
Investigating the functional organisation of human visual cortex using ultra-high resolution fMRI



Ms Sarah Finnegan MSc BSc(Hons)

St Catherine's College

University of Oxford

A thesis submitted for the degree of

Doctor of Philosophy

Trinity 2016

Investigating the functional organisation of human visual cortex using ultra-high resolution fMRI

**Sarah Finnegan
St. Catherine's College**

D.Phil. Thesis. Trinity Term 2016

Abstract

Current thinking suggests that specialised modules process visual information in a hierarchical manner, using local circuitry in order to maximise efficiency both in terms of wiring costs and stimulus coverage (Reichl et al. (2012)). The resulting organisation has been shown to contain structure in the form of stripes, columns and pinwheels, which in animal models have been linked to functional segregation and specificity. In human cortex, post-mortem investigations have assisted in the visualisation of two such key features: ocular dominance columns (ODC) within V1, and a stripe system within V2 (Adams et al. (2007); Hockfield et al. (1990)). However, functional observations in humans have until recently been beyond the investigatory scope of *in-vivo* methodology, and as such, a role of these networks has yet to be conclusively determined.

In the small number of instances of *in-vivo* investigations of human ODC and V2 stripes (Cheng et al. (2001); Yacoub et al. (2001); Nasr et al. (2016)), data have been acquired for a small number of carefully selected participants over long scan durations. I aimed to overcome these limitations and explore the functional similarities further, employing a novel, ultra-high resolution fMRI sequence to do so.

I measured the cortical response to monocular stimulation and recorded a robust response within V1. However, the regular and repeating functional patterns of ODCs were not observed. Using multivariate techniques I concluded, based on robust classification, that reliable monocular signals were present but that they were subtle and difficult to differentiate from noise. I then investigated the segregation of colour, form and motion within V2, where I found evidence for spatially segregated signals in response to colour and motion, but not to form. I hypothesised that the form stimulus was sub-optimal in driving the neural population of the associated stripes. Based on a limited number of samples, activity in response to colour and motion stimulation conformed on average to the neuroanatomical profile of the V2 stripe system. I suggest that my results offer encouragement for *in-vivo* investigations of small-scale functional organisation in visual cortex.

Acknowledgements

I would first like to thank my supervisors, Andrew Parker and Holly Bridge. To Andrew, who has provided a limitless source of knowledge of the visual system, with a keen insight for every question, I offer my thanks. Thanks also to Holly for her constant support and guidance in my pursuit of ultra-high functional resolution. Holly demonstrates the positive role that a scientist can play in the community, and has allowed me to develop as an independent researcher in the exciting environment that is FMRIB.

I would also like to extend my thanks to my college adviser Kerry Walker, whose can-do attitude encouraged me to apply for a college scholarship. For granting me this scholarship I offer my gratitude to St Catherine's College. And to my friends here at college, it has been a fantastic and fun-filled experience. Laura O'Connor, Michael Grange and Victoria Sainsbury, I thank you in particular for providing enough cultural conversation to stretch my mind in a different direction and enough gin to ensure that no cognitive improvement was permanent.

To my lab mates and colleagues both past and present at DPAG and FMRIB Ivan Alvarez, Gaelle Coullon, Stephanie Larcombe, Jackson Smith, Kimmy Tang-Wright, Miaoqing Shi and Nina Reislev. Our social gatherings of pot-luck dinners, hikes in the countryside and well earned Friday pub outings have made my experience of Oxford truly enjoyable. A special thank you to my excellent colleague Betina Ip who has provided constant positivity and true resourcefulness. Thanks also to my wonderful lab-mate and one time flatmate Imogen Large. A kindred spirit from the start, a well timed wry eyebrow has brightened my day on many occasions.

To James Bennett, who started off as my tree-climbing colleague and ended as my tree-climbing partner in crime. Your support has been invaluable, whether it be tirelessly coaching me through endless MATLAB problems or providing me with tea and jokes at crucial moments.

Finally and most importantly, my thanks are given to my family, Eamonn Finnegan, Susan Finnegan and Helen Finnegan who have always believed in me. Through the love and support of my family I have had the opportunity to dream so much bigger than I could have dared. I cannot thank you enough.

Contents

1	Introduction	1
1.1	Abstract	2
1.2	Introduction	3
1.3	An overview of the functional organisation of V1	10
1.4	An overview of the functional organisation of V2	14
1.5	Using ultra-high resolution fMRI to investigate the hypercolumn	18
1.6	Aims and achievements	20
2	General methods	25
2.1	Abstract	26
2.2	Experimental sessions	26
2.3	Functional MRI and the origin of the BOLD signal	28
2.4	Retinotopic methods	34
2.5	Cortical surface based methods	48
3	Advanced methods	53
3.1	Abstract	54
3.2	Introduction	55

3.3	Methods	63
3.4	Results - Part 1	69
3.5	Discussion - Part 1	77
3.6	Conclusion - Part 1	78
3.7	Further fMRI protocol development	78
3.8	Discussion - Part 2	80
3.9	Conclusions	82
4	Investigating Ocular Dominance Columns	83
4.1	Abstract	84
4.2	Introduction	85
4.3	Methods	93
4.4	Results	103
4.5	Discussion	119
4.6	Conclusion	122
5	The V2 stripe system	124
5.1	Abstract	125
5.2	Introduction	125
5.3	Chapter aims	131
5.4	Methods	133
5.5	Results	140
5.6	Discussion	164

5.7	Conclusions	171
6	Multi-Voxel Pattern Analysis of ocular dominance columns and V2 stripes	172
6.1	Abstract	173
6.2	Introduction	174
6.3	Methods	179
6.4	Results 1: Ocular Dominance Columns	188
6.5	Results 2: V2 segregation of colour, form and motion	196
6.6	Results 3: Segregation of colour, form and motion in other visual areas	218
6.7	Discussion	226
6.8	Conclusion	232
7	General discussion	233
7.1	Abstract	234
7.2	Introduction	235
7.3	Major findings	236
7.4	Future experiments	245
7.5	Conclusions	249
8	Appendix	250
	Bibliography	257

Chapter 1

Introduction

1.1 Abstract

In this chapter I introduce the concept of functional modules in human visual cortex. Firstly I provide an overview of the human visual system and the modular organisation included therein. I then describe the small-scale feature specific processing networks known as hypercolumns. Hypercolumn networks have been identified within several different visual areas, and in this thesis I focus on two systems in particular; the ocular dominance columns of V1, and the stripe system of V2. I discuss briefly the literature of the *in-vivo* investigation of these networks, which I then expand upon in Chapters 4 and 5. Finally, I introduce the use of magnetic resonance imaging (MRI) techniques as a non-invasive method of *in-vivo* investigation. I conclude by summarising the major aims and achievements of this thesis.

1.2 Introduction

Despite the fact that visual scenes are complex, multidimensional and information rich, we perceive the visual environment holistically, and confidently navigate our surroundings.

Cortical modules are thought to play a key role in simplifying the process of translating retinal electrical potentials into meaningful content. Using hierarchical processing, simple features such as orientation and contrast are processed within distinct cortical units before being combined in higher visual areas to form complex constructs, such as faces or scenes. The dense network of connections to higher cortical areas are also bi-directional, and top-down mechanisms are thought to interact with early visual cortex to modulate attention.

At points along this pathway, visual information is divided into specific processing streams. Segregation can occur through structural division, for example within the LGN, where depending on ocular input, projections of retinal ganglion cells synapse with different layers, or alternatively, via functional division. One such example of structured functional organisation is retinotopy, whereby visual information is organised such that it topographically mirrors an object's spatial position on the retina. This organisational method is repeated across different visual areas and the change in visual field representation can be used to reliably differentiate neighbouring regions (Serenio et al. (1995)). In other cases delineations can be made on the basis of feature specificity, where clusters of cells demonstrate common receptive field properties. This channelling centres on features that include: ocular input, colour, orientation, distance and contrast. Unique combinations of these features can be used to define objects in association with their position on the retina (Bartels (2009)).

The presence of modular stimulus preference, where large areas of cortex respond preferentially to a stimulus group, is more frequently associated with higher visual areas. For example, V4 has been identified as responding strongly to colour stimulation, (Gallant et al. (2000); McKeefry and Zeki (1997)), while V5 has been shown to respond to visual stimuli modulated to contain motion (Kolster et al. (2010); Shipp and Zeki (1989); McKeefry and Zeki (1997)). Within early visual cortex, such modularised feature-based processing has also been identified, but

at a much smaller scale. Despite the identification of topological feature maps across a wide range of mammalian species and visual areas, the functional significance of their presence remains unclear (Wilson and Bednar (2015)). In contrast, the retinotopic organisation of visual space, and the modular basis of feature processing, within and across human cortical regions has received considerable research attention (Wandell et al. (2005, 2007); Sereno et al. (1995); Dougherty et al. (2003); Wandell and Winawer (2011); Larsson and Heeger (2006); Tootell et al. (1995); Kanwisher et al. (1997); Downing et al. (2006); Bartels (2009)).

The disparity between the information available regarding these organisational methods of feature processing can be explained in terms of the size difference of the investigatory targets. Nevertheless, small-scale feature maps within visual areas offer a unitary example with which to probe network dynamics, wiring optimisation techniques and the effect of visual deprivation and damage at a local and system wide level. Their investigation is of paramount importance, and this thesis focuses on addressing the methodological constraints that have limited this field of research. To this end, the topological organisation of feature based processing within two members of this network, V1 and V2 are investigated.

1.2.1 An overview of the human visual system

In the human visual system, light waves transmit information to the retina by stimulating the retinal photoreceptors. The decussation of neural fibres serving the nasal retinal hemifield facilitates the transfer of electrical potentials across the optic chiasm. The optic tract, a large white matter bundle, continues this transfer and terminates within the Lateral Geniculate Nucleus (LGN). At this point monocular afferents synapse within different layers of the LGN, thus maintaining the segregation of each eye. The LGN, which consists of six layers, separates ocular input as follows: layers 2, 3 and 5 receive input from the contralateral retina, while layers 1, 4 and 6 receive their input from the ipsilateral retina. A second division within the LGN segregates cell groups belonging to the magno (M) and parvo (P) pathways. Cells of the M-pathway are associated large receptive fields and short response latencies. This is in contrast

to cells of the P-pathway, which are smaller and slower than their M counterparts (Maunsell et al. (1999)). Within the LGN, layers 1 and 2 receive their principle afferents from the M-pathway, emanating from the parasol retinal ganglion cells. Layers 3-6 receive their input from the P-pathway, which arises from the midget retinal ganglion cells (Callaway (2005)).




Figure 1.1: A diagram illustrating the flow of visual information from the retina to the striate cortex (V1). Visual information from the left or right visual field falling upon the nasal hemifields of the retina (A) is subsequently routed across the optic chiasm to the contralateral hemisphere. As a result of this decussation of neural fibres, the left visual field is processed within the right LGN, a preference reversed for the left LGN. The termination of afferents from the left and right retina into different layers of the LGN (B) maintains a segregated representation of the input of each eye. The afferents of each of the six layers of the LGN, which send their major projections to the striate cortex (V1), continue this ocular segregation (C)

Leaving the LGN, the strongest projections in the primate are received by layer 4C of V1, the first and largest cortical area in the occipital cortex (Briggs and Usrey (2007); Hubel and Wiesel (1972)) (Fig 1.1). These visual afferents, which connect to early visual cortex, arise largely from the contralateral visual hemifield, a bias that is less pronounced in higher cortical areas (Tootell et al. (1998b)). Within the occipital cortex a network of cortical areas have been identified (Fig 1.2). Features particular to a region can be used to facilitate its identification and include: stimulus specific processing preferences, cytoarchitectural features, and boundary

points between ordered representations of external visual space.




Figure 1.2: Visual areas across the cortex. Adapted from Wandell et al. (2005). Visual areas are shown here for four orientations of the partially inflated right hemisphere of one subject (AB). Visual information enters V1 from the LGN (Fig1.1). From V1 visual afferents project to higher visual areas, progressing initially through V2 and V3 before two major pathways develop. The “what” pathway, which has been identified as processing colour and form, includes V4 and VO-1. The “where” pathway encompasses areas V3a and hMT+/(V5), and processes motion and depth information.

Within visual areas V1-V3 cells demonstrating receptive field preferences for basic features such as colour, contrast, orientation and absolute disparity have been identified (Cumming and Parker (1999); Neri et al. (2004); Nienborg et al. (2004); Riesenhuber and Poggio (2000, 2002); Zhang et al. (2014)). Feeding forwards from V3, visual information is thought to follow one of two processing streams, commonly known as the “where” or “vision for action” and the “what” or “vision for perception” streams (Neri et al. (2004); Riesenhuber and Poggio (1999); Mishkin and Ungerleider (1982)). As a major contributor to dorsal “where” pathway, V5 responds preferentially to motion components of visual stimuli (Dumoulin et al. (2000); Orban et al. (2003)). In contrast V4, a ventral module, is involved in the processing of form and colour, which are typically considered to be “what” components (Gallant et al. (2000); Ungerleider and Haxby (1994)). The progression from posterior to anterior visual cortex is

associated with the increasing specialisation of cell response properties, lending support to the theory of hierarchical visual processing across the visual cortex.

1.2.2 Functional modules in visual cortex

Functional organisation in the visual cortex takes several forms. In order to organise neural signaling between, and within modules, visual information is arranged as a topographic map. Within these maps spatially proximal neurons process adjacent points of visual space in a point-to-point translation of the retina known as retinotopy (Wandell et al. (2007)). Each module contains its own representation of visual space, and the structure provided by retinotopic organisation acts as a common reference frame across modules, facilitating the efficient processing of visual information.

The division of cortical space into modules has been hypothesised to act as an efficiency strategy. These modules act as nested processors, avoiding the long distance cortical wiring that would violate optimisation principles. Within an external visual scene spatial content is highly relevant for decoding object identity. It is therefore likely that neurons with adjacent receptive fields will have strong temporal correlations, and their spatial proximity therefore acts to reduce cortical wiring cost (Reichl et al. (2012); Wandell et al. (2005); Swindale et al. (2000); Wandell et al. (2002); Kaas (1997)).

Similarly, wiring conservation strategies may also be in place within modules, giving rise to further internal structure (de Haan and Cowey (2011); Modha and Singh (2010)). This internal organisation within cortical modules is the particular focus of this thesis. The relationship between the presence of visual functional organisation, feature segregation and behavioural adaptation is largely unknown. In the instance of human research, methodological constraints have limited the *in-vivo* investigation of small-scale functional organisation. The way in which networks communicate, organise information and change with experience is key to visual processing. Exploring these networks within, and between cortical modules will further our understanding of feature processing and the role of top-down attentional modules.

1.2.3 Hypercolumns within visual cortex

Investigation of structured functional networks within visual areas *in-vivo* in humans is limited by the spatial resolution of data collected when using techniques such as fMRI. Electrophysiological investigations and post-mortem staining have, for decades, provided evidence for the presence of small-scale functional organisation within visual areas (Adams et al. (2007); Hubel and Wiesel (1969); Livingstone and Hubel (1988); Sincich and Horton (2002); Sincich et al. (2003); Ts'o et al. (2001)). Based on this work, there is evidence to suggest that these small-scale networks demonstrate functional preference in their receptive field properties.

The regular and repeating pattern made by these functional networks within a visual area has been termed a hypercolumn. Hubel and Wiesel, when first defining the hypercolumn in their electrophysiological investigation of cat cortex (Hubel and Wiesel (1974)) sought to “extend the concept of the (cortical) column, and consider a complete array of columns as a small machine that looks after all values of a given variable” (Fig 1.3). The extent to which visual information is specific to, and segregated by these features is not yet fully understood. The hypercolumn may be key to keeping related processing circuitry local, and provide vital clues as to how functional specialisation is organised within modules and across the cortex (Van Hooser (2007)).




Figure 1.3: Relationships among pinwheels, ocular dominance columns and blobs in primary visual cortex of the macaque monkey, as revealed by optical imaging (Adapted from Grinvald et al. (2004)). (A) Ocular dominance columns (ODCs). Dark bars indicate a column of right eye dominant neurons, Scale bar=500 μ m. (B) Black lines indicate the borders of the ODCs, overlaid onto an angle map of orientation preference known as orientation pinwheels. The pinwheel centers are marked here with circles. (C) Cytchrome oxidase blobs, which represent a collection of neurons with demonstrable chromatic sensitivity. (D) A composite of CO blobs, orientation pinwheels and ODCs overlaid to show their relationships. (E) depicts the full organisation of these hypercolumn structures, based around the pinwheel model. Black lines mark the borders between ODCs. White ovals represent groups of neurons responsible for color perception (blobs). The pinwheels are formed by neurons involved in the perception of shape, with each color marking a column of neurons responding selectively to a particular orientation in space.

This thesis focuses on two hypercolumn networks. The ocular dominance columns (ODCs) of V1, and the V2 stripe system.

ODCs have been extensively characterised in non-human neurophysiological investigations of the cat, ferret and macaque (Hubel and Wiesel (1969); Horton and Hocking (1996); Hubel and Wiesel (1972, 1974)). In contrast, investigations of response properties of the V2 stripe system in non-human primates have been faced with methodological difficulties. V2 is positioned deep within the lunate sulcus of the macaque, a species often considered to be the closest experimental homologue to humans. The difficulty in recording within this sulcus, combined with the variable uptake of metabolic stain has made the V2 stripes difficult to characterise. In work characterising ODCs in human cortex, post-mortem staining of subjects who sustained damage to, or loss of one eye at an early age has revealed a similar columnar organisation to that observed within the macaque (Adams et al. (2007)). V2 stripes have also been reconstructed

via staining techniques applied to human cortex, although only partially, in a small section of V2 (Hockfield et al. (1990)).

In-vivo investigations of human ODCs and V2 stripes have been limited by the small spatial scale of the target features. I attempt to address this restriction in this thesis. In the following sections I provide a brief introduction to the ODCs and V2 stripes, as well as the cortical areas within which they can be identified.

1.3 An overview of the functional organisation of V1

V1 is the largest visual cortical area. The surface area of this region can vary up to three times in size across individuals, a measurement which correlates with the size of the optic tract and LGN (Schwarzkopf et al. (2011); Dougherty et al. (2003)). On average, the area of human V1 representing between 2° and 12° of visual space has been estimated to occupy 1470mm^2 of cortical space (Dougherty et al. (2003)).

V1 is widely believed to be homologous across mammals, and can be uniquely identified as fulfilling the following criteria (Van Hooser (2007)):

- Containing a retinotopic map of space
- Receiving a major projection from the LGN
- Possessing a dense band of myelin within layer 4

V1 contains a dense band of myelination known as the stria of Gennari. The termination of this band can be used to identify the border between V1 and V2 (Fig1.4 **A**). V1 is situated around the calcarine sulcus (Fig1.4 **B**), where the representation of the visual field is spatiotopically reversed and divided along the horizontal meridian by the two banks (Fig1.4 **C**). The retinal information of the fovea is contained at the occipital pole, while anterior points of V1 represent increasingly eccentric positions of visual space.




Figure 1.4: Defining the position of human V1. The position of V1 is highlighted in three panels. Using structural MR techniques or via myelin or Nissl staining, a dense band of myelin known as the stria of Gennari (SG) can be visualised (A). The stria of Gennari is unique to V1, and its termination marks the boundary points between V1 and V2. In panel B, the right hemisphere is orientated to expose ventral V1. V1 can be identified by its position (B) straddling the banks of the calcarine sulcus (CS). Alternatively, using fMRI retinotopic mapping techniques, V1 can be identified by the phase change in neuronal response to visual field location (C). A retinotopic map is shown here (C) projected onto a ventral view of the right hemisphere. The inset contains the same retinotopic data projected onto a flattened patch of the occipital cortex where both ventral and dorsal V1 can be more clearly visualised. The response to the presentation of a periodic stimulus is used to identify regions of V1 responsive to the upper (green), mid (blue) and lower (red) visual field (C). V1 is marked by a dotted white line. V1 and V2 share a common upper and lower visual field representation, the phase change from which can be used to distinguish the two. Panels A and B, adapted from Balaram et al. (2014)

1.3.1 Ocular Dominance Columns

The geniculate afferents projected via the optic radiation to V1 are segregated into columns serving principally the left or right eye. These monocular columns, known as ocular dominance columns are arranged in a banded pattern across V1 (Fig 1.5). The columns, which are orientated perpendicular to the calcarine sulcus, form a continuation of the optic radiations received from the layers of the LGN (Adams et al. (2007); Horton and Adams (2005); Hubel and Wiesel (1969)). V1 itself can be subdivided into 11 laminar divisions, from pial surface to white matter (1, 2, 3, 4A, 4B, 4C α , 4C β , 5A, 5B, 6A and 6B) (Balaram et al. (2014); Tootell et al. (1998b)). Monocular dominance can be observed through the cortical depth, but is most distinct in the input layer of 4C.




Figure 1.5: A schematic of ocular dominance columns overlaid onto a right human cortex. The subject lost one eye one year prior to death. During this time, the columns serving the remaining eye remained metabolically active. Metabolism is carried out in the mitochondria, which uptake the applied cytochrome oxidase stain, colouring the active cells dark. Adapted from Adams and Horton (2006)

Ocular dominance columns form a part of the processing pathway by which monocular cues are integrated to become binocular, necessary in the processing of depth or disparity. Normal binocular cues contain strong disparity information, where depth is encoded as a function of absolute distance from the observer and the relative distance between two objects in a scene and the observer. Single cell recordings in animals and fMRI investigations in humans suggest that relative depth calculations are performed later within V2 (Hubel and Livingstone (1987); Neri et al. (2004); Thomas et al. (2002)), V3 (Adams and Zeki (2001); Backus et al. (2001); Bridge and Parker (2007); Felleman and Van Essen (1987)) and V5/MT (Kolster et al. (2010); Maunsell and van Essen (1983); Neri et al. (2004)). These later modules rely upon the accurate balance of input of each eye's representation of the visual scene. The importance of monocular integration in this process suggests that ocular dominance columns may play a key role in depth perception in humans. However, assigning a functional role of the columns has been difficult, largely due to methodological challenges that arise as a result of their small size.

Suppositions of the role of ODCs have been made by examining the effect of alterations to the ocular balance on the ocular dominance column structure (Adams and Horton (2006); Adams

et al. (2007); Hubel and Wiesel (1963)). In patients who experience the enucleation of one eye at a young age, differences have been observed at post mortem in the size of the deprived eye's ODC compared to those of the non-deprived eye. (Adams et al. (2007); Horton and Hocking (1996, 1998)). This finding builds upon the evidence of the experimental animal literature (Hubel and Wiesel (1972, 1963)), in which a consistent change in ODC organisation followed periods of visual deprivation or enucleation. It has been suggested that for people with squint, some of whom experience depth perception deficit, this period of deprivation, albeit brief, occurs during a critical period of development, and is sufficient to alter the ODCs.

In this relationship the nature of cause and effect remains unclear, does a period of deprivation result in an over-expression of one column, allowing one eye to dominate and therefore affect binocular processing? Or, because of the changes to binocular processing, does one eye dominate and become overly represented in ocular dominance structure? The resulting effect appears to be a loss in visual acuity, possibly due to the loss of coherent binocular integration.

The role of ODCs in binocular processing is uncertain. Firm conclusions are difficult to draw due to the added challenge of the presence of large variations in hypercolumn presentation between, and even within species (Adams et al. (2007); Horton and Adams (2005); Livingstone (1996)). The literature from *in-vivo* human investigations is also somewhat sparse. The small size of the columns has, until recently, limited fMRI investigations. Where *in-vivo* investigations have been carried out, only a small number of ODC have been resolvable in participants who fulfill specific morphological criteria (Cheng et al. (2001); Yacoub et al. (2007)).

1.3.2 Investigating ocular dominance columns *in-vivo*

A detailed discussion of the *in-vivo* investigatory literature of human ODCs is provided within Chapter 4. Here I provide a brief summary.

At the inception of this thesis, two groups had successfully imaged small patches of ODCs in human cortex (Cheng et al. (2001); Yacoub et al. (2007)). Both groups selected their partic-

ipants based upon specific morphological constraints and, to achieve a reasonable SNR, the groups acquired data over two hours of scan time per session. I aimed to overcome these limitations, which prevent the wider use of fMRI, to investigate the role of ODC *in-vivo* in health and disease.

More recently, another group has claimed to have successfully imaged ODCs. For one subject Nasr et al., show evidence of a full reconstruction of the ODCs of one eye (Nasr et al. (2016)). However, limitations pervasive of the previous investigations remain. A single subject underwent long scan sessions, and the ODCs of only one eye were identified.

The development of reliable and robust imaging methods would provide the opportunity to investigate the effect of visual experience on the expression of ODC. These investigations could be repeated across multiple sessions over a period of years. Theoretically, changes to ODC of patients undergoing therapy for squint could also be monitored, and the cause-effect relationship of ODCs and depth perception further examined.

In the wider context, taking a target feature such as ODCs, for which there is considerable histological and electrophysiological evidence provides a testable prediction. In this case a null result can be assumed to be have a methodological explanation. Building on the development of a successful investigatory technique, feature processing within other visual areas can be investigated.

1.4 An overview of the functional organisation of V2

The area of cortical space occupied by V2 is on average 1115mm². V2 can vary in size between 1-3 times across different people and is slightly smaller than V1, from which V2 receives its major efferent connections (Dougherty et al. (2003); El-Shamayleh et al. (2013)). V2 curves around V1, dividing into ventral and dorsal components on either sides of the calcarine banks (Fig1.6 **B**). The map of visual space within V2 reverses direction from V1 (Fig1.6 **C**). Moving anterior from the boundary between the two areas moves from the vertical meridian of either

the upper or lower field towards the horizontal axis (Wandell et al. (2007)). The point of angular reversal can be used to find the boundary between V1 and V2, and V2 and V3.




Figure 1.6: The position of human V2 is highlighted in three panels. (A) The dense band of myelin, known as the stria of Gennari (SG), which terminates at the boundary between V1 and V2 can be used to identify the posterior border of V2. In panel B the right hemisphere is orientated to expose ventral V2. V2 can be identified by its position, running along the upper (V2d) and lower banks (V2v) of the calcarine sulcus, surrounding V1. Using fMRI retinotopic mapping techniques V2 can also be identified by the neuronal population response to specific visual locations. The colours shown in (C) indicate a preference for a visual location for upper (green), mid (blue) and lower (red) visual field, and, by using the phase reversal of visual location preference from upper to lower visual field, the boundaries of V2 can be identified. The inset of panel C shows the same data projected onto a flattened occipital patch of cortex. Displaying the data in this way enables both dorsal and ventral V2 to be visualised more clearly. Panels A and B, adapted from Balaram et al. (2014)

1.4.1 The V2 stripe system

V2 consists of nine layers (1, 2, 3A, 3B, 4, 5A, 5B, 6A and 6B) (Balaram et al. (2014)). The larger receptive fields found within V2 result in the convergence of many V1 neurons onto a single V2 neuron. As a result of the pooling of ocular inputs, the neuronal response properties of V2 are predominantly binocular (Roe (2003)).

The network of projections between V1 and V2 are revealed when injecting V1 with an anterograde tracer. The strongest projections from V1 are found to terminate within layer 4 of V2, with extensions continuing vertically through layers 2 and 3 (Horton and Adams (2005); Tootell et al. (1983)). These connections are thought to arise from at least three afferent V1 pathways. The afferent neurons converge in V2 into particular functional groups, forming bands of neurons with demonstrable receptive field properties responsive to colour, object borders and

binocular disparity/motion (Fig 1.7) (Hubel and Livingstone (1987); Shipp et al. (2009); Sincich and Horton (2005); Sincich et al. (2010); Thomas et al. (2002); Briggs and Usrey (2007); Ts'o et al. (2001)). The innervation provided to V2 from V1 is believed to remain predominantly separate, although there is evidence for short range horizontal connections between clusters, particularly within different cortical layers. The arrangement of these clusters form a series of stripes which are orientated orthogonal to the V1 border and encircle V1 as a corona.




Figure 1.7: A composite of macaque V1 (dots) and V2 (stripes), which have been stained for the metabolic enzyme cytochrome oxidase. Examples of thick, thin and pale stripes are labelled. In situ the stripes are organised in repeating patterns of thick-pale-thin-pale. Adapted from Sincich and Horton (2005)

The parallel system of stripes consists of three types; thick, thin and pale, organised: pale-thin-pale-thick. This arrangement results in twice as many pale stripes as thick or thin. In humans V2 stripes are between 1-3mm in width, and can run for 1.5-3cm across the cortical sheet (Hockfield et al. (1990)).

It has been suggested that the V2 stripe system encodes information regarding motion/disparity (thick), colour (thin) and form (pale), forming separate and parallel feature based processing streams. While functional specialisation for colour, form and motion/disparity has been recorded in animal investigations of stripe type using electrophysiological and optical imaging techniques, (DeYoe and Van Essen (1985); Felleman and Van Essen (1987); Gegenfurtner et al. (1996); Livingstone and Hubel (1987); Shipp and Zeki (2002); Tootell and Hamilton (1989);

Roe and Ts'o (1997, 1995)) no such finding has yet been shown for all three stripe types in the case of humans despite histological evidence for stripe presence (Hockfield et al. (1990)).

The tripartite system theory treats V2 as a parallel processing unit, segregating afferent inputs from V1 and feeding them forward to higher visual modules for feature interrogation. However, strong feedback connections from V4 and V5 have been recorded (Shipp et al. (2009)), leading to the suggestion that V2 could also play a role in the process of feature binding. The mechanism and pathways of feature binding has yet to be identified, and strong feedback networks in V2 could implicate the region in attentional regulation and complex feature perception (Shipp et al. (2009)).

Despite these numerous investigations no conclusive role for the V2 stripes in visual perception has been determined. The difficulty in reaching consensus of the role of the stripes in V2 is likely due in part to their variable expression in different species. In squirrel monkeys almost all examined cortices contain clear V2 stripes, while in the macaque, their expression is more variable. In addition, the position of V2, deep within the lunate sulcus makes electrophysiological recording difficult (Adams and Horton (2003); Horton and Adams (2005)), and the unfolding and flattening of the cortical surface more prone to distortions and damage.

1.4.2 Investigating V2 stripes *in-vivo*

Research of the V2 stripe system in humans has been limited, with only one functional investigation linking stripe organisation and functional specialisation (Nasr et al. (2016)). A detailed discussion of the literature regarding the *in-vivo* investigation of the V2 stripe system is provided within Chapter 5. Here I provide a brief summary.

In their recent investigation, published after the collection of data for this thesis, Nasr et al., utilise ultra-high resolution fMRI to functionally image stripe-like patterns responsive to colour and disparity stimuli within V2 (Nasr et al. (2016)). The group did not attempt to image the stripe subset thought to be responsive to form stimuli. The experimental stimuli used by the

group contained combinations of visual features. Specifically, the stimulus designed to activate colour responsive neurons also contained a motion component, while the disparity stimulus included a form component. The group utilised a subtraction analysis, within which activity to a control stimulus was subtracted from their stimulus of interest, leaving only stimulus specific activity. However, there is evidence to suggest that some stripes respond to combinations of features (Shipp et al. (2009)). In this case, the design may have artificially increased the number of stripes responding to a particular stimulus condition.

The paradigm implemented by Nasr et al., required several scan sessions, with colour and disparity stimulation occurring on different days. I suggest that a protocol with sufficient power to stimulate the entire stripe system within one scan session would provide compelling evidence for functional stripes within V2. Additionally, the use of a form condition would address whether functional specificity for colour, form and disparity/motion stimuli exists within V2. With this as a foundation, the extent to which functional processing is truly segregated within V2, and the role of V2 in the process of feature binding could be further explored.

1.5 Using ultra-high resolution fMRI to investigate the hypercolumn

Research in the field of non-invasive *in-vivo* investigations in humans has been considerably advanced by the introduction of functional magnetic resonance imaging (fMRI) (Ogawa et al. (1990)). The repeatable, and non-invasive nature of fMRI makes it a useful investigatory tool with which to probe functional organisation in the human brain.

fMRI relies upon the assumption that functional activity in the brain can be coupled to a locus of high metabolic activity. Utilising the differing magnetic susceptibility of oxygenated, and deoxygenated blood as an endogenous marker, the neurometabolic response to a stimulus can be measured. The resulting statistical maps can be overlaid onto a structural image of the brain for the purposes of localisation.

With current MRI technology it is possible to acquire functional data with a volume of between 3-6mm³ in resolution. While this is sufficient to investigate many features of the cortex, in order to segregate small-scale functional organisation such as the ODC of V1 and the V2 stripes from other local activity, data must be acquired at ultra-high resolution. The width of ODC in humans has been recorded on average between 0.8-1mm (Adams et al. (2007)), while the V2 stripes vary between 1-3mm in width (Adams et al. (2007); Hockfield et al. (1990)). In the case of the smallest functional target, voxels would be required to have a maximum width of 0.5mm to improve the chances of separating neighbouring signals.

Despite continued advancements in the field, acquiring functional data at this resolution is at the cutting edge of MR technology. Signal is key to fMRI investigations, and the linear relationship between voxel size and the signal to noise ratio (SNR) is a major limiting factor for high spatial resolution.

At high magnetic field strength, the issue of signal is partially addressed. High field normally denotes a magnet with field strength greater than 4 Tesla. In Oxford, a 7T MRI facility is available, and the signal from the increased magnetic field strength is two times that of 3T. This figure is further increased by the suppression of large vessel signals, increasing the relative signal from the small vessels, which are more proximal to the sites of neuronal activity. Considering these facts working at 7T provides a supra-linear increase in signal compared to only a linear increase in noise. At 7T therefore, functional contrast is greater and more specific than at 3T (Olman and Yacoub (2011); Lutti et al. (2013); van der Zwaag et al. (2009); Yacoub et al. (2001)). These features of high field imaging allow the size of voxels to be reduced while SNR is maintained at a reasonable level.

To investigate the role of small-scale functional networks in visual processing it is clear that there is a need for a reliable and robust ultra-high resolution imaging method. Developing an ultra-high resolution sequence applicable to any brain area, with sufficient signal to allow data to be acquired within a reasonable time frame poses a significant challenge. Chapter 3 discusses in detail key considerations in the development of an ultra-high resolution fMRI protocol.

1.6 Aims and achievements

1.6.1 Chapter 3: Advanced methods

In this chapter I aimed to develop and validate a reliable and robust ultra-high resolution imaging sequence. The paradigm was required to fulfil the following criteria

- **Resolution** The fMRI sequence should acquire data at sufficiently high resolution to differentiate adjacent functional activity patterns.
- **Coverage** The sequence must provide good cortical coverage, sufficient to image the region of interest.
- **SNR** The data acquired must have sufficiently high SNR to achieve good data quality within a time frame tolerable by non-expert subjects.

13 ultra-high resolution fMRI protocols were designed, of which 6 were tested. Two different types of image acquisition techniques, the 2D and 3D gradient echo planar imaging were investigated. Data were assessed both for visual quality and for percentage signal change. It was determined that 3D sequences, while more vulnerable to artifact, offered more opportunity to apply advanced imaging techniques such as acceleration, and therefore to acquire data within a reasonable time frame. One 0.7mm isotropic 3D EPI sequence fulfilled the above criteria and was selected. This protocol was used in the acquisition of data for Chapter 4 and was subsequently updated, reducing the field of view and TR. This development improved the temporal resolution and reduced the difference between the effective and nominal resolution. The data of Chapter 5 were collected with this updated sequence.

1.6.2 Chapter 4: Ocular dominance columns of V1

In this chapter I aimed to investigate the monocular architecture of V1 using ultra-high resolution fMRI at high field. The sequence developed in Chapter 3 was used for this investigation. I designed and implemented a monocular display for the scanner environment, enabling monoc-

ular stimulation. Three subjects were recruited and viewed a monocular high contrast, flashing checkerboard in a block design consisting of fixation-left eye-right eye stimulation over a period of 30 minutes. Subjects took part in two or three functional scans in order to assess the reliability between scan sessions.

The aims for the chapter were as follows:

- To establish that robust activity within V1 was achievable using a monocular stimulation paradigm.
- To determine whether periods of monocular stimulation differentially activated populations of voxels within V1.
- To determine whether monocular activity was repeatable across multiple sessions.
- To determine whether any monocular activity patterns could be considered as existing as part of a larger functional network, as would be predicted by functional neuroanatomical evidence (Adams et al. (2007)).

I was able to show robust but variable activity within V1 for all subjects to a simple checkerboard versus fixation stimulation. Voxels demonstrating consistent ocular preference were identifiable, but these patterns were sparse and lacked repeatability across sessions. The columnar structure of ocular dominance columns was not revealed in this investigation. I suggest that this was as a result of insufficient spatial resolution. To address this the fMRI sequence was updated, improving both the temporal resolution and effective spatial resolution of data acquired. This updated protocol was then used in the examination of the V2 stripe system.

1.6.3 Chapter 5: The V2 Stripe System

In this chapter I investigated the segregation of colour, form and motion within the V2 stripe system. The V2 stripe system, being larger (1-3mm in width) than ODC, was a good test case for the updated fMRI protocol. Eight subjects were scanned in 2 or 3 functional scan sessions.

The aims for the chapter were as follows:

- To establish that the cortical activity in response to the visual stimulus could be robustly detected in V2.
- To establish that activity in response to colour, form and motion was detectable in modules known to demonstrate preference for these stimuli.
- To quantify the efficacy of each stimulus in driving the cortical response of V2.
- To determine the extent to which activity patterns in response to colour, form and motion stimulation overlap within V2.
- To determine the extent to which activity in response to stimulus specific stimulation was repeatable over multiple sessions.
- To create a spatial profile of any observed activity patterns using the following measures: (i) the width of clusters with stimulus preference, (ii) the periodicity of clusters with stimulus preference.

In all cases activity within V2 was robust, and there were clear repeatable regions of functional preference for colour and motion. These regions were largely separable. In some cases this preference was elongated along the expected stripe axis. The response to form was weaker and in some cases absent within V2. The incompleteness of the pattern made frequency analysis somewhat unreliable, but of those activity patterns available cluster width was found to be on average between 1-3 mm.

1.6.4 Chapter 6: Using multivariate techniques to examine ocular dominance columns and the V2 stripes

In this chapter I applied MVPA techniques to data from both Chapters 4 and 5. I asked whether as a population, activity in response to a stimulus of interest could be identified as distinct from that of a control stimulus. I then asked whether two stimuli of interest were distinguishable in their activation patterns, and whether this was restricted to certain visual areas. Finally, I investigated whether the spatial organisation of reliable stimulus encoding voxels resembled

those patterns identified in neuroanatomical literature (Hubel and Wiesel (1969); Sincich and Horton (2002)).

The aims for each data set were the same:

- To establish whether a linear classifier could discriminate between a control stimulus and fixation.
- To investigate whether the total number of voxels affected classification accuracy in the simple case of control versus fixation.
- To assess whether a classifier could distinguish between stimuli of interest.
- To assess whether activity patterns projected onto the cortical surface demonstrated the spatial patterns indicative of ODC columns and V2 stripes (Adams et al. (2007); Hubel and Livingstone (1987); Hubel and Wiesel (1969); Tootell et al. (1998b)).

For the analysis of the data from both chapters activity in response to a control stimulus was distinguishable from that in response to fixation at above chance levels. This finding was not altered by varying the number of voxels included in the classification analysis. Using a pair-wise classification analysis, left eye and right eye stimulation was significantly distinguishable within V1 but not V2. This was also found to be the case for colour and control, and motion and control stimulation within V2 for 6 of 8 subjects. Form stimulation, which had elicited a weaker response as recorded in Chapter 5, was not classifiable on average from control stimulation at above chance levels. Voxels contributing most strongly to the classification result were projected onto the cortical surface in order to examine their spatial arrangement. For the data of Chapters 4 and 5 no clear structural arrangement that resembled ODC in the case of Chapter 4, or stripes in the case of Chapter 5 was identified. I conclude from this investigation that activity in response to the stimuli of interest were classifiable significantly above chance levels. This indicates that these stimuli elicited sufficiently unique activation patterns (excluding form within V2) to be considered distinct.

1.6.5 Summary

The human visual system is known to contain small-scale functionally specific repeating units. In this thesis I sought to investigate the activity of two particular units, the ODCs of V1, and the stripe system of V2. At the time of data collection for this thesis the only examples of ultra-high resolution imaging of small-scale functional organisation required long scan times, expert volunteers and restricted fields of view. I aimed to overcome these limitations, and, in doing so sought to develop a robust methodological technique with which to probe the units' functional properties. The long term objective of this work was to establish a role for these functional units in the wider network of visual processing.

Chapter 2

General methods

2.1 Abstract

This chapter provides detail of experimental methods common to Chapters 3, 4 and 5. Firstly, I describe the process of projecting, programming and calibrating visual stimuli in the scanner environment. I then provide information regarding the fundamentals of the major methodological technique of this thesis, magnetic resonance imaging (MRI). I demonstrate the basis of the blood-oxygenation-level dependent (BOLD) signal and basic statistical analysis. Finally, the technique for delineating visual borders of the cortical surface, and converting these surface regions of interest into a volume is shown. These 3D regions of interest were used to isolate activity within visual areas of the brain.

2.2 Experimental sessions

2.2.1 Participant recruitment

Participants of the study were recruited via direct response to posters placed in the Department of Physiology Anatomy and Genetics at the University of Oxford, and also by word of mouth. Those selected met the required 7T MRI safety criteria by having no irremovable metal in or about their person, and had no history of epilepsy or claustrophobia. Additionally, they had the required fluent command of English, as well as normal or corrected to normal vision (based on the participants optical prescription). Before booking the scan sessions, each participant was provided with an information sheet and consent form (Appendix, items 1 and 2), and immediately prior to the start of the scan session there was opportunity to discuss the project and raise and concerns. They were then shown the experimental stimulus, were briefed on any tasks and informed of the session duration. Stress was placed on the need to remain still and maintain fixation at all times, and participants were reminded of this throughout. Prior to commencing the scan, they were also informed that they would be free to leave at any point. Each participant was reimbursed £25 per scan session.

2.2.2 Visual presentation

A 7 Tesla MRI scanner was used to collect experimental data for the work detailed in Chapters 3, 4 and 5. Retinotopic data were collected using a 3 Tesla MRI scanner for Chapter 4, and in Chapter 5 a 7 Tesla MRI scanner was used. The change in scanner type was necessary due to the decommissioning of the 3T scanner.

2.2.3 Visual projection in the 7T scanner environment

Stimuli were back projected using an Eiki LC-XL100 projector with custom throw lens onto an acrylic screen mounted on the patient bed. A chromatic calibration was carried out using a PR650 photometer. Mean luminance of projected white, black, mid-grey, red, green and blue light was 1140cd/m², 2.48cd/m², 580cd/m², 162cd/m², 898cd/m² and 79.7cd/m² respectively. These measurements were utilised by the inbuilt calibration toolbox of PsychoPy, a stimulus presentation software package, in order to modulate the stimulus delivery of Chapter 5. In this experimental paradigm, chromatic rather than luminance modulation was key. The head coil arrangement within the scanner resulted in a restricted and curved visual field of view (Fig 2.1), which, at its highest and widest points measured 12° by 23°.




Figure 2.1: A pictorial representation of the field of view visible to a participant lying within the scanner. The participant lies in the supine position on the scanner bed and views the screen directly above their head. The screen is reflected in a set of mirrors attached to the head coil. The total distance to the screen measured 60cm, and the visible area of the screen measured 12° by 23° .

2.2.4 Visual projection in the 3T scanner environment

Stimuli were back projected onto a screen placed behind the scanner, and were viewed by the participant through a small rectangular mirror placed above the head. The luminance of projected white light and mid-grey was measured using a PR650 photometer and was 356 cd/m^2 and 49.5 cd/m^2 respectively. Gamma values were not recorded, and so projected luminance is reported pre-correction. The screen display area measured 52cm by 32.5cm, with a viewing distance of 205cm. Taking into account restrictions due to the head coil this gave a maximal display of 9.06° by 14.46° by degrees.

2.3 Functional MRI and the origin of the BOLD signal

Since its inception in 1990, fMRI has attracted interest within the field of neuroscience for its potential to non-invasively investigate the temporal dynamics of neural networks. Research by

Seiji Ogawa and colleagues in rat cortex was key in the development of this method (Ogawa et al. (1990)). The most common method of fMRI, among others, is blood oxygenation level dependent (BOLD) imaging. This thesis makes use of BOLD fMRI as a repeatable, non-invasive, indirect measure of cortical activity. The appeal of fMRI over other methods, such as electro-encephalogram (EEG) or magneto-encephalogram (MEG) is the potential to achieve high spatial resolution. The temporal resolution, although poorer than EEG or MEG, is sufficient for the experimental paradigms of this thesis.

2.3.1 The origin of the BOLD signal

BOLD imaging draws a link between activity levels, energy consumption and the requirement of the cortex for oxygen (Attwell and Iadecola (2002); Logothetis et al. (2001); Ogawa et al. (1990); Logothetis (2002)). The human brain consumes 20% of total oxygen input, using 3.5ml per 100g of brain tissue each minute when at rest. Oxygen consumption (CMRO₂) in grey matter is four times greater than that of white matter. The BOLD signal therefore, is principally a feature of grey matter (Logothetis (2002)). When a brain area becomes active, the consumption of oxygen initially outstrips the supply, and the level of deoxygenated blood in a region increases (2.2). Deoxygenated blood is paramagnetic, meaning that it disrupts any magnetic resonance (MR) field applied across it. This causes a decrease in MR signal typically known as the initial dip, which lasts between 2-4 seconds (Duong et al. (2003); Hu and Yacoub (2012); Logothetis (2003)). This period is generally too short to be imaged using fMRI techniques, and so is not modelled. Within a short time frame the cerebral blood flow (CBF) increases above the level required to replace the oxygen used for glucose metabolism, of which 90% is aerobic (Logothetis and Pfeuffer (2004)). The increase in oxygenated blood decreases the fraction of deoxygenated blood present in the vessels and so the MR signal increases. This increase reaches a plateau between 6 and 12 seconds before returning to baseline at a similar rate, with a small undershoot at around 14 seconds (Fig 2.3) (Buxton et al. (2004); Logothetis (2003, 2002)).




Figure 2.2: A pictorial representation of the change in blood flow in response to stimulation. Box 1 depicts blood flow at rest. Boxes 2 and 3 show the change in blood flow in response to increased activity in surrounding tissue. In box 2, oxygen requirements in the surrounding tissue increase. Deoxyhaemoglobin (dHb) concentration increases as a result. Box 3 shows the increase in cerebral blood flow (CBF) and subsequent increase in the concentration of oxygenated blood (HbO_2). This results in a rise in MR signal.

The waveform produced by this BOLD response is termed the haemodynamic response function (HRF) (Fig2.3). The HRF can be convolved with stimulus onset times in order to predict BOLD response, and therefore the predicted activity patterns. Statistical analysis of areas with activity patterns similar to the predicted response allows researchers to investigate by proxy, both the temporal and spatial patterns of brain activity.




Figure 2.3: The HRF. Following activation there is initially a drop in BOLD signal (initial dip). This is due to the rate of consumption of O_2 being larger than that of replenishment. Because of the temporal sampling rate of fMRI the initial dip is rarely resolvable (**M1**). After 2-4 seconds the CBF increases (**M2**), supplying more HbO_2 . This produces an increase in signal, which after reaching a plateau between 6-12 seconds (**M3**), returns to baseline, with a small undershoot after approximately 14 seconds (**M4**).

Theoretical estimates suggest that within primate cortex up to 14% of cortical energy is expended by action potentials. However, post-synaptic potentials are the largest consumer, using 74% of cortical energy (Attwell and Iadecola (2002)). This requirement is so great due to the large number of synapses per neuron, and the high metabolic cost of restabilising ionic concentration gradients via ion channel pumps across the synapse (Arthurs and Boniface (2002)). An experiment seeking to investigate the neural underpinnings of the BOLD signal carried out simultaneous electrophysiological recordings and fMRI in non-human primate cortex. The groups findings suggested that post-synaptic potentials, known as local field potential (LFPs) were more closely and reliably related to the BOLD signal than were multiple unit activity (MUA), arising as a result of action potentials (Arthurs and Boniface (2002); Ekstrom (2010); Logothetis and Pfeuffer (2004); Logothetis (2002, 2003)). These results suggest that BOLD signal represents a summation of local neuronal activity rather than firing rate (Attwell and Iadecola (2002)).

2.3.2 Spatial specificity of the BOLD response

Spatial specificity of the signal is key to resolving small-scale functional structures with BOLD imaging techniques. Within a single cortical column of around 0.8mm in width that extends from the white matter to the surface there can be 10^5 neurons present (Arthurs and Boniface (2002)). The BOLD response measured within one voxel is the sum of the haemodynamic response of all of the responsive neurons contained within the voxel. These columns are supplied by a fine network of capillaries, which are in turn supplied with oxygenated blood from larger vessels. Gradient echo (GE) sequences, which are a common BOLD acquisition method, and are utilised in this thesis, are sensitive to signal from both small and large vessels. When looking to resolve activity from within a single column it is the small vessels ($<8\mu\text{m}$) that are closest to the column structure. The large vessels ($>8\mu\text{mm}$), which supply and remove blood over wider areas of cortex can smooth these activity patterns by making the patterns appear larger and more distal (Logothetis (2002); Weisskoff et al. (1994)). One particular advantage of working at higher magnetic field strength is that the relative signal contribution from the vasculature favours the small vessels. This is due to the linear increase in signal arising from the large vessels, compared to the supralinear increase in signal from the small vessels, when moving to higher MR field strength. As a result, the recorded signal is more specific to the small vessels, and the underlying functional architecture (Logothetis (2008)).

2.3.3 MRI pre-processing procedure

Prior to any statistical analysis, all data underwent a standard pre-processing procedure. A loose brain extraction filter removed non-brain tissue, such as the neck and skull. The data were then motion corrected using FSL MCFLIRT. MCFLIRT takes the middle time series volume as a reference point and calculates a linear identity transformation between this volume and adjacent volumes. Initially a coarse 8mm search is carried out for motion parameters, and is followed by two 4mm searches with tighter tolerances (Jenkinson et al. (2002)). A high-pass temporal filter set at 90 seconds removed slow drift resulting from physiological noise. At this point any

smoothing parameters were applied before the data were submitted for statistical modelling.

2.3.4 General linear model analysis

This thesis makes use of the general linear model (GLM) for fMRI statistical analysis. The GLM models the response of each voxel as a linear combination of one or more predictors and calculates an estimate of the residual error.

The equation for this is expressed as:

$$y(t) = a \times x(t) + b + e(t), \quad (2.1)$$

In this equation $y(t)$ is the data observed from a single voxel. a is the parameter estimate of $x(t)$, which describes the model, or design matrix and contains the explanatory information such as timing onsets. The variable b is a constant, typically the rest condition and $e(t)$ describes the noise or error term, which the GLM assumes is Gaussian.

When the GLM is implemented using the FSL toolbox for first level analysis, i.e. a single session, the data are first demeaned, and as a result the baseline condition is not modelled. When conducting higher-level analysis, i.e. combining multiple sessions, a term is inserted into the model to account for the mean. The GLM creates a model based upon the timing information provided, which is applied to each voxel independently, and so is a form of univariate statistical analysis. Once a model has been determined, a goodness of fit analysis is carried out in order to assess how well the proposed model fits the data. If this model only includes one contrast e.g. on versus off, then t-tests are used to calculate the goodness of fit. If however, multiple contrasts are used such as left and right and off, then the fit is measured using F-contrasts. The goodness of fit is described by the parameter estimates (PE). To convert between a PE and a t-statistic the PE is divided by its standard deviation. If the t statistic falls below the required p significance on a standard distribution then the null hypothesis is unlikely.

When performing multiple t-tests, the likelihood of a false positive is increased. To account for this, corrections for multiple comparisons are often used. Corrections for multiple comparisons are often made using a cluster-based threshold. This method estimates the number of statistically independent clusters in the activity map. Clusters are defined as contiguous voxels, the voxel-wise statistics of which are above a chosen threshold. This method is less conservative than a Bonferroni correction, which uses a voxel based correction method. In an ultra-high resolution data set, voxel based correction methods are likely to be overly conservative due to the very large number of voxels. In Chapter 5 a minimum cluster size threshold of 10 voxels was set to compensate for this. In Chapter 4, this method was not applicable as patterns of interest were likely to consist of only one or two voxels, and so no correction for multiple comparisons was applied. Instead, significance was assessed by the z-statistic and repeatability across sessions. To calculate a z-statistic, the population mean is subtracted from the observed value and this result is then divided by the population standard deviation. The p-value is the probability that the null-hypothesis has been incorrectly rejected. The resulting statistical maps were visualised by projecting the data onto the underlying anatomy as a colour coded map, with a relevant significance value for each voxel.

2.4 Retinotopic methods

In the primate brain, visual cortex contains multiple topographic maps of the retina (Dougherty et al. (2003); Saygin and Sereno (2008); Sereno et al. (1995, 1994)). These maps, which are duplicated across each cortical hemisphere, predominantly contain information regarding contralateral visual space. The measurement procedure of these maps is termed retinotopy. Each visual area possesses its own retinotopic map, containing the upper, middle and lower visual fields. The boundary points between these visual areas coincide with either the vertical or horizontal visual meridian. Neighbouring areas contain mirror-reversed maps of visual space, and with an appropriate visual stimulus paradigm these phase changes can be utilised to reliably delineate visual areas from one another.

2.4.1 Phase encoded techniques

Employing retinotopic mapping techniques, in order to identify specific visual areas, is preferable to using standard space coordinates. The human visual cortex is highly convoluted, and displays considerable inter-subject variability in folding patterns. Estimates of variability suggest that the early visual cortex can vary 2-3 times in size between subjects (Andrews et al. (1997); Toro et al. (2008)). Although stereotactic coordinates systems are based upon the averaging of large number of subjects data, and will probabilistically overlap with the retinotopic maps, individual subject retinotopic maps are specific, reliable and not difficult to obtain.

This thesis utilises phase-encoded retinotopic mapping (DeYoe et al. (1996); Engel et al. (1997); Sereno et al. (1994)). This technique is based on the assumption that as a stimulus moves across visual space, neurons with receptive fields in this region will be activated. The periodicity of a visual stimulus such as a rotating wedge is therefore mirrored in the activation waveform of the underlying neurons. To determine the position of the preferred retinotopic location, the phase shift of the fundamental frequency is calculated with respect to the waveform of the stimulus at zero degrees. In order to achieve this, the fast Fourier transform of the time series of each voxel is calculated. This calculates the correlation between the response waveform with Sine and Cosine waveforms of different frequencies. The resulting vector of complex numbers contains a real and imaginary component describing this fit:

$$phase = atan2(Y, X), \quad (2.2)$$

Where Y and X are the real and imaginary components of the fast Fourier transform respectively. To account for haemodynamic lag it is typical to reverse the direction of the stimulus for some runs. The imaginary component, Y is then multiplied by minus one. The final estimate of the phase is the average of the stimulus moving clockwise and anticlockwise.

2.4.2 Experimental design

Within this thesis two retinotopic mapping techniques were used. The first, using a 3 Tesla MRI scanner employed a standard black and white flashing checkerboard in a wedge and ring design (Fig 2.4). The second set of retinotopic maps, acquired for the investigation of V2 stripes used a 7 Tesla MRI scanner. The curved field of view in the 7T scanning environment limited the area over which the stimulus was visible in upper and lower visual field. In addition to the standard checkerboard wedge and ring design, temporal and chromatic modulation was included (Fig 2.4). Stimulus parameters were varied in order to optimise activation across the higher visual areas of V4 and V5. In each experimental session participants were requested to maintain a centrally fixated gaze, but to attend to the stimulus as it moved through the visual field.

- Stimulus for Stimulus Design 1

The presentation rate of 4Hz was chosen to be closely in line with the optimal temporal frequency for driving visual cortical neurons (Foster et al. (1985); Sereno et al. (1995)). **Wedge:** (Fig 2.4. Polar angle) A rotating 45° wedge consisting of a high contrast black and white checkerboard was used to map polar angle. The checkerboard flashed at 4Hz and consisted of 33 segments, each scaled for eccentricity. The wedge moved through 24 positions around a central fixation dot of 0.5° with a period of 36 seconds. Rotation always began from the position of 3 o'clock. **Ring:** (Fig 2.4. Eccentricity) An expanding or contracting ring was used to map eccentricity. The ring contained a high-contrast checkerboard flashing at 4Hz and was presented around a central fixation point of 0.5°. The ring moved through 24 positions with a period of 36 seconds.




Figure 2.4: Two wedge and ring designs were used for retinotopic mapping of polar angle and eccentricity. Polar angle was measured using a black and white rotating wedge (A1) or red and green or blue and yellow in (A3). Eccentricity was measured using an expanding or contracting ring of the same colour pairs as for polar angle (C1 and C3). The cortical preference to a particular phase position is mapped using a colour wheel. Examples of a polar angle map and eccentricity map are shown in (B), where the visual border of V1 is marked by a dotted white line.

• Modifications for retinotopy design two

The stimulus configuration used for the second retinotopy experiment was adapted from the Stanford VistaSoft toolbox (Amano et al. (2009)). **Wedge:** (Fig 2.4. Polar angle) A rotating 45° wedge consisting of a high contrast coloured checkerboard was used to map polar angle. The wedge moved through 24 positions around a central fixation point of 0.5° over a period of 36 seconds. The wedge contained 33 segments that were scaled for eccentricity, and moved with a temporal frequency of 2Hz. The wedge contained three vertical spokes, and the direction of motion along the longest axis of each spoke changed randomly every 2-3 seconds. Rotation always began from the position of 3 o'clock. The wedge contained either red and green or blue and yellow checks, which were alternated for each functional run. The inclusion of chromatic information in pairs of opponents was designed to maximally activate colour pathways across

the cortex (Gegenfurtner and Sharpe (2001); Li et al. (2014); Shapley and Hawken (2011)).

Ring: (Fig 2.4. Eccentricity) Eccentricity was mapped using an expanding or contracting ring containing a high contrast coloured checkerboard. Motion and colour parameters were as described above for the wedge stimulus.

Stimulus duration Each cycle lasted 36 seconds and each run consisted of 6 cycles. Each session contained between 7 and 9 functional runs, of which two were always ring stimulus and the remainder were wedge. Each alternate stimulus run was reversed in direction in order to compensate for haemodynamic lag.

2.4.3 Data analysis

The middle functional volume of each session was registered to a high-resolution structural scan using a linear transformation with 6 degrees of freedom. To ensure a good registration, further refinement was conducted manually. The data were processed using the FreeSurfer toolbox. FreeSurfer calculates the Fourier transform of the fMRI time series for each voxel and creates a vector of phase and amplitude corresponding to the real and imaginary components. Significance is calculated by dividing the squared amplitude of the signal at the fundamental frequency (signal) by the sum of squared amplitudes of all other frequencies (noise). The corresponding p values are calculated from this ratio and returned by FreeSurfer at $-\log_{10}(p)$. Analysis was restricted in this case to voxels displaying activity at the stimulus frequency at a threshold greater than 4 ($p < 0.0001$) for all subjects. The statistical phase-encoded retinotopic maps were projected onto the cortical surface using the transformation matrix created by the registration of the middle functional volume to the high resolution structural scan. In order to visualise data located in the sulci, the data were displayed on the inflated and flattened cortical sheet, as described in general methods surface reconstruction.

2.4.4 Retinotopic mapping protocol

Retinotopic data were collected in coronally oriented slices, acquired using a 2D EPI sequence (TR-2000ms, TE-30ms, 2mm isotropic resolution).

2.4.5 Identifying visual areas: retinotopic maps from design one

The results shown in Figures 2.5 and 2.6 show the predicted pattern of phase boundaries. V1 is situated on the banks of the calcarine sulcus, and the cut, placed along the fundus of the calcarine sulcus splits the visual field into dorsal and ventral components. The upper visual field is situated on the ventral surface and the lower visual field on the dorsal surface. Clear visual field preference is observable for all subjects, and the mid-point of the upper or lower visual field preference was marked as the boundary between V1 and V2. The maps produced in response to a radially moving ring show a robust representation of eccentricity in all participants. A gap is observable in the retinotopic map of Subject 2 on the ventral bank, although the dorsal bank representation is clear. The region of cortex falling between the representation of upper and lower field demonstrated, on average, no observable visual field preference. This area encodes the foveal region of visual space, and stimulation at the point of fixation remained constant. It was predicted that no activity would be observed in this region due to the lack of phase-modulation.




Figure 2.5: Retinotopic data collected on the 3T MRI scanner is presented for polar angle in three subjects on flattened cortical surfaces. Subject ID's remain consistent across chapters. Voxels displaying activity in response to a wedge shaped flashing checkerboard presented in 24 locations around a central dot. Lower, mid and upper visual field are coloured red, blue and green respectively. The upper border of V1 is marked by the phase transition of upper and lower visual fields back to the horizontal meridian




Figure 2.6: Retinotopic data acquired using the 3T MRI scanner is presented for eccentricity in three subjects on flattened cortical surfaces. Subject IDs remain consistent across chapters. Foveal regions responding to central visual field are depicted as dark purple. Radial to this, voxels demonstrate preference to more eccentric positions of the flashing checkerboard. The stimulus was presented around a central dot in 24 locations in the form of a ring. The upper border of V1 is marked by the phase transition found using polar angle mapping

2.4.6 Identifying visual areas: retinotopic maps from design two

Eight subjects took part in retinotopic mapping for the project investigating the V2 stripe system (Fig2.7 and 2.8). Maps of eccentricity were not obtained for Subject 3. This was due to on-going hardware issues with the scanner at this time. The subject was unable to attend any further sessions. V1 was delineated as described above, and the confluence between V1 and V2 can be separated at the boundary between the upper and lower vertical meridians. The termination of V2 and the beginning of V3 is marked by the horizontal meridian. V1-V3 can be divided into dorsal and ventral portions situated either side of the calcarine sulcus. Moving anterior across dorsal V2 and V3 the phase reversal is marked by the transition between lower and horizontal meridians. In contrast moving anterior across ventral V2 and V3 the transition is between upper and horizontal meridians. V4 can be identified on the ventral surface, anterior to ventral V3 and contains a full representation of visual space, moving posterior to anterior from upper to lower visual field. V5 is part of a larger cluster of motion selective modules that make up the hMT+ complex in humans. V5 contains a complete retinotopic map and is relatively isolated from other regions of retinotopic cortex, positioned on the flattened cortex dorsal and anterior to V4 (Amano et al. (2009); Kolster et al. (2010)). Utilising these criteria, boundaries for retinotopic V1, V2, V3, V4 and V5 were identified. In most cases the foveal gap described above is preserved, indicating good fixation. The maps produced in response to a radially moving ring show a robust representation of eccentricity in all participants except Subject 2. The eccentricity maps for Subject 2 appear to contain considerable noise and do not resemble the clear maps seen for other subjects. The restriction of the visual field along the vertical axis did not affect the ability to delineate retinotopic V1-V5 from the 7T data set. Area V5 demonstrated more variability in its presentation than other visual areas. Specifically, the position of V5 along the dorsal/ventral cortical axis and the orientation of the retinotopic vertical axis differed across subjects. This inconsistent orientation of upper and lower retinotopic representation across participants was also observed by Huk et al.,. Similarly to Huk et al., the position of V5 consistently fell within a single sulcus (Huk et al. (2002)), and, for all subjects the area was anterior to and distinct from retinotopic V1,V2,V3 and V4. Inconsistencies in

retinotopic representation across subjects and hemispheres have also been observed in electrophysiological recordings of the macaque (Van Essen et al. (1981)). Despite this variability across subjects, Huk et al., found evidence to suggest that within subjects the maps of V5 were consistent. This suggests that for the purpose of this work, in which no cross subject averaging took place, the definition of V5 was sufficient.




Figure 2.7: Retinotopic data is presented for polar angle in Subjects 1-4 in response to a moving wedge containing coloured checks and local motion. Data were collected using the 7T MRI scanner. Subject IDs remain consistent across chapters. The wedge was presented in 24 locations around a central dot. Voxels with preference for stimuli presented within lower, mid and upper visual field are coloured red, blue and green respectively. Retinotopically delineated regions include the dorsal and ventral portions of V1, V2, V3, V4 and V5. Delineations were carried out as described above




Figure 2.8: Retinotopic data is presented for polar angle in Subjects 5-8 in response to a moving wedge containing coloured checks and local motion. Data were collected using the 7T MRI scanner. Subject IDs remain consistent across chapters. The wedge was presented in 24 locations around a central dot. Voxels with preference for stimuli presented within lower, mid and upper visual field are coloured red, blue and green respectively. Retinotopically delineated regions include the dorsal and ventral portions of V1, V2, V3, V4 and V5. Delineations were carried out as described above




Figure 2.9: Retinotopic data is presented for eccentricity in Subjects 1-4. Eccentricity data for Subject 3 was not collected due to reasons outlined above. Data were collected using the 7T MRI scanner. Subject IDs remain consistent across chapters. Foveal regions responding to central visual field are depicted as dark purple. Radial to this, voxels demonstrate preference for more eccentric positions of the moving, coloured checkerboard presented around a central dot in 24 locations in the form of a ring. The borders of retinotopic V1, V2, V3, V4 and V5 are marked by the phase transition found using polar angle mapping response to a moving wedge containing coloured checks and local motion.




Figure 2.10: Retinotopic data is presented for eccentricity in Subjects 5-8. Data were collected using the 7T MRI scanner. Subject IDs remain consistent across chapters. Foveal regions responding to central visual field are depicted as dark purple. Radial to this, voxels demonstrate preference for more eccentric positions of the moving, coloured checkerboard presented around a central dot in 24 locations in the form of a ring. The borders of retinotopic V1, V2, V3, V4 and V5 are marked by the phase transition found using polar angle mapping response to a moving wedge containing coloured checks and local motion.

2.5 Cortical surface based methods

2.5.1 Considerations when using surface based techniques

The presentation and interpretation of fMRI data can be carried out either in the native 3D volume space, or on a 2D surface. Some cortical features are best understood in the context of the 2D cortical sheet rather than the folded 3D volume. Consider the example of the V2 stripes. The stripes run orthogonal to the border of V1 encircling the region as a corona. Cortical surface reconstruction and inflation can allow functional patterns, that would otherwise be obscured by sulcal folds, to be visualised. In this way, the relationship between orientation of activation patterns and other visual areas is clear. If the stripes were spread over a series of slices following cortical folding patterns, as would be the case in the 3D volume, the full architecture would be obscured. Careful consideration should be taken before implementing a surface based technique, as the registration and inflation steps required induce an unavoidable distortion factor. This distortion is caused by the differing curvature factors between the cortical surface and a sphere, termed the “mapmakers problem” (Fischl et al. (1999)). Where possible, quantitative analysis was carried out in the 3D volume in order to avoid additional registration steps, and the effect of distortions.

2.5.2 An overview of the surface based analysis pipeline

In order to examine activity patterns on the cortical surface, an accurate segmentation of the grey-white matter boundary was first carried out. The surface was then unfolded and inflated before a series of cuts were introduced, allowing the surface to relax into a flattened patch. This procedure allows the entire surface to be visible from one view. The FreeSurfer toolbox (Dale et al. (1999); Fischl et al. (1999)) is commonly used for this kind of analysis. The toolbox contains an automated series of steps for surface reconstruction requiring little manual intervention. Some of the major steps undertaken by the FreeSurfer pipeline are described in Figure 2.11.




Figure 2.11: A brief illustration of some of the major FreeSurfer surface reconstruction, inflation and flattening steps. Panel A.1 shows the original T1 weighted structural scan and the segmented grey-white matter boundary. The hemispheres are separated and the ribbon covered with a triangular mesh before being smoothed (A.2), and inflated (A.3, A.4). Panel B illustrates the flattening process, from cut placements (B.5, 1-5), the extraction of the occipital lobe (B.6) and the final relaxed surface as a flattened cortical sheet.

2.5.3 Surface reconstruction

The FreeSurfer toolbox takes as its input a T1 weighted structural image. First intensity variations, due to any magnetic field inhomogeneities in the structural image, are corrected and normalised. Non-brain voxels such as the skull and neck are then removed, and the boundary between the grey and white matter boundary is computed. The output is a segmentation file that contains the grey matter ribbon (Fig 2.11. A.1.). The two hemispheres are then separated at the mid line and the grey matter ribbon is covered with a mesh of triangular tessellation. This

fitted mesh is then deformed to form a smooth representation of the grey-white border with a minimal cost function (Fig2.11 A.2.) (Dale et al. (1999)).

2.5.4 Inflation and flattening

The reconstructed surface from the first stage of analysis undergoes several more processing stages in order to be successfully inflated and flattened. First the geodesic distances are calculated from the original and folded surface. The cortical sheet is then projected onto the target surface and assigned a normal vector field with consistent orientation, before distortions are minimised (Fig2.11 A.3., Fig2.11 A.4.). When creating a flattened surface the same processes are carried out as described above, with the inclusion of a series of cuts to the cortical surface (Fig2.11 B.5-6). The target structure is also changed from a sphere to a flat sheet. When flattening the occipital cortex, as in this thesis, cuts are applied down the fundus of the calcarine sulcus (Fig2.11 B.5.), separating the upper and lower visual field. A series of radial cuts, applied anterior to the calcarine sulcus enable the occipital cortex to be separated from the rest of the brain (Fig2.11 B.6.) and relaxed into a sheet (Fig2.11 B.7.) (Fischl et al. (1999)).

2.5.5 Converting surface labels to volumes of interest

A custom retinotopic label conversion pipeline was designed in order to carry out regional analysis in the 3D volume. The pipeline combined FreeSurfer and FSLmaths toolbox commands. Restricting analyses where possible to the 3D volume avoids the registration and translational distortions associated with the surface based approach.




Figure 2.12: A brief illustration of the label conversion process. In (A) retinotopic data is projected onto the flattened cortical surface. Using the phase boundaries as a guideline, a retinotopic region is delineated (B). This label, defined on the flattened sheet, is then converted into a 3D volume (C), before being registered to the functional data (D).

Label conversion processing stages were carried out as follows:

1. The first retinotopic scan was averaged over time to produce a high quality functional template. This template was subsequently registered using a linear registration with 6 degrees of freedom, to the high resolution structural scan. This registration was refined manually.
2. Retinotopic regions were projected onto the cortical surface (Fig2.12 A) and defined as described previously. The border of each region was marked on the flattened retinotopic map, and then saved as a label (Fig2.12 B).
3. Each label was expanded to fill the grey matter ribbon, becoming a retinotopic volume (Fig2.12 C). The left and right retinotopic volumes were added together to create one volume for each retinotopic region (Fig2.12 C).
4. The registration matrix used in Step 1 was converted from FreeSurfer .dat to FSL .mat format.
5. The structural scan and associated retinotopic volumes were reoriented to the FSL standard coordinate system.

6. The first functional scan acquired during the experimental paradigm was averaged over time to produce a high quality functional template.
7. The transformation matrix created in Step 1 and transformed in Step 4, was inverted. This inverted transformation matrix was used to register the retinotopic volumes to the each Subjects high quality functional template (Fig2.12 **D**).

The above steps enabled the identification of subject specific retinotopically defined visual areas. These volumes were accurately registered to both the high-resolution T1-weighted structural scan, and to the high quality functional template of each subject. From this point, region specific analysis could be conducted.

Chapter 3

Advanced methods

3.1 Abstract

Until recently, it has been beyond the scope of fMRI to acquire data with sub-millimetre resolution. Now however, cutting edge sequence and hardware design have facilitated *in-vivo* investigations of functional columns within human visual cortex. This relatively new field of fMRI imaging relies upon the application of advanced MR techniques to acquire ultra-high resolution data. Each technique has an associated set of costs of varying impact depending upon the study aim or functional target. Taking this into consideration, my investigations required an fMRI protocol capable of resolving functional organisation between 0.8mm and 3mm in width. In this chapter I describe some of the advanced MR techniques considered in the development of this protocol. I also discuss the process of selection of the ultra-high resolution fMRI protocol and its subsequent refinement. Finally, I outline some of the limitations of the selected protocol and discuss their potential impact on the work conducted in Chapters 4 and 5.

3.2 Introduction

Since the development of fMRI by Seiji Ogawa et al., in 1990, the pursuit of increasingly specific functional correlates of neuronal activity has driven MRI hardware and acquisition techniques (Logothetis and Pfeuffer (2004); Ogawa et al. (1990)). As a result, with appropriate hardware and careful sequence and paradigm design, it is now possible to acquire functional data at sub-millimetre resolution. This development offers exciting opportunities for the investigation of small-scale functional modules. Until recently post-mortem staining procedures, tracer techniques, optical imaging and electrophysiological recordings have provided the only means of investigation for small-scale functional activity. The invasive nature of these methods has placed limitations on their applicability for human studies. This restriction of the functional investigations to non-human models such as the mouse, cat and non-human primate provides only inferences for human visual processing. It is clearly advantageous to develop non-invasive, *in-vivo* ultra-high resolution techniques in order to investigate the functional activity of small-scale functional modules in humans. However, the design of sequences capable of acquiring ultra-high resolution data is still at the forefront of fMRI technology, and the best scanning parameters are subject to discussion. The following sections outline the advantages and costs for a range of parameters. The aim of this being to establish an optimal protocol facilitating investigation into functional columns in human visual cortex.

3.2.1 The effect of MR field strength

A key limiting factor in the acquisition of ultra-high resolution fMRI data is the signal to noise ratio (SNR). SNR decreases linearly with voxel volume, making it more difficult to detect relevant signals at high resolution. However, the decrease in SNR resulting from reduced voxel volume can be offset by a linear increase in available signal when moving from 3T to 7T (Olman and Yacoub (2011)). This linear increase, resulting from greater field strength, becomes supralinear when taking into account the additional signal contribution from the microvasculature (Yacoub et al. (2001)). The proportional increase in signal contribution from the microvas-

culature ($<50\mu\text{m}$), which is more local to the sites of neural activity, arises as a result of the shortened relaxation rate of blood ($T2^*$). Because of their diameter the coherence timescale of the large vessels ($>150\mu\text{m}$) is shorter, increasing signal dephasing and subsequently reducing their contribution as a signal component. This suppression of large vessel contribution reduces the spatial displacement of signal due to large draining veins (Polimeni et al. (2010); van der Zwaag et al. (2009); Yacoub et al. (2001)).

The proportionally greater contribution from grey matter results in BOLD contrast that is greater and more specific to neural activity at 7T than at 3T (Olman and Yacoub (2011); Lutti et al. (2013); van der Zwaag et al. (2009); Yacoub et al. (2001)). This is expressed as an increase in the contrast to noise ratio (CNR), which describes the image SNR, magnitude of the BOLD contrast and temporal noise (Olman and Yacoub (2011)). In this instance image SNR refers to the mean signal divided by the mean standard deviation, and is calculated on a voxel wise basis. The magnitude of the BOLD contrast is the amplitude of the stimulus specific signal against a baseline condition. Finally, the temporal noise refers to unmodeled signal components remaining after GLM fitting has been carried out. In addition to increases in SNR the noise contribution is also amplified as a result of increased contributions from physiological processes and static field fluctuations. These effects can act as limiting factors for spatial and temporal sensitivity (Jorge et al. (2013); Olman and Yacoub (2011)). This effect can be offset by reducing voxel volume. Triantafyllou et al., showed that by reducing voxel volume at high field, the ratio between thermal and physiological noise contributions could also be reduced and translated into improved SNR (Triantafyllou et al. (2005)).

At high magnetic field strength (7T) available signal is increased beyond the level achievable at lower field strength (3T). However, without careful sequence design the effects of noise and distortion can offset this SNR boost. The following sections describe some of the MR techniques under consideration, with the aim of selecting an optimum combination with which to acquire data at high spatial resolution without sacrificing SNR.

3.2.2 An introduction to 2D and 3D EPI

One of the first considerations when designing an fMRI sequence to acquire data at sub-millimetre resolution is the acquisition method. Three common acquisition methods include Spin-Echo (SE) (Constable et al. (1994)), Gradient And Spin Echo (GRASE) (Oshio and Feinberg (1991)), and Gradient Echo (GE) (Ogawa et al. (1990)). Since the advent of BOLD imaging the high sampling speed and sensitivity to signal change has made the GE EPI a commonly used acquisition technique (Poser et al. (2010)). In this thesis, two types of GE sequence (2D and 3D) are considered (Fig 3.1). In a 2D sequence a slice excitation gradient is first applied across the region of interest. Signal is then acquired using an oscillating read and blipped phase encode gradient (Poser et al. (2010)). The repetition time (TR) is the time taken to acquire a single slice multiplied by the number of slices ($TR_{\text{slice}} * N_{\text{slices}}$). In a 3D GE acquisition a thick slab of tissue is excited, and signal is acquired from the region of interest as a whole rather than in a series of slices. This difference is achieved by changing the way in which data is represented in k-space.

k-space is a matrix in which spatial information regarding the acquired data is temporarily stored. Each point in k-space is a complex number, which, after all data points have been acquired is transformed via the inverse Fourier transform into the MR image. In a 3D sampling scheme, k-space is sampled and transformed in 3D. For each excitation a given k_x - k_y plane of k-space is acquired with a different k_z increment. k_z is advanced through the use of a secondary phase-encode direction, which causes points along the k_z axis to precess at different frequencies and replaces the slice selective gradient used in 2D methods (Jorge et al. (2013); Poser et al. (2010)).




Figure 3.1: A schematic of 2D and 3D acquisition strategies. In a 2D EPI, data are acquired in a series of slices. A 3D EPI acquires a slab of data, replacing the slice selective gradient with a second phase encode gradient, which is used to move through k_z of k -space. No data is encoded into k -space in the 3D sequence until the entire slab has been read out. In 2D sequences each ky - kx plane is acquired independently for each slice.

At high resolution, more data points must be acquired per slice. When using 2D acquisition techniques this results in longer single volume acquisition times. As the acquisition time of each slice increases, slice dephasing, which describes the loss of signal homogeneity across slices becomes more problematic (Jorge et al. (2013)). To compensate it is common for 2D sequences to increase the slice thickness, thereby reducing the number of overall data points collected. This strategy, when employing non-isotropic voxels to image an elongated target functional organisation relies upon careful slice placement (Figure 3.2). Incorrect slice placement in regions of sulcated cortex can result in a loss of image contrast and prevent the full organisation of the feature from being fully reconstructed.




Figure 3.2: Adapted from Cheng et al., (Cheng et al. (2001)). A schematic of two slices through a region of cortex are shown to demonstrate the importance of slice position when imaging a functional organisation with orientated activation. Here, a slice (A) is placed along the lower bank of the calcarine sulcus (CS), such that columns run along the high-resolution imaging plane. In this case the orientation of the structure is captured accurately. In (B) The slice is placed higher and columns appear only as blobs in the resulting image.

In addition to longer read out times, users of the 2D acquisition technique face challenges when attempting to pursue thinner slices. Thinner slices, while benefiting from a reduction in the effect of slice dephasing, suffer more from slice profile effects and a reduction in SNR due to the smaller voxel size (Miller et al. (2006)). Slice profile effects describe the blurring of signal across slices and result from an insufficient number of frequency steps. Frequency steps are required to distinguish between neighbouring points, and as a result, where the number of steps are too low, tissue contrast can be poor. The number of frequency steps is determined by the bandwidth, which can be increased, but is associated with higher signal deposition levels (SAR). The value that the SAR is permitted to reach is limited for participant safety (Poser et al. (2010)). In comparison, 3D sequences, which do not acquire data in steps of slices, experience no such issues; as in a 3D sequence the entire region of interest is stimulated and collected in one volume. However, because no data is read out until the entire volume has been stimulated, 3D sequences are more vulnerable to motion and physiological effects of pulsation.

In a comparison of 2D and 3D acquisition methods using voxel sizes of 1.5mm, 2mm and 3mm isotropic, Poser et al., concluded that 3D sequences offer superior opportunities to achieve high resolution data than 2D (Poser et al. (2010)). 3D EPI sequences are more readily able to achieve small voxel sizes, and at high resolution benefit from increased SNR where noise is driven by

thermal rather than physiological sources (Lutti et al. (2013); van der Zwaag et al. (2009)). When examining the temporal SNR in both 2D and 3D EPI sequences at 1.5mm isotropic at 3T the 3D sequences always outperformed the 2D sequences. Three major limitations of 2D EPI were highlighted when seeking to achieve high resolution data 1) SAR safety limits are quickly reached, 2) The imperfect slice profile of high resolution 2D EPI can reduce the effective resolution and affect SNR, 3) The slice-select gradient amplitude required to achieve thin slices ($<2\text{mm}$) can quickly exceed the system limits of the scanner (Lutti et al. (2013); Poser et al. (2010)).

3.2.3 Acceleration methods

In addition to the choice of acquisition method there is also the opportunity to apply acceleration techniques, which can be used to shorten acquisition time. This not only reduces temporal blur but also the influence of physiological fluctuations, and allows more data points to be collected within a given time frame. When working at 7T acquisition time is particularly important because of instabilities in the B0 field. Instabilities can occur as a result of the electromagnetic interaction between the magnetic field and the person within the scanner (Sled and Pike (2000)). This distorting effect is greater at 7T and k-space must be covered quickly in order to prevent signal loss and blurring (Uğurbil et al. (2003)). A major advantage of 3D acquisition techniques is that a number of the acceleration methods can be applied in two dimensions, speeding data acquisition further. The following section outlines some of the major options for acceleration techniques.

- **Partial Fourier**

k-space, also known as frequency space, contains all of the information acquired during a single TR. The Fourier transform is used to move between frequency space and image space. The time taken to acquire an image is dictated by how many points of k-space need filling. One feature of k-space is its conjugate symmetry (Fig3.3), and by taking advantage of this property, it is possible to acquire a proportion of k-space and fill the rest in via mirroring. The reduction in

the number of data points resulting from the truncation of k -space is associated with a reduction in SNR.




Figure 3.3: Using partial Fourier to fill k -space. k -space is represented here before (1), and after (2) partial Fourier has been applied. k -space is organised such that it contains conjugate symmetry, $a+ib$ in the top right corner is the mirror of the bottom left. Likewise $c+id$ in the top left corner is the mirror of the bottom right. This symmetry can be used to speed acquisition by acquiring only part of k -space (1), while the rest is filled by mirroring the acquired lines onto remaining lines (2).

• Parallel imaging

When k -space is read some lines can be skipped and zero filled. While this can speed the image acquisition, the reduction of data points containing relevant information decreases the image SNR overall. The reduction in sampling time can also lead to errors in k -space reconstruction, which reveal themselves as aliasing (Heidemann et al. (2012)). In 3D sequences the second phase-encode direction allows entire k_x - k_y planes to be skipped. In a 2D sequence an acceleration factor of 2 will result in a time saving of 10-15%. This time saving is improved in a 3D sequence, where acquisition time is reduced linearly as a proportion of the non-acquired data (Poser et al. (2010)). The efficacy of this technique relies upon the accurate filling of the non-acquired k -space lines. However, phase-errors caused by variations in resonant frequency, flow and motion can introduce inaccuracies in the reconstruction process. Because the time taken to traverse k -space is longer when more data-samples are collected, these phase-errors are more likely to accumulate and reduce SNR. It is standard procedure to estimate the SNR

cost as being proportional to the square root of the number of points acquired (Olman and Yacoub (2011)). In this thesis parallel imaging was achieved using Generalized Autocalibrating Partial Parallel Acquisition (GRAPPA), whereby the correction for the data undersampling occurs prior to the Fourier transformation. This is in contrast to SENSitivity Encoding (SENSE) where the correction occurs after the Fourier transformation has occurred.

- **Reduced field of view**

Reducing the number of data points required per TR can reduce overall volume acquisition time. This is equivalent to reducing the number of slices or the volume acquired. Reducing acquisition time reduces inter-voxel dephasing, where a voxel contains spins in different states of relaxation, and as a result experiences internal signal interference. Inter-voxel dephasing is greater at high-field because of the shorter $T2^*$ of blood, and the reconstructed image may therefore contain distortion or drop-out (Heidemann et al. (2012, 2009)).

3.2.4 Summary

While there is no current gold standard sequence for achieving sub-millimetre voxel size, it is clear that there are a number of options available to tailor to the individual case. The aim of this chapter has been to create a protocol suitable for the examination of small-scale functional patterns with elongated orientation i.e. columns and stripes. Following the acquisition of the first data set, outlined in Chapter 4, the protocol was reviewed and subsequently updated. Key in this review was the relationship between nominal and effective resolution. The nominal resolution describes the voxel size that one dials into the scanner, and the effective resolution is the true resolution of the acquired data. Ideally these values would be the same. However, depending upon the choice of sequence parameters, the actual resolution of the data acquired can be reduced. The disparity between effective and nominal resolution can be increased as a consequence of some acceleration techniques, particularly those that blur voxels along the phase-encode direction. I conclude by summarising my findings and present the selected protocol for use in the acquisition of data in the investigation of the V2 stripe system.

3.3 Methods

The work for this chapter was carried out by myself (SF), under the supervision of Holly Bridge (HB) and in close collaboration with Stuart Clare (SC), Karla Miller (KM), Mark Chiew (MC) and Nadine Graedel (NG) at FMRIB. Contributions were as follows: sequence design and programming (NG and MC); testing and validation (SF, HB, NG and MC); discussion and development (SC, KM, MC, NG, SF and HB).

3.3.1 Participants

One participant (female) aged 25, with normal vision. The participant fulfilled the screening criteria for scanning at 7T (Chapter 2 - General methods).

3.3.2 Stimuli

Visual stimuli were programmed using the Matlab based toolbox Psychtoolbox. Visual stimuli were projected as described in Chapter 2 onto a mid grey screen (580cd/mm^2). The stimulus consisted of a 15° by 15° black and white checkerboard, with a white fixation dot of 0.5° presented at the centre of the screen. The contrast of the checkerboard was reversed at a rate of 6Hz. The luminance of black and white was measured to be 1.25cd/m^2 and 1140cd/m^2 respectively. Contrast was high at 0.99. Two different visual conditions were used, each presented in a block of 24 seconds as follows (i) checkerboard: the flashing checkerboard was presented in the centre of the screen along with a fixation point, (ii) fixation only: no checkerboard stimulus, with only the fixation point present on the screen.

3.3.3 Experimental session

The participant was briefed outside the scanner environment (see Chapter 2). The participant was reminded to remain still and maintain fixation on the fixation point at all times during the scan. Head position was stabilised and supported by cushion placement. Protocol development took place across four scan sessions using the same participant. Scans 1-3 lasted 600 seconds, while scans 4-6 lasted each between 171-190 seconds. Within each run, for scans 1-3, 12 blocks of visual stimulation were presented, while scans 4-6 contained 4 blocks. Stimulus order remained constant as follows: fixation-checkerboard-fixation-checkerboard. The 3D sequences (4-6) were tested as part of a larger group of potential protocols with considerably more online updating required, and so were run for a shorter period of time than the 2D sequences (1-3). This accounts for the differing duration of scans sets 1-3 and 4-6.

3.3.4 Protocols tested

13 potential protocols were designed, of which 6 were developed, tested and validated (Table 3.1). Scans 1-3 were all 2D sequences of varying resolution, while scans 4-6 were 3D sequences. The resolution of the 3D sequences was fixed at 0.7mm isotropic and acceleration parameters were varied. The effects of key parameters such as voxel size, TR, Partial Fourier and acceleration were examined for each test. Reasons for protocol dismissal included an unacceptably long TR measure and exceeding scanner limits of signal deposition (SAR). All protocols were tested using a 7T Siemens Magnetom scanner with a Nova Medical 32 channel head coil.

	Scan	Voxel size (mm)	TR(s)	TE(ms)	Slices	OS(%)	PF	Ac
2D	1	0.50 x 0.50 x 3.00	3.00	29	40	-	7/8	4
	2	0.75 x 0.75 x 3.00	2.00	29	33	-	-	4
	3	1.00 x 1.00 x 3.00	3.00	29	40	-	-	3
3D	4	0.70 x 0.70 x 0.70	2.45	31	48	-	6/8	4/2
	5	0.70 x 0.70 x 0.70	4.13	29	48	33.3	6/8	4/1
	6	0.70 x 0.70 x 0.70	5.95	54	48	-	-	4/1

Table 3.1: A summary of trialled scan parameters. Trials for both 2D and 3D scans were carried out for a range of TR, TE, slices, percentage oversampling (OS), Partial Fourier (PF) and Acceleration (Ac). For 3D sequences acceleration was applied to one or both of the phase-encode directions. Where no data is present, this technique was not utilised.

3.3.5 Data preprocessing

Data were analysed using FSL 5.0.1, MATLAB and associated custom scripts combining `fsfmats` and `fsfstats` commands. Data were pre-processed according to the steps outlined in Chapter 2. Briefly, data were corrected for motion using MCFLIRT (Jenkinson et al. (2002)) and were high pass filtered (90 seconds) in order to remove low frequency noise and slow-drift. No spatial smoothing was applied to the data. For each run, one condition of interest was modelled using the GLM implementation in FSL FEAT (i) checkerboard > fixation. To determine the significance of BOLD responses a cluster based threshold of $z > 2.3$ was applied to the statistical maps with a cluster significance of $p < 0.05$. The z-statistic describes the standardised difference between an observed statistic and a null distribution, assuming normal variance. To calculate a z-statistic the population mean is subtracted from the observed value, and this result is then divided by the population standard deviation. The p-value is the probability that the null-hypothesis has been incorrectly rejected. A cluster-based threshold is one possible approach to control for the number of false positives arising from the multiple comparisons of voxels in the brain. Cluster based thresholding takes into account the likely non-independence of signal arising from one voxel from its neighbours. In this method, an initial threshold is applied at

the voxel level and produces groups of supra-threshold voxels. Secondly, an extent-threshold is applied and measures the likelihood that voxels with particular z-stats would be clustered into groups of particular sizes.

Scan quality was assessed qualitatively by eye, checking for distortions, contrast and noise. V1 is known to demonstrate strong activity in response to high contrast stimuli such as checkerboards. The activity in this region was quantified and used as a marker for signal strength. A probabilistic mask of left and right V1, taken from the Juelich Histological Atlas, was registered to each functional scan (Fig.3.4) using a linear transformation with 6 degrees of freedom (dof), a 180° search angle and a nearest neighbour interpolation. A probabilistic threshold of 40 was applied to the registered V1 mask. The registered V1 masks of left and right hemisphere were used to restrict the calculation of percentage BOLD change to within the region of interest using FSLs FeatQuery tool.

3.3.6 Measuring percentage BOLD change

Featquery is implemented in the FSL toolbox and can be used to compute a range of statistics, including percentage BOLD change. Percentage BOLD change can be used to estimate the efficacy of a stimulus in eliciting activity from within a region of the brain. Where a stimulus is known to be robust in its capacity to elicit a response Featquery can also be used to compare signal strength between fMRI acquisition methods. Percentage BOLD change is calculated by scaling the parameter estimate (PE/COPE) (i.e. the scaling factor between an estimated square wave with a height of one and a data fitted square wave), by 100 times the peak-to-peak height of the regressor (EV) (number of contrasts of interest). This value is divided by the mean signal over time of the functional data. The variability of this result is expressed in terms of the standard deviation of the percentage BOLD change using the contrast of parameter estimate. These measurements were made for voxels within a region of interest, which in this case was V1 in the left and right hemispheres.




Figure 3.4: Calculating activity within V1. A sagittal, coronal and horizontal slice of one scan showing the registered V1 mask for both the left (blue) and right (red) hemispheres. Within this mask, activity to the presented stimulus was measured to quantify the sensitivity of the protocol to cortical activity.

3.3.7 Measuring CNR

Contrast to noise provides a quantitative measure of data quality. This can be calculated either as temporal CNR, or as image CNR. In this case image CNR is calculated to provide a measure of tissue contrast and the potential influence of static noise on signal detection (Magnotta et al. (2006)).

The mean signal was calculated from within a cluster of 250 grey matter voxels. These were manually extracted from within the region of interest defined here as V1. A further 250 white matter voxels were extracted from white matter tracts across the same slices used in the collection of grey matter voxels. Finally, 250 voxels containing no brain, aliasing or ringing components were collected. These masks were applied to the registered data to extract identical grey matter, white matter and air regions of interest across all scans.

CNR was calculated as follows:

$$CNR = \frac{Mean_{GM} - Mean_{WM}}{SD_{air}}, \quad (3.1)$$

Where $Mean_{GM}$ and $Mean_{WM}$ refer to the mean grey and white matter signal within the respective regions of interest. SD_{air} refers to the standard deviation of the signal recorded

within the air region of interest.

3.3.8 Protocol review and development

Protocol development occurred in two phases. Firstly, a suitable fMRI protocol was developed and used to investigate ocular dominance columns within V1. Subsequently, calculations were made by NG and MC to suggest that the effective resolution of the protocol could be improved by reducing the field of view. The change in the size of the field of view reduced the TR from 4.128 seconds to 3.60 seconds, improving the temporal sensitivity of the sequence.

3.4 Results - Part 1

The following section presents the data analysis of three 2D EPI sequences and three 3D EPI sequences. Data from each protocol was assessed qualitatively by eye, and quantitatively in terms of percentage BOLD change and CNR. Comments regarding the visual quality of the data are based upon the un-masked images, i.e. no V1 masks have been applied. Measures of percentage BOLD change however, are restricted to the probabilistic masks of V1. The results acquired using 2D EPI sequences are discussed first, followed by a review of the data acquired using the 3D EPI sequences. I highlight key visual features of interest and compare the costs and benefits of each sequence. Finally, I consider possible improvements to the selected protocol with specific relevance to effects of protocol choices on signal blurring.

3.4.1 Results of 2D sequence testing

In this investigation three 2D sequences, differing in spatial resolution, were assessed for their efficacy to detect signal change in response to visual stimulation. The field of view of all sequences was sufficient to cover the visual cortex, which was the region of interest for this investigation. Partial Fourier was applied in the highest resolution sequence in order to keep the TR within the bounds of blood T2* decay. In each case, acquisition was accelerated by a factor of 3 or 4, depending upon the sequence.

Scar	Voxel size (mm)	TR(s)	TE(ms)	Slices	OS(%)	PF	Ac
1	0.50 x 0.50 x 3.00	3.00	29	40	-	7/8	4/1
2	0.75 x 0.75 x 3.00	2.00	29	33	-	-	4/1
3	1.00 x 1.00 x 3.00	3.00	29	40	-	-	3/1

Table 3.2: The parameters for the three 2D sequences are shown here. Spatial resolution is measured in mm. TR, TE, slices, percentage oversampling (OS), Partial Fourier (PF) and Acceleration (Ac) are also recorded. Where no data is present this technique was not utilised.

In Figure 3.5 a sample slice of each of the three 2D EPI sequences is shown. These are accompanied by an overlay of the activity in response to visual stimulation. The area of cortex from which signal is detected visibly increases when moving from the highest ($0.5 \times 0.5 \times 3 \text{ mm}^3$) to the lowest resolution ($1 \times 1 \times 3 \text{ mm}^3$) image set (Fig 3.5). The noise, which can be seen outside of the brain, particularly in the coronal plane as a grey haze, also reduces visibly when using a lower resolution sequence. Tissue contrast is most distinct in Sequence 3 ($1 \times 1 \times 3 \text{ mm}^3$). This observation is reflected by the CNR recorded for the three sequences. The CNR recorded for Sequence 1 is 1.3, a value slightly higher than that of Sequence 2 (1.1). The CNR for Sequence 3 is more than double that of Sequence 1 or 2 (Table 3.3). The visible signal strength differences is not reliably mirrored by the recorded percentage BOLD change within V1. Measured across both hemispheres, activity recorded within V1 during the application of Sequences 1, 2 and 3 demonstrated high variability.




Figure 3.5: 2D sequences Sample slices and time series. Sample sagittal, coronal and horizontal slices from three 2D sequences at varying resolution ($0.5 \times 0.5 \times 3 \text{ mm}^3$, $0.75 \times 0.75 \times 3 \text{ mm}^3$, $1 \times 1 \times 3 \text{ mm}^3$, labelled as 1-3 respectively). The data were cluster corrected ($z=2.3$, $p<0.05$). A heat scale, which saturates when $Z>7$, indicates the significance of activity in response to a flashing checkerboard.

Sequence 1 recorded a 2.2% (+/-2.7%) change in percentage BOLD and activated 7% of V1 voxels, while Sequence 2 detected a 2.1%, (+/- 2.6%) change and reliably activated 20% of V1 voxels. The noise in these two sequences is visible when comparing the coronal view of Sequences 1 or 2 with Sequence 3. In Sequence 1, although the activity clusters are more significant than Sequence 2 (as indicated by the heat scale), the area of activity is visibly more restricted. Sequence 3, during which clusters of highly significant activity were present, demonstrated a percentage BOLD change of 2.7% (+/- 2.2%) and activated 25% of V1 voxels.




Figure 3.6: Percentage BOLD change in VI for a single subject in response to visual stimulation, calculated for three 2D sequences. Three different spatial resolutions were tested. The spatial resolution of Sequences 1, 2 and 3 was $0.5 \times 0.5 \times 3 \text{ mm}^3$, $0.75 \times 0.75 \times 3 \text{ mm}^3$ and $1 \times 1 \times 3 \text{ mm}^3$ respectively. Error bars indicate standard deviation of the BOLD response across the contrast of parameter estimate (COPE)

Sequence	CNR
1	1.3
2	1.1
3	2.8

Table 3.3: CNR measured between grey and white matter for a manually defined region of interest. Shown here for each 2D sequence

3.4.2 Results of 3D sequence testing

3D sequences are a relatively new MR acquisition method. They offer advantages over 2D sequences for their potential to speed data acquisition, but are vulnerable to motion and physiological artifacts. Three 3D sequences, each with the same spatial resolution but varying in MR acceleration parameters, were assessed for their efficacy to detect signal change in response to visual stimulation. The field of view of all sequences was sufficient to cover the visual cortex, which was the region of interest for this investigation. Partial Fourier was applied to two of the three sequences and for one sequence, the effect of acceleration in the second phase-encode direction was tested.

Scar	Voxel size (mm)	TR(s)	TE(ms)	Slices	OS(%)	PF	Ac
4	0.70 x 0.70 x 0.70	2.45	31	48	-	6/8	4/2
5	0.70 x 0.70 x 0.70	4.13	29	48	33.3	6/8	4/1
6	0.70 x 0.70 x 0.70	5.95	54	48	-	-	4/1

Table 3.4: Parameters tested for three 3D sequences. Spatial resolution is measured in mm. TR, TE, slices, percentage oversampling (OS), Partial Fourier (PF) and Acceleration (Ac) are also recorded. Where no data is present, this technique was not utilised.

Visually inspecting the activity recorded by three 3D fMRI protocols, there is improvement in CNR when compared to the 2D sequences (Table 3.5). All three sequences produced activity that was well restricted to the grey matter, as determined by visual inspection. The effect of additional acceleration in the second phase encode direction is apparent in Sequence 4 where, as highlighted in Figure 3.9, a region of artifact can be seen. As a result of this the tissue contrast is poorer in Sequence 4 than either Sequences 5 or 6, as illustrated by a CNR of 5.7.

The best tissue contrast is observed using Sequence 6. A CNR of 17.3 was recorded (a result almost double that of Sequence 5) and this finding resulted from the effect of the long TR. However, despite good tissue contrast, the long TR makes this sequence less sensitive to signal change. This effect is highlighted by the smaller area of cortex active in response to visual stimulation (Fig 3.7). In Sequence 5, 6/8 partial Fourier has been applied bringing the TR within the decay period of $T2^*$, while allowing sufficient time to show good tissue contrast, evidenced by a CNR of 8.8. This improved sensitivity is apparent in Figure 3.7, where the largest and most significant activity coverage is visible. Because of the different slab placement and head position of the participant for each sequence the area common to all scans was reduced. Analysis was conducted in the acquired image space and was not affected. However for display purposes the sagittal image is not shown in Figure 3.7.




Figure 3.7: Sample coronal and horizontal slices from three 3D sequences acquiring the same resolution data with varying acceleration techniques. The data were cluster corrected ($z=2.3$, $p<0.05$). Significance is indicated here with a heat scale bar where saturation occurs at when $Z>7$. Right and left hemispheres are shown as (**R**) and (**L**) respectively.

The activity to visual stimulation recorded by Sequence 5 is demonstrably the greatest (11.0%, $\pm 13.7\%$) of all three sequences, activating 63% of V1 grey matter voxels. The data recorded during Sequence 6 produced the weakest signal change (6.2%, $\pm 8.0\%$). Despite good tissue contrast only 8% of V1 voxels were reliably active. The data of Sequence 4 also demonstrated weaker activity when compared to Sequence 5, possibly due to the effect of the distortions on acquisition sensitivity. Signal change was recorded to be 8.2%, ($\pm 10.7\%$) and 24% of V1 voxels were active (Fig 3.9). All sequences demonstrated high variability in their percentage signal change when measured in terms of percentage BOLD change. However, the visual inspection and CNR supported the trend of mean percentage BOLD recorded in Figure 3.8




Figure 3.8: 3D sequence signal strength within region of interest VI, measured as percentage signal change for three sequences. Sequence 4, 5 and 6 all acquired data at the same spatial resolution but applied different acceleration factors (Table 3.4). Error bars indicate standard deviation of the BOLD response across the contrast of parameter estimate (COPE).

Sequence	CNR
4	5.7
5	8.8
6	17.3

Table 3.5: CNR measured between grey and white matter for a manually defined region of interest. Shown here for each 3D sequence

For the three 3D sequences data were acquired from within a restricted field of view. This kept the sampling time for the ultra-high resolution volume within the decay period of blood $T2^*$. Although data is read from within a region of interest (ROI), the entire brain continues to be stimulated during the excitation period. This effect, known as wrap or aliasing, occurs when signal from outside the ROI is misappropriated and reconstructed within the ROI. The effect of this is shown in Figure 3.9, Box 1. To compensate for this Sequence 5 employs slice oversampling where an additional volume of the brain is collected, used as a buffer zone for wrap, and is then discarded. Box 2 highlights the distorting effect of increasing the acceleration in the second phase encode direction, as tested in Sequence 4. Such distortion could result in the displacement of signal away from the locus of activity, or reduce the sensitivity to activity within regions of interest. This is supported by the results of percentage BOLD calculations,

where Sequence 4 was found to demonstrate the lowest percentage BOLD change of the three tested sequences.




Figure 3.9: Sample distortions. When acquiring a partial field of view, the slices sitting either side of the 3D slab, not included in the acquisition field of view can wrap back into the acquired data slab. This produces images such as in Box 1, taken from data acquired using Sequence 4. Box 2 highlights a region of distortion that has occurred as a result of acceleration in the first and second phase encode directions. This technique was used to speed acquisition time in Sequence 4.

3.5 Discussion - Part 1

3.5.1 Acquiring data using 3D EPI sequences offers greater opportunity for ultra-high resolution

Of the three 2D sequences tested, Sequence 3, which acquired data at $1 \times 1 \times 3 \text{ mm}^3$, produced the strongest activity and highest CNR. However, this difference was not reliably different from Sequences 1 and 2, as recorded by the measures of percentage BOLD signal change (Fig 3.6). Although $1 \times 1 \times 3 \text{ mm}^3$ resolution data is considerably higher in-plane than the current standard of $1.2 \times 1.2 \times 1.2 \text{ mm}^3$ - $1.5 \times 1.5 \times 1.5 \text{ mm}^3$ resolution data, only the largest of the target structures could still be resolved. Additionally, this sequence does not overcome limitations caused by the partial volume effects outlined previously. The visible noise and lower CNR measured by Sequences 1 and 2 is low, and suggests poor SNR, which would likely make the reliable reconstruction of small-scale organisation difficult.

When comparing the results of the 2D sequences it is important to keep in mind that the 3D sequences were acquired over a shorter period of time, and therefore the number of data points and signal values are not directly comparable. Taking this into consideration, it remains clear that the percentage signal change and raw signal strength of the 3D sequences is on average greater than the 2D sequences. Of the three 3D sequences tested, Sequence 5 produced the strongest activity measured by percentage signal change, (11.0%). Attempts to accelerate the data acquisition by employing acceleration in the second phase-encode direction produced some distortions (Fig3.9), and may have resulted in a reduced sensitivity to BOLD change. The percentage BOLD change recorded during the application of the sequences is lower than other investigations of percentage BOLD change in the visual cortex (Pfeuffer et al. (2002); Qing et al. (2010)). This investigation has attempted to reach a resolution higher than previously attained, and the signals recorded here are sufficient to record the signals of interest for the proposed investigation. The percentage BOLD change was subsequently recorded using Sequence 5 for one full run (10 minutes). Sequence 5 collected data at a spatial resolution of

0.7x0.7x0.7mm³, using a partial Fourier of 6/8, and an acceleration factor of 4/1. The TR of this sequence was 4.128 seconds, TE was 29ms, and a slice over-sample of 33.3% prevented the data either side of the slab from aliasing back onto the region of interest. The percentage BOLD change recorded for this protocol was 12.3%, (+/- 10.3%). The similarity to the percentage BOLD change recorded during the shorter measurement period suggests that the sequence was able to detect activity across short and long stimulation periods.

3.6 Conclusion - Part 1

Of the six sequences tested here, Sequence 5 demonstrated the greatest potential to acquire data at ultra-high resolution within a reasonable TR. The key aims of this investigation were 1) to avoid the effects of partial voluming, and 2) to achieve a resolution sufficient to capture even the smallest target organisation (ODC range from 0.8-1mm in width), and these were achieved. This sequence was therefore selected and the data for Chapter 4 were acquired using this sequence.

3.7 Further fMRI protocol development

3.7.1 Considering effective and nominal spatial resolution

Subsequent to the investigation of the small-scale organisation of ocular dominance columns in Chapter 4, further refinement of the fMRI protocol was carried out. Nadine Graedel and Mark Chiew, two of the major contributors in the development of this protocol, carried out a series of simulations. They investigated the combined effect of partial Fourier and field of view on the effective resolution of the acquired data (Figure 3.10). It was concluded that the blurring effect of partial Fourier, and the time taken to read all of the data points within the field of view could be reduced. As shown in Figure 3.10, when taking the T2* of grey matter in the visual cortex to be 40ms (Peters et al. (2007)) the application of 6/8 partial Fourier to a field of view of 200mm

resulted in an effective enlargement voxel size from $0.7 \times 0.7 \times 0.7 \text{ mm}^3$ to $1.4 \times 1.4 \times 1.4 \text{ mm}^3$. It is clear that some acceleration is necessary, as illustrated by the data acquired using Sequence 6, where the lack of acceleration caused the TR to be too long. A compromise was reached to reduce the field of view, thereby decreasing the number of data points acquired per TR. This slightly reduced the difference between the effective and nominal resolutions (Fig 3.10). The TR was also reduced from 4.128 seconds to 3.6 seconds. The effective resolution was then estimated to be around $1.3 \times 1.3 \times 1.3 \text{ mm}^3$, with improved temporal sensitivity. The true impact of effect resolution is difficult to quantify within the experimental literature, as it is usual to report nominal rather than effective resolution. This is partly due to the differing approaches of calculating the effective resolution, and a lack of consensus on the true effect of blurring on the efficacy of signal detection. In this work the FWHM of the magnitude of the complex point spread function was taken to define effective resolution.




Figure 3.10: Estimated effective versus nominal resolution. Results of simulations carried out by FMRIB physicists Mark Chiew and Nadine Graedel. The red dotted and solid lines indicate the change in effective resolution when applying a partial Fourier of 6/8, assuming the $T2^*$ of grey matter to be 40ms. Where partial Fourier (PF) has been applied, the resulting blur causes the effective resolution to move away from the nominal resolution (black line). This effect is shown to interact with the size of the field of view. Black box indicates the resolution at which data were acquired.

3.8 Discussion - Part 2

In order to investigate the small-scale functional organisation of ocular dominance columns and V2 stripes, an ultra-high resolution fMRI protocol was developed. The aim of this chapter was to develop a sequence to acquire data at a resolution of 1mm or higher, with a temporal sensitivity sufficient to avoid signal distortion and blurring.

3.8.1 The importance of resolution

In several recent high-resolution investigations of visual functional organisation, 2D EPI sequences have been employed. These investigations utilised the technique of creating one high-resolution in-plane and one low resolution slice thickness in order to achieve reasonable SNR (Cheng et al. (2001); Yacoub et al. (2001); Uğurbil (2012)). An alternative to the 2D EPI sequence is the 3D EPI sequence, in which a volume, rather than a series of slices are acquired across a region of interest. In a 3D sequence acceleration methods can be applied to the two phase-encode directions, and slice profile effects, common when working with thin 2D slices, are not an issue. These factors make attaining small, isotropic resolution data points more readily achievable. Isotropic resolution is particularly advantageous when imaging functional targets that may not be captured within a single slice. This is often the case in highly sulcated regions, such as the visual cortex. Strategies employed to circumvent this issue have typically relied upon recruiting participants with a particular cortical morphology (Cheng et al. (2001); Yacoub et al. (2007)). This is untenable when recruiting patient populations, and so an alternative to this approach, the isotropic resolution, must be investigated.

The major limitation for achieving a sub-millimetre isotropic resolution is SNR. This is however, less of a limiting factor when working at 7 Tesla, where the high field strength coupled with the supra linear increase in signal arising from small vessels affords not only greater SNR, but a more specific signal localisation (Logothetis and Pfeuffer (2004); Peters et al. (2007)).

3.8.2 A discussion of the current findings

In close collaboration with members of the FMRI physics team I investigated the effect of different fMRI parameters on achieving an ultra-high resolution fMRI sequence. The 2D sequences were found to elicit weaker BOLD changes within V1 in response to a flashing checkerboard than any of the 3D sequences. The 3D sequences were tested for a range of acceleration methods, maintaining the same spatial resolution. The importance of maintaining a TE shorter than the rate $T2^*$ decay of the grey matter (40ms) was highlighted by the lower sensitivity of Sequence 6 to signal change than Sequence 5 (Peters et al. (2007)). When working at 7T, an optimum TE has been suggested to be ~ 30 ms (van der Zwaag et al. (2009)). This is within the dephasing time of grey matter and at this point there is the largest signal separation for grey matter and large vessels (Yacoub et al. (2001)). This makes signal more specific to the grey matter. Of the 3D sequences, Sequence 5 elicited the largest percentage signal change (11.0%), making use of a 6/8 partial Fourier with acceleration only in one phase-encode direction. The data collected using Sequence 5 had good tissue contrast, while the 33% slice oversampling avoided the aliasing visible in the data collected using Sequence 4 and 6.

Further refinement of the accepted protocol took place after the protocol had been used to investigate ocular dominance columns. At this point the field of view was reduced. This lessened the difference between the effective and nominal resolutions, and decreased the TR from 4.128 seconds to 3.6 seconds. It was predicted that this would improve the sensitivity of the protocol to small activity patterns.

The calculation of percentage BOLD provides one method for measuring signal change. However, Triantafyllou et al., in their investigatory work of sequence parameter optimisation modelled the contribution of thermal and physiological noise to tSNR separately. I suggest that this may be useful for future sequence optimisation work. The addition of physiological noise monitoring could effectively regress certain predictable noise components from the signal, which may have reduced the percentage BOLD signal change here. This may provide a more accurate assessment of the sensitivity of each sequence in detecting relevant stimulus related signal.

3.9 Conclusions

Six fMRI protocols were tested each using gradient echo EPI. Three sequences employed a 2D acquisition technique, while three were collected using a 3D sequence design. A range of voxel resolutions and acceleration techniques were explored before a 3D sequence of $0.7 \times 0.7 \times 0.7 \text{ mm}^3$ resolution was selected. Data collected when the participant viewed a flashing checkerboard stimulus displayed reasonable percentage signal change. The signal change was well restricted to the grey matter, and no distortions were observed. The sequence was able to acquire data at a reasonable TR (4.128 seconds prior to modification, 3.6 seconds after modification), and provided reasonable coverage across the visual cortex (200mm prior to modification, 150mm after modification).

Chapter 4

Investigating Ocular Dominance Columns

4.1 Abstract

In this chapter, I investigate whether the small-scale functional network of ocular dominance columns can be imaged *in-vivo*. Ocular dominance columns refer to bands within V1 where the monocular cell population is particularly dense. The columns are arranged in an interlocking pattern, within which, the right eye serving neurons are on average separate from those of the left eye. I present data suggesting that voxels encoding reliable monocular preference were identifiable, but that there was limited evidence in support of the large-scale functional network of ocular dominance columns. I suggest that the spatial resolution of the data may not be sufficiently high to reliably separate neighbouring columns. To address this, I propose to carry out further development of the fMRI protocol and suggest that validation should be carried out on a larger functional target.

4.2 Introduction

Ocular dominance columns are just one example of many such networks of columns and stripes that exist across the cortex. The role of such functional units within the wider visual hierarchy is unclear, and little evidence exists to show that stripe/column presence or absence can be reliably correlated with a behavioural advantage. Questions therefore remain, and concern: the advantage of stripes and columns, the role that they play in visual processing, and whether they can be used as a model with which to probe the organisation of feature processing across the visual cortex (Horton and Adams (2005)). Such investigations of human functional architecture have remained particularly limited, largely because of a lack of suitable *in-vivo* methodology.

4.2.1 What, and where are the ocular dominance columns?

Hubel and Wiesel first presented evidence for the visual0 systems cortical columns in 1962 (Hubel and Wiesel (1962)). Their compelling electrophysiological investigations of ocular preference in anaesthetised cats was built upon in 1969, when the group visualised in the macaque what we now know as ocular dominance columns (ODC) (Hubel and Wiesel (1969)). ODC have now been identified in a range of mammals, including cat, ferret, macaque and human (Katz et al. (1989)). These columns, found in primary visual cortex, encode on average the neural representations of either the left or right eye. Their ocular preference is dependent upon the input layer received by the column from the laminae of the lateral geniculate nucleus (LGN). The LGN sends postsynaptic neurons via the optic radiation to V1 layer 4C, where the greatest density of monocular cells are located (Horton et al. (1990); Horton and Hedley-Whyte). Although monocular preference is most strongly observed in layer 4C β , their profile is maintained to a variable extent throughout the cortical layers (Fig 4.1) (Cheng et al. (2001); Horton and Adams (2005); Horton and Hedley-Whyte).




Figure 4.1: A diagram illustrating the flow of visual information from the retina to the striate cortex (V1). Visual information from the left or right visual field falling upon the nasal hemifields of the retina (A) is subsequently routed across the optic chiasm to the contralateral hemisphere. As a result of this decussation of neural fibres, the left visual field is processed within the right LGN, a preference reversed for the left LGN. The termination of afferents from the left and right retina into different layers of the LGN (B) maintains a segregated representation of each eyes input. The afferents of each of the six layers of the LGN, which send their major projections to the striate cortex (V1), continue this ocular segregation (C). While this ocular dominance is most strongly observed in Layer 4 of V1, monocular preference is maintained throughout the cortical thickness

Cell populations displaying monocular preference can be detected using single cell recordings or optical imaging techniques (Horton et al. (1990); Horton and Adams (2005)). However, these techniques are limited by the small spatial samples acquired. A wider region of V1 can be investigated in situations of ocular deprivation and subsequent metabolic staining. The associated reduction in the metabolism of pathways involving the deprived eye results in a differential expression of the marker for metabolism, Cytochrome Oxidase (CO). CO remains present in the mitochondria for several days after death, and can be stained for using a cytochrome oxi-

dase stain or by using the autoradiography technique of 2-deoxyglucose injection (Horton and Hocking (1998)). The columns associated with the non-deprived eye uptake the CO stain in greater concentrations than the columns of the deprived eye. These bands of high and low concentrations of CO produce patterns of light (deprived) and dark (normal) stripes as shown in Figure 4.2 (Horton and Hedley-Whyte; Sincich et al. (2003); Sincich and Horton (2002)).




Figure 4.2: A schematic of the pattern of ocular dominance columns in the human right striate cortex, adapted from (Adams and Horton (2009)). The subject became blind in one eye prior to death. Staining the cortex for cytochrome oxidase reveals the difference in metabolic enzyme content between the columns of the lost eye (pale), and the remaining eye (dark). Where the cortex is flattened (A), the entire network can be observed. The blind spot contains no columns, and is shown as a white region in the centre of the patch. Panel (B) shows an overlay of the ocular dominance columns, as they would be organised in situ. The columns are orientated perpendicular to the calcarine sulcus.

4.2.2 A functional role for the ocular dominance columns

ODCs are absent in the mouse, rat, squirrel, rabbit and sheep, and are present in the cat, ferret, macaque and human (Adams and Horton (2009)). It has been suggested that if ODCs were functionally relevant, their presence would have asserted an evolutionary pressure. This theory is rebutted by the columns variable presence in the squirrel monkey, where between 30-50% of squirrel monkeys lack ODCs altogether. This lack of columnar organisation has not been found to be linked to any visual deficit (Adams and Horton (2006, 2003); Horton and Hocking (1996); Horton and Adams (2005); Livingstone (1996)).

It has been suggested that in humans, interactions between ocular dominance columns, in the

form of short-range lateral inhibitory connections arising from layer 4C, play a role in disparity sensitivity (Katz et al. (1989)). This controlled binocular integration, together with other visual cues such as relative size or motion, facilitates our perception of depth. Depth is a key feature of the human visual experience, and when ocular input is interrupted to one eye during the critical development period (1-7 weeks post-natal for macaque) the preference of single cells in V1 can be shifted (Hubel and Wiesel (1963); Jaffer et al. (2012)). The resulting cortical maps of ODCs, stained for CO, are expanded in favour of the non-deprived eye. This change, first recorded in the cat, and stained for in the macaque, has been shown to also occur in humans with dense cataract (Adams et al. (2007)). Dense cataract is a condition where the ocular media become opaque. Although vision is often corrected, the early loss of visual input can be associated with an uncorrectable loss of visual acuity. Despite this correlation, linking ODC architecture to perception has not yet revealed a systematic relationship.

4.2.3 Investigating ocular dominance columns *in-vivo* in humans

The inability to reach cross species consensus in the role of ODCs highlights the importance of *in-vivo* human investigation. Until recently, investigations focused upon human ODCs have relied upon post mortem staining techniques (Adams et al. (2007); Horton and Hocking (1996)). Examination of several flattened human cortices by Adams et al., revealed human ODC width to be between 0.8-1mm, 1.7 times wider than in the macaque (Adams and Horton (2009); Adams et al. (2007); Horton and Hocking (1996)).

Post mortem investigations are fundamentally limited in several key ways 1) sources of error can be introduced during the processes of unfolding and fixation, where there is the possibility of distortion and damage of the cortex, 2) the availability of suitable post-mortem brains is infrequent and 3) regardless of improvements in methodological accuracy, inferring a functional role for the columns is beyond the scope of post-mortem techniques. The solution to the above outlined challenges can be found in the *in-vivo* investigation of human ODCs. When data for this chapter were collected, two particular groups were key in having documented monocular

columnar activity in V1 using fMRI (Cheng et al. (2001); Yacoub et al. (2007)). These investigations produced activity maps as shown in Figure 4.3, and followed a particular procedure that I outline and discuss below.




Figure 4.3: Adapted from Yacoub et al., (Yacoub et al. (2007)) who demonstrate ocular dominance columns using gradient echo functional imaging. A single slice was acquired in the posterior right hemisphere. The inter-hemispheric fissure (IHF) is highlighted. Voxels in red and blue satisfied the criteria for activation by left eye stimulation greater than activity in response to right eye stimulation, or right eye greater than left respectively. Columns are highlighted by location of yellow and blue arrows and were found to be on average 1mm apart. The paradigm was repeated on three different days (A-C). Voxels consistently active in two of three sessions are shown in (D), while voxels active in all three days are shown in (E). Scale bar is 1mm

• Issue 1: Spatial resolution

In fMRI the signal recorded is directly proportional to the size of a voxel. As a result of the small size of voxels required to image ODCs, SNR can be low, leading to difficulties in the detection of relevant signals. To overcome this issue, the above investigations used elongated voxels ($0.75 \times 0.75 \times 2 \text{ mm}^3$ or $0.5 \times 0.5 \times 3 \text{ mm}^3$). In one plane the voxels were sufficiently small to differentiate the monocular signals of right and left eye serving columns, while the elongated voxel depth sampled signal across a wider area, maintaining reasonable SNR.

• Issue 2: Slice placement

A slice consisting of one high-resolution plane and two low-resolution planes is vulnerable in areas of high curvature such as the visual cortex. In such cases, structures of interest may fail to be captured by the high-resolution plane. This effect, in which activity patterns fall within a thick slice and cannot be correctly reconstructed, is called partial voluming. To overcome this,

these groups took advantage of the flat banks of the calcarine sulcus. As shown in Figure 4.4 (A), the slice is applied tangentially to the cortical sheet at the superior or inferior banks of the calcarine sulcus, where ODCs are revealed as a series of stripes. Placing the slice higher or lower (B) reveals a dotted pattern of monocular activity that may be indistinguishable from noise.




Figure 4.4: Adapted from (Cheng et al. (2001)). In (A) the high resolution imaging plane of a slice is applied parallel to a flat region of calcarine sulcus. A series of ODC stripes can be seen in the resulting images. In (B) the high-resolution imaging plane is applied inferior to the lower bank of the calcarine sulcus, bisecting the bank of the calcarine sulcus. In the resulting images the pattern of response of ocular dominance columns is less distinct, lacking the elongated profile only a pattern of monocular dots can be seen

• Issue 3: Participant selection

The degree of curvature along the calcarine sulcus can lead to partial voluming, in which a loss of contrast prevents the accurate reconstruction of features of interest. In order to overcome this, both groups selected participants with flat banks of the calcarine sulcus. This places a limitation on participant numbers, which is reflected by the small subject numbers used by both groups.

• Issue 4: Field of view

Because of the highly specific region of interest both Cheng et al., and Yacoub et al., acquire only three or fewer slices of data.

• Issue 5: Acquisition time

To achieve sufficient SNR, both Cheng et al., and Yacoub et al., acquired data over periods of

~2 hours. This is demanding for even an experienced and motivated participant and would be unlikely to be tolerated by a naive or patient population.

4.2.4 A summary of current methodology and the proposed investigation

The techniques outlined above present a number of limitations 1) data acquisition is restricted to within regions of cortex that are particularly flat, 2) only 2-3 slices can be acquired and 3) participant numbers are limited by this morphological requirement and the long scan times necessary to gain sufficient SNR. Taken together, this experimental design is not widely applicable and fails to completely reconstruct a large area of ODCs.

I proposed to address these limitations here. In Chapter 3, I worked closely with members of the FMRI physics team to develop and test an ultra-high resolution fMRI sequence (further details in Chapter 3). This protocol was designed to acquire data at ultra-high isotropic spatial resolution ($0.7 \times 0.7 \times 0.7 \text{ mm}^3$). I proposed to overcome issues (1) and (5) by acquiring functional data at 7 Tesla. When working at high magnetic field strength SNR is boosted and voxel size can be reduced without large signal sacrifices. Using isotropic resolution, the effect of partial voluming is reduced, and constraints placed upon subject recruitment can therefore be lifted. The SNR associated with this protocol is also high enough to allow for a shorter scan period, reducing the potential for subject motion arising from discomfort.

More recently, and following the completion of this work, another group has been successful in imaging ocular dominance columns using an ultra-high resolution fMRI sequence at 7T (Nasr et al. (2016)). The group acquired data at a spatial resolution of 0.8mm isotropic and were successful in distinguishing monocular columns within V1 for one subject. While this provides good context for our work, it remains to be shown that ODCs can be reconstructed widely and reliably. In addition, similar to the earlier works of Cheng et al., and Yacoub et al., multiple lengthy scan sessions were required (Cheng et al. (2001); Yacoub et al. (2007)). I aimed here to surpass the results of Nasr et al., both in resolution and by reducing the required scan duration.




Figure 4.5: Adapted from Nasr et al. (2016). A series of bands of voxels demonstrating monocular preference can be seen in V1 of one subject. A solid black line indicates the anterior border between V1 and V2

4.2.5 Chapter Aims

In this chapter, I investigated the monocular architecture of V1 using ultra-high resolution fMRI at high field. I present investigatory results and determine whether any recorded monocular signals could be considered as part of the wider functional network of ocular dominance columns. The aims for the chapter were as follows.

- To establish that in response to monocular stimulation, robust activity within V1 was achievable.
- To determine whether periods of monocular stimulation differentially activated populations of voxels within V1.
- To determine whether monocular activity was repeatable across multiple sessions.
- To determine whether any monocular activity patterns could be considered as existing as part of a larger functional network, as would be predicted by functional neuroanatomical evidence (Adams et al. (2007)).

4.3 Methods

4.3.1 Participants

3 subjects aged 24-26 years (1 female) with normal or corrected to normal vision were recruited via standard methods. All participants fulfilled the screening criteria for scanning at 7T (see Chapter 2 - General methods).

4.3.2 Creating a monocular display

In order to identify cortical regions with demonstrable ocular preference, it was necessary to create a stimulus paradigm in which the temporal stimulation of each eye could be separated. Adapting a method of dichoptic visual stimulation described by Schurger (Schurger (2009)) a solid piece of black card was inserted between the screen and the head coil (Fig 4.6). Running a second piece of black card down the length of the head coil enabled the division of the screen into two halves to be continued. This second divider extended to the mirrors and ensured the participant lying within the scanner saw only the left half of the screen with the left eye, and likewise the right side by only the right eye.




Figure 4.6: A schematic diagram of the dichoptic screen display adapted from Schurger., (Schurger (2009)). The design is shown as a side profile through the scanner bore. The bore divider ran from the screen to the mirrors in two parts and completely divided the screen in two. The grey region of the screen illustrates what the participant can actually see when looking out of the head coil.

The stimuli were positioned 7° lateral from the centre of the screen on either side of the bore divider. This allowed maximum symmetric coverage for the visual stimulus when taking into account the curved field of view. The eyes, when diverging in this manner may have focused on a slightly different retinotopic location. In order to stimulate adjacent columns it was important for each eye to focus upon the same retinotopic coordinate. To ensure this was the case, participants wore a custom-made pair of prism glasses (Fig 4.7). A pair of plastic framed sports goggles was fitted with a clip on the nasal bridge. A set of plano lenses clipped into the nasal bridge and were coated with a pair of 7° Fresnel prisms. Fresnel prisms, although commonly associated with greater levels of diffraction than plano prisms, are thin allowing them to fit easily under the head coil. The prisms, applied base out, facilitated the participant in converging and fusing the images presented either side of the bore divider.




Figure 4.7: Prism goggles design. A pair of plano lenses with 7° Fresnel prisms (C), were attached to set of plastic framed sports goggles (A) via a clip (B2 to B1).

4.3.3 Stimuli

Visual stimuli were programmed using the Matlab based toolbox Psychtoolbox. Visual stimuli were projected as described in Chapter 2 onto a mid grey screen (580cd/m^2). The stimuli were presented either side of the bore divider 7° from the horizontal mid point and always consisted of two 8° black boxes, each containing a central black fixation cross of 1° (Fig 4.8.). In two of the three stimulus conditions a 7.5° checkerboard, whose contrast was reversed at a rate of 4Hz was presented to either the left or right eye. The luminance of black and white was measured to be 1.25cd/m^2 and 1140cd/m^2 respectively. Contrast was high at 0.99. Three different visual conditions were used, each presented in a block of 24 seconds. In all conditions the fixation cross and box were present. The possible conditions were as follows: (i) *left eye*, the flashing checkerboard was presented within the left box, the right box contained only the fixation cross; (ii) *right eye*, the flashing checkerboard was presented within the right box, the left box contained only the fixation cross, and; (iii) *fixation only*, no checkerboard stimulus, only fixation point and boxes were presented to both eyes.




Figure 4.8: The stimulus as shown during the right eye condition. A high contrast checkerboard of 7.5° , flashing at a rate of 4Hz is presented to the right eye for 24 seconds. The left eye sees only a fixation cross (1°) and box (8°). A box and fixation cross are also visible to the right eye.

4.3.4 Experimental sessions

All participants were briefed outside of the scanner environment (for full procedure see Chapter 2). Participants were reminded to remain still at all times during the scan. Head position was stabilised and supported by cushion placement. Before each functional scan, participants were presented with the static fixation cross and box. Participants were given time to practise fusing the two images before the functional experiment began. Instructions to keep the two images fused at all times and to remain fixated onto the fused central crosses were provided verbally. At the end of each functional scan participants were allowed a rest period. In each session, three or four functional runs, each lasting 600 seconds were collected. Each run consisted of eight blocks of left and right eye stimulus presentation and nine blocks of the fixation only condition. Stimulus order remained constant as follows: fixation-left-right-fixation. Additional scans including field maps and a high-resolution ($0.7 \times 0.7 \times 0.7 \text{ mm}^3$) T1-weighted structural scan were collected as described in Chapter 2. Each participant took part in two or three sessions. These sessions were collected for the main experiment, with successive sessions taken as repeats. One additional session was used to collect retinotopy scans (for further details see Chapter 2). Each

session took place on separate days.

4.3.5 Assessing eye dominance

Two measures of eye dominance were taken (Rice et al. (2008)). Participants were asked to raise their finger at arms length in front of their face and use it to cover an item of their choosing. They were then instructed to close one eye at a time and to report which eye had been closed when the item became visible to them. In the second method the participants made a ring with their index finger and thumb. Holding their hand at arms length, the participants were asked to focus on an object through the ring and keeping the object in view, to bring their hand to their face. In the first method, the dominant eye is determined as the one closed and in the second method, the dominant eye is determined as the eye through which the object remains visible. No inconsistencies were recorded between the two methods.

4.3.6 fMRI data acquisition

Gradient echo-planar images (GE EPI) were collected at $0.7 \times 0.7 \times 0.7 \text{ mm}^3$ resolution (TR=4.128 seconds, TE=29ms, 200mm FOV, GRAPPA=4\1) using a Siemens Avanto 7T scanner (FM-RIB) equipped with a 32-channel coil. Slices were acquired in the coronal plane. Slab placement was kept constant based upon the following criteria: the slab was orientated 0 to the vertical axis and the posterior slab edge was placed at the occipital pole. A high resolution $0.7 \times 0.7 \times 0.7 \text{ mm}^3$ T1-weighted whole head anatomical scan was collected during the first session.




Figure 4.9: Slab placement and coverage. The functional slab (B) is overlaid onto the T1 weighted structural image (A). The slab is placed to cover the occipital pole and 48 slices anterior to this. The slab was orientated 0° to the vertical axis (here the slab has been registered to the structural scan) and the posterior slab edge was placed at the occipital pole. Placement was prescribed on an individual basis following the acquisition of a short localiser scan

4.3.7 fMRI data analysis

Data were analysed using FSL 6.0, MATLAB and associated custom scripts combining `fslmaths` and `fslstats` commands. Data were pre-processed according to the steps outlined in Chapter 2. In brief, data were corrected for motion using MCFLIRT (Jenkinson et al. (2002)) and were high pass filtered (90s) in order to remove low frequency noise and slow-drift. The raw data sets were processed twice, in one instance spatial smoothing was set to 0mm FWHM and in the second, spatial smoothing was set at 1mm FWHM. Individual functional runs were registered to the first run of the first session using a linear transformation with 6 degrees of freedom (dof), a 180° search angle and a nearest neighbour interpolation in order to preserve local spatial correlations. This first run was motion corrected and averaged across time to produce a high quality functional template, to which all other scans could be registered.

For each run, five conditions were modelled using the GLM implementation in FSL FEAT,

(i) left > fixation, (ii) right > fixation, (iii) left > right, (iv) right > left, and (v) left & right > fixation. Conditions (iii) and (iv) were used to calculate a differential response. Typically, a voxel's activity represents a population response from the neurons within that voxel. It is likely that within a single voxel there will be neurons responsive to both left and right eye stimulation. By subtracting the activity in response to each eye from one-another it is possible to identify voxels showing a greater, but not necessarily singular response to stimulation in one eye. The 3 or 4 functional runs collected during a session were then combined into a single session for each individual subject using a fixed effects analysis. Data are presented as uncorrected p-statistics ($p < 0.05$).

4.3.8 Independent Component Analysis

When imaging small-scale functional activity patterns of interest may be confounded by noise. Noise can arise in a structured manner from physiological sources such as breathing. Independent component analysis (ICA) decomposes a dataset in a model free manner into a series of statistically independent spatial and temporal maps. By examining the maps and associated time courses produced by ICA it is possible to exclude those components arising from such noise sources from further analysis (Smith et al. (2004); Woolrich et al. (2009)). This is particularly useful in a block design where a peak of interest should arise with constant periodicity within the BOLD signal. ICA was run using FSL's MELODIC package on each functional scan before higher-level statistics analysis. The MELODIC package identifies components that are relevant and non-Gaussian relative to the residual fixed-effects within session variation. The components are ordered by decreasing amounts of uniquely explained variance. The list of components was examined, and those not occurring at a frequency equivalent to a 72 second period (left + right + fixation block length), or as a harmonic of this were excluded from further analysis by regressing these components from the signals of interest.

4.3.9 Defining the region of interest

Retinotopic V1 and V2 were defined upon the cortical surface as labels before being transformed into volume space as described in Chapter 2. Higher-level analysis was restricted to voxels within these regions of interest, thereby reducing the number of required statistical comparisons.

4.3.10 Statistical maps

Statistical analysis was conducted on each voxel's denoised time series to produce two types of maps. The first, using the total number of voxels active in response to visual stimulation in either eye (v), was used to calculate how many voxels were responsive to the checkerboard stimulus. The second, used conditions (iii) and (iv) to calculate the number and spatial organisation of voxels responsive preferentially to left or right eye stimulus presentation. This was carried out for each subject and each session for V1 and V2 separately. Reproducibility was then assessed via a conjunction analysis. The conjunction analysis imposes an "AND" operation on each voxel between sessions one and two. A voxel must be active in all sessions to be included in the resulting statistical map. These maps were generated using the data with and without spatial smoothing. The resulting statistical maps were projected onto the cortical surface using the transformation matrix created by registering the high quality functional image to the high resolution structural scan. The total number of voxels displaying ocular preference was expressed as a percentage of those that responded to left or right eye stimulation more than fixation at statistically relevant levels.

4.3.11 A measure of reproducibility

In order to assess the repeatability of the activity patterns, the data were subjected to two measures. The first quantified the total number of active voxels during each condition per session in V1. This gave a measure of similarity of the total number of active voxels. The

second measure quantified the repeatability of ocular preference across the sessions for each voxel (Machielsen et al. (2000)). Repeatability in terms of total number of voxels was defined as follows:

$$N_s = \frac{SV_{\min}}{\sum_{i=1}^S V_i}, \quad (4.1)$$

The session with the smallest number of active voxels was expressed as a percentage of the mean number of voxels active across all sessions, N_s . Where S was the total number of sessions, V_i is the number of voxels in session i . $V_{\min}$ is the number of voxels in the smallest member of the set V_i .

Repeatability in terms of the number of spatially repeatable voxels was defined as follows:

$$N_o = \frac{SV_o}{\sum_{i=1}^S V_i}, \quad (4.2)$$

The number of voxels overlapping in all sessions, N_o , was defined as a percentage of the mean number of voxels active across sessions. Where S was the total number of sessions, V_i is the number of voxels in session i . V_o is the number of overlapping voxels active in all sessions.

4.3.12 Voxels with switching ocular preference

In addition to assessing the repeatability of voxels, the percentage of voxels that switched ocular preference was also calculated. The ocular dominance maps were binarised and multiplied. In the case that a voxel was active in both conditions it would remain. Where the voxel was active in only one session the multiplication with a corresponding voxel containing a zero would result in its removal.

4.3.13 Investigating monocular activity in V2

The total number of voxels displaying ocular dominance within retinotopic V2 was calculated and expressed as a percentage of the total number of voxels responsive to stimulation. If the signals present in V1 were not simply a result of noise it would be predicted that the percentage of voxels demonstrating ocular preference would be lower in V2 than in V1.

4.4 Results

In this section I first present two sample components from the independent component analysis. I then show that the cortical response to the checkerboard stimulus was robust but varied across sessions. I present the statistical maps of ocular dominance for each subject, and discuss the repeatability of these activity patterns across sessions. Finally, I examine whether monocular signals can be detected within V2.

4.4.1 Independent Component Analysis

In order to remove obvious noise components from the dataset independent component analysis (ICA) was carried out as a pre-processing step. ICA was run using FSL's MELODIC package on the time series of each functional run. Statistical analysis was then re-run for each denoised time series. Components of the MELODIC output from each run were included on the basis of whether they: (i) occurred at the correct frequency of 0.014Hz, or as a harmonic of this frequency or (ii) arose from the correct anatomical location. The frequency of interest is determined by dividing the sampling rate (1Hz or 1 cycle per second) by the stimulus cycle duration (72 seconds).

The two sample components typifying these criteria are shown in Figure 4.10. The noise component does not occur with consistent periodicity and has a high frequency (Fig 4.10. Box 1). This is inconsistent with the predicted activity pattern. By examining the associated statistical maps it is clear that the majority of this signal arises from the cerebellum rather than the visual cortex. The ringing pattern of activation and inhibition suggests that this component is an artifact, and it was therefore excluded from further analysis. The signal component has a slower and more consistent time course and a more defined peak in the power spectrum (Fig 4.10. Box 2). The frequency of this peak is 0.014Hz, equivalent to a cycle every 72 seconds. This is consistent with the period of the presentation of fixation, left and right stimuli in three 24-second blocks. The statistical map associated with this component also has its greatest contribution

from the visual cortex. The time course and location of this component makes it likely to be a signal component and it was retained for further analysis.




Figure 4.10: Two ICA components, associated power spectra and time courses. Box 1 contains a noise component. The statistical map shown here has a large level of activity within the cerebellum. The location of the signal combined with the ringing pattern of activation, and its high frequency makes this component likely to be an artifact. In contrast, Box 2 shows a signal component. The activity map is well restricted to the grey matter within the visual cortex. The time course of this component also has a frequency of 0.014Hz, equivalent to 72 seconds, the sum of the fixation-left-right stimulation period.

4.4.2 Results of eye dominance assessment

The results of the eye dominance tests suggest that all subjects were right eye dominant. There were no contradictions across test type. These tests are absolute and do not provide a measure of the extent of dominance.

4.4.3 The total number of voxels showing stimulus specific activity

First the number of voxels responsive during the presentation of a high contrast checkerboard stimulus was calculated (Fig 4.11). The total was expressed as a percentage of retinotopically defined grey matter V1 voxels. This was used to estimate how many voxels were visually responsive and therefore could be predicted to demonstrate ocular dominance. In the case of low activity a subject may have complied poorly with the instructions. Activity in response to checkerboard presentation periods in either the left or right eye stimulation (condition (v)) was summed and compared to the fixation only condition. The percentage activity within V1 was calculated for each session, and where more than two sessions were collected the total percentage of voxels active in the conjunction of session pairs is reported.




Figure 4.11: Number of Voxels Responsive to Visual Stimulation within V1. The graphs here display activity in response to left and/or right eye stimulation in all subjects. Data is shown for individual sessions, and for activity repeatable across sessions (C), with the relevant session number shown in parenthesis. In the case of more than two sessions, the conjunction of all three sessions is also shown. All data is expressed as a percentage of the total number of voxels available as recorded from subject specific retinotopically defined V1 grey matter masks.

In the single sessions of Subjects 1 and 3 the percentage of responsive voxels was high, with 50-80% of the total number of V1 voxels responding to the stimulus. Subject 3 has the highest repeatability across sessions (48%), while for both Subjects 1 and 2 this figure is lower (35% and 21% respectively). This result suggests that signal strength across days was variable, even to a robust stimulus.

4.4.4 The total number of voxels showing monocular preference

The number of voxels displaying monocular preference at above threshold levels was calculated. This figure was expressed as a percentage of the total number of active V1 voxels. The percentage of voxels displaying monocular preference was calculated for each session. Where more than two sessions were collected, the total percentage of voxels active in the conjunction of session pairs is reported, in addition to the percentage of voxels repeatable across all sessions.




Figure 4.12: Assessing the percentage of voxels displaying preference to left or right eye stimulation within V1. The level of activity is expressed as a percentage of the total number of visually responsive voxels. This was calculated for each single session and for the conjunction of these single sessions. The combination of sessions is indicated within parentheses.

Within V1 of Subject 1, 37% of voxels showed statistically significant right eye dominance, while 20% were left eye dominant. This trend of right eye bias was observable on average

across days and for all subjects. However, there was considerable variability in the total number of active voxels across subjects. Activity in response to a checkerboard stimulus for Subject 1 was greatest in session 3 (Figure 4.11), where 74% of the number of V1 voxels were responsive. However, of these voxels, the number displaying significant ocular preference was lower, 14% and 23% of voxels demonstrated right and left eye preference respectively. In the conjunction analysis of session 1 and 2 for Subject 1, 32% of voxels were right eye dominant, while only 10% were left eye dominant. In the conjunction of other sessions the percentage of repeatable voxels is lower.

Subject 2 demonstrated the lowest number of voxels active in response to the checkerboard stimulus (Fig 4.11). In session 1, 33% showed statistically significant right eye dominance, which was balanced by the 27% of voxels that were left eye dominant. This ocular balance was not maintained in session 2, where 48% of voxels were right eye dominant compared to 32% displaying left eye dominance. With similar results to Subject 1, few voxels passed the conjunction requirement. Of those voxels that were active in both sessions, the number of right eye responsive voxels was double that of the left eye (22% versus 11%).

Subject 3 demonstrated the highest percentage of monocular activity of any single session (61% right eye dominant) and in session 2, there was an almost equal representation of each eye (24% right eye dominant, 23% left eye dominant). This was not repeatable and is reflected by the sparse pattern of voxels remaining following the conjunction analysis. For Subject 3 only 1% and 9% of the total number of active voxels reliably demonstrated left and right eye preference respectively. For all subjects on a single session basis the number of voxels displaying both left and right eye dominance was high when compared to those that pass the conjunction threshold.

To further investigate the right eye dominance commented on above and illustrated by Figure 4.12, the mean and standard deviation of ocular response per hemisphere was calculated across sessions. The aim was to establish the relative contribution of contralateral and ipsilateral hemisphere response to visual stimulation. In V1 the afferent connections from the contralateral visual field outnumber those of the ipsilateral visual field. The prediction could therefore be made that within the left hemisphere the response to right eye stimulation would be greater

than to the left eye. The inverse of this trend would then be apparent within V1 of the left hemisphere. In Figure 4.12, the ipsilateral response is shown first for all subjects. Only for Subject 2 does the number of voxels responsive to monocular stimulation in the contralateral visual hemifield outnumber the response to ipsilateral stimulation. Subjects 1 and 3 demonstrated activity greater to right visual field stimulation in both the contra and ipsilateral hemispheres. A two way ANOVA found no significant ($p < 0.05$) difference between responses to contralateral and ipsilateral visual field stimulation in the left and right hemisphere. This suggests that the variability of ocular response across sessions within each hemisphere was high and not predictive of a consistent ocular bias.




Figure 4.13: Assessing the percentage contribution of contralateral and ipsilateral eye response within V1 of the left and right hemisphere (LH and RH). The mean level of activity across all sessions is expressed as a percentage of the total number of visually responsive voxels. Error bars represent the standard deviation of monocular activity across each session.

4.4.5 Ocular dominance maps and quantification

Maps were created in order to investigate the spatial relationships of voxels demonstrating significant ocular preference. Within the statistical maps, voxels displaying significantly greater activity for stimulation in the left eye than right eye are always depicted as blue. Voxels displaying statistically greater activity for stimulation in the right eye than that for the left are

shown always as red. Activity maps are shown projected onto the cortical surface. Data were collapsed through the cortical depth between the mid point of the grey matter and the boundary of the grey and white matter. Taking data from below the mid cortical depth removes the potential contribution of surface vessels and large draining veins. Activity maps are shown for single sessions, and for the conjunction of these sessions. The maps shown in Figures 4.14-4.15 are unsmoothed. Activity maps representing ocular dominance were expected to fulfil two predictions: 1) activity in response to left or right eye stimulation would elicit comparable levels of activity within V1, and 2) activity patterns in response to left and right eye visual stimulation would interlock in elongated clusters of alternating ocular preference. The clusters were expected to be one or two voxels in width and to vary in length.

Although single session activity was high, the lack of reliable monocular preference across all three sessions of Subject 1 resulted in sparse statistical conjunction maps (Fig 4.14). However, including only the two most repeatable sessions (1-2), the resulting conjunction maps do contain some voxels displaying ocular dominance. The interlocking pattern of left and right eye responsive voxels is not observed for any subject. In session 3, a portion of the ventral visual field in the right hemisphere appears to change its monocular preference. This confirms the calculations made previously, in which it was found that ocular preference did not demonstrate reliable bias across sessions. The activity patterns consist of large clusters of highly significant voxels without apparent reliable hemifield or eccentricity bias.

Unlike Subject 1, Subject 2 demonstrated fewer large clusters of highly significant voxels. The patterns were also patchier when viewed on the cortical surface. Across two sessions, the number of repeatable voxels is also less than for Subject 1, 22% of voxels displayed right eye preference compared to 11% that were left eye responsive. The patterns across days were not repeatable for any subject, and changes in ocular preference can be clearly observed near the occipital pole of the left hemisphere in Subject 3. For no subject or session are the clusters orientated reliably orthogonal to the border of V1.




Figure 4.14: Monocular activation across the cortical surface, shown for V1 in Subject 1. Differential activity of voxels displaying statistically greater activity for left eye stimulation than for right eye (blue heat scale) and right greater than left eye (red heat scale). Statistical maps are shown as z-scores, scaled in colour by their significance. Individual session data is shown here (Sessions 1-3), as are the activity patterns of voxels repeatable across combinations of sessions. Activity patterns shown to be repeatable in two sessions are also shown for sessions 1 and 2, sessions 1 and 3, and sessions 2 and 3.




Figure 4.15: Monocular activation across the cortical surface, shown for V1 in Subject 2 and Subject 3. Differential activity of voxels displaying statistically greater activity for left eye stimulation than for right eye (blue heat scale) and right greater than left eye (red heat scale). Statistical maps are z-scores, scaled in colour by their significance. Individual session data is shown here (Sessions 1-2), as well as activity patterns repeatable for both sessions.

The predicted banding pattern was not observable for any subject or session. However, voxels with reliable ocular preference were found for all subjects in the conjunction across single ses-

sion data. In order to exclude small between session mis-registrations, the data were smoothed by 1mm. The resulting conjunction maps are shown in Figure 4.17.

4.4.6 Repeatability

The conjunction analysis offers an indication of repeatability via spatial overlap, but does not provide a measure of how similar the activity levels in each session were. To account for this, repeatability was assessed in two ways. Firstly, the total number of voxels active was expressed as a percentage of the least active session. Secondly the total number of overlapping voxels was expressed as a percentage of the total number of active voxels (Machielsen et al. (2000)). The results of these calculations were expressed as a percentage of the total number visually responsive voxels. This quantifies how likely it is for a single voxel to retain its ocular preference across sessions.




Figure 4.16: *Measuring The Repeatability of Activity Patterns.* Two measures of repeatability are shown here. The largest bar represents the percentage of voxels that were visually responsive in both sessions. The middle bar quantifies the repeatability of the number of active voxels across sessions as a percentage of the total number of active voxels. The smallest bar represents the repeatability of the number of voxels that were active in both sessions, expressed as a percentage of those that survived the conjunction analysis. The percentage of voxels that switched preference between sessions is also included. Scale bar = 1%.

To explore whether the low levels of spatial repeatability across single sessions arose as a result of small cross session mis-registration, the analysis was re-run with the application of a 1mm spatial smoothing kernel. All subjects display improvements in repeatability measures following the 1mm smoothing across ocular preference and session. The spatial maps of the re-run data are shown in Figures 4.17- 4.18. Visually examining the spatial maps, few new areas of cortex become active following the smoothing. However, the clustered organisation

of the activity is clearer. This suggests that the lack of voxels present in the conjunction analysis is not simply a result of cumulative registration errors. Following the application of the 1mm smoothing kernel, 25.6%, 24.3% and 9.4% of voxels displaying significant activity to the stimulus also had a reliable ocular preference in V1. This was a 4-8% percent increase in the number of voxels that passed the conjunction analysis.




Figure 4.17: Monocular activation across the cortical surface, shown for V1 in Subject 1, following a 1mm smoothing kernel. Differential activity of voxels displaying statistically greater activity for left eye stimulation than for right eye (blue heat scale) and right greater than left eye (red heat scale). Statistical maps are shown as z-scores, scaled in colour by their significance. Individual session data is shown here (Sessions 1-3), as are the activity patterns of voxels repeatable for across combinations of these sessions. Activity patterns shown to be repeatable in two sessions are also shown for sessions 1 and 2, sessions 1 and 3, and sessions 2 and 3. The data have been smoothed with a 1mm smoothing kernel.




Figure 4.18: Monocular activation across the cortical surface, shown for V1 in Subject 2 and Subject 3, following a 1mm smoothing kernel. Differential activity of voxels displaying statistically greater activity for left eye stimulation than for right eye (blue heat scale) and right greater than left eye (red heat scale). Statistical maps are z-scores, scaled in colour by their significance. Individual session data is shown here (Sessions 1-2), as well as activity patterns repeatable for both sessions. The data have been smoothed with a 1mm smoothing kernel.




Figure 4.19: Two measures of repeatability are shown here. The largest bar represents the percentage of voxels that were visually responsive in both sessions. The middle bar quantifies the repeatability of the number of active voxels across sessions as a percentage of the total number of active voxels. The smallest bar represents the repeatability of the number of voxels that were active in both sessions, expressed as a percentage of those that survived the conjunction analysis. The percentage of voxels that switched preference between sessions is also included. Scale bar = 1%. The data have been smoothed with a 1mm smoothing kernel.

4.4.7 Ocular preference within V2

To investigate whether noise artifacts contributed to apparent ocular preference, the percentage of voxels demonstrating ocular preference within V2 was calculated and compared to V1.




Figure 4.20: Ocular dominance is expressed as a percentage of the total number of active voxels in V1 and V2 for all subjects. The number of active V1 and V2 voxels were comparable. The percentage of voxels demonstrating ocular preference in V1 is higher than that for V2 in all subjects and all conditions.

The percentage of voxels demonstrating ocular preference was greater in V1 than V2 in all subjects, on average by 17.4%. This result was predicted based upon neuroanatomical evidence, which suggests that V1 contains a greater density of monocular neurons than V2 (Roe (2003)). The preference for right eye stimulation is maintained in V2, although across sessions this was not found to be significant at $p < 0.05$. The percentage of voxels demonstrating left eye preference was 11 times greater in V1 than V2 for Subject 1. This difference was less for right eye responsive voxels, where there were 3.3 times more voxels in V1, than in V2. Subject 2 displayed a similar, if not reduced trend. Voxels with left eye preference accounted for 3.6 times as many active voxels in V1 than V2. Right eye preference was higher, in agreement with the results of Subject 1. However, almost an equal proportion of voxels in V1 (5.3%) and V2 (4.6%) demonstrated preference for the right eye. Left eye preference in Subject 3 accounted for 0.01% of the total number of voxels responsive to the visual stimulus in V1 and 0% in V2. The percentage of right eye voxels was 10.6 times greater in V1 than within V2.

4.5 Discussion

The principle aim of this chapter was to investigate whether monocular signals could be reliably identified within V1, and whether demonstrable spatial regularity in these signals was present.

4.5.1 Recorded activity in V1 is robust but variable in response to a flashing checkerboard

The first aim of this investigation was to establish the percentage of V1 that demonstrated activity in response to a flashing checkerboard. In their 2003 paper, Dougherty et al., calculated the average surface area of V1 for 2-12° of visual space for a single hemisphere to be 1470mm². Taking the average thickness to be 2.5mm (Dougherty et al. (2003)), the volume of V1 across both hemispheres can be estimated to be 7350mm³. To convert this into voxels the cortical volume was divided by the voxel volume (0.7x0.7x0.7=0.343mm³), giving 21428 voxels. The stimulus used in this paradigm occupied between 0-7.5° of visual space, and therefore the number of active voxels can be expected to be fewer than those reported by Dougherty et al.,. The total number of voxels activated in this study varied between 10491 and 17802. As a percentage of grey matter this result was comparable with that of Cheng et al.,. In their investigation of ODCs the group noted that 68-78% of voxels were activated by a checkerboard stimulus. Here the percentage of active V1 voxels was lowest in session 1 for Subject 2 (29.8%), and varied between 53.6% and 81.6% for the remainder of subjects and sessions. This finding highlights the variability of the recorded response, even to a robust stimulus such as flashing checkerboard.

4.5.2 Monocular signals were observable but were not reliable within V1

Monocular signals were recorded within V1 for each subject. However, there was insufficient evidence to determine conclusively whether these signals were part of the functional archi-

texture of ocular dominance columns. In their 2001 paper, Cheng et al., reported that around 31.7% of the total active voxels displayed ocular preference (Cheng et al. (2001)). This figure was approached or surpassed by 5 of the 7 sessions across the subjects for right eye activity. Activity in response to left eye stimulation was weaker than this figure for 6 of 7 sessions. Following the conjunction analysis, the number of voxels demonstrating monocular preference was reduced for all subjects and sessions. No subject matched the reported percentage of Cheng et al., (68.8%) of voxels with repeatable monocular preference. This finding could have resulted from small registration errors, which may have produced a systematic shift of activity patterns. To investigate this further, a 1mm smoothing kernel was applied to the data and the analysis re-run. The additional smoothing only marginally improved repeatability across sessions, suggesting that registration errors were not the cause of the low across session repeatability.

4.5.3 Monocular signals did not elicit the predicted spatial pattern of ocular dominance columns

The spatial patterns observable in Figures 4.14-4.15 do not follow the predicted structure of ocular dominance columns. Clusters of eye specific activity were predicted to occupy roughly similar areas of cortex, and to interlock in bands on average of 1mm in width. In this investigation, activity patterns were either large, lacking the predicted seed region of high significance, or were very small and dispersed across V1. The potential for mis-registration was considered and was found not to be a major contributing factor to the lack of functional architecture. Although the application of a 1mm smoothing kernel did allow more voxels to pass the conjunction analysis, the spatial patterns did not demonstrate the spatial characteristics of ocular dominance columns.

Two possible explanations for this finding can be suggested: 1) An insufficient number of samples were collected to overcome the inherent noise of ultra-high resolution data; 2) The spatial resolution of the fMRI protocol was not sufficient to segregate neighbouring monocular columns. In their investigations both Cheng et al., and Yacoub et al., acquired data over a period

of up to 2 hours (Cheng et al. (2001); Yacoub et al. (2007)). This provided many more samples over which to average and reduce the contribution of unstructured noise. Additionally, both groups used specialised monocular presentation equipment. In the present investigation a low cost stimulation design was created, in which subjects wore prism goggles to fuse two stimuli presented either side of a black dividing board. The divide of the screen was complete, and so it is unlikely that stimulation was erroneously applied to both eyes. However, the task of fusing the two images was effortful, and could have affected attention to the stimuli or promoted excessive blinking and non-compliance from the participants.

4.5.4 Monocular signals were not exclusive to V1

The quantification of monocular activity within V2 revealed apparent monocular signals. While on average monocular signal contributions were larger within V1, right eye stimulation elicited comparable high activity for some sessions. The density of monocular neurons in V2 is lower than in V1 (Roe (2003)), and if, in this investigation, the response to monocular stimulation was significant in V2, this finding could reflect the contribution of noise to this data set. However, this trend was not found to be significant (t-test $p < 0.05$), and with only two or three samples per subject it is difficult to assess the true significance of this result.

4.5.5 A comparison with other investigations of monocular signals within V1

Using an ultra-high isotropic resolution fMRI sequence the spatial limitations, substantial effects of partial voluming and strict selection criteria of participants were overcome.

However, despite these measures, the signal collected here was not sufficiently strong to reliably differentiate monocular signals. In their recent work, Nasr et al., presented results in which a wide area of V1 was shown to contain monocularly responding columns in one subject. The group used similar MR parameters to those used here. Data were acquired at $0.8 \times 0.8 \times 0.8 \text{mm}^3$

resolution, with a slightly shorter TR (4 seconds), 7/8 partial Fourier and an acceleration factor of 4. The group sacrificed the size of their field of view to achieve ultra-high resolution, (54 slices, 153.6 x 153.6mm) (Nasr et al. (2016)). In the current investigation 48 slices were acquired over a field of view of (200x200mm). The group also boosted their contrast to noise (CNR) by averaging their data across sessions rather than carrying out a conjunction analysis, which is a stricter measure of reliability. Consistent with my own conclusions in Chapter 3, the group also note that T2* blurring and partial Fourier effects may have caused the true spatial resolution to be increased beyond that recorded by the estimated voxel size. I comment further on the implications of T2* blur and partial Fourier effects in Chapter 3, where I attempt to reduce their contribution to the effective resolution.

4.6 Conclusion

Taken together, these results suggest that reliably differentiating true monocular signals from noise was not possible. The findings outlined above may reflect the following limitations: 1) the real resolution of the fMRI data acquired may be lower than calculated; 2) the SNR may not be high enough; 3) univariate analysis, which is known to be limited in data sets where dimensionality i.e. the number of voxels, is high, may not be able to classify stimulus specific activity with sufficient confidence. These limitations could be overcome by: 1) revisiting the design of the fMRI protocol and adjusting the field of view accordingly; 2) acquiring data over a longer period of time to provide more samples over which to improve SNR; 3) using multivariate analysis to investigate the data with a method capable of distinguishing between subtle signals.

At the present time resolving ocular dominance columns over a wide area of V1 in multiple subjects may not be possible. No group has yet achieved this and where small sections of left and right eye ODCs have been revealed, long scan times, small field of view and meticulous participant selection criteria have all been implemented.

I conclude that although monocular signals can be detected within V1 for three subjects, the reliability of the activity patterns is low and cannot be linked to the ocular dominance organisation of V1.

Chapter 5

The V2 stripe system

5.1 Abstract

In this chapter, I investigate the functional specialisation of V2. The neurons of V2 are organised into bands that demonstrate, on average, a particular receptive field preference for colour, form, motion or disparity. These functionally specific bands repeat across V2, and are orientated orthogonal to the border of V1. This selectivity for different visual features has contributed to the suggestion that the network of stripes forms part of a parallel feature-processing stream. The evidence to date, for functional organisation within V2, has until recently only been acquired using non-human models. Considering this, I ask here if feature selectivity can be identified in humans using non-invasive ultra-high resolution fMRI. I also examine whether the organisation of activity mirrors the neurophysiologically defined stripe pattern characteristic of V2. I provide evidence in support of feature segregation between colour and motion, and find that on average, activity clusters fall within neurophysiological estimates of stripe width and spacing. However, I suggest that these clusters form only a small part of the larger functional architecture of the V2 stripe system.

5.2 Introduction

This chapter builds upon the work carried out in Chapter 4, within which I sought to resolve ocular dominance columns in three subjects. I concluded that there was insufficient evidence to confirm the presence of populations of voxels with reliable monocular preference. This finding may be attributable to fMRI parameter choices, with particular relevance to their effect on the true spatial resolution of the data acquired. To address these concerns the protocol was reviewed and refined in Chapter 3. I now ask whether, with an updated fMRI protocol it is possible to identify functional specialisation within the human V2 stripe system. V2 stripes are predicted to be larger on average than ocular dominance columns, and so should represent an easier target structure to identify.

5.2.1 An introduction to the V2 stripe system

When macaque extrastriate cortex is stained for the metabolic enzyme cytochrome oxidase (Fig 5.1), a series of stripes orientated orthogonal to the border of V1 is revealed within V2. A similar organisation is apparent when using markers such as CAT-301 or Gallyas silver stain (Sincich et al. (2010); Roe and Ts'o (1997); Tootell et al. (1995); Tootell and Hamilton (1989)). These stripes can be categorised as belonging to one of three different types; thick, thin and pale. In macaques, between 26 and 34 cycles of stripes, organised; pale-thick-pale-thin have been observed per hemisphere (Sincich and Horton (2005); Tootell et al. (1983)). V2 stripes process binocular information, and each stripe is on average 1-1.5mm in width (Ts'o et al. (2001); Roe and Ts'o (1997)).




Figure 5.1: Macaque V1 patches (bottom) and V2 stripes (top) stained for cytochrome oxidase are shown here as a composite of sections. Arrows indicate the position of pale and thin stripes, while brackets highlight the thick stripes. Adapted from Sincich and Horton (2005)

In humans V2 is located immediately adjacent to V1 and from this area the majority of afferent connections are received. V1 receives its major input projections from the lateral geniculate nucleus (LGN) in two streams, the magno (M) and parvo (P), which are associated with particular receptive field properties. Within V2 cell clusters have been observed to demonstrate specialisation in their receptive field properties. These mirror those found earlier in the visual cortex and suggest that the division of M and P pathways continues beyond V1. This segrega-

tion is hypothesised to form the basis of the supposed structured functional organisation within V2.

5.2.2 Functional specialisation and stripe connectivity profiles

The functional specialisation of cells within V2 arises as a result of their connectivity profile with V1. This is evidenced by layer 4B of V1, which receives input from the M-pathway and contains orientation and motion selective cells. These cells send preferential inputs to the thick stripes of V2, within which about two thirds of cells are tuned for binocular disparity. Cells with similar disparities are clustered together suggesting the potential for further modular specialisation within the stripes themselves (Livingstone and Hubel (1987)). The interblobs and blob cells of V1 receive their input from the P-pathway. The interblob cells, which are concentrated in layer 2/3, are tuned to particular orientations with only 10% of orientation specific cells demonstrating colour-opponency. In contrast, over 80% of blob cells show colour-opponency. From layer 2/3 the interblob cells project to the pale stripes of V2, while the blob cells project to the thin stripes, where their respective response properties to orientation and colour are maintained (Livingstone and Hubel (1987)). The outputs of V2 stripes to higher visual modules are not exclusive, but do demonstrate broad specificity in their projections to specific modules. The retrograde tracer wheat-germ agglutinin horseradish peroxidase (HRP-WG), when injected into V5, is only found to label cells within the thick stripes of V2 (Shipp and Zeki (1989)). Similarly, when placing the injection site into V4 two projection patterns are observed. In the first, one stripe per cycle in V2 is found to contain labelled clusters of cells. These clusters are predominantly located in the thin stripe, with some limited overlap with the neighbouring pale stripes. In the second projection type, two stripes per cycle are found to contain labelled cells, which are centered on the pale stripes.




Figure 5.2: In the tripartite model of connectivity, inputs to V1 from the LGN are segregated by magno, parvo and konio cellular input type. These inputs remain largely segregated within V1, and form the basis of functional specialisation in V2. Depending upon stripe type, afferents are sent from V2 to V5 and V4. These projections are not exclusive, and labelled cells are frequently found in neighbouring stripes. While many of these afferents connect directly to V4 and V5, a large number of these projections will also pass through V3. Adapted from Sincich and Horton (2005)

V5, a hub of motion processing, receives major projections from the thick stripes. The pale and thin stripes meanwhile, project preferentially to V4, a module associated with colour and form. This segregation suggests that the V2 stripe system forms a small part of a larger parallel feature-based processing stream (Levitt et al. (1994); Shipp and Zeki (2002); Sincich et al. (2010); Ts'o et al. (2001)).

The simplistic tripartite model of connectivity has been challenged by investigations of interconnectivity between the cortical layers. The dendrites of cells in layer 4B of V1, where many M connections terminate, also reach up into layers 2 and 3 (Livingstone and Hubel (1987)).

Investigations by Sincich et al., into the validity of M and P based pathway segregation found that from V1 onwards, horizontal connections are a prominent feature and potentially introduce cross talk between the two pathways (Levitt et al. (1994); Sincich and Horton (2002, 2005); Sincich et al. (2010)). Although horizontal connections between same stripe type could be considered necessary for maintaining a topographically organised representation of feature space, Malach et al., found that neurons within the pale stripes projected equally as often to thin

and thick as to other pale stripes (Malach et al. (1994)). Making paired injections into adjacent thick and pale stripes, Sincich and Horton noted that a third of the V1 projection neurons were double labelled (Sincich and Horton (2005, 2002)). It is possible that the unstained neurons could represent distinct and segregated projections to only one stripe type within V2, or as the group believes to be more likely, have terminal arbors too small to resolve. In comparison, when applying a similar strategy to pale and thin stripes, a small number of double labelled neurons resulted out of thousands of single labelled cells (Sincich and Horton (2005); Sincich et al. (2010)). While a shared role of thick and pale stripes has been suggested based on the sensitivity in both stripe types to orientation, direction sensitivity is only observed in thick stripes, and chromatic sensitivity only in thin stripes (Shipp and Zeki (2002)).

Despite the potential for shared processing across stripes, stripe specific stimulus preferences are repeatedly observed during electrophysiological recordings (Felleman et al. (2014); Hubel and Livingstone (1987); Levitt et al. (1994); Lim et al. (2009)). To explain these different findings it has been suggested that functional specialisation may change as a function of cortical depth. layers 3A and 3B have been hypothesised to be the most specific in feature preference, while layers 2 and 6 more commonly respond to features not typical of their stripe type (Shipp and Zeki (2002)). The difference in specificity as a function of cortical depth has been linked to feedback mechanisms of higher visual areas. Considering this evidence, it is possible that the role of the V2 stripe system may be extended from basic feature processing and segregation, to potentially directly contributing to the binding of different features (Shipp et al. (2009)). In order to investigate the implications for feature processing in human models, the development of *in-vivo* techniques is of paramount importance.

5.2.3 Identifying the V2 stripe system *in-vivo*

The role of the V2 stripe system in humans, and the extent to which each stripe represents a distinct feature-processing channel is unclear. Metabolic markers have successfully revealed at least three classes of functional stripe: thick, thin and pale, and the arrangement of these is

consistent with that reported in the macaque (thick-pale-thin-pale). Stripe width is scaled with the increase in cortical area and is on average 1-3mm, and 1.5-3cm in length (Hockfield et al. (1990)). Periodicity can range from between 4-12mm (Adams et al. (2007); Burkhalter and Bernardo (1989); Hockfield et al. (1990)).

At the time of this investigation it was not known whether the V2 stripes, identified histologically, were associated with any particular functional specificity in humans. Seymour et al., employed machine-learning techniques to examine feature selectivity as a population response in early visual cortex. Pairs of different visual features, colour and form, colour and motion, and disparity and motion, were found to elicit sufficiently distinct activation patterns within V2 to be considered segregated (Seymour et al. (2009, 2010); Seymour and Clifford (2012)). Seymour et al., also examined the spatial pattern of voxels contributing strongly to the decision making process. The group projected voxels with high weight values ($>2SD$) onto the cortical surface and found no structured spatial organisation. Where clusters of voxels with similar stimulus preferences were predicted to be present, activity patterns were scattered over V2 with no apparent functional organisation. However, this finding could be explained by the spatial resolution of the acquired data. Seymour et al., acquired data with a spatial resolution of 2mm isotropic (Seymour et al. (2009, 2010)) and 1.5mm isotropic (Seymour and Clifford (2012)). The average human V2 stripe width falls between 1-3mm, therefore a proportion of the stripes would likely not be resolved in their entirety, or at all. The group also showed that features could be encoded together within V2. However, referring again to the spatial resolution of the acquired data, multiple stripes could fall within a single voxel and be interpreted as a conjunctions of features.

More recently, Nasr et al., have provided evidence for functional segregation within V2. The group used disparity and colour stimulation to target neuronal groups within the thin and thick stripes (Nasr et al. (2016)). No attempt was made to identify regions responsive to form stimulation. The activation maps shown in their 2016 paper are compelling (Fig 5.3). Visually, the pattern produced is orientated as would be predicted; largely perpendicular to the border of V1. The width of the activation patterns also fell within the predicted bounds for the V2

stripes. However, the group acquired data for each stimulus type across different days and it is unclear whether the time series were concatenated. Ordinarily, noise is modelled on a session-by-session basis. In this way session specific noise, such as scanner variability, thermal noise and motion are accounted for. Concatenating the time series, while improving CNR, may have introduced elements of un-modelled noise into the analysis by not taking this into account.




Figure 5.3: An activity map produced in response to colour stimulation projected onto the flattened right hemisphere of one subject. Solid and dashed lines indicate the borders of V2, with V1 and V3 respectively. The scale bar, shown as a white bar is 1cm. Nasr et al. (2016)

5.3 Chapter aims

Single unit recordings from non-human primate V2 have demonstrated repeatable, stripe dependent stimulus preference (Roe and Ts'o (1997); Sincich and Horton (2005)). The human cortex, despite promising similarities in the uptake of metabolic markers, remains enigmatic where direct functional comparisons are concerned. In this chapter, I aimed to identify populations of neurons with specific stimulus preference to colour, form and motion within V2. I also aimed to show that activation patterns in response to colour, form and motion can be spatially segregated. Unlike the paradigm used by Nasr et al., data for all stimulus conditions were acquired within one session, with a second session collected to verify the repeatability

of the observed patterns. Data were acquired at $0.7 \times 0.7 \times 0.7 \text{ mm}^3$ resolution compared to the $1 \times 1 \times 1 \text{ mm}^3$ resolution data of Nasr et al., representing a 66% improvement in spatial volume (Nasr et al. (2016)). The aims of this chapter were as follows:

- To establish that the cortical activity in response to the visual stimulus was robust and could be detected in V2.
- To establish that activity in response to colour, form and motion could be detected in cortical modules predicted to demonstrate preference for these stimuli.
- To quantify the efficacy of each stimulus in driving the cortical response of V2.
- To determine the extent to which activity patterns in response to colour, form and motion stimulation overlap within V2.
- To determine the extent to which activity in response to colour, form or motion stimulation was repeatable over multiple sessions.
- To create a spatial profile of any observed activity patterns using the following measures:
(i) the width of clusters with stimulus preference, (ii) the periodicity of clusters with stimulus preference.

5.4 Methods

5.4.1 Participants

8 subjects aged 24-35 years (1 male) with normal or corrected vision were recruited via standard methods. All participants fulfilled the standard 7 Tesla screening criteria (see Chapter 2 - General methods for further details).

5.4.2 Stimuli

Visual stimuli were programmed using the Python based program PsychoPy (Peirce (2007, 2009)) with support from Professor Jonathan Peirce at the University of Nottingham. Visual stimuli were projected onto a mid grey screen (580cd/m^2) and subtended 15° by 13° of visual angle. Five different visual conditions were used (Fig 5.4). Each condition was presented in a block of 25 seconds with all, excluding the fixation condition, containing 100 Gabor patches (0.5°). Spatial frequency was set as 3 cycles per degree. Conditions were as follows (i) control: static, achromatic random Gabor patches with a life span of 0.20s (ii) colour: Gabor patches also contained chromatic modulation (iii) form: Gabor patches were arranged to create radial global form (iv) motion: Gabor patches also contained local motion (6Hz, 3 cycles per degree) and (v) fixation: no Gabor patches, only fixation point present (0.5°). In order to achieve isoluminant stimuli, a calibration was carried out on each colour channel of the projector with a PR650 spectrophotometer. For chromatic stimuli each Gabor patch was assigned a random vector in DKL colour space (Peirce (2009)). The colour vector was sinusoidally modulated with the phase of the Gabor along the isoluminant plane. At random intervals within each stimulus block the field of Gabor patches would double in contrast for 0.2 seconds. Participants were instructed to respond to this change by pressing a button on a game pad. There were between 2 and 3 contrast changes per block.




Figure 5.4: Experimental stimuli for control, colour, form, motion and fixation conditions. The fixation dot was present in all conditions. Gabor patches were present in all conditions excluding the fixation only condition. The Gabor patches were modulated either chromatically, or to include radial form or motion. Each of the 100 Gabor patches subtended 0.5° of visual angle.

5.4.3 Experimental sessions

Participants were first briefed outside of the scanner environment, where they were also screened for any MRI contraindications. During each session, three or four functional runs, each lasting 630 seconds were collected. Participants were asked to return for one or two further sessions. These sessions took place on separate days. During the first session, the data slab position was recorded and used to guide future slice placement. Head position was stabilised and supported by cushion placement. Additional scans, including a high resolution ($0.7 \times 0.7 \times 0.7 \text{ mm}^3$) T1-


weighted structural scan and field maps were also collected, as detailed in Chapter 2. Before each functional run participants were instructed to remain still, to maintain fixation onto the central dot subtending 0.5° of visual angle and to attend to the stimulus.

5.4.4 Data analysis stage 1: time series analyses

The data were analysed using FSL 6.0, MATLAB and associated custom scripts with combined `fslmaths` and `fslstats` commands. The data were pre-processed according to the steps outlined in Chapter 2. In brief, data were corrected for motion using MCFLIRT (Jenkinson et al. (2002)) and were high pass filtered (90 seconds) to remove low frequency noise and slow-drift. A 1mm FWHM spatial smoothing kernel was applied to the data. The first run of the first session was motion corrected and averaged across time to produce a high quality functional template, to which all other scans could be registered. Individual functional runs were registered to the high quality functional template using a linear transformation with 6 degrees of freedom (dof), a 180° search angle and a nearest neighbour interpolation in order to preserve local spatial correlations. For each run, four conditions were modelled using the GLM implementation in FSL FEAT, (i) colour >control; (ii) form >control; (iii) motion >control; (iv) mean activity across colour, form, motion and control conditions >fixation. A fixed effects analysis was performed on the functional runs collected during each single session. The resulting statistical maps were cluster corrected ($z > 2.3$), and a minimum cluster size of ten voxels was applied. The resulting statistical maps contained voxels demonstrating preference to colour, form or motion stimuli when compared to the control stimulus, (a field of static, achromatic Gabor patches). Voxels within each session satisfying these conditions were assessed for their repeatability via a conjunction analysis, in which a voxel was considered reliably active if it exhibited significant activity in all sessions.

1. Pre-processing - motion correction, high-pass filtering,
1mm FWHM spatial smoothing kernel applied

2. Lower level analysis - analysis of significant responses for the contrast of interest



3. Registration - single runs registered to a high quality functional template

4. Higher level analysis - fixed effects analysis performed across single runs

5. Clustering - cluster threshold ($z > 2.3$) and minimum cluster size of 10 voxels imposed

6. Conjunction analysis - an AND operation is performed, a voxel must be active in both sessions to be considered reliable




Figure 5.5: The first stages of data analysis are shown here in six steps. These steps identified reliable areas of stimulus specific activity across the visual cortex.


5.4.5 Data analysis stage 2: masking regions of interest

This second stage of analysis isolated activity patterns to within the grey matter of cortical areas V2, V3, V4 and V5. A full description of this process is described in Chapter 2, and here a brief summary is provided. Firstly, cortical areas V2, V3, V4 and V5 were defined using the phase boundaries of subject specific retinotopic maps. The labels demarcating these regions were transformed from the surface of a high-resolution inflated cortical surface to a 3D volume. The labels were binarised and registered to a high quality functional template of each subject. The registered retinotopic labels were used to mask the voxels responsive to any Gabor-containing stimulus. This procedure produced retinotopically restricted masks of voxels known to respond to visual stimulation. These masks were applied to the condition specific activity maps, from which the total number of active voxels in each visual area, and the repeatability of the spatial

maps, could be computed.


1. Transform retinotopic boundaries into the 3D volume

2. Register the label to the high quality functional template




Mask of retinotopic V5

3. Apply these boundaries to functional data - mask activity in response to any visual stimulation



Mask of activity to visual stimulation in V5

4. Apply these boundaries to functional data - mask activity in response to visual stimulation in response to a specific stimulus



Activity to motion in V5

Figure 5.6: The number of voxels demonstrating stimulus specific activity was calculated per visual area for both single session and conjunction data sets. This was achieved first by transforming retinotopic boundaries into a 3D volume (a). Activity in response to any visual stimulation was then masked with the registered retinotopic labels (b). These retinotopically defined areas of activity were subsequently used to mask stimulus specific activity masks (c).

5.4.6 Data analysis stage 3: assessing repeatability

In this third stage of analysis, the repeatability of the observed functional patterns was assessed. Repeatability provides a measure of confidence in the functional clusters, and was assessed using two measures. The first quantified the total number of active voxels for each condition and session. The second quantified the reproducibility of each individual voxel (Machielsen et al. (2000)). Together these measures answered whether a given stimulus repeatedly activated a particular cortical area, and when doing so, produced the same spatial pattern of activation.

Repeatability was assessed across the sessions for each voxel (Machielsen et al. (2000)). The session with the smallest number of active voxels was expressed as a percentage of the mean number of voxels active across all sessions, N_s . Repeatability in terms of total number of voxels was defined as follows:

$$N_s = \frac{SV_{\min}}{\sum_{i=1}^S V_i}, \quad (5.1)$$

Where S was the total number of sessions, V_i is the number of voxels in session i . $V_{\min}$ is the number of voxels in the smallest member of the set V_i .

The number of voxels overlapping in all sessions, N_o , was defined as a percentage of the mean number of voxels active across sessions. Repeatability in terms of the number of spatially repeatable voxels was defined as follows:

$$N_o = \frac{SV_o}{\sum_{i=1}^S V_i}, \quad (5.2)$$

Where S was the total number of sessions, V_i is the number of voxels in session i . V_o is the number of overlapping voxels active in all sessions.

5.4.7 Data analysis - stage 4: profiling activity patterns

In the fourth stage of analysis, data from both single sessions and the conjunction analysis were registered to a high-resolution T1-weighted structural scan. The total number of clusters, the width of each cluster and spacing of each cluster for each condition within V2 was quantified using the FreeSurfer toolbox and custom shell scripts. The first stage of this process required the high quality functional template of each subject to be registered to each subject's high-resolution T1 weighted structural image. The transformation matrix of this registration was then used to register and project the single session and conjunction data onto the cortical surface. The data within the cortical ribbon were collapsed between the mid point of the

cortical depth and the boundary between the grey and white matter. The peak vertex between these two points was projected onto the cortical surface. Sampling between these two cortical depths avoided large surface vessels and potential spurious signals (Shipp et al. (2009)). Once projected onto the cortical surface, measures of cluster width and inter-cluster distance were taken. The width of each cluster was measured across each cluster's centre of gravity, orthogonal to the major axis of the cluster. The distance between each cluster's centre of gravity and that of its two nearest neighbours was recorded to give a measure of inter cluster distance. These measurements were carried out for each condition separately and include single session and conjunction data.




Figure 5.7: Cluster width was defined as the Euclidean distance between the edges of the cluster when passing through the centre of gravity orthogonal to the major axis of the cluster. Cluster distance was defined as the Euclidean distance between a clusters centre of gravity and its two nearest neighbours centres of gravity.

5.5 Results

Drawing upon neurophysiological evidence, it was hypothesised that activity in response to colour, form and motion would produce spatially separate patterns within V2 (Shipp and Zeki (2002); Ts'o et al. (2001)). Before exploring this prediction further, it was necessary to confirm that the stimuli were able to strongly drive cortical activity. Voxels lacking demonstrable response to any of Gabor-containing stimuli were excluded from future analysis.

5.5.1 The total number of active voxels in each visual area

Retinotopically defined visual areas can vary up to two times in size between individuals (Song et al.; Wandell et al. (2007)). To account for this, activity was expressed as a percentage of the total number of grey matter voxels previously isolated during retinotopic stimulation, rather than as the absolute number of active voxels (Fig 5.8). All subjects demonstrated strong activity in response to visual stimulation. The visual area driven most by stimulus activity varied across subjects. Activity in V2 in response to any of the Gabor-containing stimuli varied from 38.61% to 77.70%, and in V3 from 48.13% to 81.17%. The level of activity in V2 and V3 was comparable across all subjects. The cortical response in V2 and V3 was not predictive of the response in V4 and V5.




Figure 5.8: The total number of voxels active in response to any Gabor-containing stimulus expressed as a percentage of the total number of visually responsive grey matter voxels. Percentages were calculated for each subject and visual area (V2-V5).

5.5.2 Determining the stimulus preference of each visual area

The results shown in Figure 5.8 suggest that the Gabor containing stimuli were able to elicit a strong response from V2 in 6 of 8 subjects. The two remaining subjects (4 and 7) were not excluded, but this result was recorded for future reference. Of particular interest to this

investigation were the modulatory effects of colour, form and motion stimulus conditions in V2.

Regions for which there is evidence of a stimulus specific preference were examined first. These areas acted as a marker of the strength of effect of a particular stimulus. It was predicted that V4 would demonstrate preference for the Gabor patches when modulated chromatically, while V5 was predicted to respond during those periods in which the patches contained local motion. V4 was also predicted to respond to the stimulation periods in which the Gabor patches were orientated to create a percept of global form. Failure to elicit a strong response within these regions provided a potential explanation for any future lack of functional activity within V2.




Figure 5.9: The extent to which colour (red), form (yellow) and motion (cyan) stimulation were able to drive the cortical response of each retinotopically defined visual area. Circles indicate individual results, and solid bars are the median response for each condition. 95% confidence interval is indicated by a blue solid line. Activation is expressed as the percentage of voxels known to respond to visual stimulation (Fig 5.8) that were preferentially modulated by either colour, form or motion.

The results shown in Figure 5.9 suggest that on average, both colour and motion strongly drive activity within V4 and V5 respectively. The median response to colour within V4 was 15.0% of

the total number of visually active voxels. In V5 the strongest cortical response was to motion, of which 8.7% of visually responsive voxels were active. On average, the response to colour stimulation periods was greater than to motion within V3. In V2 the response to colour and motion was comparable. Activity in response to the form condition was weak in all visual areas, including V4, where the strongest response was predicted.

Each individual's response to stimulus specific activity was also plotted in order to further investigate whether activity to colour stimulation in V4, and to motion in V5 was predictive of strong responses in V2 and V3 on a subject specific basis. In 7 of 8 subjects, a preference for colour in V4 and motion in V5 was observed as predicted. Only in Subjects 3 and 4 does the form condition clearly drive any visual areas cortical response.




Figure 5.10: The extent to which colour (red), form (yellow) and motion (cyan) stimulation were able to drive the cortical response of retinotopically defined V2, V3, V4 and V5 for each subject. Activation is expressed as the percentage of voxels shown to respond to visual stimulation, that were also preferentially modulated by either colour, form or motion

5.5.3 Activation maps

Voxels demonstrating reliable activity across sessions were projected onto the flattened cortical surface for each individual. Figure 5.11 illustrates the position of visual boundaries on the inflated surface and how, through the process of unfolding and flattening, the maps of visual areas can be visualised.




Figure 5.11: Retinotopic maps shown on inflated and flattened left and right cortical surfaces. Visual boundaries are shown by dotted white lines. The cut to enable unfolding and flattening was placed along the Calcarine Sulcus, shown by the position of an asterisk.




Figure 5.12: The results of conjunction analysis shown here projected onto flattened cortical surfaces for Subjects 1-3. Visual area boundaries are shown as white dotted lines. Where reliable activity in response to colour, form or motion stimulation was recorded, voxels are coloured red, yellow and cyan respectively. Where activity in response to form and motion, colour and form or colour and motion overlaps, the combination is shown as green, orange or pink respectively. Where an area responded to all conditions the region was coloured dark blue.




Figure 5.13: The results of conjunction analysis shown here projected onto flattened cortical surfaces for Subjects 4-6. Visual area boundaries are shown as white dotted lines. Where reliable activity in response to colour, form or motion stimulation was recorded, voxels are coloured red, yellow and cyan respectively. Where activity in responsive to form and motion, colour and form or colour and motion overlaps, the combination is shown as green, orange or pink respectively. Where an area responded to all conditions the region was coloured dark blue.




Figure 5.14: The results of conjunction analysis shown here projected onto flattened cortical surfaces for Subjects 7-8. Visual boundaries are shown as white dotted lines. Where reliable activity in response to colour, form or motion stimulation was recorded, voxels are coloured red, yellow and cyan respectively. Where activity is responsive to form and motion, colour and form or colour and motion overlaps, the combination is shown as green, orange or pink respectively. Where an area responded to all conditions the region was coloured dark blue.

The maps of Figure 5.12-5.14 demonstrate that colour and motion were able to strongly activate areas V4 and V5 respectively in the majority of subjects, as predicted by the region of interest analysis (Fig5.9). The form condition elicited a weaker response across 6 of 8 subjects. The activation patterns of the visually stimulated area of V2 are relatively sparse in the majority of subjects. However, where activation to colour and motion is present there is a clear spatial

segregation between them. The contrasts of colour and control, and motion and control were selected for this reason. Rather than directly comparing colour activity to motion activity, this method enabled voxels expressing conjunction for the two stimuli to also be examined. Stripe-like patterns were also visible in cases such as the right hemisphere of Subject 5, where adjacent clusters demonstrating activation to colour and motion can be observed, with little overlap. These activity clusters were orientated orthogonal to the visual border of V1 as predicted by neurophysiological data (Hockfield et al. (1990)). Similar patterns are visible in 5 of the 8 subjects. The segregation of colour and motion can also be observed to continue over the anterior border of V2 into V3 in 6 of the 8 subjects. In some cases, these patterns are also elongated in a manner similar to that recorded by Nasr et al., (Nasr et al. (2016)).

5.5.4 Percentage overlap across conditions

To quantify the specificity of response of any voxel to a particular stimulus, it was necessary to calculate the degree of separation of activity patterns within V2. This was measured as the percentage overlap of activity patterns in response to colour, form and motion stimulation. This figure was then contrasted with the activity in V3, where it was predicted that there would be a larger proportion of overlapping voxels. A sign test was performed on the difference between the percentage overlap calculated for V2 and V3. The sign test assesses whether the paired differences could have arisen from a population with a median of 0.

Subject	Visual areas	
	V2	V3
1	1.11	5.45
2	10.25	3.58
3	13.50	15.55
4	0.00	3.16
5	7.15	3.79
6	0.00	5.80
7	0.00	0.71
8	1.69	2.82

Table 5.1: The percentage of voxels demonstrating activity to more than one stimulus condition. This percentage was calculated for V2 and V3

The average percentage of voxels with overlapping stimulus preference was 4.21% and 5.11% in V2 and V3 respectively. For 6 of 8 subjects, the proportion of voxels demonstrating activity to more than one condition was greater in V3 than V2. However, percentage overlap between the two areas was not found to be significantly different in a paired t-test ($p < 0.05$, $\text{dof}(n-1)$).

The degree of separation of activity patterns in V2 was predicted to be larger than V3. This predication was not supported here. Nasr et al., noted that functional clusters extended over the border of V2 into V3 (Nasr et al. (2016)). This is a feature that can also be observed in Figure 5.14. This finding suggests that there may be some functional organisation in place within V3 to maintain feature segregation.

5.5.5 Characterising cluster profile

To describe the activity clusters present in each single session and following the conjunction analysis, the clusters were first labelled and counted (Table 5.2). A cluster was defined as a number of contiguous voxels, touching along any vertex edge. The width across each labelled

cluster was then measured as the distance between the cluster boundary vertices across the centre of gravity, orthogonal to the major axis of the cluster (Fig 5.7).

Within the central 1° of visual eccentricity the population receptive field size of neurons within V2 are estimated to be on average 0.5° (Harvey and Dumoulin (2011)). The relative similarity in size of Gabor patches (0.5°) to the size of the receptive field size within foveal V2, means that fewer voxels within V2 would be likely to respond to the stimulus. Additionally, the fixation point itself, which occupied 0.5° was always constant at the center of fixation. These two factors would likely mean that fewer voxels would be predicted to be active within the foveal region of V1-V3. To account for this, inter-cluster distances were not measured across this region. The average cluster width is recorded for each subject and condition in Table 5.2, while Figure 5.15 shows the distribution of cluster widths for all subjects. To examine the inter-cluster distance, which according to neuroanatomical estimates would fall between 4-12mm, the distance between each cluster's centre of gravity and that of its two nearest neighbours was measured.

Cluster width (mm)/SD	Subject							
	01	02	03	04	05	06	07	08
Colour	1.50	1.43	1.69	1.33	1.29	1.20	1.67	1.21
	+/-	+/-	+/-	+/-	+/-	+/-	+/-	+/-
	0.37	0.70	0.69	0.34	0.53	0.50	1.46	1.87
	(2)	(11)	(25)	(7)	(35)	(14)	(20)	(9)
Form	1.19	1.59	1.24	0	2.44	0	0	0
	+/-	-	+/-	-	-	-	-	-
	0.45	0	0.72	0	0	0	0	0
	(12)	(1)	(9)	(0)	(1)	(0)	(0)	(0)
Motion	1.29	1.33	1.77	0.91	1.72	1.01	0	1.55
	+/-	+/-	+/-	+/-	+/-	+/-	-	+/-
	0.36	0.65	0.81	0.31	0.88	0.35	-	0.95
	(22)	(24)	(19)	(4)	(18)	(4)	(0)	(48)

Table 5.2: Width and standard deviation (SD) of clusters following conjunction analysis for each stimulus type. The number of clusters measured is recorded and shown in brackets. Where no width is shown, no clusters were present. Where no standard deviation is shown, only one cluster was present.

The width of neurophysiologically defined stripes in the macaque is estimated to be between 1-1.5mm in width. In humans this estimate is larger, between 1-3mm. The average cluster width recorded here falls within that estimation, although the number of form clusters was too few to determine a reliable distribution.




Figure 5.15: The frequency of clusters with a particular cluster width, measured in millimetres is shown here across for all subjects. Dotted lines indicate predicted minimum and maximum average stripe width of the thick, thin and pale stripes based upon neurophysiological estimates.

On average the inter-cluster distance for each condition also fell between the neurophysiological estimates. However, as shown by the standard deviation there was considerable variation in cluster distances. The predicted inter-cluster distance was between 2-6mm for the form condition to reflect the pale stripes frequency, which occur twice as often as the thin or thick stripes.




Figure 5.16: Spacing of clusters (mm), with error bars indicating the standard deviation of this measure. Where no data is present, there were no clusters present. Dotted lines indicate predicted minimum and maximum inter-cluster distance based upon neuroanatomical estimates.

5.5.6 Examining the repeatability across sessions

The number of voxels demonstrating repeatable stimulus preference within V2 was measured and expressed as a percentage of the total number of voxels active in response to any of the Gabor-containing stimuli. Two measures of repeatability are reported. Firstly, the similarity in the total number of active voxels on each day was calculated (Fig 5.17) and secondly, the total number of overlapping voxels was expressed as a fraction of voxels active across all sessions (Fig 5.18).




Figure 5.17: Repeatability (%) was measured by taking the session with the smallest number of active voxels and expressing it as a percentage of the mean number of voxels active across all sessions




Figure 5.18: Repeatability (%) in terms of the number of spatially repeatable voxels was calculated by taking the number of voxels overlapping voxels and presenting it as a percentage of the mean number of voxels active across sessions

Together the two measures of repeatability assessed whether activity to a given stimulus elicited similar levels of activity in each session, and whether the same population of voxels was active in both cases.

Colour, form and motion stimulation periods elicited highly repeatable levels of activity within V2 for the majority of subjects. It is possible that the low repeatability of the form activity is a feature of the sparseness of activity. The extent to which activity in response to a particular

stimulus was repeatable was not predictive of responses to the other stimuli. For all subjects and stimulus conditions, the repeatability at the level of a single voxel was much lower. This suggests that different populations of voxels responded in each session. To further examine the spatial repeatability of activity patterns in V2 across sessions, activity maps for each session were produced and are shown in Figures 5.19 and 5.20. A summary of these maps is shown here for three subjects, with the remainder included in Appendix Figure 4.

The graphs shown in Figures 5.17 and 5.18, and associated activity maps highlight how activity patterns demonstrate variability across sessions. Although on average, the total number of voxels was 50.65-66.42% repeatable, depending upon the stimulus type, only 6.87-17.48% of the same voxels were active across all sessions.




Figure 5.19: Maps of spatial repeatability across sessions for V2 were created by projecting single session data onto the cortical surface for each subject. The conjunction of these sessions is shown first for each hemisphere for Subjects 1 and 3, and is followed by subsequent sessions in order of their acquisition.




Figure 5.20: Maps of spatial repeatability across sessions for V2 were created by projecting single session data onto the cortical surface for each subject. The conjunction of these sessions is shown first for each hemisphere for Subject 5, and is followed by subsequent sessions in order of their acquisition.

The use of a conjunction analysis could be considered to be a conservative method of examining areas of common activity across sessions. The conjunction method requires each voxel to be active above threshold in all sessions in order to be considered reliable. In their recent work, Nasr et al., used averaged data across sessions to identify voxels with consistent stimulus activity (Nasr et al. (2016)). This method benefits from the improved CNR of including more samples, and may be an alternative way to explore this data set.




Figure 5.21: Maps of average activity patterns across session within V2 shown here projected onto the flattened cortical surface. Data is shown for the right and left hemisphere of all subjects. Where reliable activity in response to colour, form or motion stimulation was recorded, voxels are coloured red, yellow and cyan respectively. Where activity in responsive to form and motion, colour and form or colour and motion overlaps, the combination is shown as green, orange or pink respectively. Where an area responded to all conditions the region was coloured dark blue

The maps produced by averaging activity across sessions contain considerably more activity than those following the conjunction analysis. However, the overlap between conditions is also increased. In some subjects, such as Subjects 4, 6 and 7, activity clusters can be observed as running orthogonal to the horizontal axis of V2.

5.6 Discussion

Activity in voxels responding to colour and motion stimulation periods were on average separable within V2. In contrast, the cortical response to the form condition was weak in 7 of 8 subjects and few repeatable clusters were identified within V2. The spatial profile of colour and motion clusters matched on average the neuroanatomical estimates made for the V2 stripe system. However, the measures of cluster width varied considerably within and between subjects. The patterns produced by these clusters were sparse, but several subjects demonstrated the orthogonal orientation to the border to V1 characteristic of the V2 stripe system. Taken together, I suggest that the clusters observed here may be a small part of the V2 stripe system, but without a more complete reconstruction of the network there is insufficient evidence to conclude this firmly.

5.6.1 The Gabor stimulus elicited strong activity within V2

In their 2001 paper, Cheng et al., report that 76% of V1 grey matter voxels responded to their checkerboard stimulation versus dark paradigm (Cheng et al. (2001)). If we assume the response of V1 and V2 are similar, in a simple on-versus-off paradigm, and take into account the differences in stimulation type, it is reasonable to draw parallels between the two data sets. Of the eight subjects who took part in this study, four demonstrated comparable activity in V2 (between 66.97 and 77.70%) as described by Cheng et al., with the remainder showing activity in the range of 58.26% and 33.18% of grey matter voxels. This finding is similar to that of Chapter 4, in which on average 34% of voxels in V2 were active in response to a flashing checkerboard. This suggests that the Gabor stimulus was able to elicit robust activity in V2. However, even with a robust stimulus, the cortical response varied across subjects.

5.6.2 Colour and motion stimulation drove activity within V4 and V5

The cortical response to Gabor patches modulated to contain colour or form was predicted to be greatest in V4, while stimuli containing a motion component was predicted to modulate the activity of V5. These predictions were based on neuroanatomical evidence from macaque, and human fMRI data, which together suggest that these regions respond preferentially to colour, form and motion. (Bartels and Zeki (2000); Orban et al. (2003); Gallant et al. (2000)). Demonstrable stimulus preference within functional modules V4 and V5 would provide compelling evidence of the efficacy of the modulated Gabor containing stimuli. These modules contain large clusters of neurons with particular receptive field preferences. The preferences of these clusters should therefore be more readily detected than the small bands within V2. Here, the median response of V4 and V5 to colour and motion supported these predictions. In contrast, the response to form stimulation was weak across all visual areas. This finding suggests that colour and motion stimuli, unlike form, elicited distinct activity patterns distinguishable from that of control.

5.6.3 The response to form stimulation was weak

In the macaque, the pale stripes, which flank both the thin and thick stripes, demonstrate specificity for form stimulation. The exclusivity of stimulus preference within the pale stripes has however, been challenged by neuroanatomical investigations. Placing paired injections into adjacent thick and pale stripes, Sincich et al., noted that one third of V1 projection neurons were found to be double labelled (Sincich and Horton (2005)). Similarly, Malach et al., found that neurons within the pale stripes projected equally as often to thin and thick stripes as to other pale stripes Malach et al. (1994). These findings suggest that the pale stripes may be less consistent in their preference to form stimulation. Additionally, any weak signal arising from the pale stripes could be masked by activity from the stronger responding neighbouring thick and thin stripes.

This evidence explains the low response to form stimulation observed in this thesis. However

here, V4, which is known to respond to stimuli containing form, was not driven by periods in which the Gabor patches were modulated to contain a radial form component. This was consistently the case for 6 of 8 subjects. The poor response to form stimulation may therefore be attributed to stimulus choices. There is evidence to suggest that the cells from layer 4C β do not respond well to flashed stimuli (Livingstone and Hubel (1987)). The short lifetime of each Gabor patch (20ms), combined with the global nature of the radial form may have been sub-optimal for driving neural populations responsive to form. Radial form has previously been used to categorise form from control stimulation periods within V2 (Ostwald et al. (2008)). However, in doing so, the group noted that this perception of global form arose through the combinations of local orientation. This finding built on the evidence of electrophysiological work conducted in the macaque by Smith et al.,. In their investigation of macaque V2 activity in response to the presentation of Glass patterns, the group found that radial form failed to elicit a response. Smith et al., suggest that V2 neurons do not summate activity from beyond their own receptive field, and that more specialised neural populations in lateral occipital cortex may be key in perceiving the global form component of Glass patterns (Smith et al. (2007)). The stimulus used in this experiment may not have provided sufficient local cue structure from which form could be identified. Ostwald et al., created a stimulus with 0.4% dot density and were able to successfully classify periods of radial global form from a control stimulus in V2 using multivariate techniques. In the work conducted for this thesis dot density was 0.07%. The receptive field size of V2 neurons vary between 0.5° at 1° eccentricity and 2° at 12° eccentricity (Dumoulin and Wandell (2008)). Taking this estimation, between 1 and 4 Gabor patches could fit within a single receptive field. A more effective stimulus could use differences in local orientation, such as the spiral stimuli used by Seymour et al., (Seymour et al. (2010)). However, the aim of this work was to create a common control stimulus that could be compared across conditions. The effect of changing dot density, either to create local or a more dense global form condition would have become difficult to disentangle from the effect of a form component.

5.6.4 The cortical response to colour and motion within V2 is separable

Colour and motion stimulation periods produced activity patterns that were distinguishable from control. The responses to the two conditions were also equivalent in the percentage of V2 voxels they activated. On average, the percentage overlap between colour and motion was low within V2. This finding was not specific to V2, and percentage overlap within V3 for the two conditions was also low. This finding supports the previous work conducted by Seymour et al., in which colour and motion stimuli were found to elicit sufficiently distinct activity patterns to be distinguished using multivariate techniques (Seymour et al. (2009)). The work conducted in this chapter goes beyond that conducted by Seymour et al., by aiming to identify stripe like patterns of voxels demonstrating these preferences. In their work, the data lacked the sufficient resolution to draw reliable conclusions regarding the spatial organisation of stimulus specific activity within V2.

5.6.5 The observable activity clusters are sparse across V2

In this chapter, the activity patterns produced in response to colour and motion stimulation periods were sparse. In their recent investigation, Nasr et al., demonstrated an interdigitated response to colour and disparity stimulation periods within V2 (Nasr et al. (2016)). However, comparing the two investigations there are several key differences. In this chapter a central aim was to collect all data within one session, using subsequent sessions for validation. In contrast Nasr et al., collected data for each condition on separate days, dedicating one scanning session each to the collection of colour and disparity data. This resulted in the collection of many more data points per condition, and as a result the data set contained more statistical power. Additionally, because the group did not attempt to create a paradigm with a common baseline, they were able to select stimuli to separately and maximally activate the populations responsive to colour and disparity. In this thesis, I compromised on the effectiveness of each individual stimulus in eliciting activity within V2 in order to create a paradigm in which all elements of the stimulus, excluding the feature of interest were directly comparable. In their

investigation, Nasr et al., included an element of form in the disparity condition, and motion in their colour condition. These additions may have driven populations of neurones with dual response properties, boosting the recorded activity. Finally, rather than the conjunction analysis used in this thesis to identify the most reliable voxels across multiple sessions, Nasr et al., concatenated their runs and took an average of time-series activity.

The methodological differences stated above provided Nasr et al., with more statistical power with which to identify subtle stimulus specific activity patterns. The potential advantage of using an average, rather than a conjunction of sessions was tested here, with the results shown in Figure 5.21. Although more voxels were found to be active in the case of averaging the data sets, the overlap between conditions was also increased. In some subject, such as Subjects 4 and Subject 7, activity clusters can be observed as running orthogonal to the horizontal axis of V2. Such organisation is suggestive of an underlying stripe pattern. However, for the majority of subjects the clarity of patterns are obscured by the greater degree of overlap between voxels responding to the two conditions. The more conservative method of conjunction analysis was utilised in this thesis in order to exclude activity patterns arising as a result of noise. In the case of averaging this is not directly controlled for. Taking this into account, the information regarding the underlying functional organisation was not changed by using an averaging technique over a conjunction analysis.

5.6.6 The profiles of available clusters conform to neurophysiological estimates

In the macaque there are estimated to be between 26 and 34 cycles of stripes in V2. The number of stripes in human V2 has yet to be conclusively determined. Because of the organisation of the stripe cycle (thick-pale-thin-pale), the number of pale stripes would be predicted to be double that of thin or thick. However, in this chapter, the number of form responsive clusters was found to be lower than for any other condition. I suggest that the number of clusters recorded is unlikely to reflect the true number of possible V2 stripes in the human for three reasons. Firstly,

the field of visual stimulation was restricted in this investigation to 15° by 13° . This restriction of the visual hemifield reduced the area of V2 receiving visual information, and therefore the area that could be predicted to demonstrate stripe activity. Secondly, the size of clusters was often small, and several neighbouring clusters could conceivably contribute to a single stripe. Finally, the number of clusters measured varied considerably across subjects and sessions.

Neurophysiological estimates of cluster width and spacing suggest that the V2 stripes are 1-3mm in width and repeat every 4-12mm in humans (Hockfield et al. (1990)). These stripes are larger in size compared to the macaque, where V2 stripes are estimated to be between 1-1.5mm in width. The clusters observed here in this chapter conform to the estimated profile of human V2 stripes, although there was considerable variation in both cluster width and frequency.

Examining the spatial maps of these data in Figures 5.19 and 5.20, the poor spatial repeatability occurs for a number of reasons. The data of Subject 1 illustrates the variability in overall activity levels across sessions. In cases such as this, few voxels will pass the conjunction analysis. In other cases, large clusters of activity in each session have some common region of overlap, and subsequent to the conjunction analysis, result in narrow or isolated clusters of activity. This is best illustrated by the activity map produced in response to the motion condition in the left hemisphere of Subject 3. For some subjects, such as Subject 5, activity patterns were well aligned and a larger number of voxels passed the conjunction analysis.

I suggest that variability across sessions directly contributed to the sparseness of activity patterns observed in the spatial maps. A major aim of this thesis was to acquire data within a time frame tolerable by naive participants. I suggest that the number of time points collected during 30 minutes of scan time may not produce data with strong enough SNR. In previous successful investigations of imaging small-scale functional activity, groups have acquired data over a period of up to 2 hours (Cheng et al. (2001); Yacoub et al. (2007); Nasr et al. (2016)).

5.6.7 Segregated activity in response to colour and motion was also observed in V3

The segregation of colour and motion observed in V2 was replicated for some subjects within V3. Although evidence for stripe organisation in V3 is lacking in the macaque, this could be due partly, to the difficulties in accessing the region (Felleman and Van Essen (1991)). A large number of V2 stripe afferents project through V3, and as a result, the segregation of colour, motion and form may also be a feature of V3. Furthermore, functional organisation has been recorded for disparity within V3/V3a in the macaque (Adams and Zeki (2001)), and human (Goncalves et al. (2015)). Recently, Nasr et al., demonstrated stripe like clusters that extended across the border of V2 into V3 in response to colour stimulation (Nasr et al. (2016)). In this chapter I show evidence of both colour and motion clusters extending over the V2-V3 border.

5.6.8 Future directions

The design and analysis of this chapter may have been overly ambitious. By attempting to acquire data for colour, form and motion within a single functional session of 30 minutes in duration, an insufficient number of samples per condition may have been collected. It is also possible that the stimulus design here may not have been able to robustly drive a sufficient number of the V2 neurons. I suggest a paradigm in which stimuli were presented with their own paired control. In this design, stimuli of interest would no longer be comparable, but the removal of a restriction of a common baseline would enable more robust stimuli to be developed. For example, activity in response to colour stimulation could be investigated using a sinusoidal grating, which would be modulated chromatically or contain global motion. This design would result in more stimulus blocks, but removing the restriction on scan duration would compensate for this.

5.7 Conclusions

Visual activity in response to Gabor patches modulated either chromatically or to contain local motion elicited spatially distinct activation patterns within V2. Cortical activity in response to colour and motion were found to be separable on average. Additionally, the clusters conformed to the predicted widths and spacing recorded in the work carried out by Hockfield et al., (Hockfield et al. (1990)). However, the pattern of these clusters was sparse, highly variable and did not conform to the interdigitated stripe profile as predicted.

Chapter 6

Multi-Voxel Pattern Analysis of ocular dominance columns and V2 stripes

6.1 Abstract

In this chapter, I reanalysed data from Chapters 4 and 5 using the multivariate method of Multiple Voxel Pattern Analysis (MVPA). In Chapter 4, I sought to identify the structured monocular system within V1 of ocular dominance columns. The evidence for monocular specificity was limited, and there was no evidence of consistent monocular organisation. In Chapter 5, I demonstrated that activity in response to colour or motion stimulation was distinguishable from a control condition. Clusters with a preference for colour or motion stimulation demonstrated on average, features consistent with V2 stripes. However, single session clusters demonstrated limited repeatability, particularly in the case of form stimulation. MVPA is a sensitive and high-powered analysis technique that has been used to extract relevant signals from noisy, information rich data sets. Here, I present evidence of the successful identification of monocular signals within V1 but not V2. When projected onto the cortical surface the pattern of voxels encoding these preferences did not show the interlocking organisation of ocular dominance columns. In the second investigation, effects of colour and motion stimuli were separable from a control stimulus at significantly above chance levels in 6 of 8 subjects. Activity patterns in response to a form stimulus were not, on average, distinguishable from control. The spatial maps of voxels reliably encoding these signals covered a large proportion of V2. However, the interdigitated pattern of V2 stripes was not observed. I conclude by suggesting that MVPA offers improved sensitivity, reflected by the successful classification of monocular stimulation within V1. However, this sensitivity is at the expense of specificity, indicated by the larger percentage overlap between colour and motion activity patterns in the V2. I suggest that spatial organisation may be investigated further by using a different implementation of MVPA, such as searchlight, a tool which includes clustered search techniques to identify probabilistic spatial relationships.

6.2 Introduction

In Chapter 4, I attempted to resolve ocular dominance columns in three subjects. Despite repeated sessions, a suitable fMRI protocol and robust activity to the stimulus, it was not possible to identify a pattern of voxels exhibiting monocular preference resembling the neuroanatomical descriptions of ODC (Adams et al. (2007); Hubel and Wiesel (1969)). In Chapter 5 I found evidence for the segregation of colour, form and motion within V2. A robust response to colour and motion was observed, while the form condition elicited a weaker response in all subjects. Some evidence for functional banding was apparent, and activity clusters demonstrated a width and frequency consistent with anatomical estimates of stripe anatomy. The clusters were however, limited in number, and did not extend over the entire area of visually stimulated V2.

When considering these findings there are a number of possible explanations for the lack of clear functional architecture. Firstly, to ensure that the scan sessions were tolerated by the non-expert volunteers recruited for the experiments of Chapters 4 and 5, scan time was restricted in duration to around 30 minutes. The number of available data samples may therefore limit the power available from each data set. As a result, relevant activity patterns may not be distinguishable from noise components. Alternatively, the stimulus design itself can be questioned. In Chapter 5, I attempted to identify populations of voxels with demonstrable preference for stimuli modulated to contain colour, form or motion. To isolate the activity of these populations I created a paradigm within which the colour, form or motion components of a field of static achromatic Gabor patches were modulated. To preserve the density of the visual stimulus, a global, rather than local, form condition was created. In the subsequent analysis, while clusters of voxels were found to respond to colour and motion, the global form stimulus did not activate specific regions of V2 in all subjects. I suggest that the choice of global rather than local form may have resulted in this finding.

This chapter considers whether monocular signals are present within the data set of Chapter 4, but are undetectable to the standard univariate analysis techniques. I also explore whether in the data set of Chapter 5, cortical activity patterns in response to form stimulation in V2 can

be identified. Furthermore, I aimed to replicate and extend the results shown in Chapter 5 with regard to the segregation of colour and motion within V2 and V3, and the preference for colour and motion stimulation in V4 and V5 respectively.

6.2.1 fMRI data sets are fundamentally multivariate

An fMRI data set typically consists of hundreds of signal time points taken from a sample population of thousands of voxels. These voxels contain information regarding the activity of tens of thousands of neurons. These neural populations are interconnected and therefore, voxel boundaries can be considered in many cases to be somewhat arbitrary. Despite this, traditional univariate analysis of fMRI data treats each voxel as a separate statistical sample (Cox and Savoy (2003); Norman et al. (2006)). Although effective for strong and discrete signals, trends across populations of voxels are ignored. As a result of this, univariate analysis may struggle to distinguish activity patterns which differ only subtly on a voxel by voxel basis (Haynes and Rees (2006)). Multivariate techniques address these issues and provide a multidimensional approach to investigating large data sets. In this context dimensionality refers to the number of voxels under consideration. Employing supervised machine-learning algorithms, MVPA creates rules based upon the unique characteristics of stimulus specific activity. To create these rules, all data points across time and space are considered. If each stimulus elicits a sufficiently unique response within the population, then these learned rules can be used to predict future cases of activity. MVPA may be particularly useful for examining ultra-high resolution fMRI data sets. In such data sets, many voxels and time points are collected but noise contributions can be high. In univariate analysis it is common to overcome this by spatially averaging across voxels (Norman et al. (2006)). However, this not only reduces the noise component, but also the signal component and the fine spatial signals of interest may be lost (Norman et al. (2006)). MVPA, by comparison does not generally apply spatial smoothing, and instead relies upon the increased sensitivity provided by the large number of samples. This technique is sensitive enough to identify patterns indistinguishable to univariate analyses (Fig6.1) (Cox and Savoy (2003); Norman et al. (2006)).




Figure 6.1: Determining stimulus type based on single or multiple voxel activity. The activity of each voxel can be visualised in N -dimensional space. Here, the activity in response to two stimuli (red and blue) is plotted for each point in time for two voxels. In case (A), the activity for each voxel can be classified independently because the response distributions (red and blue Gaussians) are clearly separable. In case (B), the distributions have considerable overlap and classification is impossible when examining the activity of a single voxel. However, by considering the activity of both voxels, a linear decision boundary can be used to separate the two conditions.

The technique of MVPA has been used in the categorisation of direction of motion (Beckett et al. (2012)), orientation (Carlson (2014); Kamitani and Tong (2005); Pratte et al. (2016)) and the perception of form arising from disparity (Ip et al. (2014)). MVPA has also been applied to examine mental imagery in higher visual areas, (Stokes et al. (2009)) and the retinotopic encoding of facial features (Henriksson et al. (2015)).

In their 2010 and 2009 papers, Seymour et al., report evidence for the conjunction and disjunction of colour and form, and colour and motion within early visual areas (Seymour et al. (2009, 2010)). The group reported voxels encoding not only feature specific information, but also those encoding conjunctions of visual features. However, the resolution at which the group acquired their functional data ($2 \times 2 \times 2 \text{ mm}^3$) could lead to a single voxel covering multiple stripe types (stripes are estimated to be 1-3mm in width). If multiple stripe types were to fall within a single voxel, then instances of stimuli containing several features could activate both stripe

types, giving the impression of preference for conjunctions of stimuli. The possibility of this finding being an artifact of insufficient spatial resolution cannot, therefore be excluded.

Additionally, the voxel patterns shown in the spatial maps of Seymour et al., are sparse, and without functional architecture. One of this chapter's aims is to demonstrate that feature specific voxels are organised in the interdigitated pattern of V2 stripes. It is clear that at the present time, open questions remain as to whether it is possible to identify functional architecture within V2 or V1 with MVPA.

6.2.2 Chapter aims

In this chapter, I applied MVPA techniques to data from both Chapters 4 and 5. In Chapter 4, I asked whether as a population, activity in response to a flashing checkerboard could be distinguished from stimulation periods in which only a fixation cross was present. I then asked whether periods of stimulation to either the right or left eye could be identified as distinct from each other.

In Chapter 5, I investigated whether activity in response to a static field of achromatic Gabor patches could be distinguished from a period of fixation. I then asked whether Gabor patches modulated to contain colour, form or motion could be distinguished from periods of presentation of a static field of chromatic Gabor patches. Finally, I investigated whether the spatial organisation of reliable stimulus encoding voxels resembled those patterns identified in neuroanatomical literature (Hubel and Wiesel (1969); Sincich and Horton (2002)). The aims for each data set were the same and are listed as follows:

- To investigate whether a classifier could distinguish between a control stimulus and fixation.
- To investigate whether the total number of voxels affected classification accuracy.
- To assess whether a classifier could distinguish between stimuli of interest. This classification task was applied to area V1 and V2 for the data set of Chapter 4, and V2, V3, V4

and V5 for the data set of Chapter 5.

- To investigate whether the functional organisation of this analysis resembled the predictions of columns and stripes based upon neuroanatomical evidence (Adams et al. (2007); Hubel and Livingstone (1987); Hubel and Wiesel (1969); Tootell et al. (1998a)).

6.3 Methods

6.3.1 An overview of a multiple voxel pattern analysis pipeline

Accurate classification of stimulus specific activity patterns relies upon the training of the MVPA algorithm to recognise the distinct features of each condition. Firstly, a time-series is divided into training and testing sets. The training data set is given to the classifier along with a label decoding the order of stimulus presentation. The classifier looks for dissimilarity between the activity patterns associated with each condition, and constructs a boundary, termed the decision boundary. If the boundary is well placed, it will succeed in generalising the rule learned during training to the test data. The higher the accuracy of the classifiers prediction, the better it has been able to determine an appropriate boundary. This suggests that the two cases of interest are more likely to be significantly different from each other.

6.3.2 Toolbox

Data pre-processing was completed using a custom set of shell scripts and FSL toolboxes (Jenkinson et al. (2002); Smith et al. (2004)). All MVPA was conducted using LibSVM, implemented in MATLAB (Chang and Lin (2011)). Surface maps were created using FreeSurfer and associated toolboxes (Dale et al. (1999)).

6.3.3 fMRI data pre-processing

Individual runs were motion corrected and high-pass filtered to remove low frequency noise (the temporal frequency was set as the duration of a complete stimulus presentation block). No spatial smoothing was applied. Each functional run was registered to a reference image using a linear registration with 6 degrees-of-freedom. This reference image was always the first functional run, averaged over time to produce a high-quality registration template. Visual manual checks were carried out to ensure that cross-session registrations were correct.




Figure 6.2: An overview of the pipeline for the MVPA process. The time series of each relevant voxel is extracted, normalised, and averaged across stimulus presentation blocks (**1**). The time series is then given to the SVM as N -blocks, along with a list of time locked stimulus labels. The classifier is trained on $N-1$ blocks and tested on the remaining block. Training and testing is cycled over all the time series blocks (**2A-C**). The average accuracy of all testing blocks is reported.

6.3.4 Feature selection

An fMRI data set can consist of thousands of voxels taken across the entire brain. While MVPA can be performed over any size area, the performance of a classifier depends heavily upon the number of variables given to it (Norman et al. (2006)). Dimensionality reduction selects a population of voxels based upon some particular criteria. A balance must be maintained between keeping as many potentially relevant voxels as possible without degrading classifier performance. In this chapter, voxels were selected from within each region of interest. To isolate visually responsive grey matter voxels, a voxel was required to have shown activity

above threshold ($z\text{-stat} > 2.3$, $p > 0.05$), to an on-versus-off contrast in the fixed effects within-subject univariate analysis. These statistical maps were then used to extract the time-series of each relevant voxel.

6.3.5 Feature selection for ocular dominance investigation

In the investigation of ocular dominance, the regions of interest were V1 and V2. As a test case, the classifier was required to distinguish between activity produced in response to monocular stimulation versus a fixation cross. Voxels responding to the on-versus-off contrast were ranked in order of their univariate analysis z -score. From this ranked population, the top 1% of responsive voxels were selected and the discriminability of left and right eye stimulation against fixation was tested. More voxels were then added, so that 10%-100% of active voxels were tested in increments of 10%. The number of voxels giving the highest classification accuracy was selected for the investigation of left versus right eye stimulation. If the classifier was able to correctly label activity at significantly above chance levels, then it was accepted that monocular stimulation produced sufficiently distinct activation patterns. This analysis was run for three subjects. The investigation was repeated for voxels within V2, an area for which no evidence for structured monocular organisation has been shown (Roe (2003)). The number of voxels selected from V2 was matched to the number in V1.

6.3.6 Feature selection for the investigation of colour, form and motion

For the investigation of colour, form and motion activity the regions of interest were V2, V3, V4 and V5. Voxels were selected from the statistical maps as described above. In order to investigate the effect of voxel number on classification accuracy, a two-way classification of a field of static achromatic Gabor patches versus a fixation cross was carried out. Voxels responding to the on-versus-off contrast were ranked in order of their univariate analysis z -score. From this ranked population, the top 1% of responsive voxels were selected and the discriminability of a static achromatic field of Gabor patches against fixation was tested. More voxels were then

added and the classification run again for 10% to 100% of the active voxels in increments of 10%..

When investigating feature specific activity in which a region of interest contained n -voxels, the most active n -voxels were selected to match the size of the smallest region of interest. The number of voxels was kept constant between regions of interest to avoid introducing bias based upon the number of dimensions available for investigation. The stimuli of interest consisted of a field of Gabor patches modulated chromatically, containing radial form, or local motion. To explore whether such segregation could also be observed in other visual areas, the same analysis was run for areas V3, V4 and V5. When producing spatial maps, 100% of active voxels were used.

6.3.7 Time-course extraction and concatenation

Within each functional run the coordinates of visually responsive voxels from the univariate analysis were used to extract the time course of voxels. Individual runs were normalised to a mean of 0 and standard deviation of 1. This helped to reduce variability across runs and sessions. To account for the haemodynamic lag, which is estimated to be between 2-4 seconds, the time series data were circularly shifted forwards by one volume. the labels were circularly shifted forwards by one volume. This shift aligned each label with its correct time-series partner. Following the circular shift, where the final label was now the first, the first label and time point were removed, as these were not matched. In Chapter 4, the TR was 4.128 seconds, while in Chapter 5, the TR was 3.6 seconds, placing both within the estimated bounds of the haemodynamic lag. Data from within individual stimulus presentation blocks were averaged together, creating a time-series of divided into stimulus blocks. Each averaged block was considered as a single trial for training or testing. The time-series data extracted from each functional run were then concatenated and treated as a single data set.




Figure 6.3: MVPA key pre-processing steps. (A) A single run is divided into a series of data volumes. Over the course of a run, different stimuli are presented, shown here as red and blue. Multiple volumes of data are collected during a single stimulus presentation block. (B) The time series is shifted forward by 1 volume in order to account for the haemodynamic delay. The first TR is then excluded. (C) The time series across each block of stimulus presentation is averaged together to give a number of blocks (n -block). (D) The blocks of data from each run and session are then concatenated

6.3.8 Classification

A linear classifier was chosen for this analysis (LibSVM (Chang and Lin (2011))). Linear classifiers compute the weighted sum of all voxel activity. This weighted sum is used in creating a boundary between two categories. An alternative to this linear approach is to use a non-linear classifier, which responds to higher-level features such as combinations, creating a separate class of activity for each case. The two methods have not been shown to differ in classification accuracy, and the results of non-linear classifiers are often harder to interpret (Norman et al. (2006)). A balance must be obtained between providing sufficient data for a classifier and not over-fitting the model, which can prevent its generalisation to new cases. A leave-one-out two-

way classification approach was employed in order to achieve a balance between over-fitting and generalisation. This allowed the classifier to examine and test the entire data set. Accuracy was defined as the ability of the classifier to determine the correct stimulus label-data block pairing. Performance was calculated as the mean classification accuracy across all iterations.

6.3.9 Statistics

Accuracies were calculated for each subject and each region of interest, and are presented as a mean accuracy across train-test iterations. Statistical significance was determined by randomising each voxel's time-series against the stimulus labels. The same leave-one-out train-test process of iterations was carried out and the mean accuracy recorded. 10000 permutations of this process were run, and a distribution of mean accuracies built. The proportion of results from the randomised data set, with accuracy greater than the classification result of the correctly ordered data set, was recorded. This was then used to calculate the p-value with which the null hypothesis could be rejected.




Figure 6.4: Calculating the significance of a result from a randomised distribution. The results of the randomised permutation analysis were used to establish a distribution of chance accuracies. The number of chance accuracies (**B**) greater than the accuracy with which the classifier could detect stimulation periods within the unscrambled data set, was then divided by the number of permutations, of which there were, in this case 10000. This gives a level of confidence for rejecting the null hypothesis that the accuracy with which a classifier could predict the correct label from an unscrambled data set could have been drawn by chance.

6.3.10 Spatial maps

Weight vectors were produced for each train-test iteration, creating a distribution of weights for each voxel. A t-test was performed across the distribution of weights for each voxel at the 95% confidence level. This tested the null hypothesis that a voxel's mean weight was equal to zero, which indicated that a voxel provided no reliable contribution to the classification. For visualising the spatial organisation of mean weights on the cortical surface, the mean weight of each voxel was multiplied by $1-p$, where p is the p -value returned by the t-test. This provides an indication of how strongly and reliably a voxel contributed to the classification. Only voxels whose mean weight differed significantly ($\geq 1SD$) from zero are shown in the spatial map.

During training of left versus right, left was assigned to -1 and right to +1. When training colour, form and motion against a static field of chromatic Gabor patches, the condition of

interest was always assigned -1. In this case, negative weights indicate the presence of the feature of interest via an increase in the BOLD signal. The decision to assign a condition a positive or negative weight is arbitrary as long as the stimulus labels and data blocks are both labeled with the same convention. When training the SVM on colour, motion or form stimulation versus a static field of chromatic Gabor patches, only voxels with negative mean weight are shown in the spatial map. The weight maps were converted back into a volume using the original voxel coordinates. The first functional run was registered to a high-resolution T1 weighted structural volume using a 6 degrees-of-freedom linear transform. The transformation matrix produced by this registration was used to project the statistical maps from a volume onto the cortical surface. To move from volume to the surface, the grey matter was collapsed, and the vertex with the peak response between the border of white matter and mid-grey was selected. Excluding the upper portion of the cortical ribbon reduced the potential contribution of surface veins.

6.3.11 Characterising cluster profile

Activity clusters were characterised in terms of 1) width, 2) elongation and 3) orientation. The clusters of Chapter 4 were characterised using measures 1 and 2, while the clusters of Chapter 5 were characterised using measures 1, 2 and 3.

The data was binarised and resampled onto a rectangular grid using a nearest neighbour interpolation. Clusters were defined using the MATLAB function `bwlabel`, which provides components connected by at least one common boundary with a unique identifier. Each cluster was fitted with an ellipse. The major axis, minor axis, and orientation of the ellipse were then used to compute estimates of the cluster width, elongation and orientation as follows:

- Cluster width was defined as the distance in millimeters across the minor axis of an ellipse fitted to the cluster.
- Elongation was defined as the ratio between the major and minor axis of the fitted ellipse. A value of 1 indicates that a cluster was not elongated, while values greater than 1 are indicative

of a cluster longer in the major than the minor direction.

- Orientation was defined in relation to a straight reference line, drawn by hand along the retinotopic V1-V2 boundary. Dorsal and ventral V2 were identified separately by placing reference lines along the upper and lower phase boundaries of V1-V2. A third line ran anterior to posterior along the calcarine sulcus to separate the two banks, forcing clusters into either the dorsal or ventral portion of V2 based on the position of their centroid. The angle of the centroid in relation to the reference line was then calculated. A reported orientation of 0° would indicate a cluster that lay directly parallel to the reference line, while 90° indicated a cluster that was oriented orthogonal to the V1-V2 border.

6.4 Results 1: Ocular Dominance Columns

This section presents the results of the reanalysis of the data collected for Chapter 4 using multivariate analysis techniques. In all cases accuracy is expressed as proportion correct. Confidence was assessed by calculating the likelihood of a classification result produced from an unscrambled data set, falling within the distribution of accuracies produced by testing on a randomised data set. Failure to exceed the confidence threshold signalled the inability of the classifier to reliably distinguish between two conditions.

6.4.1 The effect of total number of voxels on accuracy

I assessed whether a classifier was able to differentiate between conditions known to be distinguishable using univariate analysis. A two-way classification was carried out using 1-100% of voxels, ranked in order of their z-score in response to the presentation of a flashing checkerboard (Fig6.5) In all cases the classifier performed significantly above chance levels for both V1 and V2. Upper and lower 95% confidence intervals were calculated as the likelihood that the classification result could have been drawn from a distribution of results produced from a scrambled data set. The highest classification accuracy across subjects in V1 was achieved using 100% of active voxels (Subject 1 = 97%, Subject 2 = 96% and Subject 3 = 90%). In V2, the highest classification accuracy was achieved using 90% of active voxels (Subject 1= 83%, Subject 2 = 78% and Subject 3 = 93%). A paired t-test was performed on the classification accuracy of 1-100% of voxels in V1 and V2 for each subject. The classification accuracy of V1 and V2 in response to fixation versus a static field of chromatic Gabor patches was not found to significantly differ ($p < 0.05$) for any subject.

The effect of including a greater or smaller number of voxels on classification accuracy for V1 and V2 was on average 8%. The greatest difference in classification result was reliably found to be between 1% and 90-100% of voxels. This result is not surprising, as the inclusion of more dimensions i.e. voxels, would refine the decision boundary, improving and stabilising classi-

fication accuracy. This hypothesis is supported by the results of all but one subject (Subject 1, V1). The results of this analysis suggests that sufficient information is provided to make an accurate classification even when including the smallest number of voxels ($n=101$).




Figure 6.5: The effect of the proportion of voxels on classification accuracy for three subjects in V1 and V2. The images (A) and (B) show the stimulus conditions. Graphs (C and D) show the effect of the proportion of voxels on the classification result of each subject. From the top 1% of responsive voxels to 100% of voxels responsive at the univariate level above threshold, all results of the classification analysis were significantly above chance. The upper and lower 95% confidence limit is indicated by the red dotted line. Chance is indicated by a solid red line

6.4.2 Distinguishing between left and right eye stimulation

Following the success of distinguishing between a flashing checkerboard and fixation, the more subtle case of left and right eye stimulation were presented to the classifier. 100% of active

voxels were given to the classifier, and as Figure 6.6 shows, the classifier was able to distinguish between the two conditions (Fig6.6. (A) and (B)) within V1. In Subjects 1, 2 and 3 left eye stimulation was classified from right eye with 80%, 75% and 81% accuracy respectively.




Figure 6.6: The results of classification of left eye (A) vs. right eye (B) stimulation in V1 and V2 for three subjects. This was carried out for 100% of visually responsive voxels (C). Significance is indicated by asterisk marks, and the upper 95% confidence limit is shown for each subject as a red line.

The same analysis was run for V2, in order to compare these results with an area known to contain no known structured monocular organisation. In all subjects the classifier performed less accurately than in V1, below the level of significance. A paired t-test found the difference between the classification accuracy of V1 and V2 to be significant at $p<0.01$. Classification accuracy ranged from 53% to 57%. The null hypothesis for V2 that the classification accuracy of left versus right eye stimulation did not differ significantly ($p<0.05$) from that produced by training and testing on a randomised data set could not be rejected. This result suggests that the monocular signals were specific to V1.

The accuracy with which the SVM classifies the data provides information regarding the distinction of stimulus classes across a population of voxels. However, whether this is driven by a large population response with functional organisation, or a small number of influential voxels remains unclear. To address this question weight vector maps were produced and are shown for

V1 in (Fig6.7).




Figure 6.7: Weight maps of voxels that reliably encode monocular preference within V1. Data is shown projected onto the flattened cortical surface. Significance is indicated by colour hue, which saturates at ($SD \geq 3$). Voxels coding for right eye stimulation are shown in red/yellow, and those for left eye stimulation are shown in blue/cyan. Both hemispheres are included for all subjects.

In all cases there is little evidence for the interlocking pattern of monocular activity characteristic of ocular dominance columns. Instead there are large clusters of voxels encoding one stimulus, interspersed with smaller patches of the opposing stimulus. These patterns are fragmented in many cases. Voxels with significant weights are present in both the upper and lower visual hemifields of V1 with no apparent bias (Fig6.7). In a small section of the right hemisphere of Subject 2 and the left of Subject 3, elongated bands of left eye preference can be seen to run from dorsal V1 towards the visual meridian. No such patterns are visible for Subject 1, or elsewhere in Subjects 2 and 3. Large clusters of voxels contributing towards the right eye decision can be seen in the left hemispheres of Subjects 1 and 3.

Measurements of cluster width and elongation are shown in Figures 6.8 and 6.9. Cluster width was recorded as the distance in millimeters across a clusters minor axis, and was expressed as a fraction of the total number of clusters. This was calculated for each subject and hemisphere in response to both left and right eye stimulation. The left eye activity clusters of Subject 1 in the left and right hemisphere were on average 1.4mm and 0.6mm in width respectively. Right eye stimulation produced clusters of on average 0.8mm and 1mm in width in the left and right hemispheres. The clusters of Subject 2 were on average larger, with width recorded as 1.2mm and 3.4mm for the left eye, and 1.2mm and 1.4mm for the right eye in the left and right hemispheres respectively. The activity clusters of Subject 3 were on average 1.6mm and 1.8mm in width in response to left eye stimulation in the left and right hemisphere respectively. Right eye stimulation produced clusters 1mm to 1.4mm in width.

Figure 6.9 describes the ratio between major and minor axis of each cluster and is expressed in terms of the frequency of elongation. In all cases, excluding the right eye clusters in the right hemisphere of Subject 1, the clusters were elongated along the major axis. For Subjects 1, 2 and 3, clusters were on average, 1.35, 1.8 and 1.45 times longer in the major axis than the minor.




Figure 6.8: Cluster width, measured for left and right eye in both left and right hemispheres for three subjects. Width in millimeters is measured across the minor axis of each cluster. The frequency of cluster width is expressed as a proportion of the total number of clusters




Figure 6.9: An indication of the shape of left and right eye clusters, measured in both the left and right hemisphere for three subjects. Expressed as a ratio, length is divided by width to give a measure of whether a cluster is elongated along one axis. The frequency of length/width is expressed as a proportion of the total number of clusters.

To examine whether activity patterns demonstrate any hemispheric bias, the total number of voxels contributing reliably to either left or right eye classification was quantified per hemisphere. This was carried out in the 3D volume rather than on the cortical surface in order to include all potentially relevant voxels. The relative contribution of voxels responsive with

preference for either left or right eye stimulation is expressed as a percentage of the total number of visually responsive voxels. A contralateral bias was found to be present for right eye stimulation periods in all subjects and for left eye stimulation in 2 of 3 subjects.

Subject	Left hemisphere		Right hemisphere	
	Percentage			
	Left stimulus	Right stimulus	Left stimulus	Right stimulus
Subject1	29.15	60.86	32.96	41.62
Subject2	13.29	27.95	26.17	38.91
Subject3	19.86	53.54	17.37	28.55

Table 6.1: The percentage of voxels reliably contributing to left or right eye stimulation in the right and left hemispheres of V1

The total number of voxels contributing reliably to either left or right eye classification was quantified per hemisphere in V2. This assessed whether the hemispheric bias observed in V1 was also present in V2. No consistent hemispheric or stimulus bias was found to be present.

Subject	Left hemisphere		Right hemisphere	
	Percentage			
	Left stimulus	Right stimulus	Left stimulus	Right stimulus
Subject1	24.93	19.03	30.36	25.68
Subject2	29.83	43.14	19.54	7.49
Subject3	54.01	8.26	31.40	6.33

Table 6.2: The percentage of voxels reliably contributing to left or right eye stimulation in the right and left hemispheres of V2

To investigate the weight structure further, the mean weight of each voxel across iterations was computed. The distribution of mean weights of all voxels is expressed in Figure 6.7. A mean

weight of zero indicated that a voxel did not provide a reliable contribution to the classification, while a high positive, or negative mean weight, suggested that the voxel contributed strongly to the classification. The direction of the weight vector determined whether the voxel contributed to the left (negative) or right (positive) eye response. The dispersion of the distribution of mean weights was quantified by calculating the coefficient of variation. The coefficient of variation is defined as the ratio of the standard deviation to the mean. Dispersion provides a measure of how far the mean value is from zero, and therefore provides an indication of the strength of the population response.




Figure 6.10: Mean weight values across permutations. Negative values indicate contributions to the left eye decision, while positive values indicate a contribution to the right eye decision. Zero weight indicates a voxel did not contribute to the classification of either eye. The coefficient of variation (C_v) was calculated, and is expressed as the percentage of voxels demonstrating an absolute mean weight greater than 1SD from zero.

The percentage of voxels with a mean weight greater than 1SD from zero was 92%, 88% and 86% for Subjects 1, 2 and 3 respectively. This finding suggests that a large percentage of the active voxels did contribute strongly to the decision boundary between left and right eye.

6.5 Results 2: V2 segregation of colour, form and motion

In this section I present the data reanalysed from Chapter 5, and investigate the segregation and conjunction of colour, form and motion within V2.

6.5.1 The effect of the total number of voxels on classification accuracy

I assessed whether a classifier was able to distinguish between conditions known to be distinguishable at the univariate level. A two-way classification was made between a stimulus containing static, achromatic Gabor patches, and a fixation only condition. I also asked to what extent the number of voxels included for classification affected the result. Voxels were first ranked in order of highest to lowest z-score. The top 1% of the ranked voxels were selected and given to the SVM before the number of voxels was increased to 10% and testing was repeated. This process of adding voxels was repeated in 10% increments until 100% of visually responsive voxels within V2 were tested (Fig6.11). In all participants classification was significantly above chance when any number of voxels were included in the test. The highest classification accuracy across subjects was achieved using 100% of voxels (Subject 1 = 98.3%, Subject 2 = 84.3%, Subject 3 = 91.7%, Subject 4 = 90.8%, Subject 5 = 96.3%, Subject 6 = 86.7%, Subject 7 = 68.8% and Subject 8 = 70.8%).

Permuting the train-testing procedure on a scrambled data set 10000 times produced a distribution of accuracy results. The 95% confidence interval of this distribution is shown as a dotted red line and is an indication of confidence. The 95% confidence interval is greater than 50%, indicating that the SVM is able to find relevant information even within a scrambled data set. This is not surprising as a simple on-versus-off classification will likely have a distinct BOLD signal and fewer voxels will need to share the correct time point-label pairings to lead to an accurate classification result.

Classification accuracy remains relatively stable regardless of the percentage of voxels used. It is likely that the first voxels included in the analysis, being drawn from a population with a high z-score, offer the most information to the SVM. Subsequent voxels refine the classification boundary but do not greatly alter it. It is therefore unlikely that a hypothesis will be erroneously accepted or rejected based upon the number of voxels included in the classification. When selecting the proportion of voxels to examine stimulus specific classification across all visual areas, the number of voxels equivalent to the size of each subject's V5 was chosen. This ensured

that no more information was available to any particular visual area.




Figure 6.11: The effect of total number of voxels on classification accuracy within V2 for all subjects ($n=8$). The images (A) and (B) show the visual conditions, between which the classifier had to distinguish. In (A) a fixation dot is present while in (B) there is a full field of static achromatic Gabor patches. The graphs (C) show the effect of the total number of voxels on the classification result of each subject. Voxel number ranged from the top 1% of responsive voxels, to 100% of voxels responsive at the univariate level above threshold. The upper and lower 95% confidence limit is indicated by the red dotted line. Chance is indicated by a solid red line.

6.5.2 Segregation of colour, form and motion within V2

Having established, in all subjects, that a linear classifier was able to distinguish between on-versus-off stimulation periods, the more challenging stimulus specific classification was carried

out. Using a two-way linear classification SVM, stimuli modulated for colour and motion were detectable within V2 for all but two subjects.




Figure 6.12: The distinction of colour, form and motion activity patterns within V2, shown here for all subjects ($n=8$). Panels (A-C) depict the stimulus conditions, with a static field of chromatic Gabor patches shown on the left, and modulations in colour, form and motion shown on the right. In C, motion was not restricted to linear horizontal motion. The graphs (D) show the accuracy of a classifier in correctly labeling an instance of stimulation. Significance is indicated by asterisk marks, and the upper confidence interval ($p=0.05$) is shown for each subject as a red line.

Using a two-way linear classification SVM, stimuli modulated for colour and motion were detectable within V2 for 6 of 8 subjects. Stimuli modulated to contain radial form remained

undetectable from a static field of chromatic Gabor patches in all but 2 subjects. This suggests that the stimulus was unable to generate sufficiently distinct activation patterns to be classified. For the remaining participants, colour and motion stimulation periods were distinguishable from the presentation of a static field of achromatic Gabor patches, with high classification accuracy. To enable further comparison with the univariate work, spatial maps of voxels that reliably encoded stimulus preference (≥ 1 SD) were produced. These flattened cortical maps are shown for all subjects who produced significantly greater than chance classification accuracy (Fig6.13 and Fig6.14). The classification accuracy in response to periods of form stimuli were significantly above chance in only two subjects, and so were not included in the spatial maps. The colour of voxels exceeding threshold were saturated in order to aid comparison with the univariate maps, which are also included in Figures 6.13 and 6.14. The region of interest is shown first to highlight the voxels that were visually responsive.

The maps produced by the multivariate analysis contain more stimulus specific voxels than the maps produced by univariate analysis. In some subjects (Subjects 4, 5 and 8) the patterns produced by these voxels appear stripe like, and are arranged as a corona. A similar organisation is described by the neurophysiological literature of macaque V2 stripes (Shipp and Zeki (2002); Sincich and Horton (2005)). However, this pattern does not extend across all of the visually stimulated area of V2, and is not seen in all subjects. Large clusters, not in keeping with the neuroanatomical descriptions of stimulus specific activation patterns of V2 can also be observed in Subjects 1, 3 and 6. In dorsal V2 of the right hemisphere in Subject 3, there are clear similarities between the univariate and multivariate activation maps. The degree of similarity of the univariate and multivariate activation patterns is quantified as percentage overlap for all subjects. The segregation of colour from motion is less distinct in the multivariate maps with more overlap between conditions than the univariate (Table 6.3).

Cluster width, elongation and orientation were calculated for all subjects (Fig6.15-6.18). The average cluster width in response to colour stimulation is indicated by the peak value in Figures 6.15-6.16. On average, both colour and motion clusters were found to be between 1-1.4mm in width. Colour clusters were on average 1.2-1.8 times longer along the major axis. This measure

was similar for motion clusters, which were on average 1.4-1.8 times longer along the major axis (Fig6.17-6.18). The orientation of these clusters did not demonstrate a consistent trend, as is indicated by Figures 6.19 and 6.20.




Figure 6.13: Flattened cortical maps of V2 for each hemisphere for Subjects 1-4. The maps of the region of interest, determined during the univariate analysis are shown first. These voxels, shown in yellow, were more responsive to visual stimulation with a field of static achromatic Gabor patches than during the presentation of a fixation cross. The univariate maps are shown next and are followed by the multivariate maps. Voxels included in the multivariate maps have a weight value indicating significant ($\geq 1SD$) and reliable stimulus encoding. Subject 2 is excluded on the basis that their classification accuracy was not significantly above chance. Due to the poor classification performance, voxels encoding form specific information (yellow) are not shown in the multivariate maps, but are shown in the univariate maps. Voxels encoding colour (red), motion (cyan) are shown for both univariate and multivariate maps. Voxels responsive to both colour and motion stimulation are shown in black.




Figure 6.14: Flattened cortical maps of V2 for each hemisphere for Subjects 5-8. The maps of the region of interest, determined during the univariate analysis are shown first. These voxels were more responsive to visual stimulation than during the presentation of a fixation cross. The univariate maps are shown next, followed by the multivariate maps. Voxels included in the multivariate maps have a weight value indicating significant ($\geq 1SD$) and reliable stimulus encoding. Subject 7 is excluded on the basis that their classification accuracy was not significantly above chance. Due to the poor classification performance, voxels encoding form specific information (yellow) are not shown in the multivariate maps, but are shown in the univariate maps. Voxels encoding colour (red), motion (cyan) are shown for both univariate and multivariate maps. Voxels responsive to both colour and motion stimulation are shown in black.




Figure 6.15: Cluster width, measured for colour and motion in both left and right hemispheres for Subjects 1-4. Width in millimeters is measured across the minor axis of each cluster. The frequency of cluster width is expressed as a proportion of the total number of clusters.




Figure 6.16: Cluster width, measured for colour and motion in both left and right hemispheres for Subjects 5-8. Width in millimeters is measured across the minor axis of each cluster. The frequency of cluster width is expressed as a proportion of the total number of clusters.




Figure 6.17: Cluster elongation, measured for colour and motion in both left and right hemispheres for Subjects 1-4. Subject 2 is excluded on the basis that their classification accuracy was not significantly above chance. Elongation is measured as the ratio of major to axis of each cluster. The frequency of length/width is expressed as a proportion of the total number of clusters.




Figure 6.18: Cluster elongation, measured for colour and motion in both left and right hemispheres for Subjects 5-8. Subject 7 is excluded on the basis that their classification accuracy was not significantly above chance. Elongation is measured as the ratio of major to axis of each cluster. The frequency of length/width is expressed as a proportion of the total number of clusters.




Figure 6.19: Colour and motion cluster orientation in relation to the border of V1-V2. Orientation is measured by the position of the centroid in relation to the reference line. 0° indicates that the cluster lay parallel to the border; while 90° indicates that the major axis of the cluster lay orthogonal to the border. Here the results of Subject 1-4 are shown, and measures of orientation are expressed as a proportion of the total number of observed clusters.




Figure 6.20: Colour and motion cluster orientation in relation to the border of V1-V2. Orientation is measured by the position of the centroid in relation to the reference line. 0° indicates that the cluster lay parallel to the border, while 90° indicates that the major axis of the cluster lay orthogonal to the border. Here the results of Subject 5-8 are shown, and measures of orientation are expressed as a proportion of the total number of observed clusters.

To investigate the weight structure further, the mean weight of each voxel across iterations was computed. The distribution of mean weights of all voxels is expressed in Figure (6.21).




Figure 6.21: Mean weight values across permutations. Negative values indicate that the voxel responded to the presence of colour or motion via an increase in the BOLD signal, while voxels with mean positive values responded to the presence of the control condition. A mean weight of zero indicates that the voxel did not contribute to the decision reliably to the classification. The coefficient of variation (Cv) was calculated, and is expressed as the percentage of voxels that demonstrated an absolute mean weight greater than 1SD from zero.

A mean weight of zero indicated that a voxel did not provide a reliable contribution to the decision boundary, while a high positive, or negative mean weight, suggested that the voxel contributed strongly to the decision boundary. The direction of the weight vector determined whether the voxel contributed to the stimulus of interest, in this case colour or motion. This is shown here by a negative weight. The dispersion of the distribution of mean weights was quantified by calculating the coefficient of variation. On average the percentage of voxels with a mean weight greater than 1SD from zero was 86% (+/-3) for colour and 83% (+/-3) for motion. This finding suggests that a large percentage of the active voxels did contribute strongly to the decision boundary between colour and a control stimulus, and motion and a control stimulus.

The number of stimulus responsive voxels present after threshold was quantified for both the uni and multivariate analysis. The percentage overlap between voxels responsive to colour and motion was also calculated. The similarity of each analysis technique was quantified by calculating the percentage overlap between the two data sets (Table6.3). On average, multivariate analysis techniques revealed 80% more stimulus responsive voxels than with univariate analysis. The number of voxels responsive to each visual stimulus is roughly equal following multivariate analysis. This is not the case for the univariate analysis, in which the response to one stimulus dominates. The patterns produced by multivariate analysis show on average 16.8% overlap between stimulus conditions compared to the 0.57% of univariate analysis.

	Multivariate			Univariate			Overlap	
	Number of Voxels		Percent	Number of Voxels		Percent	Similarity	
			Overlap			Overlap		
	Colour	Motion		Colour	Motion		Colour	Motion
1	5734	5766	18.48	8	449	0.02	0.00	3.17
2	-	-	-	-	-	-	-	-
3	2460	2332	15.19	1518	1404	2.98	34.54	43.56
4	1675	1784	15.09	319	176	0.10	0.02	0.67
5	3095	3249	18.58	1509	861	0.26	25.41	14.1
6	2407	2336	17.07	175	86	0.00	12.52	2.71
7	-	-	-	-	-	-	-	-
8	1931	2569	16.38	107	599	0.08	0.39	2.27
Average			16.8			0.57		

Table 6.3: Quantifying the number of voxels displaying reliable stimulus preference to colour and/or motion stimulation within V2. The number of voxels encoding both motion and colour is shown as a percentage overlap. Two subjects were excluded due to low classification accuracy. The similarity of activity patterns produced by univariate and multivariate analysis techniques is expressed as the percentage overlap between the two. This is quantified in the final column.

To explore whether voxels with low z-score were driving the increase in the number of voxels demonstrating activity to more than one stimulus, the analysis was re-run with only the top 50% and 10% of responsive voxels. The values of 50% and 10% were selected to investigate the effect of removing low z-score contributors, while keeping sufficient voxels to be able to visualise stripe like patterns. Percentage overlap was calculated, and spatial maps were produced. This comparison revealed that the spatial overlap of the two conditions actually increased when the number of voxels included for analysis was reduced. On average, percentage overlap increased from 16.8% to 71.3% when 50% of voxels were included. This figure was reduced to 67% when 10% of voxels were included.




Figure 6.22: Flattened cortical maps of V2 for each hemisphere of Subjects 1-4, using 100%, 50% and 10% of available voxels. Voxels included have a weight value indicating significant ($\geq 1SD$) and reliable stimulus encoding voxels. Subject 2 is excluded on the basis that their classification accuracy was not significantly above chance. Voxels encoding colour (red), motion (cyan) or their conjunction (black) are shown.




Figure 6.23: Flattened cortical maps of V2 for each hemisphere of Subjects 5-8 using 100%, 50% and 10% of available voxels. Voxels included have a weight value indicating significant ($\geq 1SD$) and reliable stimulus encoding voxels. Subject 7 is excluded on the basis that their classification accuracy was not significantly above chance. Voxels encoding colour (red), motion (cyan) or their conjunction (black) are shown.

The weight structure thus far has provided information regarding voxels encoding stimulus specific activity when compared to a static field of chromatic Gabor patches. In this comparison,

where a voxel does not encode, for example, colour, it remains unclear if this is because the voxel does not respond in a stimulus specific way, or whether it actually encodes a preference for motion or form. A classification analysis was performed, to investigate this further, and distinguished colour stimulation periods from motion stimulation. This provided information as to whether any voxels positively encoded one stimulus specific condition, while responding negatively to another.

In 5 of 8 subjects, colour and motion stimulation periods were sufficiently distinguishable to be classified significantly above chance (Fig6.24). These were not the same subjects who showed high classification accuracy in the stimulus versus a static field of chromatic Gabor patch classification. The data of Subjects 1 and 3 were classified at significantly above chance when distinguishing between stimulus and a static field of chromatic Gabor patches, but were not distinguishable above the 95% confidence interval when classifying colour stimulation periods against motion stimulation periods. Conversely, the data of Subject 7, whose stimulus specific activity was not classifiable significantly above chance when compared to a static field of chromatic Gabor patches, were classifiable when stimuli of interest were compared to each other. In order to investigate the spatial organisation of voxels encoding this information, cortical maps were produced (Fig6.25). Significance was indicated by a heat scale. This method was chosen to identify whether the most significant voxels were arranged into the predicted stripe organisation.




Figure 6.24: The segregation of colour and motion activity patterns within V2, shown here for all subjects ($n=8$). Panel (A) depicts the stimulus conditions, with Gabor patches chromatically modulated shown on the left, and for motion shown on the right. Local motion within each Gabor patch could occur in any direction. The graph, (B) shows the accuracy with which a classifier could correctly label an instance of stimulation. Significance is indicated by asterisk marks, and the upper confidence interval ($p > 0.05$) is shown for each subject as a red line.




Figure 6.25: Flattened cortical maps of each hemisphere contain an illustration of activity within V2 of voxels encoding directly for colour or motion stimulation. Voxels significantly ($\geq 1SD$) responsive to colour (red) and motion (blue) are shown using a heat scale. Subjects 1-3 are excluded on the basis that classification accuracy in the multivariate analysis was not above the 95% confidence interval.

Although many of the clusters overlap with regions shown in Figures 6.13 and 6.14 to be stimulus encoding, new areas of active cortex are present in Figure 6.25.

6.6 Results 3: Segregation of colour, form and motion in other visual areas

In this section, I investigated the segregation and conjunction of colour, form and motion across the other regions of the visual cortex.

6.6.1 Segregation of colour, form and motion in V3

V3 is located anterior to V2 and receives direct input from this region. However, little is known regarding the functional organisation of colour, form or motion processing in V3. Repeating the process of classifying colour, form and motion with voxels of V3, a linear classifier was able to accurately distinguish colour from a static field of chromatic Gabor patches in 5 of 8 participants. Periods of stimulation containing radial form were distinguishable from a static field of chromatic Gabor patches in 3 of 8 participants, while motion stimulation was classifiable in 7 of 8 participants. This finding is similar to that of V2, in which colour was classifiable in 6 of 8 subjects, form was classifiable in 2 of 8 subjects and motion was classifiable in 6 of 8 subjects. However, unlike the results of V2, the subjects for whom colour and motion were classifiable were not consistent. The bias towards the classification of motion suggests that V3 may be more tuned to respond to motion stimulation than colour.




Figure 6.26: Results of a SVM classification of colour (A), form (B) and motion (C) against a control stimulus of static achromatic Gabor patches, in V3 for all subjects. Significance was estimated by comparing the classification accuracy to the distribution of accuracy results produced from a randomised data set. This is indicated by asterisk marks, while the 95% confidence interval ($p=0.05$) is shown for each subject as a red line.

6.6.2 Segregation of colour, form and motion in V4

There is evidence to suggest that V4 responds strongly to colour and form stimulation (Gallant et al. (2000); McKeefry and Zeki (1997); Bartels and Zeki (2000)). In order to test this, a linear classifier was trained to distinguish stimulation periods modulated to contain colour, form and motion. For 3 of 8 subjects, colour was distinguishable from a control stimulus of static achromatic Gabor patches. Stimulation periods modulated to contain radial form or local motion were accurately classified in only 2 of 8 participants. The ability to classify stimulation periods containing radial form has thus far been poor, and is continued here. The poor classification of colour stimulation periods within V4 was surprising, and suggests that the contribution of noise to the data set was high.




Figure 6.27: Results of a SVM classification of colour (A), form (B) and motion (C) in V4 for all subjects. Significance was estimated by comparing the classification accuracy to the distribution of accuracy results produced from a randomised data set. This is indicated by asterisk marks, while the 95% confidence interval ($p=0.05$) is shown for each subject as a red line.

6.6.3 Segregation of colour, form and motion in V5

As predicted, the classifier was able to detect periods of moving stimulation within V5 for 6 of 8 subjects. A further participant was just classifiable. Stimuli modulated to contain chromatic or radial form were only distinguishable from a static field of chromatic Gabor patches in 2 of 8 subjects. This result concurs with the results of the univariate analysis, in which V5 was found to strongly respond to motion stimulation.




Figure 6.28: Results of a SVM classification of colour (A), form (B) and motion (C) in V5 for all subjects. Significance was estimated by comparing the classification accuracy to the distribution of accuracy results produced from a randomised data set. This is indicated by asterisk marks, while the 95% confidence interval ($p=0.05$) is shown for each subject as a red line.

6.6.4 Investigating the spatial organisation of the segregation of colour, form and motion in V2, V3, V4 and V5

Nasr et al., recently provided evidence to suggest that the stripes characterised in V2 may also extend over the border and continue into V3 (Nasr et al. (2016)). In order to investigate the spatial organisation of weight vectors across the cortical surface spatial maps were produced. Classification was carried out for each area separately and used 100% of visually responsive voxels active above the standard threshold ($z=2.3$, $p<0.05$) to a static field of chromatic Gabor patches versus fixation condition. Voxels with high weight vectors ($\geq 1SD$) in response to colour or motion stimulation were projected onto the cortical surface for each visual area (Fig6.29).

At the confluence of V2 and V3, continuations of colour and motion activity can be observed for several subjects. Bands of voxels arranged in elongated stripes can be seen within V3 for Subjects 1, 3, 5 and 8. Similar to previously reported observations in V2, V3 also contains large clusters of voxels encoding colour and motion stimulation periods. V4 and V5 were predicted to preferentially encode colour and motion stimulation periods respectively.

There is limited agreement when comparing the accuracy with which a colour or motion stimulus could be decoded from a static field of chromatic Gabor patches, with the spatial maps. In the spatial map of Subject 1 (6.29, the activity within V5 is sparse, yet the classifier was able to detect motion periods from a static field of chromatic Gabor patches with nearly 90% accuracy (6.29. Similarly, Subject 3 seems to show considerable motion activity within V5, but was just classifiable above the 95% confidence interval in the same classification pairing.




Figure 6.29: Flattened cortical maps of V2, V3, V4 and V5 for each hemisphere using 100% of active voxels. Voxels included have a weight value indicating significant ($\geq 1SD$) and reliable stimulus encoding voxels. Subjects 2 and 7 are excluded on the basis that their classification accuracies were not above the 95% confidence interval. Voxels encoding colour (red), motion (cyan) or their conjunction (black) are shown.

		Percentage overlap per visual area							
		V2		V3		V4		V5	
Subject		Colour	Motion	Colour	Motion	Colour	Motion	Colour	Motion
	1	17.47	17.56	17.48	17.14	16.80	17.38	17.13	16.64
	2	-	-	-	-	-	-	-	-
	3	16.99	16.10	16.57	15.17	16.19	15.63	17.48	15.81
	4	16.52	17.60	15.32	16.65	15.64	15.40	17.36	16.53
	5	16.17	16.97	15.53	16.65	13.83	16.59	15.76	17.69
	6	16.44	15.95	16.02	15.54	13.60	16.49	17.46	16.86
	7	-	-	-	-	-	-	-	-
	8	13.45	17.52	12.26	17.52	12.75	18.01	18.31	18.51

Table 6.4: The percentage contribution of stimulus specific activity to each visual area. The total for each subject is calculated as the number of voxels active in response to colour and /or motion stimulation periods, and is expressed as a percentage of the total number of visually responsive voxels for each visual area.

The proportion of stimulus specific voxels contributing to the total number of active voxels was calculated for each visual area (Table 6.4). It was found that the responses to colour or motion stimulation did not support the respective predicted bias for colour and motion in V4 and V5.

The similarity between univariate and multivariate analysis was calculated, in order to examine the differences in stimulus specific activity each of visual area (Table6.5). The results of the univariate and multivariate analysis suggested that colour more reliably stimulated V4, and similarly, motion more reliably stimulated voxels within V5. On average, 28.25% of voxels encoding colour information in V4 were active in both univariate and multivariate analysis. In V5, 18.45% of voxels were found to be active in both univariate and multivariate analysis methods (Table6.5). This is in contrast to the response to colour stimulation within V5, where less than 1% of colour encoding voxels were found to repeat between univariate and multivariate techniques. This finding suggests that although noise contributions were high, within this, relevant and specific feature signals were present.

		Percentage activity per visual area							
		V2		V3		V4		V5	
Subject		Colour	Motion	Colour	Motion	Colour	Motion	Colour	Motion
	1	0.00	3.17	0.00	7.11	0.00	12.76	0.00	0.00
	2	-	-	-	-	-	-	-	-
	3	34.54	43.57	38.51	44.69	43.38	43.71	0.40	42.67
	4	0	0.67	0.09	0.86	0.00	2.05	0.00	0.00
	5	25.40	14.1	35.48	10.07	46.76	1.69	0.00	30.04
	6	12.5	2.71	17.74	4.68	79.00	0.89	0.00	37.99
	7	-	-	-	-	-	-	-	-
	8	0.39	2.27	0.09	0.39	0.36	0.00	0.00	0.00

Table 6.5: The percentage overlap between the results of univariate and multivariate analysis. The total for each subject is calculated as the number of voxels active in the univariate analysis which overlapped with the multivariate analysis, and is expressed as a percentage of the total number of visually responsive voxels for each visual area.

6.7 Discussion

6.7.1 Monocular signals are reliably detectable within V1 but functional organisation is not

The technique of multivariate analysis provided sufficient sensitivity to detect monocular stimulation within V1 (Fig6.6). This finding suggests that monocular signals were present within the data of Chapter 4, but were subtle and may have been indistinguishable from noise components. This is supported by the lack of functional organisation observable in the weight vector maps of voxels encoding ocular preference (Fig6.7). The clusters were, on average, 2.1mm wide, a value larger than neurophysiologically defined ODC (Adams et al. (2007)). The clusters were also only marginally elongated and the interlocking pattern of ocular dominance columns, for which evidence was first provided functionally by Cheng et al., was not clearly observable in any subject's spatial map (Cheng et al. (2001)). The sparsely organised clusters of high weight may represent part of an underlying functional organisation, but this is masked by voxels of lower weight. This finding highlights the benefit of considering the population of voxels as a whole, and suggests that while monocular signals are present, they are not strong enough to be visualised as part of a larger functional network.

This finding builds on the work conducted by Schwarzkopf et al., in which the group used the preference of cytochrome oxidase blobs for low spatial frequencies as a proxy to investigate monocular activity within V1 (Schwarzkopf et al. (2010)). The group hypothesised that a linear classifier would successfully distinguish the eye to which a low spatial frequency stimulus was presented, but would not detect this difference when high spatial frequencies were presented. High spatial frequencies are thought to stimulate the binocular population of V1 neurons, which are not known to demonstrate any regular spatial organisation. In contrast, the CO blobs are concentrated in the centre of the ocular dominance columns, and this spatial segregation of left and right eye CO blobs should facilitate the classifier. The groups hypothesis was supported by their findings of robust classification success of ocular input in V1 to low, but not high spatial

frequencies. The work conducted in this thesis extends the work conducted by Schwarzkof et al., in terms of spatial resolution. Schwarzkof et al., acquired their data at $1.5 \times 1.5 \times 1.5 \text{ mm}^3$ resolution, which they then smooth with a 4mm Gaussian smoothing kernel. In the work conducted for this thesis no spatial smoothing was applied, and therefore the relationship between the monocular signal, and the underlying functional organisation was maintained. This enabled the spatial organisation of these monocular signals to be investigated. The specificity of these monocular signals to V1 in this thesis is supported by the lack of detection of monocular signals within V2, despite robust classification of visual stimulation versus fixation (Fig6.6). This finding is novel and interestingly, Schwarzkopf et al., did detect monocular signals within V2, although this is not discussed further in their work. The population of neurons with monocular preference is not known to demonstrate any structured organisation (Roe (2003)). The detection of ocular presentation during low spatial frequencies, but not high spatial frequencies in V2 suggests that the classifier response may not be exclusively based on structured monocular organisation.

For all subjects, activity in response to right eye stimulation was larger than for the left. It could be suggested that this resulted from one eye receiving more visual stimulation. If the board were placed in such a way that a proportion of the visual stimuli became visible to both eyes then the additional information would fall upon the temporal retina of the contralateral eye. These voxels would be erroneously recorded as responsive to the contralateral eyes stimulation. However, in this experiment, the division of the visual field with solid black boards was complete. Additionally, a visual buffer zone was in place either side of the board in which no stimulation occurred. I suggest that conclusions on this apparent bias, drawn from the results of only three subjects, cannot be made. I also note that the bias towards a larger contralateral response to right eye stimulation than left is not shown for all subjects, and therefore could be random.

6.7.2 Stimulus specific signals of colour and motion are detectable within V2

In support of the findings made in Chapter 5, colour and motion stimulation periods produced sufficiently distinct activation patterns to be uniquely classified within V2 in 6 of 8 subjects. These results are in agreement with other investigations of colour and motion within V2, where the two types of stimulation were found to be uniquely detectable (Seymour et al. (2010, 2009)). In this chapter a greater proportion of voxels were found to respond to colour and motion stimulation within V2 than in Chapter 5. The multivariate analysis revealed that 16.17% and 16.95% of voxels responded to colour or motion stimulation respectively. While in Chapter 5, 3.49% of voxels responded to colour stimulation, and 3.04% were responsive to motion.

Also consistent with Chapter 5, form specific stimulation periods were undetectable from a static field of chromatic Gabor patches above the 95% confidence interval levels in 6 of 8 subjects. The inability of a high-powered analysis technique such as MVPA to detect form specific stimulation periods suggests that this stimulus was not optimal in activating the neural population of V2. Neural populations within V2 have been shown to respond to form stimuli (Sincich and Horton (2005)), but, their small receptive fields (0.5° - 2°) (Dumoulin and Wandell (2008)), combined with the global nature of the form condition may have, in this experiment, produced a weak response. Radial form has previously been used with some success to categorise form from control stimulation periods within V2 (Ostwald et al. (2008)). However, in doing so, the group noted that the perception of global form arose largely through the combinations of local orientation. In my stimulus design, these orientation cues may not have been sufficiently dense or have a long enough life span (20ms) to elicit a similar response.

6.7.3 Spatial maps of stimulus encoding voxels within V2 are inconclusive

Cluster width, elongation and orientation were calculated for all subjects. On average, both colour and motion clusters were found to be within the neurophysiological boundaries of V2

stripe width (Hockfield et al. (1990)), and were slightly elongated. The V2 stripes have been shown, via staining studies, to demonstrate consistent orientation, orthogonal to the border of V1-2 (Hockfield et al. (1990)). In this thesis, the orientation of the colour and motion clusters did not demonstrate this consistent trend. These results extend the findings of Seymour et al., who found clusters of colour and motion voxels within V2, but did not characterise the shape or orientation of the activity patterns (Seymour et al. (2009)).

Using MVPA, 80% more voxels were found to encode stimulus preference within V2 than with univariate analysis. Of these voxels, there was 16.8% overlap between stimulus conditions, compared to the 0.57% after univariate analysis. A certain number of these additional voxels may simply respond to the addition of visual information rather than to feature specific cues. Reducing the number of voxels to the top 50% and then 10% did not cause the response patterns of colour and motion to separate, and instead the degree of overlap increased. This increase may reflect the response preferences of the top 50% and 10% of voxels. Voxel number was restricted firstly by ranking voxels in descending order of z-score, after which the lower 50% were removed. Removing 50% of the voxels may have made it more difficult for the SVM to correctly identify stimulus specific activity. This is also true when including only the top 10% of voxels. The remaining voxels may be modulated by the addition of more, rather than specific visual information. This, combined with the loss of voxels to aid in feature discrimination, could be driving the increased percentage overlap reported here. This finding highlights how high z-scores are not necessarily indicative of feature specific activity, and how, with the addition of voxels, feature specific information is more likely to be identified.

The lack of a coherent functional organisation in these spatial maps is not necessarily surprising. Multivariate analysis examines small stimulus differences across a population of voxels. Individually, voxels may not show sufficient stimulus specific activation to be classified. The importance of interpreting the data as a population response could therefore explain the discrepancy between the lack of clear functional architecture (Figs 6.13-6.14), and high classification accuracy (Fig6.12).

If the high classification accuracy were simply to be explained as a population response, it

would be expected that individual voxels would not reliably have high weight values. However, as illustrated by the surface maps, many of the voxels are highly significant. This is confirmed by the coefficient of variation, measured in Figure 6.21, which showed that the dispersion of the distribution of mean weights was significantly greater than zero for more than 80% of the active voxels. Based on this evidence I suggest that this method of MVPA while useful for classification is not best suited to examine the spatial organisation of the classification weights, and that a technique such as Searchlight may be preferable.

In 5 of 8 subjects, colour and motion stimulation periods were distinguishable in a two-way classification analysis. Many voxels were shown to significantly encode a particular stimulus ($\geq 4SD$), which suggests that the SVM was not only using the addition of feature information to create accurate decision boundaries. The spatial patterns produced by these voxels were not clearly organised into the radial patterns of V2 stripes. However, looking at the most significant clusters, these do in some cases appear to arise from an elongated seed. Whether there is functional organisation underlying these robust activation patterns remains unclear.

6.7.4 The ability to detect stimulus preference in other visual areas is variable

In recent work, conducted by Nasr et al., stripes of voxels responding to chromatic modulation could be seen to extend from V2 into V3 (Nasr et al. (2016)). In the present work, examining the confluence of anterior V2 and ventral V3 continuations of colour and motion activity can be observed for several subjects. Bands of voxels arranged in elongated stripes can be seen within V3 for Subjects 1, 3, 5 and 8. Additionally, V3 contains large clusters of voxels encoding colour and motion stimulation periods in a similar way to observations made in V2.

Based upon previous work conducted by Seymour et al., V4 and V5 were predicted to preferentially encode colour and motion stimulation periods respectively (Seymour et al. (2010, 2009)). In the current investigation, colour was found to be distinguishable from a static field of chromatic Gabor patches in 3 of 8 subjects within V4. The response of V4 was found to be less

specific than predicted and voxels encoding motion stimulation were identified within V4 in all subjects. Additionally, there was low correspondence between the classification accuracy and the clarity of the activity maps. Motion stimulation periods were robustly detectable within V5 and could be classified in 7 of 8 participants. However, there was low correspondence between the classification accuracy and the clarity of the activity maps.

The segregation of colour and motion information across the ventral and dorsal streams of visual processing is less apparent following the application of MVPA. This lack of segregation is highlighted by the equal proportion of voxels encoding colour and motion stimulation within V4 and V5 respectively. It is possible that this lack of specificity reflects the classifier detecting “more” information i.e. the addition of a chromatic or motion element to the static, achromatic Gabor control stimulus.

However, when examining the reliability of voxel activity between univariate and multivariate techniques, those voxels encoding colour in V4 and motion in V5 were found to be more consistently active across the different analysis techniques. This finding suggests that the voxels encoding colour in V4 and motion in V5 do so with greater reliability. The dorsal/ ventral division of motion/colour respectively may be obscured by the increased number of voxels. The low prevalence of voxels exceeding the 95% confidence interval within V4 could be reflective of the definition of confidence in this investigation. Here, confidence was defined differently from the previous works. Seymour et al., define confidence as the likelihood that the result of a classification fell within the distribution of accuracies of all subjects (Seymour et al. (2009, 2010)). Here, confidence is based upon the likelihood that an individual's mean accuracy could have arisen from a scrambled data set when permuted 10000 times. This permutation analysis is not only more rigorous in terms of number of samples, but also provides a measure of the noise within each subject's data. When a “chance” distribution is not centred at 50% there are other factors influencing the decision. This is less clearly addressed by previous work, and here, highlights the possibility that noise may have played a confounding factor in the investigation of spatial relationships. In this case, the combined weight vectors of the voxel population may be more informative than the spatial organisation of the same population.

6.7.5 Future directions MVPA searchlight

In this investigation I used a linear SVM to investigate both the classification accuracy of stimulus specific conditions and to produce spatial maps of high weight vectors. The spatial maps of this analysis did not conclusively reveal the structured organisation of ocular dominance columns or V2 stripes. I propose that to further investigate this, the tool of searchlight analysis could be of use (Kriegeskorte et al. (2006)). Searchlight MVPA uses cluster-based machine learning algorithms. Neighbouring voxels are treated as containing related information, and features are pooled across a sphere of a specified size. Each voxel is visited in a similar manner and from this, maps of probabilities of spatial relationships are built. This analysis technique may be suitable for investigating hidden spatial correlations.

6.8 Conclusion

In summary, my findings suggest that stimulus specific features are detectable within both V1 and V2 at above chance levels. In cases where univariate analysis did not find stimulus specific signals it is possible to use the high-powered technique of MVPA to reveal hidden patterns. Whether these findings can be related to underlying functional architecture remains an open question. I suggest that the results of the univariate analysis of Chapter 5, although revealing fewer activity clusters, conformed more to the spatial pattern of the V2 stripe system. However, the findings of this investigation suggest that the noise contributions are potentially confounding in this work and, I suggest that this could be reduced by including further, or longer scan sessions.

Chapter 7

General discussion

7.1 Abstract

In this chapter, I review the major work conducted for this thesis. I discuss the implications of my findings and their contribution to the literature regarding human functional organisation within the visual cortex. I also address the limitations met within this thesis, and propose improvements to facilitate future work.

Based on the data collected for Chapter 4, I suggest that there is evidence in support of the functional segregation of monocular cues within V1. I also propose that the data collected for Chapter 5 provides evidence in favour of the parallel processing of colour and motion within V2. I argue that this information is best accessed through the application of advanced multivariate analysis techniques. Multivariate analysis is an alternative, high-powered and sensitive technique that can be used to interrogate noisy data sets. I conclude however, that this advantage does not necessarily extend to the investigation of spatial relationships.

The univariate analysis of the spatial relationship of colour and motion processing within V2 met with some success. On average, the width of the observed activity matched the neurophysiological estimates of the V2 stripes. However, I suggest that drawing firm conclusions and comparisons with animal literature, regarding the spatial arrangement of sub-millimeter functional activity, is not yet achievable with current fMRI technology.

I conclude by offering some proposals regarding future research, and stress the importance of continued investment in ultra-high resolution fMRI research.

7.2 Introduction

In this thesis I investigated the small-scale activity of functional networks within human visual cortex. I focused particularly on the monocular segregation of visual cues within V1, and the parallel processing streams of colour, form and motion within V2. It is proposed that these systems exist as part of a larger hierarchical visual processing network. This network makes use of regular and repeating functional units, that maintain localised circuitry and modular processing in order to minimise a cost function (Horton and Adams (2005); Reichl et al. (2012)). Possible cost functions have been suggested, focusing largely around the minimisation of wiring costs and stimulus coverage and continuity (Reichl et al. (2012)). The discovery of these units was hailed as a breakthrough and seemingly simplified the task of understanding cortical computation.

Until recently the small size of the functional activity of these networks has placed them beyond the investigatory scope of fMRI. However, developments in the field of fMRI sequence techniques, and the increasing availability of ultra-high field MRI facilities, have now created new possibilities for the investigation of small-scale functional activity in humans. The acquisition of sub-millimeter isotropic resolution data, however, remains at the limit of what is currently achievable.

I sought to overcome this challenge, and developed an ultra-high resolution fMRI protocol with which to investigate these small functional, modular units of visual cortex. I stipulated that the protocol must be capable of acquiring data with good SNR within a reasonable time frame and be tolerated by non-expert volunteers.

I divide this chapter into the following major sections:

- Major findings
- Future experiments
- Conclusions

7.3 Major findings

7.3.1 Monocular signals within V1 were detectable using multivariate analysis

Using both univariate and multivariate analysis techniques, monocular signals were investigated within V1. In Chapter 6 a linear classifier, trained to distinguish between periods of left and right eye stimulation, was successful in reliably classifying periods of monocular stimulation within V1. The classification accuracy ranged between 75%-81%.

These results extend the findings of Chapter 4, in which univariate analysis revealed monocular signals for all subjects, but these signals were not repeatable across sessions. To ensure that the low spatial repeatability did not result from mis-registration across scan sessions, I applied a 1mm smoothing kernel to the data. The application of spatial smoothing had been avoided in order to maximise the separation of neighbouring signals. A 1mm kernel was selected as a compromise, boosting SNR, reducing the potential for mis-registration, and remaining within the bounds of the target functional activity width. Following the application of the smoothing kernel the percentage of repeatable voxels increased only by 2.15%, suggesting that the low spatial repeatability did not result from mis-registration.

The low spatial repeatability of the monocular signals within V1, could be indicative of poor SNR. In such a case, the noise components may present as spurious active voxels, many of which would not repeat across sessions and therefore would lower the percentage repeatability across days. This effect should have been limited, but may not have been excluded by the pre-processing step of ICA. To explore this further I investigated whether monocular signals could also be recorded within V2. Within V2, the monocular cell population is not known to demonstrate any structured organisation, and therefore, the percentage of voxels demonstrating monocular preference would be predicted to be lower (Roe (2003)). In Chapter 4, I confirmed this prediction, and found that the percentage of monocular activity was significantly higher ($p < 0.05$) within V1 than V2.

On average 28.7% of voxels responded to either left or right eye stimulation within V1, compared to 11.3% of voxels within V2. The percentage of voxels with demonstrable monocular activity in V1 was lower than the 68.8% of grey matter voxels reported by Cheng et al., (Cheng et al. (2001)). However, Cheng et al., acquired a much smaller region of interest, consisting of just three slices of data. The group also used thicker slices, thereby boosting their SNR and, because of the smaller number of voxels, employed a voxel based statistical correction method. These more stringent statistical criteria in addition to the boost in SNR, reduced the contribution of noise within the data set. Another recent paper, which could be considered to be methodologically more similar to the work conducted in this thesis acquired only one session of ocular dominance data and so no measure of repeatability is available for comparison (Nasr et al. (2016)).

The larger monocular response within V1 compared to V2 suggests that the signals may indeed be part of the ocular dominance network. However, the 11.3% of V2 voxels demonstrating monocular preference required further consideration. Two possible explanations could be proposed for this finding. Firstly, in addition to true monocular signals, the data set could contain a high noise component, presenting in V2 as apparent monocular response. Alternatively, true structured monocular signals could also be present within V2. I addressed this question in Chapter 6, where I asked whether monocular stimulation periods could be accurately classified within V2 using machine learning techniques. I provided evidence to suggest that monocular stimulation periods could not be reliably detected within V2. The classifier performed with accuracy between 53-57%. This was found to be below the 95% confidence interval required to reject the null hypothesis, indicating that this result could not be distinguished from a result arising from a scrambled data set.

Taken together, these findings suggest that true monocular signals were specific to V1, but that a high noise component was present, and contributed to the low repeatability of monocular signals across sessions.

7.3.2 The organisation of monocular signals into ocular dominance columns was not supported by this study

Ocular dominance columns can be identified by their regular, repeating banded pattern of monocular preference. In humans these bands are between 0.8-1mm (Adams et al. (2007)), and are orientated perpendicular to the calcarine sulcus. In this thesis, I attempted to extend the work conducted by Cheng et al., and Yacoub et al., (Cheng et al. (2001); Yacoub et al. (2007)), by identifying ocular dominance columns across a wide region of V1, and in multiple subjects.

I found insufficient evidence to conclusively ascribe observed monocular signals to be part of the functional architecture of ODC. Despite high classification accuracy in decoding monocular signals within V1, the spatial maps produced in Chapter 6 did not reveal the consistent banding pattern of ODC. This finding was supported by the results of the univariate analysis in Chapter 4, where activity clusters did not conform to the predicted orientation or width of ODC. The lack of repeatable spatial organisation in Chapter 4 could be predicted by the finding of low spatial repeatability across sessions. However, the lack of coherent spatial organisation in Chapter 6, following the success of the classification analysis was somewhat surprising.

Based on these findings, I suggest that the strength of multivariate analysis lies in its population based approach. Multivariate classifiers analyse data sets with higher dimensionality than univariate methods, and are more sensitive to weak signals in noisy data sets. A high classification accuracy does not therefore imply that there is any spatial relationship within the data, only that when treated as a population, monocular signals can be identified.

7.3.3 Conclusions on ocular dominance columns

I conclude from the investigations of Chapters 4 and 6, that monocular signals were present within the data, but that these signals were not sufficiently robust within individual voxels to be examined spatially.

At the time of data collection for this project no published literature had shown evidence of the application of ultra-high resolution fMRI to the investigation the fine spatial signals of ODCs. The difficulty with which these small-scale functional patterns such as ODCs are localised, has led Kriegeskorte and Bandettini to suggest that univariate analysis may not be appropriate for analysing the massive data sets collected from ultra-high resolution investigations(Kriegeskorte and Bandettini (2007)). Traditional univariate analysis penalises fine grain signals, and seeks to identify smoother, wider reaching activity patterns . In investigations where the target activity is not so close to the spatial limit of the voxel size, this can often be compensated for by smoothing the data, thus improving the functional CNR. However, in cases such as ocular dominance columns, neighbouring signals are of interest, and even a small degree of smoothing can blur relevant signals into each other and prevent their distinction.

Previous successful investigations of ocular dominance within V1 have attempted to overcome these limitations by employing larger voxel volumes, and have acquired very few slices of data. These choices, combined with long scan periods of ~ 2 hours per session, boosted SNR and resulted in the successful differentiation of neighbouring monocular signals (Cheng et al. (2001); Yacoub et al. (2007)). In their recent investigation, Nasr et al., were the first to demonstrate regions of monocular preference across a wide area of V1 (Nasr et al. (2016)). The group acquired their data using $0.8 \times 0.8 \times 0.8 \text{ mm}^3$ voxels, which were 66% larger than the data collected for this thesis. The group also used long scan times, and only acquired data for one eye, of one subject, in one session. It therefore remains unclear how the “columns” of each eye interact and, whether this finding is repeatable across other participants.

I conclude this investigation of ODC by suggesting that it may not be possible, with current fMRI technology, to image ODC over a large area of V1.

7.3.4 Colour and motion signals were separable within V2, and conformed to neuroanatomical estimates of stripe width

The evidence of Chapters 5 and 6, suggest that colour and motion stimulation periods elicit distinct activation patterns within V2.

In Chapter 5, I found that within V2, 2.2% of voxels responded on average to colour stimulation, while 2.3% of voxels were active during motion periods. Of those voxels responsive to colour or motion stimulation, only 4.21% of voxels responded to both. This implies a high degree of separation. I further explored this finding in Chapter 6, where I used machine learning techniques to ask whether additional information could be extracted from the data set. The linear classifier, trained to distinguish periods of colour or motion stimulation from control, was successful in 6 of 8 participants. This finding is in agreement with previous work, in which colour and motion generated sufficiently distinct activation patterns to be distinguishable (Seymour et al. (2009)). The V2 stripe system is thought to spatially segregate the processing of colour, form and motion/disparity using a system of thick, thin and pale bands of functionally specific neuronal clusters. Evidence for feature segregation in humans has, until recently remained limited, largely due to the stripes small size. While Seymour et al., provided promising evidence to suggest that activity patterns in response to colour and motion were distinct, the spatial component of these signals was not accessible due to limitations based on spatial resolution.

V2 stripes demonstrate constant periodicity, running thick-pale-thin-pale. On average, the stripes are between 1-3mm in width, and are orientated perpendicular to the border of V1. The functional preferences of the stripes to motion/disparity, form or colour, has been measured *in-vivo* in the macaque using electrophysiological, and optical imaging techniques (Foster et al. (1985); Roe and Ts'o (1995); Sincich and Horton (2005)). In the work conducted for Chapter 5, I was able to identify distinct clusters of voxels that responded to either colour or motion stimulation, and conformed on average to the predicted width and spacing of V2 stripes. For some subjects, small patches of clear stripe like activity was visible (Fig7.1). This was supported

by the work conducted in Chapter 6, where again on average colour and motion clusters conformed to neuroanatomical estimates of stripe width. Measures of elongation also suggested that the clusters were between 1.2-1.8 times longer along the major axis than the minor.




Figure 7.1: A section of stripe like activity within dorsal V2, of the right hemisphere, of Subject 5. The activity map results from the conjunction of three sessions of data. The direction of activity can be observed to run anterior to posterior, orthogonal to the border of V1. Activity in response to colour and motion is shown in red and cyan respectively

In Figure 7.1, a small section of activity for one subject is shown, highlighting the elongated profile of some activity clusters. The activity in response to colour (shown in red) directly abuts the activity in response to motion (blue), and is orientated orthogonal to the border of V1. Further repeating patches of colour and motion specific activity continue laterally across V2. This activity was one of few examples of clearly visible stripes like patterns. The remainder of activity could be considered as patchy, partial reconstructions of stripes.

Compared to the work of Nasr et al., these spatial patterns are sparse. This can be partially explained by the reduced visual field of view available to participants due to head coil restrictions within the scanner. Within the stimulated area, I explored whether additional spatial information could be gained from examining the spatial maps of the multivariate analysis, a technique associated with greater sensitivity. The maps produced using this technique contained a much larger percentage activity (80%) than the univariate analysis, but percentage overlap was higher (16.8%, compared to 0.57%), and no greater amount of stripe like activity was apparent. No observed correlation was present between subjects demonstrating high classification accuracy, and those with more stripe like patterns.

I suggest that linear MVPA analysis may not be suitable for extracting the spatial information of a data set. This is particularly clear when examining the control conditions of colour and motion in V4 and V5 respectively. These regions are known to demonstrate preferences for these stimuli types, and are considered to be key members of the ventral and dorsal processing streams respectively (McKeefry and Zeki (1997); Bartels and Zeki (2000); Gallant et al. (2000); Kolster et al. (2010); McKeefry et al. (1997)). I selected these areas as markers of analysis success. In the univariate analysis, the dorsal/ventral pathways for motion and colour were maintained and clear. However, within the multivariate maps, this distinction of activity pathways was not as apparent.




Figure 7.2: A comparison of activity along the ventral stream of the right hemisphere of Subject 5 using univariate and multivariate methods. Activity in response to colour and motion is shown in red and cyan respectively. The overlap between colour and motion stimulation is shown in black. Borders between visual areas are shown by dotted white lines.

In Figure 7.2 it is clear that although more activity is present following multivariate analysis, this activity does not add to the stripe like patterns observed in the univariate analysis. Instead, proportion of overlap between the two conditions is increased, resulting in a loss of clarity in the activity in response to colour in the ventral stream.

I suggest that the univariate and multivariate analysis techniques both support the segregation of motion and colour signals within V2. The univariate analysis provided more useful spatial information, while the multivariate analysis confirmed that for 6 of 8 subjects, the signals in response to colour and motion were sufficiently distinct to be classified.

7.3.5 A stripe system for V3

For several subjects, colour and motion stimulation elicited spatially segregated activity within V3.

Initial consideration was given to the possibility that these patterns could have followed the path of blood vessels. Firstly I examined the univariate analysis volumes and found no large clusters of voxels active around blood vessels. Secondly, I concluded that by sampling data from between the mid-grey and white matter I was able to avoid the large surface vessels.

The functional organisation of V3 is poorly explored in non-human models, largely due to the methodological challenges of accessing the region, which is highly convoluted and thin. Some evidence has been presented in favour of some functional organisation of stripe like nature within V3 of the macaque (Sincich et al. (2003); Tootell et al. (2004)). However, little information is available for similar functional divisions in humans. In their 2015 work, Goncalves et al., provided evidence for a functional network for disparity processing within V3A (Goncalves et al. (2015)). Further work in favour of functional segregation within V3 was presented by Nasr et al., in which, columns of colour and disparity responsive voxels were found to cross the V2-V3 border (Nasr et al. (2016)). A similar finding was observed in this thesis, where colour and motion responsive voxels, organised in stripe like patterns, could be seen to cross the V2-V3 border. This finding was presented at OHBM 2015.




Figure 7.3: A section of stripe like activity, extending across the border of V2 and into V3. Here, two examples are shown: the left hemisphere of ventral V2 of Subject 3, and dorsal V2 of the right hemisphere of Subject 5. The direction of activity can be observed to run anterior to posterior away from the border of V1. Activity in response to colour and motion is shown in red and cyan respectively

I propose that these stripes may present a continuation of the functional segregation of colour and motion/disparity within V2. These stripes appear to be similar in their spatial profile as those observed within V2. Work conducted to investigate potential feature segregation in humans *in-vivo* would be an interesting extension to this project.

7.4 Future experiments

7.4.1 MR protocol development

This thesis has pushed the boundaries of achievable functional spatial resolution, and has relied upon the development of a novel ultra-high resolution protocol. In collaboration with members of the FMRI team, (detailed in Chapter 3) I designed and tested an ultra-high resolution fMRI protocol. This was subsequently updated in a review process that took place after the data collection of Chapter 4.

At high resolution, the relative contribution of noise within a single voxel rises. The associated reduction in SNR can be further added to by certain acceleration techniques. The protocol developed for this thesis employed partial Fourier (PF) and parallel imaging. PF utilises the conjugate symmetry of k-space to speed acquisition. However, small reconstruction errors due to signal dephasing and motion introduce signal blur across phase-encode direction of the voxel. Partial parallel imaging, in which lines or planes of k-space are skipped and zero filled, results in a reduced SNR caused by a loss of data point information (Heidemann et al. (2012)). Ordinarily, such contributions could be reduced by smoothing across several voxels. However, in this thesis the difference between the target functional activity, and the spatial resolution of the data was not sufficiently different for smoothing to be applied.

The process of protocol development highlighted the disparity between effective and nominal resolution of the acquired fMRI data. The nominal resolution describes the information dialled into the scanner, while the effective resolution takes into account various sources of blur. This often means that the effective resolution of the data is lower than supposed.

For this thesis a read out time close to the $T2^*$ of the underlying grey matter was selected. Shortening the read out time reduces the inter-voxel dephasing and can limit the disparity between these two values. This technique has been employed before with success (Yacoub et al. (2005)). However the blurring introduced by the application of partial Fourier reduced the actual spatial resolution of the data, potentially beyond the limit with which ODC could be

reliably imaged. After analysing the data of Chapter 4, I concluded that this difference may have contributed to the difficulty in identifying ODC within V1. Following the review process, I presented evidence to show that decreasing the field of view and TR reduced the difference between the effective and nominal resolution by 65%. Despite this change, however, the SNR of the data may not have been sufficient to robustly identify sub-millimeter activity.

Alternative approaches to ultra-high resolution data acquisition include the measurement of cerebral blood volume (CBV) and blood flow (CBF) in combination with traditional BOLD measures. The combination of these methodologies may permit the stratification of the BOLD response both temporally, i.e. to investigate the initial dip, known to be more local to the sites of neuronal activity, and spatially i.e. to localise and exclude the contribution of vessels.

Two alternative methods to consider are the 3D Gradient and Spin Echo method (3D-GRASE), and 2D Spin Echo (SE). In their comparison of 3D-GRASE and 2D SE-EPI data acquisition, Kemper et al., recorded 60% more blurring along the second phase-encode direction in the 3D sequence than the 2D (Kemper et al. (2015)). Temporal SNR, meanwhile was higher in the 3D sequence, and CNR was comparable between the two. Spin-echo sequences have been championed by Yacoub et al., in their high resolution acquisition of ODC (Yacoub et al. (2007)). The group note that data acquired using gradient echo (GE), while generally higher in its CNR, is more sensitive to draining veins than SE. The group conclude by suggesting that SE, which acquires more spatially specific and repeatable data, may be better placed to image small-scale functional activity (Yacoub et al. (2007); Kemper et al. (2015)). Kemper et al., suggest that the GRASE imaging, which combines the approaches of GE and SE could offer a potential way to combine the benefits of the two techniques (Kemper et al. (2015)). Finally, with better hardware and coil design, alternative combinations of acceleration parameters may also be considered.

7.4.2 Layer specific analysis

The spatial specificity afforded by ultra-high resolution fMRI, has brought human *in-vivo* functional investigations, one step closer to probing the functional circuitry of the brain.

A particularly exciting new application of ultra-high resolution fMRI techniques is the potential for the investigation of laminar specific activity. The human visual cortex is around 2.5mm in thickness (Dougherty et al. (2003)). Ordinarily, using voxels typical of current fMRI analysis, between 1.5-3mm isotropic (Goense et al. (2016)), one or two voxels will span the cortical depth. These voxels will almost certain contain signals from outside the grey matter, and the signal within these voxels will therefore be affected by partial voluming effects due to the high degree of curvature within the visual cortex. As a result, these voxels are unable to provide information regarding the distinction between afferent and efferent activity patterns. Characterising the laminar profile of the activity patterns is also of interest for investigating how CVB and CBF differ across the layers of cortex. At increasing field strength the blood T2 and T2* are shortened, and the relative sensitivity to veins, parenchyma and arteries is different. In addition to this effect, the orientation of vessels to the main magnetic field enhances the vessels susceptibility to signal loss. These considerations raise the question of whether the haemodynamic profile of voxels of differing cortical depth should be modeled differently (Goense et al. (2016)).

In Chapters 4 and 5 I discussed how features including ODC and the V2 stripes are most prominent in their input layer. In V1 the strongest monocular activity is observed within Layer 4C - a feature that is maintained to a lesser extent through the cortical depth. The connections of V2 are more complex, and functional specialisation has been hypothesised to change as a function of cortical depth. Layers 3A and 3B are thought to be the most specific in their feature preference, while Layers 2 and 6 more commonly respond to features not associated with their stripe type (Shipp and Zeki (2002)).

The difference in specificity as a function of cortical depth has been linked to feedback mechanisms of higher visual areas. I suggest that future high resolution work would benefit from attempting to spatially segregate these signals within V1 and V2 for two reasons. Firstly, as already discussed in Chapter 4, monocular signals are easily overpowered by noise. Noise can arise from many sources, but restricting analysis to the most specific layer (4C), would increase the specificity of the activity. Secondly, as I have suggested in Chapter 5, the segregation of

input and output layers in V2 could form part of a feedback system, indicative of the process of feature binding (Shipp et al. (2009)). This is both interesting and potentially confounding. When seeking to investigate feature specific signals the potential contribution of layers receiving horizontal connections from other stripes, or feedback from higher visual areas, may blur the signal from feature specific layers.

Layer specific investigations have already yielded promising results. In their investigation of contextual feedback in V1, Muckli et al., partitioned the cortical ribbon into 6 layers (Muckli et al. (2015)). The group used $0.8 \times 0.8 \times 0.8 \text{ mm}^3$ resolution voxels and MVPA techniques, which I have shown afford good sensitivity to weak signals. Within V2, Nasr et al., highlighted the importance of sampling activity patterns from mid to deep portions of the cortical depth. Superficial layers are known to be more affected by draining veins, which are likely to blur and add noise to signals (Nasr et al. (2016)). In the work conducted for this thesis, the cortical depth was collapsed between the mid and the deepest cortical depth to avoid the contribution of draining veins.

I conclude by suggesting that investment in layer specific signal processing would provide high yields for the spatial localisation of feature specific processing.

7.4.3 MVPA - searchlight

Ultra-high resolution data sets pose a challenge for traditional univariate analysis techniques. Small voxels are effected to a greater extent by the blurring effect of fMRI acquisition techniques and by haemodynamic processes. Additionally, the greater number of statistical samples may not be best interpreted by traditional multiple comparison analysis methods (Kriegeskorte and Bandettini (2007)).

Information theory based analysis is an increasingly popular method with which to interrogate large, noisy information-rich data sets. In information theory analysis, populations of voxels are treated as a macro-code, and local signals instead of being penalised, as is often the case

in activation-analysis techniques, are explored. I employ one variant of this technique in this thesis known as MVPA. By doing so I was therefore successful in segregating the previously indistinguishable signals of monocular activity in V1.

In this thesis, I decoded the population based response but was not able to extract meaningful information regarding the spatial components of the data. An alternative method of MVPA known as searchlight, examines local correlations between voxels by placing spherical “searchlights” on the centre of each voxel (Kriegeskorte et al. (2006)). Spatial statistical maps are then produced and express the relationship between each voxel time-course and that of its neighbours (Etzel et al. (2013)). This technique has the potential to extend the work conducted in Chapter 6, and may be more successful in extracting spatial signals from the data sets of both Chapters 4 and 5.

7.5 Conclusions

fMRI is a powerful tool when used to investigate the representation of visual information in the brain. At high field SNR, which is key to decoding the information held within voxels, is increased linearly. This linear increase becomes supralinear when taking into account the additional signal contribution from the microvasculature ((Yacoub et al., 2001)). This enables voxel volume to be reduced. However as spatial resolution increases, so do the effects of technical challenges such as sensitivity to motion, the blurring effects of acceleration techniques and the susceptibility to distortion. I conclude that at the present time, without long scan durations and expert volunteers, it is not possible to acquire robust and repeatable functional correlates of small-scale activity within the human visual cortex. However, I suggest that my results, particularly with regard to the segregation of colour and motion within V2, offer a promising indication of the potential application of sub-millimeter fMRI.

Chapter 8

Appendix



[01/04/13]

Letter to Volunteers
Version 1, 1st May 2013

Re: Investigation of visual cortical structures with high field fMRI

CUREC Ref:

We would like to invite you to take part in a research project. This letter provides some information to help you decide whether to do so. Please take time to read this information carefully and discuss it with friends, family or your GP if you wish. If there is anything that you do not understand, or if you would like more information, please ask us. Please take time to consider whether you wish to take part.

What is the purpose of the research?

The way in which we combine the input from our two eyes to view the world as one complete picture is of great interest to researchers. By studying how this process occurs in healthy volunteers we hope to be able to gain information about what happens in people who do not have good vision in both eyes. In the first stages of visual processing the input from each eye remains separate and is segregated into columns within an area of the brain called the primary visual cortex. Although we know this type of columnar organization exists, there are few studies that have imaged these columns in living human brains. Using our **3 or 7** Tesla Magnetic Resonance Imaging (MRI) scanner we can see the brain in action at very high resolution. By looking at the brain response to each eye separately we hope to show these ocular dominance columns in the healthy brain and further understand how the brain puts together the information from the two eyes which can impact processes such as depth perception.

Why have I been invited?

You have been invited to take part in this research because you are healthy, aged 18-60, have good binocular vision.

Do I have to take part?

No. It is up to you to decide if you want to take part in this study. We will describe the study and go through this information sheet with you to answer any questions you may have. If you agree to take part, we will ask you to sign a consent form and will give you a copy for you to keep. However, you would still be free to withdraw from the study at any time, without needing to give a reason. This would not affect the standard of care you would receive.

What will the study involve?

The study will take place at the FMRI Centre at the John Radcliffe Hospital. The study would involve having a series of magnetic resonance scans over a period of 45-60 minutes. We estimate that the whole visit should last around 90 minutes. You may also be asked to attend further MRI sessions, up to four times. To undergo magnetic resonance scanning you would be asked to lie still on a table inside the MRI scanner. During the scan you will be asked to look at pictures on a screen behind your head keeping your eyes very still, and also to indicate by pressing a button on a box in your hands if you have noticed something change in the picture. These picture sections will occur in short blocks of around 5 minutes and you will be able to rest your eyes in between blocks. For some periods of the scan (around 15 minutes) you will be allowed to lie still and relax with your eyes closed. You will be given specific instructions before each scan so that you know what you need to do.

If you agree, we would ask you to come to the FMRI Centre for up to four visits. There is an area near the scanner where an accompanying person can wait. On arrival, one of our research team would meet you to describe what is to be done and answer any questions you may have. We would also ask you to sign a consent form, of which you would be given a copy for you to keep. However, you would still be free to leave the study at

FMRIB Centre,
John Radcliffe Hospital,
Oxford,
OX3 9DU



Prof Andrew Parker
Direct Line: 01865-272504
E-Mail: andrew.parker@dpag.ox.ac.uk

VOLUNTEER CONSENT FORM

Version 2, 24 May 2012

Investigation of visual cortical structures with high field fMRI

CUREC ref:

*Please initial
each box*

- | | | |
|---|---|--------------------------|
| 1 | I confirm that I have read and understand the information sheet dated _____ for the above study. I have had the opportunity to consider the information, ask questions and have had these answered satisfactorily. | ☐ |
| 2 | I understand that my participation is voluntary and that I am free to withdraw at any time, without giving any reason, without penalty or affecting my legal rights. | ☐ |
| 3 | I understand that personal data collected during the study will be securely held in lockable cabinets or anonymised on computers and used solely for the purpose of contacting me about this or future studies, or for monitoring or audit of the study by designated individuals from the University of Oxford. I give permission for these individuals to see my personal data. | ☐ |
| 4 | I agree for data collected in this study to be given to researchers, including those working outside of the EU, to be used in other research studies. I understand that any brain images of me that leave the Centre will be fully anonymised so that I cannot be identified. | ☐ |
| 5 | I understand that if I have any concerns or complaints about the study, I can contact the Principal Investigator or the Research Sponsor (Clinical Trials and Research Governance) who is not directly involved in this project. | ☐ |
| 6 | I understand that this project has been reviewed by, and received ethics clearance through, the University of Oxford Central University Research Ethics Committee | ☐ |
| 7 | I understand that this is a research scan and is not useful for diagnosis. | ☐ |
| 8 | I agree to take part in the above study. | ☐ |

Name of Volunteer	Date	Signature

Name of person taking consent	Date	Signature



VOLUNTEER 7T MRI SAFETY SCREENING FORM

Volunteer name: _____

Date of birth: _____ Weight: _____ kg Height: _____ m

Please carefully check the following. Some items can interfere with MR examinations, and may be hazardous to your safety. Your answers will be kept strictly confidential. Clearly mark your answer with a circle.

IF YOU HAVE ANY QUESTIONS THEN PLEASE ASK US BEFORE YOUR SCAN			
Do you have a heart pacemaker or pacing wires?	YES	NO	
Have you had any heart surgery?	YES	NO	
Have you had any surgery to your head (including eyes/ears/brain), neck or spine?	YES	NO	
Do you have any implanted devices (e.g. programmable hydrocephalus shunt, nerve stimulator, cochlea implant, aneurysm clip)?	YES	NO	
Have you had any operation involving metallic pins / plates / screws / wires?	YES	NO	
Have you had any other surgical procedure of any kind? (Please write them here)	YES	NO	
Have you ever sustained any injuries involving metal to the eyes or other part of the body?	YES	NO	
Have you ever had a serious accident (e.g. road traffic accident, explosion injury, shooting, shrapnel injury)?	YES	NO	
Have you ever had a fit or blackout, or do you suffer from epilepsy or diabetes?	YES	NO	
Do you have any of the following (if yes please circle):			
Dentures with metal	Hearing aid	Artificial Limb or Prosthesis	
Body piercing/jewellery	Nitro patch	Tattoos	
FOR WOMEN OF CHILDBEARING AGE:	Do you have an IUD (coil)?	YES	NO
	Is there any chance that you could be pregnant? It is essential that you do not enter the magnet room if there is a chance you may be.	YES	NO
Do you understand that this is a research scan and is not useful for diagnosis?	YES	NO	
Have you removed your watch, bankcards, spectacles, hearing aids, keys, coins, jewellery, hairgrips and eye make-up?	YES	NO	


IMPORTANT: NO METAL OBJECTS TO BE TAKEN INTO THE MAGNET ROOM


Volunteer/Guardian signature _____ Date of study _____

Screened by (Print name) _____ Signature _____

Version August 2011

Date of Next Review: August 2012








Figure 8.1: Maps of spatial repeatability across sessions for V2 were created by projecting single session data onto the cortical surface for each subject. The conjunction of these sessions is shown first for each hemisphere for Subjects 2, 4, 6, 7 and 8, and is followed by subsequent sessions in order of their acquisition.

References

- Adams, D. L. and Horton, J. C. (2003). Capricious expression of cortical columns in the primate brain. *Nature neuroscience*, 6(2):113–114.
- Adams, D. L. and Horton, J. C. (2006). Ocular dominance columns in strabismus. *Visual neuroscience*, 23(5):795.
- Adams, D. L. and Horton, J. C. (2009). Ocular dominance columns: enigmas and challenges. *The Neuroscientist*, 15(1):62–77.
- Adams, D. L., Sincich, L. C., and Horton, J. C. (2007). Complete pattern of ocular dominance columns in human primary visual cortex. *The Journal of Neuroscience*, 27(39):10391–10403.
- Adams, D. L. and Zeki, S. (2001). Functional organization of macaque v3 for stereoscopic depth. *Journal of Neurophysiology*, 86(5):2195–2203.
- Amano, K., Wandell, B. A., and Dumoulin, S. O. (2009). Visual field maps, population receptive field sizes, and visual field coverage in the human mt+ complex. *Journal of neurophysiology*, 102(5):2704–2718.
- Andrews, T. J., Halpern, S. D., and Purves, D. (1997). Correlated size variations in human visual cortex, lateral geniculate nucleus, and optic tract. *The Journal of Neuroscience*, 17(8):2859–2868.
- Arthurs, O. J. and Boniface, S. (2002). How well do we understand the neural origins of the fmri bold signal? *Trends in neurosciences*, 25(1):27–31.
- Attwell, D. and Iadecola, C. (2002). The neural basis of functional brain imaging signals. *Trends in neurosciences*, 25(12):621–625.
- Backus, B. T., Fleet, D. J., Parker, A. J., and Heeger, D. J. (2001). Human cortical activity correlates with stereoscopic depth perception. *Journal of neurophysiology*, 86(4):2054–2068.
- Balaram, P., Young, N. A., and Kaas, J. H. (2014). Histological features of layers and sublayers in cortical visual areas v1 and v2 of chimpanzees, macaque monkeys, and humans. *Eye and brain*, 2014(6 Suppl 1):5.

- Bartels, A. (2009). Visual perception: converging mechanisms of attention, binding, and segmentation? *Current Biology*, 19(7):R300–R302.
- Bartels, A. and Zeki, S. (2000). The architecture of the colour centre in the human visual brain: new results and a review. *European Journal of Neuroscience*, 12(1):172–193.
- Beckett, A., Peirce, J. W., Sanchez-Panchuelo, R.-M., Francis, S., and 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.
- Bridge, H. and Parker, A. J. (2007). Topographical representation of binocular depth in the human visual cortex using fmri. *Journal of vision*, 7(14):15–15.
- Briggs, F. and Usrey, W. M. (2007). A fast, reciprocal pathway between the lateral geniculate nucleus and visual cortex in the macaque monkey. *The Journal of neuroscience*, 27(20):5431–5436.
- Burkhalter, A. and Bernardo, K. L. (1989). Organization of corticocortical connections in human visual cortex. *Proceedings of the National Academy of Sciences*, 86(3):1071–1075.
- Buxton, R. B., Uludağ, K., Dubowitz, D. J., and Liu, T. T. (2004). Modeling the hemodynamic response to brain activation. *Neuroimage*, 23:S220–S233.
- Callaway, E. M. (2005). Structure and function of parallel pathways in the primate early visual system. *The Journal of physiology*, 566(1):13–19.
- Carlson, T. A. (2014). Orientation decoding in human visual cortex: new insights from an unbiased perspective. *The Journal of Neuroscience*, 34(24):8373–8383.
- Chang, C.-C. and Lin, C.-J. (2011). Libsvm: a library for support vector machines. *ACM Transactions on Intelligent Systems and Technology (TIST)*, 2(3):27.
- Cheng, K., Waggoner, R. A., and Tanaka, K. (2001). Human ocular dominance columns as revealed by high-field functional magnetic resonance imaging. *Neuron*, 32(2):359–374.
- Constable, R. T., Kennan, R. P., Puce, A., McCarthy, G., and Gore, J. C. (1994). Functional nmr imaging using fast spin echo at 1.5 t. *Magnetic resonance in medicine*, 31(6):686–690.
- Cox, D. D. and Savoy, R. L. (2003). Functional magnetic resonance imaging (fmri) brain reading: detecting and classifying distributed patterns of fmri activity in human visual cortex. *Neuroimage*, 19(2):261–270.
- Cumming, B. G. and Parker, A. J. (1999). Binocular neurons in v1 of awake monkeys are selective for absolute, not relative, disparity. *The Journal of neuroscience*, 19(13):5602–5618.

- Dale, A. M., Fischl, B., and Sereno, M. I. (1999). Cortical surface-based analysis: I. segmentation and surface reconstruction. *Neuroimage*, 9(2):179–194.
- de Haan, E. H. and Cowey, A. (2011). On the usefulness of what and where pathways in vision. *Trends in cognitive sciences*, 15(10):460–466.
- DeYoe, E. A., Carman, G. J., Bandettini, P., Glickman, S., Wieser, J., Cox, R., Miller, D., and Neitz, J. (1996). Mapping striate and extrastriate visual areas in human cerebral cortex. *Proceedings of the National Academy of Sciences*, 93(6):2382–2386.
- DeYoe, E. A. and Van Essen, D. C. (1985). Segregation of efferent connections and receptive field properties in visual area v2 of the macaque. *Nature*, 317(6032):58–61.
- Dougherty, R. F., Koch, V. M., Brewer, A. A., Fischer, B., Modersitzki, J., and Wandell, B. A. (2003). Visual field representations and locations of visual areas v1/2/3 in human visual cortex. *Journal of Vision*, 3(10):1–1.
- Downing, P., Chan, A.-Y., Peelen, M., Dodds, C., and Kanwisher, N. (2006). Domain specificity in visual cortex. *Cerebral cortex*, 16(10):1453–1461.
- Dumoulin, S. O., Bittar, R. G., Kabani, N. J., Baker, C. L., Le Goualher, G., Pike, G. B., and Evans, A. C. (2000). A new anatomical landmark for reliable identification of human area v5/mt: a quantitative analysis of sulcal patterning. *Cerebral Cortex*, 10(5):454–463.
- Dumoulin, S. O. and Wandell, B. A. (2008). Population receptive field estimates in human visual cortex. *Neuroimage*, 39(2):647–660.
- Duong, T. Q., Yacoub, E., Adriany, G., Hu, X., Uurbil, K., and Kim, S.-G. (2003). Microvascular bold contribution at 4 and 7 t in the human brain: Gradient-echo and spin-echo fmri with suppression of blood effects. *Magnetic resonance in medicine*, 49(6):1019–1027.
- Ekstrom, A. (2010). How and when the fmri bold signal relates to underlying neural activity: the danger in dissociation. *Brain research reviews*, 62(2):233–244.
- El-Shamayleh, Y., Kumbhani, R. D., Dhruv, N. T., and Movshon, J. A. (2013). Visual response properties of v1 neurons projecting to v2 in macaque. *The Journal of Neuroscience*, 33(42):16594–16605.
- Engel, S. A., Glover, G. H., and Wandell, B. A. (1997). Retinotopic organization in human visual cortex and the spatial precision of functional mri. *Cerebral cortex*, 7(2):181–192.
- Etzel, J. A., Zacks, J. M., and Braver, T. S. (2013). Searchlight analysis: promise, pitfalls, and potential. *Neuroimage*, 78:261–269.

- Felleman, D. J., Lim, H., Xiao, Y., Wang, Y., Eriksson, A., and Parajuli, A. (2014). The representation of orientation in macaque v2: Four stripes not three. *Cerebral Cortex*, page bhu033.
- Felleman, D. J. and Van Essen, D. C. (1987). Receptive field properties of neurons in area v3 of macaque monkey extrastriate cortex. *Journal of Neurophysiology*, 57(4):889–920.
- Felleman, D. J. and Van Essen, D. C. (1991). Distributed hierarchical processing in the primate cerebral cortex. *Cerebral cortex*, 1(1):1–47.
- Fischl, B., Sereno, M. I., and Dale, A. M. (1999). Cortical surface-based analysis: Ii: inflation, flattening, and a surface-based coordinate system. *Neuroimage*, 9(2):195–207.
- Foster, K., Gaska, J. P., Nagler, M., and Pollen, D. (1985). Spatial and temporal frequency selectivity of neurones in visual cortical areas v1 and v2 of the macaque monkey. *The Journal of physiology*, 365(1):331–363.
- Gallant, J. L., Shoup, R. E., and Mazer, J. A. (2000). A human extrastriate area functionally homologous to macaque v4. *Neuron*, 27(2):227–235.
- Gegenfurtner, K. R., Kiper, D. C., and Fenstemaker, S. B. (1996). Processing of color, form, and motion in macaque area v2. *Visual neuroscience*, 13(01):161–172.
- Gegenfurtner, K. R. and Sharpe, L. T. (2001). *Color vision: From genes to perception*. Cambridge University Press.
- Goense, J., Bohraus, Y., and Logothetis, N. K. (2016). fmri at high spatial resolution: Implications for bold-models. *Frontiers in Computational Neuroscience*, 10.
- Goncalves, N. R., Ban, H., Sánchez-Panchuelo, R. M., Francis, S. T., Schluppeck, D., and Welchman, A. E. (2015). 7 tesla fmri reveals systematic functional organization for binocular disparity in dorsal visual cortex. *The Journal of Neuroscience*, 35(7):3056–3072.
- Grinvald, A., Bonhoeffer, T., Vanzetta, I., Pollack, A., Aloni, E., Ofri, R., and Nelson, D. (2004). High-resolution functional optical imaging: from the neocortex to the eye. *Ophthalmology Clinics*, 17(1):53–67.
- Harvey, B. M. and Dumoulin, S. O. (2011). The relationship between cortical magnification factor and population receptive field size in human visual cortex: constancies in cortical architecture. *The Journal of Neuroscience*, 31(38):13604–13612.
- Haynes, J.-D. and Rees, G. (2006). Decoding mental states from brain activity in humans. *Nature Reviews Neuroscience*, 7(7):523–534.

- Heidemann, R. M., Ivanov, D., Trampel, R., Fasano, F., Meyer, H., Pfeuffer, J., and Turner, R. (2012). Isotropic submillimeter fmri in the human brain at 7t: Combining reduced field-of-view imaging and partially parallel acquisitions. *Magnetic resonance in medicine*, 68(5):1506–1516.
- Heidemann, R. M., Trampel, R., Ivanov, D., Fasano, F., Pfeuffer, J., and Turner, R. (2009). Isotropic sub-millimeter fmri of v5 in human at 7t. *Neuroimage*, 47:S196.
- Henriksson, L., Mur, M., and Kriegeskorte, N. (2015). Faciotopy a face-feature map with face-like topology in the human occipital face area. *Cortex*, 72:156–167.
- Hockfield, S., Tootell, R., and Zaremba, S. (1990). Molecular differences among neurons reveal an organization of human visual cortex. *Proceedings of the National Academy of Sciences*, 87(8):3027–3031.
- Horton, J. and Hedley-Whyte, E. T. Mapping of cytochrome oxidase patches and ocular dominance columns in human visual cortex. *Philosophical Transactions of the Royal Society of London. Series B, Biological Sciences*.
- Horton, J. and Hocking, D. (1998). Effect of early monocular enucleation upon ocular dominance columns and cytochrome oxidase activity in monkey and human visual cortex. *Visual neuroscience*, 15(02):289–303.
- Horton, J. C. and Adams, D. L. (2005). The cortical column: a structure without a function. *Philosophical Transactions of the Royal Society of London B: Biological Sciences*, 360(1456):837–862.
- Horton, J. C., Dagi, L. R., McCrane, E. P., and de Monasterio, F. M. (1990). Arrangement of ocular dominance columns in human visual cortex. *Archives of Ophthalmology*, 108(7):1025–1031.
- Horton, J. C. and Hocking, D. R. (1996). Intrinsic variability of ocular dominance column periodicity in normal macaque monkeys. *The Journal of neuroscience*, 16(22):7228–7339.
- Hu, X. and Yacoub, E. (2012). The story of the initial dip in fmri. *Neuroimage*, 62(2):1103–1108.
- Hubel, D. and Wiesel, T. (1963). Shape and arrangement of columns in cat’s striate cortex. *The Journal of physiology*, 165(3):559–568.
- Hubel, D. and Wiesel, T. (1969). Anatomical demonstration of columns in the monkey striate cortex. *Nature*, 221(5182):747–750.
- Hubel, D. H. and Livingstone, M. S. (1987). Segregation of form, color, and stereopsis in primate area 18. *The Journal of Neuroscience*, 7(11):3378–3415.

- Hubel, D. H. and 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.
- Hubel, D. H. and Wiesel, T. N. (1972). Laminar and columnar distribution of geniculo-cortical fibers in the macaque monkey. *Journal of Comparative Neurology*, 146(4):421–450.
- Hubel, D. H. and Wiesel, T. N. (1974). Sequence regularity and geometry of orientation columns in the monkey striate cortex. *Journal of Comparative Neurology*, 158(3):267–293.
- Huk, A. C., Dougherty, R. F., and Heeger, D. J. (2002). Retinotopy and functional subdivision of human areas mt and mst. *The Journal of Neuroscience*, 22(16):7195–7205.
- Ip, I. B., Bridge, H., and Parker, A. J. (2014). Effects of spatial and feature attention on disparity-rendered structure-from-motion stimuli in the human visual cortex. *PloS one*, 9(6):1–10.
- Jaffer, S., Vorobyov, V., and Sengpiel, F. (2012). Effects of different forms of monocular deprivation on primary visual cortex maps. *Visual neuroscience*, 29(4-5):247–253.
- Jenkinson, M., Bannister, P., Brady, M., and Smith, S. (2002). Improved optimization for the robust and accurate linear registration and motion correction of brain images. *Neuroimage*, 17(2):825–841.
- Jorge, J., Figueiredo, P., van der Zwaag, W., and Marques, J. P. (2013). Signal fluctuations in fmri data acquired with 2d-epi and 3d-epi at 7 tesla. *Magnetic resonance imaging*, 31(2):212–220.
- Kaas, J. H. (1997). Theories of visual cortex organization in primates. In *Extrastriate cortex in primates*, pages 91–125. Springer.
- Kamitani, Y. and Tong, F. (2005). Decoding the visual and subjective contents of the human brain. *Nature neuroscience*, 8(5):679–685.
- Kanwisher, N., McDermott, J., and Chun, M. M. (1997). The fusiform face area: a module in human extrastriate cortex specialized for face perception. *The Journal of neuroscience*, 17(11):4302–4311.
- Katz, L. C., Gilbert, C. D., and Wiesel, T. N. (1989). Local circuits and ocular dominance columns in monkey striate cortex. *The Journal of Neuroscience*, 9(4):1389–1399.
- Kemper, V. G., De Martino, F., Vu, A. T., Poser, B. A., Feinberg, D. A., Goebel, R., and Yacoub, E. (2015). Sub-millimeter t2 weighted fmri at 7 t: comparison of 3d-grase and 2d se-epi. *Frontiers in neuroscience*, 9(163):1–14.

- Kolster, H., Peeters, R., and Orban, G. A. (2010). The retinotopic organization of the human middle temporal area mt/v5 and its cortical neighbors. *The Journal of Neuroscience*, 30(29):9801–9820.
- Kriegeskorte, N. and Bandettini, P. (2007). Analyzing for information, not activation, to exploit high-resolution fmri. *Neuroimage*, 38(4):649–662.
- Kriegeskorte, N., Goebel, R., and Bandettini, P. (2006). Information-based functional brain mapping. *Proceedings of the National Academy of Sciences of the United States of America*, 103(10):3863–3868.
- Larsson, J. and Heeger, D. J. (2006). Two retinotopic visual areas in human lateral occipital cortex. *The Journal of Neuroscience*, 26(51):13128–13142.
- Levitt, J. B., Kiper, D. C., and Movshon, J. A. (1994). Receptive fields and functional architecture of macaque v2. *Journal of Neurophysiology*, 71(6):2517–2542.
- Li, M., Liu, F., Juusola, M., and Tang, S. (2014). Perceptual color map in macaque visual area v4. *The Journal of Neuroscience*, 34(1):202–217.
- Lim, H., Wang, Y., Xiao, Y., Hu, M., and Felleman, D. J. (2009). Organization of hue selectivity in macaque v2 thin stripes. *Journal of neurophysiology*, 102(5):2603–2615.
- Livingstone, M. and Hubel, D. (1988). Segregation of form, color, movement, and depth: anatomy, physiology, and perception. *Science*, 240:740–749.
- Livingstone, M. S. (1996). Ocular dominance columns in new world monkeys. *The Journal of neuroscience*, 16(6):2086–2096.
- Livingstone, M. S. and Hubel, D. H. (1987). Psychophysical evidence for separate channels for the perception of form, color, movement, and depth. *The Journal of Neuroscience*, 7(11):3416–3468.
- Logothetis, N. K. (2002). The neural basis of the blood–oxygen–level–dependent functional magnetic resonance imaging signal. *Philosophical Transactions of the Royal Society of London B: Biological Sciences*, 357(1424):1003–1037.
- Logothetis, N. K. (2003). The underpinnings of the bold functional magnetic resonance imaging signal. *The Journal of Neuroscience*, 23(10):3963–3971.
- Logothetis, N. K. (2008). What we can do and what we cannot do with fmri. *Nature*, 453(7197):869–878.
- Logothetis, N. K., Pauls, J., Augath, M., Trinath, T., and Oeltermann, A. (2001). Neurophysiological investigation of the basis of the fmri signal. *Nature*, 412(6843):150–157.

- Logothetis, N. K. and Pfeuffer, J. (2004). On the nature of the bold fmri contrast mechanism. *Magnetic resonance imaging*, 22(10):1517–1531.
- Lutti, A., Thomas, D. L., Hutton, C., and Weiskopf, N. (2013). High-resolution functional mri at 3 t: 3d/2d echo-planar imaging with optimized physiological noise correction. *Magnetic resonance in medicine*, 69(6):1657–1664.
- Machielsen, W., Rombouts, S. A., Barkhof, F., Scheltens, P., and Witter, M. P. (2000). Fmri of visual encoding: reproducibility of activation. *Human brain mapping*, 9(3):156–164.
- Magnotta, V. A., Friedman, L., and BIRN-FIRST (2006). Measurement of signal-to-noise and contrast-to-noise in the fbirn multicenter imaging study. *Journal of Digital Imaging*, 19(2):140–147.
- Malach, R., Tootell, R. B., and Malonek, D. (1994). Relationship between orientation domains, cytochrome oxidase stripes, and intrinsic horizontal connections in squirrel monkey area v2. *Cerebral Cortex*, 4(2):151–165.
- Maunsell, J. and van Essen, D. C. (1983). The connections of the middle temporal visual area (mt) and their relationship to a cortical hierarchy in the macaque monkey. *The Journal of neuroscience*, 3(12):2563–2586.
- Maunsell, J. H., Ghose, G. M., Assad, J. A., McADAMS, C. J., Boudreau, C. E., and Noerager, B. D. (1999). Visual response latencies of magnocellular and parvocellular lgn neurons in macaque monkeys. *Visual neuroscience*, 16(01):1–14.
- McKeefry, D., Watson, J., Frackowiak, R., Fong, K., and Zeki, S. (1997). The activity in human areas v1/v2, v3, and v5 during the perception of coherent and incoherent motion. *Neuroimage*, 5(1):1–12.
- McKeefry, D. and Zeki, S. (1997). The position and topography of the human colour centre as revealed by functional magnetic resonance imaging. *Brain*, 120(12):2229–2242.
- Miller, K. L., Smith, S. M., Jezzard, P., and Pauly, J. M. (2006). High-resolution fmri at 1.5 t using balanced ssfp. *Magnetic resonance in medicine*, 55(1):161–170.
- Mishkin, M. and Ungerleider, L. G. (1982). Contribution of striate inputs to the visuospatial functions of parieto-preoccipital cortex in monkeys. *Behavioural brain research*, 6(1):57–77.
- Modha, D. S. and Singh, R. (2010). Network architecture of the long-distance pathways in the macaque brain. *Proceedings of the National Academy of Sciences*, 107(30):13485–13490.
- Muckli, L., De Martino, F., Vizioli, L., Petro, L. S., Smith, F. W., Ugurbil, K., Goebel, R., and Yacoub, E. (2015). Contextual feedback to superficial layers of v1. *Current Biology*, 25(20):2690–2695.

- Nasr, S., Polimeni, J. R., and Tootell, R. B. (2016). Interdigitated color-and disparity-selective columns within human visual cortical areas v2 and v3. *The Journal of Neuroscience*, 36(6):1841–1857.
- Neri, P., Bridge, H., and Heeger, D. J. (2004). Stereoscopic processing of absolute and relative disparity in human visual cortex. *Journal of Neurophysiology*, 92(3):1880–1891.
- Nienborg, H., Bridge, H., Parker, A. J., and Cumming, B. G. (2004). Receptive field size in v1 neurons limits acuity for perceiving disparity modulation. *The Journal of neuroscience*, 24(9):2065–2076.
- Norman, K. A., Polyn, S. M., Detre, G. J., and Haxby, J. V. (2006). Beyond mind-reading: multi-voxel pattern analysis of fmri data. *Trends in cognitive sciences*, 10(9):424–430.
- Ogawa, S., Lee, T.-M., Kay, A. R., and Tank, D. W. (1990). Brain magnetic resonance imaging with contrast dependent on blood oxygenation. *Proceedings of the National Academy of Sciences*, 87(24):9868–9872.
- Olman, C. A. and Yacoub, E. (2011). High-field fmri for human applications: an overview of spatial resolution and signal specificity. *The open neuroimaging journal*, 5(1):74–89.
- Orban, G. A., Fize, D., Peuskens, H., Denys, K., Nelissen, K., Sunaert, S., Todd, J., and Vanduffel, W. (2003). Similarities and differences in motion processing between the human and macaque brain: evidence from fmri. *Neuropsychologia*, 41(13):1757–1768.
- Oshio, K. and Feinberg, D. A. (1991). Grase (gradient-and spin-echo) imaging: A novel fast mri technique. *Magnetic resonance in medicine*, 20(2):344–349.
- Ostwald, D., Lam, J. M., Li, S., and Kourtzi, Z. (2008). Neural coding of global form in the human visual cortex. *Journal of Neurophysiology*, 99(5):2456–2469.
- Peirce, J. W. (2007). Psychopyspsychophysics software in python. *Journal of neuroscience methods*, 162(1):8–13.
- Peirce, J. W. (2009). Generating stimuli for neuroscience using psychopy. *Frontiers in neuroinformatics*, 2(10):1–8.
- Peters, A. M., Brookes, M. J., Hoogenraad, F. G., Gowland, P. A., Francis, S. T., Morris, P. G., and Bowtell, R. (2007). T2* measurements in human brain at 1.5, 3 and 7 t. *Magnetic resonance imaging*, 25(6):748–753.
- Pfeuffer, J., van de Moortele, P.-F., Yacoub, E., Shmuel, A., Adriany, G., Andersen, P., Merkle, H., Garwood, M., Ugurbil, K., and Hu, X. (2002). Zoomed functional imaging in the human brain at 7 tesla with simultaneous high spatial and high temporal resolution. *Neuroimage*, 17(1):272–286.

- Polimeni, J. R., Fischl, B., Greve, D. N., and Wald, L. L. (2010). Laminar analysis of 7t bold using an imposed spatial activation pattern in human v1. *Neuroimage*, 52(4):1334–1346.
- Poser, B. A., Koopmans, P. J., Witzel, T., Wald, L. L., and Barth, M. (2010). Three dimensional echo-planar imaging at 7 tesla. *Neuroimage*, 51(1):261–266.
- Pratte, M. S., Sy, J. L., Swisher, J. D., and Tong, F. (2016). Radial bias is not necessary for orientation decoding. *NeuroImage*, 127:23–33.
- Qing, G., Zhang, S., Wang, B., and Wang, N. (2010). Functional mri signal changes in primary visual cortex corresponding to the central normal visual field of patients with primary open-angle glaucoma. *Investigative ophthalmology & visual science*, 51(9):4627–4634.
- Reichl, L., Heide, D., Löwel, S., Crowley, J. C., Kaschube, M., and Wolf, F. (2012). Coordinated optimization of visual cortical maps (i) symmetry-based analysis. *PLoS Comput Biol*, 8(11):e1002466.
- Rice, M. L., Leske, D. A., Smestad, C. E., and Holmes, J. M. (2008). Results of ocular dominance testing depend on assessment method. *Journal of American Association for Pediatric Ophthalmology and Strabismus*, 12(4):365–369.
- Riesenhuber, M. and Poggio, T. (1999). Hierarchical models of object recognition in cortex. *Nature neuroscience*, 2(11):1019–1025.
- Riesenhuber, M. and Poggio, T. (2000). Models of object recognition. *Nature neuroscience*, 3:1199–1204.
- Riesenhuber, M. and Poggio, T. (2002). Neural mechanisms of object recognition. *Current opinion in neurobiology*, 12(2):162–168.
- Roe, A. W. (2003). Modular complexity of area v2 in the macaque monkey. *The primate visual system*, pages 109–138.
- Roe, A. W. and Ts'o, D. (1995). Visual topography in primate v2: multiple representation across functional stripes. *The Journal of Neuroscience*, 15(5):3689–3715.
- Roe, W. A. and Ts'o (1997). The functional architecture of area v2 in the macaque monkey. In *Extrastriate cortex in primates*, pages 295–333. Springer.
- Saygin, A. P. and Sereno, M. I. (2008). Retinotopy and attention in human occipital, temporal, parietal, and frontal cortex. *Cerebral Cortex*, 18(9):2158–2168.
- Schurger, A. (2009). A very inexpensive mri-compatible method for dichoptic visual stimulation. *Journal of neuroscience methods*, 177(1):199–202.

- Schwarzkopf, D. S., Schindler, A., and Rees, G. (2010). Knowing with which eye we see: Utrocular discrimination and eye-specific signals in human visual cortex. *PloS one*, 5(10):e13775.
- Schwarzkopf, D. S., Song, C., and Rees, G. (2011). The surface area of human v1 predicts the subjective experience of object size. *Nature neuroscience*, 14(1):28–30.
- Sereno, M. I., Dale, A., Reppas, J., Kwong, K., et al. (1995). Borders of multiple visual areas in humans revealed by functional magnetic resonance imaging. *Science*, 268(5212):889.
- Sereno, M. I., McDonald, C. T., and Allman, J. M. (1994). Analysis of retinotopic maps in extrastriate cortex. *Cerebral Cortex*, 4(6):601–620.
- Seymour, K., Clifford, C. W., Logothetis, N. K., and Bartels, A. (2009). The coding of color, motion, and their conjunction in the human visual cortex. *Current Biology*, 19(3):177–183.
- Seymour, K., Clifford, C. W., Logothetis, N. K., and Bartels, A. (2010). Coding and binding of color and form in visual cortex. *Cerebral cortex*, 20(8):1946–1954.
- Seymour, K. J. and Clifford, C. W. (2012). Decoding conjunctions of direction-of-motion and binocular disparity from human visual cortex. *Journal of neurophysiology*, 107(9):2335–2341.
- Shapley, R. and Hawken, M. J. (2011). Color in the cortex: single-and double-opponent cells. *Vision research*, 51(7):701–717.
- Shipp, S., Adams, D. L., Moutoussis, K., and Zeki, S. (2009). Feature binding in the feedback layers of area v2. *Cerebral cortex*, 19(10):2230–2239.
- Shipp, S. and Zeki, S. (1989). The organization of connections between areas v5 and v2 in macaque monkey visual cortex. *European Journal of Neuroscience*, 1(4):333–354.
- Shipp, S. and Zeki, S. (2002). The functional organization of area v2, i: specialization across stripes and layers. *Visual neuroscience*, 19(02):187–210.
- Sincich, L. C., Adams, D. L., and Horton, J. C. (2003). Complete flatmounting of the macaque cerebral cortex. *Visual neuroscience*, 20(06):663–686.
- Sincich, L. C. and Horton, J. C. (2002). Divided by cytochrome oxidase: a map of the projections from v1 to v2 in macaques. *Science*, 295(5560):1734–1737.
- Sincich, L. C. and Horton, J. C. (2005). The circuitry of v1 and v2: integration of color, form, and motion. *Annu. Rev. Neurosci.*, 28:303–326.
- Sincich, L. C., Jocson, C. M., and Horton, J. C. (2010). V1 interpatch projections to v2 thick stripes and pale stripes. *The Journal of Neuroscience*, 30(20):6963–6974.

- Sled, J. G. and Pike, G. B. (2000). Correction for b1 and b0 variations in quantitative t2 measurements using mri. *Magnetic resonance in medicine*, 43(4):589–593.
- Smith, M. A., Kohn, A., and Movshon, J. A. (2007). Glass pattern responses in macaque v2 neurons. *Journal of Vision*, 7(3):5–5.
- Smith, S. M., Jenkinson, M., Woolrich, M. W., Beckmann, C. F., Behrens, T. E., Johansen-Berg, H., Bannister, P. R., De Luca, M., Drobnjak, I., Flitney, D. E., et al. (2004). Advances in functional and structural mr image analysis and implementation as fsl. *Neuroimage*, 23:S208–S219.
- Song, C., Schwarzkopf, D. S., and Rees, G. Variability in visual cortex size reflects trade off between local orientation sensitivity and global orientation modulation. *Nature communications*, 4.
- Stokes, M., Thompson, R., Cusack, R., and Duncan, J. (2009). Top-down activation of shape-specific population codes in visual cortex during mental imagery. *The Journal of Neuroscience*, 29(5):1565–1572.
- Swindale, N. V., Shoham, D., Grinvald, A., Bonhoeffer, T., and Hübener, M. (2000). Visual cortex maps are optimized for uniform coverage. *Nature neuroscience*, 3(8):822–826.
- Thomas, O. M., Cumming, B. G., and Parker, A. J. (2002). A specialization for relative disparity in v2. *Nature neuroscience*, 5(5):472–478.
- Tootell, R. and Hamilton, S. L. (1989). Functional anatomy of the second visual area (v2) in the macaque. *The Journal of neuroscience*, 9(8):2620–2644.
- Tootell, R., Silverman, M. S., De Valois, R. L., and Jacobs, G. H. (1983). Functional organization of the second cortical visual area in primates. *Science*, 220(4598):737–739.
- Tootell, R. B., Hadjikhani, N., Hall, E. K., Marrett, S., Vanduffel, W., Vaughan, J. T., and Dale, A. M. (1998a). The retinotopy of visual spatial attention. *Neuron*, 21(6):1409–1422.
- Tootell, R. B., Mendola, J. D., Hadjikhani, N. K., Liu, A. K., and Dale, A. M. (1998b). The representation of the ipsilateral visual field in human cerebral cortex. *Proceedings of the National Academy of Sciences*, 95(3):818–824.
- Tootell, R. B., Nelissen, K., Vanduffel, W., and Orban, G. A. (2004). Search for color center (s) in macaque visual cortex. *Cerebral Cortex*, 14(4):353–363.
- Tootell, R. B., Reppas, J. B., Kwong, K. K., Malach, R., Born, R. T., Brady, T. J., Rosen, B. R., and Belliveau, J. W. (1995). Functional analysis of human mt and related visual cortical areas using magnetic resonance imaging. *The Journal of Neuroscience*, 15(4):3215–3230.

- Toro, R., Fox, P. T., and Paus, T. (2008). Functional coactivation map of the human brain. *Cerebral cortex*, 18(11):2553–2559.
- Triantafyllou, C., Hoge, R., Krueger, G., Wiggins, C., Potthast, A., Wiggins, G., and Wald, L. (2005). Comparison of physiological noise at 1.5 t, 3 t and 7 t and optimization of fmri acquisition parameters. *Neuroimage*, 26(1):243–250.
- Ts'o, D., Roe, A. W., and Gilbert, C. D. (2001). A hierarchy of the functional organization for color, form and disparity in primate visual area v2. *Vision research*, 41(10):1333–1349.
- Uğurbil, K. (2012). The road to functional imaging and ultrahigh fields. *Neuroimage*, 62(2):726–735.
- Uğurbil, K., Toth, L., and Kim, D.-S. (2003). How accurate is magnetic resonance imaging of brain function? *Trends in neurosciences*, 26(2):108–114.
- Ungerleider, L. G. and Haxby, J. V. (1994). what and wherein the human brain. *Current opinion in neurobiology*, 4(2):157–165.
- van der Zwaag, W., Francis, S., Head, K., Peters, A., Gowland, P., Morris, P., and Bowtell, R. (2009). fmri at 1.5, 3 and 7 t: characterising bold signal changes. *Neuroimage*, 47(4):1425–1434.
- Van Essen, D., Maunsell, J., and Bixby, J. (1981). The middle temporal visual area in the macaque: myeloarchitecture, connections, functional properties and topographic organization. *Journal of Comparative Neurology*, 199(3):293–326.
- Van Hooser, S. D. (2007). Similarity and diversity in visual cortex: is there a unifying theory of cortical computation? *The Neuroscientist*, 13(6):639–656.
- Wandell, B. A., Brewer, A. A., and Dougherty, R. F. (2005). Visual field map clusters in human cortex. *Philosophical Transactions of the Royal Society of London B: Biological Sciences*, 360(1456):693–707.
- Wandell, B. A., Dumoulin, S. O., and Brewer, A. A. (2007). Visual field maps in human cortex. *Neuron*, 56(2):366–383.
- Wandell, B. A., El Gamal, A., and Girod, B. (2002). Common principles of image acquisition systems and biological vision. *Proceedings of the IEEE*, 90(1):5–17.
- Wandell, B. A. and Winawer, J. (2011). Imaging retinotopic maps in the human brain. *Vision research*, 51(7):718–737.
- Weisskoff, R., Zuo, C. S., Boxerman, J. L., and Rosen, B. R. (1994). Microscopic susceptibility variation and transverse relaxation: theory and experiment. *Magnetic Resonance in Medicine*, 31(6):601–610.

- Wilson, S. P. and Bednar, J. A. (2015). What, if anything, are topological maps for? *Developmental neurobiology*, 75(6):667–681.
- Woolrich, M. W., Jbabdi, S., Patenaude, B., Chappell, M., Makni, S., Behrens, T., Beckmann, C., Jenkinson, M., and Smith, S. M. (2009). Bayesian analysis of neuroimaging data in fsl. *Neuroimage*, 45(1):S173–S186.
- Yacoub, E., Shmuel, A., Logothetis, N., and Uğurbil, K. (2007). Robust detection of ocular dominance columns in humans using hahn spin echo bold functional mri at 7 tesla. *Neuroimage*, 37(4):1161–1177.
- Yacoub, E., Shmuel, A., Pfeuffer, J., De Moortele, V., Adriany, G., Andersen, P., Vaughan, J. T., Merkle, H., Ugurbil, K., Hu, X., et al. (2001). Imaging brain function in humans at 7 tesla. *Magnetic Resonance in Medicine*, 45(4):588–594.
- Yacoub, E., Van De Moortele, P.-F., Shmuel, A., and Uğurbil, K. (2005). Signal and noise characteristics of hahn se and ge bold fmri at 7 t in humans. *Neuroimage*, 24(3):738–750.
- Zhang, D., Wen, X., Liang, B., Liu, B., Liu, M., and Huang, R. (2014). Neural associations of the early retinotopic cortex with the lateral occipital complex during visual perception. *PloS one*, 9(9):e108557.