single Gaussian with means fixed at 0°. Data were fitted using a non-linear mixed effects model, and included spatial frequency and participant as fixed effects. There was a significantly better fit of the data by the double Gaussian model compared to the single Gaussian model, (Log Likelihood Ratio = 21654.92 , p <.001).

All five parameters for the double Gaussian model, with their respective coefficients, standard error across the fixed effects, contributed significantly to the sum of two Gaussians model. These are reported in Table 4.8. The peak widths which represented the precision of the reported trajectories were very similar for both the aperture and carrier cue directions, and there was a greater peak height at the 0° located Gaussian (a1 = 0.488) than the 180° Gaussian (a2 = 0.297). This indicates that more trials were reported in the aperture cue direction than the carrier cue direction, although the difference between peak heights was less than observed in the foveal data.

Table 4.8. Table reporting the parameters of the sum of two Gaussians model, that best fit the frequency distribution of reported motion trajectories in the Opposing Motion condition, when viewed at 10 degrees eccentricity. Parameter a1 refers to the height of the Gaussian centred at 0°, and a2 the Gaussian centred at 180°. Parameter c1 refers to the width of the Gaussian centred at 0°, and c2 the width of the Gaussian centred at 180°. Parameter d refers to the intercept of the model, representing the noise in the frequency distribution data.

Parameter	Coefficient	Std. error	Std. dev. across Spatial Frequency	Std. dev. across participants	t value	df	*p* value
a1	0.488	0.080	0.185	0.160	6.130	15074	<.001*
c1	29.211	1.646	0.002	10.523	17.741	15074	<.001*
a2	0.298	0.023	0.029	0.128	12.924	15074	<.001*
c2	29.679	1.631	0.004	10.373	18.199	15074	<.001*
d	0.031	0.010	0.024	0.021	3.009	15074	0.003*

I assessed whether there was a significant effect of spatial frequency by fitting the model with both spatial frequency and participant as fixed effects, and comparing the goodness of fit to a reduced model where only participant was included as a fixed effect. This found a significant reduction in goodness of fit when spatial frequency was not included as a fixed effect (Log Likelihood Ratio =

17833.27, $p < .001$).. This demonstrates a significant effect of spatial frequency, where there is a significant difference in the proportion of trials reported in the carrier and aperture cue directions across the Opposing Motion Gabor stimuli presented at different spatial frequencies.

I explored this significance difference between spatial frequencies further, running this analysis at each of the spatial frequencies separately. These model comparison are reported in Table 4.9, where a significant p value and a positive Log Likelihood Ratio indicates a significantly better fit of the double Gaussian than the single Gaussian model. There was a significantly better fit to the data using the sum of two Gaussians model at every spatial frequency, and the height of the Gaussian centred at the carrier cue increased with spatial frequency, demonstrating more trials reported in the carrier cue direction at higher spatial frequencies.

Table 4.9 Table describing the sum of Gaussians model fitted to the frequency distribution of reported direction trajectories to the Opposing Motion Gabor stimulus, fitted at each spatial frequency tested. The relative proportion of trials reported in the carrier and aperture cue directions can be approximated by the heights of the Gaussians fitted to the trajectories reported in the aperture cue direction (a1) and the carrier cue direction (a2). The table also reports whether the trajectory data was significantly better fit by the sum of Gaussians model than the single Gaussian model.

Spatial Frequency (cpd)	a1	a2	AIC	Log Likelihood	df	Log Likelihood Ratio	*p* value
0.30	0.719	0.322	-8252.692	4137.346	11	4118.529	<.001*
0.49	0.750	0.222	-8204.959	4113.480	11	2994.282	<.001*
0.81	0.544	0.353	-8579.591	4300.796	11	4402.615	<.001*
1.33	0.358	0.350	-8340.894	4181.447	11	4491.283	<.001*
2.19	0.327	0.374	-6081.939	3051.970	11	3182.312	<.001*
3.60	0.287	0.214	-7734.293	3878.147	11	2382.356	<.001*

Orthogonal Motion condition at 10 degrees eccentricity

The trajectories reported when viewing the Orthogonal Motion condition at the fovea were reported between 0^o and 90^o, with the location of the highest frequency trajectories falling between 0^o and 90^o, but shifting towards 90^o with increasing spatial frequency (Figure 4.8, bottom row). To investigate whether there was a significant effect of the conflicting carrier cue on perceived motion direction in the Orthogonal Motion condition, I compared data fitted with a single Gaussian model

with a mean fixed at 0°, and data fitted with a single Gaussian model where the mean is not fixed. I fitted both models using non-linear mixed effects modelling, and included spatial frequency and participant as fixed effects. There was a significantly better fit to the data when the mean of the Gaussian was allowed to vary, (Log Likelihood Ratio = 14479.41, $p < .001$), compared to the fixed mean model.

The parameters of this model, with the varying mean, are reported in Table 4.10, with their respective coefficients, and standard deviation across the fixed effects. Each of the four parameters contributed significantly to the Gaussian model, and the mean was substantially higher than 0°, at 58.50° (SE = 7.80°), demonstrating a significant effect of the conflicting carrier cue, and with a mean trajectory reflecting a greater relative strength of the carrier than the aperture cue, as the mean trajectory was greater than 45° , the midway point between the aperture cue direction at 0° and the carrier cue direction at 90° .

Table 4.10 Table reporting the parameters of Gaussians model that best fit the frequency distribution of reported motion trajectories in the Orthogonal Motion condition when viewed at 10 degrees eccentricity. Parameter a1 refers to the height of the Gaussian, parameter b1 to the mean of the Gaussian, parameter c1 refers to the width of the Gaussian and parameter d refers to the intercept of the model, representing the noise in the frequency distribution data.

Parameter	Coefficient	Std. error	Std. dev. across Spatial Frequency	Std. dev. across participants	t value	df	*p* value
a1	0.431	0.036	0.061	0.167	12.004	15075	<.001*
b1	58.496	7.798	18.072	16.307	7.501	15075	<.001*
c1	35.061	4.152	0.008	26.842	8.445	15075	<.001*
d1	0.052	0.006	0.008	0.034	8.423	15075	<.001*

I assessed whether there was a significant effect of spatial frequency, by fitting the model with both spatial frequency and participant as fixed effects, and comparing the goodness of fit to a reduced model where only participant was included as a fixed effect. This found a significant reduction in goodness of fit when spatial frequency was not included as a fixed effect, (Log Likelihood Ratio = 13872.73, $p < .001$). This demonstrates a significant difference in the proportion of trials reported in the carrier and aperture cue directions at the different spatial frequencies tested.

I explored this further by fitting the Gaussian model with the free mean to the data separately at each spatial frequency, using an nlme with participant as the fixed effect, for each of the six spatial frequencies. I ran a Bonferroni corrected one sample t test on the Gaussian peak locations (mean, b1 parameter), to test whether the mean location was significantly different from 0^{o}, which would demonstrate a significant effect of the conflicting carrier cue on reported motion trajectory. These t tests are reported in Table 4.11, where using the corrected alpha level of 0.0083, the fitted Gaussians peaks were significantly higher than 0° for five of the six spatial frequencies, barring the lowest spatial frequency tested (0.3 cpd), demonstrating a significant effect of the carrier cue on reported motion trajectories in the Orthogonal Motion condition when viewed at 10 degrees eccentricity. The mean trajectory increased with spatial frequency.

Table 4.11. Table displaying the results of Bonferroni corrected one sample t tests run on single Gaussian models fitted to reported direction trajectory frequency distributions for the Orthogonal Motion Gabor when viewed at 10 degrees eccentricity, for each spatial frequency tested.

Spatial Frequency (cpd)	b1	Std.	t value	df	*p* value
0.30	14.678	13.400	2.683	6	0.0437
0.49	50.577	11.433	10.836	6	0.001*
0.81	62.088	6.586	23.091	6	<.001*
1.33	66.444	7.226	22.522	6	<.001*
2.19	73.783	6.688	27.023	6	<.001*
3.60	60.038	24.945	5.895	6	0.002*

Discussion

At the fovea, there are a significant number of trials reported in the carrier cue direction in both the Opposing and Orthogonal Motion conditions. This is demonstrated by the significantly better fit of the sum of two Gaussians model to the frequency distribution of reported motion trajectories compared to both a single Gaussian model centred solely on the aperture cue direction, in the Opposing Motion condition, and the significantly better fit of the Gaussian model with the free and not fixed mean in the Orthogonal Motion condition. This directly contradicts Shapiro *et al.*'s 'feature blur' hypothesis (Shapiro *et al.*, 2011), who proposed that motion was solely reported in the aperture cue direction at fovea. This is due to the increased relative strength of the aperture cue relative to the carrier cue because of the large displacement in position caused by the many-frame

display. This also suggests that in Chapter 4.2, there may have been trials in the Opposing Motion condition reported in the carrier cue direction, which we could not detect because of the 2afc methodology.

In the periphery a significant number of trials have been reported in the carrier cue direction in both the Opposing and Orthogonal Motion conditions, with a similarly better fit to the reported trajectories in the Opposing Motion condition by the sum of two Gaussians model, and in the Orthogonal Motion condition by the Gaussian model with the free mean. The skewed shape of the Orthogonal condition when plotted across spatial frequencies (Figures 4.9 and 4.10), is the result of the changing peaks and widths of the single Gaussians at each eccentricity. In Chapter 2.3 , there was a significant reversal of motion for the conflicting cue conditions, where at high spatial frequencies significantly more trials were reported in the carrier cue direction than the aperture cue direction. This effect is replicated here, with the mean trajectories in the Orthogonal Motion condition when viewed peripherally reported above 45° for all but the lowest spatial frequency. Similarly, in the Opposing Motion condition when viewed peripherally, the heights of the Gaussians of the frequency distribution of reported motion trajectories at the carrier and aperture cue directions, indicating the number of trials reported in these directions, were higher for the carrier cue direction than the aperture cue direction at the higher spatial frequencies.

In this chapter I have reported an effect of spatial frequency on reported motion direction at both fovea and 10 degrees eccentricity, for both the Opposing and Orthogonal Motion conditions. This was demonstrated statistically, where there was a significant improvement in the goodness of fit of the model of the distribution of reported trajectories when spatial frequency was included as a fixed effect, which enabled the model to vary the parameters for each spatial frequency. To investigate the relationship between spatial frequency and reported motion direction in the Opposing Motion condition, I extracted heights of the Gaussians fitted at the carrier and aperture cue directions. The proportion of these Gaussians heights is displayed in Figure 4.6, which shows that the relative proportion of trials reported in the carrier cue direction increased with spatial frequency, as predicted by my spatial scaling hypothesis, at both fovea and 10 degrees eccentricity. To investigate the relationship between spatial frequency and reported motion direction in the Orthogonal Motion condition, I extracted the means of the Gaussian model fitted to the reported motion trajectories. This indicates that motion is reported in a diagonal trajectory, where the trajectory represents the relative strength of the aperture and carrier cues, with trajectories closer to 0° representing a greater strength of the aperture cue, and trajectories closer to 90° representing a greater relative strength of the carrier cue. The mean trajectory increased with spatial frequency for both the fovea

and peripheral viewing conditions, demonstrating an increased weighting of the carrier cue with increasing spatial frequency.

These data support my spatial scaling explanation of the change in perceived motion for conflicting cues where, at both the fovea and 10 degrees eccentricity, reported motion is biased towards the aperture cue direction at low spatial frequencies and the carrier cue direction at high spatial frequencies.

4.4 Discussion

In both of the two conflicting cue experiments presented in this chapter, I found that reported motion was consistent with the spatial frequency dissociation for carrier and aperture cues, proposed in Chapter 3 using Gabors containing either a carrier or an aperture cue. At the 4 viewing locations tested in the visual field, motion is biased towards the aperture cue direction at lower spatial frequencies and towards the carrier cue direction at higher spatial frequencies. In Chapter 4.2 the observed reversal in reported motion direction was well predicted by the overlapping but displaced spatial frequency sensitivity curves from Chapter 3, with motion reversal occurring for the spatial frequencies where sensitivity to the aperture cue was low but remained high for the carrier cue. My spatial scaling hypothesis also explains the reported reversals of motion from the aperture cue direction at fovea, to the aperture cue direction in the periphery, as due to the shift of the spatial frequency sensitivity curves towards lower spatial frequencies with increasing eccentricity. Stimuli that fall under the aperture cue dominant part of the spatial sensitivity function at fovea, will fall under the carrier cue dominant section in the periphery.

I suggest that the foveal contribution of the carrier cue is very small, compared to the aperture cue contribution, particularly at longer frame presentations , which is why it was excluded in Shapiro's 'feature blur' hypothesis (Shapiro *et al.*, 2011). However, while the proportion of trials reported in the carrier cue direction is smaller at fovea than periphery, we see the same relationship with reported motion direction and spatial frequency in both the stimulus conditions, and in the periphery. This suggests that the same mechanism for integrating the carrier and aperture cues is operating foveally and peripherally, contrary to Shapiro's *et al.* (2011) suggestion.

One caveat in my comparison of the spatial scaling data to the reported motion in the conflicting cues, is that, in the spatial scaling experiment with Carrier and Aperture conditions reported in Chapter 3, performance across all spatial frequencies, at each eccentricity, is higher in the Carrier condition. If performance were an accurate measure of the relative strength of the carrier and aperture cues, this would suggest that performance at all eccentricities should be predominantly reported in the carrier cue direction. This is not the case, as can be seen in the reported motion

direction for both the two conflicting cue experiments. This suggests that assessing proportion correct in the 2afc task may not be the best measure of the relative strength of the two cues. If models that predict the proportion of trials reported in the carrier and aperture cue directions in the Opposing and Orthogonal conditions are required, then mapping the contrast sensitivity functions for the carrier and aperture conditions may provide a better estimate.

This overall bias to the aperture cue, seen in even these two-frame Gabors, suggests that while the amount of displacement has been matched between the carrier and aperture cues, in the single cue and the conflicting cue conditions, there is still a difference in the relative strength of the two cues. This bias could be due to the nature of the cues; the carrier cue is produced from phase displacements of the sinusoid grating, which is modulated in contrast by the Gaussian aperture, whereas the aperture cue is produced from displacements of the centroid of the Gabor. Alternatively, the bias could reflect an internal weighting of phase and position cues. Reverse phi stimuli produce the phase and position cues from the same features in the stimulus, where the displacement of the features produces the position cue, and the simultaneous reversal of polarity of these features produces the phase cue. At the fovea, reverse phi stimuli can consistently be perceived in either the phase or position cue direction, depending on the size of the displacement of the features between frames (Chubb & Sperling, 1989; Anstis & Rogers, 1975). This suggests that potential motion mechanisms extracting motion from phase and position cues (putative phase and position systems) are both able to dominate the perceived motion at the fovea, and that the bias in the aperture cue direction at the fovea for Gabor based conflicting cues is a result of the different features in the stimuli, where the relative strength of the aperture cue is higher than the relative strength of the carrier cue. This could be demonstrated experimentally by manipulating the stimuli to adapt the relative strength of the two cues. Decreasing the displacement of the Gaussian aperture or changing the aperture profile could reduce the relative strength of the aperture cue. Producing motion reported predominantly in the carrier cue direction at the fovea for Gabor based conflicting cue stimuli would support this interpretation, i.e. that the current dominance of the aperture cue at fovea reflects a difference in the features producing the carrier and aperture motion cues.

This difference in the features producing the carrier and aperture cues in the conflicting cue Gabors could also explain the decline in the use of the carrier cue at the very highest spatial frequencies tested, which partially contradicts my spatial scaling explanation. In both conditions, but especially the Orthogonal condition, performance at the highest spatial frequency at each eccentricity was near-chance, where for many participants Gaussian models were unable to converge, and the peaks at the carrier and aperture cues only emerged when data was plotted in aggregate across the participants. The chance performance at the highest spatial frequencies in Chapter 4.2 could reflect

a return to chance rather than an increase in reports in the aperture cue direction. However, in Chapter 4.3 at the second highest spatial frequencies tested for the stimuli viewed in the periphery, there is a decrease in the relative proportions of trials reported in the carrier and aperture cues in the Opposing Motion condition, and in the mean trajectory in the Orthogonal condition, at both fovea and 10 degrees eccentricity. This decrease does reflect a relative increase in the aperture cue condition at the highest spatial frequency, and could potentially reflect the loss of acuity to the grating based carrier cue, as well as the residual acuity for detecting the Gabor centroid necessary for the aperture cue displacement. This suggests that whilst the spatial scaling explanation does provide an explanation for the main effects seen in the data – the increase in reported motion in the carrier cue direction at higher spatial frequencies, and in periphery – the differences in the features producing the carrier and aperture cues could result in a preserved acuity for the aperture, but not carrier, cue at the highest spatial frequencies.

The change between reported motion trajectory with spatial frequency, and eccentricity, for the conflicting cues Gabors is strong evidence suggesting that there are two competing mechanisms processing motion. Demonstrations of the orthogonal Gabor, 'the Break of the Curveball' which won first prize in the 2009 'Best Illusion of the Year Contest' (Best illusion, 2009), show that when the orthogonal Gabor is presented on a dark or light background, that the diagonal trajectory when viewed peripherally disappears. This change in background luminance means that the Gabor is no longer the same mean luminance as the background, and as such the Gabor itself is now a low spatial frequency phase cue. The lack of diagonal trajectory demonstrates that when the Gaussian displacement and the changing phase of the grating are processed by the same mechanism, motion is perceived in the direction of the dominant phase cue, the overall displacement of the Gabor on each frame, and not in a diagonal trajectory. This explanation can also be applied to the 'Rotation Reversals' stimulus presented in Shapiro *et al.* (Shapiro *et al.*, 2011). This is further evidence that the diagonal trajectory in orthogonal Gabors viewed peripherally reflects the activity of two competing mechanisms, one processing phase and the other processing position cues. Furthermore, this suggests that the position system is using a different mechanism than the phase system – even following a potential full or half wave rectification – since the abolition of the diagonal trajectory, once the mean luminance of the Gabor and background differ, demonstrates that when the phase and position cues in the Gabor are processed by the same mechanism, that there is no change in location of the perceived motion trajectory or global change in location of the Gabor.

The Gabor stimulus was ideal for testing the spatial scaling hypothesis as I could investigate the spatial sensitivity of the carrier and aperture cues separately, in the Carrier and Aperture stimulus conditions presented in Chapter 3, and then measure the relative strength of the two cues in the

Opposing and Orthogonal conditions. These stimuli could also be used more generally to investigate the mechanisms governing the integration of carrier and aperture cues. These Gabor stimuli are ideal for this logical next set of experiments, because the relative strength of the carrier and aperture cues can be manipulated by separately varying the temporal frequency of the grating and Gaussian aperture (Heller *et al.*, 2021), or by the amount of displacement of these components. Unlike other conflicting cue stimuli, the carrier and aperture cues can also be presented in isolation, in Gabors containing solely a phase change in the grating or position displacement of the aperture. This could be essential in developing models of cue integration, as this enables the measurement of both the absolute sensitivity to the cues, and the relative sensitivity, when cues are presented in conflict.

This avenue, investigating cue integration, should also include the assessment of motion perception when the cues in the Gabor are presented congruently, with carrier and aperture cues signalling motion in the same direction. Given the difference in the spatiotemporal sensitivity functions of the carrier and aperture cues, these congruent cues will have different relative especially if the cues are produced from the same features in the image. It is simpler to build a mechanism that treats conflicting and congruent cues the same, integrating between the direction and distance vector of the two inputs without regard to their original difference in trajectory, but this would reduce overall sensitivity. This congruency of cues could be tested in a 'phi' rendering of the reverse phi stimulus, where the polarity is not reversed, and both the phase and position systems would interpret the displacement of the stimulus as motion in the same direction. This was done by Bours *et al.* (2009) who compared phi and reverse phi random dot displays. They found no difference in performance across displacement size or temporal interval, or dot density. Bours *et al.*'s results could be interpreted as a 'winner takes all mechanism', and suggest that whether cues are competing, or congruent, motion is reported in the direction of the cue with the greatest relative strength. The Gabors used in my experiment could be considered as a similar congruent stimulus, which would suggest that the motion discrimination performance assessed for the carrier cues in Gabors at low spatial frequencies is unreliable data.

Both the displacement of the spatial sensitivity curves, and the significant different in displacement threshold gradient with eccentricity, between the Carrier and Aperture conditions, are strong evidence of a psychophysical dissociation between the phase and the position cues. The bias in reported motion directions with spatial frequency and eccentricity for the conflicting cue conditions further supports this interpretation. There may be two separate systems which process the phase and position cues. The difference in spatial frequency sensitivity of these two systems may be reflected in their respective physiological substrates, and this is explored in Chapter 7.

Chapter 5: Perceived motion in reverse phi stimuli

5.1 Introduction

Using Gabor stimuli, I have shown that there is a psychophysical dissociation between the use of the spatial position of the aperture (aperture cue) and phase displacement of the carrier (carrier cue) for perceiving motion direction: reported motion is biased towards the aperture cue direction at low spatial frequencies, and towards the carrier cue direction at high spatial frequencies. This bias is seen to some extent when viewed foveally, and increasingly so when displayed at increasingly eccentric locations in the periphery. These results are potentially evidence of a psychophysical dissociation between two systems; one processing phase cues and biased towards higher spatial frequencies, and one processing position cues and biased towards lower spatial frequencies. However, the Gabor stimuli contain a confound: the change in phase of the carrier produces a corresponding change in the spatial position of the carrier features. Therefore the Gabor results could be driven by the difference in scale of the image characteristics - the position cue of the carrier features compared to the position cue of the aperture - rather than reflecting a dissociation between a change in phase and a change in spatial position. It is impossible to use the motion discrimination results of the Gabor experiments presented in Chapters 3 and 4 to derive conclusions about the possible presence of a motion mechanism extracting phase cues from the carrier cue, as it is impossible to separate the change in phase of the carrier from the change in position of the features in the carrier. However, if there are in fact separate motion processing systems, extracting motion from different cues in the visual input, then spatial frequency driven reversals of perceived motion should occur for other constructions of conflicting cue stimuli, where phase and position cues occur without confounds.

One class of conflicting cue stimuli that could be employed to validate the psychophysical dissociation between the phase and position systems are the 'reverse phi stimuli'. In these stimuli, the displacement of a feature with a simultaneous reversal in contrast produces a conflicting phase and position cue, where the phase cue is in the opposite direction to the displacement of the feature (Anstis, 1970). This reversal of polarity can be considered as reversing the orientation of the luminance in a space-time plot. If a system determines motion direction through the orientation of luminance in a space-time plot, as posited by the motion energy detector mechanism (Adelson & Bergen, 1985), then this reversal of polarity predicts a reversal in perceived motion direction. However, if a system processes motion through tracking the displacement of features in an image, the proposed position system, then the perceived direction of motion should not be disrupted by this change in contrast.

Reverse phi stimuli are excellent conflicting cue stimuli, because the phase and position cues in the stimulus embody the design of pure phase and position cues; the phase cue is unattached to the physical displacement of an object, such as the pedestal stimuli of Van Santen & Sperling (Van Santen & Sperling, 1984). The position cue changes identity between frames, similarly to the quasi-grating stimuli of Lu & Sperling (Lu & Sperling, 1995b), where the identity of features changes on each frame, and to the 'MF1F' of Baro & Levinson (Baro & Levinson, 1988). A range of reverse phi stimulus constructions have been demonstrated, from individual figures (circles; Anstis & Rogers, 1975: bars; Wehrhahn, 2006), to arrays of random dots (Sato, 1989; Edwards & Nishida, 2004; Bours *et al.*, 2009), to quasi-grating stimuli where bars are arranged sequentially to produce the appearance of a square wave grating with a small duty cycle, but only the bars vary in polarity and the background is fixed (Chubb & Sperling, 1989).

Reverse phi stimuli show the same bias in reported motion direction with spatial frequency as I reported for the carrier and aperture cue directions in the conflicting cue Gabors in Chapter 4; towards the phase cue at high spatial frequencies, and towards the position cue at lower spatial frequencies. The spatial frequency bias is dependent on the construction of the reverse phi stimuli. Whilst Anstis & Rogers (1975) reported 10 minutes of arc as the critical value above which motion was perceived in the position cue direction, Smith & Ledgeway (2001) using high-pass filtered reverse phi RDKs of varying element size found a cyclical relationship motion reported in the phase and position cue direction, due to aliasing of frequencies within their stimuli (Smith & Ledgeway, 2001). This also demonstrates a typical confound in reverse phi stimuli, where either the displacement size is varied, or the spatial frequency of the features: Smith & Ledgeway's experiment is one of the few to vary both systematically. For the purposes of this chapter, we are using a fixed relationship between the spatial frequency of the features (grating bars or square), and the displacement amount, where the change in the threshold displacement is the variable of interest.

When displayed foveally, motion was reported in the position cue direction for low spatial frequencies, and in the phase cue direction for high spatial frequencies (Gorea, 1994; Wehrhahn, 2006; Chubb & Sperling, 1989; Benton *et al.*, 1997). Reverse phi stimuli also show a reversal of reported motion direction between fovea and periphery (Chubb & Sperling, 1989; Edwards & Nishida, 2004), which is similar to the biased motion in conflicting cue Gabor stimuli (Shapiro *et al.*, 2010) where motion was reported in the position cue at fovea, and biased towards the phase cue in periphery. I explained these results in the Gabor stimuli as a consequence of the change in spatial sensitivity functions between fovea and periphery, where this change in sensitivity is different for the carrier and aperture cues. This similarity between reverse phi stimuli and Gabor stimuli is

suggestive that the reverse phi stimuli are also processed by two psychophysically distinct systems that process the phase and position cues.

In this chapter, I investigate the reported motion direction in reverse phi stimuli, across ranges of spatial frequency and at several locations in the visual field. If the spatial frequency dissociation reported for the carrier and aperture cues in Gabor stimuli in Chapter 3, accurately represents a spatial frequency dissociation in the phase and position systems, then this result predicts that reported motion in reverse phi stimuli will be in the position cue direction at low spatial frequencies and in the phase cue direction at high spatial frequencies. This would provide further evidence for my spatial frequency based dissociation of the phase and position systems.

5.2 Reverse phi grating at fovea

Reverse phi stimuli contain phase and position cues signalling motion in opposing direction. At the fovea, changes in spatial frequency of varying constructions of reverse phi stimuli, have shown the same trend in results: motion reported in the position cue direction at low spatial frequencies, and in the phase cue direction at high spatial frequencies. Chubb & Sperling (Chubb & Sperling, 1989) and Benton *et al.* (Benton *et al*, 1997) used viewing distance to manipulate the retinal spatial frequency of reverse phi gratings. They reported that at close viewing distance, where the spatial frequency of the grating was low, motion was reported in the position cue direction, and that at far viewing distances, when the spatial frequency was high, motion was reported in the phase cue direction. They reported perceived motion solely for two viewing distances, 2 and 8 metres, which correspond to gratings of 4.17 cpd and 16.67 cpd. These results are similar to those of Anstis & Rogers (Anstis & Rogers, 1975), who reported that, below a critical value of 10 minutes of arc, displacements of a contrast reversing displaced circle were reported in the phase cue direction, and that at displacement values larger than 10 minutes of arc, motion was reported in the position cue direction. These results sample the range of spatial frequency too sparsely, and without reporting individual data: this experiment addresses both these points.

In this section, I will investigate perceived motion for reverse phi gratings at 7 spatial frequencies between 1 cpd and 24 cpd, displacing the stimuli by a proportion of the wavelength of the stimulus to fix the displacement value in constant ratio to the spatial frequency and size of the stimulus. Characterising the relationship between reported motion direction and spatial frequency across a range of spatial frequencies is necessary for validation of my spatial frequency dissociation, and for the design of further reverse phi based experiments, since current results sample the range too sparsely, and without referencing the proportions of motion in either direction.

Method

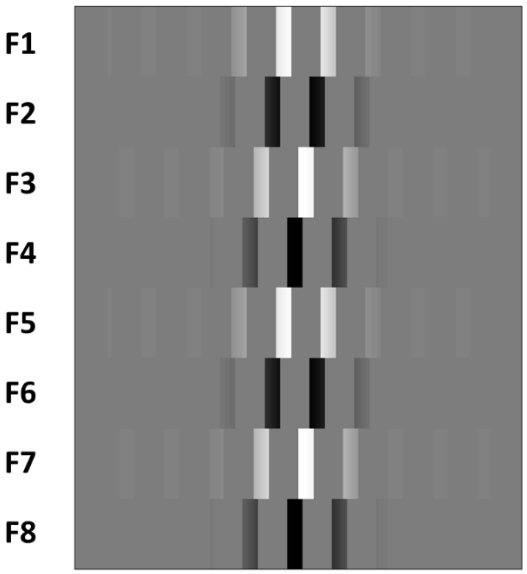


Figure 5.1. Figure displaying cross-sections of the 8 frames of the stimulus; a quasi-square grating, with a duty cycle of 1/3, with the bars ('on' portion of the grating) reversing in polarity on each frame. The grating was weighted with a 2D Gaussian function, with a standard deviation of 1.5 times the grating wavelength. This ensures that the same number of cycles are visible for stimuli presented at different spatial frequencies.

Ethics

This study was approved by the Medical Sciences Inter-divisional Research Ethics Committee (IDREC) at the University of Oxford in accordance with the procedures laid down by the University for Ethical Approval of all research involving human participants, under reference R51336/RE001.

Participants

Three participants (1 male); aged 24, 25 and 27, with normal or corrected to normal vision.

Apparatus

Stimuli were generated in MATLAB 2016, and presented on a 22-inch CRT monitor (Mitsubishi Diamond Pro 2070, 1600 × 1200 pixels, 100Hz refresh rate), with a maximum luminance of 81.32 cd/m^2 and gamma corrected using the ColorCAL MK II colorimeter (Cambridge Research Systems). Participants were seated 180cm from the screen, with the head supported by a chin rest and the right eye occluded by an opaque patch, and the room was dimly lit.

Stimuli

I produced reverse phi quasi-gratings similar to those of Chubb & Sperling (Chubb & Sperling, 1989). I used square gratings of duty cycle 1/3 to set the periodicity of the grating, but set the ‘off’ portions of the grating (2/3 of each cycle), to the background luminance. The ‘on’ portions of the grating

reversed in polarity on consecutive frames from maximum luminance to minimum luminance. The grating was also displaced by 90° of the fundamental frequency on each frame, producing an overlap in position between 'on' bars on consecutive frames. Gratings were weighted by a 2D Gaussian function, with a standard deviation of 1.5 times the grating wavelength. This ensures that the same number of cycles are presented for stimuli presented at different spatial frequencies. The starting phase and grating polarity was randomised for each trial. Eight frames were displayed on each trial, at a temporal frequency of 2.5Hz. The gratings were displayed at 7 spatial frequencies (1,2,3,5,8,14 and 24 cpd), at 50% contrast. Each grating was oriented horizontally, so that phase displacements produced motion in the vertical plane. A cross-section of the stimulus can be seen in Figure 5.1.

Procedure

Stimuli were displayed at the fovea, with a 0.15° by 0.15° white fixation cross-presented between trials. Each trial was followed by a blank screen set to the background luminance, where participants reported the perceived vertical motion of the grating, using the arrow keys on the keyboard, and then initiated the next trial. The task used a one-interval-forced-choice procedure, where participants reported whether the stimulus moved up or down.

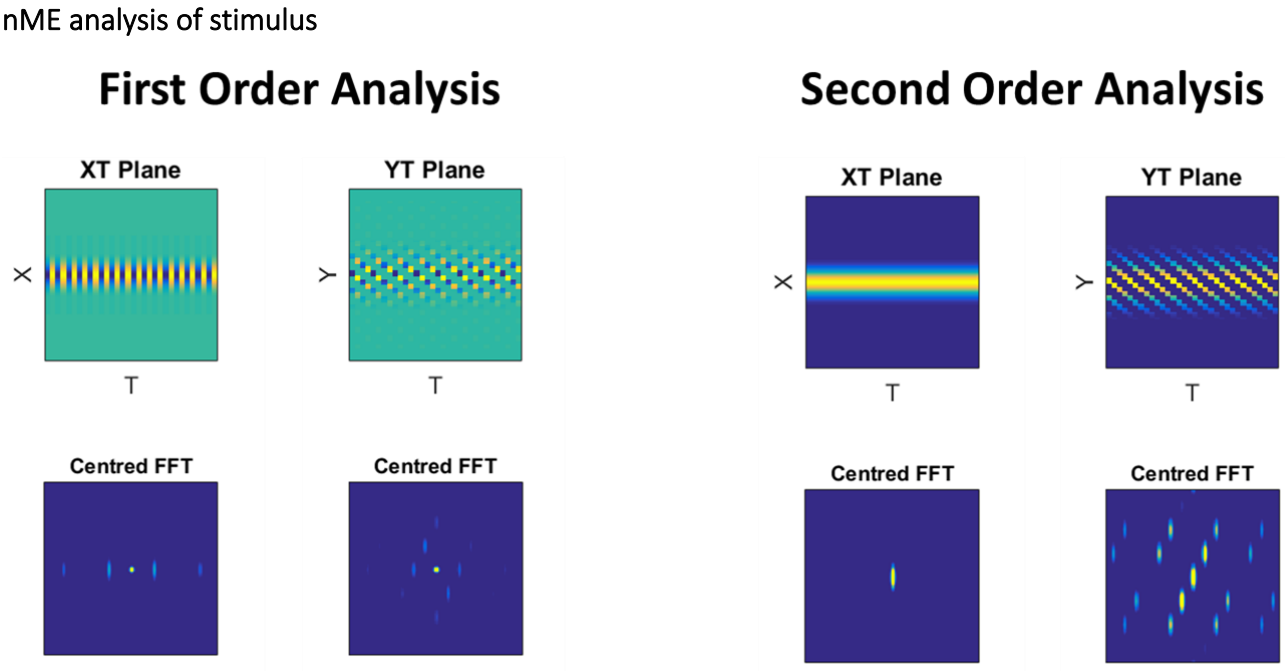


Figure 5.2. Figure displaying the Fourier analysis conducted on the 50 frame rendering of the reverse phi grating used in this chapter, similar to Chubb & Sperling (1989). The left panel displays the first order analysis of the stimulus, and the right panel the second order analysis. The first row across both panels displays the XT and YT projections of the stimulus, the bottom row displays the centred Fast Fourier Transformations (FFT) of the above projection. The XT planes had a value of -0.036 under both the first and second order analyses, whereas the first order analysis of the YT plane had an nME value of -0.079, indicating downwards motion, and the second order analysis of the YT plane had a value of 0.070, indicating upwards motion.

To verify the phase and position cue presence within the stimulus I applied normalised motion energy analysis, described in Chapter 2 section 2.4 , to 50 frame renderings of the stimulus. Using the 50 frames as the third dimension, time, this produced a three dimensional representation of the stimulus. I then produced the XT and YT projections (first row in Figure 5.2); for the first order analysis this is accomplished by taking the mean along the x-dimension to produce values for the YT plane, and along the y dimension to produce values in the XT plane. For the second order analysis (right hand side of Figure 5.2), the same procedure is followed but for a rectified version of the three dimensional representation of the stimulus, where I first removed the zero-frequency component (the mean luminance), and then full wave rectified all values (taking the absolute value), before producing the projections.

I then ran a 2D fast Fourier transform on the first and second order XT and YT planes, weighted by a Hanning Window to remove edge artefacts. The output of the Fourier transform is displayed in the bottom row of Figure 5.2. The centred fast Fourier transforms display the amplitudes of the detected frequencies in the image, with the figure arranged so that the centre of the square refers to the lowest spatial frequencies, and the angle of the signals from this origin refer to the orientation of the frequencies in the XT or YT input. The brighter the pixel, the greater the amplitude of the detected frequency. A motion signal would be reflected in an asymmetric distribution of motion energy across the vertical axis. The output of the central fast Fourier transform can be quantified as normalised motion energy, by summing the power in each quadrant of the Fourier output and comparing the power in quadrants 1 and 3 with quadrants 2 and 4.

The nME analysis of the stimulus found small values in the first and second order XT planes (-0.036), as expected, since the stimulus was displaced in the vertical direction. These values likely reflect artefacts between the 50 frames of the stimulus cube. The nME analysis of the YT planes is more informative for assessing the direction of the stimulus. The nME analysis of the first order YT plane produced a value of -0.079, downward motion, and a value of 0.070 for the analysis of the second order YT plane. Critically, this analysis confirms that the phase cue is in the opposite direction to the position cue, which is in the displacement direction; if the stimulus was displaced in the opposite direction we would see a positive value in the nME analysis of the first order YT plane, and a negative value in the nME analysis of the second order YT plane.

Results

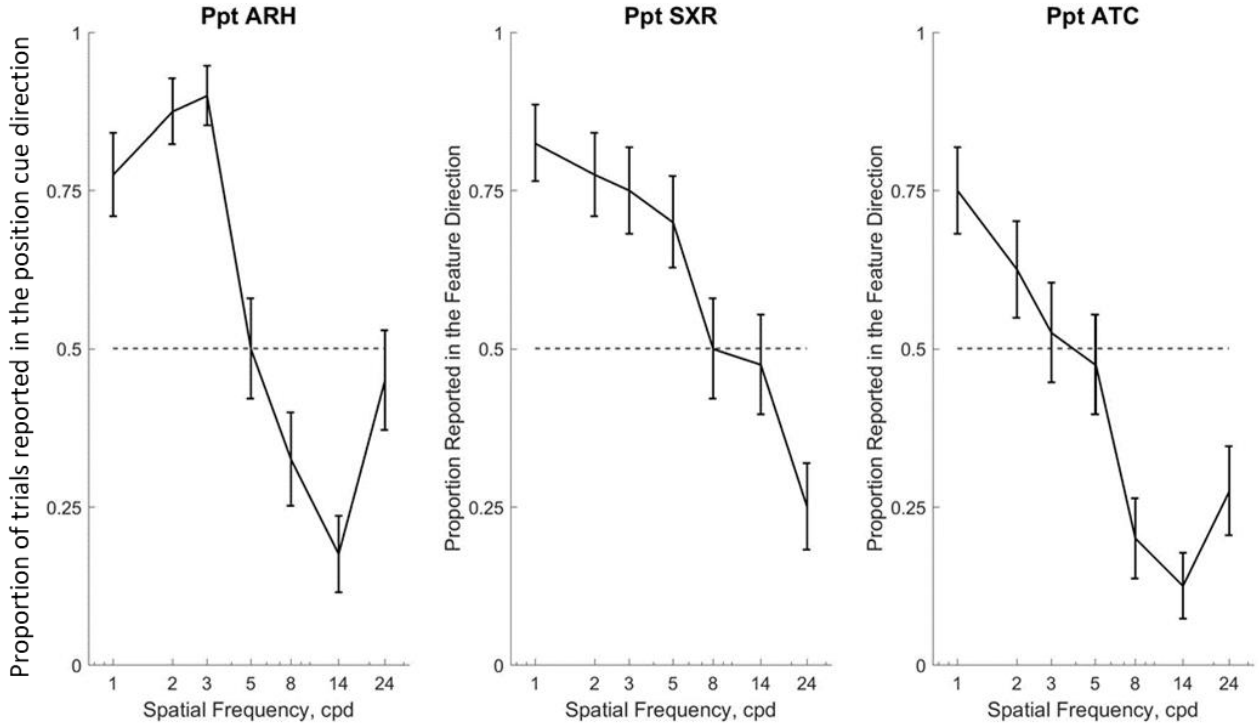


Figure 5.3. Figure displaying the reported motion direction of reverse phi gratings across 7 spatial frequencies, for the three participants. For reference, Chubb & Sperling (1989) reported that participants perceived reverse phi gratings at presented at 4.17 cpd in the position cue direction, and displayed at 16.67 cpd in the phase cue direction. Chubb & Sperling (Chubb & Sperling, 1989) displayed stimuli at a 3.75Hz, and used a duty cycle of 1/4 rather than 1/3.

Participants reported the reverse phi grating in the position cue direction at low spatial frequencies, and in the phase cue direction at high spatial frequencies, (Figure 5.3). The data are best described by a 4 parameter logistic function, where the four parameters fit the slope and threshold of the data, and the upper and lower asymptotes. I fitted mean reported direction by spatial frequency, using the logarithm in base 10 of spatial frequency, with the 4 parameter logistic function, using the nlme package in R (version 4.1.3, Pinheiro *et al.*, 2013), to include participant identity as a random effect. While there was no significant contribution of the slope parameter to the model (p = 0.1332), I preferred the four parameter logistic function rather than the three parameter logistic function, which does not contain the lower asymptote. Including the lower asymptote and using the 4 parameter logistic function significantly improved the goodness of fit of the model, (LLR = 4.9627 p = .0436).

Table 5.1 Table displaying the coefficients for the 4 parameter logistic function fitted to the data in section 5.2. Coefficients for the slope and threshold are expressed in log 10 cpd units, where the threshold is equivalent to 5.86 cpd in linear units, and the lower and upper asymptote are expressed in units of proportion correct, where > 0.5 corresponds to motion reported in the position cue direction, and < 0.5 to motion reported in the phase cue direction.

Parameter	Coefficient	Std. deviation	Variation across participant level	df	t value	*p* value
Slope	11.660	7.343	0.001	15	1.588	0.1332
Threshold	0.768	0.061	0.069	15	12.516	<.0001
Lower asymptote	0.293	0.064	0.073	15	4.546	0.0004
Upper asymptote	0.712	0.049	< 0.001	15	14.445	< .0001

There was a significant contribution to the model from the upper asymptote ($t(15) = 14.445$, $p < .001$), lower asymptote ($t(15) = 4.546$, $p < .001$) and threshold ($t(15) = 12.516$, $p < .001$), but no significant difference in slope ($p = .1332$). The values of the upper and lower asymptote suggest that performance will asymptote at 71.12% and 29.29%, demonstrating an approximately symmetric function with equal performance in the phase and position cue directions. The inflection point, at 50% performance is at 5.86 cpd, with a standard deviation of 1.15 cpd (values in table used log10 transform of spatial frequency). Finally, there was deviation across all 4 parameters at the participant level, demonstrating considerable individual variation of the fitted function to each participant.

Discussion

Participants reported the motion in the reverse phi grating in the position cue direction at low spatial frequencies, and in the phase cue direction across the higher spatial frequencies. Reported motion was modelled by a four parameter logistic function, which estimated that the reversal point is 5.86 cpd, with motion at lower spatial frequencies reported in the position direction, and at higher spatial frequencies, in the phase cue direction. This is consistent with the results of Chubb & Sperling (Chubb & Sperling, 1989) who used a very similar stimulus, but manipulated spatial frequency by viewing distance, and reported that reverse phi gratings of 4.17 cpd would be reported in the position cue order direction, and 16.67 cpd in the phase cue direction.

This spatial frequency driven bias for reported motion is also consistent with the bias in reported motion seen in conflicting cue Gabors, when viewed in the periphery. In Chapter 4, I demonstrated that, when conflicting cue Gabors were displayed in the periphery, motion was reported in the aperture cue direction at low spatial frequencies, and in the carrier cue direction at high spatial frequencies. The most substantial limitation in the use of Gabor stimuli to investigate whether there is a differential response of the visual system to phase and position cues, is that Gabor stimuli are not 'pure' phase and position cue stimuli: the phase shift of the carrier produces a corresponding displacement of the features of the carrier. Therefore any Gabor containing a phase shift of the carrier contains congruent phase and position cues. This limits the use of Gabor stimuli to investigate phase cues in isolation, but results can still be used to draw comparisons with other stimuli as long as this limitation is made explicit.

Another significant limitation of the use of conflicting cue Gabors is that reported motion in the carrier and aperture cue directions was not symmetrical in the distance from chance performance; even when motion was predominantly reported in the carrier cue direction, there was a substantial number of trials reported in the aperture cue direction. I suggest that this was due to a relatively higher strength of the aperture cue, which was maintained across spatial frequencies and eccentricity. This could be explained by the different cue types in the stimulus, where the carrier cue was dependent on the phase of the grating, whereas the aperture cue was likely determined by the change in location of the centroid of the Gabor. However, this asymmetry in performance was not seen in the reported motion for the reverse phi grating. Across participants, performance was relatively symmetric between the phase and position cue directions, with the model of reported motion direction suggesting asymptotes at 29.3% for the phase cue direction at 71.2% for the position cue direction. The near-symmetry of the results could be the result of the greater similarity in the cue construction in reverse phi stimuli, where the phase and position cues originate from the same feature in the image. As a result of this symmetry in performance, I can extract displacement thresholds which will correspond to performance that is equally distant from chance, such as displacement thresholds at 25% and 75% proportion correct, where correct is operationalised as motion in the position cue direction.

Another important difference between performance in reverse phi stimuli and the conflicting cue Gabors, is that, for reverse phi stimuli, this spatial frequency driven reversal in reported motion can occur at the fovea (Edwards & Nishida, 2004; Wehrhahn, 2006). In conflicting cue Gabors, I demonstrated in Chapter 4 that reported motion for foveal viewing did follow the same spatial frequency biases, but was not sufficient to drive a reversal of motion from the aperture cue direction to the carrier cue direction (see Figure 4.2). This restricted the ability to compare conflicting cue

motion between the fovea and periphery, where the extraction of displacement values corresponding to peak motion reported in the carrier, and aperture cue direction only included data collected from peripheral conditions. The reversal of motion at the fovea seen in the reverse phi stimuli means that I can potentially map changes between fovea and periphery in the phase and position cue directions using similar stimuli and analyses.

In Chapter 3, I measured proportion correct across varied spatial frequency in Gabor stimuli containing solely a carrier or aperture cue, and extracted a displacement threshold, relating to 75% correct. These displacement thresholds for the carrier and aperture cues were used to support a proposed psychophysical dissociation between two motion systems, one processing phase cues and the other processing position cues. In the reverse phi data reported in this chapter, the symmetry in the proportion of trials reported in the phase and position cue directions, and this reversal at the fovea, suggest that reverse phi stimuli may also be informative in supporting this psychophysical dissociation between phase and position cues, and that motion perception in reverse phi stimuli should be explored across the periphery.

5.3 Reverse phi square stimulus across eccentricity

In Chapter 3, further support for the proposed psychophysical dissociation between the phase and position systems came from the significant difference in the gradients of displacement thresholds for the carrier and aperture cue with eccentricity. If the displacement gradients calculated for the conflicting cue Gabor stimuli in Chapter 4 represent the underlying behaviour of two psychophysically distinct motion processing systems, then this dissociation should also be visible across stimuli containing phase and position cues but with different constructions. I found that the displacement gradients for the aperture cue in the Gabor stimuli closely resembled those for static acuity (Strasburger *et al*, 2011), and the displacement gradients for the carrier cue in the Gabor stimuli resembled gradients collected for unreferenced motion, in stimuli that attempt to produce motion absent of position cues (Levi *et al.*, 1984; Whitaker *et al.*, 1992).

Conflicting cue stimuli will have different displacement gradients with eccentricity in comparison to pure cue stimuli, since the reported motion represents a relative strength of one cue exceeding the other, rather than just the displacement threshold of that cue in isolation. However, the same relationship with eccentricity should exist, as the relative strength of the two cues will be produced in response to the spatiotemporal sensitivity of the underlying systems. My experiment at fovea (section 5.2) demonstrated approximately symmetrical performance for the reverse phi grating in the phase and position cue directions, indicating that I could extract displacement thresholds for the performance in the phase and position cue directions, such as at 25% and 75% performance. If the

displacement thresholds reported in Chapter 3 represent the change in spatial sensitivity of two systems separately processing phase and position cues, then I would expect to see similar gradients for the displacement thresholds in the phase and position cue direction in the reverse phi stimuli.

To produce displacement gradients across eccentricity for the reverse phi stimuli requires a stimulus that can produce consistent reversals of motion when viewed in the periphery. Reverse phi stimuli in periphery have solely been reported in the phase cue direction, for grating stimuli (Chubb & Sperling, 1989), and random dot kinematograms (Edwards & Nishida, 2004). Maruya *et al.* demonstrated that a motion-defined motion reverse phi grating stimulus, that could be manipulated by displacement phase to produce motion reported in the phase or position cue direction, was solely reported in the phase cue direction when viewed peripherally (Maruya *et al*, 2003). These findings could be used to suggest a dominance of the phase system in the periphery, similar to early assertions that the periphery is “specialised for motion perception" (Parsons, 1927). However, these results could also be explained by the parameters used to present the stimuli. In Chapter 3, I suggest that the displacement thresholds in the aperture cue condition increased more rapidly with eccentricity than the displacement thresholds in the carrier cue condition. If these stimuli are producing selective activity in putative systems processing phase and position cues separately, then this would indicate that to produce a similar response in the phase and position systems, position cue stimuli must be presented with substantially larger displacements than the increase needed for phase cue stimuli. If so, presenting the reverse phi stimuli with larger displacements in the periphery may restore the perception of motion in the position cue direction. This increase in necessary displacement size presents an issue for testing in periphery, where the increased size of the reverse phi grating stimulus used in section 5.1 would mean that the stimulus would be presented covering a large extent of eccentricity and reducing the precision of the viewing location.

The ‘reversal point’ values between grating based reverse phi stimuli and single figure reverse phi stimuli suggest that stimulus construction can be used to bias motion in the phase or position cue direction. Anstis & Rogers used a single disc figure which dissolves across multiple frames into a displaced photographic negative (Anstis & Rogers, 1975). They reported that motion was seen in the displacement (position cue) direction at large displacements, equivalent to low spatial frequencies, and in the phase cue direction at small displacements, equivalent to high spatial frequencies. They suggest that the reversal point value was 10 minutes of arc. In my stimulus construction, a grating displayed at 1.5 cpd would be displaced by 10 minutes of arc on each frame, and my critical value of 5.86 cpd, corresponds to 2.56 minutes of arc displacement on each frame. This difference, where my stimulus has a smaller critical displacement value, suggests that for the grating stimuli, compared to single figure stimuli, a higher spatial frequency is needed to produce motion perceived in the phase

cue direction. This suggests that the position cue is effectively stronger in the single figure stimuli, and weakened by the presence of multiple features. This is consistent with Boulton & Baker's Gabor micropattern arrays, where motion is reported in the carrier cue direction for dense arrays of Gabors, and in the aperture cue direction for sparse arrays of Gabors (Boulton & Baker, 1993). Switching from a grating based reverse phi stimulus to a single figure stimulus should therefore enable the presentation of smaller displacements, but still produce a robust motion percept in the phase and position cue directions. The change from grating stimulus to single cue also dramatically reduces the size of the total stimulus.

To obtain displacement thresholds for the reverse phi stimulus across periphery, I produced a single figure stimulus, a three-frame displaced square, inspired by the dissolving circle of Anstis & Rogers (Anstis & Rogers, 1975), and the single bar stimuli used by Wehrhahn (Wehrhahn, 2006) and Chubb & Sperling (Chubb & Sperling, 1989). I collected the displacement thresholds for the phase and position cue directions separately, using staircases to find thresholds, rather than assessing the entire logistic function of reported direction by displacement size.

Method

Participants

Three participants took part in this experiment, all female, aged 22 to 26, with normal or corrected to normal vision.

Apparatus

The experiment was conducted in a dimly lit room. Stimuli were presented on a CRT monitor (Mitsubishi Diamond Pro 2070, 1600 × 1200 pixels), and gamma corrected using the ColorCAL MK II colorimeter (Cambridge Research Systems). Viewing distance was maintained with a chin rest. Viewing distance was modified for the different eccentric conditions to ensure that when displacement was very small, in the phase cue conditions at the lower eccentricities tested, stimuli could be accurately rendered. Due to the different ranges of sizes used in the staircases for the phase and position cue directions, at the two highest eccentricities different viewing distances for the Phase and Position conditions were required. As a result, I ran the Phase and Position conditions in separate session for all eccentricities. Viewing distances are reported below in Table 5.2.

Table 5.2 displaying the viewing distances (in centimetres) used for the staircases at each eccentricity (in degrees) and for the two conditions.

Eccentricity	5	10	15	20	25
Phase cue threshold	125	125	105	105	85
Position cue threshold	125	125	105	85	57

Viewing was monocular, with the participant's left eye covered by an opaque eyepatch. A matte black carboard mask was applied to the edges of the CRT monitor to minimise edge artefacts, and to provide a location for a physical fixation point on right side of the display. Stimuli were displayed to the right eye in the temporal hemifield. A physical fixation spot was applied to the black mask, using a small white sticker with 0.25cm diameter. Fixation was not monitored, but participants were encouraged to abort trials in which their fixation strayed. This happened in less than five trials per session per participant.

Stimulus

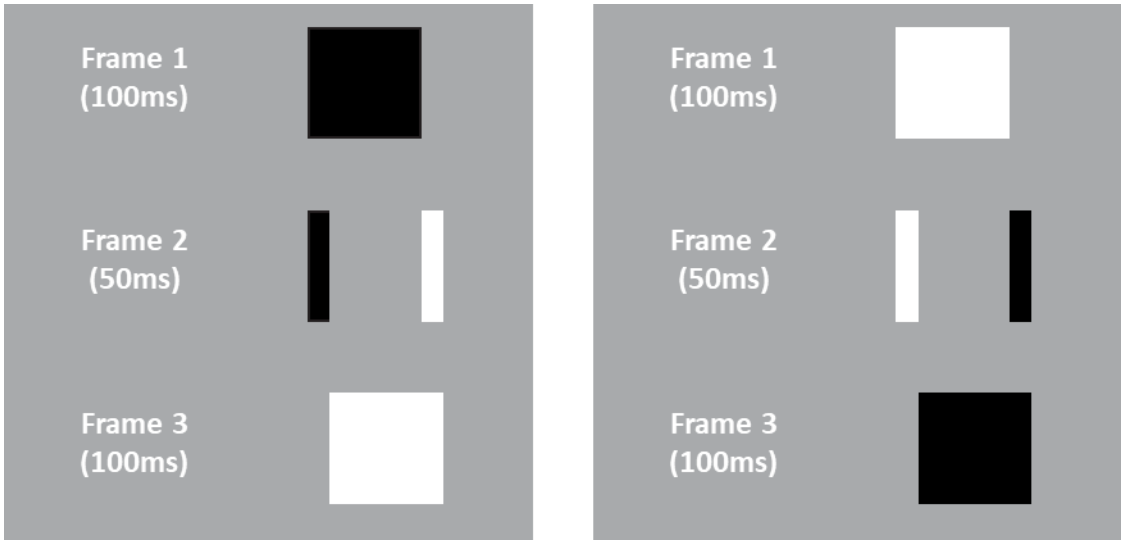


Figure 5.4 Schematic of the reverse polarity square used in section 5.3. The stimulus could transition from low luminance to high (left panel) or high luminance to low (right panel). In both panels the square is displaced by a quarter width to the right, with Frame 1 displaying the square in its initial position, Frame 2 consisting of two quarter-width strips at the original left-most quarter width position, and the righthand strip occupying the quarter-width location that the square will subsequently occupy in Frame 3, and Frame 3 displaying the square in its final location.

The stimulus was a reversing polarity square, displaced vertically by a quarter of the square's width (see Figure 5.4). The stimulus consisted of three frames. The first and last frames consisted of the square, displaced by a quarter of its width and reversed in polarity. The central frame was inspired by the 'dissolve' in Anstis & Rogers' demonstration of the reverse phi stimulus (Anstis & Rogers, 1975), and contained two strips with the same height as the square, and width equal to one quarter of the square's width. These quarter-width strips were located so that the 'overlap strip' covered the portion of Frame 1 which would be uncovered in Frame 3, and the second quarter-width strip, the 'displacement strip', was located so that it covered the area that the square was displaced to in Frame 3. These two strips were located three-quarters of the square width apart, and matched to the polarity of Frame 1 (overlap strip) and Frame 3 (displacement strip). In pilot testing, the inclusion of this central frame with the simultaneous display of the two borders improved the sensation of smooth motion in the periphery.

Squares were presented at maximum (81.33 cd/m^2) and minimum luminance values on a background of gamma corrected mean luminance. For half of the trials, the square was black on Frame 1 and reversed polarity to white on Frame 3, and vice versa.

Procedure

Following pilot testing, it was apparent that a very large range of stimulus sizes would be needed to cover the range of reported motion, from the phase cue direction at small stimulus sizes, to the position cue direction at large stimulus sizes. Instead of running a method of constant stimulus procedure, as in section 5.1, I decided to find the 25% and 75% displacement thresholds separately, running separate blocks of trials for the phase and position arms of the logistic function, for each of the five eccentricities tested.

To obtain the thresholds, I used 2-up-1-down and 3-up-1-down staircases, which converged at 70.7% and 79.4% (or 20.7% and 29.4% for phase cue threshold), where I can interpolate the 25% and 75% thresholds. To find the 25% and 75% thresholds, I used very different starting sizes for the stimuli, corresponding to stimulus sizes that have estimated performance of 10% and 90%, from a pilot experiment. To obtain the thresholds for the position cue direction, I started with a large stimulus, where the stimulus width was increased with incorrect answers and decreased with correct answers, with 'correct' referring to responses in the position cue direction. To obtain the thresholds for the phase cue direction, I started with a small stimulus, and used the same procedure, except where 'correct' now refers to stimuli reported in the phase cue direction. Participants completed 3-5 repeats of each staircase, at each eccentricity for each task type, and staircases ended after 10 reversals. Staircases were interleaved between the two staircase types.

In each trial, I used a two-alternative forced choice procedure, where participants reported the interval in which the motion was perceived in the upward direction. Each trial consisted of two intervals, each containing three frames of the stimulus. One interval contained upward motion and the other downward motion; half the trials in the experiment contained downward motion followed by upward motion, and the other half upward motion followed by downward motion. Each trial lasted 750 ms, with Frames 1 and 3 displayed for 100 ms, Frame 2 for 50 ms, and a blank screen of mean luminance was displayed between the first and second interval for 500 ms. Trials were followed by a blank screen at mean luminance, which was displayed until participants reported the interval in which upward motion was perceived. Participants initiated the next trial using the spacebar.

Participants completed 3-5 staircases for each eccentricity and threshold type, across three sessions, which each lasted up to two hours. Each session was organised sequentially by viewing distance, with participants completing the sessions in near-periphery first.

nME analysis of stimulus

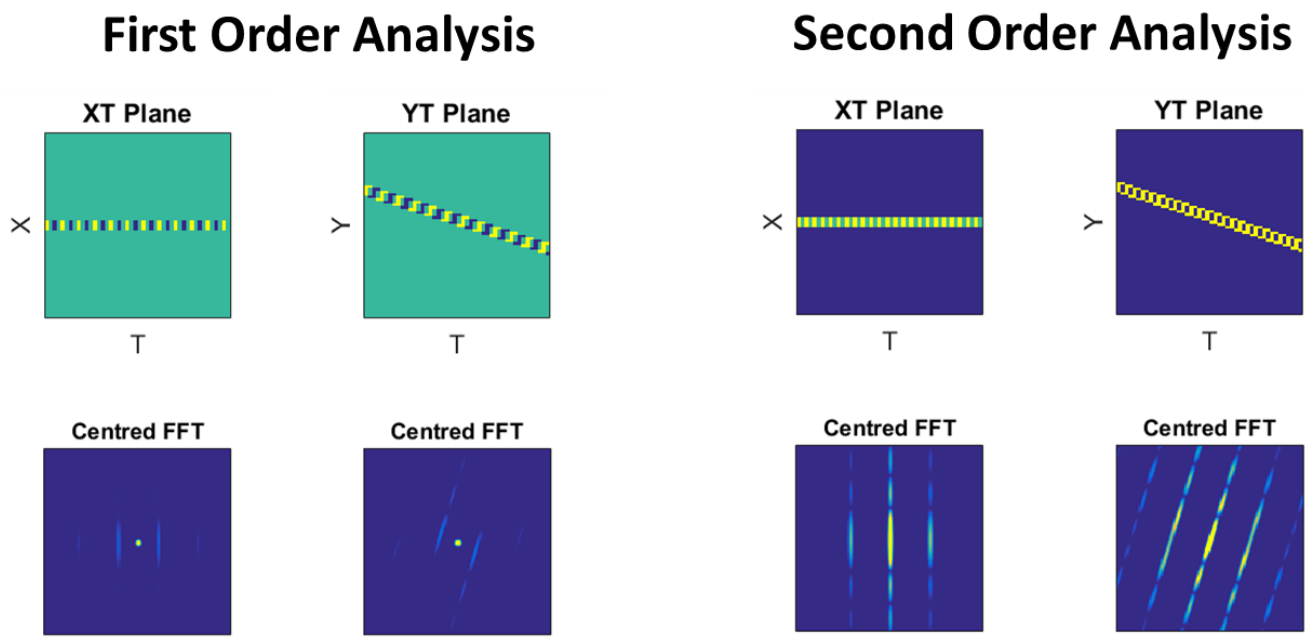


Figure 5.5. Figure displaying the Fourier analysis conducted on the 50 frame rendering of the reverse phi squares stimulus. The left panel displays the first order analysis of the stimulus, and the right panel the second order analysis. The top row across both panels displays the XT and YT projections of the stimulus, and the bottom row displays the centred Fast Fourier Transformations (FFT) of the above projection. The XT planes had a value of -0.020 under the first order analysis, and -0.003 under the second order analysis, whereas the first order analysis of the YT plane had an nME value of -0.120, indicating downwards motion, and the second order analysis of the YT plane had a value of 0.082, indicating upwards motion.

To verify that the stimulus used contained conflicting phase and position cues, I ran the normalised motion energy analysis described in Chapter 2.4 for a 64-frame version of the stimulus. The nME analysis of my reverse phi squares stimulus, displayed in Figure 5.5, demonstrates that the phase and position cues in the stimulus are in opposing directions, with a negative YT values in the first order analysis of the stimulus, and a positive YT value in the second order analysis of the stimulus (see Figure 5.5, or Table 5.3). The diagonal motion energy is clearly visible in both the centred (Figure 5.5, bottom row), whereas the motion energy in the XT planes is substantially more symmetric, with the small deviations (≤ -.020) the result of artefacts or minimal aliasing between frames.

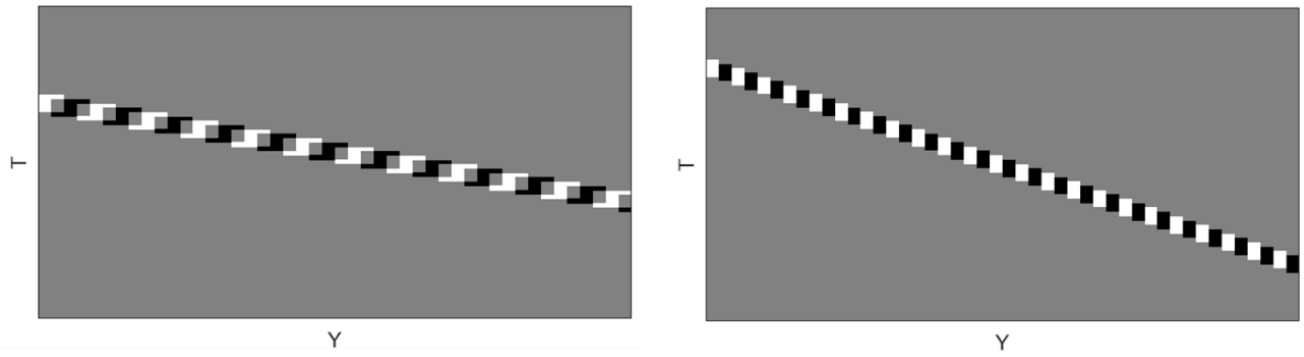


Figure 5.6. Figure displaying the YT projection of reverse phi squares stimulus used in this chapter (left), with the YT projection of a comparison stimulus that contains only the first and third frames of this stimulus, excluding the central 'borders' frames (right). These stimuli were compared using the output of the nME analysis, to ensure that the addition of the borders frame had not introduced artefacts that would interfere with the conflicting phase and position cues in the stimulus.

I also ran the analysis on a simplified version of the stimulus that excluded the 'borders' frame as a comparison, to verify that the addition of the central frame with the borders did not produce artefacts that would disrupt the conflicting phase and position cues. Cross-sections of the two stimuli, reverse phi squares and the comparison stimulus, are displayed in Figure 5.6.

I ran the same first and second order nME analysis for a comparison stimulus (Figure 5.6, right), with no central borders frame, so produced from the first and third frames of my reverse squares stimulus. This stimulus produced a similar pattern of results, with the XT and YT projections and their respective centred FFTs represented in Figure 5.7, and nME values reported in Table 5.3. Critically, this stimulus also had conflicting phase and position cues, although with a different ratio between the strength of the first and second order nME YT values. Compared to the simplified stimulus without the borders frame my stimulus had a greater first order motion energy, but less second order motion energy.

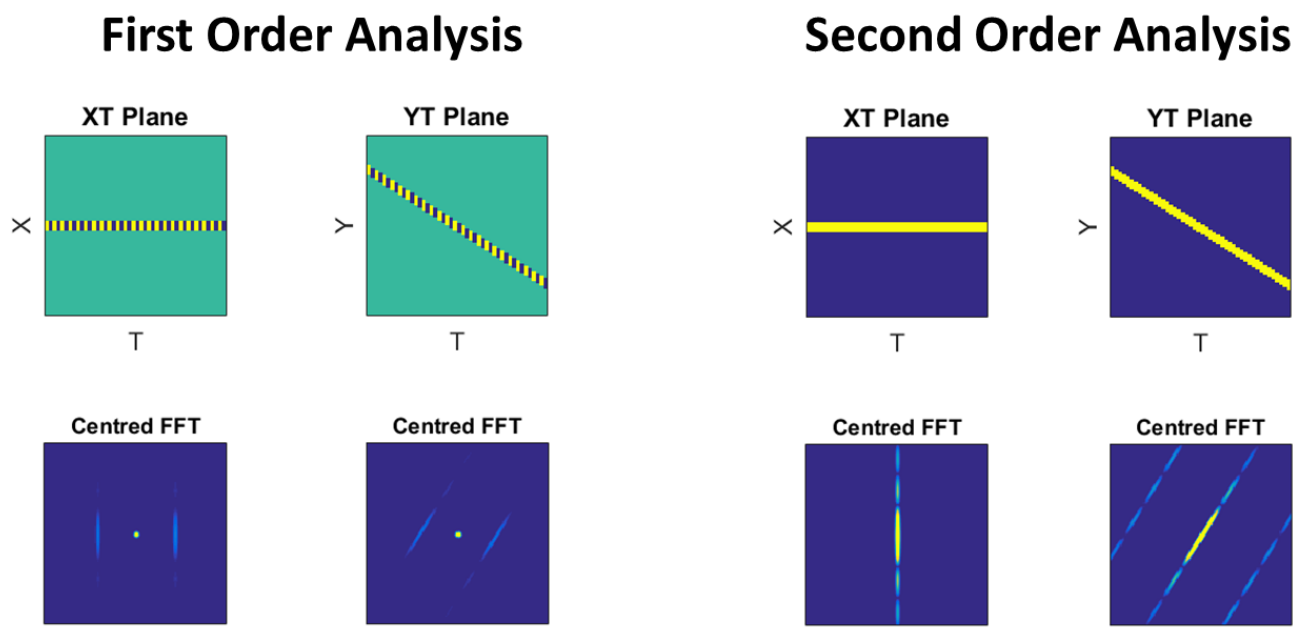


Figure 5.7. Figure displaying the Fourier analysis conducted on the 50 frame rendering of a comparison stimulus for the reverse phi squares stimulus, which contains just the first and third frames; no borders frame. The left panel displays the first order analysis of the stimulus, and the right panel the second order analysis. The first row across both panels displays the XT and YT projections of the stimulus, the bottom row displays the centred Fast Fourier Transformations (FFT) of the above projection. The XT planes had a value of -0.017 under the first order analysis, and -0.006 under the second order analysis, whereas the first order analysis of the YT plane had an nME value of -0.076, indicating downwards motion, and the second order analysis of the YT plane had a value of 0.207, indicating upwards motion.

Table 5.3. Table displaying the first and second order normalised motion energy values for this stimulus and a comparison stimulus which did not contain middle frame with the quarter-width borders.

	1st Order YT	2nd Order YT
This stimulus (with borders)	-0.1205	0.0825
Simplified stimulus (no borders)	-0.0759	0.2072

Results

Analysis strategy / Issues with staircases & the psychometric solution

Thresholds were extracted from each staircase by averaging across the final three values. However, comparing the threshold to the standard deviation and by visual inspection, a reasonable proportion of staircases did not converge; 14.74% in the phase cue direction and 8.20% in the position cue direction. If staircases do not converge then they cannot be included in threshold estimation. My solution to the poor convergence of staircases was to change tactic for the analysis and instead of

using the interpolated thresholds from the staircases which did successfully converge, to use a new analysis method.

As a result of the staircase method, the data consisted of a large range of stimulus sizes tested for each condition, with very few trials for each size. However, pooling data across a collection of trials, enabled the transformation of the data into a set of estimated performance values across a reduced range of stimulus sizes. I grouped data by number of trials, using 20 trials per datapoint, which produced approximately 12 pairs of displacement size and performance per condition (SE = 4.7).

Because I had used staircases converging at 70.7% and 79.4% (and 20.7 and 29.4% for phase cue threshold), much of the data was clustered around the threshold region. The MATLAB implementation of Wichman & Hill's four parameter psychometric function (Wichman & Hill, 2001), 'psignifit' was robust and able to fit psychometric functions for each participant across both Phase and Position conditions and the five eccentricities tested.

I fitted these psychometric functions separately for each participant, for each of the two stimulus conditions and five eccentric locations. This can be seen in Figure 5.8, where I display the data for participant EM, with the psychometric functions plotted using the aggregated staircase data. From these psychometric functions, I interpolated the 75% displacement thresholds for the phase and position cue conditions; the reverse phi and phi stimulus conditions.

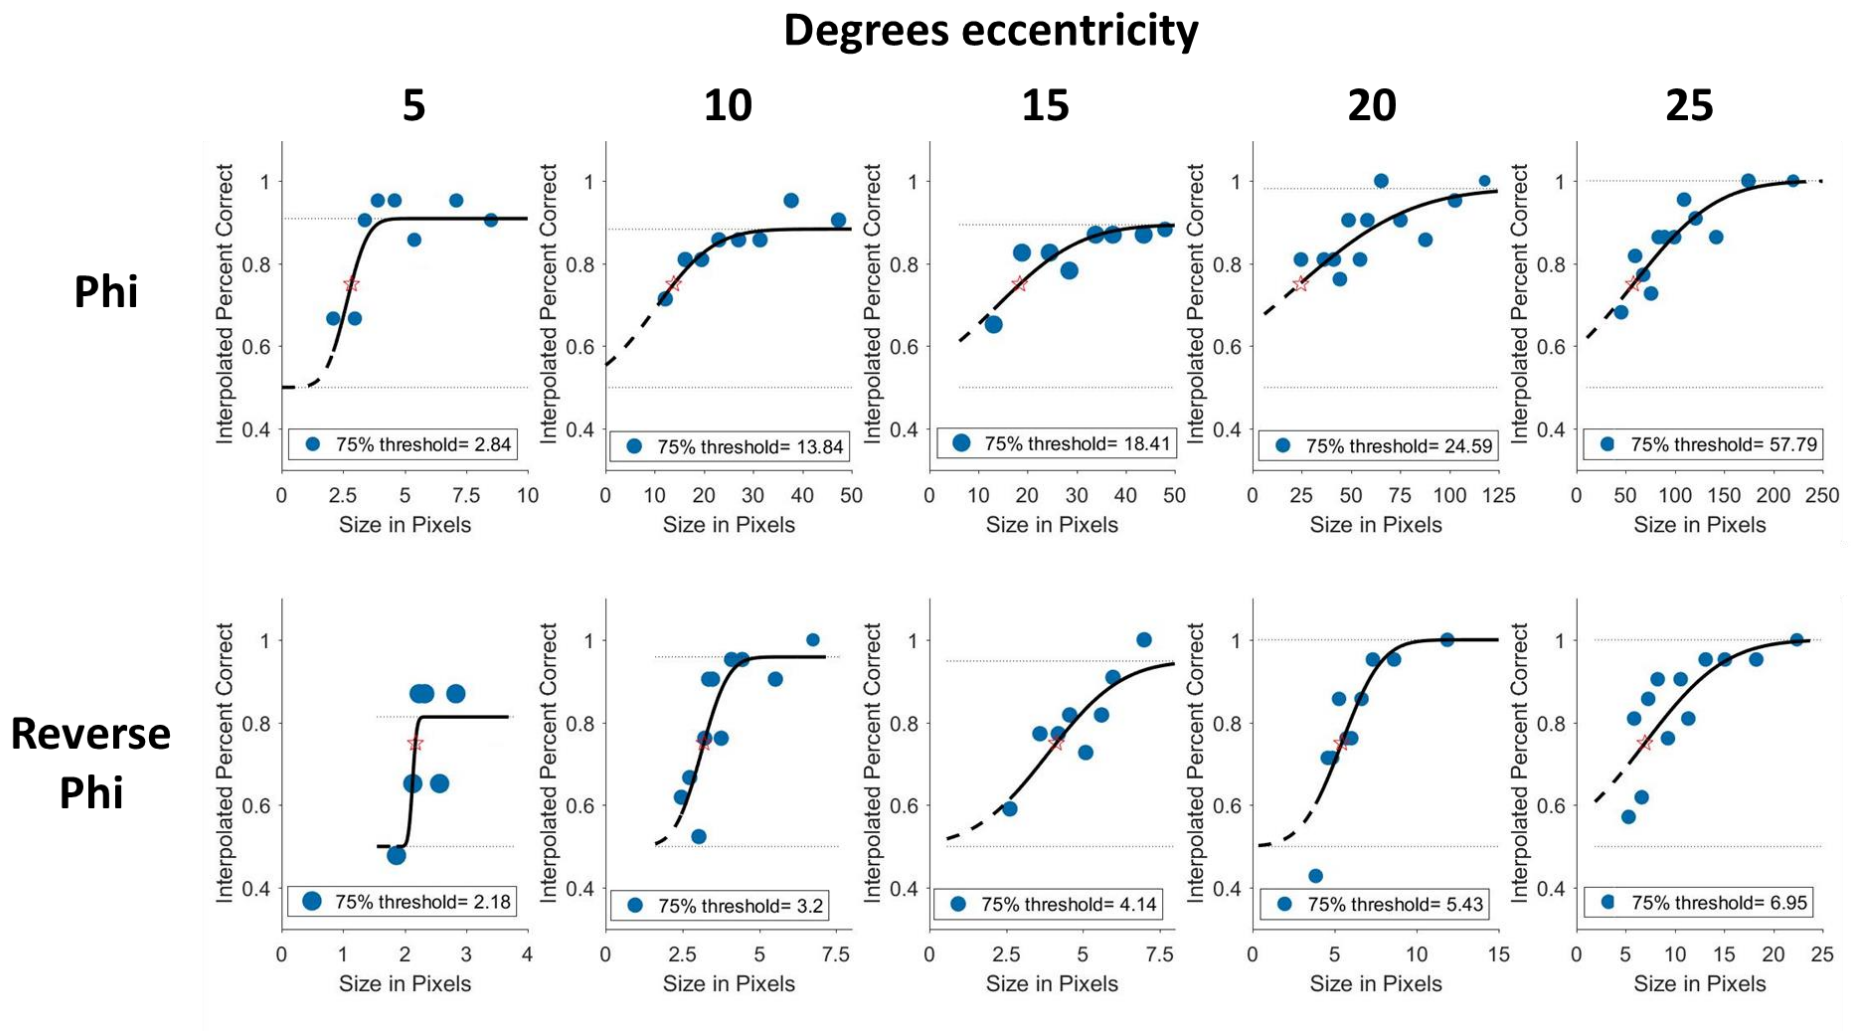


Figure 5.8. Figure showing the data for participant EM. These figures show the staircase data collected separately for Phi and Reverse Phi stimuli, at 5 eccentric locations. The data was grouped into sets of 20 trials, and psychometric functions were fitted to the mean proportion correct and mean displacement size. The top row displays the data from the phi condition, and the bottom row the data from the reverse phi condition, with the columns displaying the data collected at 5, 10, 15, 20 and 25 degrees eccentricity, respectively. The interpolated 75% thresholds are marked by the red star, and value reported in the figure legend.

Displacement thresholds

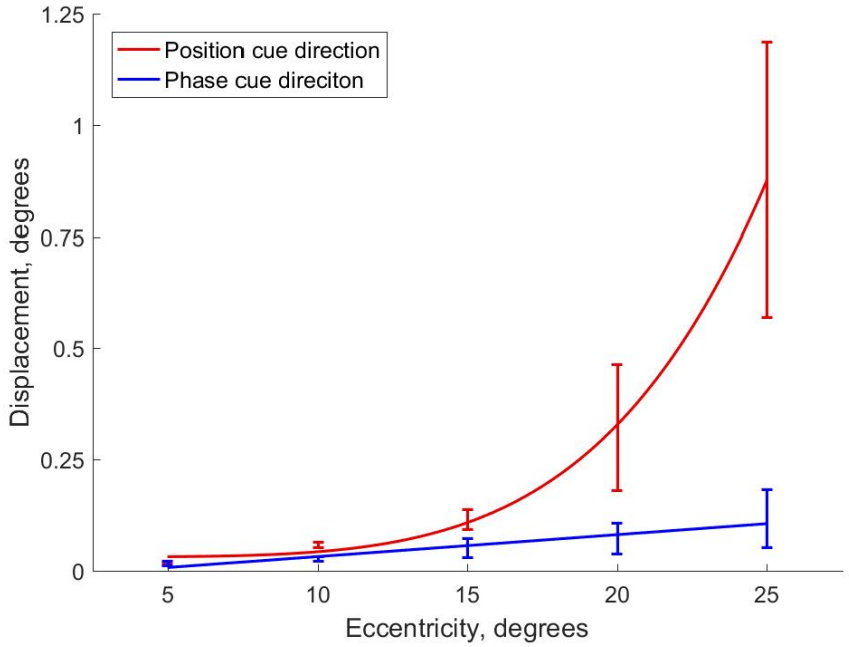


Figure 5.9. **Figure showing the interpolated displacement thresholds for the phase and position cue direction threshold conditions, across the five eccentric locations tested for the reverse phi three-frame square stimuli, with data from the three participants.**

In both the phase and position cue conditions, the extracted displacement thresholds increased with eccentricity. These were best described by a two parameter power function, between the logarithm of the displacement threshold and the logarithm of the eccentricity values.

$$\text{Log10(threshold)} \sim a*\text{log10(eccentricity)}^b$$

I fitted the data using a nonlinear mixed effects model, using the nlme package in R (version 4.1.3, Pinheiro *et al.*, 2013). I included threshold condition (phase or position) and participant as random effects, which could vary across the a and b parameters of the fitted model.

Table 5.4. Table displaying the parameters and coefficients for the model describing the change in displacement thresholds with eccentricity.

Parameter	Coefficient	Std. error	Std. error condition	Std. error participant	df	t value	*p* value
a	0.903	0.168	0.200	0.177	23	5.376	<.0001
b	1.709	0.141	0.000	0.000	23	12.144	<.0001

I used this model to test whether thresholds significantly varied between conditions by producing two variants of this two parameter model; one that included just participant as a random effect, and one that included both participant and condition as random effects. The model had a significantly better fit to the data with condition included as a random effect (AIC = 0.335) , compared to a model that only included participant as a random effect (AIC = 35.476), (Log Likelihood Ratio = 39.1406 , $p <$ .001). This demonstrates that there was a significant effect of condition on participants' displacement thresholds. The a parameter varied significantly between conditions, with a coefficient of 1.080 in the position cue condition, and a coefficient of 0.726 in the phase cue condition. The b parameter did not vary between the Phase and Position conditions. Since only the a parameter varied between conditions, I conclude that the difference in the a parameter is driving the significant difference between the condition. The a parameter controls the steepness of the function, and so I conclude that the displacement thresholds in the Position condition increase significantly more rapidly with eccentricity than do the displacement thresholds in the phase cue condition.

Discussion

Using a novel reverse phi stimulus displayed at 5 locations across the periphery participants reported motion both in the phase cue direction, with small sized stimuli with small displacements, and in the position cue direction, for larger stimuli with bigger displacements. I used staircases to converge on performance around the 25% and 75% performance thresholds on a hypothetical logistic function, which corresponded to displacement thresholds in the phase and position cue directions. There was a significant difference in the increase in displacement thresholds with eccentricity, where the displacement thresholds in the position cue direction increased significantly more than the displacement thresholds in the phase cue direction. These results are evidence that in the periphery motion can be perceived in the forward direction in reverse phi stimuli, provided the stimuli are sufficiently large.

These results show the same relationship as the displacement thresholds in the single cue Gabor stimuli reported in Chapter 3. In Chapter 3 where motion was produced either through the change in phase of the carrier (Carrier condition) or displacement of the Gaussian Aperture (Aperture condition), and there a greater increase in displacement thresholds with eccentricity for the aperture than the carrier cue. The Carrier condition contains both a phase and a position cue, with the change in phase producing a change in spatial position of the carrier features, whereas the Aperture condition contains solely a position cue, due to the rotation of the carrier removing phase cues between frames. However, in the conflicting cue Gabors, containing both carrier and aperture cues in conflicting directions, the reported direction of motion was similar to that in the reverse phi

stimuli: reported in the phase cue (or carrier cue for Gabor stimuli) direction at high spatial frequencies, and in the position cue (aperture cue for Gabor stimuli) direction at low spatial frequencies. I have not compared directly between the Gabor and reverse phi experiments, since the Gabor data were collected on single cue stimuli, where the displacement thresholds reflect the direction discrimination of the carrier or aperture cue. The reverse phi data reflect the displacement threshold, where the relative strength of the phase or position cue has exceeded the relative strength of the opposing cue. However, the Gabors and reverse phi displacement thresholds share two properties; the similarity in the relationship between the two cue types, smaller displacement thresholds in the phase cue direction and larger in the position cue direction, and the similarity in the appearance of the fitted functions, where the increase in displacement size is much more rapid for the position cue. These suggest that the reverse phi displacement thresholds are similarly determined by two competing separate motion systems, one processing phase cues and one processing position cues.

There is a rapid increase in the size of the displacement thresholds between 15 and 25 degrees for the position cue, but not the phase cue. This could be a consequence of testing at greater eccentricities: the reverse phi experiment covered a wider range of eccentricities. If I had tested motion discrimination of the Gabors in the far periphery I might have observed a similar increase in the rate of change for displacement thresholds in the aperture cue. Given the consistent gradient for the phase cue displacement thresholds in the reverse phi stimuli, this could be further evidence for a psychophysical dissociation between the systems' processing phase and position. This greater increase in the position cue thresholds between 15 and 25 degrees eccentricity could also be a result of my stimulus design. In my reverse phi stimulus, I included a central 'border frame', where the outermost edges of the squares on the first and last frames of the stimulus were presented in isolation, as two separate rectangular strips. For a system tracking the change in position of a stimulus, the border frames potentially have a larger disruptive effect. The border frame changes the identity of the stimulus from a single square, to two separated rectangles, and back to a single square percept. Apparent motion stimuli that change identity in terms of feature colour or texture, or even shape, have high performance (Cavanagh *et al.*, 1989), but the splitting and regrouping of the stimulus may impair performance more than a change in stimulus appearance, because this would affect the constancy of the features as belonging to a single entity. This splitting and regrouping of the square stimulus might produce a greater impairment in the far periphery than the near periphery, due to the increasing positional uncertainty with eccentricity (Hess & Field, 1993). The effect of increased positional uncertainty in the position cue direction would be that larger displacements are needed to reach the 75% direction discrimination threshold, consistent with the

data. Motion derived from phase cues can occur at a finer spatial grain than discrimination of position (Exner, 1875; Foster *et al.*, 1989). This suggests that the effects of positional uncertainty in the periphery may affect the position system and not the phase system. Positional uncertainty may be partially responsible for the larger displacement thresholds in the position cue direction than in the phase cue direction, but the presence of the border frame may be having a disproportionate impact on the Position condition, increasing the size of displacement needed for motion discrimination. This may explain the increasing rate of change in the position cue thresholds that is not seen in the more steady rate of change in the phase cue thresholds. This is a limitation of the current data set. Further investigations could be carried out to determine whether the displacement thresholds in the position cue direction are disproportionately affected by the stimulus design, either by replicating the experiment with a different reverse phi stimulus, or extending the Gabor experiment to collect displacement threshold data at greater eccentricities. A replication could also sample motion at the fovea with the same stimulus used in the periphery, to provide a more complete description of perceived motion in conflicting cue stimuli across the visual field.

5.4 Discussion

The aim of the reverse phi experiments presented in this chapter was to investigate the proposed psychophysical dissociation in spatial frequency sensitivity between putative phase based and position based systems for processing motion.

In this chapter, I have presented two experiments assessing reported motion in reverse phi stimuli. The data demonstrate that reverse phi stimuli are reported in both the phase and position cue directions and show a similar relationship with spatial frequency when viewed both at the fovea (Figure 5.3) and in the periphery (Figure 5.8), where reported motion is biased towards the phase cue direction at high spatial frequencies, and the position cue direction at low spatial frequencies. This is similar to the reported motion direction in conflicting cue Gabor stimuli viewed peripherally, where motion is reported in the aperture displacement direction at low spatial frequencies, and in the direction of the phase shift of the carrier at high spatial frequencies (Figure 4.2).

Reported motion direction can be used to extract displacement thresholds corresponding to 75% proportion correct, for the phi and reverse phi directions, and plotted against eccentricity (Figure 5.9). There was a significantly greater increase in displacement thresholds with eccentricity for motion reported in the position cue direction, than for displacement thresholds for motion reported in the phase cue direction. The change in displacement thresholds with eccentricity for the phase and position based cues mirrors the findings reported in Chapter 3, where displacement thresholds

increased much more rapidly with eccentricity for Gabors containing an aperture displacement (position cue), then for Gabors containing a phase shift of the carrier (phase cue, and corresponding change in position of the carrier features) (Figure 3.3).

This dissociation was proposed using displacement thresholds from Gabor stimuli, which attempted to dissociate phase and position cues for motion, but were limited in the confounding presence of both phase and position cues in the Carrier condition. Reverse phi stimuli avoid this confound, but contain both cues simultaneously, signalling motion in conflicting directions. This means that the reverse phi displacement thresholds potentially represent the increased relative strength of one system over the other, and so highlighting the regions of difference between the two systems rather than characterising the sensitivity of the two systems in isolation. These regions of difference may be more critical for describing the relative functions of the two systems, and for correlating these psychophysical dissociations between the phase and position systems to potential physiological substrates. This strategy of finding differences between the two systems has been employed for differences in chromatic sensitivity, where motion standstill in chromatically defined gratings (Lu *et al.*, 1999), indicates that the system extracting motion from changes in phase is insensitive to chromatic information, but the sensitivity of the system extracting motion from changes in position remains high. This suggests that the position system may be constructed from input from pβ retinal ganglion cells (Ruppertsberg *et al.*, 2003).

The experiments in this chapter investigate the spatial frequency sensitivity of reverse phi stimuli, demonstrating spatial frequency dissociations in the use of these cues for determining motion direction that could be explained by the presence of two separate systems processing phase and position cues. As systems processing motion stimuli, the temporal frequency variations in sensitivity are also critically important in defining the function of these two systems. It is known that the sensitivity for position cues declines rapidly with increasing temporal frequency (Lu & Sperling, 1995), while phase based motion discrimination can occur at very high temporal frequencies, as demonstrated in the 'magniphi' stimulus (Steinman *et al.*, 2000). I have suggested that the discrimination of motion from position cues is produced by the same system that processes the location and identity of static stimuli, such that the motion capabilities of this system are a by-product and not the primary role. Support for this comes from the similarity in gradients of the position cue displacement data in Chapter 3 and the gradients of displacement for the acuity of static gratings (Strasburger *et al.*, 2011). Using a conflicting cue Gabor stimuli, Heller *et al.* reported a bias for low temporal frequencies for the aperture cue and high temporal frequencies for the carrier cue (Heller *et al.*, 2021). Reverse phi stimuli show a similar bias, where both gratings and random dot kinematograms are reported in the position cue direction at low temporal frequencies and the phase

cue direction at high temporal frequencies (Gorea, 1994; Edwards & Nishida, 2004). These reverse phi stimuli can be viewed through a stationary aperture, and so produce continuous position cues without a large displacement in the stimulus location, which occurs with multi-frame presentations of Gabor stimuli.

It is difficult to produce conflicting cue stimuli containing phase and position cues from similar features. This is necessary in determining what properties of dissociations in reported motion direction, or motion discrimination thresholds, in conflicting cue stimuli are determined by the phase and position systems, and not by variation between the phase and position cues within the stimuli. The phase and position elements of the missing fundamental grating have different contrast and spatial frequency and the conflicting cue Gabors use different image features, the grating and the Gaussian aperture. In the reverse phi stimuli, the phase cues are objectless, and the position cues have no Fourier components, and the two cues are produced from the same image feature. Reverse phi stimuli do not have the potential position cue confound in Gabor stimuli, where phase displacements of the grating could be processed by the position system, as the features in the Gabor will be displaced by the change in phase. The reverse phi gratings can be viewed through an aperture, so can be displayed with multiple frames, without producing a large bias between the phase and position cue due to changes in the location of the stimulus, or confounding temporal frequency and the location change, as occurs in Gabor stimuli presented across multiple frames. The reverse phi stimulus may therefore may be the best stimulus for mapping psychophysical dissociations between the phase and position systems, particularly for the spatial and temporal frequency dimensions. However, to more convincingly argue for the presence of two separate systems for processing phase and position cues, these data need to be supported by experiments using stimuli that contain solely a phase or position cue that could be used to characterise the systems independently.

Chapter 6: Residual motion abilities following lesions of striate cortex

Declaration. The data collected for Experiments 2B and 3D were collected in collaboration with Professor Howard Hock and Dr. Sang Hong, of Florida Atlantic University, using experimental code and procedures that I designed. Field tests of the scotoma and structural imaging of the lesion were provided by these collaborators.

In this chapter, I explored the residual motion capabilities of a patient with a striate cortex lesion, across a series of conflicting motion cue stimuli. This chapter aimed to replicate and extend the results of Azzopardi & Hock (Azzopardi & Hock, 2011), who reported that a patient with a lesion of striate cortex, producing residual vision mediated predominantly the midbrain pathway, could discriminate motion using phase but not position cues.

6.1 Introduction

The midbrain and cortical pathways

The cortical area 'V5/MT', located on the posterior bank of the superior temporal sulcus in the primate brain, is referred to as the 'cortical area of motion perception' (Zeki, 1980). Signals from the retina reach V5/MT through two pathways; the cortical and the midbrain pathway (see Figure 6.1). These two pathways have vastly different sizes and compositions, with approximately 90% of retinal ganglion cells projecting to LGN, the first stage of the 'cortical' pathway, and are formed predominantly from pα and pβ RGCs (Perry & Cowey, 1984). The remaining 10% of RGCs project to superior colliculus, the first stage along the midbrain pathway, but this projection is formed of predominantly by pγ RGCs with a small proportion of pα RGCs, and no pβ RGCs (Perry *et al.*, 1984).

Despite receiving a smaller retinal projection, the midbrain pathway is active in normal observers. fMRI of the superior colliculus when visual stimuli were displayed reports substantial activity when moving stimuli are displayed (Mikellidou *et al.*, 2018), at contrasts as low as 5% (Schneider & Kastner, 2005), and with spatiotopic organisation (DuBois & Cohen, 2000). Activity of both the midbrain and cortical pathways in normal observers can also be seen in the two peaks of activity in V5/MT as a response to motion stimuli, where the first peak may be driven by the midbrain pathway and the second by the cortical pathway (Ffytche *et al*, 1995; Stevens *et al.*, 2007). Following lesions to V1, which sever the cortical pathway, there is residual activity in the midbrain pathway (macaque; Rodman *et al.*, 1990, Warner *et al.*, 2015, humans; Barleben *et al.* 2014; Tran *et al.*, 2019), as well as

residual motion perception (macaques; Gross *et al.*, 2004, humans; Azzopardi & Cowey, 2001, Sahraie *et al.*, 2008), indicating that the midbrain pathway may be sufficient to mediate some motion perception in the absence of the cortical pathway.

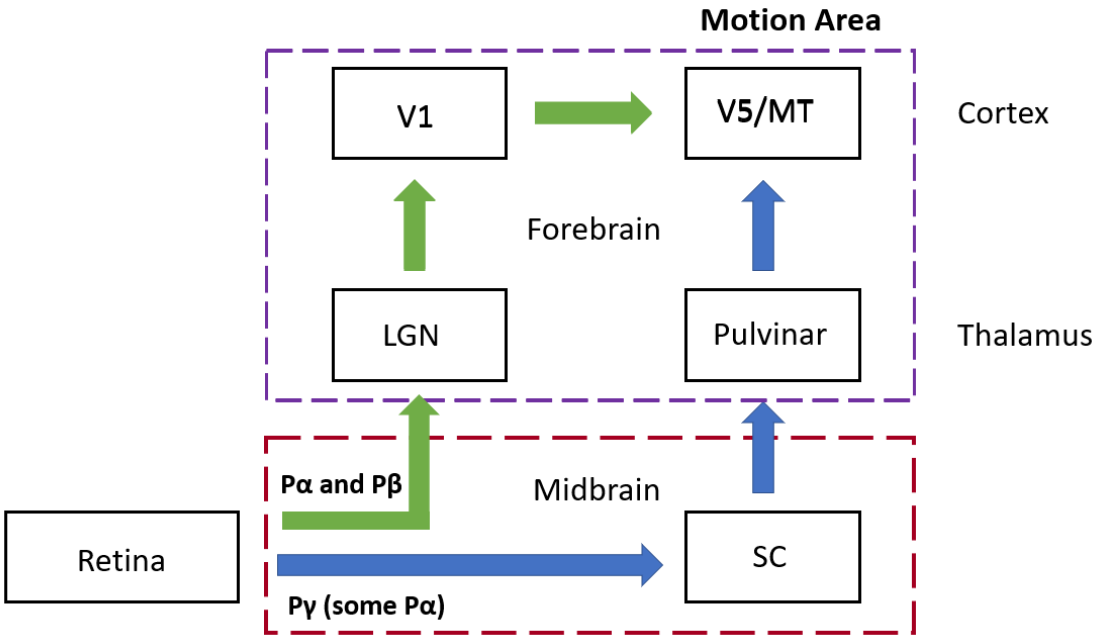


Figure 6.1 : A schematic of the major projections to V5/MT; the cortical projection of pα and pβ RGCs to the LGN, through V1, and to V5/MT, and the midbrain projection of pγ and some pα RGCs to the superior colliculus, which projects to the pulvinar and to V5/MT.

The relative contribution to motion perception from the cortical and midbrain pathways is currently unclear. Ffytche *et al.* suggested that the division of labour between the cortical and midbrain pathway is assigned by stimulus speed, suggesting that the midbrain pathway preferentially processes fast moving, and the cortical pathway, slow moving stimuli, using a drifting checkerboard stimulus that contains both phase and position cues (Ffytche *et al.*, 1995). An alternative hypothesis was put forward by Azzopardi & Hock, who suggest that the activity in the midbrain pathway could reflect the response to phase cues. They observed that for patient GY, who had a V1 lesion and thus no cortical pathway to V5/MT, GY reported motion solely in the phase cue direction (Azzopardi & Hock, 2011).

In normal observers there is a psychophysical dissociation in response to temporal frequency, where sensitivity to position cues declines rapidly with increasing temporal frequency, but remains high for phase cues (Lu & Sperling, 1995). For stimuli that contain both phase and position cues, motion is biased in the phase cue direction at high temporal frequencies, and in the position cue direction for low temporal frequencies (MFG; Georgeson & Harris, 1990; Reverse Phi: Benton et al., 1997). If, as

suggested by Azzopardi & Hock (Azzopardi & Hock, 2011), the midbrain pathway can process phase but not position cues, presenting stimuli at high temporal frequencies should produce activity in the midbrain pathway in normal observers. Using a version of the phi phenomena (Wertheimer, 1912) as a 'pure' phase cue, and a similar stimulus containing both phase and position cues, Pizlo *et al.* (Pizlo *et al.*, 2008) reported preferential activity in the superior colliculus and pulvinar, but not in the cortical pathway, for the pure phase stimulus. This suggests that a lack of sensitivity for position cues in the midbrain pathway respectively, could account for both the results of Ffytche *et al.* (Ffytche *et al.*, 1995) and Azzopardi & Hock (Azzopardi & Hock, 2011).

In this chapter, I investigate whether there is a qualitative difference between normal peripheral motion perception, mediated by both the midbrain and cortical pathways, and the residual motion perception following a lesion of V1, with respect to the sensitivity to phase and position cues for motion. A difference in residual motion perception would support the hypothesis that the midbrain pathway is solely responsive to phase cues, and provide a possible neural substrate for the 'phase system', which I have identified psychophysically in this thesis.

Residual vision following V1 lesions

Patients with lesions of V1 present an opportunity to study the motion properties of the scotoma, where vision is predominantly mediated by the midbrain pathway. These findings can be compared to the motion perception of the sighted field, or of normal observers, to investigate whether the differences in motion perception between the scotoma and sighted field reflect qualitative differences in motion perception, whereby motion in the scotoma is driven solely by phase cues, and no longer responsive to position cues.

Motion detection & discrimination

The detection of moving stimuli in the scotoma of human patients has been demonstrated consistently in the earliest studies of the effects of V1 lesions, (Holmes & Lister, 1916; Riddoch, 1917). Motion detection has been shown for simple stimuli (bars; Barbur *et al.*, 1993; Azzopardi & Cowey, 2001: circles; Blythe *et al.*, 1986; Blythe *et al.*, 1987; Weiskrantz *et al.*, 1995; Riddoch, 1917; Barbur *et al.*, 1980), and for more complex stimuli, such as sinusoid gratings (Wessinger *et al.*, 1999; Azzopardi & Cowey, 2001; Magnussen & Mathiesen, 1989), Gabors (Sahraie *et al.*, 2008), and random dot kinematograms (Azzopardi & Cowey, 2001; Alexander & Cowey, 2008/9).

Motion detection could be accomplished due to sensitivity to temporal frequencies; direction discrimination provides more conclusive evidence that patients are perceiving motion. Direction discrimination has been reported for simple stimuli in many patient studies (bars; Azzopardi & Cowey, 2001; Barbur *et al.*, 1993: circles; Blythe *et al.*, 1986). For more complex stimuli, there is more variability across reported patient abilities. Some studies have reported no direction discrimination abilities in V1-lesioned patients for sinusoidal gratings (Azzopardi & Cowey, 2001), yet at lower spatial frequencies (0.2 cpd), Morland *et al.* (Morland *et al.*, 2004) reported three out of eight V1-lesioned patients could discriminate direction (Morland *et al.*, 2004), and Perenin reported direction discrimination in all five patients, on 0.3cpd gratings (Perenin, 1991). This reduction in spatial frequency could be responsible for increased performance by operating within a parameter space more optimal for the midbrain pathway. Alternatively, at such low spatial frequencies half-cycles of gratings could be perceived as bars, where the direction discrimination of bars is typically much better (Azzopardi & Cowey, 2001).

The stimuli tested in these experiments contain congruent phase and position cues. To assess whether the midbrain pathway is selectively processing phase cues, it is essential to investigate motion discrimination abilities to stimuli containing solely phase or position cues, or conflicting phase and position cues.

Motion without phase cues

If motion can be discriminated in the scotoma from stimuli containing solely position cues, then this would be robust evidence against my theory. Presently, the only 'pure' position cues which have been presented in the scotoma are chromatic stimuli, set to isoluminance (Alexander & Cowey, 2013; Guo *et al.*, 1998). Isoluminant stimuli contain no phase cues in the luminance domain, since they contain no difference in luminance across their extent, so any perceived motion must be due to tracking of the position of features within the stimulus.

Motion without position cues

A possible stimulus to differentiate between motion using phase or position cues in the blind field is a random dot kinetogram (RDK). These consist of a field of randomly positioned dots displaced on each frame. Some percentage of these are displaced in the 'signal' direction, while the remainder move in random trajectories, 'noise'. At 100% coherence, all the dots move in the same trajectory. To minimise the tracking of the position of individual dots, dots usually have a limited 'lifetime',

where, after a specified number of frames, they are removed, and a new dot is inserted at random into the field. The RDK contains a strong phase cue, which increases with coherence, but has minimal position cues due to the number of elements. Multiple elements are hard to track and the random positioning of the dots means that on consecutive frames there are many potential matches for individual dots, as differing combinations of displacement and trajectory could be used to map each dot to any other in the field.

Patients can discriminate between static and moving RDKs in the scotoma (Azzopardi & Cowey, 2001; Alexander & Cowey, 2013). However, direction discrimination, which is the true test of motion perception rather than temporal frequency detection, is inconsistent across patients and studies. Some studies report that patients are unable to discriminate the direction of RDK stimuli, even at high velocities (Azzopardi & Cowey, 2001; Barton & Sharpe, 1997), while other studies have reported accurate direction discrimination (Benson *et al.*, 1998; Giaschi *et al.*, 2003). Imaging in human patients has shown that RDKs do produce activity in V5/MT, but that it resembles activity seen in V1 when displayed in the sighted field, suggesting that integration between the local motion vectors of each dot has not occurred (Ajina *et al*,. 2015). This could explain the ability to discriminate static from moving RDKs, but not discriminate direction (Azzopardi & Cowey, 2001).

Patients with V1 lesions have been shown to achieve direction discrimination in RDKs, following substantial training. 'Visual Restoration Training' (VRT) studies, where patients complete months-long training regimens at home, have reported chance performance on direction discrimination tasks in RDKs on initial testing, but high levels of performance following substantial training (RDKs; Cavanagh *et al.*, 2015; Huxlin *et al.*, 2009; Vaina *et al.*, 2014: Gabors; Ajina *et al.*, 2021). This improved performance on RDKs could reflect an increased ability of V5/MT to integrate between motion vectors for the specific trained region. This explanation is supported by evidence that the enhanced direction discrimination generalises to other forms of motion stimuli, but not to other non-trained areas in the scotoma (Das *et al.*, 2014; Ajina *et al.*, 2021). This indicates that motion carried purely by phase cues can be discriminated in the scotoma, but also highlights the need for simpler stimuli than the RDKs.

Motion with conflicting cues

Conflicting motion stimuli could be used to investigate whether motion perception in the scotoma is driven by phase or position cues, with reported direction revealing which cue is relatively stronger for that presentation. Azzopardi & Hock (2011) produced a reverse phi rendering of the line motion illusion (Hikosaka *et al.*, 1993), where expansion or contraction of a rectangular stimulus is produced

by adding or subtracting a rectangle of opposite polarity to the base rectangle. This produces a stimulus with a change in position, with the direction of the expansion or contraction, and a phase cue in the opposite direction. Patient GY reported motion in the position cue direction in his sighted field, and in the phase cue direction in his blind field for matched eccentric locations (Azzopardi & Hock, 2011). This suggests either that motion perception in the scotoma is produced from phase cues and is blind to position cues, consistent with my theory, or that the removal of V1 has changed the balance of phase and position systems in the scotoma, with the phase cues perceived as relatively stronger.

Replicating this result over a wider range of stimulus sizes and contrasts would provide compelling evidence that there is a qualitative difference between vision in the sighted field and scotoma, whereby the midbrain pathway processes solely phase and not position cues.

In this chapter, I present data from patient HS, who has a quadrantanopia, of the upper left field following an ischemic stroke that lesioned the lower right bank of his striate cortex. These data were collected across two week-long visits to Oxford, and several sessions of remote follow-up experiments, in collaboration with Professor Hock, and Dr. Hong of Florida Atlantic University. Across these sessions, we aimed to characterise patient HS's residual vision in his scotoma, and investigate perceived motion in a range of conflicting cue stimuli. We also conducted a range of experiments investigating the patient's capacity for 'blindsight', performance in the absence of awareness, which are not reported here.

6.2 Method

Patient

Our patient was HS, a 75 year old male who had suffered an ischemic stroke of the posterior cerebral artery 3.5 years prior to testing. This resulted in damage to the lower right bank of the calcarine sulcus, producing quadrantanopia, of the upper left field (Figure 6.2, left and central panels). Clinical field tests carried out one year before the first patient visit to Oxford revealed a dense quadrantanopia, with some damage extending below the meridian and with minimal macular sparing. The quadrantanopia was densest between 5° and 25° left of the vertical meridian (Figure 6.2, right panel).

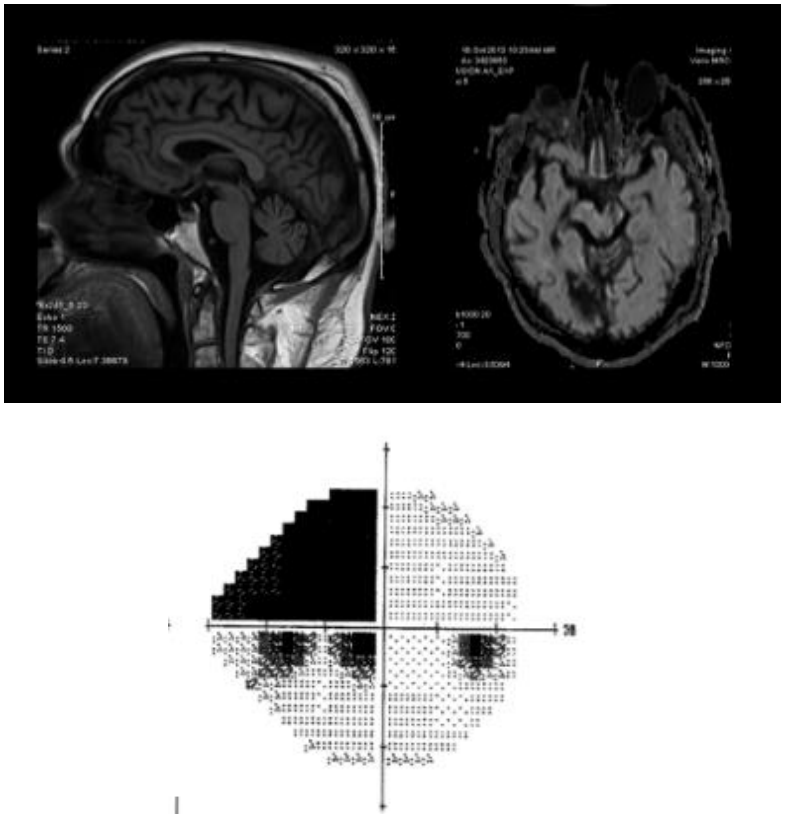


Figure 6.2. Left; Sagittal and axial sections of the structural MRI showing the damage to HS's striate cortex, which occupies the lower right bank of the calcarine sulcus. Right; A Central 24-2 field test of HS's right eye. There is a clear scotoma of the upper left visual field, as expected from imaging, and some areas of impaired vision below the horizontal meridian, but less dense.

Oxford Testing conditions

All experiments bar the reverse phi RDKs (Experiment 2B) and the replication of the rpLMI at different sizes (Experiment 3D), were carried out in Oxford, in October 2018.

These experiments were approved by the Medical Sciences Inter-divisional Research Ethics Committee (IDREC) at the University of Oxford in accordance with the procedures laid down by the University for Ethical Approval of all research involving human participants, under reference R51336/RE001.

Tests were carried out in a dimly lit room. The patient was seated 57cm away from the screen, with the head supported by a chin rest and right eye occluded by an opaque patch. 2 metre square boards were positioned either side of the CRT screen and covered with matte grey felt, to eliminate reflections.

Stimuli were generated in MATLAB 2014, and presented on a 22-inch CRT monitor (Mitsubishi Diamond Pro 2070, 1600 × 1200 pixels, 100 Hz refresh rate) gamma corrected using the ColorCAL MK II colorimeter (Cambridge Research Systems).

The patient's eye movements were monitored by closed circuit television, by the second experimenter. Trials with eye movements were aborted and inserted randomly amongst the remaining trials to ensure all trials were presented.

Static dot stimuli and missing fundamental gratings were viewed through a 5° square aperture and the Missing Fundamental Gratings were viewed through vertical window of 1° by 5°. All other experiments were viewed through a 10° square aperture. All apertures were centred at 11.57° horizontal and vertical from the fovea (diagonal of 16.35°), in the upper left visual field, to avoid overlap with the sighted field or less dense areas of the scotoma. The entire screen was covered by a sheet of matte black card to eliminate screen artefacts. Experiments in the sighted field were tested at the equivalent location in the upper right visual field, through the same aperture as trials in the scotoma.

Florida Testing conditions

Experiments 2B and 3D were carried out by Professor Howard Hock and Dr Sang Hong, at Florida Atlantic University, in January and February 2019.

Tests were carried out in a dimly lit room. The patient was seated 57cm away from the screen, with the head supported by a chin rest and right eye occluded by an opaque patch.

Stimuli were generated in Oxford in MATLAB 2014, and sent to the collaborators in Florida. There, the stimuli were presented on a 4K LCD monitor. As in Oxford, a 10° square aperture of matte black card was fixed to the screen centred at 11.57° horizontal and vertical from the fovea (diagonal of 16.35°), in the upper left visual field.

Fixation was monitored manually via closed-circuit tv, and trials where the patient perceived that their gaze had deviated from fixation were aborted and inserted randomly into the remaining trials to ensure all trials were presented.

General procedure

A 2IFC procedure was used for all experiments bar Experiment 4A (MFG), in which each of two alternative stimuli, which moved in opposite directions, were presented in randomized order in every trial, and the patient was asked to report the interval in which the stimulus moved in a specified direction (e.g. upwards) together with a rating indicating their confidence in their response. The start of each interval was indicated by a tone of 200 ms duration and the end of the

trial by a tone of 200 ms duration an octave higher. Following the two intervals of the stimulus (or one interval for the MFG), a blank interval of the background luminance was displayed until the patient had reported both the interval and their confidence rating.

Before and after stimulus presentation a blank screen was presented set to the background luminance of the task. For all experiments bar reverse phi line motion illusion (rpLMI) at different sizes, and the phi and reverse phi RDKs at different speeds, the patient also reported their awareness of the stimuli, using the Perceptual Awareness Scale (Overgaard *et al.*, 2006). Awareness responses are not discussed, for brevity. The patient carried out 20 trials at a time, with breaks between each block.

Static dots and random dot kinematograms followed the design and procedure of Azzopardi & Cowey (2001), and the line motion stimuli and controls followed the design and procedure of Azzopardi & Hock (2011).

6.3 Experimental Methods and Results

6.3.1 Static dots & bars

In this section I describe the effect of contrast on HS's detection of static stimuli.

Experiments in Section 6.3.1 were carried out solely in HS's scotoma. Informal comparisons in the sighted field at matched eccentricity found that HS could easily detect bar and RDK static stimuli when presented at 10% of threshold contrast from the scotoma.

Experiment 1A: Detection of Static Dot Displays

Method

Procedure

I used a 2afc procedure, with one interval containing the stimulus, and the other interval was blank. Each frame was presented for 1000 ms, and the trial was followed by a blank frame set to minimum luminance until the patient's responses were stored and the next trial was initiated. The patient verbally reported in which interval the stimulus was displayed.

Stimuli

0.5° diameter dots were distributed randomly on a dark background at a density of 0.414 dots/degree. Dots were displayed on a background of minimum luminance, and the dots displayed at 10%, 55% or 100% of maximum luminance (8.13, 44.73, 81.32 cd/m^2), and were redrawn for each

trial. Stimuli followed the same parameters Azzopardi & Cowey, (Azzopardi & Cowey, 2001), and there were 40 trials per contrast level, giving a total of 120 trials. These were conducted in one session, and in separate blocks increasing in contrast.

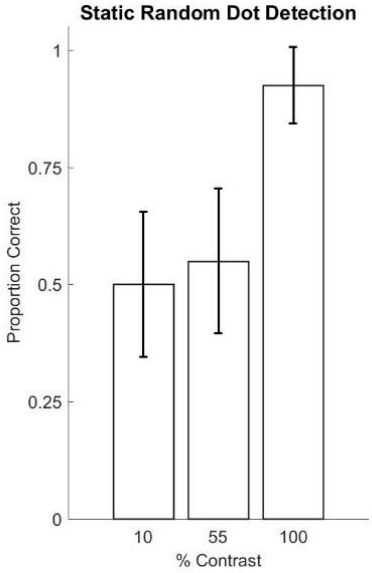


Figure 6.3 Bar graph demonstrating detection of static display of random dots, presented at three contrast levels in HS's scotoma.

Results

Patient HS achieved high proportion correct results for trials presented at higher contrast, as shown in Figure 6.3. A chi-square test for independence was conducted between the distribution of correct and incorrect responses for each contrast level for the static RDK experiment. There was a statistically significant difference in the number of correct and incorrect responses across contrast levels in the static RDKs, $x^2(2) = 19.191$, $p < .001$, Cramer's V = .400. To calculate whether the proportion of correct and incorrect trials differed from chance, I ran a chi square goodness of fit at each contrast level. There was no significant difference from chance at 10 or 55% contrast ($p \geq .537$), but performance was significantly different from chance at 100% contrast, x^2 (1) = 28.900, $p < .001$, demonstrating detection performance significantly higher than chance.

Experiment 1B: Detection of Static Bars

Method

Procedure

I used a 2afc procedure, with one interval containing the stimulus, and the other interval was blank. Each frame was presented for 1000 ms, and the trial was followed by a blank frame set to minimum

luminance until the patient's responses were stored and the next trial was initiated. The patient verbally reported in which interval the stimulus was displayed.

Stimuli

A 1° horizonal bar extended fully across the 5° square aperture, on a dark background. The bar was displayed at 2.1, 9.1 or 20 % of maximum luminance (1.71, 7.40, 16.26 cd/m^2). The static bars were used to assess awareness, with 160, 240 and 200 trials of each contrast respectively, giving a total of 600 trials. These were carried out during the same afternoon, where blocks contained one contrast, but were interleaved in order and different lengths. The patient took breaks between contrast levels, and when fatigued. No block lasted longer than 80 trials without a break.

Results

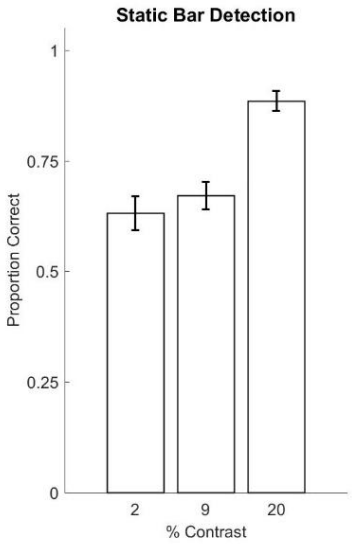


Figure 6.4 Bar graph demonstrating detection of static bars, presented at three contrast levels in HS's scotoma.

Patient HS showed an increase in proportion correct responses for detecting static bars with increased contrast, as shown in Figure 6.4. A chi-square test for independence was conducted between the distribution of correct and incorrect responses for each contrast level for the static bar experiment. There was a statistically significant difference in the number of correct and incorrect responses across contrast levels, x^2 (2) = 36.692, $p < .001$, Cramer's V = .247. To calculate whether the proportion of correct and incorrect trials differed from chance, I ran a chi square goodness of fit at each contrast level. These were significantly different at each contrast level, 2% contrast; x^2 (1) = 11.025, p =.001, 9% contrast; $x^2(1) = 28.017$, $p < .001$ and 20% contrast $x^2(1) = 118.580$, $p < .001$, demonstrating detection performance significantly higher than chance.

6.3.2 RDKs

In this section I describe HS's direction discrimination for RDKs, rendered as phi and reverse phi stimuli.

Experiment 2A: Direction Discrimination in RDKs

Method

Stimuli

Three experiments were run using random dot kinematograms (RDKs). Across these three experiments I used 0.5° diameter dots distributed randomly on a minimum luminance background (~0 cd/m^2), at a density of 0.414 dots/degree, with the dots displayed at maximum luminance (81.32 cd/m^2). The dots were displaced on each frame, with 100% coherence and infinite dot lifetime, to produce 4 speeds; 4 deg/s, 20 deg/s, 32 deg/s and 64 deg/s.

- **Motion detection task;** Dots moved upwards in one interval, and were static in the other.
- **Direction discrimination of transient RDKs;** Dots moved upwards in one interval and downwards in the other. The dots were initially displayed as static, and then began moving for the remainder of the 1000 ms interval. This includes a motion onset cue.
- **Direction discrimination of persistent RDKs;** Dots moved upwards in one interval and downwards in the other. The dots were moving for the entirety of each 1000 ms interval.

Procedure

Each frame was presented for 1000 ms and was not redrawn between frame 1 and frame 2 of the interval, to eliminate onset and offset cues. 20 trials were presented for each of the 4 levels of speed, producing 80 trials per task. I used a 2afc procedure. In the dot detection task one interval contained the moving RDKs and the other the same RDK display, but static. In the direction discrimination tasks one interval contained upward and the other downward motion. The patient verbally reported in which interval the stimulus was moving (detection), or moving upward (direction discrimination).

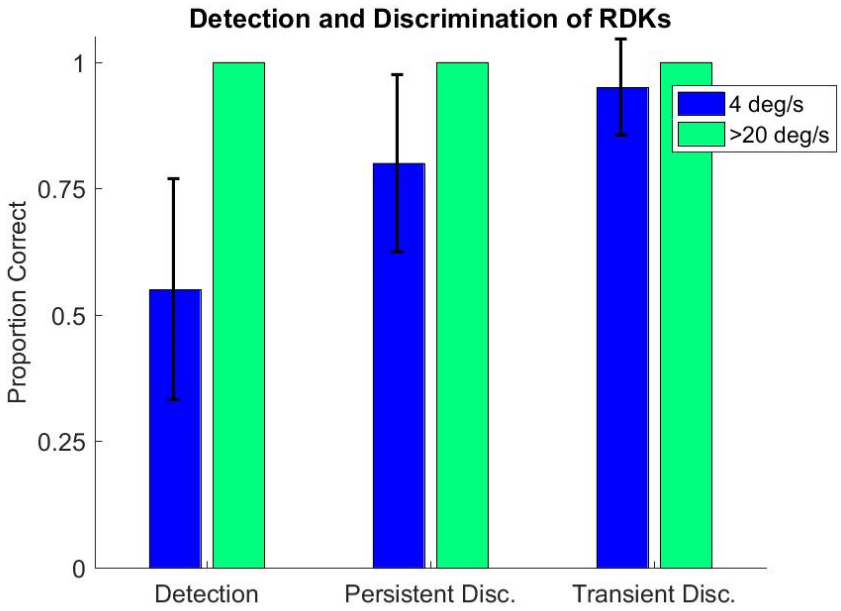


Figure 6.5. Bar chart presenting HS's performance across three tasks; detection of motion, direction discrimination of persistent dots, and direction discrimination of transient dots. 4 speeds were presented for each task (4, 20, 32 and 64 deg/s) but performance was at ceiling for all speeds ≥ 20 deg/sand so are presented together.

Results

Patient HS had substantially lower proportion correct for trials where RDKs drifted at 4 degrees per second, than for trials when the dots drifted at speeds of 20 degrees per second or higher, and this pattern was seen for detection of moving vs. static dots, and for the direction discrimination of RDKs both when dots were presented persistently or transiently (Figure 6.5). Responses were cast as a contingency table on the proportions of correct and incorrect trials, across the three tasks and four speed levels. There was an overall effect of speed across the three tasks, x^2 (3) = 44.602, $p < .001$, Cramer's V = .431, and a significant effect of speed, in two of the three task types; Detection, x^2 (3) = 30.423, $p < .001$, Cramer's V = .617, Persistent Discrimination, x^2 (3) = 12.632, p =.006, Cramer's V = .397, but not in Transient Discrimination, x^2 (3) = 3.038, p =.386, Cramer's V = .195.

I explored the significant effects of speed for each task, evaluating the squared adjusted residuals as a chi square statistic with 1 degree of freedom. In the Detection task only the lowest speed, 4 deg/s, was significant, x^2 (1) =5.5^2 $p < .001$, and all higher speeds were not significant, p =.0719. Similarly, in the Persistent Discrimination task, there was a significant difference in the distribution of responses compared to expected only for 4 deg/s, $x^2(1) = 3.6^2$, $p < .001$, and not significant for higher speeds, p = .2301. For the Transient Discrimination task, there was no significant difference across the 4

speeds ($p \geq .0891$). Across the tasks, there was a significant deviation from the expected distribution for 4 deg/s (x^2 (1) = 6.7^2. $p < .001$) and for higher speeds (x^2 (1) = 2.2^2, $p = .0278$).

Performance was at ceiling for trials of 20 deg/s and above, with all trials reported in the correct direction. To assess whether RDKs for each of the three tasks displayed at 4 deg/s resulted in a performance greater than chance, I ran a chi square goodness of fit for all of the three tasks, and assessed significance using a Bonferroni corrected alpha level to account for multiple comparisons (0.05/3 = 0.0167). In the Detection task, performance at 4 deg/s was not significantly different from chance ($p = .655$), but in the Persistent Discrimination task, performance was significantly above chance performance ($x^2(1) = 7.200$, $p = .007$), and also in the Transient Discrimination task, ($x^2(1) = 16.200$, $p < .001$).

Experiment 2B: Reverse phi RDKs

Method

Stimuli

0.5° diameter dots were distributed randomly at a density of 0.414 dots/degree. 50% of dots were at maximum contrast, and 50% at minimum contrast, and displayed on a mean luminance background. The dots had 100% coherence and infinite dot lifetime. In the phi condition, dots retained polarity on each frame, and in the reverse phi condition, dots changed polarity on each consecutive frame. Stimuli were displayed for 1000 ms, and at 4 speeds; 4, 20, 32 or 64 deg/s. Phi and reverse phi stimuli were interleaved within blocks, and different speeds were presented in separate blocks.

Procedure

A 2afc forced choice procedure was used, with the patient reporting in which interval the motion was upwards. Awareness was measured on a two point scale, with the patient reporting ‘aware’ or ‘unaware’, but this is not reported here. There were 60 trials carried out per condition, with trials carried out across three sessions, which were carried out within one week. Different speeds were presented in separate blocks.

Results

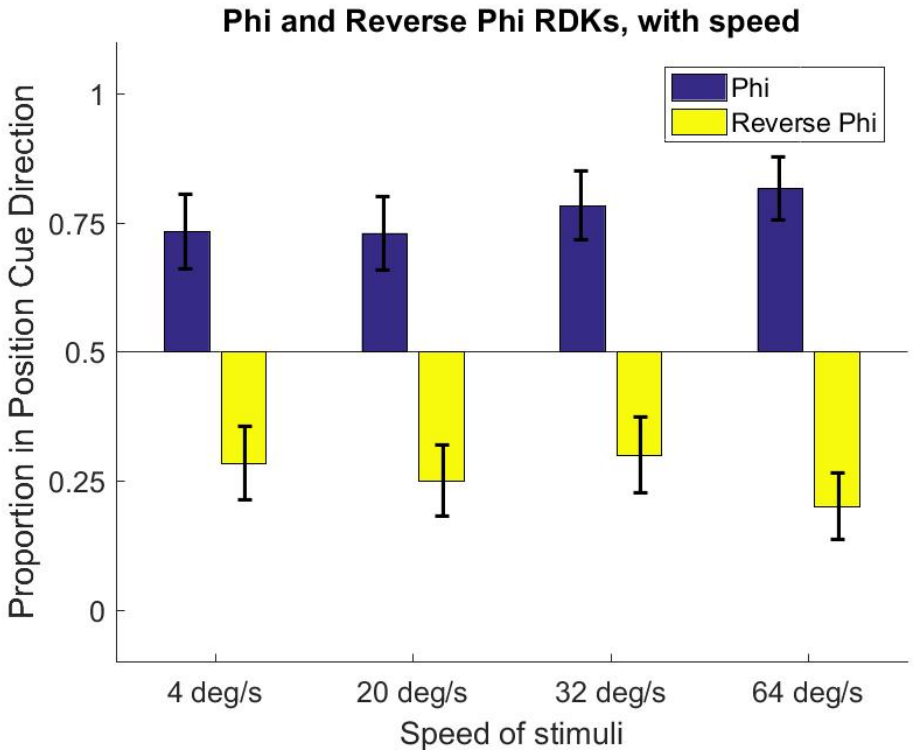


Figure 6.6. Bar chart illustrating HS's performance on phi and reverse phi renderings of random dot kinematograms (RDKs), at 4 different speeds. For the phi condition, both phase and position cues signal motion in the same direction. However, for the reverse phi condition, the phase and position cues signal motion in opposing directions; performance >0.5 corresponds to motion reported in the position cue direction, and performance < 0.5 corresponds to motion reported in the phase cue direction.

Patient HS was able to discriminate the direction of both phi and reverse phi RDKs, presented at all 4 speeds (Figure 6.6). Patient HS reported motion in the expected direction for the phi rendered RDKs, where both the phase and position cues signalled motion in the same direction, but showed a reversal of reported motion direction in the reverse phi RDKs. This can be seen in Figure 6.6, where data were plotted in terms of proportion reported in the position cue direction. Patient HS reported the reverse phi RDKs substantially below chance levels, indicating that he was reporting motion using the phase, and not the position, cue. I ran a chi square test of independence on the proportions of correct and incorrect trials, across the two stimulus types and four speed levels.

There was no overall effect of speed on the distribution of correct and incorrect responses, $x^2(3) = .634$, $p = .889$, Cramer's V = .036, and this was also not significant within the phi RDKs, $x^2(3) = 1.677$, $p = .642$, Cramer's V = .084, or within the reverse phi RDKs, $x^2(3) = 1.827$, $p = .609$, Cramer's V = .087.

To assess whether performance was different from chance, I ran a chi square goodness of fit for each task at each of 4 speeds, and assessed significance using a Bonferroni corrected alpha level

(0.05/ 8 = .0063). All conditions were significantly different from chance performance ($p \leq .002$), with all phi RDKs reported in the forwards direction, and all reverse phi RDKs in the reversed direction, with specific chi square and significance values reported in table 6.1.

Table 6. 1. Table displaying the direction discrimination results for phi and reverse phi RDK stimuli.

Task		4 deg/s	20 deg/s	32 deg/s	64 deg/s
Phi	x^2 value	13.067	13.067	19.267	24.067
	p value	<0.001*	<0.001*	<0.001*	<0.001*
Reverse Phi	x^2 value	11.267	15.000	9.600	21.600
	p value	0.001*	<0.001*	0.002*	<0.001*

I also analysed the performance scores, translated the proportion correct for both phi and reverse phi conditions onto the same 0.5 to 1 axis, and ran a chi square goodness of fit for each task at each of 4 speeds. There was no significant difference between stimulus condition ($p = .448$), and, when analysed by speed level, no significant difference between conditions at any of the four speeds ($p \geq .297$).

6.3.3 Line Motion Illusion

In this section I used the phi and reverse phi line motion stimuli used in Azzopardi & Hock (2011) to investigate whether patient HH exhibited the same phase cue dependency in his scotoma as patient GY, when tested on comparable stimuli.

Experiment 3A: Phi and reverse phi line motion illusion

Method

Stimuli

For all experiments, individual rectangles were 1.33° by 4°, with the extended rectangular stimulus at 1.33° by 8°. These are slightly smaller than those used in Azzopardi & Hock (Azzopardi & Hock, 2011), who used 2° by 6° rectangles, but with the same aspect ratio.

Line motion illusion (LMI)

The LMI is a two-frame stimulus constructed from two rectangles of the same luminance. There are two conditions: expansion and contraction (see Figure 6.7). In expansion, the first frame consists of a single rectangle, and, in the second frame, the second rectangle is presented immediately above or below the first rectangle, creating a percept of the rectangle extending upwards or downwards respectively. The contraction condition consists of the same procedure in the reverse order; the two rectangles are presented adjacently, and either the top or bottom rectangle is removed on the second frame. The percept produced is of a single rectangle contracting downwards or upwards respectively. The LMI is presented on a background of minimum luminance (0 cd/m^2) and rectangles presented at maximum luminance (81.32 cd/m^2).

Reverse phi line motion illusion (rpLMI)

The rpLMI is identical to the LMI, except that the second rectangle, which is either added on the second frame (expansion condition), or removed on the second frame (contraction), is of the opposite polarity of the base rectangle (see Figure 6.8).. The rpLMI is presented on a background of mean luminance (40.66 cd/m^2). 50% of rpLMI trials have the base rectangles at maximum luminance, and the additional rectangle at minimum luminance, and 50% the opposite.

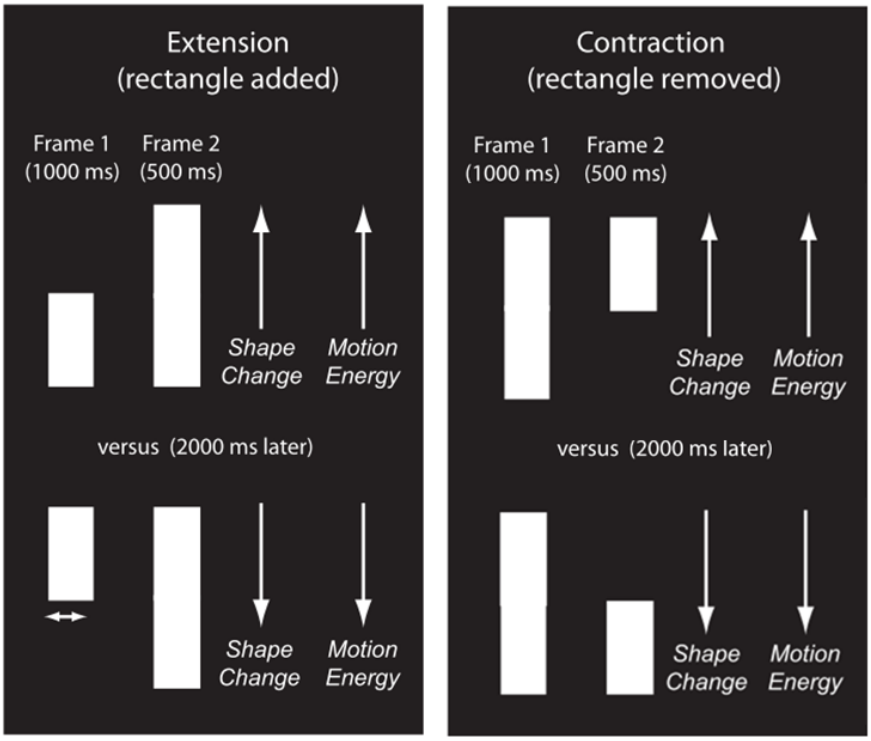


Figure 6.7 Schematic of the phi LMI, adapted from Azzopardi & Hock (2011), where 'shape change' refers to the position cue, and 'motion energy' to the phase cue. The top row of both conditions demonstrates the interval containing upward motion, and the bottom row of both conditions represents the interval containing downward motion.

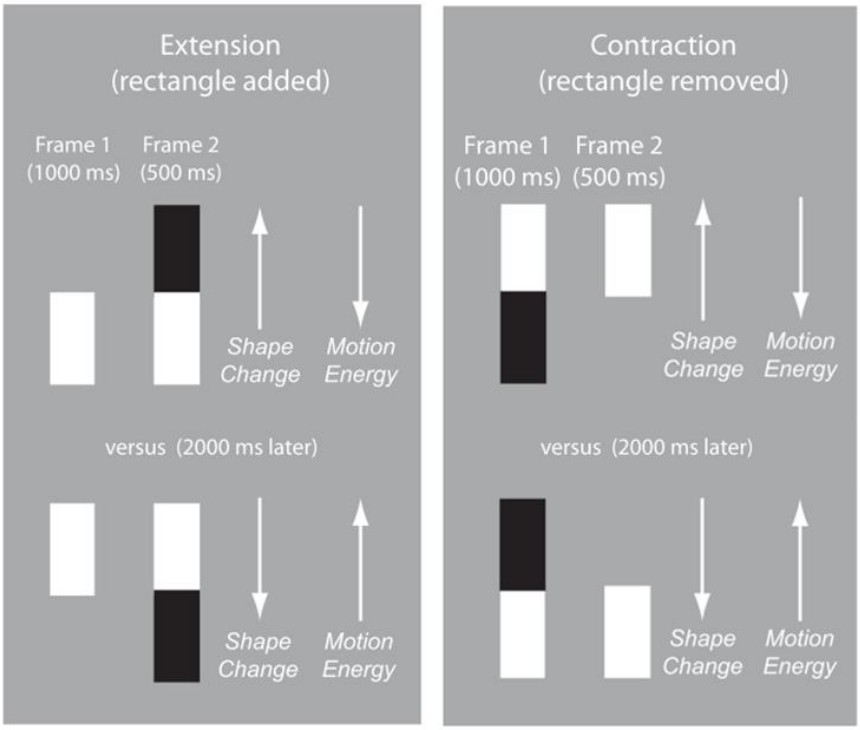


Figure 6.8 Schematic of the reverse phi LMI, adapted from Azzopardi & Hock (2011), where 'shape change' refers to the position cue, and 'motion energy' to the phase cue. The top row of both conditions demonstrates the interval containing upward motion in the position cue direction, and downward motion in the phase cue direction, and the bottom row of both conditions represents the interval containing downward motion in the position cue direction and upward motion in the phase cue direction.

Procedure

For phi and reverse phi LMI, each trial lasted 5000 ms; frame 1 was presented for 1000 ms, frame 2 for 500s, and there was a 2000 ms inter-interval blank screen to prevent apparent motion signals between intervals 1 and 2. In both phi and reverse phi LMI, I used a 2afc procedure in which an expanding and a contracting stimulus were presented in randomised order, and the patient was asked to report in which interval he perceived upward motion.

HS completed 60 trials each for the expansion and contraction of the phi rendering of the LMI, and 80 trials each for the expansion and contraction of the reverse phi rendering of the LMI.

Results

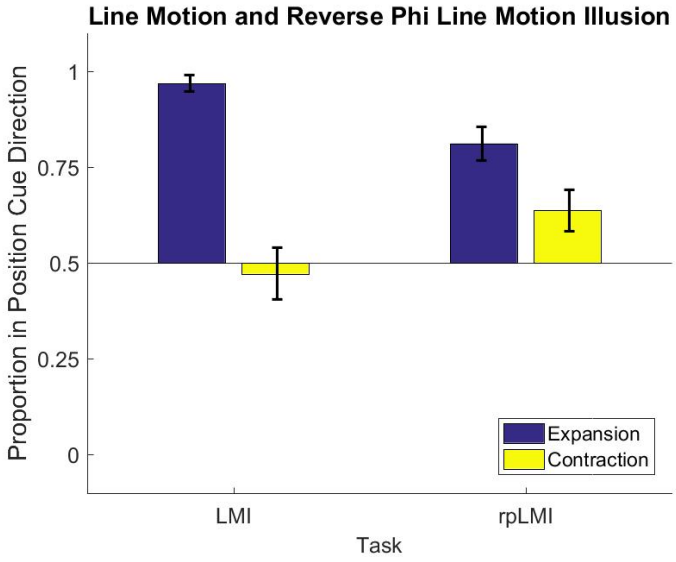


Figure 6.9. Bar chart illustrating HS's performance on the expansion and contraction trials of the phi and reverse phi versions of the Line Motion illusion. For the phi condition both the phase and position cue signal motion in the same direction. However, for the reverse phi condition the phase and position cues signal motion in opposing directions; performance >0.5 corresponds to motion reported in the position cue direction and performance < 0.5 corresponds to motion reported in the phase cue direction.

Patient HS reported the expansion trials of both the phi and reverse phi rendered versions of the LMI in the position cue direction (Figure 6.9). However, for the contraction trials HS reported motion in neither the phase cue or the position cue direction for the phi rendering of the LMI, but in the position cue direction for the reverse phi rendering of the LMI (Figure 6.9). I ran a chi square test of

independence on the proportions of trials reported in the position cue direction, across the two stimulus types (phi and reverse phi) and the two motion types (Expansion and Contraction).

There was a significant difference in the overall distribution of correct and incorrect responses across the two stimulus types and motion types, $x^2(1) = 34.788$, $p < .001$, Cramer's V = .352. There was a significant difference between the contraction and expansion conditions in the Line Motion Illusion, $x^2(1) = 38.311$, $p < .001$, Cramer's V = .565, and between the contraction and expansion conditions in the reverse phi rendering of the Line Motion Illusion, $x^2(1) = 6.144$, $p = .021$, Cramer's V = .196.

To assess whether performance was different from chance, I ran a chi square goodness of fit for the phi and reverse phi LMI conditions for each of the two motion types (Expansion and Contraction), and assessed significance using a Bonferroni corrected alpha level (0.05/ 4 = .00125). In the phi version of the Line Motion Illusion, performance was significantly different from chance in the expansion condition, $x^2(1) = 57.246$, $p < .001$, but not in the contraction condition, $p = .164$. In the reverse phi version of the LMI, performance was significantly different from chance in the expansion condition, $x^2(1) = 31.250$, $p < .001$, but not in the contraction condition, when using the Bonferroni corrected significance level, $p = .014$.

As HS was reporting motion in the position direction for rpLMI, I did not formally assess his performance in the sighted field. However, I ran a few informal trials and found that he had ceiling performance on the LMI and reported motion in the position cue direction, for both contraction and expansion on the reverse phi condition. There was no apparent difference in contraction and expansion trials.

Experiment 3B: Low contrast reverse phi LMI

I replicated this experiment presenting the stimuli at 12% contrast, with 40 trials each for the expansion and contraction conditions for reverse phi renderings of the LMI. To assess whether performance was different from chance, I ran a chi square goodness of fit for expansion and contraction conditions, and found that performance was at chance for all conditions, with 55% of trials in the correct direction in both expansion and contraction conditions, $p = .527$.

Experiment 3C: Control LMI tasks

Method

Stimuli

Control tests for shape cues and attention were carried out using the same size rectangles as the control. The Shape Control used rectangles of paired high and low luminance displayed on the mean

luminance screen, and the attentive control used rectangles of maximum luminance presented on a minimum luminance screen.

Shape Control

Frames consisted of a single rectangle, or the extended rectangle where the single rectangle was now presented adjacently with a second rectangle of opposite polarity (Figure 6.10). Patient HS was asked to report in which interval the shape extended in length. In each trial, one interval is a contraction and one an expansion, to match the direction of the phase cue across trials. The position cue was also matched across intervals, but signalling motion in the opposing direction to the phase cue. Therefore, the task cannot be accomplished by using motion direction, whichever cue is used. The change in shape between frames is either upwards or downwards. The first frame was presented for 1000 ms, and second for 500 ms.

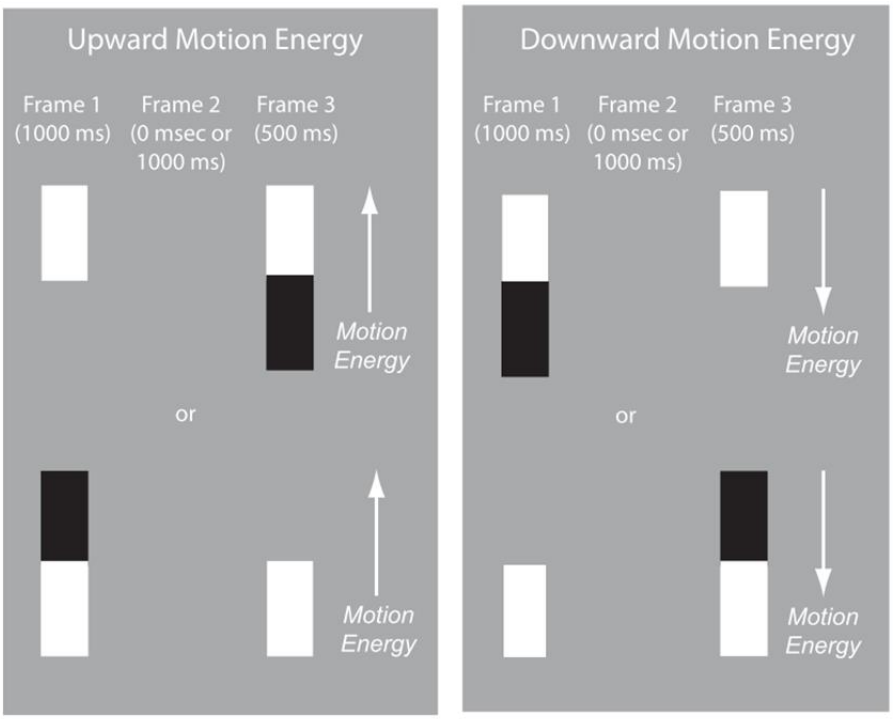


Figure 6.10. Schematic adapted from Azzopardi & Hock (2011) describing the shape change control task, where patient HS was asked to report in which interval the shape extended. In the left and right hand panels the top row displays the extending stimulus, and the bottom row displays the contracting stimulus. In the left panel, where motion energy refers to the phase cue, both the upper and lower rows have the phase cue signalling motion in the upward direction, and the position cue signalling motion in the downward direction. Likewise in the right panel, the phase cue is signalling motion in the downwards direction for both the upper and lower rows, and the position cue is signalling upward motion for both.

Attentive Tracking

Each interval consists of a single frame, with a single rectangle presented either in the upper or lower position. Each frame was presented for 1000 ms, with an inter-frame interval of 2000 ms, to avoid inducing apparent motion. Frames were presented on a mean luminance background, with rectangles on each trial having the same luminance. 50% of trials had rectangles presented at either minimum (0 cd/m^2) and 50% at maximum luminance (81.32 cd/m^2). The stimuli could change from an upper to lower presentation ('downward'), or vice versa ('upward'), (see Figure 6.11).

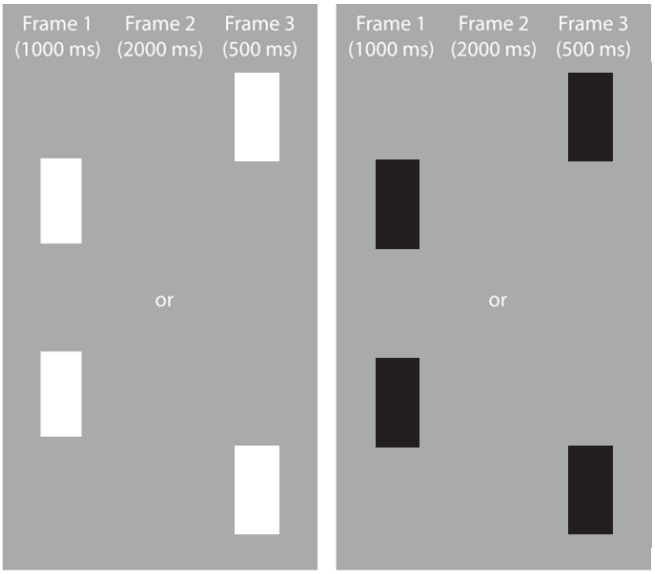


Figure 6.11. Attentive Tracking control, in high luminance rendering (left panel) and low luminance rendering (right panel). In both panels the shape was displaced upward in the frames in the top row, and downward in the frames in the bottom row.

Procedure

80 trials were presented for the shape and Attentive Tracking controls, with half in the downwards and half in the upwards direction. I used a 2afc procedure, where the patient reported in which interval he perceived the shape to increase in length (shape change) or upward motion (Attentive Tracking).

Results

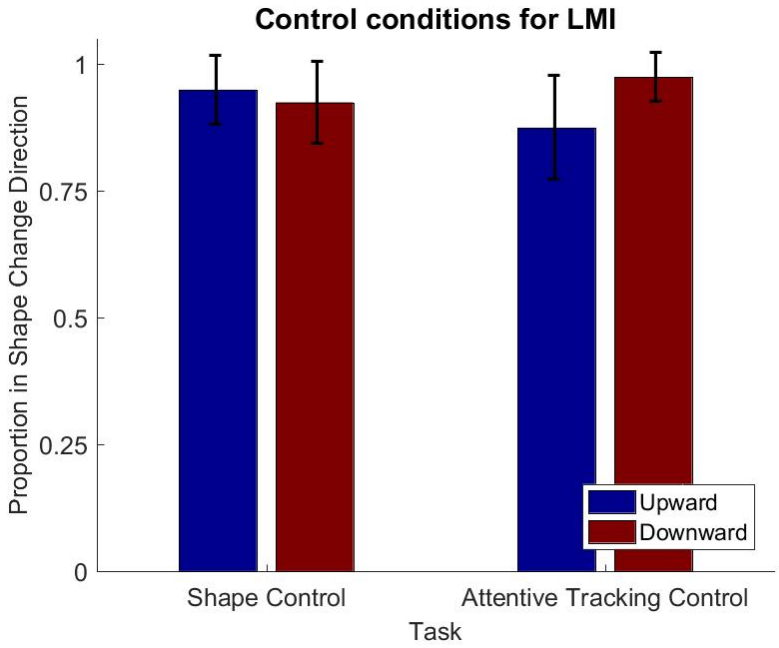


Figure 6.12. Bar charts displaying HS's performance on the two control tasks for the Line Motion Illusion, presented in his scotoma.

Patient HS had high performance for both the upward and downward versions of the Shape Change and the Attentive Tracking control tasks (see Figure 6.12). The Shape Control and Attentive Tracking control were analysed separately. For both control conditions, I ran a chi square test of independence to compare performance on the upward and downward conditions.

There was no significant difference between performance in the upward and downward trials for the Shape Control (p =.644) or the Attentive Tracking control (p =.090). To assess whether performance was different from chance, I ran a chi square goodness of fit for the control conditions for each of the directions, and assessed significance using a Bonferroni corrected alpha level (0.05/ 4 = .00125). Performance is significantly above chance in the Shape Control ($x^2(1) = 61.250$, $p < .001$) and in the Attentive Tracking control ($x^2(1) = 57.800$, $p < .001$).

Experiment 3D: Reverse Phi LMI at Different Sizes

Stimuli

I investigated whether the size of the rpLMI would affect HS's performance. Interleaved contraction and expansion trials of the rpLMI were presented with each of the 4 sizes tested in a separate block of trials. There were 20 trials per condition, a total of 80 expansion and 80 contraction trials. The

ratio of the original rpLMI stimulus (rectangles; 4° by 1.33°) was preserved. Stimuli heights were; 1, 2.34, 3.7 and 5.05 degrees, with stimulus widths; 0.167, 0.39, 0.62, 0.84 degrees.

Procedure

4 sizes of the rpLMI were presented in separate blocks, with the contraction and expansion trials interleaved. I used a 2afc procedure where the patient reported in which interleaved trial the stimulus moved upward.

Results

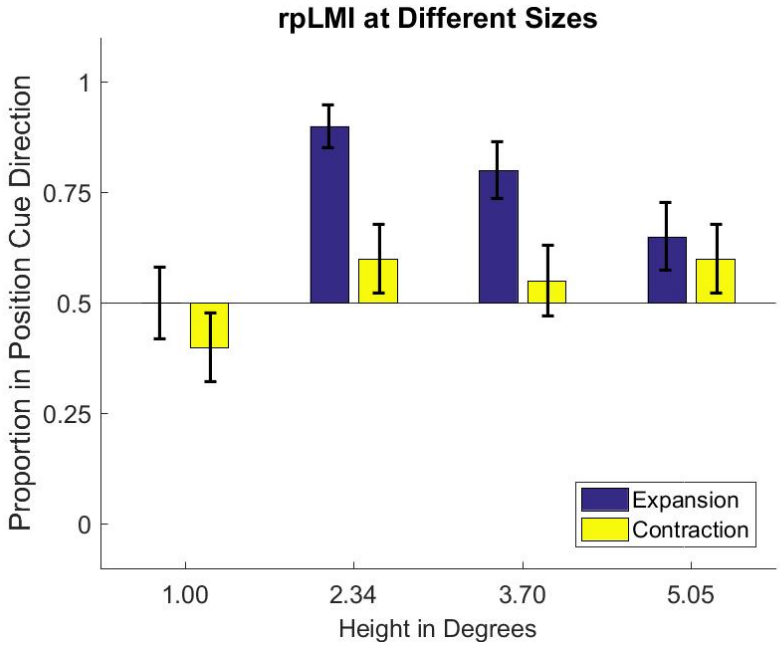


Figure 6.13. Bar graph displaying HS's performance across 4 sizes of the reverse phi rendered Line Motion Illusion, for expansion and contraction conditions. The phase and position cues signal motion in opposing directions; performance > 0.5 corresponds to motion reported in the position cue direction, and performance < 0.5 corresponds to motion reported in phase cue direction.

When the reverse phi rendering of the Line Motion Illusion was presented at different sizes, patient HS showed varying results. At the smallest size, with rpLMI at a height of 1 degree of visual angle, patient HS's performance was at chance for both the expansion and contraction conditions. At the three larger sizes (2.34, 3.70 and 5.05 degrees of visual angle), patient HS was at or near chance performance for the contraction condition, but above chance for the expansion condition. In these three conditions patient HS reported motion in the position cue direction (see Figure 6.13). I ran a chi square test of independence on the proportions of correct and incorrect trials, across the two stimulus types and four size conditions.

There was a significant effect of size on performance for the expansion condition, $x^2(3) = 17.934$, $p < .001$, Cramer's V = .335, but no significant effect of size in the contraction condition, $p = .229$.

To assess whether performance was different from chance I ran a chi square goodness of fit for the expansion and contraction conditions for each of the stimulus sizes, and assessed significance using a Bonferroni corrected alpha level (0.05/ 8 = .0063). In the expansion condition performance was at chance for the smallest size, $p = 1.00$, and the largest size, $p = .058$, but above chance for 2.34 degrees ($x^2(1) = 25.6$, $p < .001$), and for 3.70 degrees, ($x^2(1) = 14.400$, $p < .001$). In the contraction condition, performance was at chance for all four stimulus sizes, $p \geq .206$.

6.3.4 Missing Fundamental Gratings

In this section I investigated whether patient HS was able to discriminate the direction of Missing Fundamental Gratings, and if he reported motion in the phase or position cue direction, and whether this varied with stimulus speed and contrast. To my knowledge this is the first instance of a patient with a lesion to V1 discriminating motion direction in Missing Fundamental Gratings.

Experiment 4A: Missing Fundamental Gratings

Method

Stimuli

The missing fundamental grating was produced by subtracting a sine wave of 0.3 cpd from a matched frequency square wave, with contrast adjusted by pi/4, to match the amplitude of the fundamental frequencies. The resultant grating was shifted by a quarter of a cycle on each frame, to produce the quadrature shift that produces the aliasing in the third harmonic, and was presented for 16 frames. The grating was displayed at two temporal frequencies, 1.25 or 3.75 Hz. The gratings were presented at 12% and 50% in the blind field, and 5%, 12% and 50% in the sighted field and modulated around half maximum luminance (40.67 cd/m^2). The gratings were displayed on a background of the same luminance. Each contrast level was displayed in a separate block.

Procedure

Motion discrimination of the gratings was tested using a 1-interval task (1IFC), in which HS was asked to report the direction of motion of the grating presented (up or down), and an awareness rating using the PAS, not reported here. There were 40 trials per condition in the sighted field, and 80 trials per condition in the blind field.

Results

Scotoma

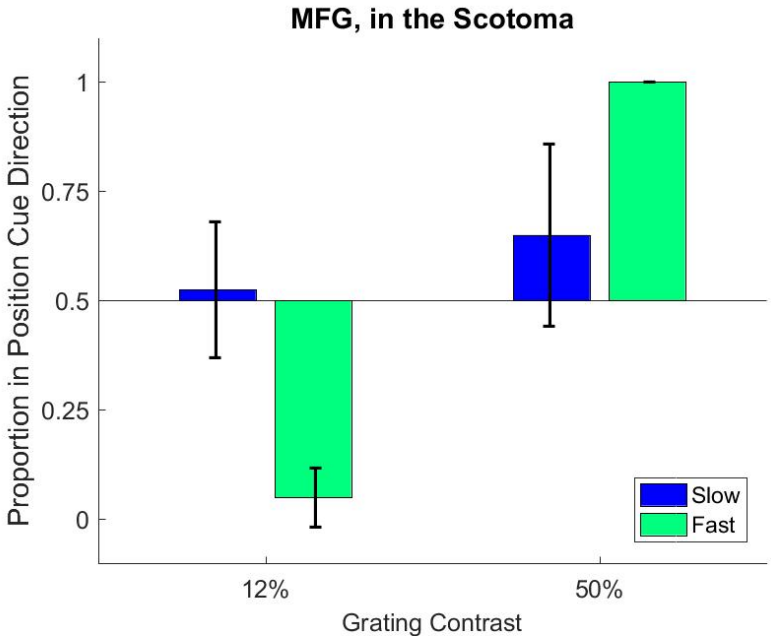


Figure 6.14. Bar graphs illustrating the direction that HS reports motion in when stimuli are presented in the scotoma, for Missing Fundamental Gratings presented at slow and fast speeds, for various contrast levels. Performance >0.5 corresponds to motion reported in the position cue direction, and performance < 0.5 corresponds to motion reported in phase cue direction.

When the MFG was displayed in patient HS's scotoma, he reported motion in the phase cue direction for the 12% contrast MFG presented at 3.75 Hz, and in the position cue direction for the 50% contrast MFG presented at 3.75 Hz. When the MFG was presented at 1.75 Hz (slow condition), patient HS reported motion neither in the phase or position cue direction, with overall performance at chance (see Figure 6.14).

I ran a chi square test of independence on the proportions of correct and incorrect trials, which were operationalised using the difference from chance, in either direction, across the two contrast levels and two speeds. There was a significant difference between performance across the fast and slow speeds in the low contrast condition, $x^2(1) = 44.058$, $p < .001$, Cramer's V =.525, and a significant difference between speeds in the higher contrast condition, $x^2(1) = 33.393$, $p < .001$, Cramer's V = .461.

To assess whether performance was different from chance, I ran a chi square goodness of fit for the 12% and 50% contrast conditions, for the two stimulus speeds, and assessed significance using a Bonferroni corrected alpha level (0.05/ 4 = .00125). In the 12% contrast condition, performance

was not significantly different from chance at slow speeds, p =.655, but was significantly different from chance at fast speeds, $x^2(1) = 64.800$, $p < .001$. In the 50% contrast condition, performance was at chance for the slow speed, p =.007, and at ceiling performance for the high speed.

Sighted Field

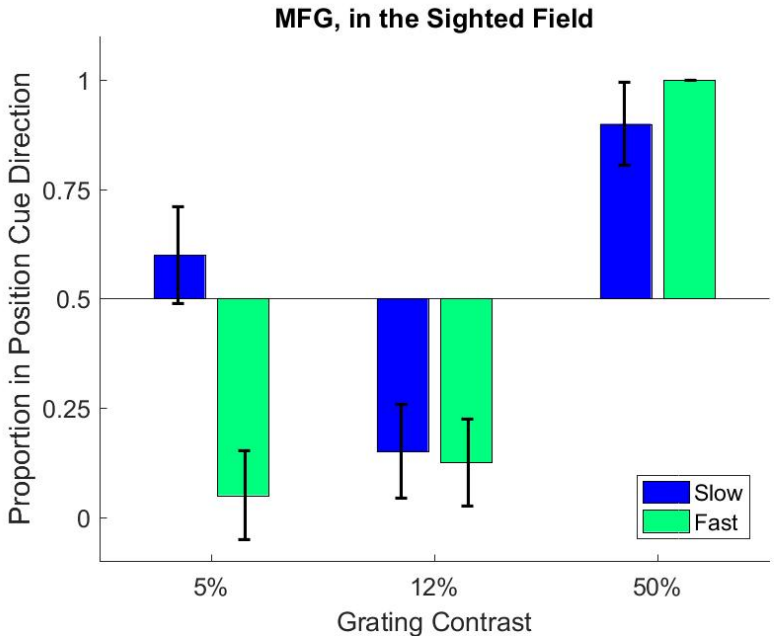


Figure 6.15. Bar graphs illustrating the direction that HS reports motion in when stimuli are presented in his sighted field, for Missing Fundamental Gratings presented at slow and fast speeds, for various contrast levels. Performance >0.5 corresponds to motion reported in the position cue direction, and performance < 0.5 corresponds to motion reported in phase cue direction.

When the MFG was displayed in patient HS's sighted field, he reported motion in the phase cue direction for the 5% and 12% contrast MFG presented at 3.75 Hz. When presented at 5% contrast and at 1.75 Hz (slow condition) patient HS's performance was at chance, with overall motion reported in neither the phase or position cue direction, although for the same slow speed, patient HS reported the MFG in the phase cue direction when displayed at 12% contrast. At the highest contrast presented, 50% contrast, patient HS reported motion in the position cue direction for both the slow (1.75 Hz) and fast (3.75 Hz) speeds (see Figure 6.15).

I ran a chi square test of independence on the proportions of correct and incorrect trials, across the three contrast levels and two speeds.

There was a significant difference between performance across the fast and slow speeds in the 5% contrast condition, $x^2(1) = 27.578$, $p < .001$, Cramer's V = .587, but no significant difference in performance with speed at the 12% contrast condition, $p = 1.000$, or in the 50% contrast condition, $p = .116$.

To assess whether performance was different from chance, I ran a chi square goodness of fit for the 5%, 12% and 50% contrast conditions, for the two stimulus speeds, and assessed significance using a Bonferroni corrected alpha level (0.05/ 6 = .0083). In the 5% contrast condition, performance was not significantly different from chance at slow speeds, $p = .206$, but was significantly different from chance at fast speeds, $x^2(1) = 32.400$, $p < .001$. In the 12% contrast condition, performance was significantly different from chance at slow speeds, $x^2(1) = 19.600$, $p < .001$, and at fast speeds, $x^2(1) = 22.500$, $p < .001$. Similarly, in the 50% contrast condition, performance was significantly different from chance at slow speeds, $x^2(1) = 25.600$, $p < .001$, and was at ceiling performance for fast speeds.

6.4 Discussion

Visual abilities of HS

I confirmed the perimeters of HS's scotoma that were first revealed by the field test in the laboratory. His experience of visual events in this region was consistent with descriptions from other V1 lesioned patients in the literature (Riddoch, 1917), remarking that I was asking him to report the direction of 'shadows moving in the dark', and that "stimuli get lost [in the scotoma] when they stay still".

HS's residual vision exhibits some properties that are typical of patients with V1 lesions, such as lower contrast sensitivity in the scotoma than the matched sighted periphery, which I assessed informally with the static bars and RDKs. HS exhibits a loss of sensitivity to low temporal frequencies, demonstrated in his performance on RDKs and the MFG. This impairment at low temporal frequencies is consistent with previous V1 lesioned patients (Sahraie *et al.*, 2008). This tuning for higher temporal frequencies is seen in single unit recordings of the pα (parasol) and smooth monostratified RGCs temporal frequency sensitivity (Crooke *et al.*, 2008), and in cells in superior colliculus (Marocco & Li, 1977), so is consistent with the classic 'processing via midbrain' explanation for residual vision in V1 patients.

HS also demonstrates residual visual abilities beyond those of other V1-lesioned patients. Most V1-lesioned patients are unable to discriminate direction on high contrast narrowband gratings, (Azzopardi & Cowey, 2001), yet HS was able to report motion in the broadband Missing

Fundamental Grating, even at 12% contrast. HS was also able to discriminate the direction of RDKs, contrary to some established findings (Azzopardi & Cowey, 2001). These results suggest that HS has either a higher level of residual motion compared to other reviewed patients, or a qualitative difference in his scotoma.

In all three of the conflicting cue stimuli tested, HS was able to report motion in both the phase cue and the position cue direction in his scotoma. On the reverse phi line motion illusion, I failed to replicate the results of Azzopardi & Hock (2011), with patient HS reporting motion in the position cue direction on the rpLMI, across all tested sizes. I modified the size and contrast of the rpLMI to attempt to change the relative detectability of the phase and position cue, but these were unsuccessful in provoking a reversal of direction. This results are contradictory to those obtained from a different patient by Azzopardi & Hock (Azzopardi & Hock, 2011), who reported that on the rpLMI stimulus with similar size and contrast, GY reported motion in the phase cue direction.

Using reverse phi RDKs, I demonstrated that HS was using the phase cue to report motion direction, as reported direction reversed in the reverse phi RDKs, consistent with the direction of the phase cue and not the position cue carried by the displacement of the dots. This is conclusive evidence that HS can perform direction discriminations using the phase cue, which suggests that the motion reported in the position cue direction in the rpLMI stimulus was due to the relative detectability of the position cue exceeding the phase cue. Finally, on the Missing Fundamental Grating stimulus, HS performed at chance for low temporal frequencies, but at higher temporal frequencies reported motion in either the phase or position cue direction in the MFG, dependent on the contrast of the grating. This response pattern with contrast was also seen in his sighted field, demonstrating that the residual vision in the scotoma showed a similar bias towards phase cues at low contrast and position cues at high contrast.

Interpretation of Results

There are two explanations for these results. First, that HS has a higher residual visual abilities than previously assessed patients, which would explain performance on the RDK and gratings. If so, then the use of both phase and position cues could be explained by the higher visual capabilities, and suggests that the midbrain pathway can process both cues, providing the cues are sufficiently visible to patients. A second explanation is that HS has qualitatively different vision in his scotoma than in the sighted field, and that this is driving the increased performance and the use of position cues in the scotoma. This qualitative difference could be due to loss of tissue in V1 that acts to decrease sensitivity in the scotoma, or could be driven by the activity of additional pathways, separate to the

cortical and midbrain pathways. These alternative explanations provide different interpretations of my results, and for interpreting the motion perception properties of the midbrain pathway. They also highlight the necessity of functional scans in patients, to verify the absence of residual tissue in V1, and the use of conflicting cue stimuli across a range of parameters to determine whether patients are unable to perceive the phase or position cue, or whether the cue is below the detection threshold.

Shape discrimination abilities in HS

HS reported motion in the position cue direction on the rpLMI, contrary to the results of patient GY (Azzopardi & Hock, 2011), but did report motion in the phase cue direction on the reverse phi RDK and low contrast MFG when the stimulus features were less detectable. GY's results could reflect a lack of access to the position cue, or could be alternatively explained in terms of contrast sensitivity. Patient HS was tested less peripherally than GY, where stimuli were presented at 11.57^{o} vertically and horizontally in the scotoma for patient HS, but at 12 degrees vertically and 30 degrees horizontally in patient GY, which may provide some preliminary explanation for this difference in contrast sensitivity.

In normal observers reported motion direction is influenced by stimulus contrast such that for conflicting cue stimuli motion is reported in the phase cue direction at low contrasts and in the position cue direction at high contrasts (Cavanagh, 1990; Georgeson & Harris, 1990). Contrast sensitivity is reduced by two log units (Barbur *et al.*, 1994) in the scotoma for V1 lesioned patients (Sahraie *et al.*, 2008), possibly due to the huge reduction in available retinal ganglion cells (10% projecting to midbrain, Perry & Cowey, 1984). GY had near chance performance on a sinusoid grating of 50% contrast (Barbur *et al.*, 1994), whilst HS was able to discriminate motion in MFG grating at 12%, and detect static bars at 2% contrast, and static dots at 10% contrast. Furthermore, my rpLMI stimulus was presented at a greater contrast than in Azzopardi & Hock (Azzopardi & Hock, 2011), where the white and black regions of the rpLMI were presented at 44.5 and 0.23 cd/m^2, whilst mine were presented at 81.3 and 0.1 cd/m^2; thus the increased physical contrast in my stimulus may have increased the availably of the position cue. I did test HS with lower contrast version of the rpLMI (12% contrast, at 54.2 and 34.8 cd/m^2) and performance was at chance. If contrast sensitivity can explain the difference in reported direction for GY and HS, then I would expect a motion to be reported in the phase cue direction when the rpLMI is displayed at an intermediate contrast between 12 and 100% for HS. Thus, the difference in reported motion direction on the rpLMI stimulus between HS and GY could be explained by a difference in contrast sensitivity of the scotoma, where the reduced contrast sensitivity in GY's scotoma biases reported motion in the phase cue direction. This explanation would conclude that the midbrain pathway is

able to process position cues, if presented at sufficiently high contrast, placing HS and GY's visual abilities on a continuum. If higher contrast sensitivity is driving HS's ability to perceive position cues, then the cause of the increased sensitivity needs to be determined. The higher contrast sensitivity could be the result of individual differences, but still reliant on the midbrain pathway, or could reflect that residual vision in HS's scotoma is mediated by tissue other than the midbrain pathway.

Spared V1 tissue

HS has a greater ability to discriminate motion than other patients in the literature, exemplified by his ability to discriminate the direction of RDKs and gratings, which other patients cannot (Azzopardi & Cowey, 2001), without substantial training on regions close to the scotoma perimeter (Huxlin *et al.*, 2009). One possibility is that this greater visual ability in the scotoma is the result of 'islands of spared [V1] tissue' (Fendrich *et al.*, 1992), which could include magnocellular, but also parvocellular and koniocellular cells. Studies using patients with structurally defined lesions of V1, with standard impaired motion abilities, have sometimes found patients with residual V1 tissue which was only revealed following functional imaging (Fendrich *et al.*, 1992; Barbur *et al.*, 1993; Morland *et al.*, 2004). These islands of residual tissue did not dramatically change patients' performances, with residual vision appearing very similar across all the patients described in Barbur *et al.* (Barbur *et al.*, 1993) and Morland *et al.* (Morland *et al.*, 2004). It is possible that HS has some small islands of residual tissue in V1, and that these account for the use of position cues by HS.

Additional pathways

There is growing evidence for direct projections from LGN to V4 and to V5/MT, which are not disrupted by V1 lesions, and may be involved in producing the residual vision in the scotoma. The presence of these additional pathways negates my experimental purpose; assessing motion perception mediated purely by the midbrain pathway.

The direct projection from LGN-V5/MT has been shown structurally in some patients with V1 lesions (Ajina *et al.*, 2015; Pedersini *et al.*, 2020) and in macaques, where retrograde tracer injected into V5/MT has shown a direct projection from LGN to V5/MT, (Sincich *et al.*, 2004). This presence of this pathway in human patients correlates with patients' overall quality of residual vision (Ajina & Bridge, 2018; Ajina & Bridge, 2019). A lesion study in the macaque has suggested a causal, not just correlational, relationship between activity in the LGN and residual motion perception. Schmid *et al.* reported residual activity in V5/MT following a lesion of V1, which was abolished along with residual motion abilities following an additional lesion of LGN (Schmid *et al.*, 2010). The stimulus used, a 2^{o} rotating chequerboard pattern, may have been below the detection threshold for the midbrain

pathway, hence the loss of V5/MT activity and reported detection despite the presumably intact midbrain pathway.

A structural connection between LGN and V4 has also been demonstrated in macaques using retrograde tracers (Cowey & Stoerig, 1989; Lysakowski *et al.*, 1988). These studies also reported that the neurons projecting from LGN to V1 and V4 came from separate populations, and thus the LGN projection to V4 would remain following a striate cortex lesion. This is supported by the presence of activity in V4 following V1 lesions (macaques; Schmid *et al.*, 2013; humans; Goebel *et al.*, 2001; Schoenfeld, *et al.*, 2002; Raposo *et al.*, 2011) , which suggest that the LGN-V4 projection may also have a functional role following lesions of V1.

The cells in LGN following a lesion of V1 and forming these extrastriate projections are predominantly koniocellular (Sincich *et al.*, 2004; Yu *et al.*, 2018). There is some psychophysical evidence that parvocellular neurons may remain partially active following V1 lesions, where the detection of chromatic targets in V1 lesioned patients has shown similar spectral sensitivity in the sighted and blind hemifields (Stoerig & Cowey, 1989; Brent *et al.*, 1994). However, chromatic discrimination between isoluminant chromatic targets was not seen in human patients or lesioned monkeys (Stoerig & Cowey, 2001; Alexander & Cowey, 2013). Residual functional parvocellular neurons could be due to residual V1 tissue, or suggest that some portion of the cortical projections from LGN may be carried by parvocellular neurons. Following lesions of V1, there is transneuronal retrograde degeneration (TRD) of many retinal ganglion cells, as they are involved in projections to V1 that are no longer active, with TRD occurring in approximately 85% pβ RGCs (Cowey *et al.*, 1989). It is possible that some of these pβ RGCs do not degenerate, because they form still-active projections, and these account for the activity recorded in V4 following V1 lesions (macaques; Schmid *et al.*, 2013: humans; Goebel *et al.*, 2001).

This body of evidence suggests we cannot conclude that the midbrain pathway to V5/MT is the sole pathway remaining following lesions to V1. However, cells in the koniocellular and parvocellular layers of the dLGN are unlikely to be responsible for the residual motion perception that we observe in the scotoma. The temporal frequency response function of the parvocellular layers in dLGN is biased towards low temporal frequencies (Levitt *et al.*, 2001), and lesions to the parvocellular layers of dLGN produce a 3–4-fold reduction in spatial acuity (Merigan &Katz, 1990), and huge reduction in chromatic sensitivity (Schiller *et al.*, 1990b). Residual vision in the scotoma mediated by parvocellular layers of the dLGN should therefore operate at relatively high spatial frequencies, low temporal frequencies, and with chromatic discrimination; the opposite to the residual vision that we see in macaques and human patients with lesions to striate cortex. In contrast lesions to the

magnocellular layers of the dLGN, that leave the parvocellular and koniocellular layers intact, produce a huge reduction in sensitivity to motion stimuli that more closely mirrors the reported visual sensitivities of patients with striate cortex lesions (Schiller *et al.*, 1990b). It was found that sensitivity to moving stimuli in macaques with lesioned magnocellular layers of the dLGN was most impaired at higher speeds (above 1 deg/s) (Merigan, Byrne, *et al.*, 1991a; Schiller *et al.*, 1990b). This is very different to the bias towards high temporal frequencies that we see in V1-leisoned patients (Sahraie *et al.*, 2008), which mirrors the high temporal frequency bias of cells in the superior colliculus (Malpelli *et al.*, 1977) and the temporal frequency functions of the pα and smooth bistratified RGCs that project to the midbrain (Crooks *et al.*, 2008a,b). This suggests any contribution of a projection from LGN to V4 or V5/MT would likely only contribute to the perception of static stimuli.

Conclusion

Patient HS, unlike patient GY, was able to discriminate motion in his scotoma using both phase and position cues. This was a failure to replicate the results of Azzopardi & Hock (2011), but poses some intriguing questions as to why HS was able to discriminate position cues so readily. One key feature of this difference between patients may be caused by HS's greater ability to detect and localise shapes in comparison to GY. Across the conflicting cue stimuli used, it was evident that HS's use of position cues in discriminating motion direction correlated with the detection of features in the visual stimulus, and that, when features could not be discriminated, HS reported motion in the phase cue direction. This was seen in the rpLMI, where HS reported motion in the position cue direction, but was also able to discriminate the change in shape on the shape change control in the rpLMI (Experiment 3C). In the reverse phi RDKs, where the size of the dots reduces their detectability, and the correspondence problem reduce the salience of the dots, HS reported motion in the phase cue direction. Finally, on the MFG (Experiment 4A), HS reported motion in the position cue direction only at high contrasts, where the features are more detectable, and in the phase cue direction at low contrast presentations of the MFG, when the features are less detectable (Georgeson & Harris, 1990). GY, who reported motion in the phase cue direction in the rpLMI, was unable to discriminate shape change and position change of large rectangular stimuli in the scotoma (Azzopardi & Hock, 2011). One interpretation of these data is that the primary difference between residual motion abilities of GY and HS is driven by the difference in ability to detect and discriminate features in the scotoma.

Differences in the residual motion abilities between HS and GY could also be attributed to differences in the active visual tissue between the two participants. GY acquired his lesion as a 7 year old, and HS as a 70-year old. In macaques, there are substantial differences in the visual abilities

between macaques lesioned in infancy and those lesioned as adults, where both macaques lesioned in infancy and adults were able discriminate the motion direction of a moving bar, but only macaques lesioned in infancy could discriminate the direction of a coherently moving field of RDKs (Gross *et al.*, 2004). This is in direct contradiction to our results with patient HS and those of GY, whereby HS was able to discriminate the direction in an RDK, and patient GY was not (Azzopardi & Cowey, 2001). Fox *et al.* (2020) suggests that differences in motion discrimination between macaques lesioned as infants and adults could be explained by macaques lesioned as infants having a more substantial midbrain pathway (pulvinar to V5/MT projection specifically), and macaques lesioned as adults having a weaker or non-functional midbrain pathway, but a functional projection from LGN to V5/MT. This is supported by diffusion training and MRI studies showing decreased neural pruning of the pulvinar to V5/MT pathway in macaques lesioned as infants compared to macaques lesioned as adults (Warner *et al.*, 2015), whereas the visual function of residual LGN neurons, which would potentially form the basis of an LGN-V5 projection, were similar across macaques lesioned as infants and adults (Yu *et al.*, 2018), and has been structurally demonstrated in human patients with V1-lesions acquired in adulthood (Ajina *et al.*, 2015). These results, if applicable to the human brain, would suggest that patient GY would have residual visual functions predominantly served by a midbrain pathway, whilst HS's residual vision may also be served by an LGN-V1 projection. A final difference is in the number of retinal ganglion cells preserved following the lesion of V1. Retinal ganglion cell loss through transneuronal retrograde degeneration (TRD) varies when monkeys are lesioned as infants or adults, with about 20% RGC TRD loss when lesioned as adults, and 70-80% RGC TRD when lesioned as infants (New world monkeys; Hendrickson *et al.*, 2015; Old world monkeys; Dineen & Hendrickson, 1981). This increased number of preserved RGCs, suggesting retained functional projections at later stages of the visual system, could account for the increased contrast sensitivity seen in HS compared to GY, and in patient HS's ability to discriminate shape. This hypothesis – a difference between the number of RGC in HS and GY's retinae - could be measured using combined Adaptive Optics with Optical Coherence Tomography (Liu *et al.*, 2017). Equally, this difference in discrimination abilities between patient HS and GY could be the result of residual V1 tissue in HS. This explanation cannot be discounted without functional imaging confirming an absence of activity in the lesioned striate cortex, although this still would not account for islands of residual tissue (Fendrich *et al.*, 1992) below the spatial resolution of the imaging system used.

This difference between HS and GY could also be the result of the extended parameters used in my experiments. It is possible that there is no qualitative difference in the residual vision of HS and GY, and that GY might have responded in the position cue direction if the stimuli had been presented at

sufficiently low spatial frequency, high temporal frequency and high contrast. These differences in potential physiology between GY and HS and in the experimental methods of the two studies, require replication of these experiments in more patients, and with a wide range of stimulus parameters. Ideally, these experiments would be accompanied with imaging, particularly the use of fMRI optimised to image both LGN and SC (Wall *et al.*, 2009). This would concurrently test whether there is sufficient activity in the dLGN-V5/MT pathway to mediate visual performance, as suggested by Ajina & Bridge (Ajina & Bridge, 2018), and whether there is residual V1 activity. This comparison would also reveal whether there is preferential activation of the midbrain pathway to phase cues, as suggested in response to comparisons of phi and beta motion in normal observers (Pizlo *et al.*, 2008).

An alternative framework to investigate the physiological substrates processing phase and position cues is the different classes of retinal ganglion cells, which may form distinct pathways separately involved in processing phase and position motion cues.

Chapter 7: Correlations between psychophysics and physiology

7.1 Introduction

Motion can be described either by a change in phase of the Fourier components of an image, or the change in spatial position of image features. Human observers can discriminate motion using either cue, as demonstrated through stimuli with conflicting phase and position cues, where the parameters of the stimulus and retinal viewing location determine which cue direction motion is perceived in. The consistent pattern of results across different conflicting cue stimuli with varied image characteristics suggests that this could be the result of two different mechanisms being used to extract the phase and position cues. The psychophysical dissociations observed can be used not only to suggest that there are separate mechanisms, but also in an attempt to characterise the mechanisms in terms of function and physiological substrates. In this chapter we refer to these two putative separate mechanisms as the 'phase' and 'position' systems.

Motion discrimination using phase or position cues responds differently to temporal frequency: motion discrimination using position cues declines rapidly with increasing temporal frequency, whereas motion discrimination using the phase cues remains high across increasing temporal frequency (Lu & Sperling, 1995). A putative phase system may be able to process phase cues in the visual input without assigning them to features (Van Santen & Sperling, 1984), but a system using position cues may rely on the detection of features (Cavanagh, 1990; Foster *et al.*, 1989). One method by which motion discrimination could be achieved is by abstracting features from their image characteristics by assigning salience maps (Lu & Sperling, 1999, isoluminant/iso-salient gratings). The introduction of blank interframe-intervals to motion stimuli abolishes motion discrimination for stimuli where motion is produced by a change in phase, but do not affect motion discrimination for stimuli where motion is produced by a change in position (Chapter 2; Georgeson & Harris, 1990). This suggests a difference in mechanism between possible phase and position systems, where the phase system integrates luminance information across time, and the position system does not. These psychophysical dissociations support a hypothesis of separate phase and position systems.

One parameter where the dissociation between phase and position cues is most apparent is spatial frequency. This is visible in reverse phi stimuli, where motion is reported in the phase cue direction at high spatial frequencies, and in the position cue direction at lower spatial frequencies (Chapter 5). This is also the case in Gabor stimuli, where motion is reported in the carrier cue direction at high spatial frequencies, and in the aperture cue direction at lower spatial frequencies (Chapter 4). When the carrier and aperture cues are manipulated separately, spatial sensitivity was higher for the phase

displacements than the changes in aperture (Chapter 3). The gradient of displacement thresholds plotted by eccentricity is significantly different between carrier and aperture cues, with a significantly greater increase in displacement thresholds with eccentricity for the aperture cues than the carrier cues. Similar differences in spatial sensitivity, and gradients of displacement threshold across eccentricity were found for displacement thresholds for apparent motion (phase and position cues) compared to spatial separation (position cue) (Foster *et al.,* 1989), and for unreferenced motion (phase cue) compared to grating acuity (Levi *et al*., 1984). These results are suggestive of a difference in the spatial acuity for phase and position cues, perhaps caused by a difference in the physiological substrates for the two systems.

The spatial acuity of a system is limited by the sampling of the input. To faithfully represent a periodic signal, the signal must be sampled at more than twice the highest frequency component of the signal (Nyquist-Shannon sampling theorem). Therefore, for all systems, there will be a resolution limit, above which the system cannot differentiate between inputs of different higher frequencies. Two different systems, with different sampling rates, will therefore have different resolution limits. This idea could be applied to the visual system, where the input is the visual array projected onto the retina, and the sampling units, the early units in the visual system, such as the photoreceptors and retinal ganglion cells. Under this theoretical framework, the difference between the acuity in the phase and position systems could represent a difference in the initial sampling of the visual input. If the sampling of the visual input is causing the difference in resolution between the phase and position systems then this would strongly suggest that there is a difference in the physiological substrates for these two systems at the early stages of the visual system; perhaps the retinal ganglion cells.

Which cell classes might form the phase and position systems?

The majority of retinal ganglion cells (RGCSs) in the primate retina are made up of two classes; 70% are pβ (midget) RGCs, and 10% are pα (parasol) RGCs (Dacey, 2000; Yan *et al.*, 2020), with and a substantial proportion of the remaining 20% of RGCs classified as pγ RGCs (Perry & Cowey, 1984), and small bistratified RGCs (Dacey & Lee, 1994). All of the pβ RGCs, the majority of pα RGCs and the small bistratified RGCs project to the dorsal lateral geniculate nucleus (Perry & Cowey, 1984; Dacey & Lee, 1994). The pγ RGCs and a small projection of the pα RGCs to the superior colliculus (Perry *et al.*, 1984; Crook *et al.*, 2008).

The pα RGCs synapse in layers 1 and 2 of the LGN, referred to as the magnocellular layers and comprising ~10% of the LGN neurons, and the pβ RGCS synapse in layers 3 to 6 of the LGN, referred to as the parvocellular layers, and comprising of ~80% of the LGN neurons (Perry & Cowey, 1984). The small bistratified RGCs (Dacey & Lee, 1994) project to the koniocellular layers of the LGN, which are 'intralaminar' layers sandwiched between the layers of the LGN structure, (Roy *et al.*, 2009), ~8% of LGN neurons (Dacey, 2000). These projections remain separate in the cortex, where the parvocellular layers of the LGN project to layer 4Cβ of V1, and the magnocellular layers in layer 4Cα of V1 (Chatterjee and Callaway, 2003).

Most of the psychophysical investigations of these neurons have focused on the pα and pβ RGCs, and their respective layers in the LGN; the magnocellular and parvocellular layers (Leventhal *et al.*, 1981; Perry *et al.*, 1984). The distinct morphology of the two classes of RGC (Polyak, 1941), and the preserved separation across projections to LGN and V1 (Perry *et al.*, 1984; Chatterjee and Callaway, 2003), is suggestive of distinct psychophysical functions. In this Chapter, I discuss data from both the RGCs and the neurons in the LGN, with the understanding that these form two parallel pathways between the retina and LGN, where differences in function between the pα and pβ RGCs are likely to be preserved between the magnocellular and parvocellular LGN layers.

The spatial frequency sensitivity of the pβ and pα retinal ganglion cells are very similar (Crook *et al.*, 1988), and similarly, the spatial frequency sensitivity ranges of the parvocellular and magnocellular LGN cells have considerable overlap with a slight but significant higher spatial frequency sensitivity in the parvocellular neurons, (Kaplan & Shapley, 1982; Levitt *et al.*, 2001). Again, there is considerable overlap between the temporal frequency sensitivity of the pα and pβ retinal ganglion cells, with a similar peak around 10Hz, (Benardete & Kaplan, 1997; Kaplan & Benardete, 2001), although the sensitivity in the pα retinal ganglion cells increases to around 16Hz with increased contrasts (Kaplan & Benardete, 2001). A greater distinction in temporal frequency has been observed at the level of the LGN, with consistent findings that magnocellular cells are tuned to higher temporal frequencies than parvocellular cells (Derrington & Lennie, 1984; Levitt *et al.*, 2001; Movshon *et al.*, 2005), although there is considerable overlap at lower temporal frequencies. However, the spatial and temporal sensitivities should be considered linked properties. When the two parameters are assessed simultaneously, responses in the pβ RGCs decline rapidly with increased temporal and spatial frequencies (Benardete & Kaplan, 1997).

More striking differences between the spatiotemporal acuity of the parvocellular and magnocellular neurons in the LGN are observed from studies that selectively lesion either the parvocellular or magnocellular layers of the LGN, and observe the changes in visual abilities of macaques. Colour and

texture discrimination were impaired following parvocellular lesions (Merigan & Eskin, 1986; Schiller *et al.*, 1990), and flicker detection and motion discrimination were impaired following magnocellular lesions (Schiller *et al.*, 1990; Merigan *et al.*, 1991a,b). Critically, for considering potential associations with the phase and position pathways, the sensitivity to spatial and temporal frequencies interacted. Merigan & Eskin (Merigan & Eskin, 1986) reported that lesioning the parvocellular layers of the LGN left macaques with "normal sensitivity for stimuli of high temporal, low spatial frequency", and subsequently reported a 3-4 fold reduction in visual acuity across all static spatial frequencies tested (Merigan *et al.*, 1991a). Conversely, lesioning the magnocellular layers of the LGN had little effect on visual acuity, but did dramatically impair temporal acuity, measured by critical-flicker- fusion thresholds (Schiller *et al.*, 1990), and the discrimination of low contrast stimuli, which was impaired significantly more at higher temporal frequencies (Merigan *et al.*, 1991b, Schiller *et al.*, 1990). Taken together, these results suggest that the parvocellular layer of the LGN has high spatial acuity which declines with increasing temporal frequencies, whilst the spatial sensitivity of the magnocellular layer of the LGN is more broadband. This indicates that acuity for static stimuli is probably higher in the parvocellular layers, but when stimuli are presented at varying temporal frequencies, the acuity of the parvocellular layer drops rapidly, whilst acuity of the magnocellular layer does not. This could result in higher acuity for medium to high spatial frequencies in the magnocellular than in the parvocellular layer, if presented at high temporal frequencies. These differences in the response to spatial and temporal frequencies suggest that the parvocellular layers in the LGN form the basis for colour and shape processing, and the magnocellular layers for motion perception.

Can we correlate RGC density and psychophysical data?

In this chapter I propose that the difference in displacement thresholds measured for the phase and position cues reflects a difference in sampling resolution of the phase and position system. I hypothesise that this difference in sampling resolution could be driven by differences in the retinal ganglion cell inputs to the phase and position systems. In Chapter 3, I reported a decline in spatial acuity in the phase and position systems, with increasing eccentricity. This decline in acuity with eccentricity is mirrored in the decline in density of the retinal ganglion cells density across the retina (see Figure 7.1; Azzopardi & Cowey (1996a), Drasdo, 1989; Watson, 2014), and this density distribution is seen too in the input layers of the LGN (Schein & de Monasterio, 1987) and in V1 (Brindley & Lewin, 1968). Critically for my hypothesis, there is a difference in the density distribution of different classes of RGCs across the retina.

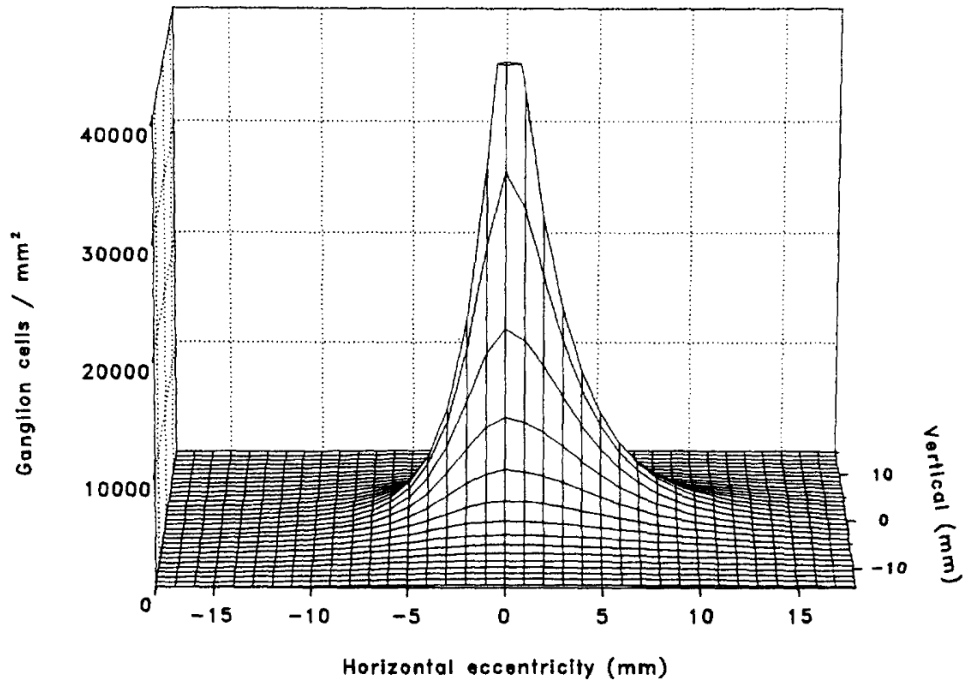


Figure 7.1. Figure from Azzopardi & Cowey (1996a), displaying the estimated distribution of total retinal ganglion cells across the macaque retina. The distribution peaks at the fovea and declines rapidly with increasing eccentricity.

The pα and pβ RGCs differ in the number of total cells for each class and their relative densities across the retina. Across the 0.7 to 1.5 million retinal ganglion cells in the human retina (Curcio & Allen, 1990), 80% are pβ RGCs and 10% are pα RGCs (Yan *et al.*, 2020). Consequently, within the pα RGCs there is a much sparser sampling of the visual field. This can be considered in terms of the Nyquist limit of spatial sampling, which states that to faithfully represent an image, the image must be sampled at more than twice the highest spatial frequency component. Whilst the spatial frequency sensitivity of individual pβ and pα RGCs is relatively similar (Crook *et al.*, 1988), the ability to represent spatial patterns is severely limited in the pα RGC class, relative to the pβ RGC class, as demonstrated by the 3-4 fold decrease in acuity for static patterns seen following parvocellular LGN lesions (Merigan *et al.*, 1991a).

The relative density of the pα and pβ retinal ganglion cell classes differs with eccentricity. In the human fovea, 90% of the total RGCs are pβ, and only 5% are pα, but there is a steeper decline in density for the pβ RGCs with eccentricity. In periphery, pβ RGCs constitute 45% of the total RGCs, whilst the pα RGCs now make up 20% of total RGCs (Dacey, 2007). These proportions vary slightly across different models of human data (Drasdo, 1989; Rovamo & Virsu, 1979), but all demonstrate a substantially steeper decline in density with eccentricity in the pβ retinal ganglion cells than in the pα retinal ganglion cells (see Figure 7.1). Watson (Watson, 2014) suggests that the “density of their [pβ RGCS] receptive fields imposes a fundamental limit on the spatial resolution of human vision”. If

the density of cells confers a limit on the spatial resolution of human vision, then this would suggest a different limit in the pα and pβ classes of retinal ganglion cells at each point on the retina, and a different gradient of this spatial resolution limit with eccentricity. If the pα and pβ retinal ganglion cells mediate different visual abilities, then this difference in spatial resolution should be apparent in these different visual functions as the limiting factor for the sampling resolution of the pathway.

Surprisingly few studies have taken this approach; correlating psychophysical data separately with the density distributions of the pα and pβ RGCs. Weymouth (Weymouth, 1958) compared acuity thresholds with the "ratio of the total number of cones divided by the total number of optic nerve fibres" as an estimate of retinal ganglion cells, and showed a clear correlation. This demonstrates early support for the basic relationship between retinal ganglion cell density and spatial sampling of the visual field. Drasdo & Thompson (1989) correlated visibility and colour recognition isopters with estimated densities for pα and pβ RGCs, where visibility detection thresholds showed reasonable correlation with the pβ RGC gradient, but recognition was steeper. This may reflect the difference between the tasks, where detection of coloured targets may have a similar sampling resolution to the low level visual system pathway, but the recognition of colour may require higher processes, such as attention and memory, which are not related to the sampling resolution of the visual system. My data, which is in the form or displacement thresholds for direction discrimination thresholds, should perform better for correlations with the cell densities than the recognition data of Drasdo & Thompson (Drasdo & Thompson, 1989).

I have demonstrated a consistent psychophysical dissociation in spatial frequency sensitivity between the phase and position cues. In chapter 3, I measured displacement thresholds for carrier and aperture cues across the visual field. These displacement thresholds, which represent direction discrimination at 75% correct, could be considered a metric of the cut-off point, or resolution limit, for spatial sensitivity in the phase and position systems. I have proposed that the phase system represents a true 'motion' system, whilst the motion abilities of the position system are a by-product of a system that processes identity and position of salient features in the visual input. These mirror the suggested functions of the parvocellular and magnocellular layers in the LGN (Merigan *et al.* 1991a,b,; Schiller *et al.*, 1990). In this chapter, I plot models of the distribution of cell density across the visual field, for the pα and pβ classes of retinal ganglion cells and parvocellular and magnocellular neurons in the LGN, and correlate these with the displacement thresholds for the carrier and aperture cues reported in chapter 3.

Classes of RGC excluded from these models

In this chapter, I consider just the cell density distributions of the pα and pβ RGCs. This is due to the availability of the data for the density distributions across the retina, and due to the well-characterised functional properties of the pα and pβ RGCs, and their respective layers in the LGN. These functional properties share striking similarities with the spatiotemporal properties of the phase and position systems, suggesting that these two classes of RGC may be separately involved in the physiological substrates of the two systems. This approach assumes that the number of RGCs is directly related to sampling and therefore to resolution; the other classes of RGC in the retina do not have sufficient numbers of RGCs to produce sampling resolutions comparable to the pα and pβ resolutions. Finally, this approach is predicated on the idea that early visual system sampling produces the limits of acuity. These two RGC classes project separately through the layers of the LGN and V1, and can be considered theoretically separate systems; including other classes, particularly the small bistratified RGCs or the pγ RGCs would complicate the models, as an explanation would need to be developed about how the resolution of the two systems may be combined. This is particularly true for the pγ RGCs, which have been implicated in the cortical blindness literature, as they form the majority of the retinal projection to the superior colliculus, and therefore may contribute to the residual motion processing observed following lesions of V1.

This chapter presents a novel application of this approach to motion thresholds, where it has previously been used solely for considering acuity for static stimuli. The models presented within this chapter are not presented as an exclusive explanation for the physiological substrates of the phase and position systems, which may receive inputs from other classes of RGC, particularly the small bistratified RGCs to the position system, and potentially the pγ RGCs to the phase system.

7.2 Method

Models of cell density

In this chapter I compared the densities of pβ and pα RGCs, and densities of the layers in the LGN, with the normalised displacement gradients. I predicted that if there are correlations between cell density and my displacement thresholds then these will be strongest between the phase cue thresholds and the pα RGC densities, and the position cue thresholds with the parvocellular LGN data.

The pα RGC density estimates include both the pα RGCS which project to the magnocellular layers of the LGN (Perry *et al.*, 1984), and the small proportion of pα RGCs which project to the superior colliculus (Perry & Cowey, 1984; Crook *et al.*, 2008; Leventhal *et al.*, 1981). In chapter 6, I discussed

the residual visual capabilities of patients with V1 lesions, which were particularly well-preserved in the motion domain (Riddoch, 1917). Motion perception abilities are retained despite the loss of shape discrimination (Azzopardi & Hock, 2011), and direction is discriminated using the phase and not position cues (Azzopardi & Hock, 2011). This residual motion could be mediated by the projection of the pα RGCs to the superior colliculus, as the directionally-responsive cells in the superior colliculus are achromatic and responsive to high temporal frequencies (Malpeli *et al.*, 1977). The superior colliculus is responsive to motion in visual observers without lesions of V1, (Schneider & Kastner, 2005; DuBois & Cohen, 1999). Mikellidou *et al.* (Mikellidou *et al.*, 2018) reported similar fMRI activity in healthy observers in the superior colliculus and the LGN to stimuli displayed at different spatial and temporal frequencies. This suggests that a similar neural substrate is processing the visual stimuli at both locations; this activity could be reflecting the pα projection to both the LGN and superior colliculus.

The relative densities with eccentricity in the ganglion cell estimates show a steeper decline of density with eccentricity in the LGN data, than the RGC data. This reflects the increased overrepresentation of the fovea and central vision which occurs between the retina and the LGN (Azzopardi & Cowey, 1996b; Silveira *et al.*, 1989; Perry & Cowey, 1985). This over-representation is caused by the allocation of additional neural tissue to processing the responses from the retinal ganglion cells at fovea, to amplify the acuity at the fovea. Azzopardi *et al.*, (Azzopardi *et al.*, 1999) demonstrated in macaques that this overrepresentation between the LGN and V1 is uneven between the parvocellular and magnocellular layers, and suggested that the overrepresentation is occurring in the parvocellular layer alone. This could also be the case between RGC and LGN, which would mean that the parvocellular layers of the LGN may exhibit a greater foveal overrepresentation than the magnocellular layers of the LGN. This suggests that, for the visual functions processed by a pathway formed from pβ RGC inputs, the density distribution of the parvocellular layers of the LGN will better represent the neural resolution of the pathway than does the density distribution of the pβ RGCs.

Obtaining data

There are several published models of both human and non-human primate density distributions for retinal ganglion cells and the neurons in the parvocellular and magnocellular layers of the LGN. Humans and non-human primates have similar visual systems, with lower order visual areas (V1, V2, V3) showing similar locations and functional responses (Orban *et al.*, 2004), although there is less similarity in higher order visual areas (Van Essen *et al.*, 2001). For my hypothesis, the similarity of the retina is of the upmost importance, as this produces the initial sampling resolution of the visual system. However, there are significant differences between human and macaque retinae. Dacey &

Petersen (Dacey & Petersen, 1992) report that whilst pβ cells are similar between macaques and humans, the dendritic field of the pα RGCs is 20% larger in the humans than in the macaques at fovea, and 95% larger than macaque pα RGCs at 5 degrees eccentricity. The size of receptive fields is inversely related to spatial frequency acuity (Rovamo *et al.*, 1985), so this difference may disrupt the relative spatial scale between human and macaques, or change the retinal coverage factor within the pα cells, which would also affect the sampling resolution (Perry *et al.*, 1984). For these reasons, I chose to use predominantly human data in my models.

There are numerous difficulties in producing estimates of retinal ganglion cells density and the density of neurons in the parvocellular and magnocellular layers in the LGN. Retinal ganglion cells are displaced from the fovea (Drasdo *et al.*, 2007), and retinae are not uniform, with significant differences between densities in nasal and temporal, superior and inferior portions of the retina (Curcio & Allen, 1990). Similarly, the LGN is a three dimensional structure, with parvocellular cells across 4 layers, and magnocellular across 2 (Connolly & Van Essen, 1984). Calculating cell counts in LGN has been achieved by some studies by separating sections of the LGN into cylinders and then flattening (Connolly & Van Essen, 1984), or modelling the 3D structure and counting with 3D representations, to avoid distortions, (Malpeli *et al.*, 1996). Different estimates employ different approaches, and use different tissue samples for the retina and LGN (Le Gros Clark , 1941; Oppel, 1967; Curcio & Allen, 1990; Malpeli, 1999). I included two models for the retina and LGN estimates, to balance these different experimental choices.

Constructing cell density models

Retinal ganglion cell data

For my estimates of pα and pβ RGCs I used data from Drasdo (1989) and Watson (2014) presented in Figure 7.2 (left panel).

In 1977, Drasdo produced a model using data from 12 human retinae, taken across three studies (Vilter, 1954; Van Buren, 1963; Oppel, 1967), and all projected onto a schematic eye (Drasdo & Fowler, 1974) before aggregating across the datasets (Drasdo, 1977). Foveal displacement of RGCs was handled by interpolating between the linear section of the curve representing RGC density, and the estimated number of RGCs at fovea (~ 33,058), given the number of estimated foveal cones (Polyak & Klüver, 1957). This produces a linear estimate of RGC density, to estimate the 'true location' of RGCs, with the interpolation set so that the cumulative total of the RGCs was equal to ~94% of central cone distribution (Polyak & Klüver, 1957). In 1989, Drasdo extended this model to account for the relative density of the pα and pβ RGCs (Drasdo, 1989), by dividing the RGC distribution into the relative density distributions of the pα, pβ, pγ and pε reported in non-human primate retinae (Perry & Cowey, 1984; Schein & de Monasterio, 1987; Schein, 1988). Therefore,

Drasdo's (1989) models of the pβ and pα RGC distributions are based on human RGC data, following the density distributions of non-human primate RGCs (Drasdo, 1989).

Watson (Watson, 2014) produced an updated estimate for total and pβ human RGC distribution. Watson used human RGC data from Curcio & Allen (Curcio & Allen, 1990, data available online), who published the RGC distributions from 6 young adult retinae. Watson corrected RGC distributions for foveal displacement using a corrected version of equation presented in Drasdo *et al.* (Drasdo *et al.*, 2007). To calculate the relative proportion of pβ RGCs, Watson applied a modified version of Drasdo *et al.*, (2007), who had calculated the pβ proportion from both psychophysical measures and anatomical measures, where Watson also added a correction for the asymmetry between ON and OFF mRGCs (Dacey & Peterson, 1992; Dacey, 1993). I calculated a pα estimate from Watson (2014)'s data, by taking the ratio between pβ and pα RGCs presented by Drasdo *et al.* (Drasdo *et al.*, 1989), and then applying this to Watson's (2014) estimate of pβ RGC density.

Lateral geniculate nucleus estimates

For my estimates of the density distributions of the parvocellular and magnocellular LGN neurons, I used the models of Schein & de Monasterio (Schein & de Monasterio, 1987), and an adjusted model, where I applied the P:M LGN ratio from Azzopardi & Cowey (Azzopardi & Cowey, 1996a) to the sum of parvocellular and magnocellular LGN neurons reported by Malpeli *et al.* (Malpeli *et al.*, 1999). These models are presented in Figure 7.2 (right panel).

Schein & de Monasterio produced their model of LGN cell distributions from the transformed raw data of Le Gros Clark (Le Gros Clark, 1941). This data was subsequently mapped by Malpeli and Baker (Malpeli and Baker, 1975), and then divided into the relative parvocellular and magnocellular cell distributions by Connolly & van Essen (Connolly & van Essen, 1984). Schein & de Monasterio (1987) calculated the estimate of total cells in LGN from an integration of these equations and found it was lower than the total suggested in Connolly & van Essen (1984) by a factor of 1.31, which they added to equations as a correction. I have used the corrected Schein & de Monasterio (1987) estimates for the distributions of parvocellular and magnocellular LGN neurons (Schein & de Monasterio, 1987).

Malpeli *et al.* (1996) used the data from a single macaque LGN, published within a larger data set (Malpeli & Baker, 1975). Malpeli *et al.* (1996) reported that the parvocellular neuron density peaked at fovea and the decline in density could be fitted with power function, similar to the other models of cell distribution the estimates for magnocellular density actually peaked at 1.8 degrees eccentricity, so the power function for magnocellular required an additional term. This produces a distribution of magnocellular neurons very different from the Schein & de Monasterio (1987)

estimates, and a distribution that suggests an uneven pattern of foveal expansion in M cells. This difference which appears only in the M cells is unlikely to be due to poor localisation of fovea, or reflecting the foveal displacement of RGC cells, as it is occurring solely in the magnocellular distribution. Instead of the Malpeli *et al.* (1996) parvocellular and magnocellular estimates, I have summed the Malpeli data to produce an estimate of the total LGNs, and applied the ratio of parvocellular and magnocellular LGN cells proposed by Azzopardi *et al.* (Azzopardi *et al.*, 1999). This produces an additional LGN estimate, which is similar in form to Schein & de Monasterio (1987), and consistent with the increase in foveal overrepresentation between the retina and the LGN (Azzopardi & Cowey, 1996b), but not a dramatic reorganisation.

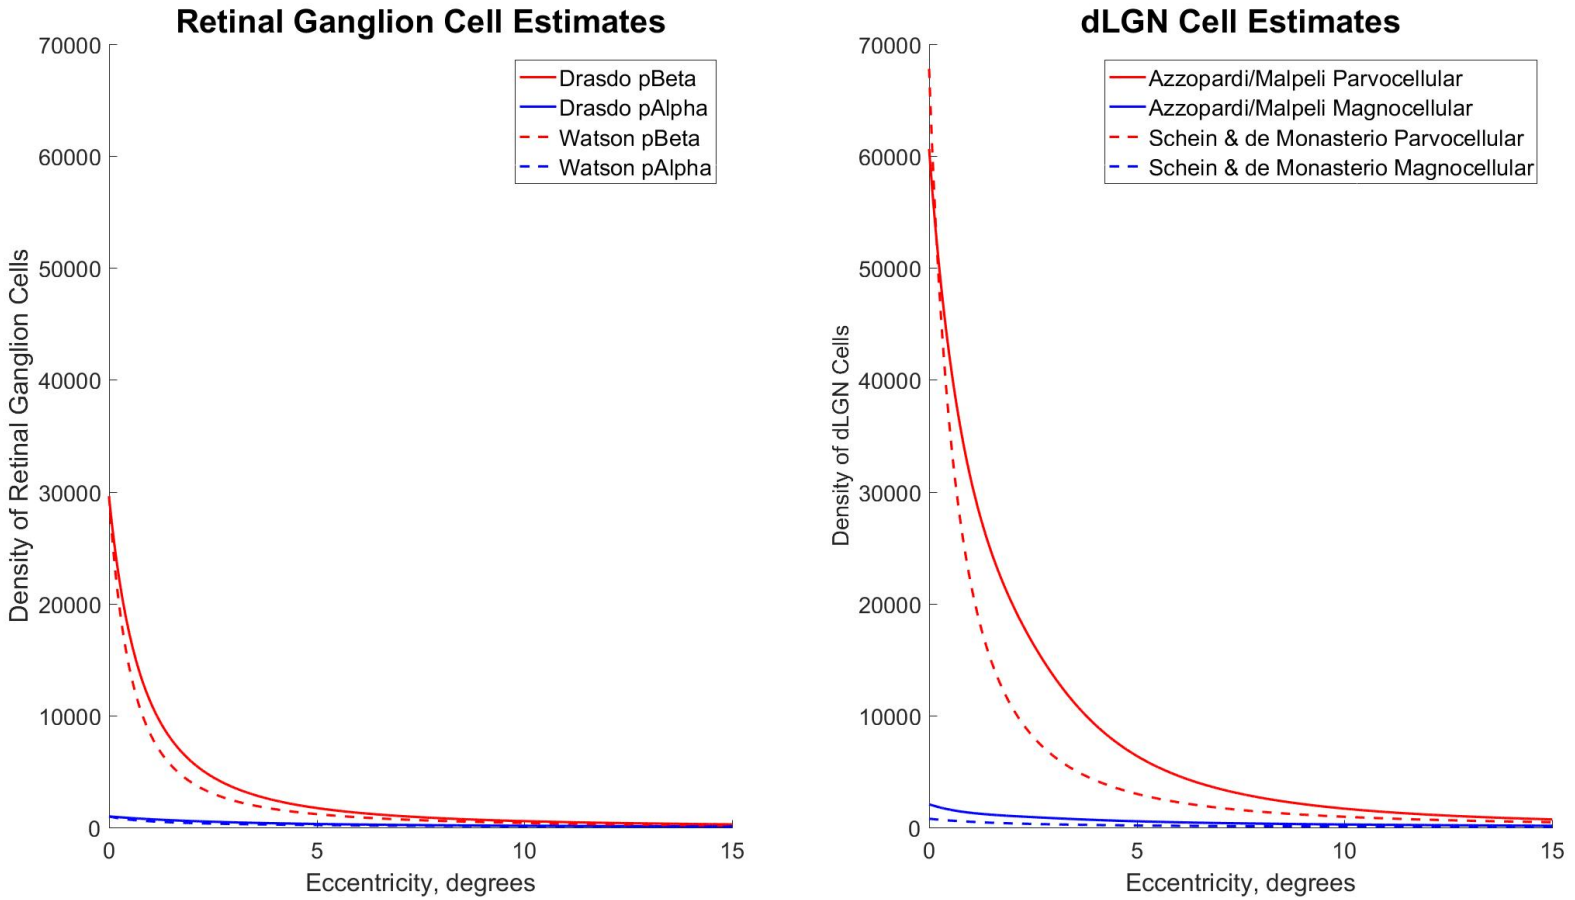


Figure 7.2. Figure displaying the estimated cell density distributions from the models of the pα and pβ retinal ganglion cell data (left), from Drasdo *et al.* (1989), Watson (2014), and the models of magnocellular and parvocellular dorsal lateral geniculate nucleus data (right), with one model produced using data from Azzopardi & Cowey, (1996a) and Malpeli *et al.* (1999), and the second model using data from Schein & de Monasterio, (1987).

Data transformations

The spatial sampling approach that I have taken here assumes that the relative density of a class of cells forming the input to a pathway can estimate the sampling resolution of that pathway. By describing the change in density of that cell type across the retina, I can predict how the spatial resolution of this pathway will change with eccentricity (Watson, 2014). I use this approach to compare two different classes of cell, the pα and pβ RGCs, and with the assumption that the relationship between cell density and sampling resolution is matched between these two classes.

This assumption is somewhat flawed, due to the significant differences in morphology between the two classes of RGC (Polyak, 1941; Perry *et al.*, 1984), which may affect sampling. Critical amongst these differences is the form of the dendritic trees of the two cell classes. The dendrites of a neuron are the branched extensions of the cell, which receive the synaptic inputs that drive the cell's activity. The differences in dendritic architecture between classes of RGCs likely reflects differences in information processing between the cells (Koch *et al.*, 1982). The dendritic trees of pβ RGCs have been characterised as "small bushy dendritic trees", and the pα RGCs as "large radiate dendritic trees" (Silveira & Perry, 1981). These differences in dendritic structure could reflect the greater spacing between pα RGC cell bodies, with the increased size of the dendritic trees increasing the surface area of the retina sampled by each pα RGC.

This potential greater overlap of dendrites in the pα than the pβ RGCs is reflected in the 'retinal coverage factor', which Perry & Cowey (1984) measured at between 1.9 and 2.3 for the pβ RGCs, and between 2 and 7 for the pα RGCs. The retinal coverage factor of approximately 2 in the pβ RGCs is consistent with my current models of cell density, where under Nyquist sampling theorem it was assumed that cells were sampling the retina by a factor of 2. The proposed retinal coverage factor between 2 and 7 for the pα RGCs (Perry & Cowey, 1984), suggests that the effective sampling resolution of the magnocellular pathway is greater than the estimated sampling resolution based solely on cell density estimates. This suggests that the sampling resolution of the magnocellular pathway is underestimated by my current models which use cell density, and that the actual sampling resolution of the pα RGCs is some factor greater than my current models, between 1 and 3.5 (Perry & Cowey, 1984).

Applying a retinal coverage factor to the pα density distribution produces a scaled density distribution of the pα RGCs. Depending on the retinal coverage factor chosen, this scaled pα density distribution may exceed the density of the pβ RGC distribution. The eccentricity at which this occurs can be referred to as the 'retinal crossover location'. This is modelled in Figure 7.3, where I assess the retinal crossover location for different scaling factors, applied to the mean pα density

distribution, and compared to the mean pβ distributions from my calculated cell density models. As scaling factors increase, they produce retinal crossover locations closer to 0, with scaling factors greater than 29.45 producing scaled pα distributions that always exceed the pβ distribution (as shown in Figure 7.4). Scaling factors below 2 result in no retinal crossover; the pβ density distribution will always exceed the scaled pα distribution. In this chapter I applied a scaling factor of 3.5, which is the mid-point between not applying a scaling factor, and the upper limit of the range of 2-7 suggested by Perry & Cowey (1984).

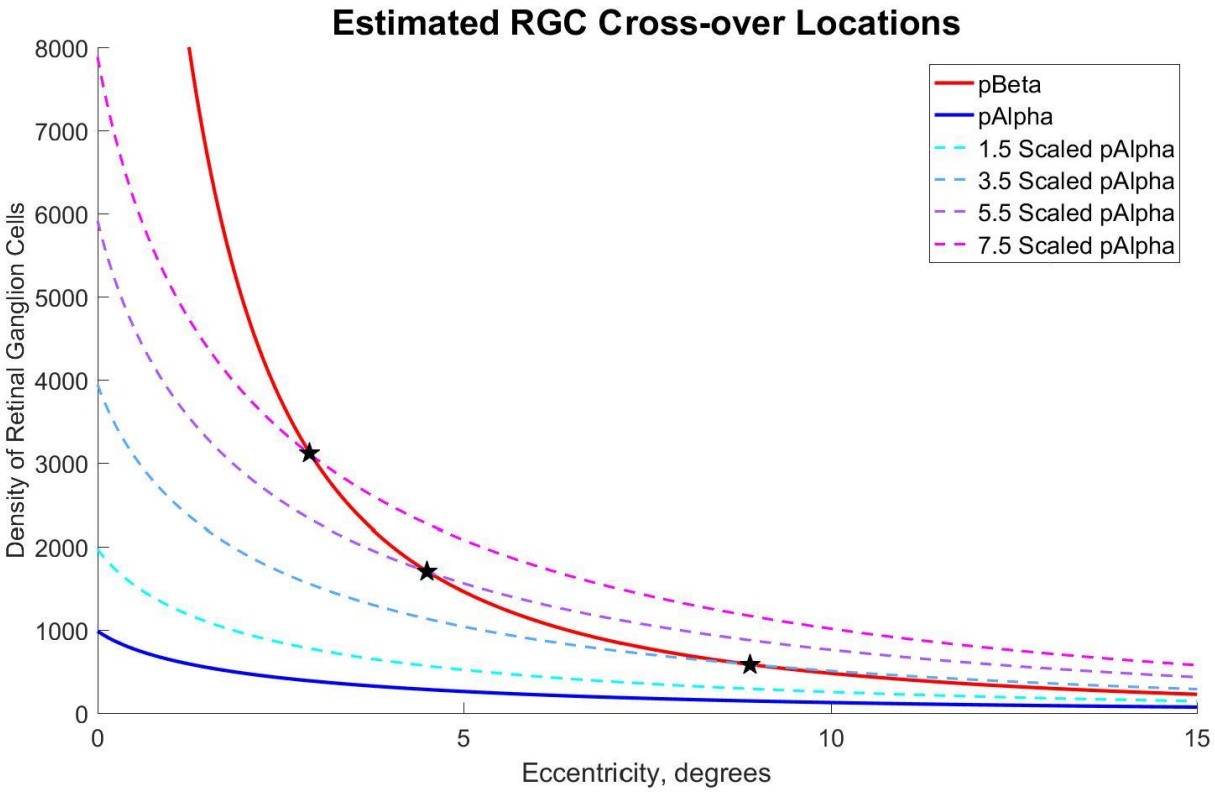


Figure 7.3. The cross-over points for pα RGC models with different scaling factors applied, where the pα and pβ shown here are calculated from the average of the Drasdo *et al.* (1989) and Watson (2014) models. The four scaled pα RGC estimates are multiplied by a scaling factor of 1.5, 3.5, 5.5 or 7.5. Black stars mark the locations where the scaled pα RGC estimates have equal density to the pβ RGC model, where 1.5 is equivalent at 30 degrees eccentricity (not shown within these bounds), 3.5 at 8.9 degrees eccentricity, 5.5 at 4.5 degrees eccentricity, and 7.5 at 2.9 degrees eccentricity.

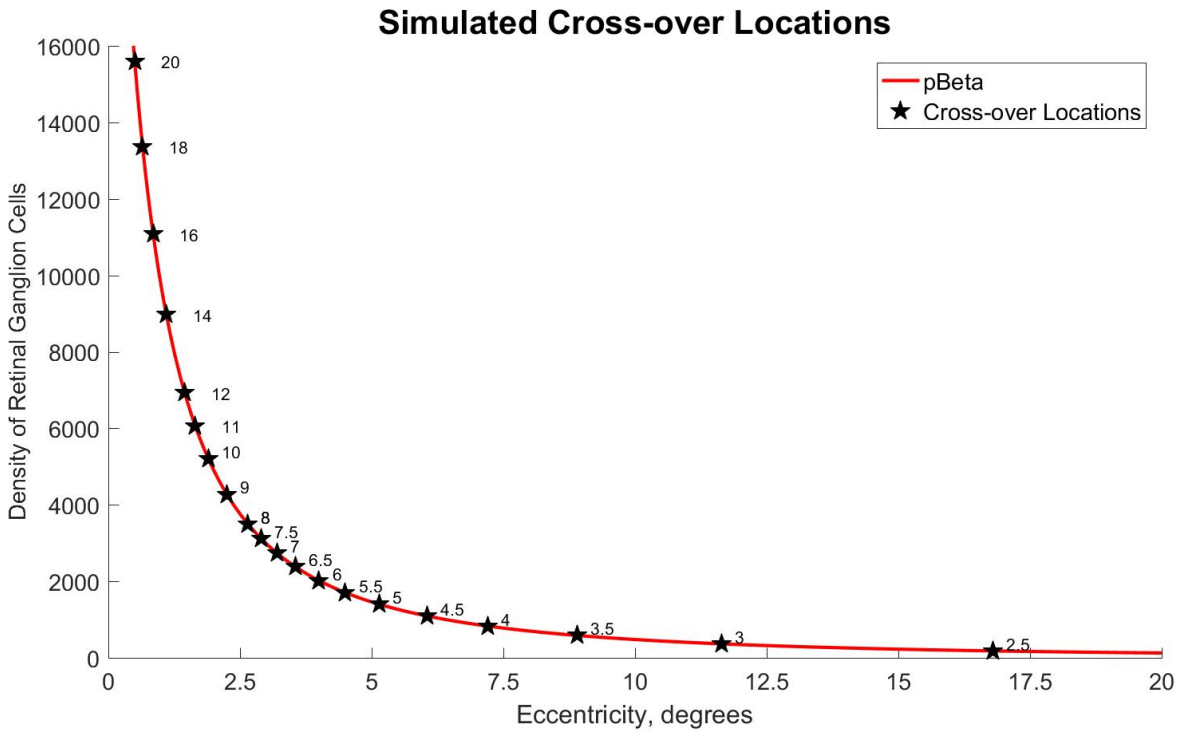


Figure 7.4. Simulated data depicting the estimated crossover locations for different retinal coverage factors, where the labelled filled stars are the retinal coverage scaling factor, plotted overlaid on the pα retinal ganglion cell density distribution, from the averaged Drasdo *et al.* (1989) and Watson (2014) models. Scaling factors below 2 result in no retinal crossover, scaling factors above 12 are not supported by psychophysical data, as the cross-over location would be in the parafoveal region.

Differences in the spatial frequency sensitivity of phase and position systems, could be explained by a difference in the spatial arrangement of the spatio-temporal detectors that form the input to the motion processing system. Two systems with different arrangements of detectors would have very different spatial sensitivity, reflecting the different purposes of the systems: a system developed to have high precision for spatial location may have less or no overlap between detectors, whereas a system detecting changes in phase may have more overlap between detectors, and also be oriented in space-time. The differences in the retinal coverage factor for pα and pβ retinal ganglion cells described by Perry & Cowey (1984) could be an implementation of this. Given the compelling rationale between retinal coverage factors, both physiological and psychophysical, I included a scaled pα RGC and scaled magnocellular dLGN model in my final selection of models to compare with my psychophysical data. To convert the density distribution models into the appropriate format for comparison with my displacement thresholds I transformed these distributions into models of inverse density, by normalising each distribution to its foveal value, following the methodology of Drasdo (1989) (Figure 7.5). If we consider these models of cell density as a proxy for spatial resolution, these inverse density models now represent the decline in spatial resolution across eccentricity.

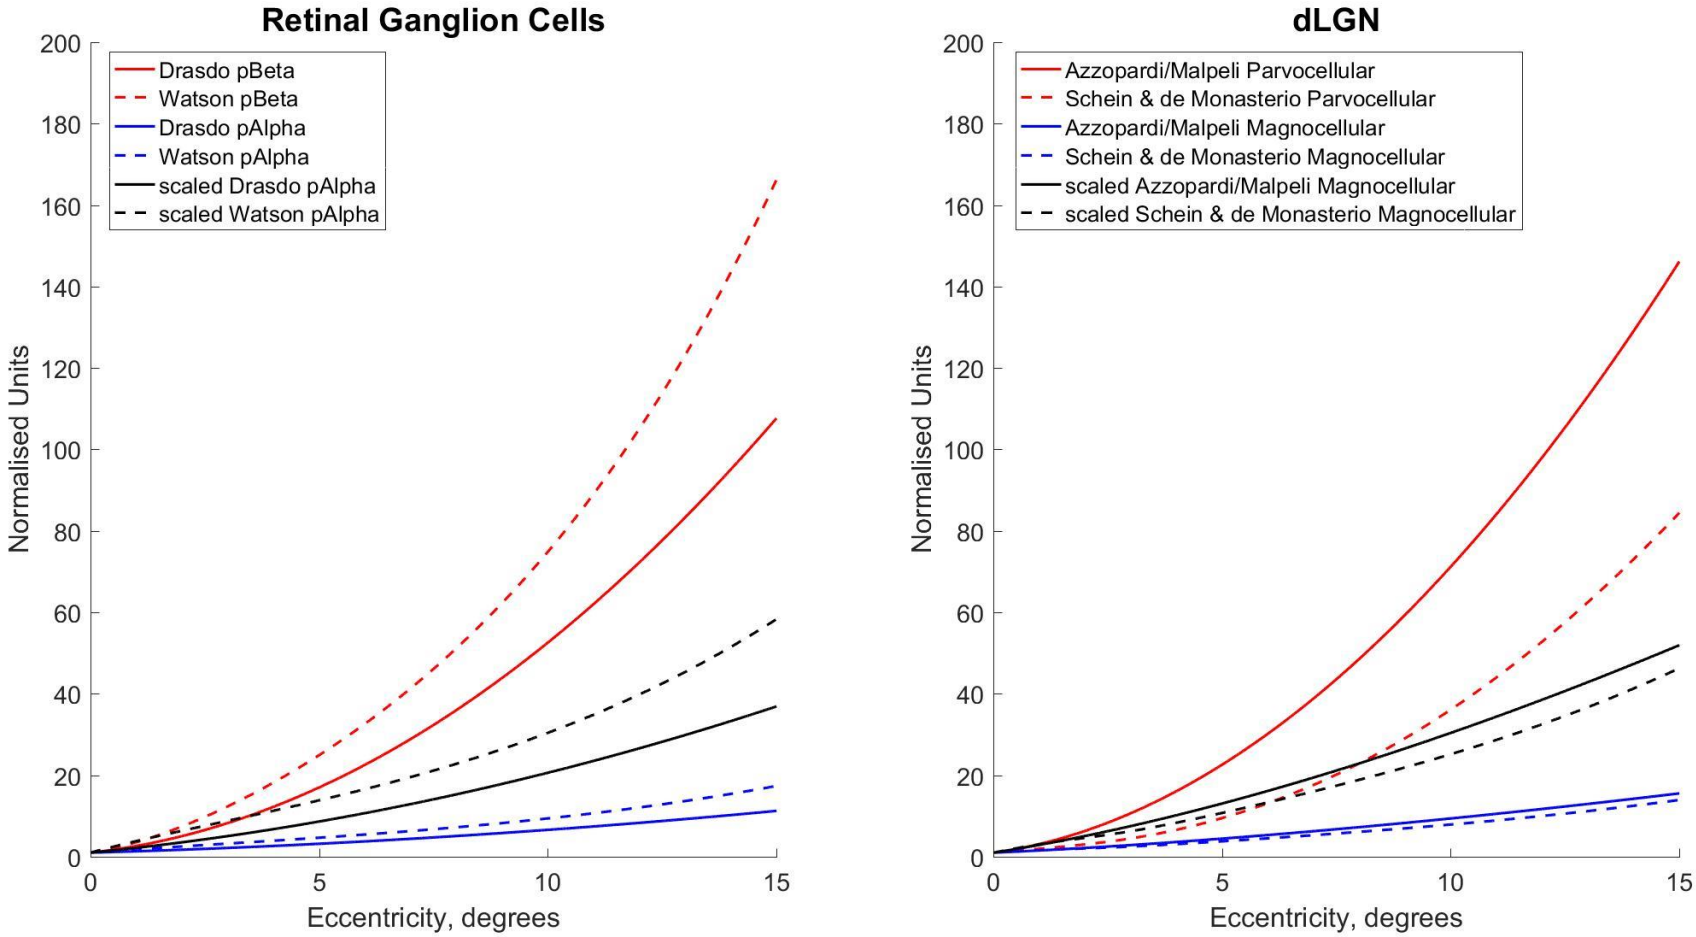


Figure 7.5. Inverse cell density distributions for the models of retinal ganglion cell data, left, and the models of the parvocellular and magnocellular layer of the lateral geniculate nucleus, right, with the scaled pα and Magnocellular dLGN distributions scaled by a scaling factor of 3.

Preparing displacement thresholds

The purpose of this chapter was to assess whether the different classes of retinal ganglion cell were convincing candidates for the physiological substrate of the psychophysically dissociated phase and position motion processing systems. The psychophysical data which I compared with these models of inverse cell data were presented in Chapter 3. These psychophysical data are in the form of direction discrimination thresholds, collected at four locations between fovea and at an eccentricity of 15 degrees (Figure 7.6, left panel). A two-frame Gabor was displaced, either by the phase displacement of the carrier of the Gabor (Carrier condition), or by the equidistant displacement of the Gaussian aperture and rotation of the sinusoid carrier (Aperture condition). Performance was measured across a range of spatial frequencies at each eccentricity, with the displacement corresponding to 75% performance, designated as the 'displacement threshold'.

These displacement thresholds increased with eccentricity, with a significantly greater increase with eccentricity for the aperture cue displacement thresholds, suggestive of different psychophysical systems for the processing of carrier and aperture cues for motion. To prepare these displacement thresholds for comparison with the cell density models, I normalised the displacement thresholds for each participant by their foveal value and then scaled the data relative to the range of values present in the RGC and LGN models, (Figure 7.6, right panel).

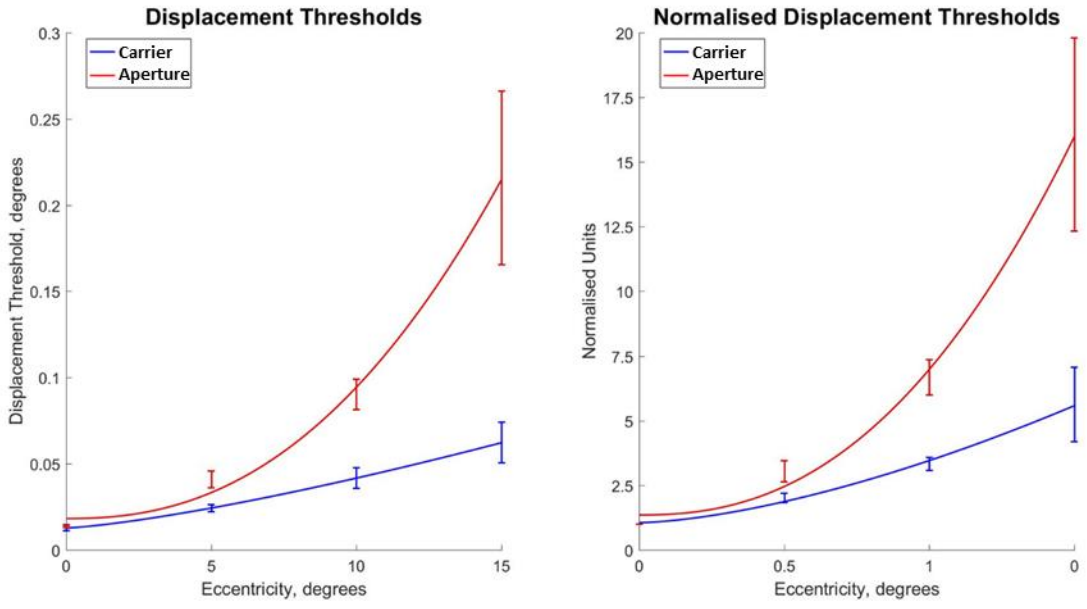


Figure 7.6. Displacement thresholds for the Carrier and Aperture conditions, reported in Chapter 3, left, and the same displacement thresholds normalised to their foveal values, right.

Whilst the carrier displacement contains both a phase cue and a position cue, in the form of the displacement of the carrier features, the aperture displacement contains no corresponding change in the phase of its Fourier components, which are only made explicit following rectification e.g. full wave rectification, as demonstrated in Chapter 2.4.1. Despite the confound of the position cue in the Carrier condition, for the purposes of this analysis the displacement thresholds for the Carrier condition are used as data representing the phase cue. These displacement thresholds may indeed be determined by a putative phase system: given that the putative phase system is proposed to have higher spatial sensitivity, displacement thresholds for the carrier cue may be determined by the phase cue, and not the congruent position cue.

7.3 Results

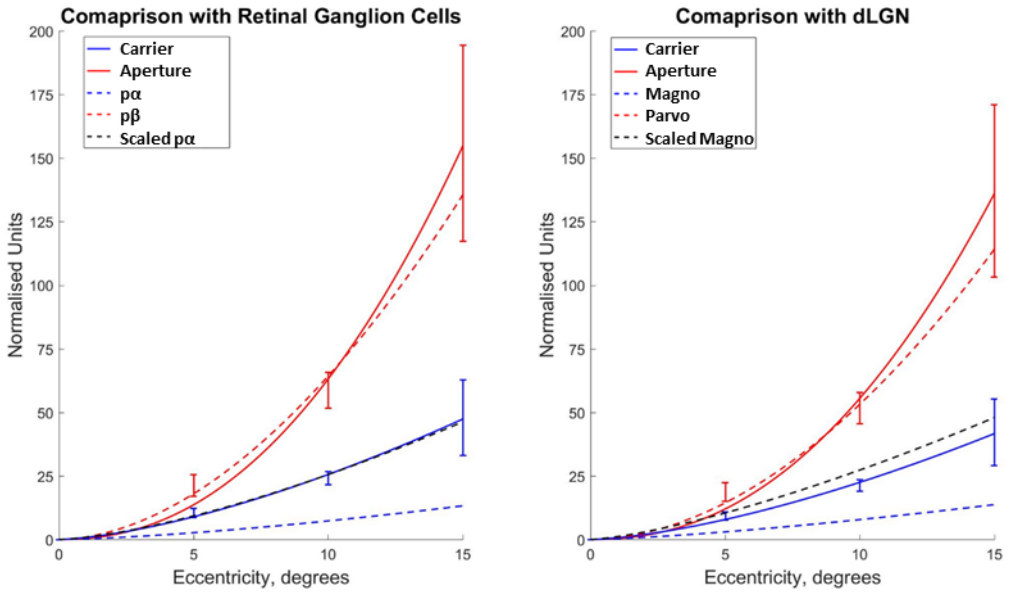


Figure 7.7. Figure displaying the displacement thresholds, and the cell density models. The dashed lines represent the mean estimates across the two models produced at both RGC and LGN layers, with the red dashed lines corresponding to the pβ and parvocellular LGN models, the blue dashed lines corresponding to the pα RGC and magnocellular LGN models, and the black dashed line corresponding to the pα RGCs and magnocellular LGN models scaled using a retinal coverage factor of 3.5.

I collated two models of cell density for the pα and pβ retinal ganglion cells, and for the cell density of the parvocellular and magnocellular neurons in LGN. For the RGC estimates, I reported the cell density in pα and pβ RGCs published by Drasdo *et al.* (Drasdo *et al.*, 1989), and constructed a second model by combining the pβ and total RGC estimates published by Watson (Watson, 2014), with ratio of pα to pβ RGCs suggested by Drasdo *et al.* (1989). For the models of LGN cell density, I used the model of Schein & de Monasterio (Schein & de Monasterio, 1987), and a combined model where I had applied the ratio of parvocellular and magnocellular LGN cell density, reported in Azzopardi &

Cowey (Azzopardi & Cowey, 1996a) to the summed distribution of parvocellular and magnocellular LGN neurons reported by Malpeli *et al.* (Malpeli *et al*, 1999). I also applied an estimate of the retinal coverage factor to produce a scaled model of the pα RGCs and magnocellular LGN neurons (see Figure 7.5).

To compare the cell density models to the displacement threshold data, the displacement thresholds were first normalised to the foveal values, and then scaled relative to the range of values present in the RGC and LGN models respectively. This produced slightly different transformations of the displacement thresholds for the RGC and dLGN comparisons. In Figure 7.7 I plot the mean of the model pairs (displayed in Figure 7.5), and overlay the normalised displacement thresholds, shown unscaled in Figure 7.6 (right panel).

To compare the psychophysical displacement data to the cell density models, I calculated the root mean squared error for each participant at each of the 4 eccentricity values, between the normalised displacement thresholds for the Carrier and Aperture conditions, reported in Chapter 3, and the normalised inverse cell density distributions, for the three model types; parvocellular, magnocellular, and scaled magnocellular, for the RGC models, and for the LGN models.

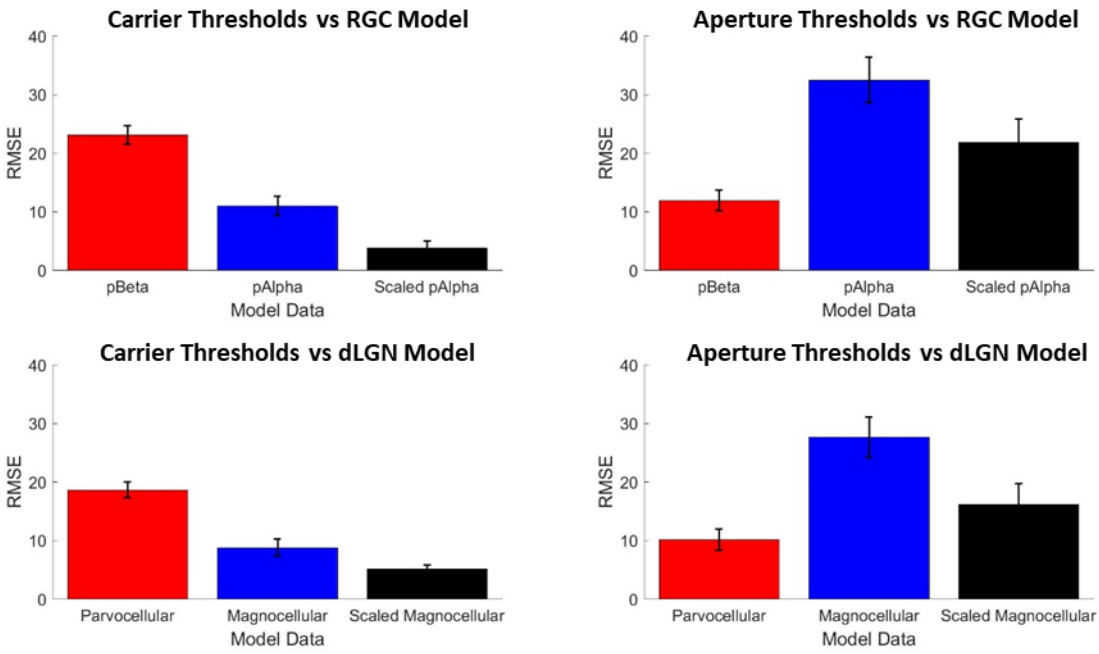


Figure 7.8. Figure displaying the root mean squared error (RMSE) values, calculated between the Carrier and Aperture condition displacement threshold data for each participant and the three models produced for both the RGC and LGN data.

The mean RMSE values for each model type compared to the normalised Carrier and Aperture displacement thresholds are plotted above in Figure 7.8. I ran a 2 (cell class; RGC vs LGN) x 3 (model; parvocellular, magnocellular, scaled magnocellular) repeated measures ANOVA on these RMSE values, for the displacement thresholds from the Carrier and Aperture conditions. The cell classes, the parvocellular, magnocellular and scaled magnocellular refer to the potential pathway in the visual system, where the RGC cells and the LGN cells form inputs at consecutive early stages of the pathway.

Carrier condition

There was a significant effect of the model type $F(2,12) = 80.516$, $p < .001$, $\eta_p^2 = .920$, where RMSE values were lowest for the scaled magnocellular models, and highest for the parvocellular models. There was no significant effect of cell class, $p = .127$, but there was a significant interaction between cell class and model type, $F(2,12) = 20.977$, $p < .001$, $\eta_p^2 = .750$, and a significant between subjects interaction, $F(1,6) = 184.335$, $p < .001$, $\eta_p^2 = .996$.

I explored these effects further through a 2(cell type; RGC vs LGN) x 2(magnocellular type; magnocellular vs scaled magnocellular) repeated measures ANOVA. There was no significant difference in cell type, $p = .452$, but there was a significant difference in magnocellular type, ($F(1,7) = 12.704$, $p = .009$, $\eta_p^2 = .645$), where RMSE was lower in both RGC and LGN scaled magnocellular estimates, and there was a significant interaction between cell type and magnocellular type, ($F(1,7) = 27.746$, $p = .001$, $\eta_p^2 = .799$). These differences reveal that there is a significantly lower RMSE value, and hence goodness of fit, for the scaled magnocellular values than the unscaled magnocellular values, but that there was not a significant improvement in the fit between RGC and LGN models (see table 7.1).

There was also a significant between subjects effect, $F(1,7) = 55.609$, $p < .001$, $\eta_p^2 = .888$.

Table 7.1 Table of the RMSE values, from comparing the Carrier condition displacement data from Chapter 3, with the models of sampling resolution calculated for the classes of RGC and layers of the LGN.

	RGC		LGN	
	Mean	Std. Error	Mean	Std. Error
Parvocellular	23.081	1.603	18.606	1.374
Magnocellular	10.954	1.639	8.736	1.439
Scaled Magnocellular	3.774	1.111	5.161	0.598

Aperture condition

There was a significant effect of cell class on the RMSE values, F(1,6) = 33.119, p = .001, η_p^2 = .826, where the RMSE values were lower for the LGN models than the RGC models. There was also a significant effect of model type, F(2,12) = 15.749, p < .001, η_p^2 = .692, where the RMSE values were lowest in the parvocellular models. There was a significant interaction between the cell class and model type, F(2,12) = 23.826, p < .001, η_p^2 = .771, and there was a significant between subjects effect, F(1,6) = 55.623, p < .001, η_p^2= .888.

I explored these results further conducting a repeated measures ANOVA between the RMSE values for the parvocellular RGC and the parvocellular LGN estimates. There was a significant difference between the RMSE values, (F(1,7), = 15.937, p = .005, η_p^2= .695), with a significantly lower RMSE for the LGN parvocellular estimates (M = 10.132, SE= 1.823), than the RGC estimates (M = 11.869, SE = 1.770). There was also a significant between subjects effect, (F(1,7) = 38.047, p < .001, η_p^2 = .845).

Table 7.2 Table of the RMSE values, from comparing the Aperture condition displacement data from Chapter 3, with the models of sampling resolution calculated for the classes of RGC and layers of the LGN.

	RGC		LGN	
	Mean	Std. Error	Mean	St. Error
Parvocellular	11.869	1.770	10.132	1.823
Magnocellular	32.431	3.885	27.618	3.418
Scaled Magnocellular	21.833	3.940	16.160	3.482

7.4 Discussion

Direction displacement thresholds for Gabor stimuli where motion is produced from the phase shift of the carrier, or the displacement of the Gaussian aperture, are significantly different across the visual field, with a significantly greater increase in displacement thresholds with eccentricity for aperture cue stimuli than for carrier cue stimuli (Chapter 3). This difference could potentially be explained by the activity of two separate motion processing pathways, one processing phase cues and the other processing position cues, produced by two distinct physiological pathways. In this chapter I propose that the psychophysical dissociation in displacement thresholds could be produced by differences in the spatial resolution of the two pathways. One limiting factor of the spatial resolution of visual pathways is the initial sampling of the retina, which takes place at the

earliest levels of the visual pathway. The sampling of the retina can be estimated from the density of receptive fields of the retinal ganglion cells, where greater cell density produces greater sampling of the retina.

I have suggested that the physiological substrate for the phase system is formed from the projections of the pα RGCs to the magnocellular layer of the LGN, and the physiological substrate for the position system is formed from the projections of the pβ RGCs to the parvocellular layer of the LGN. In this chapter, I produced models of cell density for these two pathways, for the retinal ganglion cell level and at the level of the lateral geniculate nucleus, using a combination of published cell density distributions, (Drasdo, 1989; Watson, 2014; Schein & de Monasterio, 1987; Malpeli *et al.*, 1999) and published ratios of cell classes (Drasdo, 1989; Azzopardi & Cowey, 1996a). I also produced a scaled model for the pα RGC and magnocellular layers of the LGN, to compensate for the increased retinal coverage factor in the pα RGCs (Perry & Cowey, 1984). These models of cell density can be considered as a metric for the spatial resolution of the different classes of RGC and different layers of the LGN.

To compare the cell density to the displacement thresholds, I normalised both the cell density and the displacement thresholds to their foveal values. I calculated the root mean square error value for each participant, using the difference between the normalised displacement threshold, and the normalised estimates from each of the three models (parvocellular, magnocellular and scaled magnocellular), for the RGC and the LGN data.

The Aperture condition displacement thresholds from Chapter 3 were significantly more similar, with smaller RMSE values, to the parvocellular density estimates, at the RGC and LGN level, than either the magnocellular or the scaled magnocellular estimates. This supports my initial hypothesis; that the position systems is produced using input from the pβ RGCs and the parvocellular layers of the LGN. There was a significantly lower RMSE for the Aperture condition displacement thresholds for the LGN parvocellular estimates than the RGC estimates. If the Aperture condition displacement thresholds are determined from the behaviour of a system driven by input from the pβ RGCs, and parvocellular layers of the LGN, then the change in displacement thresholds with eccentricity will reflect the topography of the higher levels of the system. There is an increase in foveal overrepresentation in the pβ projections to the LGN, (Azzopardi & Cowey, 1996b), which could explain the improved goodness of fit between Aperture condition displacement thresholds and the density estimates of the parvocellular layers of the LGN compared to the pβ RGC density estimates.

The Carrier condition displacement thresholds were significantly more similar to the scaled magnocellular and magnocellular density estimates, with smaller RMSE values, at the RGC and LGN level, than the parvocellular estimates. The RMSE values were significantly lower, indicating a greater goodness of fit, for the scaled magnocellular data than the unscaled data. This supports my initial hypothesis, that the phase system is produced from the pα RGCs and the magnocellular layers of the LGN, with the scaling accounting for the overlapping receptive fields of the pα RGCs. There was not a significant improvement in the goodness of fit between the scaled RGC and LGN models, which could reflect the proposed absence of further foveal overrepresentation between the pα RGCs and the magnocellular layers of the LGN (Azzopardi & Cowey, 1996b).

The significant improvement in goodness of fit of the Carrier condition displacement thresholds to the scaled data is partially driven by my choice of retinal coverage factor, where a retinal coverage factor between 1 and 3.5 would have improved goodness of fit, but potentially not significantly. Perry & Cowey (Perry & Cowey, 1984) suggested a range of retinal coverage factors for the pα RGCs between 2 and 7, as a result of measurement taken at 3 locations each in the nasal and temporal hemifields of the macaque retina. It is probable that the retinal coverage factor will vary between humans and macaques, since Dacey & Petersen (Dacey & Petersen, 1992) suggest that the size of the pα RGCs are 20-95% larger in humans. If the precision of Carrier condition displacement thresholds can be improved, either by increased sampling across eccentricity, or by greater precision within participants, then this method could be used to estimate the retinal coverage factor for human RGCs.

Across the participants the relationship between displacement thresholds from the Carrier and Aperture conditions and the estimated cell density distributions supports my hypothesis that the physiological substrate for the position system is initially produced from the pβ input to the parvocellular LGN layers, and that the physiological substrate for the phase system is initially produced from the pα input to the magnocellular LGN layers, when scaled by retinal coverage factors.

Limitations of the approach and data

There are substantial limitations to this approach, both in the collection of the psychophysical data, and the assumptions about the visual system that underpin this approach.

Firstly, this method is purely correlational and reliant on the large difference in density distributions of the pα and pβ RGCs. There were significant between subjects effects for both the phase and position cue displacement thresholds, as reported in Chapter 3, and this variance between participants indicates that this approach may only be useful for the comparison of psychophysically obtained displacement thresholds between pathways with large differences in their sampling resolution. If correlations could be made for individuals participants, using individualised models of cell density distributions, estimated from measurements taken by retinal imagining, and these models have a better fit with individual participants' displacement thresholds, this would be a more convincing correlation.

Secondly, this chapter assess correlations for motion discrimination but sampling at a single temporal frequency. This method has been applied successfully to spatial vision, where the spacing of pβ RGCs has been correlated with acuity in human participants, demonstrated by Anderson *et al.*, (Anderson *et al.*, 1991), and Watson (Watson, 2014) with data from Rossi and Roorda (Rossi and Roorda, 2010). However, I am investigating motion perception, which in two dimensions; space and also time. The psychophysical data used in this chapter are the displacement thresholds presented in Chapter 3 where I assessed the direction discrimination of the carrier and aperture cues in a two-frame Gabor. In this chapter I presented stimuli at a relatively low temporal frequency as I wanted to maximise the sensitivity of both the phase and position systems, where the spatial sensitivity of the position system declines with increasing temporal frequency, whereas the spatial sensitivity of the phase system remains high (Gorea, 1995). The temporal frequency that stimuli are presented at can bias reported motion direction, where Heller *et al.* (Heller *et al.*, 2021) have shown that in an orthogonal Gabor, reported motion direction is biased towards the carrier cue at high temporal frequencies and the aperture cue at low temporal frequencies. My choice of a low temporal frequency may therefore be resulting in lower spatial acuity for the Carrier condition, producing larger displacement thresholds and so a steeper gradient of normalised displacement thresholds with eccentricity than if the stimuli were presented at a higher temporal frequency.
Replicating my original direction discrimination experiment at a range of temporal frequencies would reveal if displacement thresholds in the Carrier condition are larger at low temporal frequencies and smaller at high temporal frequencies. This would demonstrate whether this cell

density comparison presented in this chapter is an appropriate method for investigating motion systems. If increasing temporal frequency decreases the size of displacement thresholds in the Carrier condition, then a more accurate comparison between the cell density of the pα RGCs and magnocellular LGN cells may be achieved by comparisons between displacement thresholds collected for Gabors presented at higher temporal frequencies.

Thirdly, this approach assumes that the pα and pβ retinal ganglion cells are identical sampling units. This is not the case, as I have shown that differences in morphology, such as the difference between the surface area of the dendrites (Polyak, 1941), affect the function of the pathway, where the increased dendrite area in the pα RGCs increases the retinal coverage factor (Perry *et al.*, 1984). We can correct for the difference in dendrite surface area, by applying a scaling factor to the pα RGC density distribution, producing the scaled magnocellular model. However, other differences between the cells may be less easily corrected for, such as the difference in the larger cell body size of the pα RGCs than the pβ RGCs, (Perry *et al.*, 1984). A further complication is that cell body size and dendritic field size increase with eccentricity, in both pα and pβ RGCs. If scaling factors were needed to compensate for these, then the values of the scaling factor would increase by eccentricity. These morphological differences between the pα and pβ RGCs should have functional correlates, and therefore invalidate the assumption of this approach, that the pα and pβ RGCs can be considered identical sampling units, and that differences between the two systems can be explained solely by the differences in the cell density distribution.

While the differences in pα and pβ RGC morphology limit the validity of this approach for correlations with psychophysical data, the models of scaled pα RGCs which I have constructed may be able to shed light on a puzzling difference between the resolution for motion and position in the periphery. In the periphery participants are able to discriminate the motion direction of stimuli whilst not being able to resolve the spatial separation of the start and end locations of the motion stimulus, (Exner 1875; Foster *et al.*, 1989). If we accept that the spatial resolution for motion discrimination is the product of a system with input from the pα RGCs, and that the spatial resolution for position discrimination, either for static or position cues for motion, is the product of a system with input from the pβ RGCs, then this result is in direct contradiction of the cell density estimates for these two systems, where across the retina the density of the pβ RGCs always exceeds the pα RGCs (Figure 7.2). However, with increasing retinal coverage factors applied to the pα RGCs the scaled density, and therefore the functional resolution, of the pα RGCs exceeds the pβ RGCs, with this crossover point occurring closer to the fovea with greater retinal coverage factors (Figure 7.4). Exner's (1875) demonstration has been suggested as evidence for the specialisation of the

periphery for motion perception (Exner, 1875; Finlay, 1982). I have demonstrated in Chapters 3 to 5 that the change in reported motion between the fovea and periphery, where motion is biased towards the position (or aperture) cue at fovea and the phase (or carrier) cue in the periphery, can be explained by a change in spatial frequency sensitivity of the phase and position systems. The difference in spatial resolution across the retina between the pβ and scaled pα RGCs is a plausible physiological explanation for my psychophysical data, and the results of Exner (Exner, 1875) and Foster *et al.* (Foster *et al.*, 1989).

Finally, the biggest limitation of this approach is in the choice of displacement thresholds for correlation with cell density estimates. The thresholds used in this chapter were from Gabor stimuli, and whilst the Aperture condition can be described as a 'pure' position cue stimulus, the Carrier condition contains both phase and position cues, and as such, the displacement thresholds cannot be considered as pure phase cues, as both a putative phase and position system could be involved in motion discrimination. Applying a pedestal to the Gabors would remove the confounding position cue in the Carrier condition, although necessitate the use of more frames. Choosing stimuli with similar image characteristics that differ only in the presence of phase or position cues is difficult. An alternative approach is to gather data across many stimuli with varied image constructions, with the aim that differences in image characteristics will be smaller than the overall difference between stimuli that selectively contain phase and position cues, as was the approach in comparing E2 values between various stimulus constructions taken in Chapter 3.

Conclusion

While there are limitations in the application of this approach to motion stimuli, there is a clear correlation between the Aperture condition displacement thresholds and the estimated spatial resolution of the parvocellular system, and between the Carrier condition displacement thresholds and the estimated spatial resolution of the magnocellular system. This is suggestive that the phase and position systems may receive a dominant input from the pα RGCs and pβ RGCs, respectively. This model should be extended in the future to account for the cell density distributions of the small bistratified RGCs, which likely contribute to the position system. With a better understanding of the different physiological components, it may be possible to produce theories about the mechanisms within these proposed system, and compare these with models of motion processing, such as filter-rectify-filter models.

Chapter 8: Discussion

This thesis investigated whether there are psychophysical dissociations between phase and position cues for human motion perception, and whether these can be explained by the activity of two separate systems: one selectively processing motion stimuli in terms of their phase characteristics and one selectively processing motion stimuli using the change in position of the image features.

An initial challenge for the thesis is that phase and position cues are a description of the image characteristics, and for most stimuli, mathematically equivalent. However, stimuli can be produced that selectively contain phase or position cues, or contain both cues but signalling motion in opposing directions. Critically, motion discrimination is possible across these stimuli, indicating that there are visual systems that can extract motion from these different cues. Various models have been proposed to suggest possible mechanisms through which these cues may be extracted from the visual input. Whilst cues can be described as phase and position cues, this does not necessitate a system performing Fourier analysis to extract the phase cues, and likewise, a system that processes position cues does not necessarily extract and identify the image features before updating the change in position with time. The aim of this thesis was not to characterise the models of the mechanisms used to separately extract phase and position cues, but to provide psychophysical evidence in support of the presence of two systems: characterising the mechanisms by which cues are extracted and processed is the logical future extension of this project.

This thesis aimed to investigate psychophysical dissociations of the phase and position cues for motion, specifically looking at differences in direction discrimination for stimuli varying in spatial frequency across the visual field (Chapters 2-5). Motion in conflicting cue stimuli has been shown to systematically be reported in the phase or the position cue direction, depending on the parameters that the stimuli are presented at. Spatial frequency and eccentricity produce consistent motion reversals, and across a range of stimulus constructions, suggesting that the reversals may be attributable to system properties rather than the image characteristics of the stimuli. In order to determine whether these reversals were the property of two putative separable systems, the aim was to test stimuli that contained motion solely through a phase or position cue, and use these data to generate testable predictions for reported motion in conflicting cue stimuli. Producing these stimuli proved challenging, and the stimuli presented in this thesis represent a compromise. Gabor stimuli were used as stimuli where the carrier and aperture could be separately manipulated, to produce stimuli with a single or conflicting motion cues. Whilst changes in the phase of the carrier of the Gabor produce both a phase cue and a congruent position cue (in the change of the position of the carrier features), the separate manipulations of the carrier and aperture were used to generate a hypothesis about changes in spatial sensitivity to these cues that could be tested in the conflicting

cues condition. Reverse phi stimuli were also used in this thesis, as although the phase and position cues could not be presented independently, the cues signal motion in opposite directions with no congruent confounding cues. Comparing the pattern of motion reversals in the reverse phi and Gabor data suggests that motion in the carrier and aperture cue directions in the Gabors may be governed by similar processes to the phase and position cues in the reverse phi data.

These psychophysical dissociations in reported motion direction, as a result of changing spatial frequency or retinal position, were used to argue that there are two distinct systems that process the phase and position cues for motion, and that these can explain qualitative differences between motion perception at the fovea and periphery. These psychophysical dissociations potentially represent distinct physiological substrates of these two proposed systems. This was discussed in terms of a cortical and midbrain projection to V5/MT, investigating the motion discrimination of phase and position cues in a V1-lesioned patient (Chapter 6), and in terms of different classes of retinal ganglion cells (RGC), where participants' displacement thresholds for the carrier and aperture manipulations in Gabor stimuli were compared to modelled density distributions of RGC classes across the retina (Chapter 7).

8.1 Overview of results

Chapter 2: Exploring stimuli for spatial and temporal frequency dissociation of phase and position cues

Chapter 2 explored dissociations in temporal and spatial frequency of grating and Gabor stimuli, to determine which would be more experimentally tractable in searching for psychophysical dissociations for phase and position cues.

In Chapter 2 I introduced a novel stimulus, the 'MF3F', consisting of interleaved frames of the Missing Fundamental Grating (MFG) and a third harmonic, matched in spatial frequency to the third harmonic of the MFG. I compared motion discrimination for the MF3F across a range of blank interframe intervals (IFI) with the Missing Fundamental Grating (MFG), which contained phase cues of the matched spatial frequency. Motion discrimination was abolished in the MF3F with IFIs greater than 20 ms, whereas in the MFG motion was reported in the phase cue direction with no IFI, and the position cue direction when IFIs were introduced, consistent with the work of Georgeson & Harris (1990). I suggested that the results indicated a difference in how motion cues were integrated across time, proposing that phase cues were analysed by a system processing the visual input in spatiotemporal coordinates, and the position cues by a system that stored information about features and their position, which was updated following a change in position and unimpaired by the

insertion of blank IFIs. However, whilst the stimulus was intended to contain solely phase cues, the zero-crossings between frames could potentially be used as a position cue for motion discrimination. An additional issue in the MF3F stimulus is that the MFG changes appearance based on contrast and spatial frequency (Campbell & Robson, 1968), limiting the use of these stimuli across spatial frequency, and across eccentricity.

A simpler method to investigate temporal integration dissociations for phase cues would be presenting a pedestalled sine wave at different spatial and temporal frequencies across the visual field, and mapping the IFIs at which motion discrimination falls to chance. This could then be compared against a representative frame, the sum of the pedestal and sine wave, displaced as a static grating on each frame. This static grating condition will contain both phase and position cues, but at the limit where the performance on the pedestalled sine wave declines to chance, motion discrimination on the static grating should remain high, as in the MFG at long IFIs, where motion discrimination is reliant on the position cue in the image. This can then be used to demarcate the region of temporal space where position cues, but not phase cues, are available. These psychophysical dissociations could be used to inform comparisons with published models of potential mechanisms for processing phase and position cues, as discussed in the Introduction. Substantial differences in the temporal domain would be evidence for a more fundamental difference in the processing between phase and position cues, supporting a model based on updated feature positions, rather than filter-rectify-filter models for processing position cues.

In Chapter 2 I also explored motion discrimination as a function of spatial frequency for a set of Gabor stimul, where motion was produced from either the phase shift of the carrier, a displacement of the aperture, or containing both cues and signalling motion in different directions, ‘conflicting cue’ stimuli. Participants showed a consistent reversal of reported motion direction in conflicting cue Gabor stimuli, reporting motion in the position cue direction at the lowest spatial frequencies, and in the phase cue direction at the mid-range spatial frequencies. A critical limitation of the Gabor stimulus is that it cannot be used to produce ‘pure’ phase cues without a congruent position cue occurring, due to the change in location of the carrier features when the phase of the carrier changes. However, the similarity of the motion reversals in the conflicting cue Gabors to reversals in reverse phi stimuli, which don’t contain this congruent position confound, suggests that motion reported in the carrier cue direction in the conflicting cue Gabors may still be the result of systems relying more heavily on changes in phase.

Finally, I ran a Fourier-transform based motion energy analysis (Shapiro *et al.*, 2010) on both classes of stimuli tested in this chapter. This analysis is currently limited to identifying the presence or

absence of phase cues in images, and is not intended as a description of potential mechanisms in the visual system, or of predicted motion experience. I discussed further extensions to the nME analysis to produce more accurate models of motion processing mechanisms, and potentially develop models that could systematically identify the presence of position cues..

Chapter 3: Spatial frequency sensitivity across the visual field

In Chapter 3, I used two-frame Gabors to assess spatial frequency sensitivity across the visual field for phase and position cues, where a differing response to spatial frequency could suggest different systems processing the phase and position cues. Motion was produced by either a phase displacement of the grating or a position displacement of the aperture. Participants reported the perceived direction of motion in a 2afc procedure, and proportion correct was measured for Gabor stimuli presented at the fovea and three locations in the periphery. At the fovea, there was very little difference in motion discrimination performance for the Gabors containing either phase or position cues. In the three peripheral locations, the parabolic curves which described motion discrimination performance plotted across spatial frequency were displaced along the spatial frequency axis towards lower spatial frequencies, with participants requiring Gabors to be presented at lower spatial frequencies to produce similar levels of proportion correct. Motion discrimination for Gabors where motion was produced from the position displacement declined more rapidly at high spatial frequencies than the phase cue condition.

From these data I extracted the spatial frequency values which corresponded to 75% correct for each participant, and expressed these as displacement thresholds, i.e. the distance that the aperture cue was displaced over, or the length of the phase shift (1/4 of a cycle) of the carrier. Displacement thresholds increased with eccentricity for both conditions, but this increase was significantly greater for the aperture cue, reflecting the greater decline in spatial sensitivity with increasing eccentricity. This result replicates those of previous experiments, which have reported steeper increases in the change in displacement thresholds with eccentricity for stimuli containing changes in position (Minimum angle of resolution), than changes in phase (Motion in light: Weymouth, 1958; Fine grain motion illusion: Foster *et al.*, 1989).

As in Chapter 2, the Carrier condition cannot be described as a ‘pure’ phase cue stimulus due to the presence of congruent changes in position of the carrier. This limits whether the displacement thresholds can be interpreted as the response of the phase system, as at different spatial frequencies either or both of a putative phase or position system could be contributing to motion

discrimination. The Aperture condition contained no confounding phase cues, as demonstrated through the nME analysis in Chapter 2, and can be considered a 'pure' position cue stimulus. The aperture cue displacement threshold gradients were similar to published thresholds collected for acuity for static gratings (Strasburger *et al.*, 2011). This supports an interpretation of a position system less reliant on the temporal information of the stimulus, as suggested in Chapter 2 from the retained motion discrimination of MFG with long inter-frame intervals.

The confound of the position cues in the Carrier condition limit the interpretation of the data in terms of the spatial sensitivity of an underlying phase system. This position cue confound seen in the Carrier condition could be eliminated by applying a pedestal to the Gabor stimulus, but this would then require the stimuli to be presented in multi-frame displays. If the Aperture stimuli were also increased to multi-frame displays this would reduce the specificity of location at the different eccentric locations. At increasing eccentricity, lower spatial frequencies are used, and as the aperture of the Gaussian is set in ratio to the carrier frequency, this would exacerbate the confound of increased retinal area covered by the Aperture condition compared to the stationary aperture in the Carrier condition. This confound only occurs if the position cue is produced from a change in position of the aperture: a better method might be to use grating features as the position cue. The E2 value for a putative phase system could be calculated from displacement thresholds could be extracted for a pedestalled grating, and E2 values for a position based system could be calculated from the displacement thresholds from a representative frame of the pedestalled grating (to match the image characteristics of the pedestal, but without having separate phase shifts for the pedestal and fundamental frequency) displaced within a stationary aperture with long inter-frame intervals (to eliminate the contribution of phase based system).

Chapter 4: Perceived motion in Gabor based conflicting cue stimuli

Motion appears qualitatively different in the fovea and periphery, a property that has led to assertions that the periphery is 'specialised for motion perception' (Exner, 1886). This difference has also been demonstrated quantitatively, by Shapiro *et al.* (2011) who demonstrated that for a Gabor with orthogonally oriented phase and position cues, motion is reported in the position cue direction at fovea, but at trajectories intermediate to the phase and position cue when viewed peripherally. Shapiro *et al.* suggested that this reflected a difference in integration between phase and position systems, that occurs in the periphery and not in the fovea, the 'feature blur' hypothesis. However, the data I presented in Chapter 3 offers a competing explanation: at greater eccentricities participants retained higher motion discrimination for the carrier cues than the aperture cues. This could drive the reversals of motion from the aperture cue direction at low spatial frequencies, to

reported in the carrier cue direction at high spatial frequencies. I refer to this as the 'spatial scaling' hypothesis.

In Chapter 4 I assessed reported motion direction for conflicting cue Gabor stimuli, in two experiments.

In the first experiment, I assessed motion discrimination for an Opposing Motion Gabor, at the fovea and at three peripheral viewing locations. At the fovea motion was reported in the aperture cue direction for the majority of trials. This is consistent with Shapiro's 'feature blur' hypothesis – no integration between the phase and position cues - but would also be consistent with the spatial scaling hypothesis if the spatial sensitivities of the phase and position system were overlapping, and the overall relative strength (or output) of the position system was greater. In the peripheral viewing locations, motion was reported in the aperture cue direction at low spatial frequencies and the carrier cue direction at high spatial frequencies, consistent with the spatial frequency dissociation reported for pure phase and position cues in Chapter 3. However, motion discrimination was assessed using a 2 interval stimulus where participants reported the interval in which motion was perceived upwards. Responses were therefore either coded as in the phase or position cue direction, and errors are not measured. Whilst performance was predominantly reported in the aperture cue direction at the fovea, as predicted by Shapiro's feature blur hypothesis, it was unclear whether the trials where motion was reported in the carrier cue direction were evidence of motion reported in the carrier cue direction, or errors. Another limitation was in the use of the 'Opposing Motion' Gabor where the carrier and aperture cues were arranged to signal motion in opposite directions. Any 'integration' between cue directions could result in the perception of a smaller displacement, which the response methodology would not detect. Therefore data could not be used to determine if integration was occurring, and so could not differentiate between Shapiro's feature blur hypothesis and my spatial scaling hypothesis.

These limitations were addressed in the second experiment in Chapter 4, where I used both an Opposing Motion Gabor and an Orthogonal Motion Gabor, and participants reported perceived motion trajectory by rotating an arrow. Stimuli were displayed across a range of spatial frequencies, and viewed foveally or peripherally, at 10 degrees eccentricity. In the periphery participants reported motion at intermediate trajectories between the phase and position cues for the Orthogonal Gabor, with motion trajectory closer to the position cue direction at low spatial frequencies and the phase cue position at higher spatial frequencies. In the Opposing Gabor, motion was reported either in the phase or position cue direction, with more trials reported in the position cue direction at low spatial frequencies, and in the phase cue direction at higher spatial frequencies.

At the fovea motion in the Opposing Gabor was reported predominantly in the position cue direction, but at higher spatial frequencies, was reported in the phase cue direction more times than would be predicted by chance. Critically, for the Orthogonal Gabor, the mean trajectory of motion was skewed in the direction of the phase cue, especially at higher spatial frequencies. This demonstrates integration between the phase and position cue that is spatial frequency dependent, as predicted by the spatial scaling hypothesis, but not by Shapiro's 'feature blur' hypothesis.

Assessing reported motion direction can differentiate between motion perceived in either cue direction, and errors, but is still limited. It is possible that integration is occurring in the Opposing Motion condition, but this pattern of results would be indistinguishable from reporting motion in a 'winner takes all' fashion, solely in the aperture or carrier cue direction. Likewise, measuring the reported direction can show integration between the carrier and aperture cue direction in the Orthogonal Motion condition, but does not measure perceived distance which may also be affected. Replicating this experiment with reported distance and direction, by adjusting the length of the arrow used to report direction, could investigate how the carrier and aperture cues are being integrated.

Chapter 5: Perceived motion in reverse phi stimuli

In Chapter 5 I explored whether reported motion in reverse phi stimuli, presented at the fovea and periphery, showed a similar pattern of motion reversal to the conflicting cue Gabor stimuli, when spatial frequency and retinal location were varied - reported in the feature displacement direction at low spatial frequencies, and in the phase displacement at higher spatial frequencies.

In the first experiment participants viewed 8 frames of a contrast reversing grating foveally. Reported motion direction showed a complete reversal in reported motion with spatial frequency: reported in the position cue direction at low spatial frequencies, and in the phase cue direction at high spatial frequencies.

The second experiment in Chapter 5 assessed the displacement thresholds for motion reported in the phase and position cues directions, using displacing squares that simultaneously reversed in polarity, viewed at 5 eccentric locations between 5 and 25 degrees eccentricity. The size of the reverse phi squares was adapted using a staircase procedure, designed to converge at 25% and 75% of the psychometric functions shown in Experiment 1. These correspond to motion reported in the phase and position cue direction respectively. The displacement thresholds showed a similar gradient of displacement thresholds with eccentricity to the displacement thresholds reported for the Carrier and Aperture condition Gabors reported in Chapter 3: a much steeper increase in the size

of the displacement thresholds for the position (or aperture) cue with eccentricity, than the increase in the size of the displacement thresholds for the phase (or carrier) cue. However, as the reverse phi stimuli contains conflicting cues, displacement thresholds cannot be directly compared to the Gabor based displacement thresholds, which assume that discrimination rests on the activity of a single cue (carrier or aperture), rather than possible competition between two cues, and possibly two separate systems.

While both experiments in this chapter measure motion discrimination in reverse phi stimuli across spatial frequency changes, comparisons are limited between the two experiments as they contained different stimuli: the reverse phi grating used at the fovea (experiment 1) likely biased perceived motion towards the phase cue as there were more cycles of the grating visible, and the three-frame squares stimulus used in periphery (experiment 2) likely biased perceived motion towards the position cue as it used a single displaced feature. Comparisons are also limited due to experimental design: at the fovea data were collected using the method of constant stimuli to sample 6 spatial frequencies, and psychometric functions were fitted to the data. In the peripheral experiment data were collected using adaptive staircases to interpolate the 25% and 75% discrimination thresholds on the theoretical underlying psychometric function, rather than sampling the full extent of the function. Obtaining the full psychometric function provides more information about the competing motion directions across eccentricity: is the shape of the psychometric function simply displaced along the spatial frequency axis, or, is the psychometric function increasingly asymmetric with eccentricity with a longer tail for motion reported in the phase cue direction. This could suggest that the reversals of motion are produced due to a decline in sensitivity to position cues at higher spatial frequencies, with the difference between sensitivity to phase and position cues increasing with increasing eccentricity, as suggested by the spatial scaling hypothesis applied to conflicting cue Gabor stimuli in Chapter 4. More thorough descriptions of the underlying psychometric function, and the change with eccentricity and temporal frequency are essential if this data are to be used to determine whether reversals of motion in reverse phi stimuli reveal the activity of two competing motion mechanisms.

A limitation that affected both the of the experiments in this chapter was the use of the 2afc procedure, where participants reported which in interval motion was perceived as upwards. Both intervals contained an upwards and a downwards cue, so the interval chosen indicates whether participants perceived the phase or position cue as a stronger upward motion. However, as with the Gabor experiments in Chapter 3, this confounds errors with choosing the non-dominant direction. The decline in motion discrimination to chance performance, will likely be asymmetric at higher and lower spatial frequencies, which could confound interpretations of performance at extreme spatial

frequencies, and interfere with the fitting of psychometric functions to the data as a fully sampled range of spatial frequencies could not asymptote at floor and ceiling (position and phase cue directions), but have additional ascending and descending arms to chance performance. Poor sampling along the spatial frequency range could then result in artificially high floor and low ceiling asymptotes, affected by these ascending and descending arms. This effect could be partially mitigated by a change in the number of directions that the reverse phi stimulus could be oriented along, such as having horizontal and vertical oriented reverse phi gratings, although errors would likely continue to cluster along the same axis, as seen in the Carrier and Opposing Motion Gabors in Chapter 2.4.3 (Figure 2.11). Alternatively, the response method could be changed, where participants could report 'unclear' motion direction, or rate their confidence in their motion direction response. This would also be useful in interpreting the midpoint of the psychometric functions seen in experiment 1: were participants perceiving motion in either the phase or position direction on trials, were the cues perceived as having matched relative strength and producing unclear motion, or was the grating perceived as stationary.

An assumption underpinning this chapter is that the description of the stimuli is mirrored in the mechanisms or systems processing the stimuli. Whilst the reverse phi stimulus can be described in terms of conflicting phase and position cues, this does not mean that the mechanisms processing the reverse phi stimuli are using this distinction, or using these descriptions of stimulus characteristics to process the stimuli. Evidence to support an interpretation of conflicting mechanisms would require data where the cues are presented in isolation, and that this data could predict participant performance when cues are presented in conflict. The phase cue could be reduced or eliminated by inserting long blank inter-frame intervals (Georgeson & Harris, 1990), or by presenting stimuli inter-ocularly (Lu & Sperling, 1995a). However, these stimulus manipulations are limited in comparison to the reverse phi data whereby differences in spatial sensitivity could be due to the new stimulus manipulation, rather than the reduction of the phase cue.

Chapter 6: Residual motion abilities following lesions of striate cortex

In Chapter 6 I explored the residual motion capabilities in the scotoma of patient HS, a previously uncharacterised patient. This chapter aimed to replicate the results of Azzopardi & Hock (2011), who reported that patient GY, who has a longstanding lesion of striate cortex, could discriminate motion using phase but not position cues. For patients with V1 lesions, such as GY and our patient HS, motion perception is occurring without the 'cortical pathway' projection to V5. The residual motion perception in the scotoma is therefore the result of other pathways, and these patients provide an opportunity to study these pathways without the interference of the 'cortical pathway'. If patients

with lesions of V1 are unable to discriminate motion using position cues, this would provide a potential physiological substrate for the system processing phase cues.

Patient HS was able to discriminate motion in both the phase and position cue direction, dependent on the stimulus parameters of the conflicting cue stimuli. One possible explanation for the differences between patient HS and GY is that patient HS has a greater ability to discriminate shape, as shown by HS's success and GY's failure on the shape control task (Azzopardi & Hock, 2011). Patient HS showed a similar bias with spatial frequency and contrast as normal observers on Missing Fundamental Gratings (Georgeson & Harris, 1990), reporting motion in the position cue direction at high contrasts (when features are more detectable), and the phase cue direction at low contrasts. Across reverse phi stimuli the size and density of the stimuli determined reported motion direction, with patient HS reporting motion in the phase cue direction for reverse phi random dot kinematograms, when stimuli are small, but in the position cue direction for the reverse phi line motion illusion, when the rectangles producing the stimulus were larger and more detectable. There is a similar relationship between the size of reverse phi stimuli and reported motion direction for normal observers, where participants report motion in the phase cue direction for small reverse phi squares, and in the position cue direction for larger reverse phi squares (Chapter 5, Experiment 2).

This ability to discriminate motion using position cues, potentially caused by the ability to discriminate shape, could be due to a difference in the residual tissue between HS and GY. HS acquired his lesion in his 70s, whereas GY acquired his lesion at aged 7. Transneuronal retrograde degeneration is substantially higher for monkeys lesioned as infants (70-80%) compared to monkeys lesioned as adults (20%), (New world monkeys; Hendrickson *et al.*, 2015; Old world monkeys; Dineen & Hendrickson, 1981). This increased number of preserved RGCs, suggesting retained functional projections at later stages of the visual system, could account for the increased contrast sensitivity seen in HS compared to GY, supporting shape discrimination. Another difference seen in infant and adult lesioned monkeys is in the relative strength, in structural and functional terms, of the midbrain pathway and LGN-V5 projections with infant-lesioned macaques having a relatively weaker LGN-V5 projection than adult lesioned monkeys , (Fox *et al.*, 2020). This would be consistent with the work of Ajina *et al.* (2015), who reported that the ability of V1 lesioned patients to detect a drifting sine wave weighted by a 2D Gaussian profile relied on an intact LGN-V5 pathway. Finally, HS may have residual V1 tissue that is mediating the use of position cues when features are detectable. Further work with HS should use functional imagining whilst viewing phase and position cues to determine whether this is the case, although this still would not account for islands of residual tissue (Fendrich *et al.*, 1992), below the spatial resolution of the imaging system used.

Chapter 7: Correlations between psychophysics and physiology

Using Gabors and reverse phi stimuli, this thesis develops a 'spatial scaling hypothesis', suggesting reported motion direction in conflicting cue stimuli could be driven by two separate mechanisms, with different spatial frequency sensitivity. Physiologically, this could be caused by a difference in the retinal distribution of detectors for these proposed phase and position systems, with a greater decrease in receptor density (and so increase in linear separation) for the position system. The density distribution of these detectors can be related to spatial resolution by treating these detectors as sampling units for visual input, where the density of these sampling units would specify the limit of spatial resolution across the retina; high densities of detectors are able to take more samples of the visual input and consequentially detect higher spatial frequencies. This spatial resolution value sets a limit across the visual field, which should be reflected across the system using these early layers of the visual system as input.

These two aspects – sensitivity to spatial frequency, and distribution across the retina – could be produced if different classes of retinal ganglion cells form the initial input for systems processing phase and position cues. In Chapter 7 I constructed models of the distribution of cell density, using published data for cell density estimates for the pα and pβ RGCs, and for estimated density distribution across the magnocellular and parvocellular layers of the LGN. Using the retinal coverage estimates published by Perry & Cowey (1984), I also produced a scaled estimated for the pα RGCs and magnocellular LGN cells, to reflect the different receptive field arrangement in pα RGCs.

I compared the displacement thresholds collected in Chapter 3 with the models of cell density for the pα and pβ RGCs, and magnocellular and parvocellular layers of the LGN, using root mean squared error (RMSE) to compare goodness of fit between the models of cell density and the displacement thresholds. The aperture cue displacement thresholds were best fit by the estimated density distribution for the parvocellular layers of the LGN and the carrier cue displacement thresholds by the scaled pα RGC and magnocellular layers of the LGN.

This model is purely correlational, and as such the similarity in displacement thresholds and cell density estimates could be coincidental. The confidence in this correlation could be increased by sampling displacement thresholds at more retinal eccentricities. Currently data is normalised to the foveal threshold, restricting comparisons with the cell density models to only three eccentric locations.

This modelling approach is unlikely to be useful for investigating the motion discrimination abilities of the midbrain pathway for the time being, because there is insufficient data available on the density distributions of the RGCs which project to the superior colliculus. The limited residual visual capabilities of patients with V1 lesions suggest that obtaining motion displacement thresholds on the necessary stimuli is highly unlikely, where many of the V1 lesioned patients in the literature are unable to discriminate motion in sinusoid gratings (Azzopardi & Cowey, 2001), limiting the use of Gabor stimuli. Furthermore, patient lesions are often non-uniform with regions of greater and lesser residual vision which will confound any collected displacement thresholds, and restrict the regions of eccentricity at which displacement thresholds could be gathered.

The modelling in Chapter 7 suggests that the motion perceived from phase cues is heavily reliant on a system that receives pα RGC input, and that motion perceived from position cues is heavily reliant on a system that receives pβ RGC input. These classes of RGC are taken as the initial sampling units for these separate systems, but this is an oversimplification of the possible inputs from other classes of RGC, the interactions between cell classes at later stages in the visual system, and the interaction between these two proposed systems, as seen by the perceived motion in conflicting cue stimuli. The model should be extended to include data from other classes of RGC, such as the smooth bistratified RGCs which have similar functional responses as the pα RGCs (Crook *et al.*, 2008).

8.2 Main contributions

Psychophysical dissociation between phase and position cues

In this thesis I demonstrate a spatial frequency dissociation between carrier and aperture cues in Gabor stimuli viewed at the fovea and in the periphery. At all eccentricities motion discrimination of the aperture cue declined at lower spatial frequencies than motion discrimination to the carrier cue (Chapter 3). This dissociation in spatial frequency sensitivity could account for the reported motion direction across two different constructions of conflicting cue stimuli: in conflicting cue Gabors (Chapters 2 and 4) and reverse phi stimuli (Chapter 5), where the phase and position cues showed the same pattern of responses as the carrier and aperture cues in the Gabors.

Displacement thresholds can also be extracted for the two conflicting cue experiments (Figure 4.3, 5.9). These three sets of displacement thresholds all show substantially higher displacement thresholds in the position cue condition, which increase with eccentricity much more rapidly than

the increase in displacement thresholds with eccentricity seen in the phase or carrier condition. The similarity in the relationship between these three sets of data, despite the differences in stimulus construction, the presence of conflicting cues in the reverse phi and Opposing Motion Gabors, and the limitation of the congruent position cue in the Carrier condition Gabors, suggests that these three stimuli may be representing a dissociation in spatial frequency between two separate systems that selectively process phase or position cues.

The consistency of this dissociation across pure and conflicting cue stimuli and retinal location, suggests that this dissociation is a function of the system processing phase and position cues and not driven by differences in stimulus parameters. The reported motion direction in these conflicting cue stimuli can therefore be explained as the product of two simultaneously active motion systems with different spatial sensitivity functions. My data provide a parsimonious explanation for reversals of reported motion for conflicting cue stimuli, across both spatial frequency and eccentricity, and suggest that the psychophysical dissociation in spatial frequency is evidence supporting the presence of two separate system processing phase and position cues independently

Explanation for changes in motion perception between fovea and periphery

There are qualitative differences in the sensation of motion between fovea and periphery, and quantifiable differences, in reported motion direction, across both conflicting cue Gabor and reverse phi stimuli. These could be parsimoniously explained by the difference in this spatial frequency sensitivity of the proposed phase and position systems.

I also provide a novel explanation for the physiological substrates underlying this phenomenon. The smaller displacement thresholds, equivalent to higher spatial frequencies, for possible phase cues across eccentricity suggest that there is a higher resolution for the phase system than the position system in the periphery. I propose that the pα RGCs are the substrates for the detectors of the phase system, and the pβ RGCs for the position system. These have different density distributions across the retina. Whilst the density distribution of pβ RGCs, always exceeds that of the pα RGCs, the arrangement of receptive fields differs between the two classes of RGC. When the retinal coverage factor is applied to the pα RGCs to account for the increased overlap of the dendritic fields (Perry & Cowey, 1984), the pα RGC density exceeds that of the pβ RGC density in the periphery, with the eccentricity at which this density is exceeded depending on the scaling factor applied. The cross-over occurs peripherally, with higher retinal coverage factors producing a ‘cross-over’ point closer to fovea. This could produce regions of the periphery where the effective coverage of the retina is greater for the pα RGCs despite the greater number of the pβ RGCs. I suggest that the classic finding of Exner (Exner, 1875), that in the periphery motion discrimination exceeds the acuity for position

discrimination, is the result of different retinal coverage factors. This suggests Exner's demonstration reflected a fundamental difference in retinal ganglion dendrite structure.

My data provide a psychophysical explanation for the qualitative change in motion sensation between fovea and periphery – a change in the relative spatial frequency sensitivity functions of the phase and position systems. I propose that this is caused by a change in the density distributions of the detectors for phase and position systems and the spatial organisation of these detectors, suggesting that the pα RGCs are the substrate for the phase system, and pβ RGCs for the position system.

The position system is not reliant on a specific temporal interval

Models of motion processing have proposed two approaches for extracting motion that cannot be described by a change in the phase of the Fourier components: filter-rectify-filter models, where various image transformations are applied to make the motion signal explicit in the Fourier domain, or salience-based models, where salient features are detected, stored in a salience map representation, and their position updated with time. This second approach could be better characterised as a system specialised for identifying and localising objects, where the ability to discriminate motion from position cues may be reliant on memory, and not the primary 'aim' of the system. These two approaches have different temporal dynamics: the first approach would have similar temporal dynamics to a model extracting motion from Fourier components, where the receptive fields for the detectors are oriented in space time. This means that temporal frequency and the temporal interval of the stimuli matters. In the second approach, the model is less reliant on motion occurring within specific intervals, with the limit imposed by memory. The second model could explain both the coherent motion discrimination in the position cue direction in MFG displays with long inter-frame intervals (Georgeson & Harris, 1990), and the ability to discriminate motion between stimuli that share salience, but not any image characteristics (Lu & Sperling, 1995b). It is plausible that both systems operate simultaneously, as suggested by Lu & Sperling's (1996) 'Three systems for visual motion perception'.

Evidence supporting the presence of a position system with different temporal dynamics to those of the phase system, is presented in Chapter 2, where I replicated Georgeson & Harris (1990)'s demonstration of sustained motion discrimination in MFG stimuli with long inter-frame intervals. In Chapter 3 I compared the change in displacement thresholds with eccentricity for phase and position cues with published thresholds, using calculated E2 values. E2 values for phase cues were similar to published E2 values for other motion stimuli, but for position cues were most similar to published

thresholds for static grating acuity. This is consistent with an interpretation of the first stage of the position system as extracting features, and subsequently responding to changes in displacement.

The physiology data modelled in this thesis support this interpretation of the position system. In Chapter 7 I reported that the displacement thresholds for the position cue were better fit by the pβ than the pα RGC distribution. The pβ RGCs project to V4 via the parvocellular layers of the LGN and V1. V4 has classically been defined as the site for processing colour and form (Zeki, 1983; Schein & Desimone, 1990). This supports my suggestion that a position cues can be extracted and processed by a system where motion is inferred from the change in position of features, where information is updated in time rather than extracted from spatio-temporal coordinates, or needing to occur within a specific temporal interval. This system could operate alongside filter-rectify-filter models, with both systems active for any stimuli that contain position cues, and with this salience-model system able to additionally process stimuli not presented within specific temporal intervals, or with stimuli where rectification is insufficient to produce available motion cues.

Proposed potential physiological substrates for the phase and position systems

I compared the changes in displacement thresholds for carrier and aperture cues (Chapter 3), with models of cell density distributions, using pα and pβ RGCs, and the cells in the magnocellular and parvocellular layers of the LGN (Chapter 7). I suggested that motion discriminated from phase cues is heavily reliant on a system that receives pα RGC input, and that motion discriminated from position cues is heavily reliant on a system that receives pβ RGC input.

The inclusion of the retinal coverage factor in the model produces a relative density distribution of the pα RGCs, that exceeds the density distribution of the pβ RGCs in periphery. These results offer a mechanistic explanation for the increased reliance on carrier or phase cues in the periphery (Chapters 4 and 5), often discussed as a 'qualitative difference in the sensation of motion between fovea and periphery'. Furthermore, these results are also consistent with data suggesting that V1 lesioned patients and non-human primates are able to discriminate phase cues, but have poor shape discrimination (Chapter 6), as a small number of pα RGCs project to the superior colliculus (Crooks *et al.*, 2008), yet no pβ RGCs do.

8.3 Future Directions

This thesis demonstrates a spatial frequency dissociation between phase and position cues, and suggests that these two cues are processed by two independent systems with initially separate physiological substrates. This thesis raises both psychophysical questions and physiological questions, that should be explored further.

8.3.1. Psychophysical

Improved stimuli to investigate psychophysical dissociations

In this thesis results were limited by the lack of a 'pure' phase cue stimulus, whereby the MF3F and a Gabor with a phase shifted carrier and stationary aperture (the Carrier condition), were used as stimuli despite the confounding presence of the change in position of the grating features. The aim in producing 'pure' phase and position cue stimuli was that these could be used to explore psychophysical dissociations, and that these dissociations could be attributed to the presence or absence of the cues, rather than differences in image characteristics. Producing pure stimuli, with similar image characteristics is challenging, and often relies on stimulus constructions that would not occur naturally, e.g. interleaved gratings of different identities (Lu & Sperling, 1995b) or drifting Kanisza squares (Seifert & Hock, 2014). This fact could be used to argue that there is little evolutionary pressure for the visual system to have developed separable phase and position systems, alongside the fact that for the vast majority of stimuli phase and position cues are mathematically equivalent. Future work using 'pure' stimuli could use either of the pure 'position' cues discussed in this thesis: the 'MF1F' in Baro & Levinson (1988), or the Aperture or Aperture + Rotation Gabor stimuli, displayed in Chapters 2-4. An 8 frame pedestalled Gabor with a phase shift of the carrier could be used as a corresponding 'pure' phase cue stimulus for the Aperture condition, and replicate the experiment in Chapter 3 to find displacement thresholds across eccentricity for these now matched pure stimuli.

Further parameters probed for psychophysical dissociations

Position cues are posited to be extracted from the visual input using filter-rectify-filter and feature correlation models. To differentiate between these two models temporal frequency is a much more informative dimension than spatial frequency. Filter-rectify-filter models suggest that position cues are processed by motion energy analysis, once the visual input has undergone some transformation to render explicit information that might otherwise be invisible to the motion energy model e.g. contrast modulated gratings. These motion energy models use spatial and temporal filters to extract information, and therefore motion must occur within a specific window in time. On the contrary,

models which correlate the change in position of features with time, or update the location of a feature in a salience map, do not have a specific temporal requirement: this could explain why motion discrimination can remain high in MFG stimuli reported in the position cue direction, even with the introduction of very long IFIs (Chapter 2, experiment 2, Georgeson & Harris, 1990). Therefore exploring temporal frequency dissociations in motion stimuli may help differentiate between different models of motion perception and produce further evidence for potential physiological substrates.

Reverse phi stimuli could also be used to investigate temporal frequency dissociations, choosing spatial frequency values at the midpoint of motion reversal for each eccentricity, and mapping the parameter space for temporal frequency dissociations. Lesion studies in the LGN of macaques suggest that functions mediated by the parvocellular layers of the LGN would declines with increasing temporal frequencies, whilst functions mediated by the magnocellular layers of the LGN are particularly sensitive for high temporal frequencies (Merigan *et al.*, 1991b; Schiller *et al.*, 1990). Therefore the suggested mediation of the position cues by a system projecting to the parvocellular layers of the LGN and phase cues by a system projecting to the magnocellular layers of the LGN would suggest that motion should be reported in the phase cue direction at high temporal frequencies. Temporal frequency of the carrier and aperture cues in Gabors can be separately manipulated (Heller *et al*., 2021): demonstrating a possible temporal frequency dissociation across two forms of conflicting cue stimuli would provide further evidence for the physiological substrates for the phase and position systems suggested in this thesis. Demonstrating this dissociation across two forms of conflicting cue stimuli would suggest that this dissociation is a property of the systems processing the motion cues, rather than a dissociation driven by image characteristics.

Additional response metrics to probe psychophysical dissociations

One limitation of the methodology used in this thesis was the reliance on reported motion direction to quantify performance. If the phase and position cues are processed by separate systems, with separate physiological substrates this should be evident in other responses. One promising methodology is eye movements, where the ocular following response has been demonstrated to follow the phase cue direction in Missing Fundamental Gratings (Sheliga *et al.* 2006) and reverse phi RDKs (Masson *et al.,* 2002). If the OFR varies systematically with reported motion direction, which was not assessed in these studies, then the OFR could be evidence that separate systems process phase and position cues, with the phase system having closer physiological ties to a system involved in reflex eye movements.

The use of reported motion trajectory also obscures differences in perceptual experience. Shapiro's 'Break of the Curveball' stimulus is not only perceived as moving in an intermediate trajectory when viewed peripherally, but the nature of the percept changes: at fovea, the phase displacement of the carrier and the aperture displacement are two distinct motion percepts, whereas when viewed peripherally, the Gabor is perceived as a coherent object. This difference in motion experience could be the result of the difference in scale between the two viewing locations, where the separate aperture and carrier displacement may be visible if we scaled the aperture size and the carrier frequency in for the peripheral viewing, and potentially also adjusted for the change in contrast sensitivity of two cues. However, this percept, of separate or integrated components, could be informative in investigations of how perceived motion changes between the fovea and the periphery, and is not characterised in reported motion direction. This could be assessed through a matching procedure where participants match the perceptual properties of two Orthogonal Gabors, presented simultaneously at the fovea and periphery.

This should also be explored further in reverse phi stimuli, to determine whether stimuli presented under parameters producing overall chance performance (the midpoint on the psychometric function) are perceived as having matched relative strength, so perceived as stationary, or with ambiguous motion, or whether on individual trials motion is perceived either in the phase or position cue direction and so across a set of trials reported overall at chance. Gathering data about the perceptual experience of motion can then inform experiments to investigate how motion is integrated.

Investigate the integration of phase and position cues

For the majority of visual stimuli, the descriptions of motion in terms of the phase of the Fourier components and the change in position of features, are mathematically equivalent and signal motion in the same direction. If there are two separate systems processing phase and position cues, how is a stable motion percept produced from the outputs of these two systems? There are two different solutions: perceived motion could be produced through the integration of signals from the phase or position system, or perceived motion could occur through a 'winner takes all' mechanism, determined by the system with the stronger output. These solutions may also vary depending on the difference in cue strength, or position on the retina, as suggested by Shapiro *et al.* (2011)'s 'feature blur hypothesis'.

Motion integration could be investigated using Gabors, comparing reported distance, direction and perception as a whole or separate parts. Heller *et al.* (2011) has produced a preliminary model using orthogonally oriented Gabors where the velocity of the phase and position cue differ: extending this

model in spatial frequency and across multiple eccentric locations could be a promising avenue in investigating the integration of phase and position cues.

Motion integration could be investigated in other conflicting cue stimuli. Hammet *et al.* (1993) reported 'perceptual transparency' in the 3f4f stimulus (Figure 2.2), viewed foveally, similar to the percept of separable components in the conflicting cue Gabor stimuli. This stimulus has a simple construction, well suited for investigation across the retina, and can also be constructed as a contrast modulated version which could be used to probe whether this integration, or lack of integration, is affected when rectification is required to extract phase cue components. Finally, the reverse phi stimulus could also be used to investigate motion integration. Whilst reverse phi stimuli cannot present the phase and position cues in isolation, performance could be compared to phi stimuli, with no contrast reversal, where phase and position cues are congruent. This could explore whether there is a cost of cues being presented in conflicting directions, either through motion discrimination thresholds, measuring contrast sensitivity thresholds, or using reverse phi RDK displays and assessing motion coherence thresholds cf. Bours *et al.* (2009).

Motion integration could also be considered in terms of 'false' integration. The perpetual diamond stimulus of Flynn & Shapiro (2018) is a diamond with narrow edges that reverse in contrast, in phase with the background. Pairing the contrast reversals of different sets of the four edges can be used to produce a strong percept of the diamond drifting vertically, horizontally, or even expanding and contracting. This is perceived as a change in position, despite the diamond remaining stationary on the retina. Similarly, Orthogonal Motion gabors also have a perceived change in position on the retina, where the reported final location of the Gabor is displaced when viewed peripherally (Lisi & Cavanagh, 2015). This suggests an interaction between the phase system, where motion is processed using local detectors, and the overall perceived position of the stimulus. This provides an opportunity to determine whether the position system discussed in this thesis, in terms of the direction discrimination of position cues, is also responsible for the general localisation of stimuli. This could be investigated by manipulating stimulus parameters that affect the motion discrimination in the position system discussed in this thesis, but that do not change the actual position of the stimulus, such as including additional Gabors in the display (Boulton & Baker, 1993). This motion induced position shift is further evidence that the position system responsible for discriminating the direction of position cues is better characterised by the salience map hypothesis. If position cues were processed by a filter-rectification-filter model and then processed by the phase system, there should be no global change in position since the phase system calculates motion signals at a local scale.

Motion integration is therefore a critical avenue to investigate in determining whether there are separable motion systems, and if so, how they interact. These experiments could be paired with neuroimaging to investigate whether there is evidence for physiologically separate phase and position systems, and when these signals are compared or combined, whereby stimulus parameters could be manipulated to produce cases of conflict and cases of integration.

8.3.2 Physiological

Correlating psychophysical abilities with physiological substrates in patients with V1 lesions

Patient HS was included in this thesis as it was thought that patients with V1 lesions offer an opportunity to characterise the motion perception abilities of the residual motion mediated solely by the midbrain pathway, and to isolate motion perception without pβ RGC contributions. The growing body of evidence for additional functional pathways active alongside the midbrain pathway suggests that it is untenable to describe residual vision following striate cortex lesions as solely mediated by the midbrain pathway. Furthermore, the failure to replicate Azzopardi & Hock (2011), against a background of heterogenous visual abilities and residual tissue, suggest that this aim is unlikely to be fulfilled with a single case study.

There is a huge range in the residual motion abilities of patients with V1 lesions. Some researchers reporting that patients are unable to discriminate motion in drifting gratings or RDKs even at high contrasts and temporal frequencies (Azzopardi and Cowey, 2001), and other researchers reporting that motion discrimination is possible in the scotoma for drifting sinusoid gratings (Morland *et al.*, 2004) and RDKs (Giaschi *et al.*, 1998). The parameters that the stimuli are presented at likely plays a large role, as can be seen in the range of abilities across spatial frequency seen in the literature with sinusoid gratings at 1 cpd producing no direction discrimination (Azzopardi & Cowey, 2001), whilst 3 of 8 V1 lesioned patients were able to discriminate motion in sinusoid gratings of 0.2 cpd (Morland *et al.*, 2004). Differing temporal frequency and contrast, along with the variation in the stimuli used, adds further variation that could explain the broad range of visual abilities in the scotoma, as could the possibility of artefacts within the stimuli (e.g. Benson *et al.*, 1998).

A further complication in the cortical blindness literature is in the heterogeneity of lesions; their extent, location and density. Diagnosing lesions through structural imaging and field tests of the scotoma is not sufficient, where functional imagining has revealed activity in residual tissue for some patients with structurally defined lesions of V1 (Fendrich *et al.*, 1992; Barbur *et al.*, 1993; Morland *et al.*, 2004). There may also be systematic differences within the pathways and neural tissue processing residual vision in these patients. This has been investigated for the detection of motion

stimuli, where patients able to detect a drifting Gabor all had structural LGN-V5 projections (Ajina *et al.*, 2015), but needs to be investigated for functional differences, and for motion discrimination, where detection may rely purely on sensitivity to temporal frequency. Patients may be differentiated based on the age of lesion acquisition (macaques: Fox *et al.*, 2020), which would add to our understanding of the relative time-course of the development of midbrain and cortical pathways for vision (Bourne & Morrone, 2017). Alternatively, post-lesion visual experiences may be a driving factor in determining patients residual visual abilities , which would support the adoption of Visual Restoration Training as a therapeutic technique.

Detailed exploration of residual visual abilities in the scotoma, correlated with functional descriptions of residual tissue and active pathways, are essential for the cortical blindness literature. Future work should to move away from treating patients as a unitary class, where greater insights about system functions will be derived from more thoroughly characterising which residual functional tissue correlates with which visual abilities.

Improving the cell density models

I used published densities of retinal ganglion cells and estimates of retinal coverage factors to produce the models discussed in Chapter 7. These showed a correlation between displacement thresholds for phase and position cues across eccentricity with scaled pα and pβ RGCs. However, these density distributions are averaged across multiple retinae, are prone to deformation in the process of mounting the retina, and in some cases I had to include non-human primate estimates of RGC density or ratios between pα and pβ RGCs, rather than human estimates.

The cell density estimates could be substantially improved by imaging living human retinae, a method that is becoming increasingly possible with developments in combined Adaptive Optics and Optical Coherence Tomography (AO-OCT: Liu *et al.*, 2017). These may provide more accurate models of cell density to compare to displacement thresholds. If participant thresholds continued to show correlations with cell density distributions, but for individual models of cell densities, this would provide a more convincing correlation. Furthermore, this would enable the comparison within an individual of their displacement thresholds and different pα and pβ RGC distributions. If the correlation between phase and position displacement thresholds with the scaled pα and pβ RGC density distributions is preserved, this would be a compelling finding in labelling these as the substrates for the phase and position systems.

Alternatively, causal models could be produced in animal models, where selective lesions of the parvocellular and magnocellular layers of the LGN (cf. Schiller *et al.*, 1990) and Merigan *et al.*, 1986)

will result in transneuronal retrograde degeneration of retinal ganglion cells. If displacement thresholds correlated with cell density distributions both before and after selective lesions, this would be compelling evidence for involvement of these different cell populations as early-stage inputs to putative phase and position systems.

Combining psychophysical stimuli with neuroimaging techniques

An alternative to looking for correlations between psychophysical thresholds and physiological substrates, is to demonstrate causal differences in patterns of activity in response to viewing phase and position cues. Dension *et al.* (2014) has reported an fMRI method to localise voxels in the parvocellular and magnocellular layers in the LGN, and Liu *et al.* (2006) has reported differing time courses of activity in V1 with stimuli designed to selectively activate the magnocellular and parvocellular pathways projecting to V1. Recorded differences in the time course and regions activated by presentation of stimuli containing coherent motion solely in the phase or position cue would provide further evidence to suggest that there are separate physiological systems involved in processing phase and position cues. These experiments would rely on stimuli matched closely for spatial and temporal frequency and spatial structure: possibly a pedestalled Gabor with a carrier displacement, and a stationary pedestalled Gabor with an aperture displacement.

The correlation within individuals of displacement thresholds and cell density distributions, or topographic activity maps, would provide convincing support for the hypothesis that phase and position cues are processed by separate systems, receiving initial input from pα and pβ RGCs. This functional mapping of the early inputs could be used to further define the physiological substrates of the proposed phase and position systems. If the pα and pβ RGCs are the system inputs, what are the physiological substrates further on, and do the phase and position systems remain separate even in V5/MT. The functional mapping could be used further to determine how conflicting stimuli are processed: do they have produce simultaneous activity of the phase and position 'maps', or map depending on whether motion is perceived in the phase or position cue direction. One possible methodology is fast fMRI adaptation paradigms, which have been used to demonstrate functionally separate but spatially overlapping patterns of activity for phase and position cue stimuli in V5 (Ashida *et al.*, 2007).

8.4 Conclusion

In this thesis I demonstrated psychophysical dissociations between phase and position cues across various stimulus constructions, that suggests that these two cues are processed by two systems with independent sampling of the visual input. Evidence to support this dissociation comes from

comparing reported motion in conflicting cue stimuli, where the proposed spatial frequency dissociation of the phase and position systems can predict reversals both with changing spatial frequency, and between foveal and peripheral viewing.

I investigated the physiological substrates for these separate phase and position systems by correlating models of cell density distributions with displacement thresholds for carrier and aperture cues, and proposed that the pα and pβ RGCs form the physiological substrates for the input to these phase and position systems respectively.

This thesis focused on characterising the initial detectors of the phase and position systems, both psychophysically and in suggesting physiological substrates. Further work is needed to better characterise the spatial frequency dissociations with 'pure' phase and position cues, to investigate other psychophysical dissociations, and to explore the potential mechanisms within these proposed motion systems.

Reference List

Addams, R. (1834). An account of a peculiar optical phænomenon seen after having looked at a moving body. London Edinburgh Philos. Mag. J. Sci, 5, 373-374.

Adelson, E. H., & Bergen, J. R. (1985). Spatiotemporal energy models for the perception of motion. Journal of the Optical Society of America A, 2(2), 284-299.

Adelson, E. H., & Movshon, J. A. (1982). Phenomenal coherence of moving visual patterns. Nature, 300(5892), 523-525.

Adelson, E. H. (1982). Some new motion illusions. and some old ones, analysed in terms of their Fourier components. Investigations in Ophthalmic Visual Science. 34, 144.

Ajina, S., & Bridge, H. (2018). Blindsight relies on a functional connection between hMT+ and the lateral geniculate nucleus, not the pulvinar. Public Library of Science Biology, 16(7), e2005769.

Ajina, S., & Bridge, H. (2019). Subcortical pathways to extrastriate visual cortex underlie residual vision following bilateral damage to V1. Neuropsychologia, 128, 140-149.

Ajina, S., Jünemann, K., Sahraie, A., & Bridge, H. (2021). Increased visual sensitivity and occipital activity in patients with hemianopia following vision rehabilitation. Journal of Neuroscience, 41(28), 5994-6005.

Ajina, S., Kennard, C., Rees, G., & Bridge, H. (2015). Motion area V5/MT+ response to global motion in the absence of V1 resembles early visual cortex. Brain, 138(1), 164-178.

Ajina, S., Pestilli, F., Rokem, A., Kennard, C., & Bridge, H. (2015). Human blindsight is mediated by an intact geniculo-extrastriate pathway. Elife, 4, e08935.

Alexander, I., & Cowey, A. (2009). The cortical basis of global motion detection in blindsight. Experimental Brain Research, 192(3), 407-411.

Alexander, I., & Cowey, A. (2013). Isoluminant coloured stimuli are undetectable in blindsight even when they move. Experimental Brain Research, 225(1), 147-152.

Allman, J. M., & Kaas, J. H. (1971). Representation of the visual field in striate and adjoining cortex of the owl monkey (Aotus trivirgatus). Brain Research, 35(1), 89-106.

Anderson, D., & Burnham, K. (2004). Model selection and multi-model inference. Second. NY: Springer-Verlag, 63(2020), 10.

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. The Journal of Physiology, 442(1), 47-64.

Anstis, S. M. (1970). Phi movement as a subtraction process. Vision Research, 10(12), 1411-IN5.

Anstis, S. M. (1996). Adaptation to peripheral flicker. Vision Research, 36(21), 3479-3485.

Anstis, S.M. (2012). The furrow illusion: Peripheral motion becomes aligned with stationary contours. Journal of Vision, 12(12), 12-12.

Anstis, S. M., & Rogers, B. J. (1975). Illusory reversal of visual depth and movement during changes of contrast. Vision Research, 15(8-9), 957-IN6.

Ashida, H., Lingnau, A., Wall, M. B., & Smith, A. T. (2007). FMRI adaptation reveals separate mechanisms for first-order and second-order motion. Journal of Neurophysiology, 97(2), 1319-1325.

Azzopardi, P., & Cowey, A. (1996a). Models of ganglion cell topography in the retina of macaque monkeys and their application to sensory cortical scaling. Neuroscience, 72(3), 617-625.

Azzopardi, P., & Cowey, A. (1996b). The overrepresentation of the fovea and adjacent retina in the striate cortex and dorsal lateral geniculate nucleus of the macaque monkey. Neuroscience, 72(3), 627-639.

Azzopardi, P., & Cowey, A. (2001). Motion discrimination in cortically blind patients. Brain, 124(1), 30-46.

Azzopardi, P., & Hock, H. S. (2011). Illusory motion perception in blindsight. Proceedings of the National Academy of Sciences, 108(2), 876-881.

Azzopardi, P., Jones, K. E., & Cowey, A. (1999). Uneven mapping of magnocellular and parvocellular projections from the lateral geniculate nucleus to the striate cortex in the macaque monkey. Vision Research, 39(13), 2179-2189.

Baker Jr, C. L., & Braddick, O. J. (1985). Eccentricity-dependent scaling of the limits for short-range apparent motion perception. Vision Research, 25(6), 803-812.

Baker Jr, C. L., & Hess, R. F. (1998). Two mechanisms underlie processing of stochastic motion stimuli. Vision Research, 38(9), 1211-1222.

Baker, C. L., Hess, R. F., & Zihl, J. (1991). Residual motion perception in a "motion-blind" patient, assessed with limited-lifetime random dot stimuli. Journal of Neuroscience, 11(2), 454-461.

Barbur, J. L., Ruddock, K. H., & Waterfield, V. A. (1980). Human visual responses in the absence of the geniculo-calcarine projection. Brain: A Journal of Neurology, 103(4), 905-928.

Barbur, J. L., Watson, J. D., Frackowiak, R. S., & Zeki, S. (1993). Conscious visual perception without VI. Brain, 116(6), 1293-1302.

Barbur, J. L., Harlow, A. J., & Weiskrantz, L. (1994). Spatial and temporal response properties of residual vision in a case of hemianopia. Philosophical Transactions of the Royal Society of London. Series B: Biological Sciences, 343(1304), 157-166.

Barleben, M., Stoppel, C. M., Kaufmann, J., Merkel, C., Wecke, T., Goertler, M., Heinze, H-J., Hopf, J-H., & Schoenfeld, M. A. (2015). Neural correlates of visual motion processing without awareness in patients with striate cortex and pulvinar lesions. Human Brain Mapping, 36(4), 1585-1594.

Baro, J. A., & Levinson, E. (1988). Apparent motion can be perceived between patterns with dissimilar spatial frequencies. Vision Research, 28(12), 1311-1313.

Barton, J. J., & Sharpe, J. A. (1997). Motion direction discrimination in blind hemifields. Annals of Neurology: Official Journal of the American Neurological Association and the Child Neurology Society, 41(2), 255-264.

Basler, A. (1909). Über das Sehen von Bewegungen. Archiv für die Gesamte Physiologie des Menschen und der Tiere, 128(3), 145-176.

Benardete, E. A., & Kaplan, E. (1997). The receptive field of the primate P retinal ganglion cell, I: Linear dynamics. Visual Neuroscience, 14(1), 169-185.

Benevento, L. A., & Standage, G. P. (1983). The organization of projections of the retinorecipient and nonretinorecipient nuclei of the pretectal complex and layers of the superior colliculus to the lateral pulvinar and medial pulvinar in the macaque monkey. Journal of Comparative Neurology, 217(3), 307-336.

Benson, P. J., Guo, K., & Blakemore, C. (1998). Direction discrimination of moving gratings and plaids and coherence in dot displays without primary visual cortex (V1). European Journal of Neuroscience, 10(12), 3767-3772.

Benton, C. P., Johnston, A., & McOwan, P. W. (1997). Perception of motion direction in luminance- and contrast-defined reversed-phi motion sequences. Vision Research, 37(17), 2381-2399.

Berkley, M. A., Debruyn, B., & Orban, G. (1994). Illusory, motion, and luminance-defined contours interact in the human visual system. Vision Research, 34(2), 209-216.

Berman, R. A., & Wurtz, R. H. (2010). Functional identification of a pulvinar path from superior colliculus to cortical area MT. Journal of Neuroscience, 30(18), 6342-6354.

Best illusion of the year contest. The break of the curveball (2009). Retrieved from http://illusionoftheyear.com/2009/05/the-break-of-the-curveball/

Bex, P. J., & Baker, C. L. (1999). Motion perception over long interstimulus intervals. Perception & Psychophysics, 61(6), 1066-1074.

Blakemore, C., & Campbell, F. W. (1969). On the existence of neurones in the human visual system selectively sensitive to the orientation and size of retinal images. The Journal of Physiology, 203(1), 237-260.

Blythe, I. M., Bromley, J. M., Kennard, C., & Ruddock, K. H. (1986). Visual discrimination of target displacement remains after damage to the striate cortex in humans. Nature, 320(6063), 619-621.

Blythe, I. M., Kennard, C., & Ruddock, K. H. (1987). Residual vision in patients with retrogeniculate lesions of the visual pathways. Brain, 110(4), 887-905.

Boulton, J. C., & Baker Jr, C. L. (1993). Different parameters control motion perception above and below a critical density. Vision Research, 33(13), 1803-1811.

Bourne, J. A., & Morrone, M. C. (2017). Plasticity of visual pathways and function in the developing brain: is the pulvinar a crucial player?. Frontiers in Systems Neuroscience, 11, 3.

Bours, R. J., Kroes, M. C., & Lankheet, M. J. (2009). Sensitivity for reverse-phi motion. Vision Research, 49(1), 1-9.

Braddick, O. (1974). A short-range process in apparent motion. Vision Research, 14(7), 519-527.

Brent, P. J., Kennard, C., & Ruddock, K. H. (1994). Residual colour vision in a human hemianope: spectral responses and colour discrimination. Proceedings of the Royal Society of London. Series B: Biological Sciences, 256(1347), 219-225.

Brindley, G. S., & Lewin, W. S. (1968). The sensations produced by electrical stimulation of the visual cortex. The Journal of Physiology, 196(2), 479-493.

Britten, K. H., Shadlen, M. N., Newsome, W. T., & Movshon, J. A. (1992). The analysis of visual motion: a comparison of neuronal and psychophysical performance. Journal of Neuroscience, 12(12), 4745-4765.

Brown, S. P., He, S., & Masland, R. H. (2000). Receptive field microstructure and dendritic geometry of retinal ganglion cells. Neuron, 27(2), 371-383.

Buchner, H., Gobbelé, R., Wagner, M., Fuchs, M., Waberski, T. D., & Beckmann, R. (1997). Fast visual evoked potential input into human area V5. NeuroReport, 8(11), 2419-2422.

Campbell, F. W., Carpenter, R. H. S., & Levinson, J. Z. (1969). Visibility of aperiodic patterns compared with that of sinusoidal gratings. The Journal of Physiology, 204(2), 283-298.

Campbell, F. W., Howell, E. R., & Johnstone, J. R. (1978). A comparison of threshold and suprathreshold appearance of gratings with components in the low and high spatial frequency range. The Journal of Physiology, 284(1), 193-201.

Campbell, F. W., & Robson, J. G. (1968). Application of Fourier analysis to the visibility of gratings. The Journal of Physiology, 197(3), 551.

Cavanagh, P. (1987). Reconstructing the third dimension: Interactions between color, texture, motion, binocular disparity, and shape. Computer Vision, Graphics, and Image Processing, 37(2), 171-195.

Cavanagh, P. (1990). Pursuing moving objects with attention. In 12th Annual Conf. CSS Pod (pp. 1046-1049). Psychology Press.

Cavanagh, P. (1992). Attention-based motion perception. Science, 257(5076), 1563-1565.

Cavanagh, P., & Peter, U. T. (2019). The vector combination underlying the double-drift illusion is based on motion in world coordinates: Evidence from smooth pursuit. Journal of Vision, 19(14), 2-2.

Cavanagh, P., Arguin, M., & von Grünau, M. (1989). Interattribute apparent motion. Vision Research, 29(9), 1197-1204.

Challinor, K. L., & Mather, G. (2010). A motion-energy model predicts the direction discrimination and MAE duration of two-stroke apparent motion at high and low retinal illuminance. Vision Research, 50(12), 1109-1116.

Chatterjee, S., & Callaway, E. M. (2003). Parallel colour-opponent pathways to primary visual cortex. Nature, 426(6967), 668-671.

Chubb, C., & Sperling, G. (1988). Drift-balanced random stimuli: a general basis for studying non-Fourier motion perception. Journal of the Optical Society of America A, 5(11), 1986-2007.

Chubb, C., & Sperling, G. (1989). Two motion perception mechanisms revealed through distance-driven reversal of apparent motion. Proceedings of the National Academy of Sciences, 86(8), 2985-2989.

Chubb, C., & Sperling, G. (1991). Texture quilts: Basic tools for studying motion-from-texture. Journal of Mathematical Psychology, 35(4), 411-442.

Connolly, M., & Van Essen, D. (1984). The representation of the visual field in parvicellular and magnocellular layers of the lateral geniculate nucleus in the macaque monkey. Journal of Comparative Neurology, 226(4), 544-564.

Coren, S. (1991). Retinal mechanisms in the perception of subjective contours: the contribution of lateral inhibition. Perception, 20(2), 181-191.

Cowey, A., & Rolls, E. T. (1974). Human cortical magnification factor and its relation to visual acuity. Experimental Brain Research, 21(5), 447-454.

Cowey, A., & Stoerig, P. (1989). Projection patterns of surviving neurons in the dorsal lateral geniculate nucleus following discrete lesions of striate cortex: implications for residual vision. Experimental Brain Research, 75(3), 631-638.

Cowey, A., & Stoerig, P. (2001). Detection and discrimination of chromatic targets in hemianopic macaque monkeys and humans. European Journal of Neuroscience, 14(8), 1320-1330.

Cowey, A., Stoerig, P., & Perry, V. H. (1989). Transneuronal retrograde degeneration of retinal ganglion cells after damage to striate cortex in macaque monkeys: selective loss of Pβ cells. Neuroscience, 29(1), 65-80.

Crook, J. M., Lange-Malecki, B., Lee, B. B., & Valberg, A. (1988). Visual resolution of macaque retinal ganglion cells. The Journal of Physiology, 396(1), 205-224.

Crook, J. D., Peterson, B. B., Packer, O. S., Robinson, F. R., Gamlin, P. D., Troy, J. B., & Dacey, D. M. (2008). The smooth monostratified ganglion cell: evidence for spatial diversity in the Y-cell pathway to the lateral geniculate nucleus and superior colliculus in the macaque monkey. Journal of Neuroscience, 28(48), 12654-12671.

Crook, J. D., Peterson, B. B., Packer, O. S., Robinson, F. R., Troy, J. B., & Dacey, D. M. (2008). Y-cell receptive field and collicular projection of parasol ganglion cells in macaque monkey retina. Journal of Neuroscience, 28(44), 11277-11291.

Cumming, B. G., & Parker, A. J. (1999). Binocular neurons in V1 of awake monkeys are selective for absolute, not relative, disparity. Journal of Neuroscience, 19(13), 5602-5618.

Curcio, C. A., & Allen, K. A. (1990). Topography of ganglion cells in human retina. Journal of comparative Neurology, 300(1), 5-25.

Curcio, C. A., Sloan Jr, K. R., Packer, O., Hendrickson, A. E., & Kalina, R. E. (1987). Distribution of cones in human and monkey retina: individual variability and radial asymmetry. Science, 236(4801), 579-582.

Cynader, M., & Berman, N. (1972). Receptive-field organization of monkey superior colliculus. Journal of Neurophysiology, 35(2), 187-201.

Dacey, D. M. (1993). The mosaic of midget ganglion cells in the human retina. Journal of Neuroscience, 13(12), 5334-5355.

Dacey, D. M. (2000). Parallel pathways for spectral coding in primate retina. Annual review of Neuroscience, 23(1), 743-775.

Dacey, D. M. (2004). 20 origins of perception: Retinal ganglion cell diversity and the creation of parallel visual pathways. The Cognitive Neurosciences. MIT Press, 281

Dacey, D. M. (2007). Physiology, morphology and spatial densities of identified ganglion cell types in primate retina. In Ciba Foundation Symposium 184-Higher-Order Processing in the Visual System: Higher-Order Processing in the Visual System: Ciba Foundation Symposium ,184,12-34. Chichester, UK: John Wiley & Sons, Ltd.

Dacey, D. M., & Lee, B. B. (1994). The 'blue-on' opponent pathway in primate retina originates from a distinct bistratified ganglion cell type. Nature, 367(6465), 731-735.

Dacey, D. M., & Petersen, M. R. (1992). Dendritic field size and morphology of midget and parasol ganglion cells of the human retina. Proceedings of the National Academy of Sciences, 89(20), 9666-9670.

Das, A., Tadin, D., & Huxlin, K. R. (2014). Beyond blindsight: properties of visual relearning in cortically blind fields. Journal of Neuroscience, 34(35), 11652-11664.

Daugman, J. G. (1985). Uncertainty relation for resolution in space, spatial frequency, and orientation optimized by two-dimensional visual cortical filters. Journal of the Optical Society Of America A, 2(7), 1160-1169.

DeAngelis, G. C., Ohzawa, I., & Freeman, R. D. (1991). Depth is encoded in the visual cortex by a specialized receptive field structure. Nature, 352(6331), 156-159.

de Graaf, T. A., & Sack, A. T. (2014). Using brain stimulation to disentangle neural correlates of conscious vision. Frontiers in Psychology, 5, 1019.

de Monasterio, F. M., Schein, S. J., & McCrane, E. P. (1981). Staining of blue-sensitive cones of the macaque retina by a fluorescent dye. Science, 213(4513), 1278-1281.

De Valois, R. L., & De Valois, K. K. (1991). Vernier acuity with stationary moving Gabors. Vision Research, 31(9), 1619-1626.

Denison, R. N., Vu, A. T., Yacoub, E., Feinberg, D. A., & Silver, M. A. (2014). Functional mapping of the magnocellular and parvocellular subdivisions of human LGN. Neuroimage, 102, 358-369.

Derrington, A. M., & Lennie, P. (1984). Spatial and temporal contrast sensitivities of neurones in lateral geniculate nucleus of macaque. The Journal of Physiology, 357(1), 219-240.

Derrington, A. M., & Ukkonen, O. I. (1999). Second-order motion discrimination by feature-tracking. Vision Research, 39(8), 1465-1475.

Dineen, J. T., & Hendrickson, A. E. (1981). Age correlated differences in the amount of retinal degeneration after striate cortex lesions in monkeys. Investigative Ophthalmology & Visual Science, 21(5), 749-752.

Drasdo, N. (1977). The neural representation of visual space. Nature, 266(5602), 554-556.

Drasdo, N. (1989). Receptive field densities of the ganglion cells of the human retina. Vision Research, 29(8), 985-988.

Drasdo, N., & Fowler, C. W. (1974). Non-linear projection of the retinal image in a wide-angle schematic eye. The British Journal of Ophthalmology, 58(8), 709.

Drasdo, N., & Thompson, C. M. (1989). Do visibility and colour recognition isopters relate to the distribution of Pα and Pβ ganglion cells of the human retina?. Ophthalmic and Physiological Optics, 9(4), 447-450.

Drasdo, N., Millican, C. L., Katholi, C. R., & Curcio, C. A. (2007). The length of Henle fibers in the human retina and a model of ganglion receptive field density in the visual field. Vision Research, 47(22), 2901-2911.

Dresp, B., & Bonnet, C. (1991). Psychophysical evidence for low-level processing of illusory contours and surfaces in the Kanizsa square. Vision Research, 31(10), 1813-1817.

DuBois, R. M., & Cohen, M. S. (2000). Spatiotopic organization in human superior colliculus observed with fMRI. Neuroimage, 12(1), 63-70.

Edwards, M., & Nishida, S. Y. (2004). Contrast-reversing global-motion stimuli reveal local interactions between first-and second-order motion signals. Vision Research, 44(16), 1941-1950.

Exner, S. (1875). Über das Sehen von Bewegungen und Theorie des zusammengesetzten Auges. aus der kk Hof-und Staatsdruckerei.

Exner S. (1886) "Ein Versuch liber die Netzhaut Peripherie als Organ der Wahrnehmung von Bewegungen" Pflugers Archiv fur die Gesamte Physiologie des Menschen und der Tiere 38 217-21

Fendrich, R., Wessinger, C. M., & Gazzaniga, M. S. (1992). Residual vision in a scotoma: implications for blindsight. Science, 258(5087), 1489-1491.

Ferrera, V. P., Nealey, T. A., & Maunsell, J. H. (1992). Mixed parvocellular and magnocellular geniculate signals in visual area V4. Nature, 358(6389), 756-758.

Ferrera, V. P., Nealey, T. A., & Maunsell, J. H. (1994). Responses in macaque visual area V4 following inactivation of the parvocellular and magnocellular LGN pathways. Journal of Neuroscience, 14(4), 2080-2088.

Ffytche, D. H., Guy, C. N., & Zeki, S. (1995). The parallel visual motion inputs into areas V1 and V5 of human cerebral cortex. Brain, 118(6), 1375-1394.

Finlay, D. (1982). Motion perception in the peripheral visual field. Perception, 11(4), 457-462.

Flynn, O. J., & Shapiro, A. G. (2018). The perpetual diamond: Contrast reversals along thin edges create the appearance of motion in objects. i-Perception, 9(6), 2041669518815708.

Foley, J. M., Varadharajan, S., Koh, C. C., & Farias, M. C. (2007). Detection of Gabor patterns of different sizes, shapes, phases and eccentricities. Vision Research, 47(1), 85-107

Foster, D. H., Gravano, S., & Tomoszek, A. (1989). Acuity for fine-grain motion and for two-dot spacing as a function of retinal eccentricity: Differences in specialization of the central and peripheral retina. Vision Research, 29(8), 1017-1031.

Fox, D. M., Goodale, M. A., & Bourne, J. A. (2020). The age-dependent neural substrates of blindsight. Trends in Neurosciences, 43(4), 242-252.

Foxe, J. J., Strugstad, E. C., Sehatpour, P., Molholm, S., Pasieka, W., Schroeder, C. E., & McCourt, M. E. (2008). Parvocellular and magnocellular contributions to the initial generators of the visual evoked potential: high-density electrical mapping of the “C1” component. Brain Topography, 21(1), 11-21.

Gabor, D. (1946). Theory of communication. Journal of the Institute of Electrical Engineers, 93, 429 - 457.

Georgeson, M. A., & Harris, M. G. (1990). The temporal range of motion sensing and motion perception. Vision Research, 30(4), 615-619.

Georgeson, M. A., & Shackleton, T. M. (1989). Monocular motion sensing, binocular motion perception. Vision Research, 29(11), 1511-1523.

Giaschi, D., Jan, J. E., Bjornson, B., Young, S. A., Tata, M., Lyons, C. J., & Wong, P. K. (2003). Conscious visual abilities in a patient with early bilateral occipital damage. Developmental Medicine and Child Neurology, 45(11), 772-781.

Girard, P., Salin, P. A., & Bullier, J. (1992). Response selectivity of neurons in area MT of the macaque monkey during reversible inactivation of area V1. Journal of Neurophysiology, 67(6), 1437-1446.

Goebel, R., Muckli, L., Zanella, F. E., Singer, W., & Stoerig, P. (2001). Sustained extrastriate cortical activation without visual awareness revealed by fMRI studies of hemianopic patients. Vision Research, 41(10-11), 1459-1474.

Goldberg, M. E., & Wurtz, R. H. (1972). Activity of superior colliculus in behaving monkey. I. Visual receptive fields of single neurons. Journal of Neurophysiology, 35(4), 542-559.

Gorea, A. (1995). Spatiotemporal characterization of a Fourier and non-Fourier motion system. Vision Research, 35(7), 907-914.

Gross, C. G., Moore, T., & Rodman, H. R. (2004). Visually guided behaviour after V1 lesions in young and adult monkeys and its relation to blindsight in humans. Progress in Brain Research, 144, 279-294.

Gross, C. G. (1991). Contribution of striate cortex and the superior colliculus to visual function in area MT, the superior temporal polysensory area and inferior temporal cortex. Neuropsychologia, 29(6), 497-515.

Guo, K., Benson, P. J., & Blakemore, C. (1998). Residual motion discrimination using colour information without primary visual cortex. NeuroReport, 9(9), 2103-2107.

Gurnsey, R., & Biard, M. (2012). Eccentricity dependence of the curveball illusion. Canadian Journal of Experimental Psychology, 66(2), 144.

Hammett, S. T., Ledgeway, T., & Smith, A. T. (1993). Transparent motion from feature-and luminance-based processes. Vision Research, 33(8), 1119-1122.

Hartline, H. K. (1949). Inhibition of activity of visual receptors by illuminating nearby retinal areas in the Limulus eye. In Federation Proceedings, 8, (1), 69.

Heller, N. H., Cavanagh, P., & Peter, U. T. (2021). Temporal integration window of the double-drift illusion: an immediate effect of the internal motion. Journal of Vision, 21(9), 2642-2642.

Helmholtz, H. V. (1860). Theorie der Luftschwingungen in Röhren mit offenen Enden. Journal für die Reine und Angewandte Mathematik, 57, 1–72

Hendrickson, A., Warner, C. E., Possin, D., Huang, J., Kwan, W. C., & Bourne, J. A. (2015). Retrograde transneuronal degeneration in the retina and lateral geniculate nucleus of the V1-lesioned marmoset monkey. Brain Structure and Function, 220, 351-360.

Hikosaka, O., Miyauchi, S., & Shimojo, S. (1993). Focal visual attention produces illusory temporal order and motion sensation. Vision Research, 33(9), 1219-1240.

Holmes, G., & Lister, W. T. (1916). Disturbances of vision from cerebral lesions, with special reference to the cortical representation of the macula. Brain, 39(1-2), 34-73.

Howard, I. P., & Rogers, B. J. (1995). Binocular vision and stereopsis. Oxford University Press, USA.

Hubel, D. H., & Wiesel, T. N. (1962). Receptive fields, binocular interaction and functional architecture in the cat's visual cortex. The Journal of Physiology, 160(1), 106.

Hubel, D. H., & Wiesel, T. N. (1965). Receptive fields and functional architecture in two non-striate visual areas (18 and 19) of the cat. Journal of Neurophysiology, 28(2), 229-289.

Hubel, D. H., & Wiesel, T. N. (1968). Receptive fields and functional architecture of monkey striate cortex. The Journal of Physiology, 195(1), 215-243.

Hughes, A. E. (2018). Dissociation between perception and smooth pursuit eye movements in speed judgments of moving Gabor targets. Journal of Vision, 18(4), 4-4.

Humphrey, N. K. (1974). Vision in a monkey without striate cortex: a case study. Perception, 3(3), 241-255.

Huxlin, K. R., Martin, T., Kelly, K., Riley, M., Friedman, D. I., Burgin, W. S., & Hayhoe, M. (2009). Perceptual relearning of complex visual motion after V1 damage in humans. Journal of Neuroscience, 29(13), 3981-3991.

Itti, L., & Koch, C. (2000). A saliency-based search mechanism for overt and covert shifts of visual attention. Vision Research, 40(10-12), 1489-1506.

Johnston, A., McOwan, P. W., & Buxton, H. (1992). A computational model of the analysis of some first-order and second-order motion patterns by simple and complex cells. Proceedings of the Royal Society of London. Series B: Biological Sciences, 250(1329), 297-306.

Julesz, B. (1971). Foundations of cyclopean perception. Chicago.

Jones, J. P., & Palmer, L. A. (1987). An evaluation of the two-dimensional Gabor filter model of simple receptive fields in cat striate cortex. Journal of Neurophysiology, 58(6), 1233-1258.

Kanizsa, G. (1955). Margini quasi-percettivi in campi con stimulazione omogenea. Rivista di Psicologia, 49, 7-30.

Kaplan, E., & Benardete, E. (2001). The dynamics of primate retinal ganglion cells. Progress in Brain Research, 134, 17-34.

Kaplan, E., & Shapley, R. M. (1982). X and Y cells in the lateral geniculate nucleus of macaque monkeys. The Journal of Physiology, 330(1), 125-143.

Kelly, D. H. (1979). Motion and vision. II. Stabilized spatio-temporal threshold surface. Journal of the Optical Society Of America, 69(10), 1340-1349.

Kinoshita, M., Kato, R., Isa, K., Kobayashi, K., Kobayashi, K., Onoe, H., & Isa, T. (2019). Dissecting the circuit for blindsight to reveal the critical role of pulvinar and superior colliculus. Nature Communications, 10(1), 1-10.

Koch, C., Poggio, T., & Torre, V. (1982). Retinal ganglion cells: a functional interpretation of dendritic morphology. Philosophical Transactions of the Royal Society of London. B, Biological Sciences, 298(1090), 227-263.

Koch, C., & Ullman, S. (1987). Shifts in selective visual attention: towards the underlying neural circuitry. Matters of intelligence: Conceptual structures in cognitive neuroscience, 115-141.

Krauzlis, R. J., Lovejoy, L. P., & Zénon, A. (2013). Superior colliculus and visual spatial attention. Annual review of Neuroscience, 36, 165-182.

Lane, R. H., Allman, J. M., Kaas, J. H., & Miezin, F. M. (1973). The visuotopic organization of the superior colliculus of the owl monkey (Aotus trivirgatus) and the bush baby (Galago senegalensis). Brain Research, 60(2), 335-349.

Laycock, R., Crewther, D. P., Fitzgerald, P. B., & Crewther, S. G. (2007). Evidence for fast signals and later processing in human V1/V2 and V5/MT+: a TMS study of motion perception. Journal of Neurophysiology, 98(3), 1253-1262.

Le Gros Clark, W. E. (1941). Observations on the association fibre system of the visual cortex and the central representation of the retina. Journal of Anatomy, 75(Pt 2), 225.

Ledgeway, T. & Hess, R.F. (2000). The properties of the motion-detecting mechanisms mediating perceived direction in stochastic displays. Vision Research, 40, 3585-3597.

Ledgeway, T., & Smith, A. T. (1994). Evidence for separate motion-detecting mechanisms for first- and second-order motion in human vision. Vision Research, 34(20), 2727-2740.

Leventhal, A. G., Rodieck, R. W., & Dreher, B. (1981). Retinal ganglion cell classes in the Old World monkey: morphology and central projections. Science, 213(4512), 1139-1142.

Levi, D. M., Klein, S. A., & Aitsebaomo, P. (1984). Detection and discrimination of the direction of motion in central and peripheral vision of normal and amblyopic observers. Vision Research, 24(8), 789-800.

Levi, D. M., Klein, S. A., & Aitsebaomo, A. P. (1985). Vernier acuity, crowding and cortical magnification. Vision Research, 25(7), 963-977.

Levitt, J. B., Schumer, R. A., Sherman, S. M., Spear, P. D., & Movshon, J. A. (2001). Visual response properties of neurons in the LGN of normally reared and visually deprived macaque monkeys. Journal of Neurophysiology, 85(5), 2111-2129.

Li, Y., Wang, S., Tian, Q., & Ding, X. (2015). A survey of recent advances in visual feature detection. Neurocomputing, 149, 736-751.

Lisi, M., & Cavanagh, P. (2015). Dissociation between the perceptual and saccadic localization of moving objects. Current Biology, 25(19), 2535-2540.

Liu, C. S., Bryan, R. N., Miki, A., Woo, J. H., Liu, G. T., & Elliott, M. A. (2006). Magnocellular and parvocellular visual pathways have different blood oxygen level–dependent signal time courses in human primary visual cortex. American Journal of Neuroradiology, 27(8), 1628-1634.

Liu, Z., Kurokawa, K., Zhang, F., Lee, J. J., & Miller, D. T. (2017). Imaging and quantifying ganglion cells and other transparent neurons in the living human retina. Proceedings of the National Academy of Sciences, 114(48), 12803-12808.

Lu, Z. L., & Sperling, G. (1995b). Attention-generated apparent motion. Nature, 377(6546), 237-239.

Lu, Z. L., & Sperling, G. (1995a). The functional architecture of human visual motion perception. Vision Research, 35(19), 2697-2722.

Lu, Z. L., & Sperling, G. (1996). Three systems for visual motion perception. Current Directions in Psychological Science, 5(2), 44-53.

Lu, Z. L., & Sperling, G. (1999). Second-order reversed phi. Perception & Psychophysics, 61(6), 1075-1088.

Lu, Z. L., & Sperling, G. (2001). Three-systems theory of human visual motion perception: review and update. Journal of the Optical Society Of America A, 18(9), 2331-2370.

Lu, Z. L., Lesmes, L. A., & Sperling, G. (1999). Perceptual motion standstill in rapidly moving chromatic displays. Proceedings of the National Academy of Sciences, 96(26), 15374-15379.

Lysakowski, A., Standage, G. P., & Benevento, L. A. (1988). An investigation of collateral projections of the dorsal lateral geniculate nucleus and other subcortical structures to cortical areas V1 and V4 in the macaque monkey: a double label retrograde tracer study. Experimental Brain Research, 69(3), 651-661.

Magnussen, S., & Mathiesen, T. (1989). Detection of moving and stationary gratings in the absence of striate cortex. Neuropsychologia, 27(5), 725-728.

Malpeli, J. G. (1999). Reversible inactivation of subcortical sites by drug injection. Journal of Neuroscience Methods, 86(2), 119-128.

Malpeli, J. G., & Baker, F. H. (1975). The representation of the visual field in the lateral geniculate nucleus of Macaca mulatta. Journal of Comparative Neurology, 161(4), 569-594.

Malpeli, J. G., Lee, D., & Baker, F. H. (1996). Laminar and retinotopic organization of the macaque lateral geniculate nucleus: magnocellular and parvocellular magnification functions. Journal of Comparative Neurology, 375(3), 363-377.

Malpeli, J. G., & Schiller, P. H. (1977). Functional differentiation of the laminae of the lateral geniculate nucleus of the rhesus monkey. In Society Neuroscience Abstract, 3, 568

Malpeli, J. G., Schiller, P. H., & Colby, C. L. (1981). Response properties of single cells in monkey striate cortex during reversible inactivation of individual lateral geniculate laminae. Journal of Neurophysiology, 46(5), 1102-1119.

Masson, G. S., Yang, D. S., & Miles, F. A. (2002). Reversed short-latency ocular following. Vision Research, 42(17), 2081-2087.

Maruya, K., Mugishima, Y., & Sato, T. (2003). Reversed-phi perception with motion-defined motion stimuli. Vision Research, 43(24), 2517-2526.

Marr, D., & Ullman, S. (1981). Directional selectivity and its use in early visual processing. Proceedings of the Royal Society of London. Series B. Biological Sciences, 211(1183), 151-180.

Marr, D., Ullman, S., & Poggio, T. (1979). Bandpass channels, zero-crossings, and early visual information processing. Journal of the Optical Society Of America, 69(6), 914-916.

Marrocco, R. T., & Li, R. H. (1977). Monkey superior colliculus: properties of single cells and their afferent inputs. Journal of Neurophysiology, 40(4), 844-860.

Mather, G., Cavanagh, P., & Anstis, S. M. (1985). A moving display which opposes short-range and long-range signals. Perception, 14(2), 163-166.

Maunsell, J. H., Nealey, T. A., & DePriest, D. D. (1990). Magnocellular and parvocellular contributions to responses in the middle temporal visual area (MT) of the macaque monkey. Journal of Neuroscience, 10(10), 3323-3334.

Maunsell, J. H., & van Essen, D. C. (1983). The connections of the middle temporal visual area (MT) and their relationship to a cortical hierarchy in the macaque monkey. Journal of Neuroscience, 3(12), 2563-2586.

Mazer, J. A., & Gallant, J. L. (2003). Goal-related activity in V4 during free viewing visual search: Evidence for a ventral stream visual salience map. Neuron, 40(6), 1241-1250.

McAdams, C. J., & Maunsell, J. H. (2000). Attention to both space and feature modulates neuronal responses in macaque area V4. Journal of Neurophysiology, 83(3), 1751-1755.

McKeefry, D. J., Watson, J. D. G., Frackowiak, R. S. J., Fong, K., & Zeki, S. (1997). The activity in human areas V1/V2, V3, and V5 during the perception of coherent and incoherent motion. Neuroimage, 5(1), 1-12.

Meiusi, R. S., Lavoie, J. D., & Summers, C. G. (1993). The effect of grating orientation on resolution acuity in patients with nystagmus. Journal of Pediatric Ophthalmology & Strabismus, 30(4), 259-261.

Merigan, WH, Byrne, CE, & Maunsell, JH. (1991). Does primate motion perception depend on the magnocellular pathway? The Journal of Neuroscience, 11(11), 3422-3429.

Merigan, W. H., & Eskin, T. A. (1986). Spatio-temporal vision of macaques with severe loss of Pβ retinal ganglion cells. Vision Research, 26(11), 1751-1761.

Merigan, W. H., Katz, L. M., & Maunsell, J. H. (1991). The effects of parvocellular lateral geniculate lesions on the acuity and contrast sensitivity of macaque monkeys. Journal of Neuroscience, 11(4), 994-1001.

Mikellidou, K., Frijia, F., Montanaro, D., Greco, V., Burr, D. C., & Morrone, M. C. (2018). Cortical BOLD responses to moderate-and high-speed motion in the human visual cortex. Scientific Reports, 8(1), 1-12.

Mohler, C. W., & Wurtz, R. H. (1977). Role of striate cortex and superior colliculus in visual guidance of saccadic eye movements in monkeys. Journal of Neurophysiology, 40(1), 74-94.

Morland, A. B., Lê, S., Carroll, E., Hoffmann, M. B., & Pambakian, A. (2004). The role of spared calcarine cortex and lateral occipital cortex in the responses of human hemianopes to visual motion. Journal of Cognitive Neuroscience, 16(2), 204-218.

Movshon, J. A. (1986). The analysis of moving patterns. Pattern Recognition Mechanisms, 117-151.

Movshon, J. A., Kiorpes, L., Hawken, M. J., & Cavanaugh, J. R. (2005). Functional maturation of the macaque's lateral geniculate nucleus. Journal of Neuroscience, 25(10), 2712-2722.

Movshon, J. A., & Newsome, W. T. (1996). Visual response properties of striate cortical neurons projecting to area MT in macaque monkeys. Journal of Neuroscience, 16(23), 7733-7741.

Murray, I., MacCana, F., & Kulikowski, J. J. (1983). Contribution of two movement detecting mechanisms to central and peripheral vision. Vision Research, 23(2), 151-159.

Mussap, A. J., & Crassini, B. (1993). Barber-pole illusions and plaids: the influence of aperture shape on motion perception. Perception, 22(10), 1155-1174.

Nakayama, K., & Tyler, C. W. (1981). Psychophysical isolation of movement sensitivity by removal of familiar position cues. Vision Research, 21(4), 427-433.

Nassi, J. J., & Callaway, E. M. (2006). Multiple circuits relaying primate parallel visual pathways to the middle temporal area. Journal of Neuroscience, 26(49), 12789-12798.

Newsome, W. T., & Pare, E. B. (1988). A selective impairment of motion perception following lesions of the middle temporal visual area (MT). Journal of Neuroscience, 8(6), 2201-2211.

Ohzawa, I., DeAngelis, G. C., & Freeman, R. D. (1990). Stereoscopic depth discrimination in the visual cortex: neurons ideally suited as disparity detectors. Science, 249(4972), 1037-1041.

Olman, C. A., Harel, N., Feinberg, D. A., He, S., Zhang, P., Ugurbil, K., & Yacoub, E. (2012). Layer-specific fMRI reflects different neuronal computations at different depths in human V1. Public Library of Science One, 7(3), e32536.

Oppel, O. (1967). Untersuchungen über die Verteilung und Zahl der retinalen Ganglienzellen beim Menschen. Albrecht von Graefes Archiv für Klinische und Experimentelle Ophthalmologie, 172(1), 1-22.

Orban, G. A., Van Essen, D., & Vanduffel, W. (2004). Comparative mapping of higher visual areas in monkeys and humans. Trends in CCognitive Sciences, 8(7), 315-324.

Overgaard, M., Rote, J., Mouridsen, K., & Ramsøy, T. Z. (2006). Is conscious perception gradual or dichotomous? A comparison of report methodologies during a visual task. Consciousness and Cognition, 15(4), 700-708.

Park, S. J., Shin, J. K., & Lee, M. (2002). Biologically inspired saliency map model for bottom-up visual attention. In International Workshop on Biologically Motivated Computer Vision, 418-426, Springer, Berlin, Heidelberg.

Parsons, J. H. (1927). An introduction to the theory of perception. Cambridge University Press.

Pasupathy, A., Popovkina, D. V., & Kim, T. (2020). Visual functions of primate area V4. Annual review of vision Science, 6, 363-385.

Pedersini, C. A., Lingnau, A., Cardobi, N., Sanchez-Lopez, J., Savazzi, S., & Marzi, C. A. (2020). Neural bases of visual processing of moving and stationary stimuli presented to the blind hemifield of hemianopic patients. Neuropsychologia, 141, 107430.

Perenin, M. T. (1991). Discrimination of motion direction in perimetrically blind fields. Neuroreport: An International Journal for the Rapid Communication of Research in Neuroscience.

Perry, V. H., & Cowey, A. (1984). Retinal ganglion cells that project to the superior colliculus and pretectum in the macaque monkey. Neuroscience, 12(4), 1125-1137.

Perry, V. H., & Cowey, A. (1985). The ganglion cell and cone distributions in the monkey's retina: implications for central magnification factors. Vision Research, 25(12), 1795-1810.

Perry, V. H., Oehler, R., & Cowey, A. (1984). Retinal ganglion cells that project to the dorsal lateral geniculate nucleus in the macaque monkey. Neuroscience, 12(4), 1101-1123.

Pierrot-Deseilligny, C., Rosa, A., Masmoudi, K., Rivaud, S., & Gaymard, B. (1991). Saccade deficits after a unilateral lesion affecting the superior colliculus. Journal of Neurology, Neurosurgery & Psychiatry, 54(12), 1106-1109.

Pinheiro, J., Bates, D., Debroy, S., & Sarkar, D. (2013). Nlme: nonlinear mixed-effects models. R Package3, 1.

Pizlo, Z., Kim, S., Talavage, T., Pizlo, F., & Steinman, R. (2008). Neural substrate of the perception of phi (pure) movement. Journal of Vision, 8(6), 128-128.

Pollen, D. A., & Ronner, S. F. (1981). Phase relationships between adjacent simple cells in the visual cortex. Science, 212(4501), 1409-1411.

Polyak, S. (1941). The retina: The anatomy and the histology of the retina in man, ape, and monkey, including the consideration of visual functions, the history of physiological optics, and the histological laboratory technique. Chicago: University of Chicago Press.

Polyak, S., & Klüver, H. (1957). The vertebrate visual system: Its origin, structure, and function and its manifestations in disease with an analysis of its role in the life of animals and in the origin of man, preceded by a historical review of investigations of the eye, and of the visual pathways and centers of the brain. Science, 127(3292), 244-244.

Poston, T., & Stewart, I. (1978). Nonlinear modelling of multistable perception. Behavioural Science, 23(4), 318-334.

Protti, D. A., Di Marco, S., Huang, J. Y., Vonhoff, C. R., Nguyen, V., & Solomon, S. G. (2014). Inner retinal inhibition shapes the receptive field of retinal ganglion cells in primate. The Journal of Physiology, 592(1), 49-65.

Raiguel, S. E., Lagae, L., Gulyàs, B., & Orban, G. A. (1989). Response latencies of visual cells in macaque areas V1, V2 and V5. Brain Research, 493(1), 155-159.

Ramachandran, V. S., & Anstis, S. M. (1990). Illusory displacement of equiluminous kinetic edges. Perception, 19(5), 611-616.

Raposo, N., Cauquil, A. S., Albucher, J. F., Acket, B., Celebrini, S., Pariente, J., Cognard, C., Loubinoux, I., & Chollet, F. (2011). Poststroke conscious visual deficit: clinical course and changes in cerebral activations. Neurorehabilitation and Neural Repair, 25(8), 703-710.

Reichardt, W. (1957). Autokorrelations-auswertung als funktionsprinzip des zentralnervensystems. Zeitschrift für Naturforschung B, 12(7), 448-457.

Riddoch, G. (1917). Dissociation of visual perceptions due to occipital injuries, with especial reference to appreciation of movement. Brain, 40(1), 15-57.

Rider, A. T., McOwan, P. W., & Johnston, A. (2014). Asymmetric global motion integration in drifting Gabor arrays. Journal of Vision, 14(8), 18-18.

Robinson, D. L., & McClurkin, J. W. (1989). The visual superior colliculus and pulvinar. Reviews of oculomotor Research, 3, 337-360.

Rodieck, R., & Watanabe, M. (1993). Survey of the morphology of macaque retinal ganglion cells that project to the pretectum, superior colliculus, and parvicellular laminae of the lateral geniculate nucleus. Journal of Comparative Neurology, 338(2), 289-303.

Rodman, H. R., Gross, C. G., & Albright, T. D. (1989). Afferent basis of visual response properties in area MT of the macaque. I. Effects of striate cortex removal. Journal of Neuroscience, 9(6), 2033-2050.

Rodman, H. R., Gross, C. G., & Albright, T. D. (1990). Afferent basis of visual response properties in area MT of the macaque. II. Effects of superior colliculus removal. Journal of Neuroscience, 10(4), 1154-1164.

Rosa, M. G., Tweedale, R., & Elston, G. N. (2000). Visual responses of neurons in the middle temporal area of new world monkeys after lesions of striate cortex. Journal of Neuroscience, 20(14), 5552-5563.

Ross, W. D. (1931). The Works of Aristotle (Volume III): Parva Naturalia.

Rossi, E. A., & Roorda, A. (2010). The relationship between visual resolution and cone spacing in the human fovea. Nature Neuroscience, 13(2), 156-157.

Rovamo, J., & Virsu, V. (1979). An estimation and application of the human cortical magnification factor. Experimental Brain Research, 37(3), 495-510.

Rovamo, J., Raninen, A., & Virsu, V. (1985). The role of retinal ganglion cell density and receptive-field size in photopic perimetry. In Sixth International Visual Field Symposium (125-130). Springer, Dordrecht.

Roy, S., Jayakumar, J., Martin, P. R., Dreher, B., Saalmann, Y. B., Hu, D., & Vidyasagar, T. R. (2009). Segregation of short-wavelength-sensitive (S) cone signals in the macaque dorsal lateral geniculate nucleus. European Journal of Neuroscience, 30(8), 1517-1526.

Ruppertsberg, A. I., Wuerger, S. M., & Bertamini, M. (2003). The chromatic input to global motion perception. Visual Neuroscience, 20(4), 421-428.

Sahraie, A., Trevethan, C. T., & MacLeod, M. J. (2008). Temporal properties of spatial channel of processing in hemianopia. Neuropsychologia, 46(3), 879-885.

Salin, P. A., & Bullier, J. (1995). Corticocortical connections in the visual system: structure and function. Physiological Reviews, 75(1), 107-154.

Salzman, C. D., Murasugi, C. M., Britten, K. H., & Newsome, W. T. (1992). Microstimulation in visual area MT: effects on direction discrimination performance. Journal of Neuroscience, 12(6), 2331-2355.

Sawides, L., de Castro, A., & Burns, S. A. (2017). The organization of the cone photoreceptor mosaic measured in the living human retina. Vision Research, 132, 34-44.

Schein, S. J., & de Monasterio, F. M. (1987). Mapping of retinal and geniculate neurons onto striate cortex of macaque. Journal of Neuroscience, 7(4), 996-1009.

Schein, S. J., & Desimone, R. (1990). Spectral properties of V4 neurons in the macaque. Journal of Neuroscience, 10(10), 3369-3389.

Schein, S. J. (1988). Anatomy of macaque fovea and spatial densities of neurons in foveal representation. Journal of Comparative Neurology, 269(4), 479-505.

Schiller, P. H., & Koerner, F. (1971). Discharge characteristics of single units in superior colliculus of the alert rhesus monkey. Journal of Neurophysiology, 34(5), 920-936.

Schiller, P. H., & Malpeli, J. G. (1977). Properties and tectal projections of monkey retinal ganglion cells. Journal of Neurophysiology, 40(2), 428-445.

Schiller, P. H., Sandell, J. H., & Maunsell, J. H. (1987). The effect of frontal eye field and superior colliculus lesions on saccadic latencies in the rhesus monkey. Journal of Neurophysiology, 57(4), 1033-1049.

Schiller, P. H., Logothetis, N. K., & Charles, E. R. (1990). Functions of the colour-opponent and broad-band channels of the visual system. Nature, 343(6253), 68-70.

Schiller, P. H., & Logothetis, N. K. (1990). The color-opponent and broad-band channels of the primate visual system. Trends in Neurosciences, 13(10), 392-398.

Schiller, P. H., Logothetis, N. K., & Charles, E. R. (1990). Role of the color-opponent and broad-band channels in vision. Visual Neuroscience, 5(4), 321-346.

Schmid, M. C., Mrowka, S. W., Turchi, J., Saunders, R. C., Wilke, M., Peters, A. J., Ye, F.Q., & Leopold, D. A. (2010). Blindsight depends on the lateral geniculate nucleus. Nature, 466(7304), 373-377.

Schmid, M. C., Schmiedt, J. T., Peters, A. J., Saunders, R. C., Maier, A., & Leopold, D. A. (2013). Motion-sensitive responses in visual area V4 in the absence of primary visual cortex. Journal of Neuroscience, 33(48), 18740-18745.

Schneider, K. A., & Kastner, S. (2005). Visual responses of the human superior colliculus: a high-resolution functional magnetic resonance imaging study. Journal of Neurophysiology, 94(4), 2491-2503.

Schoenfeld, M. A., Noesselt, T., Poggel, D., Tempelmann, C., Hopf, J. M., Woldorff, M. G., Heinze, H.J., & Hillyard, S. A. (2002). Analysis of pathways mediating preserved vision after striate cortex lesions. Annals of Neurology: Official Journal of the American Neurological Association and the Child Neurology Society, 52(6), 814-824.

Seifert, M., & Hock, H. S. (2014). The independent detection of motion energy and counterchange: flexibility in motion detection. Vision Research, 98, 61-71.

Seiffert, A. E., & Cavanagh, P. (1998). Position displacement, not velocity, is the cue to motion detection of second-order stimuli. Vision Research, 38(22), 3569-3582.

Shadlen, M., & Carney, T. (1986). Mechanisms of human motion perception revealed by a new cyclopean illusion. Science, 232(4746), 95-97.

Shapiro, A., Lu, Z. L., Huang, C. B., Knight, E., & Ennis, R. (2010). Transitions between central and peripheral vision create spatial/temporal distortions: A hypothesis concerning the perceived break of the curveball. Public Library of Science One, 5(10), e13296.

Shapiro, A. G., Knight, E. J., & Lu, Z. L. (2011). A first-and second-order motion energy analysis of peripheral motion illusions leads to further evidence of “feature blur” in peripheral vision. Public Library of Science One, 6(4), e18719.

Sheliga, B. M., Kodaka, Y., FitzGibbon, E. J., & Miles, F. A. (2006). Human ocular following initiated by competing image motions: evidence for a winner-take-all mechanism. Vision Research, 46(13), 2041-2060.

Shioiri, S., & Cavanagh, P. (2000). Nonlinearity in color space measured by apparent motion. Perception & Psychophysics, 62(6), 1182-1190.

Silveira, L. C. L., & Perry, V. H. (1991). The topography of magnocellular projecting ganglion cells (M-ganglion cells) in the primate retina. Neuroscience, 40(1), 217-237.

Silveira, L. L., Picanço-Diniz, C. W., Sampaio, L. F. S., & Oswaldo-Cruz, E. (1989). Retinal ganglion cell distribution in the cebus monkey: a comparison with the cortical magnification factors. Vision Research, 29(11), 1471-1483.

Simoncelli, E. P., & Heeger, D. J. (1998). A model of neuronal responses in visual area MT. Vision Research, 38(5), 743-761.

Sincich, L. C., Park, K. F., Wohlgemuth, M. J., & Horton, J. C. (2004). Bypassing V1: a direct geniculate input to area MT. Nature Neuroscience, 7(10), 1123-1128.

Skottun, B. C. (2015). On the use of spatial frequency to isolate contributions from the magnocellular and parvocellular systems and the dorsal and ventral cortical streams. Neuroscience & Biobehavioral Reviews, 56, 266-275.

Smith, A. T. (1994). Correspondence-based and energy-based detection of second-order motion in human vision. Journal of the Optical Society of America A, 11(7), 1940-1948.

Smith, A. T., Hess, R. F., & Baker, C. L. (1994). Direction identification thresholds for second-order motion in central and peripheral vision. Journal of the Optical Society of America A, 11(2), 506-514.

Smith, A.T. & Ledgeway, T. (1998). Sensitivity to second-order motion as a function of temporal frequency and eccentricity. Vision Research, 38, 403-410.

Smith, A. T., & Ledgeway, T. (2001). Motion detection in human vision: a unifying approach based on energy and features. Proceedings of the Royal Society of London. Series B: Biological Sciences, 268(1479), 1889-1899.

Smith, A. T., Singh, K. D., Williams, A. L., & Greenlee, M. W. (2001). Estimating receptive field size from fMRI data in human striate and extrastriate visual cortex. Cerebral Cortex, 11(12), 1182-1190.

Snowden, R. J., & Hess, R. F. (1992). Temporal frequency filters in the human peripheral visual field. Vision Research, 32(1), 61-72.

Solomon, J. A., & Sperling, G. (1994). Full-wave and half-wave rectification in second-order motion perception. Vision Research, 34(17), 2239-2257.

Standage, G. P., & Benevento, L. A. (1983). The organization of connections between the pulvinar and visual area MT in the macaque monkey. Brain Research, 262(2), 288-294.

Steinman, R. M., Pizlo, Z., & Pizlo, F. J. (2000). Phi is not beta, and why Wertheimer's discovery launched the Gestalt revolution. Vision Research, 40(17), 2257-2264.

Stevens, B., Allen, N. J., Vazquez, L. E., Howell, G. R., Christopherson, K. S., Nouri, N., Micheva, K.D., Mehalow, A.K., Huberman, A.D., Stafford, B., Sher, A., Litke, A.M., Lambris, J.D., Smith, S.J., John, S.W.M. & Barres, B. A. (2007). The classical complement cascade mediates CNS synapse elimination. Cell, 131(6), 1164-1178.

Stoerig, P., & Cowey, A. (1989). Wavelength sensitivity in blindsight. Nature, 342(6252), 916-918.

Strasburger, H., Rentschler, I., & Jüttner, M. (2011). Peripheral vision and pattern recognition: A review. Journal of Vision, 11(5), 13-13.

Toet, A., & Koenderink, J. J. (1988). Differential spatial displacement discrimination thresholds for Gabor patches. Vision Research, 28(1), 133-143.

Tran, A., MacLean, M. W., Hadid, V., Lazzouni, L., Nguyen, D. K., Tremblay, J., Dehaes, M., & Lepore, F. (2019). Neuronal mechanisms of motion detection underlying blindsight assessed by functional magnetic resonance imaging (fMRI). Neuropsychologia, 128, 187-197.

Tse, P. U., & Hsieh, P. J. (2006). The infinite regress illusion reveals faulty integration of local and global motion signals. Vision Research, 46(22), 3881-3885.

Tse, P. U. (2006). Neural correlates of transformational apparent motion. Neuroimage, 31(2), 766-773.

Ukkonen, O. I., & Derrington, A. M. (2000). Motion of contrast-modulated gratings is analysed by different mechanisms at low and at high contrasts. Vision Research, 40(24), 3359-3371.

Ungerleider, L. G., & Desimone, R. (1986). Cortical connections of visual area MT in the macaque. Journal of Comparative Neurology, 248(2), 190-222.

Vaina, L. M., Soloviev, S., Calabro, F. J., Buonanno, F., Passingham, R., & Cowey, A. (2014). Reorganization of retinotopic maps after occipital lobe infarction. Journal of Cognitive Neuroscience, 26(6), 1266-1282.

Van Buren, J. M. (1963). Trans-synaptic retrograde degeneration in the visual system of primates. Journal of Neurology, Neurosurgery, and Psychiatry, 26(5), 402.

Van Essen, D. C., Lewis, J. W., Drury, H. A., Hadjikhani, N., Tootell, R. B., Bakircioglu, M., & Miller, M. I. (2001). Mapping visual cortex in monkeys and humans using surface-based atlases. Vision Research, 41(10-11), 1359-1378.

Van Santen, J. P., & Sperling, G. (1984). Temporal covariance model of human motion perception. Journal of the Optical Society of America A, 1(5), 451-473.

Van Santen, J. P., & Sperling, G. (1985). Elaborated Reichardt detectors. Journal of the Optical Society of America A, 2(2), 300-321.

Vilter, V. C. (1954). Nouvelle conception de relations synaptiques dans la photoperception par les cones retiniens. Comptes Rendus des Seances de la Societe de Biologie, 3–4, 148: 220–222.

Virsu, V., & Rovamo, J. (1979). Visual resolution, contrast sensitivity, and the cortical magnification factor. Experimental Brain Research, 37(3), 475-494.

Virsu, V., Rovamo, J., Laurinen, P., & Näsänen, R. (1982). Temporal contrast sensitivity and cortical magnification. Vision Research, 22(9), 1211-1217.

Wall, M. B., Walker, R., & Smith, A. T. (2009). Functional imaging of the human superior colliculus: an optimised approach. Neuroimage, 47(4), 1620-1627.

Wallisch, P. (2007). Principles of spatial and temporal integration within the spectral receptive field of MT neurons. PhD Thesis. The University of Chicago.

Walls, G. L. (1942). The vertebrate eye and its adaptive radiation. Cranbrook Institute of Science.

Warner, C. E., Kwan, W. C., Wright, D., Johnston, L. A., Egan, G. F., & Bourne, J. A. (2015). Preservation of vision by the pulvinar following early-life primary visual cortex lesions. Current Biology, 25(4), 424-434.

Watanabe, T., & Oyama, T. (1988). Are illusory contours a cause or a consequence of apparent differences in brightness and depth in the Kanizsa square? Perception, 17(4), 513-521.

Watson, A. B. (2014). A formula for human retinal ganglion cell receptive field density as a function of visual field location. Journal of Vision, 14(7), 15-15.

Wehrhahn, C. (2006). Reversed phi revisited. Journal of Vision, 6(10), 2-2.

Weinmann, M. (2013). Visual features—From early concepts to modern computer vision. Advanced Topics in Computer Vision, 1-34.

Weiskrantz, L., Barbur, J. L., & Sahraie, A. (1995). Parameters affecting conscious versus unconscious visual discrimination with damage to the visual cortex (V1). Proceedings of the National Academy of Sciences, 92(13), 6122-6126.

Wertheimer, M. (1912). Experimentelle Studien uber das Sehen von Bewegung. Zeitschrift fur Psychologie, 61.

Wessinger, C. M., Fendrich, R., & Gazzaniga, M. S. (1999). Variability of residual vision in hemianopic subjects. Restorative neurology and Neuroscience, 15(2-3), 243-253.

Westheimer, G. (1982). The spatial grain of the perifoveal visual field. Vision Research, 22(1), 157-162.

Westheimer, G. (1979). Scaling of visual acuity measurements. Archives of Ophthalmology, 97(2), 327-330.

Weymouth, F. W. (1958). Visual sensory units and the minimal angle of resolution. American Journal of Ophthalmology.

Whitaker, D., Mäkelä, P., Rovamo, J., & Latham, K. (1992). The influence of eccentricity on position and movement acuities as revealed by spatial scaling. Vision Research, 32(10), 1913-1930.

Wilson, H. R. (1991). Psychophysical models of spatial vision and hyperacuity. Spatial Vision, 10, 64-81.

Wohlgemuth, G. (1911). On the after-effect of seen movement (British Journal of Psychology, Monograph supplument 1). Cambridge University Press.

Xiao, K., Fu, C., Karatzas, D., & Wuerger, S. (2011). Visual gamma correction for LCD displays. Displays, 32(1), 17-23.

Yan, W., Peng, Y. R., van Zyl, T., Regev, A., Shekhar, K., Juric, D., & Sanes, J. R. (2020). Cell atlas of the human fovea and peripheral retina. Scientific Reports, 10(1), 1-17.

Yu, H. H., Atapour, N., Chaplin, T. A., Worthy, K. H., & Rosa, M. G. (2018). Robust visual responses and normal retinotopy in primate lateral geniculate nucleus following long-term lesions of striate cortex. Journal of Neuroscience, 38(16), 3955-3970.

Zeki, S., & Shipp, S. (1988). The functional logic of cortical connections. Nature, 335(6188), 311-317.

Zeki, S., Watson, J. D., & Frackowiak, R. S. (1993). Going beyond the information given: the relation of illusory visual motion to brain activity. Proceedings of the Royal Society of London. Series B: Biological Sciences, 252(1335), 215-222.

Zeki, S. (1973). Colour coding in rhesus monkey prestriate cortex. Brain Research.

Zeki S. (1978). Uniformity and diversity of structure and function in the monkey prestriate visual cortex. Journal Physiology, 277:273–290.

Zeki, S. (1980). A direct projection from area V1 to area V3A of rhesus monkey visual cortex. Proceedings of the Royal Society of London. Series B. Biological Sciences, 207(1169), 499-506.

Zeki, S. (1983). The distribution of wavelength and orientation selective cells in different areas of monkey visual cortex. Proceedings of the Royal Society of London. Series B. Biological Sciences, 217(1209), 449-470.