Global Shape Processing in Human Visual Cortex

Samuel John Lawrence

PhD

University of York

Psychology

August 2016

Abstract

The research described in this thesis used a combination of neuroimaging, brain stimulation and psychophysical methods to explore the representation of global shape in human visual cortex. Global shape processing mechanisms integrate over local orientation information to form an abstracted representation of an object's shape; an important processing step for accurate object recognition in the complex visual scenes of everyday life. This thesis presents evidence for a global shape processing stream in human lateral occipital cortex comprising brain areas LO1, LO2 and object-selective lateral occipital complex (LOC). First, brain responses to shape were shown to be more global and less dependent on task in LO2 compared to LO1. Second, a global shape aftereffect was used to demonstrate that global representations of object shape are formed by a series of integrative mechanisms, where each successive mechanism integrates over outputs from the preceding mechanisms to form increasingly global shape representations. Third, brain stimulation was used to determine whether area LO2 is causally important for global shape adaptation and shape discriminations, however brain stimulation had no effect on shape perceptions meaning conclusions about a causal role for LO2 in global shape processing could not be drawn. Finally, neural tuning to radial frequency, one possible dimension of a representational shape space, was modelled using Gaussian neural modelling. Tuning towards radial frequencies that are processed globally was localised to lateral occipital cortex, and shifted from local to global frequencies in LO2 and LOC. Overall, this thesis provides novel contributions to the current understanding of how the analysis of an object's shape facilitates object recognition, and the nature and locus of global shape representations in the human brain.

List of contents

List of figures

List of tables

Acknowledgements

First I would like to express my deepest gratitude to my supervisor and mentor Tony Morland. Thank you for your incredibly attentive and dedicated supervision. I will carry the lessons learned from your mentorship with me forever. Thanks also for nurturing such a happy and friendly lab environment for me to work in. Thank you to my TAP members Alex Wade and Tom Hartley for your insightful discussion and valuable feedback on the thesis. Special thanks to Alex for giving me my first job, and inspiring me to work in science full time.

Thank you to everyone in my lab for your collaboration and friendship. Thanks to Andre' for your technical assistance and patience (and many drinks). If I have absorbed a fraction of your knowledge and generosity while working with you I will be in good shape for the future. Thanks to Richard for your companionship, encouraging me to better myself, and for the Richardisms. Thanks to Holly for joining the lab to save Rich and I from ourselves, and for all the baked goods. Thanks to Rachel, always the voice of reason and yet delightfully fun. Thanks to Bruce for teaching me psychophysics, and for all the fun Thursdays that went on too late. Thanks also to honorary lab mascot Junior, and all the infectious silliness and joyfulness you have brought to us. The last three years have been a blast and I will miss you all.

Thank you to Mum, Dad, Jess, and the rest of my family for your unconditional love and support that allowed me to pursue a PhD and overcome its challenges.

Thank you also to all my other friends and colleagues from YNiC and the psychology department. There are too many of you to name you all, but special thanks to Milena, Hinke, James, Alex, girl Sam, Marc, Babs, Becca, Scott and Greta. I wish I could have found more time to spend with you all.

Finally, I would like to thank Browns of Heslington for their always excellent and affordable sandwiches that have sustained me over three years.

Declaration

This thesis contains original work carried out by the author, Samuel J. Lawrence, under the supervision of Prof. Antony B. Morland and has not been submitted for any other degree.

The data in Chapter 3 were collected in conjunction with Richard J. W. Vernon, a fellow PhD student in our laboratory, and were presented at the *Society for Neuroscience*, 2014 as:

Lawrence, S. J. D., Vernon, R. J. W., Gouws, A. D., Silson, E. H. & Morland, A. B. Responses to radial frequency patterns in lateral occipital visual field maps LO1 and LO2 during shape and orientation discriminations.

The data in Chapter 4 were presented at the *European Conference on Visual Perception*, 2015 as:

Lawrence, S. J. D, Keefe, B. D., Vernon, R. J. W., Brown, H. D., Wade, A. R., McKeefry, D. J. & Morland, A. B. (2015). Characterising shape aftereffects using composite radial frequency patterns. *Perception, 44*, 226-227.

And at the *Society for Neuroscience*, 2015 as:

Lawrence, S. J. D, Keefe, B. D., Vernon, R. J. W., Brown, H. D., Wade, A. R., McKeefry, D. J. & Morland, A. B. A potential extrastriate locus for adaptation to composite radial frequency patterns.

And were publish in the *Journal of Vision* as:

Lawrence, S. J. D., Keefe, B. D., Vernon, R. J. W., Wade, A. R., McKeefry, D. J., & Morland, A. B. (2016). Global shape aftereffects in composite radial frequency patterns. *Journal of vision*, 16(7), 17-17.

The data in Chapter 5 were presented at the *Vision Sciences Society*, 2015 as:

Lawrence, S. J. D, Keefe, B. D., Vernon, R. J. W., Gouws, A. D., Brown, H. D., Wade, A. R., McKeefry, D. J. & Morland, A. B. (2015). Effects of TMS to occipital face area on the

perception of face viewpoint cued only by shape changes in the external contour of the face. *Journal of Vision*, *15*(12), 432-432.

The retinotopic mapping and functional localiser data in Chapter 6 were collected and analysed in conjunction with Richard J. W. Vernon. The data in Chapter 6 were presented at the *Yorkshire Vision Network,* 2016 as:

Lawrence, S. J. D, Vernon, R. J. W., Keefe, B. D., Gouws, A. D., Wade, A. R., McKeefry, D. J. & Morland, A. B. Radial frequency tuning in human visual cortex.

And at the *European Conference on Visual Perception*, 2016 as:

Lawrence, S. J. D, Vernon, R. J. W., Keefe, B. D., Gouws, A. D., Wade, A. R., McKeefry, D. J. & Morland, A. B. Radial frequency tuning in human visual cortex.

Chapter 1: Introduction and literature review

1.1: The hierarchy of shape perception

Visual perception is the culmination of a neural analysis of light photons that enter the eye and fall upon the retina. The retinal photoreceptors responsive to this light respond to a very small, localised portion of the visual field (Zeki, 1993). A complex, cortical analysis of this local visual information generates a perception of the visual field which we can understand and interact with.

Early work examining the primate visual system has established that information originating from the retina is fed through a hierarchical cortical processing stream (Felleman & Van Essen, 1991). The first stage of cortical processing is the primary visual cortex (V1) which, like the retina, contains neurons with small receptive fields (Hubel & Wiesel, 1968; Zeki, 1993). Neurons in this area are also retinotopically organised, meaning the spatial organisation of inputs from the retina is maintained. As such adjacent neurons in V1 are responsive to adjacent portions of the visual field. At this early stage of processing, neurons respond to simple visual features such as line orientation and spatial frequency (Hubel & Wiesel, 1968). These types of cells are called simple cells, and they respond optimally to a line with a preferred orientation that falls within their receptive field. Multiple simple cells feed into a complex cell, which remains selective for orientation but responds to a larger area of the visual field than simple cells. The receptive fields of these types of cells are well described by Gabor patches (grating stimuli, the luminance across which is defined by a sinusoid), which are popular stimuli for visual experiments. Together, simple and complex cells provide an example of the hierarchical and integrative nature of the visual system.

Beyond V1, monkey research has shown that extrastriate visual areas combine outputs from V1 (and, depending on their position in the visual hierarchy, other antecedent visual areas) to integrate across visual field location and form a more complex representation of the visual field. For example area V2 combines inputs from V1 to represent more complex visual features such as angles (Hegdé & Van Essen, 2000; Ito & Komatsu, 2004). Progressing further downstream from early visual cortex, representations become increasingly complex. Area V4 is the first area where

representations of complex shapes are formed. V4 cells respond strongly to curved contours (Pasupathy & Connor, 1999), are tuned to the shape of those contours (Pasupathy & Connor, 2001), and together provide a population code of complex shapes (Pasupathy & Connor, 2002). In the monkey brain, therefore, extrastriate cortex appears to be the locus of shape analysis of objects in the visual field, with the first representations of object shape emerging in V4.

Real visual scenes are often complicated and contain many discrete objects, some of which may occlude each other. The detection and recognition of real objects, therefore, requires a representation of those objects that is abstracted beyond the local orientation information processed in early visual cortex, which might be uninformative about an object's identity in the presence of, for example, occlusion. This process may be facilitated by an adherence to the Gestalt laws of perceptual grouping. According to these laws, the general principle that portions of the visual field which share visual properties such as proximity, similarity, continuity and closure are likely to belong to the same object is utilised to facilitate automatic object identification (Wertheimer, 1923). Such representations are described as more global, because they are less dependent on lower-level visual features such as an object's location in the visual field. For example, representations of complex objects such as faces are more global than representations of simple features such as orientation, shown by perceptual face aftereffects, where adaptation to a face affects face perception throughout the visual field, whereas adaptation to orientation only affects orientation perception in the adapted location (Melcher, 2005). This demonstrates that representations of orientation are tied to the stimulus' position in the visual field, but representations of faces are not. As such, face representations are considered more global than orientation representations. Forming global representations of objects that are invariant to transient visual features such as lighting, viewpoint and occlusion is thought to facilitate successful object recognition in real visual scenes.

Simple shapes have been shown to be the simplest visual stimuli to facilitate the formation of global stimulus representations (Wilkinson, Wilson, & Habak, 1998). As such, the formation of a global representation of an object's shape is thought to be an important step towards object recognition (Loffler, 2008; Wilson & Wilkinson, 2015). The neural mechanisms underpinning global shape processing have been interrogated

psychophysically, however the neural locus of these processes in humans remains unclear. Given the progressively complex response properties of neurons through the visual hierarchy, it is likely that such a level of processing occurs somewhere in extrastriate visual cortex, as is the case in monkeys. In the following section, behavioural studies exploring human shape processing are reviewed. This will in turn be followed by a review of visual areas in the human brain that may be implicated in global shape processing.

1.2: Psychophysical characterisation of shape processing mechanisms

1.2.1: Contour closure and curvature

Psychophysical research has identified which visual features are most important for engaging global-scale shape processing mechanisms. In particular, the closure of a contour creates a pop-out effect where the stimulus is perceived as a complete object, which in turn facilitates global processing of the object. This has been demonstrated in contour integration – a task where one must detect the presence of a series of collinearly aligned Gabor patches embedded in a field of distractor Gabors with random orientations (see Figure 1.1A). In such an experiment closed contours are considerably easier to detect that open ones (Kovács & Julesz, 1993). Closed contours can be detected with greater spacing between aligned elements than open contours, and the space integrated across is greater than receptive field sizes in V1, suggestive of a global integration mechanism (Loffler, 2008). This indicates that such integrative mechanisms which link aligned elements together to extract a contour are sensitive to and enhanced by global structure.

A second important feature for global shape processing is curvature. In particular, points of maximum curvature appear to be highly informative about object shape. For example, marking just the points of maximum curvature on the outline shape of a cat and connecting them is sufficient to generate the perception of a cat (Attneave, 1954, see Figure 1.1B). Moreover, when just the corners of a shape are cued, compelling illusory contours arise in order to form the cued shape (Kanizsa, 1976, see Figure 1.1C).

This implies a predisposition of the visual system to link points of maximum curvature together to form closed shapes.

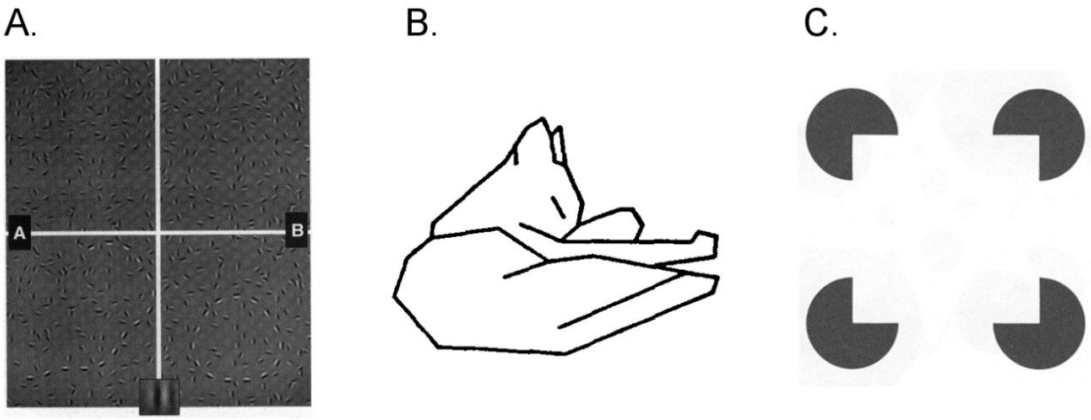

Figure 1.1: Stimuli from the literature demonstrating the important of contour closure and points of maximum curvature for global shape processing. (A) Stimuli from Kovács & Julesz (1993). An open contour of aligned Gabor patches embedded amongst distractor Gabors (left) is more difficult to detect than a closed contour (right). The contrast of the contours is increased in the bottom two panels for demonstration purposes; the top two panels show the stimuli as they appeared in the experiment. (B) An example stimulus from Attneave (1954) showing that the outline of a cat can be accurately described using only the points of maximum curvature, connected by a straight line. (C) An example stimulus from Kanizsa (1976) demonstrates an illusory square cued only by points of maximum curvature.

Additionally, psychophysical research has shown the presence of curvature allows for a greater sensitivity to shape features. For example, observers are better at discriminating aspect ratio in circles and ellipses compared to squares and rectangles (Zanker & Quenzer, 1999). The explanation provided by the authors for this effect was that an enhanced sensitivity to curvature facilitated improved discriminations for circular shapes, but this was not present in rectangular shapes. Moreover, discrimination of circles was considerably better than reported for isolated curvature arcs (Foster, Simmons, & Cook, 1993), again demonstrating the importance of contour closure for shape perception.

Similarly, Morgan (2005) compared performance on an area discrimination task to predicted performance given previous accuracy in height and width discriminations of ellipses and rectangles. Performance for rectangles was worse than predicted, however performance was better than predicted for ellipses. This provides further evidence for a

specialised sensitivity to curvature, which allows for more accurate shape discriminations. Overall, psychophysical evidence indicates that global shape mechanisms are engaged by closed contours and are particularly sensitive to curvature.

1.2.2: Radial frequency patterns

Given the sensitivity of shape mechanisms to contour closure and curvature it is unsurprising that a class of stimuli called radial frequency (RF) patterns, closed contours defined by their curvature, have proven to be powerful stimuli for probing global shape processing mechanisms. RF patterns are defined by a sinusoidal modulation of the radius of a circle as a function of polar angle (Wilkinson, Wilson, & Habak, 1998; reviewed in further detail in Chapter 2). Manipulating the frequency, amplitude and phase of the sinusoid allows for precise control over the shape of the stimulus (Figure 1.2).

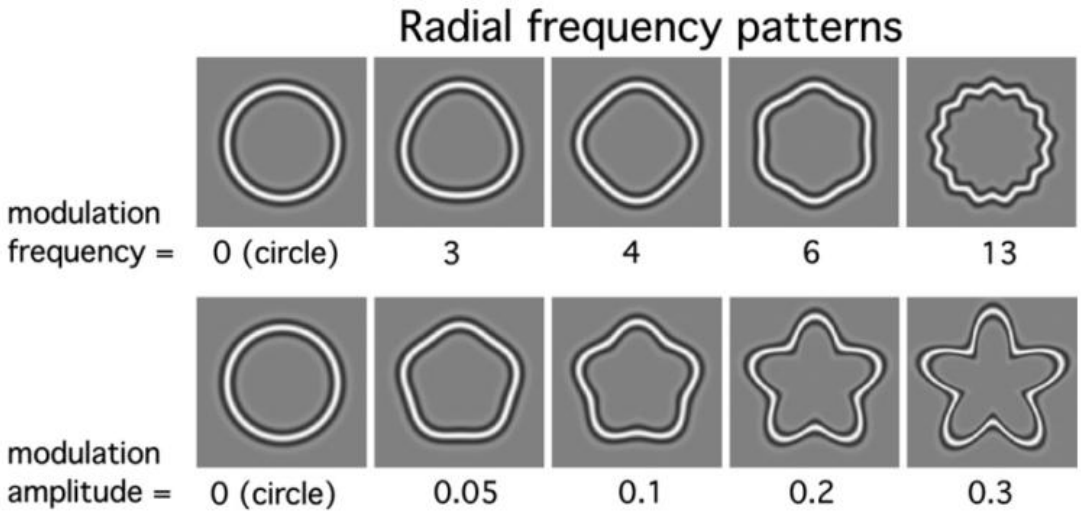

Figure 1.2: A series of RF patterns defined by a sinusoidal modulation of a circle's radius as a function of the polar angle. The frequency of the sinusoid determines the number of lobes in the shape (top row). The amplitude of the sinusoid increases the prominence of those lobes (bottom row). The orientation of the shapes is controlled by the phase of the sinusoid. Figure from Loffler (2008).

Wilkinson et al. (1998) measured detection thresholds for deformation from circularity (i.e. discriminating a lobed RF pattern from a circle) in RF patterns for a range of frequencies. Observers were extremely proficient at this task, with detection thresholds reaching the hyperacuity range (i.e. deviations from circularity that were smaller than

the spacing of retinal photoreceptors were detectable). These thresholds were shown to be stable across a range of stimulus sizes and spatial frequencies (defined by the width of the Gabor-like luminance profile that defines the shape; see Figure 1.2). Performance on this task was better than the detection of deformation from straightness in a sinusoidally modulated line, demonstrating that RF pattern discrimination cannot be explained by local orientation or curvature analysis. RF patterns therefore engage global-scale shape processing mechanisms which provide a hyper sensitivity to deformation from circularity.

Many other studies have since confirmed that RF patterns are processed by global mechanisms. For example Hess, Wang, & Dakin (1999) measured deformation from circularity thresholds for closed RF patterns and for stimuli where the contour was fragmented into separate segments. Thresholds were elevated (performance was worse) for fragmented contours compared to closed shapes. This provides additional evidence that RF patterns are processed by global shape mechanisms, as the sensitivity of these mechanisms to curvature was dependent of the global structure of the stimulus. However, in fragmented stimuli contour segments were arranged in such a way that perceptual completion of a closed shape was impossible. Their result, therefore, could either be due to the incompleteness of the contour, the arrangement of contour fragments, or both.

Addressing this, Loffler, Wilson, & Wilkinson (2003) showed that occluding part of the contour elevated deformation from circularity thresholds, even when the same global arrangement of contour segments was maintained. Therefore contour closure is important for engaging global shape processing mechanisms, consistent with perceptual pop-out effects of closed contours (Kovács & Julesz, 1993). However, threshold elevation was dependent on which part of the contour was occluded. In Loffler et al.'s (2003) study, threshold elevation was weakest when points of minimum curvature were occluded, and strongest when points of maximal curvature were occluded. Therefore points of maximum curvature are particularly important cues for deformation from circularity, and may be critical for global processing. This is consistent with Attneave's (1954) proposal that points of maximum curvature are the most informative points on an object’s contour for object recognition.

While evidence supports global processing of RF patterns, the engagement of global mechanisms may be dependent on stimulus frequency. Jeffrey, Wang, & Birch (2002) replicated Wilkinson et al.'s (1998) result, finding that detection of RF patterns was better than predicted by the analysis of local, low-level features. This was determined by comparing subjects' performance for whole RF patterns to a simulated normal distribution derived from each subject's performance for the same task in a stimulus comprising only one cycle of contour modulation. Compared to this distribution, detection of whole RF patterns was significantly improved for low frequency RFs (two-lobed RF2 and four-lobed RF4 patterns), but not high RFs (RF8 and RF16 patterns). It was therefore proposed that there are at least two mechanisms involved in detecting deformation from circularity in RF patterns; one which globally pools information for low RFs, and one which analyses local curvatures for high RFs.

This highlights a distinguishing property between low and high RFs. Multiple RFs can be summed together to create a composite RF pattern, allowing one to observe how different RFs interact with each other. In Figure 1.3, it can be seen that summing an RF3 pattern with another low RF distorts the global shape considerably. However, when the RF3 pattern is combined with a high RF, the global shape of the RF3 patterns is still easily detectable. This demonstrates how low RF detail may be more informative about global shape structure, while high RF information is more indicative of textural or other high frequency detail about the contour. This may provide some explanation for why global pooling of information seems to only occur in low RFs (Jeffrey et al., 2002). This was examined experimentally by Bell et al. (2007) who found that in an RF3+24 composite pattern the detection of the RF3 component was independent of and unimpaired by the presence of the RF24 component. Therefore the analysis of global shape is independent of local contour noise, suggesting independent mechanisms for the processing of low and high RFs.

In addition to the further interrogation of shape encoding mechanisms, composite RF patterns can be used to investigate the perception of real-world objects. Multiple RFs can be summed together to form a realistic portrayal of complex, real-world objects such as faces (Wilson et al., 2000; Wilson, Loffler, & Wilkinson, 2002; Wilson & Wilkinson, 2002) or animals (Alter & Schwartz, 1988; Vernon et al., 2016). This has allowed researchers to investigate how the analysis of global shape may be employed

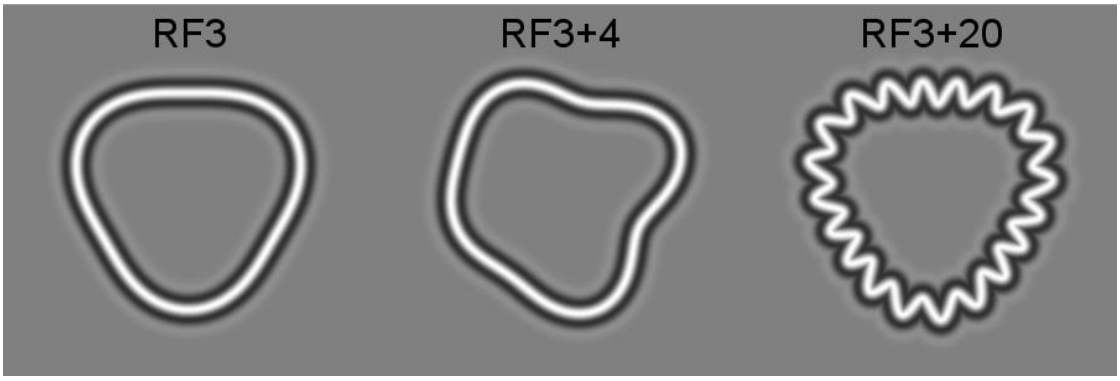

Figure 1.3: An RF pattern with a frequency of 3 (RF3) and its composites when summed with an RF4 (RF3+4) and RF20 (RF3+20) pattern. The global shape of the RF3 pattern is distorted in the RF3+4 pattern, but remains clearly visible in the RF3+20 pattern.

in biologically relevant tasks such as face viewpoint perception (Daar & Wilson, 2012). Daar and Wilson (2012) achieved this by replicating the face viewpoint aftereffect in synthetic face stimuli generated by combining multiple RF components. In the face viewpoint aftereffect, prolonged exposure to a face rotated to face away from the observer causes a subsequently neutral face to be perceived as rotated in the opposite direction (Fang & He, 2005; perceptual adaptation effects are reviewed in the following section). Daar and Wilson (2012) showed that adaptation to just the shape of a face outline comprising several RF components was sufficient to induce the aftereffect in a full face test stimulus, showing that shape encoding mechanisms are involved in face viewpoint aftereffects. In this way, composite RF patterns can be used to interrogate how shape processing mechanisms are applied to the perception of real-world, biologically relevant stimuli

Overall RF patterns represent a well-defined and easily manipulated class of closed contour shapes. These stimuli have been demonstrated to engage global-scale shape processing mechanisms, which integrate local information across the visual field. Moreover RF patterns can be combined to create composite stimuli that reflect real-world objects, allowing for the analysis of shape information that is important for biologically relevant stimuli.

1.2.3: Perceptual adaptation

Perceptual adaptation is a phenomenon which has allowed researchers to examine the neural properties of sensory systems using behavioural measures. Adaptation is a

feature of perceptual mechanisms where prolonged exposure to a stimulus with a perceptual property lying on some dimension results in the altered perception of subsequent stimuli along the same dimension (Thompson & Burr, 2009). For example, in the motion aftereffect, first described by Addams (1834), after prolonged exposure to a downward-flowing waterfall, a stationary rock face appeared to move upward.

It is generally accepted that adaptation effects are due to changed sensitivities of the neurons responsive to the adapting stimulus following adaptation (see Figure 1.4). This results in an altered overall ratio of sensitivities across multiple channels within a perceptual mechanism, which has consequences for subsequent perceptions. In the motion aftereffect described by Addams (1834), neurons responsive to the downward-motion of the waterfall adapted and reduced their sensitivities. As a result, when Addams shifted his gaze to a stationary rock face, upward motion neurons were more sensitive relative to downward motion neurons, resulting in the perception of upward motion. This neural adaptation provides benefits for the visual system, allowing for better discrimination around the adaptor as perceptual mechanisms adjust their sensitivities to suit the present environment (Thompson & Burr, 2009). The advantages of perceptual adaptation can be felt in a dark environment in which it is at first difficult to see. With time visual sensitivity to a dark environment improves. This is due to the adaptation of light-sensitive perceptual mechanisms, allowing for improved vision in mesopic light levels.

As well as being beneficial for vision, adaptation effects are a powerful tool for researchers to infer the response properties of neural mechanisms using behavioural measures. An example of this is the McCollough effect; a colour aftereffect where adaptation to a vertical red grating causes a vertical grey grating to appear green, while adaptation to a horizontal green grating causes a horizontal grey grating to appear red (McCollough, 1965, see Figure 1.5). The orientations of the gratings means that if colour encoding neurons are not selective for orientation, they will be exposed to both red and green, resulting in no aftereffect. This aftereffect, therefore, demonstrates that at least some proportion of colour sensitive neurons in visual cortex are also selective for orientation. In this way, measuring which stimulus manipulations preserve or abolish a visual aftereffect allows researchers to measure which stimulus properties the adapted mechanism is sensitive or invariant to.

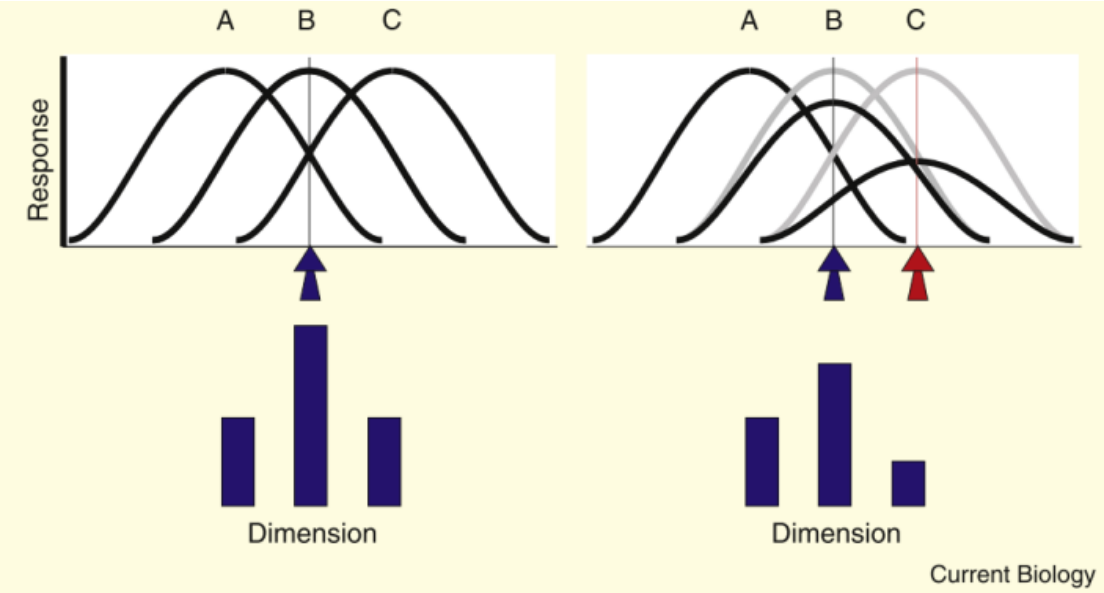

Figure 1.4: A schematic representation of the effect of perceptual adaptation on a hypothetical three-channel neural mechanism. In the left panel, the middle channel is maximally responsive to a stimulus presented at position B (blue arrow), generating the correct percept. However in the right panel, after adaptation to a stimulus at position C (red arrow), the sensitivities of the right and middle channels are reduced. This results in an imbalanced ratio of relative sensitivities between the three channels where channel A is more active than channel C. This results in a different response to the stimulus at position B compared to no adaptation (left panel), leading to an altered perception. Figure from Thompson and Burr (2009).

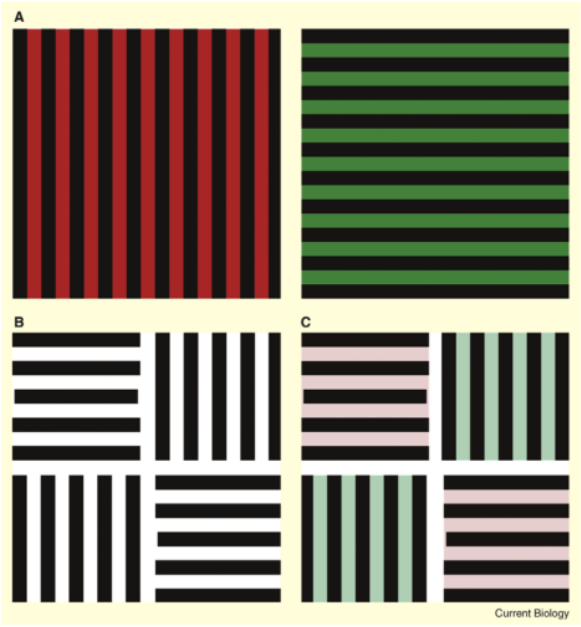

Figure 1.5: The McCollough effect. Adaptation to a red vertical grating (left of panel A) causes vertical black and white gratings (B) to appear green (C). Adaptation to a green horizontal grating (right of Panel A) causes horizontal black and white gratings (B) to appear red (C). Figure from Thompson and Burr (2009).

1.2.4: Shape adaptation

Research has shown that the perceived shape of an object is also an adaptable feature (Suzuki & Cavanagh, 1998; Suzuki, 2001, 2003). In these aftereffects prolonged exposure to a shape causes warped perceptions of subsequently presented similar shapes. For example, adaptation to a squashed circle causes a normal circle to appear squashed in the opposite direction (Suzuki & Cavanagh, 1998). The experimental manipulation of shape aftereffects allows for the response properties of neural shape encoding mechanisms to be measured. For example, shape aftereffects remain present when a test stimulus is presented with a different size to the adaptor (Suzuki, 2001), implying that shape processing mechanisms are not sensitive to stimulus size.

The response properties of shape encoding mechanisms were further investigated using shape frequency and amplitude aftereffects by Gheorghiu and Kingdom (2006, 2007) and Gheorghiu et al. (2009). In the shape frequency aftereffect, the perceived number of undulations along a sinusoidal line was manipulated by adaptation (see Figure 1.6A). In the shape amplitude aftereffect, the perceived prominence of those undulations was manipulated by adaptation (se Figure 1.6B). These aftereffects could not be explained by adaptation to local orientation or curvature and therefore recruit global shape processing mechanisms (Gheorghiu & Kingdom, 2007). The shape frequency aftereffect was significantly reduced when the adaptor and test stimuli had reversed contrast polarities, and was also sensitive to blurring of the luminance profile that defined the stimulus (Gheorghiu & Kingdom, 2006). This therefore suggests that shape processing mechanisms are sensitive to stimulus contrast and luminance, which may be surprising given their apparent invariance to stimulus size (Suzuki, 2001).

While both properties of the same stimulus, shape frequency and amplitude have been shown to be separately adaptable, where frequency and amplitude aftereffects act on the chord and the sag (maximum deviation from linearity) of the curve, respectively (Gheorghiu & Kingdom, 2007). This suggests at least two distinct mechanisms for shape processing; one tuned to curvature frequency and one tuned to curvature amplitude. These mechanisms appear not to be entirely independent, however, as amplitude aftereffects are tuned to the frequency of the adaptor, where amplitude aftereffects are reduced when the frequency of the test stimulus is different to the

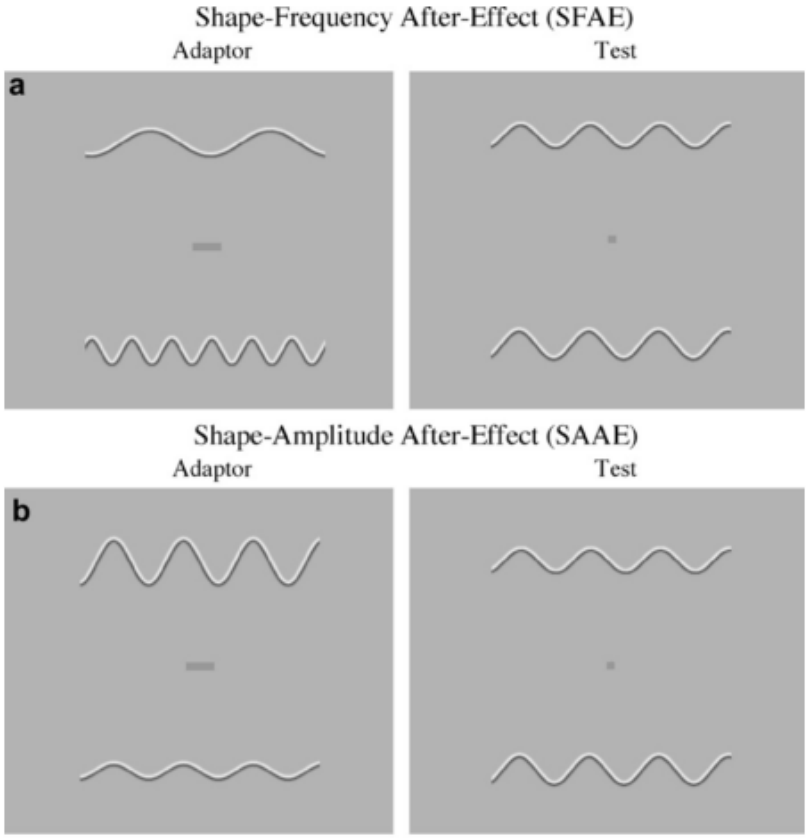

Figure 1.6: The shape frequency and shape amplitude aftereffects. (A) Adaptation to low frequency shape (top of left-hand panel) causes a test stimulus (right) to be perceived as higher frequency (bottom of left-hand panel) and vice versa. (B) Adaptation to a high amplitude shape (top of left-hand panel) causes a test stimulus (right) to be perceived as lower amplitude (bottom of left-hand panel) and vice versa. Figure from Gheorghiu and Kingdom (2007).

adaptor (Gheorghiu & Kingdom, 2007). Therefore amplitude aftereffects do not transfer well across shapes, suggesting different amplitude-tuned mechanisms for different shapes. In general, shape and frequency aftereffects demonstrate that shape processing is a complex process, involving multiple mechanisms that may possess tunings to different shape features such as frequency and amplitude.

Anderson et al. (2007) further explored shape aftereffects using RF patterns. It was demonstrated that adaptation to a high amplitude RF pattern caused a low amplitude RF pattern (circle or near-circle) to be perceived as a higher amplitude stimulus with contour modulations in antiphase to the adaptor (see Figure 1.7). This shape aftereffect was shown to be tolerant to changes in stimulus contrast and size, but not shape (e.g. adaptation to an RF5 pattern had no effect on the perceived shape of an RF4 pattern). This suggests that shape processing mechanisms are selective for shape, but not

contrast or size. This is inconsistent with the degree of contrast sensitivity for shape frequency aftereffects (Gheorghiu & Kingdom, 2006). This may be due differences in the stimuli used for the two studies. Gheorghiu and Kingdom (2006) measured shape aftereffects in open contours (Figure 1.6), whereas Anderson et al. (2007) used closed RF patterns (Figure 1.7). Since closed contours have a perceptual pop-out effect not present in open contours (Kovács & Julesz, 1993) which engages global processing mechanisms (Loffler, 2008), it seems likely shape aftereffects in open and closed contours involve separate mechanisms. This therefore suggests that shape aftereffects in RF patterns are more invariant to low-level features such as contrast than aftereffects in open contours, implying they involve more global shape processing mechanisms. In Chapter 4 this issue will be clarified by presenting psychophysical evidence for multiple mechanisms in RF pattern shape aftereffects with differing sensitivities and invariances to various stimulus properties.

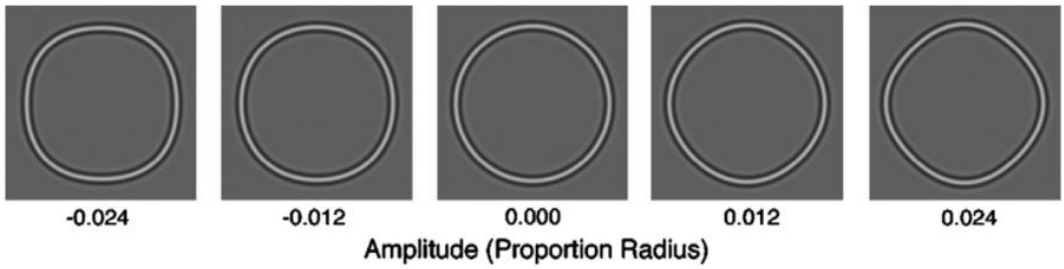

Figure 1.7: Example of a shape aftereffect in RF patterns. Adaptation to a high amplitude RF pattern (left) causes a circle (middle) to be perceived as a higher amplitude RF pattern with modulations in antiphase to the adaptor (right). Figure from Anderson et al. (2007).

Overall, shape adaptation experiments provide evidence for neural sensitivity to shape frequency and amplitude. Both these shape features are well captured by RF patterns, and shape adaptation studies show that RF patterns engage more global processing mechanisms than open contours. These studies demonstrate how shape adaptation and RF patterns can be used to investigate the properties of neural mechanisms involved in global shape processing; techniques that will be utilised in this thesis.

1.3: Models of shape perception

1.3.1: A curvature-based model of RF pattern perception

The numerous studies measuring shape detection (Hess et al., 1999; Jeffrey et al., 2002; Loffler et al., 2003; Wilkinson et al., 1998) and adaptation (Anderson et al., 2007; Bell et al., 2010; Bell & Kingdom, 2009) in RF patterns has culminated in the development of a biologically plausible model of RF pattern perception (Poirier & Wilson, 2006, 2010). This model proposes how extrastriate visual cortex may integrate over V1 outputs to form a global representation of shape. According to the model, this process is achieved in five stages (also see Figure 1.8):

1. Orientation-selective filters in V1 provide local information about the orientation of an object's contour in a specific part of the visual field.
2. A second set of filters, oriented orthogonally to those in stage 1, identify the object centre.
3. Responses from a third set of filters, extending out from the object centre, identify the number of contours and their radii.
4. For each contour identified in stage 3, curvature is assessed by a series of differently oriented filters arranged along a curved path. These responses are combined multiplicatively, such that a maximal response occurs across the filters that follow the contour along a curved path.
5. Finally, a population code across the filters from stage 4 provides a representation of the object's shape represented by curvature responses as a function of orientation about the object centre.

This model provides a biologically plausible method through which the extrastriate visual system might integrate across local, low-level V1 responses to form complex representations of shape and curvature. In a computational simulation, the stages proposed in this model have been demonstrated to be sufficient to recover the shape of one or multiple RF pattern inputs (Poirier & Wilson, 2006). Moreover, representing shape by size invariant features such as curvature relative to object centre is consistent with size invariances reported for shape detection (Wilkinson et al., 1998; Wilson & Wilkinson, 2002) and adaptation (Anderson et al., 2007; Suzuki, 2001).

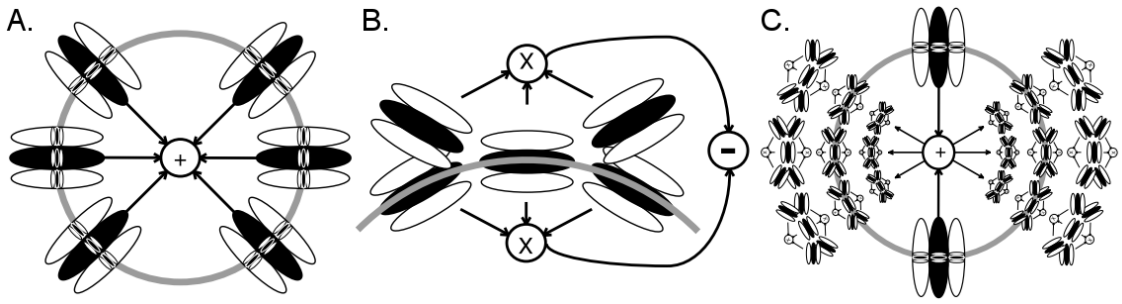

Figure 1.8: A schematic representation of the filters involved in shape encoding in Poirier and Wilson's (2006, 2010) model. (A) A first set of orientation-selective filters identify local orientation of the contour and an orthogonal set of filters identify the object centre (stages 1 and 2, see text). (B) After the number of contours and their radii are computed (stage 3, see text), a set of oriented filters analyse curvature (stage 4, see text). (C) Finally these filters provide a population code for shape as curvature responses relative to the object's centre (stage 5, see text). Figure from Poirier and Wilson (2010).

1.3.2: A multidimensional representational shape space

As previously discussed, curvature and shape encoding mechanisms such as those presented in Poirier and Wilson's (2006, 2010) model are highly adaptable (Anderson et al., 2007; Gheorghiu & Kingdom, 2006, 2007; Gheorghiu et al., 2009; Suzuki, 2001). This has led to the suggestion that shapes are represented in a multidimensional shape space centred around a template of a prototypical mean shape (Loffler, 2008). In such a system, shape adaptation effects would be due to shifting neural sensitivities along one dimension of shape space after adaptation. For example, adaptation to a high amplitude RF pattern causes a circle to be perceived as a higher amplitude RF pattern with modulations in antiphase to the adaptor (Anderson et al., 2007, see Figure 1.7), suggesting neural adaptation along a radial amplitude dimension of shape space.

This method of representation is consistent with neurophysiological recordings of shape representations in monkey visual cortex. Neural responses in macaque inferior temporal cortex (IT) represent deformation of shape along a particular dimension (Kayaert, Biederman, & Vogels, 2003; Kayaert et al., 2005). Responses to rectangles and triangles increased monotonically as the shape was deformed to become further away from a prototypical mean in shape space (Kayaert et al., 2005). Additionally, multidimensional scaling demonstrated that IT neurons represented stimuli according

to two orthogonal dimensions of deformation – convexity and concavity – each extending out from a prototypical mean shape (see Figure 1.9A).

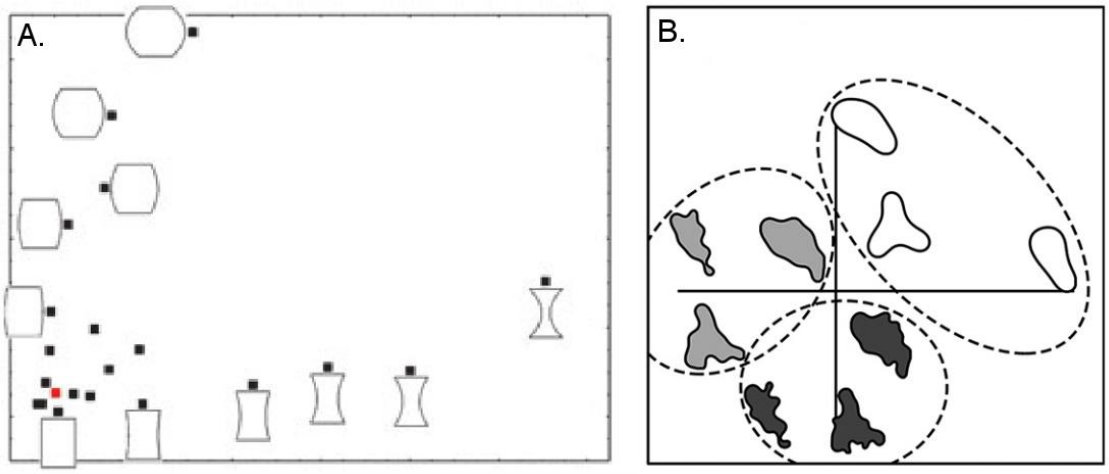

Figure 1.9: The results of a multi-dimensional scaling analysis performed on neural responses to different shape stimuli from two studies. (A) Macaque IT neurons represent stimuli along two dimensions – convexity (bottom left to bottom right) and concavity (bottom left to top left), centred around a prototypical rectangle (bottom left). Figure from Kayaert et al. (2005). (B) Human LOC represents shapes according shape complexity, defined by the number RF components within a stimulus. The low (white), mid (light grey) and high (dark grey) shapes are clustered together. Figure from Vernon et al. (2016).

A similar result has been found in the human lateral occipital complex (LOC); an object-selective area of human visual cortex which is generally considered to be a functional homologue of macaque IT (Grill-spector, Kourtzi, & Kanwisher, 2001; Malach et al., 1995). Measured with functional magnetic resonance imaging (fMRI), LOC responses to RF patterns increase monotonically with stimulus amplitude, or as they move further away from a circle in shape space (Rainville, Yourganov, & Wilson, 2010). This is consistent with proposed models of RF pattern perception, where RF patterns are represented by their deviation from a circular template (Loffler, 2008; Poirier & Wilson, 2006, 2010) and responses to shape deformation in macaque IT (Kayaert et al., 2003, 2005). Moreover, recent fMRI evidence shows that human LOC responds to shape complexity, defined by the number of RF components within a stimulus (Vernon et al., 2016, see Figure 1.9B), where responses were similar to stimuli with similar shape complexity. This provides further evidence that shapes may be

represented according to their position in a multidimensional shape space and that this space may be represented in human LOC.

In addition to lower-level shape stimuli, evidence suggests that higher-level, real-world objects might be represented in multidimensional space. Similar to LOC responses for RF patterns, the fusiform face area (FFA) – an area of visual cortex selective for face stimuli (Kanwisher, McDermott, & Chun, 1997) – shows responses that increase as a face stimulus, defined by a combination of multiple RF components, is morphed along a dimension in face identity space to become further away from a mean face (Loffler et al., 2005). As such, a multidimensional approach could be a general principle underlying stimulus representation.

Overall, research into RF pattern detection, discrimination and adaptation suggests these shapes may be represented by their curvature relative to the object centre. Poirier and Wilson (2006, 2010) provide a biologically plausible model for how this may be achieved by the visual system. Additionally, shapes may be represented in a multidimensional shape space. This is consistent with neural evidence for the representation of RF patterns, other shapes and high-level face stimuli defined by a combination of multiple RF components. RF patterns are therefore not only a useful stimulus for investigating the nature of shape processing, but also how these mechanisms may be applied to biologically relevant stimuli.

1.4: Possible neural loci for global shape processing in humans

1.4.1: Area V4

Models of shape processing often hypothesise area V4 as the most likely locus for global shape processing (e.g. Loffler, 2008; Poirier & Wilson, 2006). This notion appears to stem from single-unit recording of monkey V4 neurons. Such studies have found that V4 neurons respond preferentially to stimuli containing curvature compared to parallel gratings or other edge stimuli (Gallant, Braun, & Van Essen, 1993; Pasupathy & Connor, 1999). Further analysis of V4 neurons have revealed that they are tuned to the type and amount of curvature in a specific retinal location (Pasupathy & Connor, 2001).

Together, these tuning profiles provide a complex population code across V4 (Pasupathy & Connor, 2002) that is invariant to object size (El-Shamayleh & Pasupathy, 2016).

Correlational evidence for monkey V4's involvement in shape processing has been confirmed by causal evidence from lesion studies. Monkeys with a lesion administered to V4 suffer a deficit in pattern and orientation discriminations (Heywood & Cowey, 1987) as well as more complex object recognition (Schiller, 1995). Human V4 may possess similar properties to monkey V4. One human patient with a lesion to V4 showed deficiencies on a range of form tasks, including detection of deformation from circularity in RF patterns, but normal early vision (Gallant, Shoup, & Mazer, 2000).

Compared to in monkey, relatively little research has used neuroimaging to investigate the role of human V4 in shape perception. However there is some convergent evidence showing human V4 responds more strongly to concentric and radial gratings than parallel ones (Wilkinson et al., 2000). Importantly, this result has since been confirmed to be due to a sensitivity to the actual presence of concentric form and shape, rather than the amount of curvature present in the stimulus by a control experiment (Dumoulin & Hess, 2007, see Figure 1.10). Additionally, in this study neural selectivity for shape increased monotonically from V1 to V4, rather than being present only in V4. Therefore sensitivity to shape may emerge gradually through visual cortex from V1 to V4.

Overall, V4 is clearly important for the processing of curvature and shape in monkeys, and there is some convergent evidence to suggest the same is true in humans. This may arise from the steady increase of shape selectivity from V1 to V4 (Dumoulin & Hess, 2007). However, research into human V4's role in shape perception is limited, and a complex population code for shape in V4 that is invariant to low-level stimulus features such as size, like that of the monkey, (El-Shamayleh & Pasupathy, 2016; Pasupathy & Connor, 2002) has not been established in humans.

1.4.2: Lateral occipital areas LOC, LO1 and LO2

In monkey, a more abstract stage of shape processing than in V4 has been identified in IT. Neurons in this region exhibit shape tuning profiles in a multidimensional shape space (Brincat & Connor, 2004; Op De Beeck, Wagemans, & Vogels, 2001). This is

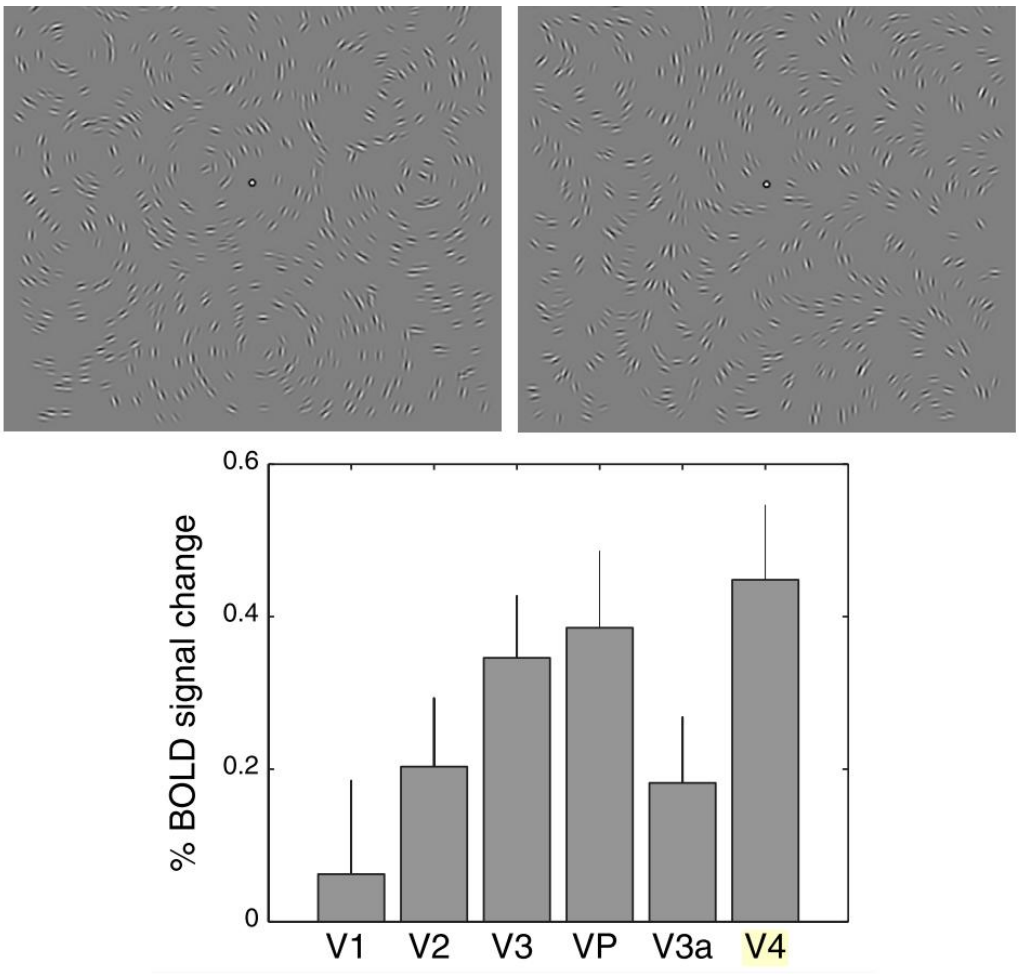

Figure 1.10: The preference in various visual areas (results at the bottom) for a field of Gabors containing concentric form and shape (top left) compared to one without form that is matched for curvature (top right). Generally, there is increasing selectivity for shape through the visual system, peaking in V4. Figure from Dumoulin and Hess (2007) .

supported by the aforementioned studies showing IT neurons adjust their responses according to shape manipulations along a given dimension (Kayaert et al., 2003, 2005, Figure 1.9A). LOC, an area of human lateral occipital cortex selective for object stimuli (Malach et al., 1995), is generally considered to be the human homologue of IT (Grill-spector et al., 2001). Aside from object stimuli, LOC also responds robustly to abstract shape and form (Kourtzi & Kanwisher, 2001; Stanley & Rubin, 2003). LOC may therefore be a more abstract and global stage of shape processing after V4, similar to IT in monkey.

Measuring neural adaptation, Kourtzi and Kanwisher (2001) demonstrated that LOC responses are abstracted beyond occlusion, treating a pair of stimuli containing

different contours but the same perceived shape (see Figure 1.11) just as similarly as it does two identical images. This suggests global shape representations, which abstract beyond an occluding object, in LOC. Moreover, evidence for a multidimensional shape space in monkey IT has been replicated in human LOC, where responses reflect shape similarity (Op de Beeck, Torfs, & Wagemans, 2008) and change parametrically according to the number of RF components and RF amplitude in composite RF pattern stimuli (Drucker & Aguirre, 2009; Vernon et al., 2016).

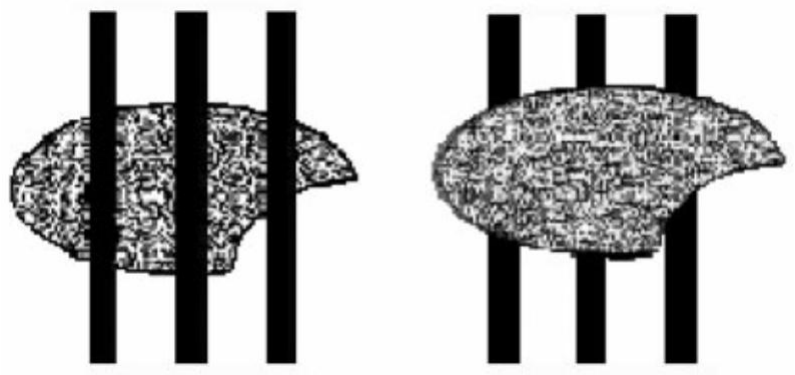

Figure 1.11: Human LOC responses to shape are identical in the presence of occlusion (left) compared to the same shape without occlusion (right), showing global shape representations in LOC. Figure from Kourtzi and Kanwisher (2001).

Typical definitions of LOC comprise a large area of cortex which responds more strongly to objects compared to scrambled versions of the same objects (Grill-spector et al., 2001; Malach et al., 1995; Op de Beeck, Haushofer, & Kanwisher, 2008). As such, the large area of cortex defined as LOC may contain submodules which perform dissociable functions. More posterior regions of LOC are selective for both objects and abstract shapes, while anterior LOC is selective only for real objects (Stanley & Rubin, 2005). This implies a gradient in the abstractness of representations where anterior LOC is more abstract, becoming increasingly selective for real-world objects. This is consistent with some level of overlap between the most posterior regions of LOC and two visual field maps called LO1 and LO2 (Larsson & Heeger, 2006; Sayres & Grill-Spector, 2008; Silson et al., 2016; Vernon et al., 2016). This overlap indicates that the most posterior, least abstract portions of LOC may also be retinotopic. Likewise, LO1 and LO2 may be involved in the analysis of object shape. Compared to V4, LO1 and

LO2 have been discovered more recently and do not have a clear homologue in monkeys, so research into their functions is more limited.

Recently, transcranial magnetic stimulation (TMS) has been used to temporarily disrupt processing within LO1 and LO2, allowing for the examination of their causal roles in visual perception (see Chapter 2 for a review of TMS). Silson et al. (2013) found that applying TMS to LO2 impaired the discrimination of amplitude modulations in RF patterns, while TMS to LO1 had no effect. On the other hand, TMS to LO1 impaired orientation discriminations of a grating stimulus, while LO2 TMS had no effect. This double dissociation indicates that computations in LO2 are critical for some aspects of shape perception, while LO1 is critical for the perception of orientation. This is also consistent with neural adaptation to orientation found in LO1 but not LO2 (Larsson, Landy, & Heeger, 2006).

One other study has examined shape processing in LO1 and LO2. Vernon et al. (2016) measured neural similarity across visual cortex to shape stimuli differing in shape complexity (defined by the number of RF components and their amplitude) and shape profile (defined by the location of the contrast describing the shape in the visual field). Early visual areas responded to shape profile, while later areas responded to shape complexity (see Figure 1.12). Both LO2 and LOC responded to shape complexity. LO1 was the only region that responded to both shape profile and shape complexity equally. Interestingly, no evidence of shape processing was found in V4, which responded more to shape profile than shape complexity. This suggests that representations in LO2 and LOC are more abstract than earlier visual cortex, as they respond to shape complexity over the physical profile of the stimuli. Moreover, LO1 may be a transitional point in the visual hierarchy where the information represented shifts from local, retinotopic information to more global representations of shape and curvature. This indicates that LO1 and LO2 may be the earliest visual areas that process global shape, with increasingly global representations in LOC that facilitate object recognition. However, the contribution of V4 in such a system is not clear.

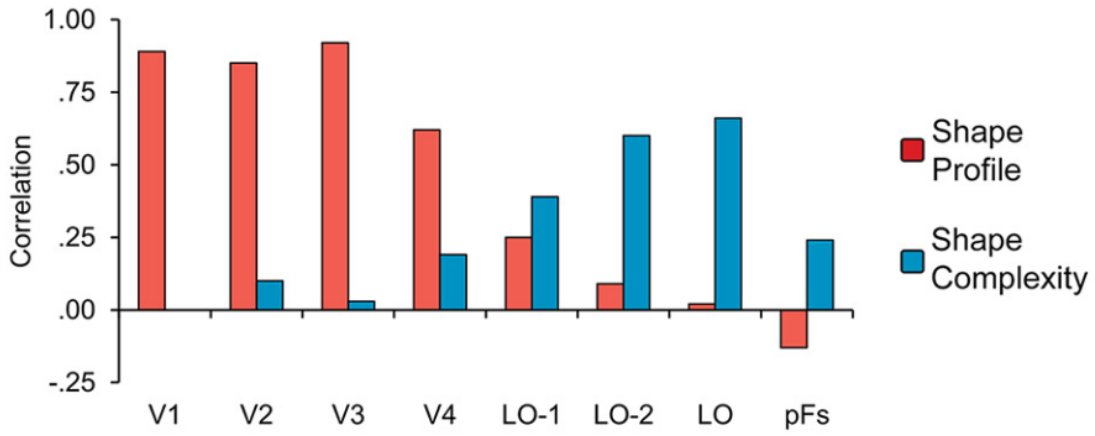

Figure 1.12: Correlations between neural responses in different areas of visual cortex and stimulus shape profile and shape complexity. Shape profile dominates in early visual areas and V4, while shape complexity dominates in LO2 and LOC (LO refers to object-selective LOC in this figure). LO1 responds to both shape profile and shape complexity evenly. Figure from Vernon et al. (2016).

1.5: Conclusions from the literature

Overall the evidence suggests that the analysis of shape is a distributed process across visual cortex. In monkeys, this process is hierarchically organised, starting with local feature extraction in V1 and V2, then a more complex analysis of contour and curvature in V4 and abstract representations of perceived shape in IT. The shape processing system in humans is similar to monkeys, with early processing in V1 and V2 and representations of global shape in LOC. However, the neural mechanisms underpinning intermediate levels of shape processing are less clear in humans. There is some evidence to suggest that human V4 shares some properties with macaque V4 (Dumoulin & Hess, 2007; Gallant et al., 2000; Wilkinson et al., 2000). However, others have not found evidence for shape processing in V4, and instead show that global shape representations emerge in LO1 and LO2 (Vernon et al., 2016). As such, the specific roles of human V4, LO1 and LO2 in the hierarchical progression of shape perception remain unclear.

1.6: Goals and outline of the thesis

The first objective of this thesis was to further characterise global shape processing mechanisms, which have received relatively little attention compared to lower- and higher-level visual mechanisms (Peirce, 2015). The second objective was to identify the neural locus of global shape processing mechanisms in humans. Thus far, most research into the neural processing of shape has been conducted on monkeys. Of particular interest were the roles of V4, LO1 and LO2; visual field maps which appear to be shape sensitive (Dumoulin & Hess, 2007; Silson et al., 2013; Wilkinson et al., 2000) and may form the basis for complex representations of objects in LOC.

Chapter 2 will review the methodologies employed for the research to follow. This will include psychophysical methods for measuring the response properties of shape processing mechanisms, neuroimaging methods for identifying retinotopic and category-selective areas of visual cortex, and fMRI-guided TMS for selectively stimulating visual areas to examine their causal roles in shape processing.

Chapter 3 describes a preliminary fMRI study which measures responses to RF patterns across visual cortex during the performance of different shape tasks. The primary goal of this chapter was to become familiar with fMRI methods, particularly using retinotopic mapping to identify key visual field maps in visual cortex. It also provided an opportunity to analyse brain responses to two tasks using a general linear model. The secondary goal was to determine whether responses to global shapes in LO2, which appear to be important for one global shape task (Silson et al., 2013), are dependent on the shape task being performed.

Chapter 4 describes the first example of a global shape adaptation effect in composite RF patterns. Composite RF patterns carry important shape information for real-world stimuli, and as such an aftereffect in composite RF patterns allows for new insights into biologically relevant shape processing. Moreover, using composite RF patterns allowed for new questions about interactions between different RF components and inferences about how RF patterns are represented in the brain. Using a shape aftereffect in composite RF patterns, a battery of psychophysical experiments was conducted to characterise global shape processing mechanisms and RF processing.

Chapter 5 describes the first attempt to impair shape discrimination and disrupt shape adaptation to composite RF patterns. TMS was used to investigate the causal roles of LO2 and occipital face area (OFA) in processing biologically relevant shape stimuli. LO2 has been shown to play a causal role in discriminating the shape of single RF patterns (Silson et al., 2013). OFA is a non-retinotopic, face-selective region of visual cortex which may process shape information that is important for face perception (Daar & Wilson, 2012; Fang, Murray, & He, 2007). As such, this experiment was designed to provide new insight into what stage in the visual hierarchy the biologically relevant shape information present in composite RF patterns is processed.

Chapter 6 measures fMRI responses to RF patterns across visual cortex and uses this data to model neural tuning to RF. Only low, rather than high RF information is informative about shape, and only low RF patterns engage global shape processing mechanisms (Hess et al., 1999; Jeffrey et al., 2002). Visual brain areas sensitive to global shape, therefore, should exhibit selective tuning to low RFs. This experiment describes the first attempt to extend the research modelling neural tuning to curvature and shape in monkeys (Pasupathy & Connor, 1999, 2001, 2002) to humans. The modelling technique introduced provides an opportunity to map out a representational shape space in the human brain. Finally, Chapter 7 will discuss the results of the research outlined above and its contribution to the field of global shape processing research.

Chapter 2: Review of methodologies

2.1: Radial frequency pattern stimuli

The research presented in this thesis used radial frequency (RF) pattern stimuli to investigate global shape processing mechanisms. To generate a RF pattern, the radius of a circle is modulated with a sinusoid as a function of the polar angle (see Figure 2.1). The radius r at polar angle θ can be defined by the following equation:

$$r(\theta) = r_0\big(1 + A\sin(\omega\theta + \emptyset)\big)$$

Where *r0* is the mean radius, A is the radial modulation amplitude, ω is the RF and $\emptyset$ is the phase of the RF pattern. This results in a broadly circular stimulus with a number of lobes equal to the frequency of the sine function. The prominence of those lobes is controlled by the amplitude, and phase controls the orientation of the pattern. In this thesis, when RF patterns are rendered, they are defined by a Gabor-like cross-sectional luminance profile (see Figure 2.1). The cross-sectional luminance profile of an RF pattern was defined by the following equation:

$$D4(d) = C\left(1 - 4\left(\frac{d}{\sigma}\right)^2 + \frac{4}{3}\left(\frac{d}{\sigma}\right)^4\right) \times exp\left(-\left(\frac{d}{\sigma}\right)^2\right)$$

Where *D4* is the fourth derivative of a Gaussian (centred on the stimulus contour), d is the distance between the pixel and the stimulus contour along the line with the shortest distance between the pixel and the contour, *C* is the contrast, *r0* is the mean radius, and σ determines the spatial frequency (σ is equal to half the peak wavelength of the D4 luminance profile). Defining stimulus contours in this way makes an easily detectable stimulus with an average luminance that is equal to the background the stimulus is presented on. One important feature of how stimuli were rendered is that the Gaussian luminance border was applied in the Cartesian, rather than polar, coordinate space. Other studies of RF patterns define the luminance profiles of their stimuli in polar coordinates (e.g. Salmela, Henriksson, & Vanni, 2016), where the luminance for each pixel is determined by its distance from the stimulus contour along the radial line according to a D4 function. This method does not account for the fact that, in Cartesian coordinates, two pairs of pixels with the same polar angle relationship can be further apart from (or closer to) each other depending on their radii. This results in a pinching

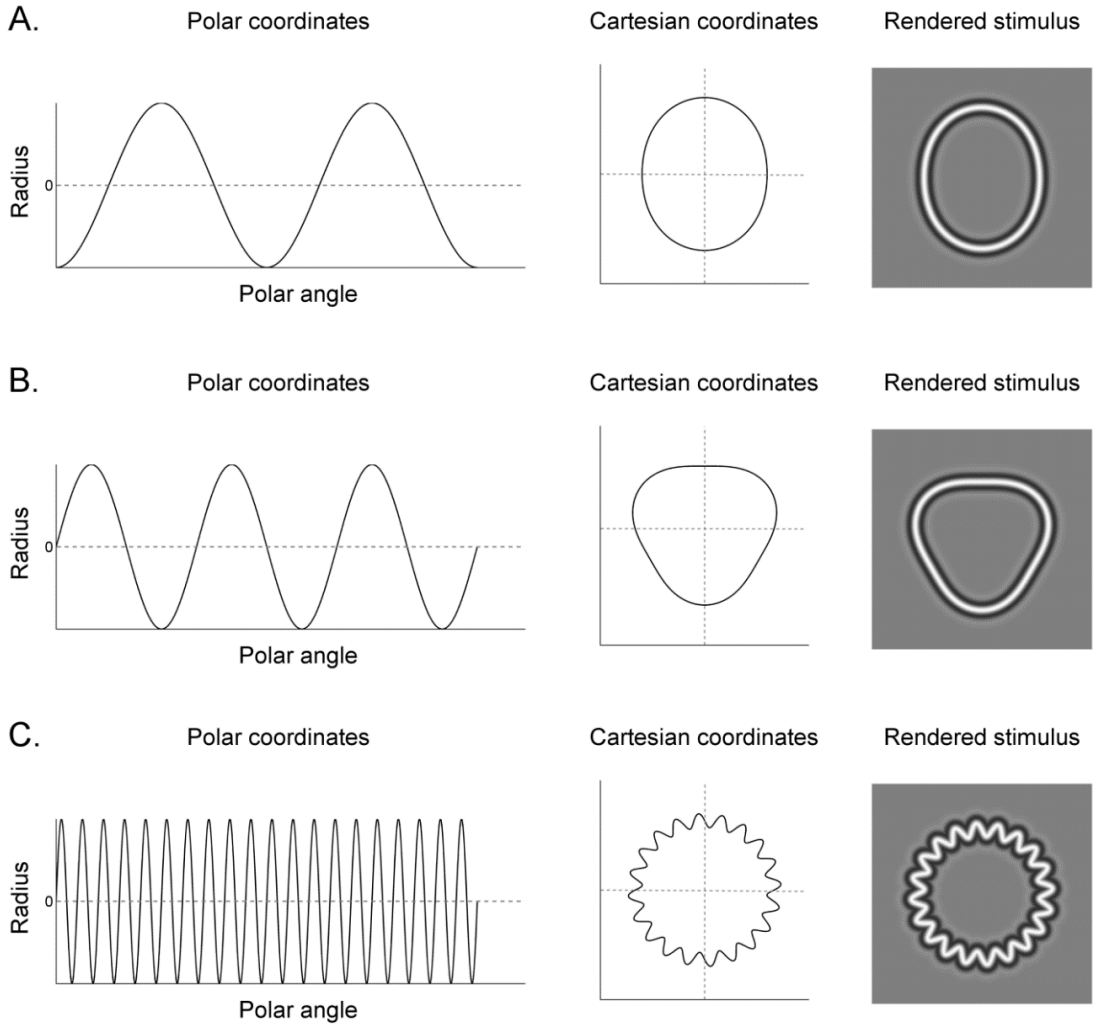

Figure 2.1: Schematics showing how an RF pattern is generated for three examples. The radius relative to polar angle is defined by a sinusoid in polar coordinates. In Cartesian coordinates, this produces a circle whose radius is modulated by the sinusoidal function, where the number of lobes and their prominence are controlled by the sine wave's frequency and amplitude, respectively. When the stimuli are rendered, the cross-sectional luminance profile of the stimulus is defined by the fourth derivative of a Gaussian, meaning the average luminance of the stimulus is equal to the background. Schematics are shown for (A) an RF2 pattern, (B) an RF3 pattern and (C) an RF20 pattern.

of the luminance profile towards the stimulus centre in high RF or high amplitude stimuli, where sharp contour modulations occur over fewer pixels towards the centre (see Figure 2.2A for an example). The consequence of this distortion is an inhomogeneous spatial frequency profile within the stimulus, as well as inconsistencies across stimuli. This distortion was corrected for by applying a luminance border in Cartesian coordinates, where the luminance of each pixel was determined by its distance from the contour along the line with the shortest distance from the contour

(see Figure 2.2A). This had the effect of removing most of the variance between different RF patterns, as shown by increased pairwise correlations between the 2D Fourier spectra of different RFs (Figure 2.2B). Stimuli were defined in this way for all experiments except the one reported in Chapter 3, for which stimuli were defined using the same methods as in the literature. This was particularly important when measuring neural responses to RF patterns, as another study has suffered from stimulus spatial frequency as a nuisance variable that dominated neural responses to RF patterns (Salmela et al., 2016).

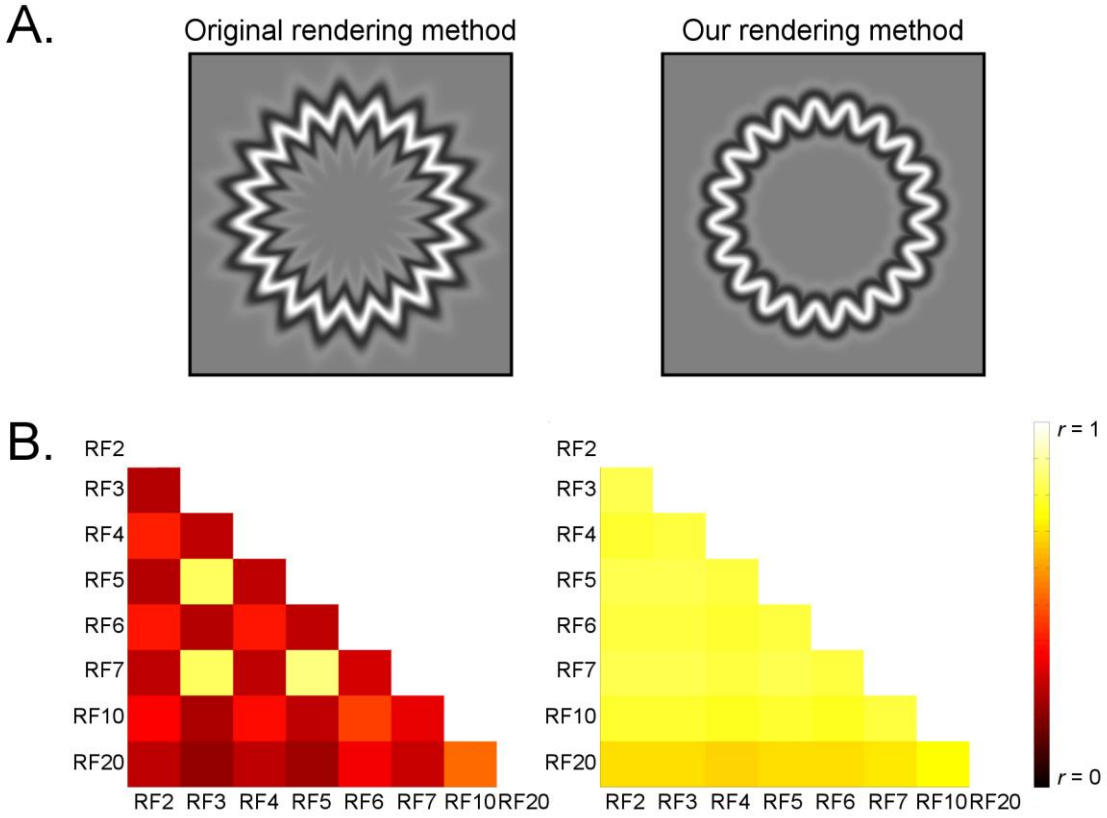

Figure 2.2: Improvements to RF pattern rendering. Panel A shows two RF20 patterns rendered using different methods. In the left stimulus, the luminance profile was defined in polar coordinates along radial line, causing pinching towards the centre of the stimulus. In the right stimulus, the luminance profile was defined in Cartesian coordinates along the line with the shortest distance to the stimulus contour, resulting in no pinching. Panel B shows pairwise correlations between the 2D Fourier spectra of several RF patterns (ranging for RF to RF20) for the original rendering method (left) and the improved method (right).

Deforming a circle with a low frequency sine function has a large effect on the shape of an RF pattern, producing clearly recognisable shapes. For example a 2 frequency RF pattern (RF2 pattern) can take on the shape of an ellipse (Figure 2.1A), while an RF3

pattern can be triangular (Figure 2.1B). High frequency modulations introduce local contour noise around the stimulus without changing its global shape. As such, high frequency RF patterns appear like circles defined by a wobbly contour (Figure 2.1C). This principle of RF patterns will be utilised in Chapter 6, which measures neural tuning to RF, expecting global shape processing regions to be more sensitive to low RFs than high RFs.

As well as simple shapes, multiple RF components can be summed to generate complex shapes representative of real-world stimuli. This method has some history in both human and computer vision literature. Fourier descriptors have been used as a method for describing and generating closed contour shapes (Zahn & Roskies, 1972). This work used the principle that any closed contour shape can be described as a combination of sine waves. Generating shapes using Fourier descriptors, therefore, is analogous to combining RF components, and has been used to generate RF-pattern like stimuli (Drucker & Aguirre, 2009). These methods can be used to generate mathematically defined stimuli that reflect real world objects. For example (Vernon et al., 2016) used Fourier descriptors to generate synthetic animals, and Wilson and Wilkinson (2002) accurately described human head outlines using a combination of relatively few RF components. As such, the use of composite RF patterns allowed for the interrogation shape processing mechanisms that may be important for perceiving real-world objects.

Figure 2.3 shows two examples of composite RF patterns comprising 2 RF components. For these patterns the radius r at polar angle θ can be defined by the following equation:

$$r(\theta) = r_0(1 + A_1 \sin(\omega_1\theta + \emptyset_1) + A_2 \sin(\omega_2\theta + \emptyset_2))$$

Where r_0 is the mean radius, A_1 and A_2 are the radial modulation amplitudes, ω_1 and ω_2 are the radial frequencies and $\emptyset_1$ and $\emptyset_2$ are the phases of the two RF components. Figure 2.3A shows an example of an RF2+3 pattern, generated by summing two sine waves with frequencies of two and three (Individual RF2 and RF3 components are showing in Figures 2.1A and B, respectively). Figure 2.3B shows an example of an RF2+20 pattern, generated by summing RF2 and RF20 components (shown individually in Figure 2.1A and C, respectively). As discussed in Chapter 1, low RFs

carry more shape information, meaning adding a low RF component to a RF pattern has a larger effect on its shape than a high RF component. Chapter 6 will utilise this to examine the neural locus of global shape processing in humans, finding areas of the brain which are invariant to high RF information which changes the contour on a local scale without affecting global shape. In addition, Chapter 4 will measure shape adaptation effects in RF2+3 patterns, exploring how global shapes and RF content might be represented in the brain.

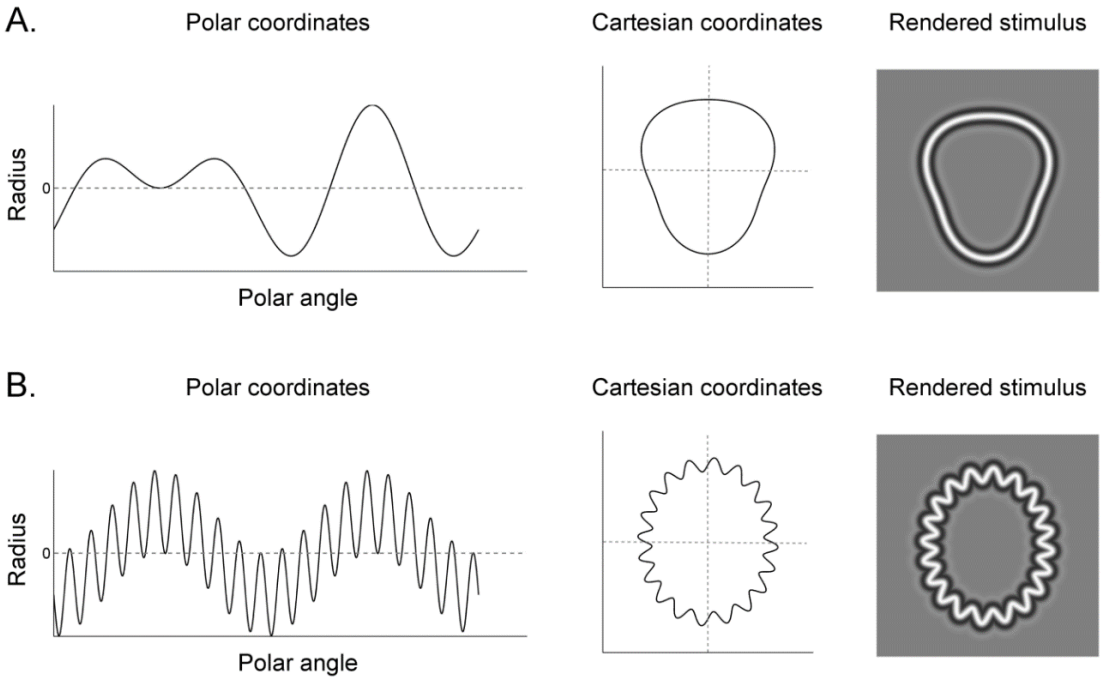

Figure 2.3: Schematics of two composite RF patterns each comprising two RF components. Panel A shows an RF2+3 pattern, generated by summing the sine functions shown in Figures 2.1A and B. Panel B shows an RF2+20 pattern, generated by summing the sine functions from Figures 2.1A and C.

2.2: Psychophysical methods

This thesis will employ psychophysical methods to measure shape discrimination of and adaptation to RF patterns. By taking precise measurements of participants' discrimination ability, tasks will be matched for difficulty, removing task difficulty as a confound variable when measuring brain responses. Shape adaptation will be used as a way of characterising the neural mechanisms

underpinning global shape processing with a battery of behavioural measures (how adaptation can be used to infer the response properties of underlying neural

mechanisms is reviewed in Chapter 1). To these ends, both staircase methods and the method of constant stimuli will be employed.

2.2.1: Staircases

Staircases are an adaptive way of adjusting the difficulty of a psychophysical task according to the subject's responses (Cornsweet, 1962). A staircase controls the magnitude of an independent variable to measure the subject's sensitivity to that variable. For example, consider an amplitude discrimination task where the subject is presented with a reference RF pattern with a fixed amplitude and must discriminate whether a test RF pattern is higher or lower amplitude than the reference. In such a task, the amplitude of the test stimulus would be controlled with a staircase function. To begin with, the amplitude of the test pattern is considerably higher than the reference and is easily discriminable from the reference (see Figure 2.4). When the subject correctly responds that the test is higher amplitude than the reference, the test amplitude is reduced according to the step size of the staircase, making test amplitude more similar to reference amplitude and the task more difficult. When the subject changes their response, responding that the test is lower amplitude than the reference, the staircase changes direction, increasing the amplitude of the test pattern (see red asterisks in Figure 2.4). These staircase reversals are coupled with a reduction in step size, so future adjustments to stimulus amplitude are smaller and the staircase converges on the test amplitude required for 50% correct performance (also known as point of subjective equivalence or PSE; the point at which the test is indiscriminable from the reference).

Different staircase functions can be used to converge on different levels of performance. The staircase shown in Figure 2.4 and described above is a one-up one-down staircase, meaning test amplitude is reduced after one higher amplitude response and increased after one lower amplitude response. Figure 2.5 shows a psychometric function measured from data from several different staircase functions. In this example, a psychometric function was generated by plotting the proportion of higher amplitude responses against the amplitude of the test stimulus and fitting a cumulative Gaussian to the data. One-up one-down staircases converge on 50% correct, providing data for the middle part of the function. Two-up one-down, three-up

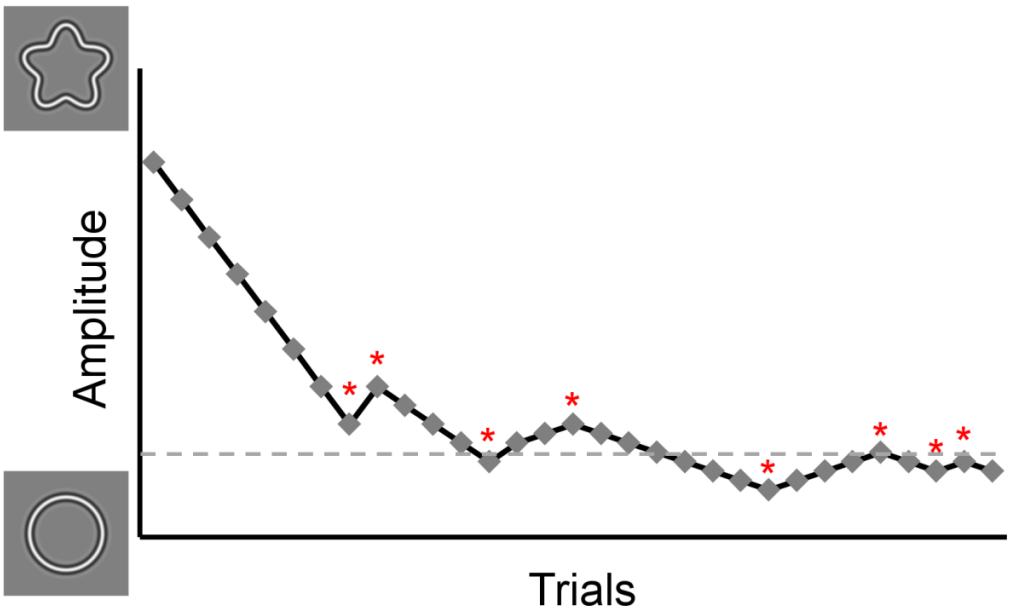

Figure 2.4: A schematic of how a staircase function controls the amplitude of an RF pattern in an amplitude discrimination task. In a task where the subject discriminates the amplitude of a test RF pattern relative to a reference pattern, the staircase controls the amplitude of the test pattern. Red asterisks demonstrate reversals where the staircase changes direction, which is coupled with a reduction in step size allowing the staircase to converge. The dashed grey line shows the amplitude of the reference stimulus. The staircase converges on chance performance (50%) correct; when the amplitude of the test RF pattern was similar to the reference pattern.

one-down, two-down one-up and three-down one-up staircases have also been used to collect data spread around the rest of the function. In this example, the up rule determines how many sequential lower amplitude responses are required for the staircase to increase test amplitude and the down rule determines how many higher amplitude responses are needed to decrease test amplitude. Essentially, a staircase converges on the stimulus level where the two reversal rules are equally likely to be met. So, in this example a two-up one-down staircase would converge on the amplitude where two lower amplitude responses are equally as likely as one higher amplitude response. The level of performance this will converge on can be determined by the equation:

$$p = 0.5^{\frac{a}{b}}$$

Where p is the proportion of correct responses (between 0 and 1), a is the smaller reversal rule and b is the larger reversal rule. Therefore, a two-up one-down (or two-

down one-up) staircase converges at $0.5^{\frac{1}{2}}$, or 0.71, or 71% correct. In this example, staircases with a high up rule (e.g. three-up one-down) begin with a lower test amplitude, while staircases with a high down rule (e.g. three-down one-up) begin with a higher test amplitude, so staircases converge towards the desired performance for test stimuli that are both lower and higher amplitude than the reference. Using this method, one can estimate the stimulus amplitude required for the subject to perform at a certain percent correct. This is called a threshold; Figure 2.5 shows 50% and 75% correct thresholds extracted from the psychometric function for demonstrative purposes.

As well as measuring discrimination and extracting thresholds for a desired level of performance, staircases can also be used to measure the effects of perceptual adaptation. For example, in the amplitude aftereffect adaptation to a high amplitude RF pattern causes a circle to be perceived as a higher amplitude RF pattern with antiphase modulations to the adapter (Anderson et al., 2007). As a result, an RF pattern with modulations in phase to the adapter is perceived as a circle, as perceptual mechanisms have become less sensitive to modulations in phase to the adaptor and more sensitive to modulations in antiphase to the adaptor. The magnitude of this effect can be measured using one-up one-down staircases. As shown in Figure 2.5, these staircases can be used to measure 50% correct performance, also known as the PSE. In the task used by Anderson et al. (2007) subjects discriminated the phase of amplitude modulations of an RF pattern which were centred around a circle. PSE therefore measured the amplitude needed for the stimulus to be perceived as a circle. As such, the difference in PSE between the adaptation condition and a control condition without adaptation is a measurement of the change in amplitude required for an RF pattern to be perceived as a circle post adaptation. The magnitude of the effect of adaptation on perception can therefore be quantified by the shift in PSE between adaptation and baseline conditions. In Chapter 4 this method will be used to measure shape aftereffects in composite RF patterns.

2.2.2: The method of constant stimuli

The method of constant stimuli is an alternative method of task design for measuring a psychometric function. In staircases, the independent variable changes adaptively

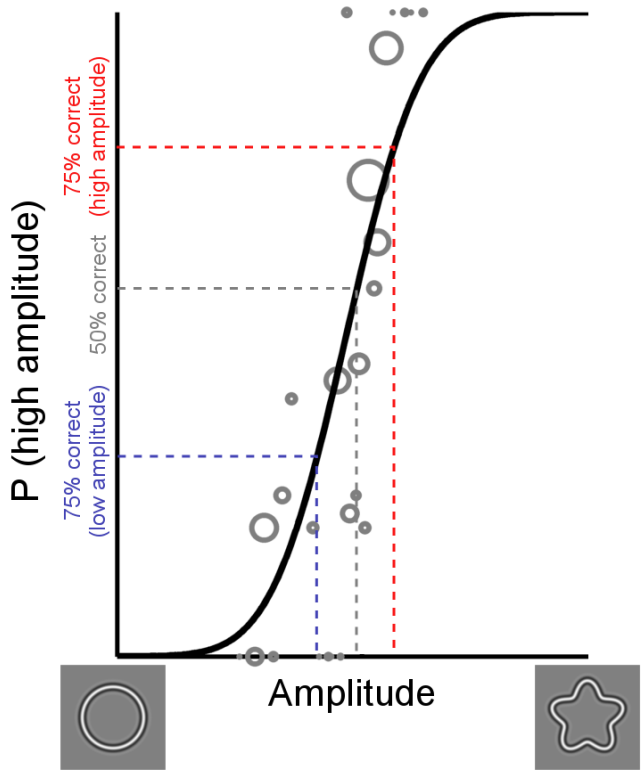

Figure 2.5: An example psychometric function generated from data from an RF pattern amplitude discrimination task using one-up one-down, two-up one-down, three-up one-down, 2-down one-up and three-down one-up staircases. The proportion of higher amplitude responses is plotted against the amplitude of the test stimulus and a cumulative Gaussian is fit to the data. The size of the data points is proportional to the amount of data acquired at that test amplitude. The Gaussian curve is used to define a 50% correct threshold (grey dashed line, this example assumes an ideal observer who performed at 50% correct when the amplitude of the test stimulus was equal to the reference with no bias towards higher or lower amplitude responses). 75% correct thresholds are also defined for low (blue dashed line) and high (red dashed line) amplitude test stimuli.

according to the subject's responses across trials whilst, in the method of constant stimuli, the stimulus levels to be presented are predetermined. Stimulus levels are chosen so that they may provide data spread across the psychometric function. In the example where a subject discriminates the amplitude of an RF pattern relative to a reference RF pattern, one could use three test stimulus amplitudes higher than the reference amplitude, three test amplitudes lower than the reference, and the reference amplitude itself. This would acquire data for a test than is indiscriminable from the reference, and stimuli with varying levels of discriminability for both lower and higher amplitude test stimuli. Once these data are acquired, a psychometric function can be plotted using the same methods as for staircases (Figure 2.5).

An advantage of the method of constant stimuli over staircases is that a large amount of data can be collected along the psychometric function with fewer trials, without trials that collect data beyond the ceiling of the psychometric function as happens in the early stages of a staircase. In addition, an accidental incorrect response from the subject in the early stages of a staircase results in a reversal which may reduce the step size, meaning the staircase will take longer to converge on the desired threshold, resulting in a large amount of data that does not contribute to the measurement of the psychometric function. However, a disadvantage of the method of constant stimuli is that pilot sessions must be conducted to judge the subject's discrimination ability and select stimulus levels that will provide informative data for their psychometric function. An adaptive staircase, however, will naturally adapt to the subject's ability and converge on their own individual performance threshold. A second advantage of staircases is that specific points on the psychometric function can be targeted with appropriate reversal rules. For example, in an adaptation experiment the PSE may be the only point of interest, so 1-up 1-down staircases can be used to collected most data around a 50% threshold. In Chapters 3, 4 and 5 both staircases and the method of constant stimuli will be used to measure shape discrimination thresholds and shape adaptation effects in RF patterns.

2.3: Neuroimaging methods

One of the objectives of this thesis was to identify the roles of different areas of human visual cortex in global shape processing. To this end, functional magnetic resonance imaging (fMRI) methods were used to localise these regions in individuals. This section will review how fMRI measures the blood-oxygen-level dependent (BOLD) response to make inferences about underlying neural activity, and how this can be used to identify both discrete maps of the visual field in retinotopic mapping, and regions of visual cortex that exhibit selectivity for one stimulus category over another in functional localiser experiments.

2.3.1: The BOLD signal

In an MRI scanner, spinning protons within the blood become aligned with the direction of the strong magnetic field of the scanner. When an image is acquired from

the scanner a radiofrequency pulse is released, which causes the protons to become aligned 90° away from the direction of the magnetic field; a higher energy state than magnetic field alignment. Following the pulse, proton magnetisation relaxes, returning to a low energy state and aligning with the magnetic field once more. As they do so, the protons release radiofrequency energy which is measured by the scanner. In structural imaging, the time taken for the protons to return to a relaxed state following a radiofrequency pulse is measured; known as a T1 contrast.

Following a radiofrequency pulse, the spins of neighbouring protons are perfectly in phase. Over time, protons naturally become out of phase as some spin faster than others, reducing the strength of the radiofrequency signal given off by the protons. This signal decay is known as T2 decay, and is measured in T2 contrast imaging. The frequency of a proton's spin is defined by the strength of the magnetic field. As such, any local inhomogeneities in magnetic field strength increase the rate at which nearby protons become out of phase. This increase in T2 decay resulting from local magnetic field inhomogeneities is call T2*, and is measured in functional imaging.

T2* is informative about brain function because it is correlated with neural activity. An increase in neural activity is followed by an increase in blood flow localised to the region of increased activity. This blood flow brings with it deoxygenated haemoglobin (deoxyhaemoglobin), which has magnetic properties. The magnetic properties of deoxyhaemoglobin create a local inhomogeneity in the strength of the magnetic field which, as described above, increases the rate of T2 decay. An increase in T2*, therefore, is a signal of the changes in local blood oxygenation that arise from neural activity, and is thus referred to as the blood-oxygen-level dependent or BOLD signal.

Critically, the relationship between neural activity and BOLD signal is largely linear (Logothetis, 2003), meaning a doubling of neural activity results in double the BOLD signal. This linearity allows for meaningful inferences about neural activity from measurements of BOLD, and is the basis of many fMRI analysis techniques (some of which are reviewed in the following sections). It is, however, important to note that BOLD is correlated with, but not necessarily causally related to, neural activity. fMRI therefore measures a correlate of neural activity and does not allow for causal

inferences about whether the computations performed within a given brain region are necessary for accurate perception or processing of a stimulus.

Since the BOLD response is underpinned by changes in blood flow, the temporal resolution of fMRI is relatively slow. Following a spike of neural activity, the BOLD response typically takes 4-6 seconds to reach its peak, and does not reach baseline until 12-16 seconds after the event. The sluggish change in BOLD following neural activity can be characterised by a hemodynamic response function (HRF, see Figure 2.6). fMRI is therefore not well suited to measuring the brain's response to events that are close together in time. This is accounted for in experimental design by ensuring events (or blocks of events) are sufficiently separated from each other in time to ensure that the source of each HRF can be determined. The strength of fMRI is in its spatial resolution. For a 3 Tesla MRI scanner, as was used for this thesis, voxels are typically 2-3 millimetres on each side. As such the BOLD response can be measured with a high degree of spatial precision, allowing for inferences regarding the precise locus of the neural response to a stimulus.

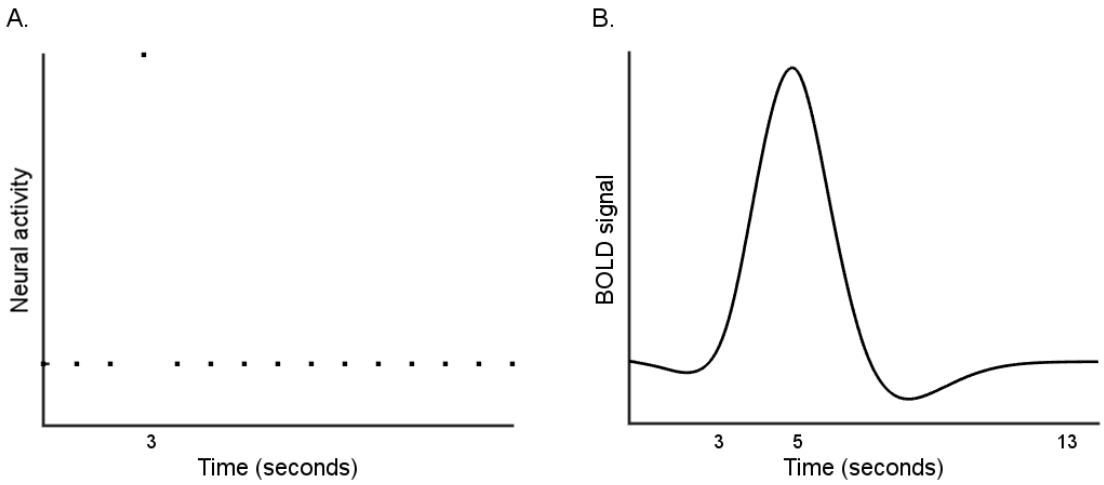

Figure 2.6: An example hemodynamic response function (HRF) that characterises the changes in BOLD signal following a hypothetical pattern of neural activity. Panel A shows the neural activity. The activity remains at baseline except for a transient spike that begins and ends at 3 seconds. Panel B shows the resulting BOLD signal from the neural activity in panel A. Due to its sluggish nature the BOLD signal peaks at 5 seconds and returns to baseline at 13 seconds. The BOLD signal in this example was calculated by convolving the neural activity in panel A with the default HRF for a TR of 3 seconds from the SPM 8 software package.

2.3.2: Retinotopic mapping of human visual cortex

As reviewed in Chapter 1, the human visual cortex contains multiple maps of the visual field. The organisation of these retinotopic maps can be measured in individuals using fMRI (Engel, Glover, & Wandell, 1997; Engel et al., 1994; Sereno et al., 1995). This is achieved by measuring the changes in BOLD associated with the passive viewing of visual stimuli explicitly designed to map out the visual field. Retinotopic mapping stimuli comprise an expanding ring stimulus designed to measure representations of eccentricity, and a rotating wedge stimulus designed to measure representations of polar angle (Wandell, Dumoulin, & Brewer, 2007, see Figure 2.7 for example stimuli, analysis and results).

Retinotopic mapping utilises the spatial precision of fMRI. Its high spatial resolution allows for a detailed measurement of how responses to visual field stimulation change across visual cortex. Responses to retinotopic mapping stimuli can be analysed using a phase encoded analysis. As a stimulus such as an expanding ring sweeps across the visual field, a voxel's response increases when the stimulus is in its receptive field and decreases when it is elsewhere. Due the shape of the BOLD signal, voxels in visual cortex yield a sinusoidal response as the stimulus sweeps across the visual field (Figure 2.7C). For an expanding ring stimulus, then, the phase of this sinusoidal response is indicative of a voxel's preferred eccentricity (Figure 2.7E). Likewise, the phase of the response to a rotating wedge indicates a voxel's preferred polar angle. In this way, the representation of the visual field in visual cortex can be mapped using retinotopic mapping.

Early retinotopic mapping experiments confirmed that the primary visual cortex (V1) lies along the calcarine sulcus and contains a complete map of the contralateral visual hemifield (Engel et al., 1994, 1997; Sereno et al., 1995). The fovea is represented in posterior V1 on the occipital pole, and has a disproportionately large representation relative to its size in the visual field; known as cortical magnification (see Figure 2.7E). The representation becomes increasingly peripheral as one moves anteriorly along the calcarine. In terms of polar angle about the centre of the visual field, the upper vertical meridian is represented on the lower bank of the sulcus, and the representation

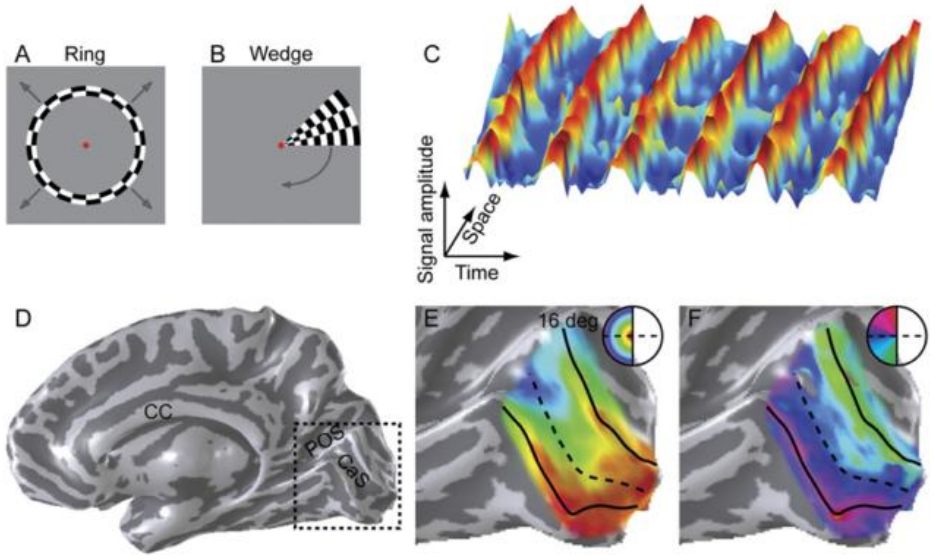

Figure 2.7: Example stimuli used in retinotopic mapping experiments and resulting data from the primary visual cortex (V1). (A) An expanding ring of a contrast reversing checkerboard stimulus begins in the centre of the visual field and gradually expands into the periphery. (B) A rotating wedge of a contrast reversing checkerboard stimulus rotates around the centre of the visual field. (C) Resulting time series data from voxels in visual cortex are analysed. The response oscillates in a sinusoidal fashion where it increases when the stimulus is in the receptive field of the voxel in question and returns to baseline when the stimulus is elsewhere. The phase of the of this sinusoidal response, combined with knowledge of the location of the stimulus as each point in time, provides information about what location in the visual field the voxel is responsive to. (D) Data are visualised on an inflated cortical surface. For right V1, data from the calcarine sulcus along the medial surface of the right occipital lobe are shown. (E) Responses in right V1 to the expanding ring stimulus. The different colours represent the eccentricity of the ring when voxels were maximally active. It can be seen that the fovea is represented at the posterior boundary of V1 on the occipital pole. As one moves anteriorly along the calcarine sulcus the representation becomes increasingly peripheral. Additionally, it can be seen that the fovea is represented by a disproportionately large amount of cortex given its size in the visual field. This is known as cortical magnification. (F) Responses in right V1 to the rotating wedge stimulus. The colours represent the position of the wedge during maximal responses and therefore the representation of polar angle across V1. It can be seen that the upper vertical meridian of the visual field is represented on the lower bank of the calcarine sulcus, and as one moves across the sulcus the representation transitions smoothly towards the lower vertical meridian represented on the upper bank of the sulcus. Figure from Wandell et al. (2007).

progresses smoothly towards the lower vertical meridian on the upper bank of the calcarine sulcus (see Figure 2.7F). As such, V1 contains a full map of the contralateral visual hemifield.

In addition to V1, retinotopic mapping can be used to identify multiple visual field maps beyond V1 in extrastriate visual cortex. These distinct regions can be delineated according to their representations of polar angle (Sereno et al., 1995). As can be seen in Figure 2.8, the polar angle representation reverses at the boundaries of visual areas. From the lower bank to the upper bank of the calcarine sulcus, V1's representation of polar angle progresses smoothly from upper to lower vertical meridian. At the upper bank, the progression of polar angle reverses in direction, signifying the border between V1 and area V2d. Across V2d, the representation progresses towards the horizontal meridian, where it again reverses to signify the border with V3d. As such, areas V2d and V3d represent the contralateral lower quarterfield. Similarly, beyond the lower bank of the calcarine sulcus areas V2v and V3v represent the contralateral upper quarterfield. In this way, visual field maps can be delineated according to their representation of polar angle.

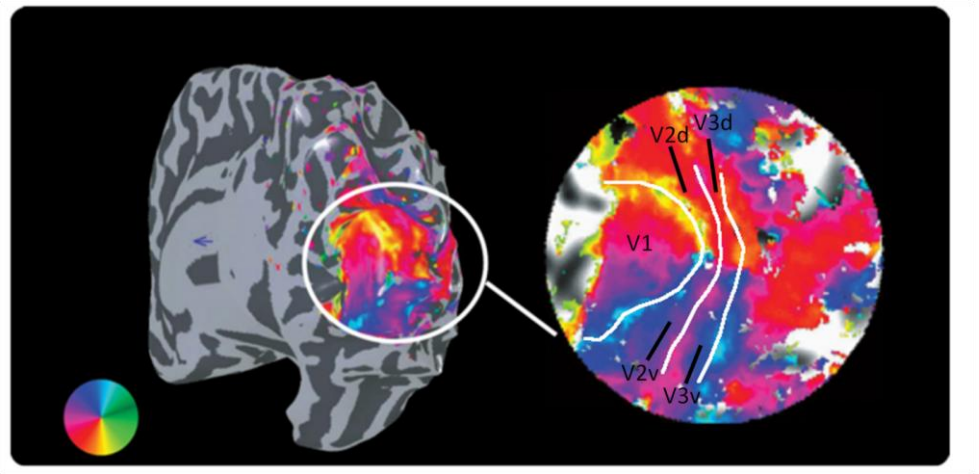

Figure 2.8: Representations of polar angle in right V1, V2 and V3. These regions are shown on an inflated cortical surface (left) and flattened cortical map (right). Visual field maps are delineated according to reversals in the representation of polar angle. It can be seen that all three areas contain a full representation of the contralateral hemifield, however V2 and V3 are split such that the more dorsal regions (V2d/V3d) represent the lower contralateral quarterfield while the ventral regions (V2v/V3v) represent the upper contralateral quarterfield. Figure adapted from Schira et al. (2009).

Several additional extrastriate visual field maps have been identified using retinotopic mapping (Amano, Wandell, & Dumoulin, 2009; Larsson & Heeger, 2006; Wandell et al., 2007), some of which are shown in Figure 2.9. Of particular interest to this thesis are areas V4, LO1 and LO2; potential candidates for global shape processing (reviewed in Chapter 1). Although it was debatable whether V4 represents a quarterfield or full

hemifield for some time, in recent years it has been established that V4 constitutes a full hemifield map extending from the border of V3v (Brewer, Liu, Wade, & Wandell, 2005; A. R. Wade, Brewer, Rieger, & Wandell, 2002; Winawer & Witthoft, 2015), and previous misconceptions were due to imaging artefacts from the nearby venous eclipse (Winawer et al, 2010). LO1 and LO2 are two hemifield maps extending from the border of V3d (see Figure 2.9). Throughout this thesis, using retinotopic mapping to identify visual field

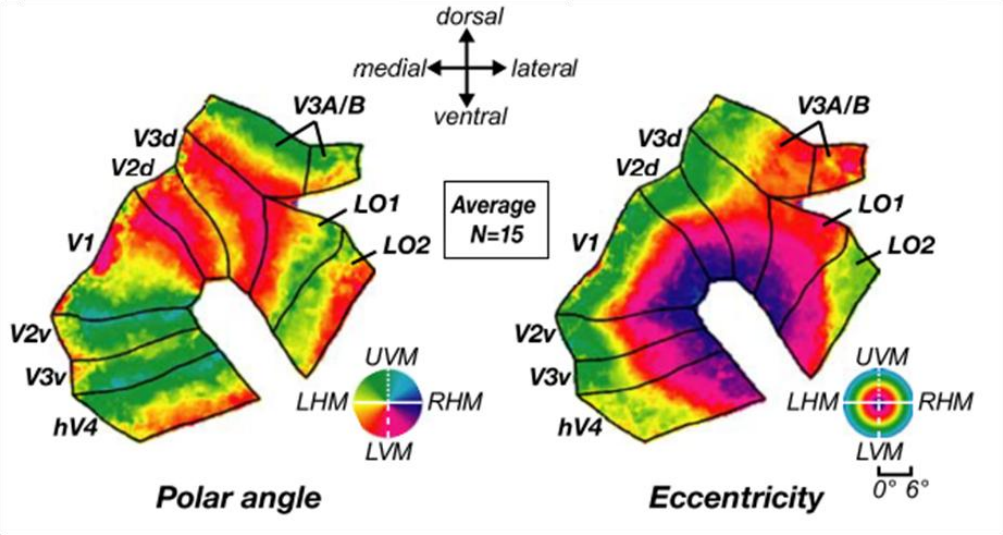

Figure 2.9: Average representations of polar angle and eccentricity in multiple visual areas from 15 participants. Beyond V3 a further five visual field maps can be identified. V3A and V3B adjoin the dorsal border of V3d at the lower vertical meridian. The border between V3A and V3B is somewhat unique in that it is identified by a phase reversal in the representation of eccentricity. Of particular importance to this thesis are areas V4, LO1 and LO2. V4 adjoins the ventral border of V3v at the upper vertical meridian. LO1 adjoins the lateral border of V3d at the lower vertical meridian and progresses to the upper vertical meridian where it borders with LO2. All the visual field maps described contain full representations of the contralateral hemifield. Figure from Larsson and Heeger (2006).

maps in individuals will be a critical step towards understanding their roles in global shape processing.

2.3.3: Localising category-selective areas of visual cortex

Beyond the visual field maps described above, the representation of retinotopy begins to break down and appears to no longer be the organisational principle of visual cortex (Op de Beeck, Haushofer, et al., 2008). Instead, later regions of visual cortex have been identified by their selectivity for different categories of visual stimuli. These regions are defined by functional localiser fMRI experiments where subjects are presented with

blocks of visual stimuli. Within each block a series of exemplars from a single visual category (such as faces, objects, scenes or a control stimulus) are sequentially presented. Category-selective regions of visual cortex are defined by a preferentially larger response to one visual category over others.

Block design experiments such as functional localisers measure the BOLD response to a series of exemplars from a stimulus category. Due to the sluggish nature of the BOLD signal, the response to individual exemplars cannot be distinguished. Instead, exemplars from the same category are presented in a block and the BOLD response across the block is an indication of a brain region's response to the presented stimulus category.

The preference of a given brain region for one stimulus category over another can be determined using a general linear model (GLM). Figure 2.10 shows an example of how a GLM can be used to analyse data from a block design experiment comparing responses to objects and scrambled objects. The stimuli alternate between blocks of objects and scrambled objects, separated by blocks of no stimuli (Figure 2.10A). A hypothetical BOLD response to these stimuli for a voxel with a preferential response to objects is shown in Figure 2.10B. One can determine whether the voxel in question shows a significantly larger response to objects than scrambled objects using a GLM with two predictors. The first predictor predicts the response for a voxel that responds only to objects (Figure 2.10C). The second predicts the response for a voxel that responds only to scrambled objects (Figure 2.10D). The predictors are then scaled to best fit the voxel's response. In this example, the BOLD response peaks at 10 with a baseline of 0, while the predictors peak at 1 with a baseline of 0. The response to objects peaks at 10, meaning the object predictor must be scaled by a factor of 10 to fit the data. The response to scrambled objects, on the other hand, peaks at 5, so the scrambled predictor is scaled by a factor of 5 to fit the data. These scalar values are called beta weights, and they define how well a predictor fits the overall response. The result of a GLM can be characterised by the following equation:

$$Y = G\beta + e$$

Where Y is the voxel response, G is a predictor, β is the scalar value that generates the best fit between G and Y and e is the left over error between G and Y. The significance

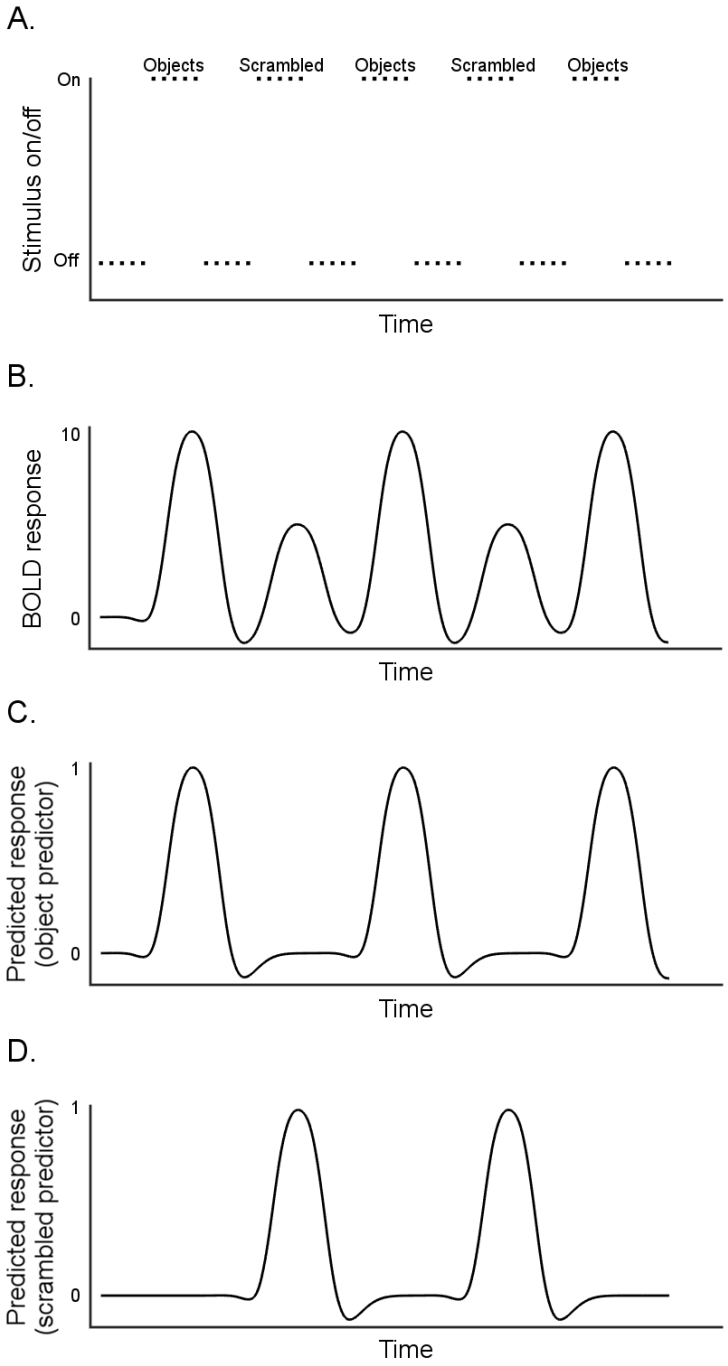

Figure 2.10: An example of how a GLM can be used to determine if a voxel responds more to one stimulus category over another. (A) Stimuli are presented in a block design – in this case blocks alternate between objects and scrambled objects, separated by blocks of no stimuli. Panel B shows a hypothetical BOLD response for a voxel that responds preferentially to objects compared to scrambled objects. A voxel's preference for objects or scrambled objects can be analysed using a GLM with two predictors. One predictor (C) predicts a response to objects only, while the other (D) predicts a response for scrambled objects only. See text for details on how a GLM can be used to statistically determine whether a voxel responds preferentially to one stimulus category over another.

of a voxel's response to a stimulus category (for example, determined by a t statistic) is proportional to $\frac{\beta}{e}$. It can be determined whether a voxel responds significantly more to one stimulus category than the other by comparing β for the two predictors.

For example, the lateral occipital complex (LOC) is an area of cortex that responds preferentially to objects over other visual stimuli (Malach et al., 1995; see Figure 2.11A). In their study, Malach et al. (1995) compared neural responses to objects against responses to texture patterns or other control stimuli. LOC consistently responded more to objects, whereas early visual cortex responses did not distinguish between objects and other stimulus categories (see Figure 2.11B). Since this study, comparing responses to objects and scrambled objects has become a standard method for localising LOC in individuals to measure the properties of object processing cortex (Grill-spector et al., 2001).

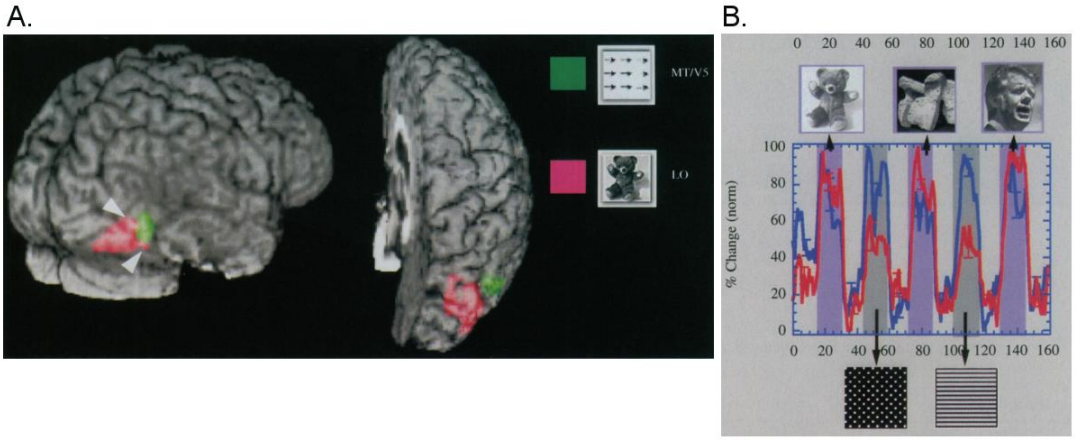

Figure 2.11: Results from an object localiser experiment identifying a region of visual cortex that responds preferentially to objects over other visual stimuli. (A) A region of object-selective cortex is shown in pink. In this thesis this region will be referred to as LOC which has become a standard label for object-selective lateral occipital cortex (Grill-spector et al., 2001), though it is labelled "LO" in this figure. Nearby motion-selective area MT/V5 is also shown in green. (B) Area LOC shows a preferential response to objects. The amplitude of the response in LOC (red) is larger for objects (top stimuli) compared to texture patterns (bottom). Conversely, responses in V1 (blue) are the same for all stimuli. Figure adapted from Malach et al. (1995).

Since Malach et al.'s (1995) first functional localiser study, the same method has been used to identify areas of visual cortex that are selective for other visual categories. For example, the fusiform face area, occipital face area (OFA) and superior temporal sulcus all respond preferentially to faces compared to other visual objects or control stimuli

(Gauthier, Skudlarski, Gore, & Anderson, 2000; Gauthier et al., 2000; Kanwisher, McDermott, & Chun, 1997). Together, these regions form a network for face processing and have been studied extensively (Haxby, Hoffman, & Gobbini, 2000). In addition to faces, regions selective to other visual categories such as scenes (Epstein & Kanwisher, 1998) and bodies (Downing et al., 2001) have also been identified. For the purposes of this thesis, functional localisers will be used to identify OFA and LOC. After they have been localised, the roles these more abstract, non-retinotopic visual areas play in global shape processing will be investigated.

2.4: Transcranial magnetic stimulation

Neuroimaging methods such as fMRI provide correlational data, meaning the causal role of a particular brain region for a given function cannot be determined. Transcranial magnetic stimulation (TMS) allows for such causal inferences by temporarily disrupting neural processing within the stimulated region. Measuring the effects of this disruption on performance in a behavioural task allows researchers to determine whether the neural computations within the targeted brain region are causally important for the behaviour measured by the task.

In TMS, a strong current is passed through a figure-of-eight coil producing a powerful, focal magnetic field projecting from the coil's centre and oriented orthogonally to the coil (see Figure 2.12). When the TMS coil is positioned against a subject's scalp, the magnetic field penetrates the scalp and induces a current in the underlying neural tissue. This current has two main effects on neural processing. In the case where the stimulated region is involved in the processing of a stimulus, TMS reduces stimulus-driven responses in those neurons responding to the stimulus. Additionally, TMS also increases the baseline fire rate of neurons not responsive to the stimulus, introducing additional noise to the neural signal (Allen et al., 2007; McKeefry et al., 2009). Together, these effects are thought to result in reduced signal to noise in the stimulated region.

There are two main approaches to the use of TMS: offline and online TMS. In offline TMS, a constant, low frequency train of TMS pulses (typically 1Hz) is applied for several minutes before the subject begins a task. In online TMS, TMS is applied during the performance of a task. This is usually done with repetitive TMS (rTMS), where a

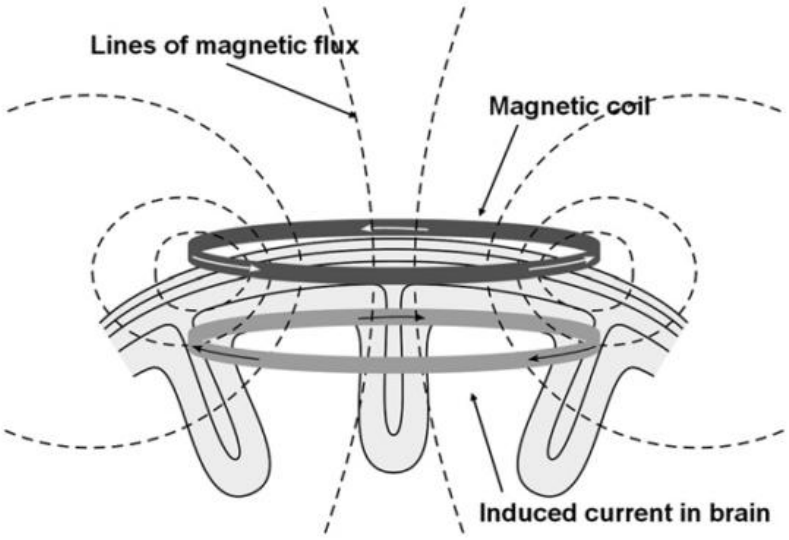

Figure 2.12: Schematic representation of the magnetic field and resulting current induced by a TMS coil. A strong current is passed through the coil (dark grey circle). This produces a magnetic field oriented orthogonally to the coil that penetrates the scalp (dashed grey lines show the lines of magnetic flux). The magnetic field induces an orthogonal current in the underlying brain tissue (light grey circle). Figure from Hallett (2007).

short train of high frequency pulses (e.g. 5 pulses at 20Hz) is applied with specific timing while the subject is engaging with a specific aspect of the task. Offline TMS has long lasting effects, and is generally used to measure the effect of TMS on general task performance (Bolognini & Ro, 2010). Online rTMS, however, has a more transient effect on neural processing, and can be used to affect neural processing of a single stimulus, or determine the time course of neural processes by varying the timing with which TMS is applied following the presentation of a stimulus (e.g. Pitcher, 2014). For the purposes of this thesis, online rTMS was applied during the perception of a stimulus, in order to disrupt the shape perception of that stimulus.

2.4.1: Applications of TMS: visual discrimination and adaptation

In the field of vision research, the effects of TMS to a particular visual brain region on visual perception are measured. One method is to measure the effect of TMS on a visual discrimination task. For example, as was reviewed in Chapter 1, Silson et al. (2013) found subjects performed worse in orientation and shape discrimination tasks when TMS was applied to LO1 and LO2, respectively. This therefore demonstrates that the neural computations performed within LO1 and LO2 are causally important for the accurate perception of orientation and shape. In this way, TMS can be used to

determine whether specific visual brain regions are employed for the discrimination of a visual feature.

In addition to discrimination, TMS can also be used to investigate the role of a visual area in visual adaptation. As reviewed in Chapter 1, adaptation is a powerful tool that allows researchers to probe the response properties of the neural mechanisms that encode the adapted feature. If TMS has an effect on perceptual adaptation effects, it implies that the stimulated region is where the adapted feature is encoded in the brain. For example, in colour aftereffects adaptation to one colour results in an aftereffect where a neutral stimulus appears to have the complimentary colour to the adapter (see the McCollough effect reviewed in Chapter 1). When high intensity TMS is applied to early visual cortex, illusory visual phosphenes are induced by the increased neural firing (Elkin-Frankston et al. , 2011). Silvanto et al. (2007) measured the effect of TMS on perceived phosphenes following colour adaptation. As expected, colour adaptation resulted in a colour aftereffect where a grey stimulus was perceived to be the complimentary colour to the adaptor. The application of TMS to early visual cortex during this visual aftereffect caused phosphenes that were the same colour as the adaptor. First, this demonstrates that TMS can interfere with the results of visual adaptation, facilitating the response of adapted neurons to reverse the effect of adaptation. Second, this demonstrates that colour is encoded in early visual cortex where TMS was applied. In this way, TMS can be combined with adaptation to determine whether the stimulated region is involved in the encoding of the adapted feature. This thesis will employ TMS as a method of investigating the roles of different visual areas in both global shape discriminations and adaptation.

2.4.2: fMRI-guided TMS

In order to measure the effect of TMS to a specific visual area on visual perception, it is essential that the targeted area is stimulated accurately and independently of neighbouring visual areas. This is particularly important for stimulating visual field maps which are relatively small and are flanked by other maps. Moreover, the organisation of visual field maps can be variable across individuals (Larsson & Heeger, 2006), meaning they must be identified in each individual before they can be stimulated. To this end, this thesis will employ the retinotopic mapping and functional

localiser methods described previously to identify the anatomical location of visual areas in each individual before stimulating them with TMS.

In addition to using fMRI to localise visual areas in individuals before stimulating them, fMRI-guided TMS can be used to measure the online location of the TMS coil relative to these anatomical locations to ensure that they are stimulated accurately. The TMS experiments described in this thesis will use the Brainsight software package and specialist hardware to measure the online location of both the subject and the TMS coil whilst the subject performs a visual task (see Figure 2.13A). An infrared camera monitors the location of a tracker attached to the subject's head and a second tracker attached to the TMS coil. First, the location of some key points on the subject's head (left and right ear, tip and bridge of the nose) are measured. These points are also identified on a structural MRI image of the subject's head, so that the Brainsight software can align the subject's head to their structural MRI image. Additionally, the location of the coil tracker relative to the orientation of the coil is calibrated so that the trajectory of the magnetic field can be monitored. The intended TMS target is also defined on the subject's structural image. As such, the trajectory of TMS relative to the TMS target can be monitored (see Figure 2.13B), so that the experimenter can make online adjustments to the TMS coil's position during the experiment, ensuring accurate stimulation throughout.

The research described in this thesis employs the methods reviewed in this chapter to characterise global shape processing mechanisms in humans. Psychophysical methods, including staircases and the method of constant stimuli, will be used to measure shape discrimination performance and shape aftereffects. Retinotopic mapping and functional localiser fMRI methods will be used to localise visual areas in individuals that may be implicated global shape processing. Once these areas have been localised, a combination of further fMRI experiments and fMRI-guided TMS will be used to interrogate the specific roles of these distinct visual areas in global shape processing.

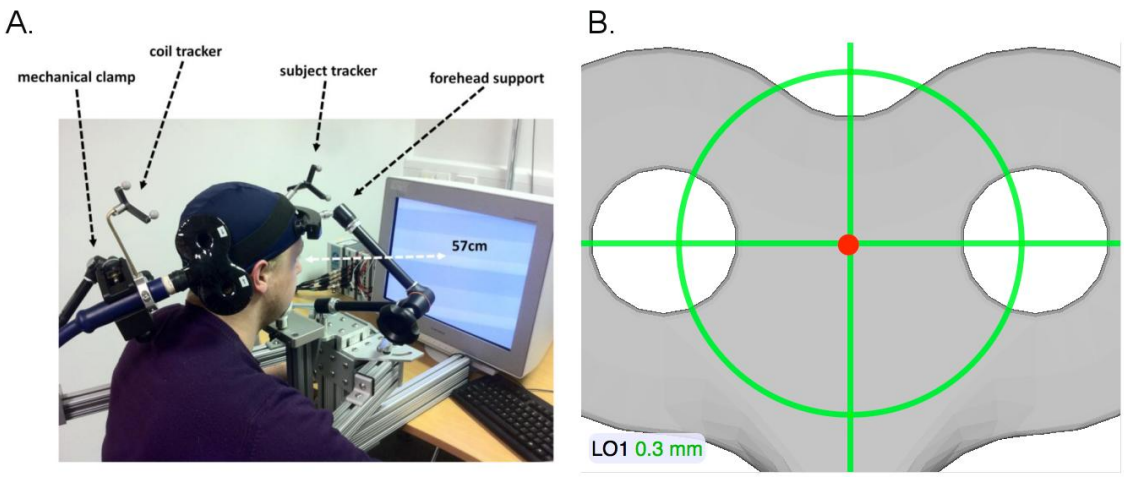

Figure 2.13: The setup of TMS equipment in the laboratory and a screenshot of the Brainsight software monitoring the location of TMS stimulation. (A) The subject's head is held in place with a forehead support, set up so that they are 57cm from a display showing visual stimuli. Trackers positioned on the subject's head and the TMS coil are monitored by an infrared camera during the experiment. A mechanical clamp is used to support and keep the TMS coil in place. Image from (Silson, 2013) (B) The Brainsight software package monitors the TMS trajectory relative to the subject's head so that the TMS target can be stimulated accurately. The image shows a screenshot of the software showing the current TMS trajectory (red dot) relative to the position of the target (green crosshairs). The name of the target (in this case LO1) is displayed in the bottom left along with the positional error of the TMS trajectory relative to the target in millimetres.

Chapter 3: Responses to radial frequency patterns in LO1 and LO2 during orientation and shape discriminations

Abstract

Radial frequency (RF) patterns are known to be processed by global, integrative shape processing mechanisms, however the neural locus of global shape mechanisms in humans is unknown. Visual field map LO2 is causally important for discriminating the amplitude of RF patterns (Silson et al., 2013) and overlaps with complex object-selective processes. This suggests global shape processing occurs in LO2, however it cannot be determined from previous research whether processing of RF patterns in LO2 is specific to amplitude discrimination or reflects more general shape processing. Responses were measured from LO2 and the immediately posterior area LO1 during orientation and amplitude discriminations of RF patterns. LO2 generally exhibited stronger responses to RF patterns than LO1, and this effect was not modulated by the task being performed. This suggests that LO2 performs general processing of global shape stimuli, implying it may be the neural locus of early global shape processing in humans.

3.1: Introduction

This chapter describes the first fMRI experiment, which measured responses from LO1 and LO2 to radial frequency (RF) patterns while subjects performed two shape tasks. The first objective for this experiment was to become familiar with the fMRI methods described in Chapter 2. This included both using retinotopic mapping to identify key visual areas within visual cortex, and using a general linear model (GLM) to compare neural responses to two experimental conditions. The second objective was to determine whether responses to global shapes in LO2 are dependent on the shape task being performed. The rationale for this is explained below.

RF patterns are a class of visual stimuli generated by deforming the radius of a circle with a sinusoidal function (Wilkinson et al., 1998). Many psychophysical experiments have demonstrated that human sensitivity to RF patterns can only been explained by integrative, global-scale shape processing mechanisms (Anderson, et al., 2007; Bell et al., 2014; Bell et al., 2010; Bell & Kingdom, 2009; Dickinson et al., 2015; Dickinson et al., 2012; Hess, Wang, & Dakin, 1999; Jeffrey, Wang, & Birch, 2002; Loffler, Wilson, & Wilkinson, 2003; Schmidtmann et al., 2012). While global shape mechanisms have be thoroughly characterised using psychophysical measurements, their locus in the human brain is largely unknown.

The extraction of global shape has been suggested to facilitate object recognition (Loffler, 2008). Regions of visual cortex which exhibit selectivity for objects may therefore be a likely candidate for global shape processing. The lateral occipital complex (LOC) is selective for objects (Malach et al., 1995) and exhibits responses to shape which are not affected by occlusion (Kourtzi & Kanwisher, 2001), suggesting some global representation of shape. The posterior portion of LOC overlaps with retinotopically defined visual field map LO2 and, to a lesser extent, LO1 (Larsson & Heeger, 2006; Sayres & Grill-Spector, 2008; Silson et al., 2016; Vernon et al., 2016). As such, LO1 and LO2 exhibit both retinotopic specificity and object selectivity, indicating they could be the transitional stage where representations of objects progress from a local, retinotopic scale to a more global scale. Since the research reported in this chapter was conducted, another study has found consistent evidence for a transition from local,

retinotopic representations to more abstract representations of shape in LO1 and LO2 (Vernon et al., 2016).

This notion is supported by evidence showing a marked difference in the sensitivities of LO2 and the immediately posterior visual field map LO1 to low level stimulus features such as orientation. Larsson, Landy and Heeger (2006) found neural adaptation to orientation in LO1, but not LO2, suggesting sensitivity to more complex visual features in LO2 compared to LO1. Moreover, a functional difference between LO1 and LO2 was confirmed using transcranial magnetic stimulation (TMS) by Silson et al. (2013). In this study, TMS was applied to LO1 or LO2 during orientation discrimination of sine-wave gratings or amplitude discrimination of RF patterns. A double dissociation was found where TMS to LO1 selectively impaired orientation discrimination, but not amplitude discrimination, while TMS to LO2 impaired amplitude, but not orientation discrimination. This demonstrates a clear difference in the processes performed by LO1 and LO2, and a causal role for LO2 in the discrimination of RF amplitude; a task which engages global shape mechanisms (Wilkinson et al., 1998).

As such, Silson et al.'s (2013) study suggests that LO2 is important for global shape processing. However, because different stimuli were used for the two tasks, it cannot be determined from this study whether the double dissociation was driven by more global shape processing in LO2 compared to LO1, or processing of RF amplitude in LO2, specifically. Their result, therefore, could either be due to differences in task, where LO1 processes orientation and LO2 processes RF amplitude, or differences in stimuli, where LO1 processes stimuli defined by local changes in contrast (gratings) and LO2 processes sparser, more global stimuli (RF patterns). In this study, functional magnetic resonance imaging (fMRI) was used to determine whether Silson et al.'s (2013) results were due to processing of orientation in LO1 and RF amplitude in LO2, or more global processing of shape in LO2 compared to LO1.

Blood oxygenation level dependent (BOLD) responses in LO1 and LO2 were measured during orientation and amplitude discriminations of a pair of three-lobed RF patterns (RF3 patterns). Stimuli always differed in both orientation and amplitude, meaning the

stimuli were identical across both orientation and amplitude tasks. This allowed for responses driven by the stimuli and responses driven by the tasks to be distinguished.

Alternative predictions based on the possible response properties of LO1 and LO2 are shown in Figure 3.1. If Silson et al.'s (2013) double dissociation was due to task selectivity, LO1 should respond more strongly to orientation discriminations than amplitude discriminations, while LO2 should respond more to amplitude discriminations (Figure 3.1A). On the other hand, if the double dissociation was due to stimulus selectivity, LO2 should show stronger responses to RF3 patterns than LO1, irrespective of the task being performed (Figure 3.1B).

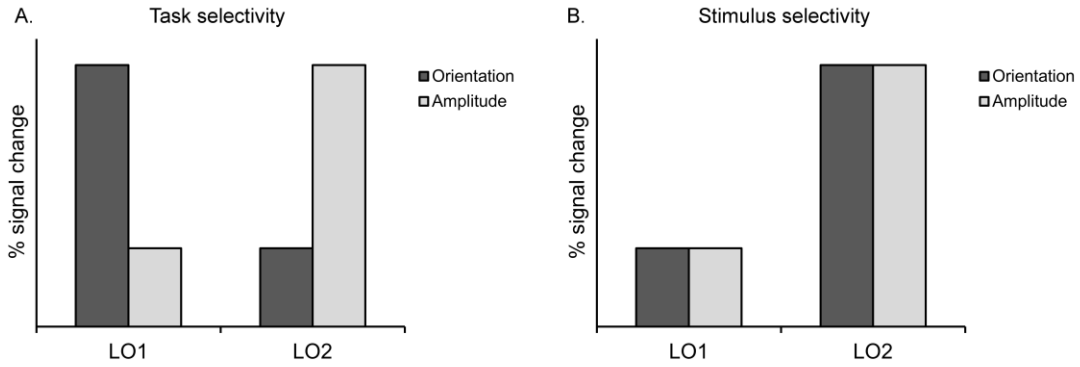

Figure 3.1: Hypothetical responses from LO1 and LO2 to discriminations of orientation (darker bars) and amplitude (lighter bars) of RF3 patterns. (A) If LO1 processes orientation and LO2 processes amplitude, LO1 should respond more to orientation discrimination than amplitude discrimination while LO2 should show the opposite. (B) If LO1 processes more local stimuli and LO2 processes RF patterns, LO2 should respond more than LO1 to RF3 patterns irrespective of task.

3.2: Methods

BOLD responses were measured during orientation and amplitude discriminations of RF3 patterns in eight subjects (1 female, mean age 25.9, age range 21-35). All subjects had normal or corrected-to-normal visual acuity and gave informed consent. The study was approved by the York Neuroimaging Centre Research Governance Committee and adhered to the declaration of Helsinki. All participants underwent one structural MRI session, one retinotopic mapping session and one main functional session, totalling three hours of scanning per subject. In addition, all subjects completed one to two hours of psychophysical testing which was used to define the stimuli for the main functional scan.

3.2.1: Preliminary data acquisition and analysis

All MRI data were acquired on a GE 3-Tesla Signa HD Excite scanner using a 16-channel half-head coil to improve signal-to-noise in the occipital lobe, except for one structural scan which was acquired with an 8-channel head coil to improve alignments between functional and structural data.

3.2.2: Structural scans

Three 16-channel T1-weighted anatomical images (TR = 7.8ms, TE = 3.0ms, TI = 600ms, voxel size = 1x1x1mm3, flip angle = 12°, matrix size 256x256x176, FOV = 25.6cm), one 8-channel T1-weighted image (TR = 7.8ms, TE = 2.9ms, TI = 450ms, voxel size = 1.13x1.13x1mm3, flip angle = 20°, matrix size 256x256x176, FOV = 29cm) and one T2*-weighted fast gradient recalled echo scan (16-channel coil, TR = 400ms, TE = 4.3ms, voxel size = 1x1x2mm3, flip angle = 25°, matrix size 128x128, FOV = 26cm) were acquired from each subject.

The three 16-channel structural scans were averaged together and divided by the T2* data to improve white-grey matter contrast and correct for the signal dropout of the 16-channel coil. The data were then automatically segmented into white and grey matter followed by manual corrections.

3.2.3: Retinotopic mapping

Each subject completed four to six rotating wedge and two expanding ring scans (TR = 3000ms, TE = 30ms, voxel size = 2x2x2mm3, flip angle = 90°, matrix size 96x96x39, FOV = 19.2cm). 90 degree counter-clockwise rotating wedges and expanding ring stimuli were high contrast (>98%, 400cdm^-2) checkerboards whose contrast reversed at a rate of 6Hz. Wedge and ring scans comprised 8 stimulus cycles (36 seconds per cycle) in which the stimulus traversed a circular region of 14.34 degrees of visual angle. Participants maintained fixation on a red central fixation cross throughout all scans. From the retinotopy data, visual field maps V1, V2, V3. V4, LO1 and LO2 were identified using published methods (Deyoe et al., 1996; Engel, Glover, & Wandell, 1997; Engel et al., 1994; Larsson & Heeger, 2006; Sereno et al., 1995; Wandell, Dumoulin, & Brewer, 2007; also reviewed in Chapter 2).

3.2.4: Tasks and stimuli

The stimuli used for the experiment were three-lobed RF (RF3) patterns. The radius r at polar angle θ was calculated using the following equation:

$$r(\theta) = r_0\big(1 + A\sin(\omega\theta + \emptyset)\big)$$

Where *r0* is the mean radius, *A* is the radial modulation amplitude, ω is the RF and $\emptyset$ is the phase of the RF pattern. All stimuli had an RF of 3 and mean radius of 2 degrees of visual angle, while amplitude and phase were variable. The cross-sectional luminance profiles of stimulus contours were defined using the same methods as others in the literature. In future chapters, stimuli will be defined using improved methods as reviewed in Chapter 2. For this study, luminance profiles were defined by the following equation:

$$D4(r) = C\left(1 - 4\left(\frac{r-r_0}{\sigma}\right)^2 + \frac{4}{3}\left(\frac{r-r_0}{\sigma}\right)^4\right) \times exp\left(-\left(\frac{r-r_0}{\sigma}\right)^2\right)$$

Where *D4* is the radial fourth derivative of a Gaussian, *C* is the contrast, *r0* is the mean radius, and σ determines the spatial frequency. All stimuli were presented at 50% contrast and the peak spatial frequency was 1.3 cycles per degree.

Subjects were presented with two RF3 patterns simultaneously (one 10 degrees of visual angle left of fixation, one 10 degrees right). One of the stimuli was a reference stimulus, which always had an amplitude of 0.2 and an orientation of 20, 40, 60, 80 or 100 degrees (chosen at random). The position of the reference was randomised so that it was left of fixation for 50% of trials and right for the other 50%. The other stimulus was a comparison stimulus which differed from the reference in both amplitude and orientation. There were four possible comparison stimuli that could be presented (shown in Figure 3.2): more counterclockwise and lower amplitude than the reference, more counterclockwise and higher amplitude than the reference, more clockwise and lower amplitude than the reference or more clockwise and higher amplitude than the reference. The comparison stimulus was randomised so that each possibility was presented on 25% of trials. For the orientation task, subjects indicated which of the two stimuli presented was more clockwise with a button press, for the amplitude task they indicated which had the higher amplitude.

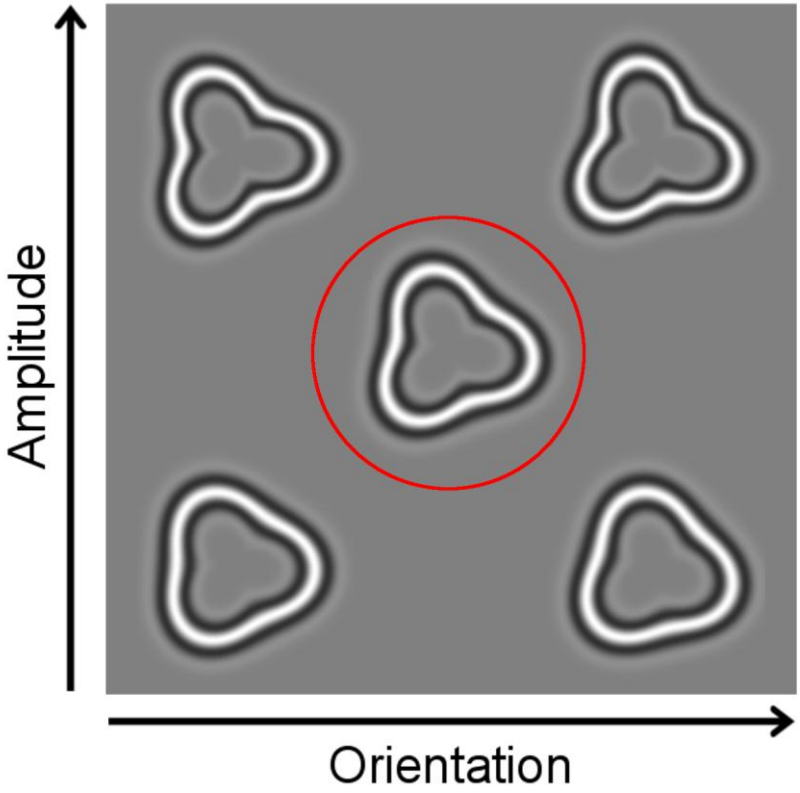

Figure 3.2: Example RF3 patterns used in the experiment. The reference stimulus (shown in the centre and encompassed by a red circle) was present on all trials. The other stimulus was one of four possible comparison stimuli. The comparison could either be counter-clockwise and lower amplitude than the reference (bottom left), contour-clockwise and higher amplitude (top left), clockwise and higher amplitude (top right), or clockwise and lower amplitude (bottom right).

3.2.5: Psychophysics

In order to equate task difficulty across orientation and amplitude discrimination tasks and across subjects, orientation and amplitude discrimination thresholds were acquired for each subject using staircasing methods (Cornsweet, 1962, reviewed in Chapter 2). Stimuli were generated and presented using Psychtoolbox 3.0 (www.psychtoolbox.com) in Matlab v7.14 (Mathworks) and presented on a gamma-corrected Mitsubishi Diamond Pro 2070SB display of 1024x768 resolution and a refresh rate of 75Hz on a mid-grey background (49cd/m2). Subjects viewed stimuli in a darkened room from a chin rest positioned 57cm from the display.

In the first session thresholds were aquired for orientation discrimination. In this session subjects indicated which of the two stimuli had the more clockwise orientation

with a key press. The orientation of the comparison was governed by a set of interleaved staircases. Pairs of 2-up 1-down and 3-up 1-down staircases were used in order to collect data spread across informative parts of the psychometric function (reviewed in Chapter 2). Staircases started at ± 20 degrees relative to the reference with an initial step size of 12 which halved on reversals 3 to 5. Staircases terminated after 14 reversals. As thresholds for amplitude discrimination had not yet been determined, the amplitude difference between reference and comparison stimuli was 0.03 for all subjects (approximately equal to the average amplitude discrimination threshold recorded by Silson et al., 2013).

In the second session amplitude discrimination thresholds were acquired. This time, the amplitude of the comparison was determined using the same interleaved staircases as for orientation thresholds. Staircases started at ± 0.13 relative to the reference with an initial step size of 0.08 which halved on reversals 3 to 5. The orientation difference between the reference and comparison was equal to ± the orientation discrimination threshold calculated from the first session.

3.2.6: fMRI

Subjects completed three 10 minute 10 second fMRI runs in which they performed orientation and amplitude discriminations of RF3 patterns (imaging parameters: TR = 2500ms, TE = 15.9ms, voxel size = 3x3x3mm3, flip angle = 90°, matrix size 96x96x45, FOV = 192mm). Subjects alternated between blocks of orientation and shape discriminations, starting with an orientation block. Stimuli were the same as those used for the psychophysics experiment, presented on a mid-grey background (200cdm^{-2}). For each subject orientation and amplitude differences between the reference and comparison stimuli were equal to their discrimination thresholds measured from the psychophysics experiments.

Subjects completed twelve 15 second blocks of orientation discrimination and twelve 15 second blocks of amplitude discrimination, separated by 10 second fixation periods, in each run (see Figure 3.3). The first block was preceded by a 10 second fixation period to allow the scanner to reach stable magnetisation. There were 8 trials in each stimulus block. A trial consisted of a 0.4 second stimulus presentation and a 1.4 second response window. No stimuli were presented during the first 0.6 seconds of a block, during

which time the fixation turned red to notify the subject that trials would be starting shortly. In addition, subjects were cued as to which task to perform during the fixation period preceding a trial block, where the orientation task was cued by a circular fixation and the amplitude task was cued by a plus-shaped fixation.

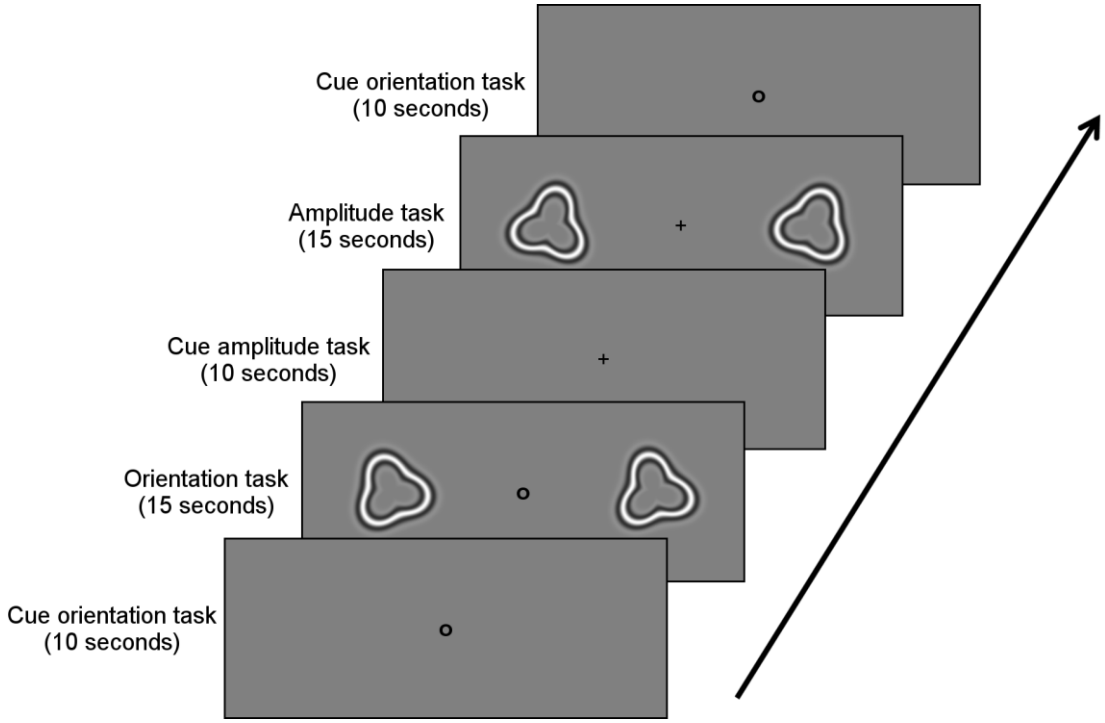

Figure 3.3 Schematic representation of the ABAB block design used in the fMRI experiment. Subjects alternated between 15 second blocks of orientation and amplitude discriminations. Blocks were separated by 10 second fixation periods. During task blocks and the preceding fixation period, the task to be performed was cued by the fixation where a circular fixation cued the orientation task and a plus-shaped fixation cued the amplitude task.

fMRI data were analysed with a general linear model (GLM) in FEAT (fMRI expert analysis tool) version 4.1 where orientation and amplitude discriminations were the regressors of interest. The first four volumes from each scan were discarded and a high pass filter with a cutoff of 50 seconds was used to remove low frequency drift. Head movements were corrected for with MCFLIRT and motion estimates were included in the GLM as regressors of no interest. Spatial smoothing with a Gaussian kernel of full width half maximum 6 millimetres was applied to the data and FILM prewhitening was used. For each subject, data from the three fMRI runs were combined using a fixed-effects analysis with cluster correction ($Z = 2.3$, $p < .05$). Percent signal change

during orientation and amplitude discriminations were then extracted from each retinotopic region of interest using FEATQUERY v4.1.

3.3: Results

3.3.1: Retinotopic mapping

Visual field maps V1, V2, V3, V4, LO1 and LO2 were successfully identified in both hemispheres of all subjects using published criteria (Larsson & Heeger, 2006; Sereno et al., 1995; Wade, Brewer, Rieger, & Wandell, 2002; Wandell et al., 2007; reviewed in Chapter 2). An example retinotopic mapping data set is shown in Figure 3.4.

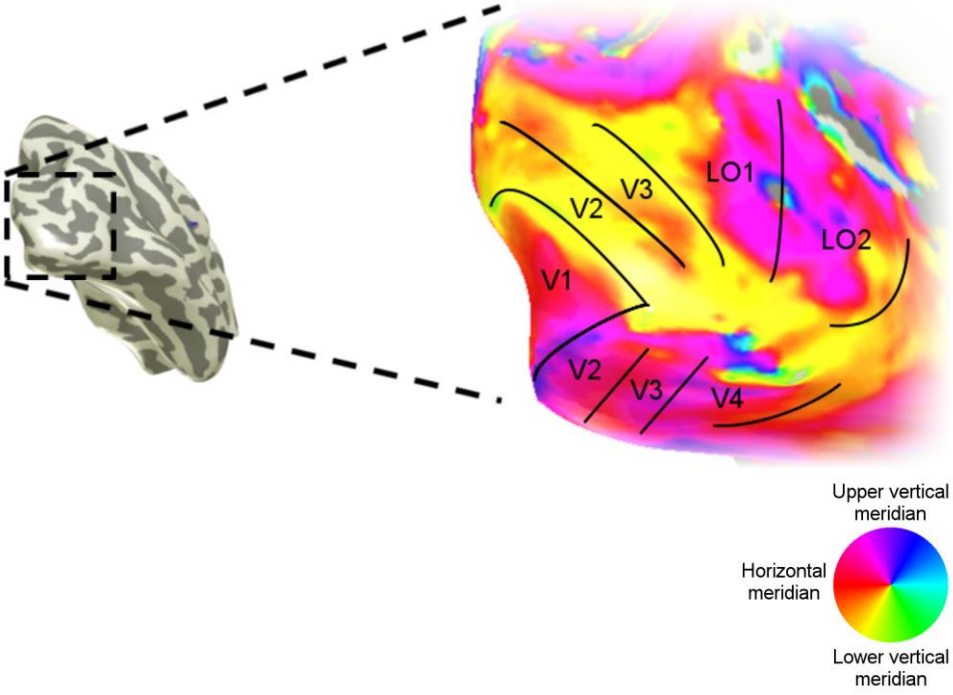

Figure 3.4: Responses to a rotating wedge stimulus in the right hemisphere of an example subject. The colour represents the polar angle that elicited a maximal response. The colour wheel (bottom right) shows which polar angles are represented by colours. Visual field map boundaries were identified by reversals in the progression of polar angle across the cortical surface. V1 was identified as a map of the contralateral hemifield on the medial surface of the occipital pole. V2 and V3 were identified as lower and upper quarterfield maps of the contralateral visual field which flanked V1. LO1 was identified as a map of the contralateral hemifield extending from the anterior border of the dorsal portion of V3, and LO2 as a second hemifield map extending from the anterior border of LO1. V4 was identified as a map of the contralateral hemifield extending from the anterior border of the ventral portion of V3.

3.3.2: Orientation and amplitude discrimination thresholds

Orientation and amplitude discrimination thresholds were measured during the psychophysics experiments. To measure discrimination thresholds for the orientation task, the proportion of trials for which the subjected reported that the comparison stimulus was more clockwise than the reference was plotted against the orientation of the comparison relative to the reference and a cumulative Gaussian was fitted to the data (see Figure 3.5). This Gaussian defined the comparison orientation required for a given subject to perform at 84% correct (equal to one standard deviation of the underlying Gaussian from the mean). For amplitude discrimination, the proportion of trials where the subject reported the comparison was higher amplitude than the reference was plotted against the amplitude of the comparison and thresholds were established using the same methods. As can be seen in Table 3.1, there were substantial differences in discrimination thresholds across individuals, highlighting the importance of using psychophysical measurements to equate task difficulty across subjects for the fMRI experiment.

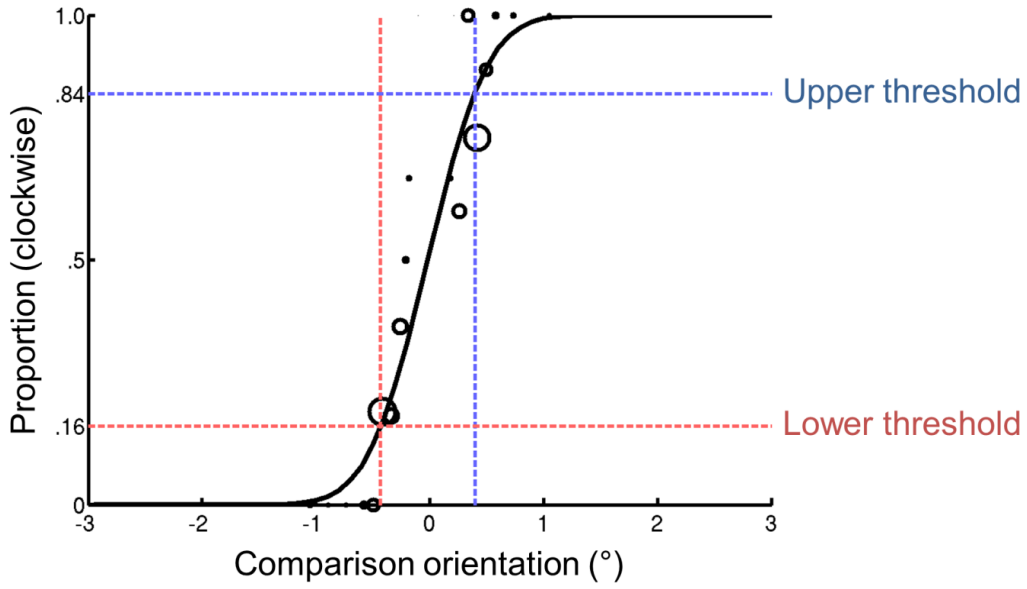

Figure 3.5: Example data from the orientation discrimination psychophysics experiment for one subject. The proportion of trials where the subject reported the comparison was more clockwise than the reference is plotted against the orientation of the comparison relative to the reference. Marker size is representative of the number of measurements taken at a given comparison orientation. A cumulative Gaussian was fit to the data to define the change in comparison orientation needed for the subject to report clockwise on 84% or 16% of trials (84% correct).

Table 3.1: Orientation and amplitude discrimination thresholds for all subjects derived from the psychophysics experiments. Average orientation and amplitude discrimination thresholds are shown at the bottom of the table with standard deviations in parentheses.

Subject	Orientation discrimination threshold (°)	Amplitude discrimination threshold
1	7.9	0.018
2	12.0	0.057
3	14.0	0.043
4	10.6	0.047
5	15.9	0.042
6	14.7	0.030
7	8.0	0.030
8	14.1	0.027
Average (SD)	12.1 (3.06)	0.037 (0.013)

3.3.3: Brain responses to orientation and shape discriminations of RF3 patterns

LO1 and LO2 show a double dissociation in their roles in orientation discriminations of sine-wave gratings and amplitude discriminations of RF3 patterns (Silson et al., 2013). As the goal of this study was to determine whether this double dissociation was due to the use of different stimuli or different tasks, a targeted region of interest (ROI) approach was used to examine the relative responses in LO1 and LO2 during orientation and shape discriminations of RF3 patterns (see Figure 3.6A).

A 2x2 repeated measures ANOVA examined the overall difference in BOLD response to RF3 patterns in LO1 and LO2 and the effect of task on the BOLD response. A significant main effect of ROI showed that responses were generally higher in LO2 than LO1 ($F(1, 7) = 12.05$, $p = .010$). There was no significant main effect of task ($F(1, 7) = 0.65$, $p = .447$), however a significant ROI x task interaction suggested some effect of task ($F(1, 7) = 13.25$, $p = .008$). To explore this further, two paired- samples t tests were conducted to examine the effect of task on responses in LO1 and LO2 separately. LO1 responses to the amplitude task were greater than the orientation task, but this

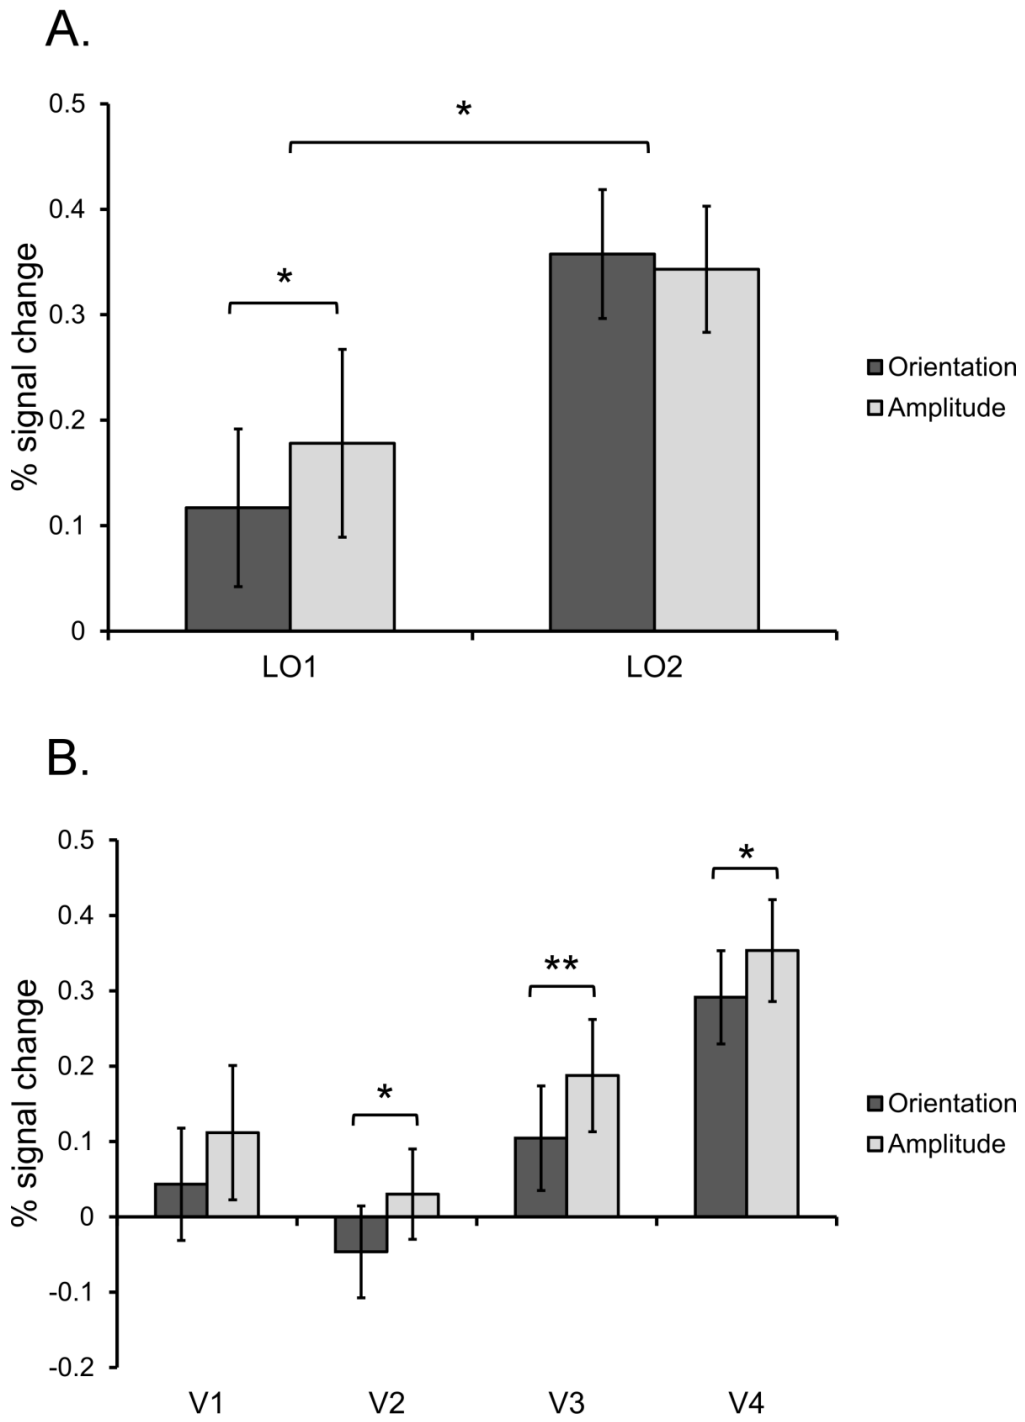

Figure 3.6: BOLD responses to orientation (darker bars) and amplitude discriminations (lighter bars) of RF3 patterns in various visual areas. (A) Responses in areas LO1 and LO2. The top * denotes a significant main effect of ROI according to a 2x2 repeated measures ANOVA where responses were larger in LO2 than LO1. The bottom * denotes a difference between responses to the two tasks in LO1 that approaches significance, but does not survive correction for multiple comparisons (paired-samples t test, see text for statistical details). (B) Responses in V1, V2, V3 and V4. Asterisks denote a significant difference between responses to the two tasks according to paired-samples t tests where * is significant at a p value of .05 and ** is significant at a Bonferroni corrected p value of .013 (see text for statistical details). In both panels error bars show standard error of mean.

difference did not reach significance with a Bonferroni corrected p value of .025 (t (7) = 2.7, p = .030). LO2 on the other hand showed no difference in response to the two tasks (t (7) = 0.65, p = .447). Overall, LO2 responded more strongly to RF3 patterns than LO1 and was not affected by which task was being performed. This suggests that the causal role for LO2 in amplitude discriminations of RF patterns (Silson et al., 2013) was due to LO2's role in processing RF patterns generally, and is not specific to amplitude discrimination.

Secondly, a 4x2 repeated measures ANOVA was used to examine responses to the two tasks in early visual areas V1, V2 and V3 and area V4 which exhibits some sensitivity to shape and curvature (Dumoulin & Hess, 2007; Wilkinson et al., 2000; see figure 3.6B). Receptive field sizes in V1 are small (Amano et al., 2009; Dumoulin & Wandell, 2008), meaning that analysing all voxels in V1 would result in the inclusion of many voxels outside of the stimulus representation. Retinotopic mapping data was therefore used to restrict V1 ROIs to voxels responsive to the portion of the visual field in which stimuli were presented (8-12° eccentricity and 11.5° either side of the horizontal meridian for polar angle). Possibly due to increasing receptive field sizes through visual cortex (Amano et al., 2009; Dumoulin & Wandell, 2008), retinotopy data did not provide the spatial specificity required to restrict ROIs beyond V1 in this manner. Therefore all voxels in V2, V3 and V4 were included for analyses (as was the case for LO1 and LO2).

Responses to RF3 patterns were larger in more anterior visual areas, resulting in a significant main effect of ROI (Huynh-Feldt correct F (1.58, 11.03) = 9.93, p = .005). Responses were largest in V4 (Figure 3.6B), consistent with a possible role for V4 in the processing of shape (Dumoulin & Hess, 2007; Wilkinson et al., 2000). There was also a significant main effect of task where, similar to LO1, responses to the amplitude task were larger in all four visual areas (F (1, 7) = 7.36, p = .030). This difference reached significance in V2 (t (7) = 3.20, p = .015), V3 (t (7) = 3.48, p = .010) and V4 (t (7) = 3.09, p = .018), but not V1 (t (7) = 1.55, p = .165), however only the difference in V3 remained significant using a Bonferroni corrected p value of .013. The ROI x task interaction was not significant (F (1, 7) = 0.29, p = .831). Overall, therefore, responses to RF3 patterns were larger in more anterior visual areas, and responses were generally larger during amplitude discriminations compared to orientation discriminations.

3.3.4: Task performance

Although preliminary psychophysical testing was used to match task difficulty, performance was higher for the amplitude task compared to the orientation task in the fMRI experiment. On average, participants performed 90% correct (SD = 5.26) for the amplitude task and 78% correct (SD = 7.51) for the orientation task, and a paired-samples t test showed this difference was significant (t (7) = 4.01, p = .005). To determine whether the results could be accounted for by this difference in task performance, performance for each task was correlated against neural responses to the same task from LO1 and LO2 (the primary areas of interest for the hypotheses). Performance on the orientation task did not correlate with responses to the orientation task in either LO1 (Pearson's r = .24, p = .560) or LO2 (r = .12, p = .777). Similarly, amplitude task performance did not correlate with responses to the amplitude task in LO1 (r = -.41, p = .320) or LO2 (r = -.253, p = .546). Therefore, although difficulty was not successfully matched across tasks, variance in task performance does not appear to be related to variance in responses from LO1 or LO2. Consequently, conclusions can be drawn from responses in LO1 and LO2 across the two tasks, despite differences in task difficulty.

3.5: Discussion

Previous research has identified a double dissociation between the roles of visual areas LO1 and LO2 in orientation discrimination of sin-wave gratings and amplitude discriminations of RF patterns (Silson et al., 2013). In Silson et al.'s (2013) study, TMS to LO1 selectively impaired orientation discrimination while TMS to LO2 selectively impaired amplitude discrimination. As RF amplitude discriminations engage global-scale shape processing mechanisms (Wilkinson et al., 1998), this suggests that LO2 may be involved in global shape processing. However, because different stimuli were used for the two tasks it cannot be determined whether LO2 performs general global shape processing or processes related to amplitude discrimination, specifically. It is therefore not clear whether this double dissociation reflects differences in task specificity or stimulus selectivity between LO1 and LO2. TMS to LO1 could have impaired orientation perception, or might have disrupted the processing of stimuli defined by local changes in luminance such as gratings. Similarly, TMS to LO2 could have either

impaired RF amplitude perception, or disrupted the processing of sparse, global stimuli such as RF patterns. To address this BOLD responses were measured from LO1 and LO2 during orientation and amplitude discriminations of a pair of RF patterns. The stimuli remained identical across both tasks, meaning it could be determined whether responses in LO1 and LO2 were driven by sensitivity to the stimuli or the task being performed.

Our results indicated that responses in LO1 and LO2 were driven by sensitivity to the stimuli rather than the task. LO2 responded more to RF patterns than LO1, suggesting a greater sensitivity to sparse, globally processed stimuli (Dickinson et al., 2012; Loffler et al., 2003; Wilkinson et al., 1998) in LO2 compared to LO1. Moreover, LO2 responses were not affected by the task being performed, demonstrating that LO2 responds robustly to RF patterns irrespective of the task being performed or the stimulus features being attended to. There was an effect of task in LO1, but it was in the opposite direction to that predicted by orientation processing in LO1 (Larsson et al., 2006; Silson et al., 2013) as responses were largest during amplitude discriminations. Furthermore this task difference was likely inherited from early visual areas which showed the same effect.

Overall, the results are consistent with the notion that more global and complex processing, which is not present in early visual areas, emerges in LO2 (Larsson et al., 2006). LO2 contains a retinotopic representation of the visual field and overlaps with the posterior portion of object-selective LOC (Larsson & Heeger, 2006; Sayres & Grill-Spector, 2008; Silson et al., 2016; Vernon et al., 2016), while the more anterior portions of LOC are not retinotopically defined. As such, LO2 may be a transitional point were representations progress from a retinotopic representation of the physical boundaries of a stimulus to something more abstract, which can be captured using category-selective localisers (Malach et al., 1995). This is consistent with models of visual perception where information about local, low level features is integrated across in extrastriate cortex to reduce the dimensionality of the information represented (Güçlü & van Gerven, 2015; Loffler, 2008; Wilson & Wilkinson, 2015), resulting in more efficient stimulus encoding, and places LO2 as a middle, transitional step in this process.

All visual areas, except for LO2, responded more to the amplitude task than the orientation task. The amplitude task required specific attention to differences in the contours defining the two stimuli. On the other hand the orientation task demanded a more global appreciation of the overall orientation of both stimuli. Selective featural attention boosts responses in visual cortex (Corbetra et al., 1989) and the effects of attention can be seen as early as V1 (Gandhi, Heeger, & Boynton, 1999; Silver, Ress, & Heeger, 2007). As such, larger responses to the amplitude task in early visual areas may be due to increased attention to stimulus contours and thus an increased demand on the retinotopic representations of those contours. Most importantly, task differences in earlier visual areas were not present in LO2, providing further evidence for a shift towards more global processing in LO2.

It has been established that early visual areas, posterior to LO2, are sensitive to orientation (Larsson et al., 2006; Silson et al., 2013). However, responses were lower during the discrimination of orientation compared to amplitude. Whilst previous studies used sinusoidal gratings to define orientation, the current study used RF patterns. This highlights a difference in the way these stimuli are processed, and suggests that while earlier visual areas are sensitive to orientation defined by local elements in grating stimuli, this is not the case for global orientation of RF patterns. As such, different sensitivities to orientation in posterior and anterior regions of visual cortex may be due to differences in local and global processing, not orientation processing per se. Overall, it seems that representations in LO2 are more global compared to LO1, resulting in stronger responses to RF patterns in LO2 than LO1, even when attention was driven towards the orientation of the stimuli.

It was surprising that early visual areas (V1, V2) did not show strong responses to the stimuli. Since the responses in Figure 3.6B show task responses versus a no stimulus baseline, one would expect early visual areas to show a strong response. A possible explanation is that the ROIs used for analyses included many voxels responsive to portions of the visual field outside of the stimuli. This is supported by the fact the responses from V1, which was constrained to only include voxels roughly within the stimulus representation, shows a slightly larger response than V2, which was not constrained. As such, the increase in response through V2 to V4 is likely due to increasing receptive field size (Amano et al., 2009), resulting in a larger proportion of

voxels within ROIs that were responsive to the stimuli. This could also explain the response difference between LO1 and LO2, as receptive fields are larger in LO2 than LO1 (Amano et al., 2009). As such, the difference between LO1 and LO2 could be due to receptive field size differences, rather than a specific preference for RF patterns in LO2 over LO1 per se. Importantly, however, the within ROI task difference present in all early ROIs and LO1 was not present in LO2. This cannot be explained by receptive field size differences as the stimuli were identical across tasks, and therefore their relationship to receptive field size was also identical across tasks. This therefore suggests a more global and less task specific response to RF patterns in LO2 compared to other visual areas.

3.6: Conclusions

In conclusion, LO2 responds robustly to RF patterns and its responses were not modulated by the task being performed. LO1 responded less strongly to RF patterns than LO2 and, counter to previous evidence (Larsson et al., 2006; Silson et al., 2013), did not show any preference for orientation discrimination. This indicates that the double dissociation between the functions of LO1 and LO2 previously identified by Silson et al. (2013) was due to global shape processing in LO2, but not LO1. This, together, with evidence of overlap between LO2 and complex object processing in LOC (Sayres & Grill-Spector, 2008; Silson et al., 2016; Vernon et al., 2016), suggests that LO2 may be the locus of early global shape processing in humans.

Chapter 4: Global shape aftereffects in composite radial frequency patterns

Abstract

Individual radial frequency (RF) patterns are generated by modulating a circle's radius as a sinusoidal function of polar angle and have been shown to tap into global shape processing mechanisms. Composite RF patterns can reproduce the complex outlines of natural shapes and examining these stimuli may allow for the interrogation global shape mechanisms that are recruited in biologically relevant tasks. In this chapter, evidence is presented for a global shape aftereffect in a composite RF pattern stimulus comprising two RF components. Manipulations of the shape, location, size and spatial frequency of the stimuli revealed that this aftereffect could only be explained by the attenuation of intermediate-level global shape mechanisms. The tuning of the aftereffect to test stimulus size also revealed two mechanisms underlying the aftereffect; one that was tuned to size and one that was invariant. Finally, evidence was shown to suggest these shape mechanisms may encode some RF information. However, the RF encoding found was not capable of explaining the full extent of the aftereffect, indicating that encoding of other shape features such as curvature are also important in global shape processing.

4.1: Introduction

Radial frequency (RF) patterns, stimuli defined by a sinusoidal modulation of a circle's radius as a function of the polar angle (Wilkinson et al., 1998), have proven to be powerful stimuli for investigating global shape processing. Importantly, RF patterns have been shown many times to engage global-scale shape processing mechanisms through a range of psychophysical tasks including shape detection (Dickinson et al., 2012, 2015; Hess, Wang, & Dakin, 1999; Jeffrey, Wang, & Birch, 2002; Loffler, Wilson, & Wilkinson, 2003; Schmidtmann et al., 2012; Wilkinson et al., 1998) and shape adaptation (Anderson et al., 2007; Bell et al., 2014; Bell et al., 2010; Bell & Kingdom, 2009). In particular, shape adaptation experiments are powerful as they use behavioural data, which can be gathered quickly and cost effectively, to measure the sensitivities of the neural mechanisms that encode the adapted feature to various stimulus properties.

While the mechanisms which process single RF patterns have been studied extensively, much less is known about the processing of composite RF patterns, where multiple RF components are summed to form a single, more complex pattern. Composite RF patterns can represent the arbitrarily complex boundaries of natural objects such as faces (Wilson et al., 2000; Wilson, Loffler, & Wilkinson, 2002; Wilson & Wilkinson, 2002) and animals (Vernon et al., 2016) in a convenient, parameterised manner. They can therefore be used to investigate the contribution of global shape mechanisms to the processing of biologically relevant stimuli.

An example of this is provided by Daar and Wilson (2012) who replicated the face viewpoint aftereffect using synthetic face stimuli generated through a combination of multiple RF components. The face viewpoint aftereffect is a perceptual aftereffect where adaptation to a face rotated away from the observer causes a neutral face to appear rotated in the opposite direction to the adaptor (Fang & He, 2005). Daar and Wilson showed that adaptation to face outlines alone was sufficient to induce a viewpoint aftereffect in a full-face test stimulus. In so doing, they demonstrated that the shape of a face outline acts as a cue to viewpoint; an example of the recruitment of shape processing mechanisms for a biologically relevant task.

Here shape adaptation was used to characterise the neural mechanisms which process composite RF patterns. This chapter will introduce a global shape aftereffect in RF2+3 patterns, composite RF pattern stimuli comprising two components with frequencies of 2 (RF2) and 3 (RF3), where manipulation of the phase of the RF3 component affects shape perceptions. This aftereffect was similar to Daar and Wilson's (2012) synthetic viewpoint aftereffect, however the use of a simpler combination of RF components allowed for the specific interrogation of the response properties of intermediate-level shape processing mechanisms, which may be recruited for biologically relevant tasks such as viewpoint perception. In addition, the use of a composite RF pattern stimulus allowed for the measurement of how adaptation to RF3 component phase transfers to different shapes that also include an RF3 component, shedding light on how RF information might be processed in the brain. The following series of experiments characterised shape aftereffects in RF2+3 patterns and the neural mechanisms which underpin them.

The results from these experiments are presented in three sections. The first uses RF patterns, which contain shape information that can accurately describe head outlines (Wilson et al., 2000) and act as a cue to face viewpoint (Daar & Wilson, 2012). Whilst the stimuli used were clearly less complex than a full head outline, they could still give the impression of an abstract head outline changing its viewpoint (see Figure 4.1 in General Methods section), and simple RF pattern stimuli have been shown to elicit responses from face-selective cortex (Wilkinson et al., 2000). It was therefore seen as important to first determine whether shape aftereffects in RF2+3 patterns could be explained by genuine viewpoint adaptation encoded by high-level, face-specific mechanisms as opposed to being a lower-level shape aftereffect. Results suggested no engagement of face-specific mechanisms and were consistent with a lower-level shape aftereffect.

In the second section it was determined whether the aftereffect could be explained by tilt aftereffects by characterising its sensitivity to stimulus size and spatial frequency. The strength of the aftereffect was measured across a range of test stimulus sizes and spatial frequencies whilst the adaptor remained constant. The aftereffect exhibited a broad, Gaussian tuning profile to size, plateauing to remain significant across the entire tested range, and complete invariance to spatial frequency. This provided strong

evidence that the aftereffect was not a manifestation of tilt aftereffects, which are tuned to spatial frequency (Ware & Mitchell, 1974), and is instead an example of a global shape aftereffect mediated by intermediate-level shape mechanisms. In addition, results were consistent with other reports of two mechanisms underlying shape aftereffects, one feature-tuned and one feature-invariant (Bell & Kingdom, 2009). It was added to these results that feature-tuned mechanisms are tuned to stimulus size but not spatial frequency.

In the third section, it was asked whether the aftereffect was more consistent with shape encoding by neural channels tuned to RF (Bell & Badcock, 2009) or curvature-based models of global shape perception (Dickinson et al., 2015; Poirier & Wilson, 2006). There was evidence to support both alternatives, where adaptation to RF3 phase transferred to test stimuli of different shapes and sizes which contained an RF3 component. However, the magnitude of this effect was not sufficient to account for the entire illusion. It was therefore concluded that global shape aftereffects in composite RF patterns constitute adaptation to both RF content and other shape features such as curvature.

4.2: General Methods

4.2.1: Observers

All observers had normal or corrected-to-normal vision and gave written informed consent. The number of observers for each individual experiment is detailed in the Results section. The experiments were approved by the York Neuroimaging Centre Research Governance Committee and adhered to the Declaration of Helsinki.

4.2.2: Stimuli

Two RF patterns (Wilkinson et al., 1998) were summed to generate RF2+3 pattern stimuli (Figure 4.1). The radius r at polar angle θ for this stimulus was calculated using the following equation:

$$r(\theta) = r_0(1 + A_1 \sin(\omega_1\theta + \emptyset_1) + A_2 \sin(\omega_2\theta + \emptyset_2))$$

Where r_0 is the mean radius, A_1 and A_2 are the radial modulation amplitudes, ω_1 and ω_2 are the radial frequencies and $\emptyset_1$ and $\emptyset_2$ are the phases of the two RF components. The component stimuli were an RF2 (ω_1= 2, A_1= 0.12, $\emptyset_1$= -90˚) and RF3 pattern (ω_2= 3, A_2= 0.05, $\emptyset_2$ was variable, where 180˚ produced a neutral, front-facing stimulus and higher or lower values produced increasingly rightward- or leftward-facing stimuli, respectively). Henceforth reference to the RF3 component phase will be normalised such that the phase required for a front-facing stimulus will be defined as 0˚ (as in Figure 4.1). These stimulus parameters resulted in an RF2+3 pattern whose shape was manipulated using the RF3 component phase (see Figure 4.1). The cross-sectional luminance profiles of stimulus contours were defined by the following equation:

$$D4(d) = C\left(1 - 4\left(\frac{d}{\sigma}\right)^2 + \frac{4}{3}\left(\frac{d}{\sigma}\right)^4\right) \times \exp\left(-\left(\frac{d}{\sigma}\right)^2\right)$$

Where $D4$ is the fourth derivative of a Gaussian, d is the distance between the pixel and the stimulus contour along the line with the shortest distance between the pixel and the contour, C is the contrast, r_0 is the mean radius, and σ determines the peak spatial frequency.

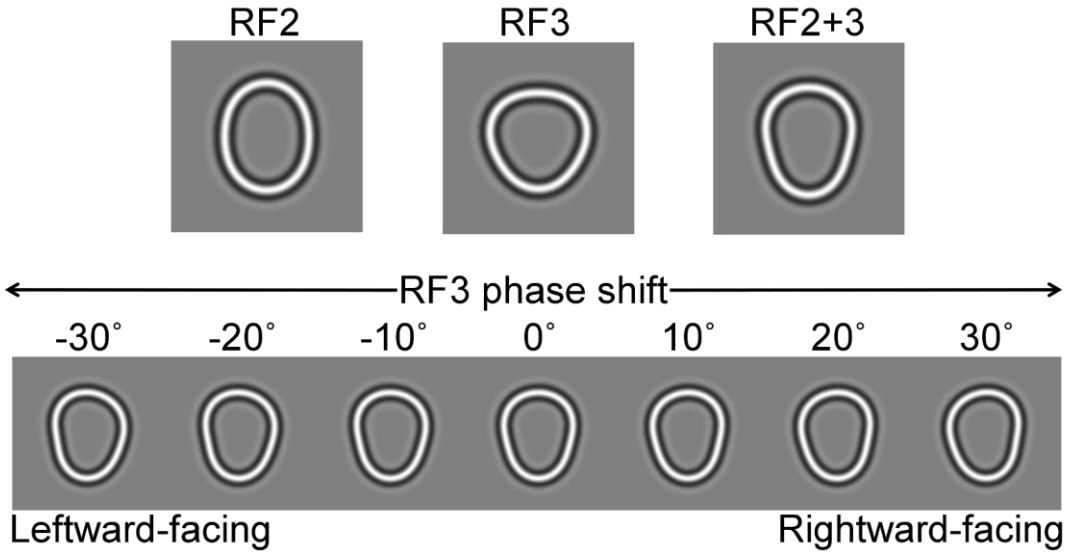

Figure 4.1. Two RF components with frequencies of two (RF2) and three (RF3) were summed to create an RF2+3 pattern. The shape of this stimulus was manipulated by the phase of the RF3 component, where a phase of 0º created a neutral, front-facing stimulus and increases and decreases in phase resulted in rightward- and leftward-facing rotations, respectively. The RF3 component phases used to create each example stimulus are listed above the stimuli. See text for specific stimulus parameters.

4.2.3: Procedure

Stimuli were generated and presented using Psychtoolbox 3.0 (www.psychtoolbox.com) in Matlab v7.14 (Mathworks) and presented on a gamma-corrected Mitsubishi Diamond Pro 2070^{SB} display of 1024x768 resolution and a refresh rate of 75Hz on a mid-grey background (49cd/m^2). Subjects viewed stimuli in a darkened room from a chin rest positioned 57cm from the display.

The following procedure was used for all experiments except for the bias free experiment (described in the following section). Any stimulus parameters that deviated from the following are specified in the Results section. The adaptor stimuli were RF2+3 patterns with a mean radius of 2.5 degrees of visual angle and peak spatial frequency of 1.26 cycles per degree. To avoid local luminance adaptation, the contrast of the adaptor modulated between -100% and 100% according to a 1Hz sinusoidal function. At the start of each trial a single adaptor stimulus was presented centrally for 30 seconds on the first trial and 5 seconds on subsequent trials. In the baseline adaptation condition, the adaptor RF3 component had a phase of 0° (front-facing), while in the adapt-rotated condition the adaptor was rotated to face away from the observer (RF3 component phase = ± 12°). The sign of rotation in the rotated adaptor was counterbalanced across participants. After a 1 second inter-stimulus interval a test RF2+3 pattern was presented for 200 milliseconds at 50% contrast in the same position as the adaptor. The task was to identify whether the test stimulus was rotated towards the observer's left or right. After the participant responded, the next trial followed a 1.5 second inter-trial interval. A pair of 1-up 1-down staircases governed the rotation of the test stimulus in order to estimate point of subjective equality (PSE). The two staircases began at 16° and -16° with an initial step size of 4º, which halved on reversals 3 to 5. Staircases terminated after 14 reversals. A small (0.5 degrees of visual angle across) central fixation cross was present through the experiment, except for during the 200ms test presentation period, and participants were instructed to maintain fixation at all times. Fixation was not present during the presentation of the test stimulus so that, in cases where the radius of the test stimulus was small (the rationale for which is explained in the Results section), fixation did not overlap with the stimulus.

4.2.4: Bias free procedure

In order to ensure shape aftereffects in composite RF patterns were due to genuine perceptual adaptation and not changes in response bias, aftereffects were also measured using a bias free adaptation design (Morgan, 2013). In this design, adaptation is measured for two stimuli simultaneously, and the observer is asked to make a relative judgement about the two stimuli. One of these stimuli is a fixed reference stimulus, which is the same as an adaptor, while the second is the true test stimulus, controlled by staircases that are centred on the reference. In the case of no aftereffect, therefore, staircases converge on the reference stimulus, meaning the observer adjusts the true test stimulus until it matches the reference (as should also be the case for a baseline condition). The magnitude of adaptation can be measured by the discrepancy between the stimulus level at which staircases converge on for the true test and the stimulus level of the reference. The effect of this paradigm is that, for any given trial, the observer cannot know which of the two stimuli is the true test and which is the reference, meaning they could not bias their responses in the predicted direction of an aftereffect (Morgan, 2013).

For bias free experiments, two RF2+3 pattern adaptors were presented simultaneously, one 4.5 degrees above fixation and the other 4.5 degrees below, followed by two test stimuli in the same locations (see Figure 4.2). Observers were asked to indicate which of the two test stimuli was more rotated away from the observer. Stimulus timings were identical to those described in the previous Procedure section. In the baseline condition both adaptors were front-facing (RF3 phase = 0°), whereas in the adapt-rotated condition the adaptors were rotated in opposite directions to face away from the observer (RF3 phase = -20° for the top stimulus and 20° for the bottom stimulus). Other stimulus parameters were identical to those described in the previous section. There were 4 possible combinations of test stimulus rotations: both left facing, both right, one left one right and one right one left. Staircases began at ± 20 ° with an initial step size of 6° which halved on reversals 1 and 2. Observers completed four staircases per trial type (one pair of one-up one-down staircases each for when the true test stimulus was the top or bottom stimulus). Due to the time consuming nature of this experiment (approximately 1.5 hours per observer per condition), subject completed this experiment over several 30 minute sessions.

PSE shifts between adaptation conditions were not predicted for test stimuli that were both rotated in the same or opposite direction to adapt-rotated adaptors (trial types 1 and 2 in Figure 4.2). This was because the effect of adaptation should be equal across the two test stimuli, meaning the observer should be able to match the rotation of the true test stimulus to the reference, which had the same rotation magnitude as adaptors. Aftereffects were predicted when the two test stimuli were rotated in the same direction as each other, but different directions relative to adapt-rotated adaptors (trial types 3 and 4 in Figure 4.2). This was because shape perception should now be different for the two test stimuli as they had a different relationship with their respective adaptors.

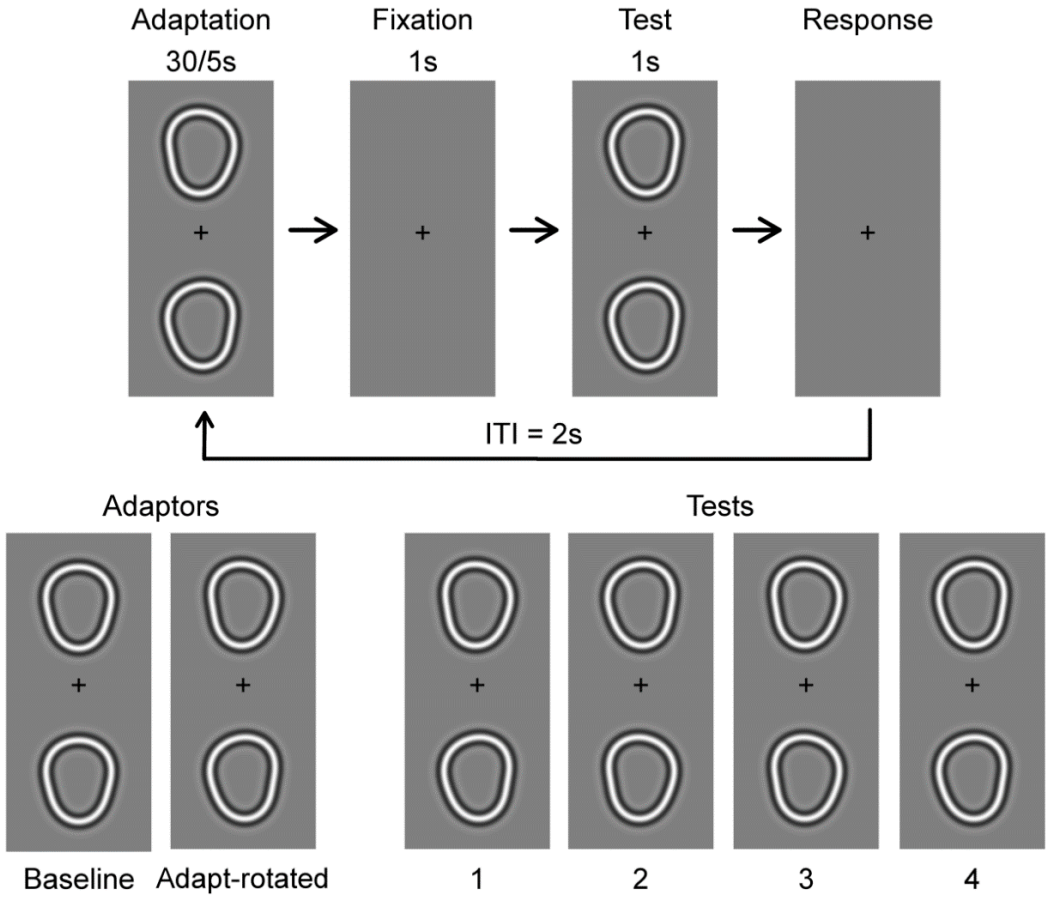

Figure 4.2: Stimulus presentation for the bias free adaptation paradigm. Two adaptors were presented simultaneously, followed by two test stimuli (top). The adaptors for the two adaptation conditions are shown in the bottom left. The four combinations of test stimuli are shown in the bottom right. No PSE shift was predicted for trial types 1 and 2, but shifts were predicted for trial types 3 and 4 (see text for more detail).

4.2.5: Data analysis

The same approach was used to analyse data from both standard and bias free adaptation paradigms. For each adaptation condition, the PSE was calculated by averaging the RF3 component phase of the test stimulus for every trial during the final 9 reversals (when step size was minimum) for both staircases (or for all four staircases per trial type in bias free experiments). The difference between PSE measurements from each adaptation condition was a measure of the strength of adaptation. For the standard adaptation paradigm, the sign of both baseline and adapt-rotated PSEs was reversed for participants who adapted to a stimulus rotated to the left in the adapt-rotated condition (where a negative PSE shift was predicted). As such, any aftereffect in the predicted direction was represented by a positive shift in PSE. For the bias free adaptation paradigm, the same was done for trials where the true test stimulus was in the same location as the leftward-rotated adaptor in the adapt-rotated condition.

4.3: Results

4.3.1: Shape versus face adaptation

First, it was determined whether manipulation of RF3 component phase in the adaptor could affect the perceived shape of RF2+3 patterns. 6 observers took part in the experiment (2 authors, 4 naïve). Average PSE measurements were significantly different between adaptation conditions ($t(5) = 9.69$, $p < .001$), showing that manipulation of RF3 component phase in the rotated adaptor altered shape perceptions (Figure 4.3A). Whilst stimuli in the current study were simple and contained considerably less shape complexity than real head outlines, it was noted that the RF3 phase manipulation shown in Figure 4.1 gave the impression of an abstract head outline changing its viewpoint. Moreover, RF patterns can elicit responses from face-selective cortex (Wilkinson et al., 2000) and can convey shape cues to face viewpoint (Daar & Wilson, 2012; Wilson et al., 2000). Therefore a series of experiments were conducted to determine whether the aftereffect recruited higher-level, face-specific mechanisms encoding the viewpoint of the stimulus, or lower-level shape encoding mechanisms.

Face processing is impaired by stimulus inversion (Freire, Lee, & Symons, 2000; Maurer, Le Grand, & Mondloch, 2002). To establish whether this processing impairment was in any way reflected in RF2+3 patterns shape aftereffects, the experiment was repeated using inverted RF2+3 patterns (RF3 component phase increased by 180º, see Figure 4.3B) in 6 observers (4 naïve). Again, a significant aftereffect was found using inverted RF2+3 patterns ($t(5) = 8.50$, $p < .001$). Importantly, there was no significant difference in the magnitude of the aftereffect in upright and inverted stimuli (Figure 4.3F). The aftereffect was therefore equivalent in both upright and inverted RF2+3 patterns and does not appear to reflect any processing impairment as a result of stimulus inversion.

High-level face aftereffects transfer across large portions of the visual field, remaining significant when stimuli are presented as far as 6 degrees of visual angle from fixation in opposite hemifields (Afraz & Cavanagh, 2008; Zimmer & Kovács, 2011), whereas shape aftereffects do not show this degree of position invariance (Melcher, 2005). This yields a simple prediction: if the aftereffect in RF2+3 patterns recruits high-level face encoding mechanisms it should transfer across retinal locations. To test this, shape aftereffects were measured in RF2+3 patterns in 6 observers (4 naïve) when adaptor and test stimuli were presented in different quadrants within the same visual hemifield (adapt 5 degrees below and right of fixation, test 5 degrees above and right, see Figure 4.3C) and when presented in opposite hemifields (adapt 5 degrees right of fixation, test 5 degrees left, see Figure 4.3D). In both cases, there was no significant aftereffect (within-hemifield: $t(5) = 1.59$, $p = .174$, cross-hemifield: $t(5) = 0.48$, $p = .649$). Additionally, aftereffect magnitudes for both experiments were significantly less than those measured using standard and inverted stimuli (Figure 4.3F). Therefore, unlike high-level face aftereffects, the RF2+3 pattern aftereffect does not transfer across large portions of the visual field, and hence likely involves lower-level shape encoding mechanisms.

Finally, it was determined whether similar effects of adaptor RF3 component phase could be found in a completely non-face-like stimulus. To do this the frequency of the RF2 component was changed to 4 for both adaptor and test stimuli, keeping all other stimulus parameters the same. Resulting RF3+4 patterns did not resemble face outlines, and the effect of RF3 component phase in the adaptor was measured in the same way

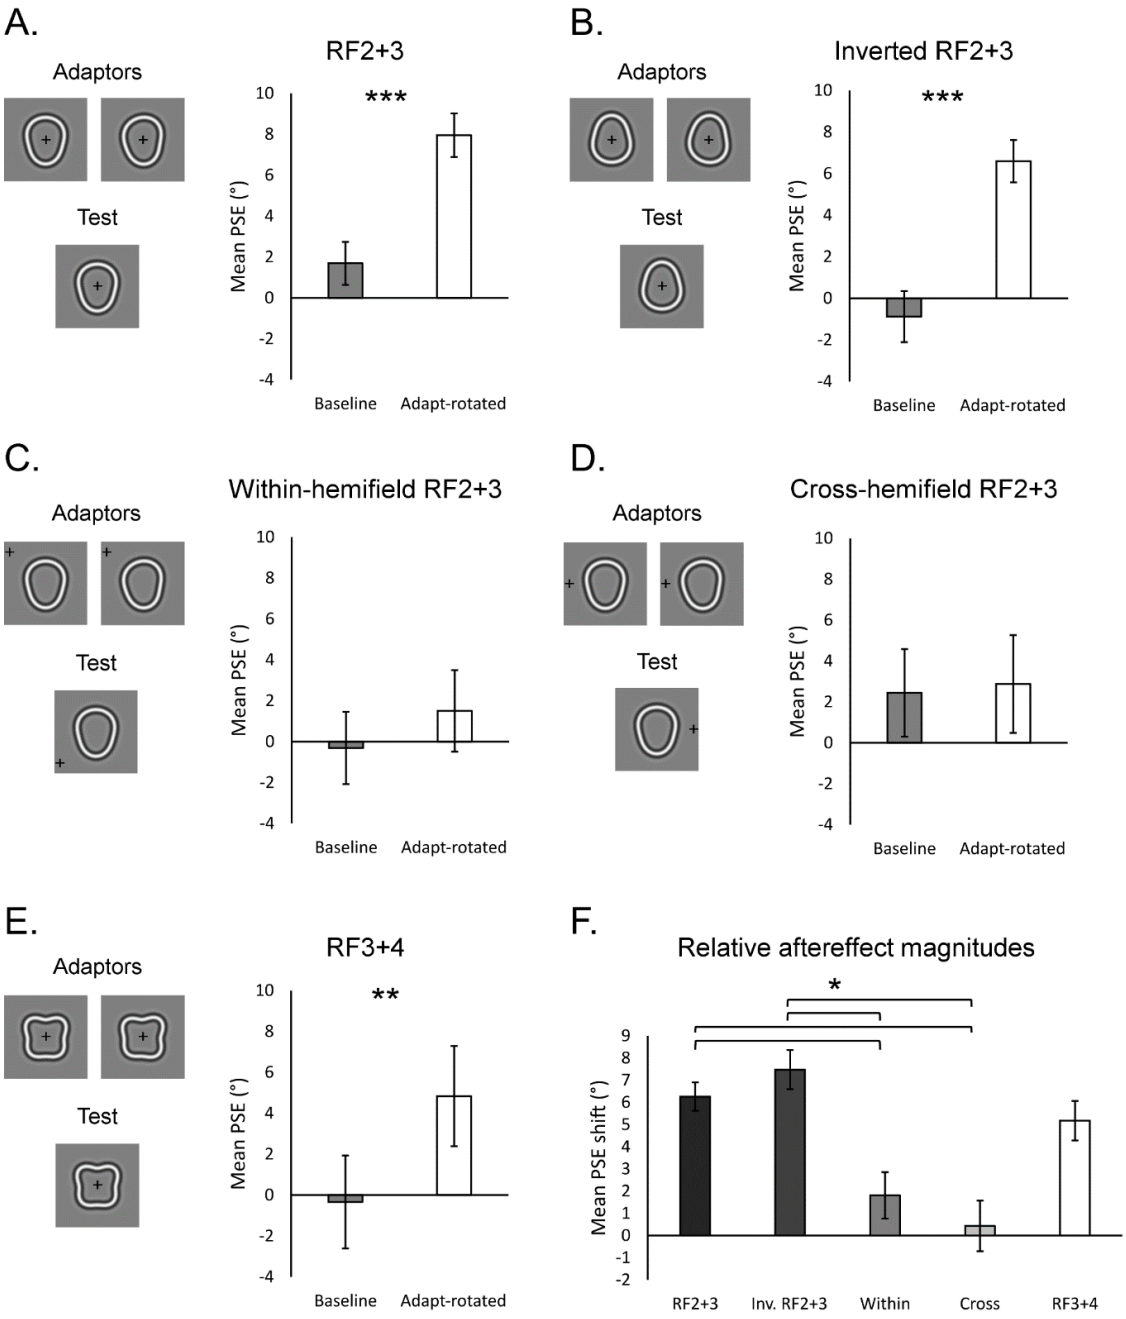

Figure 4.3: Results from experiments characterising the processing mechanisms involved in shape aftereffects in RF2+3 patterns. Panels A-E show average PSE measurements for the baseline and adapt-rotated conditions with grey and white bars, respectively, and asterisks denote a significant difference according to a paired-samples t test (*** = $p < .001$ and ** = $p < .01$). Results are shown for (A) upright RF2+3 patterns, (B) inverted RF2+3 patterns, (C) RF2+3 patterns where adaptor and test stimuli were presented in separate quadrants within the same visual hemifield, (D) RF2+3 patterns where adaptor and test stimuli were presented in opposite hemifields and (E) non-face-like RF3+4 patterns. (F) The average aftereffect magnitude (adapt-rotated PSE minus baseline PSE) is plotted for each of the experiments shown in panels A-E. A repeated-measures ANOVA revealed a significant main effect across experiments ($F(4, 20) = 12.33$, $p < .001$) and asterisks denote significant differences according to Bonferroni post-hoc tests at the $p < .05$ level. In all panels error bars show standard error of mean.

as for RF2+3 patterns (Figure 4.3E). In 6 observers (4 naïve), a significant aftereffect was found using RF3+4 patterns (t (5) = 4.95, p = .004), and the magnitude of this effect was not significantly different from standard or inverted aftereffects (Figure 4.3F). The effect of adaptor RF3 component phase, therefore, is applicable to non-face stimuli such as RF3+4 patterns, and is not specific to shapes that resemble face outlines.

Together, these experiments indicate that the aftereffect in RF2+3 patterns is processed by lower-level shape mechanisms, not higher-level, face-specific mechanisms. However, shape aftereffects may constitute a manifestation of local tilt aftereffects (Dickinson et al., 2010) rather than true global shape processing. The next series of experiments determined whether the aftereffect could be explained by tilt aftereffects.

4.3.2: Size and spatial frequency tuning

An important aspect of global shape aftereffects is their tolerance to changes in stimulus size between adaptor and test (Anderson et al., 2007; Suzuki, 2001). When stimulus size is changed, the global shape is constant across adaptor and test whilst each occupies a separate portion of the retinal image. The preservation of a shape aftereffect under these conditions is evidence that it cannot be explained by a simple manifestation of local tilt aftereffects, and that it engages global shape mechanisms which may represent shape by size-invariant stimulus features (Dickinson et al., 2015; Poirier & Wilson, 2006).

The global shape aftereffect in RF2+3 patterns was measured when adaptor and test stimuli had different sizes to ensure it exhibited the same resilience to changes in test stimulus size as other reported shape aftereffects. The mean radius of the test stimulus was 2.1 degrees of visual angle larger than the adaptor (sufficiently large so that there was no contour overlap between adaptor and test stimuli) and the spatial frequency was scaled appropriately using an inverse power law (Figure 4.4A). In 6 observers (4 naïve), there was a significant difference between PSE measurements from the two adaptation conditions (t (5) = 10.26, p = <.001), demonstrating that the aftereffect shows a tolerance to changes in stimulus size similar to other shape aftereffects.

It has been shown that experienced psychophysical observers can voluntarily shift their psychometric functions when instructed to bias themselves (Morgan et al., 2012). It was therefore argued that some perceptual aftereffects could reflect changes in

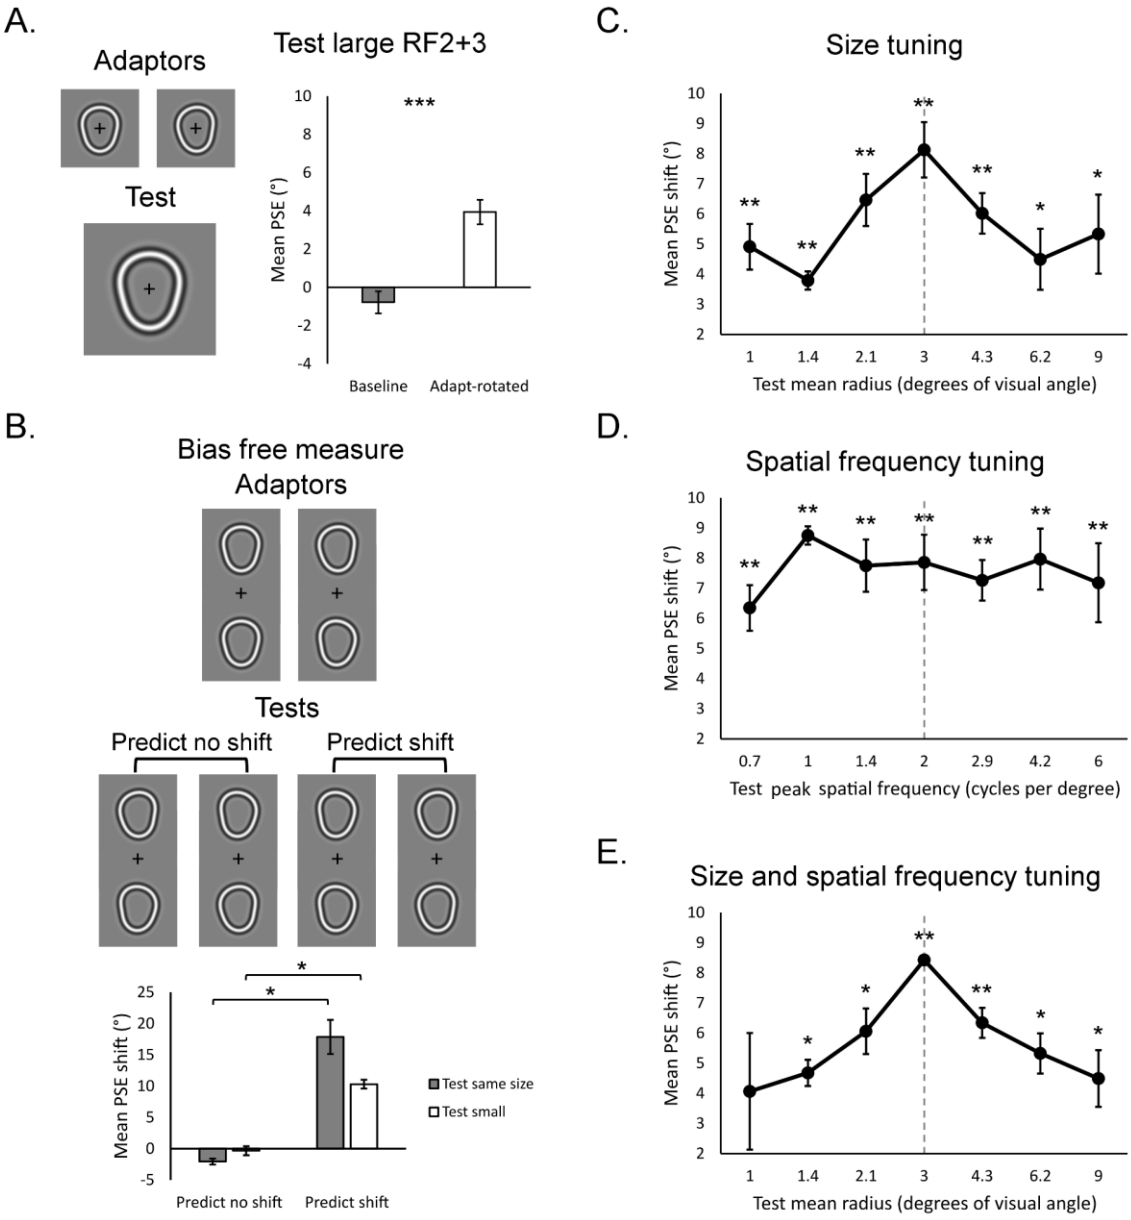

Figure 4.4: Results from experiments assessing the sensitivity of shape aftereffects in RF2+3 patterns to test stimulus size and spatial frequency. (A) Average PSE measurements are plotted for the baseline (grey bar) and adapt-rotated (white bar) conditions when the test stimulus was larger than the adaptor. *** denotes a significant difference at the $p < .001$ level according to a paired-samples t test. (B) Average PSE shifts measured using a bias free paradigm (Morgan, 2013, see text for further details) are plotted for predict shift and predict no shift trials for both same size and smaller test stimuli. * denotes a significant difference at the $p < .05$ level according to a paired-samples t test. Panels C-E show average PSE shifts (adapt-rotated PSE minus baseline PSE) for (C) a range of test stimulus sizes, (D) a range of test stimulus spatial frequencies and (E) a range of test stimulus sizes with scaled spatial frequency. In panels C-E grey dashed lines indicate the size/spatial frequency of the adaptor and asterisks denote a significant aftereffect according to a one-sample t test, where * denotes a significant aftereffect at the $p < .05$ and ** denotes a significant aftereffect using a Bonferroni corrected p value of $p < .007$. In all panels error bars show standard error of mean.

response bias as opposed to changed perception. This is particularly important when changing stimulus features between adaptor and test as any residual aftereffect could be due to response bias. As a control experiment, the shape aftereffect in RF2+3 patterns and the effect of stimulus size were replicated using a bias free adaptation paradigm (Morgan, 2013).

3 observers (1 naïve) completed two bias free experiments. In the first experiment, adaptor and test stimuli were the same size. In the second experiment, test stimuli were smaller than the adaptors (the size difference was proportionally equal to that used in the test large experiment shown in Figure 4.4A and spatial frequency was scaled in the same way).

The PSEs measured from shift-predicted trials and no-shift-predicted trails were significantly different for both same size ($t(2) = 6.79$, $p = .021$, see Figure 4.4B) and small test stimuli ($t(2) = 7.43$, $p = .018$). Therefore both the basic aftereffect and the effect of stimulus size were replicated using a bias free adaptation paradigm. It can thus be concluded that global shape aftereffects in RF2+3 patterns and their relative tolerance to changes in stimulus size are due to genuine changes in perception, not changes in response bias. Due to the time consuming nature of the bias free paradigm (approximately 1.5 hours per participant per condition), the remaining experiments reverted to the standard adaptation paradigm described in the Procedure. However, the consistency of results obtained using bias-free and standard paradigms was taken as evidence that the following experiments measure genuine perceptual effects.

Whilst the aftereffect remained significant with larger and smaller test stimuli, its magnitude was nevertheless reduced compared to the standard aftereffect. This is consistent with other reported effects of stimulus size (Anderson et al., 2007) and suggests that shape aftereffects may exhibit some broad tuning to stimulus size which cannot be appreciated without testing across a broader range. To address this issue, and explore the receptive field characteristics of the neural mechanisms underpinning global shape aftereffects, the strength of the aftereffect across a range of size and spatial frequency manipulations was measured.

First, the aftereffect was measured across a broad range of test stimulus sizes while the adaptor size remained constant. The adaptor was presented with a mean radius of 3

degrees of visual angle, and test stimuli were presented over a logarithmic range of mean radii from 33% to 300% of the adaptor with 7 levels. All test stimuli were presented with a peak spatial frequency of 2 cycles per degree. For 1 observer (an author) PSEs were measured for all test stimuli in a single session using interleaved staircases. For the 4 remaining observers (naïve), data were acquired across 4 sessions. In 3 of the 4 sessions, PSEs were measured for a pair of test stimuli (one smaller than the adaptor, one larger) using interleaved staircases. In the remaining session, PSEs were measured for the test stimulus which was the same size as the adaptor. The ordering of the 4 sessions and which test stimuli were paired together were counterbalanced across participants.

A 1x7 repeated measures ANOVA revealed a significant main effect of test stimulus size (Huynh-Feldt corrected F (4.9, 19.6) = 2.97, p = .038), demonstrating some level of tuning to test stimulus size. As can be seen in Figure 4.4C, the tuning profile of the aftereffect across test sizes was broadly Gaussian, peaking when adaptor and test were the same size, and plateaus after a stimulus size difference of ~50% or 200%. One-sample t tests confirmed that the aftereffect was significantly above 0 at the $p < .05$ level for all test sizes, where 5 of 7 aftereffects survived a Bonferroni correction for multiple comparisons (see Figure 4.4C). The aftereffect, therefore, is systematically reduced by increasing differences in adaptor and test stimulus size until it plateaus, at which point it remains significant in stimuli with size differences of up to a factor of 3.

Tilt aftereffects have been shown to be tuned to stimulus spatial frequency, where aftereffects reduced by half when adaptor and test spatial frequency differed by a factor of 2, and halved again when they differed by a factor of 4 (Ware & Mitchell, 1974). Therefore, if the aftereffect is driven by tilt aftereffects it should show a similar tuning profile to the cross-sectional spatial frequency of the test stimulus. The tuning of the aftereffect was measured across test stimulus spatial frequencies in 4 observers (3 naïve) using the same methods as for size. Adaptor peak spatial frequency was 2 cycles per degree and test stimuli were presented across a logarithmic range with 7 levels as previously explained. All stimuli were presented with the same mean radius (3 degrees of visual angle). The main effect of the test spatial frequency was not significant (Huynh-Feldt corrected F (4.4, 13.1) = 1.96, p = .157) and the tuning profile shown in Figure 4.4D is relatively flat. Again, all aftereffects were confirmed to be significant (all

survived a Bonferroni correction for multiple comparisons, see Figure 4.4D). It therefore seems that the spatial frequency relationship between adaptor and test stimuli has no effect on the magnitude of the aftereffect.

Finally, the tuning of the aftereffect was measured when test stimulus size and spatial frequency varied together in 3 observers (2 naïve). The test stimulus sizes and peak spatial frequencies from the previous two experiments were combined, such that test size and spatial frequency were paired by an inverse power law. Whilst the main effect of test stimulus size and spatial frequency did not reach significance (Huynh-Feldt corrected $F(2.3, 4.7) = 3.60$, $p = .111$), this was likely due to weak statistical power from testing only 3 observers and the tuning profile shown in Figure 4.4E is similar to that for size tuning in Figure 4.3C (confirmed by a 2-way repeated measures ANOVA which found no main effect when comparing the two tuning profiles; $F(1, 2) = 0.05$, $p = .851$). The aftereffect was significant for all test stimuli except for the smallest one; 2 of the 6 significant aftereffects survived a Bonferroni correction (see Figure 4.4E). Therefore the broad size tuning of the aftereffect replicated when both size and spatial frequency changed together. Moreover, the similarity of the tuning profiles across the two experiments serves to confirm that test stimulus spatial frequency has no effect on the aftereffect.

Overall, global shape aftereffects in RF2+3 patterns are relatively tolerant to changes in stimulus size, similar to other reported shape aftereffects, and this is due to genuine changes in perception rather than response bias. However, when measured over a broader range of test stimulus sizes, the aftereffect showed a broad Gaussian tuning profile. The aftereffect was also completely invariant to the cross-sectional spatial frequency of the test stimulus across the tested range, and therefore cannot be explained as a manifestation of tilt aftereffects.

4.3.3: Shape versus RF processing

It has been argued that global shape perception may be performed by a bank of narrow-band RF tuned channels which analyse the RF content of an object (Bell & Badcock, 2009; Salmela et al., 2016). Alternatively, global shape processing may involve an encoding of the shape profile through some form of curvature analysis (Dickinson et al., 2015; Poirier & Wilson, 2006). The following series of experiments were designed

to determine whether global shape aftereffects in RF2+3 patterns are more consistent with RF processing, where RF3 component phase is the adapted feature, or shape processing, where the shape profile of the adaptor is the adapted feature.

First, the RF2+3 pattern shape aftereffect experiment was repeated but the RF2 component was removed from adaptor and test stimuli (Figure 4.5A). Detection thresholds for a target RF are elevated when the target is paired with a mask RF of neighbouring frequency to form a composite RF pattern stimulus (Bell et al., 2009). If RF3 component phase is the adapted feature in RF2+3 pattern shape aftereffects, this should be masked by the presence of the RF2 component in RF2+3 patterns. As such, the magnitude of the aftereffect should be larger in RF3 patterns than RF2+3 patterns. Alternatively, if shape profile is the adapted feature, the aftereffect might be roughly the same in RF3 patterns, as the RF3 phase manipulation has a similar effect on points of maximum curvature in both RF3 and RF2+3 patterns, which are proposed to be the most important shape features in curvature-based models of global shape perception (Dickinson et al., 2015; Poirier & Wilson, 2006).

The aftereffect in RF3 patterns was significant ($t(5) = 7.68$, $p < .001$) in 6 observers (4 naïve). More importantly, the magnitude of the aftereffect was almost the same as in RF2+3 patterns, and the difference between them was not significant ($t(5) = .45$, $p = .671$, Figure 4.5B). Therefore it appears that the presence of the RF2 component has no masking effect on adaptation to RF3 phase, supporting shape over RF adaptation.

Next, it was investigated whether adaptation to RF3 phase could transfer across different shapes. If the aftereffect is driven by the attenuation of RF-tuned channels, adaptation to RF3 phase should affect the perception of different shapes so long as they also contain an RF3 component. As such the aftereffect was measured in RF2+3 patterns following adaptation to an RF3 pattern (Figure 4.5C). The aftereffect was significant under these conditions ($t(4) = 5.79$, $p = .004$) in 5 observers (3 naïve). The effect of adaptation to RF3 phase, therefore, did transfer to a different shape to affect the perception of an RF2+3 test, supporting RF adaptation.

However, RF3 pattern adaptors and RF2+3 tests were presented in the same location and were the same size. Hence the observed aftereffect could be due to a manifestation of local tilt aftereffects, which have been argued to be an important and often

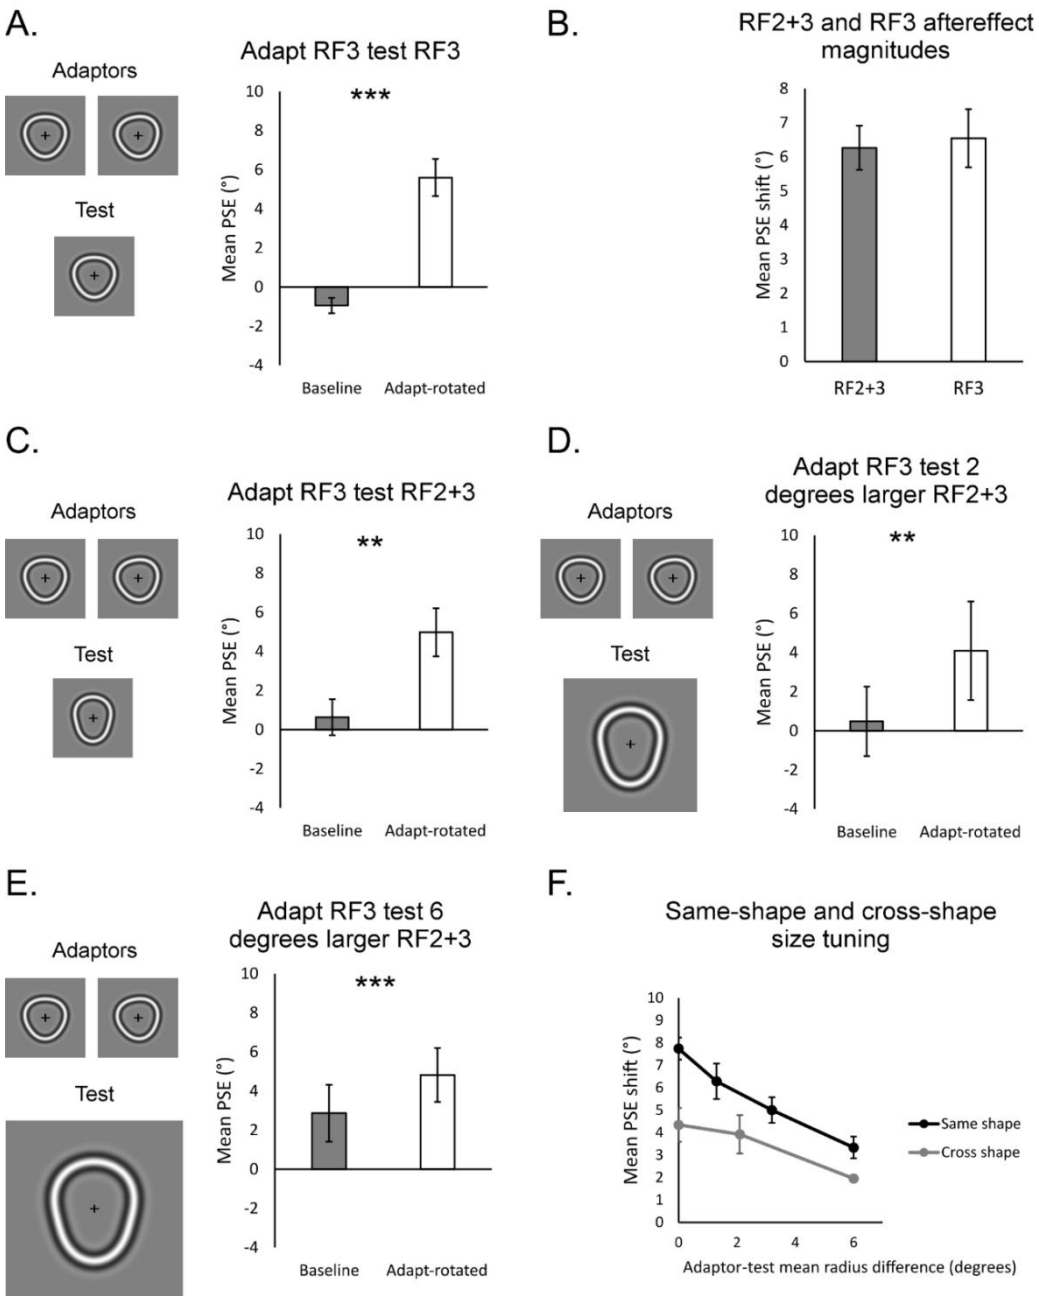

Figure 4.5: Shape aftereffects assessing the perception of RF3 and RF2+3 patterns following adaptation to RF3 patterns. In panels A and C-E average PSE measurements for the baseline and adapt-rotated conditions are shown with grey and white bars, respectively, and asterisks denote a significant difference according to a paired-samples t test (*** = $p < .001$ and ** = $p < .01$). (A) Data are shown for RF3 pattern adaptor and test stimuli. (B) The magnitude of the aftereffect shown in panel A (grey bar) is compared with the standard aftereffect in RF2+3 patterns (white bar) derived from the data shown in Figure 4.3A. In panels C-E data are plotted for (C) RF2+3 patterns following adaptation to RF3 patterns, (D) RF2+3 patterns which were 2.1 degrees larger than RF3 pattern adaptors, (E) RF2+3 patterns which were 6 degrees larger than RF3 pattern adaptors. (F) Data from panels C-E are replotted to show cross-shape aftereffect magnitudes as a function of adaptor-test stimulus size differences (grey line) alongside the same function for same-shape aftereffects in RF2+3 patterns (black line, data originally shown in Figure 4.3E). In all panels error bars show standard error of mean.

overlooked component of global shape aftereffects (Dickinson et al., 2010), rather than RF adaptation. To address this, it was determined whether this cross-shape aftereffect exhibited the same degree of size invariance as same-shape effects. Two experiments were conducted where the adaptor was an RF3 pattern and the test was an RF2+3 pattern, which was larger than the adaptor.

For the first experiment the test stimulus was 2.1 degrees of visual angle larger in mean radius than the adaptor and its spatial frequency was scaled (as in the original test large experiment shown in Figure 4.4A). The aftereffect was again significant ($t(4) = 4.60$, $p = .010$, Figure 4.5D) in 5 observers (3 naïve). In the second experiment the test stimulus was 6 degrees larger than the adaptor and its spatial frequency was scaled (equal to the maximum size difference tested in same-shape aftereffects). The aftereffect was also significant under these conditions ($t(3) = 25.84$, $p < .001$), Figure 4.5E) in 4 observers (2 naïve). Moreover, the effect of stimulus size is similar in the cross-shape aftereffect to in the same-shape aftereffect, where large size differences reduce the strength of the aftereffect to approximately 50% (Figure 4.5F). This therefore indicates some explicit adaptation to RF3 phase which transfers across both shape and large differences in stimulus size. However, as shown in Figure 4.5F the cross-shape aftereffect is clearly reduced in magnitude compared to the same-shape aftereffect. RF3 phase adaptation, therefore, cannot account for the entire illusion, and thus the contribution of other, more shape-specific mechanisms cannot be ruled out.

4.4: Discussion

This chapter presents evidence for a novel class of global shape aftereffect, where manipulation of the phase of a single RF component within a composite RF pattern adaptor affects subsequent shape perceptions. This was demonstrated in RF2+3 patterns where manipulation of the RF3 component phase in the adaptor affected the perceived shape of the test stimulus. Using a range of experiments, the nature of this aftereffect was characterised as well as the shape processing mechanisms which underpin it.

The aftereffect was confirmed to be a lower-level shape aftereffect as opposed to a higher-level face aftereffect. There was no apparent effect of stimulus inversion on the

aftereffect and, unlike high-level face aftereffects (Afraz & Cavanagh, 2008; Zimmer & Kovács, 2011), it did not transfer across large portions of the visual field. Moreover, the same RF3 component phase manipulation was used to produce a similar shape aftereffect in RF3+4 patterns, showing the same effect can be induced using non-face-like shapes. Together these results indicate that shape aftereffects in RF2+3 patterns are consistent with the attenuation of shape encoding mechanisms, as opposed to higher-level, face-specific mechanisms.

Shape aftereffects in RF2+3 patterns showed tolerance to changes in stimulus size, remaining significant in stimuli which were larger or smaller than the adaptor. This is consistent with other reported shape aftereffects (Anderson et al., 2007; Suzuki, 2001) and demonstrates global processing as the contours defining adaptor and test stimuli occupied separate portions of the retinal image. It was possible that residual aftereffects measured after stimulus size changes were due to response bias rather than perceptual adaptation (Morgan et al., 2012), however this was ruled out by replicating the aftereffect and the effect of test stimulus size using a bias free adaptation paradigm (Morgan, 2013). Shape aftereffects in RF2+3 patterns, therefore, tap into the same global shape mechanisms as other shape aftereffects in the literature, and the experiments described add that such shape aftereffects reflect genuine changes in perception and cannot be accounted for by simple manifestations of response bias.

Although aftereffects were still present in stimuli of different sizes, they were reduced in magnitude, consistent with other reported effects of stimulus size (Anderson et al., 2007). This suggests some degree of tuning to size, which may not be appreciable when only testing a few stimulus sizes. To address this, the aftereffect was measured across a broad range of test stimulus sizes. The aftereffect exhibited a broad Gaussian tuning profile, peaking when adaptor and test were the same size and declining with increasing size differences until it plateaued, remaining significant in up to size differences of a factor of three. The aftereffect also showed complete invariance to the cross-sectional spatial frequency of the test stimulus across the tested range. Moreover, when test size and spatial frequency were varied together the resulting tuning profile was similar to that for size changes alone, confirming that spatial frequency contributes nothing to the aftereffect. Together, these results demonstrate that the aftereffect cannot be explained by a manifestation of tilt aftereffects, which are tuned to spatial

frequency (Ware & Mitchell, 1974). Therefore, shape aftereffects in composite RF patterns cannot be mediated by high-level face-specific or low-level orientation-specific mechanisms, and instead likely engage intermediate-level shape mechanisms.

The size-tuning profile of the aftereffect constituted both a broad, Gaussian portion and a plateau. This is consistent with the results of Bell et al. (2014), who measured a shape aftereffect in single RF patterns across a range of orientation differences between adaptor and test and found a similar tuning profile with a plateau. Moreover, they provided evidence for two mechanisms, one that was selective for luminance polarity and captured the orientation-tuned portion of the aftereffect, and one that was not selective for luminance polarity and captured the plateau. These results suggest that the feature-tuned shape mechanism described by Bell et al. (2014) is also tuned to size. In addition, the results indicate that both of these mechanisms are invariant to spatial frequency, implying that the feature-tuned mechanism is invariant to some low-level features. Therefore, both of these mechanisms seem to be involved in some intermediate-level of shape processing, although one is higher-level than the other. Overall, the feature-tuned and feature-invariant components of the aftereffect imply increasingly abstract shape representations, consistent with dimensionality reduction (Wilson & Wilkinson, 2015) and increased complexity of the feature space in which shapes are represented (Güçlü & van Gerven, 2015) through the visual system.

Next, it was asked how the intermediate-level shape mechanisms which mediate the aftereffect might encode shape information. Global shape may be represented as a decomposition of the shape profile into its RF components (Bell & Badcock, 2009; Salmela et al., 2016). Alternatively, shape may be represented by the shape profile determined through some form of curvature analysis (Dickinson et al., 2015; Poirier & Wilson, 2006). The global shape aftereffect was used to disambiguate between these two possibilities. The presence of the RF2 component appears to have no masking effect on adaptation to RF3 component phase in RF2+3 patterns as the aftereffect was of equal magnitude in RF3 patterns and RF2+3 patterns. This provided evidence for shape over RF adaptation as RF components have been shown to have masking effects on the detection of neighbouring RFs when presented in a composite stimulus (Bell et al., 2009). However, this interpretation relies on the assumption that the masking effects found by Bell et al. for a shape detection paradigm should replicate in an adaptation

paradigm, which may not be true. If Bell et al.'s results were due to increased noise in the masked RF channel, rather than suppression of the masked RF channel, one may not expect this to affect adaptation as the relative overall weighting of different channels might remain the same. In order to determine whether shape aftereffects in RF2+3 patterns were due to shape or RF adaptation, therefore, further experiments were required.

In addition, evidence was also provided to support adaptation to RF3 component phase as adaptation to an RF3 pattern changed the perception of an RF2+3 pattern test stimulus. Moreover, this cross-shape aftereffect exhibited a similar tuning profile to test stimulus size as same-shape effects, remaining significant in stimuli with size differences of 6 degrees. Therefore cross-shape aftereffects cannot be accounted for solely by tilt aftereffects and imply the attenuation of RF-tuned channels (Bell & Badcock, 2009). It should also be noted, however, that cross-shape aftereffects were always smaller in magnitude than same-shape aftereffects. As such, RF3 phase adaptation is not sufficient to account for the entire illusion, which is likely mediated by a combination of adaptation to both RF3 phase and other shape features such as curvature.

Finally, this study highlights the possibility that mid-level shape adaptation could contribute to face aftereffects which are often assumed to involve face-specific processes. RF patterns convey shape information, which can describe head outlines (Wilson et al., 2000), act as a cue to face viewpoint (Daar & Wilson, 2012), and elicit responses from face-selective cortex (Wilkinson et al., 2000). The shape aftereffect in RF2+3 patterns shows some similarities to face aftereffects in the literature. For example, the face distortion aftereffect (Webster & MacLin, 1996) exhibits a similar tuning to stimulus size to this aftereffect, where it is reduced but still significant in stimuli with size differences of up to two octaves (Zhao & Chubb, 2001). These similar results imply that face distortion aftereffects could, at least in part, be explained by shape adaptation. Neural adaptation to faces may also involve shape processing mechanisms. For example, shape adaptation could contribute to neural adaptation to face viewpoint (Fang et al., 2007), particularly those effects localised to lateral occipital areas which are known to be shape-selective (Kourtzi & Kanwisher, 2001; Malach et al., 1995). Moreover, neural adaptation to faces in the ventral visual stream has been

shown to be invariant to stimulus size, but not to manipulations which affect object shape such as viewpoint (Andrews & Ewbank, 2004). This pattern of results could be consistent with size-invariant global shape mechanisms, implying that shape mechanisms may contribute to neural face adaptation effects. Overall, the results indicate that changing stimulus size is likely not sufficient to isolate face-specific mechanisms, as shape processing mechanisms are tolerant to stimulus size changes. Instead, translation across large portions of the visual field could be used to eliminate the contributions of mid-level shape processing mechanisms.

4.5 Conclusions

In conclusion, shape aftereffects in composite RF patterns cannot be explained by high-level, face-specific or low-level, orientation-specific mechanisms, and are instead mediated by intermediate-level shape mechanisms. These mechanisms appear to comprise at least two processing stages. The first is tuned to lower-level stimulus features such as luminance and orientation (Bell et al., 2014), and the experiments described add that it is also tuned to stimulus size, but not spatial frequency. The second is invariant to these features and may be the stage at which shapes are represented by size invariant features such as curvature relative to the centre (Dickinson et al., 2015; Poirier & Wilson, 2006) or RF content (Bell & Badcock, 2009; Salmela et al., 2016). Evidence was shown to support both of these alternatives, and suggest that global shape aftereffects in composite RF patterns are mediated by the attenuation of RF-tuned channels as well as other global shape mechanisms.

Chapter 5: Effects of TMS to LO2 and OFA on adaptation to and discrimination of RF2+3 patterns

Abstract

RF2+3 patterns are processed globally (Chapter 4) and therefore may be an apt stimulus for identifying the neural locus of global shape processing. Here transcranial magnetic stimulation (TMS) was used to disrupt neural processing in areas LO2 and occipital face area (OFA) during shape discriminations of RF2+3 patterns following shape adaptation or in the absence of adaptation. Following the results from Chapter 3, LO2 is a candidate for global shape processing. OFA preferentially responds to face outlines defined by radial frequency (RF) components, and therefore may process biologically relevant shape information in composite RF pattern stimuli. However, there was no clear effect of TMS to LO2 or OFA on adaptation to or discrimination of RF2+3 patterns. This could indicate that the global processing of RF2+3 patterns does not occur in LO2 or OFA, however potential confounds of local stimulus processing or perceptual learning reducing the efficacy of TMS challenge this conclusion.

5.1: Introduction

Chapter 4 characterised a shape aftereffect in composite radial frequency (RF) patterns. By manipulating the size and spatial frequency of the test stimulus it was demonstrated that the shape of RF2+3 patterns is processed by global-scale shape mechanisms. Therefore, identifying where in the human brain these stimuli are processed may shed light on the neural locus of global shape processing. While there is evidence that single RF patterns may be processed in LO2 (Chapter 3; Silson et al., 2013) and V4 (Wilkinson et al., 2000), the neural processing of composite RF patterns has not been investigated.

For this chapter transcranial magnetic stimulation (TMS) was used to disrupt the processing of RF2+3 patterns with the goal of identifying where in the human brain they are processed. Previous research has utilised TMS to disrupt the effects of perceptual adaptation (e.g. Burton et al., 2009; Cattaneo & Silvanto, 2008; Silvanto et al. 2007; Silvanto, Muggleton, & Walsh, 2008; Silvanto & Pascual-Leone, 2008), demonstrating a causal relationship between processing within the stimulated brain region and encoding of the adapted feature. TMS can also be used to impair the discrimination of a visual feature (e.g. McKeefry et al., 2008; Pitcher et al., 2007; Silson et al., 2013), demonstrating a causal importance for processing within the stimulated brain region and accurate perception (see Chapter 2 for a full review of TMS and its applications). Here it was investigated whether TMS to different areas of extrastriate visual cortex could disrupt adaptation to and impair shape discriminations of RF2+3 patterns.

One area that may process the shape of RF2+3 patterns is visual field map LO2 (Larsson & Heeger, 2006). LO2 exhibits robust responses to single RF patterns (Chapter 3) and is causally important in discriminating RF amplitude (Silson et al., 2013). LO2 is also a good candidate for the early stages of global shape processing as it overlaps with object-selective cortex, which performs complex processing of objects (Malach et al., 1995; Sayres & Grill-Spector, 2008; Silson et al., 2016; Vernon et al., 2016). While LO2 clearly plays a role in the processing of single RF patterns and may perform shape-related processes, its role in processing composite RF patterns, which may carry important shape information for biologically-relevant stimuli, is unknown.

As discussed in Chapter 4, composite RF patterns can describe the complex boundaries of natural shapes such as faces (Wilson et al., 2000; Wilson, Loffler, & Wilkinson, 2002; Wilson & Wilkinson, 2002). As such, they may be processed by regions of the brain which also respond to more complex, natural stimuli such as real faces. One candidate is the occipital face area (OFA), an area of lateral occipital cortex defined by a selective response to faces over scrambled faces or other objects (Isabel Gauthier et al., 2000; Kanwisher et al., 1997; Pitcher, Walsh, & Duchaine, 2011).

As shown in Chapter 4 (Figure 4.1), the shape manipulation of RF2+3 patterns is somewhat comparable to a face outline changing its viewpoint, and may therefore capture some of the RF information which can accurately describe face outlines (Wilson et al., 2000; Wilson et al., 2002; Wilson & Wilkinson, 2002). Indeed, composite RF patterns have been shown to carry information about face viewpoint as adaptation to an RF pattern head outline can induce a viewpoint aftereffect in full face test stimuli (Daar & Wilson, 2012). OFA is implicated in adaptation to face viewpoint (Fang et al., 2007) and is causally important for face viewpoint discrimination (Kietzmann et al., 2015). Moreover, OFA is responsible for the early stages of face processing (Haxby et al., 2000; Pitcher et al., 2007) and, relative to other face-selective regions, contains a stronger representation of face outlines defined by a combination of RF components compared to internal face features of full faces (Nichols, Betts, & Wilson, 2010). Consequently, it seems that RF2+3 patterns could be a shape that is well suited for processing in OFA.

While RF2+3 patterns are clearly simpler stimuli than faces and, as demonstrated in Chapter 4, are not processed by face-specific mechanisms, it is important to note that while OFA may be face-selective it is not face-specific. Indeed, TMS evidence shows that OFA plays a causal role in the discrimination of orientation (Silvanto et al., 2010) and symmetry (Bona, Cattaneo, & Silvanto, 2015) in non-face stimuli. Therefore, OFA appears to process visual features that are salient in faces, such as symmetry, in non-face stimuli. Based on OFA's sensitivity to face viewpoint, RF face outlines and features of non-face visual stimuli, it could be a candidate for the processing of shape in RF2+3 patterns.

Two experiments were conducted to investigate the effects of TMS to LO2 and OFA on shape perception of RF2+3 patterns. The first experiment replicated the shape aftereffect in RF2+3 patterns described in Chapter 4, and examined the effect of TMS on the perception of the test stimulus following shape adaptation. This allowed for the effect of TMS on both shape adaptation and shape discrimination to be measured. There was no effect of TMS on adaptation, but a marginal effect of TMS on discrimination. Therefore a second experiment was conducted to explore this effect further, where the effect of TMS on shape discriminations of RF2+3 patterns was measured in the absence of adaptation.

5.2: Experiment one: shape adaptation and discrimination

5.2.1: Methods

The first experiment measured the effect of TMS to LO2 and OFA on both shape adaptation to and discrimination of RF2+3 patterns. 6 subjects were tested in the experiment (five males, one female, mean age = 27.7, age range = 21- 47). All subjects had normal or corrected-to-normal visual acuity, no history of psychological or neurological impairment and gave written informed consent. The experiment was approved by the York Neuroimaging Centre Research Governance Committee and adhered to the declaration of Helsinki.

Retinotopic mapping was used to identify LO1 and LO2 and a functional localiser to identify OFA in all subjects. Subsequently, these regions were stimulated with TMS while subjects performed shape discriminations of RF2+3 patterns following shape adaptation. LO1 was the control TMS site as it is a neighbouring visual field map to LO2 (Larsson & Heeger, 2006), the stimulation of which has no effect on shape discriminations of RF patterns (Silson et al., 2013) and shows weaker responses to RF patterns than LO2 (Chapter 3).

5.2.1.1: General MRI data acquisition

All MRI data were acquired on a GE 3-Tesla Signa HD Excite scanner using a 16-channel head coil to improve signal-to-noise in the occipital lobe, except for one

structural scan which was acquired with an 8-channel head coil to improve alignments between functional and structural data and for coregistration between the subject's head and structural MRI data in TMS sessions. Structural and retinotopic mapping measurements were acquired using methods identical to those described in Chapter 3 and reviewed in Chapter 2.

5.2.1.2: Functional localiser

OFA was defined using a functional localiser (TR = 3000ms, TE = 30ms, voxel size = 2x2x2mm^3, flip angle = 90°, matrix size 96x96x39, FOV = 19.2cm). In a block-design fMRI experiment, subjects viewed blocks of greyscale faces, objects, scrambled faces and scrambled objects (stimuli shown in Figure 5.2 in the Results section) while they performed a one-back task. Stimuli were presented using the same MRI presentation equipment described in Chapter 3. There were 24 blocks of stimuli in each run (6 blocks per stimulus category). Stimulus blocks were 12 seconds long, with 12 stimuli presented in a block (stimuli on for 0.8 seconds with a 0.2 seconds interstimulus interval). Stimulus blocks were separated by 12 second fixation periods. During each block 11 unique stimuli were presented, with one (randomly selected) repeated for the one-back task. Total run length was 9 minutes 45 seconds and each subject completed two runs. Blocks were presented in a fixed, pseudo-randomised order, but the order in which individual images within stimulus categories were presented was randomised for each run.

The stimuli were 66 PNG images of faces and 66 PNG images of easily recognisable objects, manually extracted from their original backgrounds. Images were converted to greyscale and their image histograms were equalised so they all had the same mean luminance (200cdm^{-2}). On average, stimuli subtended 4x4 degrees of visual angle, though aspect ratios were somewhat variable. Scrambled stimuli were generated by splitting each image into 10 rows and columns. All squares lying within the convex hull of the face or object were randomly permutated and rotated. This generated scrambled images which contained all the local details and the same coarse outline as unscrambled stimuli but were unrecognisable. After scrambling, a Gaussian smoothing filter (SD = 1 pixel) was applied to smooth hard edges.

Data from the functional localiser were analysed using a general linear model (GLM) in mrVista where faces, objects, scrambled faces and scrambled objects were the regressors of interest. Between and within scan head movements were corrected for using mrVista's motion correction algorithms. The first 3 volumes of each scan were discarded to allow the scanner to reach stable magnetisation. A high pass filter with a cut off of 60 seconds was used to remove low frequency scanner drift from the data. The Boynton gamma function was used to convolve the GLM. To identify OFA, results from both faces vs scrambled faces and faces vs objects contrasts were considered. The contrast that yielded a cluster of face-selective voxels in an area of cortex that was most accessible with TMS was used to define OFA.

5.2.1.3: Identification of TMS targets

TMS was applied to the hemisphere in which LO1 and LO2 were most readily identifiable and organised in such a way that they could be targeted independently with TMS. This resulted in TMS being applied to the left hemisphere in three subjects and the right hemisphere in three subjects. Where possible, TMS was applied to the centres of gravity of LO1 and LO2. For two subjects, LO1 and LO2 centroids were too close together to be targeted independently (< 7.5 millimetres). For these subjects, the vector between LO1 and LO2 centroids was calculated and the LO1 target was moved along that vector until it was 7.5 millimetres from the LO2 centroid. For a third subject, due to cortical folding, it was not possible to apply TMS to LO1 and LO2 centroids independently, even though there were > 7.5 millimetres apart. For this subject TMS targets were manually positioned elsewhere within LO1 and LO2 such that independent targeting would be possible. For OFA, a faces vs scrambled faces contrast yielded a cluster of face-selective voxels which was most easily targetable for TMS for five subjects, and a faces vs objects contrast was used for one subject. Example retinotopy and localiser data sets with accompanying regions of interest are shown in Figure 5.2 in the Results section.

5.2.1.4: Psychophysical measurements of shape discrimination thresholds and aftereffects

The staircasing psychophysics procedures described in Chapter 2 were used to measure shape adaptation and discrimination thresholds. This allowed the definition of stimuli for the TMS experiment which were tailored to each individual's

discrimination ability and adaptation effects. Stimuli were generated and presented using Psychtoolbox 3.0 (www.psychtoolbox.com) in Matlab v7.14 (2012a x64, Mathworks) and presented on a gamma-corrected Mitsubishi Diamond Pro 2070^{SB} display with a refresh rate of 75Hz. Subjects viewed stimuli monocularly with their dominant eye from a chin rest placed 57cm from the display.

The stimuli were RF2+3 patterns generated using the methods described in Chapter 4 and reviewed in Chapter 2. Stimuli had a mean radius of 2.5 degrees of visual angle and the peak spatial frequency of the luminance border was 1.26 cycles per degree. Stimuli were presented 4 degrees of visual angle left or right of a black central fixation cross (0.5 degrees across). Stimuli were always presented in the contralateral hemifield to the hemisphere that would be stimulated with TMS, so that the retinotopic representation of the stimulus was in the stimulated hemisphere.

Subjects performed two psychophysics sessions. In the first the adaptor was a neutral, front-facing stimulus (baseline condition). In the second the adaptor was rotated to face 12° away from the observer (adapt-rotated condition; see Figure 5.1). The rotated adaptor was always rotated to face away from fixation (-12° in the left hemifield, 12º in the right). An RF2+3 pattern adaptor was presented for 30 seconds on the first trial, and 5 seconds on subsequent trials. To avoid local luminance adaptation, the contrast of the adaptor modulated between -100% and 100% by a 1Hz sinusoidal function. After a 1 second interstimulus interval, the test stimulus was shown for 200 milliseconds at 50% contrast. Subjects responded as to whether the test stimulus was rotated towards the left or the right with a key press. The next trial followed a 1.5 second intertrial interval.

The rotation of the test stimulus was controlled by interleaved pairs of 3-up 1-down, 2-up 1-down and 1-up 1-down staircases, acquiring data spread along the psychometric function. Staircases began at ± 16° with a step size of 3° which halved on reversals 3 to 5. Data from the first two reversals were discarded. Resulting data were used to identify the test stimulus rotation required for subjects to perform at 50% correct (PSE). The rotation required for subjects to perform at 84% correct was taken as the discrimination threshold. PSEs and thresholds were used to generate test stimuli for the TMS experiment that were tailored to each subject's individual adaptation effects and discrimination ability.

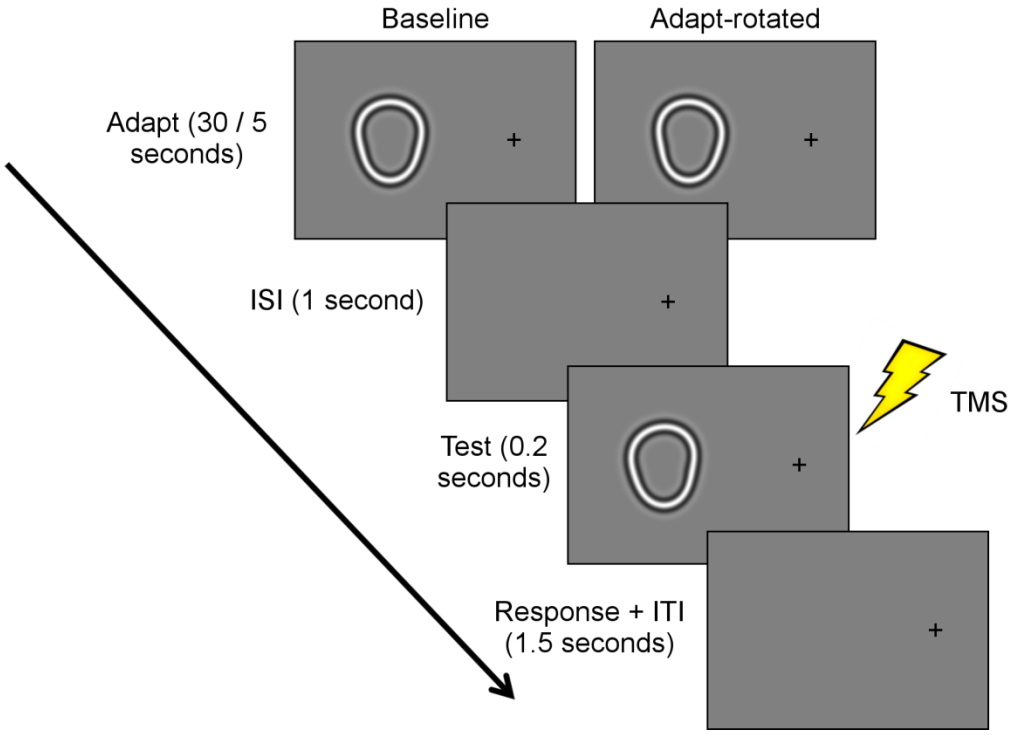

Figure 5.1: Stimulus presentation for the psychophysics and TMS experiments. Stimuli were presented 4 degrees of visual angle from fixation. In the TMS experiment, TMS was applied coincident with the test stimulus (see text for stimulation parameters).

5.2.1.5: TMS experiment

For the TMS task subjects viewed the same adapting stimuli as in the psychophysics experiment, displayed using the same equipment (Figure 5.1). The rotation of the test stimulus was determined by the subject's threshold measured from the psychophysics experiment. In baseline TMS sessions, test stimulus rotations were a threshold (measured from the baseline psychophysics session) away from zero (i.e. test rotation = ± threshold). For adapt-rotated TMS sessions, test stimulus rotations were a threshold away from the PSE measured in the adapt-rotated psychophysics session (test rotation = PSE ± threshold). One subject's PSE from the adapt-rotated psychophysics session was greater than the magnitude of the rotation in the adaptor (> 12º, subject 6 in Table 5.2). For this subject test stimuli for adapt-rotated TMS sessions were centred on the rotation of the adaptor (test rotation = 12º ± threshold).

Subjects completed baseline and adapt-rotated adaptation sessions for four TMS conditions: no TMS, control TMS (LO1), LO2 TMS and OFA TMS. Each session comprised 120 trials. 50 trials were conducted for left- and right-facing test stimuli. The remaining 20 trials were catch trials where the test stimulus was a neutral, front-facing stimulus. This was to ensure that adaptation had the expected perceptual effect in

adapt-rotated TMS conditions and data from catch trials were not used for any other analyses. On average, subjects responded in the predicted direction given the rotation of the adaptor on 97.5% of catch trials (SD = 1.95) in adapt-rotated TMS conditions. The sequence of trials was pseudorandomised so that no more than three of the same trial type could occur in sequence. After 60 trials there was a pause in the experiment, where the experimenter switched over to a fresh TMS coil to prevent coils from overheating. Therefore subjects adapted for 30 seconds on the 1st and 61st trials and for 5 seconds on all other trials. All other aspects of the experimental procedure and stimuli were identical to those used for the psychophysics experiment.

The location of the TMS coil was tracked in real time using the TMS Brainsight system (Rogue Research) to ensure TMS sites were stimulated accurately. Data from any trials where the location of TMS was greater than 2mm from the target were removed before analysis. TMS was applied using a 50mm Figure of 8 biphasic coil powered by a Magstim Rapid 2 stimulator (Magstim). A train of four pulses, separated by 50ms was applied over 200ms, coincident with test stimulus onset, at 70% maximum stimulator output (2.6 Tesla). TMS was applied coincident with stimulus onset as past research has shown this can affect both discrimination ability (Silson et al., 2013) and the effects of adaptation (Silvanto et al., 2007; Burton et al., 2009). The TMS protocol adhered to published safety guidelines (Wassermann, 1998; Rossi et al., 2009).

5.2.2: Results

5.2.2.1: Identification of LO1, LO2 and OFA

LO1 and LO2 were successfully identified in at least one hemisphere of all subjects. Results from the functional localiser scan also revealed a patch of face-selective cortex in an area of lateral occipital cortex inferior and anterior to LO2 (see Figure 5.2). The centroid coordinates of LO1 and LO2 match closely those provided by Larsson and Heeger (2006) and OFA coordinates closely match peak voxel coordinates published by (Rossion et al., 2003; see Table 5.1).

5.2.2.2: Shape discrimination thresholds and aftereffects

A cumulative Gaussian was fit to subjects' data from the psychophysics experiment using the methods described in Chapter 2 (see Figure 5.3). This was used to derive PSEs and 84% correct thresholds, which are shown for all subjects in Table 5.2. It is

clear that both thresholds and PSEs vary across subjects, emphasising the importance of tailoring stimuli to each subject's discrimination ability and perceptual adaptation for the TMS experiment.

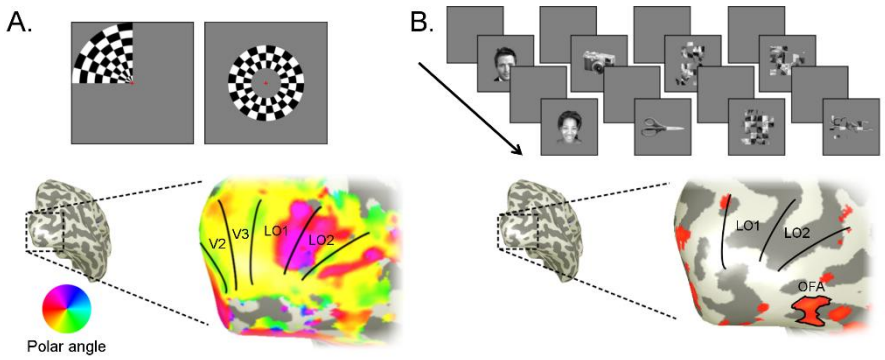

Figure 5.2: Results from retinotopic mapping and functional localiser procedures for a single representative subject. (A) Checkerboard wedge and ring stimuli were used to map the visual field. LO1 and LO2 were identified using the criteria described by Larsson and Heeger (2006) and reviewed in Chapter 2. (B) A functional localiser including blocks of faces, objects, scrambled faces and scrambled objects was used to identify OFA. OFA was identified as a patch of face-selective cortex in lateral occipital cortex anterior and inferior to LO2. Results from a faces vs objects contrast are shown in this example.

For half of the subjects, stimuli were presented in the left visual hemifield, meaning a left-facing adaptor was used for the adapt-rotated condition, causing a negative aftereffect (see Figure 5.3, subject 3). In order to be able to compare across all subjects, the sign of baseline and adapt-rotated PSEs was reversed for these subjects (as was also done in the psychophysics experiments in Chapter 4). PSEs measured from the two adaptation conditions were significantly different from each other ($t(5) = 8.20$, $p < .001$), but there was no significant difference between thresholds (t (5) 0.28, $p = .793$). The shape aftereffect in RF2+3 patterns from Chapter 4 was therefore replicated, and adaptation had no effect on discrimination ability.

5.2.2.3: Effects of TMS on shape adaptation and discrimination

First, each subject's PSE was estimated for every TMS session. The proportion of right-facing responses for low and high threshold test stimuli (catch trials excluded) were plotted against test stimulus rotation and a vector was drawn between these two points (see Figure 5.4A). This vector was used to estimate the PSE test stimulus rotation at which subjects should perform at 50% correct. Average PSEs for each session are

Table 5.1: Talairach coordinates (mm) of left (n = 3) and right (n = 3) LO1, LO2 and OFA centroids (mean and range are shown) from experiment one and right LO1, LO2 and OFA from experiment two (n = 7). Published coordinates of LO1 and LO2 centroids from Larsson and Heeger (2006) and OFA peak voxel coordinates from Rossion et al. (2003) are also included for comparison.

		LO1			LO2			OFA		
		X	Y	Z	X	Y	Z	X	Y	Z
Experiment one	Right	27	-87	6	32	-82	3	38	-78	-8
	Range	21, 30	-92, -83	2, 10	27, 37	-88, -79	0, 5	35, 43	-81, -75	-10, -4
	Left	-26	-88	3	-28	-84	1	-35	-82	-8
	Range	-34, -15	-96, -79	0, 6	-36, -21	-94, -77	-5, 4	-42, -22	-92, -75	-10, -5
Experiment two	Right	26	-85	0	34	-80	1	43	-74	-7
	Range	23, 30	-90, -79	-9, 8	29, 38	-84, -73	-7, 6	34, 49	-80, -65	-18, 0
Larsson and Heeger (2006)	Right	32	-89	2.6	38	-82	0.6	–	–	–
	Range	24, 38	-98, -81	-8, 13	32, 46	-89, -72	-13, 12	–	–	–
	Left	-31	-90	1.4	-38	-83	-0.1	–	–	–
	Range	-37, -26	-101, -82	-12, 11	-43, -31	-92, -67	-13, 11	–	–	–
Rossion et al. (2003)	Right	–	–	–	–	–	–	38	-80	-7
	Range	–	–	–	–	–	–	27, 48	-93, -69	-19, 0
	Left	–	–	–	–	–	–	-34	-81	-14
	Range	–	–	–	–	–	–	-37, -29	-83, -78	-18, -11

plotted in Figure 5.4B. A 4x2 repeated measures ANOVA was used to determine the effects of TMS condition and adaptation condition on PSE. The main effect of adaptation was significant ($F(1, 5) = 188.26$, $p < .001$), showing that the shape aftereffect in RF2+3 patterns was replicated in the TMS experiment. However, there was no significant effect of TMS condition (Huynh-Feldt corrected $F(1.22, 6.10) = 1.34$, $p = .302$) and no significant interaction (Huynh-Feldt corrected $F(1.32, 6.61) = 0.51$, $p = .552$). Overall, therefore, the shape aftereffect was replicated in the TMS experiment, however neither TMS to LO2 nor OFA had any effect on the magnitude of the aftereffect.

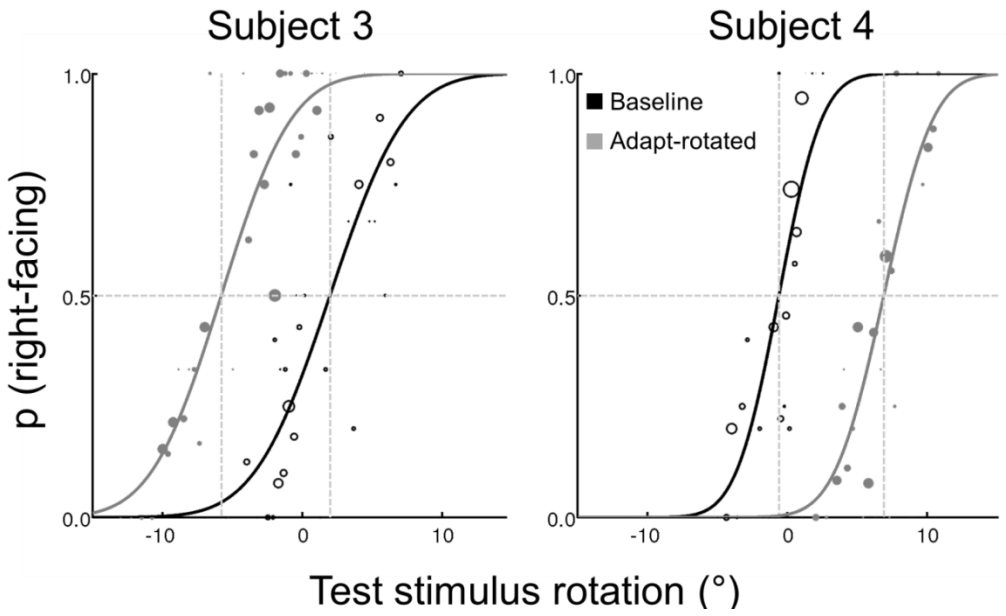

Figure 5.3: Results from the psychophysics experiment for two subjects. The proportion of right-facing responses is plotted against the rotation of the test stimulus for baseline (black) and adapt-rotated (grey) conditions. The relative sizes of data points are directly proportional to the amount of data acquired for a given test stimulus. Subject 3 adapted to a left-facing stimulus in the adapt-rotated condition, resulting in a negative PSE shift. Subject 4 adapted to a right-facing stimulus, resulting in a positive PSE shift.

Table 5.2: PSEs and 84% correct thresholds for all subjects measured from the psychophysics experiment for baseline and adapt-rotated conditions. Asterisks identify the 3 subjects for whom adapt-rotated adaptor stimuli were left-facing. PSEs for these subjects have had their signs reversed. Group means are shown at the bottom with standard deviations in parentheses.

Subject	Baseline		Adapt-rotated	
	PSE	Threshold	PSE	Threshold
1	0.67	4.08	8.12	4.24
2*	3.49	3.69	10.72	3.80
3*	-1.94	4.25	5.82	3.94
4	-0.63	2.41	6.85	2.78
5	-2.27	5.92	7.99	5.06
6*	-0.63	4.82	13.42	5.76
Group mean (SD)	-0.22 (2.1)	4.19 (1.17)	8.82 (2.79)	4.26 (1.04)

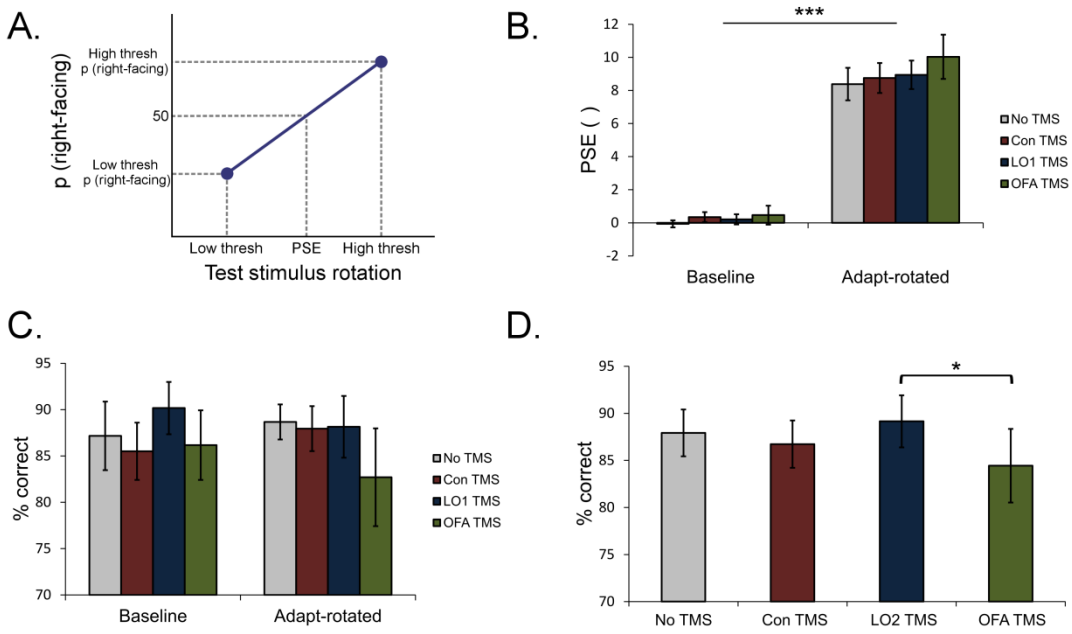

Figure 5.4: Results from the TMS experiment. (A) A schematic of how PSE was estimated from data acquired from TMS sessions. The proportion of right-facing responses was plotted against test stimulus rotations for low and high threshold stimuli and a vector was drawn between the two points to estimate PSE. (B) Average PSE estimates across adaptation and TMS conditions. *** denotes a significant main effect of adaptation showing a replication of the shape aftereffect (see text for statistical details. (C) Average discrimination performance across adaptation and TMS conditions. (D) Average discrimination performance, collapsed across adaptation conditions for each TMS condition. * denotes a notable difference between LO2 and OFA TMS conditions, but this was not evaluated statistically due to the lack of a main effect (see text for statistical details). In all panels, error bars show standard error of mean.

Next, discrimination performance (Figure 5.4C) was calculated by averaging the proportion of trials where subjects reported left-facing for low threshold stimuli and right-facing for high threshold stimuli (catch trials again excluded). A 4x2 repeated measures ANOVA was coducted to determine the effects of adaptation condition and TMS condition on shape discrimination. There was no effect of adaptation on discrimination (F (1, 5) = 0.05, p = .826), showing that, similar to the psychophysics experiment, adaptation did not affect discrimination ability. The main effect of TMS approached significance, but was not significant when a correction was applied for breaking the assumption of sphericity (Huynh-Feldt correct F (1,27, 6.36) = 3.51, p = .104). There was no significant interaction (F (3, 15) = 0.68, p = .577).

As there was no effect of adaptation on discrimination ability, performance was averaged across the two adaptation conditions for each TMS condition to get a clearer picture of the effect of TMS on discrimination (Figure 5.4D). Overall TMS had a marginal effect on discrimination ability. Subjects performed worse when TMS was applied to OFA compared to LO2, implying a greater role of OFA in detecting the shape of RF2+3 patterns than LO2. However, neither TMS to OFA nor LO2 resulted in significantly different discrimination performed from either control TMS condition.

Since TMS had no effect on the magnitude of shape aftereffects, and adaptation had no effect on discrimination performance, the only effect found was a marginal effect of TMS on discrimination performance. Therefore a second experiment was conducted to examine this effect more closely. The following experiment measured the effects of TMS to LO2 and OFA on shape discrimination of RF2+3 patterns in the absence of adaptation.

5.3: Experiment two: Shape discrimination

5.3.1: Methods

Experiment two measured the effects of TMS to LO2 and OFA on shape discriminations of RF2+3 patterns in the absence of adaptation. 7 subjects (4 males, 3 females, mean age = 29.1, age range = 23 - 48) with normal or corrected to normal vision and no history of psychological or neurological impairment were tested. 4 of these subjects had also taken part in experiment one. All subjects gave written informed consent. The experiment was approved by the York Neuroimaging Centre Research Governance Committee and adhered to the declaration of Helsinki.

5.3.1.2: Identification of TMS targets

LO1, LO2 and OFA were identified using the same methods as for experiment one. All areas were identifiable in the right hemisphere for all subjects and were targetable with TMS. Therefore stimulated the right hemisphere of all subjects was stimulated as right OFA appears to be more consistently implicated in face processing than left OFA (see Pitcher et al., 2011 for a review). To ensure that all areas could be stimulated independently, TMS targets were positioned a full centimetre apart from each other. For 6 of 7 subjects, the centroids of all areas were more than 1 centimetre apart and

could be stimulated independently. For the 7th subject LO1 and LO2 centroids were too close together, so the vector between LO1 and LO2 centroids was measured and the LO1 target moved along that vector until it was 1 centimetre away from the LO2 target. For OFA, a faces vs scrambled faces contrast revealed a patch of face-selective cortex that was most accessible with TMS for 2 subjects, and a faces vs objects contrast was used for the other 5.

5.3.1.3: Psychophysical measurements of shape discrimination thresholds

As in experiment one, a psychophysics experiment was conducted to measure shape discrimination thresholds from each subject. The method of constant stimuli was used to measure thresholds (described in Chapter 2). The same task from experiment one was used (Figure 5.1), except the adaptor was replaced with a neutral, front-facing reference stimulus presented for 200 milliseconds, and the intertrial interval was extended to 2 seconds. Subjects judged the rotation of a test RF2+3 pattern relative to the reference pattern.

Each session comprised 70 trials. 7 different test stimuli were presented in a session (10 trials per stimulus). The rotations of the 7 test stimuli were linearly spaced around 0°. Pilot sessions were used to establish the test stimulus range that best suited each subject and would allow for data spread evenly around their psychometric function. For example, if a 10° range was used the test stimulus rotations were -10°, -6.67°, -3.33°, 0°, 3.33°, 6.67° and 10°. All other methodological and stimulus parameters were identical to those used for experiment one.

Once the ideal test stimulus range was identified for a subject, 5 sessions were completed with that range, the data from which were averaged. The average data from the total of 350 trials were then used to model a psychometric function. This was used to estimate a 75% correct threshold; equal to that used by Silson et al. (2013).

5.3.1.4: TMS experiment

fMRI-guided TMS was used to accurately stimulate areas LO1 (control), LO2 and OFA during shape discriminations using the same methods and equipment as experiment one. Stimulus presentation was the same as in the method of constant stimuli psychophysics experiment described above. The application of TMS was coincident with the presentation of the test stimulus and stimulation parameters were the same as

in experiment one. Each TMS session comprised 200 trials. There was a pause in the experiment after trials 80 and 140, where the experimenter applied a fresh TMS coil.

5.3.2: Results

5.3.2.1 Identification of LO1, LO2 and OFA

LO1, LO2 and OFA were successfully identified in the right hemisphere of all subjects using the same methods as experiment one. The average centroid coordinates of all these areas for experiment two are displayed in Table 5.1 alongside those from experiment one.

5.3.3.2: Shape discrimination thresholds

The data from the method of constant stimuli psychophysics experiment were used to identify 75% correct discrimination thresholds for each subject. The average proportion of right-facing responses (averaged across 350 trials) was plotted against test stimulus rotation and a cumulative Gaussian was fit to the data (see Figure 5.5). This Gaussian was used to estimate the test stimulus rotation required for subjects to perform at 75% correct. In the TMS experiment, test stimuli were presented with a rotation of ± threshold. Thresholds for all subjects are displayed in Table 5.3.

5.3.3.3: Effect of TMS on shape discriminations

For each TMS condition, the average % correct across the 7 subjects was calculated (Figure 5.6). A one-way repeated measures ANOVA revealed that TMS had no effect on shape discriminations (Huynh-Feldt corrected $F(2.46, 14.75) = 0.35$, $p = .754$). Moreover, it is can be seen from Figure 5.6 that, on average, subjects' performance was close to 75% correct (equal to the threshold used to define test stimuli) for all conditions. Therefore, in the absence of adaptation, neither TMS to LO2 nor OFA had any effect on subjects' shape discrimination ability. Also, the difference in performance between TMS to LO2 and OFA from experiment one did not replicate. Overall the results suggest that neither LO2 nor OFA play a causal role in discriminating the shape of RF2+3 patterns.

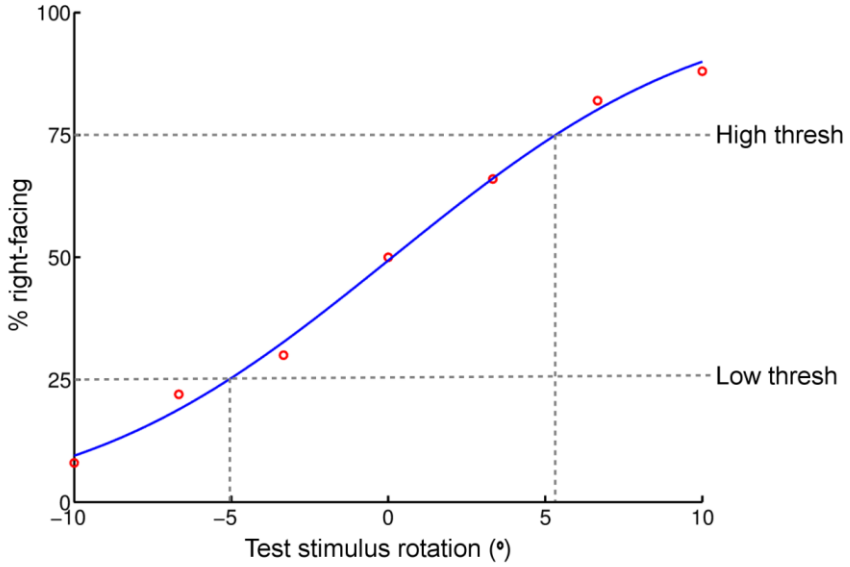

Figure 5.5: Results of the method of constant stimuli psychophysics experiment for one example subject (subject 2 in Table 5.3). The proportion of right-facing responses, averaged over 350 trials, was plotted against the rotation of the test stimulus for the 7 stimuli presented. A cumulative Gaussian was fit to the data to estimate the test stimulus rotation required for the subject to perform at 75% correct. This value was used as the threshold to define test stimuli for the TMS experiment, where stimuli were presented at ± threshold.

Table 5.3: 75% correct thresholds for all subjects measured from the method of constant stimuli psychophysics experiment. The group average is displayed at the bottom with standard deviation in parentheses.

Subject	75% correct threshold (°)
1	1.46
2	5.21
3	3.88
4	7.40
5	2.16
6	4.00
7	2.21
Average (SD)	3.76 (2.01)

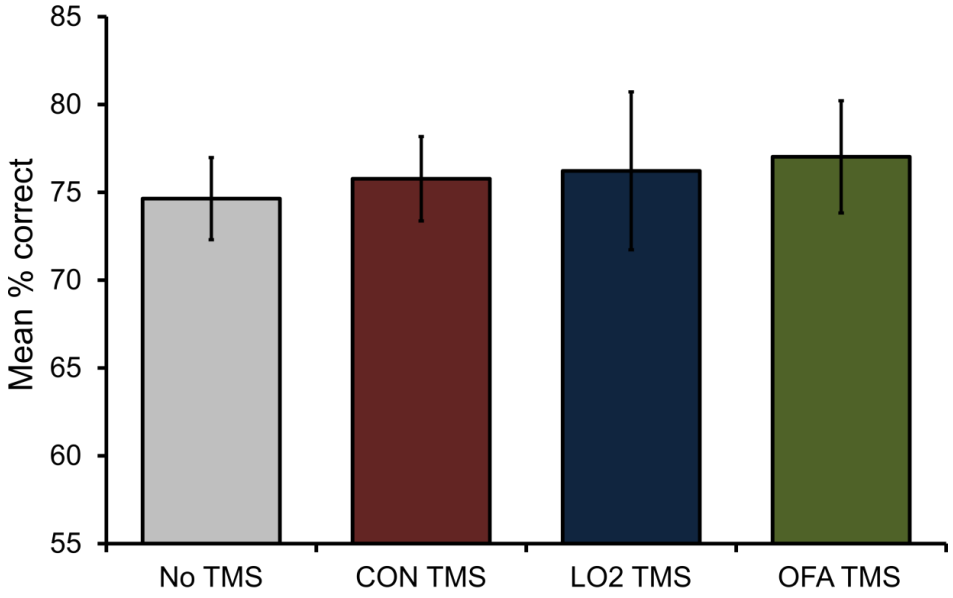

Figure 5.6: Average shape discrimination performance for no TMS, control TMS (LO1), LO2 TMS and OFA TMS conditions. Error bars show standard error of mean.

5.4: Discussion

This chapter measured the effect of TMS to areas LO2 and OFA on the shape aftereffect reported in Chapter 4 and shape discrimination of RF2+3 patterns. In the first experiment, TMS had no effect on the magnitude of the shape aftereffect. This suggests that the shape of the adaptor was not encoded in LO2 or OFA. However, relieving the effects of adaptation requires TMS to act in a state-dependent manner, either by facilitating the activity of adapted neurons, or suppressing unadapted neurons, or both (Silvanto et al., 2007; Silvanto & Pascual-Leone, 2008). The effect of TMS on perceptual adaptation can be inconsistent, and there are reports of TMS affecting all neural populations equally regardless of their activity state following adaptation (e.g. Burton et al., 2009). In such a case, one would not expect TMS to relieve the aftereffect as a uniform effect of TMS would result in the relative sensitivities between adapted and unadapted neural populations remaining the same. As such, no effect of TMS on adaptation may not preclude shape encoding of RF2+3 patterns in LO2 and OFA, so long as TMS affects shape discrimination.

There was a marginal effect of TMS on shape discrimination in the first experiment, where performance was worse when TMS was applied to OFA compared to LO2. However, neither LO2 nor OFA TMS conditions resulted in performance that was different from LO1 TMS or no TMS control conditions. To better understand the effect

of TMS on shape discrimination, a second experiment was conducted to measure discrimination in the absence of adaptation. The effect of TMS on shape discrimination from experiment one did not replicate; none of the TMS conditions had any effect on performance. Therefore there was no evidence for a clear causal role for LO2 or OFA in discriminating the shape of RF2+3 patterns.

It is possible that TMS had no effect on discrimination because the task could be solved on a local scale, by computations performed in earlier visual areas than the ones stimulated. Although RF2+3 patterns are clearly processed globally (see Chapter 4), there is also a more local component to their processing which was only eliminated when the sizes of adaptor and test stimuli were changed. This local component may have been sufficient to compensate for any global processing impairment as a result of TMS to later visual areas. Indeed, in Chapter 3 early visual areas were more engaged by a task that could be solved on a local scale (amplitude task) than a more global task (orientation task). This demonstrates that earlier visual areas can be engaged by stimuli that are processed globally if attention is drawn to a local element of the stimulus. The shape changes subjects were asked to discriminate were subtle, particularly at the threshold levels stimuli were presented at, and may have required attention to a local part of the contour rather than the whole stimulus. This potentially highlights a necessity for removing the contributions of local processing to find effects of TMS to more global processing regions such as LO2 and OFA. For example, Silson et al. (2013) randomised the orientations of reference and test stimulus to find an effect of TMS to LO2 on RF amplitude discriminations, removing local contour changes as a cue. It is possible that changing the sizes of adaptor (or reference for experiment two) and test stimuli could eliminate the contribution of local processing and result in an effect of TMS to LO2 or OFA on shape discrimination performance.

One area, which could not be stimulated due to its anatomical positon on the ventral surface of the occipital lobe that might encode the shape of RF2+3 patterns is V4. Macaque V4 encodes shape and curvature information (Pasupathy & Connor, 1999, 2001, 2002; Roe et al., 2012) and neuroimaging evidence shows that human V4 also responds to curvature (Dumoulin & Hess, 2007; Wilkinson et al., 2000). Moreover, V4 responded strongly to RF patterns in Chapter 3. Residual processing of curvature in V4 may have contributed to a lack of effect of TMS to LO2 or OFA on shape discrimination

in these experiments. This highlights a challenge of using TMS to probe neural functional organisation; one must choose a limited number of areas to be stimulated, and cannot control for or investigate the contributions of unstimulated brain regions.

An additional challenge of using TMS is that effects can vary depending on the subjects' familiarity with the task or stimulus. The efficacy of TMS is reduced as a function of increased perceptual learning (Neary, Anand & Hotson, 2005). Since extensive psychophysical testing was used to define the stimuli for TMS experiments, it is possible that this process resulted in perceptual learning which reduced the potential effect of TMS. Although TMS effects have been found in stimuli defined through psychophysical procedures before (Silson et al., 2013), the problem of perceptual learning may have been compounded in these TMS experiments. In a single TMS session, a total of two different test stimuli were presented (excluding catch trials in experiment one). Therefore, subjects were only required to learn to recognise a specific pair of test stimuli to succeed at the task. This learning could have reduced the effect of TMS due to perceptual learning of the limited range of test stimuli. It is therefore possible that randomising the rotation of a reference stimulus and defining test stimuli relative to the reference, increasing the total number of different test stimuli, could limit perceptual learning and restore TMS effects.

One positive of the study is that, for both TMS experiments, the anatomical positions of ROIs were highly consistent with those published in the literature (shown in Table 5.1). Moreover, any trials from TMS sessions where the locus of TMS was 2mm or more from the ROI centroid were removed from analyses, ensuring only trials where TMS was delivered accurately were included. One can therefore be confident that the null results are not due to errors in the way ROIs were localised or human error affecting TMS accuracy.

Ultimately, the results of this chapter are relatively uninformative. While TMS can provide important insights into the causal role of visual areas for visual tasks (McKeefry et al., 2009; Silson et al., 2013), null results can be challenging to interpret. The null result could equally be due to no causal role of the stimulated regions in the task used, residual processing in unstimulated regions, or methodological choices which impacted the potential efficacy of TMS. In the next chapter, fMRI will be used to

explore the sensitivity of many more visual areas to shape and curvature, with an aim to identifying where in the visual hierarchy a shift from local to global shape processing occurs.

5.5: Conclusions

In conclusion, there was no clear effect of TMS to LO2 or OFA on shape adaptation or shape discrimination in RF2+3 patterns. This may indicate that LO2 and OFA do not perform global shape processing of RF2+3 patterns, however reported inconsistencies in how TMS effects adaptation (Burton et al., 2009) and potential confounds of local processing and perceptual learning (Neary et al., 2005) make it difficult to draw this conclusion. Consequently, the neural locus of global shape processing in humans remains unclear. Responses to three-lobed RF patterns suggest LO2 could be a good candidate for early global shape processing (Chapter 3), however the TMS results cannot determine whether LO2 or OFA perform the global processing of RF2+3 patterns described in Chapter 4. The following chapter will use a broader approach to examine fMRI responses to a variety of RF patterns throughout visual cortex.

Chapter 6: Neural tuning to radial frequency patterns reveals global shape processing is localised to the lateral occipital cortex

Abstract

Radial frequency (RF) patterns engage global-scale shape processing mechanisms, however the neural locus of global shape processing remains unclear. This chapter describes an experiment that measured neural responses to a large range of RFs, and modelled neural tuning to RF in early, lateral and ventral occipital cortex. Tuning to RF was modelled by a Gaussian neural model defined in RF space. Psychophysical evidence shows that low RFs are processed globally, while high RFs are processed locally. Selective tuning to low RFs was localised to lateral occipital cortex, and emerged in visual field maps LO1 and LO2 and persisted through object-selective lateral occipital complex (LOC). This was true across two experiments, where the second experiment used control stimuli that were matched for contrast energy and spatial frequency. To better understand the processes underpinning neural tuning to RF in different visual areas, RF tuning profiles were correlated with low-level (measuring stimulus contrast energy) and shape (measuring stimulus deviations from circularity) predictors. RF tuning in early visual areas correlated highly with stimulus contrast energy, however only LO2 and LOC correlated with stimulus shape over contrast energy. Results indicate that global shape processing of RF patterns is localised to lateral occipital cortex, particularly LO2 and LOC, and demonstrate that modelling neural tuning to shape dimensions can be used to successfully map out a representational shape space in the human brain.

6.1 Introduction

Chapter 4 used shape adaptation in composite radial frequency (RF) patterns to provide evidence for global shape processing and neural encoding of RF. Chapter 5, however, was unsuccessful in identifying the neural locus of these mechanisms using transcranial magnetic stimulation to disrupt global shape processing. The neural locus of global shape processing therefore remains unclear. This chapter will use functional magnetic resonance imaging (fMRI) to measure neural tuning to RF across a wide range of RFs, with an aim to identifying selective tuning to low RFs which are known to be processed globally (Hess, Wang, & Dakin, 1999; Jeffrey, Wang, & Birch, 2002; reviewed in Chapter 1).

To measure tuning to RF, responses to RF patterns were modelled using a Gaussian neural model defined in RF space. Classically, this approach has been used to measure population receptive fields (pRFs) in visual cortex and provides a more accurate measure than traditional retinotopic mapping methods (Dumoulin & Wandell, 2008). As well as pRF mapping, this method has been used to successfully measure neural tuning to numerosity (Harvey et al., 2013), suggesting it may be an effective method for capturing the neural representation of stimulus dimensions. Indeed, the neural mechanisms generally accepted to underpin perceptual adaptation assume Gaussian-like tuning profiles to stimulus dimensions that could be modelled in this way (Thompson & Burr, 2009; reviewed in Chapter 1).

Recordings from macaque monkeys show that neurons in V4 and inferotemporal cortex (IT) are tuned to shape (Cadieu et al., 2007; Gallant et al., 1993; Pasupathy & Connor, 1999, 2001, 2002; Roe et al., 2012) . Importantly, neurons in IT respond parametrically to the configuration of composite RF patterns in a low dimensional shape space (Op De Beeck et al., 2001) indicating that shape representations have dimensional properties that could be modelled using the methods of Dumoulin and Wandell (2008) and Harvey et al. (2013). Moreover, responses in human lateral occipital cortex (LO) reflect shape similarity in terms of number of RF components and RF amplitude (Drucker & Aguirre, 2009; Vernon et al., 2016). Therefore, modelling sensitivity to the RF dimension may be an effective way of identifying shape tuning that could facilitate global shape processing in humans.

Likely regions that might process global shape are LO and ventral occipital cortex (VO). In LO, object-selective responses in the lateral occipital complex (LOC; Malach et al., 1995) may be facilitated by sensitivity to global shape (Kourtzi & Kanwisher, 2001) and emerge from an increasing sensitivity to shape in visual field maps LO1 and LO2 (Chapter 3; Silson et al., 2013; Vernon et al., 2016). In VO, area V4 shows preferential responses to RF patterns (Wilkinson et al., 2000) and concentric global shapes over unconnected curvature arcs (Dumoulin & Hess, 2007). It was therefore predicted that tuning to RF indicative of global shape processing might be found in these regions.

One study has attempted to measure neural sensitivity to RF and radial amplitude using a neural similarity analysis (Salmela et al., 2016). Counter to their predictions, this study did not find any correlation between neural responses and stimulus RF in LO or VO. However, the manipulations of frequency and amplitude used also generated changes in low-level stimulus features such as contrast and spatial frequency. Stimulus contrast was an issue because increases in RF or amplitude increase the total path length of the contour, resulting in a greater amount of contrast energy that defines the stimulus. Spatial frequency was an issue because stimulus luminance profiles were defined in polar coordinates along the radial line which, as reviewed in Chapter 2, creates an inhomogeneous cross-sectional spatial frequency profile within the stimulus. Moreover, this issue is not independent of amplitude or RF, as increases in both are accompanied by more rapid changes in radius towards the centre, resulting in greater pinching of the spatial frequency profile towards the stimulus centre. The result of this is that linear manipulations of RF or amplitude were accompanied by nonlinear changes in spatial frequency across stimuli (see Figure 6.1). Indeed, the neural responses from this study were generally dominated by spatial frequency, which may have obscured any representation of RF in LO and VO. RF did correlate with areas V2, V3, V3A/B and IPS0, however responses in these areas also correlated with stimulus contrast and spatial frequency and, since RF is confounded by these features, it cannot be determined whether correlations with RF account for any variance not accounted for by contrast and spatial frequency. It therefore remains to be seen whether true RF tuning can be identified in LO and VO using stimuli that control for low-level features and a novel analysis technique where tuning in RF space is explicitly modelled.

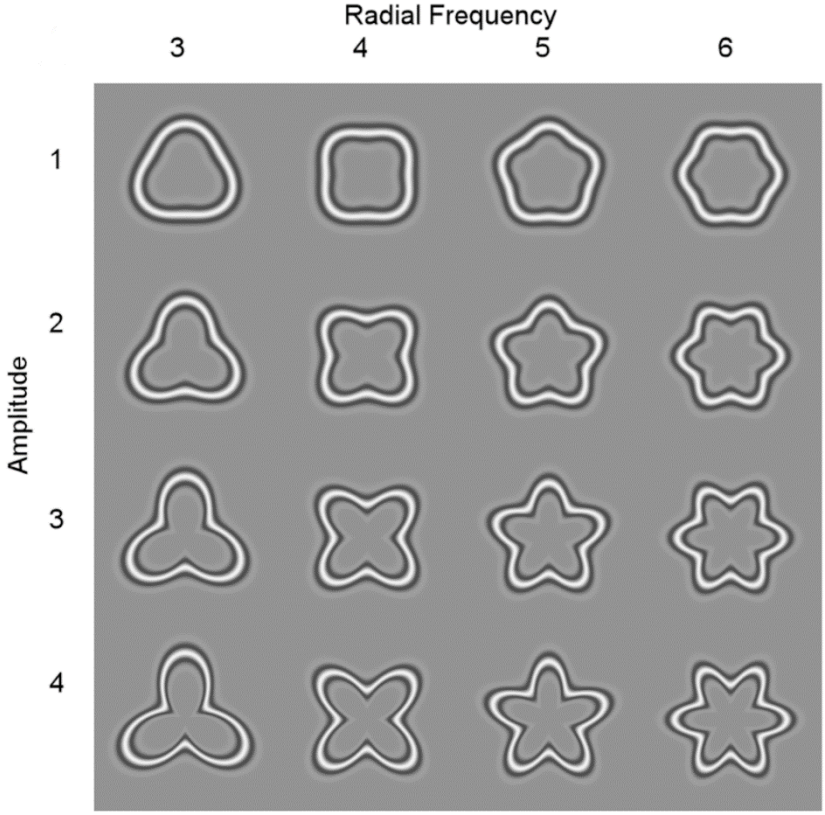

Figure 6.1: Stimuli used by Salmela et al. (2016) to measure neural responses to RF and amplitude. Luminance profiles were defined in polar coordinates according to the distance between the pixel and the contour along the radial line. As reviewed in Chapter 2, this creates pinching of the spatial frequency profile towards the centre of the stimulus. This results in unmatched spatial frequency between stimuli that differ in RF or amplitude. This was corrected for in the stimuli shown in Figure 6.2 (in the Methods section below) by defining the luminance profile in Cartesian coordinates according to the distance between the pixel and the contour along the line with the shortest distance to the contour (reviewed in Chapter 2). Figure from Salmela et al. (2016).

RF tuning was measured in early visual areas V1, V2 and V3, lateral areas LO1, LO2 and object-selective LOC and ventral areas V4, VO1 and VO2 in two experiments. The first experiment measured tuning to single RF patterns with frequencies between 2 and 20 to identify areas that showed selective tuning to low RFs. Homogenous spatial frequency profiles within stimuli were ensured by defining luminance profiles according to the distance of each pixel from the contour along the line with the shortest distance from the contour (reviewed in Chapter 2 and see stimuli in Figure 6.2A in the Methods section below). This had the effect of removing the variance between different RFs in terms of Fourier power spectra (demonstrated in Chapter 2, Figure 2.2B). However, there were still differences between stimuli in terms of the total contrast

energy within each stimulus as changes in RF are naturally accompanied by changes in contrast energy due to changes in contour path length. There were also differences between stimuli in terms of spatial frequency. Although the stimulus rendering method used ensured a homogeneous cross-sectional spatial frequency profile within each stimulus, changes in RF are naturally accompanied by changes in spatial frequency across the whole image due to changes in the total number of contour modulations. These low-level features were controlled for in a second experiment, which measured tuning to the same RFs as the first experiment, however a high RF component was added to all stimuli (see Figure 6.2B in the Methods section below). This had the effect of equating contour path length and therefore contrast energy between stimuli. It also applied the same high frequency contour modulations to all stimuli, reducing the variance in spatial frequency between stimuli. This second experiment measured RF tuning independent of sensitivity to contrast energy and spatial frequency. In addition, any tuning to RF from the first experiment that was preserved in the second experiment would indicate a global representation of RF that is not dependent on contrast energy and is invariant to the addition of high frequency contour detail that does not affect global shape.

6.2 Methods

Voxel-wise tuning to RF patterns in 7 was measured observers (4 males, 3 females, mean age = 28.1, age range = 24-38). All subjects had normal or corrected-to-normal visual acuity and gave informed consent. All participants completed one structural MRI session, one retinotopic mapping session, one LOC localiser session and two RF tuning sessions, totalling 5.25 hours of scanning per participant. The study was approved by the York Neuroimaging Centre Research Governance Committee and adhered to the declaration of Helsinki.

6.2.1 General MRI data acquisition

All MRI data were acquired on a GE 3-Tesla Signa HD Excite scanner using a 16-channel head coil to improve signal-to-noise in the occipital lobe, except for one structural scan which was acquired with an 8-channel head coil to improve alignments between functional and structural data. Structural and retinotopic mapping

measurements were acquired using methods identical to those described in Chapter 3 (also reviewed in Chapter 2).

6.2.2 LOC localiser

To identify object-selective LOC subjects completed three 8 minute fMRI scans (TR = 3000ms, TE = 30ms, voxel size = 2x2x2mm3, flip angle = 90°, matrix size 96x96x39, FOV = 19.2cm). Subjects viewed sixteen blocks of objects and sixteen blocks of scrambled objects in an ABAB block design experiment. Blocks were 15 seconds long with one image presented per second (images were on for 0.8 seconds with a 0.2 second interstimulus interval). Subjects fixated on a central red fixation cross throughout the scans and performed a one-back task where there could be zero, one or two repeats in a block. All stimuli were presented centrally on a mid-grey background (200cdm^{-2}).

The stimuli were 225 PNG images of easily recognisable objects which were manually extracted from their original backgrounds. Images were converted to greyscale and their image histograms were equalised. On average, stimuli subtended 4x4 degrees of visual angle. Scrambled stimuli were generated by splitting each image into 20 rows and 20 columns. All squares lying within the convex hull of the object were randomly permutated and rotated. This generated scrambled images which contained all the local details of the objects and the same coarse outline but were unrecognisable. After scrambling, a Gaussian smoothing filter (standard deviation = 1 pixel) was applied to smooth hard edges.

LOC localiser data were analysed with a general linear model (GLM) in FEAT (fMRI expert analysis tool) version 5.0 where objects and scrambled objects were the regressors of interest. The first three volumes from each scan were discarded and a high pass filter with a cutoff of 60 seconds was used to remove low frequency scanner drift from the data. Head movements were corrected for with MCFLIRT and motion estimates were included in the GLM as regressors of no interest. Spatial smoothing with a Gaussian kernel of full-width half-maximum 4 millimetres was applied to the data and FILM prewhitening was used. For each subject, data from the three fMRI runs were combined using a fixed-effects analysis with cluster correction ($Z = 5$, $p < .05$).

The results of an objects vs scrambled objects contrast were transformed into grey matter coordinates using mrVista and visualised on an inflated cortical surface. An

LOC region of interest (ROI) was then manually drawn around and constrained to significant voxels on the lateral surface of the occipital lobe.

6.2.3 RF pattern stimuli

To measure tuning to RF two sets of RF patterns were generated for use in two experiments. The first experiment used RF2, 3, 4, 5, 6, 7, 10 and 20 patterns (see Figure 6.2A). The radius r at polar angle θ was calculated using the following equation:

$$r(\theta) = r_0\big(1 + A\sin(\omega\theta + \emptyset)\big)$$

Where *r0* is the mean radius, A is the radial modulation amplitude, ω is the RF and $\emptyset$ is the phase of the RF pattern. All RF patterns had the same mean radius (2.5 degrees of visual angle) and amplitude (0.1), while RF and phase were variable.

The second experiment used composite RF patterns where stimuli from the first scan (except for the RF20 pattern, which was unchanged) were combined with an RF20 component to produce stimuli with similar levels of local contour modulations (see Figure 6.2B). For these stimuli, the radius r at polar angle θ for this stimulus was calculated using the following equation:

$$r(\theta) = r_0(1 + A_1\sin(\omega_1\theta + \emptyset_1) + A_2\sin(\omega_2\theta + \emptyset_2))$$

Where r_0 is the mean radius, A_1 and A_2 are the radial modulation amplitudes, ω_1 and ω_2 are the radial frequencies and $\emptyset_1$ and $\emptyset_2$ are the phases of the two RF components. For all stimuli the first RF component was identical to that used in the first experiment, and the second RF component had an RF of 20, with equal amplitude and phase to the first RF component. The cross-sectional luminance profile of stimulus contours was defined by the radial fourth derivate of a Gaussian, using the methods reviewed in Chapter 2 that ensure a homogeneous cross-sectional spatial frequency. All stimuli were presented at 50% contrast and the peak cross-sectional spatial frequency of stimulus contours was 2 cycles per degree.

6.2.4 Stimulus localiser

In order to identify which voxels in visual cortex responded to the stimuli, subjects completed a 290 second fMRI scan in which RF patterns were presented in an on-off block design (TR = 2000ms, TE = 30ms, voxel size = 2x2x2.5mm^3, flip angle = 90°, matrix

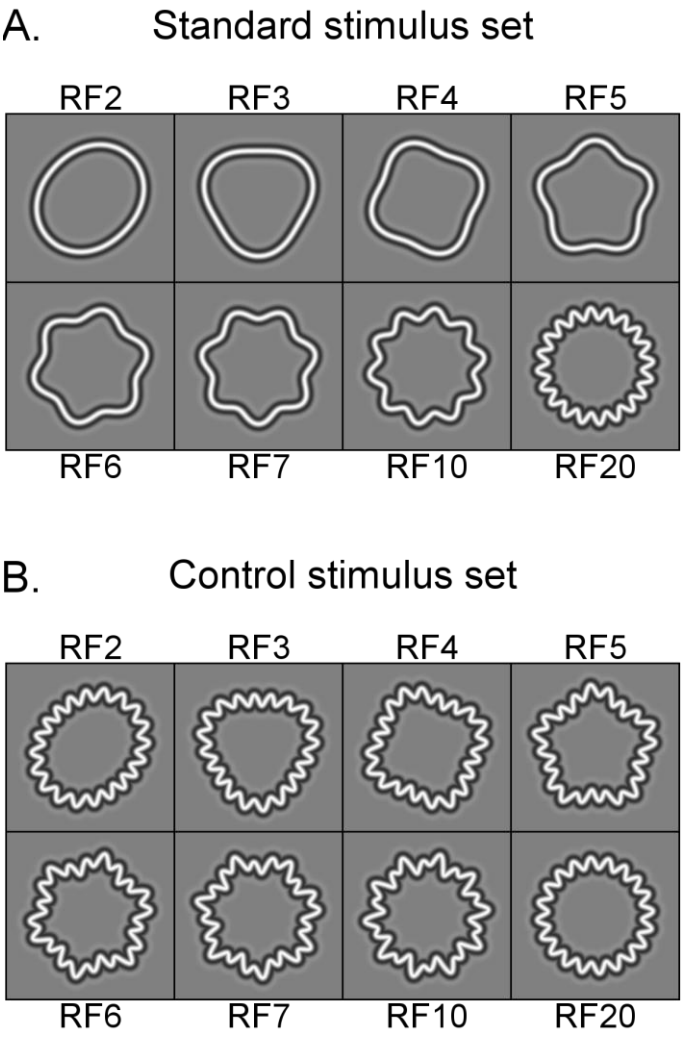

Figure 6.2: Stimuli used for the RF tuning experiments. (A) The first experiment used single RF patterns with frequencies between 2 and 20. (B) In the second experiment an RF20 component was added to all stimuli except the RF20 pattern so that all stimuli contained an RF20 component. This removed most of the variance between stimuli in terms of contrast energy and spatial frequency and introduced local, high frequency contour modulations that did not affect global shape for low RF patterns.

size 96 x 96 x 26, FOV = 19.2cm). Five second blocks of standard stimuli and five blocks of control stimuli were presented in a pseudo-random order. Stimulus blocks were 14 seconds long and separated by 14 second fixation periods. During stimulus blocks RF2, 3, 4, 5, 6, 7 and 10 patterns were presented twice each at a rate of 1Hz (on for 0.8 seconds with a 0.2 second interstimulus interval) in a random order with randomised phase. The first stimulus block was preceded by a 10 second fixation period to allow the scanner to reach stable magnetisation.

Stimuli were presented centrally on a mid-grey background (200cdm^{-2}). Subjects maintained fixation on a black central fixation cross and performed a contrast reversal detection task where they indicated when a stimulus was presented with negative contrast (-50%). 10% of stimuli were pseudo-randomly selected for contrast reversal (one stimulus out of every $\frac{N\ stimuli}{10}$ was randomly selected for contrast reversal).

Data from the stimulus localiser were analysed with a GLM in FEAT version 5.0 where standard stimuli and control stimuli were the regressors of interest. The first five volumes were discarded and a high pass filter with a cut off of 56 seconds was used to remove low frequency drift. Head movements were corrected for with MCFLIRT and motion estimates were included in the GLM as regressors of no interest. FILM prewhitening was used and data were cluster corrected ($Z = 2.3$ $p < .05$).

The results of an all stimuli vs baseline contrast were transformed into grey matter coordinates using mrVista. Retinotopic and LOC ROIs were then constrained to voxels that showed a significant response to RF patterns in the stimulus localiser.

6.2.5 RF tuning experiment: stimulus presentation

All subjects completed two fMRI scan sessions which measured tuning to RF patterns (the imaging parameters were identical to those used for the stimulus localiser). The first session comprised eight 294 second runs. The second session also comprised eight 294 second runs, plus one run of the stimulus localiser. RF patterns were presented in 6 second blocks. During a block the same RF was presented six times (0.8 seconds on with a 0.2 second interstimulus interval) at different orientations (the 6 RF patterns had phases of 0°, 60°, 120°, 180°, 240° and 300° and were presented in a random order). A sweeping paradigm was used where RF blocks were presented in ascending order (i.e. RF2, 3, 4, 5, 6, 7, 10), followed by a long baseline period consisting of four 6 second blocks of RF20 patterns (see Figure 6.3). Next RF blocks were presented in descending order from RF10 to RF2, followed by another baseline period. This stimulus presentation was repeated a second time within a single fMRI run. Stimulus presentation was preceded by a 10 second fixation period to allow the scanner to reach stable magnetisation, and followed by a 20 second fixation period to capture the full profile of the BOLD response's return to baseline. Participants maintained fixation on a central black fixation cross and performed the same contrast reversal task as for the

stimulus localiser. On average participants performed 97% correct (SD = 2), showing attention was maintained throughout the experiment.

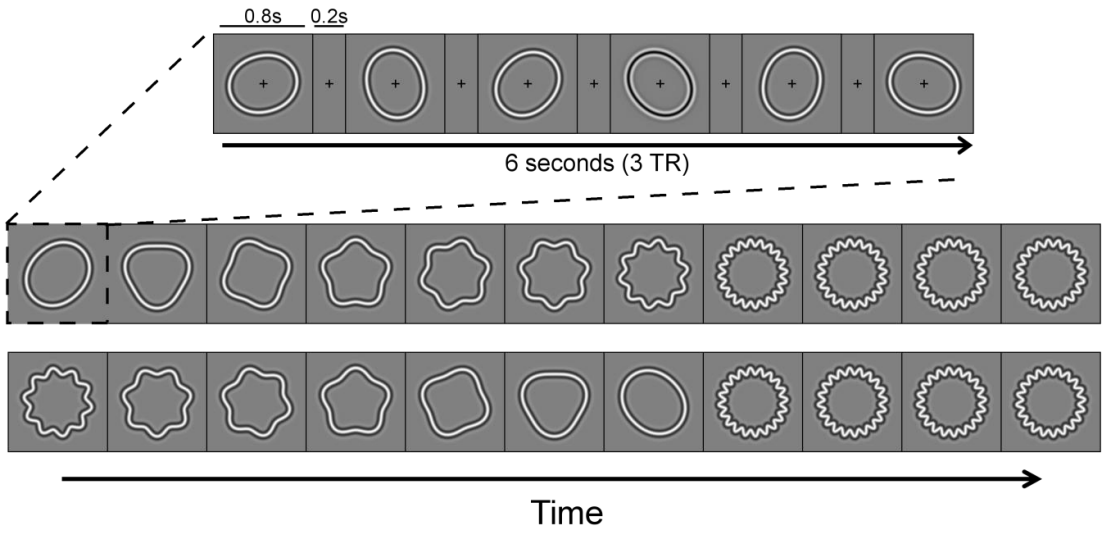

Figure 6.3: Order of stimulus presentation for the RF tuning experiment. A sweeping paradigm was used where RF blocks were presented in ascending order, followed by four baseline (RF20) blocks, then in descending order followed by four baseline blocks. This stimulus presentation was repeated a second time within a single fMRI run. During each stimulus block, the same RF pattern was presented at different orientations. One example of a stimulus with reverse contrast is shown (top, fourth from the left), which subjects identified with a button press.

6.2.6 Modelling RF tuning

The same methods were used to model tuning to RF for both RF tuning expeirments, but the two data sets were modelled separately. Within and between scan head movements were corrected for and data from the eight runs within a scan session were averaged using mrVista (data from two runs of the control experiment were lost for one participant due to a scanner fault). The rest of the modelling was conducted using custom code written in Matlab, and mrVista was used for visualising data on inflated cortical surfaces.

First, linear trends were removed from the average time series to remove low frequency drift from the data. A modelling technique was then used to define voxel-wise tuning to RF which was similar to the methods of Harvey et al. (2013). A voxel's tuning to RF was modelled using a Gaussian neural model defined in RF space (see Figure 6.4). This Gaussian was used to generate a predicted timecourse where the predicted response at each point in time was equal to the amplitude of the Gaussian for

the RF which was presented at that time in the fMRI experiment. This prediction was then convolved with a hemodynamic response function (SPM 8 default HRF for a TR of 2 seconds).

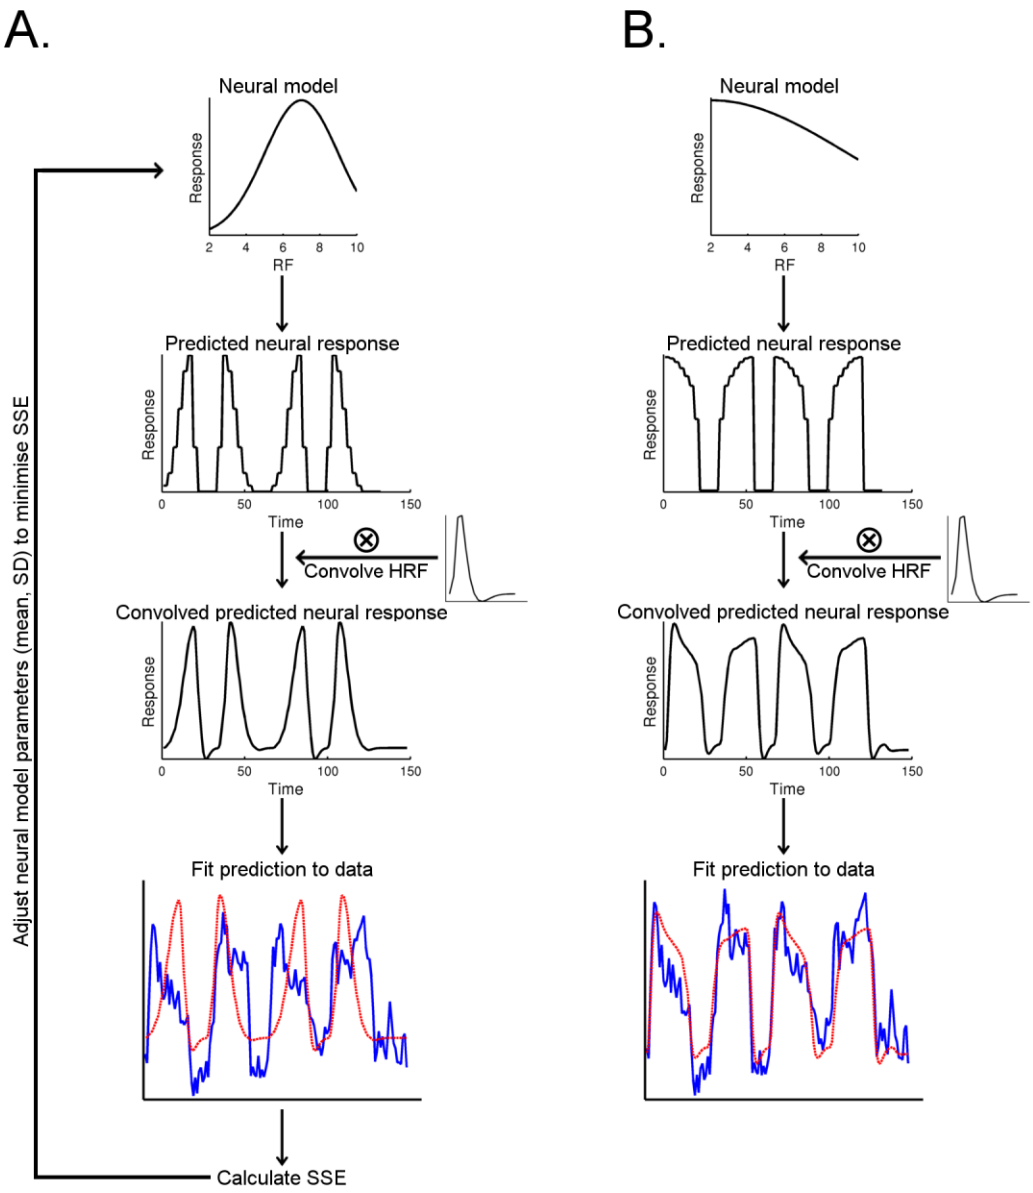

Figure 6.4: Schematic representation of the RF tuning modelling process. (A) A Gaussian neural model was used to generate a predicted neural response where the response amplitude at each point in time was equal to the amplitude of the neural model for the RF presented at that time in the fMRI experiment. This prediction was convolved with an HRF and fit to the average time series for a voxel to calculate the sum of square errors (SSE) between the data and the prediction. Nonlinear optimisation was used to optimise the mean and standard deviation (SD) of the neural model until the SSE was minimised. (B) The result of the modelling process is shown for the example data in panel A.

The sum of squared errors between the prediction and the average time series was calculated for each voxel. The mean and standard deviation of the neural model were then optimised using nonlinear optimisation until the sum of squared errors was minimised. This algorithm first optimised the Gaussian mean between limits of 2 and 20, using a fixed standard deviation of 25. Next the standard deviation was optimised between limits of 1 and 50, using the best fitting mean identified by the previous step. The Gaussian mean and standard deviation which produced the best fitting prediction were then taken as a voxels preferred RF and tuning width, respectively.

6.3 Results

6.3.1 Identification of regions of interest

Tuning to RF was measured in nine areas of visual cortex. These ROIs were identified using retinotopic mapping, LOC localiser and stimulus localiser experiments (see Figure 6.5 for an example data set from each). First, three early ROIs, V1, V2 and V3 were identified with retinotopic mapping. Second, three LO ROIs, LO1 and LO2 were defined retinotopically, and LOC was defined by an LOC localiser. Third, three VO ROIs V4, VO1 and VO2 were identified, all defined retinotopically. Finally, all ROIs were constrained to only include voxels which responded significantly to RF patterns in the stimulus localiser scan to ensure that all analyses only considered voxels within the stimulus representation. All nine ROIs were identified in the left and right hemispheres of all subjects, and further analyses were collapsed across hemispheres. For one subject no voxels in left VO2 showed significant responses to the stimulus localiser, so only voxels from right VO2 were used for this subject.

6.3.2 Tuning to RF (standard stimuli)

To measure voxel-wise tuning to RF, each voxel's preferred RF and tuning width was modelled with a Gaussian tuning to RF, with an aim of identifying regions of selective tuning to low RFs, which are known to be processed globally (Hess et al., 1999; Jeffrey et al., 2002). First, it was ensured that predictions generated from the RF tuning modelling process could explain a reasonable amount of variance in a voxel's response to RF patterns. Figure 6.6A shows two example voxels for which the model explained most of the variance, one with a preferred RF of 2 and one which preferred RF20,

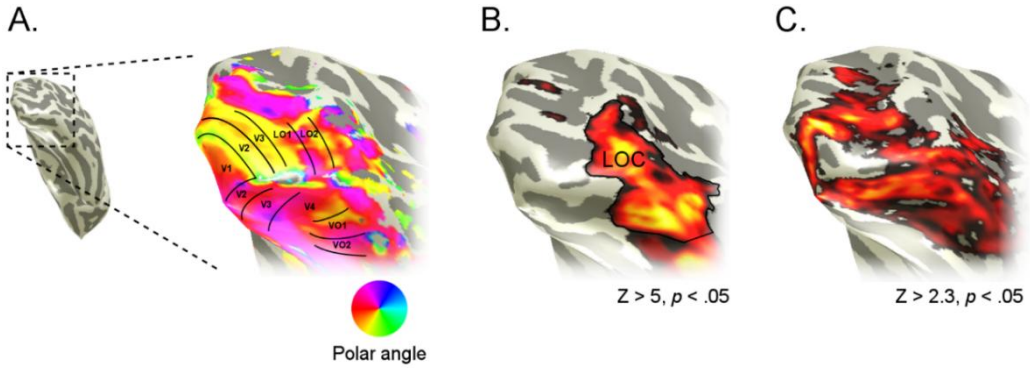

Figure 6.5: Example data sets from the retinotopic mapping, LOC localiser and stimulus localiser experiments for a single subject. (A) Retinotopic mapping was used to identify ROIs V1, V2, V3, LO1, LO2, V4, VO1 and VO2. (B) An objects vs scrambled objects localiser was used to identify LOC as a cluster of object-selective voxels in lateral occipital cortex. (C) A stimulus localiser experiment was used to identify voxels in visual cortex that responded to the stimuli. All ROIs were constrained to only include these voxels to ensure that all analyses only considered voxels within the stimulus representation.

demonstrating that the modelling process can explain most of the variance in at least some voxels and that some voxels tuned to both ends of the RF distribution that was presented.

Figure 6.6B shows all voxels for all subjects in which the model explained at least 20% of the variance. Using this threshold, it is evident that modelling a Gaussian sensitivity to RF predicted a large amount of the variance in neural responses to RF patterns for a large number of voxels. The average variance explained across subjects for all voxels within the stimulus representation (identified with the stimulus localiser) was 30.4% (SD = 0.05). Moreover, on average neural models explained more than 20% of the variance in 67.3% (SD = 9.63) of voxels within the stimulus representation.

From Figure 6.6B it can be seen that most voxels preferred high RFs (reds). However, in most subjects there is a clear cluster of voxels in left and right LO that shows tuning to low RFs (blues). To quantify this pattern of data, mean preferred RF (averaged across all voxels within the stimulus representation) was measured for each ROI across subjects (Figure 6.7A). A 1x9 repeated measures ANOVA revealed a significant main effect of ROI ($F(8, 48) = 25.55$, $p < .001$), showing that mean preferred RF varied across ROIs. Selectivity to RF20 patterns was predicted in V1 because it is the stimulus with the most contrast energy (see Figure 6.1), while regions which process global shape should be selective for lower RFs. Planned simple contrasts were therefore used to

compare the mean preferred RF in V1 to all other ROIs to identify where representations of RF shifted away from those in V1. The mean preferred RF was significantly lower in LO1 (p = .041), LO2 (p = <.001) and LOC (p < .001) compared to V1 with no other significant differences (VO2 approached significance, p = .052). This indicates a greater sensitivity to lower RFs was expressed LO, and first emerged in LO1.

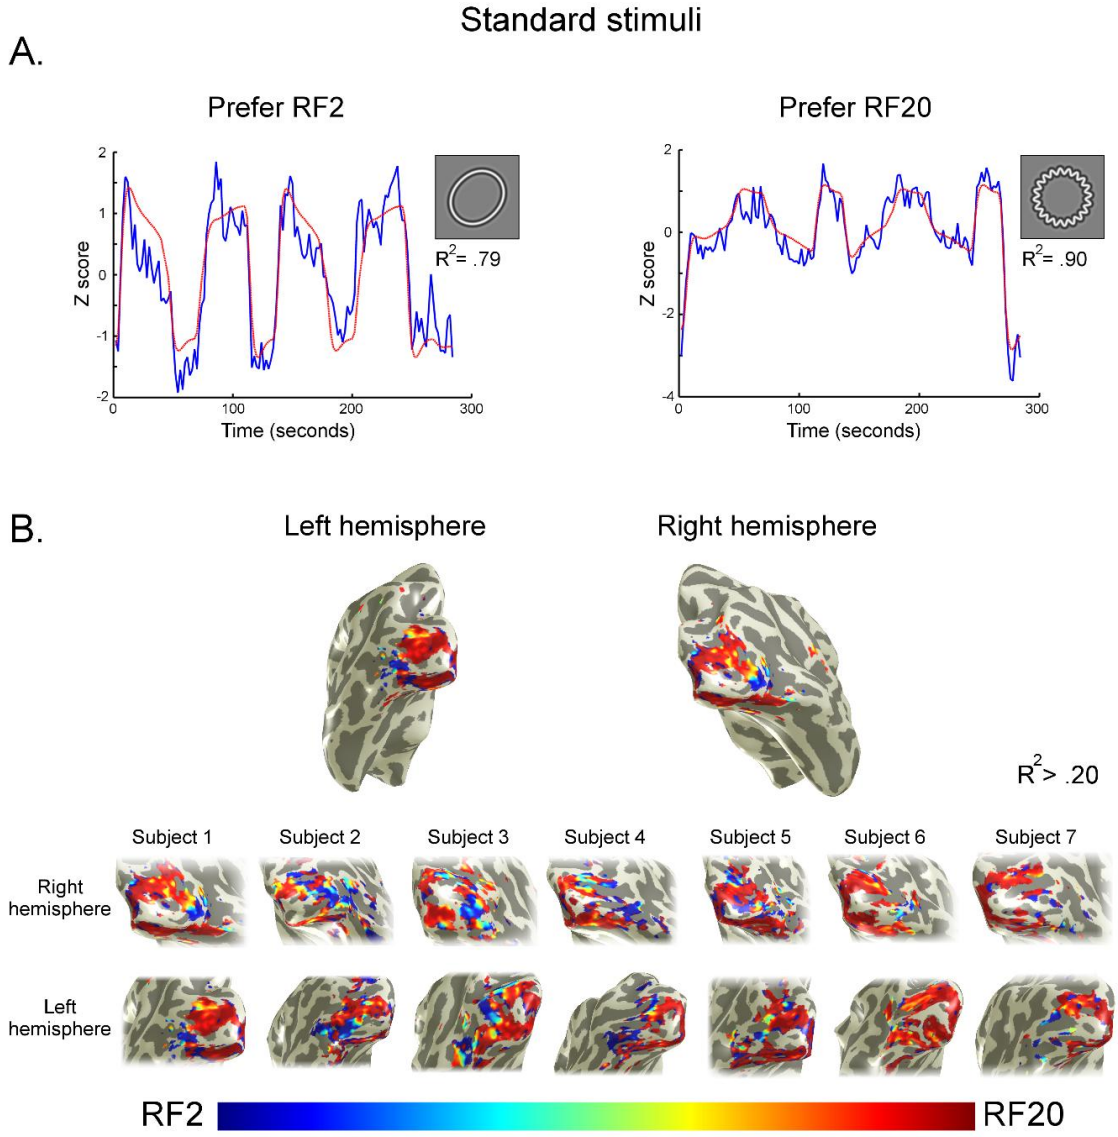

Figure 6.6: Results of modelling voxel-wise tuning to RF for standard stimuli. (A) Model fits are shown for two example voxels with preferred RFs of 2 (left) and 20 (right). The average time series (averaged across 8 runs) from the fMRI experiment is in blue and the prediction generated by a modelled Gaussian sensitivity to RF is in red. (B) Preferred RF is shown for all subjects and all voxels in which the modelled RF tuning explained at least 20% of the variance in their response to RF patterns. Low RFs are shown in blues, mid RFs are shown in cyans and yellows, high RFs are shown in reds.

Although the mean preferred RF was significantly lower in LO ROIs than V1, it was still higher than would be expected for a selective response to low RFs (approximately 10 in LO2 and LOC, see Figure 6.7A). However, Figure 6.6B shows a cluster of voxels in LO which show selectivity for low RFs, suggesting some low RF tuning that was not well captured by taking the mean across all voxels within an ROI. Voxel-wise tuning to RF was therefore examined in each ROI to determine how many voxels in each were selective for low RFs. Figure 6.7B plots the proportion of voxels within each ROI which showed a preference for each RF between RF2 and RF20, averaged across subjects. All early and VO ROIs showed a clear bias towards high RFs, however LO ROIs showed a different pattern, with a larger proportion of voxels preferring RFs within the range that are globally processed (< RF8, Jeffrey et al., 2002). This indicates that responses in early and VO ROIs were dominated by the RF20 pattern, the stimulus with the most contrast energy, while LO ROIs, particularly LO2 and LOC, showed a different pattern of response which appears to be driven by the global processing of low RFs.

Additionally, the plots in Figure 6.7B demonstrate that the ROI mean preferred RFs shown in Figure 6.7A were heavily influenced by the larger proportions of voxels that preferred RFs at the edges of the presented RF distribution (RF2 and RF20). The values plotted in Figure 6.7A, therefore, should only be interpreted as showing an overall difference in RF tuning between V1 and LO ROIs, but are not an accurate indication of the RF tuning across the population of voxels within those ROIs, which can be better appreciated in Figure 6.7B.

Finally, it was determined how voxel tuning width varied with preferred RF. Changes in RF have a greater effect on the shape of an RF pattern for low RFs compared to high RFs (see Figure 6.1). One would therefore predict that voxels tuned to low RFs should have narrower tuning widths than those tuned to high RFs. Moreover, in neural tuning to numerosity increases in preferred numerosity are coupled with a linear increase in tuning width (Harvey et al., 2013) and, as in Weber's law, this is likely a general principle of neural tuning to stimulus dimensions. According to Weber's law, the difference between a primary and secondary stimulus required for them to be perceived as different is a constant ratio of the primary stimulus (Stevens, 1957). As such, as the magnitude of a stimulus feature increases we become less sensitive to it, meaning the contrast between two stimuli needs to be greater in order to be a

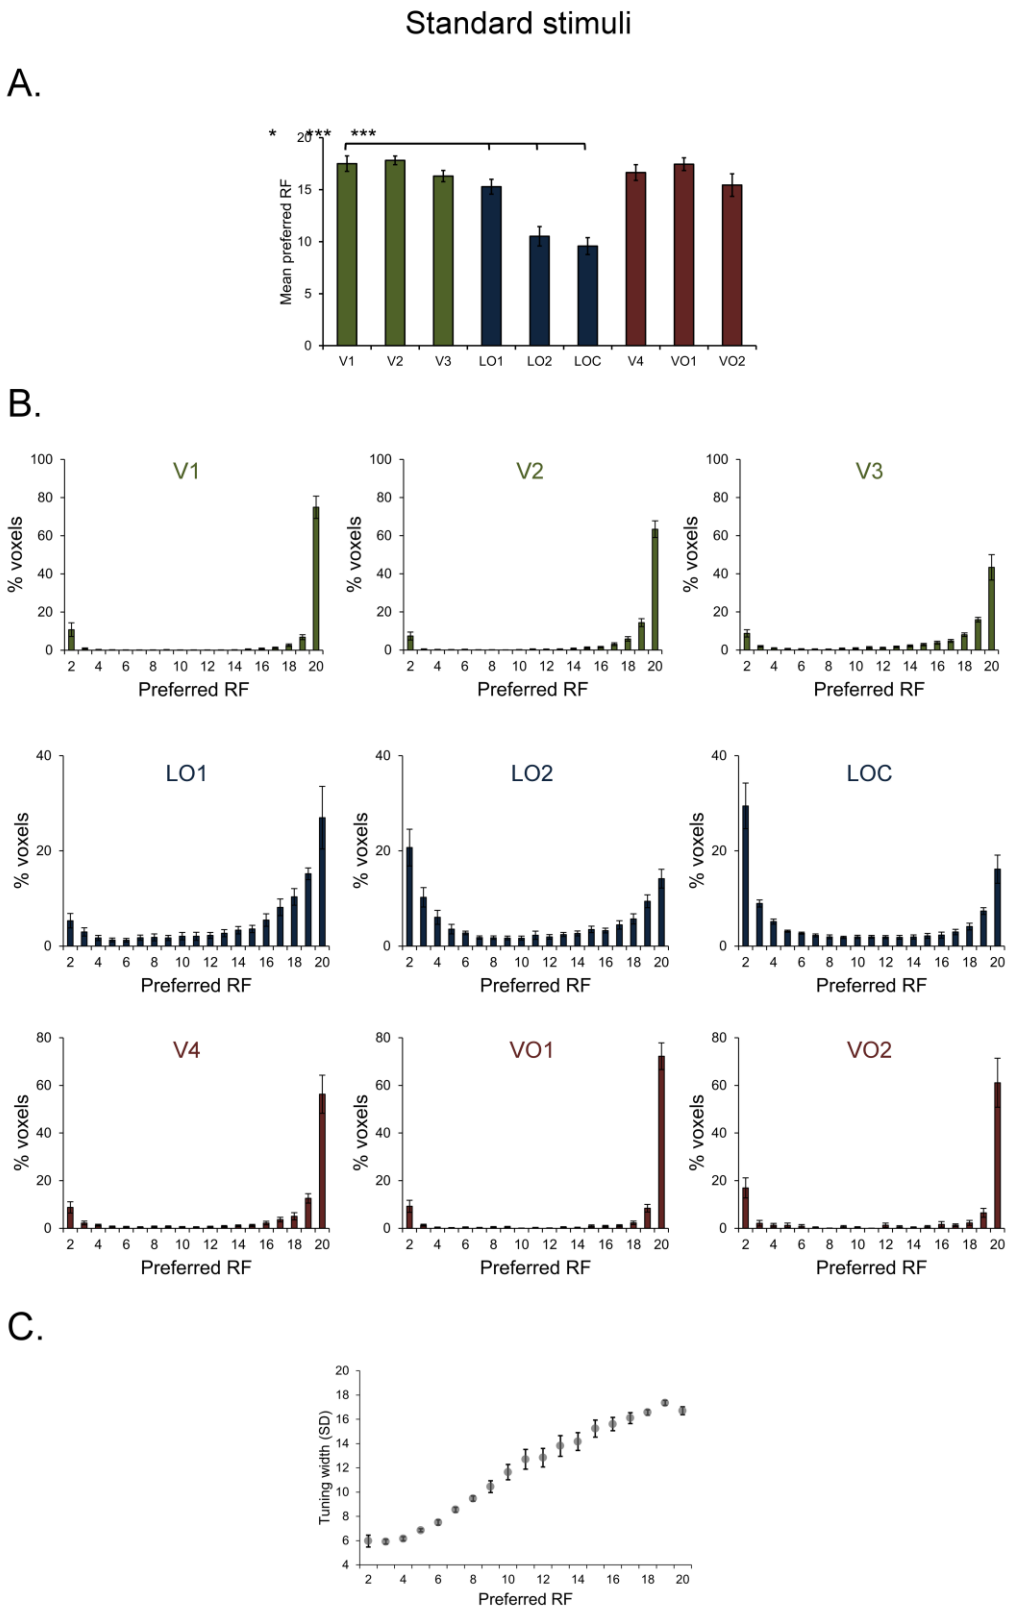

Figure 6.7: Summary statistics of voxel-wise tuning to RF in standard stimuli. (A) Average preferred RF for early (green), LO (blue) and VO ROIs (red), collapsed across voxels and averaged across subjects. (B) The proportion of voxels with a preference for each RF between RF2 and RF20 for early (green), lateral (blue) and ventral ROIs (red), averaged across subjects. (C) Voxel tuning width plotted against preferred RF, collapsed across all voxels from all ROIs and averaged across subjects.

perceivable difference. This may be reflected neurally by increasing tuning widths with increasing magnitude of an encoded stimulus feature. One would therefore expect to see larger tuning widths in voxels with higher preferred RFs if measurements reflected true tuning to RF. Collapsed across all voxels from all ROIs and averaged across subjects, preferred RF correlated significantly with tuning width (Pearson's $r = .989$, $p < .001$), and it can be seen from Figure 6.6C that tuning width increases monotonically with preferred RF. Neural tuning to RF, therefore, obeys a similar law to numerosity tuning, where preferred RF and tuning width are strongly correlated, increasing the credibility of RF as a genuine representational dimension.

6.3.3 Tuning to RF (control stimuli)

Using single RF patterns, it was identified that tuning to low RFs was localised to LO, particularly LO2 and LOC. However, in the standard stimulus set increases in RF were also associated with increases in contrast energy and spatial frequency. Tuning to low RFs could therefore be influenced by a greater sensitivity to stimuli with less contrast energy or lower spatial frequency rather than true global shape processing of low RF patterns. That is, the potential for responses in extrastriate areas to contrast energy and spatial frequency to be more compressive could influence measures of preferred RF. Therefore a second experiment was conducted using a control stimulus set where the control RF (RF20) was present in all stimuli (Figure 6.2B). This had the effect of removing most of the variance in contrast energy across stimuli (Table 6.1) and introduced local, high frequency contour noise which had little effect on the global shape of low RFs. In order to demonstrate global shape processing in LO2 and LOC, tuning to low RFs in these regions should survive this stimulus manipulation.

Again, a large portion of the variance in the data was explained for at least some voxels (Figure 6.8A). The average variance explained across subjects for all voxels within the stimulus representation was 26.9% (SD = 0.06). On average, neural models explained more than 20% of the variance in 60.1% (SD = 12.2) of voxels within the stimulus representation. Models consistently explained less variance in the control experiment compared to the first experiment, resulting in a significantly lower average variance explained ($t(6) = 2.49$, $p = .047$, paired-samples t test) in the control experiment. There was also a lower proportion of voxels for which more than 20% of the variance was

Table 6.1: The perimeter (pixels, degrees of visual angle) of each RF stimulus for standard and control stimulus sets. The standard deviation (bottom) across control stimuli was reduced by 86% compared to standard stimuli. As stimulus perimeter was directly proportional to contrast energy, most of the variance in contrast energy between stimuli was therefore removed in control stimuli.

RF	Standard stimulus perimeter (pixels, degrees of visual angle)	Control stimulus perimeter (pixels, degrees of visual angle)
2	624, 16.8	1030, 27.7
3	631, 17	1038, 28
4	650, 17.5	1044, 28.1
5	656, 17.7	1045, 28.1
6	667, 18	1051, 28.3
7	687, 18.5	1058, 28.5
10	757, 20.4	1086, 29.2
20	1025, 27.6	1025, 27.6
Standard deviation	133, 3.6	19, 0.5

explained in the control experiment, which approached significance ($t(6) = 2.28$, $p = .063$). Although the drop in modelling performance in the control experiment was significant, it was fairly small in magnitude (30.4% to 26.9% for average variance explained and 67.3% to 60.1% for proportion of voxels with more than 20% variance explained). It therefore appears that, despite a drop in modelling performance, neural models still explained a reasonable amount of variance in responses to RF patterns in the control experiment.

Figure 6.8B shows that, in general, there were very few voxels tuned to high RFs (reds) for control stimuli. This should be the case, because RF20 was now present in all stimuli and throughout the entire experiment. In terms of RF20, therefore, there was zero contrast between stimuli in the stimulus regime. As such, there was only contrast within the RF spectrum between RF2 and RF10, so tuning to RF20 would indicate an issue with the modelling procedure. Instead, these voxels showed tuning to middle RFs (yellows and cyans), which was likely because the RF10 stimulus (RF10+20 pattern) had the longest perimeter (Table 6.1) and therefore the most contrast energy out of the control stimuli. This suggests that responses to high RFs in the first experiment were driven by stimulus contrast energy, as the tuning of these voxels

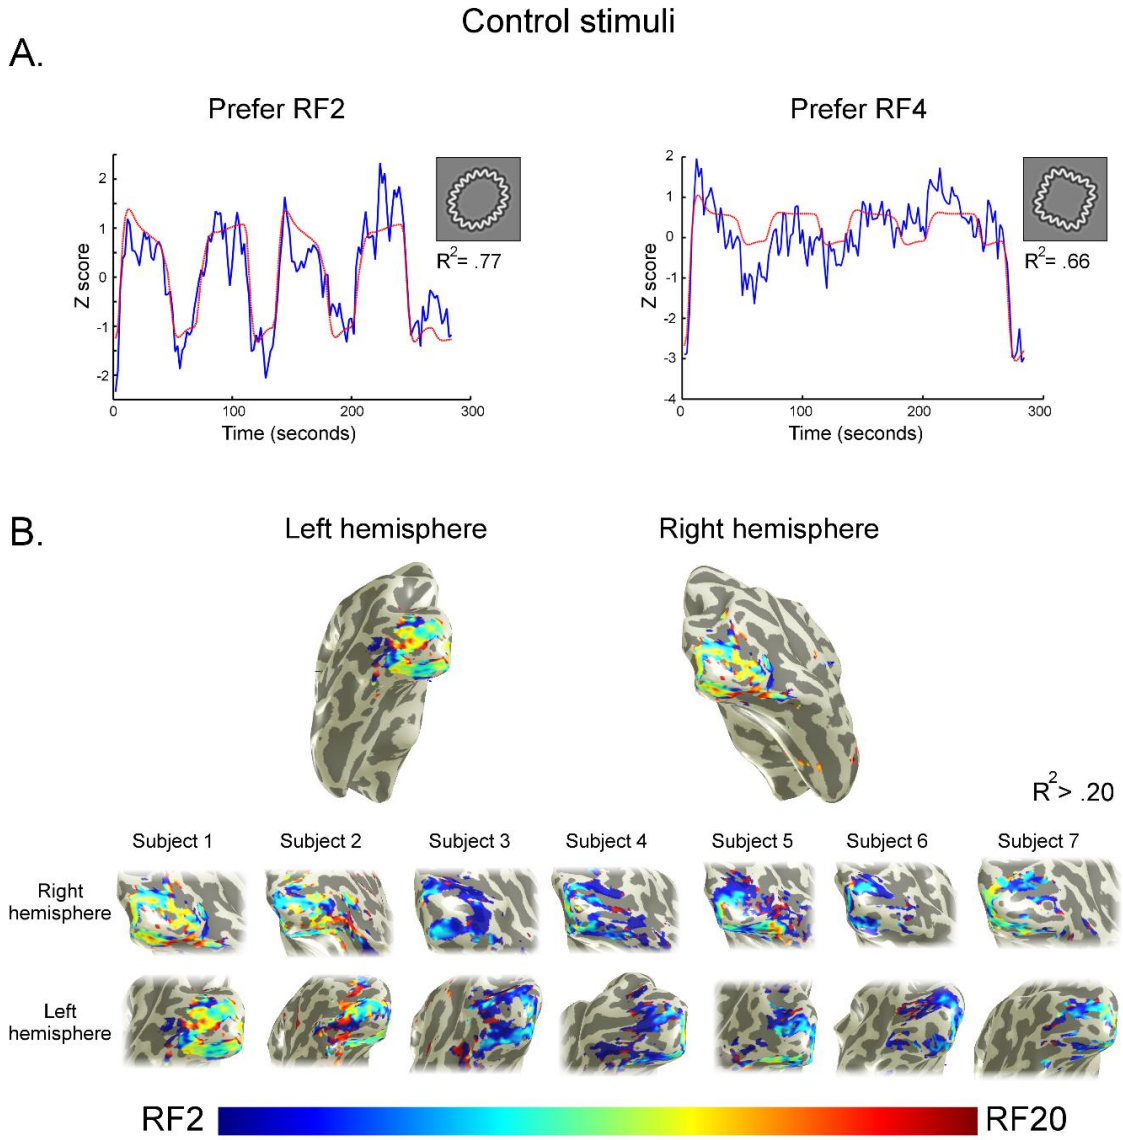

Figure 6.8: Results of modelling voxel-wise tuning to RF for control stimuli. (A) Model fits are shown for two example voxels with preferred RFs of 2 (left) and 4 (right). The average time series (averaged across 8 runs) from the fMRI experiment is in blue and the prediction generated by a modelled Gaussian sensitivity to RF is in red. (B) Preferred RF is shown for all subjects and all voxels in which the modelled RF tuning explained at least 20% of the variance in their response to RF patterns. Low RFs are shown in blues, mid RFs are shown in cyans and yellows, high RFs are shown in reds.

shifted towards RF10 in the control experiment. Importantly, tuning to low RFs (blues) remains present for control stimuli and again appears mostly localised to LO.

These data were quantified using an ROI analysis of mean preferred RF (Figure 6.9A). A 1x9 repeated measures ANOVA showed that mean preferred RF significantly varied across ROIs ($F(8, 48) = 3.98$, $p = .001$). As in the first experiment, preferred RF was lower in lateral ROIs than early and ventral ROIs, however preferred RF was only

significantly lower than V1 in LO2 (p = .041, LO1 and LOC approached significance, LO1: p = .056, LOC: p = .060). Mean preferred RF was also significantly lower in V3 compared to V1 (p = .044), suggesting that tuning to lower RFs emerged earlier in visual cortex for the control stimuli compared to standard stimuli.

In early ROIs, voxel-wise preference for RF was more evenly distributed around early to mid RFs (Figure 6.9B), likely because most of the variance in contrast energy between stimuli was removed (Table 6.1). In LO ROIs, RF tuning was once again biased towards lower RFs, and this emerged more prominently in LO1 compared to the first experiment. The tuning to low RFs in lateral ROIs was therefore replicated, indicating that the processing of these stimuli is sufficiently global to survive the introduction of local, high frequency contour noise. VO ROIs also showed some tuning to low RFs, which was not present for standard stimuli. This suggests that VO may perform some shape processing, however this only became evident once low-level stimulus properties such as contrast energy and spatial frequency were controlled for.

Preferred RF was once again positively correlated with tuning width (Figure 6.9C), and this correlation approached significance (Pearson's r = .421, p = .073). However, the relationship between preferred RF and tuning width was only present between RF2 and RF10 – i.e. the range in which there was contrast between stimuli in terms of RF. Since tuning RFs above 10 was not expected, the breakdown of the relationship between preferred RF and tuning width after RF10 adds credibility to the modelling procedure. When considering only data points between RF2 and RF10 the correlation between preferred RF and tuning width was highly significant (Pearson's r = .975, p < .001). Overall, therefore, the relationship between preferred RF and tuning width was replicated in the control experiment.

6.3.4 Contrast and shape predictors of neural responses to RF patterns

There was a clear dissociation in responses to RF patterns between LO ROIs and other ROIs. It was predicted that responses in early visual cortex would be driven by low-level stimulus features such as contrast energy and spatial frequency, while responses in global shape processing areas should be driven by the shape of lower RF patterns. In order to determine which stimulus features predominately drove ROI responses, two predictors for neural responses to RF patterns were generated based on different

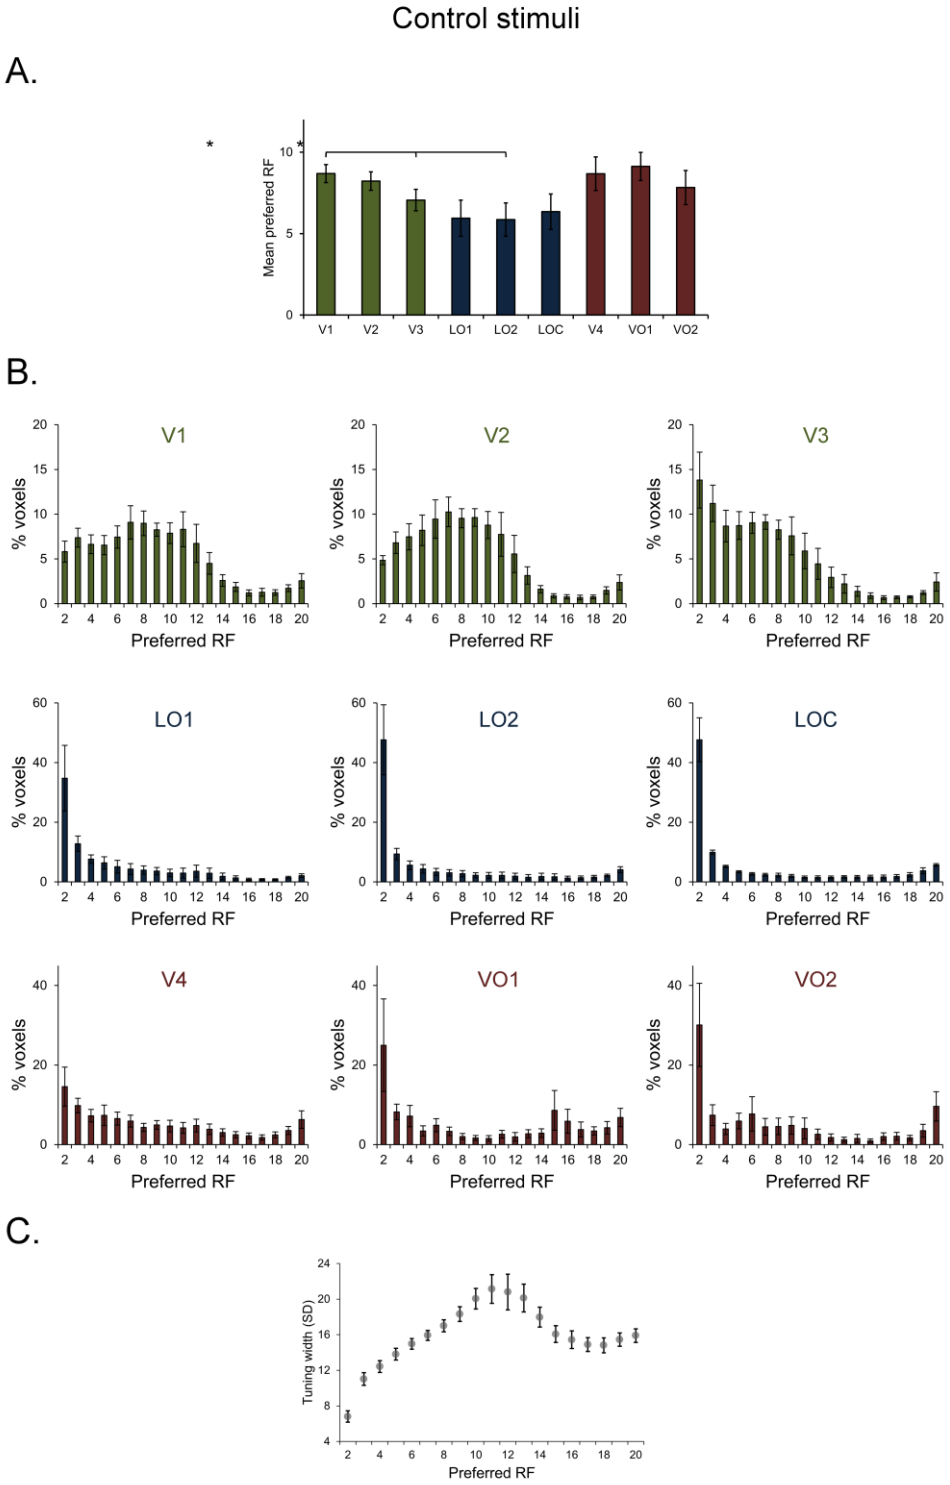

Figure 6.9: Summary statistics of voxel-wise tuning to RF in control stimuli. (A) Average preferred RF for early (green), LO (blue) and VO ROIs (red), collapsed across voxels and averaged across subjects. (B) The proportion of voxels with a preference for each RF between RF2 and RF20 for early (green), lateral (blue) and ventral ROIs (red), averaged across subjects. (C) Voxel tuning width plotted against preferred RF, collapsed across all voxels from all ROIs and averaged across subjects.

stimulus features. The first was a low-level predictor that measured the perimeter of a stimulus, which was directly proportional to the total amount of contrast energy in the

stimulus where a longer perimeter was coupled with higher contrast energy. The second predictor was designed to identify which brain regions were performing global shape analysis of RF patterns. Shape processing research has converged on the notion that shapes are represented in multidimensional shape space (Kayaert et al., 2005; Op De Beeck, Wagemans, & Vogels, 2001; Pasupathy & Connor, 2001, 2002; Vernon et al., 2016), centred on a prototypical mean shape (Anderson et al., 2007; Kayaert et al., 2005). RF patterns, in particular, appear to be represented by their deviation from a circular template (Anderson et al., 2007; Loffler, 2008; Poirier & Wilson, 2006). For the shape predictor, therefore, deviation from circularity was measured in all stimuli (see Figure 6.10). This was done by measuring the degree of overlap between the convex hull of an RF pattern and the best fitting ellipse to that convex hull. The best fitting ellipse was used because it is a more flexible template than a pure circle, which allows for variable aspect ratios. The convex hull was used because it mostly ignored high frequency contour detail in the control stimuli (Figure 6.10B), which were expected from any brain regions performing global shape processing.

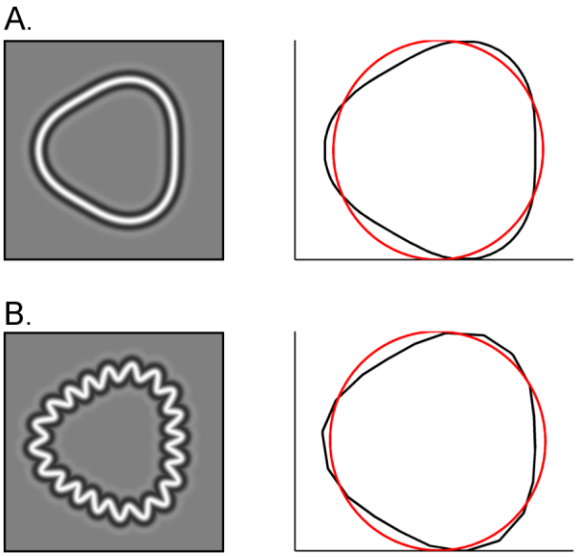

Figure 6.10: Schematic showing how the shape predictor was measured for each stimulus. For an RF pattern (left), the degree of overlap between the convex hull (black line on the right) and the best fitting ellipse to that convex hull (red line on the right) was measured. Panel A shows an RF3 pattern from the standard stimulus set, and panel B shows the same from the control stimulus set. It can be seen that the convex hull mostly ignores the high frequency contour details in control stimuli.

To measure whether ROI responses were driven by the contrast energy or shape of RF patterns, contrast energy and shape predictors were computed for each of the RF

patterns presented (RF2, 3, 4, 5, 6, 7, 10 and 20). These were then correlated against the proportion of voxels tuned to those RFs for each subject (Figure 6.11). Responses to RFs not presented in the experiment were not included for this analysis. Partial correlations were used to ensure that the two predictors accounted for separate portions of the variance in neural responses. A Fisher Z transformation was applied to subjects' correlations so that group level statistical analyses could be performed.

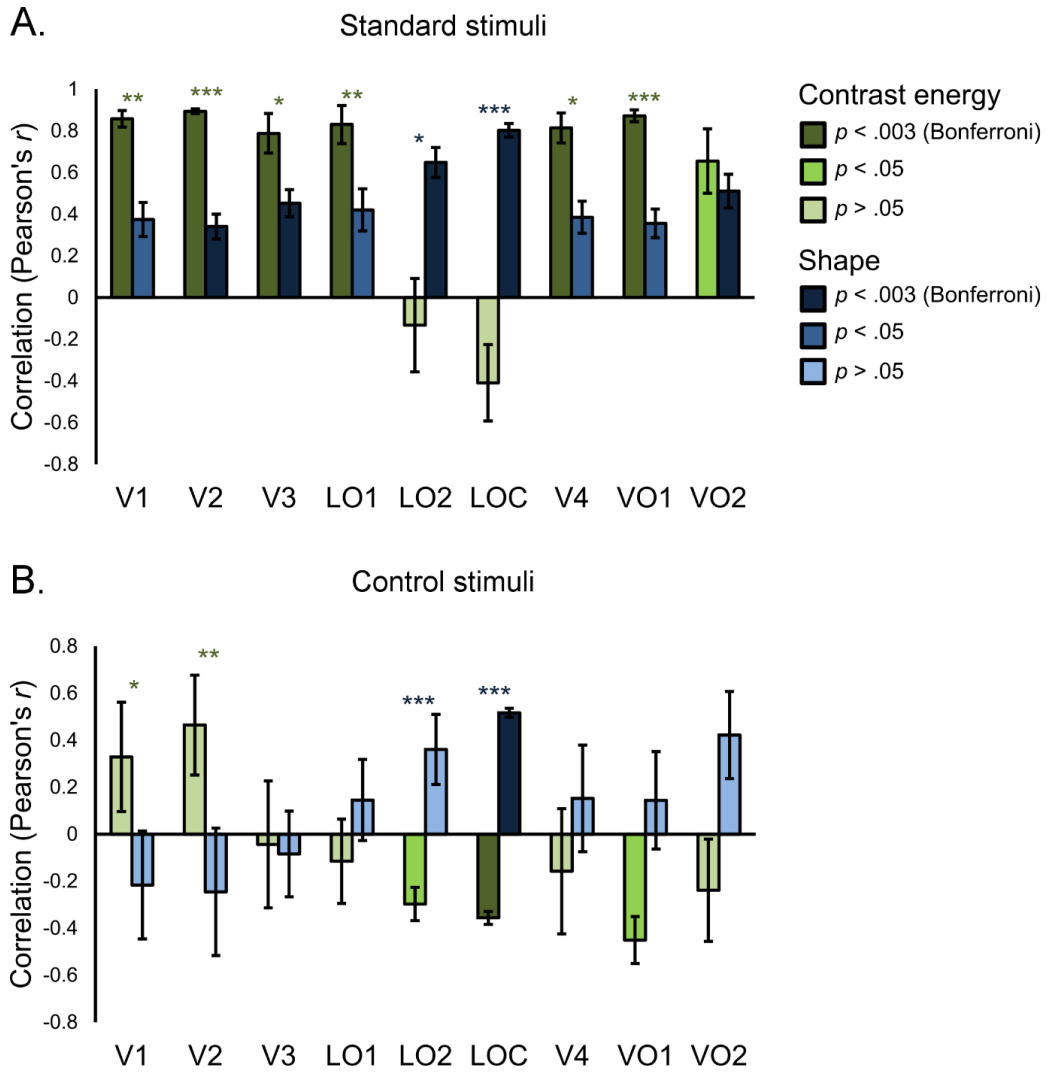

Figure 6.11: Correlations between contrast energy (green bars) and shape (blue bars) predictors for each RF and the proportion of voxels tuned to those RFs for each ROI, averaged across subjects for (A) standard stimuli and (B) control stimuli. The shading of bars indicates the strength of the correlation (one sample t tests), where light bars are not significant, darker bars are significant at the $p < .05$ level, and the darkest bars survive a Bonferroni correction for multiple comparisons ($p < .003$). Asterisks indicate a significant difference between contrast and shape correlations (paired-samples t tests) for an ROI where * = $p < .05$, ** = $p < .01$ and ** = $p < .006$ (Bonferroni-corrected p value). All statistical tests were performed on Fisher Z transformed correlations. Green asterisks show a higher contrast correlation; blue asterisks show a higher shape correlation. Error bars show standard error of mean.

For standard stimuli, the contrast energy predictor correlated significantly with neural responses in all ROIs except LO2 and LOC (Figure 6.11A). While the shape predictor also correlated with these regions, the correlation for the contrast energy predictors was significantly higher than the shape predictor in all ROIs except LO2, LOC and VO2. Of these regions, correlations for the shape predictor were significantly higher than the contrast energy predictor in LO2 and LOC, while there was no difference between predictors in VO2. VO2 responses, therefore, may have been driven by both contrast energy and shape in equal measure. Importantly, only in LO2 and LOC were responses driven by shape, but not contrast energy. These regions, therefore, do not show sensitivity to stimulus contrast energy and appear to be involved in more global shape processing.

For control stimuli, correlations were unsurprisingly weaker as the majority of variance between stimuli in terms of low-level differences was removed (Figure 6.11B). The contrast energy predictor did not significantly positively correlate with any ROIs, though it did correlate negatively with LO2, LOC and VO1. The shape predictor only significantly correlated with LOC, suggesting that representations of global shape only survived the control stimulus manipulation in LOC but not LO2. As such, representations of shape may be more global in LOC than LO2. Importantly, only responses in LO2 and LOC were driven significantly more by stimulus shape than contrast energy (as was also the case for standard stimuli), while V1 and V2 were driven significantly more by contrast than shape. The results from the first experiment were therefore replicated, showing evidence for global shape processing only in LO2 and LOC. In addition, the results from the control experiment suggest that shape representations in LOC may be more global than in LO2.

6.4 Discussion

In this chapter voxel-wise tuning to RF was measured in early visual cortex, LO and VO. Across two experiments using standard RF patterns and RF patterns controlled for low-level differences. LO, particularly areas LO2 and LOC, showed a biased tuning towards low RFs. In these regions, most voxels were tuned to RFs within the range that are processed on a global scale (< 8, Jeffrey et al., 2002), suggesting they may perform global analysis of shape. Moreover, correlating tuning profiles of these areas to RF

against stimulus contrast energy and shape revealed that their responses were driven primarily by shape over contrast energy. As expected, early visual areas were tuned to stimuli with the most contrast energy, and their responses were not dependent on shape. Areas in VO generally responded to stimulus contrast energy when it varied over a large range in the stimuli. Controlling for low-level stimulus features revealed some tuning towards lower RFs in VO, however their responses did not correlate with stimulus shape as strongly as LO2 and LOC. Overall, only LO showed a consistent preference for stimuli that are processed globally, and exhibited responses that were driven by stimulus shape.

Our results are consistent with other research into the roles of LO2 and LOC. LO2 responded strongly to RF patterns in Chapter 3, and LO2 is causally important for radial amplitude discrimination (Silson et al., 2013) and responds to shape complexity (Vernon et al., 2016). Similarly, LOC responds to global shape (Kourtzi & Kanwisher, 2001) and responds to the number and salience of RF components within a stimulus (Drucker & Aguirre, 2009; Vernon et al., 2016) and shape similarity (Op de Beeck, Torfs, et al., 2008). Evidence is therefore converging on the notion that the integration of local, low-level information to generate more abstract feature spaces for shape representation (Güçlü & van Gerven, 2015; Wilson & Wilkinson, 2015) occurs in LO. In particular, there was evidence for a gradient in how global shapes are represented are through LO. LO1 showed some tuning to low RFs, but only when variance in low-level features between stimuli was removed. LO2 and LOC showed robust tuning to low RFs, but only LOC responses were well predicted by the shape predictor in the control experiment, suggesting its representations were more global than those in LO2. Shape representations, therefore, appear to become increasingly global from LO1 through to LOC. As such, abstract representations of objects in LOC (Malach et al., 1995) might be predicated on a preliminary analysis of shape performed in LO1 and LO2.

Our results are also consistent with a biologically plausible model of shape perception where RF patterns are represented by deviations from a circular template (Anderson et al., 2007; Loffler, 2008; Poirier & Wilson, 2006, 2010). Shape-responsive areas showed the strongest preference for RF2 patterns (ellipses) and their responses were well explained by the shape predictor which measured stimulus circularity. An ellipse template may be more flexible than a circle as it allows for variable aspect ratios, and

the results presented suggest that the ellipse may be a particularly important shape for shape processing regions. In addition, only LO2 and LOC responses were well explained by the shape predictor that measured stimulus circularity. As such, LO2 and LOC may be the neural locus of circularity-based representations of RF patterns that have been hypothesised in psychophysics literature.

Surprisingly, there was little evidence of shape processing in V4. Macaque V4 is tuned to shape and curvature (Cadieu et al., 2007; Pasupathy & Connor, 1999, 2001, 2002; Roe et al., 2012) and contains shape representations that are invariant to object size (El-Shamayleh & Pasupathy, 2016). Therefore one might have expected to find shape tuning that was invariant to low-level stimulus features in human V4, however this was only found in lateral areas. This is consistent with some other research that has identified size invariant representations of objects and shapes in LOC (Grill-Spector et al., 1999; Sawamura et al., 2005). In addition to this, shape representations in LO may also be invariant to contrast energy and spatial frequency, consistent with the more abstract, feature invariant shape processing mechanism identified psychophysically in Chapter 4. These results therefore suggest that the global shape processing mechanisms identified behaviourally in Chapter 4 likely reside in LO. Other research has also found clear responses to shape in LO, but mixed results from V4 (Vernon et al., 2016). Therefore, though human V4 is often assumed to be a functional homologue of macaque V4, lateral visual field maps LO1 or LO2 might be a better homologue for the shape processing attributes found in macaque V4.

Our results contradict one other study that measured representations of RF (Salmela et al., 2016), which may be due to differences in methodologies between the two studies. First, it was ensured that the cross-sectional spatial frequency of the stimuli was homogenous. This was important as research has shown that low-level stimulus properties can account for a large proportion of variance in neural responses to objects or scenes when they are not controlled for (Andrews et al., 2015; Coggan et al., 2016). Indeed, in Salmela et al.'s (2016) study neural responses to RF patterns were mostly dominated by stimulus spatial frequency. Additionally, a second experiment controlled for low-level differences between stimuli, which was important as LO1 and VO primarily responded to these features in the first experiment. In this study, LO2 and LOC showed tuning to RF across both experiments - including even when stimuli

varied in low-level features, which Salmela et al. (2016) did not find. This could be because neural responses to the inhomogeneous spatial frequency profiles of their stimuli masked any potential responses to RF.

Shape processing research has converged on the notion that shapes are represented in a multidimensional shape space (Kayaert et al., 2005; Op De Beeck, Wagemans, & Vogels, 2001; Pasupathy & Connor, 2001, 2002; Vernon et al., 2016). In the past, explicit attempts to measure neural tuning to curvature or shape have only been performed on monkeys (Pasupathy & Connor, 2001, 2002). This chapter presented the first attempt to explicitly model neural tuning to a possible dimension of shape space in humans using Gaussian neural modelling. This modelling technique, therefore, provides an opportunity by which the work measuring neural tuning to shape in monkey neurons can be extended to humans by measuring population tuning to shape in fMRI voxels.

Modelling Gaussian sensitivity profiles to stimulus dimensions may be a generally effective method for mapping feature spaces within the brain. This approach has been shown to be highly effective for measuring pRFs (Dumoulin & Wandell, 2008) and numerosity tuning (Harvey et al., 2013), and this study demonstrates that it can be used to measure tuning to object shape. This study mapped one possible dimension of shape space (RF), however evidence suggests that shapes are represented along multiple dimensions (Kayaert et al., 2005; Op De Beeck, Wagemans, & Vogels, 2001; Pasupathy & Connor, 2001, 2002; Vernon et al., 2016). Future research could vary both radial amplitude and RF using a two-dimensional Gaussian neural model. Further, this could be expanded upon to measure tuning to additional dimensions such as number of RF components using multidimensional models. In so doing, one could map out a representational shape space within the human brain.

It is worth noting that, while these methods may allow for a deep understanding of the neural representation of RF patterns, they may not generalise to other types of shapes. The current study was successful in modelling responses to RF patterns with a Gaussian model in RF space, however RF tuning could have been underpinned by tuning to some other aspect of curvature or shape that is correlated with RF. For example, Dickinson et al. (2015) identified that sensitivity to the polar angle subtended by adjacent points of maximum curvature (which is correlated with RF) was a better

model for the perception of rectified RF patterns (where the radius was modulated with a rectified sinusoid), which could not be accounted for by sensitivity to RF. As such, tuning to RF itself may only be applicable to RF patterns. Whilst RF patterns are a powerful stimulus for investigating shape perception, future research should consider tuning to other shape or curvature features that can account for responses to both RF patterns and other shapes.

6.5 Conclusions

In conclusion, the results in this chapter show that tuning to low RFs that are globally processed is localised to LO. This was most prominent in areas LO2 and LOC, and these were the only areas where responses were driven by stimulus shape over contrast. These results point to LO as the neural locus of human global shape processing, where shape representations become increasingly abstract and global through visual field maps LO1 and LO2 and object-selective LOC. Finally, the modelling technique used provides a method by which further dimensions of shape space may be mapped in the human brain, providing insight into the neural computations that underpin global shape and object perception.

Chapter 7: Overall discussion

7.1 Overview of objectives

The research described in this thesis used a combination of psychophysical methods, functional magnetic resonance imaging (fMRI) and transcranial magnetic stimulation (TMS) to characterise global shape processing mechanisms in humans. Global shape processing mechanisms integrate over a large portion of the visual field to generate an abstract representation of the overall shape of an object. This is an important step for recognising objects in complex visual scenes and recognising the same object under different lighting conditions, from difference viewpoints, or in the presence of occlusion (Loffler, 2008; Wilson & Wilkinson, 2015).

Most research into visual perception and object recognition has focused on early, low-level processing or high-level processes. For example, the response properties of neurons in early retinotopic cortex (e.g. V1) have been thoroughly researched (Felleman & Van Essen, 1991; Hubel & Wiesel, 1968; Zeki, 1993). Similarly, a large portion of the lateral occipital cortex that is selective for object stimuli, termed the lateral occipital complex (LOC; Malach et al., 1995), has been identified as important for object recognition and investigated extensively (Grill-spector et al., 2001). In comparison, intermediate levels of visual processing, between the local analyses of early visual cortex and high-level object representations, have received less attention, despite being an important component towards a fully realised model of object recognition (Peirce, 2015). Consequently, the first objective of this thesis was to further characterise the neural mechanisms that underpin intermediate levels of vision where global shape is first extracted.

The second objective of this thesis was to identify the neural locus of global shape representations in the human brain. Of particular interest was the roles of extrastriate visual field maps, V4, LO1 and LO2, which appear to respond to curvature, form and shape (Dumoulin & Hess, 2007; Silson et al., 2013; Vernon et al., 2016; Wilkinson et al., 2000). A secondary interest was the roles of higher-level, category selective cortex, such as occipital face area (OFA) and LOC. Complex representations of faces and objects in

these regions could be underpinned by representations of global shape, which may be first established in shape-sensitive visual field maps.

7.2 Summary of empirical chapters

7.2.1 Global shape processing in LO2 is not task specific

Chapter Three (the first empirical chapter) described an fMRI experiment designed to determine whether global shape representations in visual field map LO2 were dependent on task. The first objective of this chapter was to become familiar with fMRI methods, particularly retinotopic mapping and delineating discrete visual field maps from retinotopic mapping data. The second objective was to determine whether LO2 processes global shapes in a task-dependent manner. LO2 plays a critical role in discriminating the amplitude of radial frequency (RF) patterns (Silson et al., 2013); a task which is known to engage global shape mechanisms (Wilkinson et al., 1998). However, it had not been established whether LO2 was similarly important for global shape processing in other tasks. It was found that LO2 responded similarly to both orientation and amplitude discriminations of RF patterns. Moreover, LO2 responded more to RF patterns than the immediately posterior visual field map, LO1. This indicated that LO2 was more sensitive to global shapes than earlier visual areas, and this was not dependent on the task being performed. LO2, therefore, may be involved in the general global processing of shapes, rather than radial modulation amplitudes, specifically.

7.2.2 Global shape processing comprises multiple, increasingly global mechanisms

Chapter Four described a series of psychophysical experiments designed to further characterise the nature of global shape processing mechanisms. A novel global shape aftereffect in composite RF2+3 patterns was described, where manipulation of the phase of the RF3 component in the adaptor affected shape perceptions. Using this aftereffect, evidence was established for two distinct mechanisms involved in global shape processing. The first mechanism was invariant to stimulus spatial frequency, but tuned to stimulus size with a Gaussian tuning profile. The second mechanism was invariant to both spatial frequency and size. This suggested that global shape

processing involves multiple mechanisms whose representations become increasingly global, abstracting beyond higher-level stimulus features than preceding mechanisms. This progression of global representations may culminate in abstract internal representations of objects, facilitating object recognition. In addition, it was identified that the shape aftereffect could partially be explained by neural encoding of RF information, supporting the notion of RF-tuned channels as a perceptual mechanism (Bell & Badcock, 2009; Salmela et al., 2016).

7.2.3 LO2 and OFA may not play causal roles in shape perceptions of RF2+3 patterns

Chapter Five described two TMS experiments designed to identify which visual brain areas were causally important for global shape adaptation and discrimination in RF2+3 patterns. The areas stimulated were LO2, and nearby category-selective area OFA, which may have processed the shape information in the stimuli that is relevant to the perception of faces (Anderson & Wilson, 2005; Daar & Wilson, 2012; Fang, Murray, & He, 2007; Wilson et al., 2000, 2002). There were no effects of TMS to LO2 or OFA on shape adaptation or discrimination of RF2+3 patterns. This may indicate that neither LO2 nor OFA were causally important for global shape processing of the stimuli. Alternatively, TMS effects may have been nullified by perceptual learning (Neary et al., 2005) or local processing of the stimuli. Overall, the TMS experiments were inconclusive about the roles of LO2 and OFA in global shape processing.

7.2.4 Tuning to globally processed RFs first emerges in LO2

Finally, Chapter Six further explored the RF tuning, which was supported by behavioural evidence from Chapter Four, using fMRI to measure voxel-wise tuning to RF across visual cortex. The study used a Gaussian neural modelling technique to model tuning profiles to RF in early, lateral and ventral visual cortex. Only lateral occipital cortex showed selective for tuning for low RFs, which are known to be processed globally (Hess et al., 1999; Jeffrey et al., 2002). Global RF tuning was most pronounced in LOC, but emerged first in LO2. Moreover, only responses in LO2 and LOC were well predicted by a shape predictor that measured the circularity of the RF pattern stimuli, suggesting these regions performed some circle-based shape analysis as has been proposed in models of shape perception (Loffler, 2008; Poirier & Wilson,

2006). As such, this provided evidence that retinotopic area LO2 is the earliest stage where global shape representations emerge. These representations may form the basis of complex object representations in LOC.

Overall, the research presented in this thesis provides novel insights into the nature of global shape processing mechanisms and their locus in the human brain. The remainder of this chapter will discuss the main contributions of this thesis to the field of shape processing research. First, the thesis' contribution to the general understanding of global shape processing mechanisms will be discussed. Second, its contribution to the understanding of visual areas involved in shape processing will be discussed, including retinotopic areas V4, LO1 and LO2 and category-selective areas OFA and LOC. Finally, possible future directions for the field of global shape processing research will be proposed.

7.3 Contributions of the thesis to global shape processing research

7.3.1 Multiple mechanisms process RF information and other shape features to provide a representation of global shape that is based on circularity

Global shape processing mechanisms integrate over local, low-level information to generate an abstract representation of the overall shape of an object. An analysis of global shape is an important step in models of object recognition (Loffler, 2008; Wilson & Wilkinson, 2015), however the mechanisms involved are not as well understood as lower- and higher-level mechanisms (Peirce, 2015). One of the main objectives of this thesis was to shed new light on the neural mechanisms that underpin global shape processing.

Chapter 4 described the first example of a global shape aftereffect in composite RF patterns. The benefit of using composite RF patterns is that they can describe the arbitrary complex boundaries of natural objects such as faces (Wilson et al., 2000; Wilson et al., 2002) or animals (Vernon et al., 2016), and therefore may convey shape information that is important for biologically relevant stimuli. Moreover, the use of

composite patterns allowed for the examination of interactions between RF components, providing new insights on how RF information might be encoded in the brain.

Firstly, it was demonstrated that the same global shape mechanisms reported to underpin other shape aftereffects process composite RF patterns. Like other shape aftereffects (Anderson et al., 2007; Suzuki, 2001), the aftereffect in RF2+3 patterns was tolerant to changes in stimulus size. Moreover, unlike low-level orientation aftereffects (Ware & Mitchell, 1974), the RF2+3 pattern aftereffect was not tuned to spatial frequency and, unlike higher-level aftereffects (Afraz & Cavanagh, 2008; Zimmer & Kovács, 2011), was not position invariant. This aftereffect, therefore, could only be explained by the attenuation of intermediate-level, shape encoding mechanisms. This was the first demonstration that the same global shape mechanisms, typically probed using single RF patterns, also appear to be involved in processing composite RF patterns.

Secondly, it was demonstrated that global shape processing comprises multiple mechanisms that are increasingly global. There was evidence for two distinct mechanisms. The first mechanism was broadly tuned to stimulus size, but was invariant to its spatial frequency. The second mechanism was invariant to both size and spatial frequency. This implies a gradual transition from local to global representations, with each successive mechanism abstracting beyond additional visual features. Such an organisation is consistent with early visual cortex, where V2 neurons integrate across orientation outputs from V1 to represent angles (Hegdé & Van Essen, 2000; Ito & Komatsu, 2004). Similarly, shape processing mechanisms might integrate across this information to form a representation of global shape. This is also consistent with models of the visual system where successive perceptual mechanisms work to reduce the number of dimensions along which information is represented (Wilson & Wilkinson, 2015), increasing the complexity of the information that can be represented by a single neuron (Güçlü & van Gerven, 2015). Indeed, a gradient in the complexity of shape representations has been identified in the human visual system by Vernon et al. (2016). Overall, this research contributes to other contemporary shape research that is converging on the notion of a gradual, integrative series of mechanisms that form

increasingly global representations of shape and, eventually, abstract object representations.

Using this shape aftereffect, evidence was also provided to show that one of the features encoded by global shape mechanisms might be the RF content of the shape. RF modulations can be used to accurately describe complex, real-world objects (Vernon et al., 2016; Wilson et al., 2002), and has been hypothesised to be an important dimension for neural representations of shapes (Bell & Badcock, 2009; Salmela et al., 2016). In the experiments in this thesis, adaptation to the phase of an RF3 component affected the perception of subsequent shapes that were different to the adaptor in both shape and size, but contained the adapted RF. This implied an explicit adaptation to RF3 phase, perhaps resulting from the attenuation of RF-tuned neural channels (Bell & Badcock, 2009). Further, there was additional evidence for neural tuning to RF in Chapter Six, where neural responses to RF patterns were well predicted by a Gaussian neural model defined in RF space. This method provided the first demonstration of how the RF dimension is mapped across visual cortex (RF tunings in specific visual areas will be discussed in the following sections). This thesis therefore provides both behavioural and neuroimaging evidence for the neural encoding of RF, and introduce a technique that allows neural tuning to RF to be measured.

Finally, the RF tuning experiment also provided evidence to support models of shape perception that are based on circularity (Loffler, 2008; Poirier & Wilson, 2006). Within brain regions that showed tuning to global RFs, most voxels preferred RF2 patterns (ellipses). Moreover, neural responses within these regions were predicted accurately by a shape predictor that measured how circular each of the stimuli were. Circular (or elliptical) shapes, therefore may be particularly important for the global analysis of shape performed in these regions. In computer vision, ellipses are used as a flexible template that can be used to reliably identify both natural and man-made objects in a visual scene (Rosin, 1999; Sobottka & Pitas, 1996). Similarly, circularity has been proposed as key to human shape perception and object recognition (Loffler, 2008). Novel evidence that global shape processing regions are particularly sensitive to the circularity of a shape stimulus was provided, in line with this hypothesis.

Overall, the research in this thesis contributes to furthering the current understanding of global shape processing mechanisms. It was identified that representations of global shape are formed through a gradual, integrative process performed by multiple, increasingly global mechanisms. Both behavioural and neuroimaging methods identified that RF may be one dimension that these shape mechanisms are sensitive to. Finally, global shape processing regions were particularly sensitive to circular shapes, and circularity might therefore be a particularly important visual feature for shape detection and object recognition.

7.3.2 Visual brain areas underpinning global shape processing

The second main objective of this thesis was to identify the neural locus of global shape processing in humans. The research mostly focussed on the contributions of extrastriate visual field maps to global shape representations. Areas such as V4, LO1 and LO2 are retinotopic (Larsson & Heeger, 2006; Wandell et al., 2007; Winawer & Witthoft, 2015) and respond to shapes (Dumoulin & Hess, 2007; Silson et al., 2013; Vernon et al., 2016; Wilkinson et al., 2000). Moreover, LO1 and LO2 somewhat overlap with object-selective cortex in LOC (Sayres & Grill-Spector, 2008; Vernon et al., 2016). These regions, therefore, may form the basis for a transition for local to global representations, feeding into LOC which represents objects (Grill-spector et al., 2001; Malach et al., 1995) but is not retinotopic (Op de Beeck, Haushofer, et al., 2008). In addition to these retinotopic areas, the roles of category-selective areas OFA and LOC were also examined, which may utilise global shape as a cue for the processing of faces and objects, respectively.

7.3.3 Little evidence for a role of human V4 in global shape processing

The properties of area V4 have been thoroughly investigated in monkeys. Such experiments have identified that monkey V4 is tuned to curvature (Pasupathy & Connor, 1999, 2001), providing a population code for shapes (Pasupathy & Connor, 2002) that is invariant to some stimulus features such as size (El-Shamayleh & Pasupathy, 2016). As a result, human V4 is often assumed to be the locus of human global shape processing in models of shape perception (Loffler, 2008; Poirier & Wilson, 2006, 2010; Wilson & Wilkinson, 2015). Surprisingly, limited research into the role of human V4 in shape processing has been conducted, however two studies have shown

that it responds preferentially to radial and concentric patterns (Wilkinson et al., 2000) and concentric form (Dumoulin & Hess, 2007).

There was little evidence for global shape processing in V4 and other areas of ventral occipital cortex. In Chapter Six, V4 did not exhibit tuning to low RFs that are globally processed (Jeffrey et al., 2002). Instead, V4 responses were well predicted by a low-level predictor that measured the contrast energy of the stimuli, suggesting V4 responses were driven by low-level stimulus features rather than shape. This seems inconsistent with the aforementioned studies showing V4 responds to concentric form, however it is important to consider the differences between stimuli and how shapes were defined across studies. In Wilkinson et al.'s (2000) study, responses in V4 were measured for radial and concentric grating stimuli (see Figure 7.1A). In Dumoulin and Hess' (2007) study, responses were measured for concentric forms embedded within a field of Gabors (see Figure 7.1B). In both cases, shapes were defined by patterns or textures. Conversely, in the stimuli used shapes were defined by a single closed contour (see Figure 7.1C). Under these conditions, V4 does not appear to process shape information. Similar results have also been found by Vernon et al. (2016), who did not find shape responses in V4 using shapes defined by a single contour. V4 would therefore appear to respond to the presence of form in some cases, but does not seem to be involved in analysing the global shape of a single object.

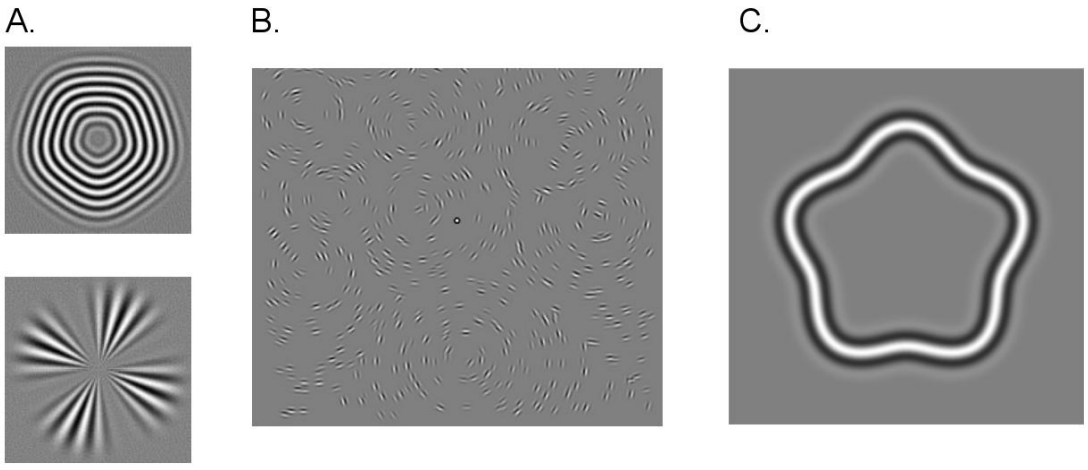

Figure 7.1: Different stimuli used to explore shape processing in V4. (A) Concentric (top) and radial (bottom) grating stimuli were used by Wilkinson et al. (2000). (B) Concentric forms embedded in a field of Gabors were used by Dumoulin & Hess (2007). (C) This thesis used RF patterns, where object shape is defined by a single closed contour.

One possible explanation for the conflicting results from V4 is that, rather than shape processing, V4 performs some form of figure-ground segregation. As shown in Figure 7.1, previous evidence for shape sensitivity in V4 has used stimuli were shape was defined by texture. Other studies have also shown that V4 is sensitive to repetitive textures (Kohleret al., 2016; Winawer & Witthoft, 2015). As well as texture, V4 is also highly sensitive to colour (Brewer et al., 2005; Wade et al., 2002, 2008). Differences in both colour and texture may be important cues for picking out an object in the visual field. Consistent with this notion, V4 has been implicated in directing visual attention for figure-ground segregation (Poort et al., 2012).

A figure-ground account of V4 could also be consistent with the research in this thesis. In Chapter Three, V4 was engaged by the presence of a shape stimulus compared to no stimulus. This could be because V4 was involved in allocating attention to segregate the stimulus from the background. Indeed, in this experiment subjects were performing a demanding stimulus-related task that required strong attention to the stimulus. In Chapter Six, V4 was tuned to high RFs. It was surmised that this was because V4 was responding to low-level stimulus features such as contrast energy, however it may instead have been responding to the increased repetitive texture about the contour in high RFs. As such, figure-ground segregation could be a function for V4 that is consistent with both the results in this thesis and the literature. If this were the case, the role of V4 in shape processing might be to pick objects out of the visual scene so that they can be analysed and identified by shape processing regions within the lateral occipital cortex.

Consequently, this research suggests that human V4 may not be a complete functional homologue of monkey V4, which is known to be tuned to shape and curvature (Pasupathy & Connor, 1999, 2001, 2002). In terms of anatomical structure, there are differences between human and monkey V4, so it follows that there might be functional differences as well. The representation of the visual field in monkey V4 is split into dorsal and ventral quarterfields, each represented in separate maps (Felleman & Van Essen, 1991), similar to V2 and V3. Human V4, on the other hand, represents the full contralateral hemifield in a single map that borders V3v (Brewer et al., 2005; Wade et al., 2002; Winawer et al., 2010; Winawer & Witthoft, 2015). A better functional homologue of the shape processing properties of monkey V4 might be LO1 or LO2.

These maps have a similar anatomical position to the dorsal portion of monkey V4, neighbouring V3d (Larsson & Heeger, 2006), and this thesis identified shape tuning in LO2 (discussed further in the next section) that is similar to curvature tuning in monkey V4 (Pasupathy & Connor, 2001). Overall, while V4 may respond to shape structures in some cases, there was little evidence that it performs the global shape processes that were the main concern of this thesis. As such, V4 does not appear to be the locus of global shape processing as is often assumed by theories of shape perception (Loffler, 2008; Poirier & Wilson, 2006, 2010; Wilson & Wilkinson, 2015), and therefore may not be a complete functional homologue of monkey V4.

7.3.4 Representations of global shape emerge in LO2

The lateral occipital cortex comprises two visual field maps, LO1 and LO2, and a large area of object-selective cortex called LOC. LO1 and LO2, which are defined retinotopically, partially overlap with LOC (Sayres & Grill-Spector, 2008; Silson et al., 2016; Vernon et al., 2016), which is not retinotopic (Op de Beeck, et al., 2008). One of the objectives of the thesis was to determine whether complex representations of objects in LOC are predicated on an analysis of global shape in LO1 or LO2.

Across multiple experiments, there was little evidence for global shape processing in LO1. In Chapter Three, LO1 responded less to global shapes than LO2. Moreover, responses in LO1 were enhanced by engaging in a task that demanded attention to a local part of the stimulus (amplitude task, which was proposed to be a more local task in this case because orientation difference between stimuli were small). This was consistent with other research showing LO1 processes low-level features of locally processed stimuli such as orientations of gratings (Larsson et al., 2006; Silson et al., 2013). Similar to other researchers, therefore, it was shown that LO1 performs more local processes than the more anterior area LO2.

For this reason LO1 was used as a control region for the TMS experiments in Chapter Five. There was no effect of TMS to LO1 on global shape adaptation or discrimination (though, there were similarly no effects of TMS in any of the stimulated regions). In Chapter Six, LO1 did not show tuning to low, globally processed RFs when low-level stimulus features were not controlled for. However, these features were controlled for, tuning to low RFs emerged in LO1. Despite a shift towards lower RF tuning in the

control experiment, unlike LO2 and LOC responses in LO1 were not better predicted by the shape predictor compared to the lower-level, contrast energy predictor. Together, the results for the RF tuning experiments suggests that LO1 may be sensitive to shape, however its representations were not sufficiently global to ignore variance in stimulus contrast energy or spatial frequency. As such, LO1 does not appear to contain global shape representations, but likely provides input to more anterior, global areas.

Our experiments indicate that LO2 is the earliest stage of the visual system where representations of global shape are formed. In Chapter Three, LO2 responded more strongly to global shapes than earlier visual areas. Moreover, responses in LO2 were stable across two different tasks, suggesting general processing of global shapes in LO2 that is not dependent on task. LO2 also showed robust tuning to low, globally processed RFs in Chapter Six. This tuning profile was preserved when high frequency contour noise was introduced to the stimuli, showing it was underpinned by global representations of low RFs. This had the effect of equalising spatial frequency across stimuli, suggesting the spatial frequency invariant shape mechanisms identified psychophysically could be in LO2.

Responses in LO2 to RF patterns were well predicted by a shape predictor that measured stimulus circularity. Additionally, a large proportion of voxels in LO2 exhibited a preference for RF2 patterns (ellipses). This suggests that LO2 may be the neural locus of circle-based shape representations (Loffler, 2008; Poirier & Wilson, 2006). As such, shape mechanisms in LO2 may use a circular template as a basis for representations of shape. Such a mechanism would be consistent with neurons in macaque inferotemporal cortex, which represent shape stimuli according to their deviation from a prototypical mean (Kayaertet al., 2005). A mechanism based on circularity could facilitate shape detection and recognition, as has been used in computer vision (Rosin, 1999; Sobottka & Pitas, 1996), and hyperacuity for global shapes (Wilkinson et al., 1998).

It is worth noting that TMS to LO2 did not affect shape perceptions of RF2+3 patterns. However, as was discussed in Chapter Five, it could not be ruled out that TMS to LO2 could affect global shape processing for different tasks or stimuli. Indeed, TMS to LO2 has been shown to impair amplitude discriminations of RF3 patterns (Silson et al.,

2013). Due to potential confounds of local processing or perceptual learning reducing the efficacy of TMS (Neary et al., 2005), it could not be concluded that LO2 plays no role in global shape processing – a conclusion that would be inconsistent with both the literature and the neuroimaging results presented in this thesis.

Overall, the neuroimaging evidence in this thesis indicates that LO2 is the earliest visual area to contain representations of global shape, and may be the neural locus of global, circular shape representations. LO1 may respond to shape features in some cases, however any representation of shape features in LO1 does not appear to be as global as in LO2. These results mirror the psychophysics experiments, which provided evidence for multiple, increasingly global shape processing mechanisms. The lateral occipital cortex, therefore, may contain the multiple shape mechanisms identified psychophysically, where shape representations become more global in more anterior areas within this region.

7.3.5 LOC contains representations of shape that are more global than LO2

As well as retinotopic visual areas, the roles of two category-selective areas in global shape processing were also explored. First, TMS was used to examine the causal role of OFA in processing shape information that may be important for face viewpoint perception (Daar & Wilson, 2012; Wilson et al., 2000). TMS to OFA had no effect on shape discrimination or shape adaptation of RF2+3 pattern stimuli. However, this method was unable to find effects of TMS to any cortical regions, including LO2, which the neuroimaging experiments suggest processes global shapes. The results regarding OFA's role in global shape processing were therefore inconclusive. There are several reasons to think OFA might be involved in global shape processing. First, OFA is causally important for discriminating shape orientations (Silvanto et al., 2010) and symmetry (Bona et al., 2015) in non-face stimuli. As such OFA would appear to process visual properties that are not specific to faces, but may be salient in face stimuli. Second, the global shape information present in RF patterns can accurately describe changes in face viewpoint (Daar & Wilson, 2012; Wilson et al., 2000), which is encoded in OFA (Fang et al., 2007). Global shape representations therefore may be useful for accurate face perception and might be processed in OFA. Overall, the results regarding

OFA were inconclusive, and there is more room to explore the relationship between global shape processing and face representations within OFA.

RF tuning in object-selective LOC was also measured. Similar to LO2, LOC was tuned to low, global RFs across both the standard and control experiments. However, this seemed to be underpinned by more global shape representations in LOC compared to LO2 as only responses in LOC correlated significantly with a shape-based predictor for both standard and control stimuli. As such, increasingly global shape representations through LO1 and LO2 appear to persist in LOC, which is in turn more global than LO2. The psychophysics experiments identified a secondary shape mechanism that was invariant to stimulus size, while the first was tuned to size. Neural adaptation in LOC to objects and shapes has also been shown to be invariant to stimulus size (Grill-Spector et al., 1999; Sawamura et al., 2005), where adaptation in more anterior portions of LOC were increasingly size invariant. As such, the size-invariant shape mechanism identified psychophysically could reside in LOC, with the size-tuned mechanisms residing in an earlier shape processing regions such as LO1 or LO2. This therefore shows that representations of shape become increasingly global through LO1, LO2 and LOC. As such, complex representations of objects in LOC are likely underpinned by a transition from local to more global analysis of shape in LO1 and LO2.

7.3.6 Conclusions on global shape processing mechanisms

Overall, these experiments show that global shape processing comprises multiple mechanisms that process increasingly abstract shape features for more global representations of shape. Shape processing begins with a local analysis of low-level features such as orientation and spatial frequency in early visual cortex (Hubel & Wiesel, 1968). Information about shape is then propagated through the lateral occipital cortex, where more anterior visual areas integrate across outputs from the preceding area to form more global shape representations. This is likely facilitated by increasing neural receptive field sizes through the visual system (Amano et al., 2009; Dumoulin & Wandell, 2008), where larger receptive fields allow neurons to integrate over larger portions of the visual field. LO1 is sensitive to some shape features, however the first global shape representations emerge in LO2. LO2 is both retinotopic and partially object-selective (Sayres & Grill-Spector, 2008; Silson et al., 2016; Vernon et al., 2016),

suggesting it is a transitional stage from local to global processing in the visual system. This likely underpins more abstract and global representations of objects in LOC, which is also object-selective but no longer retinotopic (Op de Beeck, et al., 2008). Together, LO2 and LOC form the basis for circular representations of shapes, which allow for a hypersensitivity to shapes (Wilkinson et al., 1998) that may underpin accurate object recognition.

7.4 Future research directions

7.4.1 Using shape adaptation to characterise RF encoding mechanisms

Shape processing research has converged on the notion that shapes are represented in a multidimensional shape space (Kayaert et al., 2005; Op De Beeck, Wagemans, & Vogels, 2001; Pasupathy & Connor, 2001, 2002; Vernon et al., 2016). It was found both behavioural and neural evidence for neural encoding of RF information, suggests RF could be one dimension of this space. However, it remains to be determined whether RF is a true dimension of shape space, or if the tuning to RF shown is underpinned by tuning to some other aspect of shape or curvature that is correlated with RF. For example, increases in RF are coupled by a decrease in the angle subtended between the lobes and sensitivity to this metric has been shown to be a better model for the perception of modulated RF patterns where the radius is modulated with a rectified sinusoid (Dickinson et al., 2015).

The psychophysics experiments in this thesis showed that adaptation to RF3 phase transferred across both stimulus shape and size, affecting the perception of larger RF2+3 patterns. This was interpreted as evidence for an RF encoding mechanism where shape perceptions were affected by the attenuation of neural channels tuned to RF3. Alternatively this could be interpreted as adaptation of some other curvature mechanism that is size invariant and not selective for shape. If the aftereffect was driven by the attenuation of RF-tuned channels, adaptation to an RF3 component should have a smaller effect or no effect, depending on the tuning width of RF channels, in stimuli without an RF3 component such as an RF2+4 pattern. By measuring the effect of adaptation to one RF component on the perception of

composite RF patterns comprising RF components varied around the adaptor, one could establish whether there are RF-tuned channels and how tightly they are tuned.

For example, consider an experiment similar to the adaptation experiments in Chapter Four, where the subject adapts to the phase of an RF6 component and is tested on a series of test stimuli ranging from RF2+3 to RF2+9 patterns. By measuring aftereffect magnitudes across the different test stimuli, one could learn about the nature of RF tuning. If the brain does not encode RF, then the aftereffects might be approximately the same across all test stimuli as adaptation of curvature or shape mechanisms that are not selective for RF affects all the stimuli equally (Figure 7.2A). If RF patterns are processed by narrow-band RF channels, as proposed by (Bell & Badcock, 2009), aftereffects should be maximal in RF2+6 patterns (including the adapted RF) and minimal for other test stimuli (Figure 7.2B). If RF channels are more broadly tuned, the aftereffect should gradually drop off as the secondary RF component moves further away from the adapted RF (Figure 2.7C). Using such an experiment, one could utilise the novel class of shape aftereffect introduced in this thesis to characterise the nature of RF encoding mechanisms.

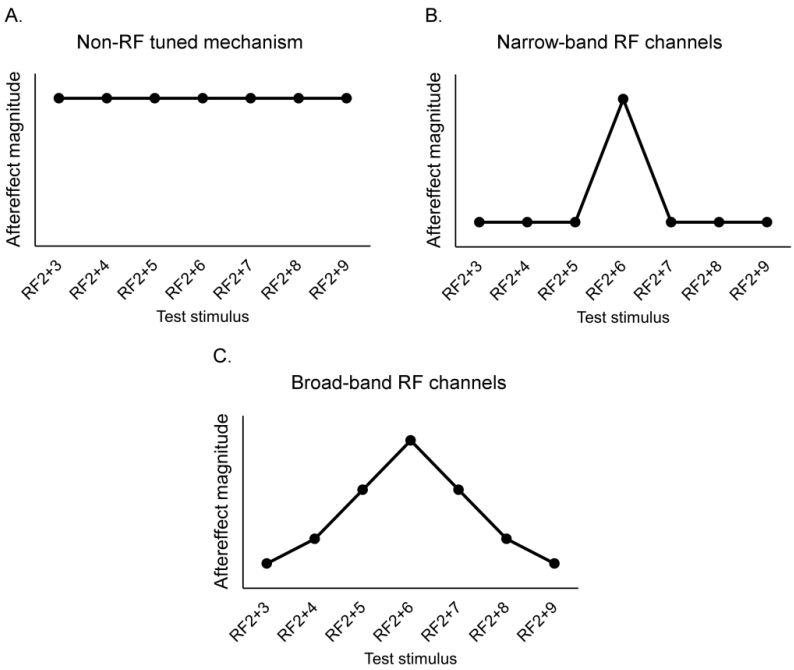

Figure 7.2: Hypothetical results for an experiment measuring shape aftereffects in test stimuli ranging from RF2+3 to RF2+9 patterns following adaptation to an RF6 component. Hypothetical results are plotted for scenarios where the aftereffect is driven by (A) a curvature mechanism that is not tuned to RF, (B) an RF encoding mechanism where RF channels are narrowly tuned and (C) an RF encoding mechanism were RF channels are broadly tuned.

7.4.2 Using TMS to establish causal roles for LO2 and LOC in global shape processing

Chapter 5 was unable to establish whether LO2 plays a causal role in global shape processing using TMS. Additionally, TMS was not applied to LOC in any experiments, which evidently contains global shape representations. Whilst neural representations of global shape were successfully identified using fMRI, it is important to also establish causal links between neural processing and behaviour to fully understand the underlying mechanisms. As discussed in Chapter Five, the choice of task may not have been appropriate for a TMS experiment on global shape processing as, although it clearly engages global mechanisms, there may have been enough local cues for local mechanisms to compensate for the effects of TMS. One task that is known to rely on global shape processing mechanisms is the detection of deformation from circularity in RF patterns (Wilkinson et al., 1998). The hypersensitivity shown by observers in this task can only be achieved by global mechanisms, as the same sensitivity is not seen for stimuli that are not processed globally (Jeffrey et al., 2002; Wilkinson et al., 1998). Combining this task with TMS could determine whether computations within LO2 and LOC underpin this hypersensitivity to global shape. For example, TMS could be applied to LO2 (or LOC) and V1 during a deformation from circularity detection task in RF2 and RF20 patterns. Based on the neuroimaging results in this thesis, one would expect a double dissociation where TMS reduces performance for RF2 patterns (global task), but not RF20 patterns (local task), while TMS to V1 should impair performance for RF20 but not RF2 patterns (see Figure 7.3). Using this method, one could demonstrate that processing in LO2 (or LOC) is critical to accurate perceptions of global shape.

7.4.3 Mapping a multidimensional representational shape space

Shape processing research has converged on the notion that shapes are represented in a multidimensional shape space (Kayaert et al., 2005; Op De Beeck, Wagemans, & Vogels, 2001; Pasupathy & Connor, 2001, 2002; Vernon et al., 2016). The representations of other feature spaces, such as visual space (Dumoulin & Wandell, 2008) and numerosity space (Harvey et al., 2013), have been successfully mapped in the human brain using Gaussian neural modelling techniques. In Chapter Six, the same technique was used to map neural sensitivity to RF. This provides an opportunity to apply the

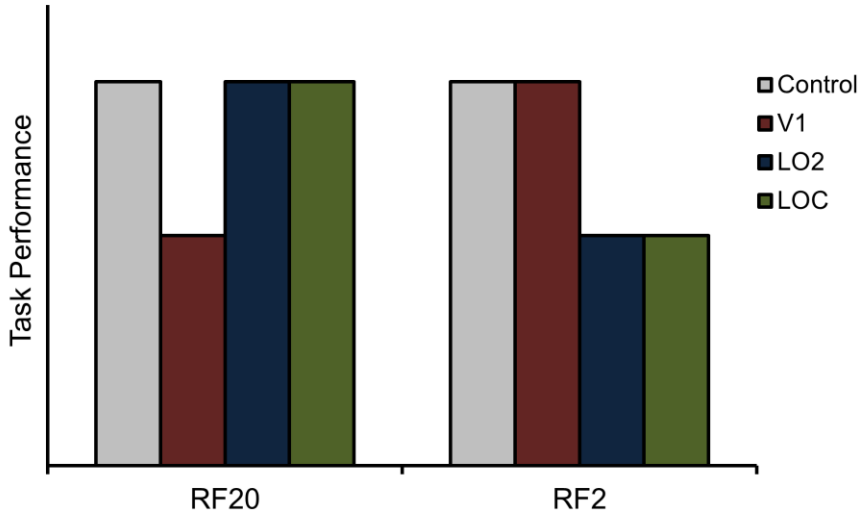

Figure 7.3: Hypothetical data for an experiment measuring performance in a detection of deformation from circularity task following TMS to V1, LO2 or LOC. Performance is shown for a control TMS condition (Grey), V1 TMS (red), LO2 TMS (blue) and LOC TMS (green). It is predicted that V1 would reduce performance for RF20 patterns (local task) but not RF2 patterns (global task), while TMS to LO2 and LOC would impair detection of RF2 but not RF20 patterns.

same technique to other possible dimensions of shape space. For example, one could further explore the neural representation of RF patterns by modelling neural tuning to other RF parameters, such as amplitude or the number of RFs within a composite pattern. These parameters could also be manipulated together, and neural tuning profiles within RF space could be modelled by a three-dimensional Gaussian where the three dimensions are RF, amplitude and number of RFs.

This method could also be applied to mapping dimensions of a more general shape space that is not specific to RF patterns. For example, the angle subtended at the centre of a shape by adjacent points of maximum curvature is thought to be important for shape representations (Dickinson et al., 2015), however no neural evidence has been established to support the psychophysical evidence. Neural tuning profiles to curvature have been thoroughly researched in the monkey (Pasupathy & Connor, 1999, 2001) to characterise a representational shape space within the monkey brain (Kayaert et al., 2005; Op De Beeck et al., 2001; Pasupathy & Connor, 2002). However, due to the practical limitations of single-unit recordings in humans, this type of research has not been replicated in the human visual system. Gaussian neural modelling provides an opportunity for the principles of shape processing research in monkeys to be applied

humans, modelling population tunings to curvature or shape in fMRI voxels. The modelling technique described, therefore, is a powerful method that researchers could use to better understand the neural representation of shapes, and potentially other feature spaces as well.

7.5 Final conclusions

In conclusion, this thesis provides a number of meaningful contributions to the field of global shape processing research. Psychophysical, neuroimaging and TMS methods were used to further characterisee the nature of global shape processing mechanisms and identified their neural locus in humans. Using a novel global shape aftereffect in composite RF patterns, global shape processing was shown to comprise multiple, increasingly global mechanisms. Using fMRI, these mechanisms were identified in the lateral occipital cortex. The first neural representations of global shape emerge in LO2, which likely integrates over outputs from earlier visual areas to form its representations. Similarly, LOC forms more global shape representations than LO2, which likely form the basis of complex representations of objects in LOC. These mechanisms appear to process shapes according to their position in a multidimensional representational shape space. From both behavioural and fMRI experiments there was convergent evidence to suggest RF might be one dimension of this space. Gaussian neural modelling was used to map the representation of the RF dimension in the human brain, identifying neural tuning to global RFs in LO2 and LOC. In the future, this technique can be applied to mapping other dimensions of shape space, providing a means to obtain a fuller understanding of how global shapes are represented to facilitate accurate object recognition.

References

Addams, R. (1834). An account of a peculiar optical phænomenon seen after having looked at a moving body. *The London and Edinburgh Philosophical Magazine and Journal of Science*, *5*(29), 373–374.

Afraz, S.-R., & Cavanagh, P. (2008). Retinotopy of the face aftereffect. *Vision Research*, *48*(1), 42–54.

Allen, E. A., Pasley, B. N., Duong, T., Freeman, R. D., & Miller, G. (2007). Transcranial Magnetic Stimulation Elicits Coupled Neural and Hemodynamic Consequences. *Science*, *317*, 1918–1921.

Alter, I., & Schwartz, E. L. (1988). Psychophysical studies of shape with Fourier descriptor stimuli. *Perception*, *17*(2), 191–202.

Amano, K., Wandell, B. A., & Dumoulin, S. O. (2009). Visual field maps, population receptive field sizes, and visual field coverage in the human MT+ complex. *Journal of Neurophysiology*, *102*(5), 2704–18.

Anderson, N. D., Habak, C., Wilkinson, F., & Wilson, H. R. (2007). Evaluating shape after-effects with radial frequency patterns. *Vision Research*, *47*(3), 298–308.

Anderson, N. D., & Wilson, H. R. (2005). The nature of synthetic face adaptation. *Vision Research*, *45*(14), 1815–1828.

Andrews, T. J., & Ewbank, M. P. (2004). Distinct representations for facial identity and changeable aspects of faces in the human temporal lobe. *NeuroImage*, *23*(3), 905–913.

Andrews, T. J., Watson, D. M., Rice, G. E., & Hartley, T. (2015). Low-level properties of natural images predict topographic patterns of neural response in the ventral visual pathway. *Journal of Vision*, *15*(7), 3.

Attneave, F. (1954). Some informational aspects of visual perception. *Psychological Review*, *61*(3), 183–193.

Bell, J., & Badcock, D. R. (2009). Narrow-band radial frequency shape channels

revealed by sub-threshold summation. *Vision Research*, *49*(8), 843–850.

Bell, J., Badcock, D. R., Wilson, H., & Wilkinson, F. (2007). Detection of shape in radial frequency contours: Independence of local and global form information. *Vision Research*, *47*(11), 1518–1522.

Bell, J., Forsyth, M., Badcock, D. R., & Kingdom, F. A. A. (2014). Global shape processing involves feature-selective and feature-agnostic coding mechanisms. *Journal of Vision*, *14*(11), 1–14.

Bell, J., Hancock, S., Kingdom, F. A. A., & Peirce, J. W. (2010). Global shape processing: which parts form the whole? *Journal of Vision*, *10*(6), 1–13.

Bell, J., & Kingdom, F. A. A. (2009). Global contour shapes are coded differently from their local components. *Vision Research*, *49*(13), 1702–1710.

Bell, J., Wilkinson, F., Wilson, H. R., Loffler, G., & Badcock, D. R. (2009). Radial frequency adaptation reveals interacting contour shape channels. *Vision Research*, *49*(18), 2306–2317.

Bolognini, N., & Ro, T. (2010). Transcranial Magnetic Stimulation: Disrupting Neural Activity to Alter and Assess Brain Function. *The Journal of Neuroscience*, *30*(29), 9647–9650.

Bona, S., Cattaneo, Z., & Silvanto, J. (2015). The Causal Role of the Occipital Face Area (OFA) and Lateral Occipital (LO) Cortex in Symmetry Perception. *Journal of Neuroscience*, *35*(2), 731–738.

Brewer, A., Liu, J., Wade, A., & Wandell, B. (2005). Visual field maps and stimulus selectivity in human ventral occipital cortex. *Nature Neuroscience*, *8*(8), 1102–9.

Brincat, S. L., & Connor, C. E. (2004). Underlying principles of visual shape selectivity in posterior inferotemporal cortex. *Nature Neuroscience*, *7*(8), 880–886.

Burton, M. P., McKeefry, D. J., Barrett, B. T., Vakrou, C., & Morland, A. A. B. (2009). Disruptions to human speed perception induced by motion adaptation and transcranial magnetic stimulation. *The European Journal of Neuroscience*, *30*(10), 1989–98.

Cadieu, C., Kouh, M., Pasupathy, A., Connor, C. E., Riesenhuber, M., & Poggio, T. (2007). A Model of V4 Shape Selectivity and Invariance. *Journal of Neurophysiology*, *98*, 1733–1750.

Cattaneo, Z., & Silvanto, J. (2008). Time course of the state-dependent effect of transcranial magnetic stimulation in the TMS-adaptation paradigm. *Neuroscience Letters*, *443*(2), 82–5.

Coggan, D. D., Liu, W., Baker, D. H., & Andrews, T. J. (2016). Category-selective patterns of neural response in the ventral visual pathway in the absence of categorical information. *NeuroImage*, *135*(April), 107–114.

Corbetra, M., Miezin, F. M., Dobmeyer, S., Shulman, G. L., & Petersen, S. E. (1989). Attentional Modulation of Neural Processing of Shape , Color , and Velocity in Humans. *Science*, *248*(4962), 1556–1559.

Cornsweet, T. N. (1962). The Staircase-Method in Psychophysics. *The American Journal of Psychology*, *75*(3), 485–491.

Daar, M., & Wilson, H. R. (2012). The face viewpoint aftereffect: adapting to full faces, head outlines, and features. *Vision Research*, *53*(1), 54–9.

Deyoe, E. A., Carmant, G. J., Bandettinit, P., Glickman, S., Wieser, J., Cox, R., … Neitz, J. (1996). Mapping striate and extrastriate visual cerebral cortex human. *Proceedings of the National Academy of Sciences*, *93*, 2382–2386.

Dickinson, J. E., Almeida, R. A., Bell, J., & Badcock, D. R. (2010). Global shape aftereffects have a local substrate: A tilt aftereffect field. *Journal of Vision*, *10*(13), 1–12.

Dickinson, J. E., Cribb, S. J., Riddell, H., & Badcock, D. R. (2015). Tolerance for local and global differences in the integration of shape information. *Journal of Vision*, *15*(3), 1–24.

Dickinson, J. E., McGinty, J., Webster, K. E., & Badcock, D. R. (2012). Further evidence that local cues to shape in RF patterns are integrated globally. *Journal of Vision*, *12*(12), 1–17.

Downing, P. E., Jiang, Y., Shuman, M., & Kanwisher, N. (2001). A cortical area selective for visual processing of the human body. *Science*, *293*(5539), 2470–3.

Drucker, D. M., & Aguirre, G. K. (2009). Different spatial scales of shape similarity representation in lateral and ventral LOC. *Cerebral Cortex*, *19*(10), 2269–2280.

Dumoulin, S. O., & Hess, R. F. (2007). Cortical specialization for concentric shape processing. *Vision Research*, *47*(12), 1608–13.

Dumoulin, S. O., & Wandell, B. A. (2008). Population receptive field estimates in human visual cortex. *NeuroImage*, *39*(2), 647–60.

Elkin-Frankston, S., Fried, P. J., Rushmore, R. J., & Valero-Cabré, A. (2011). From qualia to quantia: a system to document and quantify phosphene percepts elicited by non-invasive neurostimulation of the human occipital cortex. *Journal of Neuroscience Methods*, *198*(2), 149–57.

El-Shamayleh, Y., & Pasupathy, A. (2016). Contour Curvature As an Invariant Code for Objects in Visual Area V4. *Journal of Neuroscience*, *36*(20), 5532–5543.

Engel, S. A., Glover, G. H., & Wandell, B. A. (1997). Retinotopic organization in human visual cortex and the spatial precision of functional MRI. *Cerebral Cortex*, *7*(2), 181–92.

Engel, S. A., Rumelhart, D. E., Wandell, B. A., Lee, A. T., Glover, G. H., Chichilnisky, E., & Shadlen, M. N. (1994). fMRI of human visual cortex. *Nature*, *369*, 525.

Epstein, R., & Kanwisher, N. (1998). A cortical representation of the local visual environment. *Nature*, *392*, 6–9.

Fang, F., & He, S. (2005). Viewer-centered object representation in the human visual system revealed by viewpoint aftereffects. *Neuron*, *45*(5), 793–800.

Fang, F., Murray, S. O., & He, S. (2007). Duration-dependent FMRI adaptation and distributed viewer-centered face representation in human visual cortex. *Cerebral Cortex*, *17*(6), 1402–11.

Felleman, D. J., & Van Essen, D. C. (1991). Distributed hierarchical processing in the primate cerebral cortex. *Cerebral Cortex*, *1*(1), 1–47.

Foster, D. H., Simmons, D. R., & Cook, M. J. (1993). The cue for contour-curvature discrimination. *Vision Research*, *33*(3), 329–341.

Freire, A., Lee, K., & Symons, L. A. (2000). The face-inversion effect as a deficit in the encoding of configural information: Direct evidence. *Perception*, *29*(2), 159–170.

Gallant, J. L., Braun, J., & Van Essen, D. C. (1993). Selectivity for polar, hyperbolic, and Cartesian gratings in macaque visual cortex. *Science*, *259*(5091), 100–103.

Gallant, J. L., Shoup, R. E., & Mazer, J. A. (2000). A human extrastriate area functionally homologous to macaque V4. *Neuron*, *27*(2), 227–235.

Gandhi, S. P., Heeger, D. J., & Boynton, G. M. (1999). Spatial attention affects brain activity in human primary visual cortex. *Proceedings of the National Academy of Sciences*, *96*, 3314–3319.

Gauthier, I., Michael, J. T., Moylan, J., Skudlarksi, P., Gore, J. C., & Anderson, A. W. (2000). The Fusiform "Face Area" is Part of a Network that Processes Faces at the Individual Level. *Journal of Cognitive Neuroscience*, *12*(3), 495–504.

Gauthier, I., Skudlarski, P., Gore, J. C., & Anderson, A. W. (2000). Expertise for cars and birds recruits brain areas involved in face recognition. *Nature Neuroscience*, *3*(2), 191–7.

Gheorghiu, E., & Kingdom, F. A. A. (2006). Luminance-contrast properties of contour-shape processing revealed through the shape-frequency after-effect. *Vision Research*, *46*(21), 3603–3615.

Gheorghiu, E., & Kingdom, F. A. A. (2007). The spatial feature underlying the shape-frequency and shape-amplitude after-effects. *Vision Research*, *47*, 834–844.

Gheorghiu, E., Kingdom, F. A. A., Thai, M. T., & Sampasivam, L. (2009). Binocular properties of curvature-encoding mechanisms revealed through two shape after-effects. *Vision Research*, *49*(14), 1765–1774.

Grill-spector, K., Kourtzi, Z., & Kanwisher, N. G. (2001). The lateral occipital complex and its role in object recognition. *Vision Research*, *41*, 1409–1422.

Grill-Spector, K., Kushnir, T., Edelman, S., Avidan, G., Itzchak, Y., & Malach, R. (1999).

Differential processing of objects under various viewing conditions in the human lateral occipital complex. *Neuron*, *24*(1), 187–203.

Güçlü, U., & van Gerven, M. A. J. (2015). Deep Neural Networks Reveal a Gradient in the Complexity of Neural Representations across the Ventral Stream. *Journal of Neuroscience*, *35*(27), 10005–10014.

Hallett, M. (2007). Transcranial magnetic stimulation: a primer. *Neuron*, *55*(2), 187–99.

Harvey, B. M., Klein, B. P., Petridou, N., & Dumoulin, S. O. (2013). Topographic representation of numerosity in the human parietal cortex. *Science*, *341*(September), 1123–1126.

Haxby, J., Hoffman, E., & Gobbini, M. (2000). The distributed human neural system for face perception. *Trends in Cognitive Sciences*, *4*(6), 223–233.

Hegdé, J., & Van Essen, D. C. (2000). Selectivity for complex shapes in primate visual area V2. *The Journal of Neuroscience*, *20*(5), RC61.

Hess, R. F., Wang, Y. Z., & Dakin, S. C. (1999). Are judgements of circularity local or global? *Vision Research*, *39*(26), 4354–4360.

Heywood, C. A., & Cowey, A. (1987). On the role of cortical area V4 in the discrimination of hue and pattern in macaque monkeys. *The Journal of Neuroscience*, *7*(9), 2601–2617.

Hubel, D. H., & Wiesel, T. N. (1968). Receptive Fields and Functional Architecture of Monkey Striate Cortex. *Journal of Physiology*, *195*, 215–243.

Ito, M., & Komatsu, H. (2004). Representation of Angles Embedded within Contour Stimuli in Area V2 of Macaque Monkeys. *The Journal of Neuroscience*, *24*(13), 3313–3324.

Jeffrey, B. G., Wang, Y. Z., & Birch, E. E. (2002). Circular contour frequency in shape discrimination. *Vision Research*, *42*(25), 2773–2779.

Kanizsa, G. (1976). Subjective contours. *Scientific American*, *234*(4), 48–52.

Kanwisher, N., McDermott, J., & Chun, M. M. (1997). The fusiform face area: a module

in human extrastriate cortex specialized for face perception. *The Journal of Neuroscience*, *17*(11), 4302–11.

Kayaert, G., Biederman, I., Op De Beeck, H. P., & Vogels, R. (2005). Tuning for shape dimensions in macaque inferior temporal cortex. *European Journal of Neuroscience*, *22*(1), 212–224.

Kayaert, G., Biederman, I., & Vogels, R. (2003). Shape tuning in macaque inferior temporal cortex. *The Journal of Neuroscience*, *23*(7), 3016–3027.

Kietzmann, T. C., Ling, S., Poltoratski, S., Konig, P., Blake, R., & Tong, F. (2015). The Occipital Face Area is Causally Involved in Viewpoint Symmetry Judgments of Faces. *The Journal of Neuroscience*, *35*(50), 16398–16403.

Kohler, P. J., Clarke, A., Yakovleva, A., Liu, Y., & Norcia, A. M. (2016). Representation of Maximally Regular Textures in Human Visual Cortex. *Journal of Neuroscience*, *36*(3), 714–729.

Kourtzi, Z., & Kanwisher, N. (2001). Representation of Perceived Object Shape by the Human Lateral Occipital Complex. *Science*, *293*, 1506–1510.

Kovács, I., & Julesz, B. (1993). A closed curve is much more than an incomplete one: effect of closure in figure-ground segmentation. *Proceedings of the National Academy of Sciences of the United States of America*, *90*(16), 7495–7497.

Larsson, J., & Heeger, D. J. (2006). Two retinotopic visual areas in human lateral occipital cortex. *The Journal of Neuroscience*, *26*(51), 13128–42.

Larsson, J., Landy, M. S., & Heeger, D. J. (2006). Orientation-selective adaptation to first- and second-order patterns in human visual cortex. *Journal of Neurophysiology*, *95*(2), 862–881.

Loffler, G. (2008). Perception of contours and shapes: Low and intermediate stage mechanisms. *Vision Research*, *48*(20), 2106–2127.

Loffler, G., Wilson, H. R., & Wilkinson, F. (2003). Local and global contributions to shape discrimination. *Vision Research*, *43*(5), 519–530.

Loffler, G., Yourganov, G., Wilkinson, F., & Wilson, H. R. (2005). fMRI evidence for the

neural representation of faces. *Nature Neuroscience*, *8*(10), 1386–1390.

Logothetis, N. K. (2003). The Underpinnings of the BOLD Functional Magnetic Resonance Imaging Signal. *The Journal of Neuroscience*, *23*(10), 3963–3971.

Malach, R., Reppas, J. B., Beson, R. R., Kwong, K. K., Jiang, H., Kennedy, W. A., … Tootell, R. B. H. (1995). Object-related activity revealed by functional magnetic resonance imaging in human occipital cortex. *Proceedings of the National Academy of Sciences*, *92*, 8135–8139.

Maurer, D., Le Grand, R., & Mondloch, C. J. (2002). The many faces of configural processing. *Trends in Cognitive Sciences*, *6*(6), 255–260.

McCollough, C. (1965). Color Adaptation of Edge-Detectors in the Human Visual System. *Science*, *149*(3688), 1115–1116.

McKeefry, D. J., Burton, M. P., Vakrou, C., Barrett, B. T., & Morland, A. B. (2008). Induced deficits in speed perception by transcranial magnetic stimulation of human cortical areas V5/MT+ and V3A. *The Journal of Neuroscience*, *28*(27), 6848–57.

McKeefry, D. J., Gouws, A., Burton, M. P., & Morland, A. B. (2009). The noninvasive dissection of the human visual cortex: using FMRI and TMS to study the organization of the visual brain. *The Neuroscientist : A Review Journal Bringing Neurobiology, Neurology and Psychiatry*, *15*(5), 489–506.

Melcher, D. (2005). Spatiotopic transfer of visual-form adaptation across saccadic eye movements. *Current Biology*, *15*(19), 1745–8.

Morgan, M. (2013). Sustained attention is not necessary for velocity adaptation. *Journal of Vision*, *13*(8), 1–11.

Morgan, M., Dillenburger, B., Raphael, S., & Solomon, J. A. (2012). Observers can voluntarily shift their psychometric functions without losing sensitivity. *Attention, Perception & Psychophysics*, *74*(1), 185–93.

Morgan, M. J. (2005). The visual computation of 2-D area by human observers. *Vision Research*, *45*(19), 2564–2570.

Neary, K., Anand, S., & Hotson, J. R. (2005). Perceptual learning of line orientation modifies the effects of transcranial magnetic stimulation of visual cortex. *Experimental Brain Research*, *162*(1), 23–34.

Nichols, D. F., Betts, L. R., & Wilson, H. R. (2010). Decoding of faces and face components in face-sensitive human visual cortex. *Frontiers in Psychology*, *1*(July), Article 28.

Op de Beeck, H. P., Haushofer, J., & Kanwisher, N. G. (2008). Interpreting fMRI data: maps, modules and dimensions. *Nature Reviews. Neuroscience*, *9*(2), 123–35.

Op de Beeck, H. P., Torfs, K., & Wagemans, J. (2008). Perceived shape similarity among unfamiliar objects and the organization of the human object vision pathway. *The Journal of Neuroscience*, *28*(40), 10111–23.

Op De Beeck, H. P., Wagemans, J., & Vogels, R. (2001). Inferotemporal neurons represent low-dimensional configurations of parameterized shapes. *Nature Neuroscience*, *4*(12), 1244–1252.

Pasupathy, A., & Connor, C. E. (1999). Responses to contour features in macaque area V4. *Journal of Neurophysiology*, *82*(5), 2490–2502.

Pasupathy, A., & Connor, C. E. (2001). Shape representation in area V4: position-specific tuning for boundary conformation. *Journal of Neurophysiology*, *86*(5), 2505–2519.

Pasupathy, A., & Connor, C. E. (2002). Population coding of shape in area V4. *Nature Neuroscience*, *5*(12), 1332–8.

Peirce, J. W. (2015). Understanding mid-level representations in visual processing vision. *Journal of Vision*, *15*(7), 1–9.

Pitcher, D. (2014). Facial expression recognition takes longer in the posterior superior temporal sulcus than in the occipital face area. *The Journal of Neuroscience*, *34*(27), 9173–7.

Pitcher, D., Walsh, V., & Duchaine, B. (2011). The role of the occipital face area in the cortical face perception network. *Experimental Brain Research*, *209*, 481–493.

Pitcher, D., Walsh, V., Yovel, G., & Duchaine, B. (2007). TMS evidence for the involvement of the right occipital face area in early face processing. *Current Biology*, *17*(18), 1568–73.

Poirier, F. J. A. M., & Wilson, H. R. (2006). A biologically plausible model of human radial frequency perception. *Vision Research*, *46*(15), 2443–55.

Poirier, F. J. A. M., & Wilson, H. R. (2010). A biologically plausible model of human shape symmetry perception. *Journal of Vision*, *10*(1), 1–16.

Poort, J., Raudies, F., Wannig, A., Lamme, V. A. F., Neumann, H., & Roelfsema, P. R. (2012). Article The Role of Attention in Figure-Ground Segregation in Areas V1 and V4 of the Visual Cortex. *Neuron*, *75*, 143–156.

Rainville, S. J., Yourganov, G., & Wilson, H. R. (2010). Closed-contour shapes encoded through deviations from circularity in lateral-occipital complex (LOC): An fMRI study. *Journal of Vision*, *5*(8), 471–471.

Roe, A. W., Chelazzi, L., Connor, C. E., Conway, B. R., Fujita, I., Gallant, J. L., … Vanduffel, W. (2012). Toward a Unified Theory of Visual Area V4. *Neuron*, *74*(1), 12–29.

Rosin, P. L. (1999). Further Five Point Fit Ellipse Fitting. *Graphical Models and Image Processing*, *61*(5), 245–259.

Rossi, S., Hallett, M., Rossini, P. M., & Pascual-Leone, A. (2009). Safety, ethical considerations, and application guidelines for the use of transcranial magnetic stimulation in clinical practice and research. *Clinical Neurophysiology*, *120*(12), 2008–39.

Rossion, B., Caldara, R., Seghier, M., Schuller, A., Lazeyras, F., & Mayer, E. (2003). A network of occipito-temporal face-sensitive areas besides the right middle fusiform gyrus is necessary for normal face processing. *Brain*, *126*, 2381–2395.

Salmela, V. R., Henriksson, L., & Vanni, S. (2016). Radial Frequency Analysis of Contour Shapes in the Visual Cortex. *PLoS Computational Biology*, *12*(2), e1004719.

Sawamura, H., Georgieva, S., Vogels, R., Vanduffel, W., & Orban, G. A. (2005). Using

Functional Magnetic Resonance Imaging to Assess Adaptation and Size Invariance of Shape Processing by Humans and Monkeys. *The Journal of Neuroscience*, *25*(17), 4294–4306.

Sayres, R., & Grill-Spector, K. (2008). Relating retinotopic and object-selective responses in human lateral occipital cortex. *Journal of Neurophysiology*, *100*(1), 249–67.

Schiller, P. H. (1995). Effect of lesions in visual cortical area V4 on the recognition of transformed objects. *Nature*, *376*(27), 342–344.

Schira, M. M., Tyler, C. W., Breakspear, M., & Spehar, B. (2009). The foveal confluence in human visual cortex. *The Journal of Neuroscience*, *29*(28), 9050–8.

Schmidtmann, G., Kennedy, G. J., Orbach, H. S., & Loffler, G. (2012). Non-linear global pooling in the discrimination of circular and non-circular shapes. *Vision Research*, *62*, 44–56.

Sereno, M. I., Dale, A. M., Reppas, J. B., Kwong, K. K., Belliveau, J. W., Brady, T. J., … Tootell, R. B. (1995). Borders of multiple visual areas in humans revealed by functional magnetic resonance imaging. *Science*, *268*(5212), 889–93.

Silson, E. H. (2013). *Functional Specialization & Parallel Processing within Retinotopic Subdivisions of Lateral Occipital Cortex*. University of York.

Silson, E. H., Groen, I. I. A., Kravitz, D. J., & Baker, C. I. (2016). Evaluating the correspondence between face- , scene- , and object-selectivity and retinotopic organization within lateral occipitotemporal cortex. *Journal of Vision*, *16*(6), 1–21.

Silson, E. H., McKeefry, D. J., Rodgers, J., Gouws, A. D., Hymers, M., & Morland, A. B. (2013). Specialized and independent processing of orientation and shape in visual field maps LO1 and LO2. *Nature Neuroscience*, *16*(3), 267–9.

Silvanto, J., Muggleton, N. G., Cowey, A., & Walsh, V. (2007). Neural adaptation reveals state-dependent effects of transcranial magnetic stimulation. *The European Journal of Neuroscience*, *25*(6), 1874–81.

Silvanto, J., Muggleton, N. G., & Walsh, V. (2008). State-dependency in brain

stimulation studies of perception and cognition. *Trends in Cognitive Sciences*, *12*(12), 447–54.

Silvanto, J., & Pascual-Leone, A. (2008). State-dependency of transcranial magnetic stimulation. *Brain Topography*, *21*(1), 1–10.

Silvanto, J., Schwarzkopf, D. S., Gilaie-Dotan, S., & Rees, G. (2010). Differing causal roles for lateral occipital cortex and occipital face area in invariant shape recognition. *The European Journal of Neuroscience*, *32*(1), 165–71.

Silver, M. A., Ress, D., & Heeger, D. J. (2007). Neural Correlates of Sustained Spatial Attention in Human Early Visual Cortex. *Journal of Neurophysiology*, *97*, 229–237.

Sobottka, K., & Pitas, I. (1996). Face localization and facial feature extraction based on shape and color information. *Proceedings of the IEEE International Conference on Image Processing*, *3*, 483–486.

Stanley, D. A., & Rubin, N. (2003). fMRI activation in response to illusory contours and salient regions in the human Lateral Occipital Complex. *Neuron*, *37*(2), 323–331.

Stanley, D. A., & Rubin, N. (2005). Functionally distinct sub-regions in the lateral occipital complex revealed by fMRI responses to abstract 2-dimensional shapes and familiar objects. *Journal of Vision*, *5*(8), 911.

Stevens, S. S. (1957). On the psychophysical law. *Psychological Review*, *64*(3), 153–181.

Suzuki, S. (2001). Attention-dependent brief adaptation to contour orientation: A high-level aftereffect for convexity? *Vision Research*, *41*(28), 3883–3902.

Suzuki, S. (2003). Attentional selection of overlapped shapes: A study using brief shape aftereffects. *Vision Research*, *43*(5), 549–561.

Suzuki, S., & Cavanagh, P. (1998). A shape-contrast effect for briefly presented stimuli. *Journal of Experimental Psychology: Human Perception and Performance*, *24*(5), 1315–1341.

Thompson, P., & Burr, D. (2009). Visual aftereffects. *Current Biology*, *19*, 11–14.

Vernon, R. J. W., Gouws, A. D., Lawrence, S. J. D., Wade, A. R., & Morland, A. B.

(2016). Multivariate Patterns in the Human Object-Processing Pathway Reveal a Shift from Retinotopic to Shape Curvature Representations in Lateral Occipital Areas, LO-1 and LO-2. *Journal of Neuroscience*, *36*(21), 5763–5774.

Wade, A., Augath, M., Logothetis, N., & Wandell, B. (2008). fMRI measurements of color in macaque and human. *Journal of Vision*, *8*(10), 1–19.

Wade, A. R., Brewer, A. A., Rieger, J. W., & Wandell, B. A. (2002). Functional measurements of human ventral occipital cortex : retinotopy and colour. *Philosophical Transactions of the Royal Society B: Biological Sciences*, *357*, 963–973.

Wandell, B. A., Dumoulin, S. O., & Brewer, A. A. (2007). Visual field maps in human cortex. *Neuron*, *56*(2), 366–83.

Ware, C., & Mitchell, D. E. (1974). The Spatial Selectivity of the Tilt Aftereffect. *Vision Research*, *14*, 735–737.

Wassermann, E. M. (1998). Risk and safety of repetitive transcranial magnetic stimulation: report and suggested guidelines from the International Workshop on the Safety of Repetitive Transcranial Magnetic Stimulation, June 5-7, 1996. *Electroencephalography and Clinical Neurophysiology*, *108*(1), 1–16.

Webster, M. A., & MacLin, O. H. (1996). Figural aftereffects in the perception of faces. *Psychonomic Bulletin & Review*, *6*(4), 647–653.

Wertheimer, M. (1923). Untersuchungen zur Lehre von der Gestalt. II. *Psychologische Forschung*, *4*(1), 301–350.

Wilkinson, F., James, T. W., Wilson, H. R., Gati, J. S., Menon, R. S., & Goodale, M. A. (2000). An fMRI study of the selective activation of human extrastriate form vision areas by radial and concentric gratings. *Current Biology*, *10*(22), 1455–8.

Wilkinson, F., Wilson, H. R., & Habak, C. (1998). Detection and recognition of radial frequency patterns. *Vision Research*, *38*(22), 3555–68.

Wilson, H. R., Loffler, G., & Wilkinson, F. (2002). Synthetic faces, face cubes, and the geometry of face space. *Vision Research*, *42*(27), 2909–2923.

Wilson, H. R., & Wilkinson, F. (2002). Symmetry perception: a novel approach for

biological shapes. *Vision Research*, *42*(5), 589–97.

Wilson, H. R., & Wilkinson, F. (2015). From orientations to objects : Configural processing in the ventral stream. *Journal of Vision*, *15*(7), 1–10.

Wilson, H. R., Wilkinson, F., Lin, L. M., & Castillo, M. (2000). Perception of head orientation. *Vision Research*, *40*(5), 459–72.

Winawer, J., Horiguchi, H., Sayres, R. A., Amano, K., & Wandell, B. A. (2010). Mapping hV4 and ventral occipital cortex: The venous eclipse. *Journal of Vision*, *10*(5), 1.

Winawer, J., & Witthoft, N. (2015). Human V4 and ventral occipital retinotopic maps. *Visual Neuroscience*, *32*, E020.

Zahn, C. T., & Roskies, R. Z. (1972). Fourier Descriptors for Plane Closed Curves. *IEEE Transactions on Computers*, *100*(3), 269–281.

Zanker, J. M., & Quenzer, T. (1999). How to tell circles from ellipses: Perceiving the regularity of simple shapes. *Naturwissenschaften*, *86*(10), 492–495.

Zeki, S. (1993). *A Vision of the Brain*. Oxford: Blackwell Scientific Publications.

Zhao, L., & Chubb, C. (2001). The size-tuning of the face-distortion after-effect. *Vision Research*, *41*(23), 2979–94.

Zimmer, M., & Kovács, G. (2011). Position specificity of adaptation-related face aftereffects. *Philosophical Transactions of the Royal Society of London. Series B, Biological Sciences*, *366*(1564), 586–95.