

# **Studying the selectivity of neuronal sub-populations within fMRI voxels**

Panagiotis Sapountzis, MSc

*Thesis submitted to the University of Nottingham for the degree of  
Doctor of Philosophy, November 2009*

## ***Abstract***

---

Functional magnetic resonance imaging (fMRI) has become a ubiquitous tool in cognitive neuroscience. The technique allows the non-invasive measurements of cortical responses, but only on the millimeter scale. Recently, two methods aimed at studying the selectivity of neuronal populations on a sub-voxel scale. The first technique, fMRI adaptation, relies on the observation that the fMRI response in a given voxel is reduced after prolonged presentation of a stimulus, and that this reduction is selective to the characteristics of the repeated stimuli. The second technique, multi-variate pattern analysis (MVPA), makes use of multi-variate statistics to recover small biases in individual voxels. This thesis compared the two techniques with the aim of studying early- and mid-level processing in the visual cortex.

Chapter 3 investigated whether adaptation and MVPA provide consistent results about the properties of the cortical areas under study. To address this question, this thesis compared the two methods for their ability to detect the well-documented orientation selectivity in early visual cortex. Using optimised experimental designs for each, this thesis found that the MVPA approach was more sensitive to small differences in stimulus orientation than the adaptation paradigm. Estimates of orientation selectivity obtained with the two methods were, however, very highly correlated across visual areas.

Chapters 4 and 5 used both techniques to investigate how local orientation signals are combined and detected in intermediate levels of visual processing. In both chapters the MVPA was more efficient in detecting differences between stimulus categories. In particular, chapter 4 used plaid stimuli, made of the linear sum of two sinusoidal gratings. We obtained weak but consistent evidence,

pointing to the direction that V2 might play a role in Fourier component integration. Chapter 5 used contour stimuli constructed from two luminance modulated sinusoidal gratings, with different orientations. Whereas, adaptation failed to reveal contour selectivity, MVPA accuracy was high in most areas tested. However, the experiment did not reveal a significant difference between the test and control conditions.

Chapter 6 investigated the encoding of spatial phase in the cortex. Phase is a fundamental aspect of spatial vision, crucial both for the extraction of local features and overall scene perception. This thesis showed that several visual areas, including the primary visual cortex, were sensitive to relative phase combinations. However, phase coherence was optimally encoded in extrastriate areas as predicted by the physiological properties of higher regions.

Parts of the work discussed in Chapter 3 were published in the following article:

**Sapountzis P, Schluppeck D, Bowtell R, and Peirce JW.** A comparison of fMRI adaptation and multivariate pattern classification analysis in visual cortex. *Neuroimage* 2009.

## ***Acknowledgments***

---

First and foremost I must thank my supervisor Jon Peirce for trusting me and giving me the chance to work on such an interesting project. He has been a fantastic supervisor who has always been available whenever I needed his help and support.

I am grateful to Denis Schluppeck for teaching me how to use the fMRI analysis tools. Without his help this work simply could not have been completed on time. He was always keen to help with any difficulties I faced with my analyses (and I did face a lot!).

Many thanks to Richard Bowtell for supporting my study. His useful comments helped me improve my work.

Thanks to Sarah Hancock for kindly accepting to participate in my experiments. Also to my office mates Ross Deas, Laura Stevens, David McGovern, Andrew Astle, Lily Xing and Francesca Rocchi for making these three years in the office far more interesting. Many thanks to the Notts vision group for support and useful discussions.

Special thanks to my family and especially my parents who have always encouraged me to follow my dreams. Without their support, especially in times of difficulty, I would not have made it this far.

Finally, I must thank Maria Makridaki for being such a patient and loving partner. These three years have been an exciting journey for both of us, with times of happiness and times of great difficulty. Without Maria this journey would simply not have been the same! This work is dedicated to her.

<b>Abstract</b> .....	<b>ii</b>
<b>Acknowledgments</b> .....	<b>iv</b>
<b>1 Introduction</b> .....	<b>1</b>
1.1 The visual system .....	3
1.1.1 Projections from retina to the brain .....	3
1.1.2 Retinotopic maps in the visual cortex .....	6
1.1.3 Early visual areas .....	7
1.1.4 The primary visual cortex (V1) .....	7
1.1.5 Visual cortical pathways .....	11
1.1.6 Extrastriate cortical areas .....	14
1.2 Principles of functional Magnetic Resonance Imaging .....	21
1.2.1 Magnetic Resonance Imaging (MRI) .....	21
1.2.2 Functional Magnetic Resonance Imaging (fMRI) .....	29
1.2.3 Limitations of the BOLD technique .....	34
1.3 Experimental techniques .....	37
1.3.1 Adaptation .....	37
1.3.2 Multi-variate Pattern Analyses (MVPA) .....	43
<b>2 General Methods</b> .....	<b>48</b>
2.1 Cortical segmentation and flattening .....	48
2.1.1 Analysis Stream .....	49
2.1.2 Comparison with other packages .....	54
2.2 Retinotopic mapping .....	57
2.2.1 Phase-encoding techniques .....	57
2.2.2 Identification of early visual areas .....	61
2.2.3 Non-retinotopic visual areas .....	65
2.3 fMRI scanning .....	66
2.3.1 Participants .....	66
2.3.2 Visual stimuli .....	66
2.3.3 Functional imaging .....	67
2.3.4 Attention control task .....	68
2.3.5 Data analysis .....	68
<b>3 Do fMRI adaptation and pattern-classification analysis measure the same thing?</b> .....	<b>74</b>
3.1.1 Methods .....	77
3.1.2 Results .....	81
3.1.3 Discussion .....	88
3.2 Does the interleaving of scans affect the performance of the methods? .....	92
3.2.1 Results .....	92
3.2.2 Discussion .....	94
3.3 Do orientation biased voxels adapt more? .....	95
3.3.1 Results .....	95
3.3.2 Discussion .....	95
3.4 High-field (7T) measurements .....	97
3.4.1 Methods .....	97
3.4.2 Results .....	98
3.4.3 Discussion .....	99
<b>4 How are mechanisms selective for complex visual patterns represented in the cortex?</b> .....	<b>102</b>
4.1 Experiment 1 .....	103
4.1.1 Methods .....	104
4.1.2 Results .....	107
4.1.3 Discussion .....	110
4.2 Experiment 2 .....	111
4.2.1 Methods .....	112
4.2.2 Results .....	113

4.2.3	Discussion .....	115
4.3	Experiment 3.....	115
4.3.1	Methods.....	116
4.3.2	Results .....	119
4.3.3	Discussion .....	122
4.4	General discussion .....	123
<b>5</b>	<b>Integration of simple contours in the visual cortex.....</b>	<b>126</b>
5.1	Experiment 1.....	127
5.1.1	Methods.....	128
5.1.2	Results .....	131
5.1.3	Discussion .....	134
5.2	Experiment 2.....	135
5.2.1	Methods.....	135
5.2.2	Results .....	137
5.2.3	Discussion .....	139
5.3	General discussion .....	139
<b>6</b>	<b>Spatial phase encoding in the cortex.....</b>	<b>141</b>
6.1	Methods .....	145
6.2	Results.....	147
6.3	Discussion .....	151
<b>7</b>	<b>General discussion.....</b>	<b>154</b>
	<b>References.....</b>	<b>162</b>

---

## **1 Introduction**

---

*'By nature, all men long to know. An indication is their delight in the senses. For these, quite apart from their utility, are intrinsically delightful, and that through the eyes more than the others. For it is not only with a view to action but also when we have no intention to do anything that we choose, so to speak, sight rather than all the others. And the reason for this is that vision is the sense that especially produces cognition in us and reveals many distinguishing features of things'* (Aristotle, The Metaphysics).

The above quote from Aristotle's Metaphysics brilliantly captures the now widely accepted link between vision and cognition. Our visual system does not just faithfully record images but it also provides the necessary information for us to behave appropriately, by giving meaning to the perceived images. The above quote is only an example of the way ancient philosophers, based on observation and rational thinking, revealed for the first time the mysteries of nature and human behaviour that up to that time were only explained through prejudice and mythology. However, many have changed since then. Advances in natural sciences have allowed the development of techniques that make possible the direct study of scientific and philosophical hypotheses. In the last few years new tools have been developed that even allow the functional imaging of the human brain.

Functional magnetic resonance imaging (fMRI) is currently the mainstay of brain imaging. Since first introduced in the early 1990s, it has a major impact on basic cognitive neuroscience research. Its use is so prominent that about eight peer-reviewed articles per day include fMRI as a keyword (Logothetis 2008). fMRI measurements typically rely on the blood oxygenation level dependent (BOLD) signal, which is used to track local increases in neural activity. One of its

limitations, in comparison with the direct recording of neuronal responses via microelectrode, is its spatial resolution. Improvements to imaging hardware and analysis techniques have provided access to higher-resolution images at improved signal-to-noise ratios (Moon et al. 2007; Yacoub et al. 2007).

However, BOLD measurements are ultimately limited in spatial resolution by, amongst other things, the spatial scale of the local vascular system. Recent fMRI studies have demonstrated new methods for revealing separate subpopulations of neurons selectively tuned for different stimuli even when these neurons are intermingled at a spatial scale that is smaller than the sampling resolution of the measurement. The first technique, fMRI adaptation, relies on the observation that the BOLD response in a given voxel is reduced after prolonged presentation of a stimulus, and that this reduction is selective to the characteristics of the repeated stimuli (adapters). The second technique, multi-variate pattern analysis (MVPA), makes use of multi-variate statistics to recover small biases in individual voxels in their responses to different stimuli. It is thought that these biases arise due to the uneven distribution of neurons (with different properties) sampled by the many voxels in the imaged volume.

Although in recent years these two techniques have been heavily used, they have never been compared explicitly, and little is known about their relative sensitivities. Thus, one major theme throughout this thesis is to compare the two techniques and evaluate their relative sensitivities. The first experimental chapter directly compares the two methods for their ability to detect the well-documented orientation selectivity in early visual cortex. In the following chapters the techniques are used to study more complex visual features. The aim of these chapters is to understand how local Fourier components are combined and detected in intermediate levels of visual processing. The following literature review

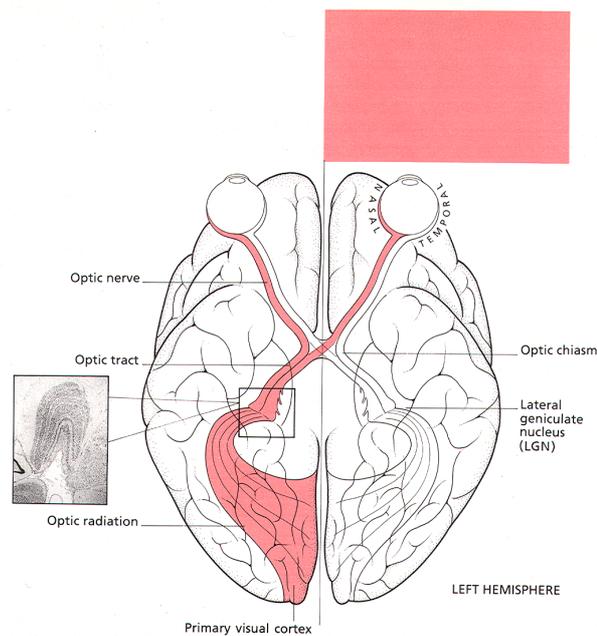
summarises known aspects of early- and mid-level vision to the degree they are of interest to this thesis.

## **1.1 The visual system**

In the first steps of visual processing an image is formed at the back of the eye, in the retina, where the photoreceptor mosaic samples it. The output of the photoreceptors is compared and combined by the retinal ganglion cells. The optic nerves, which are made up of the axons of the ganglion cells convey the visual information to the brain. The optic nerves cross at the optic chiasm; after they pass the optic chiasm, the axons of the retinal ganglion cells are known collectively as the optic tract. Most axons of the optic tract terminate on cells in the lateral geniculate nucleus (LGN). Axons of postsynaptic cells in the LGN form the optic radiations, which terminate in the primary visual cortex (V1). In addition to the primary visual cortex, other surrounding cortical regions are also associated with visual processing (e.g. V2, V3, V4) as described in more detail below.

### **1.1.1 Projections from retina to the brain**

The retinal image after the initial stages of processing is conveyed to each cerebral hemisphere by the axons of retinal ganglion cells, which cross over at the optic chiasm. As a result, those fibres that originate at the nasal part of the retina cross over to the opposite hemisphere, while those that originate at the temporal retina, continue to the same side of the brain (see Figure 1.1). This means that the fibres from the temporal retina of the left eye and from the nasal retina of the right eye pass to the left cerebral hemisphere, which therefore looks at the right visual field. By contrast, fibres from the nasal retina of the left eye and the temporal retina of the right eye pass over to the right hemisphere, which therefore looks at the left visual field. Thus, the right visual field projects to the left cerebral hemisphere, while the left visual field projects to the right cerebral hemisphere.



*Figure 1.1 Connections from retina to the cerebral hemispheres. Inset to the left shows the lateral geniculate nucleus (LGN) (from Zeki 1993).*

Beyond the optic chiasm the visual pathway is called the optic tract. Most of the optic tract fibres terminate at the lateral geniculate nucleus (LGN). In the monkey, 90% of the retinal ganglion cells send their axons to the LGN layers (Perry et al. 1984). In human, the LGN is a complex, six-layered structure. One noteworthy feature of this structure is the specificity of connections between the retina and the LGN. The inputs from the two eyes to the LGN are segregated, so that the optic nerve fibres coming from the eye on the same side of the LGN in question terminate in layers 5, 3 and 2, whereas the fibres coming from the eye on the opposite side of the LGN in question terminate in layers 6, 4 and 1. This is a very detailed, point-to-point projection from the retina, so that adjacent points on the retina project on adjacent points in each layer of the LGN. Moreover, the layers are organised in precise registration in terms of retinal representation. For example, if the cells at a point A in layer 6 receive input from a particular point in the left retina, the cells in point B in the layer below will receive their input from the corresponding point in the right retina. Thus, for each point in every layer

representing a particular point in one retina, there is a corresponding point in the other layers representing the corresponding point in the other retina.

There is also considerable regularity in the distribution of axons from the retinal ganglion cells within the LGN (Connolly and Van Essen 1984; Livingstone and Hubel 1988). The four upper layers, called parvocellular layers, contain neurons with small cell bodies. The two lower layers, called magnocellular layers, contain neurons with large cell bodies. Recently, a third type of cells with very small cell bodies called koniocellular (from the Greek 'κόvio' meaning dust) have been also discovered (Hendry and Reid 2000). Because koniocellular cells have been studied less than the first two types they are not further discussed in this thesis.

The different types of LGN neurons receive input from different types of retinal ganglion cells. Parvocellular layers receive their input from midget retinal ganglion cells and magnocellular layers from parasol cells. Dacey and Petersen (1992) showed that the signals from the two types of retinal ganglion cells form two segregated visual streams. The pathway that begins with the midget ganglion cells and projects to parvocellular layers of LGN is known as the parvocellular pathway, whereas the pathway that begins with the parasol cells and projects to the magnocellular layers of the LGN is known as the magnocellular pathway. These two pathways carry different types of information to the brain. The parvocellular pathway includes more than the 70% of the retinal ganglion cells and draws its input mostly from the central visual field. Damage in the parvocellular pathway impairs perception of finer spatial frequencies and colour (Merigan et al. 1991; Schiller et al. 1990). The magnocellular pathway carries information about coarser spatial frequencies and motion (Holodniy et al. 1993; Schiller et al. 1990). The cortical locations these two pathways project to and the

degree to which they remain segregated are discussed in more detail in section 1.1.5.

The axons of the LGN cells on which the optic nerve fibres terminate, travel in the optic radiation to terminate in the visual areas of the occipital cortex at the back of the brain. The primary visual cortex (V1) is often called the striate cortex because a striation is visible in anatomical sections through this region. In addition to V1, more than 30 other 'extrastriate' cortical areas are implicated in visual processing (Felleman and Van Essen 1991). The projection from each LGN to the ipsilateral cortex takes place in an orderly, point-to-point manner. The consequence of this is twofold. First, the cortex of each hemisphere receives signals from the opposite half of the visual field and second, adjacent retinal points are mapped at adjacent points in this cortex, just as they are in the LGN. This organised retinal projection on the visual cortex is known as retinotopic mapping.

### **1.1.2 Retinotopic maps in the visual cortex**

The discovery and analysis of cortical visual areas is a major accomplishment of neuroscience. The size and anatomical location of the visual areas may vary among individuals (Engel et al. 1997); therefore, the ability to define the exact anatomical location and size of these areas for different subjects is of great importance in experimental studies of the human brain. These visual areas have been extensively studied in animals using invasive techniques, such as electrophysiology. In animal models, more information is available to identify visual areas (Felleman and Van Essen 1991): (a) retinotopic maps, (b) differences in cellular architecture between areas, (c) connectivity between cortical regions and (d) regularities in the receptive field properties of neurons. However, in the living human brain fewer criteria are available to delineate visual areas. The main data available for partitioning visual areas are functional MRI (fMRI)

measurements of visual field maps. There are limited data on connectivity and cellular architecture, and there are almost no data on neuronal receptive fields (Wandell and Wade 2003). Reversely, the existence of these maps turns out to be very useful for fMRI in that they allow the various visual areas to be delineated conveniently and consistently for an individual subject. These maps are typically measured using stimuli that cause a travelling wave of neuronal activity in the visual cortex (DeYoe et al. 1996; Engel et al. 1997; Engel et al. 1994; Sereno et al. 1995). The actual experiment used to generate retinotopic maps is described in section 2.2.

### **1.1.3 Early visual areas**

The following sections discuss the representation in early visual areas in the occipital cortex and the interconnections between them. Early occipital areas have been extensively studied in animal models and many of their neuronal properties are well known. However, it has not been easy to discern what is achieved in each cortical visual region and how the outcomes of the analyses at each level are brought together.

### **1.1.4 The primary visual cortex (V1)**

The primary visual cortex (V1) is located within the calcarine sulcus in the occipital lobe at the back of the head and it is the first cortical region to process visual information. The distinctive properties of V1 cells have been studied heavily since the seminal works of Hubel and Wiesel (1962, 1968).

#### ***Cell types***

According to Hubel and Wiesel, cells in V1 can be arranged in a hierarchy of simple, complex and hypercomplex cells. Whereas simple cells show linear summation in their responses, complex cells do not. As Hubel writes, '*For the*

*most part, we can predict the responses of simple cells to complicated shapes from their responses to small-point stimuli* (Hubel 1988). Hypercomplex cells (also called end-stopped cells) are sensitive to stimulus length (Hubel and Wiesel 1965). If the length of the stimulus is too long their response will be significantly reduced, perhaps even to zero.

### **Orientation selectivity**

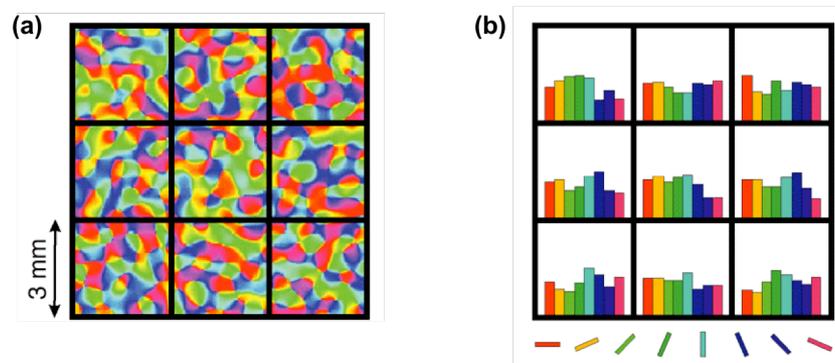
Another feature of cortical neurons Hubel and Wiesel realised, was that they are orientation selective, responding more to stimuli in some orientations than others (Hubel and Wiesel 1968; 1962). Figure 1.2 shows a simulation of how different orientation-selective V1 neurons might respond to a real scene. According to this simulation orientation-selective neurons might to be used to detect the ‘edges’ of a visual scene, each at a different orientation.



*Figure 1.2 Illustration of how orientation-selective simple cells might respond to a real image. Note that each cell type might to be used to detect edges at its preferred orientation. To generate this basic illustration the original image was convoluted with a Gabor kernel. The size of the kernel was approximately equal to the receptive field of a typical V1 simple cell.*

Orientation selective neurons are spatially organised in the cortex. The preferred orientation of neurons varies in an orderly way that depends on the cell's position in the cortex. Orientation selective neurons are arranged in orientation columns, each containing cells of a single preferred orientation. Successive columns have a preferred orientation slightly advanced from the column before (Hubel and Wiesel 1977; Obermayer and Blasdel 1993).

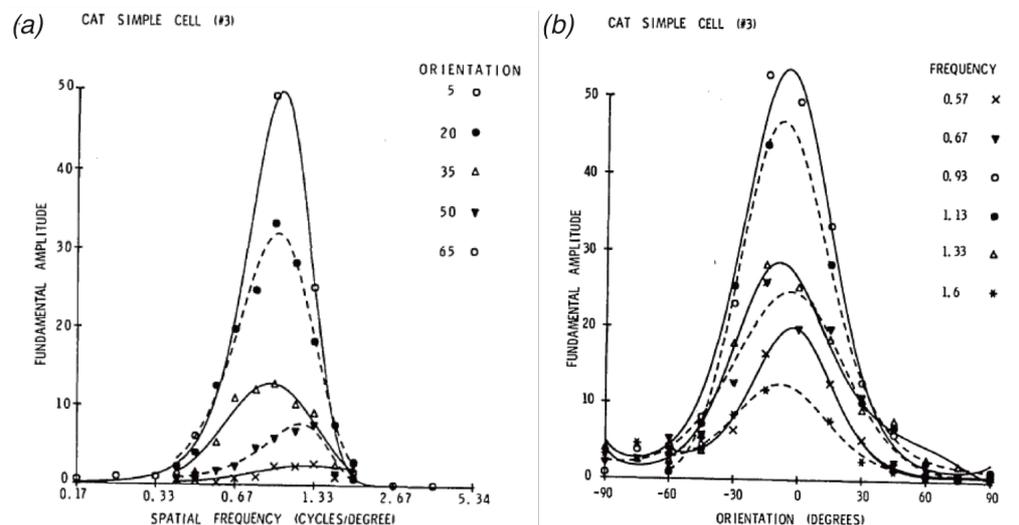
Figure 1.3 shows a synthetic map in which different colours code orientation-selective cells clustered into 'columns' of different orientation preference. The size of these columns in the macaque is around  $500\ \mu\text{m}$  in diameter (Bartfeld and Grinvald 1992; Obermayer and Blasdel 1993). Although, orientation-selective mechanisms have been extensively studied in animals, only recently the study has been extended to the human cortex using mainly fMRI. Chapter 3 compares the different techniques used to study orientation selectivity in the human cortex.



*Figure 1.3 (a) Synthetic orientation tuning map in which different colours code orientation selective columns of different orientation preference. The overlaid grid corresponds to a 2-dimensional projection of conventional  $3\times 3\times 3\text{mm}^3$  fMRI voxels. (b) Histograms showing the amount of orientation selectivity in each voxel (from Boynton 2005).*

### ***Other dimensions of cortical analysis***

Cortical cells are also selective to spatial frequency (Campbell et al. 1969; Webster and De Valois 1985). A cell's spatial frequency tuning curve, like orientation tuning, has a distinctive band-pass shape (Figure 1.4). Movshon et al. (1978) found a strong positive correlation between the narrowness of tuning of cells in the cat across the dimensions of spatial frequency and orientation. Cells that are narrowly tuned for spatial frequency tend to be narrowly tuned for orientation and vice versa. Experiments measuring the responses of cells across a wide range of spatial frequencies and orientations showed that striate cell's receptive fields can be modelled by a 2-dimensional Gabor function (Webster and De Valois 1985). Physiological data combined with psychophysical studies suggested that orientation and spatial frequency selective neurons may act as a set of spatial frequency filters, each set performing a crude Fourier analysis of images at different meridians (orientations) (Graham 1977).



*Figure 1.4 Responses of a simple cell, from cat's striate cortex to gratings of various (a) spatial frequencies and (b) orientations. This cell responds to only a narrow range of spatial frequencies and orientations (from Webster and De Valois 1985).*

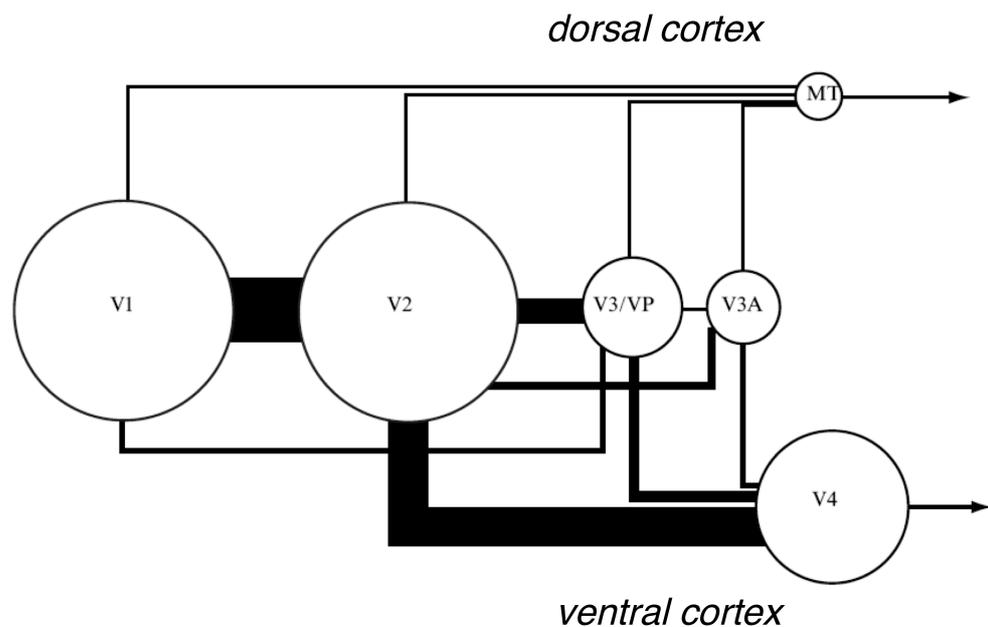
Another important step in the analysis striate cells perform is contrast normalisation (Albrecht and Geisler 1991; Heeger 1992). This step introduces a compressive nonlinearity in their responses to stimulus contrast. The same way retinal mechanisms of light adaptation discard information about ambient illumination (Derrington and Lennie 1982), the striate cortex, through a mechanism of contrast adaptation, adjusts its so that the information passed to other areas is relatively independent of the absolute level of contrast in the image. Striate cells also process perceptual features like binocular disparity, motion selectivity, and colour, however, discussion of these goes beyond the aims of this thesis.

In summary, striate cortex, appears to capture local structure and reduce redundancy from regions of extended common structure in images. Also, it performs a very important analysis step by normalising signals for contrast, so that the ratio of neuronal responses remains constant across a wide range of contrast levels (Lennie 1998). The ideas discussed above suggest that one of V1's roles might be to construct an economic and elaborate analysis of images.

### **1.1.5 Visual cortical pathways**

Striate cortex projects directly to at least seven other areas (Felleman and Van Essen 1991). A large projection goes to V2, which in turn projects to V3/VP, V4 and MT. Intercortical connections are illustrated in Figure 1.5. Macaque visual areas are shown in proportion to their size, and the connections between them are drawn proportionally to the relative number of fibres involved. One striking feature in this illustration is that the main body of connections leaving V1 is directed, through V2 and V4, ventrally to the temporal lobe. Not only the interconnections but also the functional properties of cortical regions suggest the segregation of visual information within the cortex into two pathways (Ungerleider and Mishkin 1982). This modularity in the cortex has been connected with the

anatomical and physiological diversity of retinal ganglion cells and neurons in LGN (see also section 1.1.1). Hubel and Livingstone (1987) suggested that the parvocellular pathway drives the temporal stream in cortex, while the magnocellular pathway provides the major input to the parietal stream in cortex. In particular, the parvocellular pathway is suggested to be the precursor of a system that in the cortex is directed ventrally into the temporal lobe and appears to be crucial for shape and object recognition. On the other hand, the magnocellular pathway is suggested to be the precursor of a system that in the cortex is directed dorsally into the parietal lobe and appears to be crucial for spatial perception, motion, depth and visuomotor performance.



*Figure 1.5 Illustration of macaque cortical areas shown in proportion to their size. The feed-forward connections between them are drawn proportionally to the relative number of fibres involved. The estimates are based on two assumptions: (1) the density of afferents entering an area is uniform, and (2) the areas to which V1 projects receive all the output from V1. For a complete discussion on these assumptions see Lennie (1998).*

Although the segregation of parvocellular and magnocellular projections in the cortex has been disputed by more recent studies<sup>1</sup> (e.g. Nealey and Maunsell 1994), evidence on the existence of two distinct cortical streams has been derived from the anatomical and physiological properties of neurons in the areas across pathways. Temporal and parietal lesions in monkeys significantly impair performance in object discrimination tasks, whereas, posterior parietal lesions do not affect visual discrimination, but instead cause deficits in tasks that require appreciating the spatial relationships among objects (see Desimone et al. 1985 for a review)

A study in the monkey by Baizer et al. (1991) provided anatomical evidence in support of the lesion studies. They injected two different types of tracers into the two streams, including areas V1, V2, V4 and MT among others. They found almost no neurons containing both tracers, suggesting that neuronal signals are directed towards either the parietal or temporal lobe. The functional properties and anatomical segregation of the two streams suggest that temporal and parietal areas receive and process, to a large degree, different information about the visual image (Goodale and Milner 1992; Ungerleider and Mishkin 1982).

Another feature evident in Figure 1.5 is that beyond V2 the size of visual areas becomes progressively smaller. For example, no area after V2 is as large as V2. This might suggest that at each level of processing much of the information required to make perceptual decisions is discarded. As Lennie (1998) writes, '*perceptual decisions are made at each level of the hierarchy; each level passes to the next the results of its analysis, but not the information used in undertaking the analysis*'. The above notion apart from a reasonable explanation about the

---

<sup>1</sup> While one branch of the magnocellular stream continues on an independent path to area MT, another branch converges with the parvocellular pathway in the upper layers of V1.

decreasing size of visual areas in the hierarchy provides a framework in which the concept of hierarchical organisation can be considered.

### **1.1.6 Extrastriate cortical areas**

Although it would be desirable to describe how later areas process visual information, not enough is understood either of the detailed physiology or the functional properties of these areas. This thesis considers how later centres combine the outputs of striate cortex in an attempt to provide insight into the computational role of extrastriate areas. Therefore, it would be useful to review some of the known physiological and functional properties of early extrastriate areas to the extent these are related to the purposes of this thesis.

#### **V2**

V2 is at least as large as the primary visual cortex and serves as an intermediate centre via which most cortical information is transmitted from V1 to the rest of the cortex. Most V2 cells are orientation selective (Zeki 1978b), though may be tuned less sharply than V1 cells (Levitt et al. 1994). However, in most respects V1 and V2 neurons are not remarkably different (Hubel and Livingstone 1987). A very important property is that cells with common preferences tend to be clustered in similar ways as in V1. Hence, in a region of around 500  $\mu\text{m}$  cells might all prefer similar orientations (Hubel and Livingstone 1987; Tootell and Hamilton 1989). Tootell and Hamilton (1989) suggested that there might also be some systematic mapping of spatial frequency. The binocular properties of V2 neurons are somehow different from those in V1. Almost all V2 neurons are driven through either eye or both eyes together and, in contrast to V1 cells, most are equally excitable through either eye (Burkhalter and Van Essen 1986). These properties attribute to V2 an important role towards binocular integration.

Does V2 have a role in combining the outputs of V1? One property that may point to this direction is that V2 neurons have larger receptive fields than V1. This means that a single cell in V2 carries less precise information about location of a stimulus in the visual field than does a V1 cell receiving information from the same part of the retina. Although this cruder analysis might suggest a feature integration role for V2 cells, there is not much evidence pointing to this direction. However, there are experiments to suggest that neurons in V2 have nonlinear integrative mechanisms that combine line segments separated by  $2^\circ$  or more. Moreover, these mechanisms respond to more complex stimuli like contours or edges (Peterhans and von der Heydt 1989; von der Heydt and Peterhans 1989). Despite these findings the functional role of V2 remains unclear to a large degree. V2 lesions in monkeys caused no measurable decrease in visual acuity, contrast sensitivity and local discriminations of luminance, colour and orientation (Merigan et al. 1993). However, V2 lesions significantly impaired performance on tasks that require grouping of distinctive elements. This might suggest that the analyses required for some low-level discrimination are undertaken by V1, leaving for V2 tasks involving more complex spatial discriminations.

### **V3**

The role of area V3 is unclear to a large degree as cells in this area usually lack distinctive visual properties. V3 projects mainly to V4 and MT, areas that also receive much input from V2, making it even harder to infer its functional role. Orientation-selective cells are very common in V3 (Felleman and Van Essen 1987; Gegenfurtner et al. 1997; Zeki 1978a), though more broadly tuned than in V2 (Baizer 1982) and respond to a wide range of spatial frequencies. Baizer (1982) and Zeki (1978a) found that receptive fields of V3 cells are larger than those of V2. These properties indicate that V3 performs a cruder analysis of the visual world than V2. For example, an orientation cell in V3 contains less

information both about stimulus orientation and stimulus width than an orientation cell in V2. In addition, electrophysiological reports indicate that approximately one-half of the neurons in macaque V3 are motion and direction selective (Felleman and Van Essen 1987; Gegenfurtner et al. 1997).

These observations, suggest that V3 might perform a coarser analysis of form especially at low contrasts, as V3 cells appear to have very low contrast sensitivities (Gegenfurtner et al. 1997). Moreover, Lennie (1998) considering that V3 provides only a small proportion of the input to V4 and MT, suggested this area might provide a spatially coarse modulation of information, connecting V1 and V2 to the parietal and temporal lobes.

## **V4**

V4 draws its input mainly from V2 and provides the major source of visual input to the inferior temporal cortex. It was Zeki (1973) who first drew attention to the colour properties of neurons in V4 and labelled macaque V4 as being the 'colour area'. Later studies questioned this and suggested that V4 is involved in many different visual tasks besides colour (Desimone and Schein 1987; Desimone et al. 1985). Since Zeki (1973) found cells clustered in columns of chromatic preference, no other grouping of cells with common properties has been observed. Although orientation-selectivity was observed in a majority of neurons (Desimone and Schein 1987; Maunsell et al. 1991), Vanduffel et al. (2002) found that orientation-selective cells in V4 were clustered irrespective of their preferred orientation.

The stimuli typically used to study cells in V1 and V2 do not drive V4 neurons with the same success. This might suggest that the optimal stimuli for V4 cells are more complex than those used to study earlier visual areas. For example, Desimone and Schein (1987) found that complex cells in V4 respond to

gratings of high spatial frequency but in any orientation. Another interesting finding of this study is that few cells in V4 that responded poorly or not at all to sine-wave gratings were found to respond well to bars and square-wave gratings, suggesting that edge sharpness may be explicitly represented in V4. Other examples of stimuli used to drive V4 neurons are radial or spiral (non Cartesian) patterns that appear to be more effective than simple (Cartesian) gratings (Gallant et al. 1993). Lesions confined to V4 modestly impair monkey's performance on simple detection and discrimination tasks, but severely disrupts performance on tasks that require monkey to distinguish forms (Merigan 1996). In summary, these findings coupled with the anatomical position of V4 (it is the gateway to temporal lobe) suggest that this area might play an important role in grouping the outputs of early visual areas, providing input to the object selective areas in the inferior temporal cortex (IT in the macaque or LOC in humans).

### ***Dorsal areas (V3A, V3B, V7)***

The dorsal surface of the human occipital cortex, extending from the posterior portion of the intraparietal sulcus forward, contains three retinotopic areas (V3A, V3B, V7; see also section 2.2.2). After cortical visual areas V3 and V4 were identified and named in macaque monkeys, another region was discovered in between them and named "V3 accessory" (V3A) (Zeki 1978a). A human V3A map was described in early cortical mapping papers (DeYoe et al. 1996); however, the first detailed study of the human V3A was from Tootell et al. (1997). The V3B map was discussed by Smith et al. (1998) and its presence was confirmed and clarified by other labs (see Press et al. 2001 for a detailed description of the area). Tootell et al. (1998) described the presence of another retinotopic area, V7. However, a detailed description of the retinotopic map in this region was given by Press et al. (2001).

Tasks involving motion and depth perception produce powerful responses in dorsal regions and particularly in V3A (Backus et al. 2001; Tootell et al. 1997). Brain damage including these areas can result in deficits when interpreting local motion signals (Vaina et al. 2003). Vinberg and Grill-Spector (2008) showed that V3A prefers motion edges, whereas, V7 responds more strongly to edges than to surfaces and random stimuli. Although, human V3A has a retinotopy that is similar to that of the macaque, in functional aspects human V3A appears quite different from its macaque counterpart; human V3A is relatively motion-selective, whereas human V3 is less so. In macaque, the situation is qualitatively reversed. V3 is reported to be prominently motion-selective, whereas V3A is less so (Tootell et al. 1997). On the other hand, V7 is most robustly activated by disconnected edges that do not form a global shape (Vinberg and Grill-Spector 2008). This extends previous findings (e.g. Grill-Spector et al. 1998) that implicated V7 in processing objects by suggesting that local edges alone are sufficient to activate this area. These selective responses to edges across multiple cues may be useful for detecting contours in the visual input.

### ***Areas in the lateral occipital cortex***

A great mystery in vision research is how humans recognise visually presented objects with high accuracy and speed despite drastic changes in their appearance caused by changes in the viewing conditions. Efforts to duplicate this ability in machines have not met much success. The lateral occipital complex (LOC) is described in the literature as an area selective to the shape of visual objects (Eger et al. 2008; Grill-Spector et al. 2000; Kourtzi and Kanwisher 2000; 2001; Kourtzi et al. 2003; Malach et al. 1995; Vinberg and Grill-Spector 2008). LOC can be usually identified in humans as the set of contiguous voxels activated more strongly by images of intact objects than by scrambled versions of the same images. In addition, LOC responds more strongly when subjects view pictures of

objects than textures, visual noise, scrambled objects, or scrambled Fourier phase information (which maintains the spatial frequency spectrum) (Grill-Spector et al. 1998; Malach et al. 1995). Kourtzi and Kanwisher (2001) found that the LOC responded selectively to perceived shape, but not to the contours that define the shape. Likewise, Stanley and Rubin (2003) found that the LOC responded to perceptually salient regions even when there were no perceived boundaries.

The LOC can be divided into at least two putative subdivisions: a dorsal region [LO (lateral occipital)] and a more ventral region [pFus (posterior fusiform)], along the posterior fusiform gyrus. Several studies suggested that different sub-regions of the LOC process different stimulus dimensions. Grill-Spector et al. (1999) and Sawamura et al. (2005) found that the anterior part of the LOC was more invariant to changes in object size and position than the posterior and dorsal part (1999; Sawamura et al. 2005). Stanley and Rubin (2005) showed that the LOC could be subdivided into a posterior, lateral part that responded equally well to abstract 2D shapes and familiar objects, and an anterior and ventral part that responded preferentially to familiar objects.

As part of the lateral occipital (LO) cortex, lateral to the dorsal portion of V3, Larsson and Heeger (2006) identified two retinotopic areas, which they named LO1 and LO2 (see also section 2.2.2). Functionally, LO1 seems to exhibit orientation-selective responses to simple grating stimuli, whereas, LO2 shows no selectivity for stimulus orientation. On the other hand, LO2 was significantly more responsive than LO1 to object stimuli (Larsson and Heeger 2006). The differences in response properties between LO1 and LO2 provide evidence for a segregation of function between the two areas, with LO1 extracting boundary information and LO2 extracting regions and representing shape. The combination of complex shape selectivity and retinotopic organisation in LO1 and LO2 suggests that these areas represent shape information within a spatial coordinate system. This may

be useful for a variety of perceptual organisation processes that rely on spatial relations in the visual image (e.g., segmentation, grouping, boundary extraction). In addition, the location of LO1 and LO2, midway between ventral and dorsal visual processing streams, makes these areas well positioned for integrating information from both the dorsal and ventral streams.

## 1.2 Principles of functional Magnetic Resonance Imaging

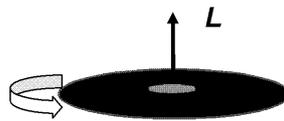
Functional Magnetic Resonance Imaging (fMRI) has revolutionised modern neuroscience, since it was introduced in 1990 (Ogawa et al. 1990). fMRI allows the non-invasive localisation and measurement of neuronal activity. This section describes its basic principles. The first part focuses on the physics related to magnetic resonance. The second part discusses the principles of the Blood Oxygenation Level-Dependant (BOLD) technique that is the most widely used method in functional imaging.

### 1.2.1 Magnetic Resonance Imaging (MRI)

To better conceptualise the process of MR imaging the procedure is divided into three steps. As a first step, the subject is placed into a magnetic field. Second, a radiofrequency pulse is applied; and third the pulse is terminated allowing relaxation of nuclei to occur.

#### ***Nuclear Magnetic Resonance***

When a subject is placed in a strong magnetic field, the atomic nuclei in the brain interact with the field. Atomic nuclei, with an odd mass number (e.g.  $^1\text{H}$ ,  $^{13}\text{C}$ ,  $^{23}\text{Na}$ ) 'spin' on their axis. They can be thought of as spinning around their axis the way the earth turns around its axis or as a vinyl record spins round (Figure 1.6). Spin or angular momentum describes the fact that each point in the record is moving with some velocity about the axis of rotation. However, the spinning velocity depends on the radial position of each point. The direction of the angular momentum vector,  $\mathbf{L}$ , is pointing along the axis of rotation, while its magnitude depends on the velocity with which the record is spinning (Jeppard and Clare 2001).

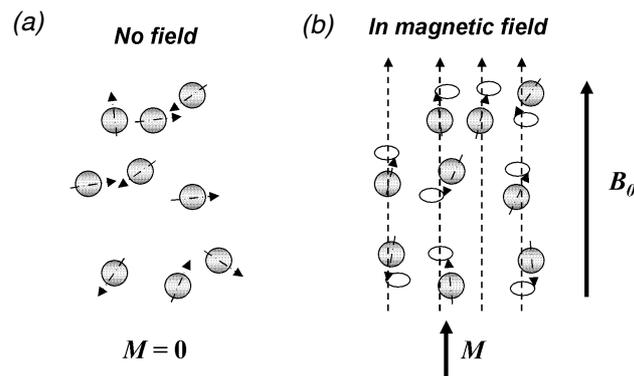


*Figure 1.6 A spinning record has an angular momentum or spin  $L$  pointing along the axis of rotation.*

The nucleus that is of most interest in MRI is the hydrogen nucleus because it is abundant in the water molecules in the brain. It also yields the strongest signal since it has one of the highest nuclear moments (Buxton 2002; Jezzard and Clare 2001). When a hydrogen nucleus enters in a magnetic field, the quantum mechanics predict that it can be found in two possible energy levels. In the lowest energy state the nucleus is aligned parallel to the external magnetic field and in the highest energy state the nucleus is aligned anti-parallel to the external magnetic field.

Any moving charge generates a magnetic field. Since the hydrogen nucleus is a spinning charged particle (the positive charge of one proton), it produces its own magnetic field, called “magnetic dipole moment” (Buxton 2002; Jezzard and Clare 2001). Consequently, each hydrogen nucleus behaves like a tiny magnet bar, which, in the presence of an external magnetic field  $\mathbf{B}_0$  tends to align itself in the direction of the field (Figure 1.7).

The magnetic dipoles, under the effect of the magnetic field, follow a compound, wobbling motion known as precession. During precession the ‘tail’ of each vector is stationary whereas the head is revolving. The same sort of motion is seen when we spin a top; the top spins around its axis, while the plane of the top precesses around the same axis. The precession frequency (Larmor frequency) depends on the type of the nucleus and on the strength of the magnetic field (Buxton 2002; Horowitz 1995; Jezzard and Clare 2001).



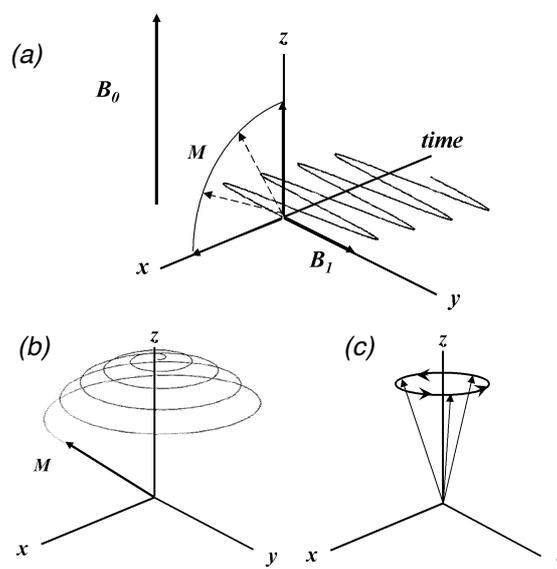
*Figure 1.7 A sample of atoms placed in a strong magnetic field. (a) In the absence of a magnetic field the spins are randomly oriented and there is no net magnetization. (b) In the presence of a magnetic field  $B_0$  the spins partially align with the field, creating a net magnetization vector  $M$ .*

In ideal conditions, all magnetic dipoles placed in a magnetic field will align themselves with the field. However, in physiological temperatures only 1 in 100,000 of the magnetic dipoles is aligned with the external field (Buxton 2002). Nevertheless, even this small alignment is sufficient enough to yield a MR signal. Eventually, slightly more dipoles align parallel to the field than anti-parallel to the field, resulting in a net magnetisation vector  $\mathbf{M}$  pointing to the same direction as the external field  $\mathbf{B}_0$  (Figure 1.7b).  $\mathbf{M}$  can be thought of as a weak, but macroscopic, local magnetic field that is the net result of summing up the magnetic fields due to each of the dipoles. It is actually the difference between dipoles aligned to the field and dipoles opposite to the field. The constant that describes how much time is required for the alignment to occur is called  $T_1$  and is different for different types of nuclei (Buxton 2002; Jezzard and Clare 2001).

### ***The Radiofrequency Pulse***

The magnetisation vector  $\mathbf{M}$  is not directly observable since it is many orders of magnitude weaker than  $\mathbf{B}_0$ . In order for a signal to be detected a brief radiofrequency pulse needs to be applied. The radiofrequency (RF) pulse is a sinusoidally-modulated electromagnetic wave. The direction of the vector  $\mathbf{B}_1$  that

describes the RF pulse is perpendicular to both  $\mathbf{B}_0$  and  $\mathbf{M}$ , and lies along the y-axis (Figure 1.8a). The effect of the RF pulse is to ‘tip’ the  $\mathbf{M}$  vector from its initial, vertical position in the z-axis in the horizontal x-y plane as shown in Figure 1.8a. But the  $\mathbf{M}$  vector does not simply tip over; it follows a complex, spiral motion (Figure 1.8b) which is actually the result of two component motions; the precession around the z-axis (Figure 1.8c) and the tipping motion from the vertical position to the horizontal x-y plane (Buxton 2002; Horowitz 1995).

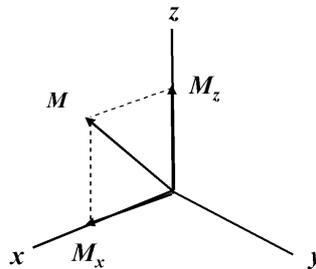


*Figure 1.8 (a) The RF pulse is a sinusoidally modulated field  $\mathbf{B}_1$  perpendicular to  $\mathbf{B}_0$  that causes the magnetisation  $\mathbf{M}$  to tip over in the horizontal plane. (b) The magnetisation vector follows a spiral motion that is the result of the tipping motion and (c) the precession around the vertical axis.*

In order to tip the  $\mathbf{M}$  vector, the frequency of the RF pulse must equal the precession frequency of the magnetic dipoles that we wish to affect. Hydrogen nuclei are affected over all the other precessing dipoles by choosing a RF pulse of equal frequency to the resonant frequency of hydrogen dipoles. Since the net effect of the RF pulse is to tip  $\mathbf{M}$  away from the vertical position, such pulses are usually described by the flip angle they produce (e.g. a  $40^\circ$  pulse or a  $90^\circ$  pulse). The flip angle is determined by the amplitude and duration of the RF pulse (Buxton 2002; Jezzard and Clare 2001).

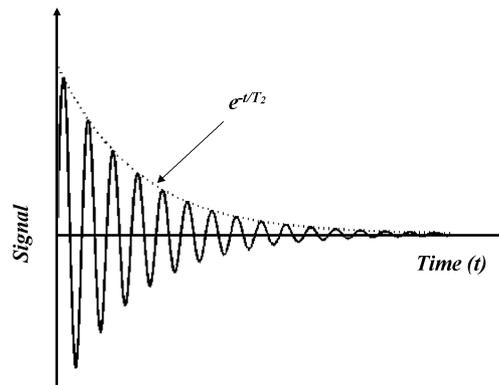
## ***Relaxation***

After the termination of the RF pulse the nuclei return to their initial state, a process known as relaxation. During relaxation the disturbed  $\mathbf{M}$  magnetisation tends to return back to its equilibrium, vertical, position (Buxton 2002; Horowitz 1995). Vector  $\mathbf{M}$  consists of a vertical ( $M_z$ ) and a horizontal ( $M_x$ ) component (Figure 1.9). During relaxation the vertical ( $M_z$ ) vector grows in magnitude and reverts back to the original  $\mathbf{M}$  vector, and the horizontal ( $M_x$ ) magnitude decays to zero.



*Figure 1.9 The  $\mathbf{M}$  magnetisation is broken down into a vertical ( $M_z$ ) and a horizontal ( $M_x$ ) component.*

In MRI it is the transverse,  $M_x$  component that represents the signal. The larger the amplitude of the transverse component, the larger the  $M_x$  component, and the greater the measurable signal. Therefore, when the  $90^\circ$  RF pulse is applied, the  $\mathbf{M}$  vector converts fully into its horizontal  $M_x$  component and the maximum signal is obtained. The signal measured after the termination of the RF pulse is an oscillating wave, which decays away under an exponential envelope. This is called free induction decay (FID) and its profile is shown in Figure 1.10. The time constant describing the decay of the signal is called T2 relaxation and is discussed in more detail below.



*Figure 1.10 The free induction decay (FID). After the  $90^\circ$  RF pulse has tipped the vertical magnetisation into the horizontal plane, the signal decays exponentially in amplitude with a time constant  $T_2$ .*

### **When it all comes together**

The time constant expressing the decay of the signal,  $T_2$  relaxation, is described by the decrease of the  $M_x$  component in the horizontal plane. Similarly,  $T_1$  relaxation is expressed by the re-growth of the  $M_z$  component along the vertical axis. In other words  $T_1$  and  $T_2$  are time constants that describe how long it takes for the amplitude of the  $M_z$  vector to recover and for the  $M_x$  vector to decay respectively (Horowitz 1995; Jeppard and Clare 2001).

$T_1$  and  $T_2$  relaxation times vary among different tissues (Buxton 2002; Jeppard and Clare 2001). This means that after the RF pulse is applied, different tissues require different amounts of time for the corresponding  $\mathbf{M}$  vectors to return to their equilibrium position. In this property lies the ability to obtain contrast between different tissues in MR images. For example, in the human brain  $T_1 \approx 700\text{ms}$  for white matter,  $T_1 \approx 900\text{ms}$  for grey matter and for cerebrospinal fluid (CSF)  $T_1 \approx 4000\text{ms}$  (Buxton 2002). Therefore, if the signal is measured at about 600ms after the termination of the RF pulse the amplitude of the vertical,  $M_z$  component will be larger for white matter than for the CSF and grey matter and the resulting MR image will show white matter bright and CSF dark (Figure 1.11). By choosing different parameters it is possible to highlight different tissues in the

brain. If the MR signal is measured at a time when the difference between the amplitude of the vertical,  $M_z$  components is maximised (as above) the images obtained are known as T1-weighted. On the other hand, if the signal is measured when the difference between the amplitude of horizontal,  $M_x$  components is maximised the images obtained are known as T2-weighted (Horowitz 1995; Jezzard and Clare 2001). Typically during a MR measurement, each time the signal has decayed away a new RF pulse is applied. The delay inserted between the application of subsequent RF pulses is called scan repetition time (TR). In functional imaging this determines the temporal resolution.



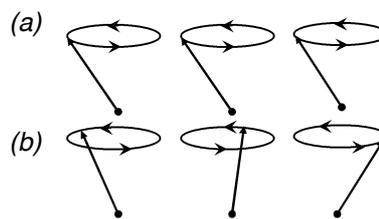
*Figure 1.11 Example of a T1-weighted image of the brain.*

### ***The T2\* decay***

During the delay inserted after each RF, the horizontal magnetisation, and thus the signal, would be expected to decrease exponentially with a time constant T2 according to the FID curve described above. However, what is usually observed is a much more rapid decay than we would expect for a known T2. This faster decay, called T2\* relaxation, has a smaller value than T2 (Buxton 2002; Horowitz 1995).

Fortunately, it is possible to correct for the effects that cause T2\* relaxation and obtain the actual T2 relaxation. But before we move on to the correction procedure let us examine the source of this faster decay. Right after the

termination of the RF pulse the precessing magnetic dipoles have the same phase and add coherently to produce the net magnetisation vector  $\mathbf{M}$ . However, over time they gradually get out of phase (Figure 1.12) and tend to cancel each other out. As a result they no longer add coherently and the signal decays away. The reason for this de-phasing, and hence for the occurrence of  $T_2^*$  decay, is magnetic field inhomogeneity (Buxton 2002). As mentioned above, the magnetic dipoles precess with a frequency (Larmor frequency) that is proportional to the strength of the magnetic field. But because the magnetic field  $\mathbf{B}_0$  is not completely uniform and the head itself is inhomogeneous, the local strength of the field is not equal for all magnetic dipoles. This means that some dipoles experience higher magnetisation than others and as a result they precess faster. This will cause dipoles to de-phase with each other and the signal to decay away because of this destructive interference.



*Figure 1.12 An example of (a) in phase and (b) out of phase precessing magnetic dipoles.*

In functional MRI (fMRI)  $T_2^*$  relaxation is of importance as the technique takes advantage of such inhomogeneities in the magnetic field. However, in anatomical MRI such effects are not desirable and it is possible to correct for them in order to obtain the actual  $T_2$  relaxation. This is accomplished by application of a  $180^\circ$  RF following the  $90^\circ$  RF pulse. At time  $t=0$ , immediately after the termination of the  $90^\circ$  RF pulse, all the magnetic dipoles are in phase. This means they all are at the same point in their cycle of precession and the resulting magnetisation vector  $\mathbf{M}$  lays along the x-axis. Due to de-phasing, at a given time after the termination of the  $90^\circ$  RF pulse, different dipoles will be at different points in their

cycle of precession. This will cause the magnetisation vector to accrue a phase  $\phi$ , with respect to the x-axis. If, however, an  $180^\circ$  RF pulse is applied at time  $t=\tau$  it will cause the magnetisation vector to flip to phase  $-\phi$ . But the magnetisation will continue to accrue phase so at time  $t=2\tau$  its phase will be zero again. This can be illustrated with an example (Figure 1.13). Imagine that the magnetisation vector  $\mathbf{M}$  is the hand of a clock that initially registers a certain time. After time  $t=\tau$  the hand has accrued a phase (angle)  $\phi$ . If at this point the hand is manually shifted symmetrically about the horizontal axis, then at time  $t=2\tau$  it will have returned back to its initial position.

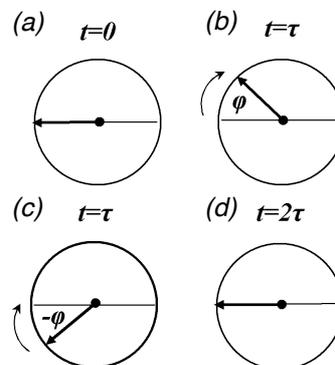


Figure 1.13 An analogy of the spin echo pulse sequence.

It is observed that the re-phasing of the magnetic dipoles generates an echo (spin echo). The time duration required for the magnetisation vector to get back in phase is known as echo time TE. In other words if the  $180^\circ$  RF pulse is applied at time  $t=\tau=TE/2$ , the echo forms at time  $t=2\tau=TE$ . The sequence of the  $90^\circ$ - $180^\circ$  RF pulses, called the spin echo sequence, is widely used in clinical MRI. During such a sequence, the parameters TR and TE are determined so as to generate T1 or T2 weighted images and, thus, emphasise contrast between different tissues.

## 1.2.2 Functional Magnetic Resonance Imaging (fMRI)

MRI techniques can also be used to study functional activity in the brain, in which case this is referred to as functional MRI (fMRI). The growth of this

technique for studying brain function has been quite remarkable over the past 15 years. fMRI is a non-invasive method that takes advantage of the coupling between neuronal activity and haemodynamics in the brain. It is this coupling that gives rise to the so-called Blood Oxygenation Level-Dependant (BOLD) technique which is the most widely used method in fMRI. Its spatial resolution of 3mm or less makes it superior to any of the other non-invasive functional imaging techniques in terms of spatial localisation. The typical temporal resolution of fMRI is about 1.5s. However, the sluggish haemodynamic response downgrades the feasible temporal spacing between trials to 5-8 seconds (Horwitz et al. 2000). This section introduces several of the principles related to the fMRI technique and considers the relationship between haemodynamics and neuronal activity.

### ***The BOLD signal***

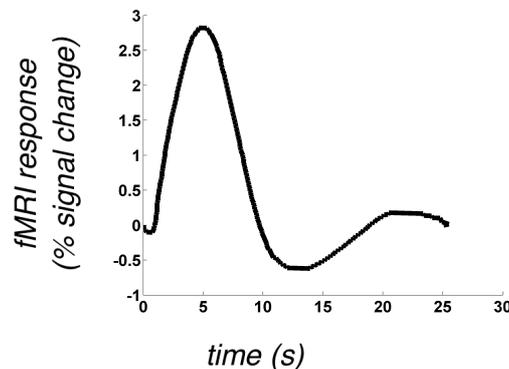
The first fMRI experiment was performed by Ogawa and colleagues (1990). In this pioneering study the brain of a mouse was imaged while breathing different levels of oxygen. It was found that the brain image of a mouse breathing 100% oxygen was rather bright, uniform and featureless (low contrast). But when the mouse breathed 20% oxygen many dark lines appeared indicating the major structures of the brain (higher contrast). This effect was later explained by taking into account the magnetic properties of oxygenated and deoxygenated blood. Deoxyhaemoglobin (red blood cells without an oxygen molecule attached to it) is a paramagnetic substance (is slightly attracted to a magnet) and introduces an inhomogeneity into the nearby magnetic field. On the other hand, oxyhaemoglobin (red blood cells with an oxygen molecule attached to it) is a weakly diamagnetic substance (is slightly repelled by a magnet) and has very little effect on the surrounding magnetic field (Heeger and Ress 2002; Wandell and Wade 2003).

When neurons are activated, there is an increased supply of oxygenated blood to the brain. As a result, the levels of deoxyhaemoglobin will fall introducing

higher field inhomogeneities. The BOLD technique is designed primarily to measure changes in such inhomogeneities. The lower the level of deoxyhaemoglobin in the blood the higher the field inhomogeneities, and the higher the BOLD signal (Heeger and Ress 2002; Wandell and Wade 2003). The scanning technique that is most commonly used in fMRI is Echo Planar Imaging (EPI) (Mansfield 1977). Techniques of this kind take account of the field inhomogeneities and therefore the BOLD effect is more pronounced. The signal measured by the BOLD technique is the  $T_2^*$  signal.

### ***The haemodynamic response function (HRF)***

If neurons are stimulated for only a brief period of time (e.g. after the appearance of a transient stimulus) the BOLD response to this stimulation is described by the haemodynamic response function (HRF), as shown in Figure 1.14. The characteristic shape of the HRF consists of three epochs, the properties of which are not fully understood.



*Figure 1.14 Haemodynamic response function (HRF) measured as the response to a brief neuronal stimulation.*

Immediately after the stimulus onset there is a brief period of 0.5–1s during which the BOLD signal slightly decreases. Although, this ‘initial dip’ is a subtle effect and is not always observed (Hoge and Pike 2001), some users believe that it may map more precisely the spatial location of neuronal activity. Subsequently,

there is an increase in the BOLD signal that peaks at about 5-8 seconds after the stimulus commences. This signal increase is due to an oversupply of oxygenated blood. There are two theories aiming to explain this increase in blood flow. The first states that the increased blood flow serves to deliver the oxygen used by the active neurons. Although there is strong evidence in support of this view, it appears that the oxygen supplied by the increased blood flow is much larger than the amount used by neurons. An alternative theory is that the increase in blood flow compensates for the glucose used by neurons. This is supported by the fact that the glucose consumption by the active neurons is directly proportional to the increase in blood flow. In any case, the metabolic mechanisms that cause the increase in blood flow are not fully understood (see Heeger and Ress, 2002 for a review). However, the result of the oversupply in oxygen is the decrease in the level of deoxyhaemoglobin in the blood that causes an increase in the fMRI signal. Finally, in the last epoch of the HRF the BOLD response returns to baseline, accompanied by a 'post-stimulus undershoot' during which the response is briefly negative.

### ***Linear systems analysis***

The haemodynamic response is not a linear function of the stimulus pattern since it includes transient features such as the post stimulus undershoot. However, the analysis and interpretation of fMRI data relies on the hypothesis that the haemodynamic response is a linear function of the average neuronal activity (Boynton et al., 1996). This assumption is known as the *linear transform model*. Although it is an approximation, the linear transform model appears to describe very well the link between neuronal activity and the fMRI signal (Boynton et al. 1996). This can be shown as follows. Since the response to a brief stimulus is described by the HRF, under the assumption of linearity, the response to a more

sustained stimulus can be predicted by convolving the stimulus pattern with the HRF, as shown in Figure 1.15.

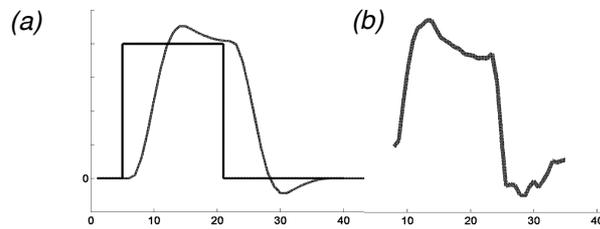


Figure 1.15 (a) Assuming that the BOLD response is linear, the response to a sustained stimulus (square) can be described by the convolution of the HRF with the stimulus pattern. (b) The measured BOLD response to the stimulus.

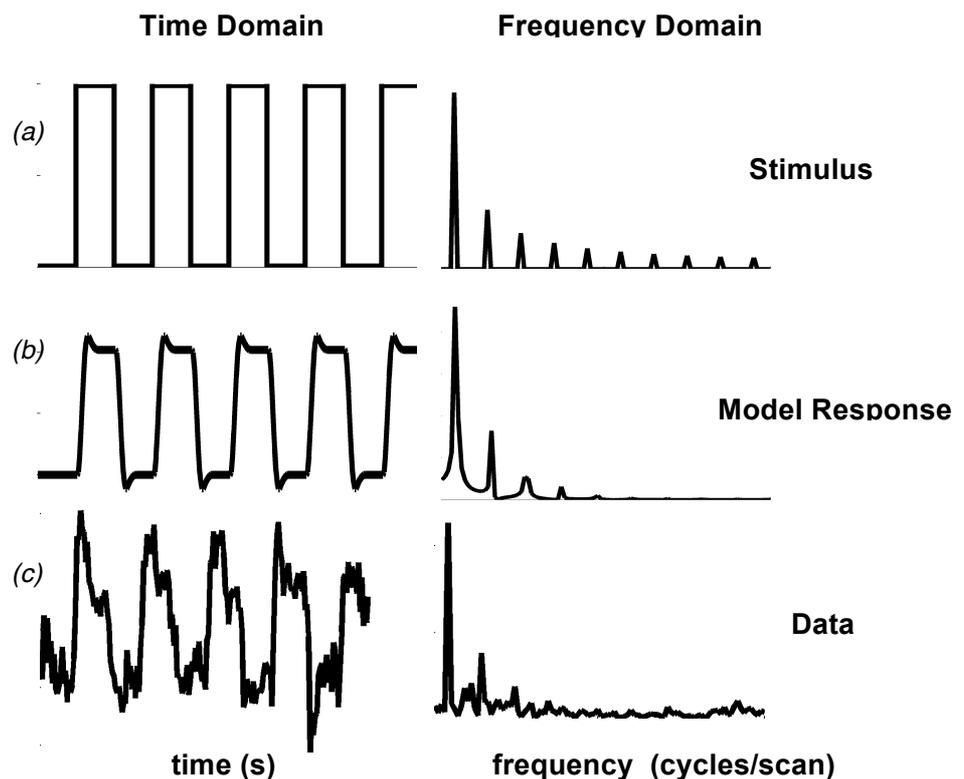


Figure 1.16 The Fourier components of (a) a blocked stimulus pattern, (b) the predicted haemodynamic response and (c) the measured BOLD signal.

One way to estimate the statistical significance of the measured fMRI signal is to compute the correlation coefficient between the measured signal and the

predicted response. The coefficient can be used as a statistical parameter. By choosing a particular threshold value it is possible to select voxels that are significantly activated.

A similar approach is based on the Fourier transform of the time course (Engel et al. 1994). If an fMRI scan consists of 5 cycles (alternations between the onset and offset of a stimulus), the Fourier spectrum of the time course peaks at a prominent fundamental frequency (5 cycles/scan) and a series of odd multiples of the fundamental frequency (Figure 1.16a). The mathematically predicted response function is smoother; therefore, the higher frequency components are not apparent (Figure 1.16b). Finally, the spectrum of the measured data contains the fundamental frequency and only one apparent harmonic (Figure 1.16c). A metric describing the amount of correlation between the activation of a voxel to the stimulus time course can be derived by comparing the amplitude of the fundamental frequency to the amplitude of the other frequencies.

This approach bears several advantages. For example, there is evidence that the haemodynamic delay may vary for different parts of the brain and this may be a problem for the standard correlation analysis. If the delay used in the HRF is not the same as the actual delay, the correlation coefficient will be reduced and some activated voxels could be missed (Buxton 2002). Although there are techniques to account for this, the method described above deals more effectively with the problem. If for example the model function is chosen to be a sine wave, the haemodynamic delay is equivalent to a phase shift of the wave. Using the Fourier transform the best fit for the delay can be easily computed.

### **1.2.3 Limitations of the BOLD technique**

The use of fMRI has undoubtedly revolutionised modern cognitive neuroscience. The main advantage is that it allows the non-invasive localisation

and measurement of neuronal activity. fMRI led modern neuroscience to a new era of research, complementary to the invasive techniques used in animal models, to study the real-time function and dysfunction of the human brain. However, there are several limitations of the BOLD technique that a user should bear in mind. The most important is that the BOLD fMRI signal is an indirect measure of neuronal activity and relies on the assumption that neuronal activity and haemodynamics are indeed linked.

A lot of research has been carried out to answer the question of whether or not the BOLD fMRI signal is loosely or tightly coupled to neuronal activity. This relationship appears to depend on a number of factors (see Heeger and Ress 2002 for a review). One is the extent to which the signal from large veins is taken into account. The BOLD signal is sensitive to the size of veins, since the relative decrease in deoxyhaemoglobin is larger in large veins than in small veins. As a result, the increase in blood flow is not very specific to the area of neuronal activation. Usually, blood is supplied to a larger area of the brain than just the activated area. Consequently, the area that shows an increase in fMRI signal might be larger than the area of neuronal activation, resulting in spatial blurring. It is possible to modify the signal acquisition in order to de-emphasise the BOLD signals from larger veins, by suppressing signals that are associated with higher flow velocities. This can be also achieved by using stronger magnetic fields. Another solution is to choose appropriate experimental protocols and data analysis methods. This way it is possible to de-emphasise the signal from large veins by forcing the blood flow and oxygenation in large veins to remain roughly constant throughout the experiment. Such a design, used to study retinotopic organisation in the visual cortex, is described in the General Methods chapter.

In an attempt to further understand the link between the BOLD signal and neuronal activity, Logothetis et al. (2001) compared simultaneously fMRI and

direct neuronal signals in anaesthetised monkeys for stimuli of different contrasts. They concluded that fMRI and neuronal signals were not simply proportional to one another. They found that there is a monotonic but non-linear relationship between the fMRI signal and neuronal responses, such that the fMRI responses increase more rapidly than neuronal responses at low contrasts, but less rapidly at high contrasts. Another important finding of the same study is that large signal differences across different recording sites might be obtained.

Not surprisingly, the strength of the connection between neuronal activity and haemodynamics has emerged as one of the most important research areas in neuroscience. To summarise the issues discussed above, the BOLD fMRI signal appears to be a good approximation of neuronal activity for some recording sites, in some brain regions, using certain experimental protocols. Basically, the fMRI signal can be seen as a qualitative (Howseman and Bowtell 1999), smoothed (Horwitz et al. 2000), and sluggish measurement of the underlying neuronal activity. But despite its limitations, the BOLD technique provides additional insights beyond other methods (see Wandell and Wade, 2003 for a comparison between fMRI and other techniques in vision). Furthermore, it has been successful in answering fundamental questions about the properties and organisation principles in the human visual cortex. Some of which are discussed in the following chapters.

### 1.3 Experimental techniques

As discussed in the previous section, BOLD measurements are ultimately limited in spatial resolution by, amongst other things, the spatial scale of the local vascular system. In many studies, the aim is to quantify the selectivity of clusters of neurons on a spatial scale much smaller than the  $3 \times 3 \times 3 \text{mm}^3$  volume of a voxel used typically in current fMRI experiments. Recent fMRI studies have demonstrated new methods for revealing separate subpopulations of neurons selectively tuned for different stimuli even when these neurons are intermingled at a spatial scale that is smaller than the sampling resolution of the measurement.

The first technique, fMRI adaptation, relies on the observation that the BOLD response in a given voxel is reduced after prolonged presentation of a stimulus, and that this reduction is selective to the characteristics of the repeated stimuli (adapters). The second technique, multi-variate pattern analysis (MVPA), makes use of multi-variate statistics to recover small biases in individual voxels in their responses to different stimuli. This section provides a general introduction to the two techniques. Further experimental details are discussed in the General Methods section and within the relevant chapters.

#### 1.3.1 Adaptation

One of the most fundamental properties of the visual system is its flexibility. The ability to adjust, or *adapt* in response to a changing environment is a fundamental principle of the visual pathways, beginning at the earliest stages and continuing into the central brain representations.

Ideas about the evolutionary value of adaptation have been motivated by two main considerations. First, the visual system must be re-calibrating itself as it maps environmental conditions onto patterns of neural activity. Self-calibration is the property of a system to change itself in response to changes in the

environment (recalibration) and to allow for error-correction within the system. Second, neurons in the visual system have a limited dynamic range for the coding of visual stimuli. Therefore, adaptation might serve as an optimisation mechanism that reduces the transmission of redundant information. This dynamic range optimisation might help to improve the visual system's performance by maximising the effective bandwidth available for the transmission of new information about the stimuli. The functional considerations of self-calibration and redundancy reduction suggest that the visual system should be viewed as a dynamic system. Adaptation mechanisms serve this dynamic system by continually operating to match the coding employed to the statistical structure of the environment (Clifford 2002).

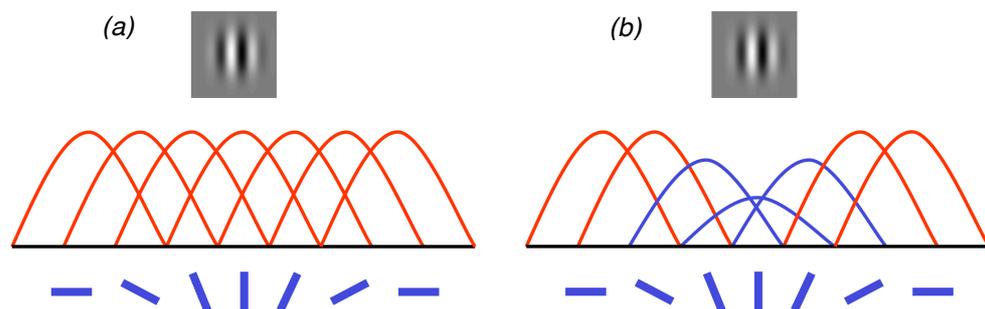
Adaptation reflects the phenomenon of reduced responses to repeated presentation of a specific stimulus. Psychophysically adaptation can be demonstrated in various ways. For example, prolonged viewing of a high contrast vertical pattern reduces the perceived contrast of a similar pattern in a subsequently viewed test stimulus (contrast adaptation) (Blakemore and Campbell 1969). Also, viewing an oriented line for a long time causes subsequent lines to appear tilted away from the adapting orientation (tilt aftereffect) (Gibson and Radner 1937; Greenlee and Magnussen 1987). Aftereffects reveal a gap between appearance and reality, and remind us that what we see is determined by how visual information is coded in the brain, and not simply by how things 'really are'. The use of adaptation visual science has been so widespread that it has been termed the 'psychophysicist's electrode'.

### ***Models of adaptation***

It is not known exactly how and why the response properties of cells in visual areas change during and after a period of exposure to an adapting stimulus. There are multiple potential neural causes of the adaptation measured psychophysically, with single-cell recording, or with fMRI. Three models have

been previously proposed in the literature (see Grill-Spector et al. 2006 for a review): (a) the *fatigue* model, according to which the amplitude of firing of stimulus responsive neurons decreases, (b) the *sharpening* model, according to which fewer neurons respond, and (c) the *facilitation* model, according to which the latency and the duration of neural activity are shortened.

The simplest model of adaptation is the fatigue model. According to this, all neurons selective to a stimulus dimension show a proportionate reduction in their response to repeated presentations of the same stimulus. As a consequence the mean population firing rate declines. One prediction from this model is that the amount of adaptation is greater in neurons that respond optimally to the adapting stimulus than in other neurons. An advantage hypothesised for such a mechanism is that the sensitivity of the system to stimuli that are different from the repeating stimulus is increased, thereby providing a mechanism for ‘novelty detection’. Reducing the firing rate might also help prevent saturation of the neural response function by increasing its dynamic range (see above). This attractively simple model can account for a large number of aftereffects resulting from adaptation (see Georgeson 2004).



*Figure 1.17 Re-tuning of cell responses after adaptation according to the ‘fatigue’ model. (a) Schematic response tuning curves for a group of orientation-selective cells before adaptation. Blue bars indicate each cell’s preferred orientation. (b) After adapting to a stimulus of a given orientation, the responses of cells tuned to similar orientations are scaled down.*

The sharpening model, initially described by Desimone (1996), suggests that adaptation results in sparser representation of stimuli. According to this model, some (but not all) of the neurons that initially responded to a stimulus will adapt to subsequent presentations of that stimulus. Critically, the sharpening model predicts that neurons adapting little or not at all to a repeated stimulus are those who are highly selective for that stimulus. Thus, repetition-related changes are viewed as a learning process, in which neuronal tuning is sharpened and as a consequence, the representation becomes sparser resulting in fewer responsive neurons in total. Such a model could explain the faster processing of repeated stimuli observed in many adaptation studies (Grill-Spector et al. 2006).

In contrast, the facilitation model predicts that repetition causes faster processing of stimuli, resulting from shorter latencies or shorter durations of neural firing. This approach might be extremely useful in modeling the coupling between neuronal activity and the BOLD signal. Given that the hemodynamic response measured by fMRI integrates over several seconds of neural activity, a shorter duration of neuronal activity results in decreased amplitude of the fMRI signal. A shorter duration of neural activity is also consistent with the earlier peaking of the fMRI response and might explain why decreases in firing rate can appear to arise after the initial visual response.

The three models differ in their predictions about whether adaptation is strongest for the preferred or the non-preferred stimuli. The sharpening model predicts that neurons showing little or no adaptation to a repeated stimulus are highly selective for that stimulus. By contrast, both the fatigue and facilitation models predict that adaptation is proportional to the initial response. Thus, neurons that respond optimally to the adapting stimulus should also show the greatest adaptation (but see voxel-wise comparisons in Chapter 3). Although it is possible that a single model could apply under all conditions, it might be the case

that different models apply in different brain regions and under different experimental conditions and designs. To distinguish between the models it is necessary that new single-cell experiments are performed and more elaborative computational models are developed.

### ***fMRI adaptation***

Adaptation has been used behaviorally (Blakemore and Campbell 1969; Bradley et al. 1988; Snowden and Hammett 1996), with single-cell recordings (Ohzawa et al. 1985; Sclar et al. 1989; Solomon et al. 2004), and more recently the method has been modified for use in conjunction with fMRI scanning. The relationship between the three methods is not straightforward. Considering, for example, the relationship between adaptation measured with electrophysiology and fMRI, one should take into account the fact that fMRI and neuronal signals are not simply proportional to one another (Logothetis et al. 2001; Logothetis 2002). In addition, large variations might be obtained across different recording sites and with different cell recording techniques.

In essence, the basis of the method is that after prolonged presentation of a particular stimulus, the BOLD response to that stimulus is reduced in areas sensitive to it, more than the response to other stimuli. In the example of orientation-selectivity, after prolonged viewing of a high-contrast grating of a particular orientation the fMRI response to a low-contrast probe of the same orientation is reduced relative to the response to a differently oriented probe. The fact that adaptation is not uniform across the domain of interest (e.g. orientation) suggests that subpopulations of neurons are responding selectively to particular subregions of the space. FMRI-adaptation has been widely used to study, orientation selectivity (Engel 2005; Fang et al. 2005; Larsson et al. 2006), perception of motion (Huk and Heeger 2002; Krekelberg et al. 2005; Tolias et al. 2001), objects (Grill-Spector et al. 1999; Kourtzi and Huberle 2005; Kourtzi et al.

2003; Sayres and Grill-Spector 2006; Vuilleumier et al. 2002), and faces (Andrews and Ewbank 2004; Grill-Spector and Malach 2001; Henson et al. 2002).

The method of fMRI adaptation is hypothesised to be able to characterize the functional properties of neural populations in a sub voxel scale (Grill-Spector and Malach 2001). The limited spatial resolution of fMRI (which is in the order of mm) places a barrier for studying the properties of groups of selective neurons, which are usually co-localised within the volume of a single voxel. Using standard imaging techniques it is impossible to access whether the measured BOLD signal is a mixture of neurons each tuned to a different stimulus category, or whether it is the outcome of the activity of a homogeneous group of neurons that share a common property. If for example, one subpopulation of neurons responds to a particular stimulus category and a separate subpopulation of intermingled neurons responds equally strongly to a different stimulus category, there will be no difference in the overall response (averaged across both subpopulations) to the two stimulus categories. This makes it difficult, in many cases, to infer from the fMRI data the properties of the underlying neurons.

The two subpopulations could, however, be distinguished using an adaptation protocol. A sampled region of tissue containing neurons selectively tuned for one stimulus category will adapt after repeated presentation of these neurons' preferred stimulus. If the same tissue contains a second, separate subpopulation of neurons selectively tuned for a different stimulus category, that tissue would also adapt to that stimulus category. Repeated presentation of one stimulus category will not, ideally, affect the responses to the other stimulus category. Selective adaptation to a particular stimulus category thus provides a measure of the stimulus selectivity of a subpopulation of neurons that is unaffected by the stimulus selectivity of other neurons in the same region of tissue.

### 1.3.2 Multi-variate Pattern Analyses (MVPA)

The study of neuronal selectivity in a sub-voxel scale using indirect methods such as adaptation requires the averaging of fMRI responses over many trials and participants. Averaging takes place in such a way that could obscure much of the information present in the spatial pattern of individual brain responses. Another characteristic is that fMRI responses are considered separately in individual voxels. Time courses are first extracted in individual voxels and then averaged together to obtain a mean response within a region of interest (ROI). Because each voxel is treated as a separate entity as far as statistical analysis is concerned, conventional fMRI analyses are considered as *univariate*.

However, fMRI data are fundamentally *multivariate*; a single acquisition in time contains information at thousands of spatial locations. In contrast to conventional univariate analyses, recent studies have shown that sensitivity of fMRI paradigms can increase significantly by taking into account patterns of activity present across many voxels at the same time (Cox and Savoy 2003; Haynes and Rees 2006; Norman et al. 2006). An example of the rationale behind pattern-classification techniques is given in *Figure 1.18*. By considering fMRI responses in many voxels simultaneously, instead of analysing one location at a time, multi-variate pattern analyses (MVPA) take into account the fine-grained information in individual voxels. Given the goal of detecting the presence of a particular mental state in the brain, the primary advantage of MVPA methods over univariate methods is increased sensitivity. Conventional fMRI analysis techniques try to find voxels showing a statistically significant response to the experimental conditions. To increase sensitivity to a particular condition, these methods spatially average across voxels that respond significantly to that condition.

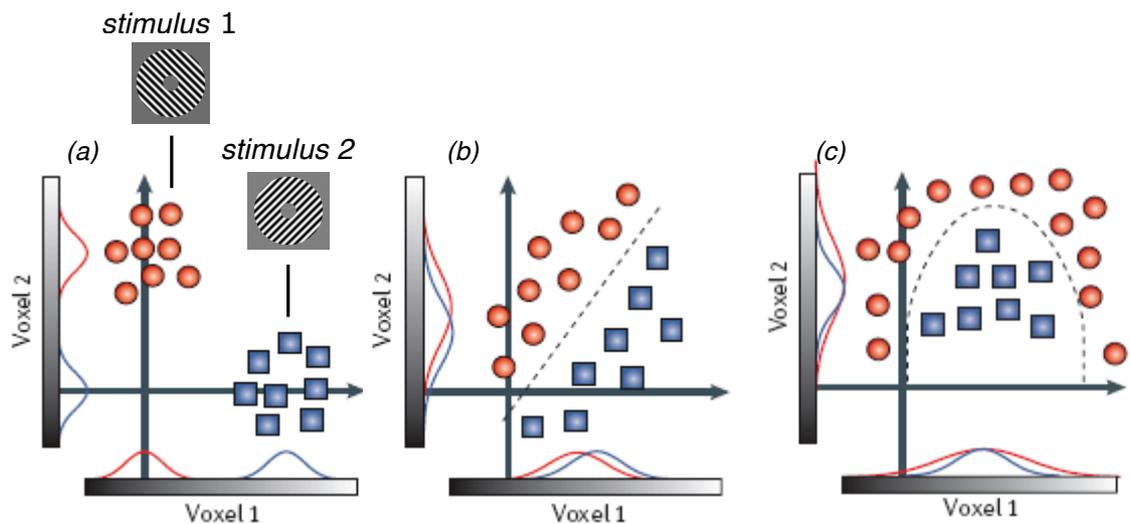


Figure 1.18 Hypothetical fMRI activity in two voxels in (a) an ideal univariate situation. The activity in voxel 2 is plotted versus the activity in voxel 1 and each point corresponds to a single fMRI measurement. Voxel 2 is more active in the first condition, but it is not active in the second condition. Because the projected distributions of the responses do not overlap, a conventional univariate analysis could easily discriminate the responses measured under the two conditions. However, in situation (b) both voxels are relatively active in both conditions and the projection of the responses overlap. In this case conventional analyses would be hard to discriminate fMRI activity. Nevertheless, data from each condition occupy a distinct region of the two-dimensional space. A linear decision boundary can be used to distinguish the response distributions. In panel (c) a linear decision boundary would not be sufficient to discriminate responses, therefore, a non-linear classification would be required (modified from Cox and Savoy 2003; Haynes and Rees 2006).

Although this approach reduces noise, it also reduces signal in two important ways: First, voxels with weaker responses to a particular condition might carry some information about the presence/ absence of that condition. Second, spatial averaging blurs out fine-grained spatial patterns that might discriminate

between experimental conditions. Like conventional methods, the MVPA approach also seeks to boost sensitivity by looking at the contributions of multiple voxels. To avoid the signal-loss issues mentioned above, MVPA methods make use of small differences in the fMRI response of different voxels thought to result from small biases in the spatial distribution of the neural subpopulations sampled by each voxel. By 'learning' the pattern of these small biases across a large number of voxels in an independent training set, multi-variate pattern analysis can successfully discriminate between stimuli in a novel set of trials.

Several reports have shown that such multivariate techniques can reliably distinguish between responses to different stimuli, where more conventional, voxel-wise univariate approaches, or signal averaging across whole ROIs could not. MVPA techniques have been used to decode the orientation of gratings (Haynes and Rees 2005; Kamitani and Tong 2005), direction of motion (Kamitani and Tong 2006), object categories (Eger et al. 2008; Haushofer et al. 2008; Haxby et al. 2001), to study visual categorisation (Li et al. 2007) and also the encoding of global form (Ostwald et al. 2008). Some of the aforementioned papers went well beyond by predicting from the fMRI patterns of activity the orientation of invisible or masked stimuli, thus, performing a kind of 'brain-reading'.

Classification performance depends a great deal on the number and choice of voxels included in the analysis (Cox and Savoy 2003; Ku et al. 2008). First, the number of voxels (features) defines the dimensionality of the problem. Classification performance decreases dramatically as the number of features exceeds the number of training points, therefore, it is necessary to choose only an appropriate subset of the total voxel population. Second, voxels that contain little information about the discrimination being made only add unrelated noise to the classifier and degrade performance. As a result, many pattern recognition applications contain a 'feature selection' step in which only a subset of voxels is

selected that contains enough information to perform the classification. Of course, the type of the feature selection procedure depends on whether the analysis is restricted to predefined ROIs or it includes locations across the whole brain.

Once a set of voxels has been selected for pattern classification, data are stored into pattern matrices corresponding to the pattern of activity across the selected voxels at a particular time in the experiment. Brain patterns are then labeled according to which experimental condition generated the pattern. This labeling procedure needs to account for the fact that the hemodynamic response measured by the scanner is delayed in time, relative to the neural event under investigation.

Once the fMRI responses are stored into pattern matrices, a subset of these labeled patterns (the train sample) are fed into a multivariate pattern classification algorithm. Based on these patterns, the classification algorithm learns a function that maps between voxel activity patterns and experimental conditions. After the classifier is trained, the next step tests the generalisation of the function. Given a new pattern of brain activity, not previously presented to the classifier (the test sample), can the trained classifier correctly determine the experimental condition associated with that pattern?

In machine learning literature there is an enormous range of classification algorithms that can be potentially used in MVPA studies (Duda et al. 2001). The majority of MVPA studies have used linear classifiers, including correlation-based classifiers (Haxby et al. 2001), neural networks without a hidden layer (Polyn et al. 2005), linear discriminant analysis (LDA) (Haynes and Rees 2006; Haynes and Rees 2005), linear support vector machines (SVM) (Cox and Savoy 2003; Kamitani and Tong 2005), and Gaussian Naive Bayes classifiers (Mitchell et al. 2003). All these classifiers compute a weighted sum of voxel activity values. This

weighted sum is then passed through a decision function, which creates a threshold for saying whether or not a category is present.

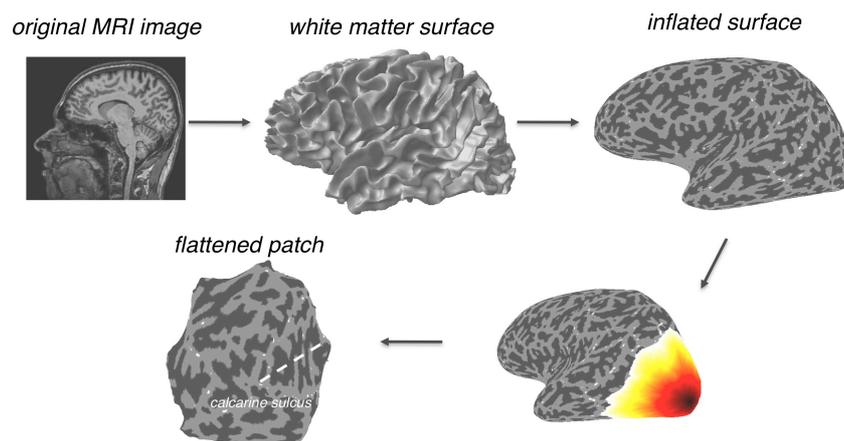
Because these techniques are relatively new and not fully understood, users should be extremely cautious about the interpretation of MVPA results (Bartels et al. 2008). Another major concern about MVPA methods is the extent to which 'brain reading' experiments are ethically acceptable in view of the potential use of this technology towards human rights violation in the future. All in all, MVPA methods have evolved extensively in the last few years and it is expected that MVPA methods will continue to evolve, as better algorithms become available in the coming years. Improvements in the spatial resolution of fMRI will make it possible to resolve even finer-grained cognitive distinctions. For all of these reasons, it is believed that MVPA has a bright future as a tool for characterising how information is represented and processed in the brain.

## 2 General Methods

The following section refers to general experimental methods, which were implemented during this thesis. These include the computational steps performed in order to segment and flatten cortical surfaces, the technique for delineating the borders of retinotopic visual areas, and the fMRI scanning settings. Further experimental details are described within the relevant chapters.

### 2.1 Cortical segmentation and flattening

Although the spatial resolution of functional EPI images is quite high in terms of spatial localisation, the images themselves do not have sufficient resolution in order to overlay functional data on them. It has become typical in fMRI studies to acquire a set of high-resolution anatomical images and overlay functional data onto these images showing cortical structure. These structural images are usually T1-weighted, with good grey-white matter contrast, and are acquired during a separate scanning session.



*Figure 2.1 Anatomical T1-weighted images are processed so that the grey/white matter surface is segmented and inflated. The area surrounding the occipital sulcus is then extracted and flattened. The inflated and flattened representations help to better appreciate spatial locations in the brain.*

However, much of the cortical surface in the structural images is obscured from view by a complex pattern of folds, making hard to interpret spatial locations and visualise functional data. This section describes the computational processing performed to extract the surface between the grey and white matter, inflate this surface and flatten the occipital cortex. Figure 2.1 summarises the analysis stream. These steps help to better appreciate spatial locations in the brain and are particularly useful for visualising retinotopic maps and delineating the borders of visual areas.

### **2.1.1 Analysis Stream**

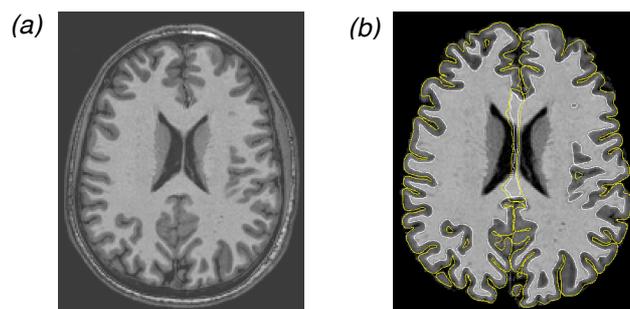
The most common representation of MR images is as 2-dimensional slices throughout the brain volume. However, using this format it is difficult to appreciate the spatial relationships between different points in the brain. Furthermore, when seen as separate slices, the position of any but the most familiar anatomic features is hard to infer. The method for surface extraction involves a series of steps starting with a T1-weighted anatomical image of the brain and resulting in two cortical surfaces for each hemisphere, one corresponding to the boundary between grey and white matter, and the other to the boundary between grey matter and CSF. The computational analysis was performed using the freely distributed software tool, SurfRelax, developed by Jonas Larsson (Larsson 2001) based on the FSL software library.

The analysis begins with a number of preprocessing steps, which include: intensity normalisation, non-brain tissue removal, filling of the ventricles and subcortical nuclei and segmenting the hemispheres from each other and from the brain stem. After preprocessing a template generated for each hemisphere is deformed onto the surface of white matter. Finally, the surface is extracted and refined after which may be inflated and flattened.

### ***Preprocessing***

MR images are susceptible to intensity variations due to magnetic field inhomogeneities. These inhomogeneities result in the same tissue having different intensities at different points. This might be a problem when image intensities are used to delineate tissues (for example finding the boundary between grey and white matter). Thus, in the first step of preprocessing image intensities are normalised using a non-parametric heuristic approach (Larsson 2001). Next, non-brain tissues (e.g. skull) are removed by means of a deformable surface algorithm similar to that described by Dale et al. (1999).

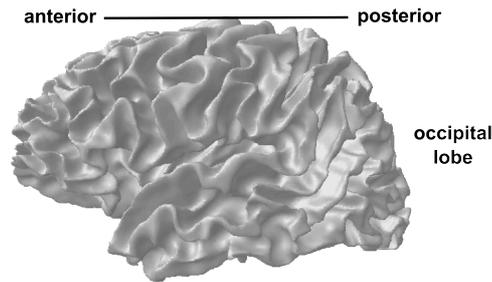
The segmentation procedure extracts the cortical surface at the boundary between grey and white matter. However, the medial white matter surface is connected to subcortical structures such as basal ganglia, ventricles and the thalamus. These structures are of no interest when generating the cortical surface. Hence, they are automatically identified and assigned the mean white matter intensity.



*Figure 2.2 (a) A slice of an anatomical T1-weighted image in horizontal view, and (b) the result of the initial preprocessing steps. During these steps the non-brain tissues (e.g. skull) are removed and the left and right hemispheres are segmented from each other.*

Then the two cortical hemispheres and subcortical structures are initially separated using an automated template fit. Predefined templates of the

hemispheres and the cerebellum are computationally deformed to match the shape of the target brain. Once the deformation process is completed the templates are smoothed by a series of closing operations and the surfaces of the templates are extracted (the result is shown in Figure 2.2b). These surfaces are then used to segment the grey/white matter boundary.



*Figure 2.3 The extracted surface between the grey and white matter for the left hemisphere. This representation helps to appreciate the spatial relation between different locations, but much of the surface is hidden from view.*

Although the templates generated in the previous step could be potentially used as a starting point for white matter extraction, they are not suitable for this purpose as they deviate substantially from the actual shape of the target brain white matter. The performance of the surface extraction algorithm is highly dependent on the shape of the initial template; therefore, a template volume with spherical topology is generated directly from the filled white matter of the target brain. An anisotropic diffusion filter smooths the template volume and the final template is extracted by thresholding. The template generated in the previous step is deformed onto the filled white matter volume by a series of topologically constrained erosions and dilations. These steps ensure that the deformation converges in such a way that the final volume is anatomically correct. Finally, the surface of the white matter is extracted from this volume using a surface extraction algorithm (Larsson 2001). The resulting surface is a good approximation of the grey/white matter boundary (the result is shown in Figure 2.3) and can be used directly as it is. However, to further improve the fit in regions

of high curvature and to generate a surface representation of the outer grey matter, an optimisation step is applied to the surface as described in the following section.

### ***Surface optimisation***

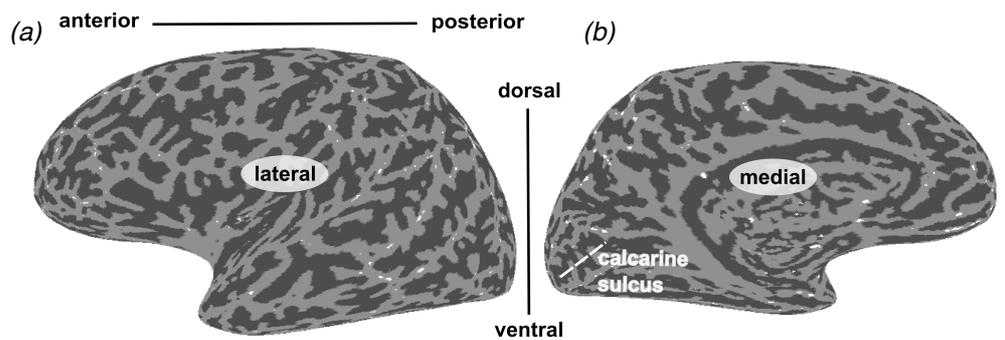
To improve the fit between the surface and the grey/white matter boundary, the surface is iteratively deformed onto the original intensity normalized MR image using a method similar to that described by Dale et al. (1999). This deformation is done in three stages. First, the surface is expanded a few millimeters toward the middle of the cortex. Second, the surface is shrunk back onto the white matter. The surface expansion taking place in the first step ensures that the subsequent shrinking captures small fragments of white matter which may have been missed in the surface extraction stage. Third, the surface is expanded from the middle of the cortex to the grey-CSF boundary.

### ***Surface relaxation and visualisation***

Although the representation generated in the final step greatly improves someone's ability to appreciate the relative spatial relations between different locations of activity, still much of the brain area is buried within sulci whose shape varies across individuals and whose form is quite complex. To make most of the brain surface visible, the surface between the grey and white matter can be inflated and flattened.

To speed up inflation and flattening and to reduce distortions, these processes are performed at multiple resolution levels. First, the surfaces are resampled to a low and to an intermediate resolution. The resampling method is similar to that described by (Wandell et al. 2000). The low-resolution surface is relaxed by iteratively moving each vertex toward the mean position of its neighbors until it is sufficiently smooth. The vertex displacement parameters used

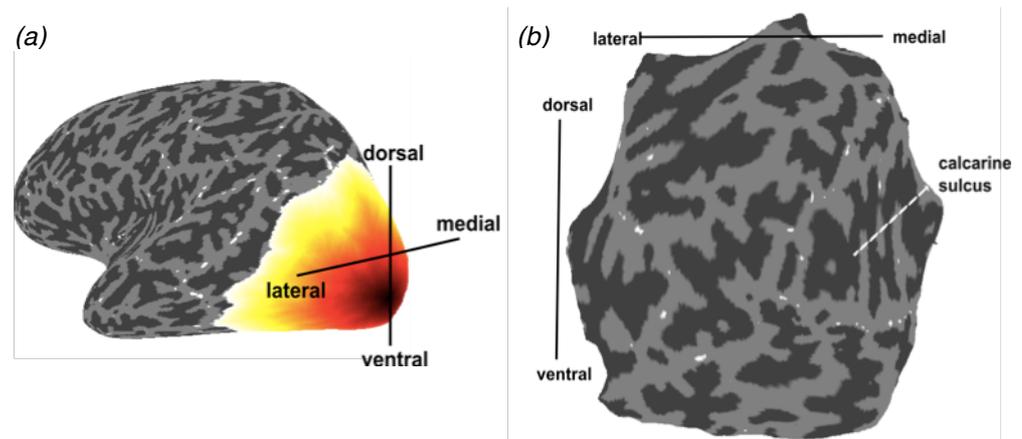
during the low-resolution surface relaxation step are then applied to the corresponding vertices in the intermediate resolution surface. The intermediate resolution surface is then relaxed under the same constraints. The displacement parameters for this surface are then applied to the original surface, and this surface is then relaxed until it is sufficiently smooth. The use of resampled surfaces makes the relaxation process quite fast, although the resampling process itself is computationally demanding. The result of the surface inflation process is shown in Figure 2.4. In the inflated representation the lighter and darker regions correspond to the major gyri and sulci respectively.



*Figure 2.4 The surface between the grey and white matter of the left hemisphere (of Figure 2.3) is inflated, and the result is shown in lateral (a) and medial (b) view. The anatomical location of calcarine sulcus is also indicated.*

Finally, certain cortical regions, including visual cortex, are visualised even more easily as a simple flattened patch of cortical sheet. For projecting the inflated surface onto a plane, cuts are manually introduced on the inflated surface. The cut surface is resampled in the same way as in the inflation process. The low-resolution surface is projected onto the x-plane as described by Wandell et al. (2000). This surface is then relaxed following a similar procedure as in the inflation step.

Figure 2.5 shows the area of the occipital cortex being (a) cut and (b) flattened. In the flattened representation, the region towards the right shows the medial cortex and the regions towards the left shows the lateral occipital cortex. Up and down are dorsal and ventral respectively. Such flattened renderings of the cortex are typically used to make it easier to see the representation along extended regions of the cortical brain (Wandell et al. 2000).



*Figure 2.5 (a) The area surrounding the occipital sulcus of the left hemisphere, shown in colour, is extracted and flattened (b). The corresponding anatomical directions and the location of calcarine sulcus are also indicated.*

For the purpose of the experiments in this thesis the aforementioned procedure was carried out for all participants. This way of visualising the cortex helps significantly appreciate the relative locations of activated voxels and delineate the early visual areas from each other.

### **2.1.2 Comparison with other packages**

The main advantage of SurfRelax over other software packages, both commercial and publicly available (Dale et al. 1999; Drury et al. 1996; Kriegeskorte and Goebel 2001; Wandell et al. 2000) is the way it deals with topological errors. A problem faced by all methods developed for cortical surface extraction is to find a way of generating surfaces with sufficient detail to capture

the fine structure of the cortex and at the same time to preserve the topology of the cortical surface. Some packages (e.g. FreeSurfer) implement a manual removal of topological errors, an approach that is laborious at best. For example, the FreeSurfer manual devotes around 50 pages describing how structural defects can be edited in the extracted surfaces. Automated handle removal algorithms can distinguish easily between handles and holes, i.e. topological errors due to the erroneous inclusion of voxels (which are corrected by removing the problematic voxels) or the erroneous exclusion of voxels (which is corrected by adding voxels). Such an automated approach that corrects for structural defects is implemented in the publicly available mrUnfold software (Wandell et al. 2000). A similar approach is used by the commercially distributed Brainvoyager package (Kriegeskorte and Goebel 2001). Kriegeskorte and Goebel (2001) deal with the issue of topological errors removal by implementing a heuristic approach that chooses the operation that requires the smallest number of voxel additions or subtractions. It is not obvious, however, that this strategy always yields anatomically correct results (see Larsson 2001 for a more detailed discussion). A different strategy for dealing with topological errors is to use an elastic deformable surface based on a template with known topology. The initial surface may be a sphere or a plane, which is subsequently deformed to fit the cortical surface, guided by local image intensities and other parameters that maintain smoothness and prevent self-intersections (Dale and Sereno 1993; MacDonald et al. 2000). A major problem with deformable surfaces is that they are prone to self-intersect, in which case collision detection algorithms are employed. Because the surfaces are typically approximated by a large number of polygons (usually triangles), collision detection can be very time consuming. For instance, the method of MacDonald et al. (2000) requires 30 hours per brain, most of which is dedicated to check for self-intersections in the evolving surface.

The method of cortical surface extraction developed by Larsson (2001) implements a combined approach. The essence of the method lies in the use of a deformable template that is voxel-based, rather than surface based. The template is deformed voxel by voxel to the white matter surface in a series of steps designed to make the deformation process able to capture fine details. The surface of the deformed template is then extracted and refined with respect to the input image. The use of a deformable volume has two principal advantages. One is that the template can capture the precise shape of the input image more easily than a deformable surface. The other is that the use of a voxel-based template obviates the need for explicit collision detection, which speeds up processing time considerably.

## 2.2 Retinotopic mapping

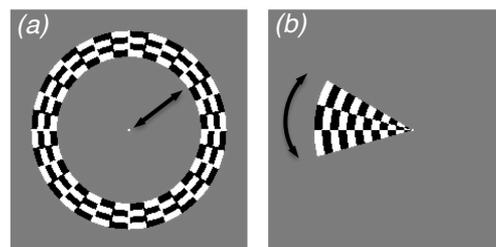
As discussed in the first chapter, the retinal signal is projected to the cortex in an organised manner. This means that adjacent points in the retina project to adjacent points in the cortex. Neurons with receptive fields in the central visual field are located in the posterior calcarine sulcus, while neurons with receptive fields in the periphery are located in the anterior portions of the sulcus. This organised registration in the cortex in terms of retinal representation is known as *retinotopic mapping*. Engel et al. (1997) reported that retinotopically organised responses extend along a 3-4 cm strip in the human cortex.

### 2.2.1 Phase-encoding techniques

Retinotopic organisation in the cortex makes it possible to delineate the borders of early visual areas. What makes it possible to distinguish retinotopic areas from each other is that the map reverses at the borders of these regions. The method for generating high-resolution retinotopic maps was described first by Engel (Engel et al. 1997; Engel et al. 1994) and later by others (DeYoe et al. 1996; Sereno et al. 1995). The mapping technique is based on checkerboard-like stimuli that change position over time and create a travelling wave of neuronal activity within retinotopically-organised areas. Such stimuli are shown in Figure 2.6. They consist of (a) an expanding and contracting ring, used to measure retinotopic organisation with respect to visual eccentricity, and (b) a rotating wedge, used to measure retinotopic organisation with respect to polar angle. The specific spatial and temporal pattern of the stimulus is not important for generating maps in early visual areas as many choices can produce satisfactory results (Dougherty et al. 2003).

The rationale of the mapping technique is based on the fact that neurons have a limited receptive field. When either the ring or wedge is seen through a

small aperture such as the receptive field of a neuron, as the stimulus progresses the image seen by the receptive field alternates between the checkerboard and the background grey field. This produces a wave of neural responses, as the stimulus enters then leaves each neuron's receptive field. If the stimulus moves with a constant velocity from fovea to periphery, the responses differ only in their phase. For example, in the case of the expanding ring, neurons with receptive fields in the fovea are phase advanced compared to neurons with receptive fields in the periphery. Hence, the phase of the temporal modulation defines the receptive field position along the dimension of eccentricity. For this reason, retinotopic mapping methods in fMRI are often called 'phase-encoded' (DeYoe et al. 1996; Engel et al. 1997; Engel et al. 1994; Sereno et al. 1995).

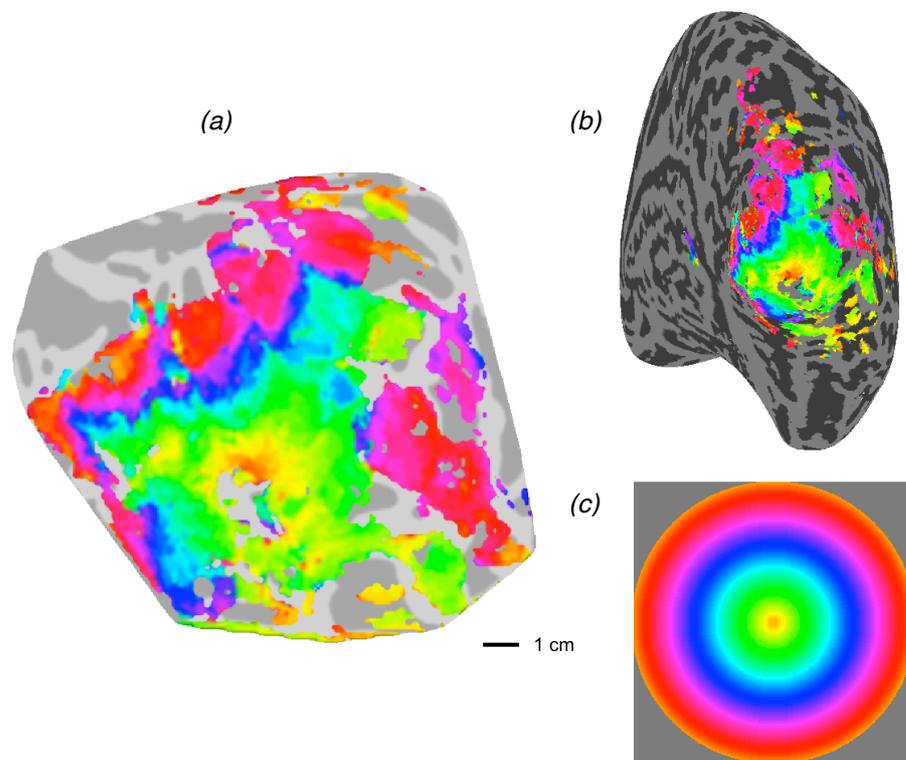


*Figure 2.6 The stimuli used to measure retinotopic organisation with respect to (a) eccentricity and (b) polar angle*

### ***Eccentricity organisation***

Measurements of eccentricity organisation using an expanding/contracting ring stimulus produce a single, large, continuous eccentricity map. Figure 2.7 shows such a map. Data are overlaid on (a) a flattened and (b) an inflated rendering of the cortex. Colour represents the phase of visual eccentricity encoded at each position according to the colour map (panel c). The way the retinal imagery is organised in the cortex with respect to eccentricity is as follows. As one moves from posterior to anterior in cortex, the representation of the visual field shifts from the centre to the periphery. This is best visualised in the inflated

depiction (panel b). The foveal representation is visible in the ventral–lateral posterior cortex, near the occipital pole; peripheral field stimuli are represented at more anterior positions forming a semicircular pattern.



*Figure 2.7 Retinotopy with respect to eccentricity on (a) a flattened and (b) an inflated representation of the right cortex, of one subject. (c) Colour represents the visual eccentricity encoded at each position, according to the colour map.*

### **Angular organisation**

Measurements of polar angle organisation using a rotating wedge stimulus are shown in Figure 2.8b. A flattened representation without the superimposed map is shown for comparison (panel a). Colour corresponds to the phase of visual angle encoded at each position according to the colour map (panel c). A plot of the response amplitude as a function of the stimulus rotation frequency is also shown (panel d). This is measured within a disk of approximately 2cm radius of cortex, originated at the point indicated by the arrow.

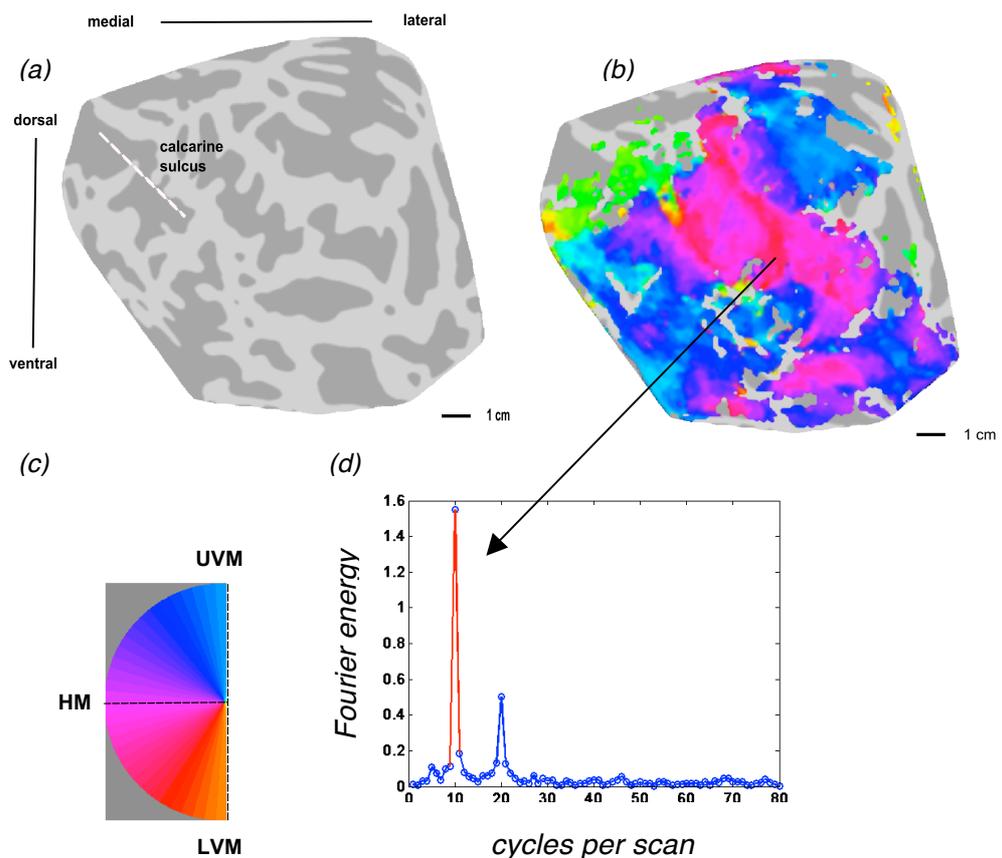


Figure 2.8 (b) Retinotopy with respect to polar angle on a flattened representation of the right visual cortex of one subject. (a) The flattened patch without the map superimposed is shown for comparison. (c) Colour corresponds to the visual angle encoded at each position according to the colour map. HM is horizontal meridian, UVM and LVM upper and lower vertical meridian respectively. (d) The response amplitude plotted as a function of the stimulus rotation frequency. This is measured within a disk of approximately 2 cm radius, originated at the point indicated by the arrow, and is presented as an assessment of the reliability of the responses (see section 1.2.2).

The response is significantly greater at the stimulus repetition frequency (10 cycles/scan, shown in red) than other temporal frequencies. The secondary peak at 20 cycles/scan is a harmonic of the fundamental frequency. These harmonics often result in large amplitudes due to the square profile of the stimulus time course. The measure of statistical significance used is a coherence metric computed as the energy of the fundamental frequency divided by the total energy (see section 1.2.2).

As with eccentricity, the angular representation of each visual field is inverted in the cortex. As one moves from the lower (ventral) to the upper (dorsal) bank of the calcarine, the representation of the visual field shifts from the upper vertical meridian (UVM) through the horizontal meridian (HM) to the lower vertical meridian (LVM).

### **2.2.2 Identification of early visual areas**

By combining the information of eccentricity and angular maps it is possible to delineate the borders of early visual areas. Primary visual cortex V1 contains a representation of an entire hemi-field, centred on the calcarine sulcus. As one moves from the middle of V1 to the V1/V2 border, the receptive field locations change from the horizontal (HM) to the vertical meridian. As one crosses the V1 border and continues into V2, the receptive field locations move from the vertical meridian back towards the horizontal meridian. The lower quadrant representation on the dorsal side of the map is V2d (dorsal). The representation of the upper quadrant on the ventral side of the map is V2v (ventral). The orderly progression of the map reverses direction again at the representation of the horizontal meridian, found at the V2/V3 border. The representation of the horizontal meridian on the dorsal side forms the boundary between V2d and V3d, and on the ventral side between V2v and V3v. Similarly, as one continues into V3, the receptive field locations move from the horizontal meridian back towards the vertical meridian.

The region enclosed by the parallel, mirror-symmetric bands of areas V1, V2 and V3 is known as the area of foveal confluence. The confluence corresponds to the central 0-2 degrees of visual field. Within this region the foveal representations of areas V1, V2 and V3 converge and with conventional fMRI resolutions it is not possible to be segmented (Dougherty et al. 2003). The foveal confluence is quite large, usually spanning approximately 2100mm<sup>2</sup>, and spreads far onto the lateral surface of the occipital cortex (Wandell and Wade 2003).

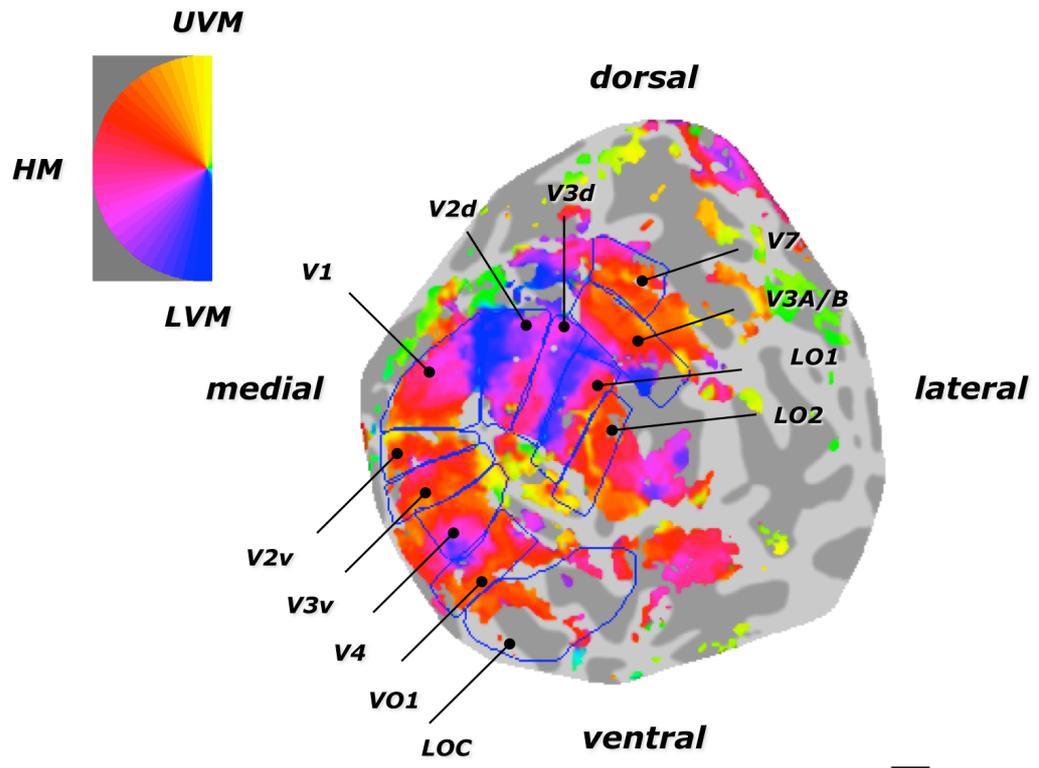
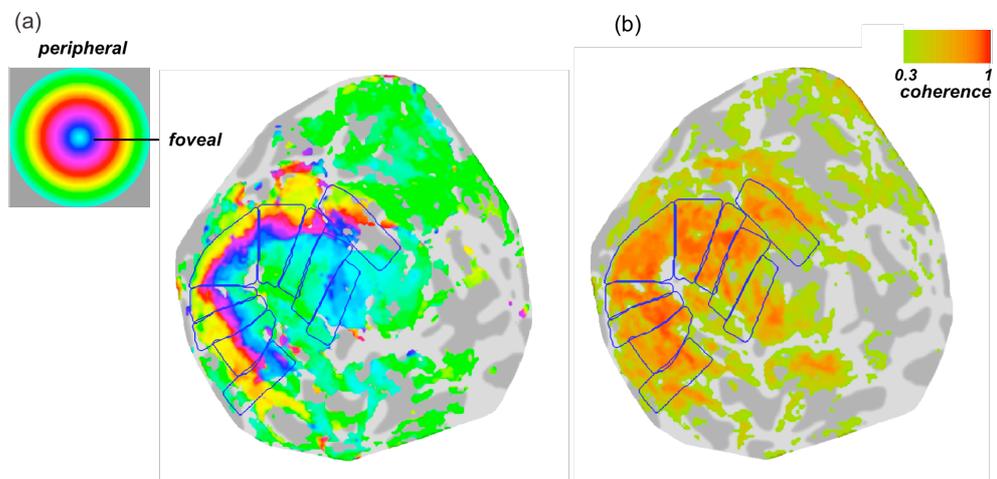


Figure 2.9 Retinotopy with respect to polar angle overlaid on a flattened patch of the right occipital cortex of one subject. Colour map codes the phase at each angular position of left visual hemi-field as indicated on the upper left. Reversals in the change of polar angle were used to identify all visual areas but LOC. LOC was mapped based on its functional properties.

Area V4 directly abuts the ventral portion of V3 and shares a common eccentricity representation with V1/V2/V3. V4's angular map spans a whole visual hemi field. The V3v/V4 boundary is formed by the representation of the upper vertical meridian (UVM) and as one continues into V4 the receptive field locations move from the upper vertical meridian towards the lower vertical meridian (LVM). The exact definition of area V4 in humans is part of an ongoing debate. Some authors label this area as hV4 to clarify that this map may not be homologous to V4 in other species. There is also controversy about what constitutes the anterior border of V4. The question is whether V4 ends prior to the lower-field representation, leaving area V4v without its corresponding V4d (Sereno et al. 1995), or whether the lower visual field representation lateral to V4v constitutes V4d, thus, forming a

complete hemi field representation, complementary to the dorsal V3a representation (McKeefry and Zeki 1997; Wade et al. 2002). In contrast, Tootell's lab suggested that on the ventral aspect there is a quarter field representation, V4v, and that the neighboring cortex belongs to a separate area termed V8, consisting of a hemi field representation (Hadjikhani et al. 1998), which is rotated relative to V4v. For the purposes of this thesis, we adopted the definition of McKeefry and Zeki (1997) and Wade et al. (2002) (for a comparison of definitions see Wade et al. 2002).



*Figure 2.10 (a) Visual eccentricity representation overlaid on a patch of flattened cortex from the right hemisphere of the same subject as in Figure 2.9. (b) Map of polar angle representation, indicating the coherence values of the BOLD responses in each voxel. Values in both panels are thresholded at coherence > 0.3.*

Areas V3A and V3B are located on the dorsal side of V3 and can be defined from each other by a common foveal representation (Press et al. 2001; Wandell et al. 2005). V3A and V3B each contain a full hemi field map such that their dorsal/anterior boundary (with V7) represents the upper vertical meridian (UVM). However, the clarity of the common foveal representation can vary significantly between subjects. As the two areas cannot be distinguished with confidence in every subject it is referred to in this thesis as a composite region, labelled as

V3AB, in keeping with previous reports (Larsson and Heeger 2006; Larsson et al. 2006; Montaser-Kouhsari et al. 2007).

In the dorsal stream, beyond V3AB, an additional retinotopic area has been identified. Area V7 contains an additional hemi field representation anterior to V3A (Press et al. 2001). V7 meridians are rotated compared to V3A meridians. In particular, the V3A/V7 boundary is formed by the representation of the upper vertical meridian (UVM) and as one continues into V7 the receptive field locations move from the upper vertical meridian towards the lower vertical meridian (LVM).

VO1 is a coarsely retinotopic area located in the ventral occipital (VO) cortex anterior and lateral to V4 (Wandell et al. 2005). The posterior portion of VO1 is adjacent to the peripheral visual field representation of hV4. The V4/VO1 boundary is formed by the representation of the lower vertical meridian (LVM) and as one moves ventrally, the representation shifts towards the upper vertical meridian (UVM).

LO1 and LO2 are the two retinotopically-organised regions in the lateral occipital (LO) cortex, lateral to the dorsal portion of V3 (Larsson and Heeger 2006). LO1 extends from the anterior boundary of V3d about halfway to V5/MT+<sup>1</sup>. Within LO1 the representation of visual polar angle progresses gradually from the lower vertical meridian (LVM) towards the upper vertical meridian (UVM). The anterior boundary of LO1 coincides with the posterior boundary of LO2. LO2 also contains an orderly representation of visual angle polar angle extending from the upper vertical meridian (UVM) towards the lower vertical meridian (LVM), which defines the lateral boundary of LO2.

---

<sup>1</sup> For the purposes of this thesis we did not perform an individual localiser scan to identify V5/MT+.

### 2.2.3 Non-retinotopic visual areas

The lateral occipital complex (LOC) is not retinotopic, but can be reliably localised with other functional techniques. LOC is described in the literature as an object selective region (Eger et al. 2008; Grill-Spector et al. 2000; Kourtzi and Kanwisher 2000; 2001; Kourtzi et al. 2003; Malach et al. 1995; Vinberg and Grill-Spector 2008). LOC can be identified as the set of contiguous voxels activated more strongly by images of intact objects than by scrambled versions of the same images. LOC localiser stimuli were presented in sixteen 16-s blocks (8 s 'on', 8 s 'off'). Sixteen images of cars, animals, boats, and abstract sculptures were shown randomly for 0.5s within each block (128 images in total, 32 of each category). In half of the blocks stimuli were intact images and in the remaining half they were chopped and scrambled versions of the same images (stimuli were courtesy of Dr. Kalanit Grill-Spector). Through the experiment subjects were instructed to covertly name the stimuli while fixating.

## 2.3 FMRI scanning

The following section refers to general experimental settings that were implemented during this thesis. Further experimental details are described within the relevant chapters.

### 2.3.1 Participants

All volunteers took part in the studies with written consent. Procedures were approved by the Medical School Research Ethics Committee of the University of Nottingham.

### 2.3.2 Visual stimuli

Stimuli were generated using the open-source package PsychoPy (Peirce 2007) and were back-projected from an LCD projector at a resolution of 1024x768 pixels to a screen sited at the feet of the subject. To control for non-linearities in the luminance profile of the display, the screen was gamma-corrected using a psychophysical procedure of 2<sup>nd</sup>-order motion-nulling (Ledgeway and Smith 1994). Subjects viewed the screen through prism goggles.

### ***Probe contrast in event-related paradigms***

Event-related adaptation paradigms (Fang et al. 2005; Larsson et al. 2006; Montaser-Kouhsari et al. 2007) rely on the observation that after prolonged presentation of a particular high-contrast adaptor, the fMRI response to a similar probe stimulus is selectively reduced compared to the response to other probe stimuli. However, the exact contrast value at which probe stimuli are presented varies between studies.

From electrophysiological studies in LGN (Solomon et al. 2004) and primary visual cortex of cat (e.g. Ohzawa et al. 1982; 1985) and macaque (Sclar et al.

1989) we know that adaptation causes a strong right-ward shift in the contrast response curve. Due to the saturating nature of this curve, the greatest difference in response between adapted and non-adapted conditions occurs for lower contrast probes (Maffei et al. 1973). Similarly, in psychophysical studies it has been shown that, although at detection threshold there is a highly selective adaptation to the spatial frequency of probe- versus adapter-stimuli (Blakemore and Campbell 1969), for higher contrast probes the tuning of adaptation is considerably broader (Snowden and Hammett 1996) and there is less impact on the apparent contrast of the probes following adaptation (Georgeson 1985). The use of low contrast probes must, of course, be traded off with the need to generate robust BOLD responses in the ROIs - the ideal stimulus is the lowest contrast for which a robust response can be measured. For the above reasons we have chosen to present probes in our event-related paradigms at lower than the adaptor contrasts.

### 2.3.3 Functional imaging

Blood-oxygen level-dependent (BOLD) cortical responses were recorded using gradient-echo (GE) echo-planar imaging (EPI) at 3T (Philips Achieva System, Philips Healthcare, Best, the Netherlands). The parameters for scanning were; voxel size  $3 \times 3 \times 3 \text{mm}^3$ , TR 1.5s, TE 40ms, flip angle  $75^\circ$ , FOV  $192 \times 192 \text{mm}^2$ , 20 slices oriented perpendicular to the calcarine sulcus. To improve signal-to-noise, functional data were acquired using a pair of surface receiver coils (Philips Flex-S Coils) positioned over occipital cortex.

At the beginning of each session, an anatomical image was obtained that covered the same volume as the functional images (T1-weighted MPRAGE, voxel size  $1.5 \times 1.5 \times 3 \text{mm}^3$ ). This 'co-planar' anatomy image was used as a proxy to register functional data to a high-resolution, whole-head anatomical image obtained in a separate session (T1-weighted 3D MPRAGE, voxel size  $1 \times 1 \times 1 \text{mm}^3$ ,

8-channel SENSE head coil) using a robust registration technique (Nestares and Heeger 2000). The high-resolution anatomical images were segmented and flattened using standard tools (SurfRelax, Larsson 2001).

### **2.3.4 Attention control task**

To control for changes in the attentional state of observers, which are known to modulate fMRI responses (Brefczynski and DeYoe 1999; Huk et al. 2001; Kastner et al. 1999; Somers et al. 1999), participants performed a demanding task at fixation. They were asked to count the number of target letters (X) appearing among a series of distractor letters (Z, L, N, T), which changed every 200ms. The duration of each letter-counting trial varied randomly between 7-14s. At the end of a sequence of letters, a fixation spot appeared for 1s prompting participants to report the number of target letters presented (1-4) by pressing one of four response buttons.

### **2.3.5 Data analysis**

Functional images were motion-corrected within and between scans using MCFLIRT (Jenkinson et al. 2002). In the presence of a localiser scan, the following analyses were performed: first, the regions of interest (ROIs) were restricted to include only voxels whose time series correlated with the stimulus epochs of the localiser scan. Specifically, ROIs were restricted only to those voxels with a coherence value greater than a certain cut-off (usually  $co > 0.3$ ). The exact coherence threshold did not appear to be critical as other thresholds (0.2, 0.3, and 0.4) resulted in similar effects.

The time series of each voxel in the restricted ROIs were then pre-processed as follows: the mean time series was subtracted and divided to convert data from arbitrary image intensity to units of percent signal change. Responses were then filtered using a high-pass boxcar kernel to remove the low-frequency

drift typical in fMRI measurements (Biswal et al. 1997a; Biswal et al. 1995; Biswal et al. 1997b; Purdon and Weisskoff 1998; Smith et al. 1999; Zarahn et al. 1997).

In event-related paradigms, the fMRI responses were averaged across all voxels within the restricted ROI. Additional band-pass filtering was then applied to the averaged time courses to remove high-frequency noise and the remaining low-frequency drift (cut-off frequencies 0.015 and 0.15Hz). The choice of the particular parameters ensured that no artefacts, such as drift or noise, were included in the time courses. At the same time it was ensured that no useful information was filtered out of the data.

### ***Adaptation***

In the cases we wanted to evaluate the amount of adaptation in different ROIs we computed a metric by fitting event-related responses with a difference of two gamma functions (Equation 1) (Glover 1999; Jezzard 2001). The amount of adaptation in each ROI was computed as the difference in the maximum values of the fitted curves (for 'same' and 'different' conditions) normalized by their sum.

$$H(t) = \left(\frac{t}{d_1}\right)^{a_1} \exp\left(\frac{-(t-d_1)}{b_1}\right) - c \left(\frac{t}{d_2}\right)^{a_2} \exp\left(\frac{-(t-d_2)}{b_2}\right) \quad (\text{Equation 1})$$

where  $d_i = a_i b_i$  defines the time-to-peak. The initial parameters for non-linear regression were,  $a_1=5.15$ ,  $a_2=12.26$ ,  $b_1=0.97\text{s}$ ,  $b_2=0.94\text{s}$ ,  $c=0.09$ .

Furthermore, following Larsson et al. (2006), we assessed the statistical reliability of adaptation by computing the response amplitude of each trial. For this analysis, we first computed a mean response vector  $\bar{R}$  by averaging the responses for all trials regardless of the trial type,

$$\bar{R} = \frac{1}{N} \sum_{i=1}^N R_i$$

where  $N$  is the number of trials and  $R_i$  the individual trials after subtracting the response to the blank probe. Then, for each trial, we computed a scalar response amplitude  $A_i$  as,

$$A_i = \frac{R_i \cdot \bar{R}}{\|\bar{R}\|}.$$

As in previous adaptation studies (e.g. Larsson et al. 2006) we estimated statistical reliability for individual subjects using a one-tailed t-test. A significant result would indicate the response amplitudes  $A_i$  to the probe that had a ‘different’ orientation to the adaptor were significantly greater than responses to probes that had the ‘same’ orientation as the adaptor.

### ***Pattern classification***

Classification performance depends on the number and choice of the voxels included in the analysis (Cox and Savoy 2003; Ku et al. 2008). For each voxel in our ROIs, we determined the stimulus-driven response in the localiser scan, computed as a t-statistic. Following Haynes & Rees (2005), we selected an unbiased sample of 100 voxels with the highest t-values (stimulus versus blank) for further analysis. To quantify the dependence of classification performance on the number of voxels used, we calculated the MVPA accuracy score on 100 permutations. In each permutation the order of the voxels included in the analysis was shuffled and the MVPA analysis was performed as described below. The mean and standard deviation of these 100 reshuffles was then computed.

The responses of the  $n=100$  voxels at  $k$  time-points were sorted in an  $n$ -by- $k$  matrix. Each column of this matrix corresponds to a feature vector  $\mathbf{x}$ , which, prior to classification, was normalised to unit Euclidean length.

**Linear discriminant analysis.** In some chapters a linear discriminant analysis algorithm (Duda et al. 2001) was used. We assumed that patterns of responses recorded under different conditions fall in multidimensional, normally distributed clusters with equal covariances. We computed the pooled covariance matrix as  $\Sigma = (\Sigma_1 + \Sigma_2)/2$ , where  $\Sigma_1$  and  $\Sigma_2$  are the individual covariance matrices, describing the spread of each cluster. As both conditions had equal prior probabilities, a minimum-error-rate classification can be achieved by use of the linear discriminant functions:

$$g_i(\mathbf{x}) = \Sigma^{-1} \mu_i \mathbf{x} - \frac{1}{2} \mu_i' \Sigma^{-1} \mu_i$$

where  $\mu_1$  and  $\mu_2$  are the means of the two clusters. As the discriminant functions are linear, the resulting decision boundary in this two-category case is a hyper-plane lying halfway between the means of the clusters.

**Neuronal network classification.** In other chapters, a neuronal network (without a hidden layer) algorithm was used to classify responses between stimulus conditions. This algorithm uses the gradient of a performance function to train a neural network, that is, to adjust the weights of the network so that a desired output is produced for a given input. The gradient is determined using a MATLAB® (©1994-2009 The MathWorks, Inc) implementation of the backpropagation technique. The classifier updates the network weights and biases in the direction in which the performance function decreases most rapidly, that is, the negative of the gradient. One iteration of this algorithm can be written as:  $\mathbf{x}_{k+1} = \mathbf{x}_k - a_k \mathbf{g}_k$ , where  $\mathbf{x}_k$  is a vector of current weights and biases,  $\mathbf{g}_k$  is the current gradient, and  $a_k$  is the learning rate. This procedure, known as conjugate gradient, is the simplest implementation of the backpropagation technique and has the advantages of good generalising and fast learning (Detre et al. 2006).

In both cases, trials were divided into  $\lambda$  groups, each group corresponding to responses collected in a pair of blocks. Data from  $\lambda-1$  of these groups were assigned to a training set and the remaining to a test set. During the training stage the classifier learned to discriminate between responses recorded under the two orientations and to define a decision boundary. Responses from the test sample were then used to assess the performance of the classification algorithm and compute the error. Classification error was evaluated using a cross-validation procedure (Duda et al., 2001) computed as the mean across  $\lambda$  leave-one-out permutations. An outline of the ‘pattern assembly’ procedure is shown in Figure 2.11.

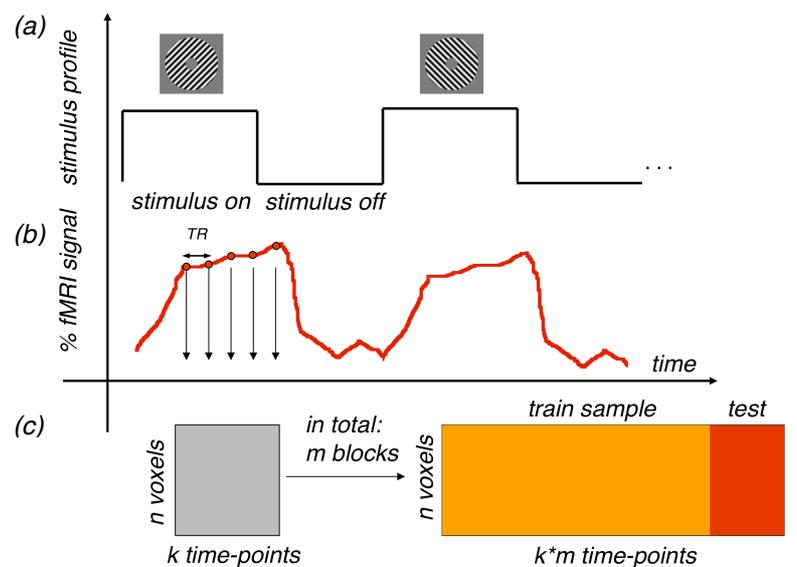


Figure 2.11 (a) The hypothetical experimental design. Two stimuli alternating in on/off blocks. (b) The hypothetical fMRI response to the stimuli. The responses in  $k$  time-points are selected from each block, delayed in time to account for the sluggish haemodynamic signal. (c) The input to the pattern-classifier. Each scan scan consists of  $m$  blocks, resulting in a total of  $k*m$  time-points. Data are divided into training and testing samples and are cross-validated.

To compute a single metric describing the classification performance in each ROI, we computed the mean accuracy at all points between the 50<sup>th</sup> and

100<sup>th</sup> voxels. This typically captures the asymptotic performance and provides a reasonable aggregate for non-asymptotic cases.

**Permutation test for classification accuracy.** To assess the statistical reliability of the multi-variate classification performance, we performed a permutation test (Efron and Tibshirani 1993). To simulate the distribution of expected classification accuracy scores under the null-hypothesis, we calculated the MVPA accuracy score on 5000 resamples. Each resample was generated by shuffling the indices assigning the responses to the two different orientation conditions, and performing the analysis exactly as described above. From the distribution of the classification accuracy values resulting from these resampled analyses, we obtained the 95% confidence interval for chance performance.

### ***3 Do fMRI adaptation and pattern-classification analysis measure the same thing?***

---

Functional magnetic resonance imaging (fMRI) has proven extremely useful in the non-invasive study of human brain function. Measurements of the blood oxygenation level dependent (BOLD) signal have been used to track local increases in neural activity in a large number of studies. These include studies investigating aspects of perception, cognition, and memory. One of the limitations of fMRI, in comparison with the direct recording of neuronal responses via microelectrode, is its spatial resolution. Improvements to imaging hardware and analysis techniques have provided access to higher-resolution images at improved signal-to-noise ratios (e.g. Logothetis 2008; Moon et al. 2007; Yacoub et al. 2007). However, BOLD measurements are ultimately limited in spatial resolution by, amongst other things, the spatial scale of the local vascular system.

In many studies, the aim is to quantify the selectivity of clusters of neurons on a spatial scale much smaller than the  $3 \times 3 \times 3 \text{mm}^3$  volume of a voxel used typically in current fMRI experiments. Orientation-selective cells in V1 of the primate, for example, are clustered into 'columns' of roughly  $500 \mu\text{m}$  in diameter (Bartfeld and Grinvald 1992; Obermayer and Blasdel 1993). Ocular dominance columns in the human primary visual cortex have a mean width of  $863 \mu\text{m}$  (Adams et al. 2007). In order to resolve differences in orientation tuning between voxels 'traditional' fMRI methods would require voxel dimensions considerably smaller than that of the column width

Recent fMRI studies have demonstrated new methods for studying the selectivity of neurons in various domains (such as orientation) without requiring that the voxel size be smaller than the resolution of the 'feature map'. These have

used either selective adaptation (Grill-Spector and Malach 2001; Krekelberg et al. 2006) or multi-variate, pattern-classification analysis (MVPA) (Cox and Savoy 2003; Haynes and Rees 2006; Norman et al. 2006).

The use of adaptation has a long history in the psychophysical study of visual processing (Blakemore and Campbell 1969; Bradley et al. 1988; Snowden and Hammett 1996). Its use has been so prominent that it has been referred to by some as the 'psychophysicists electrode'. There is, however, still a debate about the exact mechanism underlying these perceptual effects (see e.g. Desimone 1996; Grill-Spector et al. 2006).

More recently, selective adaptation effects have been demonstrated with fMRI. Adaptation can be selective for stimulus orientation (Engel 2005; Fang et al. 2005; Larsson et al. 2006), direction of motion (Huk and Heeger 2002; Krekelberg et al. 2005; Tolias et al. 2001), various higher order properties of objects (Grill-Spector et al. 1999; Kourtzi and Huberle 2005; Kourtzi et al. 2003; Sayres and Grill-Spector 2006; Vuilleumier et al. 2002) and faces (Andrews and Ewbank 2004; Grill-Spector and Malach 2001; Henson et al. 2002). In general, the method relies on the observation that after prolonged or repeated presentation of a particular stimulus, the BOLD response in areas sensitive to that stimulus, is selectively reduced compared to the equivalent response to other stimuli. The methodological details of the above studies vary enormously. It is possible that the mechanisms underlying the observed reduction in BOLD signal may differ between studies and may not reflect the changes measured in psychophysics or single-unit physiology experiments. In the example of orientation-selectivity, after prolonged viewing of a high-contrast grating of a particular orientation, the fMRI response to a probe of the same orientation is reduced relative to that for a differently oriented probe. The fact that adaptation is not uniform across different orientations is thought to reflect tuning in the underlying neural mechanisms.

It should be noted that there is some debate about the degree of selectivity demonstrated by the selective adaptation method in early visual areas. Boynton and Finney (2003) found no selective adaptation in V1. They suggest that this may have been caused by a) the responses of untuned neurons in V1 and V2, b) the fact neurons in these areas do not adapt or, c) the fact that a low spatial frequency was used for the stimulus ( $0.25c/\theta$ , which would result in only a fraction of a single cycle being presented to most V1 receptive fields). Fang et al. (2005) attribute Boynton and Finney's data to the timing of their stimulus; they found that using a prolonged adaptation period resulted in significant orientation-selective adaptation in all areas tested, although the effect was still stronger in V3 and V4. Larsson et al. (2005) use a lower contrast probe stimulus in testing orientation selectivity and find no significant difference between the visual areas in adaptation index. The choice of probe contrast may well contribute to the previous weak selective adaptation in V1 found by Boynton and Finney (2003) and Fang et al. (2005). In this study we have followed Larsson et al. (2005) in using probes of 10% Michelson contrast for reasons discussed in section 2.3.2.

Multi-variate pattern analysis (MVPA) methods, instead, make use of small differences in the fMRI response of different voxels thought to result from small biases in the spatial distribution of the neural subpopulations sampled by each voxel. By 'learning' the pattern of these small biases across a large number of voxels in an independent training set, multi-variate pattern analysis can successfully discriminate between stimuli in a novel set of trials. Several reports have shown that such multi-variate techniques can reliably distinguish between responses to different stimuli, where more conventional, voxel-wise univariate approaches, or signal averaging across whole regions of interest could not. MVPA techniques have been used to decode the orientation of gratings (Haynes and Rees 2005; Kamitani and Tong 2005), direction of motion (Kamitani and Tong 2006), object categories (Eger et al. 2008; Haushofer et al. 2008; Haxby et al.

2001), to study visual categorisation (Li et al. 2007) and also the encoding of global form (Ostwald et al. 2008).

It has already been demonstrated in separate studies, that fMRI adaptation and multi-variate techniques are capable of revealing orientation-selective responses in early visual areas. The aim of this study was to compare whether the results from the two methods are in agreement on their measurement of orientation tuning in early visual cortex. The optimal procedures for the two paradigms differ; notably, the MVPA method benefits from data acquired in a blocked design, whereas an event-related design is optimal for adaptation methods. Here we compare the two methods, each with optimal designs, for data acquired in equal periods of time. Two questions were used to frame this comparison. First, do areas that show strong orientation-specific adaptation also show high classification performance? In order to test this we compared, for a number of visual areas, the pattern classification accuracy and selectivity of adaptation from interleaved scans in a single session. Second, we wanted to know which method was more sensitive in detecting subtle orientation differences of stimuli. To measure this we reduced, in successive scanning sessions, the orientation difference between the two gratings in both adaptation and MVPA scans.

#### **3.1.1 Methods**

##### ***Participants, stimuli and procedure***

Three experienced volunteers took part in the study. Subjects participated in five scanning sessions; one session to acquire high-resolution anatomical images, one session to measure retinotopic organisation in the visual cortex, and three sessions to measure responses to gratings differing in orientation by  $90^\circ$  ( $\pm 45^\circ$ ),  $50^\circ$  ( $\pm 25^\circ$ ) and  $25^\circ$  ( $\pm 12.5^\circ$ ).

**Visual stimuli.** Stimuli consisted of oriented sinusoidal gratings (spatial frequency 1.5cycles/°). The gratings were presented in an annulus (inner radius, 2°; outer radius, 8°) whose edges were smoothed by a Gaussian kernel (std dev of 0.083° on the inner edge, 0.333° on the outer edge). The spatial phase of the gratings was randomised every 6 frames (100ms) to prevent retinal afterimages. Grating orientations were +/-45° (session 1), +/-25° (session 2) and +/-12.5° (session 3).

**Scanning session.** The degree of orientation-selective adaptation and the performance of the pattern-classification algorithm were determined from separate, interleaved scans in the same session. An event-related design was used to measure the degree of selective adaptation. An adapting high-contrast stimulus was presented for a prolonged period, followed by a brief, low contrast probe. A block design was used to measure the performance of the MVPA (Figure 3.2).

At the beginning of each functional scanning session, we ran a localiser scan. This was followed by four adaptation scans, and three MVPA scans, which were interleaved.

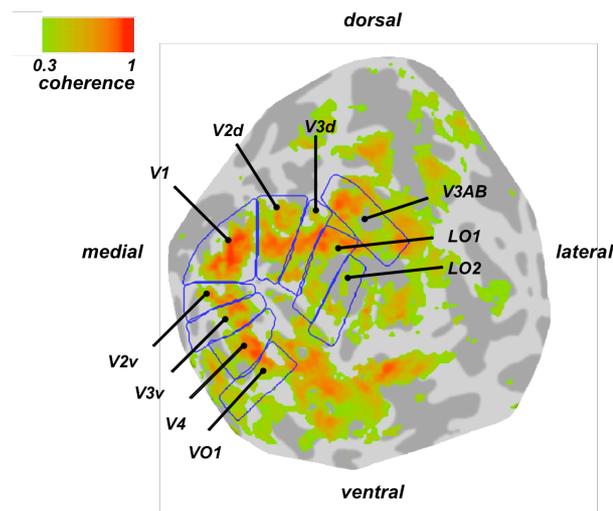
**Localiser scan.** The purpose of the localiser scan was to identify voxels in the ROIs that responded to visual stimulation at the spatial location of the patterns. Stimuli were the two oriented gratings (see above) presented at high contrast (90% Michelson), alternating at 0.5Hz for 15s followed by presentation of a blank screen. A fixation point was present throughout. Each localiser scan consisted of 8 such blocks. The responses evoked by the localiser stimuli are shown in Figure 3.1.

**Event-related adaptation scans.** The event-related fMRI adaptation protocol (Fang et al., 2005; Larsson et al., 2006; Montaser-Kouhsari et al., 2007)

### 3. Comparison of fMRI adaptation and MVPA

---

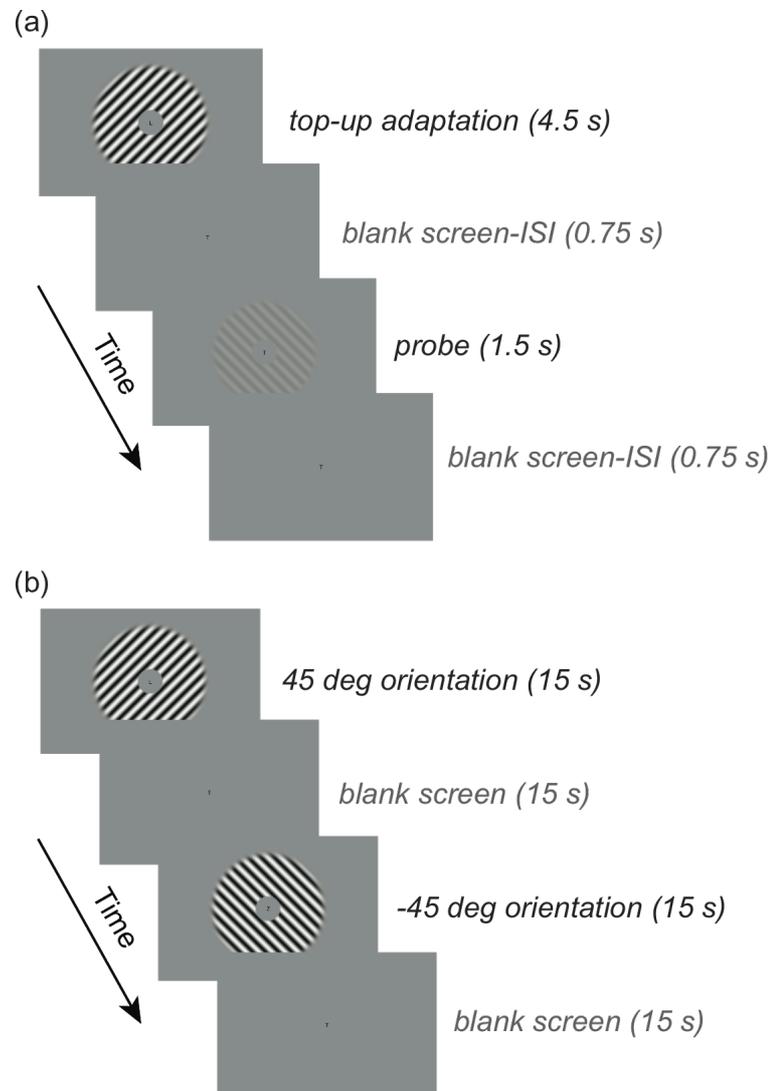
is shown in Figure 3.2a. Participants were initially adapted to a high-contrast grating (90% Michelson) at one of the two orientations for that session for 30s. In each subsequent trial, adaptation was maintained by presenting a ‘top-up’ adaptor for 4.5s. There followed a blank screen for 0.75s and the probe stimulus was then presented for 1.5s. Probes were (a) ‘same’, a 10 % Michelson grating at the orientation of the adaptor (b) ‘different’, an equivalent grating at the other orientation for that session, and (c) a blank screen (mean luminance). These conditions were equally common and randomly chosen. Each trial ended with a 0.75s presentation of a blank screen, giving a total duration of 7.5s. Each scan consisted of 30 such trials (10 in each condition). In each scanning session we ran four adaptation scans, two for each adapter orientation.



*Figure 3.1 Coherence map elicited by the localiser stimuli. Activation corresponds to coherence values greater than 0.3. Data are from the right hemisphere of one subject (SH).*

**Block-design MVPA scans.** In the MVPA scans (Figure 1b) the two oriented gratings used in the particular session (at 90% Michelson contrast) were alternated with epochs of blank screen (mean luminance) with a period 30s (15s ‘on’, 15s ‘off’). Each scan consisted of 10 blocks, 5 for each orientation.

**Attention control task.** To control for changes in the attentional state of observers a letter counting task was performed at fixation (see section 2.3.4).



*Figure 3.2 (a) An event-related design was used to measure the degree of selective adaptation, with a prolonged presentation of an adapting high-contrast stimulus, followed by a brief, low contrast probe. (b) A block design was used to measure responses for pattern classification.*

#### **Data analysis**

Regions of interest (ROIs) were restricted to include only voxels whose time series correlated with the stimulus epochs of the localiser scan (coherence,  $c > 0.3$ ,

phase  $0 < \varphi < \pi$ ). This ensured that the voxels included in adaptation and pattern-classification analysis were selected from the same overall population, but were chosen independently from either the adaptation or the MVPA measurements.

**Adaptation.** Responses to individual trials were extracted from the average ROI time course by selecting a 18s window starting 3s after the onset of the adaptor. The average response to the blank trials (which captures the response to the adaptor alone) was subtracted from each trial. Trials of each type were then averaged and the resulting event-related time courses were adjusted to zero baseline. An adaptation index was computed in each ROI as described in section 2.3.5

**Pattern-classification.** In each scanning session, we obtained data in 30 blocks (3 blocked scans, 10 blocks per scan). From each block we extracted the responses at 10 separate time-points (over 15s), delayed by three TR's to account for the haemodynamic lag. There were therefore a total of 300 time-points (or repeated 'examples' of each response) for each voxel. Trials were divided into 15 groups, each group corresponding to responses collected in a pair of blocks, 1 from each orientation of the original dataset. Data from 14 of these groups were assigned to a training set and the remaining to a test set.

### 3.1.2 Results

#### ***Comparison of results across visual areas***

We sought to examine the relationship between orientation-selective adaptation and multi-variate pattern-classification analysis (MVPA) across visual areas, by considering the responses to two gratings with a large orientation difference ( $\pm 45^\circ$ ). fMRI adaptation and MVPA were evaluated using an event-related and a blocked design respectively, which were carried out in interleaved order in a single scanning session. The probe-related modulations in fMRI signal

during the adaptation sequence are shown for one subject, and averaged across participants for eight retinotopically-defined areas (Figure 3.3a,b). The degree of selective adaptation for each area can be seen as the difference between the responses to the probe that had the 'same' orientation (shown in light grey) as the adapter, versus the 'different' orientation (dark grey). This adaptation effect is thought to reflect orientation selectivity. There was a substantial difference between the responses to the two conditions in 'lower' visual areas (V1, V2, V3 and V4). The adaptation appears less selective in 'higher' order areas (V3AB, LO1, LO2 and VO1).

Figure 3.3c,d plots the performance of the pattern classifier based on linear discriminant analysis for the same visual areas in one subject (c) and across subjects (d). Classification accuracy is plotted against the number of voxels included in the analysis. In areas V1, V2 and V3, even when classification is computed only for the single voxel, we found classification accuracy on average to be better than chance. In these areas, classification performance increased monotonically as more voxels were included in the analysis and reached an asymptote after 10-20 voxels. In higher visual regions V4, V3AB and LO1, classification accuracy was lower but significantly greater than chance (based on a permutation test, see section 2.3.5). The results of adaptation and MVPA for all subjects are summarised in Table 1.

For each method a single selectivity index was determined; the contrast between 'same' and 'different' orientations for the adaptation study, and an aggregated performance index for the MVPA (see section 2.3.5). Figure 3.4 compares these selectivity indices directly in eight retinotopic visual areas averaged across the three participants. There was a strong correlation ( $r=0.85$ ,  $n=8$ ,  $p<0.05$ ) between classification accuracy and the selective adaptation across regions. Area V4 deviated somewhat from this pattern; it had a greater selectivity to orientation, as

### 3. Comparison of fMRI adaptation and MVPA

measured by the selective-adaptation index, than would have been predicted by the MVPA performance. The V4 data-point does, however, fall within the area of 95% confidence from the regression analysis, and so its reliability is unclear.

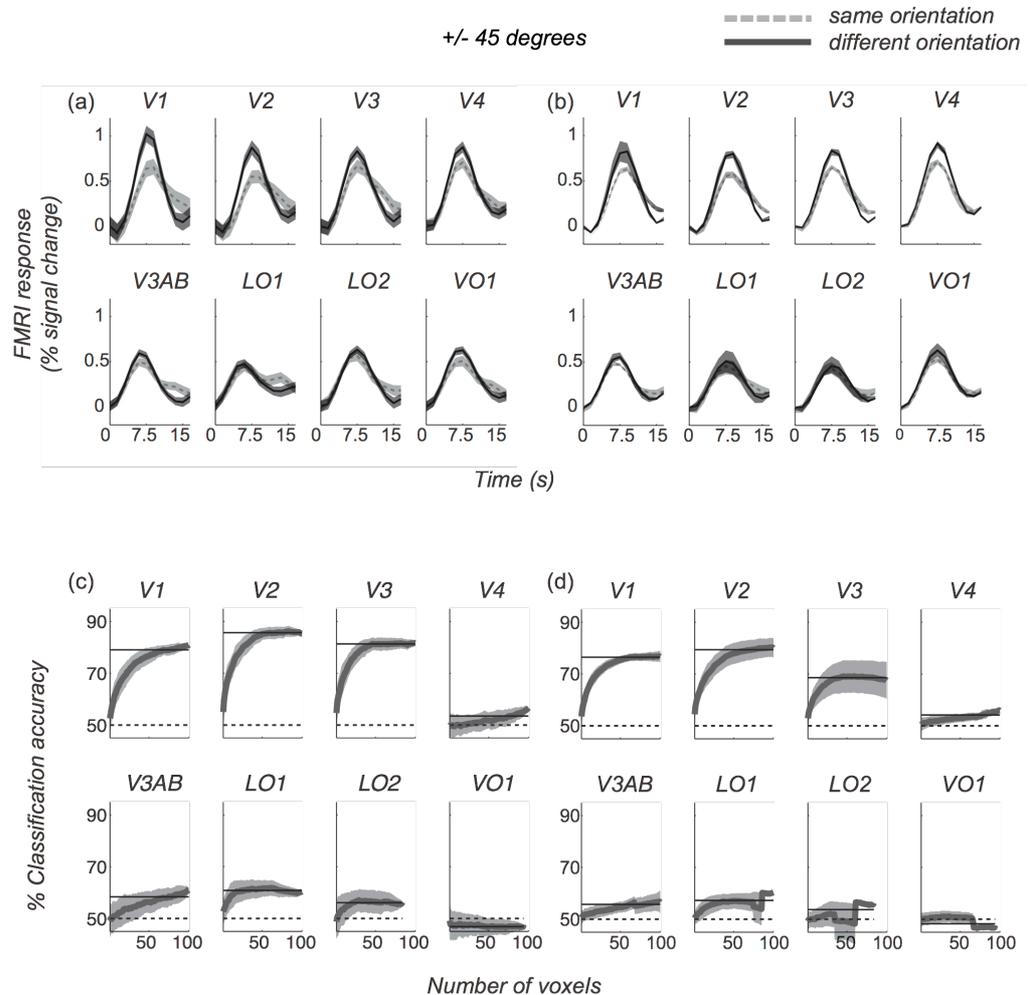


Figure 3.3 (a) Event-related modulations in fMRI signal during the adaptation sequence for one subject (JWP) and (b) averaged across subjects. The black line indicates the 'different' orientation condition; the grey line shows the 'same' orientation. Responses are averaged over 40 trials for each condition. The response to the blank condition was subtracted, to account for the response to the adapting stimulus. (c) MVPA performance versus number of voxels included in the analysis for one subject (JWP). Shaded regions are standard deviations computed over 100 reshuffles. (d) MVPA performance versus the number of voxels averaged across subjects. Shaded regions represent  $\pm 1$  s.e.m. across subjects. The dashed line shows classification accuracy based on chance (50%).

### 3. Comparison of fMRI adaptation and MVPA

The grey solid line indicates the index used to estimate classification performance (see section 2.3.5).

Table 3.1. Response amplitude differences (in units of % fMRI signal change) and pattern-classification accuracies for individual subjects by condition. P-values are shown in parantheses. Adaptation, p-values were estimated using a one-tailed, unpaired-samples t-test ( $df=37$ ). MVPA p-values are estimated from a permutation test conducted on the data for each individual (5000 resamples).

#### A. Effects for +/-45° comparison in all areas

	V1	V2	V3	V4
<b>Adaptation</b>				
JWP	<b>0.58 (0.012)</b>	<b>0.38 (0.035)</b>	0.27 (0.068)	<b>0.28 (0.039)</b>
DS	0.19 (0.051)	<b>0.23 (0.035)</b>	0.2 (0.068)	0.2 (0.132)
SH	<b>0.30 (0.041)</b>	<b>0.35 (0.013)</b>	<b>0.36 (0.020)</b>	<b>0.26 (0.024)</b>
<b>MVPA</b>				
JWP	<b>0.79 (0.000)</b>	<b>0.86 (0.000)</b>	<b>0.81 (0.000)</b>	0.53 (0.089)
DS	<b>0.75 (0.000)</b>	<b>0.80 (0.000)</b>	<b>0.61 (0.000)</b>	0.54 (0.068)
SH	<b>0.75 (0.000)</b>	<b>0.80 (0.000)</b>	<b>0.61 (0.000)</b>	<b>0.55 (0.026)</b>

	V3AB	LO1	LO2	VO1
<b>Adaptation</b>				
JWP	<b>0.22 (0.03)</b>	0.12 (0.172)	0.17 (0.093)	<b>0.23 (0.040)</b>
DS	0.002 (0.50)	0.15 (0.166)	0.01 (0.480)	-0.07 (0.420)
SH	<b>0.17 (0.017)</b>	<b>0.17 (0.023)</b>	0.07 (0.208)	<b>0.19 (0.045)</b>
<b>MVPA</b>				
JWP	<b>0.58 (0.000)</b>	<b>0.61 (0.000)</b>	<b>0.56 (0.009)</b>	0.47 (0.858)
DS	0.52 (0.145)	0.50 (0.449)	0.41 (0.998)	0.53 (0.150)
SH	0.53 (0.154)	0.59 (0.449)	0.41 (0.998)	0.53 (0.151)

#### B. Effects for +/-25° in all orientation-selective areas

	V1	V2	V3	V4
<b>Adaptation</b>				
JWP	0.40 (0.120)	0.30 (0.165)	0.32 (0.100)	<b>0.21 (0.034)</b>
DS	0.00 (0.430)	0.07 (0.240)	0.08 (0.210)	-0.06 (0.350)
SH	0.02 (0.386)	0.05 (0.260)	<b>0.24 (0.045)</b>	0.09 (0.360)
<b>MVPA</b>				
JWP	<b>0.65 (0.000)</b>	<b>0.64 (0.000)</b>	<b>0.64 (0.000)</b>	0.51 (0.317)
DS	<b>0.57 (0.003)</b>	<b>0.61 (0.000)</b>	<b>0.56 (0.011)</b>	0.45 (0.966)
SH	<b>0.69 (0.000)</b>	<b>0.71 (0.000)</b>	<b>0.63 (0.000)</b>	<b>0.56 (0.008)</b>

#### C. Effects for +/-12.5° in all orientation-selective areas

	V1	V2	V3	V4
<b>Adaptation</b>				
JWP	-0.09 (0.400)	0.02 (0.480)	0.12 (0.26)	0.10 (0.300)
DS	0.11 (0.720)	0.01 (0.380)	0.06 (0.150)	0.03 (0.400)
SH	-0.43 (0.020)	-0.21 (0.060)	-0.23 (0.010)	-0.13 (0.790)
<b>MVPA</b>				
JWP	<b>0.60 (0.000)</b>	<b>0.63 (0.000)</b>	<b>0.63 (0.000)</b>	<b>0.58 (0.002)</b>
DS	<b>0.57 (0.004)</b>	0.54 (0.056)	<b>0.57 (0.003)</b>	0.52 (0.224)
SH	<b>0.56 (0.004)</b>	0.54 (0.054)	<b>0.57 (0.002)</b>	0.52 (0.223)

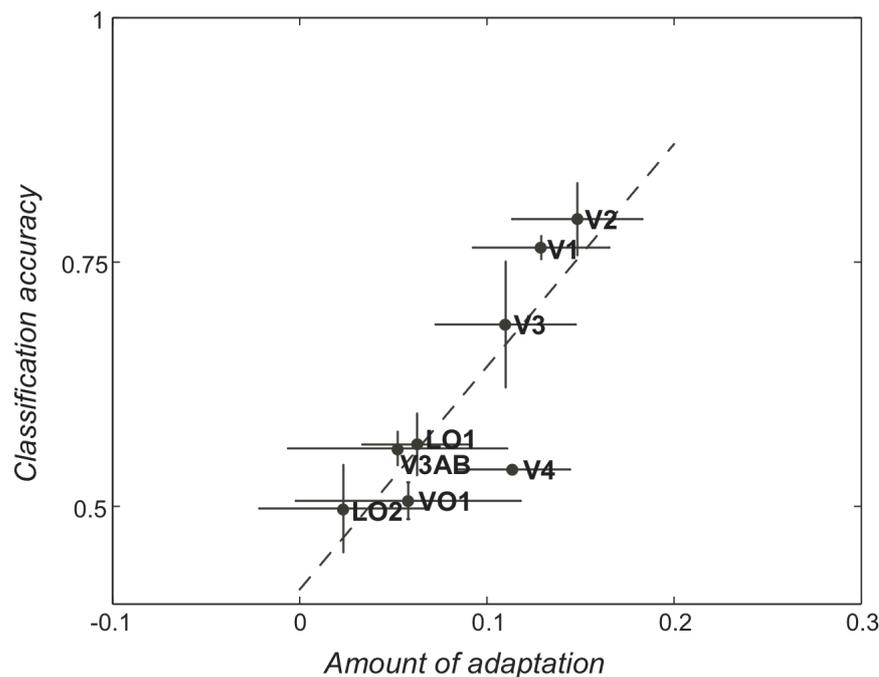


Figure 3.4 Classification accuracy plotted against the amount of adaptation across visual areas. Data are from three subjects. Error bars represent  $\pm 1$ s.e.m. The dashed line indicates the regression line ( $r=0.85$ ,  $n=8$ ,  $p<0.05$ ) computed from the averaged data across subjects.

### Comparison of results with decreasing orientation differences

Next, we tested the sensitivity of adaptation and MVPA across a range of orientation differences, for areas that showed strong orientation selectivity. For this purpose we ran two additional scanning sessions with exactly the same procedure as above, but with smaller separation between grating orientations. In one session we tested the responses to  $\pm 25^\circ$  oriented gratings and in another to  $\pm 12.5^\circ$  gratings. As the separation between orientations is reduced one would expect a drop in selective adaptation as well as lower classification accuracy. In the limit this must result in a failure to discriminate responses between stimulus categories.

The probe-related fMRI signal modulations are shown for the group average, for the  $\pm 25^\circ$  (Figure 3.5a), and the  $\pm 12.5^\circ$  (Figure 3.5b) pairs of

orientations. There was still a robust selective adaptation for orientations of  $\pm 25^\circ$ , although differences in most cases were not statistically significant (see Table 4.1). However, when separation between orientations was decreased even further (to  $\pm 12.5^\circ$ ) the adaptation was not sufficiently selective to reveal any difference between the two probes. MVPA performance is also shown for the  $\pm 25^\circ$  (Figure 3.5c), and the  $\pm 12.5^\circ$  (Figure 3.5d) oriented gratings. As the difference in orientation between target stimuli was decreased, classification accuracy also decreased but remained above chance performance even at the smallest separation (see also Table 1).

Figure 3.6 shows a summary of the data for area V1. Adaptation indices and MVPA accuracy scores were averaged across participants. Note that, since the chance level for MVPA accuracy is 0.5, and the selective adaptation index should be 0 at chance, these metrics are plotted on separate Y axes (for both, 1.0 is the maximum possible value). Surprisingly, the  $\pm 25^\circ$  orientations did not cause any less selective adaptation than the  $\pm 45^\circ$  condition although it fell to zero as separation became smaller ( $\pm 12.5^\circ$ ). The effect was also more variable between subjects at lower separations (note the size of the error bars). In contrast, classification accuracy (shown in black) falls monotonically as separation decreases, but was remarkably consistent between subjects and remained above chance performance even for small differences between target stimuli.

### 3. Comparison of fMRI adaptation and MVPA

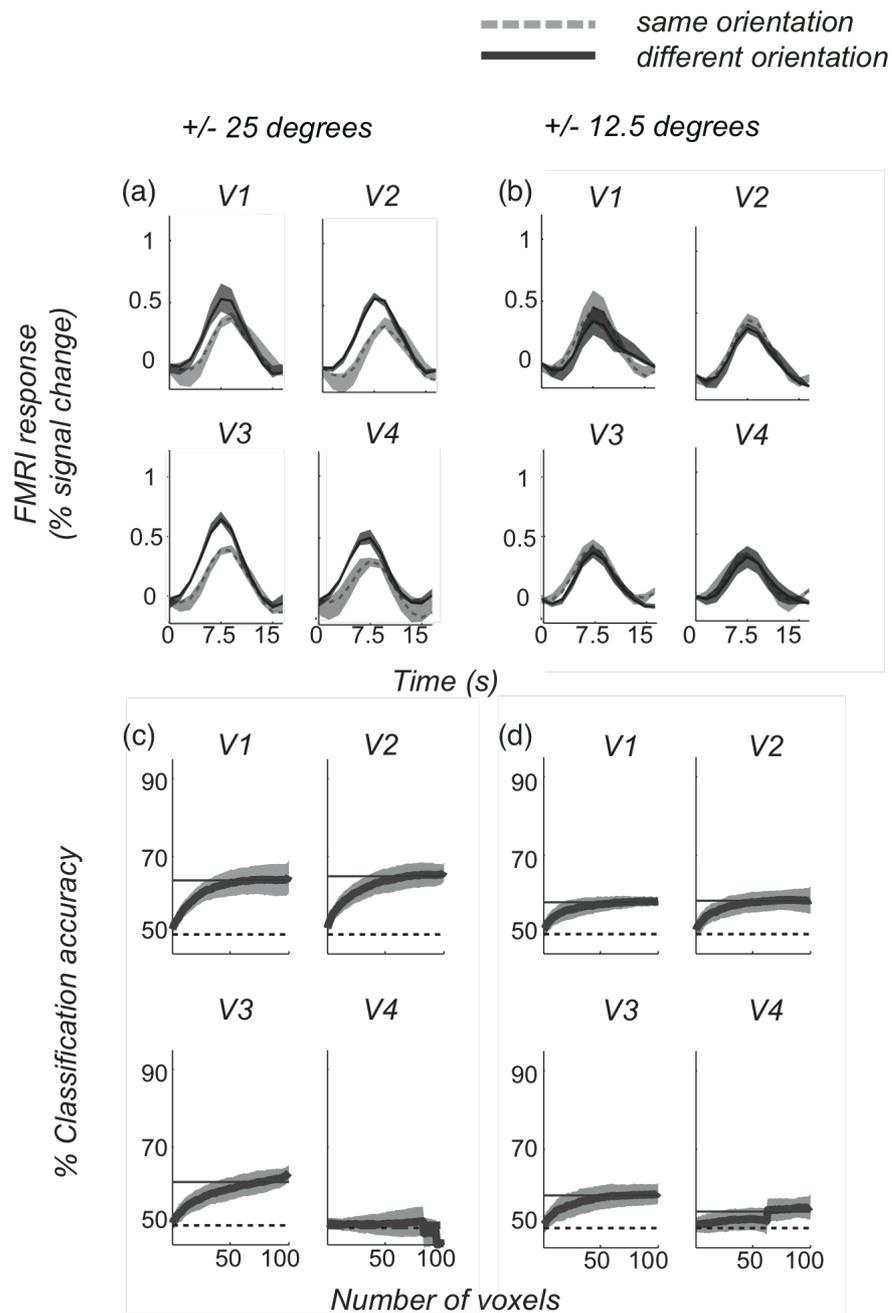


Figure 3.5 fMRI modulations in response to probe stimuli during the adaptation sequence for the (a)  $\pm 25^\circ$  and (b)  $\pm 12.5^\circ$  conditions. Performance of pattern classification versus number of voxels for the (c)  $\pm 25^\circ$  and (d)  $\pm 12.5^\circ$  conditions. Data are shown for the group average. Shaded regions represent  $\pm 1$  s.e.m. across subjects. Same conventions as in Figure 3.3.

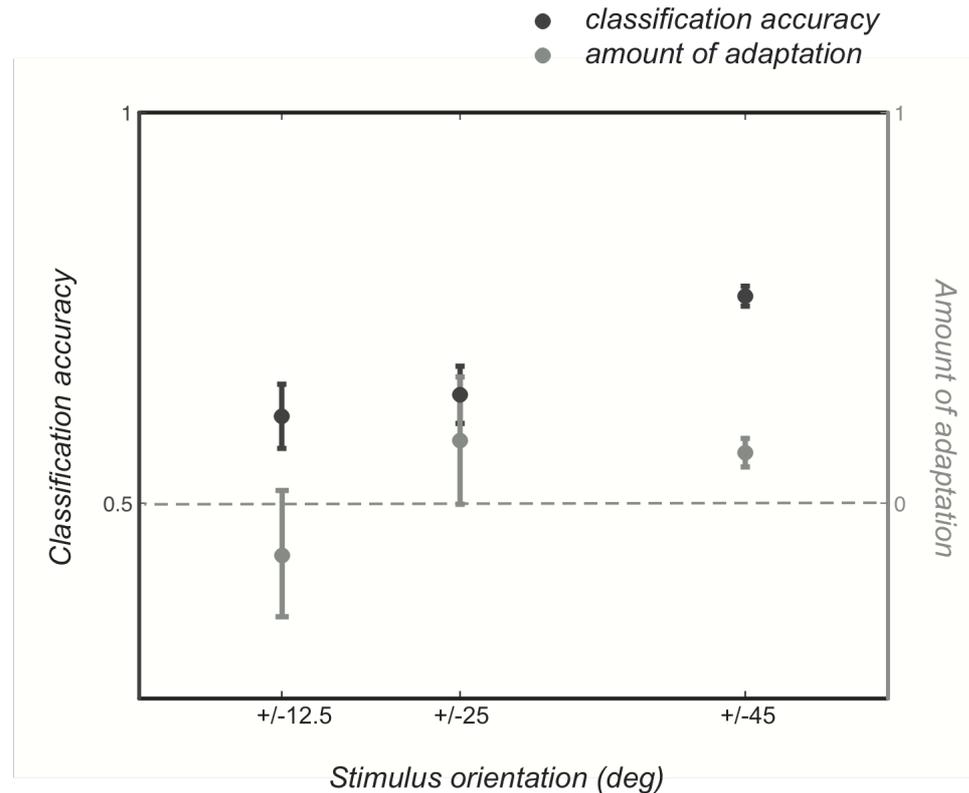


Figure 3.6 Amount of adaptation (grey) and classification performance (black) plotted against separation in stimulus orientation. Data are V1 responses averaged across three subjects. Error bars represent  $\pm 1s.e.m.$

### 3.1.3 Discussion

A wide range of fMRI studies have used selective adaptation or multivariate, pattern-classification analysis (MVPA) methods to show the selectivity of neurons on a sub-voxel scale. However, it is not known whether the two methods provide consistent results about the properties of the cortical areas under study, nor is it known which technique is the more sensitive. To address these questions, we compared the two methods directly for their ability to detect the well-documented orientation selectivity in early visual cortex. First, we considered results obtained with the two techniques using stimuli with large orientation differences. Second, we reduced the difference in stimulus orientations, to determine the dependence of each technique on orientation differences.

Both methods were clearly capable of revealing orientation selectivity in early visual areas (V1, V2, V3). For the MVPA method this has been shown previously by several studies using both support vector machine (Kamitani and Tong, 2005) and linear discriminant analysis (Haynes and Rees, 2005) and the pattern of results has been rather robust between groups. Kamitani and Tong found a diminishing trend of orientation preference across V1-V4, but no selectivity in MT+. Similarly, Haynes and Rees obtained higher classification accuracy in V1 than in V2 and V3.

For adaptation a number of groups have also shown that orientation selectivity can be demonstrated but the pattern of data has been more variable, and seems critically dependent on the choice of experimental parameters. In particular the duration of the adaptation period has a clear effect, with short-durations (1s) failing to show orientation selectivity in area V1 (Boyton and Finney, 2003). Fang et al (2006) used both short (1s) and long (5s) duration adaptation periods and show that with prolonged adaptation, the effects in V1 become measurable, although they also found greater effects in extrastriate cortex. Larsson et al (2006) used long periods of adaptation (4s) as well as a lower-contrast probe and found roughly equal adaptation indices across V1, V2, V3 and V4. We have used similar parameters with identical findings.

We found less selectivity to orientation in later visual areas (V3AB, LO1, LO2 and VO1), using both MVPA and adaptation methods. The correlation between the results across visual areas was high ( $r=0.85$ ,  $n=8$ ,  $p<0.05$ ) indicating that the two methods are in strong agreement, at least in the domain of orientation specificity. This may not have been the case and increases our confidence in both methods. It also potentially informs our understanding of orientation selectivity in the areas studied. Electrophysiology studies have demonstrated that orientation selectivity is a common feature of early visual areas. This is well-documented for

V1 and V2, but has been less explored in V3 (for review see Lennie, 1998). fMRI adaptation experiments had suggested a degree of selectivity, but this may have been simply a result of adaptation in earlier areas, resulting in reduced input to V3 (Larsson et al., 2006). The fact that we find high MVPA performance, as well as orientation-selective adaptation in this area, increases the confidence that human V3 does indeed code for stimulus orientation. Area V4 deviated furthest from the regression line but still fell well within its 95% confidence limits.

Similarly finding an absence of orientation-selective adaptation alone, would not rule out the presence of orientation-selective mechanisms; these may be present, but not susceptible to adaptation. This possibility is less likely when a second independent method (MVPA) also finds a lack of selectivity.

It should be noted that this pattern of results depended on our choice of parameters in the adaptation method, as discussed above. We chose to use low contrast probes because these are known to produce robust selective adaptation effects in previous psychophysical (Georgeson, 1985; Snowden, 1996), neurophysiological (e.g. Maffei et al, 1973) and fMRI (Larsson, 2006) studies. However, had we chosen a shorter adaptation period, or higher-contrast probes the correlation between the methods may have been weaker.

In order to determine how each technique depended on stimulus orientation, we performed the experiment with smaller orientation differences. MVPA performance remained above chance for all pairs of orientations tested and was remarkably consistent between participants. Selective adaptation failed to distinguish stimuli with smaller separations; for a 50° separation it failed in one of the three individuals, for a 25° separation it failed in all participants. This is in agreement with previous data from Fang *et al.* (2005). Their data show that, although an adaptation effect can be measured to probe stimuli as little as 7.5° from that of the adapter, the effect is not selective; the responses to such probes

are statistically indistinguishable from that to probes matching the adapter in orientation. Since we are interested in the degree to which these methods can separate the underlying neuronal populations, we consider this a failure of the method at this orientation difference.

It should be noted that these results may not be mirrored in other domains of visual selectivity. For instance, measuring the degree of selectivity to spatial frequency, direction of motion or faces may give very different results if the neurons that code these dimensions in a particular area adapt strongly, but are only weakly clustered. Clearly however, in the case of orientation selectivity measurements, the multi-variate pattern analysis was rather more sensitive than the selective adaptation measure, although the two methods were in close agreement for most visual areas.

#### **3.2 Does the interleaving of scans affect the performance of the methods?**

A limitation with the study might be the fact that the MVPA and adaptation scans were run interleaved in the same session. While this would have reduced between-scan variability, there is the concern that carry-over adaptation from the adaptation scans could bias the MVPA and vice versa. For example interleaving would, affect the performance of the MVPA method because analysis included time-points following both directions of adaptation. That is, the adaptation may have introduced noise into the measurements of the responses that the MVPA was relying on for classification. Conversely, the use of both high-contrast orientations in the MVPA runs would have washed out the adaptation effect (in effect by inducing adaptation to both orientations) which would have reduced performance in the subsequent adaptation scan. To address these concerns we collected three further sessions of fMRI data on one of the participants, with the two adaptation orientations and the MVPA data being collected on three separate days.

##### **3.2.1 Results**

Figure 3.7 shows data from the three further sessions, collected for one participant (JWP) in the lowest orientation separation ( $\pm 12.5^\circ$ ). In these sessions data for MVPA, for adapt-clockwise and adapt-anticlockwise were collected on three separate days, so that none could affect the other. Data were then analysed in exactly the same manner as the original dataset, which is also presented for comparison. Panels a, c plot adaptation and MVPA responses, collected in the original interleaved manner. Panels b, d plot responses with data being collected on separate days.

The pattern of results was the same; for this orientation difference there was no evidence of selective adaptation (if anything there was actually a greater

### 3. Comparison of fMRI adaptation and MVPA

response in V1 to the 'same' orientation after adapting), but there remained a significant, albeit weaker rate of success in discriminating the two orientations for the MVPA.

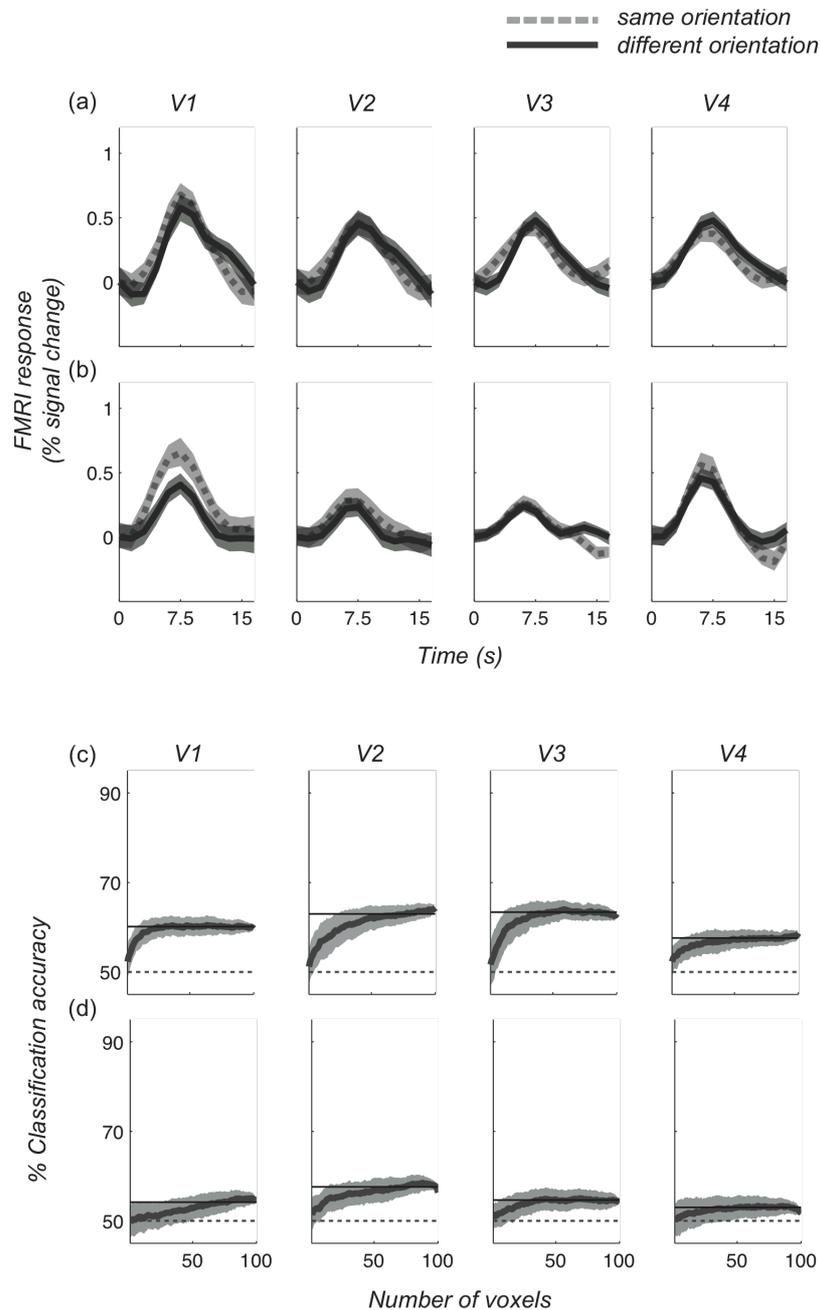


Figure 3.7 FMRI modulations in the lowest ( $\pm 12.5^\circ$ ) orientation separation. (a), (c) adaptation and MVPA accuracy collected in the original interleaved manner. (b), (d) same as above but with data being collected on different days.

#### **3.2.2 Discussion**

In order to maximize the similarity between methods in the main experiment, data were collected for the experiment in an interleaved manner, with alternating runs of adaptation and MVPA data collection. To check that this interleaving procedure did not itself affect the results, for example by cross-adapting the subject to both orientations, resulting in a weak adaptation index, we ran three further sessions for one participant (JWP) in the lowest orientation separation ( $\pm 12.5^\circ$ ). The data confirmed that, at the lowest orientation separation, the data still show no significant selectivity of adaptation but do show MVPA classification accuracy above chance. This increases our confidence that the effect did not result from the interleaved nature of the scans. There is marginally less MVPA accuracy than in the main study, which could be attributed to variability between days; since these data are now collected on different days of scanning, session-to-session variability now becomes a potential new confound.

#### 3.3 Do orientation biased voxels adapt more?

We performed a further analysis to determine whether there was any consistent relationship between the methods on an individual voxel level; whether a voxel that adapted strongly was also a strong driver of the MVPA classification. In our analyses, the former was characterised by the adaptation index calculated for individual voxels, the latter ultimately depended on the t-statistic of orientation preference for each voxel. In particular, voxels with higher t-values are more biased towards a particular orientation. If there was a correlation between orientation preference and the amount of adaptation, voxels with higher t-values should adapt more.

##### 3.3.1 Results

The results of this analysis are shown in Figure 3.8. Data are shown from a single, representative subject (JWP), for two ROIs (V1 and V4) and for all three orientation separations. This figure plots the adaptation index (from the adaptation scans) as a function of orientation preference (from the MVPA scans). Filled symbols label the subset (100 voxels) included in the MVPA analysis (based on their response during an independent localiser scan). Light grey represents voxels preferring clockwise stimuli; dark grey represents voxels preferring anticlockwise stimuli. There was no apparent relationship between the methods on a voxel-by-voxel basis.

##### 3.3.2 Discussion

We performed a complementary analysis to determine whether voxels that are more biased towards a particular orientation adapt more strongly. We found no such a correlation in any of the conditions tested. There is, in fact, no *a priori* reason that there should be a correlation between the methods on a voxel-wise basis, despite the fact that the methods agree in their analyses of the ROIs as a

### 3. Comparison of fMRI adaptation and MVPA

whole. The adaptation index is a relative measure of the responses to identical stimuli pre- and post-adaptation and does not necessarily depend on the amplitude of the response to that stimulus. Conversely, the MVPA is dependent on voxels differing in their actual responses to different stimuli.

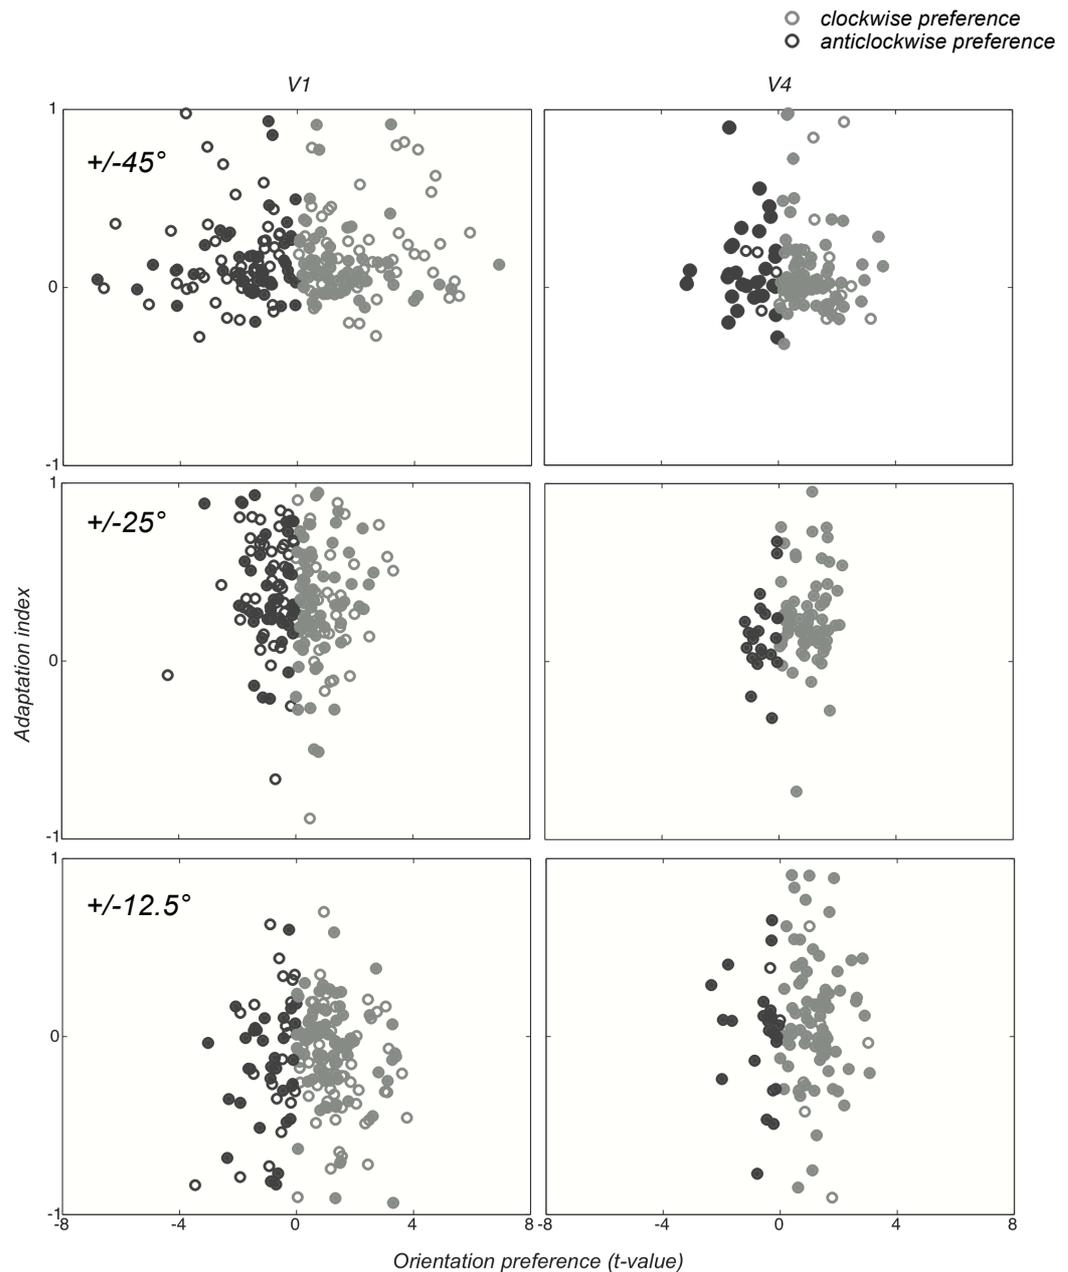


Figure 3.8 Voxel-wise analysis from a single participant (JWP), for two ROIs (V1 and V4) and for all three orientation separations. Filled circles represent voxels included in the MVPA.

The data presented in this section are in agreement with the sharpening adaptation model, which predicts that neurons showing little or no adaptation to a repeated stimulus are highly selective for that stimulus (see section 1.3.1). Although fMRI data need to be combined with single-cell recordings to conclude on the model of adaptation that better applies in this case, the results presented here contribute to this direction.

#### **3.4 High-field (7T) measurements**

Improvements to imaging hardware and analysis techniques have provided access to higher-resolution images at improved signal-to-noise ratios. High field experiments suggest that imaging with resolutions in the hundreds of microns range should be possible, and have reliably mapped columns in primary visual cortex with ocular and even orientation selectivity (e.g. Moon et al. 2007; Yacoub et al. 2007).

As 7T imaging is still rather in its infancy, we performed a pilot experiment to investigate how results from different magnetic fields compare. In principle, both methods should benefit from high-field fMRI. As pattern-classifiers rely on feature selectivity and spatial inhomogeneity of feature-selective responses, high spatial resolution should improve the performance of the MVPA by reducing averaging between columns for instance. Adaptation could benefit simply from the improved signal-to-noise ratio. We also re-ran the voxel-wise analysis described in the previous section to determine whether and how it differs.

##### **3.4.1 Methods**

**Scanning parameters.** High-resolution data were collected on a 7T (Philips Achieva) scanner and the scanning region was restricted to include primarily area V1 (voxel size  $1 \times 1 \times 1.5 \text{ mm}^3$ , TR 2.5s, TE 25ms; FOV 128mm x 128mm x 24mm; 16 slices; 1-shot EPI).

**Scanning session.** At the beginning and the end of the scanning session, we ran two localiser scans. This was followed by four adaptation scans (of 90 dynamics each), and two MVPA scans (of 120 dynamics each).

**Pattern-classification.** In each scanning session, we obtained data in 20 blocks (2 blocked scans, 10 blocks per scan). From each block we extracted the responses at 5 separate time-points, delayed by two TR to account for the haemodynamic lag. There were therefore a total of 100 time-points (or repeated 'examples' of each response) for each voxel. Trials were divided into 10 groups, each group corresponding to responses collected in a pair of blocks, 1 from each orientation of the original dataset. Data from 9 of these groups were assigned to a training set and the remaining to a test set.

#### 3.4.2 Results

The probe-related modulations in fMRI signal obtained at high field are shown in Figure 3.9a. Data are from area V1 of one subject, collected at the highest orientation separation ( $\pm 45^\circ$ ). We obtained a substantial orientation-selective adaptation effect in area V1. Figure 3.9b plots the performance of the pattern classifier based on linear discriminant analysis. As in previous sections, classification accuracy is plotted against the number of voxels included in the analysis. Because of the larger TR (2.5s) used in this study, fewer time-points were available for the pattern-classification algorithm. As the number of voxels included in the analysis depends on the number of training samples available, we were forced to include only 60 voxels. We obtained significant MVPA accuracy but less compared to the 3T data.

Figure 3.10 plots adaptation index versus orientation preference t-value for 300 high-resolution V1 voxels. The adaptation index was computed for each voxel as it was for an ROI and, basically, describes how strongly it adapts (see

General Methods). The orientation preference t-value was computed as the weighted difference between the responses to the two orientations, and captures the orientation bias of each voxel. If there was a correlation between the two, voxels with high orientation preference should also adapt a lot. However, as with the 3T data no such relationship was obtained.

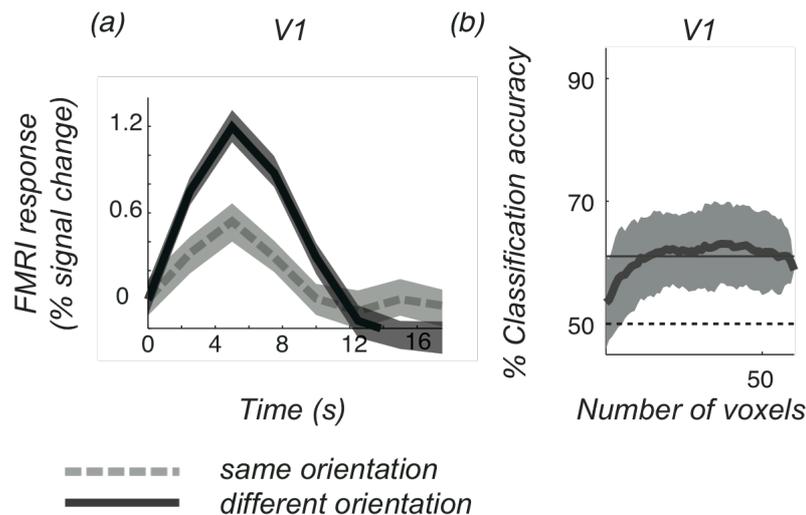
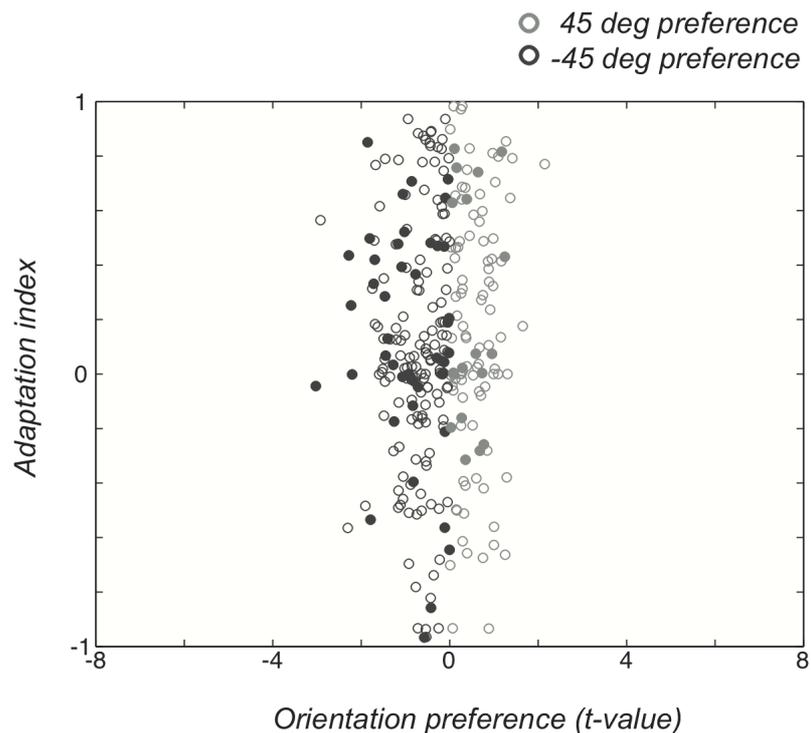


Figure 3.9 (a) High-field fMRI modulations in response to probe stimuli during the adaptation sequence for the  $\pm 45^\circ$  condition averaged over 40 trials. (b) Performance of pattern classification versus number of voxels included in the analysis. Data are from area V1 of one participant (JWP). Shaded regions represent  $\pm 1s.e.m.$  Same conventions as in Figure 3.3.

### 3.4.3 Discussion

We performed the adaptation and MVPA studies at high magnetic field (7T) to benefit from the high spatial resolution and improved signal-to-noise ratio. Data were collected for a single participant and the field of view included predominantly area V1. Strong selective adaptation and significant pattern classification effects were obtained. However, because of the small data set the 7T data cannot be directly compared to the 3T. In particular, selective adaptation was stronger at 7T; however, this might be due intersession variability. More data on different participants need to be collected. MVPA accuracy was significant but reduced

compared to 3T. As discussed above due to the largest TR at 7T less time points were available (100 at 7T; 300 at 3T). As the performance of linear discriminant analysis depends strongly on the number of time points available, it is critical to compare performance between the two field strengths using the same number of examples.



*Figure 3.10 Voxel-wise analysis for data obtained at high-field. Open circles are the 300 V1 voxels. Filled circles are voxels included in the MVPA. Data are for one participant (JWP), collected at the largest orientation separation ( $\pm 45^\circ$ ).*

Taking advantage of the improved spatial resolution at 7T we repeated the voxel-wise analysis discussed in the previous section. We wanted to test whether voxels biased towards a particular orientation, also adapt stronger to this orientation. As with the 3T data no particular correlation was found. Interestingly, high-resolution voxels exhibited smaller orientation preference t-values. All in all, despite that fact that 7T fMRI is still in evaluation we obtained comparable results

### 3. Comparison of fMRI adaptation and MVPA

---

to 3T. This increases our confidence that high magnetic field strength can provide a reliable tool for studying activity in the visual cortex.

#### ***4 How are mechanisms selective for complex visual patterns represented in the cortex?***

---

Having directly compared fMRI adaptation and multi-variate pattern analyses (MVPA) in the previous chapter, this chapter uses both techniques to investigate how local orientation signals are combined and detected in intermediate levels of visual processing. A great deal is known about the initial steps of visual processing, however, it is not well understood how the outputs of V1 are combined in order to represent more complex visual features. The simplest stimulus with which to study the combination of V1 outputs is probably the plaid, the linear sum of two (or more) sinusoidal gratings. Although plaids have been more typically used to study the combination of motion signals (Huk and Heeger 2002; Movshon et al. 1986; Rust et al. 2006) they can also be used to study the conjunction of form signals (Carandini et al. 1997a; Georgeson and Meese 1997; Peirce and Taylor 2006). Peirce and Taylor (2006) presented psychophysical evidence that the visual system has mechanisms sensitive to the overall form of compound patterns rather than to their components. The rationale of their experiments was that if the visual system has plaid detecting mechanisms, then these mechanisms would adapt more strongly to a plaid than to the components alone. Using a novel method of compound adaptation, participants simultaneously adapted to sinusoidal plaid patterns separated across the two visual hemi-fields. The two fields contained identical grating components (A, B, C, D) rearranged into different plaid patterns (AB, CD vs AC, BD) (Figure 4.1).

The aim of the experiments presented in this chapter was to use fMRI adaptation and MVPA to identify the cortical location of plaid-selective mechanisms. In the first experiment we used an event-related adaptation paradigm (Fang et al. 2005; Larsson et al. 2006; Montaser-Kouhsari et al. 2007).

After prolonged adaptation to high-contrast patterns, sensitivity was tested using low-contrast probes. In the second experiment we used a blocked adaptation paradigm similar to that of Huk and Heeger (2002). In the third experiment we used a multi-variate pattern analysis (MVPA) algorithm (Cox and Savoy 2003; Haynes and Rees 2006; Norman et al. 2006) to discriminate between the responses to a plaid and its components.

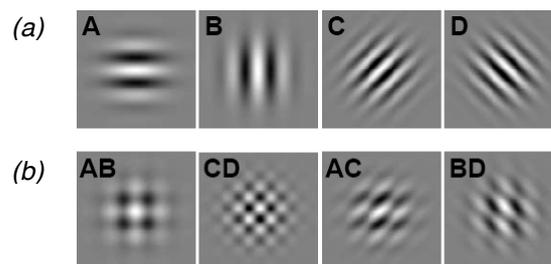


Figure 4.1 (a) Simple component gratings are combined (b) to construct plaid stimuli (from Peirce and Taylor, 2006).

### 4.1 Experiment 1

Event-related adaptation paradigms (Fang et al. 2005; Larsson et al. 2006; Montaser-Kouhsari et al. 2007) often have the advantage of being closer to their psychophysical and electrophysiological counterparts. They rely on the observation that, after prolonged presentation of a particular stimulus, the BOLD response in areas sensitive to that stimulus, is selectively reduced compared to the response to other stimuli.

In this study, adaptation to a plaid (AB) was compared to adaptation to its components (alternating gratings, A and B) (Figure 4.1). Participants were adapted to the plaid in one visual hemifield and to the components in the other (Figure 4.2). Critically, exposure to the component gratings was equal in both hemifields. The design relies on the assumption that areas selective to plaids would adapt stronger to the plaid stimuli than to the components.

To maximise the sensitivity of the design we replaced the letter-counting task at fixation with a dot-counting task that took place, not at fixation, but was distributed across the envelopes of the stimuli being presented. This way we forced subjects to attend to the stimuli instead of diverting attention away from them. However, the task had a heavy attentional load at all times, and so still served to minimise the potential for confounding variations in attention to the stimuli.

### 4.1.1 Methods

#### *Participants, procedure and stimuli*

Three experienced volunteers took part in the study. The event-related fMRI adaptation protocol is shown in Figure 4.2. At the beginning of each scan participants passively viewed the adapting stimuli for 30s. The adaptors were the component gratings alternating in one hemi-field and the plaid in the other. Critically, each component was presented at 50% contrast for half the time, but either presented in or out of phase with each other. This resulted in the plaid being presented at full-contrast for half the time and the two half-contrast gratings for all the time. This way the exposure to components in terms of duration and intensity was equal in both visual hemi-fields<sup>1</sup>. In each trial, adaptation was maintained by presenting 'top-up' adaptors for 4.5s, followed by a blank screen for 0.75s. Then, the probe stimuli were presented in both hemi-fields for 1.5s. Four probes were shown, (a) component A, (b) component B, (c) plaid AB or (d) blank screen. These conditions were equally common and randomly chosen. The

---

<sup>1</sup> An alternative way of controlling for contrast is to present stimuli on both sides at 100% with the plaid being presented throughout the adaptation period. This way the pixel intensities are matched in both hemi-fields. fMRI measurements and psychophysical experiments within the lab indicated similar results for both stimulus presentation ways.

Michelson contrast of probe stimuli was 40%<sup>1</sup>. Each trial ended with a 0.75s presentation of a blank screen, giving a total duration of 7.5s.

**Scanning session.** Each scanning session consisted of a localiser scan at the beginning of the session and 4 adaptation scans. Between adaptation scans, we alternated the side on which the plaid adaptor was presented, i.e. the plaid adaptor was presented on the left hemi-field in 2 scans and on the right in the remaining 2. This controlled for the possible side biases that participants might (and in fact did) have in BOLD responses. Each scan consisted of 60 trials, 15 trials for each probe stimulus.

**Visual stimuli.** The spatial frequency of component stimuli was 0.75 cycles/° oriented at 0° or 90°. All stimuli were presented in a Gaussian envelope with a standard deviation of 0.33 (such that the stimuli had a diameter of 5° at the point at which it fell below 1% contrast). The spatial phase of the stimuli was randomised every 2 frames to prevent retinal afterimages.

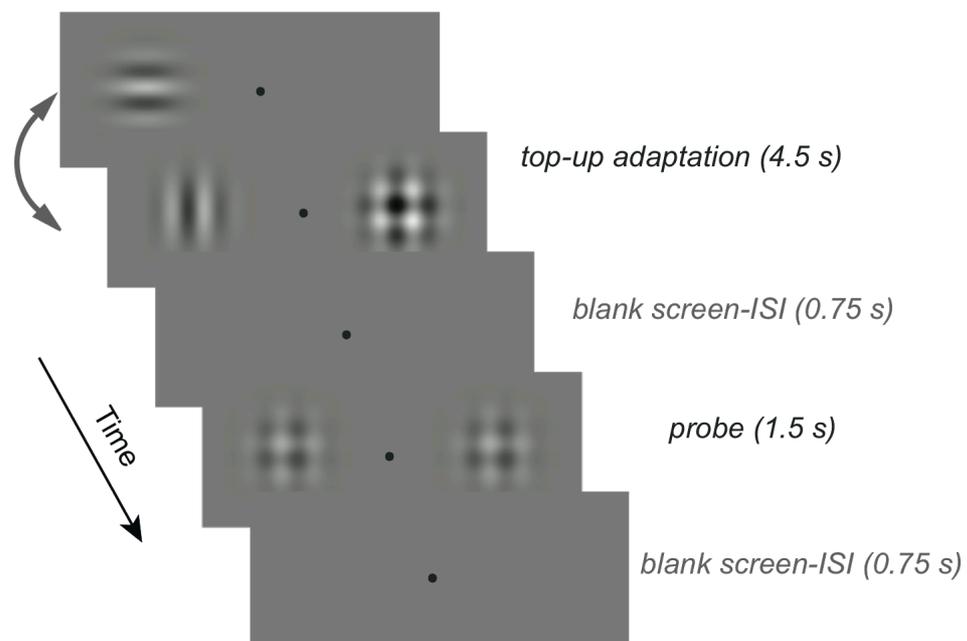
**Localiser scan.** At the beginning of each scanning session we performed a localiser scan. The purpose of the localiser was to identify voxels in the ROIs that responded to visual stimulation at the spatial location of the patterns. The stimulus was a 100% Michelson contrast plaid, presented for 18s in one visual hemi-field and for 18s in the other hemi-field. Each localiser scan consisted of 10 such blocks. A fixation point was presented throughout. The responses evoked by the localiser stimulus are shown in Figure 4.3.

**Attention control task.** To control for changes in the attentional state of observers, they performed an attention-demanding task. To increase the sensitivity of the design, instead of a task at fixation used in the previous chapter,

---

<sup>1</sup> The intensity of probe stimuli was chosen based on a series of psychophysical experiments within the lab.

the task was distributed across the envelopes of the stimuli being presented. Participants were asked to count the number of faint (50% contrast) red coloured dots appearing on the region the stimuli were drawn. The duration of each dot-counting trial varied randomly between 7-12s. Within each trial, 1-4 dots were presented at random positions within the envelope of the stimuli, either on the left or the right side of the screen. At the end of each dot-counting trial, the fixation turned red for 1s prompting participants to report, within this brief period, the number of dots counted by pressing a button. A fixation point was presented throughout the experiment.



*Figure 4.2 Event-related procedure. Each trial consisted of a top-up adaptation period for 4.5s, followed by a blank screen for 0.75s, succeeded by probe presentation for 1.5s. Trials ended with presentation of the blank screen for 0.75s. Total duration of each trial was 7.5s.*

#### **Data analysis**

Regions of interest (ROIs) were restricted to include only voxels whose time series correlated with the stimulus epochs of the localiser scan. Specifically, ROIs were restricted only to those voxels with a coherence  $c_0 > 0.25$  and a response

phase  $0 < \phi < \pi$  for the left hemisphere, and  $\pi < \phi < 2\pi$  for the right hemisphere (localiser stimuli initiated on the right visual hemi-field).

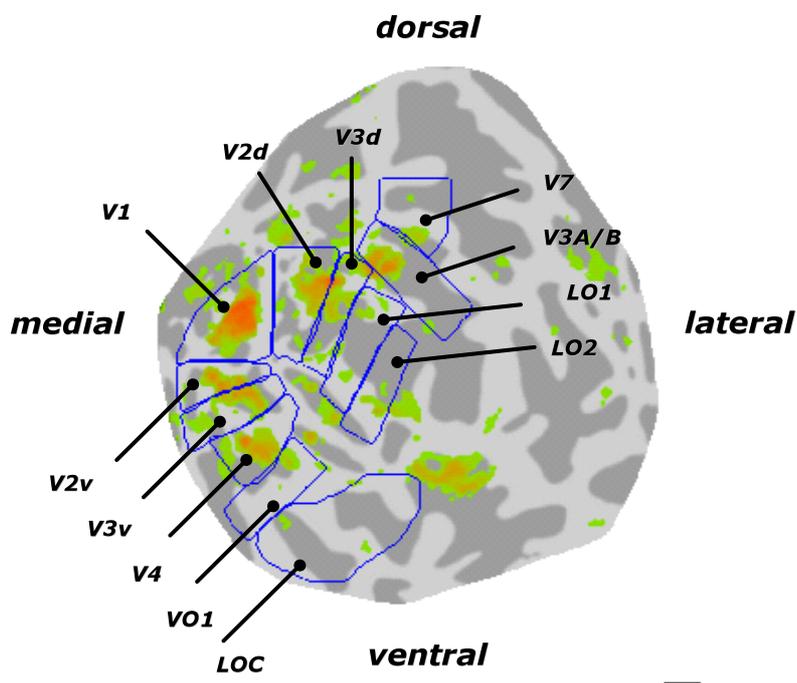
Responses to individual trials were extracted from the average ROI time course by selecting a 15s window initiating 3s after the onset of the adaptor. The average response to the blank trials (which captures the response to the adaptor alone) was subtracted from each trial. Trials of each type were then averaged and the resulting event-related time courses were adjusted to zero baseline by subtracting the mean of the first two time points (before the onset of the probe stimulus).

#### 4.1.2 Results

**Responses to the localiser stimulus.** Responses evoked by the localiser plaid stimulus are shown in Figure 4.3. As revealed by the coherence activation map, responses were stronger in areas V1, V2, V3, V4 and V3AB. The pattern of responses was generally consistent across subjects. As plaid localiser patterns were presented consecutively on the left and the right horizontal meridians, patches of activation appeared in the midst of V1 and on the borders of areas V2 and V3, consistent with the expected retinotopic representation. There appeared also to be reliable responses dorsal and lateral to LOC, possibly where V5 (MT+) is located. Because an individual mapping scan was required to accurately define V5 we did not include this area in the analysis. An exploratory analysis was performed based on retinotopic and localisation maps to include these activated voxels. However, no plaid selectivity was obtained in this region, possibly because of the fact that the stimulus phase was jittered, leaving no consistent motion signal, and the data are not considered further here.

**Event-related fMRI responses.** Event-related fMRI responses to the plaid probe are shown in Figure 4.4 for three subjects. Data are shown in nine visual

areas for two of the participants and in eight regions for one participant (LOC-mapping data were not available for subject DS). In general, fMRI signal modulations were greater than 0.5% in early visual areas. Responses in lateral-dorsal areas (LO1, LO2) were weak or variable depending on the participant. Lateral ventral areas (VO1 and LOC) that did not show much activation in the localiser scan also produced inconsistent responses across subjects.

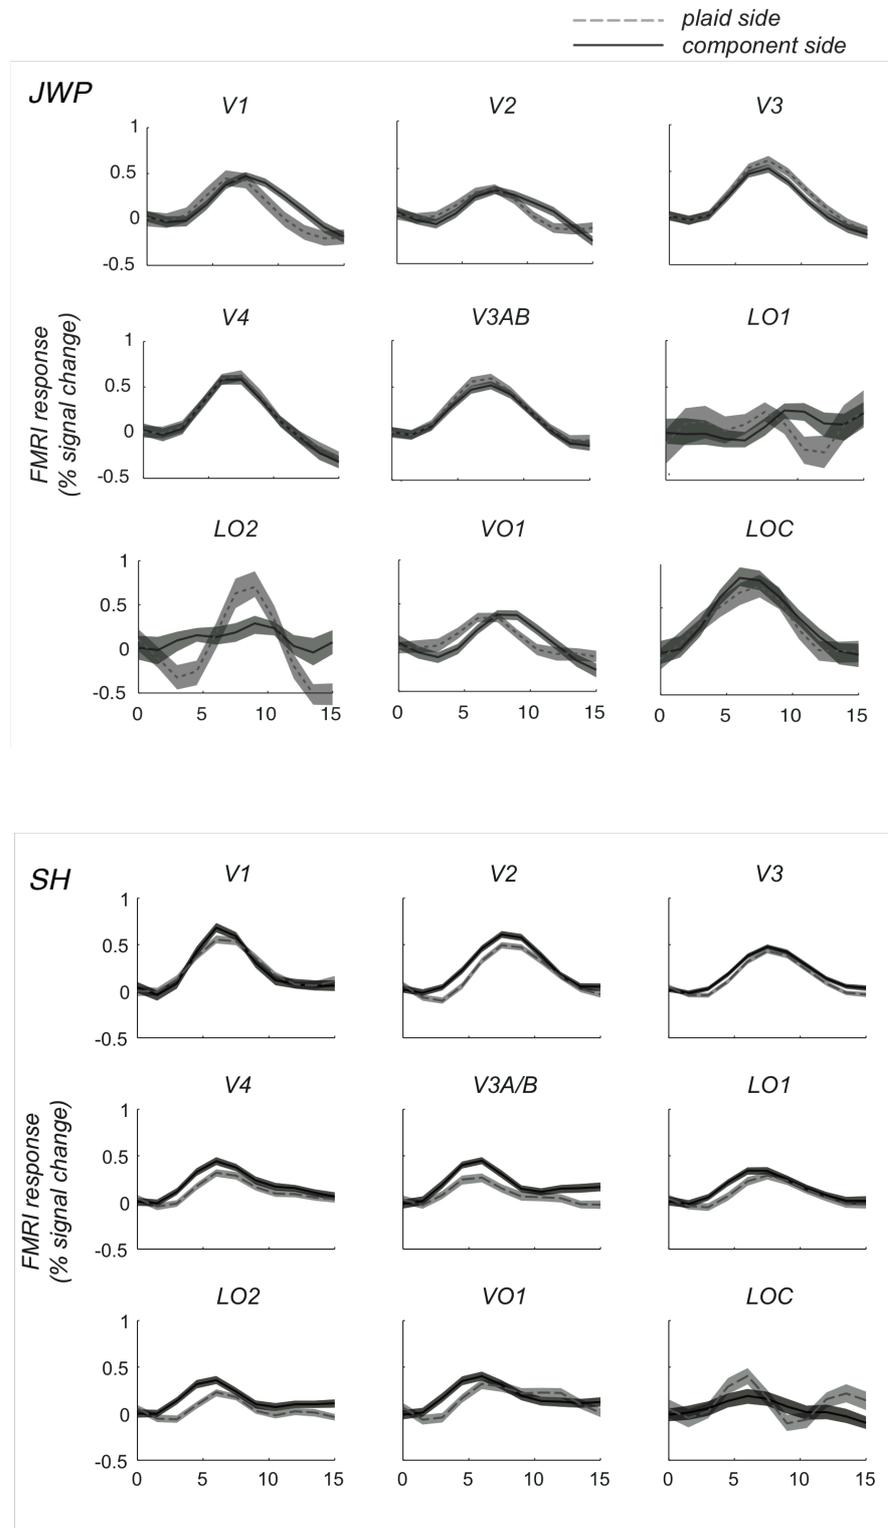


*Figure 4.3 Coherence map elicited by the localiser plaid stimulus. Activation corresponds to coherence values greater than 0.25. Adaptation data were analysed in ROIs restricted to include only those voxels activated in the localiser scan. Data are from the right hemisphere of one subject (SH).*

The degree of selective adaptation in each area can be seen as the difference between the response to the plaid probe presented at the side adapting to the plaid adapter (shown in light grey), versus the response to the plaid probe presented at the side adapting to the components (dark grey). If an ROI exhibits plaid selective adaptation then response shown in light grey would be suppressed compared to the response shown in dark grey. Subtle plaid-selective adaptation effects were evident in some areas in some subjects but these effects were not

#### 4. Mechanisms selective for complex patterns

consistent across participants. Overall, no consistent plaid-selective adaptation effect was evident across participants in any of the areas tested.



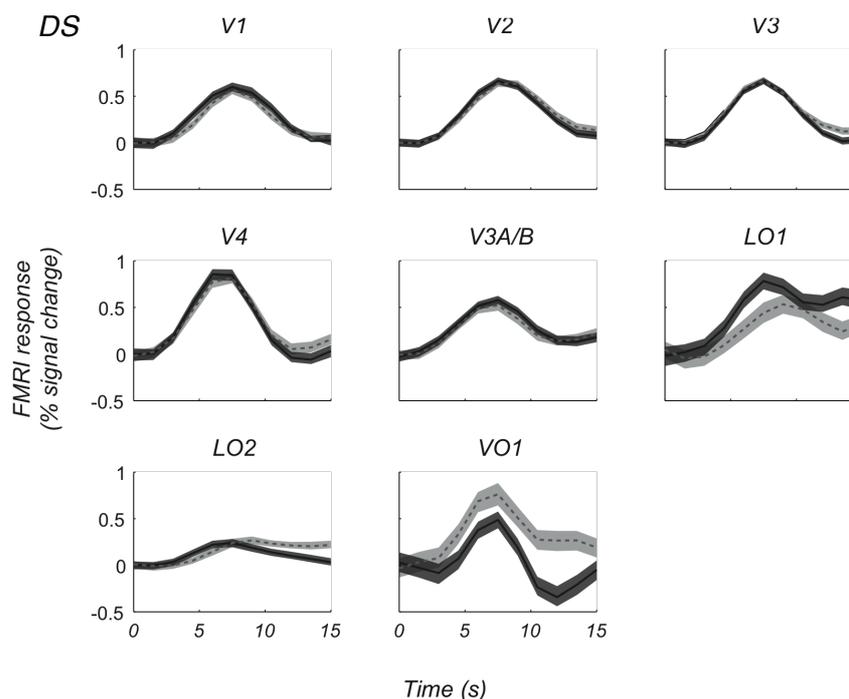


Figure 4.4 Plaid adaptation fMRI time courses are shown for three subjects (JWP, SH, DS) in order of presentation. The response to the plaid probe on the side adapting to the plaid is shown light grey and on the side adapting to the components in dark grey. Time courses initiated 3s after the onset of the adaptor stimulus. They were adjusted to zero baseline by subtracting the mean of the first two time points (before the onset of the probe stimulus). LOC mapping data were not available for subject DS. Shaded regions are standard errors of the mean.

#### 4.1.3 Discussion

In this experiment an event-related adaptation paradigm was used. Responses to low-contrast probes were measured after prolonged adaptation to high-contrast plaid and grating patterns. Subtle plaid-selective adaptation effects were evident in some instances but these were not replicated across participants. Although we modified the attention control task to enhance sensitivity, we did not obtain consistent plaid adaptation effects across subjects in the visual areas studied. The paradigm described here is the last of a series of plaid adaptation experiments performed during the first two years of the PhD. Over this period we used a number of different experimental parameters (stimulus contrasts, timing

parameters) and attention control tasks. In none of these experiments did we succeed in generating consistent plaid adaptation effects. One possible explanation is that the design itself was not sensitive enough to give rise to an effect. In the event-related paradigm described above brief, low-contrast probes generated the signal of interest. It is possible that the brevity and low-intensity of the probe stimuli was not sufficient to produce a measurable effect.

For this reason, we tried a different experimental paradigm; namely, a blocked adaptation design. This paradigm despite being dissimilar to the psychophysics is much simpler. Blocked paradigms have a longer history in fMRI research, and have the advantage of being more straightforward in their design and their interpretation. This experiment is discussed in the following section.

### **4.2 Experiment 2**

In the second experiment we used a blocked adaptation paradigm similar to that of Huk and Heeger (2002). The rationale was the following. In one block two plaids (AB, CD) were alternating in one visual hemi-field, and in a second block four plaids (AB, CD, AC, BD) were alternating in the other visual hemi-field. Again, the experimental design ensured that each component grating (e.g. A) was presented for the same amount of time in both locations. Therefore, adaptation to component gratings was the same in both hemi-fields. However, in the ‘two-plaids’ adaptation condition each plaid is presented for half of the duration in each block, whereas, in the ‘four-plaids’ condition each plaid is presented for a quarter of the duration. Consequently, any putative detector for the plaid would be expected to adapt more during the ‘two-plaids’ condition, since its preferred stimulus was present for a greater proportion of the time.

### 4.2.1 Methods

#### *Participants, procedure and stimuli*

Two experienced volunteers participated in the study. Plaid stimuli (Figure 4.1) were made of sinusoidal gratings (spatial frequencies 0.75 and 1.5cycles/°) oriented at four different orientations (0°, 90° and +/-45°). Stimuli occupied 5° degrees of visual angle. The edges of the stimuli were smoothed by a Gaussian mask. The spatial phase of the gratings was randomised every 5 frames (83.3ms) to prevent retinal afterimages.

In this experiment no separate localiser scan was performed. Instead we averaged all blocked scans and used this to select activated voxels so that the data themselves acted as their own, unbiased localiser. Each scanning session consisted of 6 blocked adaptation scans. In half of the scans the 'two-plaid' condition was presented in the left visual field and in the other half in the right visual field. Each scan consisted of 10 blocks, 30s each, resulting in a total duration of 300s (200TRs) per scan. In each block one of the conditions (two- or four-plaid) was presented in one visual field while a blank screen (mean luminance) was presented in the other visual field. For example, in scans where the two-plaid adaptor was on the left, this would be presented for 30s with nothing in the right visual hemifield and then the four-plaid would be presented in the right hemifield for 30s with nothing being presented in the left. As a result there was always some visual stimulus being presented to the participant, rather than having periods where no stimulus was presented, interleaved with periods of both presented at the same time. A letter counting task at the centre of fixation controlled for attention.

### 4.2.2 Results

Blocked fMRI responses for one subject are shown in Figure 4.5. Each trial consisted of 30 seconds of stimulus onset alternated with 30 seconds display of the blank screen (at mean luminance). The degree of selective adaptation in each area can be seen as the difference between the response to the 'four-plaid' condition (light grey), versus the response to the 'two-plaid' condition (dark grey). In principle, this difference should be greater at the end of the trial when adaptation was stronger. Data are from a single scanning session and averaged over 30 trials. Blocked time-courses were normalised so that each time course started at 0 and its maximum value was 1. This was done to better appreciate the relative signal decrease for the two conditions; since not all plaids in the four-plaid condition were present in the two-plaid adaptor, the initial response to the stimulus might differ, and this would be of no interest to the current experiment. However, responses obtained without normalising for the initial response yielded similar results.

Both conditions caused a robust modulation in the BOLD signal on the stimulus onset. Subsequently, fMRI responses gradually decreased during the end of the trial, presumably resulting from adaptation. This suppression was observed during both conditions and ranged in amplitude from 0.25 to 0.5 of normalised percent signal. After the end of stimuli presentation, the BOLD responses returned to baseline. Responses in higher regions were noisier, probably due to the small number of voxels activated in these regions. However, no consistent evidence of plaid-selective adaptation was obtained in any of the areas tested; the time-courses for the two conditions appeared to overlap.

#### 4. Mechanisms selective for complex patterns

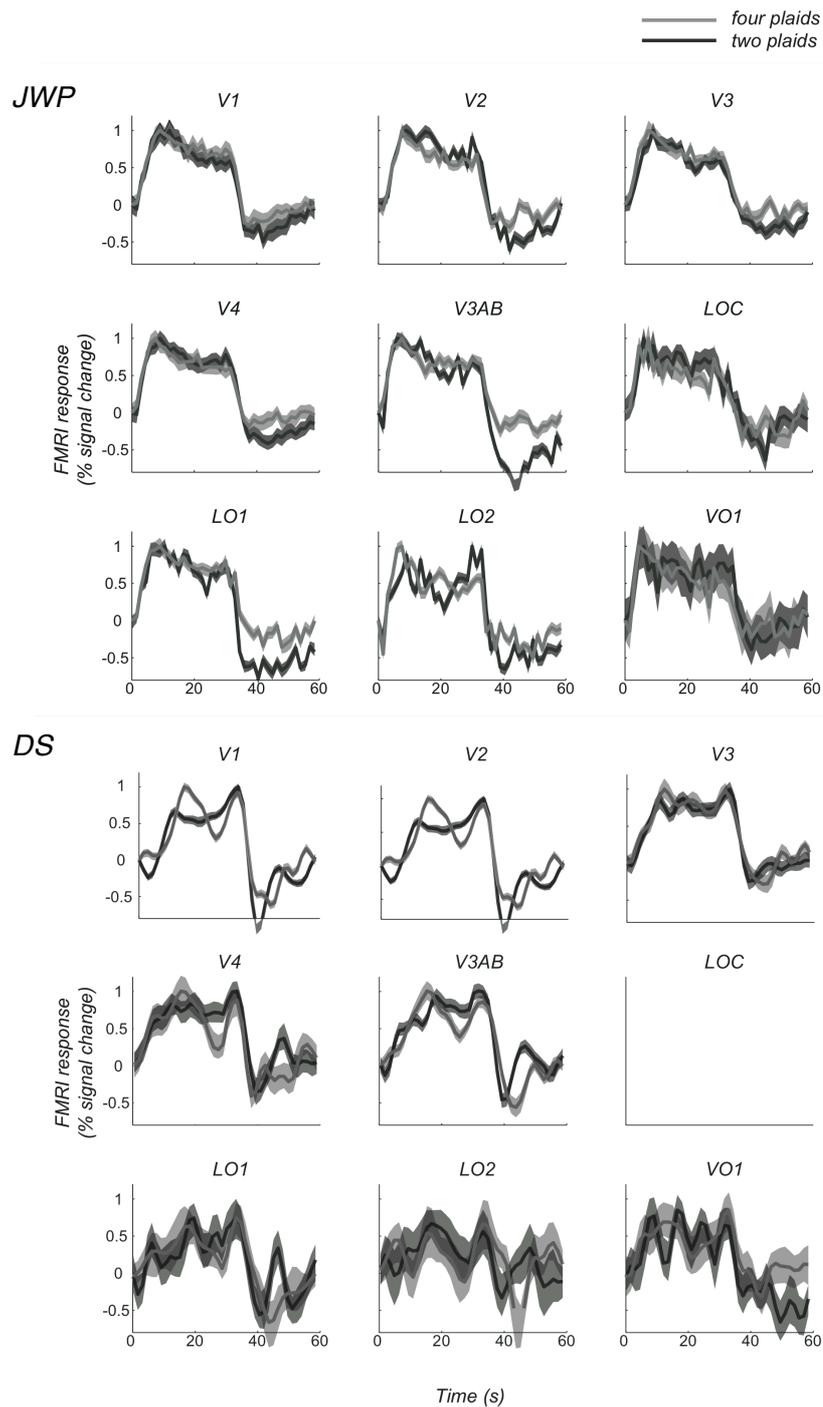


Figure 4.5. Blocked fMRI responses for the ‘two-plaids’ (dark grey) and ‘four-plaids’ (light-grey) condition, in nine visual areas. Data are from two subjects (JWP, DS), obtained during a single scanning session, and averaged over 30 trials. LOC data for subject DS were not available. Time-courses were normalised so that each time course starts from 0 and its maximum value is one. Shaded regions are standard errors of the mean across the trials.

### 4.2.3 Discussion

Blocked designs are extensively used in fMRI research. In such a design different conditions are alternated in separate blocks. Usually a rest condition is interleaved between testing trials, during which the BOLD signal returns to baseline, introducing a maximum amount of variability. The robust modulations usually evoked in a blocked design improves significantly the signal to noise ratio. As a result blocked designs offer considerable statistical power (Donaldson and Buckner 2001). If a condition consists of repeated presentation of a stimulus category it typically results in adaptation in regions sensitive to the stimulus. Such a blocked adaptation paradigm is more sensitive to the adapted response near the end of the trial. To make the most of the blocked adaptation paradigm we used blocks of long duration (30s).

However, data collected for two participants did not generate consistent plaid adaptation effects. The reason for that might be that although block adaptation designs generate strong and reliable signals, they produce weak psychophysical effects because high-contrast stimuli are used throughout. It is also possible that in conventional ROI analyses, where time-courses are averaged together across voxels to obtain a mean ROI response, subtle signal modulations in voxels that do adapt are lost. As shown in Chapter 3, pattern-classification analyses are more sensitive in detecting even small differences in stimulus parameters. Therefore, we performed a third experiment in which we used a pattern classifier to attempt to discriminate between responses to a plaid and its components.

### 4.3 Experiment 3

Considering the increased sensitivity of multi-variate pattern classification analyses (MVPA) described in the previous chapter, we ran a third experiment in

which a plaid stimulus was presented in one block and the components of that plaid in another block. Because the accuracies obtained with a parametric linear discriminant classifier were low, we decided to use a non-parametric neuronal network (without a hidden layer) algorithm to discriminate responses between the two conditions. The experiment relies on the hypothesis that visual areas containing mechanisms selective to complex stimuli, should exhibit higher classification accuracy than the rest.

In fact, a second mechanism may lead to success on the part of the classifier in this experiment. We know that many cortical cells exhibit cross-orientation suppression; responses to a stimulus with optimal orientation are partially suppressed by superposition of a second grating with orthogonal orientation (Carandini et al. 1997b; Carandini et al. 2002). This effect is thought to be caused by a divisive gain control from suppressive lateral connections of a normalisation pool (Heeger 1992). The effect of cross-orientation suppression could result in responses to the plaid stimulus in certain voxels or areas being reduced. Consequently, MVPA accuracy in these areas would be high, though, not due to plaid selectivity. To examine the relative contributions of these two effects we ran the analysis multiple times using: a) the set of voxels as in the previous chapter for orientation selectivity, b) only the voxels that responded more to the plaid than to the components, or c) only the voxels that responded less to the plaid than to the components. The procedure is described in more detail below.

#### **4.3.1 Methods**

##### ***Participants, stimuli and procedure***

Three experienced volunteers took part in the study. Each scanning session consisted of 6 blocked scans. The paradigm is shown in Figure 4.6. In one block

#### 4. Mechanisms selective for complex patterns

---

the two components, presented at 50% Michelson contrast, were alternating every second. In the other block the plaid stimulus, presented at 100% Michelson contrast, was alternating with the blank screen every second. This way it was ensured that exposure to components was the same in both blocks. The two conditions were alternated with epochs of blank screen (mean luminance) with a period 30s (15s 'on', 15s 'off'). Each scan consisted of 8 blocks.

**Visual stimuli.** Component stimuli consisted of sinusoidal gratings oriented at  $0^\circ$  or  $90^\circ$  (spatial frequency 1.5cycles/ $^\circ$ ). The plaid was made of the superposition of the two components. Stimuli were presented in an annulus (inner radius,  $2^\circ$ ; outer radius,  $14^\circ$ ) whose edges were smoothed by a Gaussian kernel (std dev of  $0.083^\circ$  on the inner edge,  $0.333^\circ$  on the outer edge). The spatial phase of the stimuli was randomised every 6 frames (100ms) to prevent retinal afterimages.

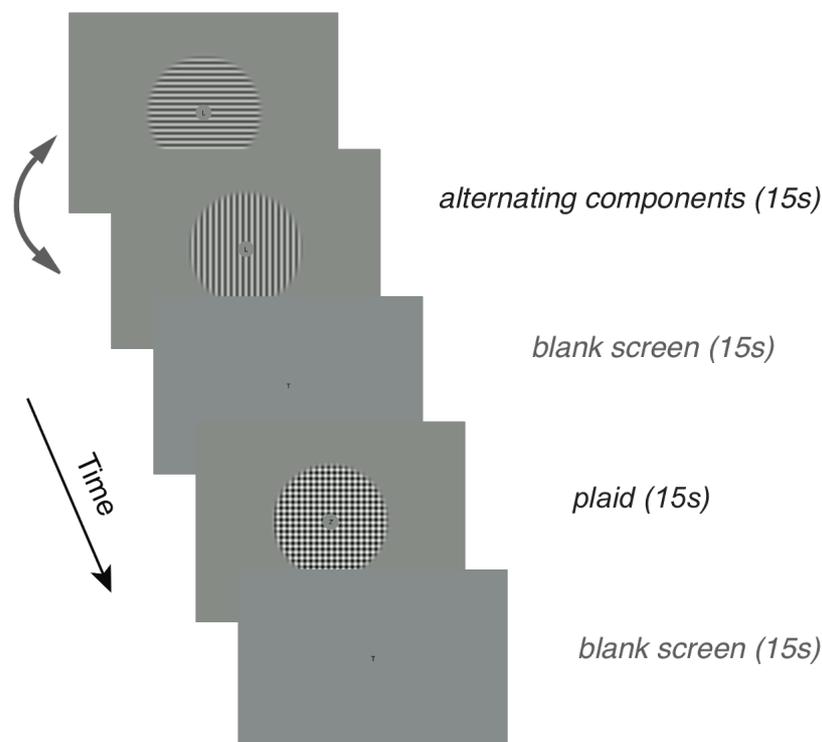


Figure 4.6 A block design was used to measure responses for pattern classification.

**Attention control task.** To control for changes in the attentional state of observers a letter counting task was performed at fixation (see section 2.3.4).

#### ***Data analysis***

Regions of interest (ROIs) were restricted to include only voxels whose time series correlated with the stimulus epochs (coherence,  $c > 0.3$ , phase  $0 < \phi < \pi$ ).

**Pattern-classification.** A linear neuronal network (without a hidden layer) algorithm was used to classify responses between stimulus conditions (see General Methods). Like the linear discriminant algorithm used in Chapter 3, this algorithm is also a linear classifier, however, it does not assume a specific model of the data points (is non-parametric). It has been shown that in certain instances non-parametric classifiers result in better accuracies (Ku et al. 2008). Because of the low accuracies obtained in this section we decided to use a non-parametric classifier. However, their disadvantage is that they require considerably more computing time. In each scanning session, we obtained data in 48 blocks (6 blocked scans, 8 blocks per scan). From each block we extracted the responses at 10 separate time-points (over 15s), delayed by three TRs to account for the haemodynamic lag. There were therefore a total of 480 time-points (or repeated 'examples' of each response) for each voxel. Trials were divided into 12 groups, each group corresponding to responses collected in two blocks, 2 from each stimulus condition. Data from 11 of these groups were assigned to a training set and the remaining to a test set to protect against overfitting. This was repeated with different blocks acting as the test set and the performances on each repeat averaged (cross-validation).

### 4.3.2 Results

#### *Mean ROI responses*

The (univariate) mean ROI responses to the two stimulus categories are shown in Figure 4.7 for all participants. The BOLD signal modulation to both stimulus conditions was greater in V1 and decreased in higher visual areas. However, there was no difference in the amplitude of the response between the two conditions.

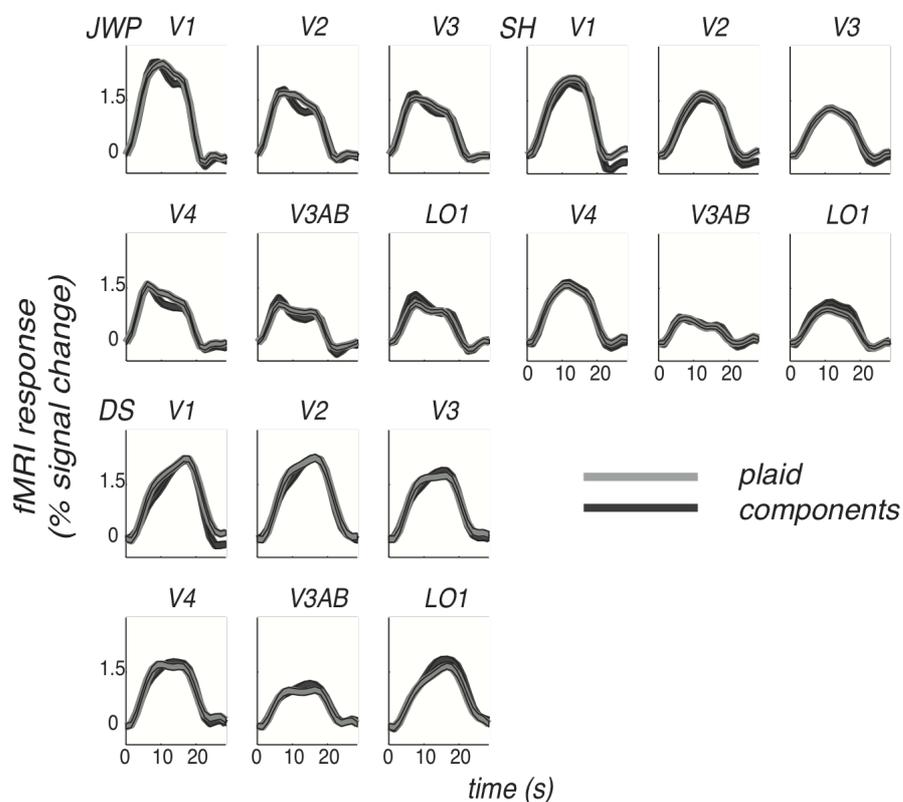


Figure 4.7 Mean ROI responses to the alternating gratings (dark grey) and the plaid (light grey) for all participants. Shaded regions are s.e.m.

#### *MVPA performance*

Figure 4.8 plots the performance of the pattern classifier, based on a neuronal network algorithm, in one participant (a) and across participants (b). Classification accuracy was computed directly on the 100 voxels, in each ROI,

included in the analysis. Performance was found to be better than chance in most areas tested. However, high classification performance, especially in early visual areas, might have resulted from cross-orientation suppression. In these regions, the responses to the plaid would be partially suppressed compared to responses to the components.

Figure 4.9 shows a histogram of t-values (in area V1) computed as the normalised difference between the responses to the two conditions. Voxels with a negative t-value responded more to the plaid; voxels with a positive t-value responded more to the components. To examine the effect of cross-orientation suppression we reran the analysis by including only voxels with negative t-values.

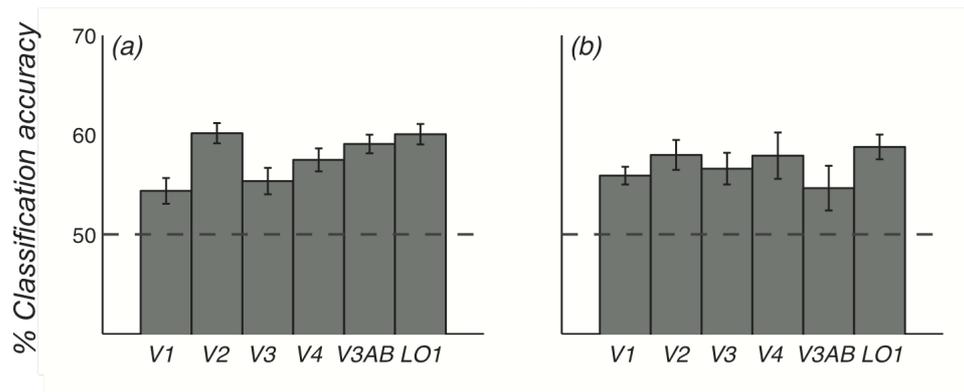


Figure 4.8 (a) MVPA performance for one subject (JWP). Error bars are standard deviations computed over 100 reshuffles. (b) MVPA performance averaged across subjects. Error bars represent  $\pm 1$ s.e.m. across the three participants. The dashed line shows classification accuracy based on chance (50%).

#### 4. Mechanisms selective for complex patterns

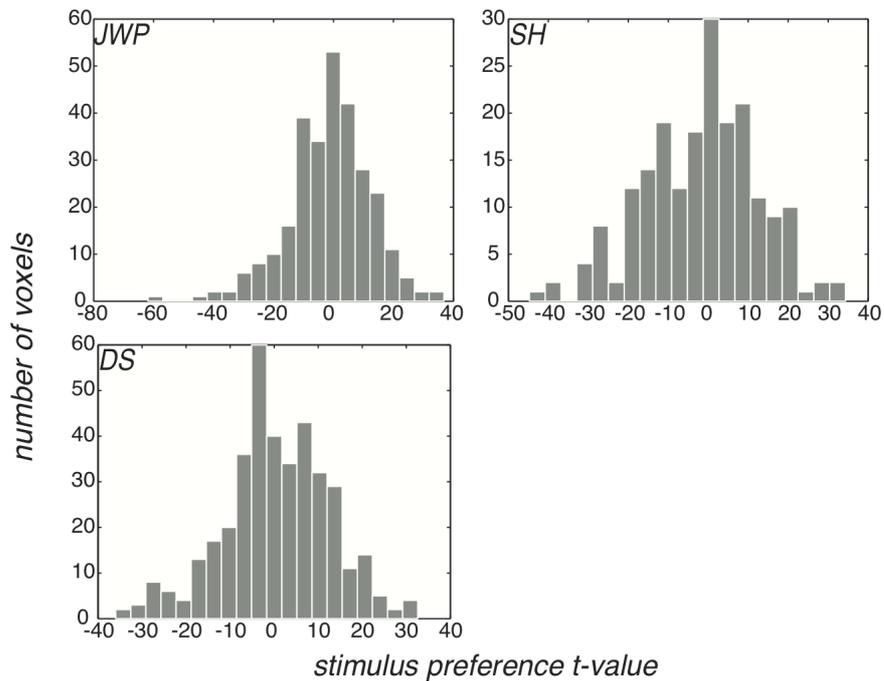


Figure 4.9 Condition preference t-values from area V1 of all participants. The same pattern was observed in the other ROIs.

Table 4.1 Pattern-classification accuracies for individual subjects. Analysis was based on voxels in each ROI that responded more strongly to the plaid. P-values (uncorrected for multiple comparisons) are shown in parentheses. P-values were estimated from a permutation test with 5000 resamples (see General Methods).

	V1	V2	V3
JWP	0.50 (0.431)	<b>0.55 (0.049)</b>	0.52 (0.253)
DS	<b>0.59 (0.000)</b>	<b>0.58 (0.002)</b>	<b>0.59 (0.000)</b>
SH	<b>0.56 (0.031)</b>	<b>0.56 (0.015)</b>	0.53 (0.127)
	V4	V3AB	LO1
JWP	0.54 (0.089)	<b>0.56 (0.022)</b>	<b>0.57 (0.010)</b>
DS	0.54 (0.101)	<b>0.56 (0.026)</b>	0.53 (0.141)
SH	0.53 (0.174)	0.51 (0.405)	<b>0.56 (0.031)</b>

The results of the analysis are shown in Figure 4.10, for one participant (panel a) and across participants (panel b). Compared to Figure 4.8, classification accuracy was somehow lower in most areas. The results (uncorrected for multiple comparisons) are summarised in Table 4.1. Classification accuracy was higher

and significant for all participants in area V2. However, after applying a Bonferroni correction for multiple comparisons, performance was significant only for one participant (DS). For comparison, Figure 4.10 plots MVPA accuracies for voxels that responded more strongly to the components. MVPA accuracies are about equal in most visual areas and also more variable compared to Figure 4.10.

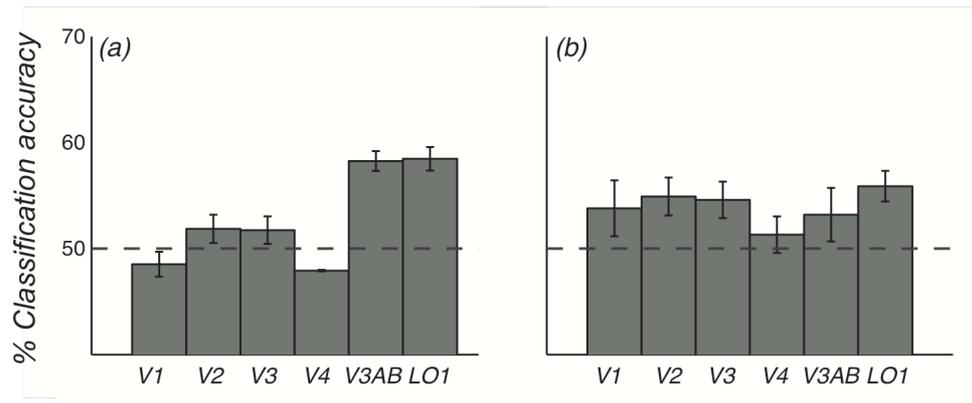


Figure 4.10 (a) MVPA performance for one subject (JWP), and (b) averaged across subjects. Analysis was based on voxels in each ROI that responded more strongly to the components. Same conventions as in Figure 4.8.

Table 4.2 Pattern-classification accuracies for individual subjects. Analysis was based on voxels that responded more strongly to the components. Same conventions as above.

	V1	V2	V3
JWP	0.49 (0.669)	0.52 (0.253)	0.52 (0.253)
DS	<b>0.56 (0.015)</b>	<b>0.55 (0.049)</b>	0.54 (0.073)
SH	<b>0.57 (0.013)</b>	<b>0.58 (0.002)</b>	<b>0.58 (0.004)</b>
	V4	V3AB	LO1
JWP	0.48 (0.719)	<b>0.58 (0.002)</b>	<b>0.58 (0.000)</b>
DS	0.54 (0.115)	0.51 (0.350)	<b>0.56 (0.026)</b>
SH	0.53 (0.174)	0.50 (0.405)	0.53 (0.115)

### 4.3.3 Discussion

In this section we used multi-variate pattern analysis (MVPA) to study complex pattern selectivity in the visual cortex. Because of the low MVPA accuracies obtained in this section we decided to use a non-parametric classifier.

It has been shown that in certain instances non-parametric classifiers result in better accuracies (Ku et al. 2008). First, we computed the mean response in each area for all participants. Interestingly, mean ROI responses were equal for both conditions, in most areas tested. Although time-courses in individual voxels are often noisy, the voxel-wise analysis we performed confirmed the results of the mean ROI analysis. The stimulus preference histograms indicated a symmetrical distribution of responses; about the same number of voxels were responding more strongly to each condition.

Then, we tested the performance of a neural network pattern classifier to discriminate responses between a plaid and its components. When all voxels in each ROI were considered, we obtained better than chance performance in most areas. To control for inhibited responses to the plaid due to cross orientation suppression, we re-ran the analysis by including only voxels that responded more strongly to the plaid. MVPA accuracies in this case were universally above chance in area V2. The analysis also indicated significant results in areas V1, V3AB and LO1 for at least two of the participants. For comparison, we ran the same analysis by including only voxels that responded more strongly to the components. The results in this case were more variable across the participants. Although MVPA accuracies in this study were weak compared, for example, to those of Chapter 3, the results might point to the direction that area V2 could perform tasks such as feature integration.

#### **4.4 General discussion**

A great deal is known about the initial steps of visual processing, however, it is not well understood how the outputs of V1 are combined in order to represent more complex visual features. In this chapter we used fMRI adaptation and multi-variate pattern analysis (MVPA) to study plaid selectivity in the human visual cortex. Experiment 1 utilised an event-related adaptation paradigm (Fang et al.

2005; Larsson et al. 2006; Montaser-Kouhsari et al. 2007), Experiment 2 a blocked adaptation paradigm (Huk and Heeger 2002), and Experiment 3 an MVPA method (Cox and Savoy 2003; Haynes and Rees 2006; Norman et al. 2006).

Both adaptation experiments failed to reveal plaid selectivity in the areas tested. Why was no plaid-specific adaptation effect evident in the visual cortex despite consistent evidence from psychophysical studies (Peirce and Taylor 2006, and several further studies in the lab)? One possibility is that the magnitude of plaid-induced adaptation was too small to be resolved by fMRI. As mentioned above, the psychophysical shift in perceived contrast was only 5%. It is possible that such subtle modulations in brain activity do not give rise to detectable blood-oxygen level-dependant (BOLD) changes. In addition, the designs themselves provide certain limitations. For example, event-related designs can generate strong perceptual effects but the measured signal is weak because of the brief, low contrast probes. On the other hand, block adaptation designs generate strong and reliable signals, but weak psychophysical effects because high-contrast stimuli are used throughout (the adaptor and probe are a single entity). Another possibility is that in conventional ROI analyses, where time-courses are averaged together across voxels to obtain a mean ROI response, subtle signal modulations in voxels that do adapt are lost. On the other hand, pattern-classification analyses are more sensitive in detecting even small differences in stimulus parameters (see Chapter 3).

In Experiment 3 we used a neural network classifier to discriminate responses between a plaid and its components. As in Chapter 3, the MVPA was more sensitive than adaptation in revealing pattern selectivity. However, the MVPA accuracies obtained in this study were lower compared, e.g. to those in Chapter 3. Furthermore, in most cases, accuracies did not exceed significance

after applying a multiple comparison correction. Nevertheless, for uncorrected P-values, MVPA accuracies were universally better than chance in area V2. V2 is ideally located to perform tasks such as feature integration. In terms of functional properties, it is known that V2 cells have larger receptive fields than V1 cells. This property, coupled with the greater overlap of receptive fields in V2, indicates a cruder analysis performed by this area. This, in turn, might suggest a feature integration role for V2 cells. Although the results we obtained are somewhat weak, they point to the direction that the role of V2 might be to perform tasks such as Fourier component integration.

Plaids have been typically used to study the combination of motion signals. However, the adaptive properties of spatial plaid-selective mechanisms are not very well understood. On the other hand, contour stimuli have been used in the literature in many different forms. For the purpose of studying mid-level vision psychophysically, Hancock and Peirce (2008) arranged simple gratings in a way that give the percept of edges. Contour stimuli have been used more extensively to study mid-level vision and the neuronal mechanisms underlying contour selectivity are a lot better understood. In the following chapter we used adaptation and multi-variate pattern classification analyses, to study the cortical mechanisms underlying contour selectivity.

### ***5 Integration of simple contours in the visual cortex***

It appears that at different levels of visual processing, signals from previous stages are combined to create a more complex and sparse representation of images. In the striate cortex, for example, the outputs of retinal ganglion cells are combined to create elongated receptive fields that typically respond to particular Fourier (orientation and spatial frequency) components (Hubel and Wiesel 1968). These are thought of by many as suitable detectors for edges or lines; locations in visual space where chromatic or luminance contrast is high. Little, however, is known about how outputs of V1 are organised and combined by the visual system in order to represent more complex visual features. Understanding and characterising these mechanisms is obviously a fundamental problem for vision science. It is possible that conjunctions of local orientation signals are detected in intermediate levels of visual processing in the form of curves or corners, thus providing the basic building blocks for what Marr (1982) referred to as the primal sketch.

Several investigators have studied how features are integrated into visual contours using curved stimuli in various different forms (Field et al. 1993; Hess and Field 1993; Loffler et al. 2003; Watt and Andrews 1982). Recent studies have reported shape aftereffects that do appear to have a basis in global shape processing rather than local orientation (Gheorghiu and Kingdom 2007; 2009; 2008; Hancock and Peirce 2008; Suzuki 2001; 2003). Particularly, Hancock and Peirce (2008) developed a psychophysical adaptation paradigm (Peirce and Taylor 2006, see also previous chapter) that was based on adapting one location to a contour and another location to the components of that contour in isolation. The Michelson contrast of all the adaptors was 100%. Using this paradigm it was ensured that both patches were adapted equally to the orientation components of

the stimuli, but only the compound patch was adapted to the global contour. Thus any difference in adaptation between the two patches should result from the presence of the contour as a global figure. After adaptation participants were probed with a high contrast (100% Michelson) compound stimulus in both locations. The orientation of the probes was varied according to a staircase procedure. Participants were asked to press one of two keys to make a 2 alternative forced choice (AFC) response indicating the side on which the stimulus appeared was most convex to the left. Hancock and Peirce found that the after-effect resulting from adaptation was greater in the side adapting to the contour than in the side adapting to the components. These results suggested the existence of detectors for conjunctions of oriented features in the visual system.

In this chapter we further investigated the nature and anatomical location of conjunction mechanisms using fMRI adaptation and multi-variate pattern analysis (MVPA). We created contours from the conjunction of two oriented gratings. The first grating was orientated clockwise or anticlockwise (at  $\pm 20^\circ$ ) and the second was mirror symmetric to the first, along the horizontal axis. In this study we ran two experiments. In the first experiment we used an event-related adaptation paradigm. The design was similar to the Hancock and Peirce (2008) protocol, modified to be used in conjunction with fMRI. In the second we used the MVPA technique to classify fMRI responses to a contour versus responses to its components in isolation.

### **5.1 Experiment 1**

In the first experiment we used an event-related adaptation paradigm (Fang et al. 2005; Larsson et al. 2006; Montaser-Kouhsari et al. 2007). In half of the scans participants adapted to a contour and in the other half to the components of that contour (Figure 5.1). They were then probed with contours of either the same or the opposite orientation. As a control condition participants were probed with a

'straight' vertical patch of the same length as the contour stimuli. This third condition was introduced to determine whether the observed effects could be caused simply by the responses of end-stopped cortical mechanisms which are known to be sensitive to stimulus length (Hubel and Wiesel 1965). Such mechanisms would respond stronger to the contour than to the components, not because they are curvature selective, but because contour stimuli are longer than the component stimuli. The design relies on the assumption that in the 'contour adaptation scans', contour selective areas would adapt stronger to contours of the same orientation than to opposite oriented contours or vertical stimuli.

In this study we also modified the timing of the event-related design. In the event-related designs described in the previous chapters presentation of the 'top-up' adaptor was followed by an inter stimulus interval (ISI) of about 0.75s. However, other studies have used rapid event-related adaptation paradigms (e.g. Grill-Spector and Malach 2001; Kourtzi and Huberle 2005), where stimuli are presented sequentially with short ISIs (0.1-0.4s). Designs with very brief ISIs may benefit from keeping the probe closer in time to the adaptor. In this experiment we kept the same design as before, with high contrast adaptors and low contrast probes, but decreased the ISI between 'top-up' adaptation and probe presentation to 0.2s.

### **5.1.1 Methods**

#### ***Participants, procedure and stimuli***

Two experienced volunteers took part in the study. The experimental protocol is shown in Figure 5.1a. Participants initially viewed passively the adapting stimulus for 30s. The adaptor was a contour in half of the scans and the components of that contour in the other half. In each trial adaptation was maintained by presenting 'top-up' adaptors for 4s. In the contour adaptation scans

the adapting stimulus was shown for only half of the adaptation duration. This way exposure to components was kept roughly constant in the 'contour' and 'component' adaptation scans. 'Top-up' adaptation was followed by a blank screen (ISI) for 0.2s. Then, the probe stimuli were presented for 1s. Four probes were shown, (a) a contour with the same orientation as the adaptor, (b) a contour with the opposite orientation to the adaptor, (c) a vertical control stimulus or (d) blank screen (Figure 5.1b). These conditions were equally common and randomly chosen. The Michelson contrast of probe stimuli was 50%<sup>1</sup>. Each trial ended with a 0.8s presentation of a blank screen, giving a total duration of 6s.

**Scanning session.** Each scanning session consisted of a localiser scan at the beginning, followed by 6 adaptation scans. In half of the scans the adapting stimulus was the contour and in the other half the components of that contour. 'Contour' and 'component' scans were interleaved. Each scan consisted of 40 trials; 10 for each probe stimulus.

**Visual stimuli.** The spatial frequency of all stimuli was 1.2 cycles/°. Contours were made of the conjunction of two oriented gratings. The first grating was oriented clockwise or counter clockwise (at  $\pm 20^\circ$ ) and the second was mirror symmetric to the first, along the horizontal axis. All stimuli were presented in a Gaussian envelope with a standard deviation of  $0.66^\circ$  (such that the stimuli had a diameter of  $4^\circ$  at the point at which it fell below 1% contrast). The spatial phase of the stimuli was randomised every 6 frames (100ms) to prevent the formation of luminance afterimages on the retina.

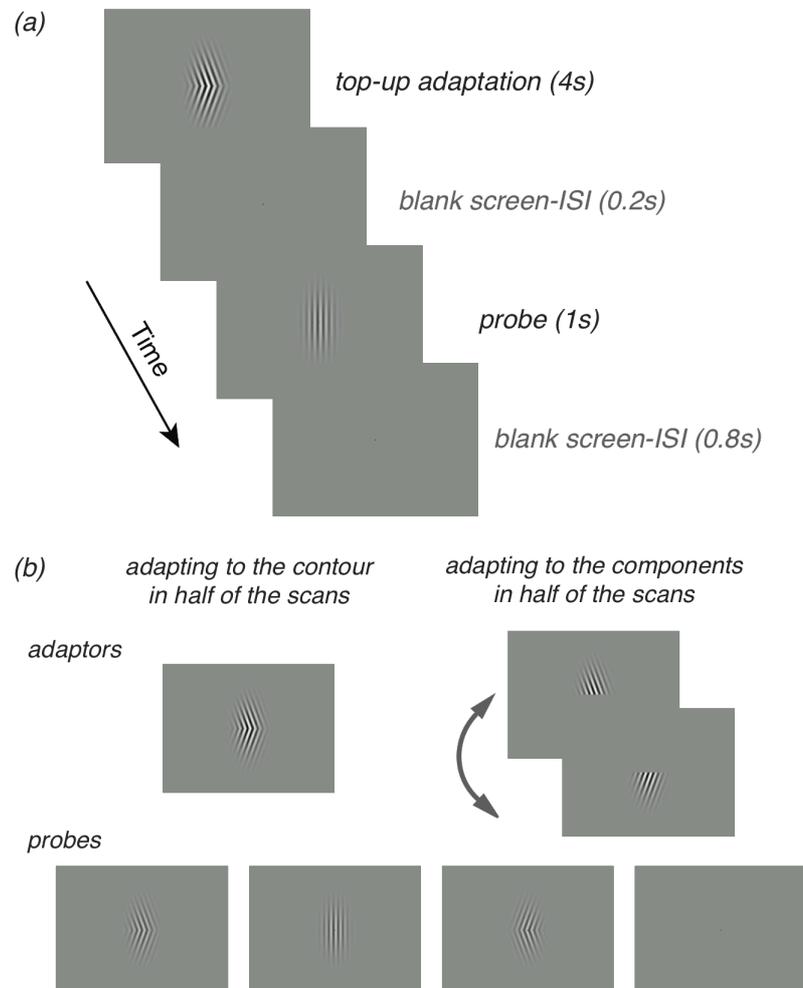
**Localiser scan.** At the beginning of each session a localiser scan was performed to identify voxels in the ROIs that responded to visual stimulation at the spatial location of the patterns. The stimulus was a left bent contour presented at

---

<sup>1</sup> The intensity of probe stimuli was chosen after a series of psychophysical experiments within the lab.

## 5. Integration of simple contours

100% Michelson contrast. Stimulus 'onset' was alternated with a blank screen every 15s. The localiser scan consisted of 8 such blocks. A fixation point was presented throughout.



*Figure 5.1 (a) Event-related procedure. Each trial consisted of a top-up adaptation period for 4s, followed by a blank screen for 0.2s, succeeded by probe presentation for 1s. Trials ended with presentation of the blank screen for 0.8s. Total duration of each trial was 6s. (b) The adaptors and probes used in the study.*

**Attention control task.** To control for changes in the attentional state of observers, they performed an attention-control task during the adaptation scans. Particularly the fixation was changing to faint red (50% contrast) 1-4 times in each attention control trial. The duration of each trial varied randomly between 7-14s. At

the end of each trial, fixation turned to green for 1s prompting participants to report, within this brief period, the number of changes in the colour of fixation by pressing a button.

### ***Data analysis***

Regions of interest (ROIs) were restricted to include only voxels whose time series correlated with the stimulus epochs of the localiser scan. Specifically, ROIs were restricted only to those voxels with a coherence  $co > 0.3$  and a response phase  $0 < \phi < \pi$ . Responses to individual trials were extracted from the average ROI time course by selecting a 10s window initiating 3s after the onset of the adaptor. The average response to the blank trials (which captures the response to the adaptor alone) was subtracted from each trial. Trials of each type were then averaged and the resulting event-related time courses were adjusted to zero baseline by subtracting the response at the first time point of average response.

### **5.1.2 Results**

Event-related responses to the three different probes are shown in Figure 5.2. This figure shows responses only from the scans in which participants were adapted to the contour stimulus. Data are shown in ten visual areas for two subjects. The degree of selective adaptation in each area can be seen as the difference between the response to the contour probe that had the same orientation as the adaptor (dashed light grey), versus the response to the contour probe that had the opposite orientation to the adaptor (grey), versus the response to the vertical patch probe (dark grey). If an ROI exhibits contour selective adaptation then the response shown in dashed light grey would be suppressed compared to the rest of the responses.

In general, signal modulations were at about 0.5% in most visual areas. Responses in area V7 were noisier due the small number of voxels activated.

## 5. Integration of simple contours

Trends towards contour-selective adaptation effects were evident in some instances but these were not consistent across participants. Overall, no consistent contour-selective adaptation effect was obtained across participants in any of the areas tested.

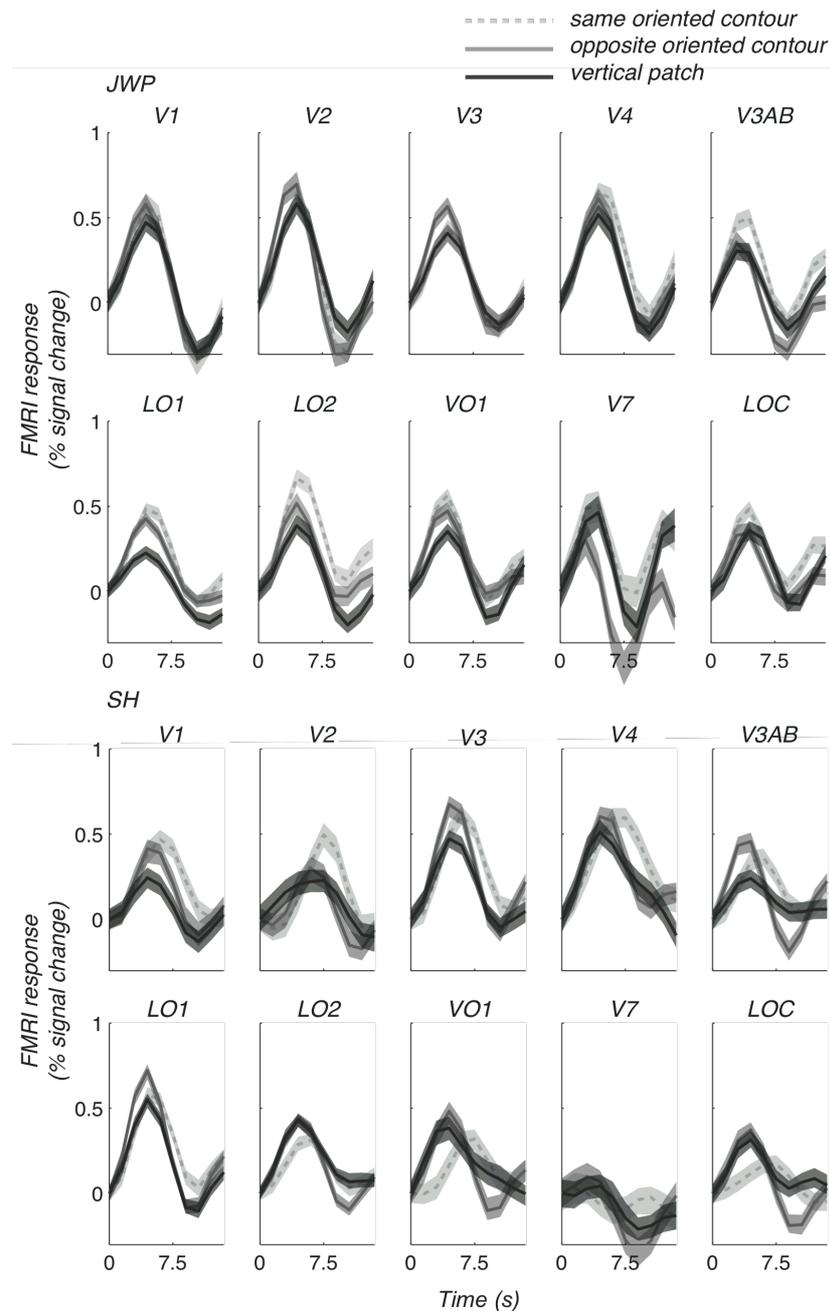


Figure 5.2 Event-related modulations in fMRI signal during the contour adaptation sequence for two subjects (JWP & SH). Time courses were averaged over 40 trials, 10 for each condition. Shaded regions are s.e.m.

## 5. Integration of simple contours

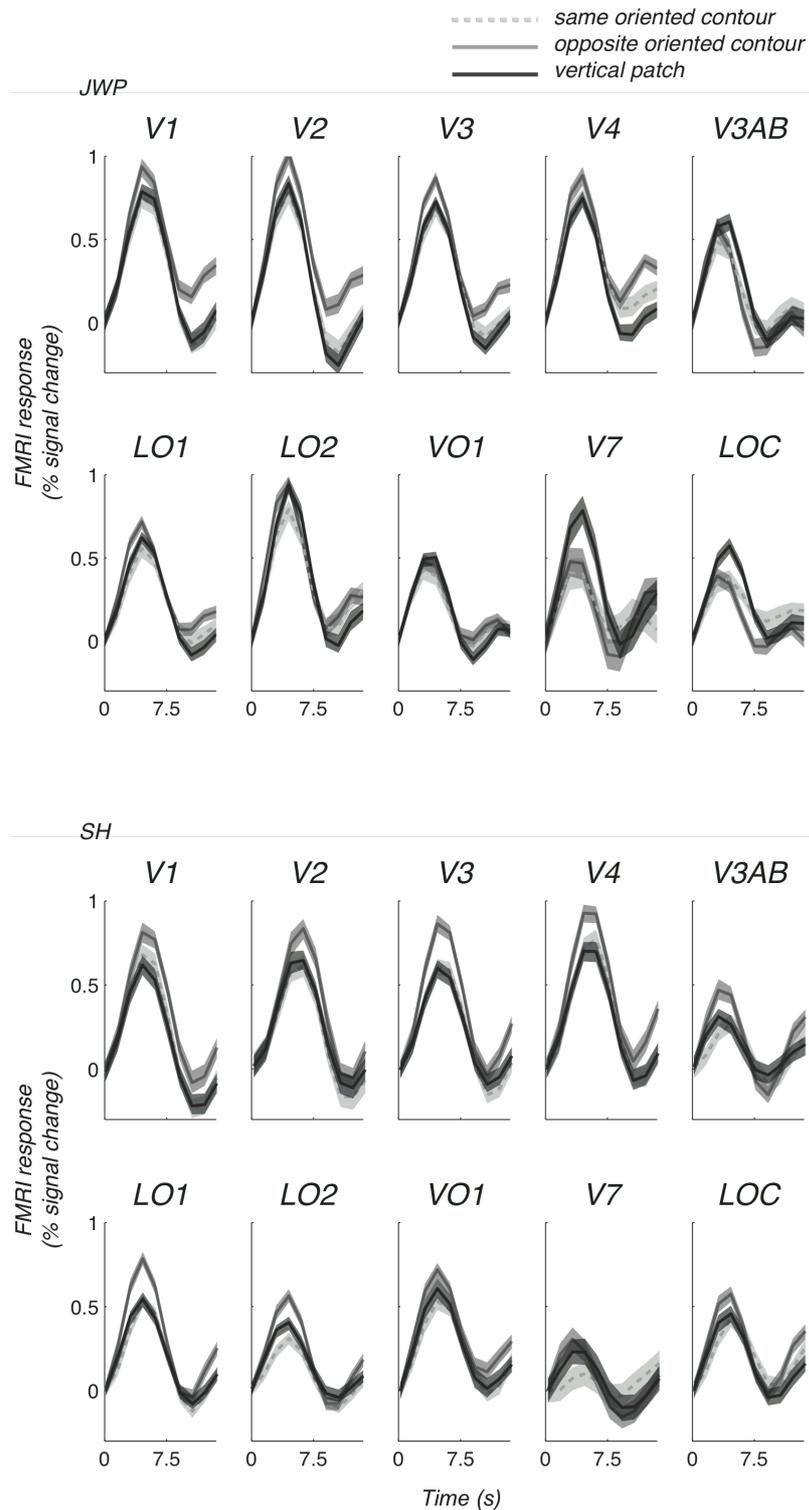


Figure 5.3 Event-related responses during the component adaptation scans for two subjects (JWP, SH). Same conventions as above.

For comparison, Figure 5.3 plots the responses during the component-adaptation scans for both subjects. For one of the participants (JWP) areas LO1

and LO2 showed reduced responses to the contour probe that had the same orientation as the adapting components. The same effect was evident for the second participant (SH) in area LO2. This is rather surprising considering the absence of such an effect in the contour adaptation scan.

### 5.1.3 Discussion

In this experiment we used an event-related adaptation design to study contour selectivity. In general, contour-selective adaptation effects were evident in some instances but these were not consistent across participants. Interestingly, in the component-adaptation scans both participants showed reduced responses to the anticlockwise oriented contour probe in area LO2. This is surprising, considering the absence of such an effect in the contour-adaptation scans. This could, however, reveal some limitations of the experimental design itself. To ensure that exposure to the components was equal in both contour- and component-adaptation scans, the adaptor in the former set of scans was presented for half the adaptation duration (e.g. in the 4s 'top-up' the timing was 1s 'off', 1s 'on', 1s 'off', 1s 'on'). It is, thus, possible that the 2s the 'top-up' adaptor was presented were not enough to retain adaptation.

A second limitation might result from the probe contrast we selected. The reasons we decided to present probes at low contrast are listed in section 2.3.2. It seems, however, that the selection of probe contrast can affect the amount of adaptation in certain visual regions. In the case of orientation selective adaptation Larsson et al. (2006) and us (chapter 3) used low contrast probes (10% Michelson) and found a robust adaptation effect in V1. Adaptation in early extrastriate areas (V2, V3, V4) was about the same as in V1. In a different study, Fang et al. (2005) used high contrast probes (100% Michelson) with all other aspects being similar to Larsson's et al. (2006) and ours. They found stronger orientation selective adaptation in extrastriate regions (e.g. V4) than in V1. It

seems, therefore, that a choice of high contrast probe could enhance the adaptation effect in higher areas. For the reasons discussed in section 2.3.2 we do not think that such a choice is appropriate for studying orientation selectivity. It might be the case, though, that for contour stimuli high contrast probes could increase the adaptation effect in extrastriate areas.

### **5.2 Experiment 2**

In the second experiment we used multi-variate pattern classification analysis (MVPA) to examine which areas can better discriminate between components and contour stimuli. To measure responses for MVPA, we used a blocked design. As in Experiment 1, to control for the potential problem of length-tuned neurons affecting the MVPA accuracy, we ran a number of scans in which stimuli were vertical patches of the same length as the contours (Figure 5.4b). The overall experiment relies on the assumption that contour-selective areas would exhibit higher classification accuracy. Furthermore, accuracy should be higher for the contour vs. components condition than for the control condition.

#### **5.2.1 Methods**

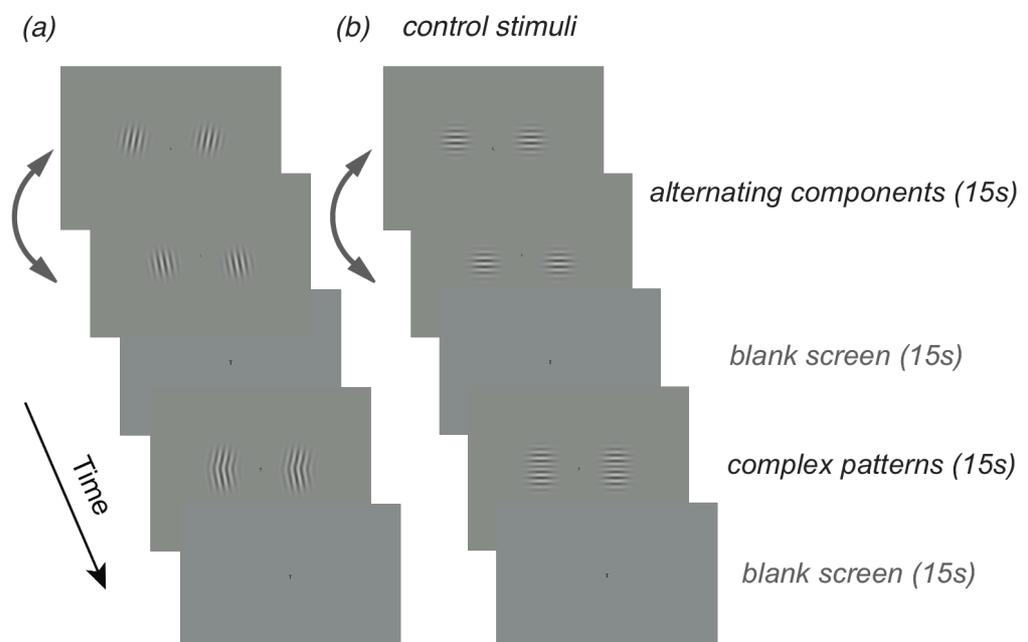
##### ***Participants, procedure and stimuli***

Two experienced volunteers took part in the study. The blocked design used is shown in Figure 5.4. Stimuli were presented for 15s alternated with blank screen for another 15s. In the 'components' block, stimuli were alternating every 0.75s. In the 'contour' block, stimuli were alternating with blank screen every 0.75s. By presenting the contour for half the block duration exposure to components was constant between the blocks.

**Scanning session.** Each scanning session consisted of 8 blocked scans; 4 contour vs. components and 4 control scans. These were interleaved. Each scan consisted of 10 blocks.

**Visual stimuli.** The spatial frequency of all stimuli was 1.5 cycles/°. Contours were made of the conjunction of two oriented gratings orientated at  $\pm 10^\circ$  for the first contour and at  $\pm 170^\circ$  for the second. Stimuli in the control scans were oriented at  $90^\circ$ . Each component occupied  $5^\circ$  degrees of visual angle in diameter. The edges of the stimuli were smoothed by a Gaussian function. The spatial phase of the stimuli was randomised every 6 frames (100ms) to prevent retinal afterimages.

**Attention control task.** A letter counting task at fixation (see Section 2.3.4) was used to control for changes in the attentional state of observers.



*Figure 5.4 A block design was used to measure responses for pattern classification. Alternating components and contours were presented sequentially in different blocks. Contour stimuli were shown for half the block duration to ensure that exposure to components was the same in both conditions.*

### **Data analysis**

To select voxels in the regions of interest (ROIs), we averaged time series across all functional scans. Next, ROIs were restricted only to those voxels with a coherence  $co > 0.25$  and a response phase  $0 < \phi < \pi$  for both hemispheres.

**Pattern-classification.** For each condition, we obtained data in 32 blocks (4 scans per condition, 8 blocks per scan). From each block we extracted the responses at 10 separate time-points (over 15s), delayed by three TR to account for the haemodynamic lag. There were therefore a total of 320 time-points (or repeated 'examples' of each response) for each voxel. Trials were divided into 16 groups, each group corresponding to responses collected in a pair of blocks. Data from 15 of these groups were assigned to a training set and the remaining to a test set to protect against overfitting. This was repeated with different blocks acting as the test set and the performances on each repeat averaged (cross-validation).

### **5.2.2 Results**

Figure 5.5 plots the performance of the pattern classifier based on linear discriminant analysis for two subjects. Classification accuracy is plotted against the number of voxels included in the analysis. MVPA performance is shown for the two conditions; the contour vs. components condition (light grey) and the vertical patch vs. components condition (dark grey). Data are presented in six areas. In higher regions only a small number of voxels were activated (less than 50), making it hard to draw safe conclusions. Classification accuracy was high (60-70% correct) in most areas tested. Performance increased monotonically as more voxels were included in the analysis and reached an asymptote after 40-50 voxels. In the second participant (SH) less than 100 voxels were activated in areas V4, V3AB, and LO1.

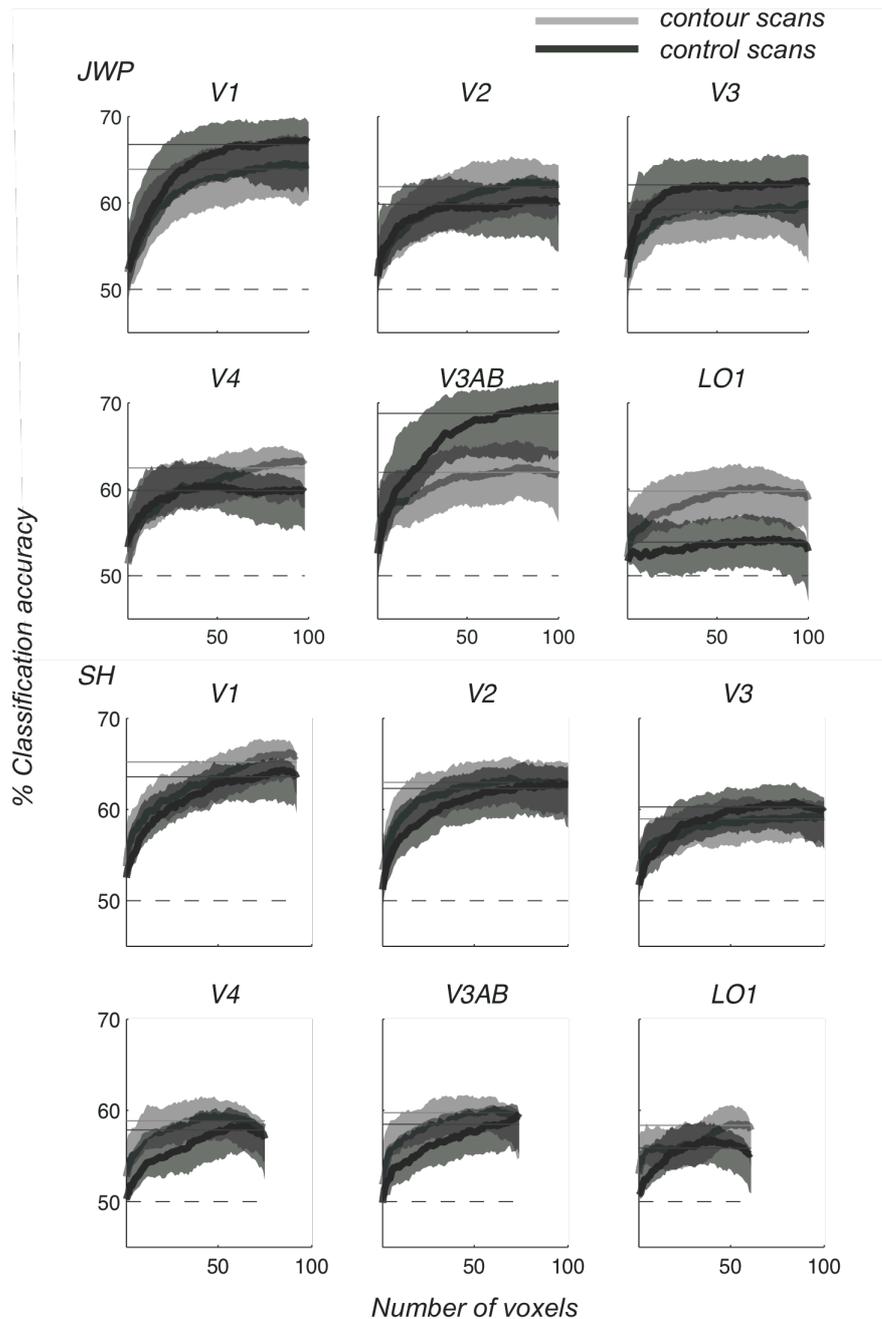


Figure 5.5 Performance of pattern classification versus number of voxels included in the analysis for two subjects (JWP & SH). The dashed line shows predicted accuracy based on chance (50%). The solid lines indicate the index used to estimate classification performance.

The pattern classification algorithm discriminated successfully between compound stimuli and components in the areas tested. MVPA accuracy was somewhat higher in V1 than in extrastriate regions. However, no significant

difference between the two conditions was obtained, with the exception of area LO1 in one of the participants (JWP). For the second participant (SH), only a small number of voxels were activated in LO1 making it hard to draw conclusions. In general, the results suggested that the cortex could successfully discriminate between the components and the complex stimuli but without revealing any particular preference for the curved stimuli.

### **5.2.3 Discussion**

We sought to measure contour selectivity in the visual cortex by means of multi-variate linear discriminant analysis. We obtained high classification accuracy (60-70% correct) in most areas tested but no difference was obtained between the two conditions, with the exception of area LO1 in one of the participants. The results suggest that the cortex could successfully discriminate between the compound and the components; however, the interpretation of the results is not straightforward. One possibility is that the significant classification accuracy we obtained in higher regions is due to integration cortical mechanisms, which are selective to the whole pattern. A different possibility, however, is that the high accuracy does not arise from pattern-integration mechanisms but from a mechanism selective to stimulus length. The data collected so far do not allow us to conclude with certainty on the nature of the mechanisms.

### **5.3 General discussion**

We performed two different experiments to study contour integration mechanisms in the cortex. In the first we used adaptation and in the second multi-variate pattern classification analysis (MVPA). The former design failed to reveal a contour adaptation effect. The latter showed high classification in most areas, however, the results are hard to interpret.

Why did the first experiment fail to reveal an fMRI adaptation effect despite consistent evidence from psychophysical studies (Hancock and Peirce 2008)? One possibility is that the probe contrast we selected was not appropriate for this study. As discussed above, the choice of probe contrast can affect the amount of adaptation in certain visual regions (e.g. Fang et al. 2005; Larsson et al. 2006). It is possible that for contour stimuli, high contrast probes could increase the adaptation effect in extrastriate areas. To test this, it would be of interest to collect adaptation data using high contrast probes (see discussion in section 5.1.3) and maybe a larger TR. An increase of the TR from 1.5s to 2s could increase the signal by about 10%, making the measurement more sensitive. It is possible that this experiment could capture contour adaptation more efficiently.

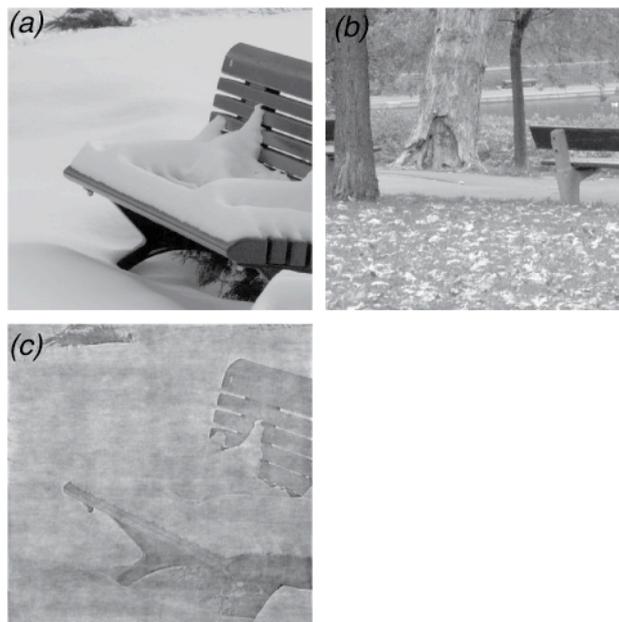
In addition, the design itself provides certain limitations. For example, event-related designs can generate strong perceptual effects but the measured signal is weak because of the brief, low contrast probes. Another possibility is that in conventional ROI analyses, where time-courses are averaged together across voxels to obtain a mean ROI response, subtle signal modulations in voxels that do adapt are lost.

Indeed, as in the previous chapters, the MVPA was more sensitive than adaptation. Classification accuracy was high in all visual areas included in the analysis. As a future study, we have redesigned the MVPA paradigm so that it generates results that are more straightforward to interpret. This new set of experiments is an extension of experiments described above. We hope that these new studies will provide insight into the nature of contour integration mechanisms in the cortex.

## ***6 Spatial phase encoding in the cortex***

---

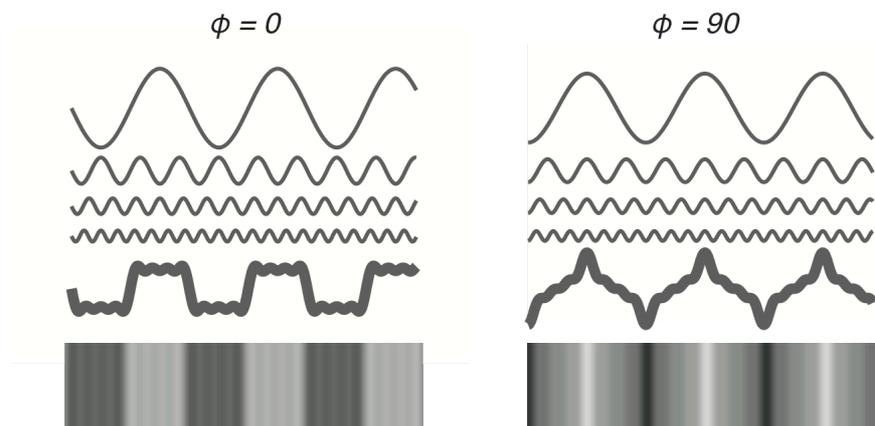
Evidence has been accumulating over the past years to support the view that visual stimuli are detected by a battery of independent mechanisms, each responsive to a narrow range of spatial frequencies and orientations. However, to fully describe a complex waveform, not only the frequency and amplitude must be specified, but also the phase of each component. Spatial phase is a fundamental aspect of spatial vision, crucial both for the extraction of local features (Burr et al. 1989) and overall scene perception. This is because, natural visual stimuli are composed of various spatial frequencies added together in a specific phase relationship (see also Figure 6.1). Therefore, for the perception of such complex stimuli, it becomes important to investigate how the visual system processes phase information.



*Figure 6.1 The phase information prevails. The hybrid image in panel (c) was composed from the Fourier phase of image (a) and the Fourier amplitude of image (b). As it can be seen it is the Fourier phase that carries most visual information.*

A straightforward way to study spatial phase is to combine two or more sinusoidal gratings of a particular relative phase into a compound waveform. Figure 6.2 illustrates such an instance, where four components are added in sine or cosine phase ( $\phi=0$  or  $90$ ). The relative phase components are added in, affects the appearance of the compound stimulus. When components are added in sine phase the resulting waveform appears edge-like, whereas, when added in cosine phase it appears line-like.

Many of the early studies of spatial frequency analysis by the visual system ignored spatial phase, or showed that phase was irrelevant in certain tasks (Graham and Nachmias 1971). However, further psychophysical (Badcock 1984a; b; Burr 1980; Burr et al. 1989; Huang et al. 2006) and physiological (Aronov et al. 2003; De Valois and Tootell 1983; Mechler et al. 2002; Pollen et al. 1988) studies provided evidence that the visual system is sensitive to the spatial phase of patterns under certain conditions.



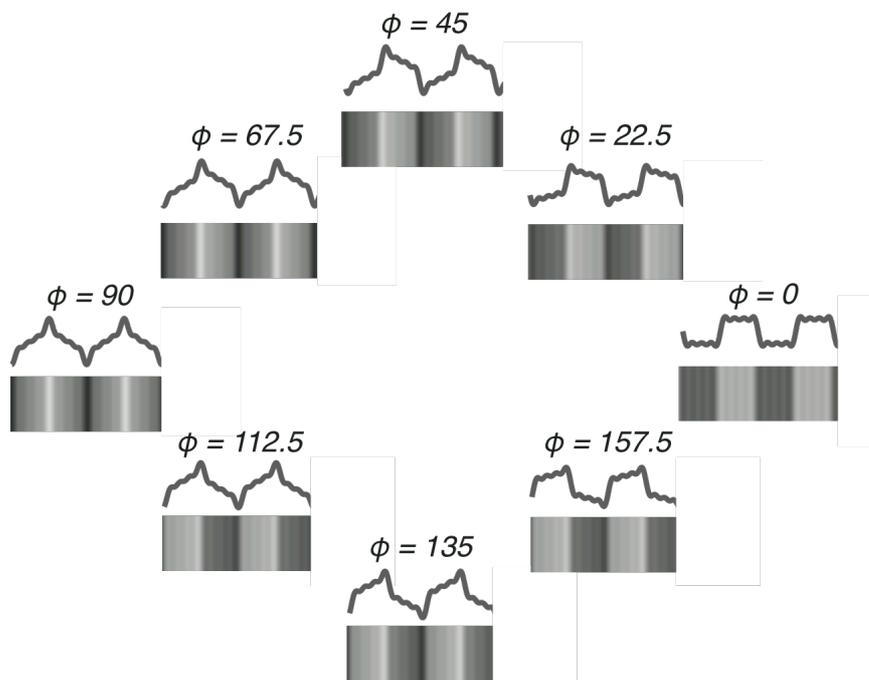
*Figure 6.2 Depending on their relative phase the sum of the same four component gratings can give rise to very different spatial compound waveforms. On the left, the components are combined in  $\phi=0$ . The harmonic components coincide at their position of maximal slopes, leading to a periodic sequence of on- and off-edges approximating a square wave. On the right, components are combined in  $\phi=90$ . The harmonic components coincide at their peaks, leading to a waveform of alternating bright and dark lines. Adapted from Mechler et al. (2002).*

For example, De Valois and Tootell (1983) used compound stimuli made of two components, added in different relative phases. For most of the complex cells studied the relative phase of the components did not affect their responses. However, a proportion (less than the half) of simple cells did respond to the relative phase of two different frequencies. Interestingly, sensitivity to relative phase was high for nearby frequencies ( $f + 2f$  or  $f + 3f$ ) but low for combinations of more widely spaced frequencies (e.g.  $f + 4f$ ) (see also Stromeyer et al. 1973). Psychophysical studies showed that phase discrimination ability is lost at high spatial frequencies ( $>15\text{cycles/}^\circ$ ) (Holt and Ross 1980). Huang et al. (2006) found no evidence of cardinal phase ( $\varphi=0, 90, 180, 270$ ) detectors in the human visual system. This was in contrast to previous studies suggesting that there exist visual detectors tuned to one of four cardinal phases (e.g. Burr et al. 1989).

Mechler et al. (2002) used compound waveforms made of 4 components with only odd harmonics being present. Their stimuli were similar to those used in our study. The responses of most V1 cells tested (they recorded responses from 137 neurons) were rather noisy and displayed low sensitivity to relative phase. However, a minority of neurons that included both simple and complex cells were highly sensitive and selective to spatial phase. Although the distribution of the preferred congruence phase in V1 was broad, with all phases represented, the population as a whole, regardless of cell class, displayed a slight bias toward lines and possibly edges. Based on this weak bias the authors suggested the existence of two classes of cortical feature detectors, one tuned to edge like and the other to line-like waveforms. Interestingly, feature preference and selectivity varied within local clusters of V1 neurons.

In the experiment discussed in this chapter we investigated spatial phase encoding in the human visual cortex. The results we obtained in Chapter 3 indicated higher sensitivity for the multi-variate pattern analysis (MVPA) in

detecting similar stimuli. Considering the small perceptual differences between stimuli created using similar relative phases, this chapter used multi-variate pattern classification analysis (MVPA) to discriminate between relative phase combinations. We created eight compound gratings (Figure 6.3). These consisted of four superimposed components ( $f + 3f + 5f + 7f$ , where  $f$  is the fundamental frequency), with relative phase,  $\phi$ , varying in eight equal steps around the phase circle  $[0, 180)$ . Individual components were added in decreasing contrasts ( $C_{3f}=1/3C_f$ ,  $C_{5f}=1/5C_f$ ,  $C_{7f}=1/7C_f$ , where  $C_{nf}$  is the contrast of the  $n$ th component). Compound waveforms created in this way have identical root mean square (RMS) contrasts, thus, equal overall energies.



*Figure 6.3 Luminance profiles of the eight, equal-energy compound gratings used in this study. The relative phase,  $\phi$ , was varied in eight equal steps around the phase circle  $[0, 180)$ . The spatial waveform of the compound gratings varied with  $\phi$ , from edge-like ( $\phi=0$ ) through line-like ( $\phi=90$ ) back to edge-like. Stimuli in the  $[180, 360)$  phase interval duplicate those in the  $[0, 180)$  phase interval, with a half-cycle shift in the compound waveforms. Therefore, stimuli in the  $[180, 360)$  interval were not included. Adapted from Mechler et al. (2002).*

### 6.1 Methods

#### ***Participants, procedure and stimuli***

Three experienced volunteers took part in the study. A block design was used to measure the performance of the MVPA (Figure 6.4). Compound stimuli were presented for 12s alternated with blank screen for another 12s. The eight compound stimuli were presented in random order.

**Scanning session.** Each scanning session consisted of 6 blocked scans; sixteen blocks per scan.

**Visual stimuli.** Compound stimuli consisted of four superimposed components; the fundamental ( $1.5\text{cycles}/^\circ$ ) and the first three odd harmonics. The intensity values of the compounds were scaled so that the Michelson contrast of the line-like waveform ( $\varphi=0$ ) was  $C_M=100\%$ , and smaller for the rest. Critically, the RMS contrast of all stimuli was the same ( $C_{RMS}=23\%$ ). Compound gratings were presented in an annulus (inner radius,  $2^\circ$ ; outer radius,  $8^\circ$ ) whose edges were smoothed by a Gaussian kernel (std dev of  $0.083^\circ$  on the inner edge,  $0.333^\circ$  on the outer edge). The spatial phase of the gratings was randomised every 6 frames (100ms) to prevent retinal afterimages.

**Attention control task.** To control for changes in the attentional state of observers a letter counting task was performed at fixation (see section 2.3.4).

#### ***Data analysis***

To select voxels in regions of interest (ROIs), time series were averaged across all functional scans. Next, ROIs were restricted only to those voxels with a coherence  $c_0 > 0.3$  and a response phase  $0 < \varphi < \pi$  in both hemispheres.

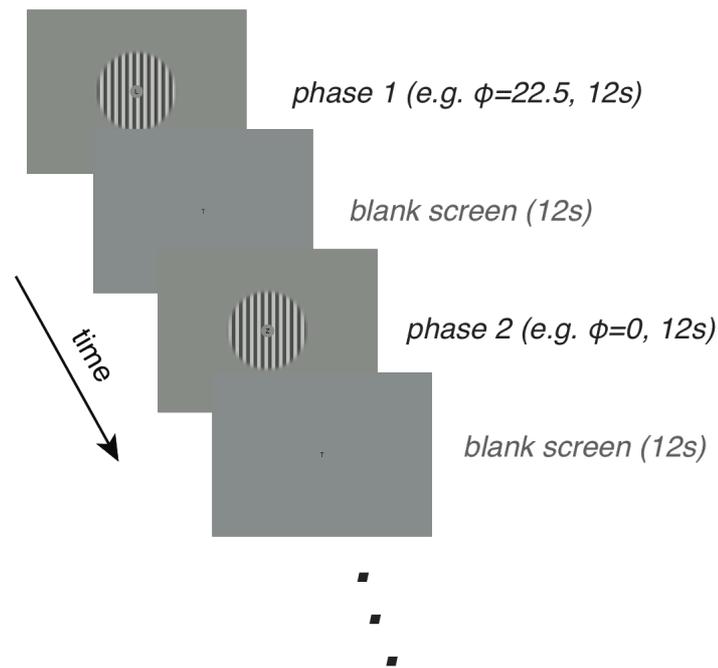


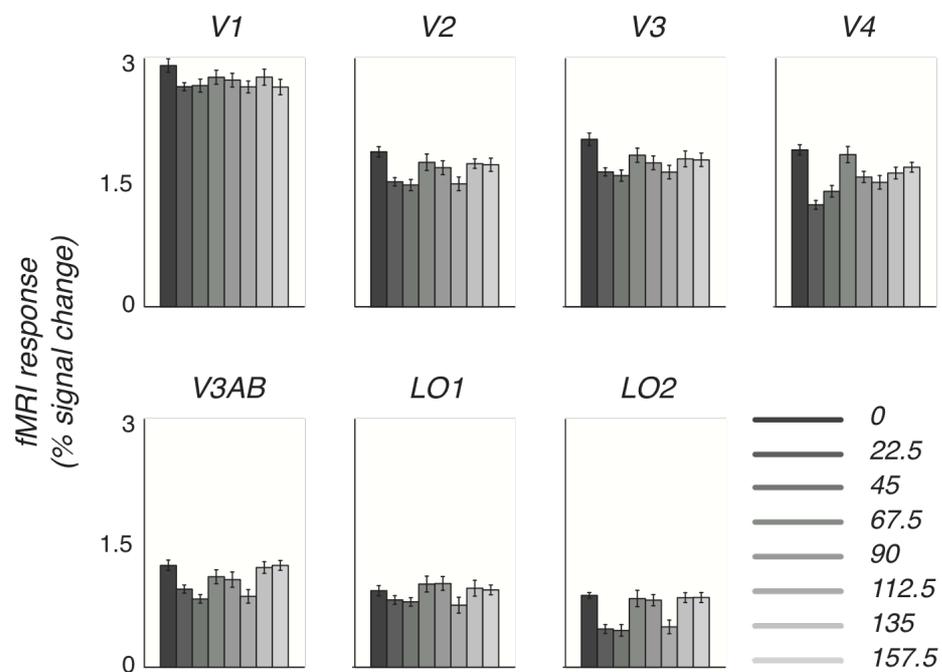
Figure 6.4 The blocked design used to measure performance of the MVPA.

**Pattern-classification.** A linear neuronal network algorithm (without a hidden layer) was used to classify responses between stimulus conditions (see General Methods). Because of the somewhat low accuracies obtained in this section we decided to use a non-parametric classifier. In each scanning session, we obtained data in 96 blocks (6 blocked scans, 16 blocks per scan) for all eight waveforms (12 blocks for each stimulus). From each block we extracted the responses at 8 separate time-points (over 12s), delayed by 3 TR to account for the haemodynamic lag. Trials were divided into 12 groups, each group corresponding to responses collected in a pair of blocks. Data from 11 of these groups were assigned to a training set and the remaining to a test set to protect against overfitting. This was repeated with different blocks acting as the test set and the performances on each repeat averaged (cross-validation).

## 6.2 Results

### *Overall responses to compound waveforms*

First, we computed the actual BOLD response elicited by different compound stimuli in each visual area. If the visual system had more neurons responding to 0 and 90 congruent phases, the responses to these particular waveforms should be higher than to the rest. The maximum values of the mean ROI responses to the eight different stimuli are shown in Figure 6.5, for the three participants.



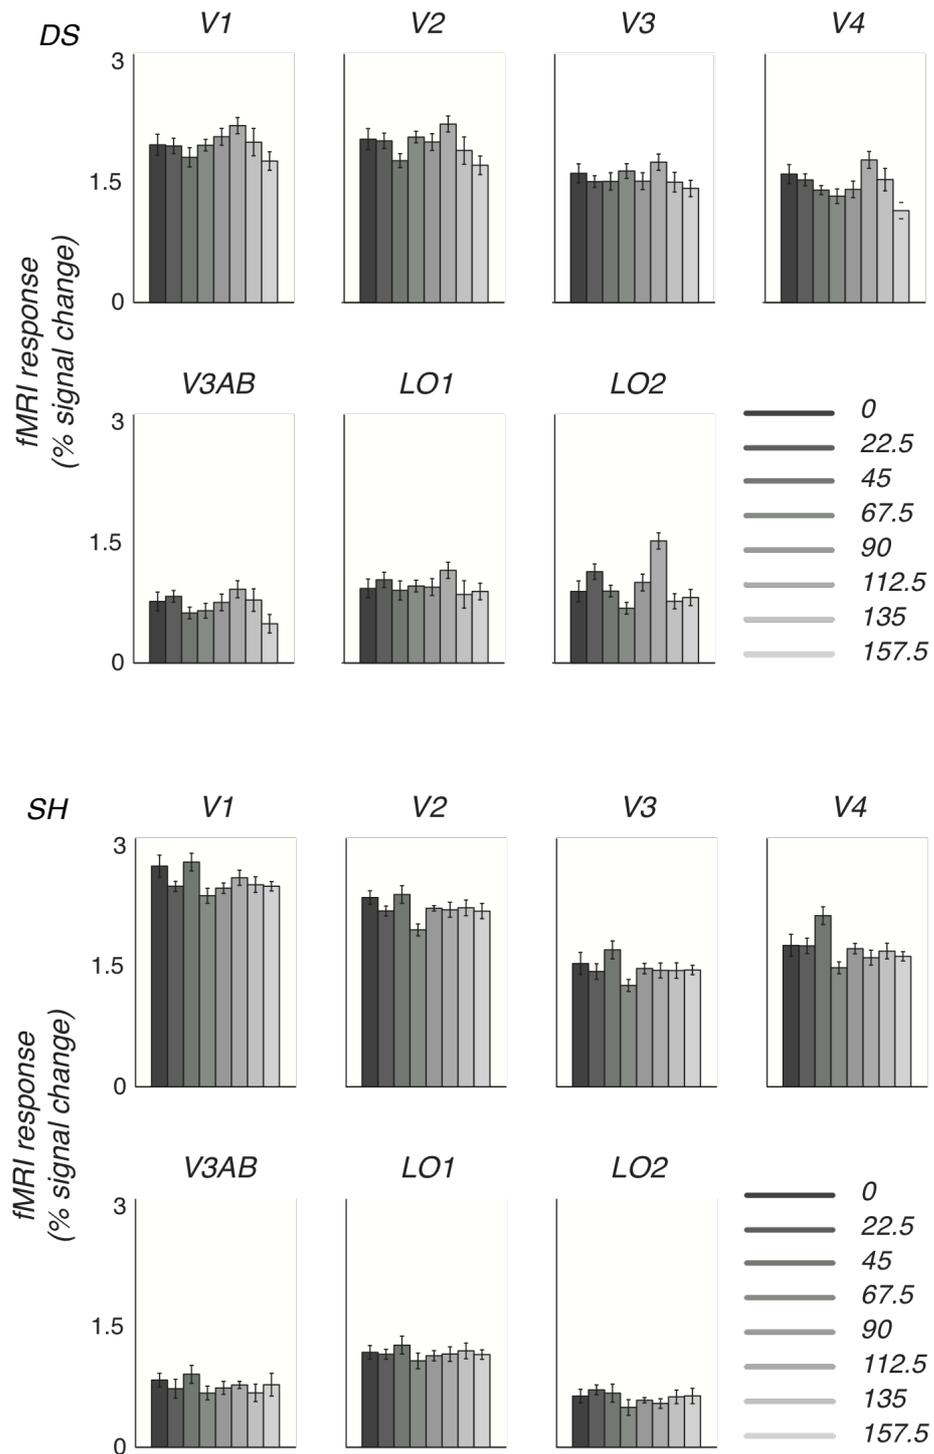


Figure 6.5 fMRI responses to the eight compound stimuli used in the study. Data are shown for three participants (JWP, DS, SH). Time courses, in each ROI, were adjusted to zero baseline by subtracting the first time point. Error bars represent  $\pm 1$  s.e.m.

Interestingly, the prevailing phase was different in different observers. Participant JWP (first panel) responded more to the 0 phase compound. In this case, responses to intermediate congruence phases (22.5, 45 and 112.5) appear to be lower, resulting in a more symmetric pattern across the phase dimension. Participants DS and SH (second and third panel) responded more to the 112.5 and 45 phase waveforms respectively.

### ***MVPA performance***

We computed classification accuracy based on a neuronal network classifier for edge- and line-like ( $\phi=0, 90$ ) compound stimuli. The results are shown in Figure 6.6 for all three participants. MVPA accuracy was high in area LO1 for two of the observers. The results for subject SH were somewhat variable. Interestingly, MVPA accuracy in area V1 of this subject was substantially below chance levels.

Next, we computed classification accuracy for all phase combinations. Figure 6.7 summarises the results. MVPA accuracies below 52% are shown in black, whereas, higher accuracies appear progressively brighter. Interestingly, classification accuracy appears to be higher in certain extrastriate areas. These include V2, V3, V3AB and LO1. However, the pattern of accuracies was somewhat variable between participants.

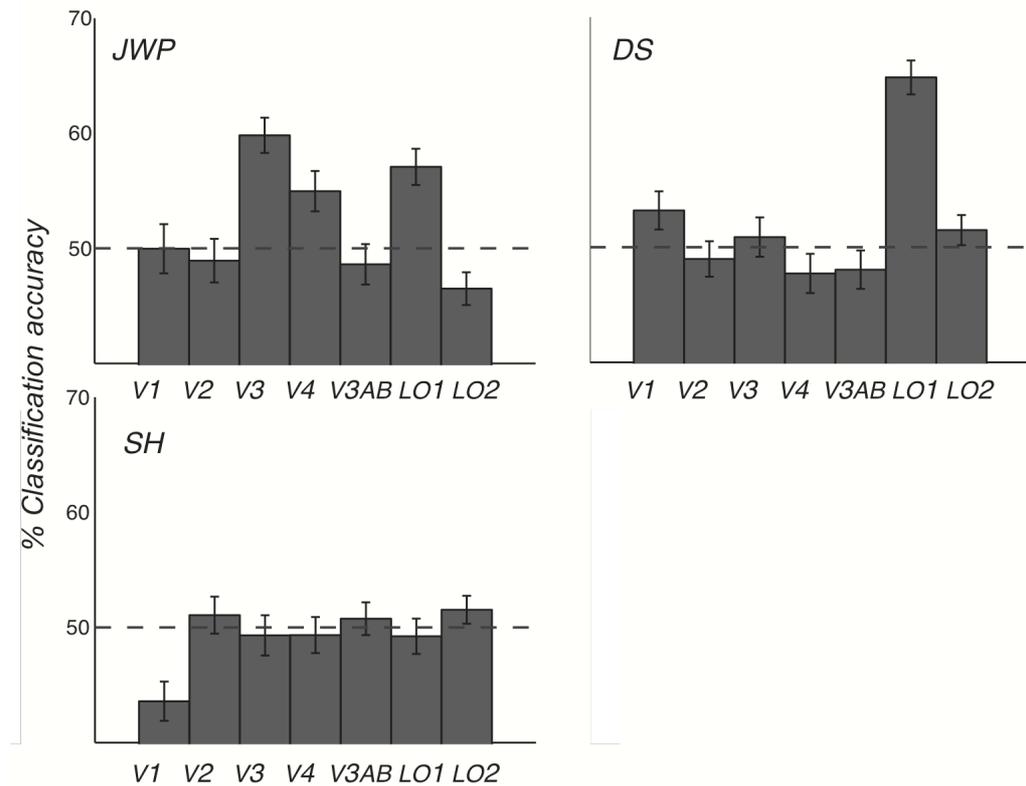
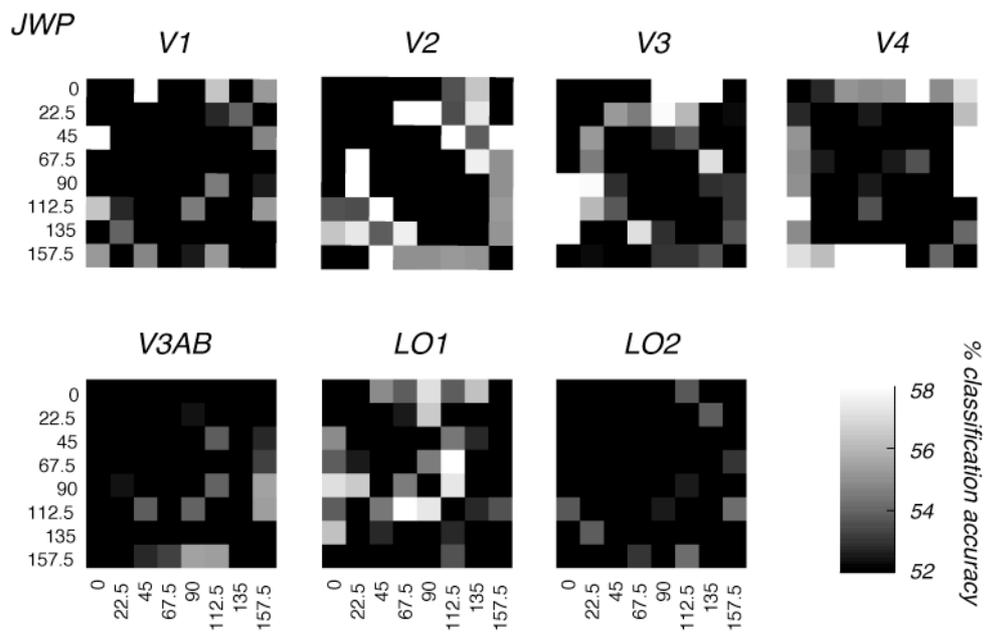


Figure 6.6 MVPA accuracy based on a neuronal network classifier for three observers. Classification was computed directly on the 100 voxels with the highest main effect in each area. The order of voxels was shuffled 100 times and MVPA accuracy was computed for each shuffle. Error bars are standard deviations computed over the 100 reshuffles.



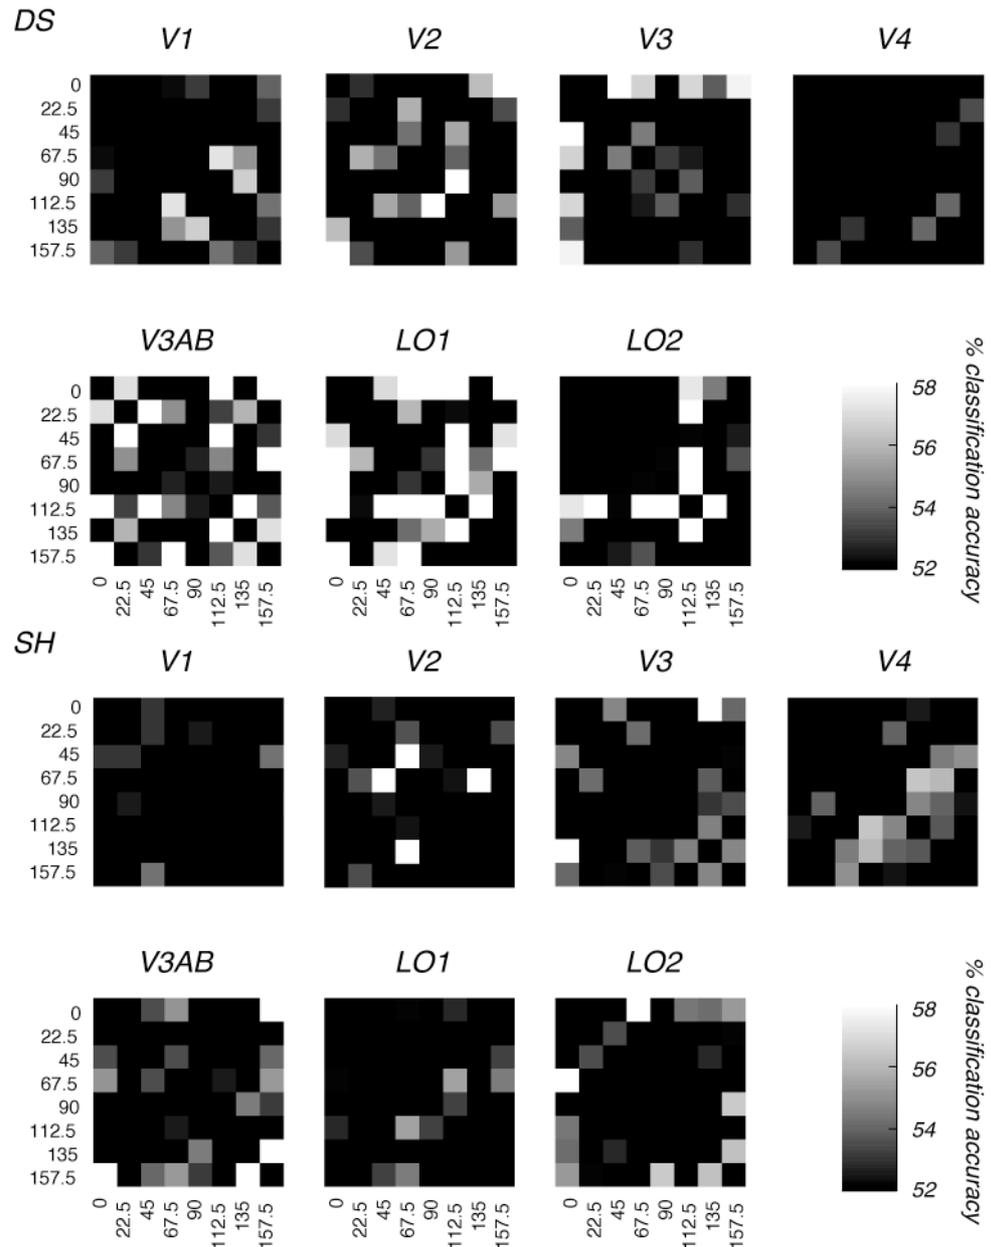


Figure 6.7 MVPA accuracy computed for all relative phase combinations. Brighter colours indicate a higher classification accuracy for the particular pair of phases.

### 6.3 Discussion

This chapter investigated the encoding of spatial phase in the cortex. Phase information is essential for perception, because important visual features like edges emerge at locations of maximal local phase coherence. It is also possible that phase congruency might provide a cue for the detection of object boundaries. Detection of relative phase requires integration of spatial frequency information

across multiple spatial scales. The pooling over different spatial frequency bands has been shown to emerge from the statistical properties of natural images. This is potentially of great importance because it has been suggested that the human visual system has evolved to optimally process stimuli from the natural environment (Simoncelli and Olshausen 2001).

Previously, spatial phase was studied psychophysically (Badcock 1984a; b; Burr 1980; Burr et al. 1989), or with electrophysiology (Aronov et al. 2003; De Valois and Tootell 1983; Mechler et al. 2002; Pollen et al. 1988). This is one of the first studies (see also Henriksson et al. 2009) that looked into phase encoding in the human cortex using fMRI. We created compound gratings from four superimposed sinusoidal components, added in different relative phase,  $\phi$ . Relative phase was varied in eight equal steps around the phase circle (0, 22.5, 45, 67.5, 90, 112.5, 135, 157.5).

First, we computed the mean response to individual stimuli in each ROI. Interestingly, different participants responded more strongly to different relative phase combinations. This phase preference in individual participants was observed in all visual areas included in the analysis. Our results are in agreement with those of Mechler et al. (2002) who found a broad distribution of relative phase the macaque V1. However, our results do not confirm their observation that there is a slight bias towards line-like ( $\phi=90$ ) and possibly edge-like ( $\phi=0$ ) stimuli. Such a bias was observed in one of our participants (JWP); however, different participants were biased towards different phases.

Next, we used an MVPA neuronal network algorithm to classify responses between pairs of different relative phase. Spatial phase stimuli were successfully discriminated in a number of visual areas. Although V1 could discriminate some phase combinations, relative phase was more efficiently discriminated in higher regions. However, the pattern of accuracies was somehow variable between

participants. Our results are in agreement with electrophysiology studies indicating that relative phase is coded as early as in V1 (De Valois and Tootell 1983; Mechler et al. 2002), but point to the direction that phase might be more efficiently represented in extrastriate areas.

All in all, our study showed that several visual areas, including the primary visual cortex, are sensitive to the difference of phases across spatial frequency bands. However, phase coherence might be optimally encoded in extrastriate areas. These findings shed more light on the encoding of phase congruence in the cortex and potentially the perception of edges and object boundaries in natural stimuli.

## ***7 General discussion***

---

Functional magnetic resonance imaging (fMRI) has become a ubiquitous tool in cognitive neuroscience. Despite its limitations, it is currently the best technique we have for studying brain mechanisms, ranging from low-level visual processes, to high-level tasks such as recognition and memory. Considering the articles speculating on potential applications, it is believed that fMRI has a bright future in basic neuroscience research. However, fMRI is not and will probably never be a mind reader, as some of the proponents of the technique suggest, nor is it a worthless and non-informative 'neo-phrenology' that is condemned to fail, as it has been occasionally argued.

The main advantage of the technique is that it allows the non-invasive measurements of cortical responses in the human brain, but only on the millimeter scale. Because a typical voxel contains many thousands of neurons with varied properties, establishing the selectivity of their responses directly is impossible. In recent years, two methods using fMRI aimed at studying the selectivity of neuronal populations on a 'subvoxel' scale have been heavily used. The first technique, fMRI adaptation, relies on the observation that the fMRI response in a given voxel is reduced after prolonged presentation of a stimulus, and that this reduction is selective to the characteristics of the repeated stimuli. The second technique, multi-variate pattern analysis (MVPA), makes use of multi-variate statistics to recover small biases in individual voxels in their responses to different stimuli. It is thought that these biases arise due to the uneven distribution of neurons (with different properties) sampled by the many voxels in the imaged volume.

Although in recent years these two techniques have been heavily used, they have never been compared explicitly, and little is known about their relative sensitivities. One major objective of this thesis was to compare the two techniques and evaluate their relative sensitivities. A second objective was to use these two techniques to study early- and mid-level processing in the human visual cortex.

The first experiment (Chapter 3) investigated whether adaptation and multi-variate pattern classification analyses (MVPA) provide consistent results about the properties of the cortical areas under study. To address this question, this thesis compared the two methods for their ability to detect the well-documented orientation selectivity in early visual cortex.

Both methods were clearly capable of revealing orientation selectivity in early visual areas (V1, V2, V3). For the MVPA method this has been shown previously by several studies (Haynes and Rees 2005; Kamitani and Tong 2005). For adaptation a number of groups have shown that orientation selectivity can be demonstrated but the pattern of data has been more variable, and seems critically dependent on the choice of experimental parameters (Boynton et al. 1996; Fang et al. 2005; Larsson et al. 2006).

This thesis found less selectivity to orientation in later visual areas (V3AB, LO1, LO2 and VO1), using both MVPA and adaptation methods. The correlation between the results across visual areas was high indicating that the two methods are in strong agreement, at least in the domain of orientation specificity. This may not have been the case and increases our confidence in both methods. It also potentially informs our understanding of orientation selectivity in the areas studied. Electrophysiology studies have demonstrated that orientation selectivity is a common feature of early visual areas. This is well-documented for V1 and V2, but has been less explored in V3 (Lennie, 1998). fMRI adaptation experiments had suggested a degree of selectivity, but this may have been simply a result of

adaptation in earlier areas, resulting in reduced input to V3 (Larsson et al., 2006). The fact that we find high MVPA performance, as well as orientation-selective adaptation in this area, increases the confidence that human V3 does indeed code for stimulus orientation.

One exception to the strong correlation between the areas was V4 (although the difference was not statistically significant), which appeared to show relatively strong orientation-selective adaptation but with rather poor MVPA performance in this stimulus domain. There are several reasons why this might be the case. It may result from the spatial organization of V4; for pattern analysis to work there must be some inhomogeneity in the spatial distribution of cells, such that an individual voxel is more or less tuned to one of the stimuli presented. Therefore, it may be that V4 has orientation-selective neurons, in agreement with electrophysiological studies in the macaque (Desimone and Schein 1987; Maunsell et al. 1991), but that these are not clustered in the same way as in V1, V2 and V3. This case is further supported by Vanduffel et al. (2002) who found orientation-selective neurons in the macaque V4 to be clustered irrespective of their preferred orientation.

In order to determine how each technique depended on stimulus orientation, the above experiment was performed with smaller orientation differences. MVPA performance remained above chance for all pairs of orientations tested and was remarkably consistent between participants. Adaptation results were in agreement with previous data (Fang et al. 2005). Selective adaptation failed to distinguish stimuli with smaller separations; for a 50° separation it failed in one of the three individuals, for a 25° separation it failed in all participants.

It should be noted that these results might not be mirrored in other domains of visual selectivity. For instance, measuring the degree of selectivity to spatial

frequency, direction of motion or faces may give very different results if the neurons that code these dimensions in a particular area adapt strongly, but are only weakly clustered. Clearly however, in the case of orientation selectivity measurements, the multi-variate pattern analysis was rather more sensitive than the selective adaptation measure, although the two methods were in close agreement for most visual areas.

Having directly compared fMRI adaptation and MVPA in Chapter 3, Chapter 4 used both techniques to investigate how local orientation signals are combined and detected in intermediate levels of visual processing. To study this the thesis used plaid stimuli, made of the linear sum of two sinusoidal gratings. Although plaids have been more typically used to study the combination of motion signals (Huk and Heeger 2002; Movshon et al. 1986; Rust et al. 2006) they have also been used to study the conjunction of form signals (Carandini et al. 1997a; Georgeson and Meese 1997; Peirce and Taylor 2006). Peirce and Taylor (2006) presented psychophysical evidence that the visual system has mechanisms sensitive to the overall form of compound patterns rather than to their components. This thesis built on this idea by using fMRI adaptation and MVPA to study plaid selectivity in the visual cortex. As in Chapter 3, the MVPA was far more sensitive in revealing plaid selectivity. Although the results we obtained were somehow weak, they pointed to the direction that V2 might play a role in Fourier component integration.

The function of V2 neurons, however, is not very clear despite the prominent location of V2 early in the visual hierarchy. It is known that V2 is essential for vision. While a lesion in V2 does not affect visual acuity and contrast sensitivity, it strongly affects a monkey's ability to perform more complex spatial tasks (Merigan et al. 1993). But up to recently no stimulus was found that excites these cells in a way that reveals their obvious contribution to vision. Perhaps this

is why V2 remains relatively unstudied compared to other regions such as V1 and MT. Recently, Ito and Komatsu (2004) suggested that V2 neurons in the macaque respond to stimuli consisting of angled lines placed in the center of each cell's receptive field. These stimuli resembled more the contours used in Chapter 5. These results point to the direction that one of the roles of V2 might be to combine local orientation signals into more complex features.

Chapter 5 built on these ideas and investigated how the outputs of V1 are organised and combined by the visual system in order to represent more complex visual features using contour stimuli. The experiments were based on the assumption that conjunctions of local orientation signals might be detected in intermediate levels of visual processing in the form of curves or corners. Several investigators have psychophysically studied how features are integrated into visual contours using curved stimuli in various different forms (Field et al. 1993; Hess and Field 1993; Loffler et al. 2003; Watt and Andrews 1982). Recent studies have reported psychophysical shape aftereffects that do appear to have a basis in global shape processing rather than local orientation (Gheorghiu and Kingdom 2007; 2009; 2008; Hancock and Peirce 2008; Suzuki 2001; 2003).

Previous fMRI studies that examined mid-level visual mechanisms have used, for example, contours of co-aligned Gabor elements (Altmann et al. 2003; Kourtzi and Huberle 2005; Kourtzi et al. 2003), stereoscopic shape contours (Kourtzi and Kanwisher 2001), random dot stereograms (Vinberg and Grill-Spector 2008), or glass patterns (Ostwald et al. 2008). In the experiments described in Chapter 5, contour stimuli were constructed from two luminance modulated sinusoidal gratings, with different orientations. We believe that, due to their simplicity, these stimuli are more appropriate for studying conjunctions of local orientation signals. In the first experiment an event-related adaptation design was used. Contour-selective adaptation effects were evident in some instances

but these were not consistent across participants. Why did the experiment fail to reveal contour adaptation despite the consistent effect obtained from the psychophysics (Hancock and Peirce 2008)? The discussion in the corresponding chapter lists several reasons for that. For example, it is known that the choice of probe contrast can affect the amount of adaptation in certain visual regions (e.g. Fang et al. 2005; Larsson et al. 2006). It is, thus, possible that high contrast probes could increase the contour adaptation effect in extrastriate areas.

In the second experiment of Chapter 5, a pattern classification analysis was used to examine which areas can better discriminate between contour stimuli and their components. To control for the potential problem of length-tuned neurons affecting the MVPA accuracy, a control condition was introduced. High classification accuracy was obtained in most areas tested but no difference was obtained between the test and control condition, with the exception of area LO1 in one of the participants. The results suggested that the cortex could successfully discriminate between the compound and the components, but the interpretation of the results was not straightforward.

The results obtained in Chapter 5, helped us reconsider the experimental design and parameters. As mentioned in the corresponding discussion sections both experimental paradigms have been redesigned. The adaptation protocol has been modified to increase its sensitivity and the MVPA paradigm has been altered so that it generates results that are more straightforward to interpret. The new set of experiments is an extension of experiments performed in this thesis. Hopefully, these new studies will provide insight into the nature of contour integration mechanisms in the cortex.

In the final experimental chapter (Chapter 6) this thesis investigated the encoding of spatial phase. Phase information appears to be essential for perception, because important visual features like edges emerge at locations of

maximal local phase coherence. Detection of relative phase requires integration of spatial frequency information across multiple spatial scales. The pooling over different spatial frequency bands has been shown to emerge from the statistical properties of natural images. This is potentially of great importance because it has been suggested that the human visual system has evolved to optimally process stimuli from the natural environment (Simoncelli and Olshausen 2001).

When fMRI responses were considered in a univariate manner, this thesis showed a broad distribution of relative phase in the cortex, without a preference to particular phase congruence across participants. Individual subjects were biased towards certain phase combinations but this bias was different in different participants. Next, we used MVPA analysis to classify responses between pairs of different relative phase. This thesis showed that several visual areas, including the primary visual cortex, were sensitive to relative phase combinations. However, we showed that phase coherence might be optimally encoded in extrastriate areas.

All in all, this thesis clearly showed that it is possible to study the selectivity of neuronal sub-populations on a scale smaller than conventional fMRI resolution. This can be achieved by means of fMRI adaptation or MVPA analyses. We showed, using a variety of stimuli, that the MVPA was more sensitive in detecting neuronal selectivity. MVPA methods have evolved extensively in the last few years and it is expected that they will continue to evolve, as better algorithms become available in the coming years. Improvements in the spatial resolution of fMRI will make it possible to resolve even finer-grained cognitive distinctions. For all of these reasons, we believe that MVPA has a bright future as a tool for characterising how information is represented and processed in the brain.

Regarding the second theme of this thesis, we clearly showed how orientation selectivity is encoded in the cortex using two different, independent

methods. This thesis also tried to investigate how local orientation and spatial frequency signals are combined in later steps of visual processing. We presented evidence that V2, as well as other extrastriate areas such as V3, V3AB and LO1 might play such a role. These experiments were the first step in this direction. The corresponding chapters suggested new experiments, which will extend the findings of this thesis and hopefully will shed more light on mid-level vision processing.

---

## References

---

- Adams DL, Sincich LC, and Horton JC.** Complete pattern of ocular dominance columns in human primary visual cortex. *J Neurosci* 27: 10391-10403, 2007.
- Albrecht DG, and Geisler WS.** Motion selectivity and the contrast-response function of simple cells in the visual cortex. *Vis Neurosci* 7: 531-546, 1991.
- Altmann CF, Bulthoff HH, and Kourtzi Z.** Perceptual organization of local elements into global shapes in the human visual cortex. *Curr Biol* 13: 342-349, 2003.
- Andrews TJ, and Ewbank MP.** Distinct representations for facial identity and changeable aspects of faces in the human temporal lobe. *Neuroimage* 23: 905-913, 2004.
- Aristotle.** *The Metaphysics*. London, England: Penguin Classics.
- Aronov D, Reich DS, Mechler F, and Victor JD.** Neural coding of spatial phase in V1 of the macaque monkey. *J Neurophysiol* 89: 3304-3327, 2003.
- Backus BT, Fleet DJ, Parker AJ, and Heeger DJ.** Human cortical activity correlates with stereoscopic depth perception. *J Neurophysiol* 86: 2054-2068, 2001.
- Badcock DR.** How do we discriminate relative spatial phase? *Vision Res* 24: 1847-1857, 1984a.
- Badcock DR.** Spatial phase or luminance profile discrimination? *Vision Res* 24: 613-623, 1984b.
- Baizer JS.** Receptive field properties of V3 neurons in monkey. *Invest Ophthalmol Vis Sci* 23: 87-95, 1982.
- Baizer JS, Ungerleider LG, and Desimone R.** Organization of visual inputs to the inferior temporal and posterior parietal cortex in macaques. *J Neurosci* 11: 168-190, 1991.
- Bartels A, Logothetis NK, and Moutoussis K.** fMRI and its interpretations: an illustration on directional selectivity in area V5/MT. *Trends Neurosci* 31: 444-453, 2008.
- Bartfeld E, and Grinvald A.** Relationships between orientation-preference pinwheels, cytochrome oxidase blobs, and ocular-dominance columns in primate striate cortex. *Proc Natl Acad Sci U S A* 89: 11905-11909, 1992.
- Biswal B, Hudetz AG, Yetkin FZ, Haughton VM, and Hyde JS.** Hypercapnia reversibly suppresses low-frequency fluctuations in the human motor cortex during rest using echo-planar MRI. *J Cereb Blood Flow Metab* 17: 301-308, 1997a.
- Biswal B, Yetkin FZ, Haughton VM, and Hyde JS.** Functional connectivity in the motor cortex of resting human brain using echo-planar MRI. *Magn Reson Med* 34: 537-541, 1995.
- Biswal BB, Van Kylen J, and Hyde JS.** Simultaneous assessment of flow and BOLD signals in resting-state functional connectivity maps. *NMR Biomed* 10: 165-170, 1997b.
- Blakemore C, and Campbell FW.** On the existence of neurones in the human visual system selectively sensitive to the orientation and size of retinal images. *J Physiol* 203: 237-260, 1969.
- Boynton GM.** Imaging orientation selectivity: decoding conscious perception in V1. *Nature Neuroscience* 8: 541 - 542 2005.
- Boynton GM, Engel SA, Glover GH, and Heeger DJ.** Linear systems analysis of functional magnetic resonance imaging in human V1. *J Neurosci* 16: 4207-4221, 1996.
- Boynton GM, and Finney EM.** Orientation-specific adaptation in human visual cortex. *J Neurosci* 23: 8781-8787, 2003.
- Bradley A, Switkes E, and De Valois K.** Orientation and spatial frequency selectivity of adaptation to color and luminance gratings. *Vision Res* 28: 841-856, 1988.
- Brefczynski JA, and DeYoe EA.** A physiological correlate of the 'spotlight' of visual attention. *Nat Neurosci* 2: 370-374, 1999.

- 
- Burkhalter A, and Van Essen DC.** Processing of color, form and disparity information in visual areas VP and V2 of ventral extrastriate cortex in the macaque monkey. *J Neurosci* 6: 2327-2351, 1986.
- Burr DC.** Sensitivity to spatial phase. *Vision Res* 20: 391-396, 1980.
- Burr DC, Morrone MC, and Spinelli D.** Evidence for edge and bar detectors in human vision. *Vision Res* 29: 419-431, 1989.
- Buxton RB.** *Introduction to Functional Magnetic Resonance Imaging. Principles and Techniques* Cambridge: Cambridge University Press, 2002.
- Campbell FW, Cooper GF, and Enroth-Cugell C.** The spatial selectivity of the visual cells of the cat. *J Physiol* 203: 223-235, 1969.
- Carandini M, Barlow HB, O'Keefe LP, Poirson AB, and Movshon JA.** Adaptation to contingencies in macaque primary visual cortex. *Philos Trans R Soc Lond B Biol Sci* 352: 1149-1154, 1997a.
- Carandini M, Heeger DJ, and Movshon JA.** Linearity and normalization in simple cells of the macaque primary visual cortex. *J Neurosci* 17: 8621-8644, 1997b.
- Carandini M, Heeger DJ, and Senn W.** A synaptic explanation of suppression in visual cortex. *J Neurosci* 22: 10053-10065, 2002.
- Clifford CW.** Perceptual adaptation: motion parallels orientation. *Trends Cogn Sci* 6: 136-143, 2002.
- Connolly M, and Van Essen D.** The representation of the visual field in parvicellular and magnocellular layers of the lateral geniculate nucleus in the macaque monkey. *J Comp Neurol* 226: 544-564, 1984.
- Cox DD, and Savoy RL.** Functional magnetic resonance imaging (fMRI) "brain reading": detecting and classifying distributed patterns of fMRI activity in human visual cortex. *NeuroImage* 19: 261-270, 2003.
- Dacey DM, and Petersen MR.** Dendritic field size and morphology of midget and parasol ganglion cells of the human retina. *Proc Natl Acad Sci U S A* 89: 9666-9670, 1992.
- Dale AM, Fischl B, and Sereno MI.** Cortical surface-based analysis. I. Segmentation and surface reconstruction. *Neuroimage* 9: 179-194, 1999.
- Dale AM, and Sereno MI.** Improved localization of cortical activity by combining EEG and MEG with MRI cortical surface reconstruction: A linear approach. *Journal of Cognitive Neuroscience* *Journal of Cognitive Neuroscience* 5: 162-176, 1993.
- De Valois KK, and Tootell RB.** Spatial-frequency-specific inhibition in cat striate cortex cells. *J Physiol* 336: 359-376, 1983.
- Derrington AM, and Lennie P.** The influence of temporal frequency and adaptation level on receptive field organization of retinal ganglion cells in cat. *J Physiol* 333: 343-366, 1982.
- Desimone R.** Neural mechanisms for visual memory and their role in attention. *Proc Natl Acad Sci U S A* 93: 13494-13499, 1996.
- Desimone R, and Schein SJ.** Visual properties of neurons in area V4 of the macaque: sensitivity to stimulus form. *J Neurophysiol* 57: 835-868, 1987.
- Desimone R, Schein SJ, Moran J, and Ungerleider LG.** Contour, color and shape analysis beyond the striate cortex. *Vision Res* 25: 441-452, 1985.
- Detre GJ, Polyn SM, Moore CD, Natu VS, Singer BD, Cohen JD, Haxby JV, and Norman KA.** The Multi-Voxel Pattern Analysis (MVPA) toolbox. In: *Annual Meeting of the Organization of Human Brain Mapping* 2006.
- DeYoe EA, Carman GJ, Bandettini P, Glickman S, Wieser J, Cox R, Miller D, and Neitz J.** Mapping striate and extrastriate visual areas in human cerebral cortex. *Proc Natl Acad Sci U S A* 93: 2382-2386, 1996.
- Donaldson DL, and Buckner RL.** Effective paradigm design. In: *Functional MRI: an introduction to methods*, edited by Jezzard P, Matthews PM, and Smith SM. New York: Oxford University Press, 2001, p. 177, 195.
- Dougherty RF, Koch VM, Brewer AA, Fischer B, Modersitzki J, and Wandell BA.** Visual field representations and locations of visual areas V1/2/3 in human visual cortex. *J Vis* 3: 586-598, 2003.
- Drury HA, Van Essen DC, Anderson CH, Lee CW, Coogan TA, and Lewis JW.** Computerized mappings of the cerebral cortex: a multiresolution flattening method and a surface-based coordinate system. *J Cogn Neurosci* 8: 1-28, 1996.

- 
- Duda OR, Hart PE, and Stork DG.** *Pattern Classification*. New York: Wiley, 2001.
- Efron B, and Tibshirani RJ.** *An Introduction to the Bootstrap*. New York: Chapman and Hall, 1993.
- Eger E, Ashburner J, Haynes JD, Dolan RJ, and Rees G.** fMRI activity patterns in human LOC carry information about object exemplars within category. *J Cogn Neurosci* 20: 356-370, 2008.
- Engel SA.** Adaptation of oriented and unoriented color-selective neurons in human visual areas. *Neuron* 45: 613-623, 2005.
- Engel SA, Glover GH, and Wandell BA.** Retinotopic organization in human visual cortex and the spatial precision of functional MRI. *Cereb Cortex* 7: 181-192, 1997.
- Engel SA, Rumelhart DE, Wandell BA, Lee AT, Glover GH, Chichilnisky EJ, and Shadlen MN.** fMRI of human visual cortex. *Nature* 369: 525, 1994.
- Fang F, Murray SO, Kersten D, and He S.** Orientation-tuned FMRI adaptation in human visual cortex. *J Neurophysiol* 94: 4188-4195, 2005.
- Felleman DJ, and Van Essen DC.** Distributed hierarchical processing in the primate cerebral cortex. *Cereb Cortex* 1: 1-47, 1991.
- Felleman DJ, and Van Essen DC.** Receptive field properties of neurons in area V3 of macaque monkey extrastriate cortex. *J Neurophysiol* 57: 889-920, 1987.
- Field DJ, Hayes A, and Hess RF.** Contour integration by the human visual system: evidence for a local "association field". *Vision Res* 33: 173-193, 1993.
- Gallant JL, Braun J, and Van Essen DC.** Selectivity for polar, hyperbolic, and Cartesian gratings in macaque visual cortex. *Science* 259: 100-103, 1993.
- Gegenfurtner KR, Kiper DC, and Levitt JB.** Functional properties of neurons in macaque area V3. *J Neurophysiol* 77: 1906-1923, 1997.
- Georgeson M.** Visual aftereffects: cortical neurons change their tune. *Curr Biol* 14: R751-753, 2004.
- Georgeson MA.** The effect of spatial adaptation on perceived contrast. *Spat Vis* 1: 103-112, 1985.
- Georgeson MA, and Meese TS.** Perception of stationary plaids: The role of spatial filters in edge analysis. *Vision Research* 37: 3255-3271, 1997.
- Gheorghiu E, and Kingdom FA.** Chromatic tuning of contour-shape mechanisms revealed through the shape-frequency and shape-amplitude after-effects. *Vision Res* 47: 1935-1949, 2007.
- Gheorghiu E, and Kingdom FA.** Multiplication in curvature processing. *J Vis* 9: 23 21-17, 2009.
- Gheorghiu E, and Kingdom FA.** Spatial properties of curvature-encoding mechanisms revealed through the shape-frequency and shape-amplitude after-effects. *Vision Res* 48: 1107-1124, 2008.
- Gibson JJ, and Radner M.** Adaptation, after-effect and contrast in the perception of tilted lines. *Journal of Experimental Psychology* 20: 453-467, 1937.
- Glover GH.** Deconvolution of impulse response in event-related BOLD fMRI. *Neuroimage* 9: 416-429, 1999.
- Goodale MA, and Milner AD.** Separate visual pathways for perception and action. *Trends Neurosci* 15: 20-25, 1992.
- Graham N.** Visual detection of aperiodic spatial stimuli by probability summation among narrowband channels. *Vision Res* 17: 637-652, 1977.
- Graham N, and Nachmias J.** Detection of grating patterns containing two spatial frequencies: a comparison of single-channel and multiple-channels models. *Vision Res* 11: 251-259, 1971.
- Greenlee MW, and Magnussen S.** Saturation of the tilt aftereffect. *Vision Res* 27: 1041-1043, 1987.
- Grill-Spector K, Henson R, and Martin A.** Repetition and the brain: neural models of stimulus-specific effects. *Trends Cogn Sci* 10: 14-23, 2006.
- Grill-Spector K, Kushnir T, Edelman S, Avidan G, Itzchak Y, and Malach R.** Differential processing of objects under various viewing conditions in the human lateral occipital complex. *Neuron* 24: 187-203, 1999.
- Grill-Spector K, Kushnir T, Edelman S, Itzchak Y, and Malach R.** Cue-invariant activation in object-related areas of the human occipital lobe. *Neuron* 21: 191-202, 1998.

- 
- Grill-Spector K, Kushnir T, Hendler T, and Malach R.** The dynamics of object-selective activation correlate with recognition performance in humans. *Nat Neurosci* 3: 837-843, 2000.
- Grill-Spector K, and Malach R.** fMR-adaptation: a tool for studying the functional properties of human cortical neurons. *Acta Psychol (Amst)* 107: 293-321, 2001.
- Hadjikhani N, Liu AK, Dale AM, Cavanagh P, and Tootell RB.** Retinotopy and color sensitivity in human visual cortical area V8. *Nat Neurosci* 1: 235-241, 1998.
- Hancock S, and Peirce JW.** Selective mechanisms for simple contours revealed by compound adaptation. *J Vis* 8: 11 11-10, 2008.
- Haushofer J, Livingstone MS, and Kanwisher N.** Multivariate patterns in object-selective cortex dissociate perceptual and physical shape similarity. *PLoS Biol* 6: e187, 2008.
- Haxby JV, Gobbini MI, Furey ML, Ishai A, Schouten JL, and Pietrini P.** Distributed and Overlapping Representations of Faces and Objects in Ventral Temporal Cortex. *Science* 293: 2425-2430, 2001.
- Haynes J-D, and Rees G.** Decoding mental states from brain activity in humans. *Nat Rev Neurosci* 7: 523-534, 2006.
- Haynes JD, and Rees G.** Predicting the orientation of invisible stimuli from activity in human primary visual cortex. *Nat Neurosci* 8: 686-691, 2005.
- Heeger DJ.** Normalization of cell responses in cat striate cortex. *Vis Neurosci* 9: 181-197, 1992.
- Heeger DJ, and Ress D.** What does fMRI tell us about neuronal activity? *Nat Rev Neurosci* 3: 142-151, 2002.
- Hendry SH, and Reid RC.** The koniocellular pathway in primate vision. *Annu Rev Neurosci* 23: 127-153, 2000.
- Henriksson L, Hyvriinen A, and Vanni S.** Representation of broadband edges and spatial phase congruency in human visual cortex. *Journal of Vision* 9: 900-900, 2009.
- Henson RN, Shallice T, Gorno-Tempini ML, and Dolan RJ.** Face repetition effects in implicit and explicit memory tests as measured by fMRI. *Cereb Cortex* 12: 178-186, 2002.
- Hess RF, and Field D.** Is the increased spatial uncertainty in the normal periphery due to spatial undersampling or uncalibrated disarray? *Vision Res* 33: 2663-2670, 1993.
- Hoge RD, and Pike GB.** *Quantitative measurement using fMRI*. New York: Oxford University Press, 2001.
- Holodniy M, Katzenstein D, Winters M, Montoya J, Shafer R, Kozal M, Ragni M, and Merigan TC.** Measurement of HIV virus load and genotypic resistance by gene amplification in asymptomatic subjects treated with combination therapy. *J Acquir Immune Defic Syndr* 6: 366-369, 1993.
- Holt JJ, and Ross J.** Phase perception in the high spatial frequency range. *Vision Res* 20: 933-935, 1980.
- Horowitz AL.** *MRI physics for radiologists*. New York: Springer-Verlag, 1995.
- Horwitz B, Friston KJ, and Taylor JG.** Neural modeling and functional brain imaging: an overview. *Neural Netw* 13: 829-846, 2000.
- Howseman AM, and Bowtell RW.** Functional magnetic resonance imaging: imaging techniques and contrast mechanisms. *Philos Trans R Soc Lond B Biol Sci* 354: 1179-1194, 1999.
- Huang PC, Kingdom FA, and Hess RF.** Only two phase mechanisms, +/-cosine, in human vision. *Vision Res* 46: 2069-2081, 2006.
- Hubel DH.** *Eye, Brain and Vision*. New York: 1988.
- Hubel DH, and Livingstone MS.** Segregation of form, color, and stereopsis in primate area 18. *J Neurosci* 7: 3378-3415, 1987.
- Hubel DH, and Wiesel TN.** Ferrier lecture. Functional architecture of macaque monkey visual cortex. *Proc R Soc Lond B Biol Sci* 198: 1-59, 1977.
- Hubel DH, and Wiesel TN.** Receptive Fields and Functional Architecture in Two Nonstriate Visual Areas (18 and 19) of the Cat. *J Neurophysiol* 28: 229-289, 1965.
- Hubel DH, and Wiesel TN.** Receptive fields and functional architecture of monkey striate cortex. *J Physiol* 195: 215-243, 1968.

---

**Hubel DH, and Wiesel TN.** Receptive fields, binocular interaction and functional architecture in the cat's visual cortex. *J Physiol* 160: 106-154, 1962.

**Huk AC, and Heeger DJ.** Pattern-motion responses in human visual cortex. *Nat Neurosci* 5: 72-75, 2002.

**Huk AC, Ress D, and Heeger DJ.** Neuronal basis of the motion aftereffect reconsidered. *Neuron* 32: 161-172, 2001.

**Ito M, and Komatsu H.** Representation of angles embedded within contour stimuli in area V2 of macaque monkeys. *J Neurosci* 24: 3313-3324, 2004.

**Jenkinson M, Bannister P, Brady M, and Smith S.** Improved optimization for the robust and accurate linear registration and motion correction of brain images. *Neuroimage* 17: 825-841, 2002.

**Jezzard P, and Clare S.** Principles of nuclear magnetic resonance and MRI. In: *Functional MRI: an introduction to methods*, edited by Jezzard P, Matthews, P.M. & Smith, S.M. . New York: Oxford University Press, 2001, p. 67-92.

**Jezzard P, Clare, S.** Principles of nuclear magnetic resonance and MRI. In: *Functional MRI: an introduction to methods*, edited by Jezzard P, Matthews, P.M. & Smith, S.M. . New York: Oxford University Press, 2001, p. 67-92.

**Kamitani Y, and Tong F.** Decoding seen and attended motion directions from activity in the human visual cortex. *Curr Biol* 16: 1096-1102, 2006.

**Kamitani Y, and Tong F.** Decoding the visual and subjective contents of the human brain. *Nat Neurosci* 8: 679-685, 2005.

**Kastner S, Pinsk MA, De Weerd P, Desimone R, and Ungerleider LG.** Increased activity in human visual cortex during directed attention in the absence of visual stimulation. *Neuron* 22: 751-761, 1999.

**Kourtzi Z, and Huberle E.** Spatiotemporal characteristics of form analysis in the human visual cortex revealed by rapid event-related fMRI adaptation. *Neuroimage* 28: 440-452, 2005.

**Kourtzi Z, and Kanwisher N.** Cortical regions involved in perceiving object shape. *J Neurosci* 20: 3310-3318, 2000.

**Kourtzi Z, and Kanwisher N.** Representation of perceived object shape by the human lateral occipital complex. *Science* 293: 1506-1509, 2001.

**Kourtzi Z, Tolias AS, Altmann CF, Augath M, and Logothetis NK.** Integration of local features into global shapes: monkey and human FMRI studies. *Neuron* 37: 333-346, 2003.

**Krekelberg B, Boynton GM, and van Wezel RJ.** Adaptation: from single cells to BOLD signals. *Trends Neurosci* 29: 250-256, 2006.

**Krekelberg B, Vatakis A, and Kourtzi Z.** Implied motion from form in the human visual cortex. *J Neurophysiol* 94: 4373-4386, 2005.

**Kriegeskorte N, and Goebel R.** An efficient algorithm for topologically correct segmentation of the cortical sheet in anatomical mr volumes. *Neuroimage* 14: 329-346, 2001.

**Ku SP, Gretton A, Macke J, and Logothetis NK.** Comparison of pattern recognition methods in classifying high-resolution BOLD signals obtained at high magnetic field in monkeys. *Magn Reson Imaging* 26: 1007-1014, 2008.

**Larsson J.** Imaging vision: functional mapping of intermediate visual processes in man. In: *Karolinska Institutet*. Stockholm: 2001.

**Larsson J, and Heeger DJ.** Two retinotopic visual areas in human lateral occipital cortex. *J Neurosci* 26: 13128-13142, 2006.

**Larsson J, Landy MS, and Heeger DJ.** Orientation-selective adaptation to first- and second-order patterns in human visual cortex. *J Neurophysiol* 95: 862-881, 2006.

**Ledgeway T, and Smith AT.** Evidence for separate motion-detecting mechanisms for first- and second-order motion in human vision. *Vision Res* 34: 2727-2740, 1994.

**Lennie P.** Single units and visual cortical organization. *Perception* 27: 889-935, 1998.

**Levitt JB, Kiper DC, and Movshon JA.** Receptive fields and functional architecture of macaque V2. *J Neurophysiol* 71: 2517-2542, 1994.

**Li S, Ostwald D, Giese M, and Kourtzi Z.** Flexible coding for categorical decisions in the human brain. *J Neurosci* 27: 12321-12330, 2007.

**Livingstone MS, and Hubel DH.** Do the relative mapping densities of the magno- and parvocellular systems vary with eccentricity? *J Neurosci* 8: 4334-4339, 1988.

---

**Loffler G, Wilson HR, and Wilkinson F.** Local and global contributions to shape discrimination. *Vision Res* 43: 519-530, 2003.

**Logothetis NK.** The neural basis of the blood-oxygen-level-dependent functional magnetic resonance imaging signal. *Philos Trans R Soc Lond B Biol Sci* 357: 1003-1037, 2002.

**Logothetis NK.** What we can do and what we cannot do with fMRI. *Nature* 453: 869-878, 2008.

**Logothetis NK, Pauls J, Augath M, Trinath T, and Oeltermann A.** Neurophysiological investigation of the basis of the fMRI signal. *Nature* 412: 150-157, 2001.

**MacDonald D, Kabani N, Avis D, and Evans AC.** Automated 3-D extraction of inner and outer surfaces of cerebral cortex from MRI. *Neuroimage* 12: 340-356, 2000.

**Maffei L, Fiorentini A, and Bisti S.** Neural correlate of perceptual adaptation to gratings. *Science* 182: 1036-1038, 1973.

**Malach R, Reppas JB, Benson RR, Kwong KK, Jiang H, Kennedy WA, Ledden PJ, Brady TJ, Rosen BR, and Tootell RB.** Object-related activity revealed by functional magnetic resonance imaging in human occipital cortex. *Proc Natl Acad Sci U S A* 92: 8135-8139, 1995.

**Mansfield P.** Multiplanar image formation using NMR spin echoes. *J Phys* L55-L58, 1977.

**Marr D.** *Vision: A Computational Investigation into the Human Representation and Processing of Visual Information.* San Francisco: W.H. Freeman, 1982.

**Maunsell JH, Sclar G, Nealey TA, and DePriest DD.** Extraretinal representations in area V4 in the macaque monkey. *Vis Neurosci* 7: 561-573, 1991.

**McKeefry DJ, and Zeki S.** The position and topography of the human colour centre as revealed by functional magnetic resonance imaging. *Brain* 120 ( Pt 12): 2229-2242, 1997.

**Mechler F, Reich DS, and Victor JD.** Detection and discrimination of relative spatial phase by V1 neurons. *J Neurosci* 22: 6129-6157, 2002.

**Merigan WH.** Basic visual capacities and shape discrimination after lesions of extrastriate area V4 in macaques. *Vis Neurosci* 13: 51-60, 1996.

**Merigan WH, Katz LM, and Maunsell JH.** The effects of parvocellular lateral geniculate lesions on the acuity and contrast sensitivity of macaque monkeys. *J Neurosci* 11: 994-1001, 1991.

**Merigan WH, Nealey TA, and Maunsell JH.** Visual effects of lesions of cortical area V2 in macaques. *J Neurosci* 13: 3180-3191, 1993.

**Mitchell TM, Hutchinson R, Just MA, Niculescu RS, Pereira F, and Wang X.** Classifying instantaneous cognitive states from FMRI data. *AMIA Annu Symp Proc* 465-469, 2003.

**Montaser-Kouhsari L, Landy MS, Heeger DJ, and Larsson J.** Orientation-selective adaptation to illusory contours in human visual cortex. *J Neurosci* 27: 2186-2195, 2007.

**Moon CH, Fukuda M, Park SH, and Kim SG.** Neural interpretation of blood oxygenation level-dependent fMRI maps at submillimeter columnar resolution. *J Neurosci* 27: 6892-6902, 2007.

**Movshon JA, Adelson EH, Gizzi MS, and Newsome WT.** The analysis of moving visual patterns. *Exp Brain Res* 11: 117-152, 1986.

**Movshon JA, Thompson ID, and Tolhurst DJ.** Receptive field organization of complex cells in the cat's striate cortex. *J Physiol* 283: 79-99, 1978.

**Nealey TA, and Maunsell JH.** Magnocellular and parvocellular contributions to the responses of neurons in macaque striate cortex. *J Neurosci* 14: 2069-2079, 1994.

**Nestares O, and Heeger DJ.** Robust multiresolution alignment of MRI brain volumes. *Magn Reson Med* 43: 705-715, 2000.

**Norman KA, Polyn SM, Detre GJ, and Haxby JV.** Beyond mind-reading: multi-voxel pattern analysis of fMRI data. *Trends in Cognitive Sciences* 10: 424-430, 2006.

**Obermayer K, and Blasdel GG.** Geometry of orientation and ocular dominance columns in monkey striate cortex. *J Neurosci* 13: 4114-4129, 1993.

---

**Ogawa S, Lee TM, Nayak AS, and Glynn P.** Oxygenation-sensitive contrast in magnetic resonance image of rodent brain at high magnetic fields. *Magn Reson Med* 14: 68-78, 1990.

**Ohzawa I, Sclar G, and Freeman RD.** Contrast gain control in the cat visual cortex. *Nature* 298: 266-268, 1982.

**Ohzawa I, Sclar G, and Freeman RD.** Contrast gain control in the cat's visual system. *J Neurophysiol* 54: 651-667, 1985.

**Ostwald D, Lam JM, Li S, and Kourtzi Z.** Neural coding of global form in the human visual cortex. *J Neurophysiol* 99: 2456-2469, 2008.

**Peirce JW.** PsychoPy-Psychophysics software in Python. *J Neurosci Methods* 162: 8-13, 2007.

**Peirce JW, and Taylor LJ.** Selective mechanisms for complex visual patterns revealed by adaptation. *Neuroscience* 141: 15-18, 2006.

**Perry VH, Oehler R, and Cowey A.** Retinal ganglion cells that project to the dorsal lateral geniculate nucleus in the macaque monkey. *Neuroscience* 12: 1101-1123, 1984.

**Peterhans E, and von der Heydt R.** Mechanisms of contour perception in monkey visual cortex. II. Contours bridging gaps. *J Neurosci* 9: 1749-1763, 1989.

**Pollen DA, Gaska JP, and Jacobson LD.** Responses of simple and complex cells to compound sine-wave gratings. *Vision Res* 28: 25-39, 1988.

**Polyn SM, Natu VS, Cohen JD, and Norman KA.** Category-specific cortical activity precedes retrieval during memory search. *Science* 310: 1963-1966, 2005.

**Press WA, Brewer AA, Dougherty RF, Wade AR, and Wandell BA.** Visual areas and spatial summation in human visual cortex. *Vision Res* 41: 1321-1332, 2001.

**Purdon PL, and Weisskoff RM.** Effect of temporal autocorrelation due to physiological noise and stimulus paradigm on voxel-level false-positive rates in fMRI. *Hum Brain Mapp* 6: 239-249, 1998.

**Rust NC, Mante V, Simoncelli EP, and Movshon JA.** How MT cells analyze the motion of visual patterns. *Nature Neuroscience* 9: 1421-1431, 2006.

**Sawamura H, Georgieva S, Vogels R, Vanduffel W, and Orban GA.** Using functional magnetic resonance imaging to assess adaptation and size invariance of shape processing by humans and monkeys. *J Neurosci* 25: 4294-4306, 2005.

**Sayres R, and Grill-Spector K.** Object-selective cortex exhibits performance-independent repetition suppression. *J Neurophysiol* 95: 995-1007, 2006.

**Schiller PH, Logothetis NK, and Charles ER.** Role of the color-opponent and broad-band channels in vision. *Vis Neurosci* 5: 321-346, 1990.

**Sclar G, Lennie P, and DePriest DD.** Contrast adaptation in striate cortex of macaque. *Vision Res* 29: 747-755, 1989.

**Sereno MI, Dale AM, Reppas JB, Kwong KK, Belliveau JW, Brady TJ, Rosen BR, and Tootell RB.** Borders of multiple visual areas in humans revealed by functional magnetic resonance imaging. *Science* 268: 889-893, 1995.

**Simoncelli EP, and Olshausen BA.** NATURAL IMAGE STATISTICS AND NEURAL REPRESENTATION. *Annual Review of Neuroscience* 24: 1193-1216, 2001.

**Smith AM, Lewis BK, Ruttimann UE, Ye FQ, Sinnwell TM, Yang Y, Duyn JH, and Frank JA.** Investigation of low frequency drift in fMRI signal. *Neuroimage* 9: 526-533, 1999.

**Smith AT, Greenlee MW, Singh KD, Kraemer FM, and Hennig J.** The processing of first- and second-order motion in human visual cortex assessed by functional magnetic resonance imaging (fMRI). *J Neurosci* 18: 3816-3830, 1998.

**Snowden RJ, and Hammett ST.** Spatial frequency adaptation: threshold elevation and perceived contrast. *Vision Res* 36: 1797-1809, 1996.

**Solomon SG, Peirce JW, Dhruv NT, and Lennie P.** Profound contrast adaptation early in the visual pathway. *Neuron* 42: 155-162, 2004.

**Somers DC, Dale AM, Seiffert AE, and Tootell RB.** Functional MRI reveals spatially specific attentional modulation in human primary visual cortex. *Proc Natl Acad Sci U S A* 96: 1663-1668, 1999.

**Stanley DA, and Rubin N.** fMRI activation in response to illusory contours and salient regions in the human lateral occipital complex. *Neuron* 37: 323-331, 2003.

- 
- Stanley DA, and Rubin N.** Functionally distinct sub-regions in the lateral occipital complex revealed by fMRI responses to abstract 2-dimensional shapes and familiar objects. *Journal of Vision* 5: 911-911, 2005.
- Stromeyer CF, 3rd, Lange AF, and Ganz L.** Spatial frequency phase effects in human vision. *Vision Res* 13: 2345-2360, 1973.
- Suzuki S.** Attention-dependent brief adaptation to contour orientation: a high-level aftereffect for convexity? *Vision Res* 41: 3883-3902, 2001.
- Suzuki S.** Attentional selection of overlapped shapes: a study using brief shape aftereffects. *Vision Res* 43: 549-561, 2003.
- Tolias AS, Smirnakis SM, Augath MA, Trinath T, and Logothetis NK.** Motion processing in the macaque: revisited with functional magnetic resonance imaging. *J Neurosci* 21: 8594-8601, 2001.
- Tootell RB, Hadjikhani N, Hall EK, Marrett S, Vanduffel W, Vaughan JT, and Dale AM.** The retinotopy of visual spatial attention. *Neuron* 21: 1409-1422, 1998.
- Tootell RB, and Hamilton SL.** Functional anatomy of the second visual area (V2) in the macaque. *J Neurosci* 9: 2620-2644, 1989.
- Tootell RB, Mendola JD, Hadjikhani NK, Ledden PJ, Liu AK, Reppas JB, Sereno MI, and Dale AM.** Functional analysis of V3A and related areas in human visual cortex. *J Neurosci* 17: 7060-7078, 1997.
- Ungerleider LG, and Mishkin M.** In: *Analysis of Visual Behavior*, edited by Ingle DJ, Goodale MA, and Mansfield RJW. Cambridge, MA: The MIT Press., 1982, p. 549-584.
- Vaina LM, Gryzwacz NM, Saiviroonporn P, LeMay M, Biefang DC, and Cowey A.** Can spatial and temporal motion integration compensate for deficits in local motion mechanisms? *Neuropsychologia* 41: 1817-1836, 2003.
- Vanduffel W, Tootell RBH, Schoups AA, and Orban GA.** The Organization of Orientation Selectivity Throughout Macaque Visual Cortex. *Cereb Cortex* 12: 647-662, 2002.
- Vinberg J, and Grill-Spector K.** Representation of shapes, edges, and surfaces across multiple cues in the human visual cortex. *J Neurophysiol* 99: 1380-1393, 2008.
- von der Heydt R, and Peterhans E.** Mechanisms of contour perception in monkey visual cortex. I. Lines of pattern discontinuity. *J Neurosci* 9: 1731-1748, 1989.
- Vuilleumier P, Henson RN, Driver J, and Dolan RJ.** Multiple levels of visual object constancy revealed by event-related fMRI of repetition priming. *Nat Neurosci* 5: 491-499, 2002.
- Wade AR, Brewer AA, Rieger JW, and Wandell BA.** Functional measurements of human ventral occipital cortex: retinotopy and colour. *Philos Trans R Soc Lond B Biol Sci* 357: 963-973, 2002.
- Wandell BA, Brewer AA, and Dougherty RF.** Visual field map clusters in human cortex. *Philos Trans R Soc Lond B Biol Sci* 360: 693-707, 2005.
- Wandell BA, Chial S, and Backus BT.** Visualization and measurement of the cortical surface. *J Cogn Neurosci* 12: 739-752, 2000.
- Wandell BA, and Wade AR.** Functional imaging of the visual pathways. *Neurol Clin* 21: 417-443, vi, 2003.
- Watt RJ, and Andrews DP.** Contour curvature analysis: hyperacuties in the discrimination of detailed shape. *Vision Res* 22: 449-460, 1982.
- Webster MA, and De Valois RL.** Relationship between spatial-frequency and orientation tuning of striate-cortex cells. *J Opt Soc Am A* 2: 1124-1132, 1985.
- Yacoub E, Shmuel A, Logothetis N, and Ugurbil K.** Robust detection of ocular dominance columns in humans using Hahn Spin Echo BOLD functional MRI at 7 Tesla. *Neuroimage* 37: 1161-1177, 2007.
- Zarahn E, Aguirre GK, and D'Esposito M.** Empirical analyses of BOLD fMRI statistics. I. Spatially unsmoothed data collected under null-hypothesis conditions. *Neuroimage* 5: 179-197, 1997.
- Zeki S.** *A vision of the brain*. Oxford: Blackwell scientific publications, 1993.
- Zeki SM.** Colour coding in rhesus monkey prestriate cortex. *Brain Res* 53: 422-427, 1973.
- Zeki SM.** The third visual complex of rhesus monkey prestriate cortex. *J Physiol* 277: 245-272, 1978a.

---

**Zeki SM.** Uniformity and diversity of structure and function in rhesus monkey prestriate visual cortex. *J Physiol* 277: 273-290, 1978b.