Feature-Specific Patterns of Attention and Functional Connectivity in Human Visual Cortex

Kirstie Holly Wailes-Newson

Doctor of Philosophy

University of York

Psychology

August 2019

Abstract

The ability to successfully allocate attention to a particular space or feature in the visual world is vital for successful day-to-day functioning. Attention refers to a narrowing of focus, with increased processing of an attended attribute at the expense of other non-attended dimensions. This attentional mechanism can modulate activity in the visual cortex and beyond. However, the full range of spatial scales at which attentional effects are evident in the visual cortex as a function of task is still relatively little understood. This thesis aimed to investigate the effects of attentional modulation across the visual cortex at several spatial scales, examining activation at the level of mean activity in individual regions-of-interest (ROIs), comparing patterns of voxel-level activity, and employing connectivity-style approaches to examine communication between multiple visual areas simultaneously.

In this thesis we examined how patterns of modulation across the visual cortex differed as a function of attentional task using a combination of visual psychophysics and functional magnetic resonance imaging (fMRI). Secondly, we assessed the generalisability of findings, initially collected with highly-controlled, low-level visual stimuli, to a similar experimental paradigm using an uncontrolled, dynamic and relatively naturalistic stimulus.

Novel findings and methods were demonstrated in this thesis.

1) We identified relatively little robust evidence of differing modulation as a function of attentional task at the univariate level, when examining mean activation in an individual visual ROI.
2) We do, however, note clear differences in patterns of attentional modulation as a function of task when comparing voxel-level patterns of activation in individual visual ROIs at the multivariate level.
3) Additionally, we assessed the communication between multiple visual regions simultaneously as a function of attentional task. Here, we identify significantly greater positive connectivity during passive viewing than during directed attention tasks. We suggest this robust finding is representative of a visual 'default mode' network, at a smaller spatial scale (millimetres) than traditional connectivity research (centimetres).
4) Finally, we find that the overall pattern of results collected with highly-controlled low-level visual stimuli, generalise relatively well to our experiment investigating attentional effects with an uncontrolled, dynamic and relatively naturalistic stimulus. In this more-naturalistic experiment, we identify little robust evidence of attentional modulation at both the univariate and multivariate levels. We also, however, demonstrate significant differences in patterns of connectivity across ROIs as a function of attentional task.

Table of Contents

List of Tables

List of Figures

specific attentional modulations averaged over all low-level visual feature modulation types. All visual ROIs show significantly lower signal modulation with blue-yellow stimulus chromaticies. All error bars are +/- 1 SEM and significance asterisks indicate Benjamini-Hochberg corrected values.

171

Figure 0.26 Feature-specific multivariate Support Vector Machine Decoding and Group-averaged voxel feature-specific weights as a function of eccentricity (6°) and polar angle (360°). Voxel-level responses with individual ROIs are modulated by attentional state. A) Overall three-way decoding accuracies. Attentional state can be decoded above chance in all ROIs except A1. Error bars reflect +/- 1 SEM. B) Two-way classification accuracies across pairwise combinations of attentional focus (orientation versus shape, orientation versus contrast and contrast versus shape). Voxel patterns in all areas differ significantly between attention to orientation and shape. Significance asterisks indicate Benjamini-Hochberg corrected values. C) Spatial backprojection analyses reveal no large-scale biases in voxel weights across location. The gray annulus reflects averaged voxel modulation at 1° intervals across 1.5-2.5° visual space. Deviations from circularity indicate positive (feature-specific) or negative (passive viewing) preferring clusters of voxels. An RFP stimulus overlay is provided for reference. Activation is thresholded at +/- 1.7 z-score ($p < .05$).

175

Figure 0.27 Colour-specific Multivariate Support Vector Machine Decoding and Group-averaged voxel feature-specific weights as a function of eccentricity (6°) and polar angle (360°). Voxel-level responses with individual ROIs are modulated by attentional state. A) Overall three-way decoding accuracies. Chromatic attentional state can be decoded above chance in all ROIs except LO-2 and A1. Error bars reflect +/- 1 SEM. B) Two-way classification accuracies across pairwise combinations of attentional focus (red-green versus blue-yellow, red-green versus achromatic and blue-yellow versus achromatic). Voxel patterns in all almost areas differ significantly between attention to red-green and achromatic-defined stimuli, with successful classifications also evident for red-green versus blue-yellow-defined stimuli and blue-yellow versus achromatic activation comparisons. Significance asterisks indicate Benjamini-Hochberg corrected values. C) Spatial backprojection analyses reveal no large-scale biases in voxel weights (chromatic preference) across location. The gray annulus reflects averaged voxel

List of Equations

Acknowledgements

I would first like to offer my thanks and gratitude to my supervisor, Professor Alex Wade. Throughout the past three years he has offered excellent guidance, enthusiasm and a lot of patience, ensuring my PhD has been highly enjoyable from start to finish. Thank you for your encouragement and the time you have taken to teach me things vital both for completion of my PhD and the future.

I would also like to thank the members of my TAP, Dr Daniel Baker and Dr Cade McCall as well as my Supervision Panel Chair Dr Rob Jenkins, for their very helpful advice and insight. You have all improved my confidence and helped to form interesting new avenues of analysis.

I would also like to thank Theo Karapanagiotidis, Dr Mark Hymers, Dr Andre Gouws, and Dr Richard Vernon for their assistance in helping me to understand complicated connectivity analyses, aligning particularly tricky brains and the creation of radial frequency patterns. Your help has most definitely helped to elevate the quality of my work and I am truly grateful.

I have had the privilege of meeting many great researchers during my PhD, ensuring these three years have been highly enjoyable. I would like to thank Marc Himmelberg, Mark Carey, Milena Kaestner, Stephanie Wake, Martin Scott and Anika Smith for your camaraderie and for helping to keep me working hard- I'm very lucky to have met you all.

Last, but not least, I would like to thank my family; my Mum and Jodie, Kyle, Zoe and Shaun for their help and support, even though you don't always understand what I do. Mum, thank you for your love, encouragement and inspiration throughout my (very long) education- I wouldn't have been able to complete this PhD (and finally get an 'adult' job!) without you. Finally, I would like to thank my husband Mladen Sormaz for your constant support, love and advice, your continual entertainment and for always being there. I am very lucky to have you by my side.

I would like to dedicate this thesis to my Dad, Gary Newson, who sadly is no longer here to see me complete my 21 years of education. Thank you so much for your hard work and dedication, which allowed me to attend university and pursue my passion for learning. Thank you for your endless love and support. You have taught me the importance of strength and bravery and I am so proud to call you and Mum my parents. I hope I have made you proud too.

Declaration

I declare that the work presented in this thesis is original and I am the sole author. This work was carried out under the supervision of Professor Alex Wade. This work has not been submitted to this or any other University for a degree. All sources are acknowledged as References.

fMRI data presented in Chapter 6 were collected with MSc project students (Bethany Adams and Rachel Gordon), who I co-supervised on the project. All of the reported data analyses, experimental design, creation of the stimuli and all other data collection were performed by myself, under supervision of Professor Alex Wade.

Data in Chapters 3 and 4 were presented in a preliminary, incomplete form at the *Vision: A Platform for Linking Circuits, Behaviour & Perception*, summer school at Cold Spring Harbor Laboratory, New York, USA (June 2017).

Data presented in Chapter 4 are included in:
Wailes-Newson, K.H. & Wade A.R. (2019). Feature-specific patterns of attention and functional connectivity in human visual cortex. *Neuroimage* (submitted).

Data presented in Chapter 5 were presented in a preliminary, incomplete form at the *Computational Neuroscience: Vision* summer school at Cold Spring Harbor Laboratory, New York, USA (July 2018).

1. Introduction

1.1 Thesis overview

Humans must constantly navigate through a complex and dynamic visual environment, yet, we possess limited resources to process that environment. Attention is a mechanism which allows us to direct these limited resources to a particular spatial location, or to a particular feature in our visual environment. This effectively reduces the amount of information we must process, at the expense of the remaining visual scene. Attention as such is a vital gain control mechanism, allowing humans to efficiently and effectively process a wealth of visual information.

The experiments in this thesis form two clear, yet interacting, strands, examining attentional modulation effects with highly controlled visual stimuli, and assessing the generalisability of these findings through research conducted with complex, uncontrolled naturalistic stimuli.

Firstly, we use psychophysical testing and functional magnetic resonance imaging (fMRI) to investigate changing patterns of activation in the visual cortex as a function of switching attentional task between low-level visual features. We examine the influence of attentional modulation at three cortical spatial scales; the region-of-interest (ROI), the voxel and the network, to identify the spatial extent of attentional modulation. These experiments are useful for furthering our understanding of the relatively little-understood signatures of featural attention at multiple spatial scales.

Second, we examine signatures of attentional modulation in response to both different low-level visual features and stimulus chromaticities, investigating the distinct representations of chromatic information in the visual cortex. Again, we examine patterns of activation at multiple spatial scales, to examine patterns of feature-specific attentional modulation and identify interactions between stimulus feature and stimulus chromaticity across the visual cortex.

Finally, we direct attentional focus towards a highly-complex, dynamic, naturalistic stimulus, to assess the extent our findings from experimentation with highly-controlled low-level visual stimuli generalise to real-world visual scenarios, and to examine the signatures of activation when attention is directed toward a relatively higher-level visual stimulus (faces). Hence, we investigate similarities in attentional modulation patterns elicited both with highly-controlled, low-level and uncontrolled, naturalistic stimuli.

We use a range of methods including visual psychophysics, functional magnetic resonance imaging, machine learning classification and connectivity analyses to help answer the following questions:

1) Is it possible to identify clear patterns of modulation when switching attention between different visual features?
2) At which spatial scale(s) do we see clear signatures of feature-specific attentional modulation?

3) To what extent do findings from highly-controlled low-level visual experiments generalise to research conducted with uncontrolled, complex relatively naturalistic stimuli?

The aim of this chapter is to introduce the key ideas related to the experiments conducted in this thesis. This includes a brief overview of the organisation of the visual cortex and the processing of low-level visual features. We also include a summary of the attentional literature and outline research regarding functional connectivity analyses and the use of naturalistic stimuli.

1.2 Human Visual System

1.2.1 Pre-cortical visual system

Visual processing originates in the retina, a layer of rod and cone photoreceptors at the back of the eye. Rod photoreceptors typically respond in low light levels, whereas cone photoreceptors (cones) respond in well-lit conditions to light of different wavelengths and are critical for the perception of colour (for a review, see Conway, 2009). Whilst many retinal ganglion cell (RGC) types exist, most common are the midget and parasol cells. Midget cells cluster around the fovea and have small, dense dendritic fields whereas parasol cells, possess larger, less dense dendritic fields. These cell types signal the beginning of segregation of visual information, forming the parvocellular (PC) and magnocellular (MC) pathways respectively.

The MC, PC and koniocellular (KC) pathways receive input from weighted combinations of these cone photoreceptors. The magnocellular pathway receives input from long- and middle- wavelength sensitive cones, possessing sensitivity to achromatic contrast, high temporal- and low spatial-frequencies. The parvocellular pathway compares responses from long- and middle-wavelength sensitive photoreceptors, providing a linear contrast response, with slow temporal- but high spatial-resolution. The KC pathway receives input from short-wavelength cones, which are sparsely arranged in the retina in comparison to long- and middle- wavelength photoreceptors, hence the pathway has a lower spatial-resolution than it's fellow pathways (Roorda & Williams, 1999).

Photoreceptors then send this information, regarding external light conditions, through their respective pathways backwards into the visual cortex through a series of retinal cells, which form the optic nerve. Information from the retina to the visual cortex undergoes a re-routing in the optic chiasm. Nasal (nearest the nose) responses from each retina cross to the contralateral (opposite) hemisphere, whist temporal (nearest the ear) responses remain on the ipsilateral (same) side. This information then enters the lateral geniculate nuclei of the thalamus (LGN) (see Figure 1.1).

The LGN are bilateral structures in the brain which receive input from each eye and contain six distinct layers. Axons of different cells types in the retina terminate in discrete layers of the LGN. Axons of midget ganglion cells terminate in the upper four (parvocellular) layers. Whereas parasol ganglion

cell axons terminate in the lower two (magnocellular) layers (Polyak, 1941). Axons from small bistratified retinal ganglion cells predominantly terminate in regions intercalated between the magnocellular and parvocellular layers, termed the koniocellular layers (see Casagrande, 1994 for a review).

1.2.2 Primary visual cortex (V1)

The information in the LGN then projects to the primary visual cortex (V1), in the calcarine sulcus through the optic radiation. V1 is composed of 6 layers, distinguished by the variation in density of neurons, axons and synapses (Wandell, 1995). Information from the LGN is transported to distinct layers of the primary visual cortex. Information from the magnocellular pathway projects to layer 4Cα and lower layer 6. Responses in the parvocellular pathway instead terminate in layer 4Cβ and the upper layer 6. Whilst less is known about the projection of information in the koniocellular pathway, it is thought projections exist from the LGN to layers 1, 3B, 4A and cytochrome oxidase (CO) blobs in the upper, superficial layers of V1 (Chatterjee & Callaway, 2003; Hendry & Reid, 2000) (see Figure 1.1). Information is then transmitted from V1 to a multitude of further visual areas, for example, inactivation of V1 activity results in a lack of response in areas V2 and V3 (Girard & Bullier, 1989; Girard, Salin, & Bullier, 1991; Schiller & Malpeli, 1977). It has been suggested these feedforward connections to higher visual areas are transmitted from the upper layers of V1, and feedback to V1 is received in its deeper layers (Rockland & Pandya, 1979).

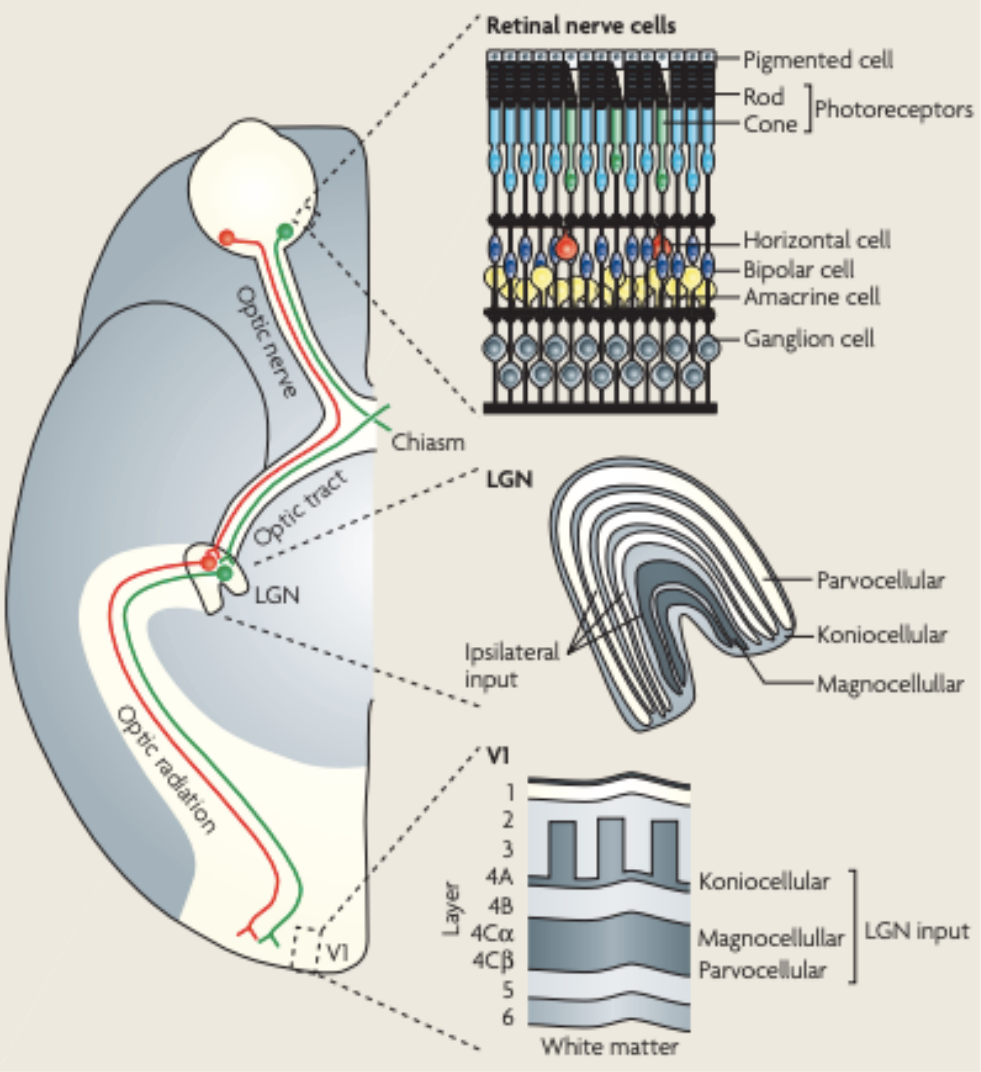

Figure 1.1 The transmission of information within the human early visual pathway. Visual information enters the eye and passes through a series of retinal nerve cells before projection to the LGN. Within the LGN, information is split into segregated layers, corresponding to the PC, MC and KC pathways. This segregated information then projects to distinct layers within V1 (Figure from Solomon & Lennie, 2007).

1.2.3 Extra-striate visual areas

Visual information travels from V1 to many areas of the visual hierarchy. For example, the motion-responsive area MT+, receives direct input from V1 (Zeki, 2004). As a general rule of thumb, the further you travel up the visual hierarchy, the more functionally-specialised visual regions become. Higher

visual areas are increasingly focused on the processing of complex visual properties such as faces and motion, and relatively detached from the processing of low-level visual features, typically reserved for early visual areas (for a review, see Grill-Spector & Malach, 2004). The extent to which we see clear functional specialisation of regions in the visual cortex is debated, with the theory of functional specialisation first suggested by Livingstone & Hubel, (1988) and Zeki, (2004; 1973). However, others suggest instead a theory of multifunctionality, where neurons in the visual cortex can encode information about multiple visual features, not possessing distinct feature-specific properties (Schiller, 1996). In short, it is likely earlier visual areas possess little strict functional segregation, instead, responsive to a multitude of low-level visual features, whereas relatively higher visual areas such as motion area MT+ and the fusiform face area (FFA) do exhibit distinct and specific feature-specific preferences (for reviews, see Born & Bradley, 2005; Kanwisher & Yovel, 2006).

1.2.4 Retinotopic organisation

Retinotopy refers to the mapping of visual input from the retina to neurons (or voxels) in the visual system. Neurons in the primary visual cortex are retinotopically organised; the centre of the visual field is represented in the fovea and the periphery is located more anteriorly in the calcarine sulcus (Holmes, 1918). The organisation of information in primary visual cortex is also dependent upon hemisphere. The left hemisphere V1 receives information from the right visual hemifield, and vice versa. Additionally, neurons located dorsally in V1 receive information from the lower quadrant of

their respective visual hemifields, whilst neurons in the ventral aspects of V1 represent the upper quadrants of the visual field. In summary, the representation of the visual field is reversed and flipped to form a retinotopic map of visual space in V1. The ability to map this organisation through fMRI was first identified by Engel et al., (1994). It is now possible to map many areas of the visual cortex retinotopically on the basis of characteristic visual field representations (see Wandell, Dumoulin, & Brewer, 2007) (see Figure 1.2).

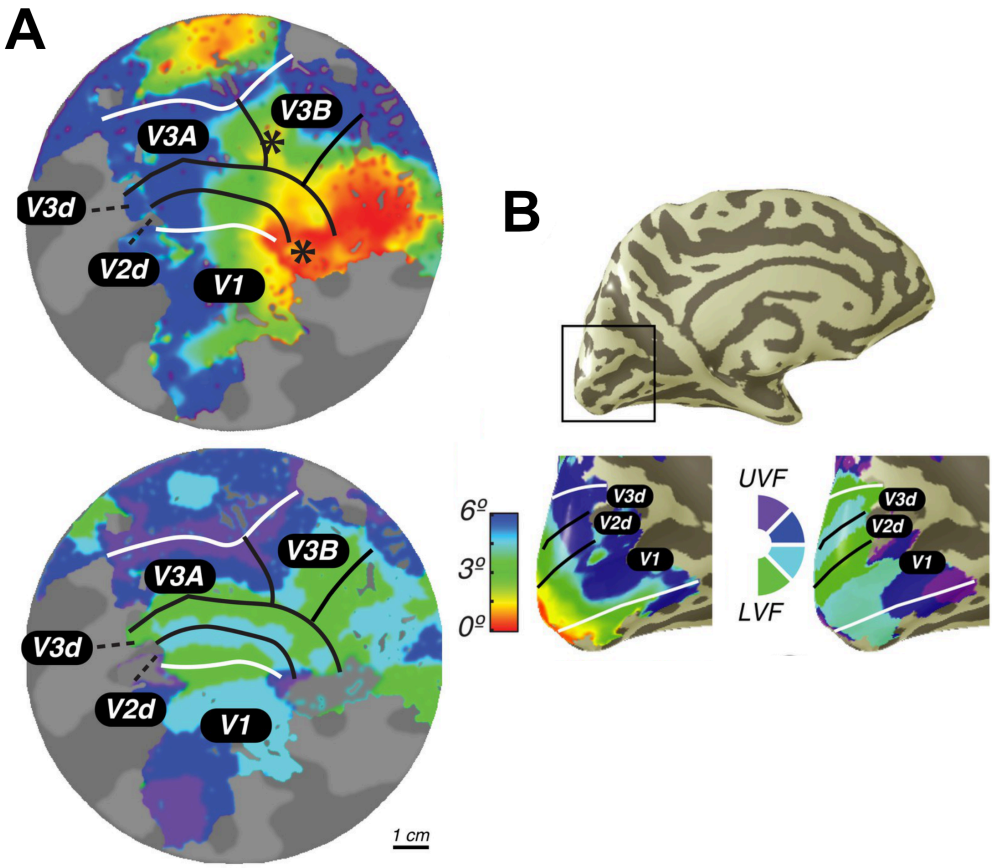

Figure 1.2 Retinotopic mapping of early visual cortex. Regions of interest (ROIs) are defined on the basis of their eccentricity (top left) and polar angle (bottom left) on flat maps (A) and the inflated cortical surface (B-left eccentricity, right polar angle) (Figure from Mackey, Winawer, & Curtis, 2017).

1.2.5 Perception of visual features

1.2.5.1 Perception of orientation

Stimulus orientation is encoded in a distributed manner across the visual cortex, with visual regions possessing subpopulations of orientation-preferring neurons. However, the extent of segregation of orientation-encoding throughout the visual hierarchy is debated. Horton & Hubel, (1981) identified cytochrome oxidase (CO) blobs and interblobs in primate V1, so called due to their visibility when V1 is stained with cytochrome oxidase, an enzyme involved in metabolism. Blobs were considered to possess weak spatial frequency and orientation tuning but respond well to isoluminant colour, whereas interblobs were highly orientation selective (Livingstone & Hubel, 1984).

Some researchers think this segregation of colour and orientation in V1 blobs and interblobs may persist further into the visual cortex, with projections into thin and interstripe regions of V2 respectively (Levitt, Kiper, & Movshon, 1994; Livingstone & Hubel, 1984). Researchers have also suggested the segregation of colour and orientation-encoding in particular visual regions, such as a non-existence of colour-selective cells in V3, V3A/B or MT+ in rhesus monkeys (for a review, see Zeki, 1978). Lesioning of area V4 has also been noted to heavily impact colour perception with little disruption to form (orientation) processing (Meadows, 1974; Zeki, 1990).

However, others question the extent to which orientation preference is independent of other low-level visual features. For example, Conway, (2001)

and Johnson, Hawken, & Shapley, (2001) provided the first evidence for orientation tuning in colour-coding V1 cells. In early visual areas V1 and V2, Engel (2005) used adaptation to demonstrate the presence of neurons tuned jointly to both colour and orientation. Single unit electrophysiology also supports this idea: Economides, Sincich, Adams, & Horton, (2011) identified only a subtle difference in orientation-selectivity between V1 blob and interblob cells, along with many previous researchers, suggesting the segregation of orientation processing in the visual cortex may not be as strict as once considered. Evidence for the existence of orientation-sensitive responses has now been identified in visual regions across the cortex, from V1, both dorsally and ventrally to higher visual areas such as LO-1 (e.g. Conway, Moeller, & Tsao, 2007; Silson et al., 2013; Sumner, Anderson, Sylvester, Haynes, & Rees, 2008).

1.2.5.2 Perception of contrast

The human visual system is highly sensitive to contrast (differences in illumination). Retinal photoreceptors possess a dynamic range with a maximum of two orders of magnitude, yet across the visual system, we are able to scale responses efficiently to process information across more than ten orders of magnitude (Ohzawa, Sclar, & Freeman, 1985).

Contrast gain control refers to the ability of the visual system to centre the limited response range of a cell around a mean level of contrast (Ohzawa et al., 1985). It begins in the retina (Shapley & Enroth-Cugell, 1984) and is strengthened along the visual hierarchy (Sclar, Maunsell, & Lennie, 1990). It

regulates the gain and integration time of the visual system on the basis of the locally prevalent contrast and the luminance of a stimulus versus the mean luminance present in a visual scene. Both gain and integration time are reduced when contrast is high, and are increased when contrast is low (Carandini et al., 2005). Contrast gain control mechanisms reduce the impact of large changes in mean contrast such as occur during eye movements.

For example, Gardner et al., (2005) provided the first evidence of a contrast gain control mechanism in early visual cortex (V1, V2 and V3) which adapts to the mean contrast level, using event-related fMRI. Neurons in V1, V2 and V3 appeared to shift their contrast response functions to centre their maximum response at the adapting contrast level (contrast gain). However, hV4 responses indicated sensitivity to the salience of contrast change rather than a providing a pure contrast representation, responding positively to both increments and decrements in contrast. This evidence suggested an ability of the visual system both to discount, slow, uninformative changes in contrast with adaptation, and to remain highly-sensitive to contrast changes which may signal important events in the environment. In support, Wang & Wade, (2011) identified evidence of attentional modulation to achromatic contrast in visual areas as early as the primary visual cortex (V1).

Hence, the processing of contrast information across the visual cortex, through contrast gain control mechanisms is a robust and well-understood phenomenon. There is a wealth of clear evidence for its existence across the entire visual cortex.

1.2.5.3 Perception of shape

As is the case with the processing of orientation and contrast, there is evidence of shape processing, in various forms, across the visual cortex. For example, Dumoulin & Hess, (2007) noted clear responses to concentric stimuli across many early visual areas (V1 to hV4), rather than shape-specific responses present solely in an isolated visual region. In support, Tschechne & Neumann, (2014) produced a recurrent computational network for the encoding of shape information, and suggested a distributed representation of shape in the visual cortex, with lower-level representations in early visual areas (e.g. V1, V2, V3) projecting to intermediate and higher levels of shape processing (e.g. hV4, and inferotemporal (IT) cortex).

Despite clear evidence for multiplexed feature processing in early visual areas, there is also clear evidence for particular regions of the visual cortex possessing distinct specialisations for shape processing. For example, much research suggests an involvement of hV4 in the encoding of shape, as oppose to purely chromatic information. hV4 has frequently been noted as an area responsive to concentric shapes (e.g. Dumoulin & Hess, 2007; Gallant, Braun, & Van Essen, 1993; Gallant, Shoup, & Mazer, 2000; Wilkinson et al., 2000), and the curvature of stimuli (first identified by Pasupathy & Connor, (2001). hV4 is considered an intermediate shape processing region, encoding shape information more complex than orientation, but at a more basic level than the processing of meaningful objects (for a review, see Loffler, 2008). In support, research has demonstrated that lesioning V4 can

profoundly negatively impact shape discrimination in primates (Merigan, 1996; Merigan & Pham, 1998; Schiller, 1995).

Additionally, the lateral occipital complex (LO) has been identified as a region of functionally segregated shape recognition (Denys et al., 2004; Grill-Spector et al., 1999; Kourtzi & Kanwisher, 2001; Malach et al., 1995). Visual object agnosia, a condition in which individuals fail to recognise the identity of visually-presented objects, is associated with LO damage (Ptak, Lazeyras, Di Pietro, Schnider, & Simon, 2014). Larsson & Heeger, (2006) identified two retinotopically-distinct regions LO-1 and LO-2, which overlap with posterior regions of LO, proposed to integrate abstract shape information from lower visual areas (Vernon, Gouws, Lawrence, Wade, & Morland, 2016). However, other research has suggested a distinct difference in the visual information LO-1 and LO-2 selectively encode. Silson et al., (2013) used transcranial magnetic stimulation (TMS) and identified a double dissociation, with LO-2 impaired in shape processing, but not orientation processing during the application of TMS to temporarily disrupt activity in the region. Conversely, TMS applied to area LO-1 disrupted orientation processing, but not shape. This suggests a specialised role for LO-2 in the processing of shape information in the human visual cortex.

1.2.5.4 Perception of colour

The perception of colour begins with the cone photoreceptors present in the retina. Retinal photoreceptors are classified into three distinct types on the basis of their sensitivity to wavelengths of light. Long-wavelength cones (L)

have a peak sensitivity at approximately 560nm, medium-wavelength cones (M) have maximum light absorption peaks at ~530nm, and around 420nm for short-wavelength (S) cones (Bowmaker & Dartnall, 1980; Schnapf, Kraft, & Baylor, 1987) (see Figure 1.3).

More recently, researchers have identified a third type of light-sensitive cell, in addition to rod and cone photoreceptors, which contains a pigment called melanopsin. These intrinsically photosensitive retinal ganglion cells (ipRGCs) were initially thought to be involved solely in the regulation of non-visual responses to photic stimuli, such as the pupillary light reflex and regulation of sleep. Yet, recent research has also demonstrated this pigment is light-sensitive. For example, Panda et al., (2005), produced light-sensitive frog eggs when injected with the genes for melanopsin, and Melyan, Tarttelin, Bellingham, Lucas, & Hankins, (2005) modified embryonic mouse neurons to produce melanopsin, which then demonstrated light-sensitive responses. Human melanopsin has been demonstrated to possess a spectral sensitivity to blue light (peaking at 479nm) (Bailes & Lucas, 2013).

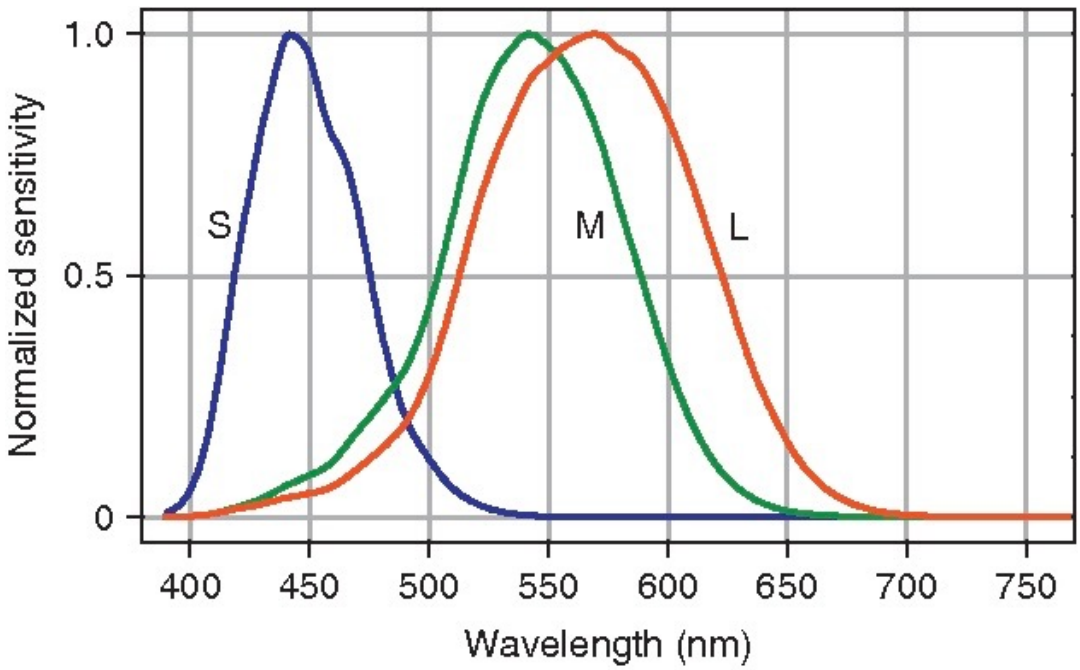

Figure 1.3 Rod and cone photoreceptor sensitivities. Retinal photoreceptors have peak sensitivities at different wavelengths of light, and comparisons of these signals provide the initial stage of colour processing within the visual system (Figure from Foster, 2010).

Chromatic information is extracted from these retinal photoreceptors and supplied to V1. Each cone photoreceptor connects to a series of retinal cells, which process information in their receptive field in a centre-surround organisation (chromatic-opponency). Midget bipolar retinal ganglion cells receive input in the centre of their receptive field from few cone photoreceptors, and input from horizontal cells in their surround, containing information about nearby cone photoreceptors (Dacey, 1996; Dacey et al., 2000). This centre-surround organisation allows for comparison of cone photoreceptor activation between the centre and surround regions of the receptive field. This comparison of L- and M-cone activation (L-M) is transmitted to the LGN and V1 via the parvocellular pathway (Sumner et al., 2008). However, information from the S-cones remains segregated, and

projects from small bistratified retinal cells to the koniocellular layers of the LGN and then to distinct layers 3B and 4A of the primary visual cortex. This pathway provides colour-opponent signals, comparing S-cone activation with L- and M-cone responses (S-(L+M)) (Casagrande, 1994; Chatterjee & Callaway, 2003).

The comparison of inputs to centre-surround receptive fields provides three colour-opponent pathways, which form the basis for the processing of colour across the visual cortex. A luminance (L+M) pathway which encodes the sum of L- and M-cone responses, is insensitive to wavelength (colour) and responds to achromatic contrast. A 'red-green' (L-M) pathway performs a chromatic comparison of signals from L- and M-sensitive cones, encoding the red-green dimension of colour space, and a 'blue-yellow' sensitive pathway (S-(L+M)), which is driven predominantly by S-cones.

The encoding of colour is considered to be distributed across the visual cortex (Seymour, Williams, & Rich, 2016). However, previous research originally identified area hV4 as a 'colour centre', based on the predominance of colour-sensitive cells detected (first proposed by Zeki, 1969, 1973) . Research also identified mm-scale colour-sensitive 'globs' in hV4, thought to reflect the specialisation of CO blobs in V1 (Conway et al., 2007). This was supported by research on cerebral achromatopsia, a loss of colour perception, which resulted from lesions in regions considered to be homologous to macaque V4 (Meadows, 1974). However, more recent research has moved away from this complete functional specialisation of

area hV4, after identifying that lesions in macaque V4 did not significantly negatively impact colour perception, suggesting other regions across the visual cortex must contribute to complete colour perception (Heywood, Gadotti, & Cowey, 1992).

1.2.5.5 Perception of faces

Despite the distributed processing of many low-level visual features across the visual cortex, the perception of relatively higher-level, complex visual stimuli (formed from combinations of low-level stimulus attributes such as orientation, colour, contrast and shape) tends to be processed in a distinct, and segregated manner in the visual cortex. For example, the fusiform face area (FFA) is considered a specialised region dedicated to the perception of faces, first identified by Kanwisher, McDermott, & Chun, (1997), as a region with significantly greater response to face stimuli than to a variety of comparison stimuli (such as scrambled faces, houses and hands). However, others argue that instead of reflecting a face-specific mechanism, the FFA provides evidence of a region responsive to fine-grain discriminations (e.g. Gauthier, Behrmann, & Tarr, 1999). In respect of this, the expertise hypothesis argues face-specific regions such as the FFA are not necessarily specialised for the processing of faces, but instead are activated when distinguishing category exemplars from one another which share a same basic configuration (Diamond & Carey, 1986).

Despite the debate regarding the precise functional segregation of the FFA, there is overwhelming evidence to support its specific response to face

stimuli (see Kanwisher & Yovel, 2006 for a review). Two additional face-responsive areas have also been identified in the visual cortex, the occipital face area (OFA) and the superior temporal sulcus (STS). The OFA was first identified by Gauthier et al., (2000), with greater responsivity to faces versus objects than identified in the FFA. The STS was originally identified as a region in macaque cortex responsive to faces, parts of faces and facial expressions (Heywood & Cowey, 1992). More recently, the STS has also been identified as responsive to dynamic faces and bodies (Allison, Puce, & McCarthy, 2000). Currently, the neuronal mechanisms behind face-specific responses are not fully understood, however, evidence clearly suggests the existence of specialised regions for the processing of complex, higher-level face stimuli in the visual cortex.

1.3 Visual Attention

1.3.1 What is visual attention?

The importance of attention as a mechanism for successful interaction with our visual world has been understood from very early in the development of Psychology as a discipline. William James reported attention as a *'[taking]' one out of what seems several simultaneously possible objects or trains of thought'* and *'a withdrawal from some things in order to deal effectively with others' (James, 1890).* Without attentional mechanisms, we would become overwhelmed by the sheer amount of information present in our visual environment. Hence, visual attention, a mechanism to weight the current important of specific aspects of a stimulus, is vital.

The primary function of attention is to optimise task performance, such as the detection of a salient stimulus in a cluttered visual environment (e.g. Serences, Saproo, Scolari, Ho, & Muftuler, 2009). There are many mechanisms controlling attentional modulation in the visual cortex, with their exact processes are typically little understood. However, attention mechanisms are generally considered to alter the activity of neurons most informative for a particular visual task. This typically results in an enhancement of processing of task-relevant information in the visual cortex and a suppression of distracting information providing greater efficiency of processing (e.g. Martinez-Trujillo & Treue, 2004; Serences et al., 2009).

More precisely, attention is a form of gain control; allowing the visual system to efficiently process information over many orders of magnitude in a limited processing range (Carandini & Heeger, 2011). Attention reshapes the distribution of activity across populations of neurons, balancing relative levels of stimulus-specific excitation and suppression (Reynolds & Heeger, 2009). This response-reshaping is thought to be dependent on the stimulus and size of attentional field. Attention directed to a relatively small region of a large visual stimulus will elicit response gain changes in activation (multiplicative changes in the output level) , whereas a small stimulus with a relatively large region of attentional focus will produce contrast gain effects (changes in the apparent input level) (Herrmann, Montaser-Kouhsari, Carrasco, & Heeger, 2010; Reynolds & Heeger, 2009) (see Figure 1.4). The effects of attentional modulation are also dependent upon the stimulus attended. Attention is a complex mechanism which does not simply up-regulate the activity of

neurons responsive to a stimulus of attentional focus. Instead, attention optimises the gain of visual cortical neurons in a flexible and adaptive manner to enable efficient interaction with task-relevant aspects of a visual stimulus (Carrasco, 2011).

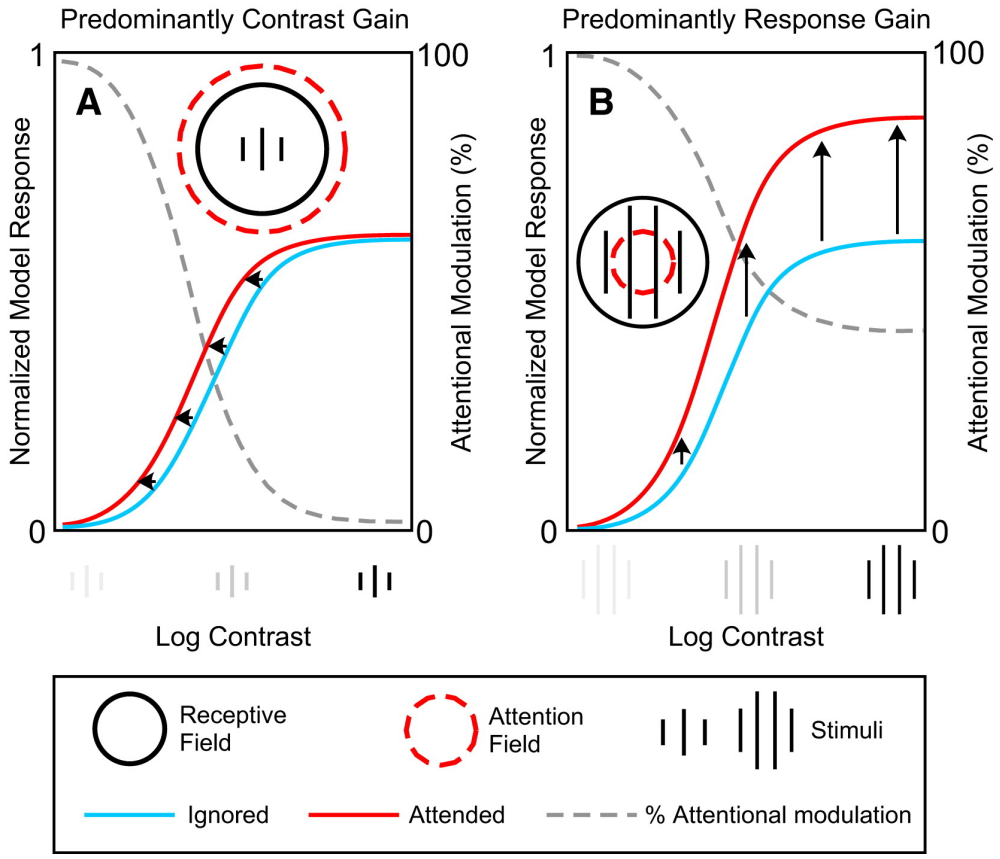

Figure 1.4 The Normalisation Model of Attention. The interaction between stimulus size and the size of the attentional field alter the form of attentional modulation. A smaller stimulus with a large attentional focus will produce contrast gain effects (left), whereas a relatively smaller attentional focus with a larger stimulus elicits predominantly response gain effects (right) (Figure from Reynolds & Heeger, 2009).

1.3.2 Spatial attention

Attention can be directed towards many aspects of a visual stimulus. In past research, spatial attention was most commonly studied. This is typically

referred to as a ‘spotlight’ of focus we apply to a visual scene (Posner, Snyder, & Davidson, 1980). Spatial attention refers to directing focus towards a particular location in visual space. A multitude of techniques have been used to demonstrate spatial attention can modulate activity in the visual cortex. For example, Tootell et al., (1998) used fMRI to identify increases in response at cortical representations of an attended visual target across multiple low-level visual areas. Similar research has identified that when two stimuli are presented simultaneously in a cells receptive field (measured in macaque V4), the neuronal response is strongly modulated by which of the two stimuli was attended (Moran & Desimone, 1985).

Additional research has also demonstrated that attention to a particular spatial location in a luminance-modulated grating produces increases in the amplitude of responses in regions of the visual cortex representing the attended location (Lauritzen, Ales, & Wade, 2010; Verghese, Kim, & Wade, 2012). Interestingly, Sumner, Tsai, Yu, & Nachev, (2006) demonstrated the robust nature of spatial attention, which can modulate activity even in response to invisible stimuli. Attention was noted to increase the perceptual strength of an invisible spatial prime, enhancing its visibility, and also independently boosted unconscious sensorimotor processes initiated by this invisible priming stimulus. Hence, spatial attention is a highly important and powerful mechanism which can produce clear modulation of activity in the visual cortex as we interact with our visual environment.

1.3.3 Featural Attention

Alternatively, attention can be directed towards a particular stimulus feature. This has been shown to alter activity in neurons encoding the attended feature. The influence of featural attention has been noted across the visual hierarchy (e.g. Kamitani & Tong, 2006). For example, fMRI BOLD responses in MT+ increase during attention toward a visual motion stimulus (e.g. O'Craven, Rosen, Kwong, Treisman, & Savoy, 1997). Additionally, increased hV4 activation has been identified in response to attention directed towards chromatic stimuli (Chawla, Rees, & Friston, 1999; Schoenfeld et al., 2007). Evidence of feature-specific attentional modulation in regions responsible for the processing of that feature is also found for higher-level, complex stimuli. For example, attention to faces and places modulates responses in the FFA and parahippocampal place area (PPA) respectively (O'Craven, Downing, & Kanwisher, 1999).

Unlike spatial attention, where focus is restricted to a single spatial location in the visual field, featural attention is global in scope. For example, attention to the colour or motion of a stimulus has been demonstrated to increase responses in the visual cortex representing all visual field locations the attended feature is present (Saenz, Buracas, & Boynton, 2002). Evidence has demonstrated the global focus of featural attention persists in early visual cortical areas, even when a restriction of attention is critical for successful task performance (Andersen, Hillyard, & Müller, 2013). This is typified in the feature-similarity gain model of Martinez-Trujillo & Treue, (2004), which suggests featural attention increases the gain of neurons

tuned to the attended feature, and decreases the response of neurons tuned away from this feature at the population level, across the visual field.

Featural attention is a robust phenomenon. Feature-specific activation in the visual cortex has been noted even in the absence of direct visual stimulation. For example, Serences & Boynton, (2007) identified activation to the attended visual feature even in regions of the visual scene which did not contain a stimulus. Xing, Ledgeway, McGraw, & Schluppeck, (2013), also demonstrated an ability to decode the contrast of a perceived stimulus from activity in early retinotopic visual areas when the stimulus had to be held in working memory (i.e. during periods when no stimulus was present). Additionally, they demonstrated the generalisation of classification from a model trained on a perceived stimulus to a purely remembered stimulus and vice versa, demonstrating the highly consistent nature of feature-specific patterns of response in early visual cortex.

Additionally, the importance of studying the contribution of feature-based attention on visual cortex activation has been highlighted by Huk, Ress, & Heeger, (2001). They demonstrated that featural attention accounted for the reported increases of MT+ activation with the perception of a motion aftereffect. Hence, consideration of the contribution of attentional effects in visual research is vital.

1.3.4 Classification of attentional state

Whilst it is widely accepted attention modulates activity in the visual cortex in areas responsive to the attended location or feature, the scale of these modulatory effects is less clearly understood. Much recent research has demonstrated, using multivariate (voxel-level) analysis techniques, that it is possible to identify the visual feature attended. For example, Kamitani & Tong, (2005) used fMRI to demonstrate it is possible to decode which of eight stimulus orientations a participant was attending to, from the pattern of activation across voxels in early visual cortical regions (V1-hV4 independently). Kamitani & Tong, (2006) also demonstrated above-chance decoding accuracies for motion direction, even during presentation of an ambiguous motion stimulus.

Since these initial findings, many other researchers have demonstrated similar classification abilities. Mannion, McDonald, & Clifford, (2009) exhibited successful classification of orientation (clockwise or anticlockwise) in a glass pattern stimulus balanced in its radial components. Brouwer & Heeger, (2009) additionally provided evidence for the decoding of stimulus colour across early visual areas. Interestingly, Sumner et al., (2008) demonstrated successful classification of colour-orientation interactions; in visual areas V1, V2 and V3, it was possible to decode luminance-, red-green- and blue-yellow-defined orientations. Furthermore, Harrison & Tong, (2009) provided evidence for the ability to predict which of two oriented gratings was held in working memory in early visual areas (V1-hV4),

demonstrating the clear existence of feature-specific patterns of attentional modulation at the multivariate level.

However, there exists debate regarding the driving forces behind these successful classification analyses. For example, Freeman, Brouwer, Heeger, & Merriam, (2011) argued that the successful classification of attended orientation was driven by large-scale coarse topographical biases in the mapping of orientation columns in V1. This viewpoint regarding large-scale biases in the underlying mapping of visual cortical regions as sufficient for classification of an attended feature has clear support in the literature (Beckett, Peirce, Sanchez-Panchuelo, Francis, & Schluppeck, 2012; Clifford, Mannion, & McDonald, 2009; Raemaekers, Lankheet, Moorman, Kourtzi, & van Wezel, 2009; Sasaki et al., 2006). In support, Op de Beeck, (2010) modelled decoding accuracies with application of various levels of spatial smoothing, to assess the impact of smoothing on the ability to detect small-scale functional organisation at the voxel level. They identified no negative impact of smoothing on the sensitivity of multivariate analyses and proposed this evidence contradicts the idea classification accuracies are reflective of sub-voxel inhomogeneities in feature preference.

However, other researchers believe successful classification is unlikely to be solely driven by coarse-scale biases in mapping amongst visual areas. Instead, they suggest important information is present in the fine-grain (voxel-level) patterns of fMRI activity. It is assumed each voxel reflects responses from many visual neurons, and sampling scale means that voxels

will sum responses of neurons with differing proportions of preference for a particular visual feature, and that these inhomogeneities in feature-preference across voxels are what contribute to above-chance classification accuracies (e.g. Kay, Naselaris, Prenger, & Gallant, 2008). For example, Mannion et al., (2009) demonstrated successful decoding of orientation in stimuli with no dominant radial component, indicating radial bias is not critical for successful decoding of attended orientation. Additionally, whilst Kamitani & Tong, (2005) and Sumner et al., (2008) did identify evidence of a weak radial bias in their responses, they argued this did not have a substantial contribution to discrimination accuracy. Finally, Kamitani & Sawahata, (2010) noted from their own simulations no basis for ruling out the possibility classification is driven to some extent by small-scale inhomogeneities across voxels. Hence, the ability to classify feature-specific attentional modulation effects at the voxel level is likely driven by both sub-voxel inhomogeneities in the organisation of feature-preference and large-scale biases in topographical mapping in visual regions.

1.4 Connectivity

1.4.1 Default mode network

Interest in cortical connectivity – the way that information is passed between cortical areas - has increased in recent years. These connectivity analyses have revealed key networks, consisting of distinct sets of cortical regions, which exhibit synchronised activity during particular tasks. The earliest such network, the so-called ‘default mode network’ (DMN) was first formally identified by Raichle et al., (2001), as a collection of regions in the medial

and lateral parietal, medial prefrontal and medial and lateral temporal cortices, who consistently decrease their activity during attentionally-demanding, complex tasks (Raichle, 2015). Typically, the network of regions forming the DMN demonstrate high metabolic activity and synchronised activation at rest, or during tasks requiring little attentional effort (for a review, see Buckner, Andrews-Hanna, & Schacter, 2008). This synchronisation of activity across a large array of cortical areas, is thought to support emotional processing, self-referential mental activity and the recollection of prior experiences, commonly termed ‘mind-wandering’ (see Raichle, 2015).

1.4.2 Task-based networks

However, there exist multiple other functionally- and anatomically- distinct networks in the cortex. For example, the dorsal attention network (DAN) is formed from regions of the intraparietal sulcus (IPS), superior parietal lobule and dorsal frontal cortex near to the frontal eye fields (see Corbetta, Patel, & Shulman, 2008). The DAN is thought to activate when attention is directed to the external world. Additionally, the ventral attention network (VAN) is considered to be predominantly right lateralised across regions of the inferior frontal gyrus, anterior insula and adjacent frontal operculum, and is suggested to automatically re-orient attention towards salient perceptual stimuli (Corbetta et al., 2008). There appears a reciprocal relationship between the activity of the default mode network and many task-based networks. For example, when the DAN is active, there is typically a simultaneous decrease in DMN activity, across many different tasks (e.g. Shulman et al., 1997). Hence, it appears that attention is able to profoundly

modulate activity across networks of cortical areas at a far larger spatial scale than changes in voxel-level activation (Spadone et al., 2015) (see Figure 1.5).

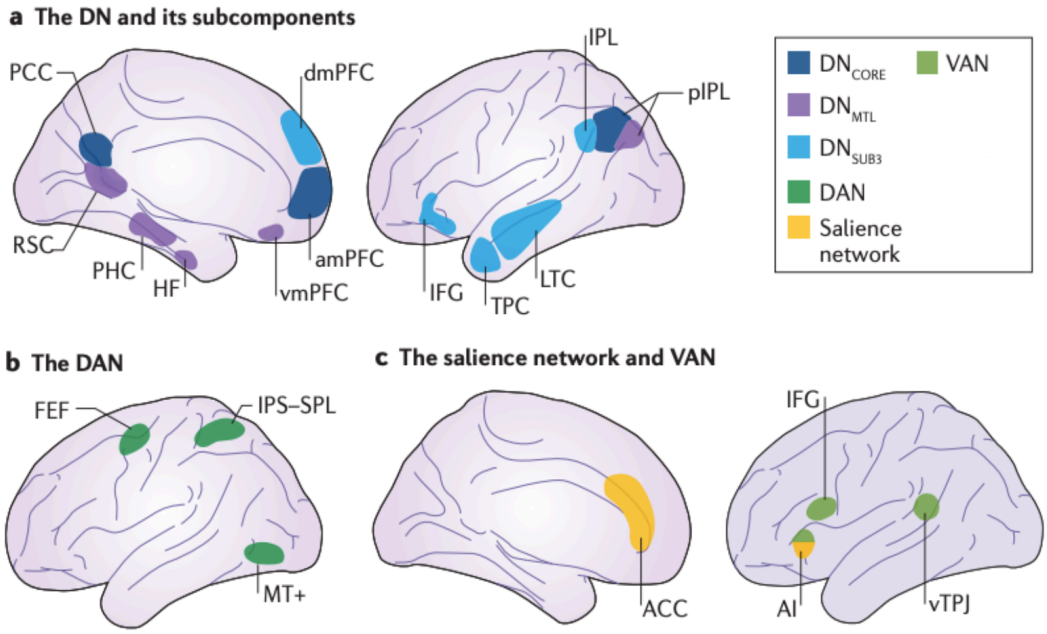

Figure 1.5 Large-scale brain connectivity networks. A) The default mode network (DN) core subsystem includes the anterior medial prefrontal cortex (mPFC), posterior cingulate cortex (PCC) and posterior inferior parietal lobule (pIPL). A second default mode subsystem centres around the medial temporal lobe (MTL) including the hippocampal formation (HF) and parahippocampal cortex (PHC). The third subcomponent of the default network extends dorsally into the dorsomedial prefrontal cortex (dmPFC) and the lateral temporal cortex (LTC). B) The dorsal attention network (DAN) is composed of regions centred around the intraparietal sulcus (IPS), the superior parietal lobule (SPL) and along the dorsal frontal cortex. C) The ventral attention network (VAN) contains a collection of ventral frontal regions such as the inferior frontal gyrus (IFG), anterior insula (AI) and the ventral temporoparietal junction (vTPJ (Christoff, Gordon, Smallwood, Smith, & Schooler, 2009) (Figure from Christoff, Irving, Fox, Spreng, & Andrews-Hanna, 2016).

1.4.3 Attentional modulation of network connectivity

Previous research has demonstrated that attention can modulate the pattern of activity in individual visual regions or across voxels. However, recent research suggests attention may also modulate the connectivity across multiple areas of cortex simultaneously. For example, Bartels & Zeki, (2004b) identified higher anatomical specificity and functional connectivity across the cortex during natural viewing of a movie stimulus versus 'rest', when participants performed no explicit task. Additionally, Hasson, Nir, Levy, Fuhrmann, & Malach, (2004) demonstrated individual brain regions 'tick together' in a synchronised fashion during passive viewing of a dynamic, complex stimulus, providing evidence for the distributed synchronized activity of the default mode network during undemanding tasks. Conversely, Fox et al., (2005) identified a decrease in BOLD signal in default mode regions during challenging, externally-directed tasks, again providing evidence for the modulation of activity in default mode and task-based networks as a function of attention.

The reciprocal relationship between the default mode and task-based networks has been termed 'functional antagonism' (see Anticevic et al., 2012). The extent of this activity change between networks as a function of externally-directed attention has been suggested to have specific benefits. For example, previous research has demonstrated that individuals possessing stronger negative connections between the fronto-parietal and visual attention networks with the default mode network had better cognitive functioning and task-switching (executive function) abilities (see Reineberg, Gustavson, Benca, Banich, & Friedman, 2018). This cognitive flexibility

(antagonism between greater positive connectivity during rest and negative connectivity during demanding tasks in the default mode network), has also been associated with higher levels of creativity (Li et al., 2017), better reading abilities in childhood and higher stress resilience (see Dajani & Uddin, 2015 for a review).

1.5 Naturalistic stimuli

Traditional visual research focuses on the use of highly-controlled low-level visual stimuli in order to probe precise aspects of visual system functioning. However, the extent this research is always a reflection of visual processes in the real-world is questionable. To tackle this issue of generalisation, recent research is beginning to investigate visual function with complex, dynamic, uncontrolled naturalistic stimuli, such as movie clips, which possess far greater similarity with real-world visual scenes than more traditional stimuli. Haxby, Connolly, & Guntupalli, (2014) argue for the importance of naturalistic visual stimuli in identifying transformation parameters to allow for modelling of low-level stimulus results to real-world scenarios. Additionally, Spiers & Maguire, (2007) note that the human brain evolved in a complex and dynamic visual world, and stress the importance of examining visual function under these conditions.

Naturalistic stimuli are also considered to offer some benefits beyond more traditional stimuli. For example, Bartels & Zeki, (2004b) demonstrate that the 'rich' stimulation provided by a naturalistic stimulus elicits greater activation of areas than typically seen with conventional stimuli. They also demonstrate a high similarity of results between two halves of an experiment conducted

with a naturalistic stimulus, indicating that despite their imprecise and uncontrolled nature, complex, dynamic stimuli are effective in investigating activation across the visual cortex (Bartels & Zeki, 2004a). In line with previous research, Russ & Leopold, (2015) also demonstrate that it is possible to create functional maps and assess aspects of functional brain organisation under natural viewing conditions. Hence, recent research demonstrates the importance of experimentation with naturalistic, dynamic visual scenes to supplement findings with more conventional low-level, highly controlled visual stimuli.

1.6 Outline of thesis

This thesis contains four experiments organised into individual empirical papers. First, *Chapter 2* provides a detailed overview of the methods used across the four experimental chapters. In *Chapter 3,* we conducted a psychophysics experiment to assess the effectiveness of radial frequency patterns as a stimulus to probe visual attention through the use of the selective versus distributed paradigm. In *Chapter 4*, we use visual psychophysics and fMRI to examine feature-specific patterns of attention and functional connectivity in human visual cortex. In *Chapter 5*, we use similar methods with chromatic stimuli to examine the feature- and colour-specific patterns of attentional modulation in the visual cortex. In *Chapter 5* we use fMRI and a complex, naturalistic stimulus to investigate differential patterns of stimulus-specific activation and connectivity when directing attention toward a dynamic, uncontrolled visual stimulus. *Chapter 6* summarises the

conclusions of these four experiments, as well as describing directions for future studies.

2. Methodologies

The following is a review of key methodologies used in this thesis that are covered only briefly in the experimental chapters themselves. This includes a definition of radial frequency patterns and isoluminance testing as well as a summary of the principles of neuroimaging, including population receptive field mapping, MT+ localisation and machine learning classification methods.

2.1 Radial frequency patterns

Here we discuss the creation of stimuli used in *Chapters 3, 4 and 5.* Radial frequency (RF) patterns are sinusoidally-modulated circular patterns, characterised by the number and size of their lobes (Ivanov & Mullen, 2012; Wilkinson, Wilson, & Habak, 1998; Wilson & Wilkinson, 1997) (see Figure 2.1). RF patterns are defined in polar coordinates, and are specified using the following formula:

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

Equation 2.1 Equation for the calculation of a RF pattern

In this equation, r is specified as a function of θ, which represents the angle around the circle's perimeter, which can be modulated sinusoidally by altering frequency (ω) which specifies the number of lobes, and amplitude (A), specifying the size of those lobes. The rotation of the stimulus is defined by

the phase (∅). The mean radius (r_0) specifies the average size of the stimulus in degrees of visual angle.

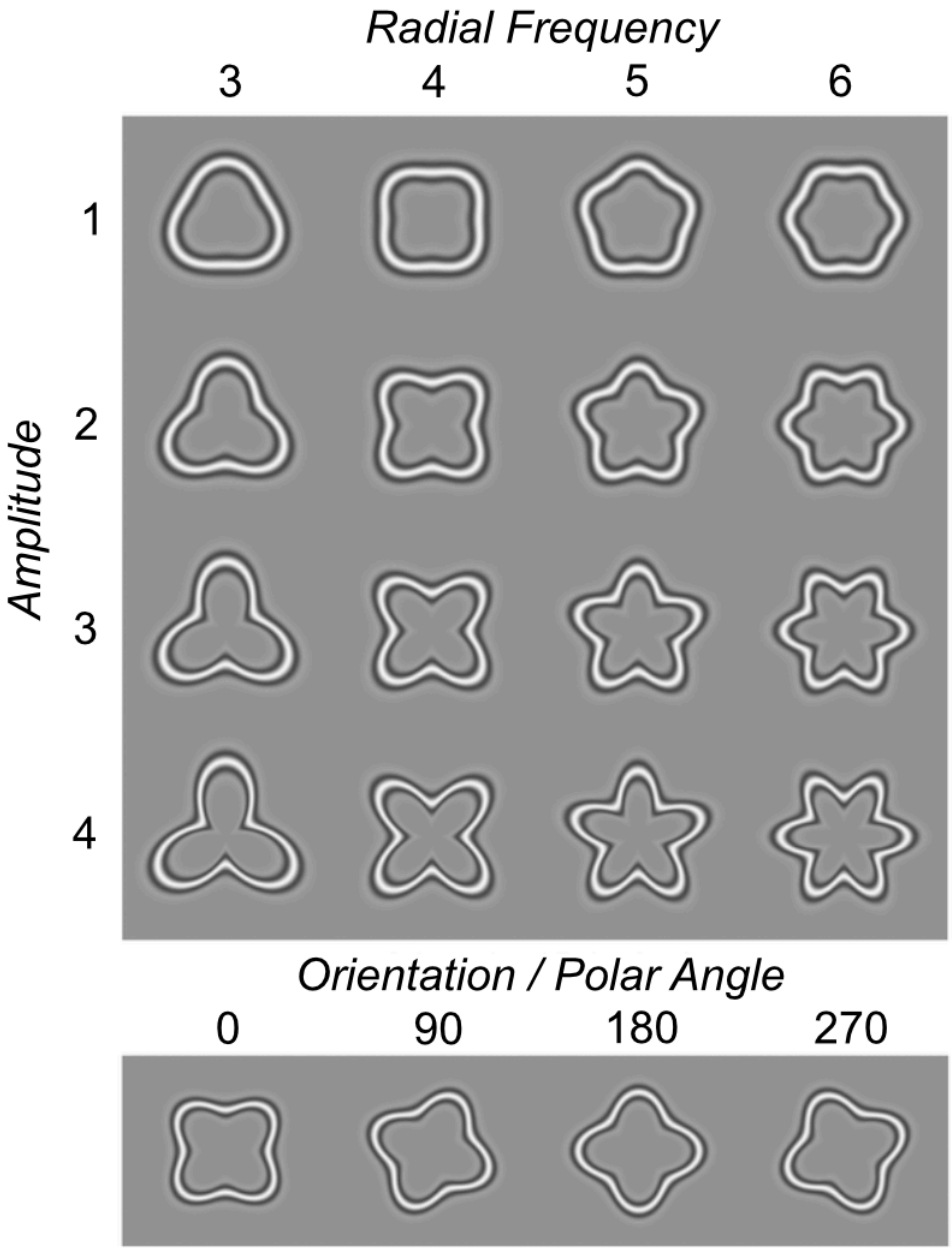

Figure 2.1 Radial frequency patterns with different radial frequencies, amplitudes and orientations (polar phase). Radial amplitude refers to modulation relative to the radius of a base circle. Within this thesis, we use three-lobed radial frequency pattern and modulate the amplitude (shape), orientation and contrast of the stimulus (Figure from Salmela, Henriksson, & Vanni, 2016).

The RF patterns in this thesis were all presented against a mid-grey luminance background. The contour of a RF pattern contains a cross-sectional luminance profile defined by the fourth derivative of a Gaussian (D4) (see Figure 2.2). This produces a Gabor-like appearance, which previous research has demonstrated is effective in eliciting activation of simple cells in V1, which respond to edges and gratings (Hubel & Wiesel, 1959). Hence, RF patterns should elicit clear activation in early visual cortical areas.

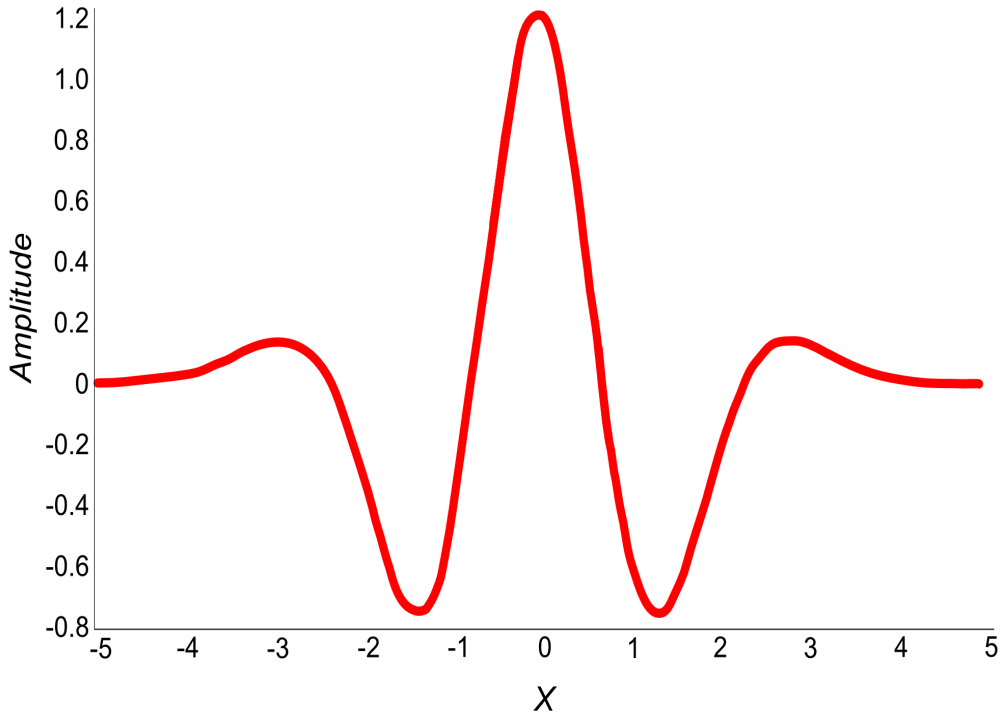

Figure 2.2 Fourth Derivative of a Gaussian (D4) used to render radial frequency patterns.

We chose to use RF pattern stimuli in *Chapters 3,4* and *5* over more traditional Gabor stimuli for several reasons. First, they allow for simple parametrical manipulation of a number of isolated stimulus characteristics

(such as orientation and contrast). Additionally, RF patterns allow for manipulation of stimulus curvature (shape), which is known to elicit activation in a number of regions of the visual cortex (see section 1.2.2.3, or see Pasupathy, 2006 for a review). Evidence also suggests visual region LO-2 is causally involved in processing of RF patterns (used in this thesis), making them a highly-relevant stimulus for the investigation of shape into relatively higher regions of the visual cortex (Silson et al., 2013).

As noted in Vernon, (2016) the rendering method used when plotting RF patterns with large deviations from concentricity produces rapid changes in radial distance, resulting in distorted shading of the RF pattern contour. To resolve this issue, we used a distance transform computed by Vernon, (2016) (see Figure 2.3). Here, for any given pixel, the shortest distance to the contour is calculated along any direction, as oppose to a single direction and then is shaped by the same D4 function. This ensures the rendering of the stimulus contour remains constant across any range of frequency and amplitude.

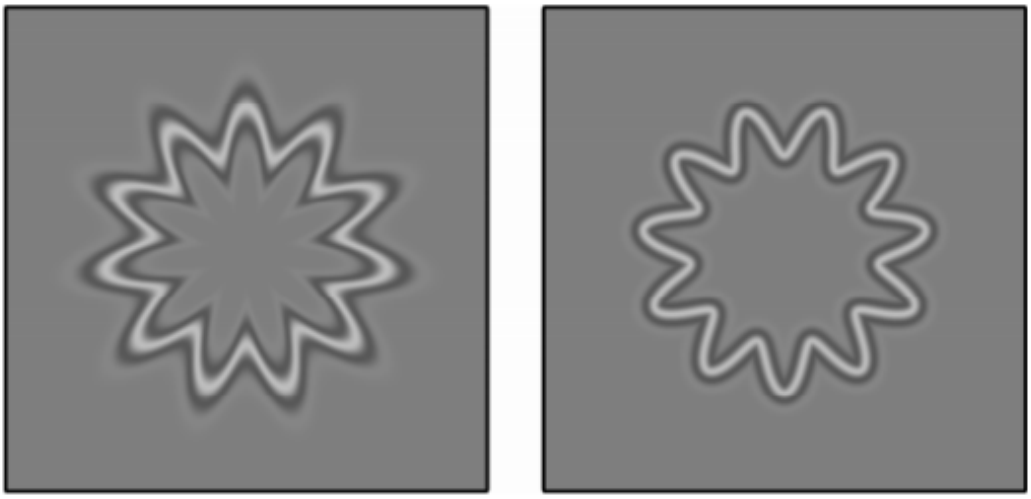

Figure 2.3 D4 Rendering Method from Vernon et al., (2016). The left image demonstrates a radial frequency pattern rendered using the fourth derivative of a Gaussian in polar coordinates. Deviations from circularity cause rapid changes in radial distance, leading to distortions in shading along the perimeter of the shape. The right stimulus demonstrates the same radial frequency pattern rendered based upon a distance transform, producing a more uniform boundary (Figure from Vernon et al., 2016).

2.2 Isoluminance testing

Isoluminance refers to chromatic stimuli in which the contributing colours (e.g. red-green) have been carefully equated in luminance so that they stimulate only colour- and not luminance-sensitive mechanisms (Anstis & Cavanagh, 1983; Lu, Lesmes, & Sperling, 1999). Humans vary in the amounts of retinal pigmentation, the ratio and sensitivities of cone photoreceptor types in the retina and random variation in neuronal responses (Roorda & Williams, 1999; Sumner, 2006). It is therefore important to calibrate chromatic stimuli perceptually to remove 'stray' luminance signals (L+M). To achieve this, in *Chapter 5*, our participants completed isoluminance (minimum motion testing).

Evidence suggests that colour and motion signals are processed in separate neural pathways (Zihl, von Cramon, Mai, & Schmid, 1991). This suggests that motion cannot be perceived in perfectly isoluminant chromatic stimuli. In isoluminance experiments, participants equate the amount of luminance in a particular chromatic channel (e.g. they equate the luminance of the red and green aspects of a L-M stimulus), until no motion is detected. We implement this by allowing participants to alter the angle 'θ' of the stimulus vector in MacLeod Boynton space - effectively changing the ratio of L to M cone excitation (see Figure 2.4).

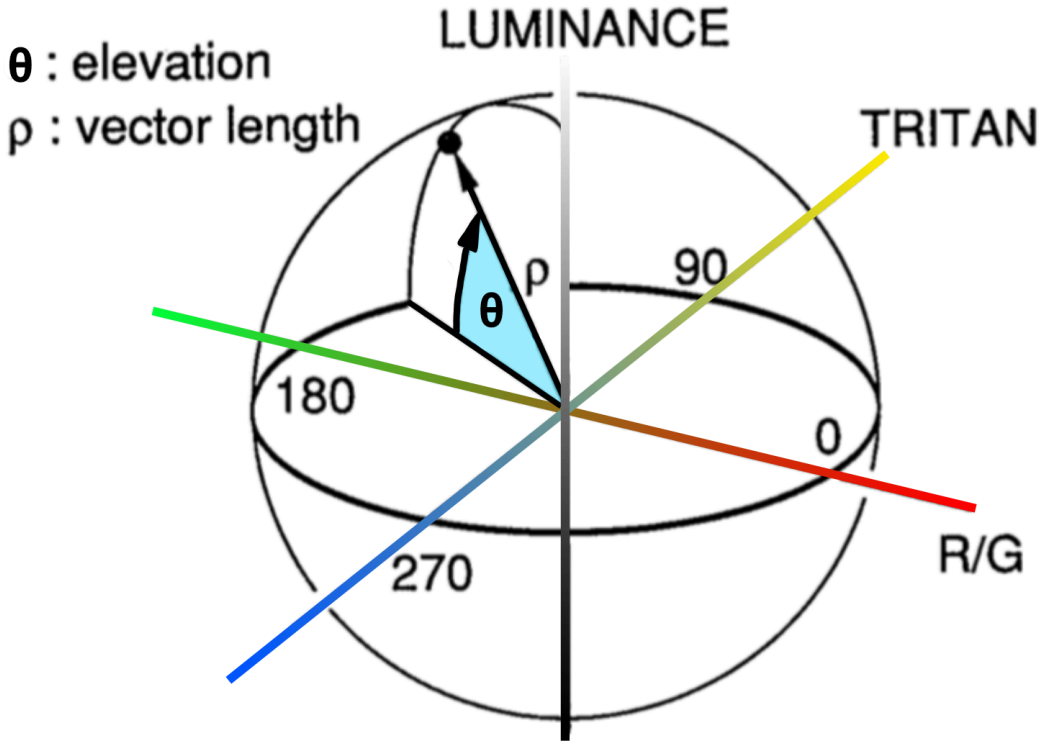

Figure 2.4 Three-dimensional colour space used to define the point of isoluminance. Alteration of the elevation parameter (θ) changes the amount of luminance present within a chromatic signal. Isoluminance testing asks participants to alter the angle of θ until a flickering/moving stimulus is no longer visible- the point at which luminance information within the two aspects of a chromatic stimulus is equated). Isoluminance testing is performed for red-green and blue-yellow (tritan) defined stimuli (Figure from Palmer, Mobley, & Teller, 1993).

During this testing, participants fixated centrally at 57cm viewing distance in a darkened room. All stimuli were presented on a mid-grey luminance background with a central white fixation cross. Stimulus colour was initially specified in LMS cone-excitation space, with matrices for the conversion from LMS to RGB values given as a product of the Stockman Sharpe 10^{o} fundamentals for the L-, M- and S-cones, and the spectral power distribution of the RGB phosphors in the stimulus screen. Participants viewed a central

annulus through a gaussian window which contained a chromatically-defined (L-M or S-(L+M)) circle (3.8° diameter) with a sine wave contour (2.33 cycles/°) (see Figure 2.5). The phase of the stimulus cycled through a range of 0- π/2 radians in 6° increments, which appeared to pulse backwards and forwards in space.

Participants altered the amount of luminance contamination, varying the angle θ (in 0.005° increments) until no or very minimal motion of the stimulus was perceived. Audio feedback was provided to participants, indicating the minimum and maximum extents of luminance contamination in the stimulus. Participants completed three sets of adjustment for each chromatic condition (L+M and S-(L+M)). The average of these three values of θ for each stimulus colour was then used to specify stimulus chromaticity in the psychophysics and fMRI experiments included in *Chapter 5.*

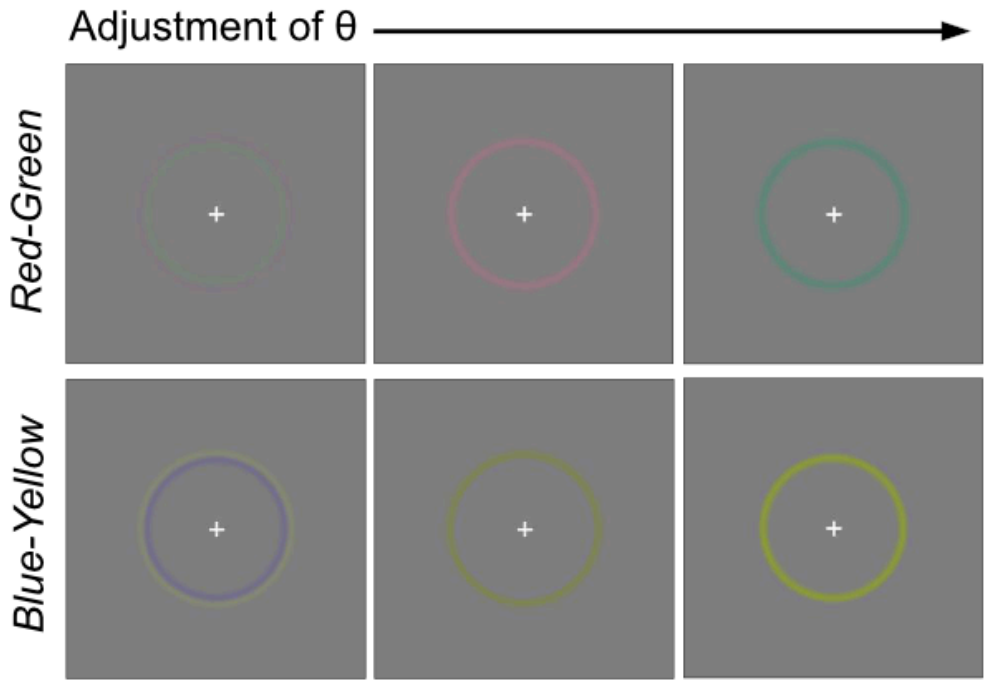

Figure 2.5 Example red-green (L-M) and blue-yellow (S-(L+M)) isoluminance stimuli. Participants adjusted the amount of luminance contamination within the chromatic signal (θ) to equate the amount of luminance within the stimulus. Stimuli were defined by a 3.8° circle with a sine-wave contour (2.33 cycles/°). The stimulus pulsed back and forth in space with altering phase, across a range of 0- π/2 radians in 6° increments.

2.3 Neuroimaging

2.3.1 The BOLD response

Functional magnetic resonance imaging (fMRI) is a non-invasive measure of blood-oxygen changes in the brain. This BOLD signal is considered a proxy for neuronal activation (see Logothetis & Wandell, 2004 for a review). When neurons fire signals, they use energy (in the form of ATP), which requires oxygen in order to be replenished (via a process of oxidative phosphorylation). Oxygen is transported around the body in a protein called haemoglobin. Oxyhaemoglobin is a haemoglobin protein possessing oxygen

molecules, conversely deoxyhaemoglobin possesses no bound oxygen molecules. These two states of haemoglobin have different magnetic susceptibility. Deoxyhaemoglobin is a paramagnetic molecule (attracted to magnetic fields), whereas oxyhaemoglobin is diamagnetic and relatively repelled from magnetic fields.

These differences produce a difference in susceptibility to the radiofrequency (RF) pulse dephasing protons; deoxyhaemoglobin molecules exhibit a greater susceptibility to dephasing of hydrogen nuclei, and an associated reduction in T_2* relaxation time, resulting in a darker colour of voxels containing deoxyhaemoglobin in an MR image. As oxyhaemoglobin is diamagnetic, it does not produce the same dephasing of protons as in deoxyhaemoglobin and these differences in relative blood oxygenation appear as signal (colour) changes in the T_2*-weighted functional magnetic resonance images. As the relative proportion of oxygenated versus deoxygenated haemoglobin increases, the blood-oxygenation level dependent (BOLD) signal increases. Hence fMRI images make the assumption increased neural firing demands increased blood flow, which is taken to indicate metabolic activity in particular regions of the brain.

2.3.2 The haemodynamic response function (HRF)

MRI does not measure neural activity directly, hence there exists a temporal lag between neuronal firing and resultant changes in relative proportions of oxygenated and deoxygenated blood. The time course of the BOLD signal is termed the haemodynamic response function (HRF). The BOLD signal

initially begins to rise 1-2 seconds after neural activity begins, as oxygen is extracted from the blood to supply the oxidative phosphorylation process and the haemoglobin becomes paramagnetic. The HRF peaks at around 4-6 seconds, in response to compensatory blood flow, with the peak denoting the primary neural response. If neural activity continues the BOLD signal will plateau. Upon removal of stimulation, the BOLD response begins to fall, with an initial undershoot (falling slightly below baseline), before returning to baseline after around 16-18 seconds. The full process can take up 30 seconds (Logothetis, Guggenberger, Peled, & Pauls, 1999) (see Figure 2.6). The haemodynamic response function shape is best typified by a combination of two gamma functions, to provide a double-gamma shaped haemodynamic response function.

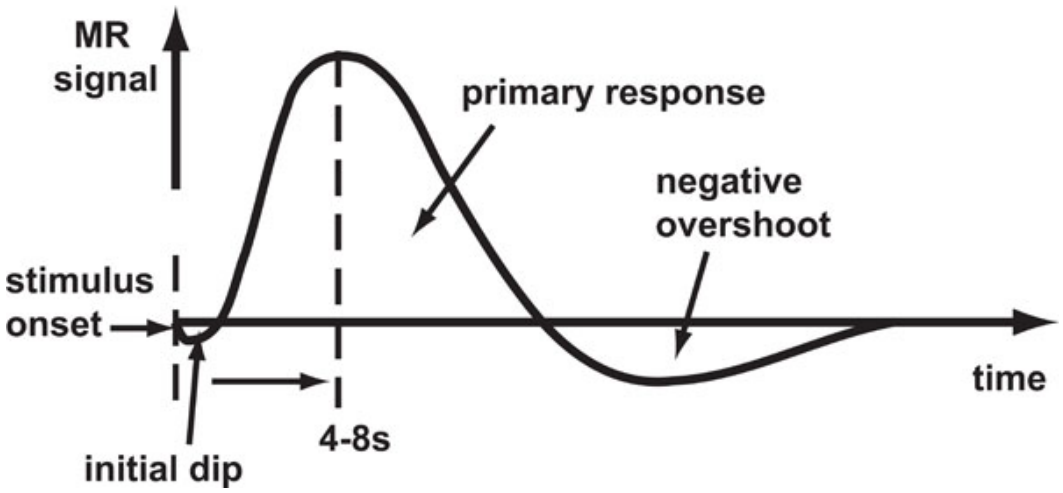

Figure 2.6 The BOLD haemodynamic response. The BOLD signal initially dips with the onset of a stimulus, then increases to a peak at 4-8 seconds. Once maximal MR signal has been achieved, the signal declines, followed by a negative overshoot then returns to baseline (Figure from Kornak, Hall, & Haggard, 2011).

2.3.3 Magnetic resonance imaging

Both magnetic resonance imaging (MRI) and fMRI measure changes in magnetic fields. With no clear magnetic field, hydrogen nuclei (protons) in a human body are randomly oriented, with no net magnetisation. However, when a participant enters a strong magnetic field (in an MRI scanner), protons align themselves along the bore of the scanner (B0) (the longitudinal direction) to reach equilibrium (they become polarised). Approximately half of these protons align parallel to the B0, with a relatively low-energy stable state. The remaining protons align antiparallel to the longitudinal direction in a high-energy, less stable state. The discrepancy between high and low stability results in a bias towards the stable, parallel alignment state, producing net magnetisation in the B0 direction. Each proton possesses a mass and spin, which provides them with angular momentum. Hence, protons are not static. Instead, they precess around a magnetic field direction (i.e. they rotate around the alignment direction, forming a cone-shaped arc). Magnetic resonance imaging reflects changes in the alignment of protons across energy states.

During an MRI scan, a short RF pulse is applied, which provides the hydrogen nuclei with energy. This causes the low-energy protons in the participant to flip to a high-energy state, which reduces the overall magnetisation along the B0 axis. By applying a 90° RF pulse specifically, approximately half of the protons are in high- and low-energy states respectively, which cancels out the magnetisation along the longitudinal direction. This also brings the spins of all protons into phase, resulting in a

net magnetisation in the transverse plane (at a right angle to the longitudinal plane).

Once the RF pulse is terminated, the protons are in an overall less stable state. Overtime, these protons return to their equilibrium state (relaxation). Protons in the high energy (unstable) state will revert to the lower-energy (stable) state over time, leading to a gradual restoration of the magnetic field in the B0 plane (T_1 recovery). T_1 recovery specifically refers to the time is takes for 63% of the magnetisation in the B0 plane to recover. However, additionally, when protons transition from high to low energy states, they also release an RF pulse at the same frequency as the stimulation pulse. This can be detected by the MRI scanner, via a receiver coil in the transverse plane, as overtime, during this transition, protons lose their spin coherence. Hence, this is termed T_2 decay, as it results in a loss of overall magnetisation in the transverse plane and refers specifically to the time it takes for transverse magnetisation to decay by 63% of its maximum value. The current induced from this decay is detected as an MR signal. T_1 relaxation time is always longer or equal to T_2 relaxation time.

T_1 and T_2 relaxation times differ dependent on in which tissue hydrogen nuclei lie. For example, in T_1-weighted images, cerebrospinal fluid (CSF) is dark as relaxation time is longer in water, yet in T_2-weighted images, CSF is bright. These differences help to provide a high-resolution spatial image.

fMRI exploits differences between oxygenated and deoxygenated blood to identify the BOLD response. Deoxygenated blood has greater magnetic susceptibility than oxygenated blood, meaning it will produce local inhomogeneities in the magnetic fields in the regions it is present. These inhomogeneities can elicit a loss of spin coherence across the protons, which can interact with the loss of phase coherence in the transverse plane. This process is known as T_2* decay and allows for measurement of BOLD response across the brain.

2.3.4 Retinotopy and population receptive field (pRF) mapping using fMRI

In *Chapters 4, 5* and *6*, we define specific regions of the visual cortex to analyse and perform subsequent regions-of-interest (ROI) based analyses. We define our ROIs for this purpose using retinotopic, population receptive field (pRF) mapping. Retinotopic mapping exploits the fact early visual areas contain explicit maps of the visual world (as described in 1.2.4) (e.g. Engel et al., 1994). Additionally, a neuronal receptive field is defined as the region of the visual field in which is it possible to influence the firing of that neuron (see Hubel & Wiesel, 1962). These receptive field preferences can be split into a location in visual space, and an associate spread. Given early visual regions are retinotopically-defined, the preferred locus of activation should vary systematically across neurons (voxels) as you travel across a visual area. Recent developments in fMRI have been able to combine the principles of retinotopic mapping and the architecture of receptive fields in order to provide a measure of a single voxels polar angle and eccentricity in visual space, through a process termed pRF mapping (Dumoulin & Wandell, 2008).

A typical voxel of the human brain will contain more than 100,000 neurons; hence the BOLD response summarises computations of populations of neurons with tuning to multiple stimulus parameters in a single voxel. pRF mapping summarises this pooled neuronal response in a voxel as a population receptive field, assuming neurons in a voxel will share similar (at least spatial) preferences due the organised topographical mapping of early visual areas (Wandell & Winawer, 2015). Hence, a pRF response for an individual voxel is considered a reflection of the mean tuning of neurons in that voxel.

Prior to the development of pRF mapping procedures, the topographical organisation of the visual cortex was identified via retinotopic mapping paradigms (Engel, Glover, & Wandell, 1997). Here, the topographic mapping of the visual field is identified through presentation of high-contrast, achromatic rotating wedges and expanding rings. As the wedge rotates for example, it will activate discrete populations of neurons as it falls in their receptive fields. This will happen multiple times with the rotation (or expansion) of the stimulus. Overtime, comparison of the phase of the stimulus and the recorded BOLD response, will produce a travelling wave of activation, providing a measure of tuning in visual space for each voxel in visual cortex. pRF mapping expands upon this traditional retinotopic mapping paradigm, by allowing investigation of both traditional (visual field location) and additional voxel parameters (such as receptive field size), using a high-

contrast checkerboard drifting bar stimulus (Dumoulin & Wandell, 2008) (see Figure 2.7).

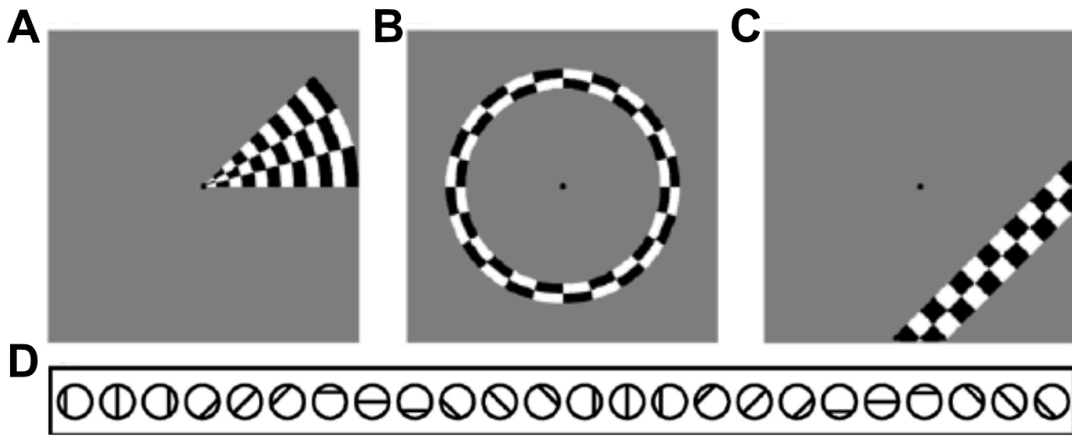

Figure 2.7 Illustration of retinotopy and population receptive field mapping (pRF) stimuli. A and B demonstrate the rotating wedge and expanding ring stimuli used in traditional retinotopic mapping used to map the topographic organisation of the visual cortex (polar angle and eccentricity). C) demonstrates the standard bar stimulus which drifts across the visual field in eight directions, detailed in D. All stimuli are high-contrast achromatic checkerboards presented on a mid-grey luminance background (Figure from Dumoulin & Wandell, 2008).

In pRF mapping, the receptive field of each voxel is defined as an isotropic Gaussian, with two visual field locations (x_0 and y_0) and a pRF size (δ). These parameters can be extracted from an fMRI timeseries in response to a traditional retinotopic mapping stimulus or a drifting bar stimulus (a high-contrast, achromatic checkerboard, size invariant bar which drifts in multiple directions across visual space). The pRF model employs a linear spatiotemporal model of fMRI response, which is defined as:

$$\gamma(t) = p(t)\beta + e$$

Equation 2.2 Equation for the spatiotemporal linear model of the fMRI response

Here, $p(t)$ refers to the predicted fMRI signal, β is the response strength scaling factor (accounting for the unknown units of fMRI signal) and e refers to noise.

The predicted fMRI signal for each voxel is calculated using a 2D Gaussian model of the population receptive field, defined below. Here, x_0 and y_0 refer to the centre of the pRF, and σ refers to the standard deviation of the Gaussian (or spread of the population receptive field).

$$g(x,y) = exp - \left(\frac{(x - x_0)^2 + (y - y_0)^2}{2\sigma^2}\right)$$

Equation 2.3 Equation for the Gaussian model of pRF used to calculate p(t)

The pRF model finds and fits the optimal pRF parameters to each voxel by minimising the residual sum of squares between the predicted and observed fMRI time course. In the first stage of this process, data are smoothed with a 5mm full-width at half-maximum (FWHM) Gaussian kernel to improve spatial correlation amongst neighbouring voxels. Next, brute force search is applied to identify the best fit for each voxel to the observed fMRI data from 100,000 timeseries predictions. In this second stage, unsmoothed voxels from the

first stage whose fits explain more than 15% of the variance in that voxels time course are retained. From these fits, x_0, y_0 and σ are estimated for each voxel. The values provide for each voxel, the eccentricity (how far from the centre of the visual field a voxel is tuned), the polar angle (radial distance preference) and receptive field size (the degrees of visual angle in space a voxel receives input from).

2.3.5 MT+ motion localiser

For some participants, we employed a MT+ motion localiser in combination with population receptive field mapping to isolate the voxels in area MT+ with a clear response to motion stimuli. The design for this MT+ localiser was adapted from Fischer, Bülthoff, Logothetis, & Bartels, (2012), Huk, Dougherty, & Heeger, (2002) and Maloney, Watson, & Clifford, (2013). The motion stimulus was composed of black and white dots with a density of 9.9dots/degree^2, with a smoothed Gaussian profile (σ = 0.04). These dots filled an annulus which extended from 0.5°-11.75° eccentricity, on a mid-grey luminance background. Dot position updated at 0.33Hz frequency, yielding a 5.3°/sec average dot speed (see Figure 2.8). The localiser consisted of a block design, with each run containing 31, 12 second blocks. Each run included five stimulus conditions; static dots (a randomly selected single frame of the motion stimulus), left-hemifield coherent radial motion (with motion restricted to the left 120° of the display, hence 66.66% of the display contained static dots), right-hemifield coherent radial motion, full-field coherent radial motion and fixation-only trials. A jittered inter-stimulus-interval (ISI) of between 3-12 seconds was included between each full stimulus cycle

(60 seconds). All participants completed two runs of the localiser to enable accurate location of MT+ post-hoc. MT+ location was identified via a full-field coherent radial motion versus static dot GLM contrast. The voxels activated in this contrast were defined as the region-of-interest for further analysis.

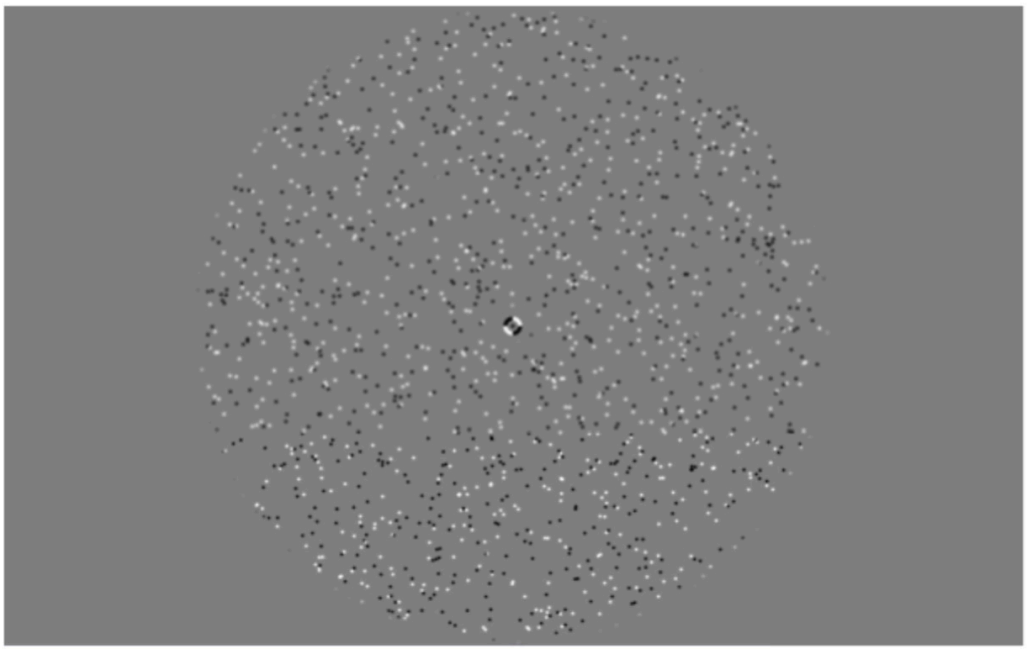

Figure 2.8 Full-field motion-localiser stimulus. Participants viewed alternating blocks of full-field motion, left- and right-hemifield motion, static motion and fixation-only. Achromatic dots were presented on a mid-grey luminance background and had a 5.3°/second dot speed. Comparisons of the full-field versus static motion conditions allowed for accurate location of visual area MT+.

2.4 Machine learning classification

2.4.1 Overview

The aim of machine learning is to accurately predict a class a single data point has originated from. Typically, classification is supervised. This refers to the fact a machine learning classifier is provided with explicitly labelled

data (i.e. data point one belongs to class A, data points 2 and 3 are members of class B). A classifier then uses this data to 'learn' how each piece of the supplied data relates to its assigned class. The goal of classification is to produce a model from this training data which then accurately predicts the class of each piece of test data, using statistical similarities it has identified about each of these classes from the test data.

Modern machine learning frequently involves very large datasets. In *Chapters 4, 5* and *6* we employ use a form of machine learning classification with multivariate (voxel-level data), with the aim of classifying unique patterns of feature-specific attentional focus. Area-level sensitivity to stimulus changes can be examined using many methods such as multivariate pattern classification analysis (MVPA) and adaptation. Researchers sought to examine the relative strengths and weaknesses of such classification methods to evaluate their most suitable applications. In a comparison of such multivariate pattern classification methods (support vector machine versus adaption), whilst both models performed similarly on typical orientation classification paradigms, the multivariate pattern analysis (SVM) method possessed greater sensitivity to smaller differences in orientation than the adaption paradigm (Sapountzis, Schluppeck, Bowtell, & Peirce, 2010). In *Chapters 4* and *5*, we present very small stimulus feature changes (close to participant's thresholds for detection), hence in this thesis, we choose to employ MVPA analyses (SVM) as oppose to adaptation paradigms.

However, here, in the consideration of participant fatigue and data collection, instead of collecting new data to test the classifier on, we split our dataset for each participant into two groups (a training and test set). In these chapters we employ leave-one-out cross validation in order to assess how our classification model performs on multiple, varying subsets of training and test data. In each cross-validation iteration, the machine learning classifier is trained on all but one of our data instances. This fitted classification model is then used to assess the accuracy of classification for the remaining data. This process is repeated for all combinations of training and test sets and we average the classification accuracies to gain an overall percentage classification accuracy score.

2.4.2 Support vector machines

In *Chapters 4,5* and *6* we use a support vector machine (SVM) classifier to identify potential differences in patterns of activation across multiple voxels during feature-specific attentional focus. Support vector machine classification is a popular and widely used form of classifier (Chang & Lin, 2011), created by Cortes & Vapnik, (1995). SVMs are capable of producing significant classification accuracies with less computational power than other comparable methods. Briefly, a support vector machine is a generalisation of a maximal margin classifier, which allows investigation of non-linear datasets and can be used for both binary and multi-class classification. However, a single support vector machine classification does not provide probability estimates. To overcome this, SVMs are frequently employed with cross-

validation techniques (see 2.4.1) or bootstrapping methods, to generate confidence intervals for the associated classification accuracies.

A support vector machine aims to find a hyperplane in an N-dimensional space (N referring to number of features in the analysis), that best separates the supplied data points into their respective classes (see Figure 2.9). Amongst many possible hyperplanes, the aim of an SVM is to identify the hyperplane with the maximum distance between the data points of the classes. This increased margin allows greater accuracy of classifying future data points which may lie slightly outside of the boundaries of the original dataset. Support vectors themselves are defined as the data points closest to (most informative in terms of the position and orientation of) the hyperplane. Regularisation can be employed to specify the size of the margins between the separating hyperplane and the data points. Low regularisation will identify the largest margin (biggest difference between the support vector and the hyperplane) even if this results in misclassifications. Alternatively, high regularisation specifies a small margin between support vectors and the hyperplane which will result in a higher classification accuracy in the training data, but which may lead to overfitting (see 2.4.4).

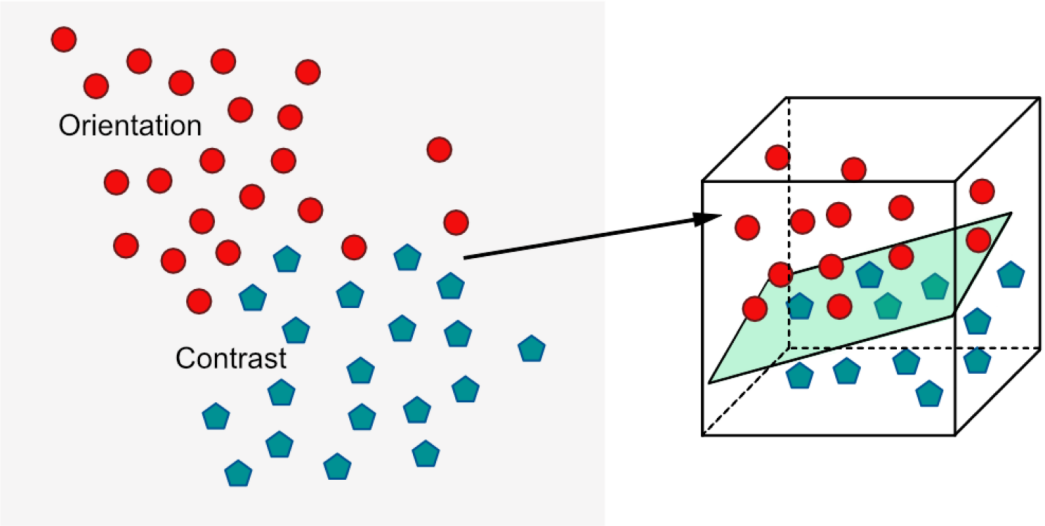

Figure 2.9 Support vector machine classification. Support vector machine classification takes an input of values across many instances for two conditions (here, fMRI voxel activation for during attention to orientation (red), and attention to contrast (blue)) (left). This data is transformed to a higher-dimensional space via a kernel trick to help fit a hyperplane (green) (right) linearly separating the data points into their respective conditions. This hyperplane is then used to assess classification accuracy on a test set of data.

Support vector machines typically employ a kernel to improve the fitting of the hyperplane for non-linear data. A kernel specifies the similarity between two sets of data and allows for transformation of data from a low-dimensional input space to a higher-dimensional space in which the data points are linearly separable. Instead of explicitly representing the original data, a kernel represents the data through pairwise similarity comparisons (a modified dot product) in this higher dimensional space.

2.4.3 Radial Basis Function Kernel

As described above, kernels transform low-dimensional data into higher-dimensional data to identify a hyperplane separating the data. There are many different types of kernel which can be used for this process. In this thesis, we employ a radial basis function (RBF) kernel in our SVM classification. The most popular RBF is the Gaussian RBF. In simple terms, the radial basis function fits each cluster of data with a radial shape. In the analysis in this thesis, this is a circle with a centre at the mean of the distribution, with a Gaussian spatially-smoothed profile. The SVM then assesses the probability another instance of data falls in the distribution for a particular category. A RBF has one free parameter, gamma, which specifies the variance of the Gaussian distribution.

2.4.4 Overfitting

Overfitting refers to situations in which a hyperplane is fit too closely to a single data set. In this situation, the model fit is heavily influenced by the precise details and noise contained in the data. When such a model is tested with a new set of data, its classification is typically very low (i.e. it does not generalise well from training data to a new test dataset). Overfitting is a common problem in machine learning classification. To overcome it, many classifiers include parameters which limit how well a model can fit the data, such as regularisation parameters (see 2.4.2). Alternatively, techniques such as leave-one-out cross validation (see 2.4.1) can be used to ensure the model tunes its hyperplane only with the original training data set, ensuring the test dataset has had no influence on the fit of the model, and this process

is repeated several times. The cross-validation technique has been suggested as the best technique for avoiding this issue in SVMs (Gholami & Fakhari, 2017). In *Chapters 4, 5* and *6* we employ this leave-one-out cross validation method to provide probability estimates and reduce the potential risk of overfitting our data.

2.5 Multiple comparisons correction

Benjamini-Hochberg correction (Benjamini & Hochberg, 1995), is a widely used method of false discovery rate correction, which controls the expected proportion of type one errors (falsely concluding a significant effect) across a range of significant values. It is a simple method which is sufficient for many exploratory experimental studies such as the ones detailed in this thesis, and the output provides a conservative estimate of significance across multiple statistical analyses, with a good balance between the discovery of statistical significance and limitation of false-positive occurrences (e.g. Noble, 2009). In comparison to methods controlling family-wise error rate, such as the highly conservative Bonferroni correction, which relates to the probability of making any type 1 errors at all, 'punishing' all p-values equally, the Benjamini-Hochberg correction, 'punishes' *p*-values according to their ranking.

Benjamini-Hochberg correction sorts the *p*-values in ascending order, then divides each observed *p*-value by its percentile rank to get an estimate false-discovery rate. Throughout this thesis, we employ Benjamini-Hochberg correction when correcting for multiple comparisons across a number of statistical tests (e.g. across the significance values for multiple one-sample

Wilcoxon signed-rank tests), and Bonferroni family-wise error correction when controlling for multiple comparisons in a related analysis (such as across post-hoc testing in a repeated measures ANOVA). These methods help to reduce the likelihood of making a type 1 error in our data analysis (i.e. concluding a significant effect when one does not truly exist).

3. Investigating Selective Versus Distributed Attentional Effects

3.1 Abstract

It has been frequently noted in the literature, that task accuracy decreases, and reaction time increases with an increased attentional load. Here, we sought to replicate these well-established findings in a selective versus divided visual psychophysics paradigm, utilising a relatively novel stimulus in the attention literature, the radial frequency pattern. Participants were directed to selectively detect changes in a single stimulus dimension (orientation, contrast or shape) or to identify changes in any of these stimulus attributes simultaneously (divided attention). Participants had significantly lower change detection accuracy and significantly longer response times when attention was divided across multiple features versus selectively focused upon one attribute. We provide substantiation for the use of the radial frequency pattern as an effective stimulus to probe visual attention.

3.2 Introduction

We are constantly exposed to complex, multi-faceted visual stimuli (e.g. objects with orientation, colour and motion information), hence a mechanism to weight the contextual importance of particular subsets of visual information is vital for efficient and successful interaction with our world (Posner, 1994; Posner & Petersen, 1990). Attention allows the selection and processing of relevant information at the expense of processing task-irrelevant information, which is filtered out. Featural attention refers to the modulation of activation

when focus is directed toward a specific visual feature (e.g. vertical orientation, red, upwards motion), across the entire visual field (Saenz et al., 2002). Featural attention is a robust phenomenon, which typically increases the activation in cortical regions responsible for the processing of the attended feature, such as increases in MT+ activation during attention to motion (Beauchamp, Cox, & DeYoe, 1997; Chawla et al., 1999; Corbetta, Miezin, Dobmeyer, Shulman, & Petersen, 1991; Huk & Heeger, 2000; O'Craven et al., 1997; Treue & Martínez Trujillo, 1999), or increased V4 activity with attention towards chromatic stimuli (Chawla et al., 1999; Liu, Slotnick, Serences, & Yantis, 2003; Schoenfeld et al., 2007).

Visual psychophysics experiments typically probe isolated and specific aspects of visual processing using highly-controlled visual stimuli such as Gabor patches. Here, we utilise a more recent visual stimulus, the radial frequency (RF) pattern (Wilkinson, Wilson, & Habak, 1998; Wilson & Wilkinson, 1997). Whilst, like the Gabor patch, the RF pattern is a highly-controlled low-level visual stimulus, it possess a relative advantage, allowing investigation into shape processing with manipulation of radial amplitude modulation and shape perimeter (for example; Wilkinson et al., 1998; Ivanov and Mullen, 2012; Bell et al., 2014; Lawrence et al., 2016). Additionally, combinations of multiple RF patterns are considered to more-closely reflect 'real-world' visual stimuli than other more traditional stimuli (Lawrence et al., 2016).

We planned to utilise these RF patterns as a method of probing the cortical regions and networks involved in encoding, processing and switching attention between different low-level visual features. However, very little research has specifically investigated attentional mechanisms using RF pattern stimuli (see Smith et al., 2009). Hence, there is little evidence to substantiate the use of RF patterns an effective stimulus to probe visual attention. As such, we conducted a psychophysics study aiming to replicate a key finding in the attentional literature, to assess the effectiveness of RF patterns as a stimulus to probe attention to low-level visual features.

The selective versus divided attentional paradigm compares participants change detection accuracy between conditions in which their attention is directed towards a single visual feature, versus a condition where participants must divide their attention across multiple visual features (Corbetta, Miezin, Dobmeyer, Shulman, & Petersen, 1990). In their original experiment, Corbetta et al., (1990) presented participants with consecutive frames of randomly dispersed bars moving leftwards or rightwards as a collective. The bars could vary in shape, colour and/or velocity between frames. Participants attention was directed toward a specific stimulus feature (for example, colour), and they were asked to make yes/no choice judgements regarding the presence/absence of a change in the *attended* stimulus between two frames (same/different). For example, a change in only the shape of the bars between frames in an ‘attend to colour’ trial should be reported as ‘same’, as no change in the attended feature occurred. In a separate condition, participants divided their attention across multiple

stimulus dimensions, making yes/no judgements to indicate the presence or absence of a change in at least one of the attended features between frames.

It is widely noted that change detection accuracy for detecting small changes in a visual feature is greater with selective attention (directed toward a single stimulus dimension) than when presented with the same stimulus, with attention divided across multiple stimulus attributes (Corbetta et al., 1991; Corbetta et al., 1990). A longer response time has also been identified in divided versus focused attention tasks (e.g. Ninio and Kahneman, 1974; Hahn et al., 2008). This deficit likely reflects the limited processing capacity humans possess; when we divide attention across multiple features, we must split or rapidly shift our attentional focus, limiting the resources we have available for the detection of change in any single feature simultaneously (Broadbent, 1958; Parasuraman, 1998).

We aimed to replicate the results of this well-established selective versus divided paradigm with RF patterns as a stimulus to probe feature-based attention, predicting lower sensitivity to change, and longer response times when attentional focus is divided across multiple stimulus attributes. We examine change detection across three low-level visual features of interest; orientation, contrast and shape. In this experiment, we used participant-specific task difficulty calibration, attempting to control the level of attention in the “attend orientation”, “attend contrast” and “attend shape” conditions, to validate this calibration method in future experiments. We hypothesised no

significant difference in loglinear d' or reaction times between the three selective attentional focus conditions. Our findings indicate participants are more accurate at detecting changes in a target feature with selective attention (directed toward a single stimulus feature) than during divided attentional focus. We find no evidence for differences in participant response (loglinear d' or reaction time) between the three selective attention conditions.

3.3 Methods

3.3.1 Participants

24 volunteer participants (all University of York staff or postgraduate students) (17 female) were recruited for the study. All participants had normal or corrected-to-normal vision. Ethical approval for the experiment was granted by the University of York, Department of Psychology ethics board. All participants provided informed consent and remained naïve regarding the specific focus of the study (differences between selective and divided attention) until after testing ended. Participants were provided with verbal instructions and had the opportunity to ask questions to clarify their understanding.

3.3.2 Design & Procedure

3.3.2.1 Stimulus presentation

Participants completed the experiment in a darkened room, at a viewing distance of 57cm from the screen. All participants completed at least one practice block (10 trials) before starting the experiment. Visual stimuli were

presented to participants on a ViewPixx 3D LCD screen system (1920 x 1080 pixels' resolution, running at 120Hz). A Shuttle XPC SZ88RG high-end graphics system with intel Core i7-4790K processor at 4.0GHz and an NVIDIA GeForce GTX970 graphics card with 4GB DDR5 Memory were used to control the experiment. All stimuli and experimental procedures were controlled by Matlab 9.2.0 (2017a) in conjunction with the Psychtoolbox 3.0.12 (Brainard, 1997; Pelli, 1997).

3.3.2.2 Threshold testing

Before assessing participants' selective versus divided attention change detection accuracies, we first had to ensure that the selective tasks were of equal difficulty, and that overall task difficulty was consistent across participants. To achieve this, we collected participants' 75% correct detection thresholds for each low-level visual feature of interest (orientation, contrast and shape).

To efficiently estimate detection thresholds, we used a Bayesian staircase procedure (Watson & Pelli, 1983). Feature-specific initial estimates of threshold (0.3 radians orientation, 30% contrast, 0.08 radial amplitude modulation/shape) were provided with 0.5-unit standard deviation. We utilised a RF pattern stimulus, a radially modulated circular contour defined by the fourth derivative of a gaussian in polar coordinates (Ivanov & Mullen, 2012; Wilkinson et al., 1998; Wilson & Wilkinson, 1997).

Participants were initially presented with a visual cue, signalling the attended feature and a central white fixation letter present across trials matched the cued feature (‘O’, ‘C’ or ‘S’). Trials commenced with a 200ms presentation of a grayscale reference stimulus against a mean luminance background. The reference RF pattern stimuli had a 2.0° average radius, 0° orientation and 50% contrast (see Figure 3.1A).

This was followed by a 200ms inter-stimulus interval with presentation of the letter-fixation cue only. A target stimulus was then presented for 200ms, with a change in the attended visual feature derived from the initial estimate and participants’ previous performance on subsequent trials. For each feature, a stimulus change could be either of two directions (with approximately equal frequency)- clockwise versus anticlockwise orientation, high contrast versus low contrast, ‘spiky’ versus smooth shape. This was followed by a central black fixation cross for 800ms, during which participants made a yes/no judgement indicating the direction of change between the reference and target stimulus (‘A’ or ‘L’ keyboard press). Participants were informed via a toned ‘beep’ if their response was correct or incorrect. Each participant completed 50 trials (alongside 10 discarded practice trials at the start of each run) to provide an estimated 75% correct detection threshold for each feature. Each feature-specific staircase was repeated three times with participants’ final detection thresholds reflecting the average of these three repetitions. Each run lasted approximately 2 minutes.

3.3.2.3 Selective versus divided

The stimulus design and procedure employed in the main selective versus divided attention paradigm was very similar to the staircase design employed above. Participants were first presented with a brief instruction screen, providing a repeat of the task instructions and required keypresses. Participants were then presented with a mean luminance grey screen with white text cueing the attentional focus condition (e.g. Attend to changes in ORIENTATION (selective) or Attend to changes in ORIENTATION or CONTRAST or SHAPE (divided). Participants were presented with the same reference RF pattern as the staircase threshold testing procedure, with a fixed 2.0° average radius, 0° orientation and 50% contrast, for 200ms. This reference RF pattern was presented with a small white letter at central fixation matching the cued featural attentional focus condition (O, C, S, or N, for orientation, contrast, radial amplitude modulation (shape) and divided attention respectively). This reference stimulus was then followed by the same central fixation letter as previous on the same mid-grey luminance background for a 200ms inter-stimulus interval. Following this, in an important distinction from the threshold staircase procedure, a target radial frequency pattern, with the potential to alter in one, multiple or no stimulus features was presented, along with the same attention-specific fixation letter (see Figure 3.1B).

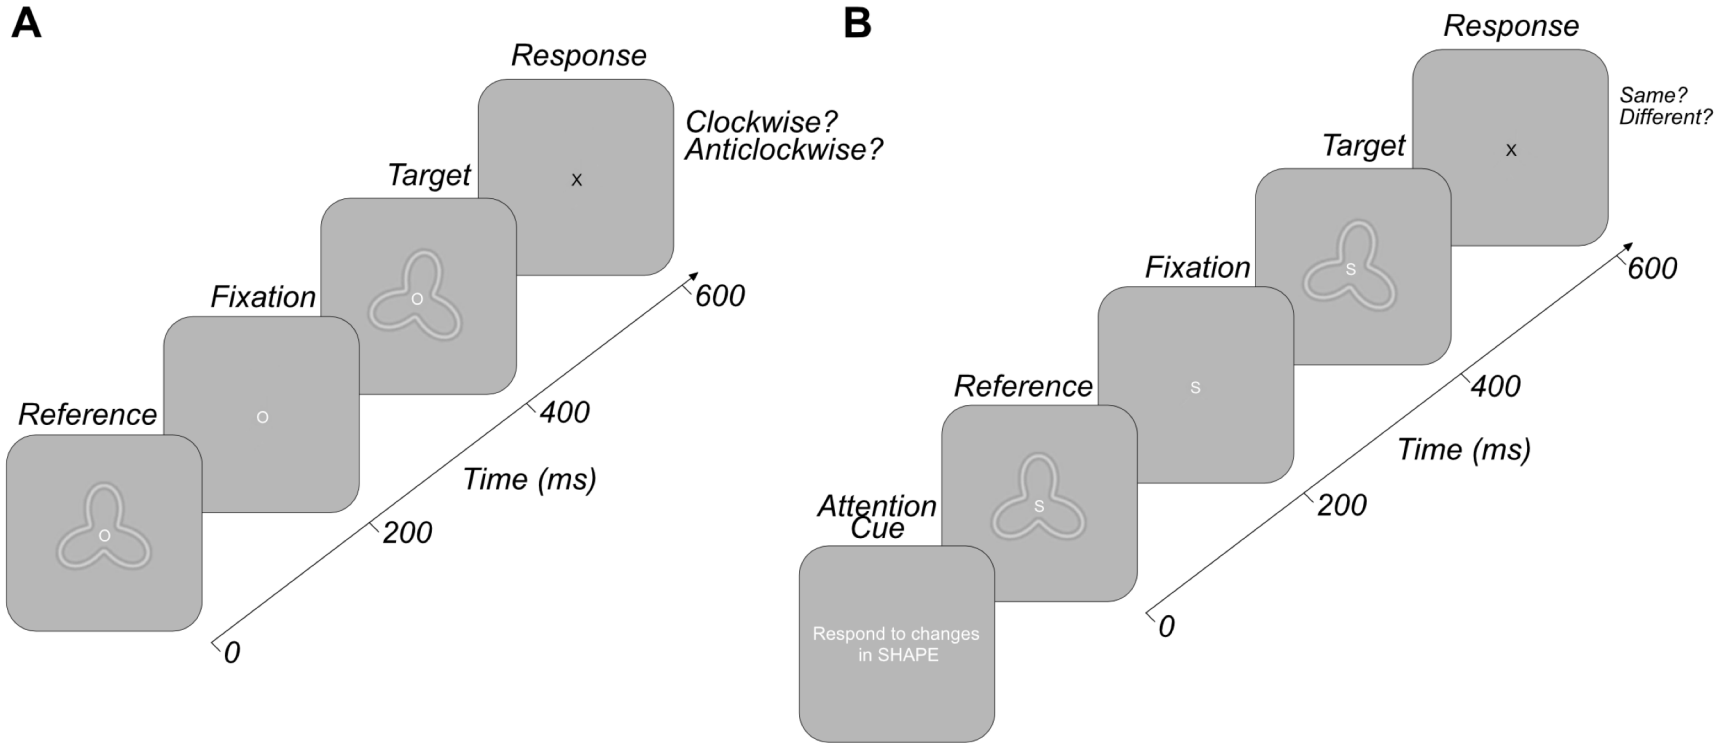

Figure 3.1 Exemplar threshold staircasing and selective versus divided trials. A) provides an example of an orientation threshold detection staircase. Participants viewed a reference radial frequency pattern and compared this to the presentation of a target radial frequency pattern which could alter in one of two directions for each visual feature (e.g. clockwise versus anticlockwise). Participants responded to indicate the direction of change. The difficulty of the task increased with every correct answer within a Bayesian staircase. B) reflects an individual trial within the selective versus distributed paradigm. Participants made a yes/no 'same/different' response based on the presence/absence of a change in the attended feature(s).

The averaged 75% correct detection thresholds of changes in each of the three stimulus features gained via previous Bayesian staircases (described above) were used in this experiment to ensure the changes in stimulus features were calibrated to each participants' level of change perception and to ensure we controlled for the level of attention across participants- maintaining an equal difficultly of task for all features and participants. These 75% correct detection thresholds were doubled for each participant, and these values specified the amount change in each stimulus feature. Unlike the threshold staircases, to help ensure the task remained a manageable difficulty, stimulus features only altered in one direction (anticlockwise orientation, increased contrast and 'spikier' shape).

The frequency of these feature changes was manipulated in two different trial types. Previous literature contains examples of constant change distributions across conditions; keeping the number of changes consistent across selective and divided blocks (e.g. Hahn et al., 2008). It also contains examples of equal response distributions; where the number of signal trials are kept constant across blocks (e.g. Corbetta et al., 1990). Both equal and unequal change conditions were tested in this experiment. The unequal change condition involved keeping the number of stimulus changes constant across selective and divided featural attention blocks. In these blocks, the probability of change in any stimulus feature was 20%, regardless of the attentional focus block. In selective attention blocks, where participants are required to respond to changes in only one stimulus feature (e.g. respond to orientation), approximately 20% of trials require a 'different' response from

the participant. However, in the divided attention condition, where participants are asked to respond to changes in any stimulus feature, participants should respond ‘different’ on approximately 60% of trials, hence the number of trials containing a signal are not equal across selective and distributed attention blocks.

Conversely, the equal change condition involved keeping the number of signal trials constant across the different attentional blocks. In the selective attention conditions, the number of potential trials to respond to was still approximately 20%, as the change in any one stimulus feature remains at 20% probability, so although 60% of trials were changing in stimulus feature, only 20% of changes matched the feature of attentional focus. However, in the divided conditions, each stimulus had a 6.66% probability of change, such that overall, the number of changes in all stimulus features (the number of trials requiring a ‘different’ response) was approximately 20%, matching the number of signal trials in the selective conditions- an equal distribution of stimulus changes requiring a response across the conditions.

Regardless of the distribution of stimulus change condition, participants viewed the target RF pattern for 200ms, and were then presented with the same single central attention-specific fixation letter for 800ms, in which they were required to make same-different responses in respect of their directed attentional focus (‘A’ (different) or ‘L’ (same) keyboard responses). The attended feature was randomly allocated in each selective attention block

and the order of equal and unequal change distribution runs was counterbalanced across participants.

Participants completed 50 trials per block, and there were 4 blocks per run (2 selective and 2 divided), hence a total of 200 trials per run (100 selective and 100 divided). At the end of each block (50 trials) participants were able to take a break and initiate the next block of testing themselves, by pressing the spacebar. In a single testing session, participants completed 4 runs (2 equal and 2 unequal), such that each participant contributed 200 selective and 200 divided trials for equal and unequal change distribution conditions respectively. Each run lasted 5 minutes, such that an entire testing session lasted approximately 30 minutes. After completing all 4 runs, participants were debriefed as to the aims of the experiment.

3.4 Results

As a result of the complete randomisation of selective feature blocks, some participants were missing data from a single selective attentional focus condition. For these participants, we created simulated response blocks via bootstrapping, selecting 50 trials with replacement from all trials for that condition concatenated across all participants. Hence, we select trials from multiple participants for a single condition and combine these to create a simulated set of block responses. This ensured every participant had response data for all three selective attentional conditions.

Data for each participant were concatenated across runs for the equal and unequal analyses separately. Across these 400 trials, data were split into individual attentional focus conditions (orientation, contrast, shape and divided). We removed data from any trial where the response time was less than 100ms, taken to be the minimum plausible human reaction time to the stimulus). This resulted in the removal of a maximum of one trial in each condition. Additionally, any trial where the participant failed to make a response was also removed from analysis.

For each condition, we then calculated the mean reaction time for (across all condition repetitions) for participant responses. We then also calculated loglinear d prime (d') for each condition (Hautus, 1995). We chose to use the calculation of loglinear d' to account for the issue of an infinite d' value in instances where the number of hits or false alarms is zero, which infrequently occurred in our dataset (as noted in Stanislaw and Todorov, 1999). This involved incrementing the number of hits and false alarms for each condition by 0.5, and the number of signal and noise trials by 1. We were then able to calculate a participant accuracy score, indicated by loglinear d' for each selective and divided attention condition.

However, in the selective versus divided attention literature, there have been two differing strategies employed regarding the distribution of stimulus change occurrences across the selective and divided attention conditions. Some researchers have employed equal numbers of possible response trials in both the selective and divided conditions, such as Corbetta et al., (1990).

However, other researchers have opted to have unequal numbers of response trials between these two condition types (e.g. Hahn et al., 2008). Hence, we first wished to identify any differences in loglinear d' or reaction time across these unequal and equal analysis pipelines.

Table 3.1 No significant difference in loglinear d prime between equal and unequal change distribution analyses for any attentional focus condition.

	Equal mean (d')	Unequal mean (d')	Statistic (test)	Adjusted *p*
Divided	1.20	1.20	-0.03 (Z)	.977
Orientation	1.70	1.63	0.09 (Z)	.977
Contrast	1.72	2.07	-2.05 (T)	.206
Shape	1.69	1.73	-0.23(T)	.977

We conducted multiple paired-sample t-tests, comparing equal and unequal loglinear d' scores for the divided, orientation, contrast and shape attention conditions individually (as displayed in Table 3.1). We repeated the same analysis performing multiple paired-sample t-tests, comparing each participant's mean reaction time in the divided, orientation, contrast and shape conditions across equal runs (see Table 3.2). Of the 8 comparisons, 3 had a non-normal distribution of one of the data conditions (Shapiro-Wilk $p < .05$). In these instances, we instead performed the non-parametric paired-sample Wilcoxon signed-rank test. For both loglinear d' and reaction time analyses, we performed a Benjamini-Hochberg correction on the resulting significance values (4) to control for the increased risk of a type one error with multiple comparisons. Statistical analysis revealed no significant difference in loglinear d' between the equal and unequal run types in any attentional focus condition ($p > .05$). We also identified no significant

differences in the mean reaction time between equal and unequal run types for any selective or divided attention condition ($p > .05$).

Table 3.2 No significant difference in mean reaction time between equal and unequal change distribution analyses for any attentional focus condition.

	Equal mean (RT)	Unequal mean (RT)	Statistic (test)	Adjusted *p*
Divided	355.22	371.41	-1.97 (T)	.243
Orientation	324.21	332.61	-0.83 (T)	.553
Contrast	321.78	341.23	-1.51 (Z)	.260
Shape	345.33	342.28	0.25 (T)	.807

As we identified no significant difference in participant accuracy or reaction time between the equal and unequal conditions, we report only the outcomes of the equal response trial distribution runs in subsequent analyses. This matches the distribution of response trials employed in the original Corbetta et al., (1990) experiment we aim to replicate, in which both the divided and selective attention conditions contained approximately 50% of trials requiring a ‘different’ response.

In this experiment, we aimed to establish if utilising a relatively novel RF pattern visual stimulus, we could replicate the well-established findings of greater accuracy and faster response during selective attention, opposed to when attention is divided across multiple stimulus features (Corbetta et al., 1991; Corbetta et al., 1990). To investigate this, for each participant, we calculated the average loglinear d’ score across orientation, contrast and shape conditions to provide a single averaged selective attention accuracy score. These averaged selective attention loglinear d’ values were then

compared with each participant's corresponding accuracy in the divided attention condition. A paired-sample Wilcoxon signed-rank test revealed participants were significantly better at detecting changes in the attended stimulus in the selective versus divided attention condition (Z (23) = 3.71, *p* <.001) (See Figure 3.2A). Additionally, we wished to identify if this difference between selective and divided attention conditions was also evident in participant's response times. As before, we calculated the average response time across orientation, contrast and shape conditions to provide a single averaged selective attention reaction time score. This was then compared with participants' reaction time in the divided attention condition. A paired-sample t-test revealed participants took significantly longer to respond in the divided attention condition than compared to the averaged selective attentional focus condition (T (23) = -4.91, *p* <.001) (See Figure 3.2B).

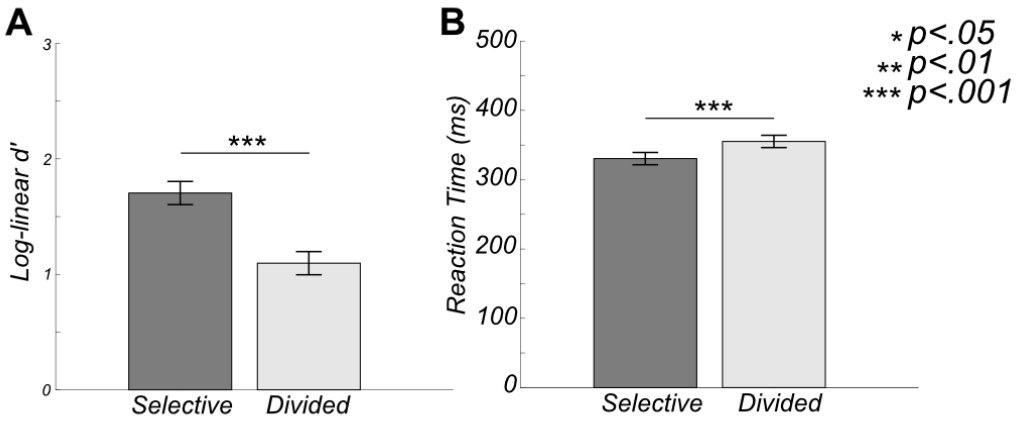

Figure 3.2 Significantly greater participant accuracy (indexed by loglinear d') (A) and faster reaction times (B) in selective than distributed attentional conditions. Error bars reflect +/- 1 standard error.

Furthermore, we wished to investigate potential differences in participant response across the three individual selective attention conditions, in order to

validate our use of participants' feature-specific detection thresholds as a method of calibrating task difficultly and associated levels of attention, for use in later fMRI experimentation. We conducted two one-way repeated measures ANOVAs (for loglinear d' and reaction time respectively) to identify potential differences in response between the three selective conditions. Analyses identified no significant difference in participant accuracy (loglinear d') across the orientation, contrast and shape conditions (F (2,46) = 0.01, p = .987) (see Figure 3.3A). For the reaction time comparison, we identify a significant difference in mean reaction time across the three selective attention conditions (F (2,46) = 4.25, p = .020). However, Bonferroni-corrected post-hoc tests reveal no significant difference between any pairwise comparison of conditions (p > .05) (see Figure 3.3B).

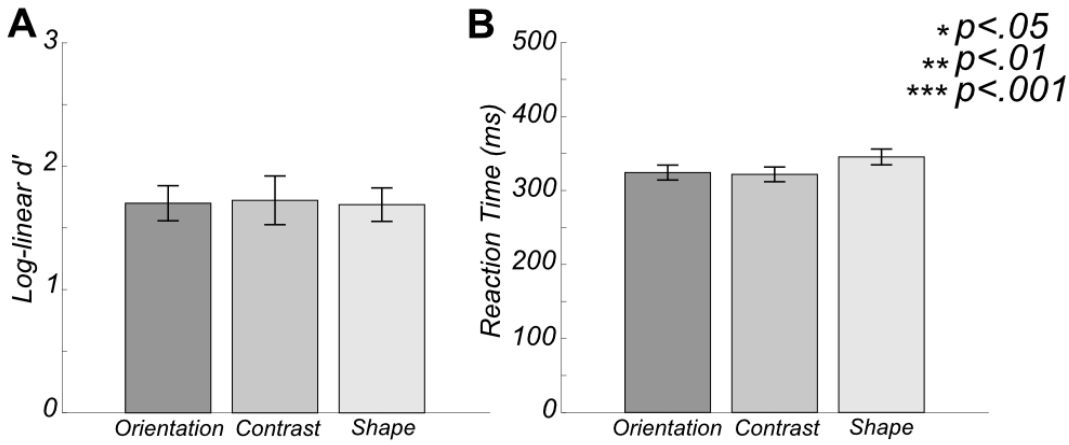

Figure 3.3 No significant difference in loglinear d' (A) or reaction time (B) between selective attention directed toward orientation, contrast or shape. Error bars reflect +/- 1 standard error.

3.5 Discussion

Here, we used a relatively novel visual stimulus, the RF pattern, to probe attentional mechanisms. We aimed to replicate robust findings of the selective versus divided attention paradigm to validate the use of the RF pattern as an effective stimulus to probe attentional modulation. We replicate these well-established findings here. Firstly, we identify significantly greater change detection accuracy across participants during selective (single-feature) attentional focus versus divided (multiple-feature) conditions. This supports our initial hypothesis and replicates the findings of a wealth of previous research (e.g. Corbetta et al., 1990, 1991).

We additionally demonstrate that participants took significantly longer to make a yes/no judgement when attention was divided across multiple stimulus features, versus when directed toward a single visual attribute. Again, these findings confirm our hypothesis and support the findings of previous literature identifying reduced sensitivity and longer response times with attention divided across multiple features of a stimulus (e.g. Ninio and Kahneman, 1974; Hahn et al., 2008). This demonstrates that our results are not reflective of a speed-accuracy trade off. Participants in the selective attention conditions were both *faster* and more *accurate* than in their responses within the divided condition. Hence, we provide support for the notion that manipulation of attentional focus can influence behavioural performance in a visual psychophysics task (Corbetta et al., 1990). Our replication of a robust finding in the attentional literature provides a clear

validation for the use of RF patterns as effective stimuli to probe visual attention processes.

Our results support the notion of attention as a mechanism with limited capacity and its division across multiple features impedes the amount of attention we are able to direct to any single visual stimulus (Broadbent, 1958). Our results provide no indication of the cortical processes involved in these behavioural differences, however previous literature demonstrates featural attention increases activation in regions of the visual cortex associated with processing that stimulus attribute (e.g. Corbetta et al., 1991; Beauchamp et al., 1997; O'Craven et al., 1997; Chawla et al., 1999; Treue and Martínez Trujillo, 1999; Huk and Heeger, 2000; Liu et al., 2003; Schoenfeld et al., 2007).

Additional research has identified increased neuronal/fMRI activity during divided attention versus when attention is directed toward a single visual feature (Corbetta et al., 1991; Herath, Klingberg, Young, Amunts, & Roland, 2001; Johnson & Zatorre, 2006; Loose, Kaufmann, Auer, & Lange, 2003; Nebel et al., 2005; Rees, Frackowiak, & Frith, 1997; Vandenberghe et al., 1997; Weerda, Vallines, Thomas, Rutschmann, & Greenlee, 2006). However, further research has failed to identify the existence of functionally-specific cortical activity in divided attention (Hahn et al., 2008). Future research could investigate potential differences in inter-regional connectivity between the selective and divided paradigms, as perhaps complex tasks involve reduced correlation between regions important for the successful

change detection of specific visual features. This decorrelation between regions has been identified to have beneficial effects on task performance in cognitive flexibility literature (e.g. Dajani and Uddin, 2015; Spadone et al., 2015; Vatansever et al., 2016; Li et al., 2017).

Additionally, we identified no significant differences in participant accuracy or reaction time across selective attentional focus conditions (attention directed to orientation, contrast or to shape). This is evidence of the success of our method of calibrating feature- and participant- specific task difficultly with previously collected participant threshold detection scores for each low-level stimulus feature. This lack of significant difference across conditions gives strength to the notion we successfully equated task difficultly across conditions, importantly suggesting the level of attention employed should be highly similar for each participant across the selective attention conditions.

In conclusion, we replicate the findings of a well-established attentional paradigm with the use of a novel highly controlled visual stimulus, substantiating it's use as an effective stimulus in probing attentional modulation. We also demonstrate the success of a method calibrating the selective versus divided task to each participant's perceptual experience. This study provides support for the use of the RF pattern stimulus and our method of calibrating task difficulty in our future fMRI studies.

4. Feature-Specific Patterns of Attention and Functional Connectivity in Human Visual Cortex

4.1 Abstract

Attending to different features of a scene can alter the responses of neurons in early- and mid- level visual areas but the nature of this change depends on both the (top down) attentional task and the (bottom up) visual stimulus. One outstanding question is the spatial scale at which cortex is modulated by attention to low-level stimulus features such as shape, contrast and orientation. It is unclear whether the recruitment of neurons to particular tasks occurs at an area level or at the level of intra-areal sub-populations, or whether the critical factor is a change in the way that areas communicate with each other. Here we use functional magnetic resonance imaging (fMRI) and psychophysics, to ask how areas known to be involved in processing different visual features (orientation, contrast and shape) are modulated as participants switch between tasks based on those features while the visual stimulus itself is effectively constant. At a univariate level, we find almost no feature-specific bottom-up or top-down responses in the areas we examine. However, multivariate analyses reveal a complex pattern of voxel-level modulation driven by attentional task. Connectivity analyses also demonstrate flexible and selective patterns of connectivity between early visual areas as a function of attentional focus. Overall, we find that attention alters the sensitivity and connectivity of neuronal subpopulations in individual

early visual areas but, surprisingly, not the univariate response amplitudes of the areas themselves.

4.2 Introduction

Attention directed toward spatial locations or visual features can influence both behaviour and physiology. In general, attention tends to increase the relative sensitivity of neurons representing the attended location or feature (e.g. Duncan and Humphreys, 1989; Corbetta et al., 1990; Luck et al., 1997; Reynolds et al., 2000; Martínez-Trujillo and Treue, 2002; Serences and Boynton, 2007), rendering participants faster and more sensitive when detecting changes at attended locations, or in pre-cued features (Corbetta et al., 1990; Posner et al., 1980; Posner & Cohen, 1984).

Given many early visual areas have been identified as having intrinsic preferences for particular features (for a review, see Kanwisher, 2010), it might seem obvious that attention towards those features would modulate activity in entire areas. For areas with high specificity for a single feature, this does seem to be the case. fMRI BOLD responses in hMT+ increase during attention toward a visual motion stimulus (Huk et al., 2001; O'Craven et al., 1997). Increased hV4 activation has been identified in response to attention to chromatic stimuli (Corbetta et al., 1991, Chawla et al., 1999; Liu et al., 2003; Schoenfeld et al., 2007) and attention to faces and places can drive robust responses in the fusiform face area (FFA) and parahippocampal place areas (PPA) respectively (O'Craven et al., 1999).

However, recent work suggests that attention acts primarily through a gain control mechanism, whose primary effect is to alter inter-neuronal noise correlation (Cohen & Maunsell, 2009). If this is correct, we might expect to see relatively little attentionally-driven change in overall activity in earlier visual areas. While individual neuronal sub-populations might become more or less correlated in their trial-to-trial responses, normalization would serve to reduce any long-term response differences (Verhoef & Maunsell, 2017).

This hypothesis is consistent with two findings from the literature: first that early visual areas exhibit relatively little overall change in activity when subjects switch low-level visual tasks (Brouwer & Heeger, 2009; Kamitani & Tong, 2006; Seymour, Clifford, Logothetis, & Bartels, 2009, 2010; Song, Rowland, McPeek, & Wade, 2011). Secondly, that attentional modulation is detected readily by EEG, which is sensitive to the level of correlated noise in large-scale neuronal populations (Di Russo, Martínez, & Hillyard, 2003; Martinez et al., 2001; Verghese et al., 2012; Wang & Wade, 2011).

If attention drives changes in the relative sensitivity and noise correlation of neuronal sub-populations, it should alter activity at the level of feature maps in each visual area. These changes might be detected using multivariate fMRI methods, which in the past, have enabled researchers to decode participants' featural attentional focus from fMRI BOLD activity (Brouwer & Heeger, 2009; Freeman et al., 2011; Kamitani & Tong, 2005, 2006; Song et al., 2011).

Finally, attention can change patterns of activity *across* areas as well as in a single region. Such changes could act at the level of input or output layers, altering not just the sensitivity of neurons performing within-area computations, but also the type of information passed between areas. In support of this hypothesis, it is clear that brain-wide changes in functional connectivity are associated with different mental states (Fox et al., 2005; Greicius, Krasnow, Reiss, & Menon, 2003) and, more specifically, previous literature has identified attentionally-driven changes in patterns of connectivity amongst regions involved in visuospatial tasks (Fox et al., 2005; Gao, Gilmore, Alcauter, & Lin, 2013; Spadone et al., 2015).

Here, we looked for evidence of differential *univariate* or *multivariate* responses in regions-of-interest (ROIs) in the visual cortex, as well as differences in inter-areal *connectivity*. In the multivariate case, we also asked if successful classification was driven by coarse-scale topographical maps or fine-scale voxel-level sensitivities.

Surprisingly, we found very few univariate differences in the regions we examined. However, voxel-level analyses revealed above-chance decoding of attentional focus in *all* visual ROIs, with successful decoding driven by fine-grain participant-specific differences in voxel activity. Connectivity analyses revealed greater connectivity between visual regions during passive viewing than during feature-specific directed attentional focus, and these feature-specific connectivity patterns changed as a function of the ROI subset examined.

4.3 Methods

4.3.1 Participants

12 participants were recruited (8 female, mean age 25 years). All participants had normal or corrected-to-normal vision and provided informed consent. Ethical approval for the study was granted by the University of York Department of Psychology and York Neuroimaging Centre ethics boards. Participants completed two 1.5-hour scanning sessions as part of this experiment, during which we collected high resolution anatomical scans, population receptive field mapping (pRF) and attentional modulation data.

4.3.2 Experimental Design

4.3.2.1 Behavioural Psychophysics

To control for difficulty and levels of attention, prior to fMRI scanning, each participant completed 1 hour of psychophysical testing, to identify their 75% correct detection thresholds for each visual feature (orientation, contrast and shape). Prior testing revealed no significant differences in difficulty (indexed by loglinear d') or reaction time across the attentional focus conditions $F(2,22) = 0.61$, $p = 0.553$, $F(2,22) = 1.36$, $p = 0.275$ (no significant Bonferroni-corrected post-hoc tests) respectively. This ensured task difficulty, and associated attentional effort was constant across attention conditions in our fMRI experiment.

Stimuli were presented on a ViewPixx monitor (120Hz, 1920x1220 pixels resolution) at 57cm viewing distance. Stimulus presentation was performed

on a Shuttle XPC SZ87RG high-end graphics system with an Intel Core i7-4790K processor at 4GHz and a NVIDIA GeForce GTX970 graphics card with 4GB DDR5 memory. All stimuli and experimental procedures were controlled by Psychtoolbox 3.0.12 routines (Brainard, 1997; Pelli, 1997).

To estimate detection thresholds, we used a Bayesian adaptive staircase design (Watson & Pelli, 1983). Feature-specific initial estimates of threshold (0.3 radians orientation, 30% contrast, 0.08 radial amplitude modulation/shape) were provided with 0.5-unit standard deviation. Our stimuli were variants of a radial frequency pattern (RFPs): radially modulated circular contours (Ivanov & Mullen, 2012; Wilkinson et al., 1998; Wilson & Wilkinson, 1997).

This stimulus allowed for an investigation of shape processing and reflected naturalistic stimuli to a greater extent than traditional Gabors (Lawrence et al., 2016; Vernon et al., 2016). Reference RF pattern stimuli had a 2.0° average radius, 0° orientation and 50% contrast (see Figure 4.1A).

Attended features were cued by a central white fixation letter present across trials (‘O - orientation’, ‘C - contrast’ or ‘S - shape’). Trials began with a 200ms presentation of a grayscale reference stimulus against a mean luminance background. The reference RF pattern stimuli had a constant 2.0° average radius, 0° orientation and 50% contrast (see Figure 4.1A). For all visual features, the testing of participant’s thresholds were modulated above and below this reference value, such that participants were responding to increments of change within the stimulus feature. For example, within the

contrast task, a reference stimulus of 50% contrast was always presented, and participants responded to relative increases or decreases in contrast from this reference value. This reference stimulus was followed by a 200ms inter-stimulus interval with presentation of the letter-fixation cue only. A target stimulus was then presented for 200ms, with a change in the attended visual feature derived from the initial estimate and individual participant's previous trial performance. For each feature, a stimulus change could be in one of two directions (with equal probability) - clockwise versus anticlockwise orientation, high contrast versus low contrast, 'spiky' versus 'smooth' shape. This was followed by a central black fixation cross for 800ms, during which participants indicated the perceived direction of change between the reference and target stimulus ('A' or 'L' keyboard press). Participants were informed via a toned 'beep' if their response was correct or incorrect. Each participant completed 50 trials (alongside 10 discarded practice trials at the start of each run) to provide an estimated 75% correct detection threshold for each feature. Each feature-specific staircase was repeated three times with a participant's final detection thresholds reflecting the average of these three repetitions. Each run lasted approximately two minutes.

4.3.2.2 fMRI Attentional Modulation

The same 12 participants who completed the psychophysics testing also completed the fMRI experiment. The fMRI experimental design followed a similar procedure to the psychophysics paradigm. However, here we wished to examine the effects of participants directing attention toward changes in individual visual features. To ensure participants maintained a constant

attentional load, they constantly monitored the stimulus for the a near-threshold-level change in an attended feature (ignoring changes in uncued features). In this respect, our fMRI experiment differed from the psychophysics- each visual feature was not modulated on every trial, and changes in multiple visual features could co-occur. The fMRI experiment was also designed such that we presented the same average stimulus (i.e. the probability of a change in each feature was constant across blocks), so that any effects we identify should be driven by differing attentional focus as oppose to stimulus-driven effects. The timing of each trial in our fMRI experiment, however, was identical to the psychophysical testing paradigm.

Prior to each attentional focus block a 1.5s visual cue directing participants' attentional focus was presented (white text; 'ORIENTATION', 'CONTRAST', 'SHAPE' or 'PASSIVE') against a mean luminance background. During passive blocks, participants were instructed to view the stimulus without directed attention and were not required to respond. As before, each trial began with a 200ms presentation of the reference RFP with a white central fixation letter matching the cued feature of attention ('O', 'C', 'S' or 'P'), followed by a 200ms presentation of the fixation cue. This was followed by a 200ms presentation of the target RFP stimulus, which could vary in no, single or multiple visual features with respect to the reference stimulus. The degree of change between the reference and target RFP was double the participant's previously-recorded feature-specific 75% averaged correct detection threshold. Each feature varied in only one direction (anticlockwise orientation, high contrast and spikier shape). For each feature, the target

stimulus differed from the reference on approximately 20% of trials, hence a constant level of visual modulation was present across blocks on average. A central fixation letter was then presented for 800ms, during which participants made a same/different response. Each trial was 1.5s duration, with a block consisting of 10 trials (15s). Each block was followed by a black central fixation cross (7.5s) allowing BOLD signal to return to baseline. Blocks were presented in a pseudo-randomised order, with the randomised four-block cycle presented four times in each run.

4.3.3 Functional Neuroimaging

4.3.3.1 fMRI data acquisition

Visual stimuli were presented using a PROpixx DLP LED projector (Vpixx Technologies Inc., Saint-Bruno-de-Montarvile, QC, Canada) with a long throw lens which projected the image through a waveguide behind the scanner bore and onto an acrylic screen. Presented images had 1920x1080 pixels resolution and 120Hz refresh rate. Participants viewed the stimulus at 57cm viewing distance in the scanner.

A Shuttle XPC SZ87RG high-end graphics system with Intel Core i7-4790K processor at 4GHz and a NVIDIA GeForce GTX970 graphics card with 4GB DDR5 memory matched to the system used in the Psychophysical testing were used to control the fMRI experiment. All stimuli and experimental procedures were controlled by MATLAB 8.5.0 (2016a) in conjunction in Psychtoolbox 3.0.12 routines (Brainard, 1997; Pelli, 1997). During scanning

behavioural responses and scanner trigger pulses were acquired using a fibre-optic response pad Forp-932 (Current Designs, Philadelphia, PA).

fMRI data were collected at the York Neuroimaging Centre using a GE 3T Excite MRI scanner (GE Healthcare, Milwaukee, WI). Structural scans were obtained using an 8-channel head coil (MRI Devices Corporation, Waukesha, WI) to minimise magnetic field inhomogeneity. Population receptive field and attentional modulation data were obtained with a 16-channel posterior head coil (Nova Medical, Wilmington, WI) to improve signal-to-noise in the occipital lobe. Two high-resolution, T1-weighed full-brain anatomical structural scans were acquired for each participant (TR 7.8ms, TE 3.0ms, T1 450ms, voxel size 1 x 1 x 1.3mm^3).

pRF scan sessions consisted of 6.5-minute stimulus presentation runs collected using a standard EPI sequence (TR 3000ms, TE 30ms, voxel size 2.0 x 2.0 x 2.5mm^3, flip angle 20°, matrix size 96 x 96 x 39). pRF parameters were obtained using procedures very similar to those described by Dumoulin & Wandell (2008) (see Figure 4.1B).

Attentional modulation scan sessions consisted of an average of 6, 6:42-minute runs, containing 134 volumes of data, including 3 dummy TRs later removed to allow for the scanner magnetisation to reach a steady-state. 39 slices were acquired in a bottom-up interleaved acquisition order (TR 3000ms, TE 30ms, field-of-view 19.2cm^3, matrix size = 96 x 96, voxel size 2.0 x 2.0 x 2.5mm^3, flip angle 90°).

During the pRF and attentional modulation scanning sessions, one 16 channel coil T1-weighted structural scan with the same spatial prescriptions as the functional scans was acquired to aid in the alignment of the functional data to the T1-weighted anatomical structural scan (TR 2100ms, TE 8.6ms, field-of-view 19.2cm^3, matrix size 512 x 512, voxel size 0.38 x 0.38 x 2.5mm^3, flip angle 90°, 39 slices).

4.3.3.2 fMRI Pre-processing

To improve grey-white matter contrast, the two T1 high-resolution anatomical scans were aligned and averaged together using the FLIRT FSL tool (Jenkinson, Beckmann, Behrens, Woolrich, & Smith, 2012). This averaged T1 was automatically segmented using a combination of FreeSurfer (Fischl, 2012) and FSL, and manual corrections were made to the segmentation using ITK-SNAP (Teo, Sapiro, & Wandell, 1997).

Functional data were analysed using MATLAB 2016a (Mathworks, MA) and VISTA software (Vista Lab, Stanford University). Between- and within-scan motion correction was performed to compensate for motion artefacts occurring during the scan session. Any scans with > 3mm movement were removed from further analysis (no attentional modulation runs were removed on the basis of excessive movement). The Vista rxAlign tool was then used to co-register the 16-channel coil T1-weighted structural scan to the 8-channel coil T1-weighted full-brain anatomical scan. We applied a manual alignment by using landmark points to bring the two volumes into approximate register, followed by a robust EM-based registration algorithm to

fine-tune the alignment (Nestares & Heeger, 2000). The final alignment was visually inspected to ensure the automatic registration procedure optimised the fit. This alignment was then used as a reference to align the functional data to the full-brain anatomical scan. These functional data were then interpolated to the anatomical segmentation.

4.3.3.3 Population receptive field mapping

To probe attentional modulation across the visual system, we focused our analysis on early visual regions with clear feature-specific preferences or organisation related to our stimulus modulations. The discrete cytoarchitecture of V1, consisting of regular blobs and interblobs with differential spatial and contrast tuning might result distinct activation patterns associated with orientation and contrast. (Horton & Hubel, 1981; Livingstone & Hubel, 1988; Song et al., 2011). Regions LO-1 and LO-2 are clearly-defined, retinotopic areas on the lateral surface of visual cortex (Larsson & Heeger, 2006) that have been identified as having a particular significance in both shape and orientation processing (Larsson & Heeger, 2006; Silson et al., 2013). We also selected two additional visual ROIs, a 'ventral' area, hV4 and a 'dorsal' area V3A/B; regions with no clear expected patterns of attentional modulation to summarise patterns of attentional modulation across the visual cortex. For example, hV4 has been implicated in contrast, orientation and shape processing (Dumoulin & Hess, 2007; Ghose & Ts' O, 1997; Sani, Santandrea, Golzar, Morrone, & Chelazzi, 2013).

We additionally noted attentional modulation driven by changes in cued task demands must involve feedback signals from higher cortical areas such as the IPS to lower-level visual regions (e.g. Di Russo et al., 2003; Bressler et al., 2008; Lauritzen et al., 2009). We therefore included the IPS as a separate ROI of particular interest in our analysis of connectivity.

pRF parameters (eccentricity, polar angle and size) were estimated for each voxel using the standard pRF model in mrVista (Dumoulin & Wandell, 2008). Following Wandell et al., (2007) we manually delineated 9 bilateral ROIs: V1, V2, V3, V3A/B, hV4, LO-1, LO-2, MT+ and IPS0 on cortical flat maps on the basis of polar angle reversals and eccentricity for each participant (see Figure 4.1D and Figure 4.1E).

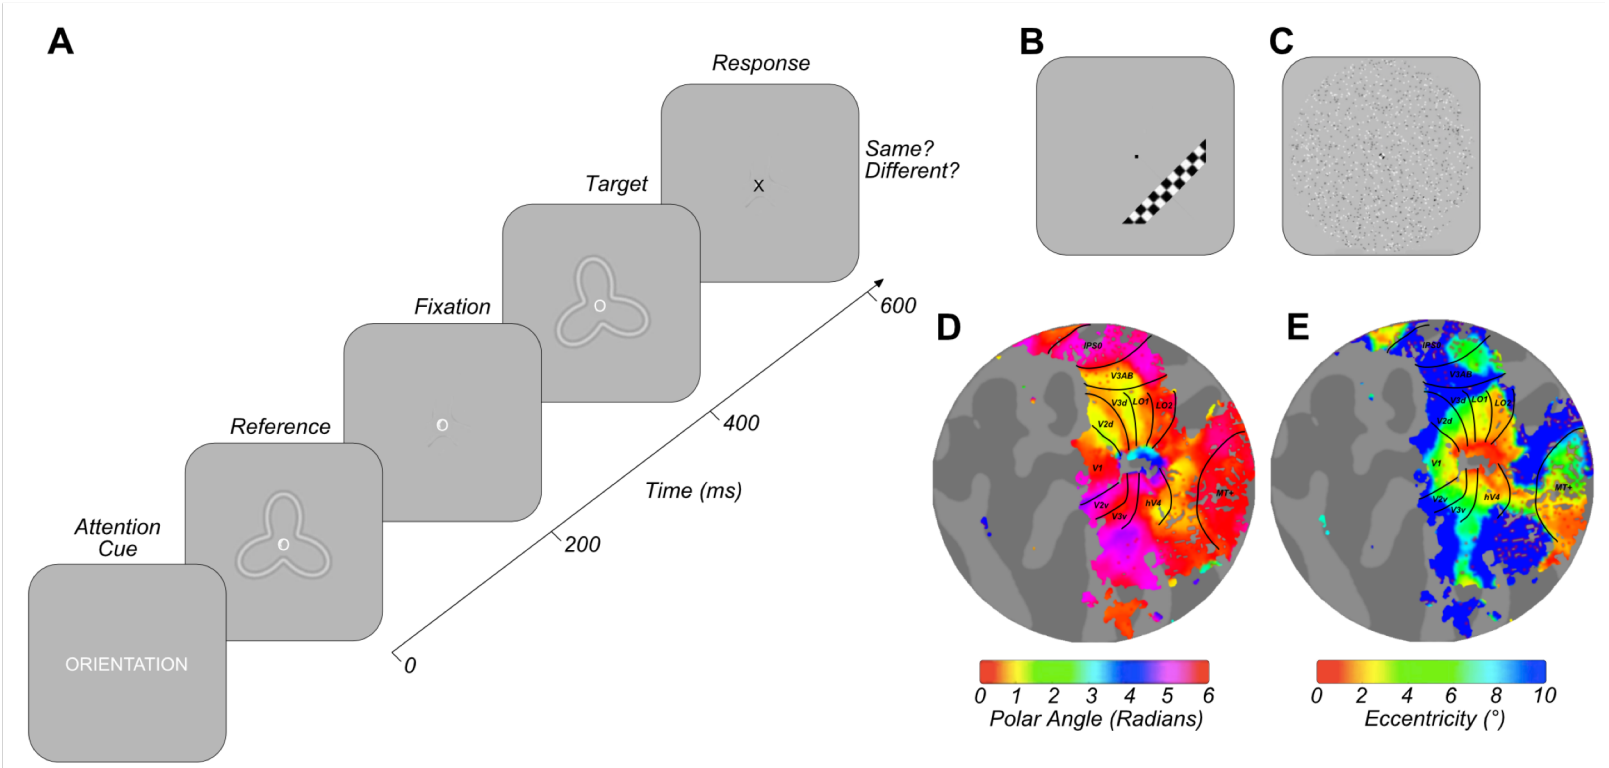

Figure 4.1 Exemplar stimuli and retinotopic maps. A) Example psychophysics/fMRI trial structure B) Example high-contrast phase-reversing checkerboard drifting bar stimulus used to gain estimates of population receptive field size, as in Dumoulin & Wandell (2008). C) Full-field motion stimulus used to help delineate MT+ ROIs, (adapted from Huk et al., 2002; Fischer et al., 2012; Maloney et al., 2013). D & E) Exemplar left hemisphere retinotopic maps with ROI border overlays presented on flattened cortical representations for one participant. Visual ROIs were defined on the basis of eccentricity (A) and polar angle (B) in accordance with the processes described in Dumoulin, Wandell & Brewer (2007).

Seven participants possessed previously-collected MT+ localiser data (the design adapted from (Fischer et al., 2011; Huk et al., 2002; Maloney et al., 2013), defining MT+ on the basis of responses to motion versus static stimuli (see Figure 4.1C). For these participants, we used these motion-defined MT+ ROIs. To ensure consistency and examine reliability across these different MT+ identification techniques, each participants' structural space was transformed to Talairach coordinates using the seven landmarks outlined in (Ryu et al., 2010). A reference spherical 8mm MT+ ROI was created from standardised Talairach coordinates centred on [LH: -47 -76 2, RH: 44 -67 0] (Dumoulin et al., 2000). All defined-MT+ ROIs overlapped with the standardised control MT+ spherical ROIs. Using the same process, a 5mm spherical control ROI of the primary auditory cortex (A1) was created for each participant with standardised coordinates centred on [LH: -49 -20 9, RH: 48 -21 10] (Lacadie, Fulbright, Arora, Constable, & Papademetris, 2008).

4.3.3.4 Attentional Modulation

General linear models (GLM) were implemented to test the contribution of stimulus condition to the BOLD time course (Friston et al., 1998). The default double-gamma haemodynamic response function (HRF) from the SPM8 toolbox was used (Penny, Friston, Ashburner, Kiebel, & Nichols, 2006) and we fit the model to an averaged time course of BOLD signal changed for each stimulus condition by minimising the sum of squared errors between the predicted timeseries and the measured BOLD response.

The first GLM analysed bottom-up stimulus feature change events during which beta weights were obtained by multiple linear regression. Events were classified according to the nature of the stimulus change occurring in a 1 TR (3s) period, regardless of attentional focus (orientation, contrast, shape, no change and multiple change events). Multiple change events reflected two different feature stimulus changes occurring in a single TR. Feature change events (orientation, contrast and shape) reflected when both trials in a TR contained a change in the same feature (e.g. orientation change, orientation change), or a feature change and no feature change (e.g. orientation change, no change). This resulted in 52 to 141 events per stimulus change condition. The second GLM analysed the contribution of attentional focus (15s blocks of directed attention to orientation, contrast, shape or passive viewing). This resulted in a vector of 24 average beta weight estimates for each voxel at the multivariate level.

Participants with >5% average variance explained across the visual ROIs were retained for further analysis (no participants were discarded on the basis of percentage variance explained).

4.3.4 Statistical analysis

4.3.4.1 Univariate Analyses

Feature-specific attentional modulation univariate beta weights (orientation, contrast and shape) were averaged (per participant) and compared with the passive beta-weight through Wilcoxon signed-rank tests for each ROI. Independently for event and attentional modulation analyses, feature-specific

univariate betas were also analysed through one-way repeated measures ANOVAs to assess stimulus-driven and attentional modulation differences in BOLD signal modulation across orientation, contrast and shape conditions.

4.3.4.2 Multivariate Analyses

In each ROI we selected the 100 voxels that explained the largest amount of variance across conditions. These multivariate beta weights were z-scored across voxels and used as input to a support vector machine (SVM) ('LIBSVM' toolbox optimised for Matlab with a radial basis function) (Chang & Lin, 2011) to decode either bottom-up stimulus change or featural attentional focus in two separate analyses using leave-one-out cross validation for each participant. We first assessed multi-class decoding accuracy, supplying orientation, contrast and shape data simultaneously using the 'one-against-one' approach to produce a single classification accuracy score for each participant (Knerr, Personnaz, & Dreyfus, 1990). For the attentional modulation data, pairwise classification was also performed assessing classification accuracy for each combination of attentional conditions (orientation versus contrast, orientation versus shape, contrast versus shape) to identify the driving forces behind successful multi-class decoding. Classification accuracies were then assessed against chance performance through one-sample Wilcoxon signed-rank tests for each ROI.

To investigate the spatial localisation of feature-specific attentional modulation across voxels, additional pairwise SVM classification was performed between each attentional focus condition and the passive viewing

data. The weighted mean of support vectors from these classifications was calculated to provide a feature-specific attention 'preference' for each voxel in an ROI. These voxel 'preference' weights were back-projected onto an interpolated polar grid (6° eccentricity, 360° polar angle, across 500 samples) to reflect voxel activation as a function of each voxels preferred visual angle and eccentricity, extracted from pRF analysis. This voxel-data was then presented visually in polar co-ordinates to reflect this representative visual space. Voxel preference for each visual feature was thresholded at +/-1.7 z-score (p <.05) and displayed to indicate the spatial distribution of attentional modulation. We averaged data over eccentricities between 1.5° and 3.5° of visual angle and plot the resulting average activation as a function of polar angle to provide an intuitive summary of spatial attentional focus. For reference, a standard RF pattern (0° orientation, 0.2 radial amplitude modulation) was overlaid. This back-projection analysis was performed for V1, V2, V3 and hV4 ROIs, as other ROIs lack a high-resolution representation of the full visual field.

4.3.4.3 Timeseries analyses

To quantify functional connectivity between ROIs, participant-specific multivariate timeseries data were extracted (grouping TRs by attentional focus condition) and underwent noise removal (fit and removal of a grand mean) to eliminate scan-to-scan differences in raw amplitude intensity. These multivariate data were then averaged across all voxels in an ROI to provide a single univariate timeseries for each attention condition. Non-parametric Kendall's tau correlations were performed for all pairwise

combinations of ROI (V1, V3A/B, hV4, LO-1, LO-2 and IPS0) – generating a correlation matrix for each condition. To assess the overall similarity of connectivity patterns, the correlation matrix for each attentional focus condition was itself converted to a vector (removing self-to-self correlations) and normalised via participant-specific global mean extraction and Fisher-z transformation. The average (Fisher-z transformed) correlation coefficient was then computed to provide a single number summarising group connectivity in each attentional task. The difference in correlations between conditions was analysed via a one-way repeated measures ANOVA with Bonferroni-corrected post-hoc tests.

Additionally, we investigated ROI-specific *connectivity patterns* as a function of attentional condition. For each of the ROIs; V1, V3A/B, hV4, LO-1, LO-2 and IPS0, we extracted the data reflecting the correlation of this ROI with all others (e.g. V1-V3A/B, V1-hV4, V1-LO-1, V1-LO-2, V1-IPS0) for every participant. We then sampled data from all participants (12 samples with replacement) with a sample reflecting a full complement of ROI-specific correlation data for each attentional focus condition and calculated the mean across these samples. To simulate noise/chance data in this analysis, we took the same 12 samples selected with replacement for each condition, and for each pairwise comparison of conditions, we switched the condition labels approximately 50% of the time, keeping ROI-ROI relationships constant and calculated the average across these scrambled condition-specific datasets. For both the observed and noise data, we calculated the root mean squared error (RMSE) distance between each pairwise combination of condition

vectors as a measure of difference in patterns of connectivity across the ROI profile of interest between different attentional modulation conditions. This process was repeated across 10,000 iterations for each ROI comparison (5 comparisons). Across all iterations, we then calculated the percentage of observed RMSEs for a pairwise comparison falling below the RMSE of the comparable simulated noise distribution. Any percentile below 5% indicates a difference in ROI-specific patterns of connectivity between two attentional modulation conditions significantly larger than predicted by chance (p <0.05).

4.4 Results

4.4.1 Univariate analysis: stimulus-driven events

We first asked whether our near-threshold stimulus modulations altered the BOLD signal at a univariate level. The one-sample Wilcoxon-signed rank test of average BOLD signal across orientation, contrast and shape conditions was significantly different from zero in almost all ROIs examined (Benjamini-Hochberg adjusted p-values reported) (V1: Z(11) = 2.27, p = 0.034, V3A/B: Z(11) = 2.35, p = 0.034, LO-1: Z(11) = 3.06, p = 0.009, LO-2: Z(11) = 2.98, p = 0.009, A1: Z(11) = -2.12, p = 0.041). The only exception was hV4 where averaged bottom-up stimulus-driven activity was not significantly different from zero (Z(11) = 1.88, p = 0.060). Overall, we find that our subtle stimulus modulations do generate a small but significant change in BOLD contrast in most visual areas while auditory cortex ROI (A1) demonstrated small but significant *negative* BOLD responses on average.

We then asked whether this BOLD modulation changed depending on the type of stimulus change. For example, BOLD signal changes in LO-1 and LO-2 have been associated with bottom-up changes in orientation and shape respectively (Larsson & Heeger, 2006; Silson et al., 2013) although the stimulus manipulations in these reports were far larger than those used in this study.

One-way repeated measures ANOVAs Benjamini-Hochberg corrected for the number of ROIs revealed no significant differences in univariate BOLD modulation between stimulus-driven changes in orientation, contrast and shape in any visual ROI ($p > 0.05$) (see Figure 4.2B). Negative BOLD responses were found across all three stimulus-driven change conditions in A1, again with no significant differences related to particular stimulus manipulations.

4.4.2 Univariate analysis: attentional modulation

In the absence of differential BOLD responses relating to stimulus change, we looked for evidence of BOLD modulation as a function of attentional task. Although the majority of visual ROIs examined (excluding V3A/B and hV4) exhibited, on average, greater positive BOLD modulation during directed attention versus passive viewing, Wilcoxon-signed rank tests revealed these trends were not significant (adjusted $p > 0.05$) (see Figure 4.2A). A1 demonstrated a trend towards greater positive modulation during passive viewing versus averaged directed attention, but again this was not significant *(*adjusted $p > 0.05$).

Although on average we found no evidence of attentionally-driven BOLD modulation, it is possible that such changes were present for individual attentional condition types. We therefore conducted one-way repeated measures ANOVAs, Benjamini-Hochberg corrected for the number of ROIs, which revealed no significant differences in attentional modulation across the three featural attentional focus conditions (orientation, contrast and shape) in across almost all visual ROIs examined ($p > 0.05$) (see Figure 4.2C). There was a single exception to this null finding: In LO-2, Bonferroni-corrected post-hoc tests revealed significantly greater attentional modulation for attention to orientation compared to either shape ($F(2,22) = 6.61$, $p = 0.034$, post-hoc $p = 0.011$) or contrast ($p = 0.026$). Again, no significant differences in attentional modulation were evident in A1 ($F(2,22) = 3.29$, $p = 0.112$).

Broadly, our univariate analysis showed that neither the bottom-up stimulus manipulations nor the top-down attentional demands had strong *differential* effects on the BOLD signal in early visual cortex. Area LO-2 was an exception, demonstrating a weak but significant differential response for attention to orientation compared to the other task conditions.

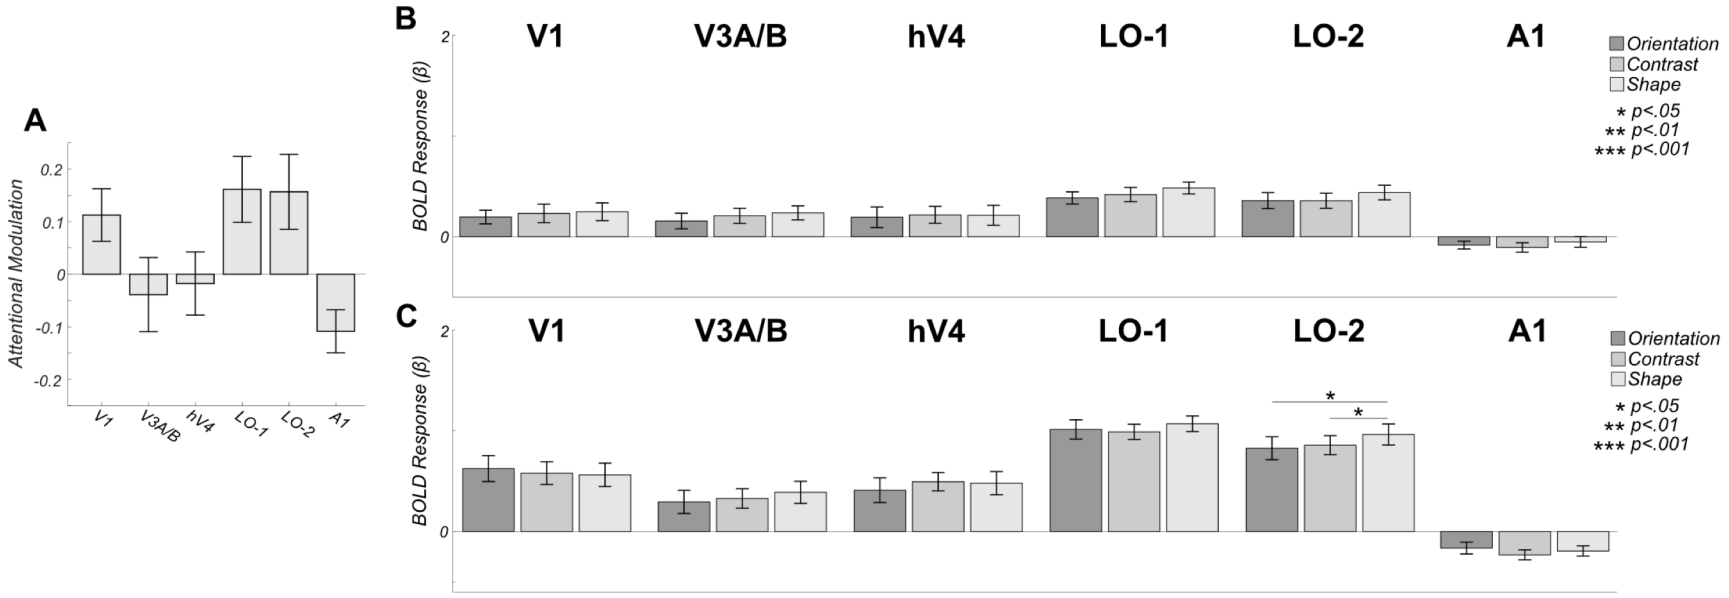

Figure 4.2 Univariate responses to both attention and stimulus modulation are weak. A) Mean BOLD modulation due to attention: positive values reflect greater BOLD amplitude during directed attention (averaged across attentional conditions). Data are shown for 5 visual ROIs: V1, V3A/B, hV4, LO-1 and LO-2 as well as a control area (auditory cortex- A1). No areas exhibit significantly larger BOLD modulation the attend vs passive comparison. B) Mean bottom-up stimulus responses to individual modulation events averaged over all attentional conditions. No areas exhibit differential responses to bottom-up stimulus modulations. C) Feature-specific attentional modulations averaged over all modulation types. Area LO-2 shows a slight increase in BOLD amplitude when subjects attend to shape. All error bars are +/- 1 SEM and significance asterisks indicate Benjamini-Hochberg corrected values.

4.4.3 Multivariate Analysis: pattern classification

The stimulus modulations that subjects detected were extremely subtle. Nevertheless, as shown in Figure 4.2B, we detected BOLD activity with an amplitude significantly greater than zero time-locked to these modulations. We therefore asked if we could decode the identity of these bottom-up responses based on the pattern of responses they elicited in each ROI. To do this, we performed a three-way multivariate pattern classification analysis on the event-related responses, simultaneously classifying orientation, contrast and shape modulations. One-sample Wilcoxon signed-rank tests versus chance (33.33%) Benjamini-Hochberg corrected for the number of ROIs (6), revealed that the type of *bottom-up* stimulus driven modulations could not be decoded at rates significantly above chance in any ROI examined (p >0.05). Mean decoding accuracy ranged from 31.81 to 34.18% (see Figure 4.3A).

The overall BOLD activity linked to attentional modulation was also significantly greater than zero (Figure 4.2C). However, in comparison to the bottom-up results, multivariate pattern classification of attentional state achieved accuracies significantly greater than chance in all visual ROIs examined (V1: $Z(11) = 2.90$, $p = 0.014$, V3A/B: $Z(11) = 3.06$, $p = 0.014$, hV4: $Z(11) = 2.75$, $p = 0.018$, LO-1: $Z(11) = 3.06$, $p = 0.014$, LO-2: $Z(11) = 2.94$, $p = 0.014$) (see Figure 4.3B). Mean classification accuracy ranged from 35.49 to 54.96% across ROIs. As expected, classification was not significantly different from chance in the auditory control ROI A1 ($Z(11)= 1.42$, $p = 0.382$).

To further analyse these significant, attentionally-driven classification patterns, we performed pair-wise pattern classification to determine which attentional states generated different voxel-level responses. Significance was assessed using one-sample Wilcoxon signed-rank tests versus chance (50%), Benjamini-Hochberg corrected for the number of comparisons across ROIs.

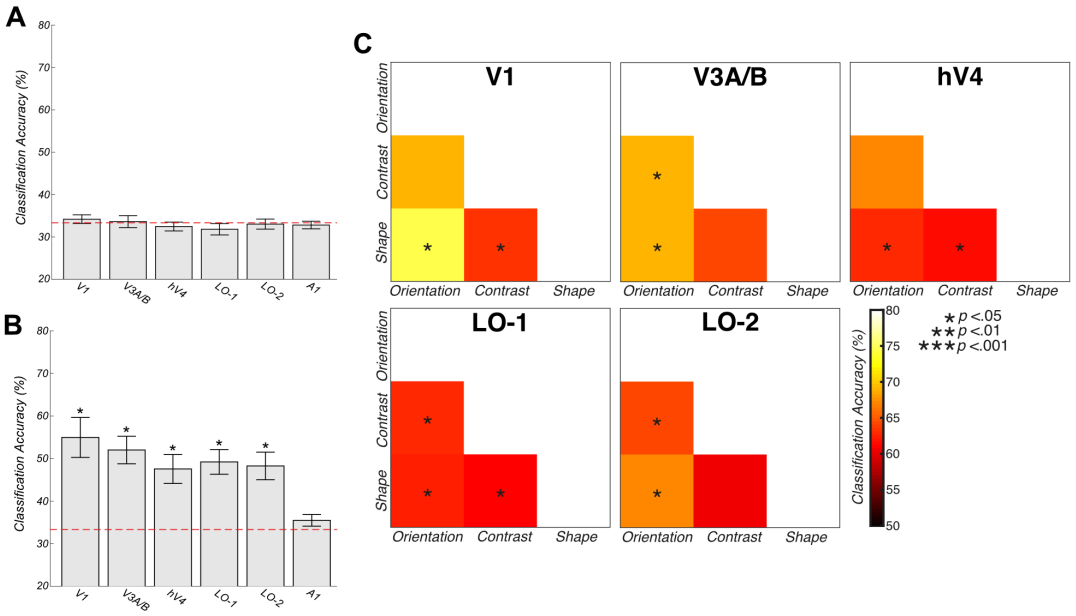

Figure 4.3 Multivariate Support Vector Machine Decoding. Voxel-level responses with individual ROIs are modulated by attentional state, but not by bottom-up changes in stimulus features. A) Overall three-way decoding accuracies. Stimulus change cannot be accurately decoded in an ROI examined. B) Attentional state can be decoded above chance in all ROIs except A1. Error bars reflect +/- 1 SEM. C) Two-way classification accuracies across pairwise combinations of attentional focus (orientation versus shape, orientation versus contrast and contrast versus shape). Voxel patterns in all areas differ significantly between attention to orientation and shape. Significance asterisks indicate Benjamini-Hochberg corrected values.

In V1, we found above-chance classification of attention to orientation versus shape and contrast versus shape. In both V3A/B and LO-1, classification accuracies were significantly above chance for orientation versus contrast and orientation versus shape conditions. In LO-1, the classification of contrast versus shape data was also significantly above chance. In hV4, successful classification of orientation versus shape and contrast versus shape above change level was identified. In LO-2, classification of both orientation versus contrast and orientation versus shape above chance level was identified (see Figure 4.3C and Table 4.1).

Table 4.1 Attention directed toward orientation, contrast and shape can be decoded through two-way classification in visual ROIs.

	Orientation vs. Contrast		Orientation vs. Shape		Contrast vs. Shape	
ROI	Z	*p*	Z	*p*	Z	*p*
V1	2.45	.051	3.06	.028*	2.94	.028*
V3A/B	2.63	.042*	2.94	.028*	2.51	.050
hV4	2.39	.055	2.75	.037*	2.99	.028*
LO-1	2.98	.028*	2.83	.034*	2.67	.042*
LO-2	2.59	.043*	2.93	.028*	2.48	.051

** p <.05, ** p<.01, *** p<.001*

4.4.4 Multivariate Analysis: spatial pattern analysis

Previous reports have shown that some types of bottom-up stimuli generate patterns of retinotopically-biased responses in individual areas at relatively large spatial scales. For example, vertical and horizontal gratings drive voxels near the vertical and horizontal midline respectively (Freeman et al., 2011; Tootell et al., 1998). Other researchers however, suggest many early-visual voxels contain more complex tuning properties than predicted by coarse-scale biases such as radial bias or cardinal orientation selectivity, and these voxels with varying preferences are intermingled in V1 (e.g. Alink,

Walther, Krugliak, & Kriegeskorte, 2017; Kamitani & Sawahata, 2010; Kay et al., 2008; Mannion et al., 2009; Pratte, Sy, Swisher, & Tong, 2016). We asked whether coarse-scale biases, or fine-grain patterns of voxel selectivity might be responsible for the top-down, attentionally-driven multivariate classification results that we found here. Subjects, might, for example, always attend preferentially to a particular part of visual space in order to solve different types of shape or orientation discrimination tasks.

To answer this question, we identified the voxels that were most informative for each type of classification decision and back-projected these into visual space. This allowed us to average spatial patterns of voxel preferences across observers. If all subjects used a common strategy (for example, attending to the vertical meridian) for a particular task, these averages would reveal a consistent non-zero response in this location. If, on the other hand, no changes in the large-scale pattern of responses was generated by attention, these maps would average to zero. Figure 4.4 shows mean values were computed across participants and thresholded (+/- 1.7 z-score, $p < 0.05$) to produce feature-specific attentional modulation maps. No significant patterns of attentional modulation are evident for any featural attentional focus across ROIs, and the 2° annulus of averaged attentional modulation revealed no clear peak of spatial attentional focus as a function of polar angle (see Figure 4.4).

We conclude that while the stimulus modulations that we use to drive attentional tasks may be too subtle to drive different voxel-level BOLD

responses, the different attentional states that subjects employ to detect these changes do select different neuronal populations in early visual areas. Our back-projection analysis indicates that these different populations are not consistent at a large-spatial scale across subjects. Our results are consistent with the hypothesis that subjects are selecting neurons from interdigitated populations that are optimal for particular tasks.

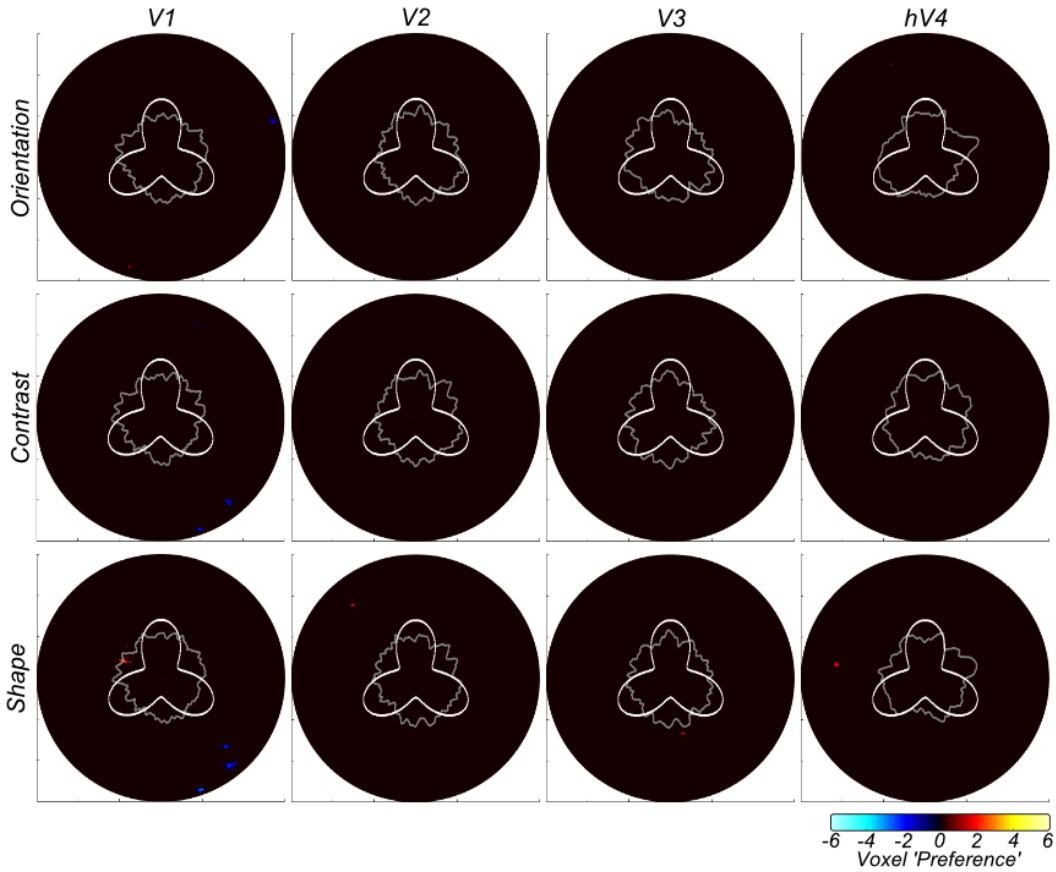

Figure 4.4 Group-averaged voxel feature-specific weights as a function of eccentricity (6°) and polar angle (360°) reveal no large-scale biases in voxel weights across location. The gray annulus reflects averaged voxel modulation at 1° intervals across 1.5-2.5° visual space. Deviations from circularity indicate positive (feature-specific) or negative (passive viewing) preferring clusters of voxels. A radial frequency pattern stimulus overlay is provided for reference. Activation is thresholded at +/- 1.7 z-score ($p < 0.05$).

4.4.5 Timeseries Connectivity Analysis

Attention had relatively little effect on time-averaged univariate BOLD responses in individual visual areas. In our final analysis, we asked whether attention altered the way that individual ROIs communicate with each other. Specifically, asked whether functional connectivity (as measured by the similarity of time courses in different ROIs) might be altered when subjects change their attentional state.

We first performed non-parametric Kendall's Tau correlations between univariate (ROI-averaged) timecourses from each pairwise combination of visual ROIs (V1, V3A/B, hV4, LO-1, LO-2 and IPS0) for each featural attention condition. These inter-ROI correlation patterns represent a 'fingerprint' for each task (see Figure 4.5A). We asked if this overall fingerprint was altered by attentional task and then examined more detailed, pairwise combinations of attentional condition through a one-way repeated measure ANOVA with Bonferroni-corrected post-hoc tests, conducted on normalised and Fisher-z transformed coefficients.

Most strikingly, our analysis revealed significantly greater positive correlation between ROIs during *passive* viewing than any attentional task condition ($F(3,42) = 11.03$, $p < 0.001$, Bonferroni-corrected post-hoc tests, orientation versus passive; $p = .003$, contrast versus passive; $p = 0.006$, shape versus passive; $p = 0.002$). On average, there were no differences in the total level of connectivity between attentional conditions although, as shown below,

individual pairs of areas do show significantly different levels of functional correlation in different attentional tasks.

To identify the specific ROI-ROI connectivity patterns driving these differences in correlation 'fingerprint' across attentional tasks, we calculated the Euclidian distance (RMSE) between pairwise comparisons of attentional task condition (averaged across a subsample of participants for each condition), for both observed and noise (scrambled attention-condition label) datasets, bootstrapped across 10,000 iterations.

Comparing patterns of connectivity associated with IPS0 (IPS0-V1, IPS0-V3A/B, IPS0-hV4, IPS0-LO-1, IPS0-LO-2) revealed a difference between orientation and shape attentional focus conditions significantly larger than expected by chance (p <.001). We also discovered significant differences in IPS0 connectivity between shape versus passive attentional focus conditions (p <0.001). hV4 connectivity across conditions revealed a similar pattern of results, with significant differences in connectivity between orientation and shape attention conditions (p = 0.004) and shape versus passive conditions (p = 0.004). There was an additional significant difference between activation patterns across hV4 connectivity between the contrast and passive conditions (p = 0.015). Analysis across LO-1-correlated ROIs revealed a significant difference in patterns of connectivity between orientation and passive conditions (p = 0.006) and we identified a significant difference between contrast and shape attentional focus conditions across LO-2-correlated ROIs (p = 0.025). However, connectivity to and from V1 (V1-

V3A/B, V1-hV4, V1-LO-1, V1-LO-2, V1-IPS0), and V3A/B (V3A/B-V1, V3A/B -hV4, V3A/B -LO-1, V3A/B -LO-2, V3A/B -IPS0), did not appear to change as a function of task ($p > 0.05$) (see Figure 4.5B).

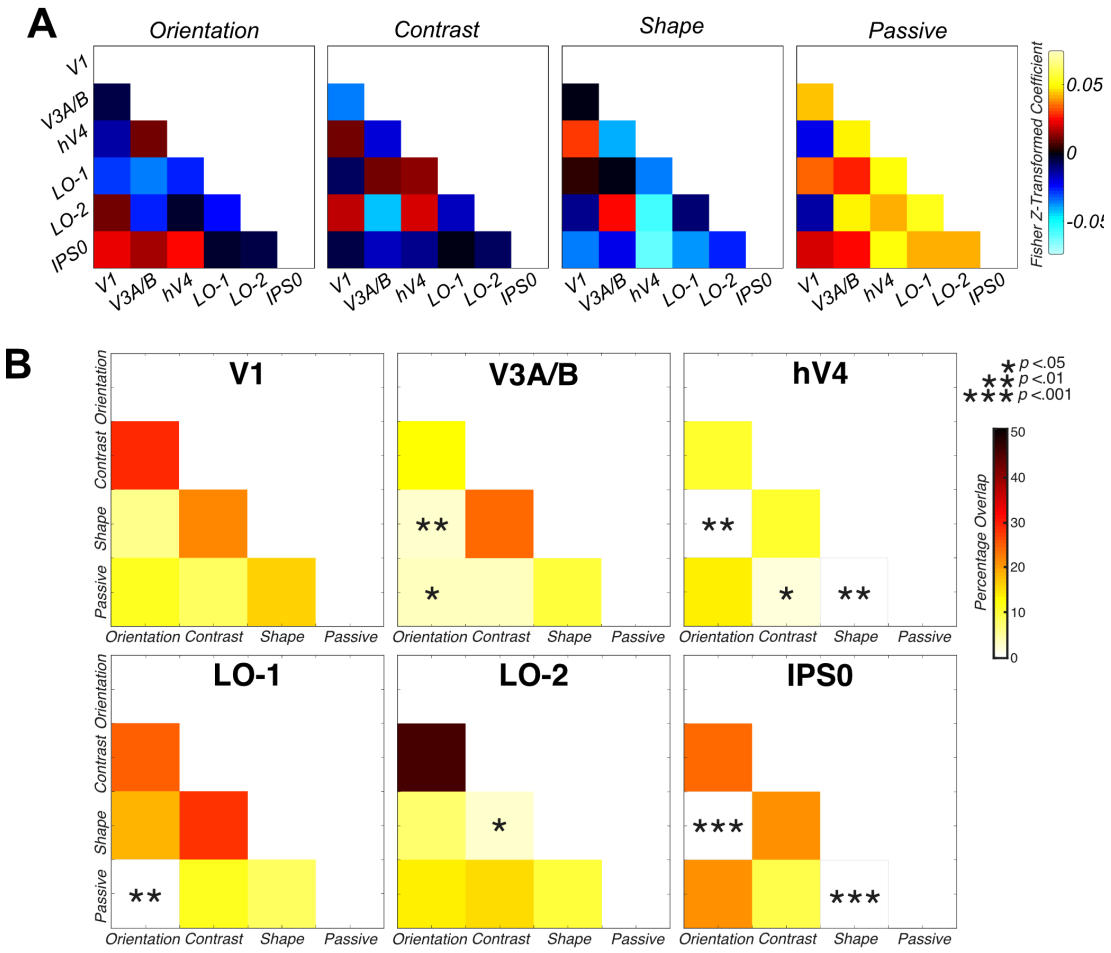

Figure 4.5 Greater overall connectivity during passive viewing compared to directed attention. A) Group feature-specific averaged attentional modulation connectivity values, indicating significantly greater mean connectivity during passive viewing than in the attentional tasks. B) Bootstrapped measures of distance between pairwise combinations of attentional focus conditions, for V1, V3A/B, hV4, LO-1, LO-2 and IPS0 ROIs. The matrices demonstrates the percentage overlap between the distribution of RMSE scores across 10,000 iterations of randomly-selected samples of the observed data and scrambled correlation 'noise' data, for each ROI across multiple pairwise comparisons. Significant overlap (less than 5%) between pairwise combinations of condition are indicated with asterisks.

4.5 Discussion

Here, we show that both *within*-area neuronal activity and *between*-area connectivity change as subjects switch between visual tasks. Perhaps surprisingly, at the univariate level we identify no differences in bottom-up BOLD signal modulation in response to differences in the visual features themselves: brief changes in orientation, contrast and shape all elicited essentially the same pattern of activation at the level of individual areas. Additionally, we identify very few areas that exhibit attentionally-driven changes in averaged BOLD signal, supporting the findings of previous attentional modulation literature (Brouwer & Heeger, 2009; Kamitani & Tong, 2006; Seymour et al., 2009, 2010; Song et al., 2011; Xing et al., 2013). This null finding is not due to neuronal response saturation: We see both robust bottom-up and top-down modulations in all visual regions in our data: areas are modulated by both subtle stimulus changes and differing attentional focus. However, these modulations are not dependent on the nature of the change. Reassuringly, we also identify no significant differences in univariate modulation between attention to orientation, contrast and shape in our control auditory cortex ROI (A1), suggesting any univariate differences we do identify are restricted to visually-responsive regions of cortex. Univariate changes in individual visual areas are therefore relatively uninformative about either bottom-up stimulus parameters or attentional state.

As an example, area LO-1 responds robustly to changes in all three stimulus parameters. This result is intriguing because visual areas are often characterised by their response to bottom-up changes in specific stimulus features and this area has been identified previously as having a causal role

in orientation processing (Larsson & Heeger, 2006; Silson et al., 2013). Our data suggest that measuring univariate BOLD activity in visual cortex as a function of either stimulus type or task may not provide a complete picture of the computations that are being performed in each area. This finding is consistent with a model of visual cortex where areas typically contain multiple, overlapping feature maps in which activity depends both on the stimulus and, to an equal degree, on the attentional state. However, from univariate analysis alone, it is unclear whether there are simply no feature-specific attentional modulation effects, or whether by reducing ROI activity to a single number, potential fine-grain differences in feature-specific patterns of attentional modulation are lost.

To answer this question, we then analysed pattern changes at the voxel level. We showed that response *patterns* in individual visual areas are highly selective for the visual task and can, therefore, be used to decode attentional state. This supports the findings of previous decoding analyses (Brouwer & Heeger, 2009; Freeman et al., 2011; Kamitani & Tong, 2005, 2006; Song et al., 2011).

Classification performance depends both on area and conditions: For example, in V1 successful classification was evident for all pairwise combinations of attentional modulation, supporting literature demonstrating the influence of attention at the earliest stage of the cortical hierarchy (Lauritzen et al., 2010; Serences & Boynton, 2007; Seymour et al., 2009, 2010; Tootell et al., 1998; Verghese et al., 2012). We believe these

successful classifications reflect allocation of top-down attentional focus as oppose to bottom-up stimulus driven effects of selection history as suggested by Awh et al., (2012) and Theeuwes, (2013). In our experiment, each attentional block was repeated the same amount of times, with the presentation order randomised across runs and participants. Additionally, changes in stimulus feature were randomised (in their 20% probability of change occurrence) hence stimulus information and task difficulty remained constant across blocks. Therefore, we believe our findings reflect voluntary allocation of top-down attentional focus, rather than any by-product of differential priming, arousal or selection history across conditions.

The more selective patterns of decoding performance in V3A/B and hV4 suggest that attentional exerts feature-specific effects along both the dorsal and ventral pathways of visual cortex. It clearly demonstrates that both streams possess neuronal populations with feature-specific preferences and that these neuronal populations are responsive to attentional demands (or else inherit attentionally-driven modulations from earlier in the visual pathway). In LO-1, successful classification was evident across all pairwise combinations of visual task, but LO-2 exhibited successful decoding of orientation versus contrast and orientation versus shape only, partial support for the original conclusions of previous literature, with differing patterns of activation across orientation and shape attention (Larsson & Heeger, 2006; Silson et al., 2013). In A1, we demonstrate no classification accuracies significantly greater than would be expected by chance, adding strength to our findings and demonstrating the specificity of our results in the visual

cortex. Overall, these findings extend those from more traditional univariate analyses, and indicates the brain is able to up- and down- weight activity of specific neurons in particular visual regions in a task-dependent manner, in agreement with many previous reports (Treue and Martínez Trujillo, 1999; Martinez-Trujillo and Treue, 2004; Serences and Boynton, 2007; Serences et al., 2009; Verghese et al., 2012).

There is some debate in the literature about the origin of the information that drives fMRI multivariate analyses. Although the ability to decode at above chance rates from a visual area demonstrates that voxel activation patterns in that area depend on the experimental condition, the spatial scale of this change is important for interpreting the results. It has been shown that certain types of information are decoded from changes in voxel activity at a far coarser scale: In particular, researchers have demonstrated that grating orientation is encoded largely by retinotopically-driven, low spatial frequency patterns of response that switch from the vertical to the horizontal midline for vertical and horizontal gratings respectively (Freeman et al., 2011). Yet, if changes occur at the level of individual voxels, they may be driven by neuronal modulations at the level of columnar-scale tuning maps: for example, it has been hypothesised that the ability to decode local radial biases from primary visual cortex is driven by selective activation of orientation-selective neurons in the orientation pinwheels (Mannion et al., 2009). Additionally, recent work has provided evidence for the complexity of voxel tuning profiles in the early visual cortex and demonstrated that experimental task design can influence the conclusion that radial bias is the

only source of orientation information in fMRI signals, for example (Pratte et al., 2016). Hence, global areal maps are unlikely to fully account for the ability to decode orientation signals in early visual cortex (e.g. Alink et al., 2017).

Although the stimuli used here varied only slightly in their physical characteristics, it is possible that MVPA performance was still driven by gross changes in attentional focus. Subjects might have adopted a spatially-driven strategy to solve different featural tasks – focusing on the very top of the radial frequency pattern, for example, to detect orientation changes. In Figure.4 we demonstrate that decoding does *not* depend solely on large-scale biases in sensitivity in visual areas. We find no evidence of a consistent spatial bias in the voxels used for different types of decoding. Although in principle, it is possible that such biases exist on a subject-by-subject (or even trial to trial) basis, our data are consistent with the hypothesis that attention selects sub-populations of relatively fine-scaled maps in individual visual areas. We suggest multivariate pattern classification shown here is derived from changes in activation at a relatively fine scale although we do not rule out coarser-scale topographically-determined responses that may vary across subjects (Kamitani and Sawahata, 2010; Op de Beeck, 2010; Freeman et al., 2011).

Our connectivity results show that visual processing is a dynamic, interactive process that is dependent on the task. While attentional effects have been noted in fMRI research since the late 1990s (Kastner, De Weerd, Desimone,

& Ungerleider, 1998; Tootell et al., 1998), investigating connectivity between regions of visual cortex as a function of task is a relatively novel approach.

Connectivity changes dramatically as a function of attentional state: Overall, we identify significantly greater average connectivity between ROIs during passive viewing than directed attentional focus. This enhanced overall connectivity during passive viewing is similar to the pattern of connectivity observed in the 'default mode' network (DMN) (e.g. Gusnard and Raichle, 2001; Raichle et al., 2001; Christoff et al., 2009), which is abolished by attentional task – and indeed the DMN includes parts of visual cortex. Once a visual task is provided, the brain switches to a more specific connectivity pattern and this pattern is task dependent. We identify different patterns of connectivity when subjects change attentional state for orientation, contrast and shape.

A previous fMRI experiment demonstrated that attention to a particular object category (faces or scenes) lead to strengthened coupling between category-selective visual areas (FFA and PPA respectively) and early visual cortex (e.g. Al-Aidroos, Said, & Turk-Browne, 2012). For example, they demonstrated attention to faces increased the proportion of intrinsic variance shared between regions of the occipital cortex and the FFA. In unpublished research, examining patterns of connectivity across the visual cortex during attention directed towards face stimuli, we also identify increased connectivity between visual areas in comparison with attention directed to low-level visual features. However, here, we identify *reduced* connectivity

between visual areas during attention to low-level visual stimulus attributes in comparison to passive viewing. We believe this reflects the nature of our experimental design. In this experiment, participants directed attention towards highly controlled low-level stimulus attributes in a challenging task, in comparison to attention directed towards relatively higher level, complex features (faces and scenes) that are relatively independent of low-level cues. Here, we also have an explicit passive viewing condition for comparison. Therefore, it is likely we see differing effects of attention on functional connectivity in the visual cortex as a reflection on the type of stimulus attended and the task employed.

We examined these different connectivity patterns across ROIs as a function of task and identify different patterns of connectivity between certain ROIs for different visual tasks. Intriguingly, V1 is *not* one of those areas; correlations between V1 and other areas do not appear to change significantly depending on visual task (although they are reduced overall compared to the passive condition). However, we identify a difference in correlation between attentional tasks in IPS0, supporting a wealth of previous literature indicating the role of IPS in the modulation of top-down attention (e.g. Di Russo et al., 2003; Bressler et al., 2008; Lauritzen et al., 2009; Buffalo et al., 2010). Areas hV4, LO-1 and LO-2 have different patterns of connectivity across ROIs in directed attention. These findings indicate different fingerprints of attention across ROIs: not all ROIs and their connections are modulated the same way in all attentional tasks. Perhaps these regions selectively disengage or de-correlate with other networks during feature-specific directed attentional

focus in order to process these features most effectively as previous cognitive flexibility research suggests (e.g. Spadone et al., 2015; Vatansever et al., 2016; Reineberg et al., 2018). Future research could seek to identify the link between specific patterns of attentional modulation across time in the visual cortex and indices of participant ability, through a network, as oppose to more traditional region-specific approach.

We believe our findings here are not a reflection of gross physiological differences between our four attentional conditions. For example, it is unlikely factors such as extraneous eye movements or changes in heart rate were responsible for differences we see in fine-grain voxel-level activation maps or patterns of functional connectivity. Our experiment recruited experienced observers, who are well-trained in maintaining a constant central fixation in 10-20 minutes of arc (Kowler, 1990). Additionally, our visual stimuli were present for 200ms, shorter than the time needed to make a visual saccade (Carpenter, 1988). Finally, if such gross-scale differences were apparent between attentional conditions, we would expect to see these differences evident at a univariate level. Instead, we identify relatively few significant univariate differences between our featural attention conditions and with passive viewing data, hence, our results likely reflect differences in attentional focus, rather than any consistent differences in gross-scale measures of arousal.

To conclude, we have used a relatively novel approach for investigating top-down attentional modulation signals in visual cortex. We show clear evidence

of attentional modulation from the earliest stage of the visual cortical hierarchy and suggest that directed attention produces local voxel-level changes in activation as oppose to reflecting global topographical organisation of visual regions. Connectivity analyses demonstrate that attention causes a strong *decorrelation* of ROI responses relative to the passive state, which appears to be mediated by top-down signals processed in specific visual regions. This paradigm is a useful tool to probe the influence of a common confound in visual neuroscience, examining activation in response to shifting attentional focus rather than stimulus driven changes.

5. Feature and Colour-Specific Patterns of Attention and Functional Connectivity in Human Visual Cortex

5.1 Abstract

Attention directed towards low-level visual features can alter the activity of neurons across the visual cortex. However, less is understood regarding the effects of attentional modulation when focus is directed towards stimuli of different chromaticities. An outstanding question remains as to whether attentional modulation effects differ in respect to the chromaticity of the attended stimulus. Here, we use functional magnetic resonance imaging (fMRI) and visual psychophysics to investigate the patterns of attentional modulation across the visual cortex in response to attention directed to low-level visual features (orientation, contrast and shape) in combination with differing stimulus chromaticity (red-green, blue-yellow and achromatic). At the univariate level, we find few clear differential top-down attentional responses in the visual areas we examine. However, we do identify significantly greater attentional modulation during attention directed towards red-green stimuli than blue-yellow. Multivariate analyses reveal a complex pattern of voxel-level modulation driven by attentional task. In addition, connectivity analyses demonstrate flexible and selective patterns of connectivity between early visual areas as a function of attentional focus.

5.2 Introduction

Visual attention is a highly important mechanism for successful interaction with our environment. Attention acts as filter, increasing the processing of task-relevant information at the expense of the remaining visual field, to allow for targeted and efficient processing of an otherwise overwhelming amount of information (e.g. Posner, Snyder, & Davidson, 1980). Attention has been demonstrated to increase activation in the visual cortex in regions associated with a particular spatial location in the visual field (Lauritzen et al., 2010; Luck et al., 1997; Moran & Desimone, 1985; Tootell et al., 1998b; Verghese et al., 2012), or to boost the activation of neurons associated with the encoding of an attended visual feature (Corbetta et al., 1990; Martínez-Trujillo & Treue, 2002; Reynolds et al., 2000; Saenz et al., 2002; Serences & Boynton, 2007). For example, colour-preferring areas such as hV4 demonstrate a significant increase in activation during attention directed towards stimulus chromaticity (e.g. Chawla, Rees, & Friston, 1999; Liu, Slotnick, Serences, & Yantis, 2003; Schoenfeld et al., 2007).

Here, we wished to expand upon our previous experiment, investigating attentional modulation effects during manipulation of the low-level visual feature attended, and also of the stimulus chromaticity. The segregation of chromatic information begins with the three cone photoreceptors types in the retina, where information is then transported via the lateral geniculate nucleus (LGN) to the initial locus of cortical processing (primary visual cortex) through distinct magnocellular and parvocellular pathways, which segregate luminance from chromatic information respectively. The

parvocellular pathway contains information from the comparison of long- and medium-wavelength sensitive cone activation (L-M), signalling the red-green dimension of colour, whereas the magnocellular pathway contains information about the sum of L- and M-cone signals (L+M), coding achromatic information (see Wandell, 1995 for a review). A further, separate S-cone (koniocellular) pathway carries information from a third type of cone photoreceptor, sensitive to short-wavelengths of light. This pathway compares S-cone activation with L and M-cone responses (S-(L+M)) to signal the blue-yellow dimension of colour space (Casagrande, 1994; Chatterjee & Callaway, 2003). Each of these pathways project to distinct layers of the primary visual cortex. The primary visual cortex (V1) has also been demonstrated to possess cytochrome oxidase (CO) blobs, considered to have weak orientation tuning, but a robust response to isoluminant colour (Livingstone & Hubel, 1984), though the extent these blobs are explicitly dedicated to colour as opposed to form vision is debated (see Economides, Sincich, Adams, & Horton, 2011).

Beyond these initial synapses in V1, the segregation of chromatic information is less clearly understood. Some suggest segregation of chromatic information may persist higher in the visual cortex, with projections from V1 colour-sensitive blobs thought to input to distinct thin stripe regions of V2 (Levitt et al., 1994; Livingstone & Hubel, 1984). Previous research has also identified area V4, as a specific locus of relatively higher-order colour processing (see Zeki, 1973). This suggestion was supported by evidence from cerebral achromatopsia; macaques with V4 lesions exhibited severe

deficits in colour perception (Meadows, 1974). However, the notion of a complete functional specialisation for chromatic tuning in V4 is also questioned, for example, other researchers found that lesioning macaque V4 did not produce significant negative impacts on colour perception (Heywood, Gadotti, & Cowey, 1992). Hence, the extent strict segregation of chromatic information exists across the visual pathway is unclear. In this experiment, we employed a relatively novel approach, identifying patterns of attentional modulation across the visual cortex in response to shifts in attentional focus, with a particular focus on the investigation of the distributed effects of attention with changing stimulus chromaticity.

Signals from the magno-, parvo- and konio-cellular pathways are combined in the first few synapses in V1. Evidence currently suggests attention exerts different influences on information in the distinct colour channels. For example, Wang & Wade, (2011) identified robust amplitude increases in response to contrast with both achromatic and S-cone-defined gratings, however, they identified no effects of attention on the amplitude or phase of SSVEP responses during attention to S-cone stimuli in any visual area examined. Additionally, Highsmith & Crognale, (2010) find no evidence of attention modulating the amplitude or phase L-M or S-(L+M) chromatic stimuli. These relatively weak attentional modulations of S-cone stimuli provides some support for the work of previous researchers who identify a generally weaker response to S-cone-defined stimuli across the dorsal visual pathway (e.g. Mullen, Chang, & Hess, 2015). For example, Liu & Wandell, (2005) examined responses to chromatically-defined motion stimuli across

the visual cortex, and identified strong responses to all chromatic stimuli in ventral area VO, but a very weak response to S-cone motion in dorsal areas V3A and MT+. Hence, previous literature predicts we should identify different patterns of attentional modulation across visual areas with chromatically-defined stimuli, in comparison to our previous experiment with a purely achromatic stimulus (see Chapter 4).

Previous literature has demonstrated the existence of interactions between low-level visual features and stimulus chromaticity on activation in the visual cortex. For example, Sumner, Anderson, Sylvester, Haynes, & Rees, (2008) found no evidence of different patterns of modulation between attended orientations for luminance, red-green or blue-yellow defined stimuli at the univariate level. However, they demonstrated successful classification of orientation for each of the three stimulus chromaticities at the multivariate level in V1, V2 and V3 early visual areas. These classification accuracies were significantly greater within-group than between-group suggesting different patterns of voxel-level attentional modulation in response to orientation defined by different stimulus chromaticities. Additionally, Seymour, Clifford, Logothetis, & Bartels, (2010) demonstrated specific pairings of colour and orientation could be decoded at the multivariate (but not univariate) level across many early visual ROIs.

The vast majority of attentional modulation studies in the human neuroimaging literature use achromatic stimuli (as noted in Wang & Wade, 2011). Here, we investigate patterns of attentional modulation to low-level

visual features and their associated chromaticities at univariate and multivariate spatial scales. We predict differing patterns of attentional modulation across the three chromatic channels, with a reduced response with attention directed towards S-cone stimuli particularly along the dorsal visual pathway. As in our previous experiment, we identify distinct feature-specific patterns of attentional modulation in multivariate and connectivity analyses. We identify significantly reduced BOLD signal modulation during attention towards S-cone versus red-green defined stimuli at the univariate level, and differences in the patterns of modulation for each chromatic channel in our multivariate and connectivity analyses. We provide support for the varying effects of attentional modulation as a function of both low-level visual feature and stimulus chromaticity.

5.3 Methods

5.3.1 Participants

12 University of York staff and students (8 females, mean age 26.92 years) were recruited for this study. All participants possessed population receptive field (pRF) mapping data, collected in a previous scanning session using standardised protocols (see 2.3.4) and two high-resolution structural scans. Each participant completed 30-minutes of psychophysical testing, and one 90-minute fMRI scanning session, completing an average of 8 attentional modulation runs. All participants had normal or corrected-to-normal vision. Participants provided informed consent and ethical approval was granted by the University of York Department of Psychology and York Neuroimaging Centre ethics boards.

5.3.2 Behavioural Psychophysics

5.3.2.1 Stimulus presentation

Stimuli were presented on a ViewPixx monitor (120Hz, 1920x1220 pixels resolution) at 57cm viewing distance. Stimulus presentation was performed on a Shuttle XPC SZ87RG high-end graphics system with an Intel Core i7-4790K processor at 40GHx and a NVIDIA GeForce GTZ970 graphics card with 4G DDR5 memory. All stimuli and experimental procedures were controlled by Psychtoolbox 3.0.12 (Brainard, 1997; Pelli, 1997).

5.3.2.2 Isoluminance testing

In order to ensure our L-M and S-(L+M) stimuli stimulated only colour and not luminance-sensitive mechanisms for each individual, all 12 participants first completed minimum motion isoluminance testing (see section 2.2) (Anstis & Cavanagh, 1983; Lu et al., 1999). Participants also completed the same isoluminance testing procedure prior to the fMRI scanning session in the scanner bore to ensure stimuli were presented with isoluminance values specific to the display.

Both the display systems used in the psychophysical and fMRI testing sessions were photometrically calibrated using a fibre-optic photospectrometer (Ocean Optics, Dunedin, FL), which measured both the gamma and spectral irradiance of the red, green and blue channels as seen by the eye. To achieve this, the tip of the fibre-optic cable was positioned to match the viewing distance and position of a human observer using a polystyrene mannequin head.

The isoluminance stimuli were initially specified in LMS cone excitation space and transformed to RGB values using the product of the Stockman Sharpe 10° cone fundamentals for the L-, M- and S-sensitive photoreceptors and the spectral power distribution of the RGB phosphors for each eye. In a darkened room, participants viewed a central Gaussian window annulus, in which a sine-wave (2.33 cycles/°, 3.8° diameter) was presented, with a cyclical phase (pulsing back and forth in space) through a range from 0 to π/2 radians in 6° increments, on a mid-grey luminance background and with a white central fixation cross. Participants fixated centrally and increased or decreased the amount of luminance contamination in the stimulus until the minimum amount of flicker was perceived. Audio feedback informed the participants if they had reached the minimum or maximum extent of L+M adjustment. Participants completed the testing procedure three times for L-M and S-(L+M) conditions respectively and an average was taken to produce the average isoluminance setting for each chromaticity used to calibrate the presentation of stimuli in the psychophysical and fMRI testing sessions.

5.3.2.3 Experimental Design

In order to gain each participants' individual threshold detection values for each combination of stimulus chromaticity (L+M, L-M and S-(L+M)) and low-level visual feature (orientation, contrast and shape), the same 12 participants initially completed a 30-minute psychophysical testing session, very similar to the screening procedure employed in our first, achromatic experiment (see Chapter 4). We utilised a RF pattern stimulus (see section

2.1) demonstrated to be an effective stimulus in probing visual attentional mechanisms (see Chapter 3).

This threshold testing used a Bayesian staircase to efficiently estimate participants' feature- and colour-specific detection thresholds (Watson & Pelli, 1983). Initial estimates of mean threshold were provided for each low-level visual feature (orientation 0.3 radians and 0.08 radial amplitude modulation (shape)) with a 0.5-unit standard deviation. Stimulus chromatic contrast was specified on the basis of standardised LMS values, to ensure equal cone excitation across the three stimulus chromaticities (L+M 5%, L-M, 4%, S-(L+M) 15%) (as in Welbourne, Morland, & Wade, 2018). These values were then multiplied by 1.5 to increase stimulus visibility. Initial estimates of contrast detection threshold were specified on the standardised LMS values specified above (5% L+M, 4% L-M and 15% S-(L+M)), which modulated above and below the baseline contrast value.

Participants completed 75 trials, and an additional 10 initial practice trials at the start of each run (discarded from analysis) to provide a 75% correct detection threshold for each combination of low-level visual feature and stimulus chromaticity. Testing order was randomised for each participant. Each participant was presented with a visual attention cue at the start of each staircase to specify the attended feature for each staircase and a central white fixation letter presented on a mid-grey luminance background throughout the run matched the staircase feature (O, C and S). Trials began with a 500ms presentation of a mid-grey luminance screen. This was

followed by a 150ms presentation of a reference radial frequency (RF) pattern stimulus. As in the previous psychophysical staircase procedure (see Chapter 4), this reference RF pattern had a 2.0° average radius, with a constant 0° orientation, colour-specific contrast (specified above) and shape (0.5 amplitude). Chromatic stimuli (L-M, S-(L+M)) were presented with participant-specific isoluminance settings. This was followed by a 600ms inter-stimulus interval, with presentation of a mid-grey luminance background. The target stimulus was then presented for 150ms with the extent of change in the low-level visual feature derived from participant's previous performance using the Bayesian staircase procedure. The target RF pattern differed from the reference in only the attended feature. The direction of change was randomly specified with approximately 50% of trials occurring in each change direction (clockwise versus anticlockwise orientation, high versus low contrast and spikier versus smoother shape) (see Figure 5.1A).

Presentation of both the target and the reference stimuli were preceded with an audible beep to help maintain participants' directed attention. Participants then made a yes/no judgement indicating the direction of change between the reference and the target stimulus ('U' or 'N' keyboard press). The participant's response initiated the next trial. A maximum 900ms 'wait' period was included, in which if a response was not recorded, the next trial commenced. Participants were informed via a toned beep if their response was correct or incorrect. At the end of 75 trials, the Bayesian staircase provided an estimate of each participants' 75% correct detection threshold

for that visual feature. 9 staircases in total were performed, one for each combination of stimulus chromaticity and low-level visual feature, with each staircase lasting approximately 3 minutes.

5.3.3 Functional Neuroimaging

5.3.3.1 fMRI stimulus display

Visual stimuli were presented using a PROpixx DLP LED projector (Vpixx Technologies Inc. Saint-Bruno-de-Montarvile, QC, Canada) with a long throw lens which projected the image through a waveguide behind the scanner bore and onto an acrylic screen. Images were presented at 57cm viewing distance with a resolution of 1920x1080 pixels and a 120Hz refresh rate.

A Shuttle XPC SZ87RG high-end graphics system with Intel Core i7-4790K processor at 4GHz and a NVIDIA GeForce GTX970 graphics card with 4GB DDR5 memory were used to control the fMRI experiment. All stimuli and experimental procedures were controlled by MATLAB 8.5.0 (2016a) in conjunction with Psychtoolbox 3.0.12 routines (Brainard, 1997; Pelli, 1997). During scanning, behavioural responses and scanner trigger pulses (used to synchronise stimulus onset) were acquired using a fibre-optic response pad Forp-932 (Current Designs, Philadelphia, PA).

5.3.3.2 fMRI data acquisition

fMRI data were collected at the York Neuroimaging Centre using a GE 3T Excite MRI scanner (GE Healthcare, Milwaukee, WI). Structural scans were obtained using an 8-channel head coil (MRI Devices Corporation,

Waukesha, WI) to minimise magnetic field inhomogeneity (see Chapter 4 for details regarding structural scan acquisition and pre-processing). pRF and attentional modulation scans were instead collected with a 16-channel posterior head coil (Nova Medical, Wilmington, WI) to improve signal-to-noise in the occipital lobe. At the beginning of both the pRF and attentional modulation scanning sessions, a single 16-channel coil T1-weighted structural scan was acquired with the same spatial prescription as the functional scans was acquired to aid alignment of the functional data to the T1-weighted anatomical scan (TR = 2100ms, TE = 8.6ms, field-of-view = 19.2cm^3, matrix size = 512 x 512, voxel resolution = 0.38 x 0.38 x 2.5mm, flip angle = 90°, 39 slices).

The same 12 participants who completed both psychophysical and isoluminance testing also completed an average of 8, 5:06 minute attentional modulation runs, containing 102 volumes of data, including 3 dummy TRs which were discarded prior to pre-processing, to allow for the scanner magnetisation to reach a steady state. 39 slices were acquired in a bottom-up interleaved acquisition order (TR = 3000ms, TE = 30ms, field-of-view = 19.2cm^3, matrix size = 96 x 96, voxel size = 2.0 x 2.0 x 2.5 mm, flip angle = 90°).

5.3.3.3 Defining regions of interest (ROIs)

Participant ROIs were defined using population receptive field (pRF) mapping data collected prior to the attentional modulation scans. Following Dumoulin & Wandell, (2008) and Wandell, Dumoulin, & Brewer, (2007) we

manually delineated 9 bilateral ROIs (V1, V2, V3, V3A/B, hV4, LO-1, LO-2, MT+ and IPS0) on the basis of polar angle reversals and eccentricity for each participant (see section 2.3.4 and Chapter 4 for a detailed explanation of pRF procedures).

Some participants also possessed previously-collected MT+ localiser data, from which MT+ was defined on the basis of responses to motion versus static stimuli (see section 2.3.5). For these participants, we used the motion-defined as oppose to pRF-defined MT+ ROIs. We ensured consistency between these differing ROI specification methods; we created a spherical MT+ ROI from standardised Talairach co-ordinates for each participant and ensured each MT+ ROI overlapped with the standardised control ROI (see Chapter 4 for full description of this process).

Finally, a bilateral primary auditory cortex (A1) ROI was defined, to use as a control region throughout experimental analysis. This was created through the same process as above, defining a 5mm spherical ROI (for each hemisphere) on the basis of standardised Talairach coordinates centred on; LH: 49 -20 9, RH: 48 -21 10 (Lacadie et al., 2008).

5.3.4 Attentional Modulation Scans

5.3.4.1 Experimental Design

Prior to attentional modulation scanning, participants completed the same minimum motion isoluminance testing as in the psychophysics session in the

fMRI scanner. This was to ensure participants settings of isoluminance were calibrated to each stimulus display setup.

The fMRI experiment followed a similar procedure as employed in Chapter 4. We used a block design, and in blocks, cued participants to attend to a specific low-level visual feature (orientation, contrast, shape) or to passively view the stimulus without attention directed towards any specific visual feature. Each 1.5 second trial in a stimulus block began with a 200ms presentation of the same reference RF pattern as used in the psychophysics testing session on a mid-grey luminance background. A white central fixation letter directed participants' attentional focus towards a single stimulus dimension (O, C, S or P). This was followed by a 200ms inter-stimulus interval, during which the white central fixation letter was presented.

A target RF pattern was then presented for 200ms, which could vary in any combination of the low-level visual stimuli (orientation, contrast or shape), or have no change. The extent of change in each of the features was set as double the participants previously collected 75% correct detection thresholds for each combination of stimulus chromaticity and low-level visual feature to ensure that task difficulty (and associated levels of attention) was consistent across attentional blocks and also across participants. Each visual feature altered in only one direction to ensure the task remained of a manageable difficulty (anticlockwise orientation, high contrast and spikier shape). For each feature, the target RF pattern differed from the reference on approximately 20% of trials, hence a constant level of visual information was

presented across blocks, helping to distinguish attentional modulation effects from any bottom-up stimulus driven activation.

A central fixation letter was then presented for 800ms, during which participants made a yes/no response. Participants were instructed to respond ‘different if the target RF pattern has changed in the attended stimulus dimension (with respect to the reference RF pattern), and to respond ‘same’ if no change in the attended stimulus feature had occurred (regardless of changes in the other stimulus dimensions) (see Figure 5.1B). Each trial lasts 1.5 seconds, and each block contained 10 trials. Each block was followed by a black central fixation cross (7.5 seconds), which allowed BOLD signal to return to baseline. The order of block presentation was pseudo-randomised, such that each of the four feature conditions (in a randomised order) were presented before repetition. The order of stimulus chromaticity was randomised. Each fMRI run contained one repetition of every low-level visual feature and stimulus chromaticity combination (12 blocks).

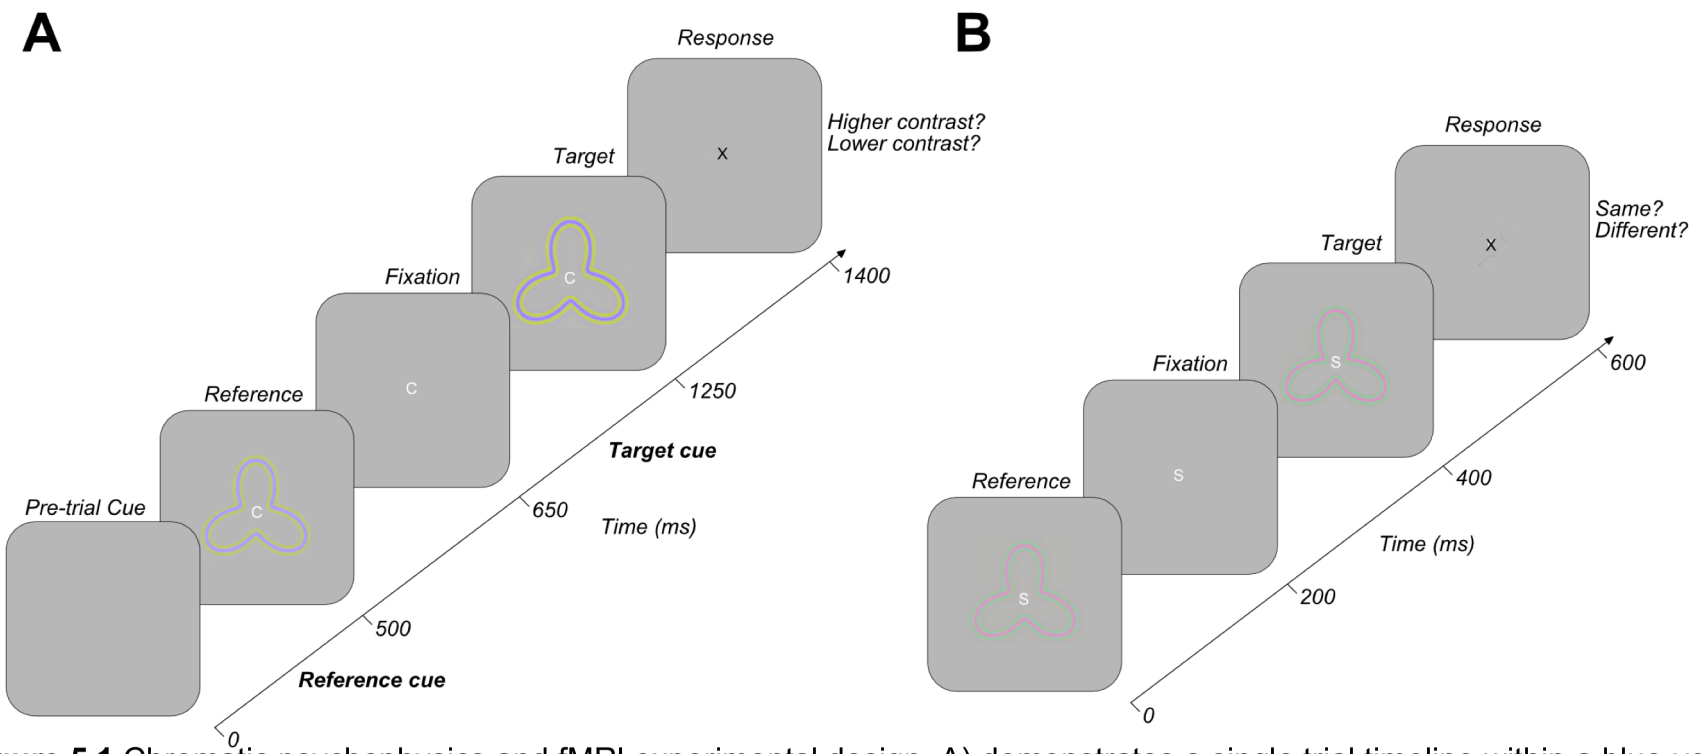

Figure 5.1 Chromatic psychophysics and fMRI experimental design. A) demonstrates a single trial timeline within a blue-yellow contrast detection Bayesian staircase procedure. Participants are cued to the onset of the reference and target radial frequency patterns with an audible beep. Participants are instructed to indicate the direction of the change within the target pattern in comparison to the reference stimulus, which can change in one of two directions on each trial. Trial difficulty increases with the number of correct responses. In B) we demonstrate a single trial structure within the fMRI attentional modulation scans. Participants indicate the presence/absence of a change in the attended feature between the target and reference radial frequency patterns. In both experiments, the cued feature of attention is indicated by the central fixation letter.

5.3.4.2 Data acquisition and analysis

Functional data were pre-processed using the same procedure as used in Chapter 4, using MATLAB 2016a (Mathworks, MA) and VISTA software (https://vistalab.stanford.edu/software/) (Vista Lab, Stanford University). Between- and within-scan motion correction first compensated for any motion artefacts occurring during the scanning session. No run exhibited >3mm movement for any participant. The mrVista rxAlign tool was then used to co-register the 16-channel coil T1-weighted structural scan to the 8-channel coil T1-weighted full-brain anatomical scan. A manual alignment was applied using landmark points to position the two volumes in approximate register. We then used a robust EM-based registration algorithm to fine-tune the alignment (Nestares & Heeger, 2000). The final alignment was checked by eye to ensure the automatic registration procedure optimised the fit. This alignment was used as a reference to then align the functional data to the full-brain anatomical scan, which were then interpolated to the anatomical segmentation of grey and white matter.

Multiple General Linear Model (GLM) analyses were then performed to test the contribution of stimulus condition to the BOLD time course (Friston et al., 1998). We used the double-gamma HRF from the SPM8 toolbox (http://www.fil.ion.ucl.ac.uk/spm/) and fit the model to an averaged time course of BOLD signal, changed for each stimulus condition by minimising the sum of squared errors (RSS) between the predicted time series and the measured BOLD response.

As in Chapter 4, we performed multiple GLM analyses in order to divide our BOLD timeseries data into multiple sets of conditions in order to investigate a wide range of hypotheses. Our first GLM analysed bottom-up stimulus feature change events. As in our previous achromatic analysis, events were classified according to the nature of the stimulus change occurring in a 1 TR (3s) period, regardless of attentional focus or stimulus chromaticity (orientation, contrast, shape, no change and multiple change events). Multiple change events reflected two or more different feature stimulus changes occurring in a single TR. Feature change events (orientation, contrast and shape) reflected when both trials in a TR contained a change in the same feature (e.g. both trials contained an orientation change) or a feature change and no feature change. This resulted in 41 to 172 events per feature change condition.

The second GLM analysed the contribution of featural attentional focus in 15 second blocks, combined across the three chromatic conditions (orientation, contrast, shape and passive). This resulted in an average of 24 beta weight estimates for each voxel at the multivariate level. The third analysed the contribution of stimulus chromaticity, combined across the three feature-specific attentional focus conditions (orientation, contrast and shape) (L+M, L-M and S-(L+M)). This produce an attention and passive dataset for each chromatic condition. This resulted in 48 beta weight estimates for each voxel at the multivariate level.

5.3.5 Statistical analysis

Participants with >5% average variance explained across the visual ROIs were retained for further analysis (no participants were discarded on the basis of percentage variance explained). In line with analysis in Chapter 4, only ROIs V1, V3A/B, hV4, LO-1, LO-2, IPS0 and A1 were retained for further analysis on the basis of previous literature providing hypotheses regarding the process of (particularly chromatic) information in these regions.

5.3.5.1 Univariate attentional modulation

For each participant, feature-specific beta weights were averaged (across orientation, contrast and shape), and compared with passive beta weights through a Wilcoxon signed-rank test for each ROI independently. For the chromatic analysis, each set of chromatic attention beta weights were compared with their respective chromatic passive beta weighted through a Wilcoxon signed-rank test for each ROI. Univariate betas (independently for feature and chromatic datasets) were analysed through one-way repeated measures ANOVAs to identify significant differences in BOLD signal modulation between either orientation, contrast and shape, or between L-M, L+M and S-(L+M) chromatic conditions.

5.3.5.2 Multivariate pattern classification

To investigate patterns of activation in ROIs, beta weights were calculated for each attentional focus predictor timeseries (for both feature and chromatic analyses respectively), for each voxel, using deconvolution to examine the fit of the HRF to the timeseries data. This deconvolution involved estimating the

haemodynamic response which best predicts a measured fMRI signal, reducing the influence of noise within the data. Voxels in each ROI were ranged by largest percentage variance explained in the GLM, and only the top 100 voxels by variance explained per ROI were retained for further analysis to ensure comparable results across ROIs of differing size.

For each ROI, the 100 top voxel beta weights were z-scored (across voxels) and fed to a 'LIBSVM' support vector machine (SVM) classifier (Chang & Lin, 2011), to decode featural attentional focus or attended stimulus chromaticity. The SVM used a radial basis function kernel and leave-one-out cross-validation procedure (across beta weights for each participant). For both the feature and chromatic analysis, we first assessed three-way decoding accuracy, supplying the classifier with information about all three data classes (orientation, contrast and shape or L-M, L+M and S-(L+M)) simultaneously. The classifier used the 'one-against-one' approach to provide a single classification accuracy score (across many cross-validations) for each participant (Knerr et al., 1990). We then assessed these multi-class classification accuracies for each participant against change performance through one-sample Wilcoxon signed-rank tests for each ROI.

Following this, we also performed pairwise classification, to pin-point the differences in patterns of BOLD signal modulation between particular classes of attentional focus. This pairwise classification assessed decoding accuracy between two conditions at a time. For the feature analysis we compared decoding of orientation versus contrast, orientation versus shape and

contrast versus shape respectively. For the chromatic analysis we compared voxel-level patterns of activation between red-green versus blue-yellow, red-green versus achromatic and blue-yellow versus achromatic.

5.3.5.3 Multivariate spatial back-projection

In addition to examining voxel-level patterns of response with changing attentional focus, we also wished to examine the spatial localisation of these attentional effects across feature- or chromatic-specific conditions. To achieve this, we ran an additional set of SVM classifications, assessing decoding accuracy between each attentional focus condition and passive viewing data. For the chromatic analysis, each chromatic attentional condition was compared with its chromatic passive counterpart. From these classifications, we calculated the weighted mean of the support vectors to provide an attention-specific 'preference' for each stimulus feature or chromaticity for each voxel in an ROI.

Following the same strategy as implemented in Chapter 4, these support vector weighted means were then back-projected onto an interpolated grid (6° eccentricity, 360° polar angle across 500 samples) representing voxel activation as a function of each voxels visual angle and eccentricity preference in visual space, extracted from pRF data. We additionally plotted an annulus of activation, which was formed from an average of voxel activations between 1.5-3.5° eccentricity, across all polar angles, to provide a clear summary of the locations of any spatially-focused attention. For reference, a standard RF pattern (0° orientation, 0.2 amplitude) was overlain. These back-projection analyses were performed for V1, V2, V3 and hV4

ROIs, as other regions of interest lacked a rich representation of the entire visual field.

5.3.5.4 Timeseries connectivity analysis

To quantify feature- or chromatic-specific connectivity between ROIs, participant-specific multivariate timeseries data (grouping TRs by attentional focus condition) and underwent noise removal (the fitting and removal of a grand-mean) to eliminate scan-to-scan differences in raw amplitude intensity. These data were then averaged across all voxels in an ROI to provide a single univariate timeseries for each attentional condition. Non-parametric Kendall's tau correlations were then performed for all pairwise combinations of ROI (V1, V3A/B, hV4, LO-1, LO-2 and IPS0), generating a correlation matrix for each attentional condition.

To assess the similarity of connectivity patterns, for each attentional focus condition, the correlation matrix was vectorised. This involved removing self-to-self correlations, and 'stacking' ROI-to-ROI correlations such that we produce a single column of correlation data as oppose to a matrix, without altering the correlation values previously computed. This vector of correlation values was then normalised through participant-specific global mean extraction and Fisher-z transformation. The average (Fisher-z transformed) correlation coefficient was then computed to produce a single number, which summarised connectivity across all ROIs in each attentional condition. The difference in correlations between conditions was then analysed via one-way

repeated measures ANOVA with Bonferroni-corrected post-hoc tests for the feature- and chromatic-specific analysis pipelines respectively.

We further investigated ROI-specific connectivity patterns as a function of attentional focus. For each of the ROIs, we extracted data reflecting the correlation of this ROI with all others for every participant. From all participants, data was then sampled with replacement, with a sample reflecting a full complement of ROI-specific correlation data for each attentional condition, and we calculated the mean across these samples. In order to simulate noise in this analysis, we took the same 12 samples selected with replacement for each condition, and for each pairwise comparison of conditions, we switched the condition labels approximately 50% of the time, keeping ROI-ROI relationships constant and calculated the average across these scrambled condition-specific datasets.

For both the observed and noise data, we calculated the root mean squared error (RMSE) distance between each pairwise combination of condition vectors as a measure of difference in patterns of connectivity across the ROI profile of interest between different attentional modulation conditions. This process was repeated across 10,000 iterations for each ROI comparison (5 comparisons). Across all iterations, we then calculated the percentage of observed RMSEs for a pairwise comparison falling below the RMSE of the comparable simulated noise distribution. Any percentile below 5% indicated a difference in ROI-specific patterns of connectivity between two attentional

modulation conditions which was significantly larger than predicted by chance (p <.05).

5.4 Results

5.4.1 Univariate: Stimulus-Driven and Attentional Modulation Analyses

5.4.1.1 Feature Analysis

We first asked whether BOLD modulation dependent on bottom-up stimulus-driven changes in low-level visual feature. As noted in Chapter 4, BOLD signal changes in many visual areas have been reported in response to bottom-up stimulus driven changes in low-level visual features such as orientation and shape, however, these stimulus manipulations are typically far larger than the near-threshold changes used in this study.

Multiple one-way repeated measures ANOVAs, Benjamini-Hochberg corrected for the number of ROIs revealed no significant differences in univariate BOLD modulation between stimulus-driven changes in orientation, contrast or shape in any ROI examined, including our control auditory cortex region (p > .05) (see Figure 5.2B).

As we identified no clear differential BOLD signal modulation in response to bottom-up changes in stimulus feature, we then sought to examine effects of feature-specific attention on modulation of fMRI activity. Whilst previous research has demonstrated the effectiveness of RF stimuli in probing attentional modulation effects in the visual cortex (see Chapters 3 and 4), we first asked whether our chromatic RF stimuli were also effective in eliciting

robust attentional modulation effects. To assess this, we compared averaged attentional activation (across orientation, contrast and shape stimuli of all chromaticities) with activation in the passive condition though paired-sample Wilcoxon signed-rank tests, Benjamini-Hochberg corrected across ROIs. These comparisons revealed averaged attention BOLD signal was significantly greater than activation during passive viewing of the same visual stimulus in all visual ROIs examined (V1: $Z(11) = 2.83$, $p = .010$, V3A/B: $Z(11) = 2.20$, $p = .034$, hV4: $Z(11)= 2.51$, $p = .018$, LO-1: $Z(11) = 2.82$, $p = .010$, LO-2: $Z(11) = 2.90$, $p = .010$). No significant difference in BOLD signal modulation between directed attention and passive viewing was identified in the control auditory cortex ROI ($Z(11) = 0.55$, $p = .583$) (see Figure 5.2A).

We next sought to examine any feature-specific differences in attentional modulation in our visual ROIs. Our previous analysis, averaging across feature-specific attentional focus conditions demonstrated a robust attentional modulation effect, but revealed nothing regarding the specificity of an attentional response versus a global mechanism boosting activation in individual areas. To examine differences in BOLD signal across featural attention task, we performed multiple one-way repeated measures ANOVAs, Benjamini-Hochberg corrected for the number of ROIs. These analyses revealed no significant difference in attentional modulation between orientation, contrast and shape conditions in V1, LO-1 and A1 ROIs ($p > .05$). In V3A/B, we identified a significant main effect of attentional modulation across conditions ($F(2,22) = 5.29$, $p = .027$), however, no significant differences between conditions were identified in Bonferroni-

corrected post-hoc tests. In hV4, analysis revealed significantly reduced BOLD signal modulation during attention to orientation, in comparison to contrast and shape conditions ($F(2,22) = 5.51$, $p = .027$, post-hoc $p = .037$ and $p = .047$ respectively). In LO-2, directed attention toward shape elicited significantly greater BOLD signal modulation than attention directed towards orientation or contrast ($F(2,22) = 10.22$, $p = .004$, post-hoc $p = .015$ and $p = .009$ respectively) (see Figure 5.2C).

In summary, our feature-specific univariate analyses demonstrated neither bottom-up stimulus manipulations nor top-down attentional demands had strong *differential* effects on BOLD signal modulation in the visual cortex. However, areas hV4 and LO-2 demonstrate relatively weak, but significant differential responses for attention to orientation (hV4) and to shape (LO-2) compared to other task conditions.

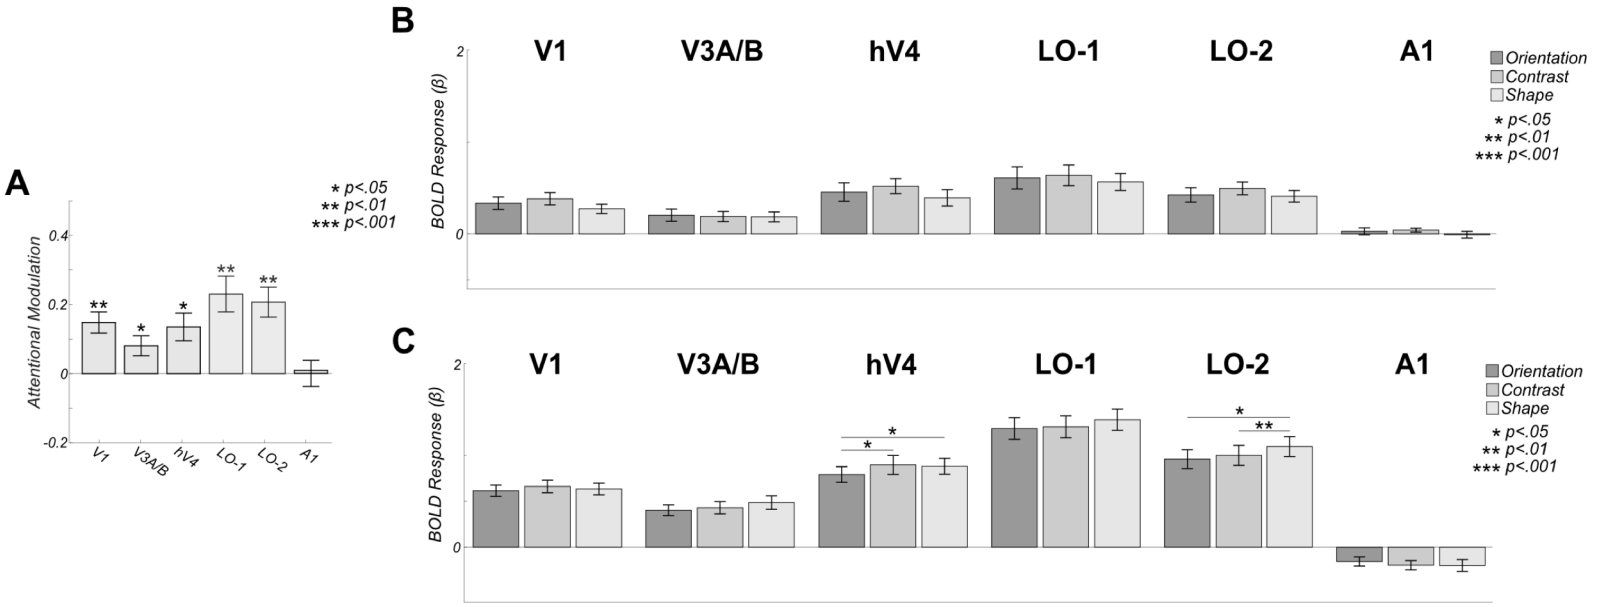

Figure 5.2 Univariate responses to both attention and stimulus modulation are weak. A) Mean BOLD modulation due to attention: positive values reflect greater BOLD amplitude during directed attention (averaged across orientation, contrast and shape attentional conditions). Data are shown for 5 visual ROIs: V1, V3A/B, hV4, LO-1 and LO-2 as well as a control area (auditory area A1). No areas exhibit a significant increase in activity in the attend vs passive comparison. B) Mean bottom-up stimulus responses to individual modulation events averaged over all attentional conditions. No areas exhibit differential responses to stimulus modulations. C) Feature-specific attentional modulations averaged over all modulation types. Areas hV4 and LO-2 show significant differences in BOLD amplitude when directing attention toward different visual features. All error bars are +/- 1 SEM and significance asterisks indicate Benjamini-Hochberg corrected values.

5.4.1.2 Chromatic Analysis

In addition to analysing feature-specific attentional modulation effects, we also wished to investigate the effects of attending to stimuli of different chromaticies to examine the extent of segregation of chromatic information in the visual cortex. Initially, to ensure our chromatic stimuli produced clear modulation of the BOLD signal, we repeated the same analysis pipeline as used for the feature-specific attention analysis. We conducted multiple paired-sample Wilcoxon comparing averaged attentional modulation (averaged across orientation, contrast and shape) with the chromatic-specific passive viewing condition, Benjamini-Hochberg corrected for the number of ROIs, independently for each stimulus chromaticity condition.

For the red-green and achromatic stimulus conditions, we identify significantly greater BOLD modulation during directed attentional focus than in passive viewing in all visual ROIs (red-green: $Z(11) = 2.43$, $p = .027$, V3A/B: $Z(11) = 2.28$, $p = .034$, hV4: $Z(11) = 2.51$, $p = .024$, LO-1: $Z(11) = 2.35$, $p = .030$, LO-2: $Z(11) = 2.67$, $p = .024$, achromatic: V1: $Z(11) = 2.60$, $p = .024$, V3A/B: $Z(11) = 2.51$, $p = .024$, hV4: $Z(11) = 2.60$, $p = .024$, LO-1: $Z(11) = 2.82$, $p = .021$, LO-2: $Z(11) = 2.98$, $p = .021$). In all three stimulus conditions, the control auditory cortex ROI reveals no significant difference in level of BOLD modulation between attention and passive viewing conditions ($p > .05$). In the blue-yellow stimulus condition, paired-sample Wilcoxon signed-rank tests revealed significantly greater BOLD signal modulation during directed attention versus passive viewing in V1 ($Z(11) = 2.12$, $p = .047$), LO-1 ($Z(11) = 2.90$, $p = .021$) and LO-2 ($Z(11) = 2.82$, $p = .021$). However, we identify no significant difference in BOLD signal modulation

during directed attention versus passive viewing in V3A/B or hV4 ($p > .05$) (see Figure 5.3A).

We then sought to identify differences in BOLD signal modulation with attention directed to different stimulus chromaticity. As in the feature-specific analysis, we conducted multiple one-way repeated measures ANOVAs, comparing BOLD signal during attention toward red-green, blue-yellow and achromatic stimuli, Benjamini-Hochberg corrected across the number of ROIs. This analysis revealed significantly greater BOLD signal modulation during attention to red-green versus blue-yellow stimuli in all visual ROIs examined V1: $F(2,22) = 7.86$, $p = .013$, post-hoc $p = .024$ V3A/B: $F(2,22) = 5.84$, post-hoc $p = .040$, $p = .032$, hV4: $F(2,22) = 5.34$, $p = .037$, post-hoc $p = .026$, LO-1: $F(2,22) = 8.25$, $p = .020$, post-hoc $p = .020$, LO-2: $F(2,22) = 10.58$, $p = .013$ post-hoc $p = .013$). In LO-2, we also identify significantly greater BOLD signal modulation during attention to achromatic stimuli versus blue-yellow stimuli (post-hoc $p = .032$) (see Figure 5.3B). Here, our analysis reveals robust evidence of attentional modulation versus passive viewing for all three stimulus chromaticities, and a relatively weaker attentional modulation to blue-yellow versus red-green stimuli across the visual hierarchy.

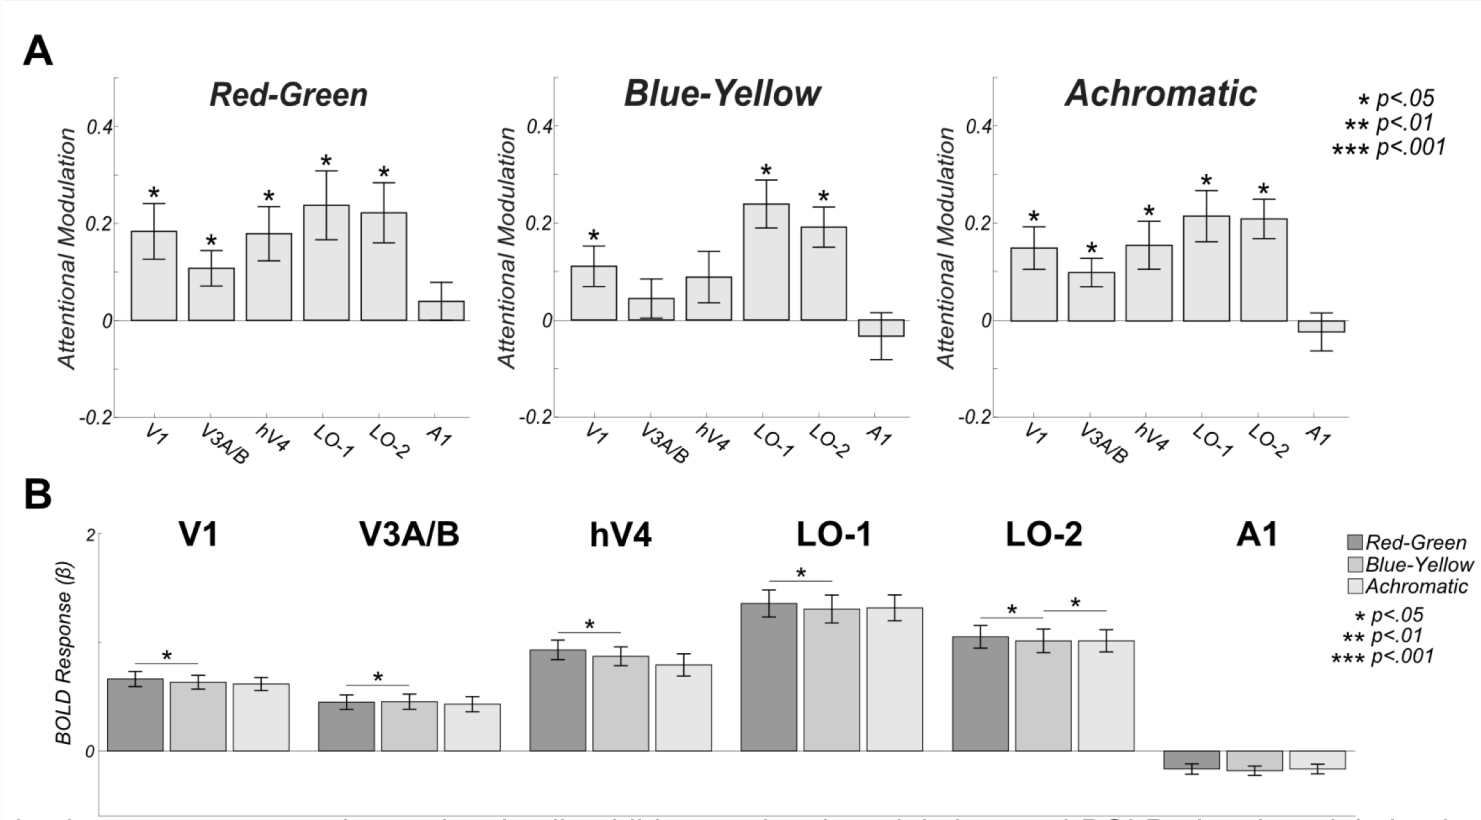

Figure 5.3 Univariate responses to chromatic stimuli exhibit attentional modulation, and BOLD signal modulation is significantly weaker with attention to blue-yellow stimuli than red-green. A) Mean BOLD modulation due to attention: positive values reflect greater BOLD amplitude during directed attention (averaged across orientation, contrast and shape attentional conditions). Data are shown for 5 visual ROIs: V1, V3A/B, hV4, LO-1 and LO-2 as well as a control area (auditory area A1). Almost all visual ROIs in every comparison exhibit significant increases in activity in the attend vs passive comparison. B) Colour-specific attentional modulations averaged over all low-level visual feature modulation types. All visual ROIs show significantly lower signal modulation with blue-yellow stimulus chromaticities. All error bars are +/- 1 SEM and significance asterisks indicate Benjamini-Hochberg corrected values.

5.4.2 Multivariate Analysis

5.4.2.1 Feature Analysis

The overall BOLD activity during attentional focus was significantly greater than passive viewing in all visual ROIs. However, we identified relatively little evidence of differential patterns of BOLD signal modulation when comparing attention to orientation, contrast and shape at the univariate level. We then sought to investigate whether we could decode attention focus through examination of fine-grain patterns of activation across voxels. To do this, we performed a three-way multivariate pattern classification analysis on the feature-attention responses, simultaneously classifying orientation, contrast and shape modulations (averaged across chromatic conditions). One-sample Wilcoxon signed-rank tests versus chance (33%) revealed it was possible to decode the feature held in attentional focus in all visual ROIs (V1: $Z(11) = 3.06$, $p = .015$, V3A/B: $Z(11) = 2.83$, $p = .015$, hV4: $Z(11) = 2.80$, $p = .015$, LO-1: $Z(11) = 2.95$, $p = .015$, LO-2: $Z(11) = 2.98$, $p = .015$). Classification accuracy ranged from 50.18 to 60.35%. In control ROI A1, classification accuracy was not significantly above chance level performance ($Z(11) = 0.31$, $p = 1.85$) (see Figure 5.4A).

To further analyse these significant classifications of attentional state, we conducted multiple pair-wise SVM decoding analyses to determine the driving forces behind successful multi-class decoding. Classification accuracies were compared with chance-level performance (50%) through one-sample Wilcoxon signed-rank tests, adjusted for multiple comparisons (number of comparisons across ROIs) using Benjamini-Hochberg correction.

In all visual ROIs, we demonstrated successful classification of attention to orientation versus. In almost all ROIs, we also identified successful classification between attention to orientation and contrast. Additional above-chance classification accuracies were demonstrated between attention to contrast and attention to shape in V1,V3A/B and LO-2 (see Figure 5.4B and Table 5.1).

Table 5.1 Attention directed toward orientation, contrast and shape can be decoded through two-way classification in visual ROIs.

	Orientation vs. Contrast		Orientation vs. Shape		Contrast vs. Shape	
ROI	Z	*p*	Z	*p*	Z	*p*
V1	2.67	.034*	3.06	.027*	2.83	.029*
V3A/B	2.47	.051	3.06	.027*	2.98	.027*
hV4	2.74	.033*	2.87	.029*	2.35	.066
LO-1	2.98	.027*	2.95	.027*	2.31	.069
LO-2	2.63	.035*	2.71	.033*	3.06	.027*

** p <.05, ** p<.01, *** p<.001*

Previous research has demonstrated that some types of stimuli can generate retinotopically-based patterns of response in individual ROIs at coarse spatial scales. For example, clear maps of orientation preference have been demonstrated in primary visual cortex, which have been suggested to influence multivariate pattern classification techniques (see Freeman, Brouwer, Heeger, & Merriam, 2011). We asked whether our stimuli may also evoke relatively large-scale biases in spatial attention, which may inform our classification analyses.

We identified the voxels most informative for each low-level visual feature classification (versus passive) and back-projected these into visual space

coordinates. These representations were then averaged across participants to reveal average spatial patterns of voxel preference across individuals. If all participants employed a consistent spatial strategy for a particular low-level visual feature task, these averaged plots would reveal a consistent non-zero response in a particular location. If no large-scale patterns of response were generated by attention in a consistent fashion across observers, these maps would average to zero. We identify no significant patterns of attentional modulation are evident for any feature of attentional focus across ROIs, with no clear peak of spatial attentional focus as a function of polar angle also identified in the annulus overlays (see Figure 5.4C).

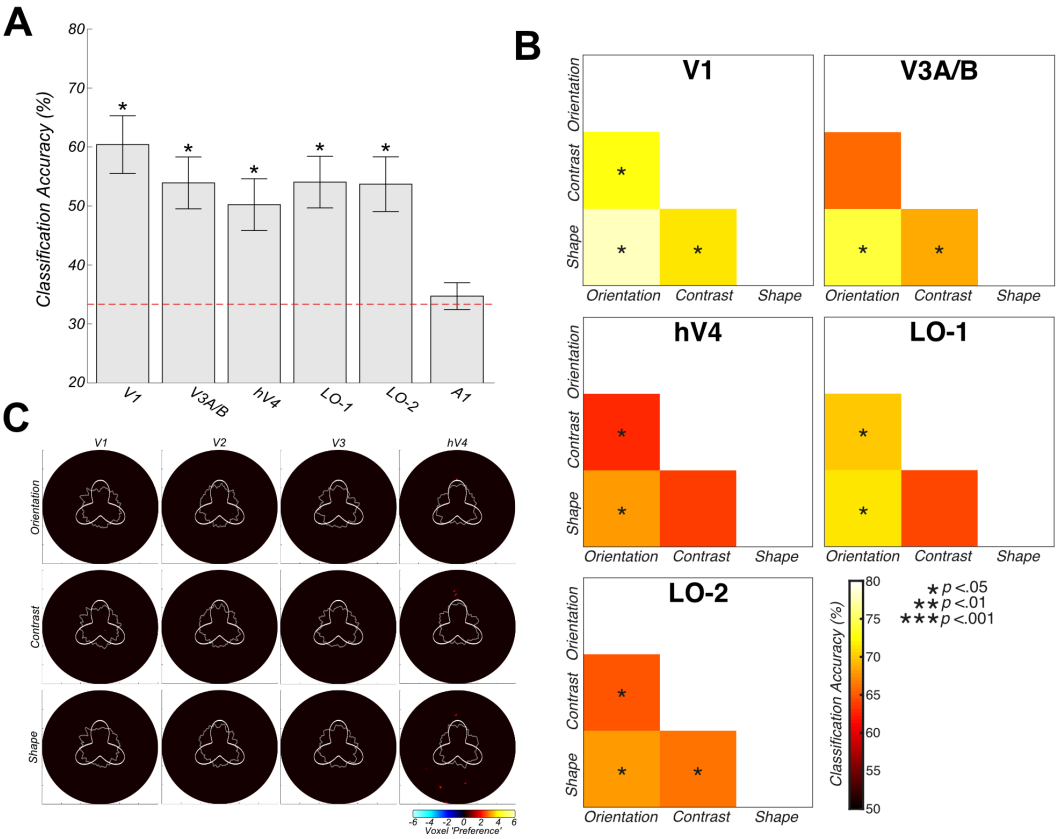

Figure 5.4 Feature-specific multivariate Support Vector Machine Decoding and Group-averaged voxel feature-specific weights as a function of eccentricity (6°) and polar angle (360°). Voxel-level responses with individual ROIs are modulated by attentional state. A) Overall three-way decoding accuracies. Attentional state can be decoded above chance in all ROIs except A1. Error bars reflect +/- 1 SEM. B) Two-way classification accuracies across pairwise combinations of attentional focus (orientation versus shape, orientation versus contrast and contrast versus shape). Voxel patterns in all areas differ significantly between attention to orientation and shape. Significance asterisks indicate Benjamini-Hochberg corrected values. C) Spatial backprojection analyses reveal no large-scale biases in voxel weights across location. The gray annulus reflects averaged voxel modulation at 1° intervals across 1.5-2.5° visual space. Deviations from circularity indicate positive (feature-specific) or negative (passive viewing) preferring clusters of voxels. An RFP stimulus overlay is provided for reference. Activation is thresholded at +/- 1.7 z-score ($p < .05$).

5.4.2.2 Colour Analysis

We performed the same analysis pipeline for our chromatic data, in order to identify differential patterns of voxel-level modulation when attention switched between three stimulus chromaticities. We performed a three-way pattern classification analysis, simultaneously supplying data from red-green, blue-yellow and achromatic attention conditions (excluding data from the passive chromatic conditions). One-sample Wilcoxon signed-rank tests versus chance (33.33%) Benjamini-Hochberg corrected for the number of ROIs, revealed above-chance classification in almost all visual ROIs examined V1: $Z(11) = 3.06$, $p = .011$, V3A/B: $Z(11) = 3.06$, $p = .011$, hV4: $Z(11) = 3.06$, $p = .011$, LO-1: $Z(11) = 2.85$, $p = .016$). In visual area LO-2 and control auditory cortex ROI (A1) classification accuracy was not significantly greater than chance (LO-2: $Z(11) = 2.27$, $p = .068$, A1: $Z(11) = 1.57$, $p = .285$). Across the visual ROIs, mean classification accuracy ranged from 42.22 to 57.64% (see Figure 5.5A).

Again, we performed further pairwise classification analyses to determine which chromatic states generated different voxel-level responses. We conducted multiple one-sample Wilcoxon signed-rank tests versus chance (50%), Benjamini-Hochberg corrected for the number of comparisons across ROIs. In almost all visual ROIs, successful classification of red-green versus achromatic stimuli was evident. Above-chance classification was also evident between blue-yellow and achromatic stimuli in V3A/B, hV4 and LO-1. In V1, hV4 and LO-1, decoding accuracies also exceed chance performance when comparing red-green activation with blue-yellow. In LO-2, not successful

classification of any pairwise comparison of stimulus chromaticity was evident (see Figure 5.5B and Table 5.2).

Table 5.2 Attention directed toward red-green, blue-yellow and achromatic stimuli can be decoded through two-way classification in visual ROIs.

	Red-Green vs. Blue-Yellow		Red-Green vs. Achromatic		Blue-Yellow vs. Achromatic	
ROI	Z	*p*	Z	*p*	Z	*p*
V1	3.07	.028*	2.59	.047*	2.36	.083
V3A/B	1.29	.065	2.75	.037*	2.76	.037*
hV4	3.07	.028*	2.95	.028*	3.06	.028*
LO-1	2.63	.047*	2.94	.028*	2.94	.028*
LO-2	1.80	.258	2.28	.088	2.28	.088

** p <.05, ** p<.01, *** p<.001*

Finally, the spatial backprojection analysis was also performed for visual ROIs V1, V2, V3 and hV4, to identify if any large-scale biases in the organisation of chromatic preference were driving our multivariate classification results. We calculated he mean chromatic-specific support vector weights for each voxel and back-projected these values into their eccentricity and polar angle coordinates, and these maps were averaged across participants. Again, we identify no large-scale biases in voxel preference for any stimulus chromaticity condition in any of the early visual ROIs examined (see Figure 5.5C).

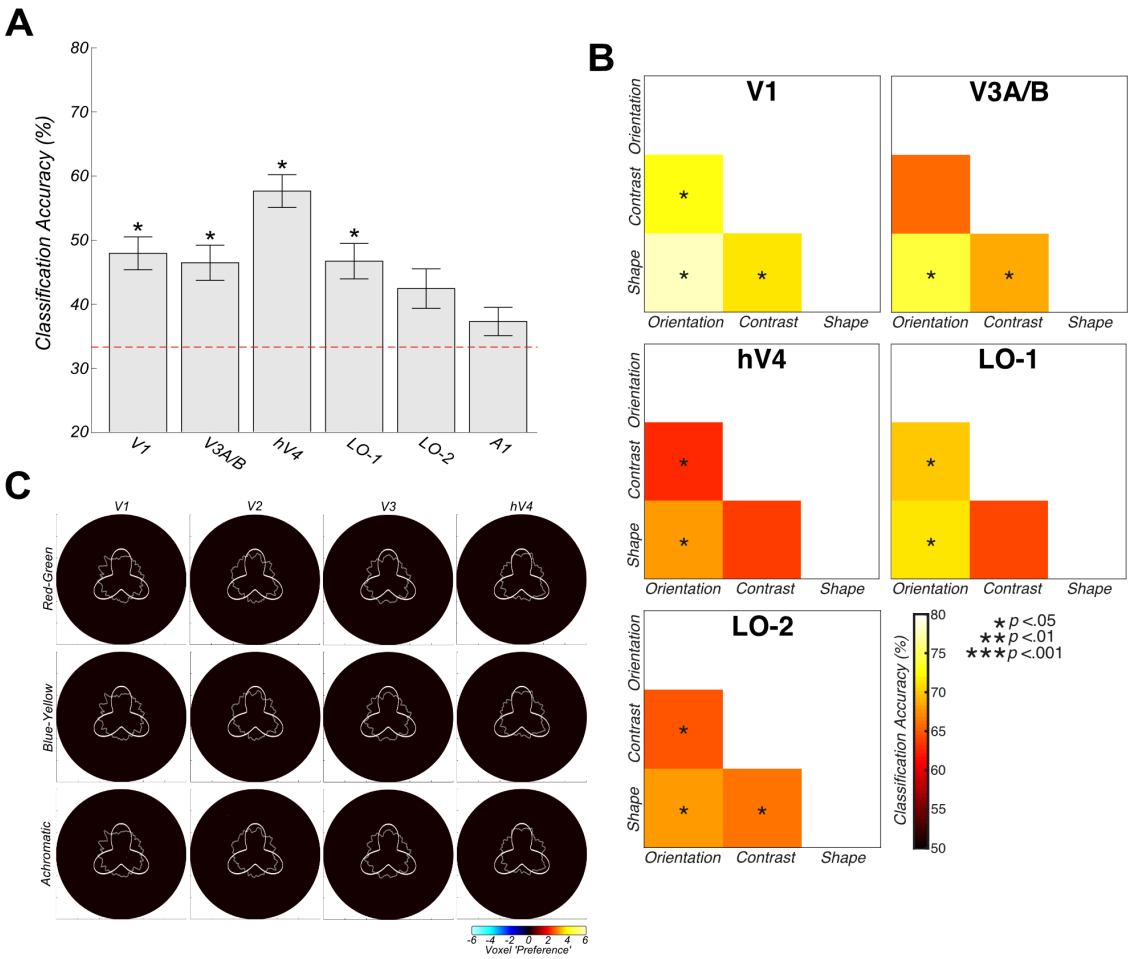

Figure 5.5 Colour-specific Multivariate Support Vector Machine Decoding and Group-averaged voxel feature-specific weights as a function of eccentricity (6°) and polar angle (360°). Voxel-level responses with individual ROIs are modulated by attentional state. A) Overall three-way decoding accuracies. Chromatic attentional state can be decoded above chance in all ROIs except LO-2 and A1. Error bars reflect +/- 1 SEM. B) Two-way classification accuracies across pairwise combinations of attentional focus (red-green versus blue-yellow, red-green versus achromatic and blue-yellow versus achromatic). Voxel patterns in all almost areas differ significantly between attention to red-green and achromatic-defined stimuli, with successful classifications also evident for red-green versus blue-yellow-defined stimuli and blue-yellow versus achromatic activation comparisons. Significance asterisks indicate Benjamini-Hochberg corrected values. C) Spatial backprojection analyses reveal no large-scale biases in voxel weights (chromatic preference) across location. The gray annulus reflects averaged voxel modulation at 1° intervals across 1.5-2.5° visual space. Deviations from circularity indicate positive (feature-specific) or negative (passive viewing) preferring clusters of voxels. An RFP stimulus overlay is provided for reference. Activation is thresholded at +/- 1.7 z-score (p <.05).

5.4.3 Timeseries Connectivity Analysis

5.4.3.1 Feature

Our univariate analysis demonstrated that overall, attention had relatively little effect on time-averaged BOLD responses in individual visual areas. In our final analysis, we asked whether attention may exert a clearer effect at a slightly coarser scale; influencing the way that individual ROIs communicate with each other. Specifically, we investigated differing functional connectivity between ROIs as a function of attentional state.

We first conducted non-parametric Kendall's Tau correlations between ROI-average (univariate) timecourses from each pairwise comparison of visual ROIs for each featural attention condition. These correlation patterns across visual ROIs (V1, V3A/B, hV4, LO-1, LO-2 and IPS0) represent a 'fingerprint' for each attentional task (see Figure 5.6A). We examined if these connectivity fingerprints were altered by attentional task, though a one-way repeated measures ANOVA with Bonferroni-corrected post-hoc tests, conducted on normalised and Fisher-z transformed coefficients. This analysis revealed significantly greater positive correlation between ROIs during passive viewing compared to attention to orientation ($F(3,42) = 5.05$, $p = .015$, post-hoc $p = .002$). No other differences in patterns of connectivity between attentional conditions were identified ($p > .05$).

Next, we asked whether differences in connectivity fingerprints between attentional task may be driven by connections with a particular ROI, rather than examining patterns of connectivity across all ROIs simultaneously. For

this, we calculated the Euclidian distance (RMSE) between pairwise comparisons of attentional task condition (averaged across a subsample of participants for each condition), for both observed and noise (scrambled attention-condition label) datasets, bootstrapped across 10,000 iterations.

Correlations with visual ROIs V1, hV4, LO-2 and IPS0 show revealed no significant differences in patterns of connectivity across all attentional states ($p > .05$). However, correlations from both V3A/B and LO-1 to all other visual areas demonstrated significant differences in patterns of connectivity between attention to contrast and passive viewing ($p = .024$ and $p = .004$ respectively). Partial correlations with V3A/B also revealed a significant difference in connectivity during attention to shape and passive viewing ($p = .037$) (see Figure 5.6B).

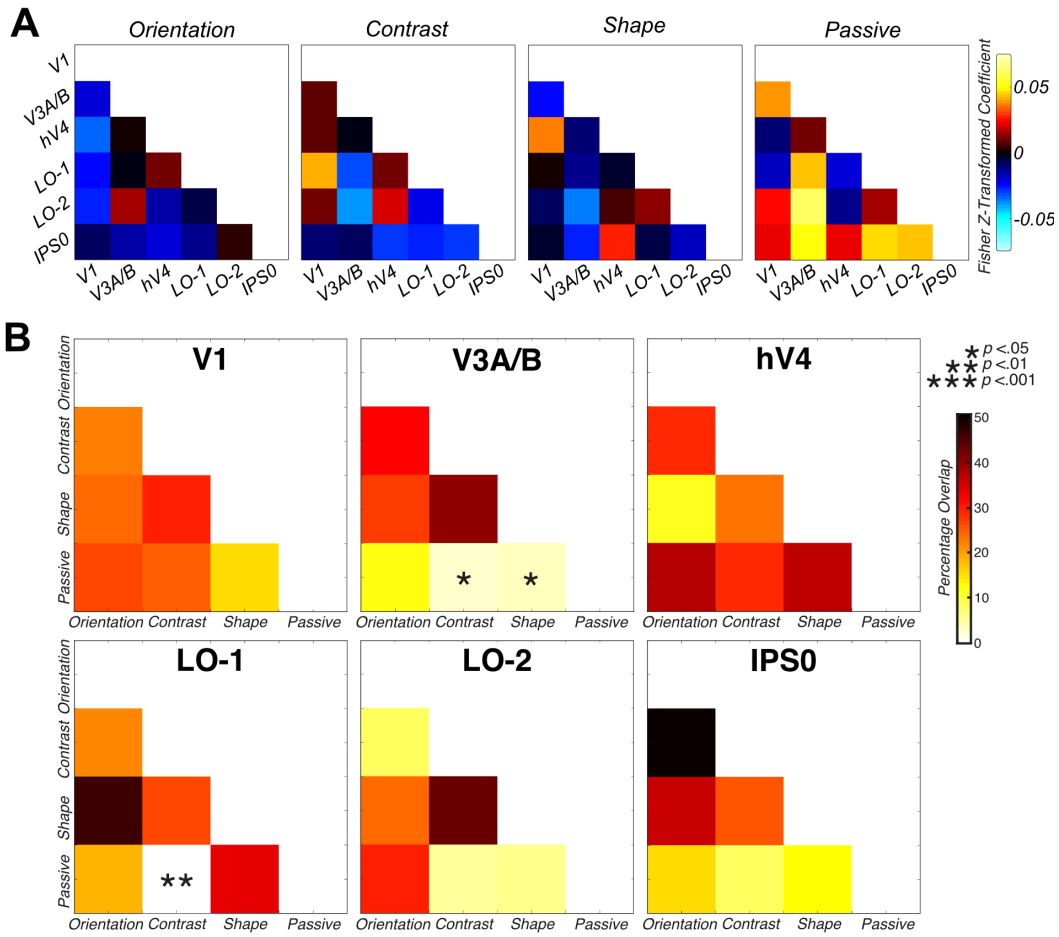

Figure 5.6 Greater overall connectivity during passive viewing compared to directed attention. A) Group feature-specific averaged attentional modulation connectivity values, indicating significantly greater mean connectivity during passive viewing in the attentional tasks. B) Bootstrapped measures of distance between pairwise combinations of attentional focus conditions, for V1, V3A/B, hV4, LO-1, LO-2 and IPS0 ROIs. The figure demonstrates the percentage overlap between the distribution of RMSE scores across 10,000 iterations of randomly-selected samples of the observed data and scrambled correlation 'noise' data for each ROI across multiple pairwise comparisons. Significant overlap (less than 5%) between pairwise combinations of condition are indicated with asterisks.

5.4.3.2 Colour

Finally, we repeated this connectivity analysis pipeline for the chromatic data subset. Firstly, we wished to investigate whether overall patterns of connectivity differ between chromatic attention and passive viewing conditions, as identified in the feature-specific analysis and in Chapter 4. We replicated the same process, performing non-parametric Kendall's Tau correlations between univariate timecourses from each pairwise combination of visual ROIs for each chromatic attention conditions (see Figure 5.7A). To assess the differences between these correlation matrices, we performed a one-way repeated measures ANOVA with Bonferroni-corrected post-hoc tests on these vectorised, normalised and Fisher-z transformed correlation coefficients. Here, we identified significant a main effect of attentional task condition on patterns of connectivity across visual ROIs ($F(5,70) = 62.19$, $p < .001$). Bonferroni-corrected post-hoc tests revealed significantly greater positive connectivity during each chromatic passive condition in comparison with all directed attention conditions ($p < .001$) (e.g. red-green passive versus red-green attention, blue-yellow attention and achromatic attention condition respectively). We also identified a significantly greater negative correlation during attention to achromatic stimuli than red-green attention ($p = .043$).

We also repeated the same partial correlation analysis as above, examining whether differences in connectivity fingerprint between attentional task condition were driven by a subset of regions of interest. We calculated the distance between pairwise comparisons of attentional task condition, for both observed and noise datasets, bootstrapped across 10,000 iterations.

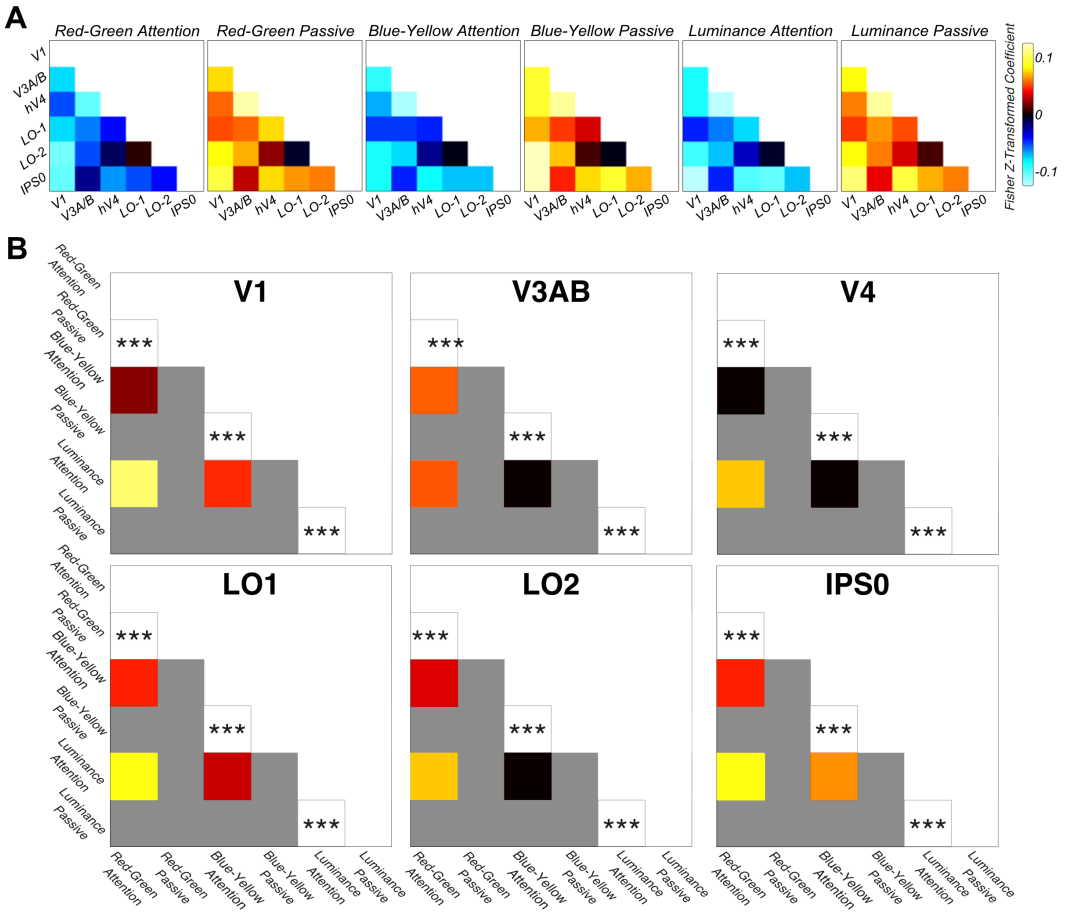

Figure 5.7 Greater overall connectivity during passive viewing compared to directed attention in all stimulus chromaticity conditions. A) Group feature-specific averaged attentional modulation on connectivity values, indicating significantly greater mean connectivity during passive viewing in the attentional tasks. B) Bootstrapped measures of distance between pairwise combinations of attentional focus conditions, for V1, V3A/B, hV4, LO-1, LO-2 and IPS0 ROIs. The figure demonstrates the percentage overlap between the distribution of RMSE scores across 10,000 iterations of randomly-selected samples of the observed data and scrambled correlation 'noise' data, for each ROI across multiple pairwise comparisons. Significant overlap (less than 5%) between pairwise combinations of condition are indicated with asterisks.

Across all visual ROI partial correlations examined, we identified significantly different patterns of connectivity between red-green attention and red-green

passive viewing (V1: p <.001, V3A/B: p <.001, hV4: p <.001, LO-1: p <.001, LO-2: p <.001, IPS0: p <.001). We also identify significant differences between blue-yellow attention and blue-yellow passive viewing (all p <.001), and between achromatic attention and achromatic passive viewing conditions (all p <.001) (see Figure 5.7B).

5.5 Discussion

Here, we demonstrate clear evidence that the human visual cortex alters both *within*-area and *between*-area activity as a function of attentional task (both when attending to low-level visual feature and to stimulus chromaticity). In support of our previous findings (see Chapter 4) at the univariate level, we identify no differences in bottom-up stimulus-driven BOLD signal modulation in response to small changes in the visual stimulus. Additionally, in support of recent literature (Brouwer & Heeger, 2009; Seymour et al., 2009, 2010; Song et al., 2011; Sumner et al., 2008; Xing et al., 2013) we identity relatively little clear evidence of differential univariate BOLD signal modulation as a function of task.

In the feature-based analysis, hV4 demonstrates a significantly reduced signal modulation when attending to orientation, which may be reflective of the fact that hV4 has been reported to contain a reduced percentage of orientation-selective neurons in comparison with early visual ROIs (V1, V2 and V3) (Zeki, 1978). LO-2 analysis also identified a significantly greater BOLD signal modulation when attending to shape, partially fitting with the suggestion of Silson et al., (2013) regarding LO-2 as a specialised shape-processing area.

In the chromatic analysis, we identify significantly reduced BOLD signal modulation during attention to S-cone defined blue-yellow stimuli than red-green in all visual ROIs examined. This weaker response during attention to S-cone defined stimuli is in fitting with previous research such as Liu & Wandell, (2005), Mullen et al., (2015) and Wang & Wade, (2011) however, are findings are not limited to the dorsal stream. Hence, in both the feature- and chromatic- analysis, overall univariate findings are relatively uninformative about feature-specific changes in the visual stimulus or attentional state. Our univariate analysis suggests the study of attentional modulation at an average, ROI-level scale may not be informative regarding the computations performed in and across visual areas a function of attentional focus.

Our univariate analysis provides little evidence the existence of feature- or chromatic-specific attentional modulation effects. However, we show that multivariate, voxel-level scale *patterns* of modulation in individual visual areas differ significantly as a function of attentional task for both the feature- and chromatic analyses. Our ability to classify featural attentional state supports previous decoding literature (see Brouwer & Heeger, 2009; Clifford, Mannion, & McDonald, 2009; Freeman et al., 2011; Kamitani & Tong, 2005, 2006; Mannion, McDonald, & Clifford, 2009; Sumner et al., 2008) and demonstrates attentional focus can modulate voxel-level patterns of activation in a distinct and feature-specific manner, supporting our findings with achromatic stimuli in Chapter 4.

We were also able to classify chromatic attention condition in almost all visual ROIs, and differential patterns of pairwise classification were evident across the ROIs examined. In all visual ROIs (excluding LO-2), we were able to classify red-green attentional data from achromatic data. This provides evidence for segregation of L-M and L+M information across the visual hierarchy. In V1, hV4 and LO-1 we also demonstrate above-chance classification of attention to red-green versus blue-yellow defined stimuli, in support of Sumner et al., (2008). These data imply that the information being used for each of these discriminations comes from different neuronal populations (even when collapsed across different low-level stimulus feature attention conditions) and again suggests some level of segregation of blue-yellow and red-green chromatic channels across the visual cortex. In V3A/B, hV4 and LO-1, it was also possible to classify blue-yellow from achromatic stimuli, suggesting again that attentional mechanisms are chromatic-specific and that different populations of neurons exist for encoding stimulus chromaticity across the visual cortex which are preferentially modulated by attentional mechanisms.

Across both our feature- and chromatic- analyses, our above-chance classifications demonstrate that attention elicits differing voxel-level patterns of modulation as a function of task, highlighting the importance of multiple spatial scales of analysis in attentional paradigms and the fine-grain nature of attentional modulation effects. The ability of the visual cortex to alter the activation of specific neurons in particular visual ROIs supports the findings of a range of previous researchers (such as Corbetta et al., 1990; Martinez-Trujillo & Treue, 2004; Martínez-Trujillo & Treue, 2002; Reynolds et al., 2000;

Saenz et al., 2002; Serences & Boynton, 2007; Treue & Martínez Trujillo, 1999).

As in our previous experiment (see Chapter 4) we demonstrate that these above-chance classification results do *not* depend solely on large-scale biases in feature- or chromatic-sensitivity in a visual area, which have been demonstrated in the past to drive classification performance in some types of task (see Clifford et al., 2009; Freeman et al., 2011 and section 1.3.4). Instead, our data are consistent with the hypotheses attention selects sub-populations of relatively fine-scaled maps in individual visual areas, although we do not rule out a coarser-scale topographic mapping which may vary across subjects (supporting Kamitani & Sawahata, 2010; Kamitani & Tong, 2005; Kay, Naselaris, Prenger, & Gallant, 2008; Mannion et al., 2009; Sumner et al., 2008).

Finally, in both our feature and chromatic connectivity analyses, we identify greater average connectivity between ROIs during passive viewing than directed attentional focus. This trend supports our previous findings with achromatic stimuli (see Chapter 4). This effect is striking in the colour analyses, where we identified significant differences between all chromatic passive viewing and all chromatic attention conditions. We suggest this positive connectivity between visual ROIs during passive viewing is reflective of 'default mode' style communication between areas (see Raichle, 2015; Raichle et al., 2001), at a cm as oppose to typically-studied multi-cm scale. We propose once a visual task is provided, particular visual regions disengage from this passive connectivity network in order to process most

effectively the attended feature, supporting the ideas of Dajani & Uddin, (2015); Li et al., (2017); Reineberg, Gustavson, Benca, Banich, & Friedman, (2018); Spadone et al., (2015) and Vatansever, Manktelow, Sahakian, Menon, & Stamatakis, (2016). Perhaps our connectivity and partial correlation findings for colour are more robust as they reflect a combining of data across all three low-level visual feature attention conditions, or maybe colour is a more salient or engaging stimulus than orientation, achromatic contrast or shape, and as such we see a clearer effect in this analysis.

In the partial feature-specific correlation analysis, we identify a significant difference in patterns of connectivity between contrast and passive in V3A/B and LO-1 and between shape and passive viewing conditions in V3A/B. These findings indicate, in support of our previous experiment (see Chapter 4) that not all ROIs and their connections are modulated in the same way in all attentional tasks. The differences we identify between this experiment and our previous achromatic experiment may reflect a by-product of combining data across stimulus chromaticity. However, we still support an overall trend of greater connectivity between visual ROIs during passive viewing than versus directed attentional focus.

In conclusion, here, we have employed a relatively novel approach for investigating attentional modulation both to low-level stimulus feature and chromaticity across the visual cortex, at a range of spatial scales. We demonstrate clear evidence of attentional modulation across the visual hierarchy and identify patterns of attentional modulation across voxels in

individual areas. We also demonstrate that attention elicits a strong *decorrelation* of ROI responses in comparison with passive viewing of a stimulus. This paradigm can reveal a wealth of information regarding the spatial extent of attentional modulation across the visual cortex in response to shifting attentional focus in oppose to more traditionally-studied bottom-up stimulus driven changes.

6. Investigating patterns of attentional modulation in the visual cortex with a naturalistic visual stimulus

6.1 Abstract

Humans constantly interact with a complex and dynamic visual environment. However, the majority of studies probing visual attention use highly controlled, low-level visual stimuli, which have little similarity with our natural environment. Here, we examined visual cortex activation with fMRI in response to shifting attentional focus, during viewing of an uncontrolled, relatively naturalistic stimulus (the animated movie *Frozen).* We identify relatively little evidence of differential patterns of attentional modulation to low-level visual features at both the univariate level and multivariate level. However, we identify distinct and robust effects of attention on patterns of connectivity between a range of visual areas. We identify significantly greater positive connectivity between visual regions when attending to faces or passively viewing a stimulus in comparison with conditions of directed attention. We also provide evidence of differing patterns of network-level connectivity in the visual cortex when directing attention towards fine-scale differences in a feature category (e.g. vertical, horizontal and diagonal orientations). We demonstrate naturalistic stimuli are effective in probing mechanisms of visual attention, identifying a synchronisation and desynchronization of visual areas a function of attentional task, and these differences emerge even at a relatively fine scale of stimulus selectivity.

6.2 Introduction

The human brain has evolved to allow for successful navigation through an immensely stimulating, complex and dynamic visual world (Spiers & Maguire, 2007). Typically, visual research focuses on using isolated, highly controlled stimuli to investigate a precise aspect of visual processing. However, such stimuli do not provide an accurate reflection of how humans interact with visual information in their external environment on a day-to-day basis. Recently, visual researchers have begun to experiment with naturalistic, real-world stimuli, to better reflect the rich visual environment we are constantly immersed in and to allow for stronger generalisation of findings to real-world scenarios. The use of dynamic naturalistic stimuli such as movies in visual research is important to ensure the findings obtained with simplified and controlled low-level visual stimuli are valid and applicable to the complex and dynamic visual world.

Previous research has shown that naturalistic, 'rich' visual stimuli can produce greater activation of the visual cortex than similar studies with relatively low-level visual stimuli (Bartels & Zeki, 2004b). Research has also demonstrated that canonical responses of specialised visual regions and patterns of connectivity are preserved with the use of naturalistic stimuli (Bartels, Zeki, & Logothetis, 2008; Bartels & Zeki, 2004b, 2004a, 2005; Russ & Leopold, 2015; Spiers & Maguire, 2007). The effectiveness of naturalistic stimuli in providing clear and consistent findings, despite their varying and uncontrolled nature has been demonstrated in a number of studies. For example, Bartels & Zeki, (2004a) demonstrated highly similar patterns of feature-specific and anatomically-localised activity across participants

viewing two halves of a film independently. Additional research has also demonstrated clear and consistent patterns of functional organisation and connectivity (even at the voxel level) through the use of naturalistic paradigms, demonstrating the effectiveness of such a stimulus in probing mechanisms of attention in the visual cortex (Bartels & Zeki, 2004b; Hasson, Nir, Levy, Fuhrmann, & Malach, 2004).

The use of naturalistic stimuli provides clear evidence of functional specificity in the visual cortex, such as the identification of colour-specific responses in hV4 (Bartels et al., 2008; Bartels & Zeki, 2004a). However, there has not been an explicit analysis of the effects of changing attentional focus during viewing of a naturalistic stimulus, with researchers instead preferring to employ reverse correlation techniques to examine connectivity between regions in an unrestricted paradigm (e.g. Bartels & Zeki, 2004b; Hasson et al., 2004). In this experiment, we wished to extend our previous findings of distinct, feature-specific patterns of attentional modulation with low-level, highly controlled visual stimuli (see Chapters 4 and 5), using a naturalistic stimulus to examine differential patterns of within- and between-area activity as a function of shifting attentional state with a dynamic and complex stimulus. We also extended our analysis to explore the effects of directed attention to a relatively higher-level visual stimulus (faces), examining functional specialisation and changing network connectivity across the visual cortex.

The fusiform face area (FFA) has been identified as a locus of face-specific processing in the visual cortex, demonstrating preferential processing of face stimuli in comparison to visual stimuli of other categories, such as scrambled faces, houses and hands (Kanwisher, McDermott, & Chun, 1997). Many researchers consider the FFA to be a key area in a specialised face-processing network across the visual cortex (termed the face-specificity hypothesis) (Kanwisher et al., 1997; Kanwisher & Yovel, 2006). However, other research suggests the FFA is a more generalized region selective for fine-grain discriminations in a category, rather than specifically selective for faces (Gauthier et al., 1999). Despite this, previous research used naturalistic stimuli to demonstrate that an overwhelming majority (94%) of movie frames eliciting the greatest FFA activation contained face images, supporting the notion of functional specialisation in this region (Hasson et al., 2004). Attention directed toward face stimuli has also been shown to modulate the response of the FFA (Wojciulik, Kanwisher, & Driver, 1998; Yi, Kelley, Marois, & Chun, 2006), boosting its activation preferentially to face versus house stimuli. Hence, we wished to examine modulation in the FFA in response to attention directed toward different visual stimuli, comparing responses when attending to faces with modulation when viewing relatively lower-level visual stimulus attributes.

Additional regions such as the occipital face area (OFA) and superior temporal sulcus (STS) have also been implicated in face-processing. The OFA region has been demonstrated to possess greater responsivity to faces versus objects (Gauthier et al., 2000), and the STS has been identified as a

region responsive to dynamic faces and bodies, and emotional expressions (Allison et al., 2000; O'Toole, Roark, & Abdi, 2002). Like the FFA, attentional modulation effects have been noted in the OFA and in the STS, where attention directed toward simulated biological motion has been demonstrated to modulate activity (Chiu, Esterman, Han, Rosen, & Yantis, 2011; Safford, Hussey, Parasuraman, & Thompson, 2010). Evidence of signal modulation in the visual cortex in response to face stimuli has been noted across the lateral fusiform gyrus (including the FFA), STS and along the middle temporal gyrus, forming a face-processing network (Bartels & Zeki, 2004a). Hence, attention directed toward face-stimuli produces widespread, yet specific activation across the visual cortex and beyond, and examination of these regions in response to attention to faces versus attention to low-level visual stimulus attributes will allow for investigation of the specialisation of processing along this face-specific network.

Further research has investigated attentional modulation during natural vision at a finer spatial scale than the level of connectivity between different functionally specialised regions. For example, Çukur, Nishimoto, Huth, & Gallant, (2013) identified voxels in the occipito-temporal and fronto-parietal cortex that shifted their tuning toward an attended category and compressed representation of categories semantically unrelated to the attended target. Additionally, Chiu et al., (2011) demonstrated differential patterns of voxel-level activity during attention to race or gender of a perceptually ambiguous face, respectively in the rFFA and bilateral OFA. Hence, we also planned to investigate the effect of shifting attentional focus at the voxel-level examining patterns of activation in individual visual ROIs. Despite the uncontrolled

nature of this naturalistic stimulus in comparison to our previous experiments (see Chapters 4 and 5), we were hopeful regarding our ability to identify similar multivariate activation patterns as identified in Chapters 4 and 5. As naturalistic stimuli have been suggested to produce an improved signal-to-noise ratio, allowing for a better chance of revealing neuronally-specific interactions (Bartels & Zeki, 2005).

In light of our previous research and the literature surrounding the use of naturalistic paradigms as effective stimuli to probe visual attentional mechanisms, we made four predictions. Firstly, we expected little difference in BOLD signal modulation at the univariate level in response to attention directed towards low-level visual features, given the findings of previous literature reflecting the relative fine-scale influence of attentional modulation (e.g. Brouwer & Heeger, 2009; Seymour, Clifford, Logothetis, & Bartels, 2009; Seymour et al., 2010; Song, Rowland, McPeek, & Wade, 2011; Sumner et al., 2008). However, we predicted univariate differences in activation when attending to faces versus attending to other visual features, in the proposed face-specialised regions (FFA, OFA and STS). We predicted differences in voxel-level patterns of activation when participants shifted their attention between different visual stimuli, and finally, we also predicted differing patterns of connectivity between visual regions as a function of attentional focus. In short, we predicted we would see evidence of *within-* as well as *between*-area differences in attentional modulation, though as in previous research, these differences may not be wholly evident at the coarse-scale level of univariate analysis.

6.3 Methods

6.3.1 Participants

14 participants, all University of York staff and students volunteered for the study (9 female, mean age 26.71 years). All participants had normal or corrected-to-normal vision. Informed consent was collected from participants prior to scanning and the study was granted ethical approval by the York Neuroimaging Centre Ethics Committee. All participants also possessed previously collected population receptive field mapping data (see section 2.2. and Chapter 4).

6.3.2 Stimulus presentation

Stimuli were projected onto a screen at the back of the scanner and were viewed via an angle mounted mirror on the head coil, at approximately 57cm viewing distance. Audio commands were received by participants via padded air-based headphones (MR Confon MkII system), that were fitted in the scanner driven by a MOTU 2408mk3 soundcard.

Stimulus presentation was controlled by Psychopy v1.9 (Peirce et al., 2019) and stimulus randomisation was controlled by MATLAB 9.4.0 (2018a) (MathWorks, Inc., Natick, Massachusetts, United States). Visual stimuli were presented using a PROpixx DLP LED projector (Vpixx Technologies Inc., Saint-Bruno-de-Montarvile, QC, Canada) with a long-throw lens which projected the image through a waveguide behind the scanner bore and onto an acrylic screen. Presented images had a 120Hz refresh rate and 1920 x 1080 resolution. During scanning behavioural responses and scanner trigger

pulses were acquired using a USB, 4-button fibre-optic response pad Forp-932 (Current Designs, Philadelphia, PA).

6.3.3 Data Acquisition

Structural and functional data were collected using a 3T Siemens Magnetom Prisma MRI scanner with a 64-channel transmit head and neck coil with a frequency of 123.26 MHz. For participants lacking a previously collected high-resolution T1 anatomical scan, a high-resolution T1 MPRAGE scan was collected (TR = 2300ms, TE = 2.26ms, flip angle = 8°, field-of-view (FOV) = 256mm^3, voxel size = 1.0 x 1.0 x 1.0 mm, 176 slices) in a 5:21 minute scan.

Functional data were collected via 150 T2*- weighted volumes, collecting 30 axial slices in a bottom-up interleaved acquisition order (TR = 2000ms, TE = 30ms, flip-angle = 80°, FOV = 19.2cm^3, voxel size = 2.0 x 2.0 x 2.5 mm) in a 5:08 minute scan. To aid alignment between the functional data and high-resolution structural images, T1-weighted (FLAIR) anatomical images were acquired (TR = 3000ms, TE = 9.1ms, flip angle = 150°, FOV = 19.2cm^3, voxel size = 0.8 x 0.8 x 2.5mm) in a 2:26 minute scan.

6.3.4 Experimental Design

Video clips (without audio) from the popular Disney animated film *Frozen* (Buck & Lee, 2013) were used as the naturalistic stimuli, selected for its clearly distinguishable low-level visual feature categories, in comparison with non-animated films.

We conducted a block design experiment in the fMRI scanner. Participants completed an average of 8 runs, each lasting 5:08 minutes, with a full scanning session lasting approximately 1 hour. Each run began with 8 dummy volumes to control for scanner magnetisation effects during which a mid-grey luminance background was presented to participants. At the beginning of each 30 second block, participants were provided with an audio cue directing their attention towards a specific visual feature (i.e. during an 'attend to vertical' block, participants should actively seek vertically-oriented stimuli). For each low-level visual feature, there were three possible dimensions to attend to; orientation (horizontal, vertical or diagonal), chromatic colour (red, green or blue) and shape (square, circular or triangular). Faces and passive viewing periods were also included. Participants were instructed to press and hold a button whenever they identified the target feature in the naturalistic stimuli. They were also instructed to release their button press if the stimulus disappeared or if they shifted their attention to a different example of the target feature and to only press the button again when they detected an example of the particular attended stimulus. During passive viewing, participants were instructed to view the stimulus without explicitly directing attention to any single stimulus feature.

The order of attention cues was semi-randomised in each run. Each run always contained an attend to faces and a passive viewing block, and at least one instance of each of the three low-level visual feature categories. There were 3 possible orders of randomisation for each run, which were

counterbalanced across participants. Additionally, the order of video clips was randomised across participants (see Figure 6.1).

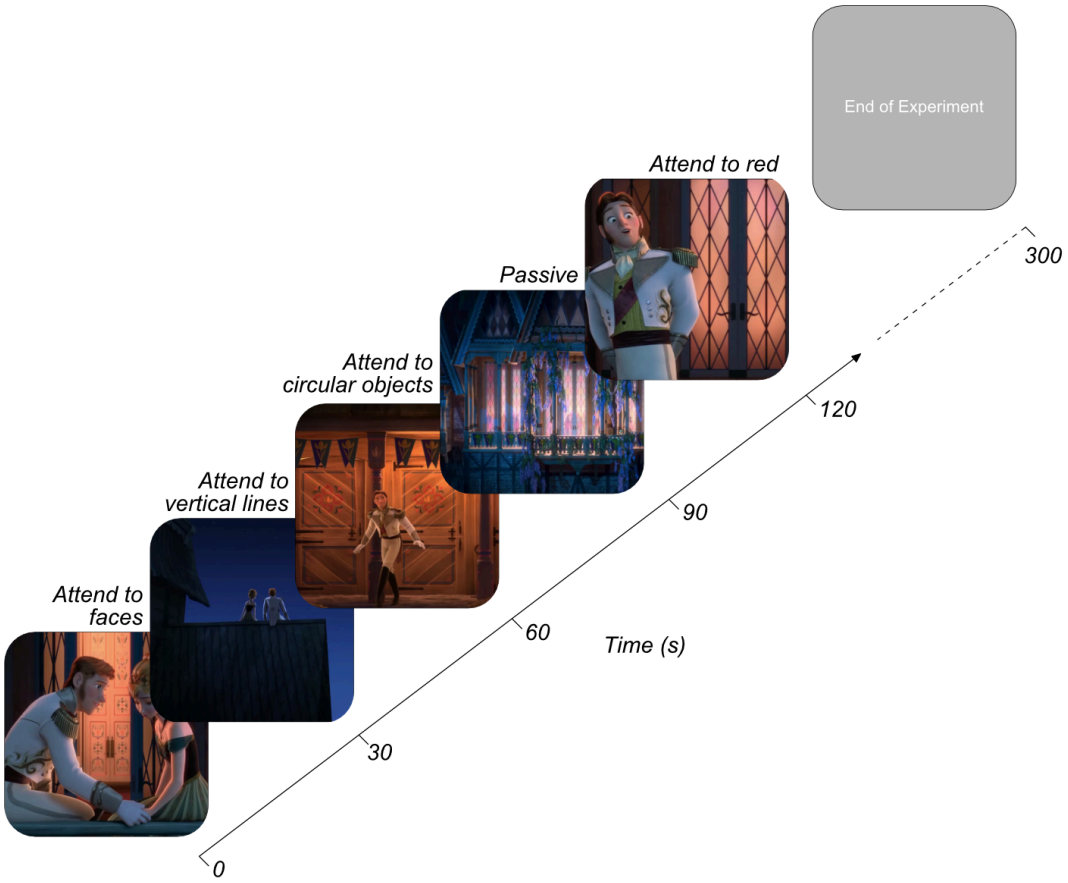

Figure 6.1 Naturalistic fMRI Attentional Modulation Experimental Design. At the start of each 30 second block an audio cue directed participants attention towards a particular stimulus feature. Participants freely viewed the Pixar movie *Frozen* (individual frames shown here for demonstration) and responded whenever they identified the feature of cued attentional focus. The order of attention cues and movie clips were randomised across participants.

6.3.5 Defining regions of interest (ROIs)

As in previous experiments, visual regions V1, V2, V3A/B, hV4, LO-1, LO-2, MT+ and IPS0 were manually delineated on the basis of polar angle reversals and eccentricity for each participant (Dumoulin & Wandell, 2008)

(for full details see section 2.2 and Chapter 4). For each participant, we produced a control auditory cortex 5cm spherical ROI from standardised Talairach coordinates centred on LH: -49, -20, 9, RH: 48, -21, 10 (Lacadie et al., 2008). In the same format, we specified three 5cm spherical ROIs in regions known for their involvement in face stimuli. The STS was defined from average Talairach coordinates stated in Vander Wyk, Voos, & Pelphrey, (2012), LH; 56.9, -42.6, 7.6, RH: 54.8, -40.9, 13.4. OFA ROIs were centred on LH: -34, -81, -14, RH: 38, -80, -7 Talairach coordinates (Rossion et al., 2003). Finally, FFA coordinates were extracted from Goffaux, Schiltz, Mur, & Goebel, (2012), LH: -38, -45, -18, RH: 37, -42, -19.

6.3.6 fMRI pre-processing

Functional data were analysed using MATLAB 2016a (MathWorks, MA) and VISTA software (https://vistalab.stanford.edu/software/) (Vista Lab, Stanford University). The first 8 seconds of each run were manually discarded prior to analysis to mitigate T1 saturation effects. Between- and within-scan motion correction was performed to compensate for motion artefacts occurring during the scan session. Any scans with > 3mm movement were removed from further analysis (no attentional modulation runs were removed on the basis of excessive movement). The VISTA rxAlign tool was then used to co-register the T1-weighted FLAIR scan to the high-resolution T1-weighted full-brain anatomical scan. We applied a manual alignment using landmark points to bring the two volumes into approximate register, followed by a robust EM-based registration algorithm to fine-rune the alignment (Nestares & Heeger, 2000). The final alignment was visually inspected to ensure the

automatic registration procedure optimised the fit. These functional data were then interpolated to the anatomical segmentation.

General linear models (GLM) were then implemented to test the contribution of each stimulus contribution to the BOLD time course (Friston et al., 1998). We used the double-gamma Boynton HRF from the SPM8 toolbox, and fit the model to an averaged time course of BOLD signal of each stimulus by minimising the sum of squared errors (RSS) between the predicted timeseries and measured BOLD response (Penny et al., 2006). In this GLM, beta weights for each condition were obtained by multiple linear regression.

The first GLM analysed attentional modulation across general low-level visual feature categories (3-feature). Data were summarised as 30 second blocks of attention to orientation, colour or shape (collapsed across the subcategories of each condition respectively, e.g. orientation was averaged across vertical, horizontal and diagonal blocks). 30 second blocks of attention to faces and passive viewing were also entered into the GLM. This resulted in between 8 and 24 beta weight estimates for across conditions at the multivariate level.

The second GLM split analysis into the individual subcategories of each low-level visual feature (3x3-feature). 30 second blocks were classified as attention to vertical, horizontal, diagonal, red, green, blue, circular, square, triangular, face or passive viewing conditions. This resulted in between 4 and 8 beta weights across conditions at the multivariate level.

6.3.7 Statistical Analysis

6.3.7.1 Univariate

We performed very similar univariate, multivariate and connectivity analysis pipelines as employed in Chapters 4 and 5. Firstly, we extracted the percentage variance explained in our GLMs from each ROI. Participants with less than 5% variance explained (averaged across all visual ROIs; V1, V2, V3A/B, hV4, LO-1, LO-2 and MT+) were removed from further analysis. This resulted in the removal of two participants from the 3-feature analysis pipeline. Hence, data from 12 participants was included in the 3-feature analysis, and data from all 14 participants was included in the 3x3-feature analysis. Attentional modulation data were then analysed through multiple one-way repeated ANOVAs, Benjamini-Hochberg corrected across the number of ROIs to assess attentional modulation differences in both signal modulation across conditions, for both the 3-feature and 3x3-feature analyses respectively.

6.3.7.2 Multivariate

To investigate voxel-level patterns of activation in each visual ROI, beta weights were calculated for each attentional focus predictor timeseries (for both the 3-feature and 3x3-feature analyses) for each voxel, using deconvolution to examine the fit of the HRF to the timeseries data. In each ROI, we selected the 100 voxels that explained the largest amount of variance in the GLM across conditions. These multivariate beta weights were z-scored across voxels and input into a weighted support vector machine (SVM) classification, employing a radial basis function kernel and leave-one-

out cross validation to decode attentional focus (Chang & Lin, 2011). We initially assessed multi-class decoding accuracy across the attention conditions, supplying data from multiple conditions simultaneously using the 'one-against-one' approach (Knerr et al., 1990). This analysis yielded a single classification score for each participant for the 3-feature analysis. For the 3x3-feature analysis we performed a separate multi-class classification for each low-level visual feature, e.g. classifying horizontal, vertical and diagonal orientations simultaneously. One-sample Wilcoxon signed rank tests were then performed to assess if classification accuracies differed significantly from chance level performance, which were Benjamini-Hochberg corrected across ROIs to control for false discovery rate.

Following this, in order to identify the driving forces behind any differential patterns of BOLD signal modulation as a functional of attentional state, we performed pairwise classification between comparisons of two attentional conditions in turn. For the 3-feature condition, weighted classification was performed between pairwise combinations of orientation versus colour, orientation versus shape, orientation versus faces, colour versus shape, colour versus faces and shape versus faces. For the 3x3-analysis, a weighted classification was performed for each of the three low-level stimulus conditions individually. For example, in the orientation analysis, we assessed classification accuracies between vertical versus horizontal, vertical versus diagonal and horizontal versus diagonal orientations independently.

6.3.7.3 Timeseries connectivity analysis

To examine featural attention-specific connectivity between ROIs, we performed the same connectivity analyses as used in Chapters 4 and 5. However, as in Chapter 5, in both the 3-feature and 3x3-feature analyses, there were not an equal number of repetitions across all attentional conditions. Therefore, we performed bootstrapping analyses to simulate data for particular conditions. We calculated the minimum and maximum number of condition repetitions (in TRs) across all attention categories. For each category, we then calculated the number of bootstrapped runs required to reach the maximum number of repetitions. For each of these required bootstrapped runs, we sampled (with replacement) half of the data across TRs for that condition and calculated the mean across those TRs to provide a simulated timeseries across all voxels for that TR. We repeated the same process across all condition repetitions for each participant to ensure every attentional condition had the same number of repetitions. The same process was performed for both the 3-feature and 3x3-feature analyses.

We then employed the same connectivity analysis as used in our previous experiments. This involved extracting each participant's multivariate timeseries data, removing noise from and averaging across all voxels in an ROI to yield a single univariate timeseries (see Chapter 4 for a detailed explanation).

Kendall's tau correlations were then performed for all pairwise combinations of ROI (V1, V3A/B, hV4, LO-1, LO-2 and IPS0), producing a correlation

matrix for each attentional condition. We then assessed the similarity of these patterns of connectivity across ROIs in each condition, though a one-way repeated measure ANOVA with Bonferroni-corrected post-hoc tests for the 3-feature and the three, 3x3-feature analyses respectively.

Next, we wished to investigate if particular ROIs were driving potential differences in connectivity across visual regions as a function of attentional task. For each ROI, we extracted data reflecting the correlation of that ROI with all other ROIs for each participant. From all participants, we then sampled data with replacement, with a sample reflecting a full complement of ROI-specific correlation data for each attentional condition, and we calculated the mean across these samples. To simulate noise, from those same samples, for each pairwise comparison of conditions, we switched the condition labels approximately 50% of the time, keeping ROI-ROI relationships constant. We then calculated the average across these scrambled condition-specific datasets.

Following this, we calculated the root mean squared error (RMSE) distance between each pairwise combination of condition vectors for both the observed and noise data. This provided a measure of differences in connectivity associated with a particular ROI across featural attention conditions. We repeated this process across 10,000 iterations for each ROI comparison. Across all iterations, we calculated the percentage of observed RMSEs for a pairwise comparison falling below the RMSE of a comparable simulated noise distribution. Any percentile below 5% indicated a difference

in ROI-specific patterns of connectivity between two attentional modulation conditions which was significantly larger than predicted by chance (p <.05).

6.4 Results

6.4.1 Univariate- 3 feature attentional modulation

We wished to examine changes in BOLD signal modulation as a function of attentional focus directed toward orientation, colour, shape or face attributes in a complex, dynamic visual stimulus. Previous research suggests some visual regions (particularly the FFA, OFA and STS) possess particular preferences for a single visual feature, and as such we predicted we may see significant differences in fMRI activity for these particular features in those ROIs (e.g. Allison et al., 2000; Gauthier et al., 2000; Kanwisher et al., 1997; Silson et al., 2013).

We first sought to identify evidence of attentional modulation with this uncontrolled, naturalistic stimulus. We calculated an average beta value for each participant across the four attentional conditions (orientation, colour, shape and faces) and performed multiple Wilcoxon signed-rank tests to assess if the average level of attentional modulation was significantly different from zero, for all ROIs examined. The resulting significance values were Benjamini-Hochberg corrected across ROIs to control for the increased risk of a false positive error. In almost all visual ROIs examined, average attentional modulation of the BOLD signal was significantly greater than zero (V1: W(11) = 3.06, p = .004, V3A/B: W(11) = 3.06, p = .004, hV4: W(11) = 3.06, p = .004, LO-1: W(11) = 2.98, p = .004, LO-2: W(11) = 3.06, p = .004,

OFA: W(11) = 2.90, p = .005). In our control auditory cortex ROI (A1), average attentional modulation was significantly below zero (W(11) = -2.98, p = .043). In both the FFA and STS, averaged attentional modulation of the BOLD signal was not significantly different from zero (p > .05).

We then performed multiple one-way repeated measures ANOVAs, Benjamini-Hochberg corrected for the number of ROIs, with Bonferroni-corrected post-hoc tests to assess differences in BOLD signal modulation during attention to orientation, colour, shape or faces in our naturalistic stimulus. Across almost all visual ROIs examined we see no significant difference in BOLD signal modulation between attention to orientation, colour and shape (p > .05) (see Figure.2A). In V3A/B, we identify significantly greater signal modulation during attention to orientation than to colour or shape respectively (F(3,33) = 52.24, p <.001, post-hoc p = .005, p = .047 respectively). In LO-2, we identify significantly increased BOLD signal modulation during attention to shape than colour (F(3,33) = 19.07, p <.001, post-hoc p = .003). In regions V1, LO-1, FFA, OFA, STS and A1, we identify no significant differences in patterns of BOLD modulation across attentional states (p > .05).

Despite, detecting relatively little evidence of univariate BOLD signal changes across our three low-level visual feature attention conditions, in V3A/B, hV4, and LO-2, we identify significantly reduced BOLD signal modulation during attention to faces than to any other visual feature (V3A/B; versus orientation, p <.001, versus colour, p = .001 and versus shape, p

<.001, hV4; ($F(3,33) = 16.39$, p <.001, versus orientation, p = .002, versus colour, p = .043 and versus shape p <.001, LO-2; versus orientation, p = .012, versus colour, p = .041 and versus shape p <.001). In our control auditory cortex ROI (A1), we identify no significant difference in BOLD signal modulation across all attentional conditions ($F(3,33) = 1.60$, p = .256) (See Figure 6.2).

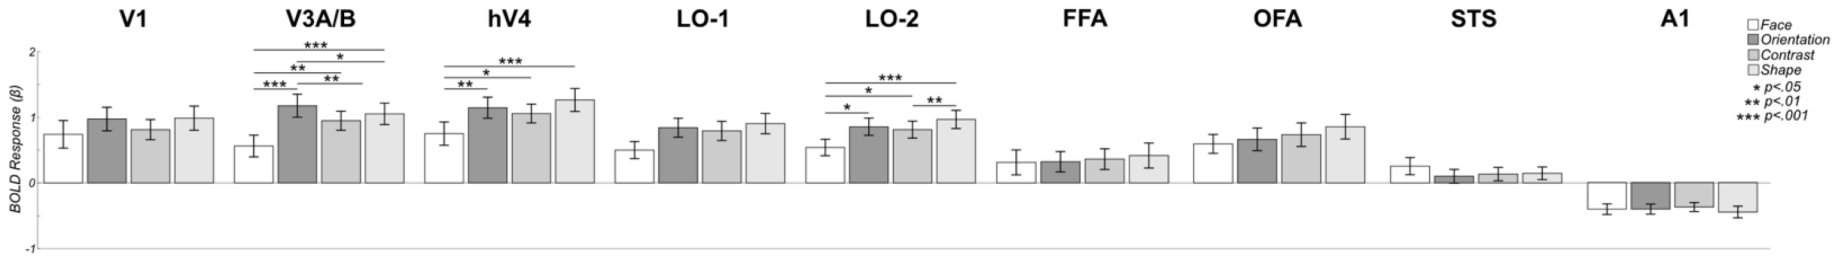

Figure 6.2 Univariate differences in attentional modulation are predominantly driven by faces. Data demonstrates feature-specific attentional modulations in BOLD signal. Data are shown for 8 visual ROIs as well as a control area (auditory cortex- A1). Regions V3A/B, hV4 and LO-2 exhibit significantly reduced BOLD signal amplitude during attention to faces in comparison to low-level visual features. Visual areas V3A/B and LO-2 show significantly increased modulation during attention to orientation and shape respectively. All error bars reflect +/- 1 SEM and significance asterisks indicate Benjamini-Hochberg corrected values.

6.4.2 Univariate- 3x3 feature attentional modulation

In our analysis of fine-scale stimulus dimensions in an individual visual feature, we again wished to assess if attention directed towards these relatively small differences in a category elicited BOLD signal modulation significantly greater than zero. For each stimulus category (orientation, colour and shape), we calculated the average beta value across the three feature dimensions (e.g. across vertical, horizontal and diagonal orientations) for each participant. We performed multiple Wilcoxon signed-rank tests; Benjamini-Hochberg corrected across ROIs to assess if levels of averaged BOLD signal attentional modulation were significantly different from zero.

Table 6.1 Almost all visual ROIs exhibit significantly greater BOLD signal modulation during attention versus an absence of signal.

	Orientation		Colour		Shape	
ROI	W	*p*	W	*p*	W	*p*
V1	3.30	.002**	3.30	.002**	3.30	.002**
V3A/B	3.30	.002**	3.30	.002**	3.30	.002**
hV4	3.30	.002**	3.30	.002**	3.30	.002**
LO-1	3.23	.002**	3.17	.002**	3.23	.002**
LO-2	3.30	.002**	3.23	.002**	3.30	.002**
OFA	3.04	.003**	3.23	.002**	3.17	.002**
FFA	2.10	.040*	2.29	.025*	2.17	.034*
STS	-0.09	.925	0.28	.778	0.28	.778
A1	-3.23	.002**	-3.17	.002**	-3.17	.002**

** p <.05, ** p<.01, *** p<.001*

In all analyses, almost all visual ROIs (excluding the STS) exhibited significantly greater BOLD signal modulation during attention versus an absence of signal modulation (0). In A1, in all conditions, averaged attentional BOLD signal modulation was significantly below zero (see Table 6.1).

Next, we performed multiple one-way repeated measures ANOVAs, with associated Bonferroni-corrected post-hoc tests to compare BOLD signal modulation across the three stimulus dimensions in each low-level visual feature category respectively.

In the orientation analysis, in V1, V3A/B, FFA, OFA, STS and A1, we identify no significant differences in BOLD signal modulation across attention to vertical, horizontal and diagonally-oriented stimulus conditions ($p > .05$). In both visual areas hV4 and LO-1, we identify significantly greater modulation to horizontal than vertically-oriented stimuli (hV4: $F(2,26) = 6.27$, $p = .018$, post-hoc $p < .001$, LO-1: $F(2,26) = 7.17$, $p = .015$, post-hoc $p = .015$). In LO-2, we demonstrate significantly reduced modulation during attention to vertical than attention to both horizontal- and diagonally-oriented stimuli ($F(2,26) = 9.09$, $p = .009$, post-hoc $p = .015$, $p = .024$ respectively) (see Figure 6.3A).

In our chromatic analysis, we identified no significant differences in BOLD signal modulation during attention directed toward red, green, or blue stimuli in any ROI examined ($p > .05$) (see Figure 6.3B). In almost all ROIs (V1, hV4, LO-1, LO-2, FFA, OFA and A1), we also identify no significant differences in fMRI activation across attention to circular, square and triangular shape conditions. In V3A/B however, we demonstrate significantly reduced modulation during attention to circular stimuli in comparison to attention to square or triangular stimulus attributes ($F(2,26) = 6.85$, $p = .004$, post-hoc $p = .040$ and $p = .019$ respectively). In the STS, we identify

significantly reduced modulation during attention directed towards triangular stimuli than to circular stimuli ($F(2,26) = 7.54$, $p = .003$, post-hoc $p = .012$) (see Figure 6.3C).

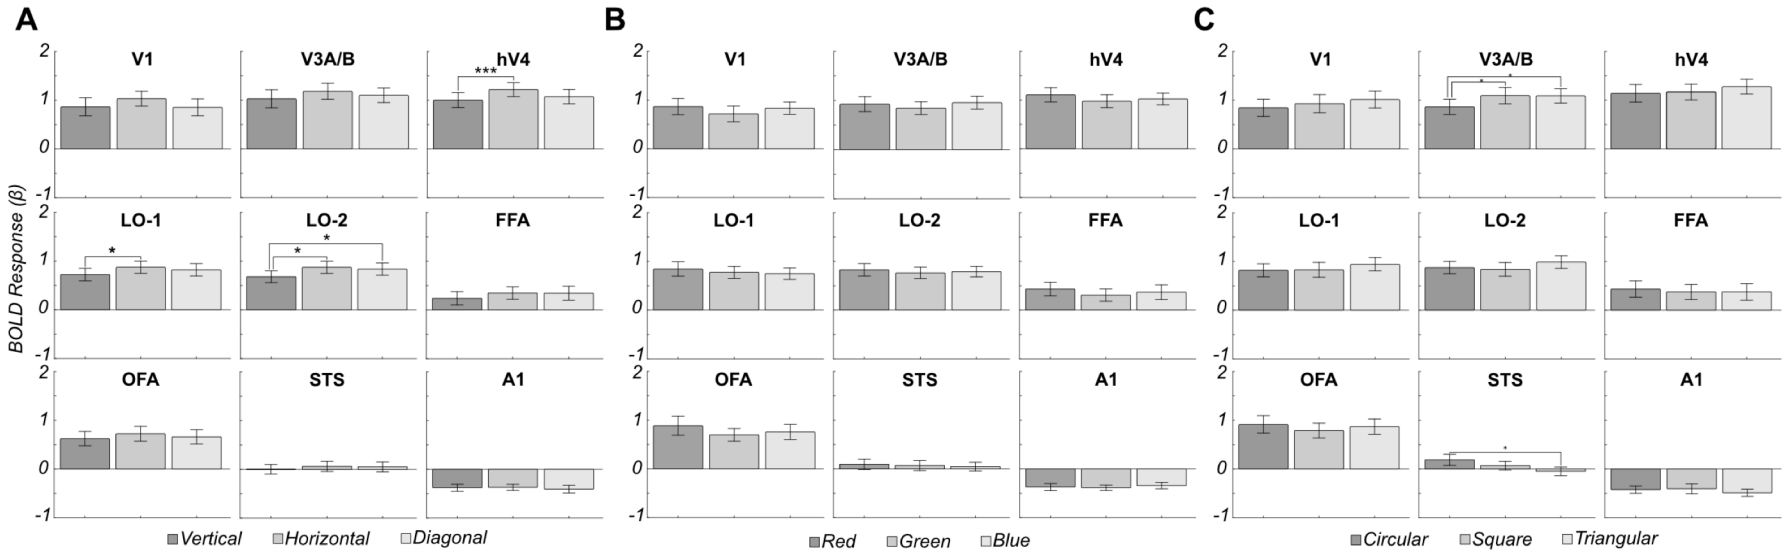

Figure 6.3 Weak evidence of feature-specific attentional modulation within low-level visual feature dimensions. Within the orientation analysis (A), in hV4, LO-1 and LO-2, we identify significantly greater modulation to horizontal than to vertically-oriented stimuli. In LO-2, evidence of significantly greater modulation during attention to diagonally- than vertically-oriented stimuli is also evident. In the chromatic analysis (B) in all 9 ROIs examined, we identify no significant difference in BOLD signal modulation during attention to red, green or blue stimuli. Within the shape analysis (C), we demonstrate significantly reduced modulation during attention to circular than to square or triangular stimuli. In the STS, significantly reduced modulation during triangular than circular attention is also evident. All error bars reflect +/- 1 SEM and significance asterisks indicate Benjamini-Hochberg corrected values.

6.4.3 Multivariate- 3 feature attentional modulation

Using a complex, dynamic, naturalistic stimulus in this experiment provided a wealth of visual information for our participants to selectively to attend to. Our univariate analyses revealed however, that we can predominantly only distinguish attention to faces, versus attention to low-level stimulus attributes in a limited subset of visual ROIs. We therefore asked, as identified in Chapters 4 and 5, if attentional modulation exerts a differential influence on voxel-level patterns of activation in visual ROIs.

We performed weighted four-way support vector machine multivariate pattern classification analysis, simultaneously classifying orientation, colour, shape and face attention data. One-sample Wilcoxon signed rank tests versus chance performance (25%), Benjamini-Hochberg corrected for the number of ROIs, revealed attentional state could be decoded at rates significantly greater than chance in almost all ROIs examined (V1: $W(11) = 2.67$, $p = .032$, V3A/B: $W(11) = 2.82$, $p = .024$, hV4: $W(11) = 2.51$, $p = .044$, LO-1: $W(11) = 2.94$, $p = .023$, FFA: $W(11) = 3.06$, $p = .023$, OFA: $W(11) = 2.98$, $p = .023$ and STS: $W(11) = 2.90$, $p = .023$). In LO-2 and A1, classification accuracies did not significantly exceed chance ($p > .05$). Mean decoding accuracy across ROIs ranged from 27.86 to 38.91% (see Figure 6.4A).

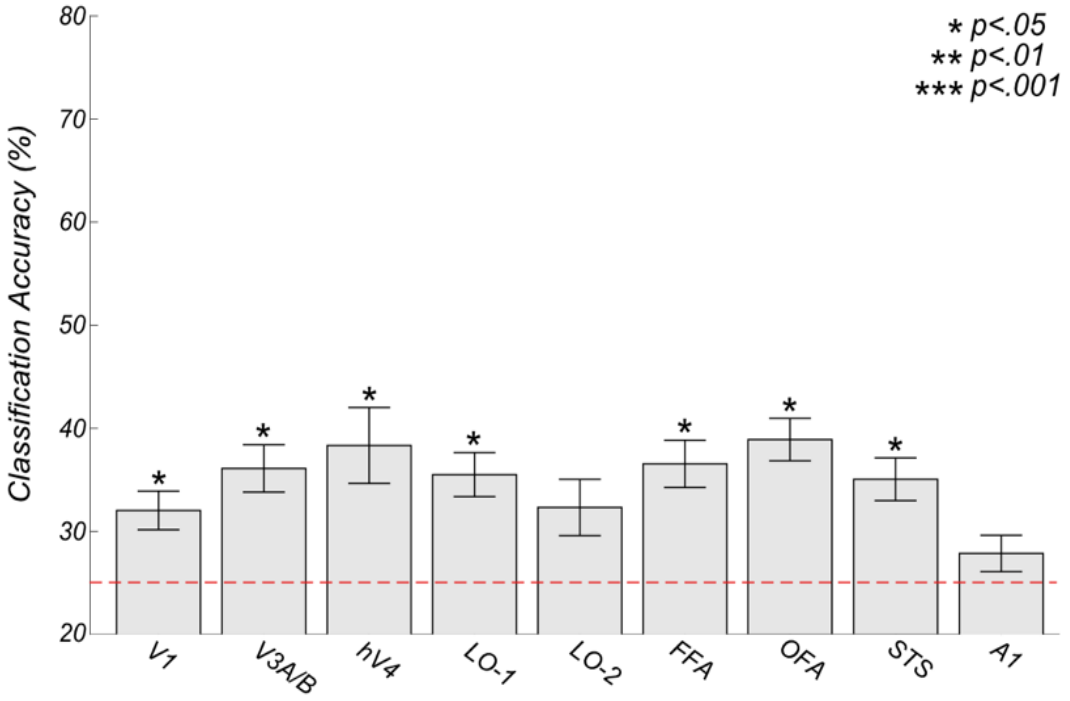

Figure 6.4 Multivariate Support Vector Machine decoding. Simultaneous classification of face, orientation, contrast and shape attention data reveal attentional state can be accurately decoded in almost all visual ROIs examined. Error bars reflect +/- 1 SEM and significance asterisks indicate Benjamini-Hochberg corrected values.

In order to identify which attentional conditions specifically were driving these above chance classifications, we then performed multiple pairwise classification analyses, comparing patterns of activation across two attentional conditions at a time. Significance was again assessed using one sample Wilcoxon signed-rank tests versus chance (50%), Benjamini-Hochberg corrected for the number of comparisons across ROIs. In V1, it was possible to decode attention to faces versus orientation ($W(11) = 3.07$, $p = .028$) and colour ($W(11) = 2.99$, $p = .031$) at rates significantly above chance (see Figure 6.5).

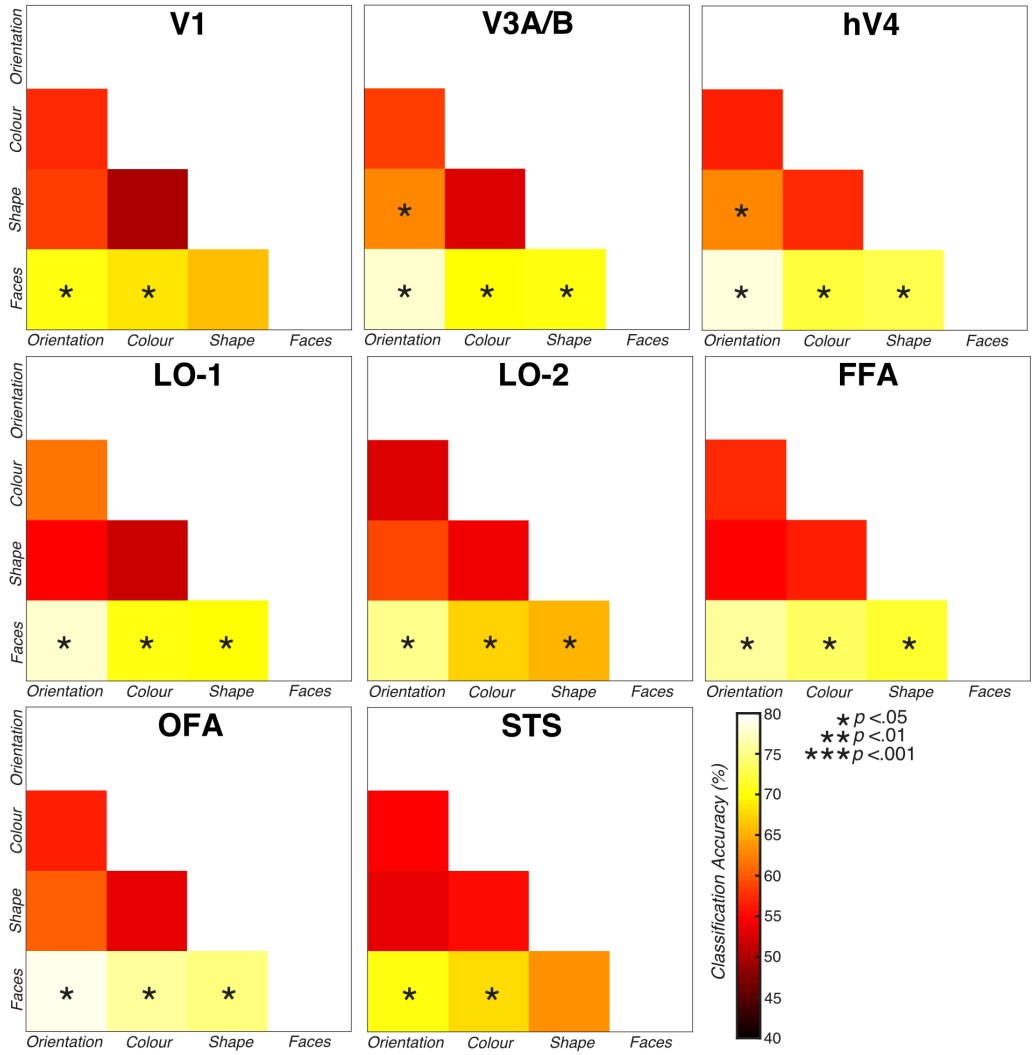

Figure 6.5 Two-way support vector machine decoding. Within all visual ROIs examined, successful classification was predominantly evident between attention to face and attention to low-level visual feature conditions. Within V3A/B and hV4, voxel patterns differ significantly between attention to orientation and attention to shape conditions. Significance asterisks indicate Benjamini-Hochberg corrected values.

In V3A/B, hV4, LO-2, FFA and OFA, successful classification was evident between attention to faces versus orientation, colour and shape respectively. In the STS, successful classification was identified between faces and

orientation and faces and colour. Finally, in both V3A/B and hV4, successful classification was evident between orientation and shape (see Table 6.2).

Table 6.2 Above-chance classification is predominantly driven by differences in multivariate patterns of activation between attention to low-level visual features and faces.

	Orientation vs. Colour		Orientation vs. Shape		Orientation vs. Face		Colour vs. Shape		Colour vs. Face		Shape vs. Face	
ROI	*Z*	*p*	*Z*	*p*	*Z*	*p*	*Z*	*p*	*Z*	*p*	*Z*	*p*
V1	2.40	.120	2.12	.228	3.07	.028*	-0.10	4.10	2.99	.031*	2.67	.062
V3A/B	1.73	.048*	2.90	.036*	3.07	.028*	0.94	1.65	3.06	.028*	3.06	.028*
hV4	1.65	.560	2.79	.047*	3.07	.02*8	1.48	.752	2.98	.031*	2.98	.031*
LO-1	2.59	.076	1.14	1.24	3.06	.028*	.075	2.117	3.07	.028*	2.95	.033*
LO-2	0.53	2.702	1.83	.417	3.07	.028*	1.58	.608	3.06	.028*	2.83	.044*
FFA	2.04	.268	1.81	.417	3.07	.028*	1.82	.417	3.07	.028*	3.07	.028*
OFA	2.20	.199	2.75	.052	3.06	.028*	1.22	1.114	3.06	.028*	3.06	.028*
STS	1.37	.865	1.58	.608	3.06	.028*	2.12	.228	3.06	.028*	2.51	.092

** p <.05, ** p<.01, *** p<.001*

Hence, the vast majority of these successful classifications at the pairwise level were driven by differences in patterns of activation between attention to faces and attention to low-level visual stimuli. To support this, a further multi-class (three-way) SVM decoding analysis was performed, classifying orientation, colour and shape data simultaneously. This identified no classification accuracies significantly above chance (33.33%) in any ROI examined ($p > .05$).

6.4.4 Multivariate- 3x3 feature attentional modulation

We then repeated the SVM analysis for each low-level visual feature category independently to investigate if was possible to detect fine-scale, voxel-level patterns of modulation which differed across dimensions in a feature category. We first performed multi-class classification and performed multiple one-sample Wilcoxon signed-rank tests assessing classification performance versus chance (33.33%). These analyses revealed no above-chance classification performance across vertical, horizontal or diagonal orientations, red, green or blue stimulus chromaticities or circular, square or triangular shapes respectively ($p > .05$). Mean classification accuracy ranged from 29.29-37.59% for the orientation analysis (see Figure 6.6A), 29.86-36.27% for attention to colour (see Figure 6.6B) and 35.57-47.71% for attention to shape conditions (see Figure 6.6C).

We again performed classification on all pairwise combinations of conditions in each feature category. In all classifications performed, we identified no

significantly above-chance (50%) decoding accuracies for any of the stimulus feature conditions ($p > .05$).

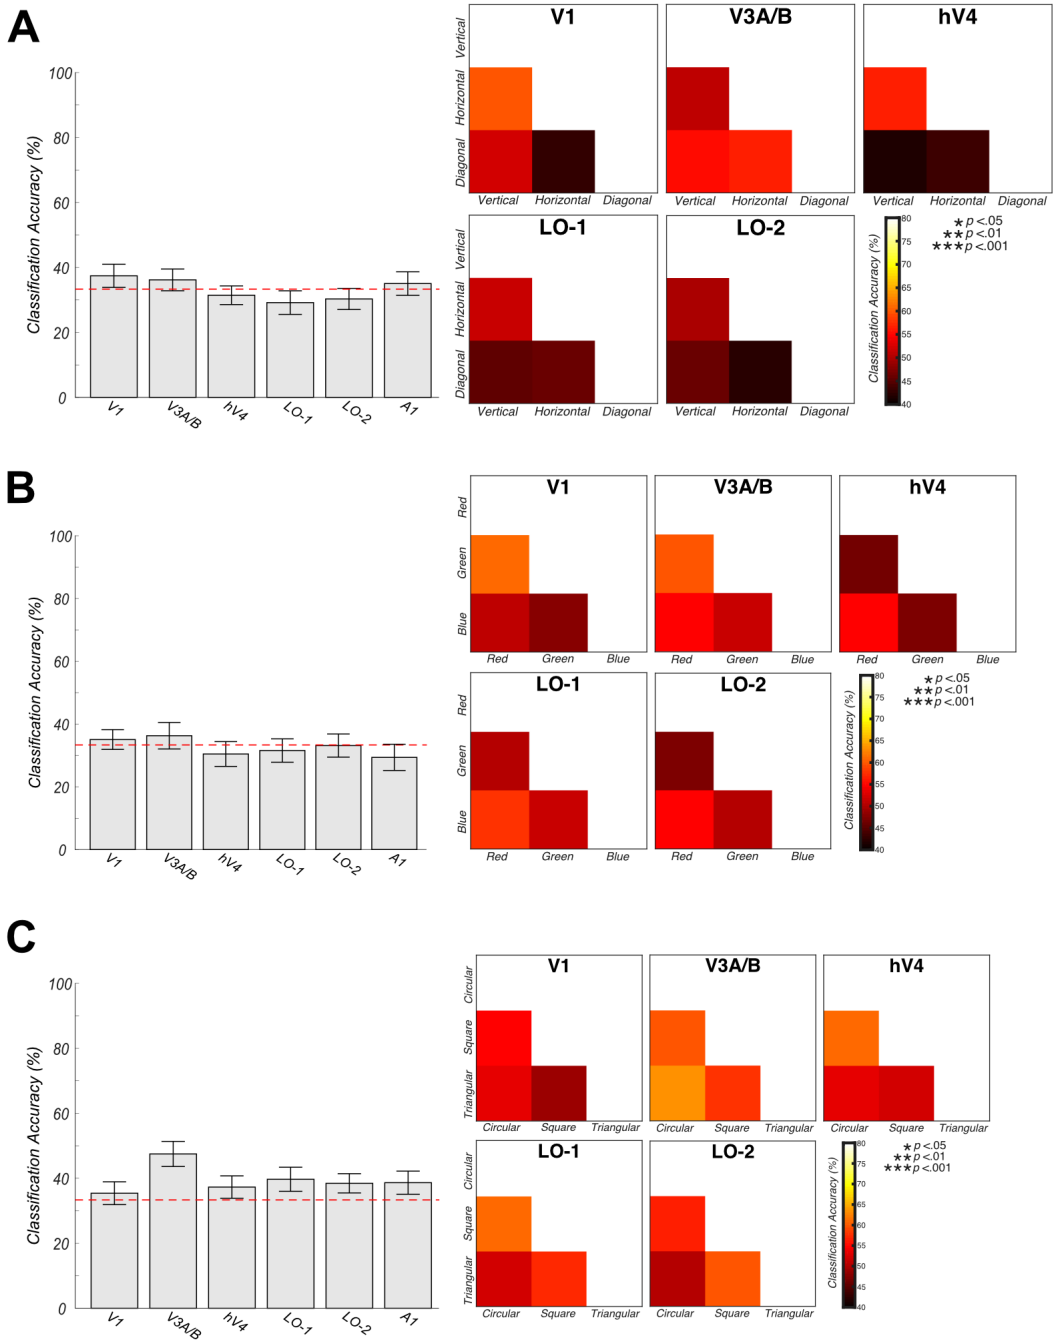

Figure 6.6 3x3 feature multivariate support vector machine decoding. Across orientation (A), colour (B) shape (C) low-level feature analysis conditions, no significant classification of attentional state was evident within multiclass or pairwise support vector machine decoding. Error bars reflect +/- 1 SEM.

6.4.5 Timeseries connectivity analysis

6.4.5.1 3-feature

Our univariate analysis demonstrated that, differences in BOLD responses across attentional conditions were driven largely by a reduction in BOLD signal modulation during attention to faces. Hence, we identified relatively little evidence of feature-specific attentional modulation at the univariate level. In our final analysis, we asked whether attention may influence the way individual ROIs interact across the visual cortex, examining changes in functional connectivity as a function of attentional state.

We performed non-parametric Kendall's Tau correlations between univariate time-courses from each pairwise combination of visual ROIs, for each participant and attentional state. We averaged these normalised, Fisher-z transformed correlation matrices across individuals to demonstrate group-level changes in connectivity between ROIs as a function of attentional task (see Figure. 6.7A). A one-way repeated-measures ANOVA with Bonferroni-corrected post-hoc tests revealed significantly greater positive correlation between ROIs during passive viewing than during all low-level visual feature conditions (orientation, colour and shape) ($F(4,56) = 45.09$, $p < .001$, post-hoc $p < .001$ for all comparisons). We also identified significantly greater positive connectivity between ROIs during attention to faces than low-level visual features (all $p < .001$). Finally, we also identified significantly greater negative correlation between ROIs during attention to colour than to orientation ($p < .001$).

We then asked whether these differences in connectivity between attentional task conditions were driven by a particular subset of ROIs, or whether patterns of connectivity were consistent across visual ROIs. We calculated the Euclidean distance (RMSE) between pairwise comparisons of attentional task condition (averaged across a subsample of participants for each condition) for both observed and noise (scrambled attention-condition label) datasets, bootstrapped across 10,000 iterations. In almost all ROIs examined, we identify similar patterns of partial correlation as a function of attentional task. Almost all regions investigated demonstrate significantly greater connectivity during attention to faces versus all low-level visual attention conditions (V1: p = .018, p = .002, p = .001, V3A/B: colour p = .014, shape p = .030, hV4: p = .008, p <.001, p <.001, LO-1: p = .009, p <.001, p = .006, LO-2: p = .011, p <.001, p <.001, IPS0: all p <.001 respectively). Almost all ROI partial correlations also show significantly greater positive connectivity during passive viewing compared with attention to low-level visual features (V1: p = .004, p = .002, p = .001, V3A/B: colour p = .002, shape p = .012, hV4: p = .012, p = .001, p <.001, LO-1: all p <.001, LO-2: p = .013, p <.001, p <.001, IPS0: p = .001, p <.001, p <.001). Almost all ROIs also show significant differences in connectivity between orientation and colour conditions (hV4: p = .012, LO-1: p <.001, LO-2: p = .001, IPS0: p = .008) (see Figure 6.7B).

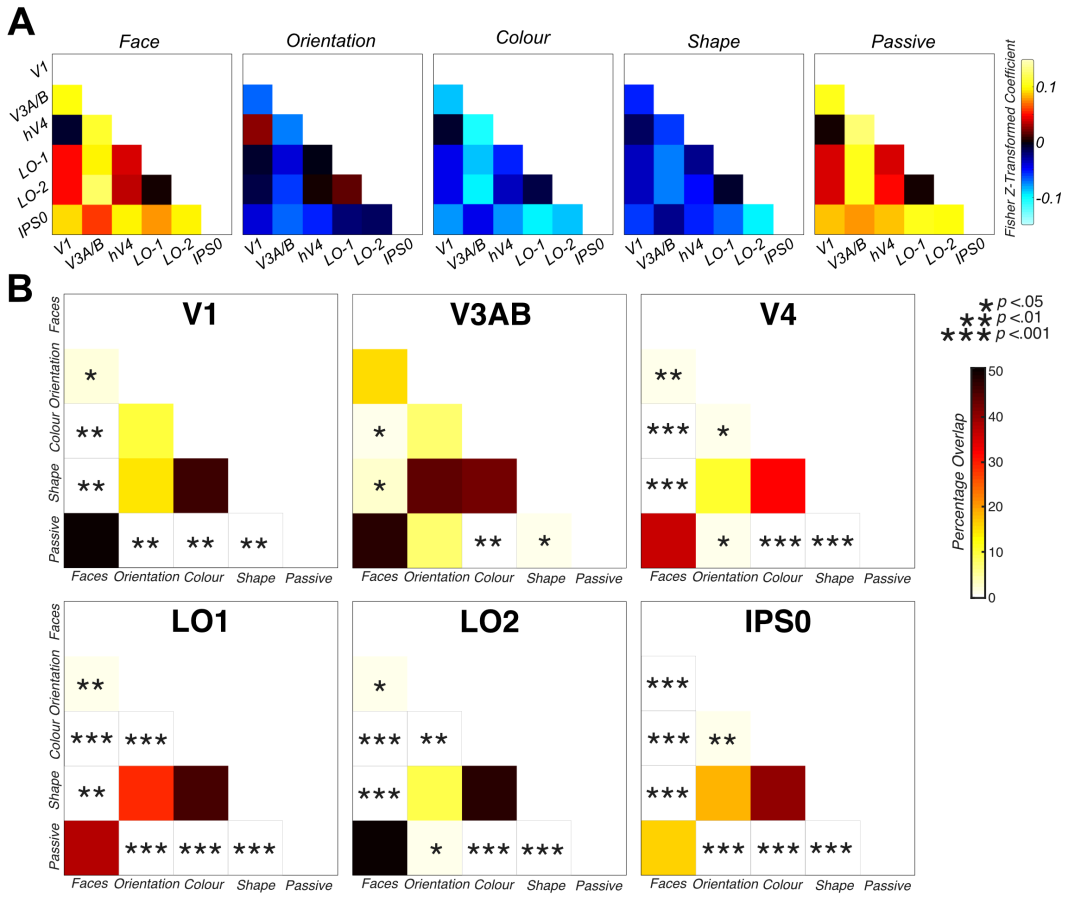

Figure 6.7 Greater overall connectivity during passive viewing and attention directed to faces than attention directed to low-level stimulus features. A) indicates group feature-specific averaged attentional modulation connectivity values. These correlation matrices indicate significantly greater positive connectivity during attention to faces and passive viewing. B) Bootstrapped measures of distance between pairwise combinations of attentional focus conditions, for V1, V3A/B, hV4, LO-1, LO-2 and IPS0 ROIs. The figure demonstrates the percentage overlap between the distribution of RMSE scores across 10,000 iterations of randomly-selected samples of the observed data and scrambled correlation 'noise' data, for each ROI across multiple pairwise comparisons. Significant overlap (less than 5%) between pairwise combinations of condition are indicated with asterisks.

6.4.5.1 3x3-feature

We also sought to examine patterns of connectivity across ROIs as a function of fine-grain differences in a low-level visual feature category (e.g. comparing patterns of connectivity during attention to vertical, horizontal and diagonally-oriented lines). We computed non-parametric Kendall's Tau correlations between ROI-averaged time courses from each pairwise combination of visual ROIs for each featural attention condition, separately for orientation, colour and shape analyses. One-way repeated measures ANOVAs with Bonferroni-corrected post-hoc tests conducted on these normalised, Fisher-z transformed correlation coefficients revealed differing patterns of connectivity as a function of attention to different stimulus dimensions in a feature category. We identified significantly greater negative connectivity during attention to diagonal than to vertical or horizontal stimuli ($F(2,28) = 16.86$, $p < .001$, all post-hoc $p < .001$) (see Figure 6.8A). We also identified significantly greater connectivity during attention to red than to green or blue coloured stimuli ($F(2,28) = 23.37$, $p < .001$, all post-hoc $p < .001$) (see Figure 6.8B). Finally, in the shape analysis, we detected significantly greater connectivity during attention to square stimuli than attention to circular or triangular ($F(2,28) = 83.27$, $p < .001$, post-hoc $p < .001$) (see Figure 6.8C).

We then asked whether these differences in patterns of connectivity were driven by correlations with particular ROIs, calculating the RMSE distance between pairwise comparisons of attentional task conditions (averaged

across a subsample of participants for each condition), for both observed and noise datasets, bootstrapped across 10,000 iterations.

For the orientation analysis, in almost all ROIs examined, we demonstrated a significantly greater positive connectivity during attention to vertical versus attention to diagonal stimuli (hV4: p = .001, LO-1: p = .005, LO-2: p <.001, IPS0: p <.001). We also identified significant negative connectivity during attention to diagonal versus horizontally-oriented stimuli in many visual areas (V1: p = .007, hV4: p = .021, LO-1: p = .015, LO-2: p = <.001). Finally, we demonstrate significant differences between vertically- and horizontally-oriented stimuli in V1 (p = .050) and LO-2 (p = .045) (see Figure 6.8A).

In the colour analysis, almost all ROI partial correlations demonstrate a significantly greater positive connectivity during attention to red versus green stimuli (V3A/B: p =.022, hV4: p <.001, LO-1: p = .005, LO-2: p = .006, IPS0: p = .004). In hV4, we also see a significant difference in connectivity between attention to green and blue stimuli (p = .019) (see Figure 6.8B).

Finally, for the shape analysis, all ROIs examined demonstrate a significant greater positive connectivity during attention to square versus circular stimuli, and also between attention to square and attention to triangular stimuli (all p <.01) (see Figure 6.8C).

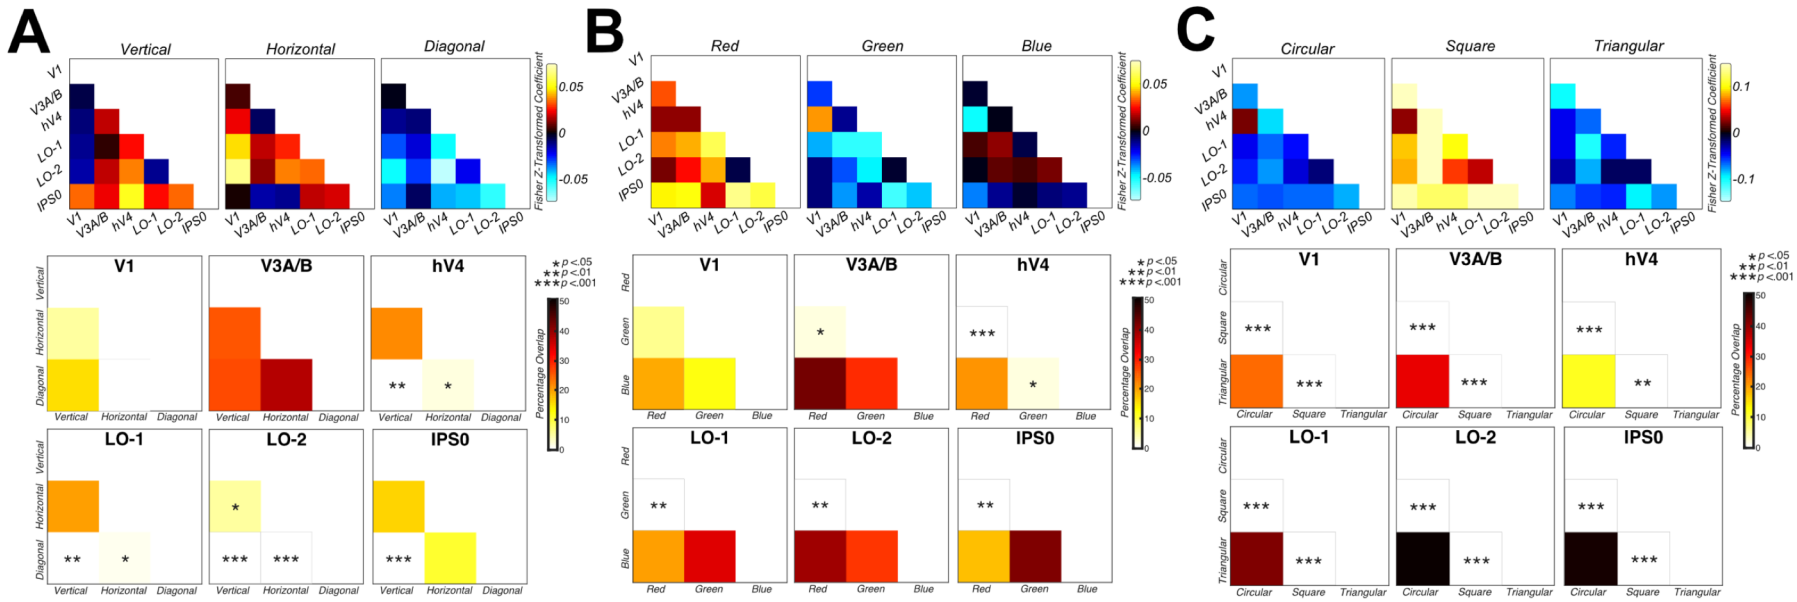

Figure 6.8 3x3 feature timeseries connectivity analyses. Group averaged feature-specific attentional modulation connectivity values, and partial correlation bootstrapped measures indicate significantly different patterns of connectivity across visual ROIs as a function of attentional task. In the orientation analysis, we identify significantly reduced correlation between ROIs during attention to diagonal than horizontal or vertically-oriented stimuli, and this different is present in connectivity with almost all visual ROIs (A). In the chromatic analysis (B), we identify significantly greater positive connectivity between ROIs during attention to red than to green or blue stimuli in almost all ROIs (excluding V3A/B). In C) we demonstrate significantly greater positive correlation between visual ROIs during attention to square than to circular- or triangular-shapes. This difference was apparent in all visual ROI RMSE partial correlations. Significant differences between pairwise combination of condition are indicated with asterisks.

6.5 Discussion

Here, we investigated patterns of attentional modulation during the viewing of a relatively naturalistic stimulus. We demonstrate little difference in univariate fMRI signal modulation across attentional tasks conditions, both at coarse (orientation, colour and shape) and relatively finer (e.g. vertical, horizontal and diagonal orientations) scales. Strikingly, we see no evidence of face-selective responses, indexed by different patterns of modulation for attention to faces than attention to low-level stimulus features, in any of the face-specific regions examined (FFA, OFA and STS).

We believe this is likely a reflection of the fact these face-specific regions have been identified traditionally through manipulation of the stimulus, rather than of attentional focus. For example Kanwisher et al., (1997) identified face-selective regions through comparison of activation elicited from passive viewing of photographs with faces compared with photographs containing common objects. Additionally, previous research demonstrating attentional modulation of the FFA employed differing paradigms to our own, such as covert attention and repetition attenuation effects (Wojciulik et al., 1998; Yi et al., 2006). Hence, we believe differences between our task design, in which participants were presented with constant information across all of our stimulus categories and were asked to overtly attend to stimulus features in a dynamic and uncontrolled stimulus, in comparison to previous attention studies presenting relatively more controlled faces in comparison to scene stimuli, may provide a suggestion as to why we fail to detect face-selective responses in these regions.

In V3A/B, hV4 and LO-2, we identify significantly reduced BOLD signal modulation during attention to faces than low-level stimulus attributes, which may reflect the relatively more complex and higher-order nature of face stimuli, with reduced processing in earlier visual areas. In V3A/B, we also identify significantly greater attentional modulation during attention to orientation, versus when attention is directed towards stimulus colour or shape. In the literature, there is no clear consensus regarding the role of V3A/B in the processing of any single visual feature, hence it's differential activation here, which we have not identified in our previous experiments (see Chapters 4 & 5) is likely a reflection of the nature of the dynamic, uncontrolled relatively naturalistic stimulus used here, which provides a far richer source of visual stimulation than our highly-controlled RF pattern stimulus.

We also identify significantly greater BOLD modulation during attention to shape versus colour in LO-2, which supports our previous findings (see Chapters 4 & 5), and provides partial support for the double-dissociation of LO-1 and LO-2, identified by Silson et al., (2013). We identify no significant changes in BOLD signal modulation in response to differing attentional task in V1 and LO-1. Overall, examining differences in signatures of attentional modulation across task at the level of individual ROIs is relatively uninformative regarding feature-specific attentional effects across the visual cortex, particularly so for the processing of relatively lower-level stimulus attributes (e.g. orientation, colour and shape).

In the 3x3 analysis, where we break down the overarching low-level stimulus categories into their three respective dimensions (e.g. red, green and blue), we again find relatively little evidence of differences in attentional modulation across the visual cortex when attending to different low-level stimulus dimensions. In our orientation analysis, we identify no evidence of differing patterns of attentional modulation when switching between attention to vertical, horizontal and diagonally-oriented stimuli in V1, V3A/B, FFA, OFA or the STS. In hV4, LO-1 and LO-2 however, we identified significantly greater BOLD modulation during attention to horizontal than to vertically-oriented stimuli. In LO-2, we also identified significantly greater signal modulation during attention to horizontal than to diagonally-oriented stimuli. These regions have been implicated in the relatively higher-level processing of form (see Loffler, 2008 for a review; Silson et al., 2013), and it is likely attending to different directions of orientation may stimulate form-selective mechanisms in some capacity, such as the processing of curvature. Orientations in this naturalistic stimulus are heavily associated with form and structure, likely to a greater extent than is the case in our comparison highly-controlled RF pattern stimulus.

We identify no significant differences in BOLD signal modulation when attention switched between red, green and blue stimuli. This again, is likely a reflection of our visual stimulus. In our previous research (see Chapter 5) we used isoluminant stimuli (which equated only the relevant chromatic, and not luminance channel) to identify a robust difference in univariate attentional modulation between red-green and blue-yellow stimuli. However, in this

experiment, our stimuli were not isoluminant, and we did not probe attention to distinct chromatic channels (e.g. L-M, S+(L-M)). Hence, our lack of difference in univariate attentional modulation between red-, green- and blue-stimuli is likely a reflection of the fact we did not probe clear and segregated channels of chromatic information in the visual cortex.

Finally, in our univariate analysis identifying differences in patterns of attentional modulation across shape dimensions, we identify no clear evidence of differential patterns of modulation across circular, square or triangular-shape attentional focus. In V3A/B, we identified significantly reduced modulation to circular than to square or triangular stimuli, and in the STS, we identified reduced modulation to triangular than to circular oriented stimuli.

Next, we sought to identify if feature-specific patterns of attentional modulation could be better distinguished when examining activation across many voxels in a particular ROI, rather than taking a relatively coarse-scale average measure to summarise activity in a visual region. We first demonstrated we could decode attentional state successfully in almost all ROIs examined, when simultaneously classifying data from attend to orientation, colour, shape and face conditions. However, further analysis demonstrated these successful classifications were predominantly driven by voxel-level differences in patterns of activation during attention to faces, versus attention to lower-level stimulus attributes (orientation, colour or shape). For example, in hV4, LO-1, LO-2, FFA and OFA, it was possible to

decode faces from orientation, colour and shape at rates significantly higher than expected by chance. In V1 and the STS, the pattern was very similar; we were able to decode multivariate voxel-level activation during attention to faces versus attention to orientation or colour. The only successful classifications not driven by differences in patterns of activation in comparison with faces, were evident in hV4, LO-1, LO-2 and IPS0 where it was possible to decode patterns of modulation across attention to orientation and colour conditions.

Overall, unlike our previous findings (see Chapters 4 & 5), and past research (e.g. Brouwer & Heeger, 2009; Kamitani & Tong, 2005, 2006; Mannion, McDonald, & Clifford, 2009; Sumner et al., 2008), we identify little evidence of our ability to decode attentional state when attending to low-level stimulus characteristics in patterns of voxel-level activation (across orientation, colour and shape). We repeated this multivariate classification analysis, also examining patterns of activation across finer-scale dimensions in an overarching stimulus category. Previous research has demonstrated the ability to decode relatively small differences in a stimulus category, such as the ability to correctly predict which of eight different stimulus orientations or directions of motion was attended (Kamitani & Tong, 2005, 2006). However, in all three low-level visual feature analyses we detect no evidence of differential patterns of voxel-level activation when we switch our focus between these relatively fine-grained differences in a low-level visual stimulus category (e.g. vertical, horizontal and diagonal orientations). These findings may reflect the uncontrolled nature our naturalistic stimuli; we did not

ask our participants to maintain central fixation, and our stimulus covered a far greater proportion of the visual field in comparison to our RF pattern stimuli. Alternatively, the previous results identified (e.g. Kamitani & Tong, 2005) may have been driven by artefactual responses linked to spatial attention or large-scale spatial biases. Attention to vertical versus horizontal central gratings might, for example, drive the vertical versus horizontal midline. Our naturalistic stimulus has no consistent spatial display of attended stimuli, and participants were free to fixate anywhere in the visual scene, therefore, our findings are unlikely to reflect any such large-scale spatial biases.

These factors, in combination with the dynamic nature of a movie stimulus may mean we do not see consistent patterns of activation across voxels, which is far likelier to result from fixed viewing of a highly controlled low-level visual stimulus. This greater variation in voxel-level activation is likely to have negatively impacted the ability of our classifier to identify consistent patterns in activation across attention to different low-level stimulus attributes. As noted in previous literature, natural stimuli do not always replicate results obtained in artificial settings (Bartels et al., 2008).

Finally, we investigated effects of attentional modulation across multiple ROIs, rather than considering any single visual region in isolation. In our 3-feature, overall visual feature category analysis, we identify relatively little evidence of attentional modulation at both the univariate and multivariate analysis levels when we remove the robust differential modulations, we see in response to relatively high-level face stimuli. Hence, as we have identified

in previous experiments (see Chapters 4 & 5), differential effects of attentional modulation in response to a naturalistic stimulus may be better characterised in measures of functional connectivity (i.e. the relative synchronisation or desynchronization of task-specific activity in particular visual regions over time). To examine this prediction, we correlated timeseries of activations for each visual feature and analysed the relative correlations between areas a function of attentional focus.

This connectivity analysis replicated a robust effect identified in our previous experiments (see Chapters 4 & 5). We demonstrated a significantly greater positive correlation between visual ROIs during passive viewing than during attention directed towards any of our low-level visual feature conditions (orientation, contrast or shape). We believe this high relative connectivity across visual areas is akin to a default-mode network type of activation, in which during rest, or during undemanding tasks, regions across the cortex possess synchronised activity, which decreases during completion of a demanding task (see Buckner, Andrews-Hanna, & Schacter, 2008, for a review).

We also identified significantly greater positive connectivity between visual ROIs during attention to faces versus all other attention conditions. These similar patterns of positive connectivity across ROIs both in the passive viewing and attend to faces conditions is likely reflective of the nature of our task. During passive viewing, participants were asked to view the stimulus, without explicitly directing attention towards any one stimulus feature.

However, faces are inherently attention-grabbing, and humans attend to faces frequently in day-to-day life for successful social interaction. Additionally, attention directed towards faces as oppose to relatively lower-level stimulus features may have resulted in more consistent eye movements, due to the inherent nature of a movie stimulus (where faces are typically the key focus of attention). These differing patterns of eye movements between faces and relatively lower-level stimuli may be reflected within both our classification and connectivity findings here. We believe in our passive viewing condition, participants' attention is likely to have been drawn towards faces in the scenes, providing a suggestion as to why patterns of connectivity in the passive and face conditions are so similar.

Our partial correlation connectivity analysis revealed that all regions of interest had significantly different patterns of connectivity across ROIs in the face and passive conditions, in comparison to attention towards low-level visual features. Additionally, in almost all visual ROIs, we also identified significantly different patterns of connectivity between orientation and colour attention conditions.

Analysis also revealed significant differences in patterns of connectivity across categories in a stimulus dimension (e.g. vertical, horizontal and diagonal orientations). In the orientation analysis, we identified significantly reduced relative connectivity during attention to diagonally-oriented stimuli in comparison with vertically- or horizontally-oriented stimuli. From our partial correlation analyses, we demonstrated these differing patterns of connectivity between diagonally oriented versus vertical or horizontal stimuli

were present in correlations with almost all visual ROIs. Finally, V1and LO-2 ROI correlations with all other ROIs exhibited significant differences in connectivity during attention to vertical versus horizontally-oriented stimuli. This provides clear evidence for differential patterns of attentional modulation in response to attending to relatively fine-grain changes in a stimulus dimension, when we study activation across several ROIs simultaneously; a result that we were unable to clearly identify at univariate or multivariate levels of analyses with this complex, naturalistic stimulus.

In the colour analysis, we also identified significantly greater overall positive connectivity during attention to red stimuli than attention to green or blue. This difference in connectivity between red- and green-attention was identified in patterns of correlation across almost all ROIs examined in our partial correlation analysis. When examining correlations across visual ROIs with hV4, we also identified significant differences in connectivity during attention to green versus attention to blue. The different patterns of correlation between hV4 and other ROIs in all chromatic conditions is reflective of previous research demonstrating a clear role for hV4 in the processing of colour (e.g. Meadows, 1974; Zeki, 1990).

Finally, in our shape analysis, we identified significantly greater overall positive connectivity during attention to square stimuli versus attention to circular, or triangular shaped features. This was reflected in the partial correlation analysis; all visual ROIs examined had different patterns of partial correlation between attention to square versus circular, or triangular-shaped stimuli.

In conclusion, here we used a relatively novel approach for investigating visual attention, using a dynamic, uncontrolled and relatively naturalistic stimulus, to assess the extent our previous findings with highly controlled, low-level visual stimuli generalise to a real-world context. We identify relatively little evidence of differential patterns of attentional modulation at both the univariate and multivariate levels when attending to low-level visual features. However, we do identify significant differing patterns of attentional modulation to a relatively higher-level face stimulus in comparison to low-level visual ROIs in many early visual areas. When examining patterns of attentional modulation across many visual ROIs simultaneously however, we identify robust differences in patterns of connectivity across attentional conditions. We demonstrate here the importance of analysing activity at multiple spatial scales to gain a more informative picture of the influence of featural attention across the visual cortex.

7. Conclusions

7.1 Overview of the thesis findings

This thesis includes four experiments which have examined the existence and spatial scale of feature-specific patterns of attentional modulation across the visual cortex. In our first psychophysical experiment, we sought to replicate the findings of the well-established selective versus distributed attentional paradigm to validate the use of the RF pattern as an effective stimulus to probe mechanisms of visual attention. The three remaining experiments investigated the existence of differential patterns of modulation across the visual cortex in response to attention directed towards different stimulus features and chromaticities in in both situations with low-level, highly controlled and dynamic, uncontrolled and relatively naturalistic visual stimuli. The four experiments described in this thesis make novel contributions in their findings and in the methods used.

Firstly, we demonstrated we were able to replicate the well-established results of the selective versus distributed attentional paradigm of Corbetta, Miezin, Dobmeyer, Shulman, & Petersen, (1990), using an RF pattern stimulus. The use of these RF patterns allowed us to not only to manipulate both the contrast and orientation of stimulus, as is the case with many traditional visual stimuli such as Gabors, but we could also manipulate the shape (‘spikiness’) of our stimulus as an effective measure to probe shape processing across the visual cortex. We demonstrated significant decreases in accuracy (indexed by loglinear d’) and increases in response time, when

participants made responses whilst their attention was distributed across multiple stimulus features, in comparison to conditions in which they were attending to changes in only a single visual feature. This simple, clear result validates the use of the relatively novel RF pattern stimulus as an effective stimulus to probe attentional mechanisms in future work.

Secondly, we investigated patterns of attentional modulation across the visual cortex when participants shifted attention between orientation, contrast, shape and passive viewing conditions in the presence of a constant RF pattern stimulus, in an fMRI scanner. At the univariate level of analysis, we identified relatively little evidence of differential patterns of attentional modulation in any visual ROI examined. We also identified no differences in BOLD signal modulation driven by changes in the stimulus features, suggesting bottom-up stimulus changes were not driving any attentional effects we identified.

However, at the multivariate level, we were able to demonstrate that attention directed towards different stimulus features produced robust, distinguishable patterns of attentional modulation across voxels in a particular visual ROI. These differential patterns of attentional modulation differed as a function of the ROI examined, suggesting attentional mechanisms do not exert the same patterns of modulations in all areas across the visual hierarchy. Importantly, we demonstrated that these successful classifications were not driven by group-level coarse scale topographical maps of the organisation of particular visual features, and we

instead suggest attentional modulation alters the fine-scale voxel-level patterns of activation in an ROI. Finally, we analysed patterns of connectivity across multiple visual ROIs simultaneously and demonstrate a significantly greater positive connectivity between ROIs during passive viewing in comparison to directed attention conditions, which we suggest is reflective of a default-mode type network activation, and regions selectively disengage from this synchronised network during attentionally demanding tasks. This identification of differential patterns of connectivity across different attentional tasks is one of the first demonstrations of this analysis in the visual literature and highlights the importance of probing multiple spatial scales of activation in the visual cortex when examining attentional effects.

In our third experiment, we used fMRI and psychophysical methods to again measure differential patterns of attentional modulation, but this time, we probed responses during attention directed towards different stimulus colours. We performed two different analyses, the first, examining differential patterns of attentional modulation across features (irrespective of the stimulus chromaticity). Here, we replicated the main findings of our previous experiment, finding relatively little difference in attentional modulation across features at the univariate level, but robust differences in patterns of BOLD signal modulation at both the multivariate and connectivity levels, which were not driven by any coarse-scale spatial maps in ROIs.

We then repeated our analysis pipeline, examining patterns of activation when attention was directed towards the chromaticity of the stimulus (in three

segregated dimensions; L+M, L-M and S-(L+M)). At the univariate level, we identified a significant difference in BOLD signal modulation between red-green and blue-yellow attentional conditions in almost all visual ROIs examined, in fitting with previous literature suggesting a relatively weaker blue-yellow signal across the visual cortex (e.g. Liu & Wandell, 2005; Mullen, Chang, & Hess, 2015; Wang & Wade, 2011). At the multivariate level however, we identify significant differences in voxel-level patterns of attentional modulation across all three stimulus chromaticities in almost all visual ROIs. Again, these differences do not appear to be driven by underlying coarse-scale organisations of chromatic information in the visual cortex. Finally, we clearly replicate our previous findings demonstrating significantly greater connectivity between visual ROIs during passive viewing versus directed attention in all three of our stimulus chromaticity conditions. This demonstrates the robust nature of our 'visual default-mode' effect and provides clear evidence of differential effects of attentional modulation in response to the attended chromaticity of the stimulus.

Finally, in our fourth experiment, we used a relatively naturalistic, dynamic and uncontrolled visual stimulus to assess the extent our findings with low-level controlled stimuli generalised to a more real-world visual environment. Our fMRI analyses demonstrated that the vast majority of differences in feature-specific attentional modulation we identified at the univariate level were driven by reduced BOLD signal modulation during attention to faces in low-level visual areas. We also identified relatively little evidence of differential attentional modulation at the univariate level when we examined

attentional modulation to relatively fine-grain dimensions with a stimulus category (e.g. vertical, horizontal and diagonal orientations).

In contrast with our previous two fMRI experiments, at the multivariate level, any successful classification of voxel-level patterns of attentional modulation were predominantly driven by comparison with attention to faces. When excluding activation associated with attention to faces, we identified almost no multivariate differences in patterns of attentional modulation across voxels in any visual area for both the overarching stimulus category (e.g. orientation versus shape) or fine-grain (e.g. vertical versus horizontal comparisons). We believe this difference in reflective of the unconstrained nature of this naturalistic experiment and the dynamic and uncontrolled nature of the stimulus.

Finally, we again identified clear evidence of significantly greater positive correlation between ROIs during passive viewing that directed attention towards low-level stimulus features. We also identified a similar pattern of positive connectivity between ROIs during attention to faces, which we believe is a reflection of a task design; it is likely during passive viewing participants' attention was drawn to the inherently salient face stimuli in the movie. We also identify significant differences in patterns of connectivity between the fine-grain stimulus dimensions, and these patterns differed to some extent as a function of the ROI partial correlations examined. Hence, even when using an uncontrolled, dynamic, rich visual stimulus, we were still

able to detect differential patterns of attentional modulation when examining patterns of connectivity at the fine-grain within-feature level.

7.2 Future work

Our investigation into signatures of featural attention across the visual cortex suggests many directions for future research. For example, our connectivity analysis identifies patterns of communication between ROIs that differ as a function of the attentional task. We have also revealed that particular ROIs decorrelate or disengage from our 'visual default mode' passive, highly correlated network when attention is directed towards a specific visual feature. Given the fact that previous research has demonstrated positive effects associated with an ability to efficiently shift between different patterns of connectivity (i.e. the shift from passive, default-mode to task-focused networks) (e.g. Dajani & Uddin, 2015; Li et al., 2017; Reineberg, Gustavson, Benca, Banich, & Friedman, 2018; Vatansever, Manktelow, Sahakian, Menon, & Stamatakis, 2016), future work should seek to investigate whether the changing patterns of connectivity across ROIs as a function of attentional task demonstrated here can predict individual participants' abilities on particular visual tasks. For example, perhaps participants who demonstrate a bigger shift in patterns of connectivity between passive viewing and attention to orientation, may have better performance on an orientation discrimination task for example.

Recent research by Semedo, Zandvakili, Machens, Yu, & Kohn, (2019) has identified a low-dimensional communication subspace, which characterises

the activity patterns relayed between early visual areas (V1 and V2). Measuring firing rate distributions (at a single neuron level) in both V1 and V2, these researchers were able to demonstrate fluctuations in V2 activity were related to a small subset of V1 neuronal activity patterns, and these patterns were distinct from the largest fluctuations shared among neurons in V1. This identification of important dimensions characterizing the types of communication between visual areas has interesting avenues for the study of attentional mechanisms.

Previous research has identified changing patterns of communication between cortical areas as a function of attention. For example, Gregoriou, Gotts, Zhou, & Desimone, (2009), demonstrated attention toward a stimulus in the joint receptive fields of the frontal eye fields (FEF) and area V4 increased the oscillatory coupling between the two visual areas in monkeys. Additionally, Oemisch, Westendorff, Everling, & Womelsdorf, (2015) identified correlations in firing rate across a sustained time window after covert attention to one of two peripheral visual stimuli in a subset of cell pairs from the anterior cingulate cortex (ACC) to dorsal prefrontal cortex (PFC), and these changing connectivity patterns between neurons carried information about the direction of attentional shift.

However, no research has attempted to use this technique of dimensionality reduction to investigate activity between visual areas as a function of attentional task. For example, the technique of reduced rank regression described by (Semedo et al., 2019) could be used to reduce activity in areas

V1 and V3A/B for example to a small set of dimensions which represent a characteristic way in which the activities of the two visual ROIs covary, and this communication subspace could be examined for changes when participants' switch attentional focus. For example, we may see different communication subspaces between visual areas as a function of feature-specific attentional task, or communication subspaces may be altered in a more general fashion by attention, with contrast to communication subspaces characterising passive viewing. Our findings in this thesis have demonstrated changing patterns of connectivity across visual ROIs as a function of attentional task. Future research should seek to extend these current findings, perhaps employing this new methodology of Semedo et al., (2019) to characterise the relationship between early visual areas in terms of the number of explanatory variables and examine how this communication varies as a function of attentional task in an extension of our current functional connectivity findings.

This relationship between the extent of differences in connectivity and task performance opens further avenues for future research. For example, attention deficit hyperactivity disorder (ADHD) is thought to affect 3.62% of boys and 0.85% of girls of school age (Ford, Goodman, & Meltzer, 2003). It is a disorder associated with inattentiveness, hyperactivity and impulsiveness. Importantly, individuals with ADHD have been demonstrated to have a reduced ability effective switch between tasks (e.g. Cepeda, Cepeda, & Kramer, 2000). Cortical signatures of this reduced task switching ability have been typically investigated in terms of executive control network

across relatively large expanses of the cortex. However, the paradigms we establish here, assessing differential patterns of univariate, multivariate and connectivity across ROIs as a function of attentional focus, may be effective in identifying differential responses in ADHD sufferers. For example, individuals with ADHD may demonstrate a reduced distinction between patterns of connectivity between passive viewing and attentional task, as we have identified in our three experiments here using participants who do not suffer from ADHD. This hypothesis is supported by recent work at a coarser scale, which demonstrated increased variability in default mode network activation and associated lower task performance in ADHD patients in comparison with controls (Mowinckel et al., 2017).

Hence, the relatively novel paradigm we have established here may be a simple and effective method for examining patterns of connectivity in patients with ADHD. The validation of our use of naturalistic stimuli in identifying these differences in patterns of connectivity suggests a potential use as an ADHD screening procedure for young children, who are likely to better engage with this dynamic and interesting stimulus. Identifying differences in connectivity across ROIs in attentional tasks, may help to reveal a potential locus for the control of task switching. For example, the IPS, an area known to be involved in the top-down control of visual attention (e.g. Bressler, Tang, Sylvester, Shulman, & Corbetta, 2008; Di Russo, Martínez, & Hillyard, 2003; Lauritzen, D'Esposito, Heeger, & Silver, 2009), may demonstrate different patterns of activity in ADHD patients comparison with healthy controls.

7.3 Conclusion

This thesis has investigated feature-specific patterns of attentional modulation across the visual cortex, using psychophysical and fMRI methods. We have established a novel paradigm, probing attentional mechanisms with a relatively constant visual stimulus, and have demonstrated our findings are not driven by bottom-up changes in our visual stimulus. We demonstrate clear evidence of differential voxel-level patterns of attentional modulation with our highly controlled, low-level RF patterns, but not with the use of a relatively naturalistic, rich, dynamic movie stimulus. In all experiments, we identify significantly different patterns of connectivity across ROIs as a function of attentional focus and demonstrate a robust visual 'default mode' network with relatively greater positive correlation across ROIs during an undemanding task. We demonstrate a generalisation of findings across low-level and naturalistic testing paradigms. We suggest this novel paradigm can be used to predict participants' performance on feature-specific tasks, on the basis of differences in patterns of connectivity and suggest a potential use for this paradigm in probing mechanisms of ADHD in the visual cortex.

References

Al-Aidroos, N., Said, C. P., & Turk-Browne, N. B. (2012). Top-down attention switches coupling between low-level and high-level areas of human visual cortex. *Proceedings of the National Academy of Sciences of the United States of America*, *109*(36), 14675–14680. https://doi.org/10.1073/pnas.1202095109

Alink, A., Walther, A., Krugliak, A., & Kriegeskorte, N. (2017). Local opposite orientation preferences in V1: FMRI sensitivity to fine-grained pattern information. *Scientific Reports*, *7*(1), 7128. https://doi.org/10.1038/s41598-017-07036-8

Allison, null, Puce, null, & McCarthy, null. (2000). Social perception from visual cues: Role of the STS region. *Trends in Cognitive Sciences*, *4*(7), 267–278.

Andersen, S. K., Hillyard, S. A., & Müller, M. M. (2013). Global Facilitation of Attended Features Is Obligatory and Restricts Divided Attention. *The Journal of Neuroscience*, *33*(46), 18200–18207. https://doi.org/10.1523/JNEUROSCI.1913-13.2013

Anstis, S., & Cavanagh, P. (1983). *A minimum motion technique for judging equiluminance*. York University.

Anticevic, A., Cole, M. W., Murray, J. D., Corlett, P. R., Wang, X.-J., & Krystal, J. H. (2012). The Role of Default Network Deactivation in Cognition and Disease. *Trends in Cognitive Sciences*, *16*(12), 584–592. https://doi.org/10.1016/j.tics.2012.10.008

Awh, E., Belopolsky, A. V., & Theeuwes, J. (2012). Top-down versus bottom-up attentional control: A failed theoretical dichotomy. *Trends in*

Cognitive Sciences, *16*(8), 437–443. https://doi.org/10.1016/j.tics.2012.06.010

Bailes, H. J., & Lucas, R. J. (2013). Human melanopsin forms a pigment maximally sensitive to blue light (λmax ≈ 479 nm) supporting activation of G(q/11) and G(i/o) signalling cascades. *Proceedings. Biological Sciences*, *280*(1759), 20122987. https://doi.org/10.1098/rspb.2012.2987

Bartels, A, Zeki, S., & Logothetis, N. K. (2008). Natural vision reveals regional specialization to local motion and to contrast-invariant, global flow in the human brain. *Cerebral Cortex (New York, N.Y.: 1991)*, *18*(3), 705–717. https://doi.org/10.1093/cercor/bhm107

Bartels, Andreas, & Zeki, S. (2004a). Functional brain mapping during free viewing of natural scenes. *Human Brain Mapping*, *21*(2), 75–85. https://doi.org/10.1002/hbm.10153

Bartels, Andreas, & Zeki, S. (2004b). The chronoarchitecture of the human brain—Natural viewing conditions reveal a time-based anatomy of the brain. *NeuroImage*, *22*(1), 419–433. https://doi.org/10.1016/j.neuroimage.2004.01.007

Bartels, Andreas, & Zeki, S. (2005). Brain dynamics during natural viewing conditions—A new guide for mapping connectivity in vivo. *NeuroImage*, *24*(2), 339–349. https://doi.org/10.1016/j.neuroimage.2004.08.044

Beauchamp, M. S., Cox, R. W., & DeYoe, E. A. (1997). Graded effects of spatial and featural attention on human area MT and associated

motion processing areas. *Journal of Neurophysiology*, *78*(1), 516–520. https://doi.org/10.1152/jn.1997.78.1.516

Beckett, A., Peirce, J. W., Sanchez-Panchuelo, R.-M., Francis, S., & Schluppeck, D. (2012). Contribution of large scale biases in decoding of direction-of-motion from high-resolution fMRI data in human early visual cortex. *NeuroImage*, *63*(3), 1623–1632. https://doi.org/10.1016/j.neuroimage.2012.07.066

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). https://doi.org/10.1167/14.11.12

Benjamini, Y., & Hochberg, Y. (1995). Controlling the false discovery rate: A practical and powerful approach to multiple testing. *Journal of the Royal Statistical Society: Series B (Methodological)*, *57*(1), 289–300.

Born, R. T., & Bradley, D. C. (2005). Structure and function of visual area MT. *Annual Review of Neuroscience*, *28*, 157–189. https://doi.org/10.1146/annurev.neuro.26.041002.131052

Bowmaker, J. K., & Dartnall, H. J. (1980). Visual pigments of rods and cones in a human retina. *The Journal of Physiology*, *298*, 501–511. https://doi.org/10.1113/jphysiol.1980.sp013097

Brainard, D. H. (1997). The psychophysics toolbox. *Spatial Vision*, *10*, 433–436.

Bressler, S. L., Tang, W., Sylvester, C. M., Shulman, G. L., & Corbetta, M. (2008). Top-down control of human visual cortex by frontal and parietal cortex in anticipatory visual spatial attention. *The Journal of*

Neuroscience: The Official Journal of the Society for Neuroscience, *28*(40), 10056–10061. https://doi.org/10.1523/JNEUROSCI.1776-08.2008

Broadbent, D. E. (1958). *Perception and Communication*. London: Plenum Press.

Brouwer, G. J., & Heeger, D. J. (2009). Decoding and reconstructing color from responses in human visual cortex. *The Journal of Neuroscience: The Official Journal of the Society for Neuroscience*, *29*(44), 13992–14003. https://doi.org/10.1523/JNEUROSCI.3577-09.2009

Buck, C., & Lee, J. (2013). *Frozen*. United States: Walt Disney Animation Studios.

Buckner, R. L., Andrews-Hanna, J. R., & Schacter, D. L. (2008). The brain's default network: Anatomy, function, and relevance to disease. *Annals of the New York Academy of Sciences*, *1124*, 1–38. https://doi.org/10.1196/annals.1440.011

Buffalo, E. A., Fries, P., Landman, R., Liang, H., & Desimone, R. (2010). A backward progression of attentional effects in the ventral stream. *Proceedings of the National Academy of Sciences of the United States of America*, *107*(1), 361–365. https://doi.org/10.1073/pnas.0907658106

Carandini, M., Demb, J. B., Mante, V., Tolhurst, D. J., Dan, Y., Olshausen, B. A., … Rust, N. C. (2005). Do we know what the early visual system does? *The Journal of Neuroscience: The Official Journal of the Society for Neuroscience*, *25*(46), 10577–10597. https://doi.org/10.1523/JNEUROSCI.3726-05.2005

Carandini, M., & Heeger, D. J. (2011). Normalization as a canonical neural computation. *Nature Reviews. Neuroscience*, *13*(1), 51–62. https://doi.org/10.1038/nrn3136

Carpenter, R. H. S. (1988). *Movements of the eyes*. Pion.

Carrasco, M. (2011). Visual attention: The past 25 years. *Vision Research*, *51*(13), 1484–1525. https://doi.org/10.1016/j.visres.2011.04.012

Casagrande, V. A. (1994). A third parallel visual pathway to primate area V1. *Trends in Neurosciences*, *17*(7), 305–310.

Cepeda, N. J., Cepeda, M. L., & Kramer, A. F. (2000). Task switching and attention deficit hyperactivity disorder. *Journal of Abnormal Child Psychology*, *28*(3), 213–226.

Chang, C.-C., & Lin, C.-J. (2011). " LIBSVM: A library for support vector machines," ACM Transactions on Intelligent Systems and Technology, 2: 27: 1–27: 27, 2011. *Http://Www. Csie. Ntu. Edu. Tw/~ Cjlin/Libsvm*, *2*.

Chatterjee, S., & Callaway, E. M. (2003). Parallel colour-opponent pathways to primary visual cortex. *Nature*, *426*(6967), 668–671. https://doi.org/10.1038/nature02167

Chawla, D., Rees, G., & Friston, K. J. (1999). The physiological basis of attentional modulation in extrastriate visual areas. *Nature Neuroscience*, *2*(7), 671.

Chiu, Y.-C., Esterman, M., Han, Y., Rosen, H., & Yantis, S. (2011). Decoding task-based attentional modulation during face categorization. *Journal of Cognitive Neuroscience*, *23*(5), 1198–1204. https://doi.org/10.1162/jocn.2010.21503

Christoff, K., Gordon, A. M., Smallwood, J., Smith, R., & Schooler, J. W. (2009). Experience sampling during fMRI reveals default network and executive system contributions to mind wandering. *Proceedings of the National Academy of Sciences*, *106*(21), 8719–8724.

Clifford, C. W. G., Mannion, D. J., & McDonald, J. S. (2009). Radial biases in the processing of motion and motion-defined contours by human visual cortex. *Journal of Neurophysiology*, *102*(5), 2974–2981. https://doi.org/10.1152/jn.00411.2009

Cohen, M. R., & Maunsell, J. H. R. (2009). Attention improves performance primarily by reducing interneuronal correlations. *Nature Neuroscience*, *12*(12), 1594–1600. https://doi.org/10.1038/nn.2439

Conway, B. R. (2001). Spatial structure of cone inputs to color cells in alert macaque primary visual cortex (V-1). *The Journal of Neuroscience: The Official Journal of the Society for Neuroscience*, *21*(8), 2768–2783.

Conway, Bevil R. (2009). Color vision, cones, and color-coding in the cortex. *The Neuroscientist: A Review Journal Bringing Neurobiology, Neurology and Psychiatry*, *15*(3), 274–290. https://doi.org/10.1177/1073858408331369

Conway, Bevil R., Moeller, S., & Tsao, D. Y. (2007). Specialized color modules in macaque extrastriate cortex. *Neuron*, *56*(3), 560–573. https://doi.org/10.1016/j.neuron.2007.10.008

Corbetta, M., Miezin, F. M., Dobmeyer, S., Shulman, G. L., & Petersen, S. E. (1991). Selective and divided attention during visual discriminations of shape, color, and speed: Functional anatomy by positron emission

tomography. *The Journal of Neuroscience: The Official Journal of the Society for Neuroscience*, *11*(8), 2383–2402.

Corbetta, Maurizio, Miezin, F. M., Dobmeyer, S., Shulman, G. L., & Petersen, S. E. (1990). Attentional modulation of neural processing of shape, color, and velocity in humans. *Science*, *248*(4962), 1556–1559.

Corbetta, Maurizio, Patel, G., & Shulman, G. L. (2008). The reorienting system of the human brain: From environment to theory of mind. *Neuron*, *58*(3), 306–324. https://doi.org/10.1016/j.neuron.2008.04.017

Cortes, C., & Vapnik, V. (1995). Support-vector networks. *Machine Learning*, *20*(3), 273–297.

Çukur, T., Nishimoto, S., Huth, A. G., & Gallant, J. L. (2013). Attention during natural vision warps semantic representation across the human brain. *Nature Neuroscience*, *16*(6), 763–770. https://doi.org/10.1038/nn.3381

Dacey, D. M. (1996). Circuitry for color coding in the primate retina. *Proceedings of the National Academy of Sciences of the United States of America*, *93*(2), 582–588. https://doi.org/10.1073/pnas.93.2.582

Dacey, D., Packer, O. S., Diller, L., Brainard, D., Peterson, B., & Lee, B. (2000). Center surround receptive field structure of cone bipolar cells in primate retina. *Vision Research*, *40*(14), 1801–1811. https://doi.org/10.1016/S0042-6989(00)00039-0

Dajani, D. R., & Uddin, L. Q. (2015). Demystifying cognitive flexibility: Implications for clinical and developmental neuroscience. *Trends in Neurosciences*, *38*(9), 571–578. https://doi.org/10.1016/j.tins.2015.07.003

Denys, K., Vanduffel, W., Fize, D., Nelissen, K., Peuskens, H., Essen, D. V., & Orban, G. A. (2004). The Processing of Visual Shape in the Cerebral Cortex of Human and Nonhuman Primates: A Functional Magnetic Resonance Imaging Study. *Journal of Neuroscience*, *24*(10), 2551–2565. https://doi.org/10.1523/JNEUROSCI.3569-03.2004

Di Russo, F., Martínez, A., & Hillyard, S. A. (2003). Source analysis of event-related cortical activity during visuo-spatial attention. *Cerebral Cortex (New York, N.Y.: 1991)*, *13*(5), 486–499.

Diamond, R., & Carey, S. (1986). Why faces are and are not special: An effect of expertise. *Journal of Experimental Psychology. General*, *115*(2), 107–117.

Dumoulin, S., & Hess, R. (2007). Cortical specialization for concentric shape processing. *Vision Research*, *47*(12), 1608–1613.

Dumoulin, S. O., Bittar, R. G., Kabani, N. J., Baker, C. L., Le Goualher, G., Bruce Pike, G., & Evans, A. C. (2000). A new anatomical landmark for reliable identification of human area V5/MT: A quantitative analysis of sulcal patterning. *Cerebral Cortex (New York, N.Y.: 1991)*, *10*(5), 454–463.

Dumoulin, Serge O., & Wandell, B. A. (2008). Population receptive field estimates in human visual cortex. *Neuroimage*, *39*(2), 647–660.

Duncan, J., & Humphreys, G. W. (1989). Visual search and stimulus similarity. *Psychological Review*, *96*(3), 433–458.

Economides, J. R., Sincich, L. C., Adams, D. L., & Horton, J. C. (2011). Orientation tuning of cytochrome oxidase patches in macaque primary visual cortex. *Nature Neuroscience*, *14*(12), 1574.

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 (New York, N.Y.: 1991)*, *7*(2), 181–192. https://doi.org/10.1093/cercor/7.2.181

Engel, S. A., Rumelhart, D. E., Wandell, B. A., Lee, A. T., Glover, G. H., Chichilnisky, E. J., & Shadlen, M. N. (1994). FMRI of human visual cortex. *Nature*, *369*(6481), 525. https://doi.org/10.1038/369525a0

Engel, Stephen A. (2005). Adaptation of oriented and unoriented color-selective neurons in human visual areas. *Neuron*, *45*(4), 613–623. https://doi.org/10.1016/j.neuron.2005.01.014

Fischer, E., Bülthoff, H. H., Logothetis, N. K., & Bartels, A. (2011). Visual motion responses in the posterior cingulate sulcus: A comparison to V5/MT and MST. *Cerebral Cortex*, *22*(4), 865–876.

Fischl, B. (2012). FreeSurfer. *Neuroimage*, *62*(2), 774–781.

Ford, T., Goodman, R., & Meltzer, H. (2003). The British Child and Adolescent Mental Health Survey 1999: The prevalence of DSM-IV disorders. *Journal of the American Academy of Child and Adolescent Psychiatry*, *42*(10), 1203–1211. https://doi.org/10.1097/00004583-200310000-00011

Foster, D. H. (2010). Chromatic Function of the Cone. In D. A. Dartt (Ed.), *Encyclopedia of the Eye* (pp. 266–274). https://doi.org/10.1016/B978-0-12-374203-2.00232-3

Fox, M. D., Snyder, A. Z., Vincent, J. L., Corbetta, M., Van Essen, D. C., & Raichle, M. E. (2005). The human brain is intrinsically organized into dynamic, anticorrelated functional networks. *Proceedings of the*

National Academy of Sciences of the United States of America, *102*(27), 9673–9678. https://doi.org/10.1073/pnas.0504136102

Freeman, J., Brouwer, G. J., Heeger, D. J., & Merriam, E. P. (2011). Orientation decoding depends on maps, not columns. *Journal of Neuroscience*, *31*(13), 4792–4804.

Friston, K. J., Fletcher, P., Josephs, O., Holmes, A., Rugg, M. D., & Turner, R. (1998). Event-related fMRI: Characterizing differential responses. *NeuroImage*, *7*(1), 30–40. https://doi.org/10.1006/nimg.1997.0306

Gallant, J. L., Braun, J., & Van Essen, D. C. (1993). Selectivity for polar, hyperbolic, and Cartesian gratings in macaque visual cortex. *Science (New York, N.Y.)*, *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.

Gao, W., Gilmore, J. H., Alcauter, S., & Lin, W. (2013). The dynamic reorganization of the default-mode network during a visual classification task. *Frontiers in Systems Neuroscience*, *7*, 34. https://doi.org/10.3389/fnsys.2013.00034

Gardner, J. L., Sun, P., Waggoner, R. A., Ueno, K., Tanaka, K., & Cheng, K. (2005). Contrast adaptation and representation in human early visual cortex. *Neuron*, *47*(4), 607–620. https://doi.org/10.1016/j.neuron.2005.07.016

Gauthier, I., Behrmann, M., & Tarr, M. J. (1999). Can face recognition really be dissociated from object recognition? *Journal of Cognitive Neuroscience*, *11*(4), 349–370.

Gauthier, I., Tarr, M. J., Moylan, J., Skudlarski, 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.

Gholami, R., & Fakhari, N. (2017). Chapter 27 - Support Vector Machine: Principles, Parameters, and Applications. In P. Samui, S. Sekhar, & V. E. Balas (Eds.), *Handbook of Neural Computation* (pp. 515–535). https://doi.org/10.1016/B978-0-12-811318-9.00027-2

Ghose, G. M., & Ts’ O, D. Y. (1997). Form processing modules in primate area V4. *Journal of Neurophysiology*, *77*(4), 2191–2196.

Girard, P., & Bullier, J. (1989). Visual activity in area V2 during reversible inactivation of area 17 in the macaque monkey. *Journal of Neurophysiology*, *62*(6), 1287–1302. https://doi.org/10.1152/jn.1989.62.6.1287

Girard, P., Salin, P. A., & Bullier, J. (1991). Visual activity in areas V3a and V3 during reversible inactivation of area V1 in the macaque monkey. *Journal of Neurophysiology*, *66*(5), 1493–1503. https://doi.org/10.1152/jn.1991.66.5.1493

Goffaux, V., Schiltz, C., Mur, M., & Goebel, R. (2012). Local discriminability determines the strength of holistic processing for faces in the fusiform face area. *Frontiers in Psychology*, *3*, 604. https://doi.org/10.3389/fpsyg.2012.00604

Gregoriou, G. G., Gotts, S. J., Zhou, H., & Desimone, R. (2009). High-frequency, long-range coupling between prefrontal and visual cortex

during attention. *Science (New York, N.Y.)*, *324*(5931), 1207–1210. https://doi.org/10.1126/science.1171402

Greicius, M. D., Krasnow, B., Reiss, A. L., & Menon, V. (2003). Functional connectivity in the resting brain: A network analysis of the default mode hypothesis. *Proceedings of the National Academy of Sciences*, *100*(1), 253–258. https://doi.org/10.1073/pnas.0135058100

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. https://doi.org/10.1016/S0896-6273(00)80832-6

Grill-Spector, K., & Malach, R. (2004). The human visual cortex. *Annual Review of Neuroscience*, *27*, 649–677. https://doi.org/10.1146/annurev.neuro.27.070203.144220

Gusnard, D. A., & Raichle, M. E. (2001). Searching for a baseline: Functional imaging and the resting human brain. *Nature Reviews Neuroscience*, *2*(10), 685.

Hahn, B., Wolkenberg, F. A., Ross, T. J., Myers, C. S., Heishman, S. J., Stein, D. J., … Stein, E. A. (2008). Divided versus selective attention: Evidence for common processing mechanisms. *Brain Research*, *1215*, 137–146. https://doi.org/10.1016/j.brainres.2008.03.058

Harrison, S. A., & Tong, F. (2009). Decoding reveals the contents of visual working memory in early visual areas. *Nature*, *458*(7238), 632–635. https://doi.org/10.1038/nature07832

Hasson, U., Nir, Y., Levy, I., Fuhrmann, G., & Malach, R. (2004). Intersubject synchronization of cortical activity during natural vision. *Science (New*

York, N.Y.), *303*(5664), 1634–1640. https://doi.org/10.1126/science.1089506

Hautus, M. J. (1995). Corrections for extreme proportions and their biasing effects on estimated values ofd′. *Behavior Research Methods, Instruments, & Computers*, *27*(1), 46–51.

Haxby, J. V., Connolly, A. C., & Guntupalli, J. S. (2014). Decoding neural representational spaces using multivariate pattern analysis. *Annual Review of Neuroscience*, *37*, 435–456. https://doi.org/10.1146/annurev-neuro-062012-170325

Hendry, S. H., & Reid, R. C. (2000). The koniocellular pathway in primate vision. *Annual Review of Neuroscience*, *23*, 127–153. https://doi.org/10.1146/annurev.neuro.23.1.127

Herath, P., Klingberg, T., Young, J., Amunts, K., & Roland, P. (2001). Neural correlates of dual task interference can be dissociated from those of divided attention: An fMRI study. *Cerebral Cortex (New York, N.Y.: 1991)*, *11*(9), 796–805.

Herrmann, K., Montaser-Kouhsari, L., Carrasco, M., & Heeger, D. J. (2010). When size matters: Attention affects performance by contrast or response gain. *Nature Neuroscience*, *13*(12), 1554–1559. https://doi.org/10.1038/nn.2669

Heywood, C. A., & Cowey, A. (1992). The role of the “face-cell” area in the discrimination and recognition of faces by monkeys. *Philosophical Transactions of the Royal Society of London. Series B, Biological Sciences*, *335*(1273), 31–37; discussion 37-38. https://doi.org/10.1098/rstb.1992.0004

Heywood, C. A., Gadotti, A., & Cowey, A. (1992). Cortical area V4 and its role in the perception of color. *The Journal of Neuroscience: The Official Journal of the Society for Neuroscience*, *12*(10), 4056–4065.

Highsmith, J., & Crognale, M. A. (2010). Attentional shifts have little effect on the waveform of the chromatic onset VEP. *Ophthalmic & Physiological Optics: The Journal of the British College of Ophthalmic Opticians (Optometrists)*, *30*(5), 525–533. https://doi.org/10.1111/j.1475-1313.2010.00747.x

Holmes, G. (1918). DISTURBANCES OF VISION BY CEREBRAL LESIONS. *The British Journal of Ophthalmology*, *2*(7), 353–384.

Horton, J. C., & Hubel, D. H. (1981). Regular patchy distribution of cytochrome oxidase staining in primary visual cortex of macaque monkey. *Nature*, *292*(5825), 762.

Hubel, D. H., & Wiesel, T. N. (1959). Receptive fields of single neurones in the cat's striate cortex. *The Journal of Physiology*, *148*(3), 574–591.

Hubel, D. H., & Wiesel, T. N. (1962). Receptive fields, binocular interaction and functional architecture in the cat's visual cortex. *The Journal of Physiology*, *160*(1), 106-154.2.

Huk, A. C., & Heeger, D. J. (2000). Task-related modulation of visual cortex. *Journal of Neurophysiology*, *83*(6), 3525–3536. https://doi.org/10.1152/jn.2000.83.6.3525

Huk, A. C., Ress, D., & Heeger, D. J. (2001). Neuronal basis of the motion aftereffect reconsidered. *Neuron*, *32*(1), 161–172.

Huk, Alexander C., Dougherty, R. F., & Heeger, D. J. (2002). Retinotopy and functional subdivision of human areas MT and MST. *Journal of Neuroscience*, *22*(16), 7195–7205.

Ivanov, I. V., & Mullen, K. T. (2012a). The role of local features in shape discrimination of contour- and surface-defined radial frequency patterns at low contrast. *Vision Research*, *52*(1), 1–10. https://doi.org/10.1016/j.visres.2011.10.002

Ivanov, I. V., & Mullen, K. T. (2012b). The role of local features in shape discrimination of contour- and surface-defined radial frequency patterns at low contrast. *Vision Research*, *52*(1), 1–10. https://doi.org/10.1016/j.visres.2011.10.002

James, W. (1890). *The principles of psychology* (Vol. 1). Macmillan London.

Jenkinson, M., Beckmann, C. F., Behrens, T. E. J., Woolrich, M. W., & Smith, S. M. (2012). FSL. *NeuroImage*, *62*(2), 782–790. https://doi.org/10.1016/j.neuroimage.2011.09.015

Johnson, E. N., Hawken, M. J., & Shapley, R. (2001). The spatial transformation of color in the primary visual cortex of the macaque monkey. *Nature Neuroscience*, *4*(4), 409–416. https://doi.org/10.1038/86061

Johnson, J. A., & Zatorre, R. J. (2006). Neural substrates for dividing and focusing attention between simultaneous auditory and visual events. *NeuroImage*, *31*(4), 1673–1681. https://doi.org/10.1016/j.neuroimage.2006.02.026

Kamitani, Y., & Sawahata, Y. (2010). Spatial smoothing hurts localization but not information: Pitfalls for brain mappers. *NeuroImage*, *49*(3), 1949–1952. https://doi.org/10.1016/j.neuroimage.2009.06.040

Kamitani, Y., & Tong, F. (2005). Decoding the visual and subjective contents of the human brain. *Nature Neuroscience*, *8*(5), 679–685. https://doi.org/10.1038/nn1444

Kamitani, Y., & Tong, F. (2006). Decoding seen and attended motion directions from activity in the human visual cortex. *Current Biology: CB*, *16*(11), 1096–1102. https://doi.org/10.1016/j.cub.2006.04.003

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: The Official Journal of the Society for Neuroscience*, *17*(11), 4302–4311.

Kanwisher, Nancy. (2010). Functional specificity in the human brain: A window into the functional architecture of the mind. *Proceedings of the National Academy of Sciences*, *107*(25), 11163–11170. https://doi.org/10.1073/pnas.1005062107

Kanwisher, Nancy, & Yovel, G. (2006). The fusiform face area: A cortical region specialized for the perception of faces. *Philosophical Transactions of the Royal Society of London. Series B, Biological Sciences*, *361*(1476), 2109–2128. https://doi.org/10.1098/rstb.2006.1934

Kastner, S., De Weerd, P., Desimone, R., & Ungerleider, L. G. (1998). Mechanisms of directed attention in the human extrastriate cortex as

revealed by functional MRI. *Science (New York, N.Y.)*, *282*(5386), 108–111.

Kay, K. N., Naselaris, T., Prenger, R. J., & Gallant, J. L. (2008). Identifying natural images from human brain activity. *Nature*, *452*(7185), 352–355. https://doi.org/10.1038/nature06713

Knerr, S., Personnaz, L., & Dreyfus, G. (1990). Single-layer learning revisited: A stepwise procedure for building and training a neural network. In *Neurocomputing* (pp. 41–50). Springer.

Kornak, J., Hall, D. A., & Haggard, M. P. (2011). Spatially Extended fMRI Signal Response to Stimulus in Non-Functionally Relevant Regions of the Human Brain: Preliminary Results. *The Open Neuroimaging Journal*, *5*(1). https://doi.org/10.2174/1874440001105010024

Kourtzi, Z., & Kanwisher, N. (2001). Representation of perceived object shape by the human lateral occipital complex. *Science (New York, N.Y.)*, *293*(5534), 1506–1509. https://doi.org/10.1126/science.1061133

Kowler, E. (1990). The role of visual and cognitive processes in the control of eye movement. *Reviews of Oculomotor Research*, *4*, 1–70.

Lacadie, C., Fulbright, R., Arora, J., Constable, R., & Papademetris, X. (2008). Brodmann Areas defined in MNI space using a new Tracing Tool in BioImage Suite. *Proceedings of the 14th Annual Meeting of the Organization for Human Brain Mapping*, *771*.

Larsson, J., & Heeger, D. J. (2006). Two retinotopic visual areas in human lateral occipital cortex. *Journal of Neuroscience*, *26*(51), 13128–13142.

Lauritzen, T Z, Ales, J. M., & Wade, A. R. (2010). The effects of visuospatial attention measured across visual cortex using source-imaged, steady-state EEG. *Journal of Vision*, *10*(14), 1–17. https://doi.org/10.1167/10.14.39

Lauritzen, Thomas Z., D'Esposito, M., Heeger, D. J., & Silver, M. A. (2009). Top-down flow of visual spatial attention signals from parietal to occipital cortex. *Journal of Vision*, *9*(13), 18.1-14. https://doi.org/10.1167/9.13.18

Lawrence, S. J., Keefe, B. D., Vernon, R. J., 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.

Levitt, J. B., Kiper, D. C., & Movshon, J. A. (1994). Receptive fields and functional architecture of macaque V2. *Journal of Neurophysiology*, *71*(6), 2517–2542. https://doi.org/10.1152/jn.1994.71.6.2517

Li, J., Zhang, D., Liang, A., Liang, B., Wang, Z., Cai, Y., … Liu, M. (2017). High transition frequencies of dynamic functional connectivity states in the creative brain. *Scientific Reports*, *7*, 46072. https://doi.org/10.1038/srep46072

Liu, J., & Wandell, B. A. (2005). Specializations for chromatic and temporal signals in human visual cortex. *The Journal of Neuroscience: The Official Journal of the Society for Neuroscience*, *25*(13), 3459–3468. https://doi.org/10.1523/JNEUROSCI.4206-04.2005

Liu, T., Scott D. Slotnick, Serences, J. T., & Yantis, S. (2003). Cortical mechanisms of feature-based attentional control. *Cerebral Cortex*, *13*(12), 1334–1343.

Livingstone, M., & Hubel, D. (1988). Segregation of form, color, movement, and depth: Anatomy, physiology, and perception. *Science*, *240*(4853), 740–749.

Livingstone, M. S., & Hubel, D. H. (1984). Anatomy and physiology of a color system in the primate visual cortex. *The Journal of Neuroscience: The Official Journal of the Society for Neuroscience*, *4*(1), 309–356.

Loffler, G. (2008). Perception of contours and shapes: Low and intermediate stage mechanisms. *Vision Research*, *48*(20), 2106–2127. https://doi.org/10.1016/j.visres.2008.03.006

Logothetis, N. K., Guggenberger, H., Peled, S., & Pauls, J. (1999). Functional imaging of the monkey brain. *Nature Neuroscience*, *2*(6), 555–562. https://doi.org/10.1038/9210

Logothetis, Nikos K., & Wandell, B. A. (2004). Interpreting the BOLD signal. *Annual Review of Physiology*, *66*, 735–769. https://doi.org/10.1146/annurev.physiol.66.082602.092845

Loose, R., Kaufmann, C., Auer, D. P., & Lange, K. W. (2003). Human prefrontal and sensory cortical activity during divided attention tasks. *Human Brain Mapping*, *18*(4), 249–259. https://doi.org/10.1002/hbm.10082

Lu, Z.-L., Lesmes, L. A., & Sperling, G. (1999). The mechanism of isoluminant chromatic motion perception. *Proceedings of the National Academy of Sciences of the United States of America*, *96*(14), 8289–8294.

Luck, S. J., Chelazzi, L., Hillyard, S. A., & Desimone, R. (1997). Neural mechanisms of spatial selective attention in areas V1, V2, and V4 of

macaque visual cortex. *Journal of Neurophysiology*, *77*(1), 24–42. https://doi.org/10.1152/jn.1997.77.1.24

Malach, R., Reppas, J. B., Benson, R. R., Kwong, K. K., Jiang, H., Kennedy, W. A., … Tootell, R. B. (1995). Object-related activity revealed by functional magnetic resonance imaging in human occipital cortex. *Proceedings of the National Academy of Sciences of the United States of America*, *92*(18), 8135–8139.

Maloney, R. T., Watson, T. L., & Clifford, C. W. (2013). Human cortical and behavioral sensitivity to patterns of complex motion at eccentricity. *Journal of Neurophysiology*, *110*(11), 2545–2556.

Mannion, D. J., McDonald, J. S., & Clifford, C. W. G. (2009). Discrimination of the local orientation structure of spiral Glass patterns early in human visual cortex. *NeuroImage*, *46*(2), 511–515.

Martinez, A., DiRusso, F., Anllo-Vento, L., Sereno, M. I., Buxton, R. B., & Hillyard, S. A. (2001). Putting spatial attention on the map: Timing and localization of stimulus selection processes in striate and extrastriate visual areas. *Vision Research*, *41*(10–11), 1437–1457.

Martinez-Trujillo, J. C., & Treue, S. (2004). Feature-based attention increases the selectivity of population responses in primate visual cortex. *Current Biology: CB*, *14*(9), 744–751. https://doi.org/10.1016/j.cub.2004.04.028

Martínez-Trujillo, J., & Treue, S. (2002). Attentional modulation strength in cortical area MT depends on stimulus contrast. *Neuron*, *35*(2), 365–370.

Meadows, J. C. (1974). Disturbed perception of colours associated with localized cerebral lesions. *Brain: A Journal of Neurology*, *97*(4), 615–632. https://doi.org/10.1093/brain/97.1.615

Melyan, Z., Tarttelin, E. E., Bellingham, J., Lucas, R. J., & Hankins, M. W. (2005). Addition of human melanopsin renders mammalian cells photoresponsive. *Nature*, *433*(7027), 741–745. https://doi.org/10.1038/nature03344

Merigan, W. H. (1996). Basic visual capacities and shape discrimination after lesions of extrastriate area V4 in macaques. *Visual Neuroscience*, *13*(1), 51–60.

Merigan, W. H., & Pham, H. A. (1998). V4 lesions in macaques affect both single- and multiple-viewpoint shape discriminations. *Visual Neuroscience*, *15*(2), 359–367.

Moran, J., & Desimone, R. (1985). Selective attention gates visual processing in the extrastriate cortex. *Science (New York, N.Y.)*, *229*(4715), 782–784.

Mowinckel, A. M., Alnæs, D., Pedersen, M. L., Ziegler, S., Fredriksen, M., Kaufmann, T., … Biele, G. (2017). Increased default-mode variability is related to reduced task-performance and is evident in adults with ADHD. *NeuroImage. Clinical*, *16*, 369–382. https://doi.org/10.1016/j.nicl.2017.03.008

Mullen, K. T., Chang, D. H. F., & Hess, R. F. (2015). The selectivity of responses to red-green colour and achromatic contrast in the human visual cortex: An fMRI adaptation study. *The European Journal of Neuroscience*, *42*(11), 2923–2933. https://doi.org/10.1111/ejn.13090

Nebel, K., Wiese, H., Stude, P., de Greiff, A., Diener, H.-C., & Keidel, M. (2005). On the neural basis of focused and divided attention. *Brain Research. Cognitive Brain Research*, *25*(3), 760–776. https://doi.org/10.1016/j.cogbrainres.2005.09.011

Nestares, O., & Heeger, D. J. (2000). Robust multiresolution alignment of MRI brain volumes. *Magnetic Resonance in Medicine: An Official Journal of the International Society for Magnetic Resonance in Medicine*, *43*(5), 705–715.

Ninio, A., & Kahneman, D. (1974). Reaction time in focused and in divided attention. *Journal of Experimental Psychology*, *103*(3), 394–399. https://doi.org/10.1037/h0037202

Noble, W. S. (2009). How does multiple testing correction work? *Nature Biotechnology*, *27*(12), 1135–1137. https://doi.org/10.1038/nbt1209-1135

O'Craven, K. M., Downing, P. E., & Kanwisher, N. (1999). FMRI evidence for objects as the units of attentional selection. *Nature*, *401*(6753), 584. https://doi.org/10.1038/44134

O'Craven, K. M., Rosen, B. R., Kwong, K. K., Treisman, A., & Savoy, R. L. (1997). Voluntary attention modulates fMRI activity in human MT–MST. *Neuron*, *18*(4), 591–598.

Oemisch, M., Westendorff, S., Everling, S., & Womelsdorf, T. (2015). Interareal Spike-Train Correlations of Anterior Cingulate and Dorsal Prefrontal Cortex during Attention Shifts. *The Journal of Neuroscience: The Official Journal of the Society for Neuroscience*,

35(38), 13076–13089. https://doi.org/10.1523/JNEUROSCI.1262-15.2015

Ohzawa, I., Sclar, G., & Freeman, R. D. (1985). Contrast gain control in the cat's visual system. *Journal of Neurophysiology*, *54*(3), 651–667. https://doi.org/10.1152/jn.1985.54.3.651

Op de Beeck, H. P. (2010). Against hyperacuity in brain reading: Spatial smoothing does not hurt multivariate fMRI analyses? *NeuroImage*, *49*(3), 1943–1948. https://doi.org/10.1016/j.neuroimage.2009.02.047

O'Toole, A. J., Roark, D. A., & Abdi, H. (2002). Recognizing moving faces: A psychological and neural synthesis. *Trends in Cognitive Sciences*, *6*(6), 261–266.

Palmer, J., Mobley, L. A., & Teller, D. Y. (1993). Motion at isoluminance: Discrimination/detection ratios and the summation of luminance and chromatic signals. *Journal of the Optical Society of America. A, Optics and Image Science*, *10*(6), 1353–1362.

Panda, S., Nayak, S. K., Campo, B., Walker, J. R., Hogenesch, J. B., & Jegla, T. (2005). Illumination of the melanopsin signaling pathway. *Science (New York, N.Y.)*, *307*(5709), 600–604. https://doi.org/10.1126/science.1105121

Parasuraman, R. (1998). The attentive brain: Issues and prospects. In *The Attentive Brain* (pp. 3–15). Cambridge, Massachusetts: Mit Press.

Pasupathy, A., & Connor, C. E. (2001). Shape representation in area V4: Position-specific tuning for boundary conformation. *Journal of Neurophysiology*, *86*(5), 2505–2519. https://doi.org/10.1152/jn.2001.86.5.2505

Pasupathy, Anitha. (2006). Neural basis of shape representation in the primate brain. *Progress in Brain Research*, *154*, 293–313. https://doi.org/10.1016/S0079-6123(06)54016-6

Peirce, J., Gray, J. R., Simpson, S., MacAskill, M., Höchenberger, R., Sogo, H., … Lindeløv, J. K. (2019). PsychoPy2: Experiments in behavior made easy. *Behavior Research Methods*, *51*(1), 195–203. https://doi.org/10.3758/s13428-018-01193-y

Pelli, D. G. (1997). The VideoToolbox software for visual psychophysics: Transforming numbers into movies. *Spatial Vision*, *10*(4), 437–442.

Penny, W., Friston, K. J., Ashburner, J., Kiebel, S., & Nichols, T. (2006). Statistical Parametric Mapping: The Analysis of Functional Brain Images—1st Edition. Retrieved February 28, 2019, from https://www.elsevier.com/books/statistical-parametric-mapping-the-analysis-of-functional-brain-images/penny/978-0-12-372560-8

Polyak, S. L. (1941). *The retina.*

Posner, M I. (1994). Attention: The mechanisms of consciousness. *Proceedings of the National Academy of Sciences of the United States of America*, *91*(16), 7398–7403.

Posner, M. I., & Petersen, S. E. (1990). The attention system of the human brain. *Annual Review of Neuroscience*, *13*, 25–42. https://doi.org/10.1146/annurev.ne.13.030190.000325

Posner, M. I., Snyder, C. R., & Davidson, B. J. (1980). Attention and the detection of signals. *Journal of Experimental Psychology*, *109*(2), 160–174.

Posner, Michael I., & Cohen, Y. (1984). Components of visual orienting. *Attention and Performance X: Control of Language Processes*, *32*, 531–556.

Pratte, M. S., Sy, J. L., Swisher, J. D., & Tong, F. (2016). Radial bias is not necessary for orientation decoding. *NeuroImage*, *127*, 23–33. https://doi.org/10.1016/j.neuroimage.2015.11.066

Ptak, R., Lazeyras, F., Di Pietro, M., Schnider, A., & Simon, S. R. (2014). Visual object agnosia is associated with a breakdown of object-selective responses in the lateral occipital cortex. *Neuropsychologia*, *60*, 10–20. https://doi.org/10.1016/j.neuropsychologia.2014.05.009

Raemaekers, M., Lankheet, M. J. M., Moorman, S., Kourtzi, Z., & van Wezel, R. J. A. (2009). Directional anisotropy of motion responses in retinotopic cortex. *Human Brain Mapping*, *30*(12), 3970–3980. https://doi.org/10.1002/hbm.20822

Raichle, M. E. (2015). The brain's default mode network. *Annual Review of Neuroscience*, *38*, 433–447. https://doi.org/10.1146/annurev-neuro-071013-014030

Raichle, M. E., MacLeod, A. M., Snyder, A. Z., Powers, W. J., Gusnard, D. A., & Shulman, G. L. (2001). A default mode of brain function. *Proceedings of the National Academy of Sciences*, *98*(2), 676–682.

Rees, G., Frackowiak, R., & Frith, C. (1997). Two modulatory effects of attention that mediate object categorization in human cortex. *Science (New York, N.Y.)*, *275*(5301), 835–838.

Reineberg, A. E., Gustavson, D. E., Benca, C., Banich, M. T., & Friedman, N. P. (2018). The Relationship Between Resting State Network

Connectivity and Individual Differences in Executive Functions. *Frontiers in Psychology*, *9*, 1600. https://doi.org/10.3389/fpsyg.2018.01600

Reynolds, J. H., & Heeger, D. J. (2009). The normalization model of attention. *Neuron*, *61*(2), 168–185. https://doi.org/10.1016/j.neuron.2009.01.002

Reynolds, J. H., Pasternak, T., & Desimone, R. (2000). Attention increases sensitivity of V4 neurons. *Neuron*, *26*(3), 703–714.

Rockland, K. S., & Pandya, D. N. (1979). Laminar origins and terminations of cortical connections of the occipital lobe in the rhesus monkey. *Brain Research*, *179*(1), 3–20.

Roorda, A., & Williams, D. R. (1999). The arrangement of the three cone classes in the living human eye. *Nature*, *397*(6719), 520–522. https://doi.org/10.1038/17383

Rossion, B., Caldara, R., Seghier, M., Schuller, A.-M., 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: A Journal of Neurology*, *126*(Pt 11), 2381–2395. https://doi.org/10.1093/brain/awg241

Russ, B. E., & Leopold, D. A. (2015). Functional MRI mapping of dynamic visual features during natural viewing in the macaque. *NeuroImage*, *109*, 84–94. https://doi.org/10.1016/j.neuroimage.2015.01.012

Ryu, S.-Y., Kwon, M. J., Lee, S.-B., Yang, D. W., Kim, T.-W., Song, I.-U., … Lee, A. Y. (2010). Measurement of precuneal and hippocampal

volumes using magnetic resonance volumetry in Alzheimer's disease. *Journal of Clinical Neurology*, *6*(4), 196–203.

Saenz, M., Buracas, G. T., & Boynton, G. M. (2002). Global effects of feature-based attention in human visual cortex. *Nat Neurosci*, *5*(7), 631–632.

Safford, A. S., Hussey, E. A., Parasuraman, R., & Thompson, J. C. (2010). Object-based attentional modulation of biological motion processing: Spatiotemporal dynamics using functional magnetic resonance imaging and electroencephalography. *The Journal of Neuroscience: The Official Journal of the Society for Neuroscience*, *30*(27), 9064–9073. https://doi.org/10.1523/JNEUROSCI.1779-10.2010

Sani, I., Santandrea, E., Golzar, A., Morrone, M. C., & Chelazzi, L. (2013). Selective tuning for contrast in macaque area V4. *Journal of Neuroscience*, *33*(47), 18583–18596.

Sapountzis, P., Schluppeck, D., Bowtell, R., & Peirce, J. W. (2010). A comparison of fMRI adaptation and multivariate pattern classification analysis in visual cortex. *NeuroImage*, *49*(2), 1632–1640. https://doi.org/10.1016/j.neuroimage.2009.09.066

Sasaki, Y., Rajimehr, R., Kim, B. W., Ekstrom, L. B., Vanduffel, W., & Tootell, R. B. H. (2006). The radial bias: A different slant on visual orientation sensitivity in human and nonhuman primates. *Neuron*, *51*(5), 661–670. https://doi.org/10.1016/j.neuron.2006.07.021

Schiller, P. H. (1995). Effect of lesions in visual cortical area V4 on the recognition of transformed objects. *Nature*, *376*(6538), 342–344. https://doi.org/10.1038/376342a0

Schiller, P. H. (1996). On the specificity of neurons and visual areas. *Behavioural Brain Research*, *76*(1–2), 21–35.

Schiller, P. H., & Malpeli, J. G. (1977). The effect of striate cortex cooling on area 18 cells in the monkey. *Brain Research*, *126*(2), 366–369.

Schnapf, J. L., Kraft, T. W., & Baylor, D. A. (1987). Spectral sensitivity of human cone photoreceptors. *Nature*, *325*(6103), 439–441. https://doi.org/10.1038/325439a0

Schoenfeld, M. A., Hopf, J.-M., Martinez, A., Mai, H. M., Sattler, C., Gasde, A., … Hillyard, S. A. (2007). Spatio-temporal analysis of feature-based attention. *Cerebral Cortex*, *17*(10), 2468–2477.

Sclar, G., Maunsell, J. H., & Lennie, P. (1990). Coding of image contrast in central visual pathways of the macaque monkey. *Vision Research*, *30*(1), 1–10.

Semedo, J. D., Zandvakili, A., Machens, C. K., Yu, B. M., & Kohn, A. (2019). Cortical Areas Interact through a Communication Subspace. *Neuron*, *102*(1), 249-259.e4. https://doi.org/10.1016/j.neuron.2019.01.026

Serences, J. T., & Boynton, G. M. (2007). Feature-based attentional modulations in the absence of direct visual stimulation. *Neuron*, *55*(2), 301–312. https://doi.org/10.1016/j.neuron.2007.06.015

Serences, J. T., Saproo, S., Scolari, M., Ho, T., & Muftuler, L. T. (2009). Estimating the influence of attention on population codes in human visual cortex using voxel-based tuning functions. *NeuroImage*, *44*(1), 223–231. https://doi.org/10.1016/j.neuroimage.2008.07.043

Seymour, K., Clifford, C. W. G., Logothetis, N. K., & Bartels, A. (2009). The coding of color, motion, and their conjunction in the human visual

cortex. *Current Biology: CB*, *19*(3), 177–183. https://doi.org/10.1016/j.cub.2008.12.050

Seymour, K., Clifford, C. W. G., Logothetis, N. K., & Bartels, A. (2010). Coding and binding of color and form in visual cortex. *Cerebral Cortex (New York, N.Y.: 1991)*, *20*(8), 1946–1954. https://doi.org/10.1093/cercor/bhp265

Seymour, K. J., Williams, M. A., & Rich, A. N. (2016). The Representation of Color across the Human Visual Cortex: Distinguishing Chromatic Signals Contributing to Object Form Versus Surface Color. *Cerebral Cortex (New York, N.Y.: 1991)*, *26*(5), 1997–2005. https://doi.org/10.1093/cercor/bhv021

Shapley, R., & Enroth-Cugell, C. (1984). Chapter 9 Visual adaptation and retinal gain controls. *Progress in Retinal Research*, *3*, 263–346. https://doi.org/10.1016/0278-4327(84)90011-7

Shulman, G. L., Fiez, J. A., Corbetta, M., Buckner, R. L., Miezin, F. M., Raichle, M. E., & Petersen, S. E. (1997). Common Blood Flow Changes across Visual Tasks: II. Decreases in Cerebral Cortex. *Journal of Cognitive Neuroscience*, *9*(5), 648–663. https://doi.org/10.1162/jocn.1997.9.5.648

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

Silson, E. H., McKeefry, D. J., Rodgers, J., Gouws, A. D., Hymers, M., & Morland, A. B. (2013b). Specialized and independent processing of

orientation and shape in visual field maps LO1 and LO2. *Nature Neuroscience*, *16*(3), 267.

Smith, P. L., Lee, Y.-E., Wolfgang, B. J., & Ratcliff, R. (2009). Attention and the detection of masked radial frequency patterns: Data and model. *Vision Research*, *49*(10), 1363–1377. https://doi.org/10.1016/j.visres.2008.04.024

Song, J.-H., Rowland, J., McPeek, R. M., & Wade, A. R. (2011). Attentional Modulation of fMRI Responses in Human V1 Is Consistent with Distinct Spatial Maps for Chromatically Defined Orientation and Contrast. *Journal of Neuroscience*, *31*(36), 12900–12905. https://doi.org/10.1523/JNEUROSCI.0580-11.2011

Spadone, S., Della Penna, S., Sestieri, C., Betti, V., Tosoni, A., Perrucci, M. G., … Corbetta, M. (2015). Dynamic reorganization of human resting-state networks during visuospatial attention. *Proceedings of the National Academy of Sciences of the United States of America*, *112*(26), 8112–8117. https://doi.org/10.1073/pnas.1415439112

Spiers, H. J., & Maguire, E. A. (2007). Decoding human brain activity during real-world experiences. *Trends in Cognitive Sciences*, *11*(8), 356–365. https://doi.org/10.1016/j.tics.2007.06.002

Stanislaw, H., & Todorov, N. (1999). Calculation of signal detection theory measures. *Behavior Research Methods, Instruments, & Computers: A Journal of the Psychonomic Society, Inc*, *31*(1), 137–149.

Sumner, P. (2006). Seeing colour. *Advances in Clinical Neuroscience and Rehabilitation*, *6*(3), 12–13.

Sumner, P., Anderson, E. J., Sylvester, R., Haynes, J.-D., & Rees, G. (2008). Combined orientation and colour information in human V1 for both L-M and S-cone chromatic axes. *NeuroImage*, *39*(2), 814–824. https://doi.org/10.1016/j.neuroimage.2007.09.013

Sumner, P., Tsai, P.-C., Yu, K., & Nachev, P. (2006). Attentional modulation of sensorimotor processes in the absence of perceptual awareness. *Proceedings of the National Academy of Sciences of the United States of America*, *103*(27), 10520–10525. https://doi.org/10.1073/pnas.0601974103

Teo, P. C., Sapiro, G., & Wandell, B. A. (1997). Creating connected representations of cortical gray matter for functional MRI visualization. *IEEE Transactions on Medical Imaging*, *16*(6), 852–863. https://doi.org/10.1109/42.650881

Theeuwes, J. (2013). Feature-based attention: It is all bottom-up priming. *Philosophical Transactions of the Royal Society of London. Series B, Biological Sciences*, *368*(1628), 20130055. https://doi.org/10.1098/rstb.2013.0055

Tootell, R. B., Hadjikhani, N., Hall, E. K., Marrett, S., Vanduffel, W., Vaughan, J. T., & Dale, A. M. (1998). The retinotopy of visual spatial attention. *Neuron*, *21*(6), 1409–1422.

Tootell, Roger BH, Hadjikhani, N., Hall, E. K., Marrett, S., Vanduffel, W., Vaughan, J. T., & Dale, A. M. (1998a). The retinotopy of visual spatial attention. *Neuron*, *21*(6), 1409–1422.

Tootell, Roger BH, Hadjikhani, N., Hall, E. K., Marrett, S., Vanduffel, W., Vaughan, J. T., & Dale, A. M. (1998b). The retinotopy of visual spatial attention. *Neuron*, *21*(6), 1409–1422.

Treue, S., & Martínez Trujillo, J. C. (1999). Feature-based attention influences motion processing gain in macaque visual cortex. *Nature*, *399*(6736), 575–579. https://doi.org/10.1038/21176

Tschechne, S., & Neumann, H. (2014). Hierarchical representation of shapes in visual cortex-from localized features to figural shape segregation. *Frontiers in Computational Neuroscience*, *8*, 93. https://doi.org/10.3389/fncom.2014.00093

Vandenberghe, R., Duncan, J., Dupont, P., Ward, R., Poline, J. B., Bormans, G., … Orban, G. A. (1997). Attention to one or two features in left or right visual field: A positron emission tomography study. *The Journal of Neuroscience: The Official Journal of the Society for Neuroscience*, *17*(10), 3739–3750.

Vander Wyk, B. C., Voos, A., & Pelphrey, K. A. (2012). Action representation in the superior temporal sulcus in children and adults: An fMRI study. *Developmental Cognitive Neuroscience*, *2*(4), 409–416.

Vatansever, D., Manktelow, A. E., Sahakian, B. J., Menon, D. K., & Stamatakis, E. A. (2016). Cognitive Flexibility: A Default Network and Basal Ganglia Connectivity Perspective. *Brain Connectivity*, *6*(3), 201–207. https://doi.org/10.1089/brain.2015.0388

Verghese, P., Kim, Y.-J., & Wade, A. R. (2012). Attention selects informative neural populations in human V1. *Journal of Neuroscience*, *32*(46), 16379–16390.

Verhoef, B.-E., & Maunsell, J. H. R. (2017). Attention-related changes in correlated neuronal activity arise from normalization mechanisms. *Nature Neuroscience*, *20*(7), 969–977. https://doi.org/10.1038/nn.4572

Vernon, R. J. W. (2016). *Shape Processing across Lateral Occipital Cortex* (Phd, University of York). Retrieved from http://etheses.whiterose.ac.uk/16777/

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. https://doi.org/10.1523/JNEUROSCI.3603-15.2016

Wandell, B. A. (1995). *Foundations of vision*. Sunderland, MA, US: Sinauer Associates.

Wandell, B. A., Dumoulin, S. O., & Brewer, A. A. (2007). Visual field maps in human cortex. *Neuron*, *56*(2), 366–383.

Wandell, B. A., & Winawer, J. (2015). Computational neuroimaging and population receptive fields. *Trends in Cognitive Sciences*, *19*(6), 349–357. https://doi.org/10.1016/j.tics.2015.03.009

Wang, J., & Wade, A. R. (2011). Differential attentional modulation of cortical responses to S-cone and luminance stimuli. *Journal of Vision*, *11*(6), 1–15. https://doi.org/10.1167/11.6.1

Watson, A. B., & Pelli, D. G. (1983). QUEST: A Bayesian adaptive psychometric method. *Perception & Psychophysics*, *33*(2), 113–120.

Weerda, R., Vallines, I., Thomas, J. P., Rutschmann, R. M., & Greenlee, M. W. (2006). Effects of nonspatial selective and divided visual attention on fMRI BOLD responses. *Experimental Brain Research*, *173*(4), 555–563. https://doi.org/10.1007/s00221-006-0403-0

Welbourne, L. E., Morland, A. B., & Wade, A. R. (2018). Population receptive field (pRF) measurements of chromatic responses in human visual cortex using fMRI. *Neuroimage*, *167*, 84–94. https://doi.org/10.1016/j.neuroimage.2017.11.022

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: CB*, *10*(22), 1455–1458.

Wilkinson, Frances, Wilson, H. R., & Habak, C. (1998a). Detection and recognition of radial frequency patterns 1. *Vision Research*, *38*(22), 3555–3568.

Wilkinson, Frances, Wilson, H. R., & Habak, C. (1998b). Detection and recognition of radial frequency patterns 1. *Vision Research*, *38*(22), 3555–3568.

Wilson, H. R., & Wilkinson, F. (1997a). Evolving concepts of spatial channels in vision: From independence to nonlinear interactions. *Perception*, *26*(8), 939–960. https://doi.org/10.1068/p260939

Wilson, H. R., & Wilkinson, F. (1997b). Evolving concepts of spatial channels in vision: From independence to nonlinear interactions. *Perception*, *26*(8), 939–960. https://doi.org/10.1068/p260939

Wojciulik, E., Kanwisher, N., & Driver, J. (1998). Covert Visual Attention Modulates Face-Specific Activity in the Human Fusiform Gyrus: FMRI Study. *Journal of Neurophysiology*, *79*(3), 1574–1578. https://doi.org/10.1152/jn.1998.79.3.1574

Xing, Y., Ledgeway, T., McGraw, P. V., & Schluppeck, D. (2013). Decoding working memory of stimulus contrast in early visual cortex. *The Journal of Neuroscience: The Official Journal of the Society for Neuroscience*, *33*(25), 10301–10311. https://doi.org/10.1523/JNEUROSCI.3754-12.2013

Yi, D.-J., Kelley, T. A., Marois, R., & Chun, M. M. (2006). Attentional modulation of repetition attenuation is anatomically dissociable for scenes and faces. *Brain Research*, *1080*(1), 53–62. https://doi.org/10.1016/j.brainres.2006.01.090

Zeki, S. (1990). A century of cerebral achromatopsia. *Brain: A Journal of Neurology*, *113 (Pt 6)*, 1721–1777. https://doi.org/10.1093/brain/113.6.1721

Zeki, S. (2004). Thirty years of a very special visual area, Area V5. *The Journal of Physiology*, *557*(Pt 1), 1–2. https://doi.org/10.1113/jphysiol.2004.063040

Zeki, S. M. (1969). Representation of central visual fields in prestriate cortex of monkey. *Brain Research*, *14*(2), 271–291.

Zeki, S. M. (1973). Colour coding in rhesus monkey prestriate cortex. *Brain Research*, *53*(2), 422–427.

Zeki, S. M. (1978). Functional specialisation in the visual cortex of the rhesus monkey. *Nature*, *274*(5670), 423–428.

Zihl, J., von Cramon, D., Mai, N., & Schmid, C. (1991). Disturbance of movement vision after bilateral posterior brain damage. Further evidence and follow up observations. *Brain: A Journal of Neurology*, *114 (Pt 5)*, 2235–2252. https://doi.org/10.1093/brain/114.5.2235