Biases in Perception Of Visual Motion

Elisa Zamboni, MSc.

Thesis submitted to The University of Nottingham
for the degree of Doctor of Philosophy

March 2017

ACKNOWLEDGMENTS

I would like to take the opportunity to thank all the people I met during this incredible journey that led me to this remarkable achievement.

Firstly, thank you to my supervisors Denis Schluppeck, Tim Ledgeway, and Paul McGraw for the support and exciting discussions throughout these years. You helped me grow, not only from a scientific point of view, but also in confidence. It has been a pleasure working with you!

Many thanks also to my advisor Ben Webb who, together with Neil Roach, gave me constructive feedbacks and suggestions over the years.

I am grateful to you and all of the Nottingham Visual Neuroscience Group: working with you felt like family!

Thank you to Chris Scholes for introducing me to eye-tracking and letting me share his lab. To Julien Besle for taking the time to go through the fMRI analysis with me, and, together with Zahra Hussain, for introducing me to the fun of Frisbee: many, many thanks!

I would like to thank my good friends from C68: I will always keep our coffee breaks, together with all the food, sport, and tv-series chats, in my heart!!! You made me feel at home, despite being a thousand miles away from it!

Thank you to Maggi and Mags that opened their house to me, helped me growing as a person and showed me the beauty of the British countryside!

Finally, a big thank you to my family for supporting me and coping with my desire of being abroad. A huge thank you to you, Pippo, for accepting my choices and being as excited as I am whenever a new challenge presents itself. You are, and always will be, my best friend and super-brother!

ABSTRACT

Perceptual decision making refers to the process of making a choice among a series of options based on sensory information. Several studies have used visual stimuli to gain an understanding of the processes involved in encoding sensory information and its decoding, leading to a perceptual decision. One popular visual modality for studying these questions is motion and the ability to discriminate between axes of motion. Several mathematical models describing the processes of perceptual decision making have been proposed – many of them are based on data from electrophysiological experiments on macaque monkeys. By directly recording neuronal activity while monkeys were presented with different visual stimuli and making categorical choices about the perceived direction of motion, scientists have been able to study how decisions are made when enough perceptual evidence is accumulated to reach a threshold.

A particularly interesting aspect of perceptual decision making is that it allows the study of situations in which the choice deviates from the physical features characterising the stimulus (e.g., a leftward motion is presented but the subject reports perceiving a rightward motion). A type of such perceptual bias is called reference repulsion: a systematic bias away from a reference when estimating the direction of motion of a stimulus. Several possible explanations of this phenomenon have been proposed: incorrect encoding of sensory information, influence of prior knowledge about the

world, response-related factors such as expectations, rewards, and response history.

The aim of this thesis was to shed light on when in the sequence of decision making such perceptual biases arise, as well as further address both sensory and higher-order factors that influence perceptual decisions of visual stimuli. We combined a series of psychophysical, eye-tracking, and neuroimaging studies, together with computational modelling approaches, to selectively look at the effect of: sensory information available during decision making, task-related sensory information processing, response modality, and also look for specific mechanisms involved in processing highly similar / dissimilar stimuli.

The findings presented in this thesis show that perceptual biases in estimates of motion direction arise at a later stage than at the encoding of sensory representation, as previously thought. In particular, we show that information present at the time of the *response* is fundamental for the bias to emerge: the presence of a reference while estimating direction of motion results in reference repulsion, but this effect is not there when the same estimate is given in the absence of a reference. Moreover, the information given by the reference at the time of response – when subjects report the perceived motion, rather than at the time of stimulus presentation – plays a crucial role in the observed perceptual bias.

These findings were used to develop a mathematical model able to describe the phenomena observed, as well as making a series of testable predictions. For example, the model could be used in future work to predict responses

when more than one reference is present, when order of presentation of target and reference is inverted, and so on.

By manipulating the modality with which subjects estimated the direction of motion of the stimuli they were presented with, it was also possible to show that a perceptual bias is observed for manual reproduction of the perceived direction, but not when the response is given by a saccadic eye movement. Finally, by looking at the brain activity recorded when performing a coarse / fine discrimination task in a functional magnetic resonance imaging (fMRI) study, we aimed at distinguishing between activity patterns encoding highly dissimilar / similar stimuli. For these analyses, we used both conventional, univariate analysis techniques, as well as a more advanced and relatively more recent multivariate approaches to the data. First, the retinotopic mapping of areas in early visual cortex and area MT was obtained through phase-encoded methods. Second, a version of the Generalised Linear Model was applied to the data measured while subjects were performing a fine / coarse discrimination task. This allowed to ensure the adequacy of tasks and stimuli used in the imaging study. I also applied the *population Receptive Field* methodology to fit a more explicit, physiologically relevant model of visual responses to the voxel-wise fMRI time series. Third, given that the spatial scale of the question we addressed in this study required aggregating sub-voxel differences in the fMRI responses during a fine versus coarse visual motion discrimination task, we employed a multivariate approach. This consisted in implementing a *forward encoding model* aimed at reducing the number of dimensions from

several hundreds (given by the number of voxels) to a much smaller set of hypothetical channels. By considering the responses in these channels as a weighted combination from many hundred voxels we re-cast the activity patterns in a physiologically relevant space to predict responses to arbitrary visual motion directions. While there were very interesting aspects to the results from these imaging experiments, the analysis was inconclusive on any task-related shifts in stimulus encoding. Possible explanations, together with alternative paradigms that can be used in future to further address this question are discussed.

1 INTRODUCTION 1

1.1 Perceptual Decision Making **1**

1.2 Anatomical and Physiological Organisation of the Visual System **17**
1.2.1 From the Retina to Early Visual Cortex 18
1.2.2 V1 25
1.2.3 V2 27
1.2.4 V3 29
1.2.5 V5 / MT and beyond 30

1.3 Reference Repulsion and Biases in Visual Motion Perception **34**

2 GENERAL METHODS 51

2.1 Psychophysics **51**
2.1.1 Visual Direction of Motion Discrimination Task and Psychometric Functions 54

2.2 Eye tracking **61**
2.2.1 Eye tracking technique 62
2.2.2 Eye-movement patterns 65
2.2.3 Eye-tracker Setup / Apparatus 67

2.3 Magnetic Resonance Imaging **76**
2.3.1 From Nuclear Magnetic Resonance to Magnetic Resonance Imaging 77
2.3.2 The Radiofrequency Pulse 79
2.3.3 Image formation in MRI 85
2.3.4 Functional MRI 90
2.3.5 Data Pre-Processing 99
2.3.6 Statistical Analysis 103
2.3.7 Functional MRI acquisition 137

3 ESTIMATES OF VISUAL MOTION DIRECTION 140

3.1 Experiment 1: Biased Estimates of Motion Direction following a Fine Discrimination Task **143**
3.1.1 Materials and Methods 143
3.1.2 Results 145

3.2 Experiment 2: Does the Bias Arise Early or Late during Information Processing? **149**
3.2.1 Materials and Methods 150
3.2.2 Results 151

3.3 Experiment 3: Is the Fine Discrimination Task Sufficient for the Bias to Arise? **157**
3.3.1 Materials and Methods 157
3.3.2 Results 159

3.4 Discussion **163**

4 MODELLING BIASES IN ESTIMATES OF MOTION DIRECTION 167

4.1 Decoding Model **177**

4.2 Late, Decision-Related Bias Model **182**
4.2.1 Modelling Performance: Reference Present 186
4.2.2 Modelling Performance: Reference Shift 187

4.3 Discussion **190**

5 VISUAL RESPONSES IN PERCEPTUAL DECISION-MAKING 193
5.1.1 Materials and Methods 198
5.1.2 Results 200

206
5.1.3 Discussion 211

6 IMAGING MECHANISMS INVOLVED IN FINE AND COARSE DISCRIMINATION TASKS 219
6.1.1 Materials and Methods 222
6.1.2 Results 225
6.1.3 Discussion 229

7 ENCODING MODELS OF VISUAL REPRESENTATIONS 233

7.1 Population Receptive Field (pRF) **233**
7.1.1 Material and Methods 234
7.1.2 Results 236
7.1.3 Discussion 242

7.2 Forward Encoding Model **244**
7.2.1 Material and Methods 246
7.2.2 Results 249
7.2.3 Discussion 253

8 GENERAL DISCUSSION 256

8.1 Summary of Findings **256**

8.2 Conclusions and Future Directions **262**

9 REFERENCES 266

10 APPENDIX – SUPPLEMENTARY MATERIAL 286

CHAPTER 1

1 Introduction

1.1 Perceptual Decision Making

My definition of *perceptual decision making* for this thesis is the ability of choosing one option amongst many – based on sensory information. Decisions are influenced by the reliability of the information upon which we need to make a choice, knowledge about past events, subjective and objective costs and benefits attached to any of the alternatives we need to choose from, as well as our emotional state and beliefs. But most importantly, perceptual decision making (Figure 1.1) requires encoding of sensory information: creating a representation of the visual stimulus that will then be used in the decision process. Once a representation of the information is built in the sensory system, this is then decoded: a strategy is applied onto the neural coding representing the sensory information in order to make sense of it and optimise the decision process. The result of the decoding process will then inform the areas involved in planning and performing a response, the output of the decision process.

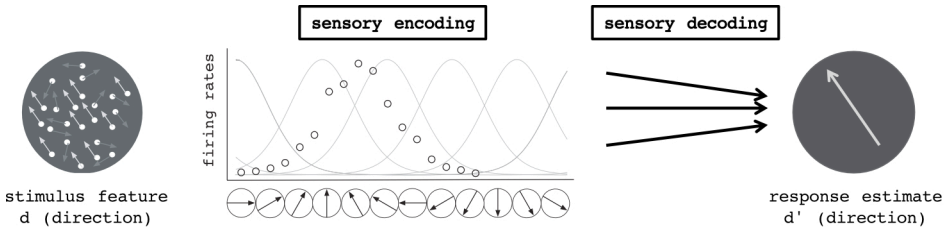


Figure 1.1 From left to right: The encoding process of a direction (d) in which a field of dots is moving, is schematically represented by a set of neuronal tuning curves. By applying a set of decoding rules, an estimate of the stimulus property d (d') is generated and used to provide a response. Adapted from Jazayeri & Movshon, 2006.

Research on perceptual decision making has focussed on the role played by each process in reaching the final outcome: understanding the neural coding used by sensory areas to represent features of the stimulus and factors that can affect such representation, leading to *biased* decisions (Gold & Shadlen, 2007). However, because this is a complex process, effects on choices of options can be due to strategies applied in order to optimise performance in one task, the expectations attached to each alternative, or even the modality used to report a choice.

In what follows I will first briefly describe the framework of one of the most popular theories in decision making, Signal Detection Theory. I will then present alternative models that have also been used to improve the understanding of how humans make use of knowledge and perceptual evidence to make a choice among alternatives. I will conclude this section with a description of models that have been developed to explain electrophysiological measurements.

It is important to note that decisions are generally made under some level of uncertainty. This is, in fact, one of the key aspects characterising *Signal Detection Theory* (SDT) (Green & Swets, 1966). The aim of SDT is to provide a quantitative framework to explain the processes underlying

decision making under conditions of uncertainty. The following provides a useful thought experiment: a participant is asked to detect a weak signal in a noisy background. The signal can either be present or absent on any given trial and the task is to respond as accurately as possible, with a *yes* or *no*. Performance on this task is limited by two factors: *internal* and *external noise* (Green & Swets, 1966; Schonhoff & Giordano, 2006). The former refers to the intrinsic fluctuations in the neural responses that can be observed even though exactly the same stimulus is repeatedly presented. External noise, on the other hand, refers to stochastic characteristics of the signal or the environment that affect the detection process. Because of those sources of noise, the task of correctly detecting a signal is probabilistic in nature: the stimulus might be present (*signal + noise*), or it might not (*noise*), and subjects might report perceiving the signal (*'yes'*), or not (*'no'*). Given these four options, the possible outcomes of a detection task can be summarised in a 2x2 table (Figure 1.2):

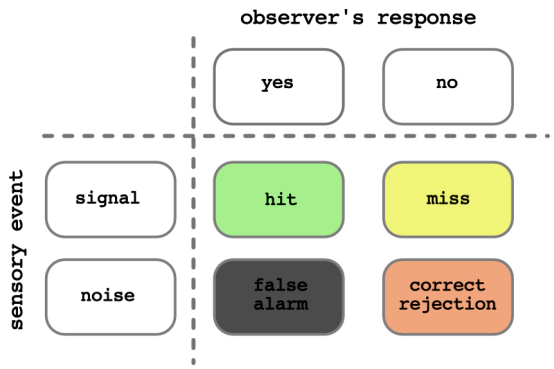


Figure 1.2 Signal detection theory framework: schematic representation of the combinations between sensory event (signal presence) and observer's response. The two outcomes in each row represent complementary events of the same probability distribution.

The '*signal + noise*' and '*noise*' states can be modelled as two probability distributions representing the internal response for each situation. In this way, the further apart these two distributions are from each other, the

higher the sensitivity in signal detection. Conversely, the closer the two distributions are along the *response* axis, the larger the overlap between them and the higher the number of errors.

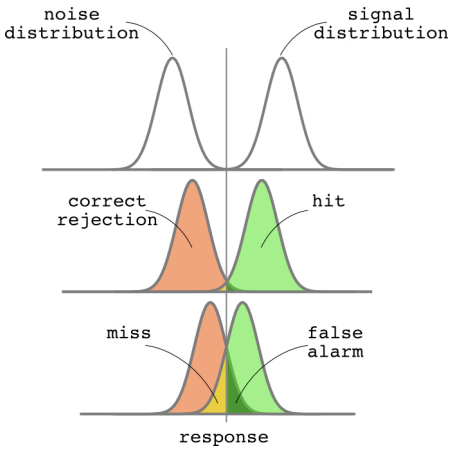


Figure 1.3 Representation of the SDT framework in probability distributions for the internal representations of "noise" (red) and "signal+noise" (green). The vertical line indicates the *criterion*, above which the subject responds "yes". The further apart the noise and signal distributions are from each other, the higher the probability of correctly detecting (hit) the presence of a signal and correctly rejecting its presence when the signal is absent. On the other hand, as the overlap between the distributions increases, the probability of false alarms (detecting a signal when is absent) and misses (rejecting a physically present signal) increases. Adapted from Gescheider, 1997

Two main factors affect the '*signal + noise*' and '*noise*' distributions: *signal strength* and a *criterion*. Intuitively, a stronger signal evokes a stronger response and this, in turn, leads to a separation between the two distributions, boosting the overall performance (see Figure 1.3).

A criterion can be thought of as a threshold value representing the willingness to report the presence of a signal. On a particular trial, if the value of the internal representation for a stimulus exceeds the criterion, then the subject responds 'yes', otherwise 'no'. Depending on the location of the criterion along the response axis, subjects will be more or less liberal in their detection: setting a low criterion means that a minimal probability of perceiving the signal is enough for producing a 'yes' response. This, in turn, results in a high proportion of false alarms. On the other hand,

choosing a more conservative approach reduces the false alarm rate, but simultaneously increases the rate of misses (see Figure 1.4). Because of the sources of noise, some kind of error will always be made: the strategy is then to adjust the criterion in such a way that the detection outcome reflects our preference for either false alarms or misses.

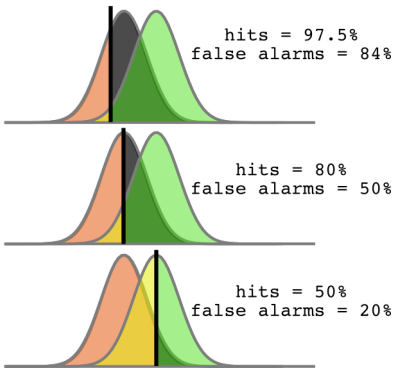


Figure 1.4 Criterion shift effects. From top to bottom: choosing a low criterion guarantees a high hit rate (almost all signals are detected), but simultaneously increases the number of false alarms (detecting a signal even when this is not physically present). When trying to reduce the rate of false alarms, the proportion of correct signal detection decreases, resulting in an increase in miss rate. Adapted from Gescheider, 1997

One of the drawbacks of SDT is that it does not take into account the accumulation of evidence over time. In order to make a decision, evidence regarding the possible alternatives often needs to be encoded and accumulated over time, until the information acquired is sufficient for an appropriate choice to be made.

An alternative to SDT is *Sequential Sampling* (Luce, 1986; Smith & Ratcliff, 2004; Townsend & Ashby, 1983; Usher & McClelland, 2001; Vickers, Caudrey, & Willson, 1971): given that the information encoded by the sensory apparatus is noisy, multiple samples of the stimulus of interest are required in order to reach a decision towards one of the possible alternatives. Sequential Sampling can be thought of as a family of models, each presenting a specific feature: evidence can be sampled in a continuous

or discrete manner, the type of accumulated evidence can vary, and so can the number of accumulating processes (see Figure 1.5).

A well-known type of sequential sampling model is the *random walk model* (Busemeyer & Townsend, 1993; Ratcliff, Smith, Brown, & McKoon, 2016). Imagine having to discriminate between two possible alternatives. The random walk model states that the evidence being accumulated consists of the difference between the features characterizing the two options. A choice is then taken when the accumulated evidence exceeds either of two possible criteria (each corresponding to one of the two alternatives). In this sense, there is a single accumulation process and the information computed is the relative difference between alternatives.

On the other hand, models known as *race models* assume that a distinct accumulation process is initiated for each possible alternative (Brown & Heathcote, 2005, 2008; Usher & McClelland, 2001). In this case, a choice is made when the level of information accumulated over time reaches one of the criterion boundaries.

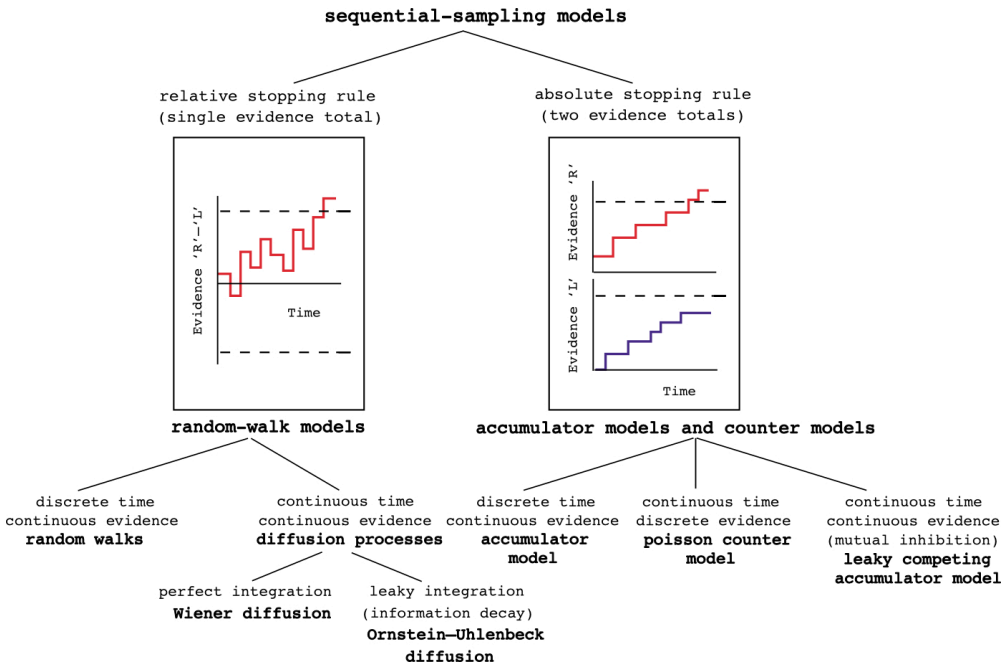


Figure 1.5 Sequential sampling models: decisions are the result of evidence accumulation up to a certain criterion. Random walk models assume a single accumulation process, whereas accumulator and counter models assume a separate process for each option and the first among these to hit a pre-determined threshold generates the final choice. Other aspects of the accumulation process determine more specific subsets of sampling models. Adapted from Smith & Ratcliff, 2004

One of the advantages of sequential sampling models is the ability to simultaneously account for reaction times and accuracy from behavioural responses. Both of these measurements are influenced by the quality of the alternatives, the instructions given in performing the decision (i.e., favouring either speed or accuracy in choice behaviour), as well as external factors, unrelated to the decision process (Gold & Ding, 2013).

Sequential sampling models highlight the relationship between activity at the neuronal level in processing information and the decision-making process: neurons fire to encode an event and this response pattern is informative of the behaviour following the event. It is known that neurons in sensory areas can be tuned to specific features (e.g., orientation selectivity for neurons in cat's primary visual cortex, see Hubel & Wiesel, 1963). This selectivity, however, does not directly account for the choice

behaviour. In fact, when making a decision, we must take into account not only the physical aspects of the options, but also factors such as history of previous choices, costs and benefits associated with each option and corresponding outcome, together with modulations due to prior knowledge and expectations. Theoretical models generally apply a decision rule to a combination of sensory and higher-order information to explain the processes involved in perceptual decision making. A key feature is the probabilistic nature of both choice and neuronal pattern encoding sensory information (Parker, 2013). In fact, presenting the same physical stimulus on more than one occasion does not necessarily lead to the same response: the choice is the product of the underlying neural activity encoding the sensory stimulus, and the patterns of activity can differ from trial to trial due to several factors (e.g., internal - properties of the sensory system - and / or external noise).

While models like those derived from the SDT appear quite abstract and are mostly aimed at explaining *behaviour*, an increasing interest in linking neuronal activity to the decision process has recently emerged. Neurophysiological studies on animals (e.g., monkeys, cats) have informed behavioural studies of perceptual decision making (Groh, Born, & Newsome, 1997; Newsome, 1989; Salzman & Newsome, 1994; Shadlen & Newsome, 2001). In a typical experiment, a monkey is presented with visual stimuli and has to perform a two alternative forced choice task while activity of single neurons from a series of brain areas is recorded. Different studies have focussed on different stages of the decision process: some were

interested in understanding the strategies involved in encoding the sensory stimulus (Newsome, 1989; Shadlen, Britten, Newsome, & Movshon, 1996), others in finding a neural correlate of the decision variable and understanding at what point in time and where in the brain a decision is made (de Lafuente, Jazayeri, & Shadlen, 2015; Shadlen & Newsome, 2001). I will first present a set of models describing neural coding of sensory stimuli, and then move to a series of studies using single cell recordings to look at the decoding and decision process.

Two main frameworks aimed at explaining the connection between behavioural performance and neural activity are: (1) *lower envelope principle* (Albright, Jessell, Kandel, & Posner, 2000; Barlow, 1995) and (2) *pooling models* (Geisler & Albrecht, 1997; Shadlen, Britten, Newsome, & Movshon, 1996; Siebert, 1965).

Imagine performing a detection task with the aim of determining whether a target is present or absent. Usually, a particular feature of the stimulus is manipulated so that different *levels* can be tested. According to the *lower envelope principle*, detection performance results from the activity of the neuron that is most selective to the feature being examined. In this scenario, when the neuron tuned to the specific feature fires, the response choice is made. On the other hand, *pooling models* generally assume that a choice results from the combination of signals in a population of neurons. One of the first studies on perceptual decision making involved tactile perception (Mountcastle, 1957): a delayed discrimination between two stimuli that differed in vibration frequency. By recording activity from

neurons in primary somatosensory cortex, a link between firing rate during sensory stimulus encoding and task performance was found: neural activity during stimulus encoding was predictive of the discrimination response on a trial-by-trial basis.

The discrimination task can be divided into several stages: processing the incoming stimuli, comparing the options, choosing one of them, and actively perform a response. By selectively looking at neuronal activity during such tasks, it has been possible to link each stage to a defined area of the brain: activity in primary somatosensory cortex (S1) reflects the encoding of the stimuli (increasing the frequency of the stimulus is accompanied by a monotonic increase in firing rate) (Romo & Salinas, 2003). Activity in secondary somatosensory cortex (S2), dorso-lateral prefrontal cortex (DLPFC), and medial and ventral premotor cortices reflects the comparison between the encoded stimuli. By modelling the firing rate of neurons in these areas as a weighted sum of the target and comparison stimuli, Romo et al., (2002) have shown that it is possible to derive what drives the neuronal response at different time points during the interval preceding the decision. During the first 100ms of the comparison stimulus presentation, firing rate of recorded neurons showed a dependency on the stimulus itself (and this was true for S1, S2, and PFC), but over the following 100ms this changed: S2 and PFC activity showed a dependency on the target stimulus. Moreover, between 200 and 300ms after the onset of the comparison, the activity in these areas resulted to be a function of the difference between the stimuli and was interpreted as reflecting the

comparison stage of the decision process. This evidence was also supported by the observation that the strength of the signal indicating the difference between the stimuli was predictive of the behavioural choice: the activity was correlated with the resulting choice behaviour in both correct and error trials.

Following the ideas developed by Mountcastle (1957), several research groups turned to the visual system to study perceptual decision making. This choice is usually justified by the fact that areas of the brain involved in processing visual information are easily accessible anatomically (occipital pole of the brain, this is true for both humans and animals), thus making recording of neural activity more suitable. Also, the organisation of the visual system and the way in which visual features are processed are well understood. In particular, visual features such as orientation (Bradley et al., 1987), contrast (Barlow, et al., 1987), speed (Liu & Newsome, 2005) and motion (Britten, et al., 1992) represent the main proxies to investigate the processes underlying perceptual decision making. Finally, compared to other sensory modalities, it is relatively straightforward to produce, control, calibrate, and deliver stimuli.

For example, Newsome et al., (1989) recorded activity from neurons in monkey area MT (middle temporal, an area of the visual system involved in the processing of moving stimuli). The animals had to fixate the centre of the screen while presented with a field of moving dots. These stimuli, also known as random-dot kinematograms (RDKs) were moving in either of two target directions: left / right of fixation point. After a certain delay, the

monkey was required to perform a saccade – a ballistic, voluntary eye movement – to a target location that corresponded to the direction of motion of the stimulus (location left / right of fixation; Figure 1.6). Importantly, the recorded neurons were selectively tuned to either of the two possible target directions of motion. Such neurons preferentially respond to one stimulus direction, say leftward, and are inhibited by motion in the opposite direction. Because of this specific response preference, the activity pattern for the two directions was expected to differ and to predict the overall performance in the task. Indeed, higher neural activity was recorded when the direction of motion within the stimulus was in the preferred, rather than the *null* direction, which in turn resulted in a saccade towards the target indicating the preferred direction. Performance, therefore, could be predicted by the firing rate of a relatively small group of neurons.

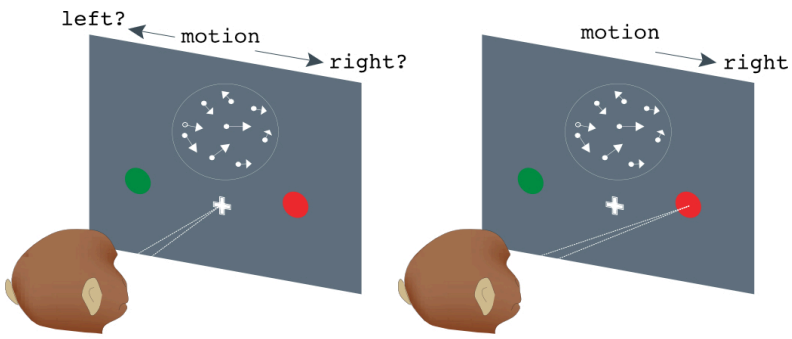


Figure 1.6 Visual motion discrimination task: a perceptual judgement based on the sensory evidence represented by the field of moving dots is provided via a saccadic eye movement to one of two possible targets. Adapted from Heekeren et al., 2008

Another interesting factor investigated by Newsome and colleagues was the effect of uncertainty on the decision process. The proportion of dots in the RDK moving coherently – together, in the same direction – was

manipulated: when a small number of dots move in the same direction among a large number of dots moving in random directions, the effective *signal strength* is low. This in turn, makes it harder to perceive the direction of the stimulus. This relationship was well described by the activity patterns recorded: higher firing rates were observed for stimuli with high coherence compared to stimuli with low coherence levels. Moreover, Shadlen & Newsome, (1996, 2001), as well as Britten, et al., (1996), reported that activity recorded from area MT represents the encoding of the sensory information, whereas the decision results from the difference in population activity computed in the *lateral intraparietal area* (LIP). This signal is then sent downstream to the *frontal eye field* (FEF) and *dorso-lateral prefrontal cortex* (DLPFC) to perform the response (here a saccade towards one of two possible target locations). The interplay between these brain areas and the specific properties represented by their activity patterns is comparable to the system described in the tactile domain by Mountcastle, Romo, and colleagues.

Experiments investigating the neural basis of perceptual decision making in human subjects followed the findings obtained from monkey's neural activity recordings. In an important early study, Heekeren, et al., (2004) presented observers with images of houses and faces while their brain activity was measured via functional magnetic resonance imaging (fMRI). The task was to judge whether each of the presented images belonged to the 'house' or 'face' category. Task difficulty was systematically manipulated: fast Fourier transforms of each image was computed to

obtain a magnitude matrix representing the average magnitude of all images and a phase matrix for each stimulus. By combining these with the average magnitude, images with varying levels of noise were generated. (Figure 1.7).

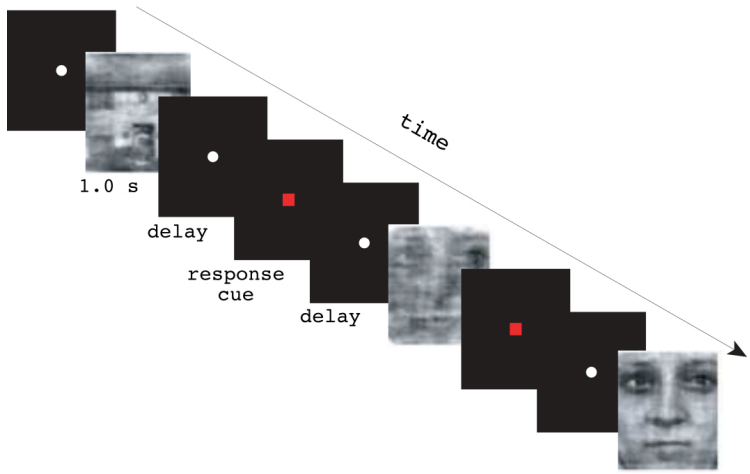


Figure 1.7 Task sequence: an image with a varying degree of noise is presented to the observer that has to categorise it as either a face or a house after a delay period has elapsed. Adapted from Heekeren et al., 2004

An important starting assumption is that a brain area known as *fusiform face area* (FFA) is involved in preferentially processing faces, whereas an area known as *parahippocampal place area* (PPA) shows enhanced activity for landscapes and buildings. The authors corroborated this idea by showing that the activity measured in these two areas reflected the encoding of the images: whenever a face was presented, FFA was active, whereas PPA showed increased BOLD signal for images of houses. Moreover, the strength of the response in each area was correlated with the amount of distortion of each visual stimulus: the higher the noise added to the images, the lower the intensity of the recorded signal. This is analogous to the encoding of visual direction of motion information described in area MT of monkeys. Another interesting finding presented by Heekeren and

colleagues consists in a correlation between the absolute difference of FFA and PPA signal and the BOLD signal measured in DLPFC (Figure 1.8A). Not only did DLPFC activity represent the encoding difference between the two image categories, but it also correlated with the amount of noise within each image: the less disrupted the image, the greater the response in this region. The authors therefore suggested that activity in DLPFC reflects the integration of sensory evidence provided by areas at earlier stages in the visual pathway and, given that behavioural performance could be predicted by BOLD signal measured in DLPFC, they also concluded that decisions are based on such activity.

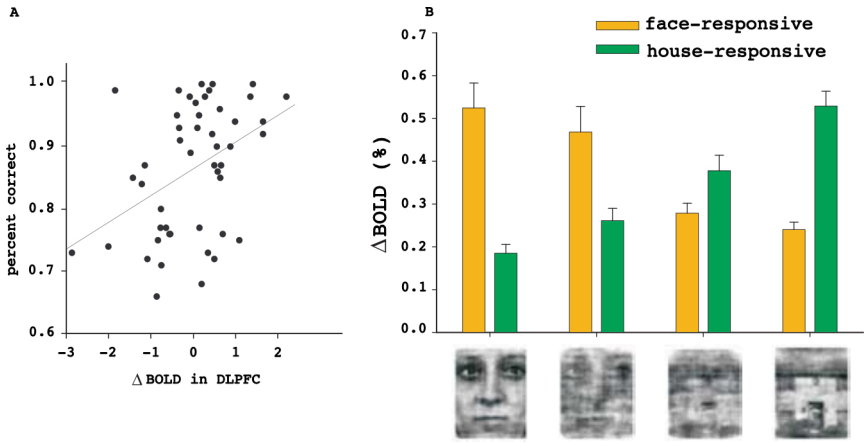


Figure 1.8 Perceptual decision-making in the human brain. A) Behavioural results showing that threshold identification of both houses and faces required approximately 45% of noise in the stimuli. B) Percentage of BOLD change response measured in areas FFA and PPA for images of faces and houses at different noise levels. Adapted from Heekeren et al., 2004

To study the cortical networks involved in perceptual decision making processes at higher temporal resolution, Siegel and colleagues (2011) used magnetoencephalography (MEG) in a series of studies. Neuronal oscillations were measured while subjects judged whether the coherent direction of motion of an RDK was upwards or downwards. The authors observed an increase in the high-frequency (gamma band) response during

the presentation of the stimulus in visual areas. Interestingly, while striate visual areas showed a sustained response throughout the display of the RDK, response in extrastriate area MT was modulated by the coherence at which the RDK was computed (the stronger the motion signal provided by the stimulus, the stronger the increase in gamma-band activity). Human area MT and its gamma-band activity therefore appear to be implicated in the perceptual decision making stage of encoding sensory evidence. In a similar task (Donner et al., 2009), where subjects had to detect the presence of coherent motion in an RDK, gamma-band activity of primary motor (M1) and premotor (PMd) areas increased in a gradual fashion (and during the delay between stimulus presentation and response). The laterality of this activity was predictive of response execution and it was suggested to action selection. Finally, by comparing activity evoked during correct and incorrect trials, Siegel and colleagues showed that responses in the beta frequency range of the DLPFC predicted accuracy of behavioural performance: increased beta-band activity correlated with correct detection for trials in which the signal was present as it was for correct rejections for trials in which the signal was absent. Thus single-unit recordings in monkeys and neuroimaging studies in humans present a common network of brain areas involved in processing the different units constituting a perceptual decision.

A yet different framework within which perceptual decision making can be interpreted is known as *embodied* decision making, or *embodied choice* (Lepora & Pezzulo, 2015). This approach highlights the importance of action in modulating decisions based on sensory information. Lepora & Pezzulo

nicely illustrate this idea by comparing three models of perceptual decision making: a *serial* model (action follows a decision), a *parallel* model (decision and action happen simultaneously), and an *embodied* model (the decision process is influenced, via feedback, by the action taken). The authors used data from a two-alternative forced choice task (2AFC) where subjects were required to make a choice between two alternatives and move a cursor onto one of the two targets during the decision process. By comparing the simulated responses from each of the models mentioned above with data obtained from the subjects, the authors concluded that the action performed plays a fundamental role in shaping the decision itself: initiating a motor response while sensory evidence is being accumulated implies a *cost* that in turn gives rise to a *commitment* factor, both modulating the decision process. The framework proposed by Lepora & Pezzulo (2015) is of relevance with regard to the data presented in Chapter 5 of this thesis, where two perceptual decision-making processes, differing mainly in response action, are compared.

1.2 *Anatomical and Physiological Organisation of the Visual System*

Vision is the process of transducing raw input signals in the form of light into information that allows us to navigate and interact in the environment. The mechanisms underlying vision have been extensively studied using a vast series of techniques, from behavioural experiments to

neurophysiological measurements, in both animals and humans. This has allowed a growing understanding of the processes involved at each stage of the visual system, relating features of the visual stimulus to behavioural responses through activity patterns measured in the brain. In particular, visual motion is a salient, low-level visual feature with a high impact on interactions with the world around us.

Objects move in the world and, as we look at them, they also move on our retinas (the most outer part of the brain located at the back of the eyes). In addition, even when objects in the visual environment are static, our eyes move continuously, making motion a fundamental feature of the interaction between perceiver and information to be perceived. In what follows, I will describe the processing of visual information in the visual cortex, presenting evidence from electrophysiology studies in animals as well as neuroimaging studies in human subjects (note that due to the similar organisation of the visual system in higher mammals such as cats, monkeys, and apes, findings from one species are usually transferred across the others – see Sereno & Allman, 1991). I will start with a summary of the information analysis carried out from the retina to the visual cortex through specific pathways. Subsequently, I will focus on the processing of visual motion information in specific areas of the visual cortex.

1.2.1 From the Retina to Early Visual Cortex

The processing of visual information starts when light hits the eyes, whose lenses focus it onto the *retina*. The retina is a layered structure

located at the back of the eye, and each layer is composed of cell types playing specific roles in extracting information from light inputs (Figure 1.9). At the far back of the retina lies the *photoreceptors* layer. These cells are divided into *rods*, active at low levels of light and specialised in processing coarse spatial information, and *cones*, processing fine spatial information and colour (Masland, 2001). A layer of *bipolar cells* transfers information processed from either rods or cones to the next layer, that of *retinal ganglion cells*. It is worth noting here the peculiar organisation of the retina: the receptors constitute the most external layer of the retina while the others develop towards the inside of the eye. This means that for the light to hit the receptors, it first needs to travel through all the wiring, to then be processed here, and afterwards sent 'backwards', layer by layer until the information is transduced by the ganglion cells and from here sent out via their axons. Aside from photoreceptors, bipolar, and ganglion cells, two other types of cells can be found to form two distinct horizontal sections: *horizontal cells* between photoreceptors and bipolar cells, and *amacrine cells* between bipolar and ganglion cells (Kolb, Linberg, & Fisher, 1992). One of the roles of these horizontal intra-layers is to allow for *lateral inhibition*: the activity of one cell inhibits that of its neighbours leading to a '*centre – surround*' organisation of the stimulated visual area. While the *visual field* corresponds to the portion of the outside world that can be captured by sight at once, the region of the retina that is stimulated by incoming light causing retinal ganglion cells to fire, is called the *receptive field* (RF). The *centre – surround* organisation of cells' RFs, observed in animals such as cats

and monkeys, as well as in humans, is fundamental in detecting contrast, that is light discontinuities that are usually processed in terms of edges. In fact, when light falls in the central region of a centre – surround organised receptive field, the cell is either excited or inhibited, and the surrounding area exhibits the opposite pattern in response to light (Kuffler, 1953).

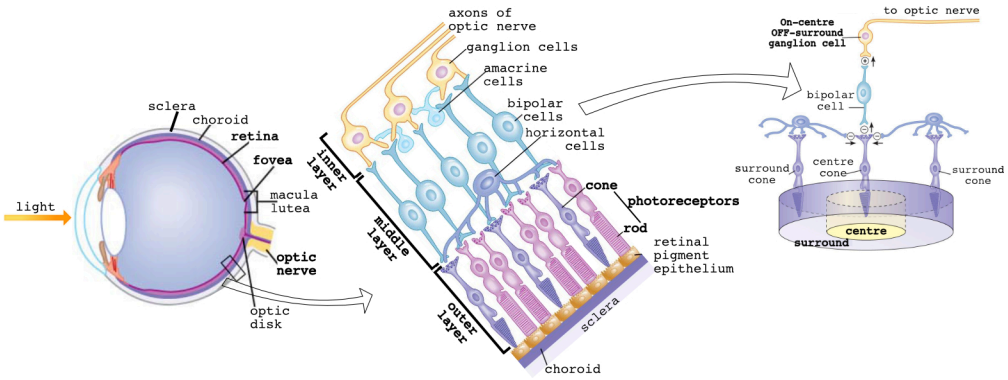


Figure 1.9 Retinal organisation and receptive fields: The light enters the eye through the pupil and hits the retina located at the back of the eyeball. A schematic representation of the layers composing the retina shows the photoreceptors at the far back, while ganglion cells constitute the most inner layer of the retina. The receptive field organisation of an on-centre off-surround ganglion cell is shown in the diagram on the right. Adapted from Pearson Education Inc., 2011

Cells with a centre – surround organisation of the receptive field are also distinguished into '*on*' and '*off*' – centre. Light hitting the centre of an ON – cell will result in an excitatory response (i.e., rapid firing activity), whereas in the case of an OFF – cell this will result in an inhibitory response (i.e., suppressed release of action potentials) (Kuffler, 1953). It is therefore clear how this specific composition aims at detecting edges and changes in light distribution across the field, even though the idea of RFs has been extended to other stimulus properties such as orientation, motion selectivity, and colour-opponency, as well as more complex combinations of low-level features such as images of faces, body parts, and objects.

Signals from receptors and bipolar layers are transduced by the retinal ganglion cells and information is sent out to the brain in form of electrical spikes through the *optic nerve* (collection of ganglion cells axons) (Kandel, Schwartz, & Jessell, 2000). A collection of bundles leaves each eye and unites at the level of the *optic chiasm*. Here information is reorganised so that the left and right portions of the visual field are kept separated (Figure 1.10). Nerves from the nasal portions of the retinas undergo *decussation* (i.e., crossover of fibres to the opposite side of the brain compared to their side of origin), leading to the left visual field being processed in the right hemisphere of the brain and the right visual field in the left hemisphere (Purves et al., 2004). Leaving the optic chiasm, the reorganised bundles of axons forming what is now called the *optic tract* reach the *Lateral Geniculate Nucleus* (LGN), whereas a small subset terminates in the *Superior Colliculus* (SC), a structure of the midbrain (Hubel, LeVay, & Wiesel, 1975).

The LGN is organised into six layers, each receiving inputs from pathways defined by cell types that process specific features of the visual input (Callaway, 2005; Wiesel & Hubel, 1966). The dorsal layers of the LGN receive inputs from the *magnocellular pathway*: projections of *parasol cells* that present large RFs originating by the pooling of information from rods and cones by horizontal and amacrine cells, respond to low contrast stimuli, and have a *transient* (high temporal frequency) firing pattern, favouring processing of motion information. On the other hand, the *parvocellular pathway* sends its projections to the ventral layers of the LGN. The retinal ganglion cells constituting this pathway are called *midget cells* and are

characterised by small RFs (the signal is passed to the parvocellular cell from the photoreceptor – rod or cone – via a bipolar cell), respond to high contrast stimuli, and have a *sustained* (low temporal frequency) firing pattern, favouring processing of fine spatial frequency information and colour (Callaway, 2005; Leventhal, Rodieck, & Dreher, 1981). A third pathway projects to cells found in between magno- and parvocellular layers of the LGN. The *koniocellular pathway* originates from *bistratified cells*, presenting response properties that lie somewhere in between those of the parasol and midget cells (Hendry & Reid, 2000).

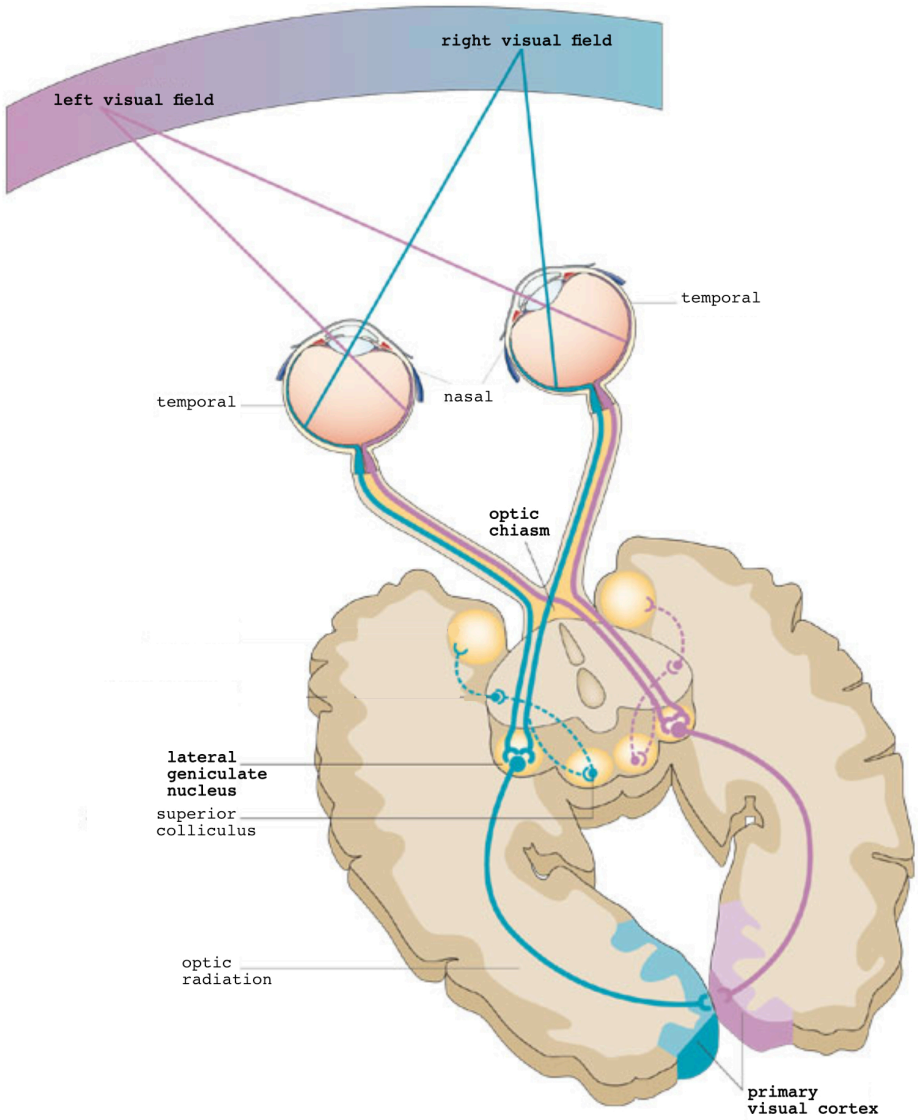


Figure 1.10 Segregation of visual information into the visual pathways: Regions of both the right and left visual fields are represented in the temporal and nasal portions of the retina. Fibres from the nasal portion of the retina decussate at the level of the optic chiasm allowing the information from the left visual field to be represented in the right visual cortex and information from the right visual filed is represented in the left visual cortex. Adapted from Hannula, et al., (2005)

The segregation into left and right visual fields is maintained through the layers of the LGN, giving rise to a *retinotopic map*: points projected close onto the retina are also represented close to each other onto the LGN (Chen, et al., 1999; Schneider, Richter, & Kastner, 2004). The topography of the retina is not the only aspect to be preserved in the LGN: the *on – off* or *centre – surround* property of the receptive fields of the ganglion cells in the retina can be found in neurons of the LGN (Hubel & Wiesel, 1961).

The *optic radiation* constitutes the group of axons leaving the layers of the LGN and projecting to the visual cortex, at the back of the brain. A subset of fibres connects the LGN to the *Superior Colliculus*, a region of the midbrain primarily involved in the processing of eye movements (Dorris, Paré, & Munoz, 1997; Lee, Rohrer, & Sparks, 1988; Mohler & Wurtz, 1977). The passing of information from the retina to the primary visual cortex through the LGN is called *retino-geniculo-striate pathway* (Schein & de Monasterio, 1987). Importantly, the retinotopic organisation observed in the LGN is preserved in the visual cortex, presenting an interesting property called *cortical magnification* (Figure 1.11). In fact, at the level of primary visual cortex (V1), the *fovea* (central portion of the retina) is overrepresented compared to the periphery: the size of the RFs changes accordingly, being small at the fovea and larger in the periphery (Daniel & Whitteridge, 1961).

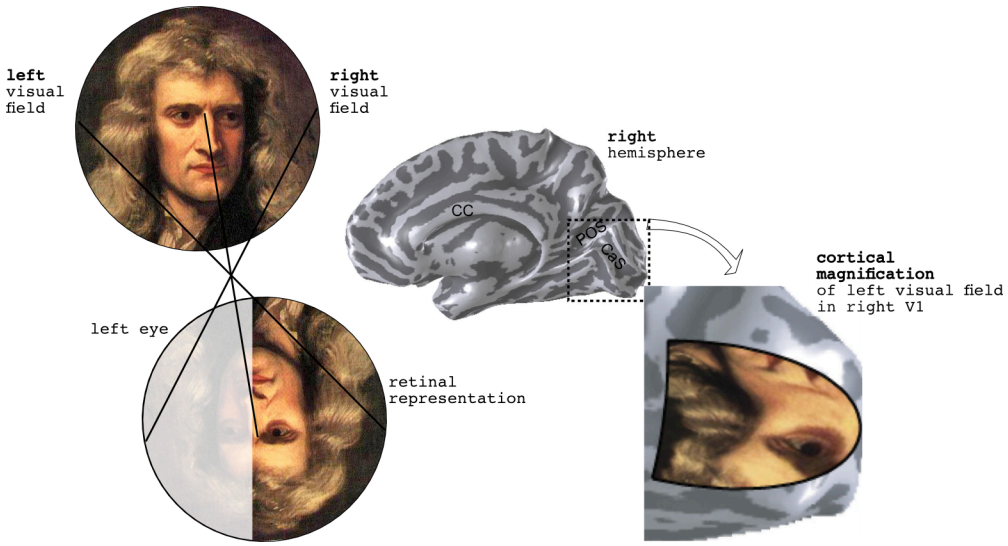


Figure 1.11 Cortical magnification: On the left the left eye retinal representation of Sir Isaac Newton with highlighted left visual field. On the right, the mapping and transformation of such portion of visual field onto the right primary visual cortex. Not only the representation is inverted, but the centre of the visual field (the area around the nose and eye of the image) is over-represented. Adapted from Wandell, Dumoulin, & Brewer, 2007

Starting from V1, a series of brain areas can be identified and distinguished according to the type of visual information they process or are

most responsive to. Given that the focus of this thesis is the processing of visual motion, a feature for which selectivity (usually described as *direction* selectivity) first emerges in area V1 in higher mammals (Hubel & Wiesel, 1968), I will now focus on the pathways underlying this mechanism.

1.2.2 V1

The primary visual cortex, located in the *calcarine sulcus* (Figure 1.12), a region of the occipital cortex of the brain, is organised into layers. Cells from each layer present specific properties that have been extensively investigated in both animals and humans (Hubel & Wiesel, 1963; Hubel & Wiesel, 1959, 1969). Hubel & Wiesel, (1968) recorded extracellular activity of neurons in primary visual cortex of monkeys while presenting a series of light patterns. They were able to identify three main types of cells:

- *simple cells*: found primarily in layer IVCα, present elongated on-off, centre-surround RFs, show orientation and direction selectivity as well as preference for stimulus position in the visual field, and are organised into functional orientation columns (i.e., neurons preferentially responding to stimuli with specific orientations are grouped together)
- *complex cells*: mainly found in layer IVB, respond to signals coming from both eyes, also show orientation and direction selectivity, but do not show any preference for the location of the stimulus in the receptive field

- *hypercomplex cells*: also called *end-stopped cells*, present a decrease in response as the physical length of the stimulus increases

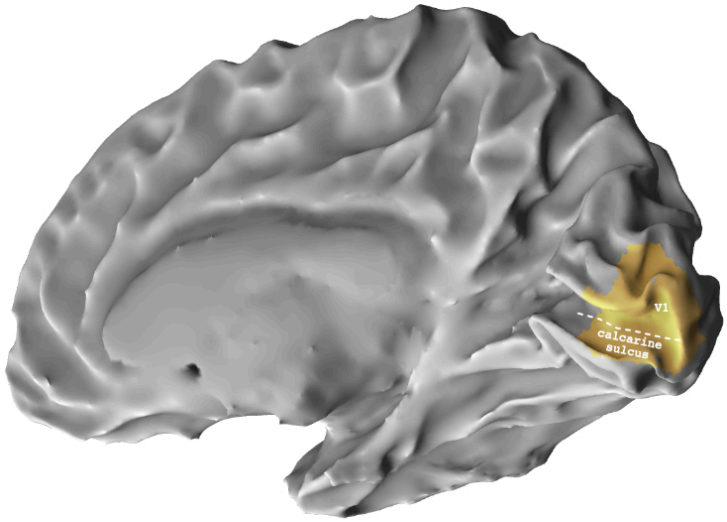


Figure 1.12 The calcarine sulcus and area V1. Grey-scale, 3D rendering of the white-matter surface of one of the participants in my experiment (cortex, right hemisphere). Highlighted area, functionally derived location of primary visual cortex, V1.

From these pioneering studies in monkeys and cats, a series of further investigations applying both electrophysiological measurements in animals (e.g., Bonhoeffer & Grinvald, 1991; Gardner, et al., 1999; Gur, Kagan, & Snodderly, 2005) and neuroimaging techniques in humans (Boynton, 2005; Freeman, et al., 2011; Yacoub, Harel, & Ugurbil, 2008), were put forward, revealing the inner structure and organisation of V1. In particular, these focussed on the processing of orientation, leading to concepts and models that now represent milestones of visual neuroscience: *orientation columns, hypercolumns,* and *ocular dominance columns* (Hubel & Wiesel, 1959, 1962, 1969, 1974, 1977). With regard to visual motion, although cells sensitive to direction have been observed for some animals (De Valois, William Yund, & Hepler, 1982; Hubel & Wiesel, 1962), a specific and clear organisation of such selectivity into columns, at least in primate area V1, has failed to emerge (Lu, et al., 2010). However, it has been shown that cells in V1 tuned to motion

project to other visual areas where a more sophisticated analysis of this feature is carried on. In particular, a direct connection between V1 and the *middle temporal area* (MT or V5) in extrastriate cortex, observed in electrophysiology and lesion studies with macaque monkeys, and later corroborated by human studies with focal brain damages (e.g., Milner & Goodale, 2008), constitutes what is generally known as the *dorsal pathway*, primarily involved in the processing of motion (as opposed to form, a feature processed in the *ventral pathway*) (Ungerleider & Mishkin, 1982). Interestingly, projections from V1 to MT can be observed both directly and indirectly, passing through areas V2 and V3 before reaching the middle temporal area (Maunsell & van Essen, 1983). It is also worth noting that the segregation between dorsal and ventral pathways derives from the distinction between parvo- and magnocellular pathways in the LGN (Dacey & Petersen, 1992).

1.2.3 V2

Area V1 projects to other extrastriate areas, and here the same retinotopic organisation that is observed in primary visual cortex and LGN emerges: the visual field is orderly represented, with objects close together in the environment being represented close together in the visual cortex. An interesting fact observed in macaques (Desimone & Gross, 1979; Gattass, Gross, & Sandell, 1981) as well as humans (Amano, Wandell, & Dumoulin, 2009; Harvey & Dumoulin, 2011; Smith, et al., 2001) by studying the properties of brain areas involved in processing visual inputs is that the

receptive fields tend to become larger as we move to areas higher up in the hierarchy (Smith et al., 2001). This property has been suggested as one of the factors that allow to explain the selectivity in processing visual motion in certain areas of the visual system (i.e., *middle temporal area MT/V5* - Andersen, 1997).

Area V2 is composed of two portions: one located ventral to V1, the other dorsal to it. Each section processes a quadrant of the contralateral visual field, respectively (DeYoe et al., 1996; Dougherty et al., 2003; Engel, et al., 1997; Sereno et al., 1995).

Cells in V2 respond to features similar to those observed in V1, but also to more complex combinations of such features (e.g., combination of orientations, Hegdé & Van Essen, 2000). Response selectivity in V2 cells has been reported to also be organised into a columnar structure, with distinctive thin, thick, or pale stripes (Livingstone & Hubel, 1984), each engaged in processing different visual properties (Roe & Tso, 1995). An interesting property that has been reported is that of direction selectivity: while a map for preferred orientation has been observed in V2 as in V1, one for preferred direction emerged from the thick stripes of this visual area in macaques (Lu et al., 2010). From this structure, projections to areas V3 and MT have been observed (Felleman & Van Essen, 1991), providing support for a pathway preferentially involved in processing visual motion information.

1.2.4 V3

V3 is also divided into a ventral and dorsal part, each located ventral and dorsal to the respective ventral and dorsal parts of V2 (Shipp, et al., 1995). Projections from ventral V2 to ventral V3 move onto a visual area called V4, whose properties are still under discussion (Bartels & Zeki, 2000; Tootell & Hadjikhani, 2001; Zeki, 1983), but generally agreed to be involved in processing colour and be part of the ventral visual pathway. Those projections from the dorsal part of V2 to the dorsal part of V3 end up in area MT/V5 (Felleman & Van Essen, 1991). Dorsal V3 contains the representation of the lower visual field, whereas ventral V3 contains that of the upper visual field. In terms of tuning properties of neurons, work on macaque monkeys showed that orientation selectivity is still present at this point, although the receptive fields of cells are broader than those found in earlier visual areas (Baizer, 1982; Gegenfurtner, Kiper, & Levitt, 1997; Zeki, 1978). From monkeys' neurophysiology an area with a specific retinotopic map, distinct from that of ventral and dorsal V3 has been reported and named V3A (V3 Accessory area –Van Essen & Zeki, 1978). The selectivity of this area, later observed also in humans, appears to be the processing of orientation and direction, although the latter was not a selective property in monkeys (V3 in monkeys resulted to be much more orientation and direction selective than area V3A – Tootell et al., 1997).

Although the response properties and structural organisation of this area are still controversial (Lyon & Connolly, 2012), the connection between V3

and the *middle temporal area* is well described (Ungerleider & Desimone, 1986) and supports the idea of a visual motion pathway.

1.2.5 V5 / MT and beyond

Area MT of macaque monkeys is located in the caudal part of the *Superior Temporal Sulcus* and presents a retinotopic map of the contralateral visual hemifield (Dubner & Zeki, 1971; Felleman & Van Essen, 1991; Van Essen, Maunsell, & Bixby, 1981). Response properties of neurons in this visual area were first investigated by recording extracellular activity in monkeys presented with moving stimuli of different shapes and sizes (Dubner & Zeki, 1971; Zeki, 1974). By systematically changing the features of the targets, two main types of cells with very specific response properties (Figure 1.13A) were observed: those that responded to any type of stimulus moving inside the receptive field of the cell, with a subset of them being also direction selective, and those that responded to a certain combination of direction of motion and shape (Zeki, 1974). Looking at the functional organisation of this area, columns of preferred direction were observed in macaque MT (Albright, Desimone, & Gross, 1984; Born & Bradley, 2005; Cao, Mingolla, & Yazdanbakhsh, 2015; Dubner & Zeki, 1971; Lisberger, Morris, & Tychsen, 1987). Interestingly, the order in which direction preferences is arranged shows a continuous sampling of directions interspersed with a sudden jump of 180 degrees from one column to the next one (Figure 1.13B). This specific setup of adjacent columns presenting preferences for opposite

directions of motion has been referred to as axis of motion columns (Albright et al., 1984; Diogo, et al., 2003).

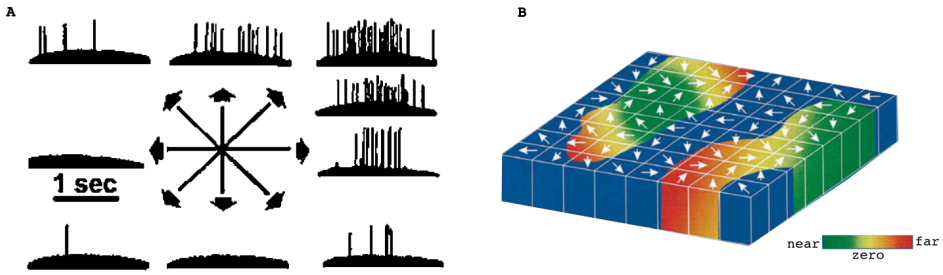


Figure 1.13 A) Neuronal responses to different directions of motion. Arrows at the centre indicate the presented direction of motion whereas plots on the outside represent spikes measured from the neuron. Preferred direction is represented by an increased amount of spikes. Adapted from Dubner & Zeki, 1971. B) Representation of the columnar organisation of area MT. Arrows indicate preferred direction of motion for neurons in each column, whereas colour indicates selectivity for disparity as indicated by the colour bar legend on the right. Adapted from DeAngelis & Newsome, 1999

Moreover, the activity in this area has been linked not only to processing of visual motion, but also to its *perception*. In a series of studies with macaque monkeys, Salzman and colleagues trained the animals in performing the RDK visual motion discrimination task shown in Figure 1.6. Activity from ensembles of MT neurons was recorded while the monkeys performed the task. The coherent motion of the RDK assumed either the direction preferred by the ensemble that was recorded, or its null direction. In order to investigate a causal link between the firing rate of MT neurons during encoding of sensory information and construction of a percept of such stimulus, the authors decided to deliver a small current to the ensemble under observation while the RDK was displayed. By using the monkey's performance as a measure of 'perception', Salzman et al. found that monkeys reported perceiving the preferred direction of the recorded ensemble of neurons more often when this was stimulated then when no stimulation was applied. This was true even when the direction of the RDK

was set to the null direction of the MT ensemble of neurons (Salzman, 1992; Salzman, Britten, & Newsome, 1990).

One key feature of neurons in area MT is the size of their receptive fields: this is larger than cells found in earlier visual areas such as V1 (Churchland, et al., 2005; Felleman & Kaas, 1984; Mikami, Newsome, & Wurtz, 1986) and has been suggested to explain the selectivity for the direction of the global (overall) motion of spatially extensive objects in the visual field. However, further studies have reported that area MT might be selective for processing complex patterns of motion (Movshon, et al., 1985) rather than simple components of direction of motion that can be computed through properties of RFs in earlier visual areas (e.g., V1 – Born & Bradley, 2005; Newsome, 1989).

Human area V5/MT has been studied by means of neuroimaging techniques such as functional magnetic resonance imaging (fMRI). Here human subjects are generally presented with moving stimuli while lying in the scanner and are required to either perform a relevant task (e.g., motion direction discrimination), or an unrelated task to ensure subjects attend to the target (Huk & Heeger, 2002). A classical contrast used in the early studies was to compare activity measured during perception of visual motion with that measured during presentation of static stimuli (Watson et al., 1993). A statistically significant increase in Blood Oxygen Level Dependent (BOLD) signal for the condition of interest (i.e., presentation of visual motion) in area MT/V5 was presented as supporting evidence of the

selectivity in processing visual motion of the *middle temporal area* in human subjects (Watson et al., 1993).

Neurons from area MT project to other fronto-parietal areas such as the *medial superior temporal* area (MST), the *lateral intraparietal* area (LIP), the *frontal eye field* (FEF), the *superior colliculus* (SC). While MST has been reported to be involved in processing 'optic flow' information, consistent with the patterns of complex image motion produced by ego-motion through the world (Duffy, 1998), LIP, FEF, and SC are involved in planning and generation of eye movements (for a review, see Kowler, 2011). The response properties of these areas have been well studied and have helped understanding the neural basis of the processes involved in perceptual decision making.

Indeed, as mentioned in the previous section, a typical task in this field of research consisted in presenting the subject (e.g., a monkey) with a field of randomly moving dots. A proportion of these dots would move coherently in one direction (e.g., upwards or downwards) and the task of the monkey was to make a saccadic eye movement in the direction it perceived the coherent motion (see Figure 1.6). The researcher monitored brain activity by recording from neurons in specific sites of the brain (i.e., MT, MST, LIP, FEF, SC) and correlated neuronal firing rate and patterns with behavioural performance of the animal. In this way, it was possible to link an event during the task with the corresponding activity in the brain area under investigation. These data were then used to inform computational models

of perceptual decision making, building a relationship between behaviour and neuronal activity.

In the next section I will focus on computational models of perceptual decision making that are based on the neural response properties of the visual system highlighted in this section.

1.3 Reference Repulsion and Biases in Visual Motion Perception

As described in Section 1.1 (Perceptual Decision Making), the process of perceptual decision making can be described as an encoding stage, followed by one or more decoding stages, and terminating in a response choice. Given this sequence of processes, it seems reasonable to think that modulations of the decision process can arise at any of the stages. Several studies have focussed on the way in which sensory information is encoded, highlighting the fact that misperception of the target can lead to *biased* decisions. In this view, the biases arise in the formation of a dependency between the sensory neuron response and the stimulus (Eagleman, 2001). On the other hand, some studies have suggested that it is the strategy used in interpreting the sensory representation encoded in the first stage that causes *biased* decisions (Gregory, 1973, 1997). In what follows I will describe a particular perceptual error, *reference repulsion*, whose origin is still under debate. I will then present studies that investigated this phenomenon from a *sensory* perspective – and others that rather suggest a top-down modulation, i.e. due to a strategy chosen to optimise performance.

Reference repulsion consists in an overestimation of the angular difference between the direction of an explicit reference and that of a motion stimulus (Rauber & Treue, 1998). This effect was observed by asking subjects to perform an identification task (Wickens, 2002): after being presented with an RDK and a static reference line (Figure 1.14), they had to report whether the field of dots moved more clockwise (CW) or anti-clockwise (CCW) with respect to the orientation of the line (Rauber & Treue, 1998). The reference could either be presented together with the moving stimulus or immediately after (Rauber & Treue, 1998).

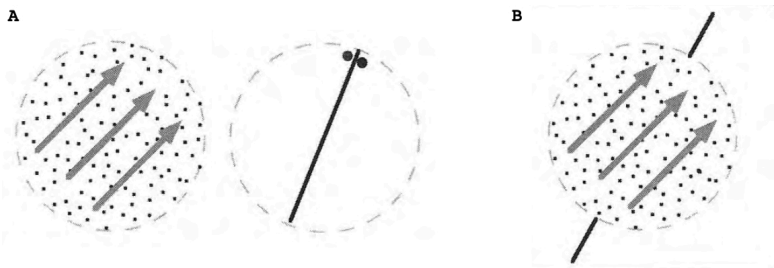


Figure 1.14 Representation of the task sequence: A) The presentation of an RDK is followed by that of an oriented line. The direction of the RDK (represented by the grey arrows) is then compared against the one of the black line in the following screen. B) In a different setup, the RDK is simultaneously presented with the oriented line and participants have compare the two directions during the same presentation. Adapted from Rauber & Treue, 1998

By plotting the angular difference between the reported direction of motion and the direction of the stimulus, as a function of the presented direction, the authors observed systematic deviations of the estimates of up to 9° for stimulus directions that differed from cardinal and diagonal axes (here deviations resulted to be close to zero, Figure 1.15A). Interestingly, when both RDK and oriented line were presented simultaneously (Figure 1.14B), the distribution of errors assumed a different pattern: subjects overestimated the angle between the stimulus direction and the horizontal axis (Figure 1.15B). Considering the circular aperture comprising the RDK,

estimates of directions pointing towards the top half of the aperture presented an upward bias, whereas estimates of directions pointing towards the bottom half of the aperture showed a downward bias. This effect has been explained in terms of neural population coding: the most informative neurons in discriminating a direction of motion against a reference are those whose tuning functions peak slightly away from the actual reference direction (Hol & Treue, 2001; Regan & Beverley, 1985).

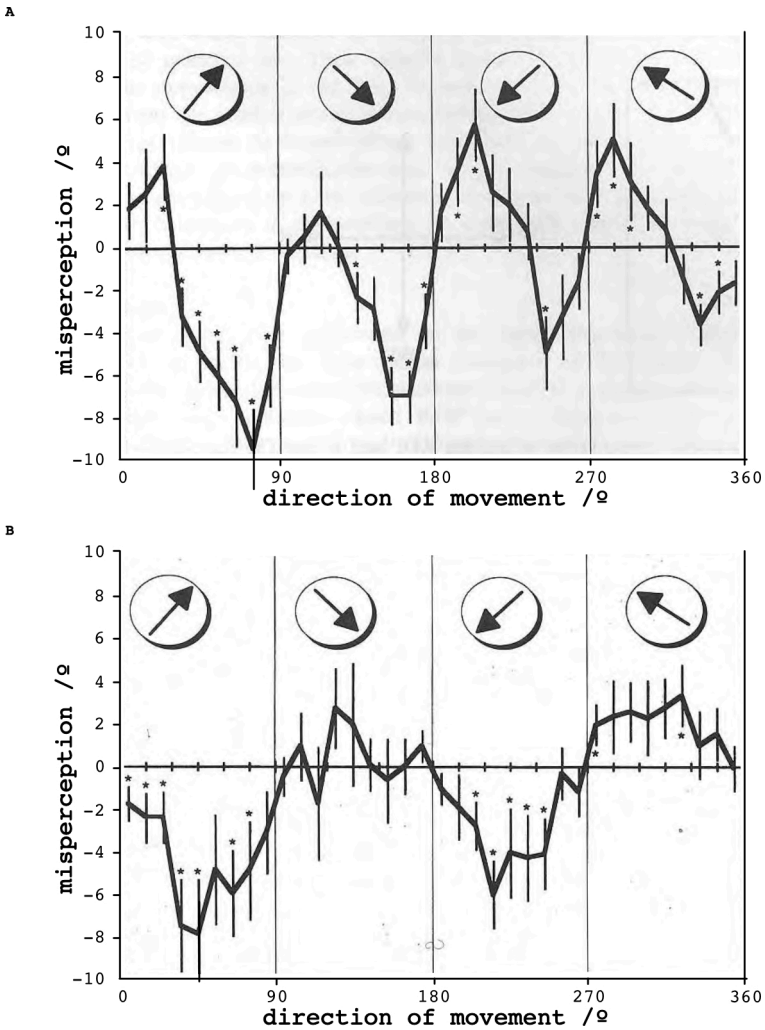


Figure 1.15 Reference repulsion effect. A) Condition where a reference follows the moving stimulus. Deviations between reported and presented directions of motion as a function of presented stimulus motion. Positive values on the y-axis indicate directions were over-estimated in a CW direction. B) Condition where the moving stimulus and the reference are presented simultaneously. Axes conventions as in A. Adapted from Rauber & Treue, 1998.

In fact, the response profile of a neuron tuned to a specific visual feature is often described as a Gaussian probability density function (Henry, Dreher, & Bishop, 1974; Hubel & Wiesel, 1968; Wandell, 1995), with the mean response corresponding to the feature value the cell is most sensitive to and the variance defined by internal fluctuations in the firing rate of the neuron. The fundamental assumption at the core of this approach is that tuning functions describing neuronal responses to different stimulus features (such as spatial orientation and direction of motion) can be readily compared in a meaningful manner. The flexibility in this type of comparison is of central importance in explaining reference repulsion effects. Thus, as nicely described by Hol & Treue, (2001), if we imagine the representation of the sensory evidence as a difference of two Gaussian distributions (Figure 1.16A), one centred at the direction of motion of the stimulus, the other at the orientation of the reference, we can easily see that the highest amount of information comes from activity encoding directions slightly off the two targets (Hol & Treue, 2001; Jazayeri & Movshon, 2006). However, when the angular difference between reference and motion is '*large*', approximately 180 degrees apart, this is not the case (see Figure 1.16B). In fact, when performing a *coarse* discrimination, such as having to decide whether

motion was up or down, the most informative neurons are those tuned specifically to the direction of motion presented (Hol & Treue, 2001).

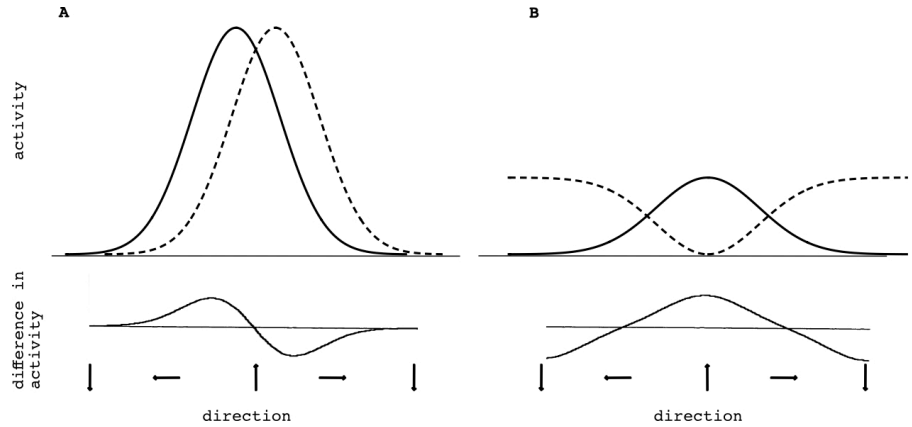


Figure 1.16 Discrimination versus detection principle. A) Representation of neural population response for upward-left (solid line) and upward-right (dashed line) motion. Discrimination is provided by the difference between the two curves (lower panel). The largest difference in response occurs around neurons tuned half the width of the curve away from the stimulus direction. B) Representation of neuronal population response to stimuli moving up (solid line) and down (dashed line). The detection of the presented direction is provided by the difference between the two responses (lower panel). The largest difference is provided by neurons tuned to the upward motion. Adapted from Hol & Treue, (2001)

These findings were used in a recent study (Jazayeri & Movshon, 2007a) to show that biases in estimates of motion direction are not due to a limitation in the sensory apparatus during the encoding stage (i.e., subjects can accurately encode the direction of a stimulus), but rather derive from strategies applied during a decoding stage, the way in which encoded information is manipulated to obtain a meaningful representation. In this context, the different principles described by Hol & Treue, (2001) represent different decoding strategies and Jazayeri & Movshon, (2007a) suggest that in discriminating two highly similar stimuli (one providing spatial orientation information, the other direction information), the most informative difference is provided by neurons tuned away from the stimulus direction. These, therefore, can be thought as assuming a higher weight (i.e., informative responses are weighted more than uninformative ones) and it

is because of this framework that estimates then result to be biased towards the direction weighted the most. This was demonstrated using the following paradigm.

A trial was divided in two stages (see Figure 1.17): first a discrimination task. Subjects were presented with a reference line whose angular position on each trial was randomly selected from a range spanning 360 degrees. Importantly, this cue was displayed for the whole length of a trial. Afterwards, a field of dots was displayed. The proportion of dots coherently moving in one direction was varied, allowing the difficulty of the task to be manipulated, and the axis of motion was randomly selected to be within 22 degrees to the left or right of the reference line. Subjects had to judge whether the dots moved clockwise or anti-clockwise with regard to the reference by pressing the corresponding key and they received feedback regarding the correctness of the response immediately following each trial. However, on a subset of trials, feedback was withheld and subjects were instead cued to perform the second part of the experiment: an estimation task. Here they were asked to estimate the direction of motion of the coherent dots by extending and aligning a bar from the fixation point at the centre of the screen using a mouse.

A variation of the discrimination task consisted in changing the angular difference between the reference line and the direction of motion of the coherent dots, resulting in a *coarse* discrimination task. Here subjects had to judge whether the signal dots moved *towards* or *away* from the reference line.

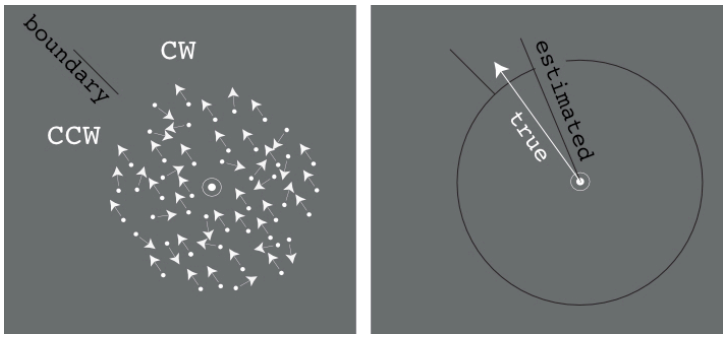


Figure 1.17 Task sequence: During the discrimination phase a field of moving dots is presented together with an oriented boundary. During the estimation phase, subjects have to report the estimate of the perceived direction of motion in the presence of the oriented boundary. Adapted from Jazayeri & Movshon, 2007a

The structure of this study is particularly clever in that it allows, on one hand, to derive the internal representation of the sensory stimulus from the behavioural performance during the discrimination task, and on the other hand to obtain a measurement of that representation.

Results of the discrimination task show a relationship between the strength of the motion signal and discrimination performance as measured via a psychometric function (Jazayeri & Movshon, 2007a). This approach allows the measurement of *sensitivity*: the minimum amount of signal change in the stimulus for the difference to be perceived. The smaller this difference, the higher the sensitivity.

Jazayeri & Movshon, (2007a) showed that with increasing coherence in the motion stimulus, performance in the discrimination task also increases. Data

from the *estimation* task showed an interesting pattern: subjects were *biased away* from the reference line. By plotting the angular difference between the reference line and the estimated direction of motion as a function of the angular difference between reference line and stimulus direction of motion, it was possible to see that the distribution of estimates deviated from the stimulus directions the closer to the orientation of the reference it got (see Figure 1.17). Moreover, this bias in estimates of perceived direction of motion changed with signal strength: the higher the coherence of the moving stimulus, the less evident the bias.

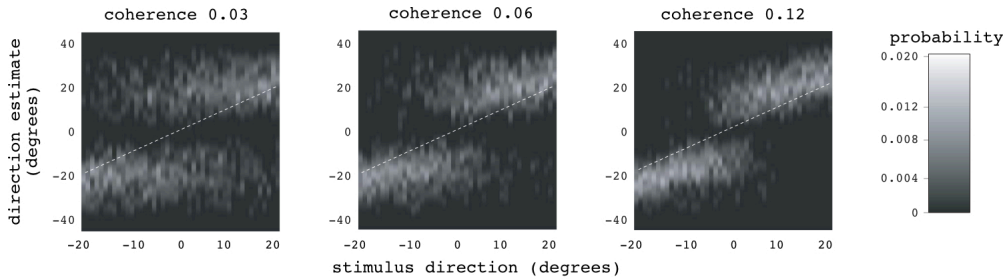


Figure 1.18 Distribution of estimation responses: estimates of motion direction are plotted as a function of stimulus direction for three coherence levels (individual panels). The white dashed line is the identity line (veridical responses). A repulsive bias is indicated by the gap in the distribution of estimates around 0 (representing trials in which boundary and motion had the same orientation / signal motion direction). Adapted from Jazayeri & Movshon, (2007a)

While this was true for estimates of motion direction following a *fine* discrimination task, the authors reported that estimates following a *coarse* discrimination task were not biased: the angular difference between estimated direction of motion and orientation of reference line closely matched that between stimulus direction of motion and reference.

The authors constructed a *decoding model* to explain the observed biases in estimates of motion direction. The model (see Figure 1.19) was composed of two stages: data from the discrimination task were used to derive the sensory representation of the moving stimulus. This sensory

information was modelled as a Gaussian probability density function, centred at the stimulus direction of motion (for any given trial) and whose variance was derived from the psychometric curve obtained from the behavioural performance in the discrimination. In particular, the steeper the psychometric curve, the narrower the internal representation, which means that the encoded information is more reliable (Gescheider, 1997; Jazayeri & Movshon, 2007a) (more on this in Chapter 4, Section 4.1.1).

The second stage of the model consisted of a *weighting operation*, whose aim was to boost the responses to the moving stimulus in order to optimise the discrimination performance. As described above, the foundation of this process lies in the work by Hol & Treue, (2001) on the *reference repulsion effect*: in a *fine* discrimination task, the most informative neurons are those whose tuning curves peak slightly away from the discrimination boundary. Thus, the brain needs to apply a specific strategy, here represented by a *weight*, in order to efficiently pool the responses to the stimulus. In particular, the *weighting profile* is recovered using the data from the study, so that the maximum of the so-called *sensory likelihood*, i.e. the product between the sensory representation and the weight itself, best predicts the estimates given during the second stage of the trial.

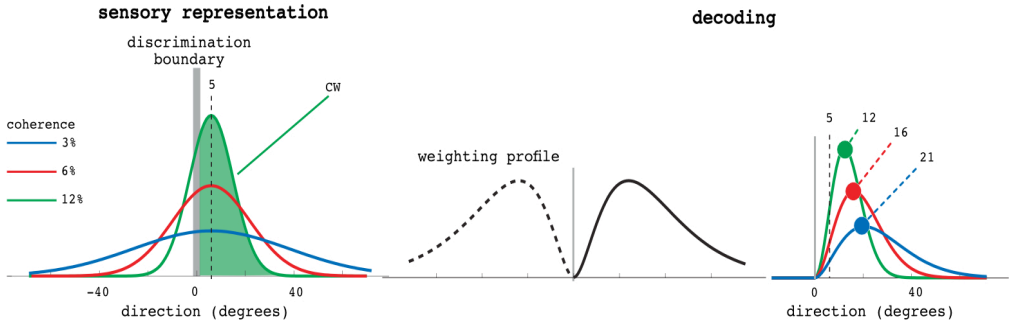


Figure 1.19 Schematic representation of the decoding model. Sensory representation encoding a stimulus moving 5° CW from the boundary for three different coherence levels. A weighting profile is applied to each of the sensory representations resulting in the shifted distributions on the right. The peak of each distribution represents the estimated direction of motion for the 5° CW stimulus. Adapted from Jazayeri & Movshon, 2007a

While this model was able to account for the effect observed in the estimates following a *fine* discrimination task, and theoretically explained why such bias is not present in estimates following a *coarse* discrimination, alternative explanations have been proposed by other researchers (e.g., Stocker & Simoncelli, 2008).

By means of a *conditional Bayesian observer model* Stocker and colleague provided an account for the *repulsion* effect observed in the estimates of sensory information. Briefly, *Bayesian theory* provides a way to combine *prior knowledge* of the world (or specific feature of interest) with *sensory data* observed in time (Stocker & Simoncelli, 2008). In general, *Bayesian Theory*, or *Bayesian Decision Theory*, is a very powerful tool: it enables prior knowledge to be combined with new sensory observations in an optimal way (from a computational point of view) (Vincent, 2015). The main assumptions upon which the theory is developed are that an observer possesses a mental representation of the way in which certain sensory events arise, that the observer uses sensory evidence to infer information regarding a certain state of the world, and that the observer possesses prior beliefs of different states of the world (Kersten, 2000).

Thus, the aim is to use the observed sensory evidence to infer states of the world. This problem translates into computing a posterior probability

(Figure 1.20A), and is also referred to as the *inverse problem* (Tarantola, 2006):

$$P(world|data)$$

The strength of Bayesian Decision Theory is to provide a solution to the *inverse problem*, consisting in multiplying the probability of an observer acquiring sensory data given a certain state of the world (*likelihood*, Figure 1.20B) by his/ her belief regarding that specific state of the world (*prior*):

$$P(world|data) \propto P(data|world) \times P(world)$$

To put it in yet different terms, appropriately weighting sensory information by prior knowledge and beliefs about a current context, environment, or state of interest.

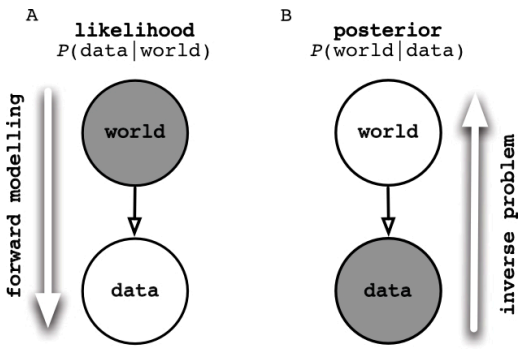


Figure 1.20 Bayesian inference representation. A) The forward model approach states that it is possible to compute the likelihood of observing certain data given the current state of the world. B) A posterior distribution consists in inferring the state of the world provided we have observed certain sensory data. This is known as the inverse problem. Grey circles in A and B show known quantities, compared to white circles indicating quantities to be deduced / inferred. Adapted from Vincent, 2015

According to the Bayesian theory, an *ideal observer* would therefore have to evaluate a piece of evidence under several possible hypotheses (priors, each representing a possible state of the world). An *ideal observer* is a *device* able to perform a certain task in an optimal way given a certain set of rules and starting information (Geisler, 2003). An optimal outcome of the decision process is then obtained by appropriately combining each hypothesis

weighted by its probability of being correct. However, several studies (Acerbi, Vijayakumar, & Wolpert, 2014; Beck, et al., 2012) have reported that this is not the case: subjects are often suboptimal in their decisions. Stocker & Simoncelli, (2008) explain this phenomenon in terms of *self-consistency*: an observer, instead of evaluating the evidence under each possible hypothesis, does so using only one prior to which he / she commits. All alternatives are then discarded and the final outcome is *biased* or suboptimal (Figure 1.21). This specific type of framework, considering priors as representing plausible hypotheses on the properties of sensory evidence during a perceptual decision task and reaching a response by combining multiple weighted alternatives, is called *Bayesian model averaging* or *optimal model evaluation* (Draper, 1995; Hoeting, et al., 1999).

Stocker and Simoncelli, however, propose a *Bayesian conditional model*: an observer evaluates the evidence under a chosen alternative, re-orienting any subsequent behaviour to that specific, chosen framework, resulting in a suboptimal behaviour.

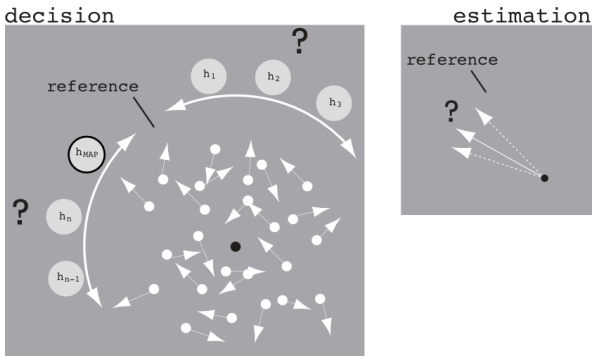


Figure 1.21 Discrimination and estimation tasks according to the Bayesian conditional model. A series of hypotheses regarding the direction of motion of the stimulus are generated and evaluated during the decision stage. The most probable hypothesis is indicated by h_{MAP} and selecting this hypothesis affects the subsequent estimation stage because the reported estimate will be based on the selected hypothesis alone rather than a combination of all possible hypotheses. Adapted from Stocker & Simoncelli, 2008

In order to explain, among others, the effect observed by Jazayeri & Movshon, (2007a), where estimates of motion direction are biased away from a reference line when following a *fine* discrimination task, Stocker & Simoncelli, (2008) use a *self-consistent conditional Bayesian model*. With no free parameters, the components of the models are: (i) a prior concerning possible alternative hypotheses regarding the state of the world, (ii) a prior on the state of the world under which we want to evaluate the observations, (iii) a likelihood function (i.e., a representation of the sensory signal corrupted by both internal and external noise).

Thus, using the experimental setup provided by Jazayeri & Movshon, (2007a), an observer is presented with a reference line with a certain orientation, followed by a field of moving dots, a proportion of which are moving coherently in one direction. The observer perceives the coherent motion directed towards the right of the reference line. In this very case, there are two possible priors about the state of the world: either the signal dots moved clockwise of the line, or anti-clockwise but noise interfered with perception and made it *look* clockwise. The strength with which one believes in each of the two hypotheses is given by the *posterior*: the probability of the signal moving in one direction given the belief in a certain hypothesis. The optimal outcome is given by choosing the hypothesis that maximises the posterior: by making a judgement in the discrimination phase, the observer decides for one of the possible hypotheses, discarding the alternatives that would otherwise act against the consistency principle (i.e., the tendency of acting in line with one's thought / decision). Estimates

of direction of motion are then biased due to the strong effect of the chosen hypothesis over the representation of the sensory signal.

This alternative model is also able to describe the effects observed in the study by Jazayeri & Movshon, (2007a), but lacks neural plausibility. While the *decoding model* proposed in the original Jazayeri & Movshon, (2007a) study gives an explanation of the biased estimation in terms of neural population mechanisms, the Bayesian model does not propose an implementation in terms of neural mechanisms involved in the process of self-consistency.

However, there are other factors that could potentially result in biased estimates of motion direction, and this thesis aims to investigate some of them.

In particular, the bias could have occurred later in the trial, rather than at the time of decoding (Rauber & Treue, 1999; Zamboni, et al., 2016). One specific aspect to consider that has a direct bearing on this is the presence of the oriented reference line. Because it was present throughout the whole trial in the Jazayeri & Movhson experiment, it is difficult to disentangle between an effect due to the decoding of the internal representation of the sensory signal (that would be consistent with the findings reported so far) and one due to the presence of the reference itself. The physical display of the line is, indeed, irrelevant to performing the estimation itself: one has to report the direction in which the coherent motion was perceived and should therefore discard any other element that might cause interference in performing the task.

To address this issue, I performed a set of experiments in which the properties of the reference line are systematically manipulated during the course of the trials. In the following chapters of this thesis I will describe three behavioural experiments and one measuring eye movements.

Chapter 2 gives a description of the behavioural psychophysical techniques, eye tracking methods, and imaging technique (fMRI) employed in conducting the experiments reported in this thesis, together with the approach used in analysing the data for each specific technique applied.

Chapter 3 investigates the point in time at which biases in estimates of motion direction arise. First, I will present data replicating the findings reported by Jazayeri & Movshon, (2007a), confirming the repulsion effect. Second, in order to understand whether such bias arises relatively early or late during information processing, a first manipulation is applied to the oriented reference line: following a *fine* discrimination task, subjects had to estimate the perceived direction of motion in the *absence* of the reference (i.e., we removed the oriented line during the estimation task). Following these results, I will present the findings from a second manipulation: subjects performed the estimation task after passively viewing the motion signal (i.e., no overt *fine* discrimination task). Moreover, the orientation of the reference line was systematically changed at the onset of the estimation task from the orientation it had at the beginning of the trial (i.e., when it was presented together with the moving stimulus). This allowed us to further address the timeline of the repulsion effect and to also see whether

it could be manipulated by overtly changing an element on the screen that should be irrelevant in performing the actual task.

In Chapter 4 I will present a computational model aimed at explaining the findings described in Chapter 3. Because we observed that manipulating information regarding the oriented reference line late in the task sequence affects the estimates of direction of motion given by the observers, we argue that the biases arise at a relatively late stage of information processing, rather than at either sensory encoding (Eagleman, 2001) or decoding (Jazayeri & Movshon, 2007a). The model proposed here supports our findings and takes into account the major role played by the reference and the information it provides in the measured behaviour.

Chapter 5 investigates the role of response modality in visual motion direction estimates. In particular, here I used eye tracking techniques to measure estimates of motion direction. Results obtained using this alternative method are compared to those presented in Chapter 3 and alternative experiments are presented for future research.

Chapter 6 presents data obtained using fMRI while subjects performed perceptual decision tasks in the scanner. Participants had to perform a *fine* or *coarse* discrimination task in a block design while being in the scanner. Here, a univariate analysis approach is used in order to assess the state of the data acquired in the scanner as well as the adequacy of both stimuli and tasks employed.

Chapter 7 presents the results from a separate study using fMRI to derive population receptive field maps of areas of the visual cortex. This part

is an introduction to a different way of looking at the data: a model describing the physical stimuli encoded during the experimental session is used to derive properties of regions of the brain. Following this, a multivariate analysis approach to the data obtained as described in Chapter 6 is presented. In addition to performing a *fine* or *coarse* discrimination task, in a set of separate scans, subjects were engaged in an attentional task while wedges of moving dots were presented. This allowed a baseline level of activation to be obtained concerning the pattern of activity related to encoding each direction of motion. This information, together with the data obtained during the *fine* / *coarse* discrimination task was then analysed using a *forward encoding model* (Brouwer & Heeger, 2009, 2011, 2013; Kay, et al., 2008; Kok, et al., 2013). The weights computed via this model for each condition were then compared and used to support the view that when judging highly similar information, a different neural encoding strategy is used than when judging highly dissimilar ones.

In Chapter 8 a general discussion will put together the findings reported in this thesis, incorporated in an appropriate theoretical framework and present possible future directions for this research in addressing questions that are still unanswered.

CHAPTER 2

2 *General Methods*

In this chapter I will describe in general the methods used in the experiments reported in this thesis. In particular, I will focus on:

- *psychophysical* methods used to perform the three main behavioural studies
- the *eye tracking* technique used to investigate the role of response modality in biases of direction of motion estimates
- *functional magnetic resonance imaging (fMRI)* methods, including general principles at the core of this imaging technique, details of how cortical surfaces were obtained from anatomical images as well as how retinotopic visual areas were identified. I will also briefly summarize analysis methods such as the *general linear model* and the principles of a *forward encoding model.*

2.1 Psychophysics

Psychophysics is an experimental method investigating the quantitative relationship between a physical property of the world and its psychological sensation. The approach has a long and distinguished history (Fechner, Adler, Howes, & Boring, 1966). Already the early psychophysicists locked onto the idea of a *threshold* (or *sensory threshold*). This is the minimum amount of sensory information a physical object has to provide for it to be

consciously perceived / detected. Thresholds can be of an *absolute* (detection) or *difference* (discrimination) type, depending on whether the amount of signal necessary for the object to be perceived is measured relative to the environment or as a difference of signal strength between multiple targets, respectively.

Detection tasks generally require subjects to give a 'Yes-No' response while presented with a signal whose strength is systematically manipulated. For example, Hecht, Shlaer, & Pirenne, (1942) measured the minimum number of photons that could be detected by the human eye. Participants were first adapted to darkness to enhance sensitivity to the weak flash of light that was then delivered for a brief amount of time (1msec) aligned to a region of the retina with high concentration of rods photoreceptors (see 1.2.1 section for details on photoreceptors). By varying the energy of the flash of light and recording the frequency with which participant reported perceiving the light (*detection*), Hecht and colleagues computed a detection *threshold*: the minimum amount of photons the human eye can detect.

Discrimination tasks, on the other hand, require subjects to compare an attribute of interest among multiple options and arrive at a final judgment. As an example, imagine having to decide whether weight *A* is heavier than weight *B* (Weber, 1834). By changing the intensity of the stimulus (i.e., the difference in weight between *A* and *B*) along a predefined range of possible values, and simultaneously recording the amount of times you report *A* as heavier than *B*, it is possible to compute the proportion of 'heavier' responses as a function of stimulus intensity (Gescheider, 1997). This

psychometric function often takes a characteristic 'S' shape and is well-described by the cumulative *normal distribution function*. The relationship between these two is shown in Figure 2.1: a y-value on the psychometric curve represents the cumulative area under a normal distribution curve (Gescheider, 1997). A discrimination task, therefore, informs upon the *just noticeable difference*: the minimum difference between two stimuli resulting in a 'yes' (here 'heavier') response 50% of the times.

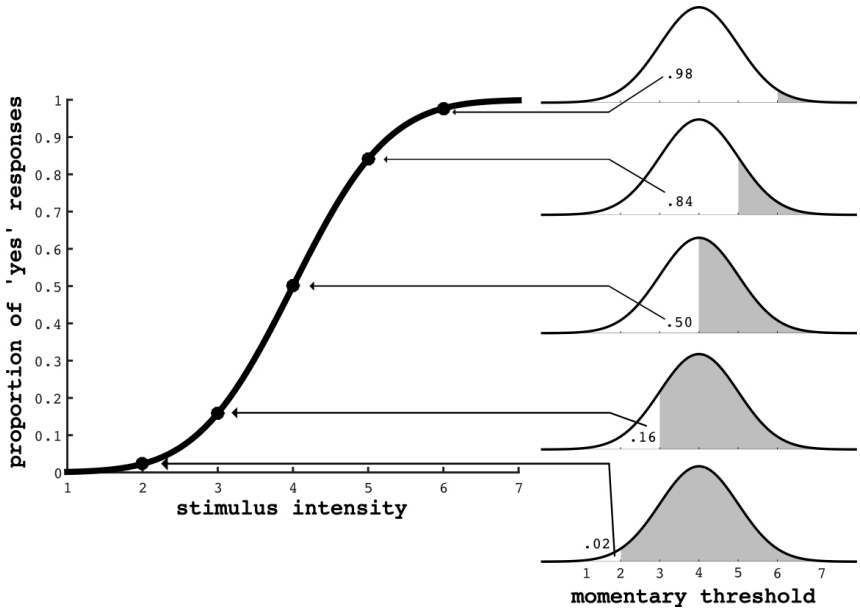


Figure 2.1 Psychometric function and signal detection: relationship between behavioural performance expressed by the psychometric curve and the probability of detecting a signal shown by the probability distributions on the right for each data point of the curve. Adapted from Gescheider, 1997

The experiments reported in this thesis will employ discrimination tasks that require participants to compare one signal against another.

2.1.1 Visual Direction of Motion Discrimination Task and Psychometric Functions

The task used in the behavioural studies described in this thesis is a visual discrimination task on direction of motion. Note that this type of discrimination task can be further divided into sub-types: a *fine* discrimination consists in comparing two stimuli that are highly similar along the feature of interest. A *coarse* discrimination, on the other hand, requires comparing stimuli whose physical property is highly dissimilar, e.g., upward vs. downward motion. For the experiments described in Chapters 3 to 5, a fine discrimination task is adopted. In the study described in Chapter 6, subjects perform both fine and coarse discrimination tasks.

Participants were presented with a stimulus (oriented reference line, also called the *reference stimulus*) and they had to judge whether the target (field of moving dots, also called the *comparison stimulus*) was moving in a direction that was either *clockwise* (CW) or *anti-clockwise* (CCW) with respect to the reference. The *method of constant stimuli* was used to manipulate the feature of interest, here the absolute angular distance between the two visual inputs. Furthermore, the strength of the information provided by the comparison stimulus varied on a trial by trial basis: the proportion of dots moving in the same direction, *coherence*, could assume one of four possible values, affecting the difficulty of the task (from easy: 25% of dots constituting the comparison stimulus moved coherently, to difficult: 4% of coherent signal).

From the judgments obtained by each participant, we computed the proportion of CW responses and plotted these as a function of the angular difference between the reference line and the moving stimulus. The relationship between responses and relative direction of motion was defined via a *psychometric curve*: Gaussian cumulative distribution function directly related to the *signal + noise* distribution assumed by *Signal Detection Theory* (see Chapter 1, Section 1.1). Analysing the data in this way allowed us to quantify changes in performance as the difference between the stimuli changed, and investigate the effect of coherence level on performance.

2.1.1.1 Subjects

In the experiments reported in Chapters 3, 5, 6, and 7, data were collected from human subjects. All participants had normal or corrected to normal acuity and gave written informed consent prior to taking part in the study they were recruited for. The minimum number of participants for each experiment was four and their ages ranged from 21 to 30 years.

2.1.1.2 Apparatus

For experiments described in Chapter 3, stimuli were presented on a cathode-ray-tube (CRT) monitor (*Mitsubishi Diamond Plus 73*) at a resolution of 1024 x 768 pixels, a refresh rate of 85 Hz and at a viewing distance of 57 cm. The display was gamma-corrected, so that luminance was a linear function of the digital representation of the image. Stimuli (used in all experiments described in Chapter 3) were generated using the MGL toolbox

(available from http://gru.stanford.edu/mgl) in MATLAB (The MathWorks, Natick, MA) on Apple Mac OS X.

2.1.1.3 Stimuli

Random-dot kinematograms (RDKs) were used in all experiments outlined in this thesis. This stimulus is widely used in studying processes of motion perception and is highly versatile in that it allows the experimenter to directly manipulate key features such as: strength of motion signal (*coherence*), direction of motion, speed of motion, etc. Here, each RDK was composed of a field of moving 'white' (147 cd m^{-2}) dots on a 'grey' (38 cd m^{-2}) background. Each dot was 0.12° in diameter, moving within a circular aperture 5° in diameter. Different coherence levels were used (4%, 7%, 13%, and 25%): *signal* dots moved at a speed of 4°s^{-1} in a predefined direction. The remaining *noise* dots were randomly re-plotted within the circular aperture. Patterns were updated on every alternate frame (42.5 Hz) to closely match presentation parameters used in previous studies (Jazayeri & Movshon, 2007a). The average dot density was 40 dots $deg^{-2}s^{-1}$. A 'black' (0.5 cd m^{-2}) reference line (length 0.5°, width 0.15°), starting 3.5° from the fixation point (0.5° diameter at the centre of the screen), was used as discrimination boundary.

2.1.1.4 Procedure

Participants were required to perform two tasks: a fine motion discrimination judgement and an estimation task. Observers had to judge

whether the RDK moved more CW or CCW with regard to the orientation of the reference line. On a subset of trials, participants were prompted to perform the estimation task: adjust a line extending from the fixation spot at the centre of the screen in the direction they perceived the coherent motion.

In general, each trial started with the presentation of the fixation point at the centre of the screen for 1 s, followed by the onset of the reference line oriented at a random angle. 0.5 s later, an RDK with a particular coherence level was displayed for 1 s. The direction of motion of the *signal* dots was randomly selected on each trial from a uniform distribution ranging between ±24.5° with respect to the orientation of the reference line. Observers were then asked to judge the direction of motion by pressing one of two keys on a standard keyboard, indicating whether the dots moved CW or CCW with respect to the reference line (*fine discrimination task*). On a subset of trials, feedback on the discrimination performance was withheld and participants were prompted to perform the *estimation task*. Here, a black circular ring was displayed (6° diameter, centred at the fixation spot) and observers reproduced the direction in which they perceived the coherent motion by manually adjusting a line that extended from the fixation point using a mouse. A trial ended when the line hit the black ring, or after 10s from the presentation of the aperture (see Figure 3.1, 3.5, 3.10, Chapter 3, Sections 3.1.1.3, 3.2.1.3, 3.3.1.3 for graphical description).

The three psychophysical experiments described in Chapter 3 present some apparently subtle but important differences in their procedures that will be addressed in the relevant sections.

2.1.1.5 Data analysis

Note that a description of the data analysis and implementation of the model is publicly available at https://github.com/schluppeck/zamboni-2016.

Psychometric curves were obtained from the results of the visual motion direction discrimination task. The proportion of times an observer reported a CW response was plotted as a function of the angular difference (in degrees) between the reference line and the coherent motion. We used non-linear least-squares (as implemented via *fminsearch* in Matlab; Nelder-Mead algorithm, Nelder & Mead, 1964) to estimate the two parameters (μ, σ) of the best-fitting Gaussian cumulative distribution function. In particular, the *sigma* parameter gives an estimate of the discrimination *threshold* for a specific participant and condition (i.e., coherence level). The higher the value of this parameter, the lower the discrimination performance (i.e., the moving signal has to be stronger and further apart from the reference line for it to be correctly discriminated).

On the other hand, data from the *estimation* task were used to compare the presented (stimulus) direction of motion with the perceived (estimated) one. To do so, we looked at the relative error by considering the angular difference between the oriented reference line and the stimulus direction of motion. This allowed to combine data across trials. Because for each

stimulus direction of motion (ranging from -24.5° to +24.5°) participants gave many estimates (depending on the overall number of trials, ~5000 on average across multiple experimental sessions), for each physical direction, an empirical distribution (histogram) of the corresponding estimated directions was constructed. To further aid visualization of all the data collected, for each participant, a *joint histogram* was computed. In these diagrams, each column consists in the distribution of estimated directions of motion for a particular, stimulus direction. Different columns correspond to physical, stimulus directions. The lightness scale, indicating the amount of estimates given for a particular stimulus direction of motion, was adjusted according to the maximum number of trials for each specific condition and subject.

To quantify deviations from a linear relationship between the two variables, stimulus and estimated directions of motion, the *dip*-statistic (Hartigan & Hartigan, 1985; Hartigan & Hartigan, 1985) was used to analyse the patterns of responses. This statistical technique allows one to look for deviations from *unimodality*: if we assume that observers are *veridical* in their estimates, then we would expect responses to be normally distributed around the *physical* direction of motion. On the other hand, hypothesising a biased observer, we would expect the distribution of estimates to *deviate* from the *stimulus* direction of motion, resulting in a *bimodal* distribution of responses (i.e., the repulsion will result half the times in a CW-direction, the other half in a CCW-direction, thus the bimodality).

We also performed a control analysis, driven mainly by the results obtained in Experiment 2 described in Chapter 3, Section 3.2.2. Here, we checked for a linear relationship between the *physical* direction of motion of the stimulus and the *estimated* direction by computing a linear correlation. Indeed, a linear relationship between stimulus direction and estimated direction of motion would be expected in the case of an unbiased observer: the axis of motion manually reproduced on the screen closely matches (allowing for a minimum of noise from action execution and recoding apparatus) the physical motion direction of the stimulus. Given that, for each angular difference between reference line and coherent motion a set of estimates was provided, it was possible to obtain a distribution of such responses for each stimulus value tested. Following the hypothesis stated above, the shape of such distributions would be expected to approximate a normal probability density function, centred around the true stimulus direction (i.e., the peak would be expected around the mean of the distribution). The *peak* of each sets of estimates was obtained via the function *findpeaks* implemented in Matlab. This is a function that allows, given a vector of data, to identify *local* maxima within such distributions. Note that this differs from finding *absolute* maxima, that is the largest value assumed by a function across its domain. Local maxima (the *peaks* here) are mathematically identified as the largest value a distribution assumes within a predefined rage (Fermat's theorem – interior extremum theorem). The versatility of *findpeaks* consists in allowing to specify the number of local maxima to be detected, together with a threshold value (i.e., a minimum

difference to be exceeded between a data point and its neighbours). The algorithm, therefore, is particularly suitable for distributions characterised by more than one peak. This function then returns the value corresponding to the local maxima satisfying the constraints provided and an index identifying the position of the peak within the input vector of data. Thus, the linearity hypothesis was tested by correlating the local maxima with the physical stimulus directions: the higher the correlation value (in absolute terms), the stronger the linear relationship between stimulus and responses. This analysis was performed on data from each subject separately, as well as grouped data (i.e., from all subjects for each condition separately).

2.2 *Eye tracking*

The role of the modality with which observers report their responses in *estimation* tasks is a point of particular interest in this thesis. In Chapter 1, I described several studies measuring single cell activity, in which macaque monkeys responded via an *eye movement* (Britten et al., 1992; Newsome, 1989; Salzman et al., 1990). These *saccades* – synchronous, ballistic movement of both eyes in between successive fixation periods can be measured with *eye tracking* (Braun & Gegenfurtner, 2016). This method allows to determine the coordinates of movement direction of the eyes to be determined and therefore a spatially and temporally accurate description of gaze position over time can be constructed.

In what follows, I will first describe the mechanics underlying the measurements of eye movements, and briefly discuss the categorisation of the different types of movements that can be identified by considering variables such as movement *latency*, *velocity*, and *acceleration*.

2.2.1 Eye tracking technique

Several different techniques for measuring eye movements are available and the most commonly used modern techniques are non-invasive. In most such *video-based eye-tracking systems*, a camera records the reflection of a source of light that hits the eye. The particular technique used in the experiments here is based on reflective properties specific to the *pupil* and the *cornea* of the eye. The eye is illuminated with an infra-red light source and the *corneal reflection* resulting from this light source is then measured (Holmqvist et al., 2011). Infra-red light is used because it allows for higher degree of control of the reflected light and because there is less interference from stimuli and room illumination in the visible part of the spectrum. Using information from both the *pupil* and the *cornea* makes the measurement and tracking of the movements more stable and reliable.

The process leading to the measured movements of the eyes can be divided into three sub-processes: (1) acquisition of the state of the eye to be tracked, (2) analysis of the acquired image to obtain information regarding pupil and corneal reflection used to effectively track the eye, and (3) estimation of the actual trajectory and gaze paths.

During image acquisition (1), the camera located in front of the observer obtains an image of the eye and the information is sent over via a low-latency link to a further analysis stage on a computer. In some setups, extracting information regarding the relative location of the eye with regard to the rest of the face might be required, especially when observers are allowed to move their heads during recording. Stage two, image analysis, can be divided into two main steps: first, the pupil and corneal reflection regions of the eye are detected and segmented from the image acquired in (1), second the centres of these two regions are computed. Pupil and corneal reflection segmentations are generally achieved through thresholding and detecting regions in the image with similar intensity values. The centres of these regions are then estimated and used to compute gaze position (3). The *gaze estimation* is achieved by tracking the distance between the computed centres of the pupil and corneal reflection, respectively, and by taking into account the fact that these two points move at different speeds when the eye moves (the pupil moves faster than the corneal reflection).

Calibrating the video-based measurements to obtain gaze positions in the stimulus space is an important step in this process. The trigonometry governing the pattern of corneal reflections with the movement of the eyes is beyond the scope of this thesis (Holmqvist et al., 2011). However, a brief description of the *calibration* phase is relevant. Here, a predefined set of points (usually between 3 and 9, spacing through the area where the stimuli will be presented) are displayed on the screen one at a time and the

observer has to foveate to each point in sequence, fixating for a brief period of time each landmark. During this short interval, data (in terms of pupil and corneal reflection centres) are recorded. Given that the coordinates of the points on the screen are known by the algorithm, as are the dimensions of the screen and the viewing distance, a function is computed relating calibration points and recorded centres positions. This function is then used to map subsequent gaze positions onto the screen. Moreover, accuracy of the calibration procedure is assessed by computing the average distance between the landmark and the gaze position. This value is subsequently compared to a pre-set threshold: whenever the error is below the predetermined accuracy value, calibration is succeeded and the experiment can start. An accuracy of 1° (i.e., the estimated error between each target and corresponding gaze position) is generally considered a reasonable value, and this was used as threshold for the study presented in Chapter 5.

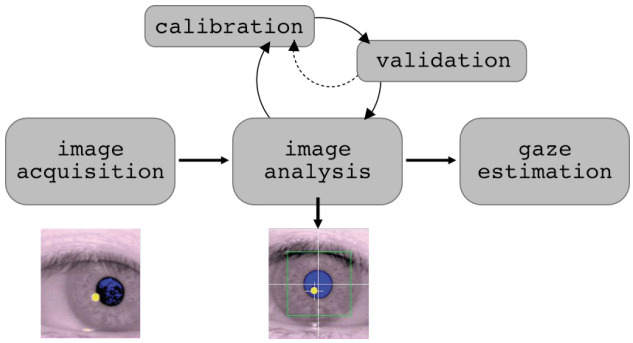


Figure 2.2 Overview of gaze estimation process. Adapted from Holmqvist et al., 2011

The rate at which the eye tracker takes images of the eyes, the *sampling frequency*, provides an important temporal limit on the technique. For the experiment described in Chapter 5, we used a 500Hz sampling frequency: 500 samples of eye position were obtained every second, leaving a 2ms

window between subsequent measurements. Such a small time interval between samples provides sufficient temporal resolution to resolve the movements we were interested in and also allows for a reliable measurement of starting and ending points of eye trajectories (Richard, Nystrom, & Kenneth, 2010). Since the eye-tracker system provides both raw and pre-processed data, it was possible to control the methodology for saccade detection and analyses of the measurements. This allows flexibility in manipulating the measurements obtained: more or less stringent constraints can be applied to the data for detecting specific events in the time series and researchers can inform customised data analysis algorithms with details of the specific experiment performed.

Two important measures that can be computed from the raw data are: instantaneous *velocity* and acceleration. The velocity is determined by the distance between two sampled data points and the time that has evolved between them (ds/dt). Acceleration is the second derivative (d2s/dt2) and can be obtained by differentiation from the *velocity* measures. These two measurements, together with parameters such as *duration* and *peak velocity*, can be used to detect saccade onsets from the raw gaze traces, as described in Section 2.2.3.3 below.

2.2.2 *Eye-movement patterns*

Eye movements, both *voluntary* and *involuntary*, are possible through an arrangement of muscle fibres, resulting in: vertical, horizontal, and torsional movement (Haslwanter, 1995). A detailed description of the oculomotor

plant and related physiology is beyond the scope of this thesis. However, it is instructive to think about the different states of the eye that are observable via eye-tracking. By combining the axes of motion and keeping track of both pupil and corneal reflection, a classification of different movement patterns emerges:

- *Fixations*: variable time interval during which the eyes are (relatively) still.

In fact, the eyes are never completely still and studies have shown that even during these "*fixation*" periods, three distinct types of eye movements can be identified:

1. *Nystagmus*: small movements of the eyes whose aim is still unclear (unspecified *duration*, less than 1' *amplitude*, and a *peak velocity* of 20'/s) (Riggs, et al., 1953)
2. *Microsaccades*: small, ballistic eye movements aimed at reposition the eyes back to a point of origin (*duration* of 10-30ms, *amplitude* of 10-40', and *velocity* between 15-50°/s) (Ditchburn, 1980)
3. *Drifts*: slow movements that make the eyes '*drift away*' from the fixation point of origin (*duration* of 200-1000ms, *amplitude* of 1-60', and *velocity* of 6-25'/s) (Martinez-Conde, Macknik, & Hubel, 2004)

When looking at more pronounced eye movements, on the other hand, these can also be categorised into different categories depending on the characteristics and purpose they serve:

- *Saccades*: very fast (ballistic) movements aimed at jumping from one 'fixation' to another (*duration* of 30-80ms, *amplitude* of 4-20°, and *velocity* of 30-500°/s). A particular characteristic of saccades is also the fact that their end point is usually uncertain (i.e., upon a decrease in velocity, there is a stage of small adjustments leading to the point at which the eyes hit the intended target). This type of movement is called *glissade* (*duration* of 10-30ms, *amplitude* of 0.5-2°, and *velocity* of 20-140°/s) (Holmqvist et al., 2011)
- *Smooth Pursuit*: slow movement of the eyes aimed at following a travelling object (*duration* is indefinite and depends on the moving object to be tracked, *amplitude* is also relative to the trajectory of the object, and *velocity* ranges between 10 and 30°/s) (Holmqvist et al., 2011)

Although many other patterns of eye movements have been reported in the literature, these are the most commonly studied.

2.2.3 Eye-tracker Setup / Apparatus

Eye movements were recorded using SR Research 1000EyeLink eye-tracker (SR Research Ltd, Ontario, Canada). This is a video-based eye-tracker (with tracking accuracy reported to be between 0.25° and 0.5°), using infrared light to acquire the position of the pupil and corneal reflection in the eyes (both eyes were traced in the study described in Chapter 5, at a sampling rate of 500Hz, although data from the left eye were then used in the analysis). The camera was located on the desk in front of the observer

and below the display screen (Desktop Mount), at a distance of approximately 50 to 55 cm from the chinrest used to minimise head movements. A host PC was used to process and store the eye movement data while stimuli were presented on a Ilyama Vision Master Pro 514 CRT monitor (screen dimensions of 40x30 cm) at a resolution of 1280x1024 pixels, a refresh rate of 85Hz (Figure 2.3). Stimuli were generated using the MGL toolbox (available from http://gru.stanford.edu/mgl) in MATLAB (The MathWorks, Natick, MA) on Apple Mac OS X.

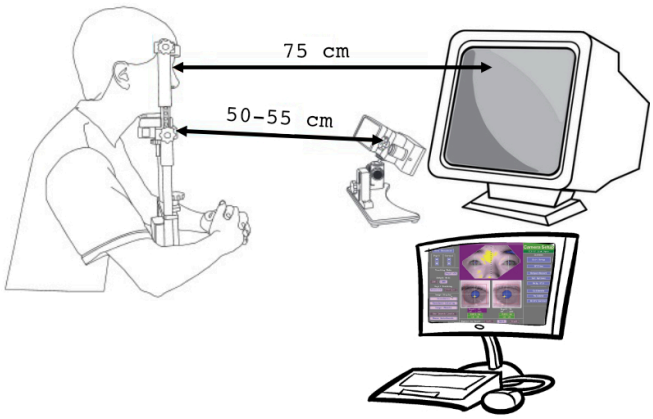


Figure 2.3 Schematic representation of eye-tracking setup in our laboratory. Adapted from Holmqvist et al., 2011

2.2.3.1 *Stimuli*

Stimuli closely resembled those described in the Psychophysics section (2.1.1.3), with the following adjustments:

- The average dot density composing the RDK stimulus was 100 dots $deg^2 s^{-1}$
- The circular aperture within which the RDKs were presented measured 7° in diameter
- The 'black' reference line started at 4° from the fixation point

- The orientations / angular directions of motion of the stimuli were restricted between 0° and 180° (upwards)
- The coherence level of the *signal* dots in the RDK was a *blocked* variable (the values used were: 7%, 15%, 25%, 50%, 80%)

While stimuli used in the behavioural studies (Chapter 3, Sections 3.1, 3.2, and 3.3) comprised of RDKs computed with coherent motion spanning 360 degrees, here the feature space was reduced in order to maximise the number of trials per condition. Moreover, it is known that tracking gaze position towards the bottom of a display (low elevation) can be problematic (Holmqvist et al., 2011): when moving the eyes downwards, eyelids and eyelashes tend to cover the pupil, causing artefacts in the traces or loss of data.

The highest coherence applied to the stimuli (80%) was used to train subjects with the task. Each participant performed an average of 900 trials with this supra-threshold stimulus, resulting in fixations during stimulus presentations that did not exceed the criterion of 1° from the central point, and confident saccadic eye movements during response (i.e., the movement was not affected by the presence of the reference on the screen).

2.2.3.2 *Procedure*

The task procedure was kept similar to the one used in the behavioural studies (see Section 2.1.1.4). Participants had their heads comfortably positioned on a chinrest and had to fixate the centre of the screen. The

experimenter, sat next to the participant, started the trial by first running the *calibration* routine (see Section 2.2.1 for more details) provided by the SR Research 1000EyeLink. This consists of presenting a sequence of 9 points, each dot sampling a specific location on the screen.

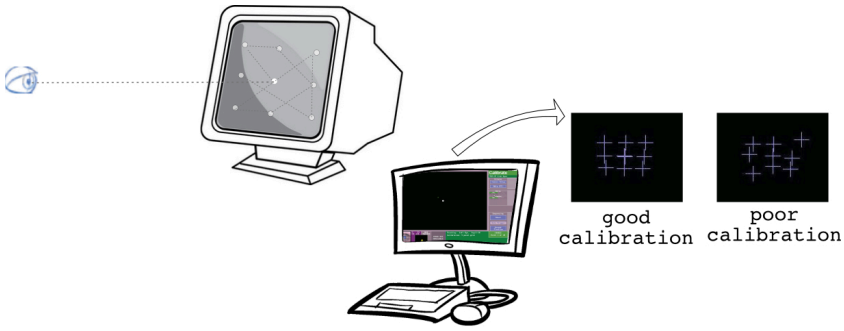


Figure 2.4 Calibration procedure: subjects are presented with one point at a time on the display screen and are asked to follow the dot. The Host PC displays the accuracy of the gaze and a summary display with corresponding average errors allows to either continue to the next stage (validation) or perform the calibration again.

Following the calibration stage, participants went through a similar routine, the *validation*. If the resulting accuracy value was within the criterion range (0.5 to 1°), the experiment started (see Figure 2.5), otherwise illumination and/or camera angle was adjusted and the calibration was repeated. Calibration and validation procedures were performed at the beginning of each set of trials.

During the actual experimental trials, subjects were asked to fixate at the centre of the screen, where a white spot was presented. After a short initial blank interval (250ms), a reference line with randomly selected orientation was displayed. 500ms later, a field of moving dots was presented: *signal* dots were moving coherently within a range of 20° clockwise or anti-clockwise of the reference line. After 2000ms, the fixation spot was removed from the screen and a black circular ring (7° diameter, centred at the fixation spot) was displayed. This was the cue for the

participant to make a saccadic eye movement in the direction they perceived the coherent motion. 750ms was allowed for the eye movement to occur, after this time elapsed a new trial started, and participants had to fixate back to the centre of the screen.

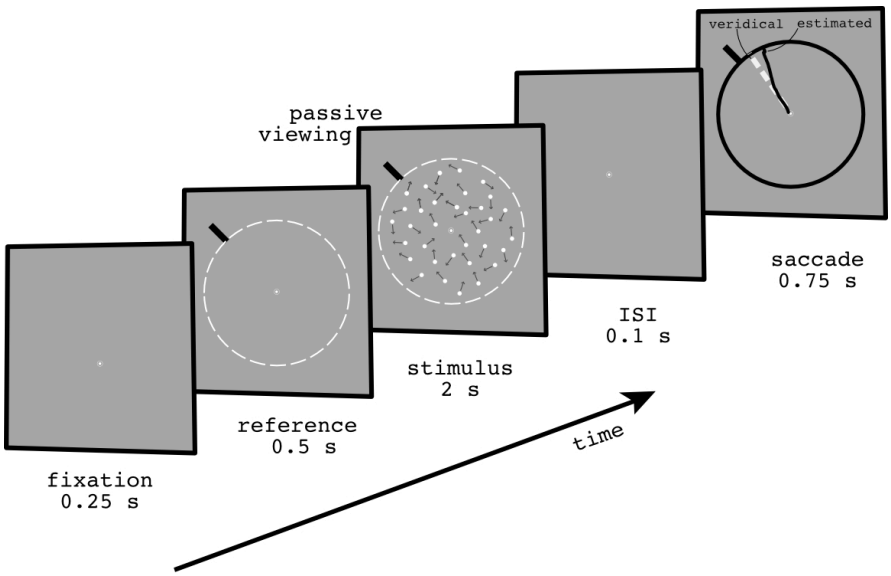


Figure 2.5 Trial sequence: after fixating the centre of the screen, subjects are presented with an oriented reference line followed by an RDK stimulus. Once the fixation point is extinguished and a circular aperture presented, subjects perform a saccadic eye movement estimating the perceived direction of motion of the RDK.

2.2.3.3 *Data Analysis*

Raw traces of eye movements, together with a record of the experimental parameters, were obtained. Gaze traces were converted from pixels to degrees of visual angle and pre-processed. First, by looking at the record of pupil size, blinks were removed from both horizontal and vertical coordinates of eye movements. Specifically, a window of 200ms around the detected blink was removed. Second, traces were mapped onto the segments corresponding to events during a trial (see Figure 2.5): (1) fixation, (2) oriented reference line presentation, (3) moving stimulus presentation,

(4) cue presentation (or *Inter-Stimulus Interval*, ISI), (5) saccadic eye movement response.

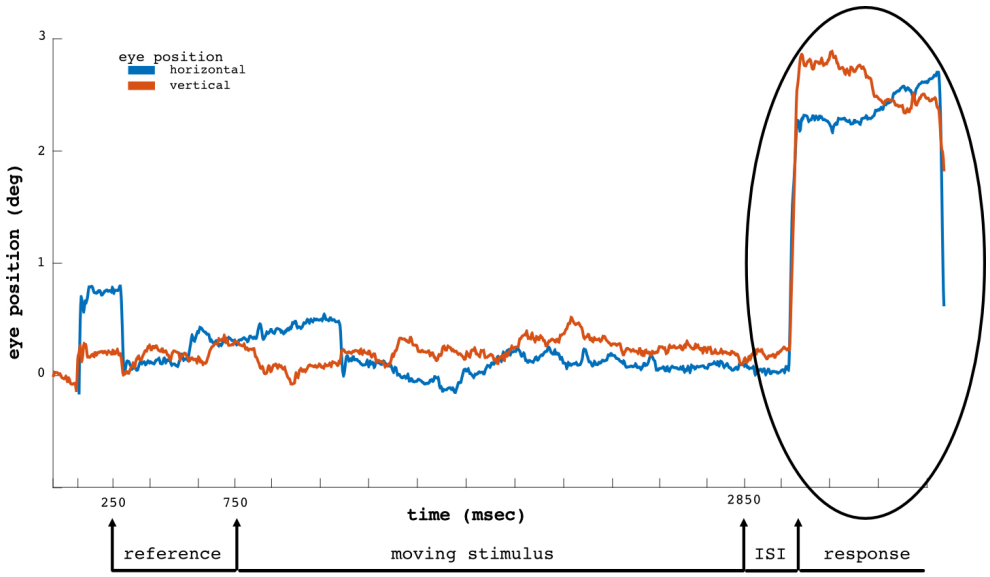


Figure 2.6 Trial sequence mapping onto eye movement trace. Horizontal and vertical gaze positions are plotted as a function of time and trial events are highlighted along the time axis. The saccadic response is identified by the ellipse surround the gaze traces.

The ultimate goal of the analysis was to compute the direction (angle) of the saccadic response for each trial, so that it was possible to directly compare the responses to the *physical* directions of motion of the presented stimuli. To do so, the following analysis pipeline was implemented (see Figure 2.7):

a) Eye movement traces (both horizontal and vertical coordinates) were smoothed using a 5-point method (Engbert & Mergenthaler, 2006) and *velocity* of eye movement was computed

b) Velocity *threshold*, obtained by computing median-based standard deviation for horizontal and vertical eye movement coordinates, was used to detect saccadic responses

c) A *saccade* was defined as that event lasting a minimum of 30ms (Holmqvist et al., 2011), during which the velocity computed in (a) exceeded the *threshold*. More specifically, the algorithm stepped through sample in sequence and set a flag on whenever the velocity traces exceeded the computed threshold and turned it off when it fell below threshold. Next, each identified segment was checked for a monotonically increasing set of measurements taking into account the time constraint previously set (i.e., *duration* of at least 30ms) (Behrens, Mackeben, & Schröder-Preikschat, 2010)

d) When these criteria were satisfied, the segment of the trace was labelled to contain a saccade and its *angular direction* was computed. This was done by transforming the Cartesian coordinates of the saccadic eye movement into polar coordinates

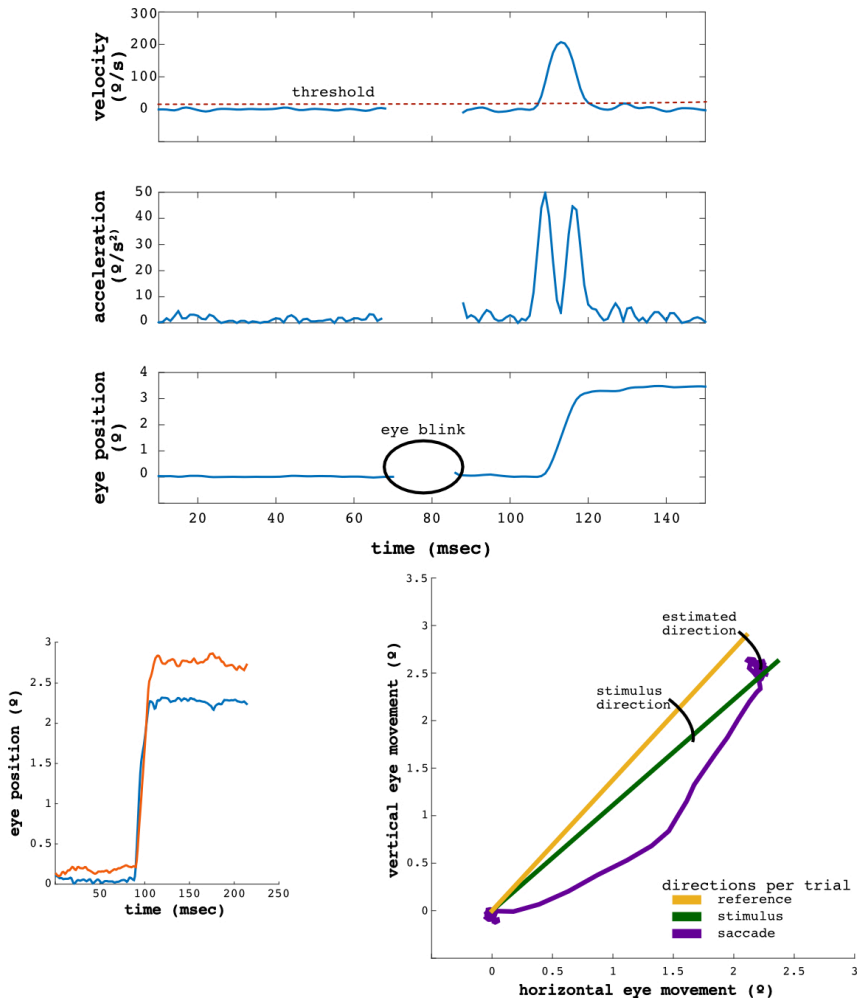


Figure 2.7 Schematic representation of the analysis algorithm. From the saccadic response highlighted in Figure 2.6, velocity and acceleration of the traces are computed, the gaze positions corresponding to the onset and endpoint of the saccade are determined and the resulting direction compared against that of the stimulus.

Angular directions could then be used in an analogous manner to that described in Section 2.1.1.5 to quantify and visualise the deviations of *perceived* (*estimated*) from the *physical* (*stimulus*) direction of motion in line with the overarching theme of the thesis. The focus of the analysis, therefore, was to look for significant differences between responses and physical characteristics of the stimulus. Mirroring the analyses in Section 2.1.1.5, I first calculated the distribution of trial-by-trial distances (in degrees) between responses and reference. These distributions were then combined

into a 2D format to aid visualisation of the overall responses. The peaks of each histogram were used to look for a linear relationship between *stimulus* and *estimated* directions of motion (this would be expected if subjects were unbiased). By presenting the data in the same format as for the experiments in Section 2.1.1, it is possible to look at the effect of response modality: the task sequence was kept the same (with few exceptions in terms of timings, see Section 2.2.3.2) and only the means by which subjects estimated the direction of perceived coherent motion varied.

In addition, a separate analysis of spatiotemporal variability of saccadic end points was performed (e.g., van Beers, 2007). Saccadic eye movements to the circular target ring were expected to show roughly the same amplitude regardless of stimulus direction. To check, for each stimulus direction, mean saccadic trajectory and saccadic endpoint variability were computed. Similarly, average amplitude, duration, and peak velocity were calculated and used to visualise them as a function of stimulus direction. This allowed me to look for distinct patterns in the data observed in previous studies – pattern especially for different saccade durations (e.g., Kubo et al., 1991; Oohira, Goto, & Ozawa, 1983): vertical saccadic eye movements are slower than movements in any other direction.

Analysis of endpoint variability required calculating the covariance between the coordinates of each response. Using bivariate mean and the covariance matrix, I then calculated the position and the axes of an ellipsoid that captures the distribution of saccadic endpoints. The directions of these axes correspond to the directions in which the distribution of points vary the

most and are given by the eigenvectors of the covariance matrix. The associated eigenvalues capture the variance of the data points in the direction defined by the associated eigenvector. Thus, it is possible to define a 95% confidence ellipse whose orientation is given by the angle between the largest eigenvector and the x-axis and the length of the major and minor axes is given by the square root of the product between the largest (smallest) eigenvalue and the critical chi-square value for a 95% confidence interval (Duda, Hart, & Stork, 2000; Hogg & Craig, 1995). The eigenvalues were also used to look for a dependency of variability as a function of coherence level in the stimulus.

In the next section, I will describe the methodological background for the final portion of my thesis work that involved taking a version of the behavioural experiment into an MRI scanner to obtain functional MRI responses in different brain areas while subjects were performing the tasks.

2.3 *Magnetic Resonance Imaging*

Magnetic Resonance Imaging (MRI) represents one of the most popular, non-invasive imaging techniques to study biological processes (Toga & Mazziotta, 2002). MR imaging evolved from *Nuclear Magnetic Resonance spectroscopy* (NMR, Bloch, 1946; Purcell, Torrey, & Pound, 1946) – some of the most important steps in this development are briefly outlined below. *Functional* MRI (fMRI), a specific application of MRI to measure physiological changes in the brain following localized neural activity, has

proven to be a powerful tool allowing to study the underlying neural basis of cognitive functions (of both high and low level). This section provides a description of the principles upon which MRI is based. Second, the basis of the most commonly measured signal in fMRI – the *Blood Oxygenation Level-Dependent or BOLD signal* – is explained. The third section is related to the study reported in Chapters 6 and 7, describing general analysis methods such as how *retinotopic maps* were constructed and *population receptive fields* (*pRF*) measured from activity recorded in visual areas. This section will also explain how analyses based on the *General Linear Model* (*GLM*) and a particular kind of *forward encoding modelling* were used.

2.3.1 From Nuclear Magnetic Resonance to Magnetic Resonance Imaging

A scanner is a piece of engineering consisting of an electromagnet whose *magnetic field* is typically in the range between 1.5 and 7 Tesla (although machines with lower, as well as higher, fields exist - Toga & Mazziotta, 2002). NMR (and MRI) relies on the impact that such magnetic fields have on *atomic nuclei* (hence the name *nuclear* magnetic resonance): by systematically modulating the behaviour of these nuclei through changes in the magnetic field it is possible to infer properties of the tissues they belong to. The most commonly used nucleus, due to its abundance in biological tissues, is *hydrogen* (^{1}H): an atom whose nucleus consists of a single proton and is therefore positively charged. Given its odd mass number, this atom possesses two fundamental properties directly related

to the imaging technique: protons are often conceptualized as *spinning* around their axis and possessing some *angular momentum* as well as a *magnetic moment* (a directional magnetic field, see Figure 2.8A). The *gyromagnetic ratio*, a constant specific for each atomic nucleus, is the factor relating the angular momentum to the magnetic moment.

When multiple nuclei spin within a confined region, a *net magnetization* (**M**) is generated as a 'sum' of single magnetic moments (Figure 2.8B).

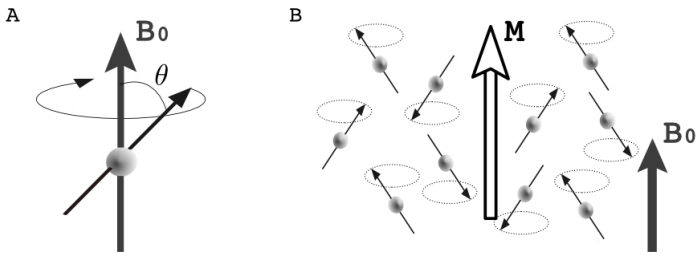


Figure 2.8 A) Schematic representation of a single spin precessing at an angle theta in a magnetic field B_0. B) Schematic representation of an ensemble of magnetic moments precesssing at equilibrium in a magnetic field B_0. Because more spins are in the low energy state than the high energy state, a resulting net magnetisation M is aligned to the magnetic field B_0. Adapted from Sanchez-Panchuelo, 2009

However, because the orientations of each magnetic moment will be randomly distributed, the resulting magnetization will be very small (Figure 2.9A). This is because forces in opposite directions will cancel each other out leaving few *net* signals to be added up. This means that **M** is normally too small to be measured reliably.

This scenario can be modified by adding an external, strong magnetic field ($\mathbf{B_0}$) to which the nuclei's magnetic moments will align (Figure 2.9B). In particular, quantum mechanics predicts that the nuclei will either align *parallel* or *anti-parallel* to the external field $\mathbf{B_0}$, that is nuclei can be found in

two possible *energy states*: a *low* (aligned) energy state and a *high* (unaligned) energy state (the *Zeeman effect*).

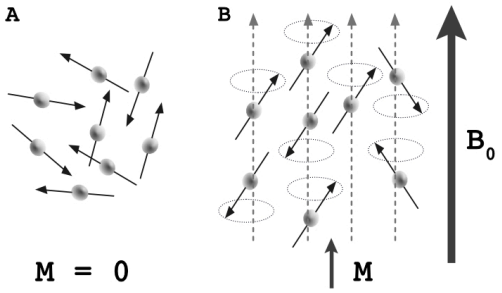


Figure 2.9 A) When no external magnetic field is applied, axes of the spins are randomly organised. B) Applying an external magnetic field causes the spins to take one of two possible states: they either align to the main field or they are antiparallel to it. Adapted from Sapountzis, 2009

The difference in energy states is the origin of the NMR signal, **M**, a small net magnetisation aligned with the external magnetic field. It is important to notice that, although aligned to the external magnetic field, the protons can still be thought of as precessing, i.e. the spinning nucleus is 'rotating' about the applied field axis at a frequency known as *Larmor frequency* (ω) that is dependent on the type of nucleus and the strength of $\mathbf{B_0}$ (Buxton, 2002; Jezzard & Clare, 2001).

The net magnetisation at the basis of the NMR signal can be viewed as a two-components vector: a *longitudinal* component (M_z parallel to the externally applied magnetic field), and a *transverse* component (M_{xy} perpendicular to $\mathbf{B_0}$). These axes define the NMR coordinate system.

2.3.2 The Radiofrequency Pulse

Once the strong, external magnetic field is applied, the net magnetisation of the system (in our case, brain tissue, although other organs can also be effectively imaged) is aligned along the longitudinal, z-

axis, but the signal is too weak to be measured ($\mathbf{B_0}$ >> **M** in terms of strength). As explained above, the nuclei are precessing about the field at the Larmor frequency, but with random *phases* – in the transverse plane, there is therefore no net-magnetisation. In this condition, the system is said to be at *equilibrium*.

To elicit and detect a signal, the system under investigation is brought out of equilibrium by the application of a *radiofrequency* (*RF*) *pulse*. By adding energy with an RF pulse, the state of the system is changed (Figure 2.10B): (1) the energy state of many spins is changed (more spins will move from a low to a high energy state), and (2) spins tend to align and the spread of phases (in the transverse plane) is reduced: they are being *focused.*

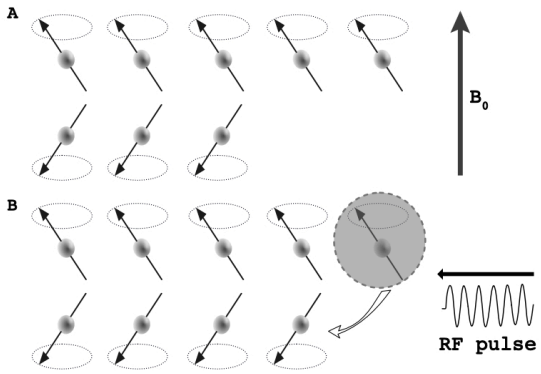


Figure 2.10 A) When an external magnetic field is applied, a greater number of spins will align to it (i.e., be in the low energy state). B) When a radiofrequency pulse is delivered at the right frequency, some spins will absorb energy and move from the low to the high energy state (spin highlighted by the grey circle). Adapted from Huettel et al., 2009

In order for this perturbation to work, the frequency of the pulse has to match that of the spins – the Larmor frequency. The RF pulse – a time-varying electromagnetic field which is often referred to as $\mathbf{B_1}$ – is a vector aligned perpendicular to $\mathbf{B_0}$ (but changing in that plane according to the Larmor frequency.) Its effect on the net magnetisation **M** is to 'tip' it over.

In this process the longitudinal magnetisation decreases, and transverse magnetisation aligned along the x-y plane increases (Figure 2.11). The amplitude and duration of the RF pulse determine the so-called *flip angle*: the amount of rotation away from vertical experienced by the net magnetisation **M**. For example, a 90° pulse brings **M** from vertical (aligned with B_0) into the transverse plane (orthogonal to B_0).

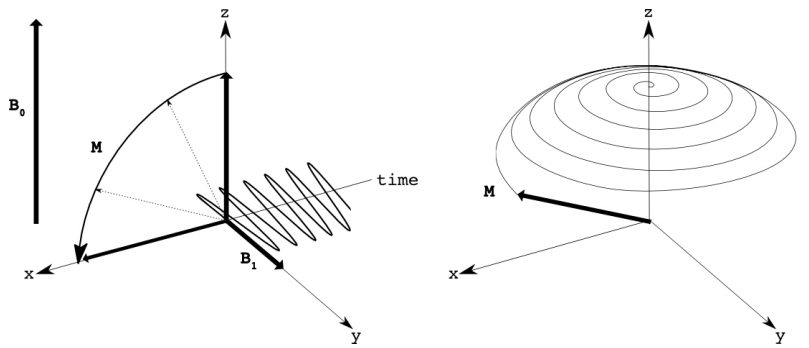


Figure 2.11 B_1 represents the applied RF pulse and is perpendicular to B_0. This causes the net magnetisation M to tip over into the horizontal plane. On the right a representation of the spiral motion followed by the net magnetisation after the RF pulse is applied. Adapted from Sapountzis, 2009

After $\mathbf{B_1}$ has been applied and the new transverse, focused magnetisation established, the system will return to its state of equilibrium. There are two separate processes that govern this return to equilibrium:

(1) the transverse magnetisation will decrease due to de-phasing in the horizontal / transverse plane

(2) the longitudinal magnetisation will return to its original status, aligned to B_0. This process is called *relaxation*: **M** is reduced through *transverse relaxation* (Figure 2.12A) – the transverse component of the net magnetisation decreases usually on the order of milliseconds. At the same time, in a process called *longitudinal relaxation* (Figure 2.12B), the longitudinal component of the net magnetisation is restored – depending

on tissue type and B_0 field strength, this process usually takes on the order of seconds. It is during these processes that the signal is created and measured.

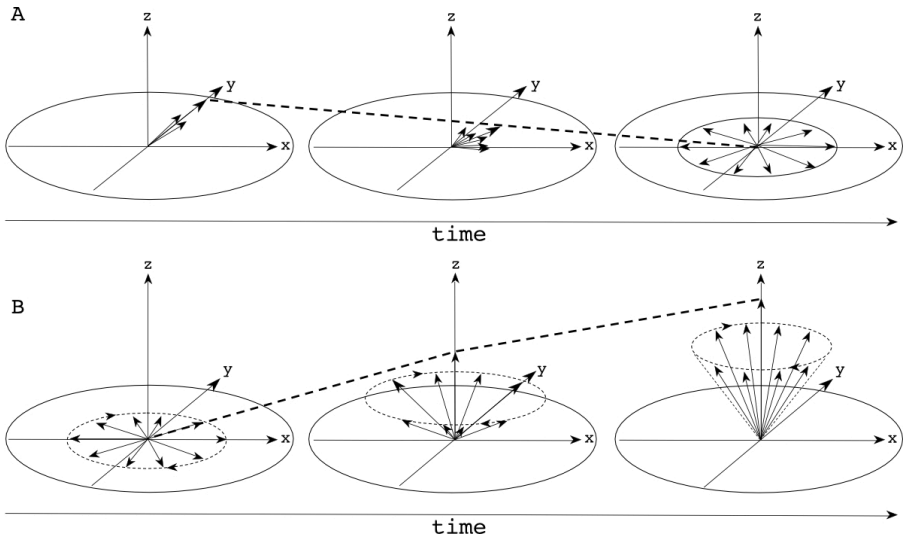


Figure 2.12 A) Schematic representation of T_2 decay. Once the net magnetisation is tipped over to the transverse plane, it starts decaying due to a lack in coherence among the spins. The dashed line across the three plots indicates the decay time. B) Schematic representation of T_1 recovery. After the RF pulse is delivered, spins start to release energy causing the net magnetisation to recover along the longitudinal axis. The dashed line indicates the time for T_1 recovery. Adapted from Huettel et al., 2009

Both longitudinal and transverse relaxations are events that evolve over time and are described by exponential processes with time constants T_1 and T_2, respectively. T_1 represents the time required by a specific tissue (or system) for its net magnetisation **M** to return to equilibrium. T_2, on the other hand, is the time constant describing the exponential decay incurred by the transverse component of the net magnetisation, a process also known as *Free Induction Decay* (FID). However, the decay of the transverse magnetisation is much quicker than described by the time constant T_2. This is due to inhomogeneities in $\mathbf{B_0}$ that affect the precession frequency, leading to a loss of coherence that adds to the loss due to spin-spin interactions. The time constant describing this event is therefore known as T_2* and presents

properties that are exploited by *functional Magnetic Resonance Imaging* (fMRI, see Section 2.3.4).

It is important to note here that these time constants vary depending on the type of tissue being studied: this specificity allows to obtain contrasts between different tissues in the acquired images (Figure 2.13). It is therefore possible to obtain images that selectively enhance the difference in the amount of longitudinal magnetisation recovered after the RF pulse is applied in different tissues, giving a so-called T_1-weighted image. Similarly, T_2-weighted images (or T_2*-weighted images), capture differences in the decay of the transverse magnetisation.

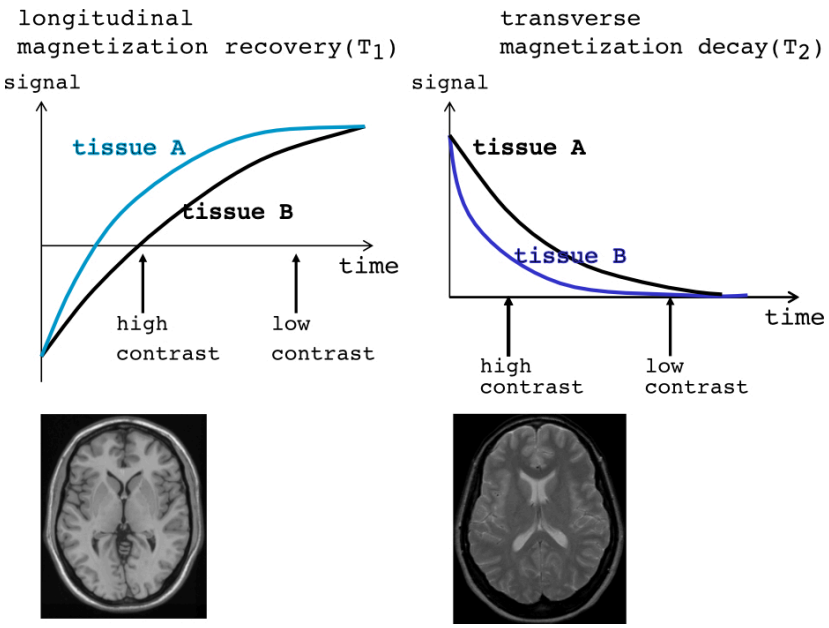


Figure 2.13 Comparison of T_1 and T_2 weighted images and the corresponding T_1 and T_2 values for different tissue types.

In general, an RF pulse is not delivered in isolation, but rather as a sequence of pulses perturbing the system. How often this excitation happens is determined by the *repetition time* (TR), that is the time interval between successive RF pulses are applied.

Another parameter that plays an important role in the acquisition of the signal is the *echo time* (TE). As briefly described above, after the system has been excited through the RF pulse, the spins are precessing in phase, but over time they get out of phase due to the inhomogeneities of the magnetic field and cause the signal to decay more rapidly than expected.

One method that allows accounting for the effects caused by the differences in strength of the magnetic field consists in applying another RF pulse that shifts the net magnetisation vector by 180° (following an initial RF pulse of 90°, Figure 2.14). In this way, the spins get back in phase with each other and generate an increase in the signal or *echo*, and the peak of the signal is determined by the T_2 decay parameter (rather than T2*).

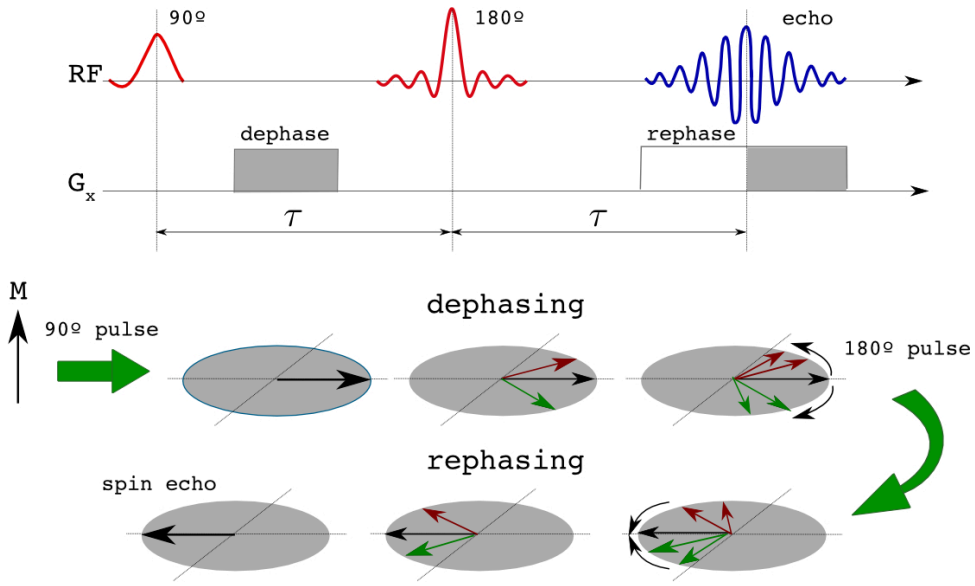


Figure 2.14 Schematic representation of the refocusing effect after applying a 180° pulse in spin echo sequence. Adapted from Sanchez-Panchuelo, 2009

TE, therefore, consists in the amount of time that passes between the RF pulse and the peak amplitude reached during the echo sequence. This specific sequence of pulses is called *spin-echo* (SE) and represents one of the imaging sequences often used in MRI, albeit mostly for anatomical images.

The main acquisition method used in functional imaging is *gradient echo* (GE) and is dependent on the T_2* parameter. A graphical description is provided in Figure 2.15 and its specifications are provided in Section 2.3.3.2.

2.3.3 Image formation in MRI

In order to obtain a 3D image of the brain, information regarding the spatial location from which the signal originates is required. This is obtained by applying *magnetic field gradients* during both generation and acquisition of the signal, so that a sequence of 2D images is recorded that can later be stacked together to obtain a 3D representation of the brain. In this multi-slice imaging scheme, the process of *slice selection* determines the plane in which the 2D information is obtained. *Phase* and *frequency encoding* allow the determination of where in this 2D plane arise.

2.3.3.1 Slice selection

Applying a gradient when the RF pulse is delivered allows to selectively excite a portion of the spins in a volume (Garroway, Grannell, & Mansfield, 1974). For example, a gradient oriented in the z-direction (G_z) will lead to the excitation of a restricted sample of spins along this specific axis, thus resulting in the so-called *slice selection* process.

2.3.3.2 Phase and Frequency encoding

Once the RF pulse is delivered and the net magnetisation tilted from the z to the x-y plane, a magnetic gradient can be applied in order to

systematically induce a manipulation of the phase at which the spins are precessing.

G_y is therefore a magnetic gradient applied along the y-axis of the transverse plane, affecting the phase of the spins excited by the RF pulse. On the other hand, a gradient applied along the x-axis (G_x) of the transverse plane allows to manipulate the *frequency* at which the spins are rotating.

Thus, G_z allows to specify the location along the z-direction of the excited area, while G_x and G_y define the x and y locations on the transverse plane by the joint encoding of frequency and phase of the measured signal.

Gradients in the x and y directions can be applied at different moments depending on the imaging sequence used (e.g., *spin echo* or *gradient echo*). For example, G_x can be applied during the formation of the echo in a spin echo (SE) sequence (see Figure 2.14), but it can also be used to generate the echo in a gradient echo (GE) sequence (Figure 2.15). The latter requires the application of a negative gradient along the x-axis to be applied when G_y is applied, both generating a dephasing in phase and frequency of the excited spins, followed by a positive gradient re-focusing the dephasing and producing an echo from which the signal is acquired.

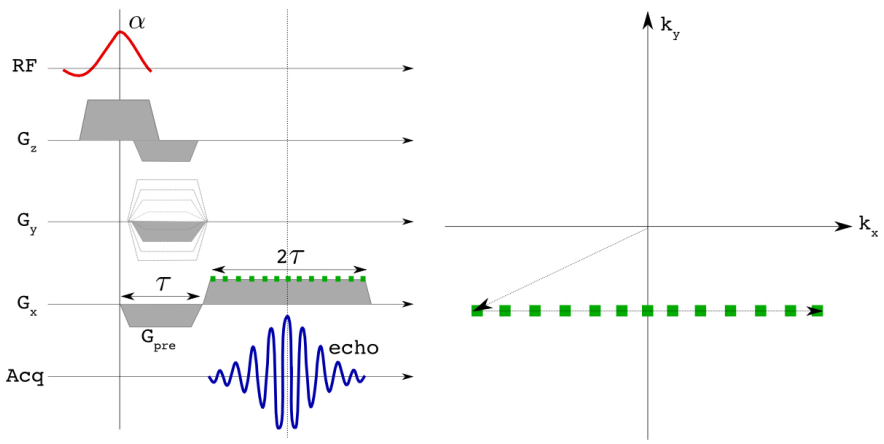


Figure 2.15 Schematic representation of gradient echo sequence. The phase encoding gradient G_y is applied generating a dephasing in the spins. When the area represented by the read-out gradient G_z matches that of G_{pre}, an echo is formed and the signal is obtained at the samples indicated by the green squares. Adapted from Sanchez-Panchuelo, 2009

2.3.3.3 *K-space and Image reconstruction*

The MR signals elicited with the different sequences are measured in the frequency domain, which is also called *k-space* (Ljunggren, 1983). In 2D k-space low spatial frequencies are represented at the centre, giving information on contrast of the image, while high spatial frequencies lie at the edges of the space and inform on fine details (e.g., tissue boundaries; see Figure 2.16). A point in k-space is defined by a k_x and k_y coordinate pair and by acquiring measurements at multiple (k_x, k_y) points, enough information filling the k-space is obtained.

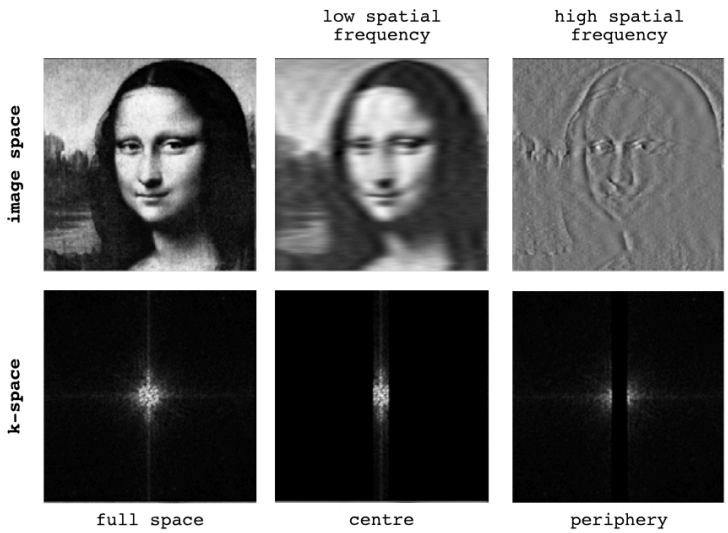


Figure 2.16 Overview of contribution of different parts of k-space to an image. The first column shows the full Fourier transform of the image in the first row. The second column shows how low spatial frequency information in the image space is provided by the centre of the k-space. The third column (right) shows that high spatial frequency information is provided by the periphery of the k-space. Adapted from Huettel, Song, & McCarthy, (2009)

The way in which the k-space is filled reflects the sequence used to acquire the signal. Echo-Planar Imaging (EPI, Mansfield, 1977) is a standard

imaging sequence where the whole k-space is filled after the application of a single RF pulse (rather than having to apply an RF pulse to acquire a portion of k-space at a time). In so doing, the total scan time is drastically reduced, making this technique suitable for measuring dynamic processes such as changes in brain activation, but at the same time the images will be more susceptible to artefacts and distortions due to variations in the magnetic field.

After the k-space is filled in, the *inverse Fourier transform* can be used to convert the information from the frequency to the conventional image domain. It is important to note that the data measured in k-space are complex numbers: they have a real and an imaginary part. In routine work, we usually only consider *magnitude* of the reconstructed images ignoring information about image phase.

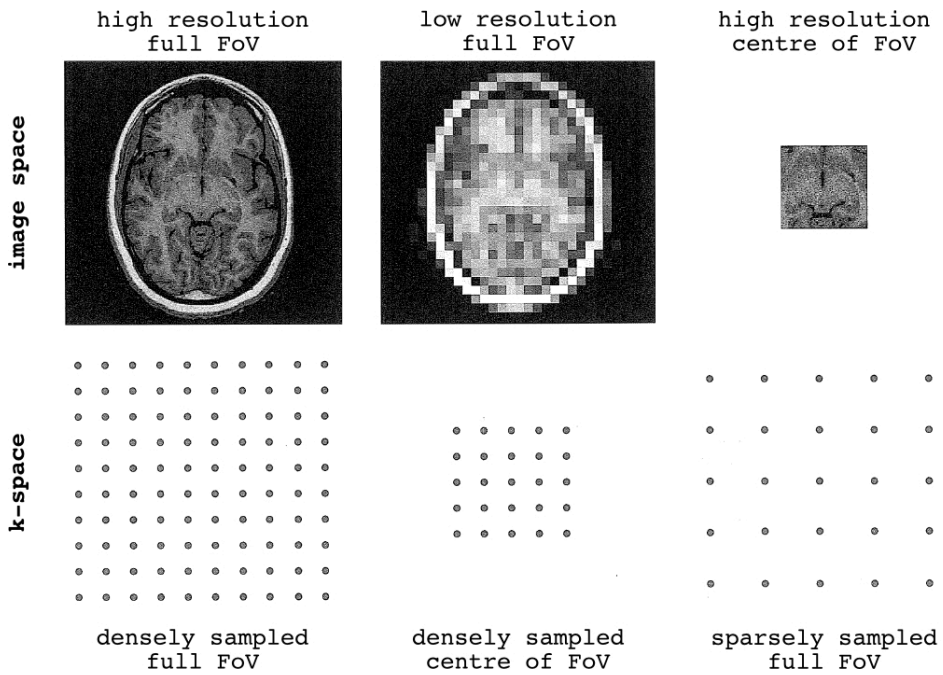


Figure 2.17 Representation of the effects of sampling in k-space. The first column (left) shows a densely sampled k-space with a large field of view (FoV) resulting in a high-resolution image. The middle column shows sampling of k-space with the same density but reduced FoV. The resulting image presents a low resolution. The third column (right) shows a large FoV but less densely sampled k-space resulting in a portion of the image at high resolution. Adapted from Huettel et al., 2009

The process of *image reconstruction* depends on the number of measurements acquired (how and where k-space was sampled), which determines the resolution and coverage of the reconstructed images (Figure 2.17).

It is important to note the relationship between points in k-space and the reconstructed images. Each representation of the data contains the same information in total, but a point in k-space has contributions from all voxels within the selected slice (Figure 2.18) and can be visualised as a linear combination of magnetisation vectors within individual voxels.

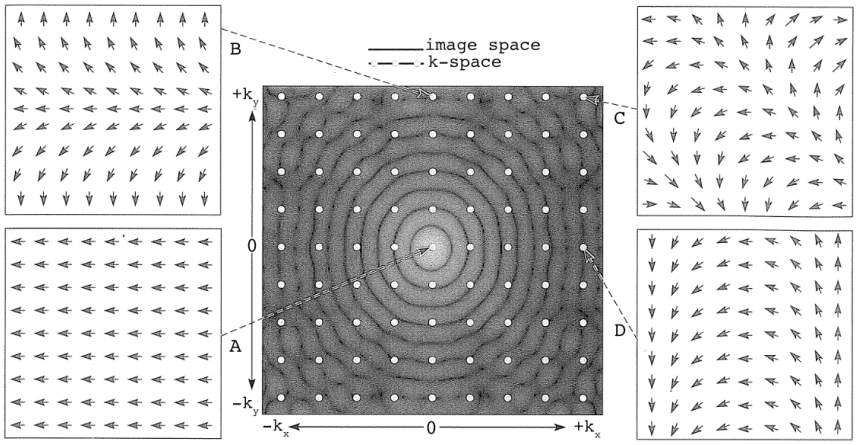


Figure 2.18 Overview of the contribution of different image locations to points in the k-space (A, B, C, D). Each k-space point is a summation of signals from the image space. The point at the centre of k-space presents maximum signal, represented by the fact that the phases (arrows) in the image space are all identical. For points in k-space where either of the dimensions (k_x or k_y) is at its maximum, the corresponding phase representation in the image space changes rapidly along this dimension (B, D). For points in k-space where both k_x and k_y are at their maximum, the phase representation in the image space changes rapidly along both dimensions (C). Adapted from Huettel et al., 2009

In multi-slice imaging, a 3D image of the object is acquired as a series of 2D slices that are later combined. For certain applications, especially high-resolution MR imaging – both for anatomy and function –, data is acquired in 3D k-space, followed by two phase encoding steps such that the excited *volume* can be spatially resolved (3D-EPI). For most published functional MRI studies, however, the multislice approach, acquiring a sequence of 2D

slices at each time point in quick succession, is the main stay (see Section 2.3.4).

2.3.4 Functional MRI

To study 'brain' function during a particular cognitive process, we need many measurements over time. By repeatedly acquiring a set of image volumes (over time) that are sensitive to changes in local blood oxygenation, functional MRI (*fMRI*) can provide important information about the areas of the brain that are involved in such processes. The data acquired with fMRI are usually T_2* weighted (gradient-echo) images obtained with the echo-planar (EPI) technique. In most studies, the temporal resolution is usually kept high, compared to the time constants of the haemodynamic changes: the time between repeated measurements, the repetition time or TR, is therefore around 2-3s. To increase coverage or spatial resolution, longer TRs are sometimes used.

The key idea underlying fMRI experiments is to link changes in measured fMRI responses to specific experimental manipulations to establish *where* in the brain those changes occur and how they are modulated by the stimuli and experimental conditions.

Functional MRI relies on differences in the magnetic properties of oxygenated and de-oxygenated haemoglobin in the blood and the changes in blood supply following changes in neuronal activity. The most commonly used technique measures this *Blood Oxygenation Level-Dependent* (BOLD) signal. In the following sections I will briefly discuss the principles underlying

this measurement and the relationship between neuronal activity and dynamic changes in magnetic properties of haemoglobin.

2.3.4.1 Blood Oxygenation Level-Dependent (BOLD) signal

The BOLD contrast mechanism was demonstrated in the early 1990s and has become the dominant technique for measuring brain activity non-invasively (Ogawa et al., 1990, 1992). The technique relies on measuring signal changes that result from a changing ratio of oxygenated to de-oxygenated haemoglobin in the blood. It is important to keep in mind that BOLD is not a direct measure of neuronal activity, but rather an indirect measure reflecting changes in cerebral blood flow (CBF), local cerebral blood volume (CBV), and changes in the metabolic rate of oxygen consumption ($CMRO_2$).

Oxygenated and de-oxygenated blood have different magnetic properties. Oxygenated haemoglobin (Hb) is *diamagnetic*, whereas de-oxygenated haemoglobin (dHB) is *paramagnetic.* Haemoglobin in these two states affects the local magnetic field differentially, especially surrounding the vascular bed around veins and venules: a relative increase in Hb tends to make the local magnetic environment more homogeneous whereas a relative increase in dHB makes the local environment less homogeneous. These small changes in the local magnetic field affect the frequencies at which different protons precess and interact. In the presence of baseline levels of dHB the exponential decay of the transverse magnetisation is relatively quicker (captured by a shorter T_2* time constant). Following a

neural response and ultimately an increase in the local amount of oxygenated blood (Hb), this decay is slowed resulting in a slightly longer T_2* time constant and overall larger responses.

In the original experiments of Ogawa and colleagues (1990), anesthetised rodents were exposed to different levels of oxygenated air while MR signals were obtained with a GE sequence. Baseline images were obtained with animals inhaling 100% oxygen. In the 'normal air' condition (21% oxygen), thin, dark lines appeared in the images. These lines were even more salient as the concentration of oxygen decreased further when rodents breathed a gas mix with 0% oxygen. The authors reasoned that the darker lines would result from the distortions of the magnetic field due to the de-oxygenated haemoglobin present in the blood vessels, and decided to test this hypothesis by scanning tubes filled with either oxygenated or de-oxygenated blood, using GE and SE sequences. Note that images obtained with GE sequences are susceptible to T_2* distortions, whereas SE images are only susceptible to T_2. As a result, the images of tubes containing oxygenated blood presented well defined black circles for both SE and GE sequences, and the same happened for images acquired with SE sequence of tubes containing de-oxygenated blood (Figure 2.19).

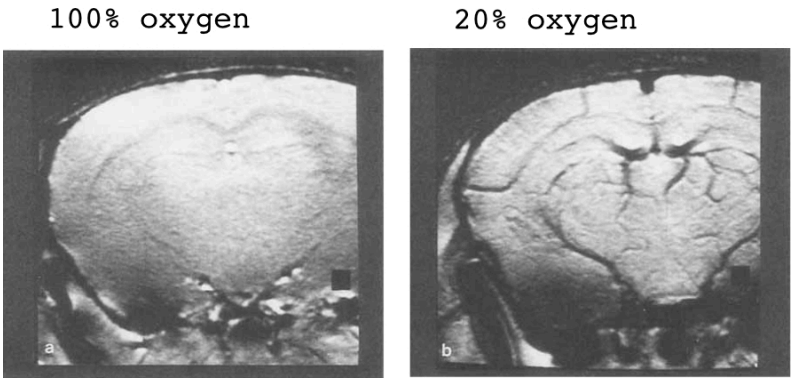


Figure 2.19 Coronal slices of a mouse brain. On the left image of the brain slice after breathing 100% oxygen. On the right a brain slice after breathing standard air (20% oxygen). Most of the black lines correspond to blood vessels. Adapted from Ogawa et al., 1990

However, GE images of de-oxygenated blood presented strong distortions: the MR signal was significantly decreased implying the detection of signal loss due to effects of de-oxygenated blood in T_2*.

2.3.4.2 The Haemodynamic Response Function (HRF)

The dynamics of how the BOLD response develops following a local increase in neural activity can be captured in the so-called haemodynamic response function (HRF). Many studies have looked at the relationship between increase in blood flow and glucose consumption in active brain regions (e.g., Fox & Raichle, 1985, 1986; Roy & Sherrington, 1890; Toga & Mazziotta, 2002). A particularly interesting set of experiments by Malonek & Grinvald (1996) used the known *optical differences* between Hb and dHB to elucidate the link between neural activity, metabolism, and the haemodynamic response. They used high-resolution optical imaging techniques to analyse the patterns of light reflected back by specific parts of the active brain. By meticulously investigating the changes in oxygenated and de-oxygenated haemoglobin during systematic stimulation of the visual

cortex in cats, it was possible to explain the elements leading to an increase in MR signal, despite an increase in oxygen consumption that should result in an increase in de-oxygenated haemoglobin concentration, a faster T_2* decay, and signal loss.

The recovered curves relating oxygenated / de-oxygenated concentration of haemoglobin as a function of time (Figure 2.20) showed that after 2s from stimulus presentation, the concentration of de-oxygenated haemoglobin reached a peak value and quickly declined thereafter, down to levels well below those measured at baseline. On the other hand, the concentration of oxygenated haemoglobin took longer to peak (approximately 5 to 6s after stimulus onset), but reached values much higher (in amplitude) than those of the de-oxygenated haemoglobin and its spatial distribution expanded well beyond the expected regions of activation.

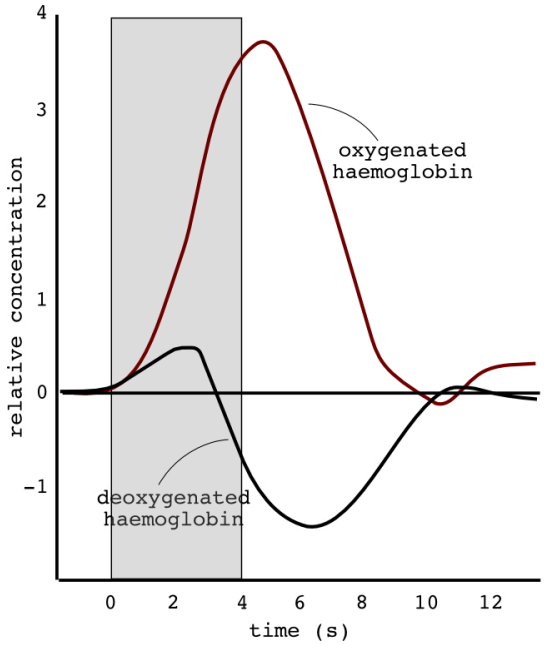


Figure 2.20 Schematic representation of changes in concentration of oxygenated and deoxygenated haemoglobin following stimulation (grey rectangle). While the concentration of deoxygenated haemoglobin increases rapidly in the 2s following stimulus onset, this declines to a minimum value at around 6s. The concentration of oxygenated haemoglobin, on the other hand, shows a slow increment following stimulus onset and a peak value reached around 5 to 6s after stimulation. The concentration then declines slowly reaching baseline after approximately 10s from stimulation. Adapted from Huettel et al., 2009

Given these observations the HRF can be divided into three periods (Figure 2.21):

- An *initial dip*: neurons encoding information start their processing by using the oxygen that is currently available.

This, in turn, leads to an increase in de-oxygenated haemoglobin and a request for oxygen supply (increase in oxygen consumption is related to a momentary decrease in concentration of oxygenated haemoglobin as well as an increase in concentration of de-oxygenated haemoglobin)

- *peak in signal*: oxygen is delivered through an increase in blood flow and this process happens at a rather coarse scale. Oxygen is delivered to a greater area than required by demand. This specific phenomenon has also been described as '*watering the entire garden for the sake of one thirsty flower*' (Turner & Grinvald, 1994).

It is therefore the overflow of oxygenated haemoglobin removing the de-oxygenated haemoglobin that generates the BOLD contrast and therefore the measured increase in MR signal

- *Undershoot*: decrease of oxygenated haemoglobin concentration back to baseline. Baseline is reached after a period of undershoot during which the signal is actually negative and this has been related to blood flow decreasing faster than the blood volume (Buxton, Wong, & Frank, 1998).

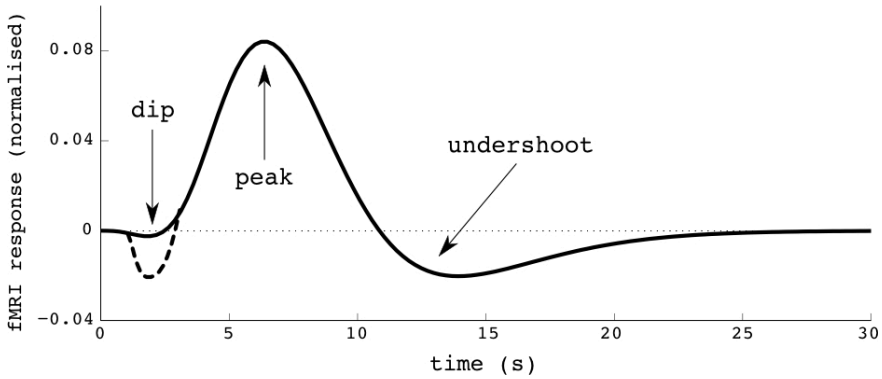


Figure 2.21 Schematic representation of the Haemodynamic Response Function (HRF). The function reaches its peak at around 6s from stimulus onset. The other two features of the function (initial dip and undershoot) are also displayed. Adapted from Sanchez-Panchuelo, 2009

The HRF is assumed to be a linear function of the underlying average neuronal activity (Boynton, et al., 1996). A sequence of stimuli will evoke the same HRF at each instance, an assumption that leads to two fundamental properties of the HRF: *scaling* and *superposition*. Scaling is related to the amplitude of activation: as the magnitude of neuronal activity changes, a proportional change in magnitude is expected in the haemodynamic response. Superposition, on the other hand, is related to the timing of activation: a sequence of stimuli will result in an overall response that is equal to the summation of HRFs evoked by each individual impulse.

These two properties were first investigated by Boynton et al. (1996): by manipulating both magnitude of contrast and duration of visual stimuli, evidence in favour of the linearity of the HRF was presented. Visual stimuli with higher contrast evoked a greater amplitude activation (consistent with the scaling property) and the recovered HRFs presented the same basic shape for each individual stimulus. On the other hand, superposition was tested by predicting responses to stimuli presented for a longer time interval from responses measured from stimuli presented for shorter intervals.

Indeed, linear superposition was observed for most stimulus durations: a response to a 24s stimulation could be predicted by adding responses to two 12s stimulation or four 6s stimulus presentation (Figure 2.22). However, an overestimation of the long stimulus responses resulted from adding together responses to very short stimuli (i.e., 3s presentation). The overestimation was explained in terms of neuronal adaptation: after a few seconds of stimulation, neuronal activation usually decreases to a slightly lower amplitude from the initial burst.

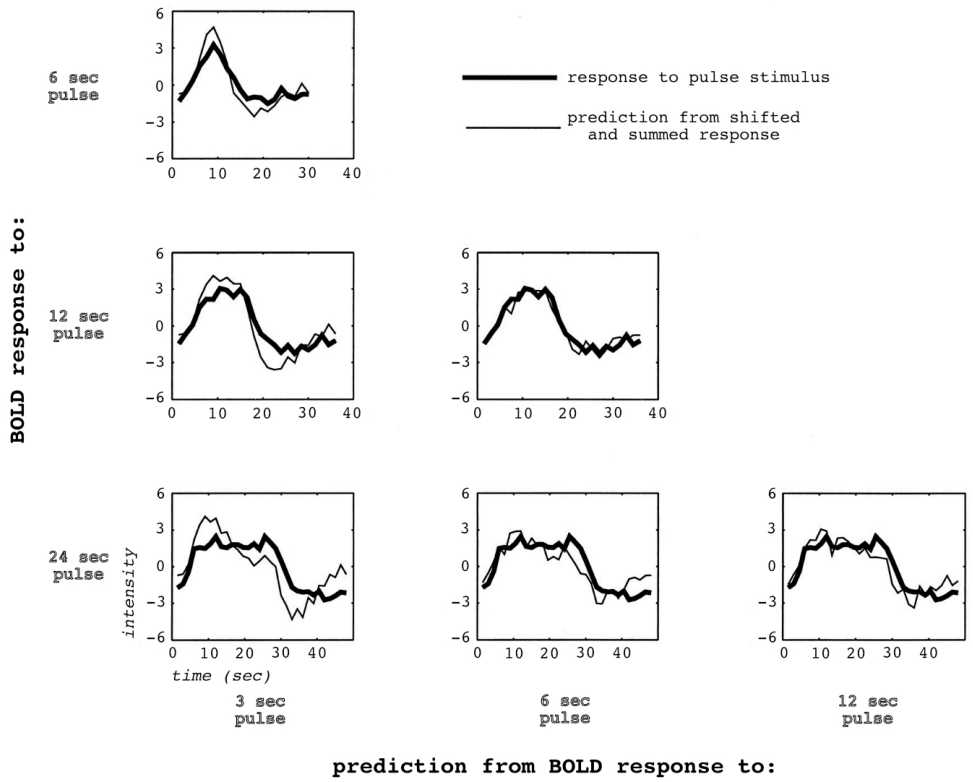


Figure 2.22 Linearity of fMRI haemodynamic response. Thick dark lines show measured responses to stimuli of different durations (3s, 6s, 12s). These responses were used to predict fMRI responses to pulses of longer durations (e.g., summation of 3s pulse response used to predict response to a 6s pulse stimulation). Adapted from Boynton et al., 1996

The fact that the haemodynamic response function approximately fulfils the requirements of the *linear transform model* (Boynton et al., 1996) has useful and important consequences: responses to a sequence of stimuli can be predicted by convolving the *time course* of the stimulus pattern with a description of the HRF. These assumptions also underlie the methods of multiple linear regression – the dominant analysis method for making sense of fMRI data.

A simple method commonly used to estimate whether the BOLD signal is correlated with a specific condition compared to noise / baseline activity / a complementary task consists in computing the correlation between the measured BOLD signal and the response predicted through convolution of HRF and scan time course. By selecting a threshold value, it is then possible

to ‘pick’ those voxels that present a correlation parameter that exceeds the threshold and ‘mark’ them as significantly activated.

An alternative method, that will be further discussed in the next section, looks at the power spectra of time course and measured data, comparing the amplitudes of the derived frequencies to obtain a coherence value indicating regions most responsive to the stimuli. This analysis technique, also called *phase-encoded* (Engel, Glover, & Wandell, 1997) presents some advantages compared to the correlational analysis, especially in dealing with delay onsets in the HRF.

The following sections will then present a brief description of classical methods for fMRI data analysis, from *General Linear Model* (2.3.6.1), to *Retinotopic Mapping* (2.3.6.2), *Population Receptive Field Analysis* (2.3.6.3). Finally, I will focus on a more advanced modelling technique, termed the *Forward Encoding Model* (2.3.6.5).

2.3.5 Data Pre-Processing

Functional data from all scanning sessions presented in this thesis underwent two stages of pre-processing (Section 2.3.5.1 and Section 2.3.5.2) in order to minimise variability in the measured signal due to head motion (biological sources) and scanner-related instabilities (high-frequency noise). Data from the anatomical scans (high-resolution and 2D inplane) were used to obtain maps of the cortical surface (Section 2.3.5.4) and flattened patches (Section 2.3.5.5). Functional and anatomical data were combined

during the alignment pre-processing step (Section 2.3.5.3) in order to register data from one subject and multiple sessions onto a common space. Pre-processing of functional data was performed using mrTools toolbox (available at http://gru.stanford.edu/doku.php/mrTools/overview), anatomical images were processed using a combination of mrTools and Freesurfer (http://surfer.nmr.mgh.harvard.edu/).

2.3.5.1 Motion Correction

Despite applying pillows and padding in trying to minimise head movements while subjects lay in the scanner, some level of motion is still detectable. This can cause distortions in the images: it is fundamental that each voxel is consistently located at the same anatomical place for each time point within a scanning session, but with head motion this consistency cannot be ensured. Motion correction is a processing step aligning each image to a reference one. A rigid-body approach (Nestares & Heeger, 2000), using six motion parameters (three for displacement and three for rotation), was employed to compensate for head movements both within and between scans for an individual session. Motion correction parameters for each subject were subsequently inspected: in one of the five participants tested this resulted in large shifts that could not be corrected. This subject was excluded from subsequent analyses.

2.3.5.2 *Temporal Filtering*

Functional data are affected by both physiological and scanner-related sources of noise (see Section 2.3.5). These can contribute to the measured signal at low and high temporal frequencies and can mask the underlying signal we are trying to measure. Temporal filtering allows to remove high and / or low frequency components in the data, without affecting the response. For the data used in this thesis, only temporal high-pass filtering was applied to remove slow temporal drift.

2.3.5.3 *Registration and Alignment*

This pre-processing step can be divided into two stages. In the first stage, an inplane anatomical image – acquired using the same slice parameters adopted for the functional data – is aligned to a high-resolution anatomical image. The whole-head, high-resolution anatomy image was acquired in separate sessions for each subject. It is useful to think of the image that will be moved / aligned as the "source" image and the target – here the whole head anatomical image – as the *destination* volume. In the second stage, the functional images (EPI, the 'new' source volumes) are registered onto the aligned inplane (the 'new' destination). In so doing, functional data can be overlaid onto a common, high resolution, anatomical space.

2.3.5.4 *Cortical Surface Extraction*

MRI images are typically shown as a 2D slice through the 3D brain volume: while displaying a high level of anatomical details, this can affect the visualisation of relationships between points across the brain. For example, points that are anatomically close may be located in different display slices and therefore rendered far away on the page.

To overcome this limitation, it is possible to extract boundaries delimiting grey and white matter, as well as those delimiting grey matter and cerebral spinal fluid (CSF), using information of image intensities from different tissue types. By adjusting an initial template to the target brain, this is then used to segment the boundary between grey and white matter (Larsson, 2001). This, together with the surface representation of the white matter of the target brain, is then used to obtain a representation of the outer grey matter structure. Provided these surfaces and information on their curvatures (i.e., sulci and gyri) are available, a representation of the cortical surface is obtained.

2.3.5.5 *Flat Maps of Cortical Surface*

3D renderings of the cortical surface can be visually appealing, but hard to interpret. By "cutting out" a patch and computationally flattening it, we can arrive at an easier way to interpret 2D map of the cortical surface. This process also allows a clearer display parts of the brain that are otherwise hidden within the sulci.

The cortical surface for each cortical hemisphere is first inflated (nearly to a sphere) in order to expose part of the cortex hidden within the sulci. While keeping the curvature information, the flattening procedure can be thought as reorganising the vertices of the surface onto a 2D plane, trying to maintain the relative distance of such vertices. Projecting the surface of a sphere onto a 2D plane requires cuts and distortions – a problem often referred to as the map-makers' problem. The approach taken here is to flatten relatively small patches of the cortical surface to minimise local distortions between vertices of the surface. This happens at two levels of resolution: at a low resolution displacement parameters obtained through rearrangement of vertices are computed. These parameters are then applied to the high resolution surface where resampling of vertices is undertaken again, until the optimal organisation is found and a flat map is obtained. At this point, the values corresponding to the curvature information in the surface are binarised to indicate whether a point belonged to a sulcus or gyrus.

2.3.6 Statistical Analysis

Different techniques of statistical analysis for fMRI data are available, with various goals:

- *Localise* brain areas that are most responsive to a specific task / stimulus
- Identify *networks* connecting multiple brain areas involved in processing specific information / performing a well-defined function

- Make *predictions* about brain states from observed data

The most commonly used technique allows to map task-specific brain areas, which is particularly helpful in increasing understanding of the general organisation of brain function as well as localising biological markers for investigated behaviours. This type of analysis generally begins by building a model / description and testing it against measured data for each voxel in the brain. Because each voxel is considered individually and assumed to be independent from others, the approach is known as *univariate analysis*. An example of univariate analysis is the *General Linear Model* (GLM), or multiple linear regression.

2.3.6.1 The General Linear Model (GLM)

The principle underlying the General Linear Model approach is to consider the data as a linear combination of *predictors* (or *regressors, explanatory variables*, *EVs*) and noise. These predictors describe signal changes over time with a particular shape but unknown magnitude. The GLM analysis consists in estimating the magnitude of the predictors: obtain how much (a 'weight' value) each function constituting the model contributes to the measured fMRI data.

The standard format of the GLM model is the following (Figure 2.23):

$$\boldsymbol{Y} = \boldsymbol{X\beta} + \boldsymbol{\varepsilon} \qquad \text{where } \boldsymbol{\varepsilon} \sim \boldsymbol{\mathcal{N}}(\mathbf{0}, \boldsymbol{V})$$

Y is a vector of length N (the number of points in the timeseries). Each Y_i corresponds to a voxel measurement at successive time points.

X is the *design matrix*, containing information about the task, as well as other information about sources that can explain variance in the data, included as nuisance regressors;

β is a vector of *regression coefficients*, the weights to be applied to each column of the design matrix;

ε is a vector of unexplained variance, generally assumed to be normally distributed with mean zero and a variance-covariance matrix **V** that strongly depends on the model used to describe the noise.

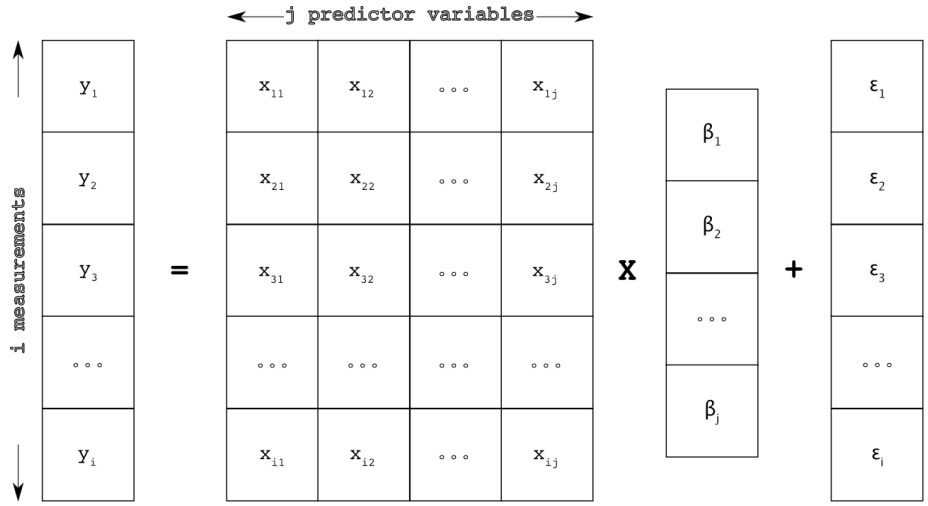


Figure 2.23 Representation of the GLM model. The measurements (y) can be described as a weighted sum of the design matrix (predictor variables) and the unknown beta weights, plus some error.

It is clear from this description that the quality of the model depends on how well the design matrix captures the explanatory variables constituting the data (i.e., how well **X** is defined, accounting for physiological and non-physiological factors), but also the ability to capture the noise structure accounting for the variance that cannot be otherwise explained.

So, while **Y** (a voxel time-series) and **X** (user defined design matrix) are known, the β weights (regression coefficients) are unknown and need to be

estimated. The most common way of achieving this is by using *Ordinary Least Squares* (OLS), which makes two main assumptions:

a) The noise term is assumed to be independent identically distributed (iid): the level of noise is independent from one time point to another during which the signal is measured
b) In a similar fashion, the explanatory variables are independent from one time point to the next

Provided the assumptions are met, the optimal estimates of the beta weights (or *best linear unbiased estimates*, a unique unbiased best estimate of the parameters) are obtained via *linear regression* that in matrix format look like this:

$$\boldsymbol{\beta_{hat}} = (\boldsymbol{X'X})^{-1}\,\boldsymbol{X'Y}$$

In order for the model to work, a few issues need to be taken into consideration:

- The BOLD response has a delayed and dispersed form (recall the shape of the HRF and the fact that its peak is delayed by approximately 6 s from the stimulus onset, likewise its spatial definition is quite broad, i.e. a larger area than the immediately active one is supplied with oxygenated blood)
- The signal measured does not only reflect the activity related to a specific function of interest, but it also includes low-frequency noise

In practice, these two issues are usually dealt with by (a) convolving the 'neural' response prediction with a haemodynamic response function with fixed / assumed parameters and (b) preprocessing the raw fMRI data to remove signal drift and other low frequency noise contributions.

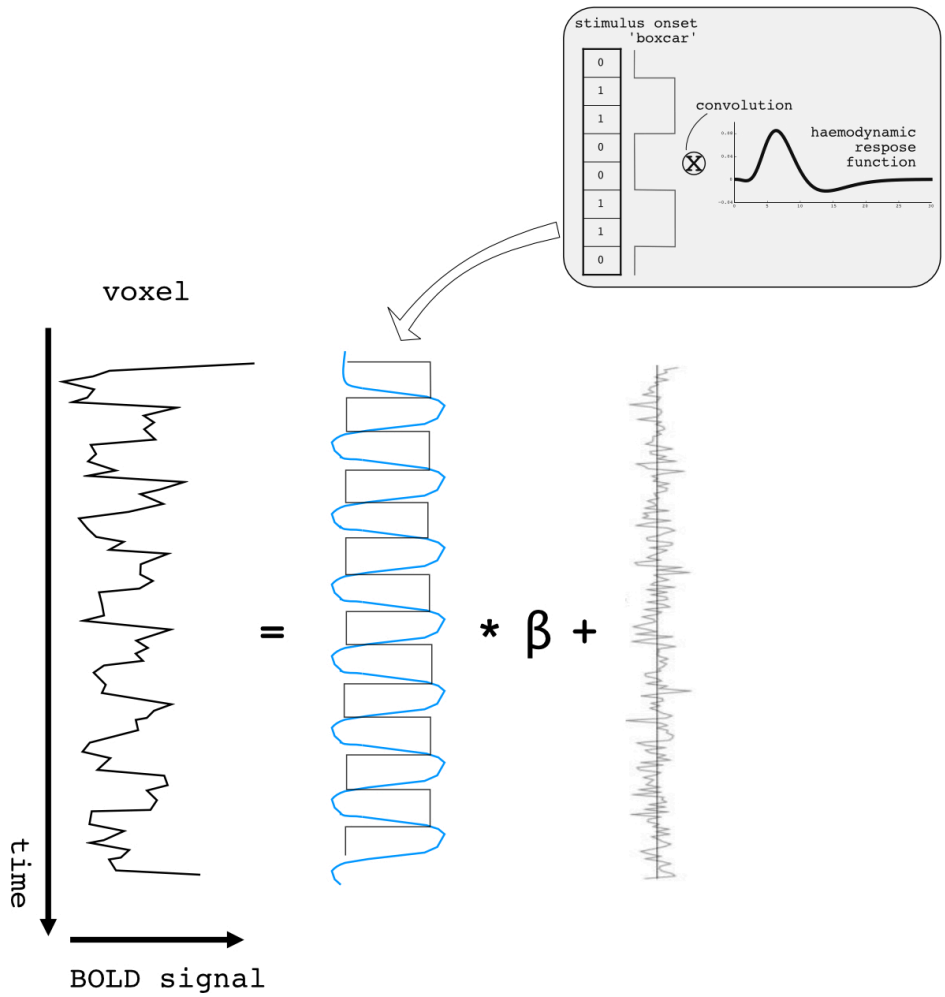


Figure 2.24 Visualisation using fMRI data. The measured BOLD signal (y) can be described as a weighted sum of the design matrix (convolution of the task sequence with the modelled HRF) and unknown beta weights, plus some noise.

Descriptive and *inferential* statistics can be performed straightforwardly in this GLM framework. For example, the *coefficient of determination* (R^2) represents a descriptive measurement of the model-fit to the data, that is how good the model is at explaining the variance in the recorded activity. Similarly, the raw beta weights can give an indication of which explanatory

variables explain variance observed in a particular timecourse. Inferential statistics within the GLM framework usually relies on computing *contrasts.* These compare the computed weights of one or more EVs against each other and test whether they are statistically different for the BOLD signal in a particular voxel. Statistical significance of a contrast can be tested against the *null hypothesis* that no difference is expected.

A more general statistical test on contrasts is the F-test. Here, instead of assessing the significance of a specific contrast against the null hypothesis, we check if any of the included contrasts is significantly different from zero. Analysis with contrasts and inferential tests based on them is one of the analysis methods that leads to the classic 'blobs' of significant activity displayed on example MR images.

While the GLM is a powerful and straightforward to implement technique, other methods, with a different focus have also become popular. In the next two sections (2.3.6.2, 2.3.6.3) we will explore *phase-encoding* (or travelling wave) techniques used for retinotopic mapping.

2.3.6.2 Retinotopic Mapping (phase-encoded technique)

Spatial information about the visual field is systematically mapped in the visual cortex (see also Section 1.2). Neighbouring points in the visual field are represented by neighbouring cells in the visual cortex resulting in a so-called *retinotopic map*. Several visual areas can be identified in the visual cortex, each representing a mapping of the visual field with discontinuities

of the representations marking the borders of different areas (Sereno et al., 1995).

Retinotopic maps are often generated using a method originally described by Engel et al., (1997). The technique adopts stimuli with well-defined parameters. Checkerboard-like patterns reverse contrast periodically while the spatial window (wedges, rings, bars) in which they are presented move at a fixed speed on the screen. The effect of this is a modulation of stimulation of the visual field both in time and space. Originally, two types of stimuli were proposed for this type of analysis, essential to map out visual space in *polar coordinates*:

a) An expanding and contracting ring used to build *eccentricity* maps - to measure r, the distance of the stimulated region of the visual field from the central fixation point (Figure 2.25A)
b) A rotating wedge used to build *polar angle* maps – to measure *theta* or polar angle, the orientation of the visual field representation with respect to the horizontal and vertical meridians (Figure 2.25B)

These stimuli represent the classical choice used in mapping visual areas of the brain, but more recent studies have shown that different configurations of stimuli can also be used (Dougherty et al., 2003; Dumoulin & Wandell, 2008).

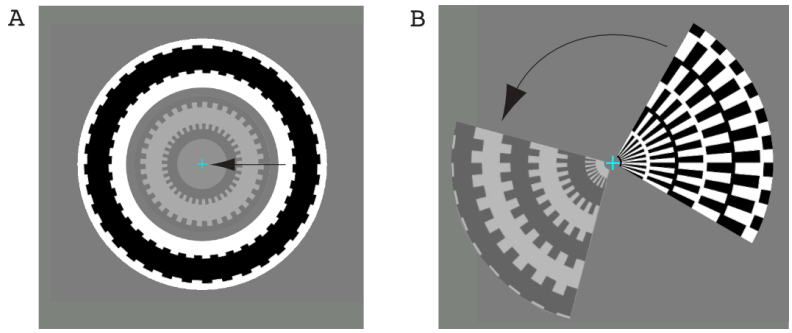


Figure 2.25 A) Concentric rings contracting (arrow). B) Wedge stimuli rotating CCW (arrow). Fixation cross is presented in cyan and shaded regions represent future stages in the moving stimuli.

At the core of *phase-encoded* retinotopic mapping, (Engel et al., 1997), is the fact that neurons represent a delimited *receptive field*. As the stimulus moves across the receptive field, neurons experience a change from the static (homogeneous) background to an alternating stimulus. Thus, as the stimulus travels across space and periodically enters and leaves adjacent receptive fields, this results in a periodic pattern of neural responses at adjacent locations in visual cortex. As the stimuli are moving at a constant speed, the phase of neural response modulation is slightly different at adjacent locations. Another way to think of this is that it sets up a travelling wave of activity across the cortical surface. It is this property of the responses that has been exploited widely to produce maps of visual areas. The processing involved in computing the actual maps can be summarised as follows (Figure 2.26):

1) fMRI timeseries in response to retinotopy stimuli (i.e., both checkerboard-like rotating wedges and expanding / contracting rings) are measured and converted to power spectrum via FFT (keeping information about phase). The power at the stimulus *repetition frequency* (relative to all frequencies in the data) captures

how well a particular voxel is modulated by the periodic stimulus. This is captured by the *coherence value*.

2) The coherence value can be used to identify most responsive regions of the brain (here the visual cortex) and obtain a first map that can be superimposed onto a high resolution, anatomical image of the brain to aid locating the active regions
3) retinotopy maps are built by taking into account the phase of the stimulus repetition and estimating temporal offset of the best fitting sinusoidal function to the timeseries. Because different locations of the visual field are stimulated at times that are consistently offset from the start of a cycle, there is a one-to-one mapping of temporal delay in the periodic signal and visual field location that was being stimulated. This can be turned into a colour-coded map of spatial organisation.

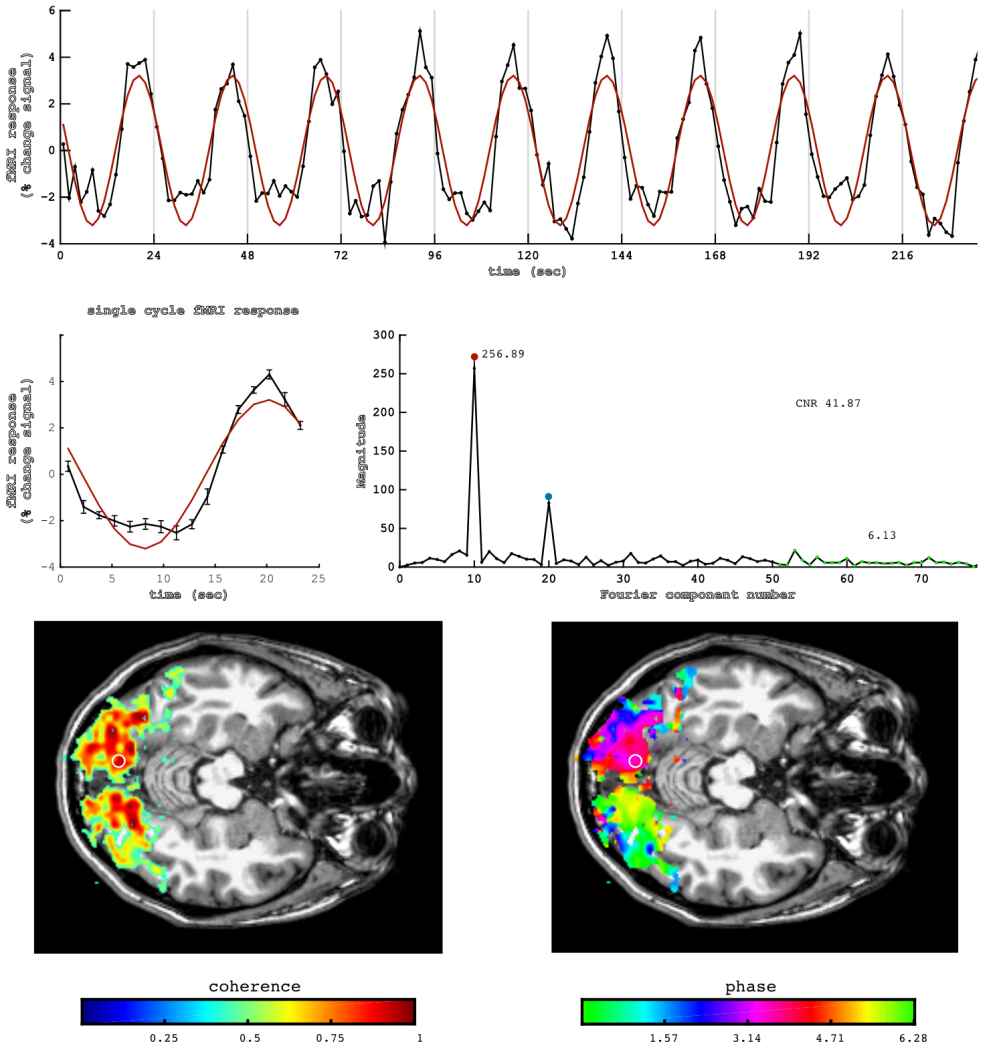


Figure 2.26 Schematic representation of the phase encoded analysis. The top row shows the response of a voxel in visual cortex (identified by the white circle on the images at the bottom row) to the retinotopy stimuli and a fitted sinusoid function oscillating at the stimulus frequency. The second row shows the Fourier decomposition of the timeseries. On the right, the peak, indicating the stimulus repetition frequency is displayed by the red circle, whereas its harmonic is represented by the blue dot. The green dots show the noise level. Bottom row: on the left display of the coherence map superimposed on the high-resolution anatomical image. On the right the phase map is presented.

Eccentricity maps of areas of the visual cortex expand from the occipital pole to more anterior regions along the calcarine sulcus (see Figure 2.27): moving from central to peripheral portions of the visual field, a distinct pattern emerges with a considerate amount of map dedicated to the foveal region (a phenomenon known as *cortical magnification*, Daniel & Whitteridge, 1961).

Polar angle maps, on the other hand, give a representation of the visual angle encoded in specific regions of the visual cortex. In particular, the visual field can be divided into *upper* and *lower* visual meridians (UVM and LVM, respectively), and *horizontal* meridian (HM) (see Figure 2.27).

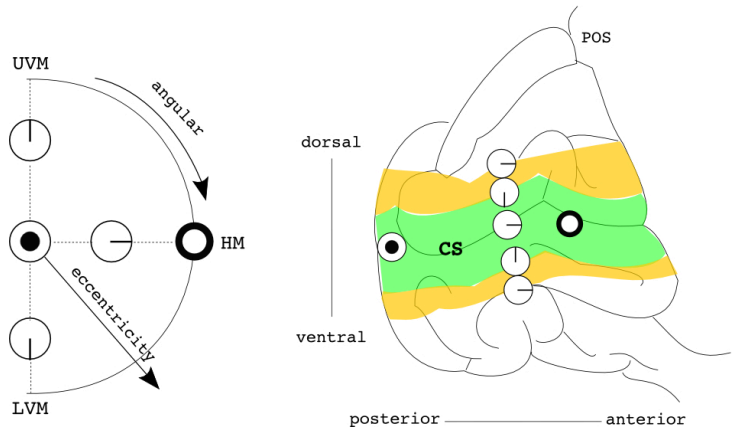
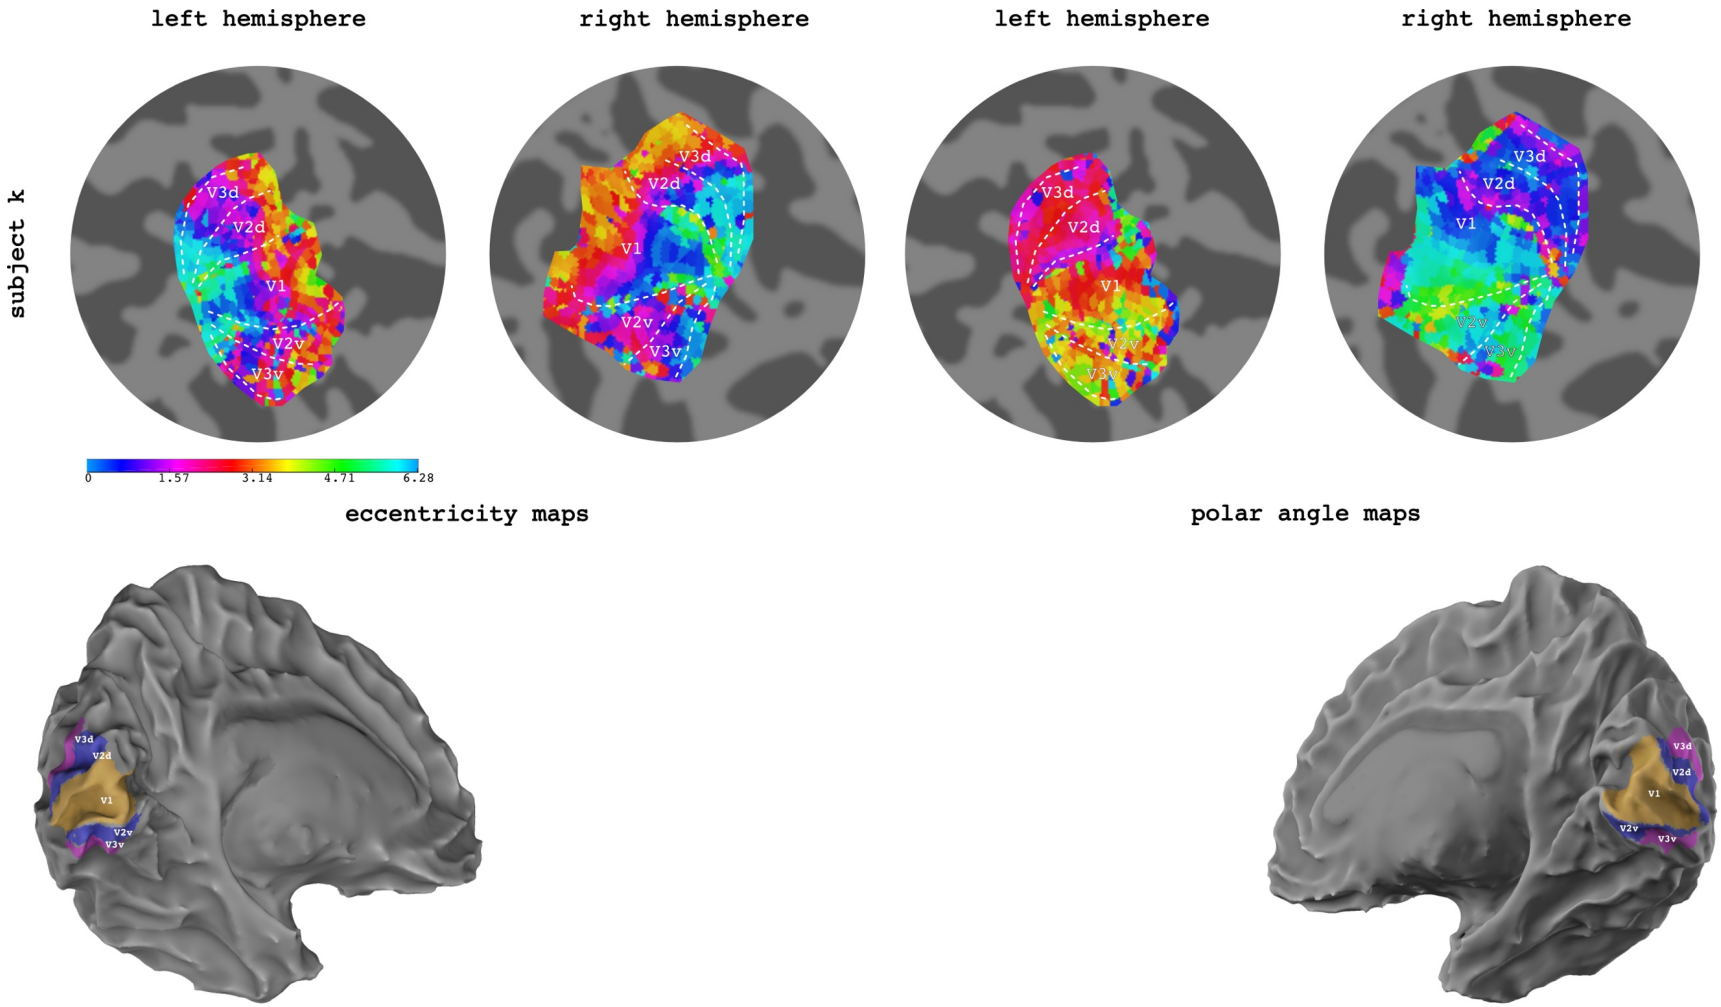


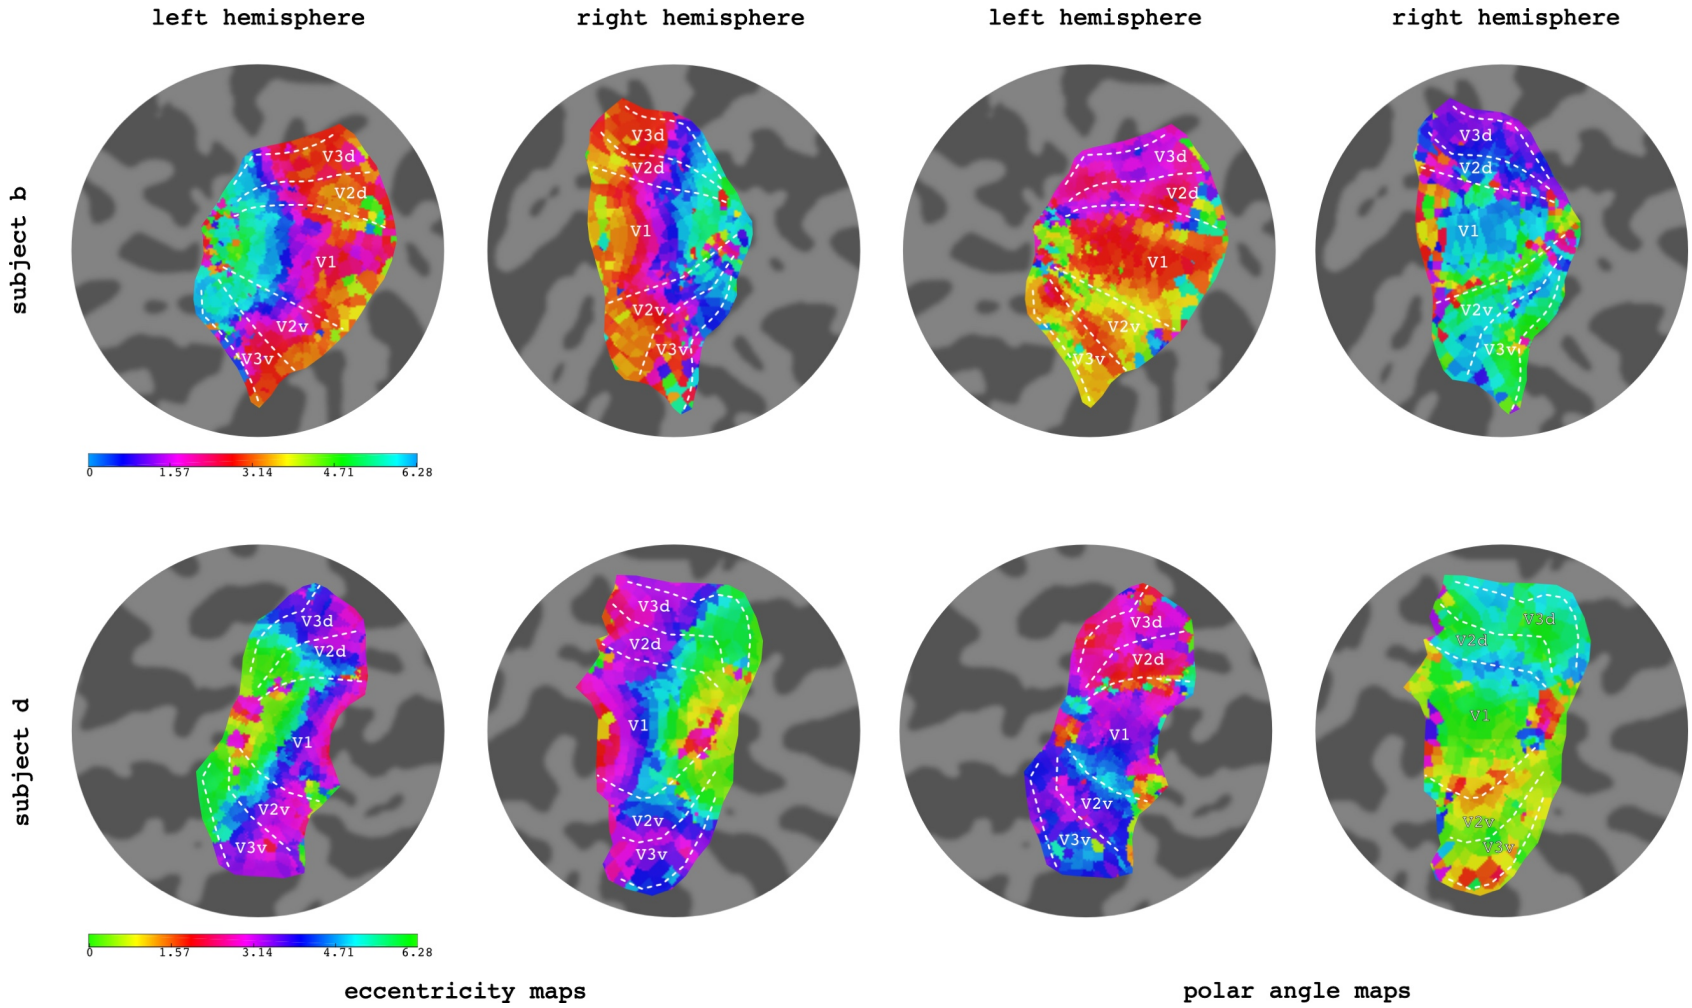
Figure 2.27 Mapping of the visual field onto the calcarine sulcus. On the left: schematic representation of the visual field with symbols indicating different locations (centre is a black dot surrounded by a white annulus, the periphery is indicated by the same symbol in opposite sign). On the right: mapping of visual field locations onto the calcarine sulcus of the left hemisphere. Green stripes indicate visual area V1, whereas V2 is indicated in yellow. Adapted from Wandell, 1999

These landmarks are used to orient along the cortical surface. HM extends along the calcarine sulcus, shifting to the LVM when moving dorsally along the cortex. Responses in the UVM are observed when moving ventrally from the HM (calcarine sulcus). This shows the well-known fact that the representation of the visual field is inverted in the visual cortex: the lower portion of the visual field is represented on the upper portion of the calcarine and vice versa. Another important aspect is the fact that each hemifield is encoded in the contralateral visual cortex (i.e., the left visual field is represented in the right visual cortex, and vice versa; see Figure 1.10 and Figure 1.11).

Polar angle maps are thus used to identify areas of the visual system (Figure 2.28) according to the following landmarks:

- V1 has a complete representation of the visual field (note that each hemisphere has a representation of the contralateral field). HM is located around the fundus of the calcarine sulcus, while UVM and LVM extend on the ventral and dorsal portions of the sulcus, respectively. The borders of V1 are identified by reversals from the representations of upper and lower meridians back towards the horizontal one: these are the edges of the primary visual cortex.
- V2 is divided into a ventral and dorsal portion, with each section representing a quarter of the contralateral hemifield: the upper visual hemifield is represented in V2v (v, ventral), whereas the lower visual hemifield is represented in V2d (d, dorsal). As it was the case for the primary visual area, similarly the borders of V2 are identified by reversals in representations from UVM back to HM for area V2v and from LVM back HM for area V2d.
- Similar to V2, V3 is also divided in a ventral and a dorsal portion, each showing the representation of a quarter of the contralateral hemifield, respectively. Thus, V3v progresses from HM to UVM, whereas V3d maps from HM to LVM.


left hemisphere
right hemisphere
left hemisphere
right hemisphere
subject k
V3d
V2d
V1
V2v
V3v
0
1.57
3.14
4.71
6.28
eccentricity maps
polar angle maps


left hemisphere
right hemisphere
left hemisphere
right hemisphere
subject b
subject d
V3d
V2d
V1
V2v
V3v
0
1.57
3.14
4.71
6.28
eccentricity maps
polar angle maps

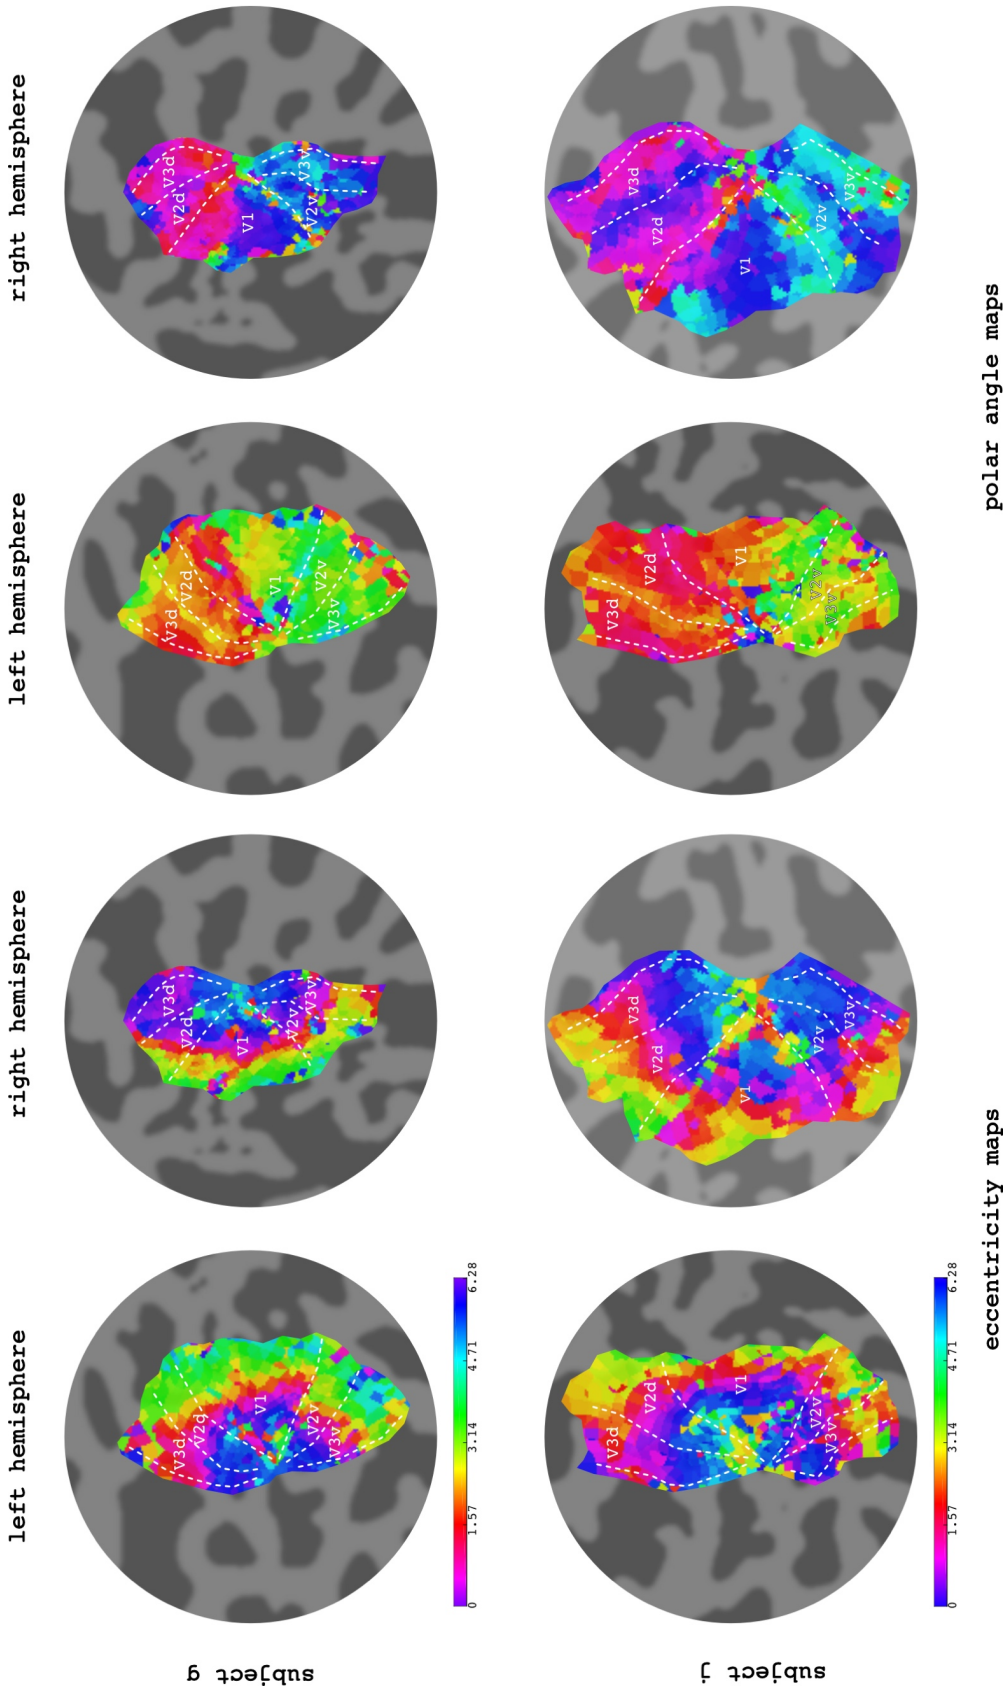


Figure 2.28 Retinotopic maps of early visual areas in both hemispheres for all subjects.

The representation of the visual field beyond areas V1-V3 changes slightly, in that different areas are found in the ventral and dorsal sections of the visual cortex and the symmetry characterising earlier visual areas progressively decreases. In particular, at the border with area V3v studies have reported an area with the representation of the whole hemifield, rather than of a quarter of it as it is for V2 and V3. The region is usually named V4, although it is still debated upon its actual structure and mapping of the visual field (Tootell & Hadjikhani, 2001).

Dorsally, from the border of V3d it has been observed that the LVM splits into two, leading to areas known as V3A/B on the dorsal side, with a full hemifield representation, and LO1 (lateral occipital) to the other end (Amano et al., 2009; Larsson & Heeger, 2006). Following LO1, another area can be identified (LO2), extending from the representation of the LVM back towards HM and then on to the UVM. At this point, the edge with another well-known visual area is observed: V5/MT with a full representation of the visual hemifield (Huk & Heeger, 2002; Kolster, Peeters, & Orban, 2010).

Given the increased selectivity in the responses of cells in higher order visual areas, *localiser* scans are usually performed in order to locate the areas of interest. As an example, relevant to the work in this thesis, Huk and colleagues (2002) obtained polar angle maps by means of rotating wedges defined by moving dots, contrasted to a background of static dots. The response to moving dots, compared to the static condition, is measured and allows to locate V5/MT.

2.3.6.3 Population Receptive Field (pRF)

An alternative analysis technique to estimate maps of the visual field is called *population receptive field* (pRF, Dumoulin & Wandell, 2008; Kay et al., 2008). Unlike retinotopic mapping, the pRF method does not rely on periodic stimuli and allows for a much wider variety of stimuli than the standard rotating wedges and expanding / contracting rings. The benefit of this type of analysis is that not only it allows to map early visual areas, but also to estimate other properties of neurons in a voxel, such as the size of the total receptive field of the ensemble of neurons (pRF size).

In general terms, a *receptive field* can be defined as a description of the way in which a unit such as a neuron (or even a region of interest, ROI) transforms a stimulus into a response. The main purpose is to formulate a link between stimulus feature and brain responses. An example of receptive field modelling is given by Kay et al. (2008). The aim of that study was to build a model able to identify natural images given a set of predictions on voxel activity patterns evoked by an independent set of images. This was achieved in two stages: first, the model estimation. fMRI activity was measured in visual areas V1, V2, and V3 while subjects viewed a large set of images (*training* set). Using these data, a model (Gabor wavelet pyramid model) describing the tuning to space, orientation, and spatial frequency was estimated for each voxel. The parameters characterising these voxels' receptive field models were based on well-known results from neuro-electrophysiology (Campbell, Cooper, & Enroth-Cugell, 1969; Hubel & Wiesel, 1959) as well as previous fMRI studies (e.g., Haynes & Rees, 2005;

Olman et al., 2004; Singh, Smith, & Greenlee, 2000). The second stage consisted in the image identification: voxel activity patterns were measured while subjects viewed a different set of images (*validation* set). The new images were then plugged into the receptive field models estimated during stage 1 in order to obtain a prediction of the voxel activity patterns for the validation set of images. The predictions from the model were then compared to the measured fMRI responses, on a voxel by voxel basis: the image linked to the highest correlation value between model prediction and measured data was taken as the decoded one (i.e., the image identified to have evoked the measured voxel activity pattern). A schematic visual representation of this two-stages approach is provided in Figure 2.29.

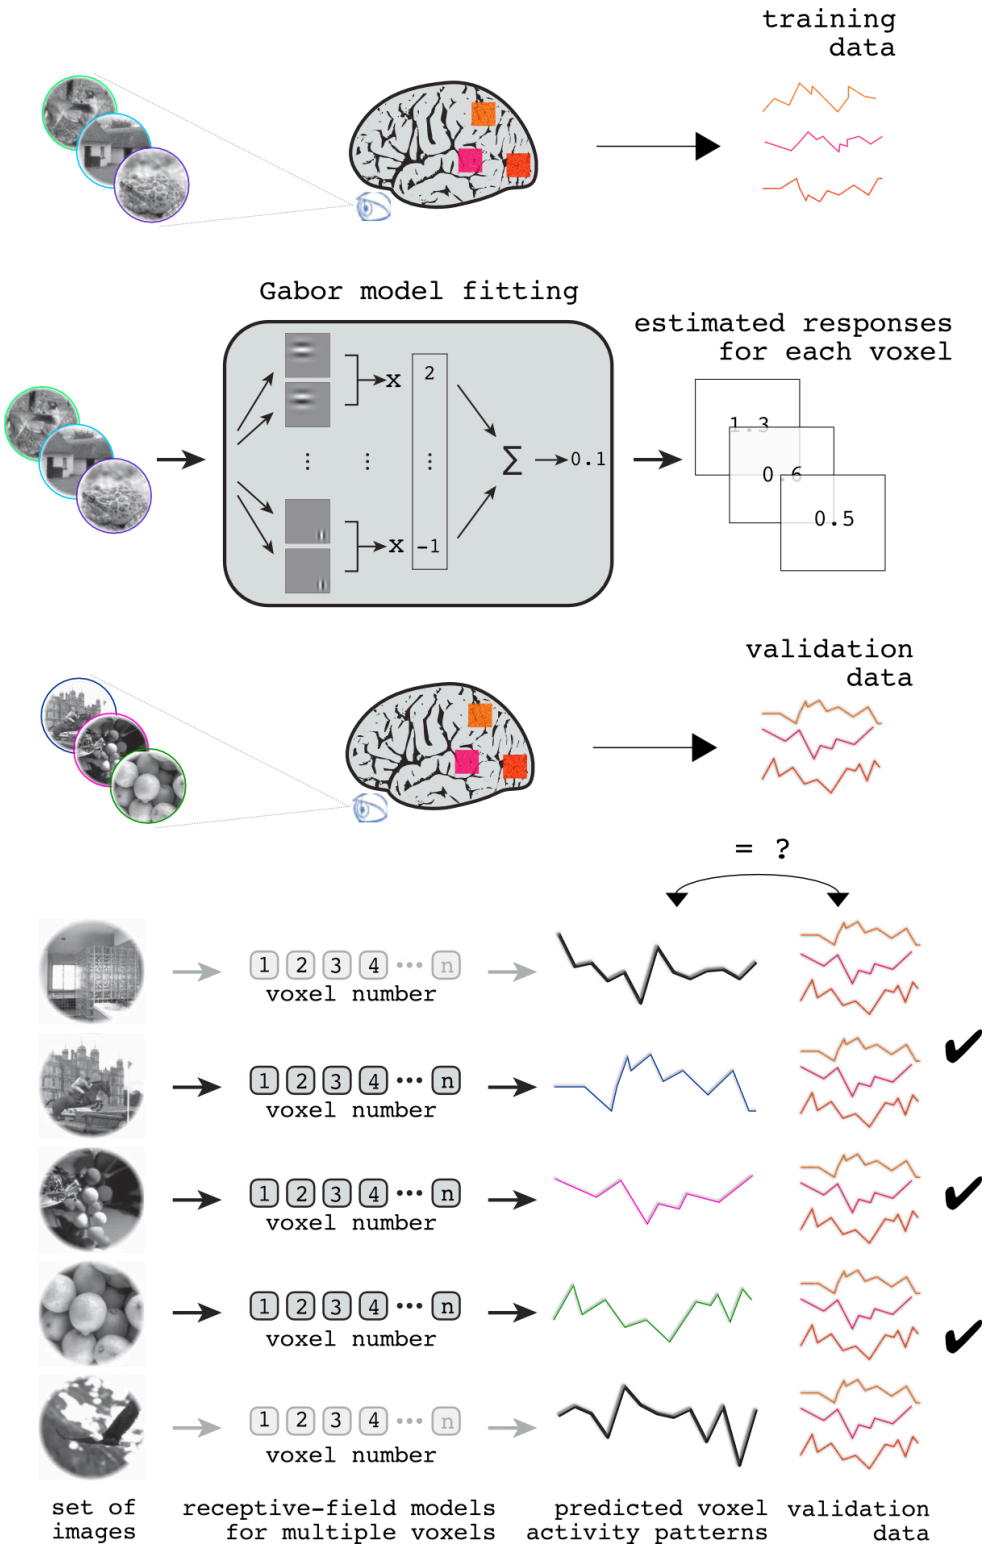


Figure 2.29 Schematic representation of Kay et al.'s model. Row one: subjects are presented with a series of images and fMRI responses are measured and later used to estimate a RF-model for each voxel tested. Row two: representation of the Gabor model where tunings along the dimensions of space, orientation, and spatial frequency are modelled. Row three: subjects are presented with a new set of images and fMRI responses are measured. The data from the validation set were compared against the activity patterns predicted through the training data and the Gabor model (last five rows). Adapted from Kay et al., 2008

The set of features taken into account by the receptive field model adopted by Kay and colleagues allows complex, natural images to be identified. However, a simpler version of receptive field model, focusing mainly on space and the retinotopic organisation of early visual areas is the population receptive field model proposed by Dumoulin and Wandell (2008). This technique was applied in a set of data collected during the retinotopic mapping sequence described in Section 2.3.6.2 and results are provided in Chapter 7, Section 7.1.2. Here, a two-dimensional Gaussian model of the population receptive field (Figure 2.30 – pRF model) is defined by three parameters: the centre of the distribution described by an (x_0, y_0) pair, and the spread of the distribution (σ, that is the actual receptive field size). Analogously as described by Kay et al. (2008), the Gaussian receptive field model is used to obtain a prediction of a voxel's timeseries: an estimate of the overlap between the pRF and stimulus (the dot product across space between stimulus description and receptive field model), then convolved with a haemodynamic response function (Figure 2.30 – prediction). The model is then fitted to the timeseries at each voxel by adjusting the parameters of the Gaussian model using non-linear least squares until the modelled timeseries best describes the measured data.

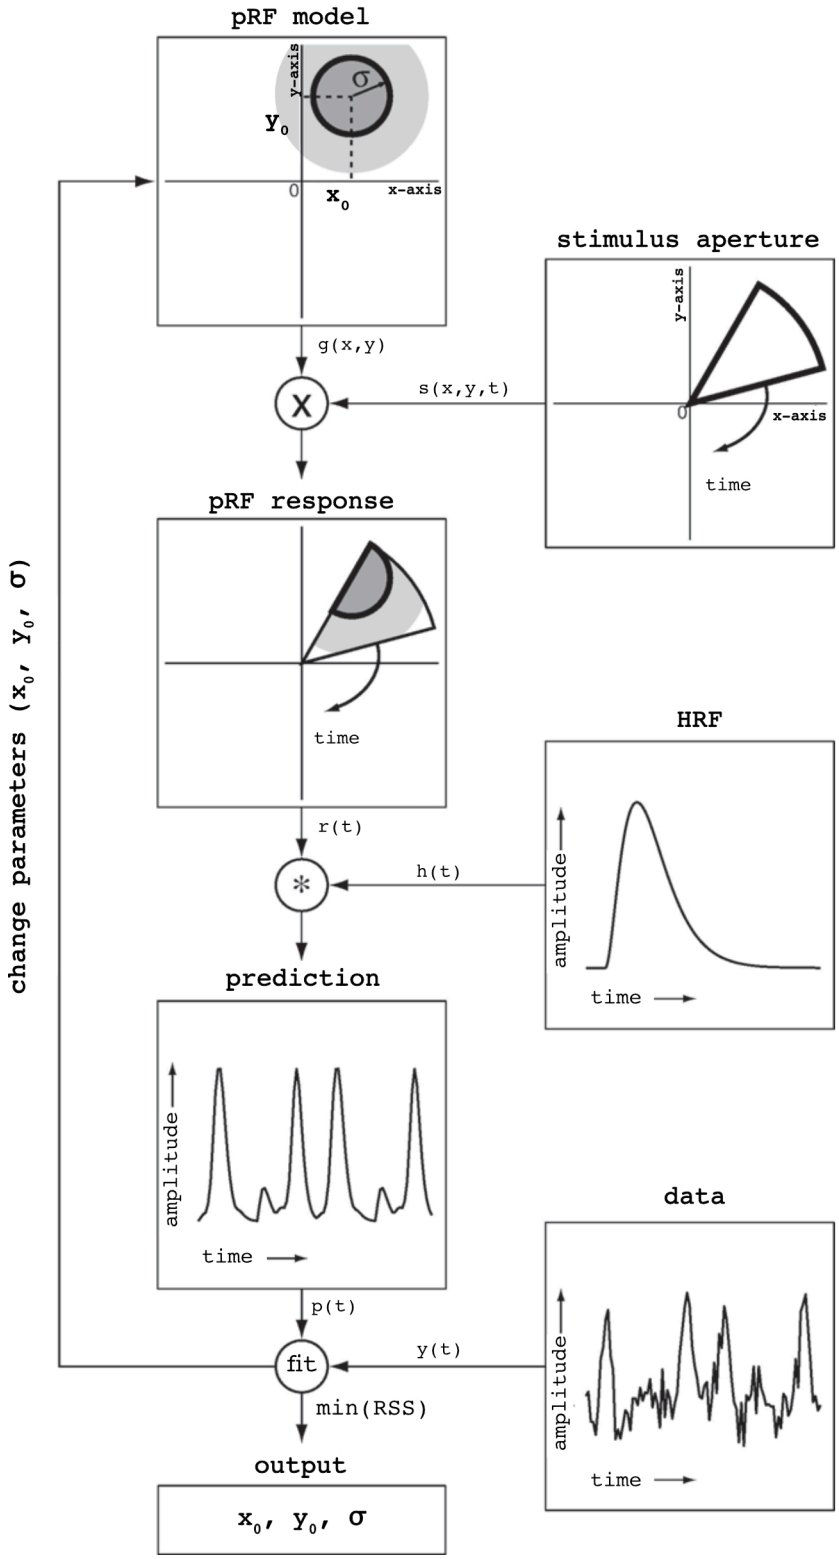


Figure 2.30 Schematic representation of the pRF model procedure. Adapted from Dumoulin & Wandell, 2008

From the fit parameters at each voxel, which describe the region of visual space that preferentially drive the fMRI response, the preferred eccentricity and polar angle, as well as additional information, such as the size of the RF, can be calculated (Dumoulin & Wandell, 2008).

2.3.6.4 Multivariate Pattern Analysis

The classical, *univariate* approach to analysing fMRI data is to apply a version of the Generalised Liner Model GLM (Section 2.3.6.1). This can be considered a *deductive* process (Gallant et al., 2010): by fixing a set of factors assumed to contribute to the observed measurement, a hypothesis is tested (i.e., the data are used to test a pre-existing theory). Over the last 10 years or so, a complementary set of approaches has become popular. These consider the responses across many different voxels simultaneously – by casting the analysis as a *multivariate* problem.

An alternative way of investigating complex systems is known as *system identification* approach (Marmarelis & Marmarelis, 1978, 2004; Wu, David, & Gallant, 2006): a framework that deals with non-linear systems (findings from neurophysiology showing that each component of the brain, from neurons (Cook et al., 2007) to broader regions of interest (Van Essen & Gallant, 1994), is a non-linear system are at the core of this approach) (Gallant et al., 2010). This is an *inductive* process: a set of different stimuli and tasks can be used allowing to formulate and test specific hypothesis once the data have been collected.

Data collection and hypothesis testing become independent and a broader selection of questions can be addressed. This inductive process can be achieved by building two different types of models: *encoding* and *decoding* models.

Encoding models (e.g., Naselaris et al., 2009; Nishimoto, Ishida, & Ohzawa, 2006) use assumptions related to features of a stimulus and the way in which these features are encoded to predict responses to arbitrary sets of stimuli. Decoding models (e.g., Friston et al., 2008; Zhang et al., 1998), on the other hand, use measured BOLD response patterns to infer the features of the stimulus / condition of the task from which it generated (Naselaris et al., 2011; Serences & Saproo, 2012). The distinction between encoding and decoding model lies in the direction in which information from different components of the models are mapped (Figure 2.31).

In general, both models present three main components: the stimulus definition (e.g., a set of images), a set of features describing the relation between the stimulus and the measured responses, and an area of the brain (one or multiple regions of interest, ROIs) where voxels are selected. Each component can be thought to identify an abstract space and that these are in turn connected via either linear or non-linear transformations (Naselaris et al., 2011).

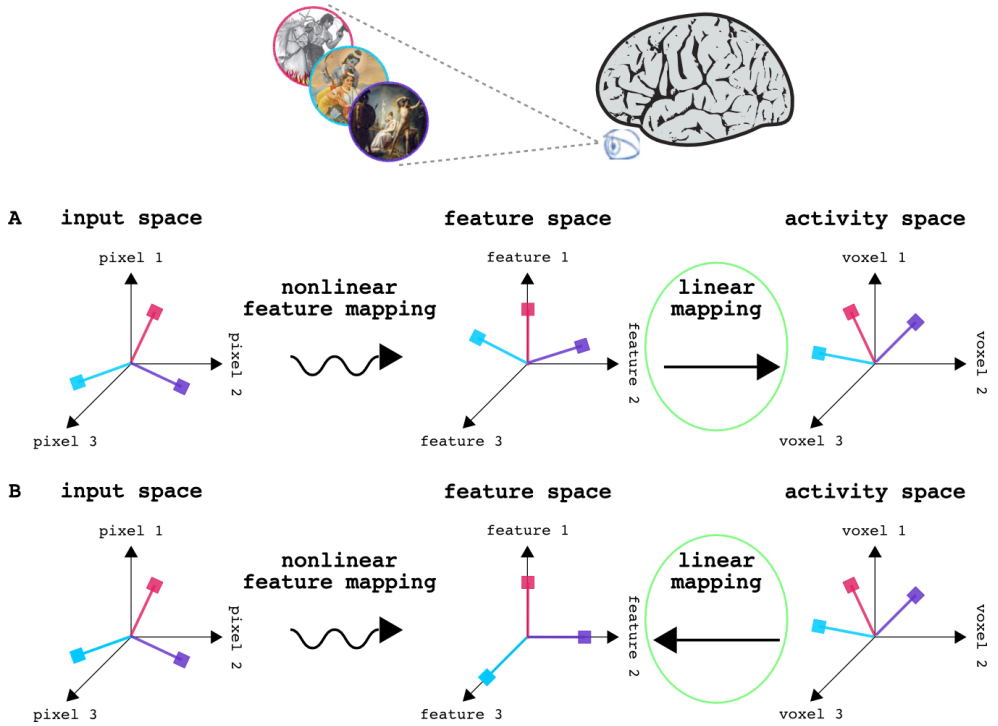


Figure 2.31 Comparison between encoding and decoding models. The first row shows a subject observing a set of images while fMRI responses are measured. A) Encoding models: the aim is to predict the set of features evoking a certain BOLD response through a linear mapping from the feature space into the activity space. B) Decoding models: the aim is to estimate the set of features that generated the measured BOLD responses through a linear mapping that goes from the activity space into the feature space. Adapted from Naselaris et al., 2011

The stimulus defines the *input space*, whose axes specify the dimensions of the stimulus (e.g., luminance values of pixels in an image). The *feature space* identifies the elements of the model that should describe the relation between stimulus and measured BOLD responses. Lastly, the *activity space* is defined by the activity recorded in the ROI (or set of ROIs) and each axis corresponds to a different voxel. While the transformation of information from the input to the feature space is non-linear (an inherent property of most processes in the brain - Wu et al., 2006, e.g. a synapse can be thought as a non-linear function transforming serial input signals into continuous current changes – Koch & Segev, 2000), the mapping between the feature and the activity space is linear (the relation between measured activity and

stimulus representation should be as simple as possible, hence the linearity assumption; Wu et al., 2006). Moreover, stimuli and features used in the model are independent of fMRI data because both are defined by the experimenter. The aim is to estimate the mapping between the activity and the feature space. The encoding and decoding models represent different approaches to this problem.

Encoding models estimate the way in which combinations of features evoke the observed BOLD responses (i.e., the mapping goes from the feature space into the activity space, Figure 2.31A – linear mapping). Decoding models, on the other hand, estimate the set of features that modulated the measured responses (i.e., the mapping goes from the activity space into the feature space, Figure 2.31B – linear mapping) (Naselaris et a., 2011).

The advantage of an encoding model over a decoding one is the possibility of obtaining a complete functional description of the ROI under investigation. This cannot be achieved using a decoding approach, because even if a feature can be perfectly decoded, this does not imply that a different feature described in a different space would not also be accurately decoded. Thus, while encoding models describe the data from the feature space, allowing to potentially fully explain the variability in the data, decoding models test individual features and to fully define the functional properties of an ROI, would require testing an infinite number of spaces (i.e., there will always potentially be a feature space that could be decoded but has yet not been tested) (Naselaris et al., 2011).

Thus, in an encoding model (Figure 2.32), for each voxel in the ROI, a weight corresponding to each feature included in the model is estimated. The weights are then applied to validate the predictions using a set of data acquired independently from the one used to derive the model parameters. If the predictions of the data, based on the feature space, are accurate, then the model well captures the mapping of stimulus onto the activity space. More plainly, the model provides a good description of the way in which stimuli are represented in the brain (Gallant et al., 2010).

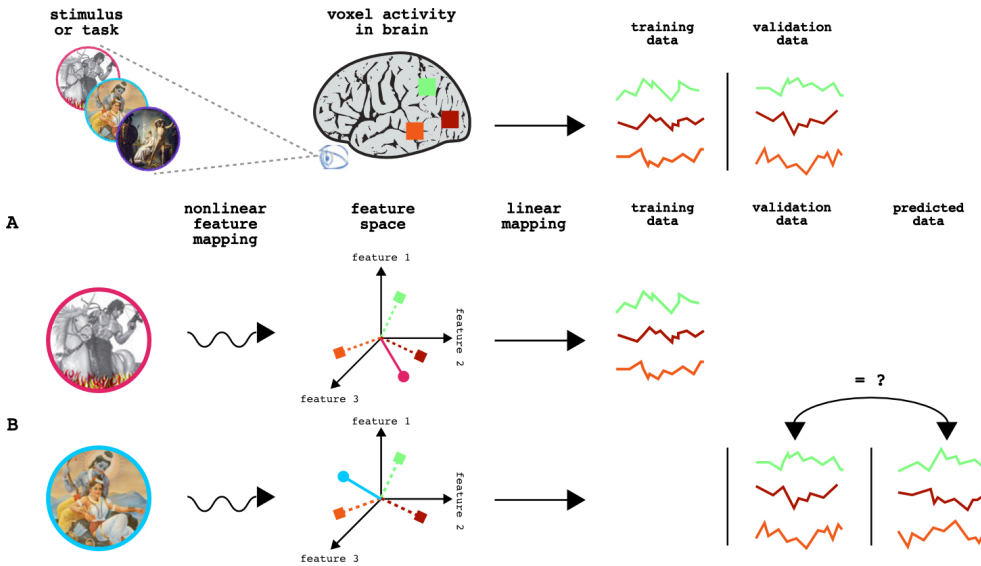


Figure 2.32 Schematic representation of the encoding approach. First row shows the recording of fMRI responses to a set of images. The activity is then divided into training and validation sets. A) Encoding models are specified using the training data. The curvy line indicates a nonlinear mapping of the stimuli into the abstract feature space. Here, dashed lines with square ends indicate estimated weights resulting in a linear mapping from the feature space into the activity space. B) Assessment of prediction accuracy: data from the validation set are compared to the activity predicted by the model in A). Adapted from Naselaris et al., 2011

In the following section I will present an example of encoding model that was then adapted and implemented in conducting the study described in Chapter 7.

2.3.6.5 Forward Encoding Model

To investigate the way in which colours are represented in the visual cortex, Brouwer & Heeger, (2009) implemented an encoding model that was later also generalised to other visual and cognitive domains (e.g., stimulus orientation – Brouwer & Heeger, 2011, motion – Kok et al., 2013, prior expectations – Kok, Failing, & de Lange, 2014, attention – Scolari, Byers, & Serences, 2012, visual working memory – Sprague et al., 2016).

By way of a brief outline, the *forward encoding model* is characterised by the following steps (Figure 2.33):

- Two separate sets of data are obtained; one is used to train the model, the other to test it
- The model consists of a description of the stimulus representation and the data in the training set are used to estimate how much each component of the stimulus description affects the stimulus representation
- The estimated weights are then used with the test data set to validate the model (i.e., estimating the activity pattern that best matches the recorded one)

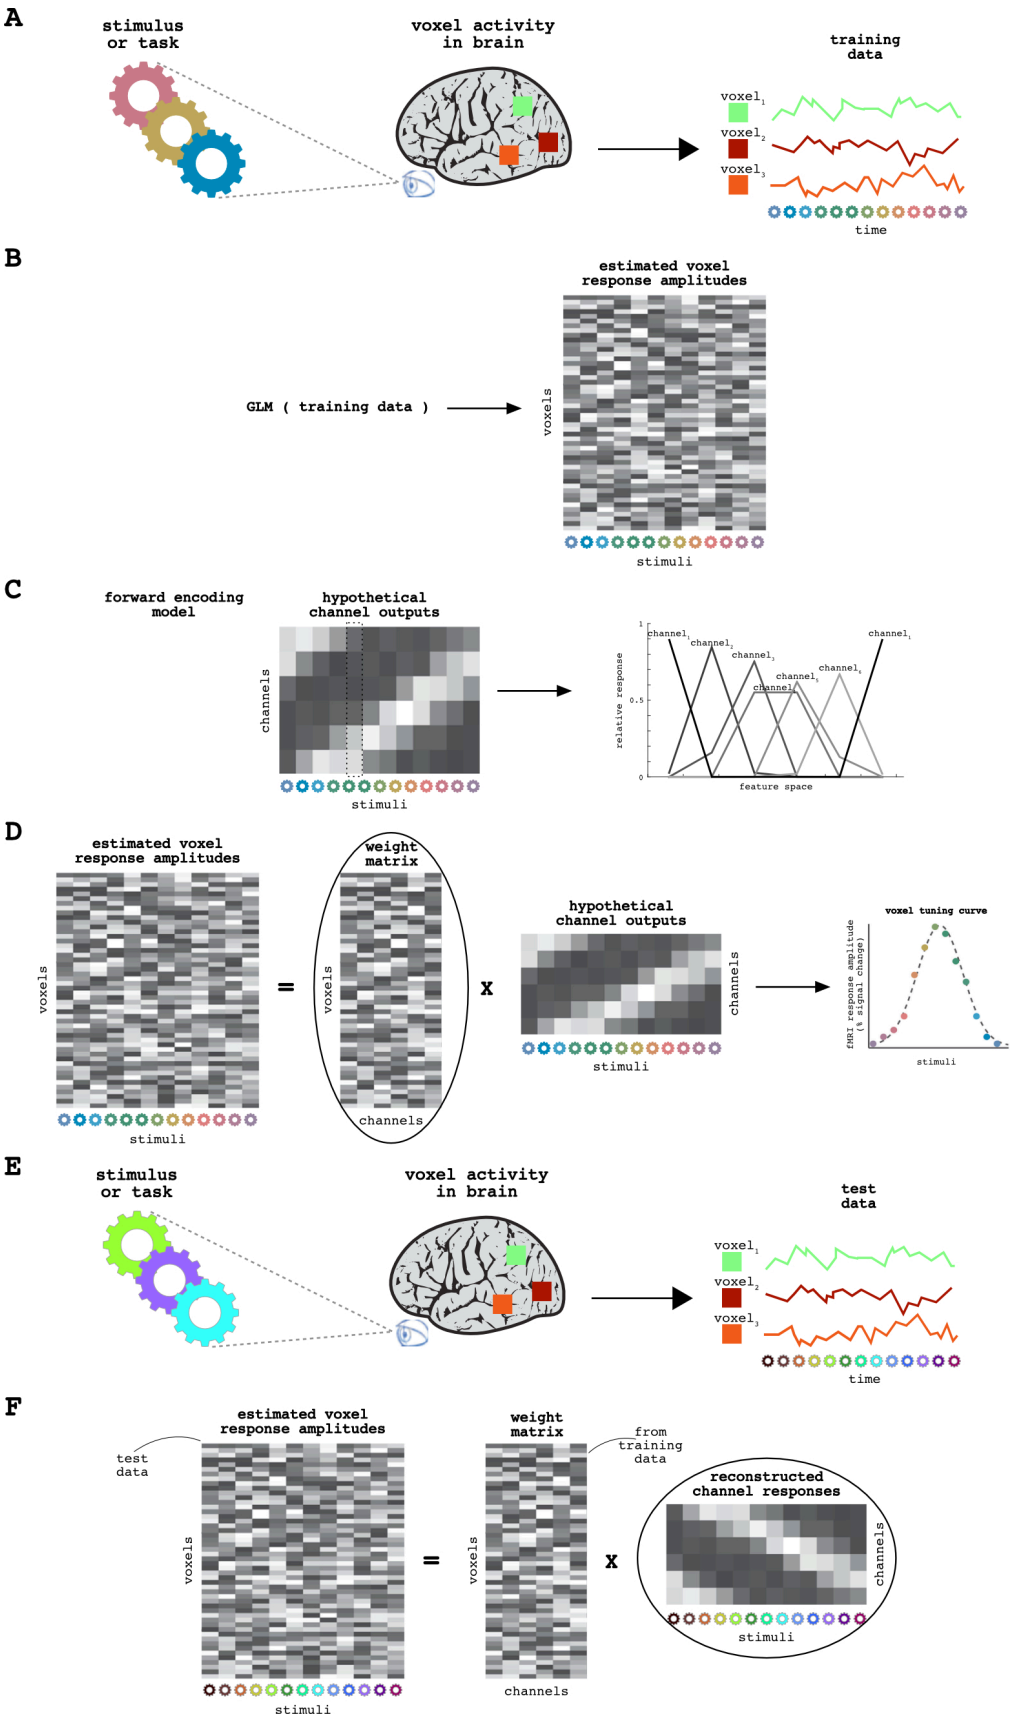

A
stimulus or task
voxel activity in brain
training data
voxel1
voxel2
voxel3
time
B
GLM (training data)
estimated voxel response amplitudes
voxels
stimuli
C
forward encoding model
hypothetical channel outputs
channels
stimuli
relative response
feature space
D
estimated voxel response amplitudes
voxels
stimuli
=
weight matrix
voxels
channels
X
hypothetical channel outputs
channels
stimuli
voxel tuning curve
fMRI response amplitude (% signal change)
stimuli
E
stimulus or task
voxel activity in brain
test data
voxel1
voxel2
voxel3
time
F
estimated voxel response amplitudes
test data
voxels
stimuli
=
weight matrix
from training data
voxels
channels
X
reconstructed channel responses
channels
stimuli

Figure 2.34 Flowchart of the forward encoding model. A) Subjects are presented with a set of stimuli and fMRI responses are recorded in different voxels of the brain. The data constitute the training set. B) Applying the GLM to the training data allows to estimate a response amplitude for each voxel, for each presented stimulus. C) A forward encoding model is built by specifying each stimulus as a combination of 6 hypothetical channels. The contribution of each channel for a stimulus is provided on the right. D) The estimated response amplitudes from the training data are used, together with the encoding model, to estimate a weight matrix. These weights allow to derive a tuning curve for each voxel indicating the contribution of the voxel to each presented stimulus. E) Subjects are presented with a second set of stimuli and the measured fMRI responses constitute the test dataset. F) The estimated response amplitudes for each stimulus in the test data are used, together with the weights derived in D) to predict the matrix of channel outputs for the new set of stimuli. Adapted from Brouwer & Heeger, 2013

As an example, consider the study conducted by Kok et al. (2013), investigating the relationship between sensory information and high-level knowledge. While systematically manipulating prior expectations related to the direction of motion of a stimulus, fMRI data were acquired and applied to a forward model. During the scanning sessions, subjects were presented with an auditory cue (two possible tones) followed by a display of a field of moving dots, a proportion of which moved coherently in one of 5 possible directions (Figure 2.34). Unknown to the subjects, the tones predicted two of the five directions with 60% probability. The task of the subjects was to estimate the perceived direction of motion (response direction) by aligning a bar segment on the screen.

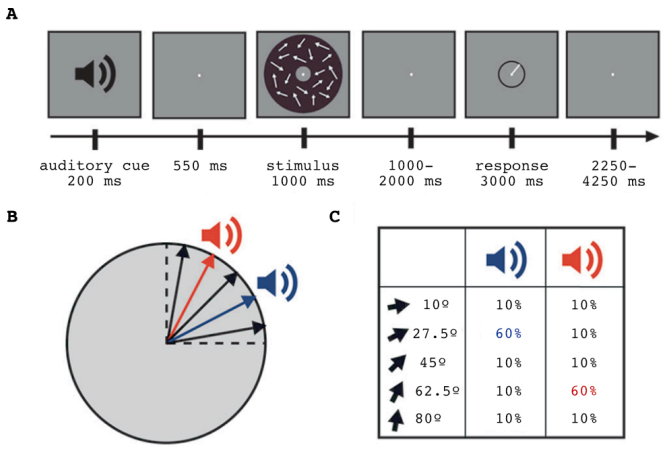


Figure 2.33 A) Task sequence: subjects are presented with an auditory cue followed by a field of moving dots. Subjects have to adjust a bar in the direction they perceived the coherent motion during stimulus presentation. B) Schematic representation of sound-direction of motion dependency. Adapted from Kok et al., 2013

The goal of the study was to investigate how sensory information and prior knowledge are integrated. Thus, a forward encoding model was used to predict the activity pattern expected during encoding of the moving stimulus. The predicted pattern was then compared to the actual stimulus direction as well as the response direction (Kok et al., 2013).

The main assumption of this specific forward model is that each voxel contains neurons that are selective for a direction of motion, thus each voxel can be defined as direction selective. A weighted sum of *hypothetical channels* or *tuning curves* can be used to model such direction selectivity: described by a set of *basis functions*, the channels are evenly spaced along the stimulus dimension (here directions from 0 to 360 degrees), and by appropriately weighting each of these N channels (usually in a number between 6 and 8; Freeman & Adelson, 1991), it is possible to reconstruct any direction of motion.

The first stage (Figure 2.35) of the encoding model consists in estimating the combination of weights applied to the channels in order to describe the response pattern of each voxel. This relationship is set as:

$$\boldsymbol{B} = \boldsymbol{W} * \boldsymbol{C} \qquad \text{(Eq. 1)}$$

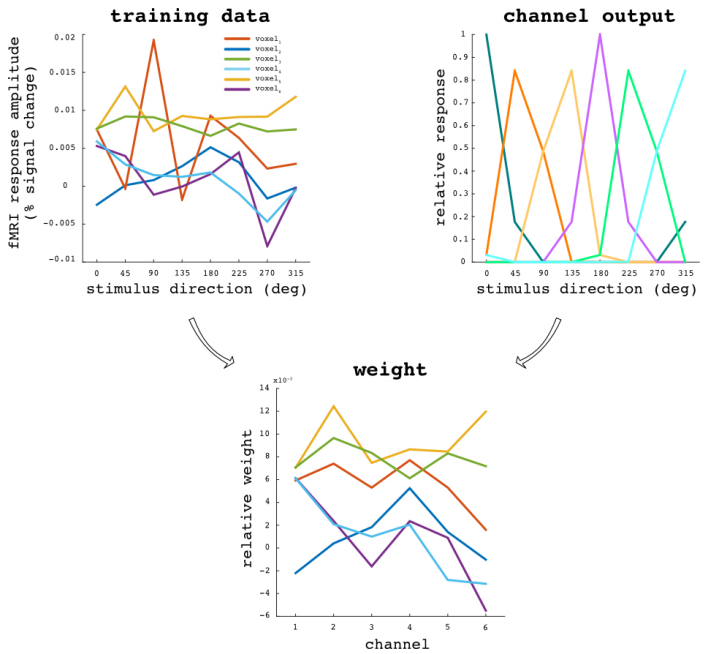


Figure 2.35 Graphical representation of weight matrix estimation stage in the forward encoding model

Where $\boldsymbol{B}$, the fMRI response amplitudes in the training set, is a matrix of size *m* x *n* (number of voxels x number of measurements); $\boldsymbol{C}$, the output of the hypothetical channels, is a matrix of size *k* x *n* (number of channels x number of measurements); and $\boldsymbol{W}$, the weights, is a matrix of size *m* x *k* (number of voxels x number of channels). Thus, to estimate how much each channel contributes to the response amplitude of a voxel, linear regression is applied:

$$\boldsymbol{W'} = \boldsymbol{B}\,\boldsymbol{C^T}(\boldsymbol{C}\,\boldsymbol{C^T})^{-1} \qquad \text{(Eq.2)}$$

In Kok et al. (2013), $\boldsymbol{B}$ represents the measured BOLD responses to a *localiser* stimulus: a scanning session used to identify visual area MT, where a fully coherent RDK stimulus was computed at 7 possible directions (the five used in the experimental runs plus one direction on each side of the range) and subjects were engaged in an attentional task around the fixation point. This highlights the fact that training data are independent from the

data used to validate the model (acquired in the *experimental* session while subjects estimated the perceived direction of motion of the stimulus).

Within this framework, it is now possible to move from the high-dimensional space determined by the response amplitude of each voxel to each trial, to a lower-dimensional space determined by the relative contribution of the six hypothetical channels to each trial. Therefore, the weights ($\boldsymbol{W'}$) estimated from the training set can now be used in combination with the voxel response amplitudes measured during the experimental session (test data set, $\boldsymbol{B_{exp}}$) to transform each measured response into a vector of six channel responses: expected channel outputs ($\boldsymbol{C'_{exp}}$). This is achieved via linear regression (Figure 2.36):

$$\boldsymbol{C'_{exp}} = (\boldsymbol{W'^{T}}\,\boldsymbol{W'})^{-1}\,\boldsymbol{W'^{T}}\,\boldsymbol{B_{exp}} \qquad \text{(Eq.3)}$$

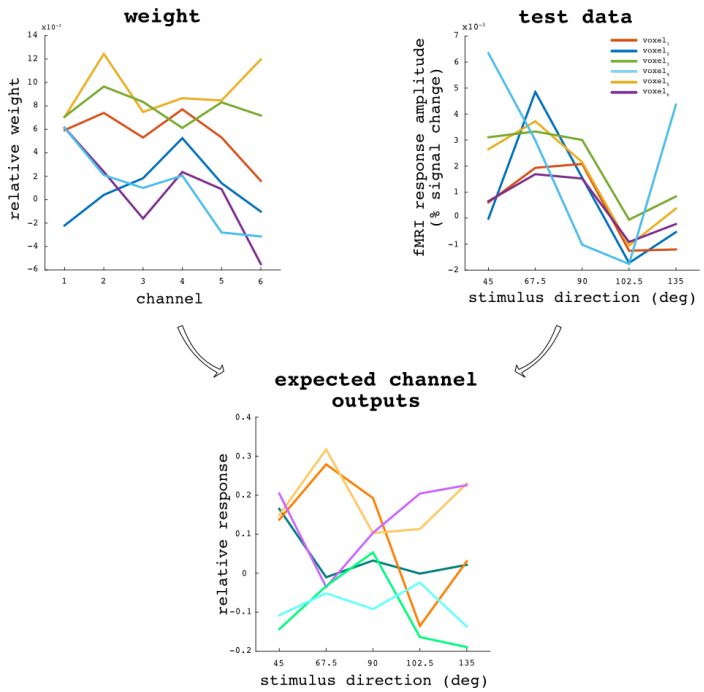


Figure 2.36 Graphical representation of the prediction of channel outputs for a set of test data, given the weights estimated from the training set.

The matrix of expected channel outputs, therefore, can be used to reconstruct (decode) the directions of the stimuli presented during the experimental session. The reconstructed direction is a weighted average of the output channels translated into the stimulus space: multiplying the channel responses by the six basis functions describing the 0-360° stimulus space, gives a response curve for each stimulus direction as a function of stimulus space (0-360° feature space). The peak of the resulting function is taken as the decoded direction of motion for a certain trial type (Figure 2.37).

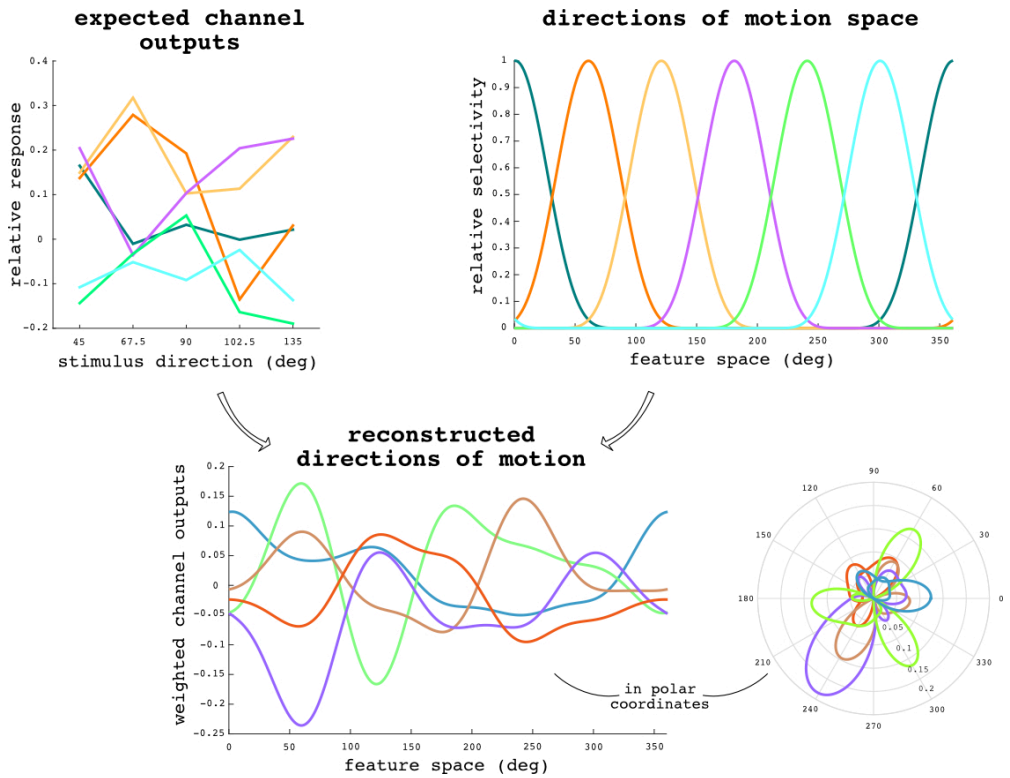


Figure 2.37 Reconstruction of the stimuli presented during the test set. The predicted channel outputs are applied to the stimulus space (0-360° direction). The peak of each of the resulting functions is taken as the decoded stimulus direction.

Kok and colleagues not only reconstructed the directions of motion presented in each trial of the experimental session of their study, but also addressed the impact of prior knowledge (cue) onto the sensory

representation of the moving stimulus. This was achieved by comparing the directions of motion reconstructed from early visual cortex (directions decoded from the forward model) with the ten trial types in the experiment (two auditory cues x five stimulus directions) and the responses provided by the subjects when estimating the perceived stimulus directions, respectively. Behavioural data (estimates of motion direction) showed an attraction effect towards the cued stimulus direction: when the auditory cue predicted a 27.5° motion, estimates were biased towards that direction and the same was true when the cue predicted a 62.5° motion. Interestingly, the directions decoded from early visual areas correlated more with the estimates given by the subjects, suggesting that top-down expectations affect sensory representation in early visual cortex.

Given the ways in which the model can be employed and the breadth of questions it allows to address, it was used in this thesis to investigate whether the sensory representation of directions of motion in areas of the visual system differs depending on the type of task a subject is performing. In particular, we asked participants to discriminate either between two highly similar (*fine* discrimination) or two dissimilar (*coarse* discrimination) visual stimuli. Data acquisition parameters and experimental designs are described in the following sections (2.3.7, 2.3.7.1, 2.3.7.2), data analysis and results are presented in Chapter 7 (Section 7.2).

2.3.7 *Functional MRI acquisition*

MR imaging data for retinotopy mapping were acquired on a Philips 3T scanner, using a 32-channel SENSE head coil and a 2D Gradient Echo EPI Sequence (TE=35 ms, TR=1500 ms, FA=75 degrees, dynamic scan time=1.5 s), with a voxel size of 3 mm isotropic and 24 slices through the brain. In one scan session, data were collected from 10 functional scans (four used wedge stimuli, two used ring stimuli, and the remaining four used bar stimuli) in order to identify visual areas: V1-V3, and MT (either using standard retinotopic mapping analysis or pRF analysis). During the same scanning session, high-resolution anatomical MPRAGE images of the whole brain were acquired and later used for registration, segmentation, and flattening. For localiser and experimental scans, data were acquired using the same scanner with a 2D Gradient Echo EPI Sequence (TE=35 ms, TR=2000 ms, FA=75 degrees), voxel size 3 mm isotropic and 34 slices through the brain. For each session, two scans were used for the localiser and four for the experimental runs (two coarse and two fine discrimination tasks).

2.3.7.1 *Localiser*

Participants underwent two localiser scans during one imaging session. The aim of these was twofold: first, they allowed to localise the extrastriate visual area V5; second, the data collected were used to train a forward encoding model.

Visual stimuli were generated using the MGL toolbox (available from http://gru.stanford.edu/mgl) in MATLAB (The MathWorks, Natick, MA) on

Apple Mac OS X, presented onto a screen located at the back of the magnet (BOLDScreen, Cambridge Research Systems), and seen by the subjects through a set of mirrors. Stimuli consisted of two circular apertures (6° in diameter) vertically aligned with the fixation point (displaced by 2° above the centre of the screen to ensure stimuli were seen fully through the mirrors) and with a horizontal offset of 2.3° from such point (see also Scolari et al., 2012). Within each aperture, RDKs were computed with dots moving at a speed of 4° s^{-1} and an average density of 150 dots deg^{-2} s^{-1}. Dots moved coherently for 2 s (1 TR) in one of 8 possible directions (across 360 degrees in steps of 45 degrees) randomly selected on each trial. A two-interval forced-choice contrast discrimination task at fixation was used to ensure subjects were paying attention during the scan sequence. Task difficulty was controlled through a staircase procedure and contrast changes and response timings were unrelated from that of the localiser stimuli (Huk, Ress, & Heeger, 2001).

2.3.7.2 Fine vs Coarse discrimination

Visual stimuli presented the same features as described in 2.3.7.1 with the following exceptions: RDKs were computed at different coherence levels (30% for fine discrimination task, and 10%, 15%, 20%, 25%, 30% selected randomly at the beginning of each trial for the coarse discrimination task) and directions of motion for the reference RDK were randomly selected from a set ranging between 45 and 135 degrees in steps of 22.5 degrees. The direction of motion of signal dots in the comparison RDK was randomly

selected to be within ±20 degrees (in steps of 10 degrees) of that of the reference for fine discrimination task blocks, whereas for half of the trials in the coarse discrimination task blocks, signal dots in the comparison RDK moved 180 degrees away from the motion direction of the reference aperture.

Subjects completed two discrimination tasks for two fMRI scans each: two blocks of a fine discrimination task on visual motion followed two blocks of coarse discrimination. During the fine discrimination task, subjects judged whether the comparison RDK moved CW or CCW with respect to the reference aperture. During the coarse discrimination task, they had to judge whether the coherent motion was in the same or opposite direction of that of the reference RDK.

CHAPTER 3[1]

3 Estimates of Visual Motion Direction

Visual perception is strongly influenced by contextual information. Visual illusions such as the *Mach bands* illusion (Mach, 1865) and the *Kanizsa square* illusion (Gregory, 1972) have been explained in terms of limitations of the sensory apparatus: a physiological correlate is at the core of these types of *misperceptions* (or perceptual errors). On the other hand, *illusions* can also derive from higher-level cognitive biases: prior knowledge about the state of the world have an impact on the way we perceive it, a typical example being the *Hollow Face* effect (Gregory, Gombrich, & Blakemore, 1973). On a similar level, it is also known that *responses* regarding perceptions can be distorted by prior expectations (Chalk, Seitz, & Seriès, 2010), reward contingencies, and the history of preceding responses (Fischer & Whitney, 2014). These examples therefore suggest that biases in visual perception can arise from different levels of information processing. A key effect for the discussion in this thesis is *reference repulsion* (Rauber & Treue, 1998), a systematic bias away from a reference when subjects estimate the direction of motion of a stimulus. This repulsion effect,

[1] Some of the work presented in this chapter has been published in the following article:
Zamboni E, Ledgeway T, McGraw PV, Schluppeck D (2016). Do perceptual biases emerge early or late in visual processing? Decision-biases in motion perception. *Proc. R. Soc. B* **283:** 2016.0263 http://dx.doi.org/10.1098/rspb.2016.0263

observed between static stimuli (e.g., two oriented lines or gratings; Regan & Beverley, 1985), a static and a moving stimulus, and two transparently moving stimuli (Hiris & Blake, 1996), has been described in terms of pure *sensory phenomenon* (e.g., Kim & Wilson, 1996; Marshak & Sekuler, 1979; Mather & Moulden, 1980). However, more recent studies have suggested that these biases could result from higher-level cognitive effects (Rauber & Treue, 1998, 1999).

Regan & Beverley (1985) investigated the processes involved in detecting and discriminating orientations of stimuli. While detection is determined by the responses of selectively tuned elements of the sensory apparatus, discrimination performance is linked to the relative response of multiple elements. In line with these findings, Hol & Treue (2001) reported that performance in the detection of motion direction is reduced when subjects undergo adaptation along a targeted direction of motion. On the other hand, discrimination of highly similar stimuli is impaired only when adaptation is applied along directions that differ by up to 50° from the ones tested.

These studies therefore suggest that different populations of neurons selectively define detection and discrimination performance: detection requires neurons tuned to the specific stimulus feature, discrimination of highly similar stimuli relies on neurons tuned away from the critical stimulus feature.

This underlying distinction between detection and discrimination of stimulus features has been suggested to be at the core of an effect known

as *reference repulsion*. Rauber & Treue (1998) illustrated this phenomenon by presenting subjects with a field of moving dots and a reference line. Participants were required to judge whether the direction of the moving stimulus was more *clockwise* (CW) or *anti-clockwise* (CCW) than the orientated line. The authors reported a systematic error in the subjects' responses: an *overestimation* of the angular difference between the moving stimulus and the static, oriented reference line. Moreover, they showed how this *bias* varies with the magnitude of angular difference: the smaller the distance between the oriented reference line and the moving pattern, the bigger the overestimation of such direction of motion. Along these lines, evidence of *motion repulsion* was reported by Hiris & Blake (1996) using two populations of moving dots within the same stimulus aperture. The authors identified this effect as a *sensory phenomenon*: inhibitory interactions between directionally selective neurons encoding the sensory stimulus cause a biased representation of the stimulus itself (see also Markshak & Sekuler, 1979; Kim & Wilson, 1996; Mather & Moulden, 1980). In replicating these findings, Rauber & Treue (1999) however suggest the possibility of it being a *cognitive* effect: subjects have a veridical representation of the sensory stimulus (i.e., the information is veridical at the encoding stage), but the observed bias arises from a set of cognitive models that are subsequently applied to such representation (as an example, see Huttenlocher, Hedges, & Duncan, 1991).

A more recent study by Jazayeri & Movshon (2007a) addressed this perceptual bias and suggested an alternative explanation, linking the

different strategies used by neural populations to either discriminate between highly similar stimuli or detect a signal. A detailed description of the paradigm is given in Chapter 1, Section 1.3.

The aim of this chapter is to examine the point in time at which this perceptual bias arises: early, during sensory encoding, or late in the task sequence, reflecting *decision-related* processes. This was achieved by applying small manipulations to the context in which subjects had to estimate the perceived direction of motion. Experiment 1 investigated whether the *reference repulsion* effect could be replicated in our observers. A replica of the fine-discrimination task proposed by Jazayeri & Movshon (2007) is implemented. In Experiments 2 and 3 the properties of the oriented reference line are manipulated during the estimation phase. These modifications later in the trial allow to investigate the impact of the discrimination boundary on the well-known biases observed in fine-discrimination tasks.

3.1 Experiment 1: Biased Estimates of Motion Direction following a Fine Discrimination Task

3.1.1 Materials and Methods

Participants performed a minimum of 2500 experimental trials (2870 trials on average), following an average of 930 practice trials.

3.1.1.1 Observers

Five subjects participated in this study (five females; age 22.2 ± 2.18 years, mean ± s.d.). Four observers were naïve to the purpose of the experiment and one was an experienced psychophysical observer. All had normal or corrected-to-normal acuity.

3.1.1.2 Apparatus and Stimuli

A detailed description of Apparatus and Stimuli is given in Section 2.1.1.2 and Section 2.1.1.3 of Chapter 2.

3.1.1.3 Procedure

A schematic representation of the experimental procedure is presented in Figure 3.1. Briefly, observers were presented with a white fixation dot at the centre of the screen for 1 s, followed by the onset of the oriented reference line. After 0.5 s, the RDK, with a randomly selected coherence level and a direction of motion also randomly selected between ±24.5° from the orientation of the reference line, was displayed for 1 s. When the RDK disappeared, the fixation dot turned orange, cueing participants to perform the fine-discrimination task: press one of two keys on a standard keyboard, indicating whether the coherent motion was more CW or CCW with respect to the oriented reference line. When subjects did not receive a feedback on this task, a black circular ring prompted them to perform the direction estimation task: extend a line 'tethered' to the fixation point in the direction they perceived the coherent motion. Once the line hit the ring, the task was

completed, the ring, line, and oriented reference were removed, and the next trial could start.

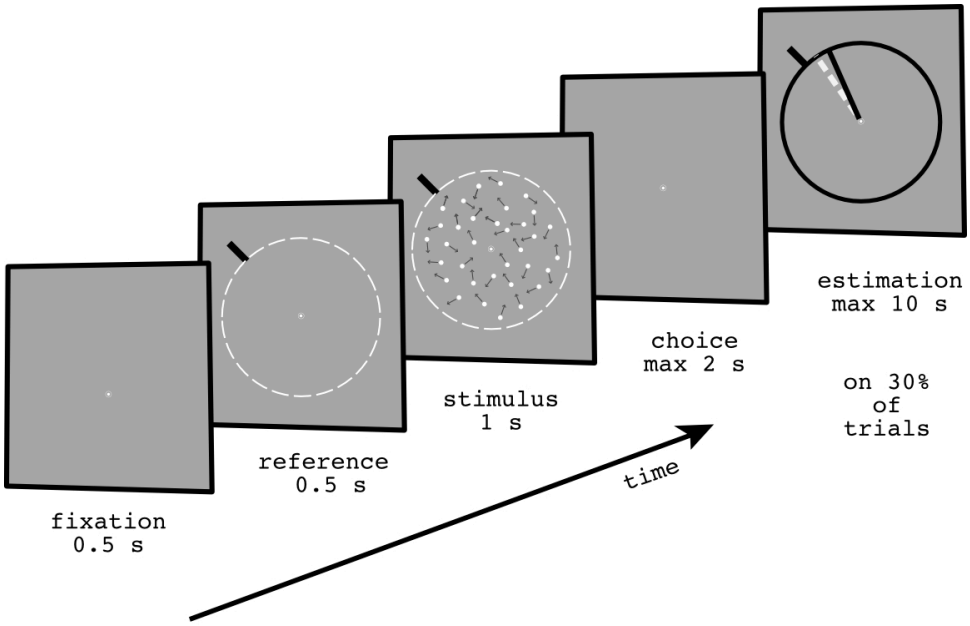


Figure 3.1 Schematic representation of a trial sequence. The oriented reference line is present throughout the whole duration of the trial. Adapted from Zamboni et al., 2016

3.1.2 Results

All subjects were able to perform the discrimination task at all coherence levels tested, and the precision of the discrimination performance was lawfully related to the amount of coherent signal present in the stimulus. The slope of the psychometric functions, obtained by fitting data with a cumulative Gaussian distribution function, become steeper as the coherence of the RDK increases (see Figure 3.2). Although statistical test failed to reach significance ($F(3) = 2.52$, $p = 0.09$), the trend described can be appreciated in Figure 10.3 in the Appendix. In addition, discrimination was more precise as the coherent signal dots moved farther away from the oriented reference line, than when the motion was closer to the boundary.

Moreover, we found no systematic direction bias (i.e., shifts in the Point of Subjective Equality, PSE) for the discrimination data ($F(3) = 0.574$, $p = 0.64$).

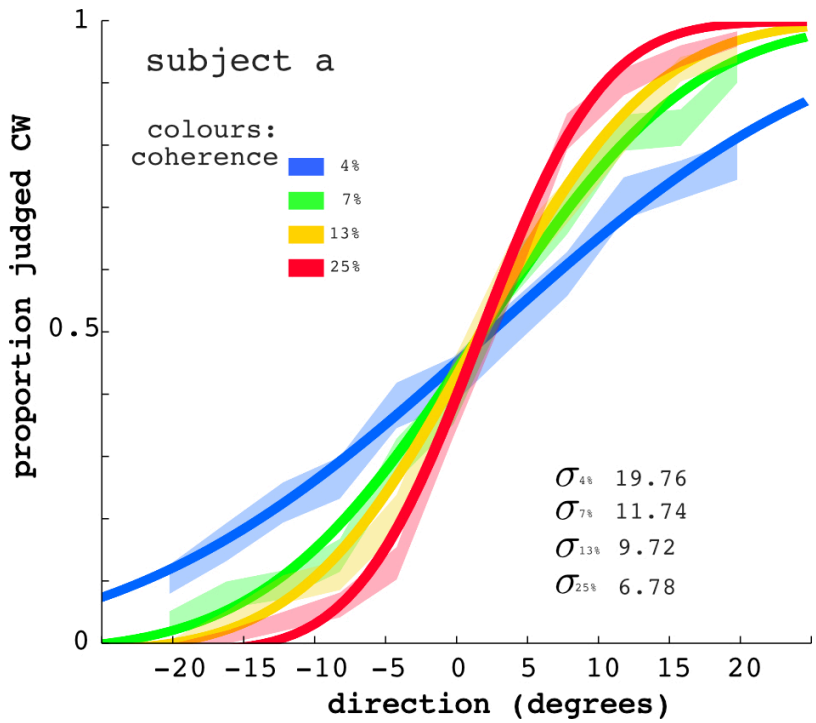


Figure 3.2 Proportion of CW judgements plotted as a function of stimulus direction. Data for one subject tested with stimuli computed at four coherence levels (different colours).

Data from the estimation task reveal an interesting pattern. Assuming a veridical subject, i.e. one that reports precisely the direction of motion presented in the task, one would expect the distribution of estimates for a given direction to be normally distributed around the true direction of motion. However, and in line with the findings presented by Jazayeri & Movshon (2007a), this is not what we found. Estimated directions of motion were characterised by a bimodal distribution, indicating a repulsive effect consistent with the discrimination judgement: when the observer's choice in the discrimination task was CW, the estimated direction of motion was CW-shifted with respect to the oriented reference line. This effect was particularly marked for trials in which the signal dots moved close to the

reference line and less so when they moved farther away from it (see Figure 3.3).

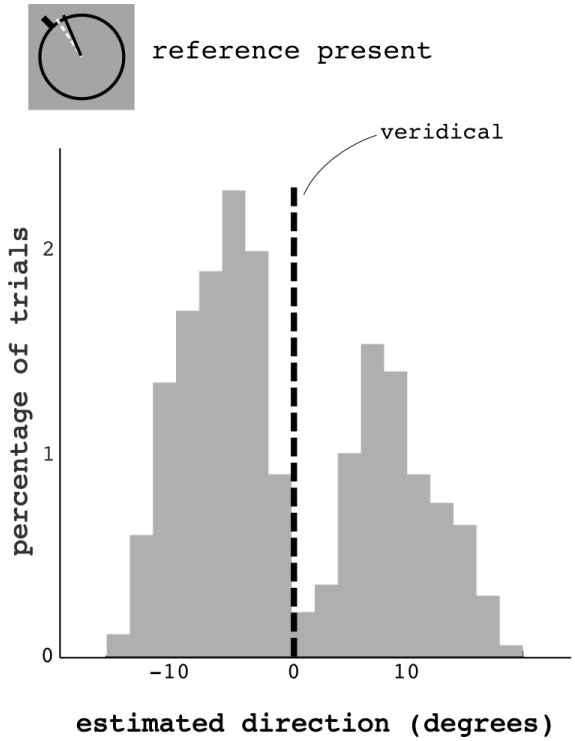


Figure 3.3 Distribution of estimated directions of motion for trials in which the direction of the stimulus was the same as the orientation of the reference line. Data presented for one subject, one coherence level.

Figure 3.4 shows data from one subject across the four coherence levels tested. Each panel of row 1 represents a joint histogram of percentage of direction estimates as a function of stimulus direction. Local maxima for the distribution of estimates for each stimulus direction are represented as red symbols. In other words, the circles represent the peaks of the histograms displayed in row 2 of Figure 3.4, that is obtained by treating the joint histograms in a column-wise manner.

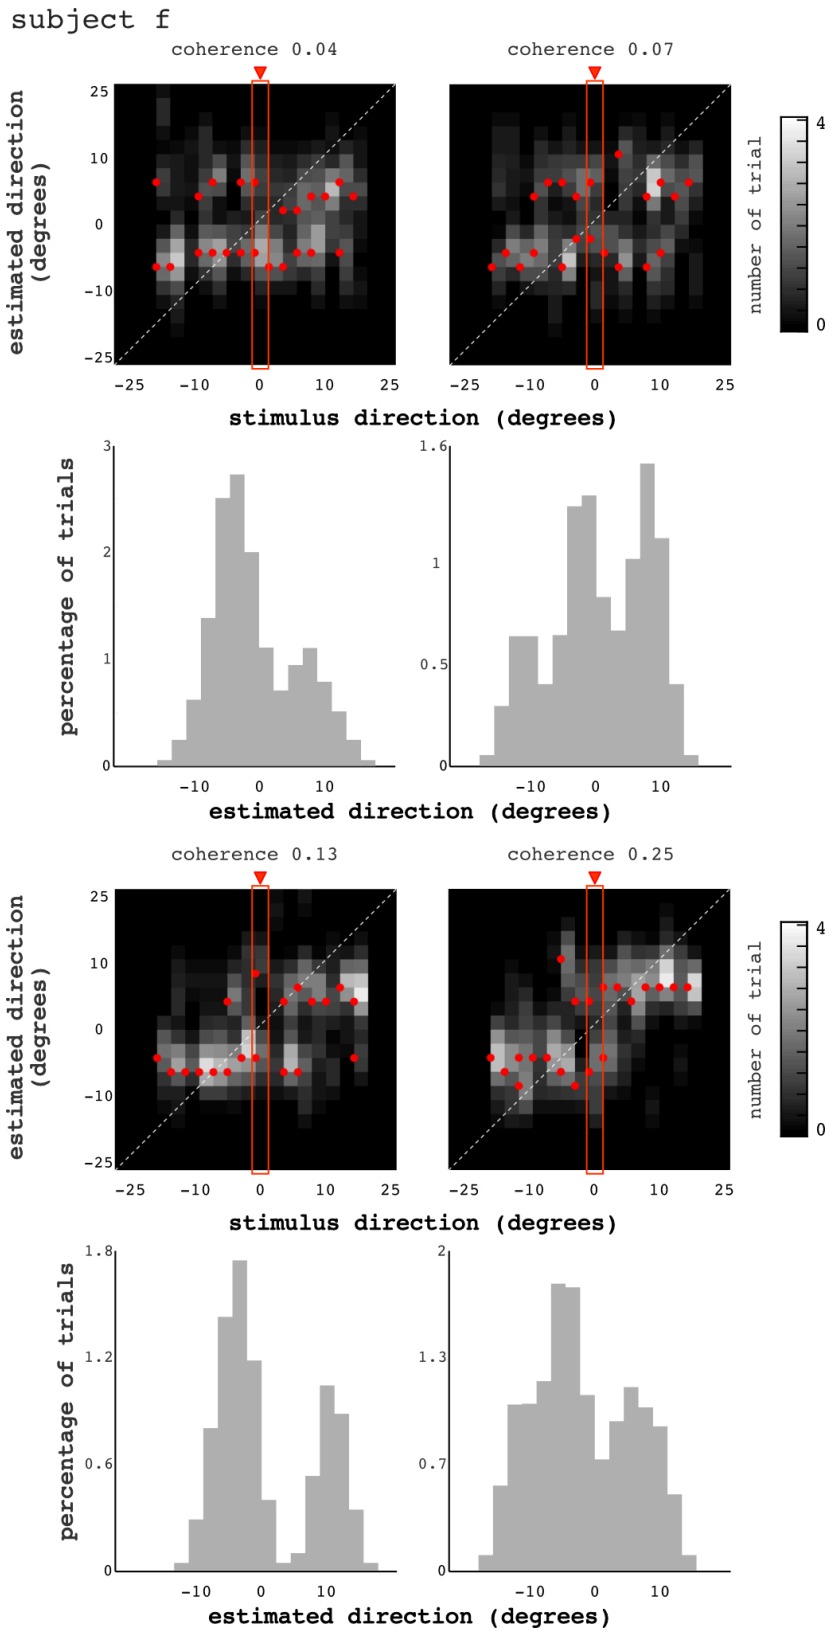


Figure 3.4 Data from one subject, four coherence levels. Plots in row one and three represent density distributions of estimated directions of motion as a function of stimulus direction. The red dots identify the peaks of column-wise distributions for each stimulus direction presented. The second and fourth rows present the histograms building the column-wise distribution highlighted by the red selection in each density plot of row one and three, respectively.

The results presented here and obtained adopting the same experimental procedure used in Jazayeri & Movshon (2007a) strongly support the previous findings. However, the biases in the reported motion directions could have occurred due to effects early in the trial, around stimulus presentation, or later in the trial, around the time of the behavioural response. The oriented reference line was present throughout the task sequence, and it is therefore not possible to determine at what point during the trial these biases arose. This aspect is addressed in Experiment 2.

3.2 *Experiment 2: Does the Bias Arise Early or Late during Information Processing?*

Because the reference repulsion effect could arise at the time of stimulus presentation (during initial encoding by the sensory apparatus), during subsequent decoding of that information for processing by higher-level visual areas, or during the decision process and formation of subject's response, a manipulation was applied to the oriented reference line to disambiguate these possibilities. Specifically, the discrimination boundary, following a fine-discrimination task, was *removed* completely during the direction estimation task. If the bias is due to either processes at the time of stimulus encoding or task-specific strategies adopted to decode sensory information, manipulations later in the task sequence, i.e. at the time of motion direction estimation, should have no effect on the biases.

3.2.1 Materials and Methods

The methods used here are essentially the same as those adopted and described in Experiment 1. Deviations from that protocol are highlighted in the corresponding sections.

Participants performed a minimum of 2500 trials (2990 trials on average).

3.2.1.1 Observers

The same group of subjects that took part in Experiment 1 participated to this study.

3.2.1.2 Apparatus and Stimuli

Stimuli were defined and presented as described in Experiment 1.

3.2.1.3 Procedure

Overall procedure is the same as described in Experiment 1, except for the presentation of the oriented reference line. The discrimination boundary was displayed only for the duration of the discrimination task and was removed as subjects engaged in the estimation task (see Figure 3.5).

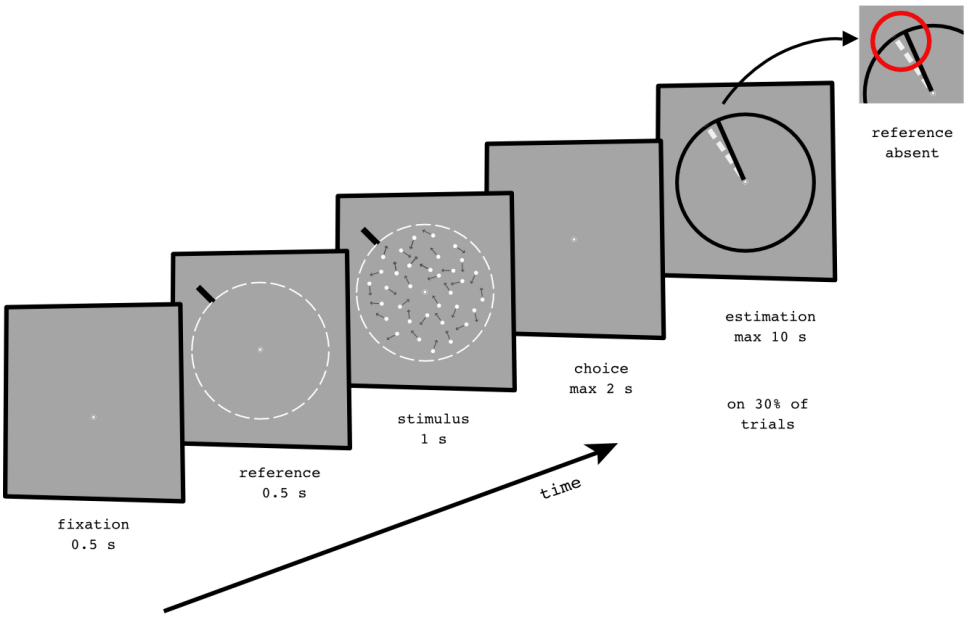


Figure 3.5 Schematic representation of a trial sequence: the oriented reference line is absent during the estimation task (highlighted in the magnification quadrant on the top right corner). Adapted from Zamboni et al., 2016

3.2.2 *Results*

The discrimination task was the same as in Experiment 1 and performance was consistent with previous results (see Figure 3.6).

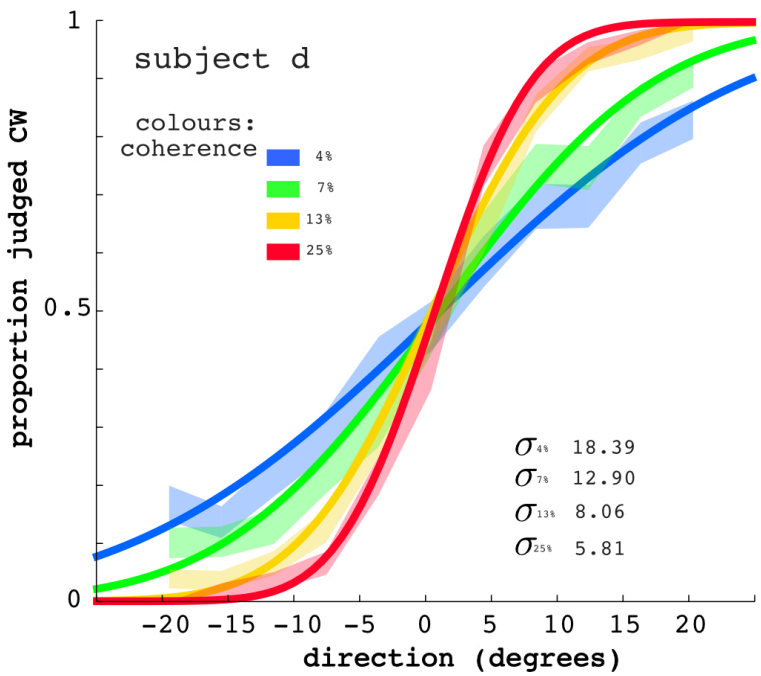


Figure 3.6 Proportion of CW judgements plotted as a function of stimulus direction. Data for one subject, four coherence levels testes (different colours).

When looking at the data from the motion direction estimation task, the pattern of responses was markedly different from that observed when the oriented reference line was present. Here, distribution of estimates differed from the bimodal one observed in Experiment 1: estimated directions of motion were centred on the stimulus direction. This is shown in Figure 3.7: each panel in row 1 and 3 represents joint histograms of estimated directions of motion for the four coherence levels tested. As for the results in Experiment 1, local maxima for each distribution are marked with red symbols. Moreover, regression lines through these points are displayed as red solid lines.

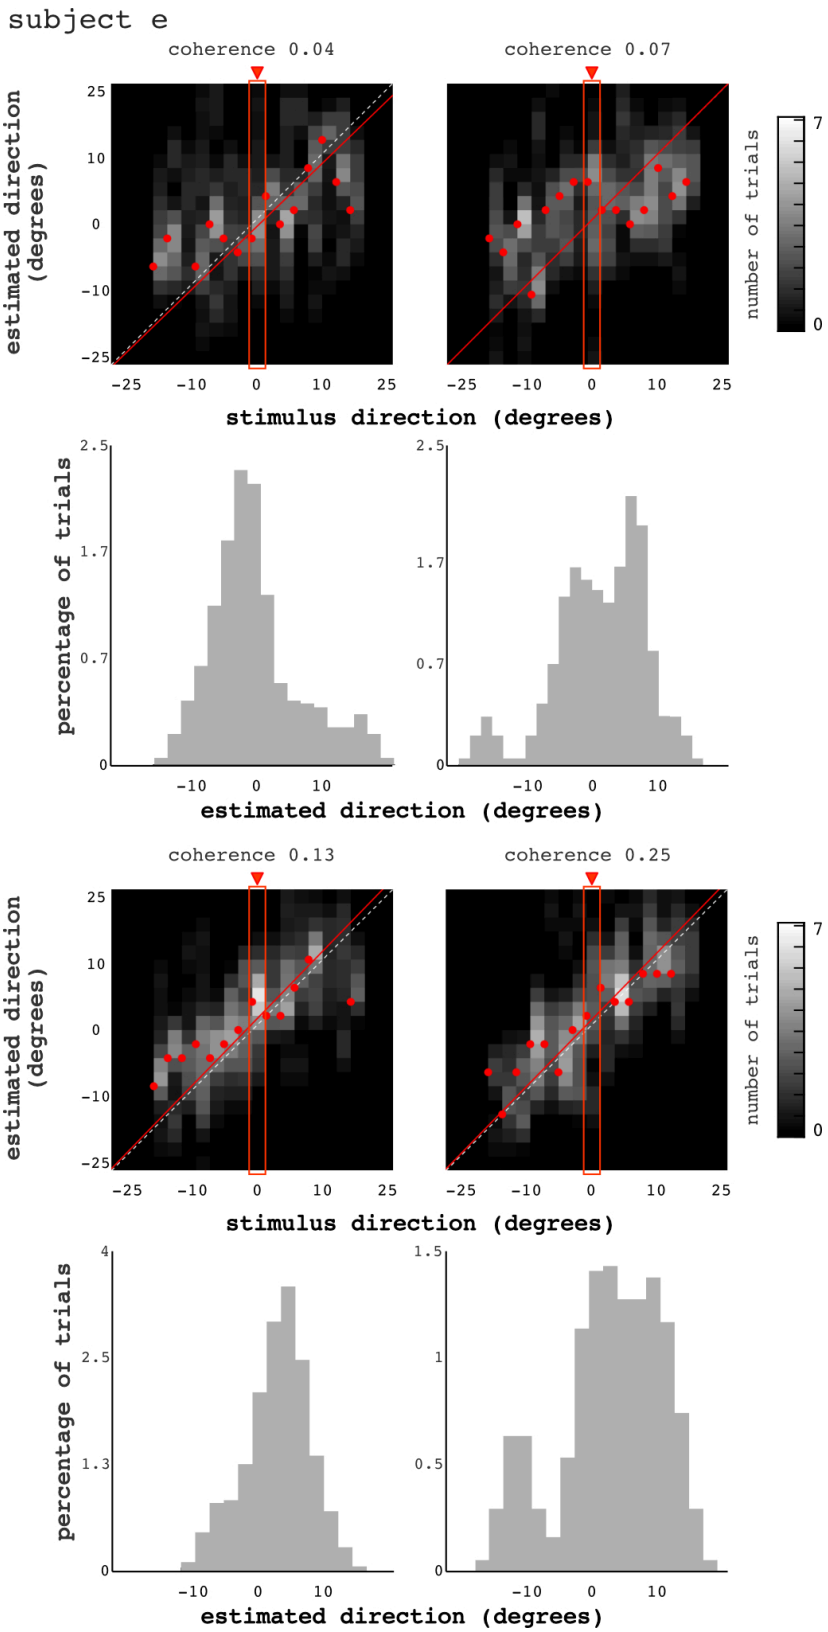


Figure 3.7 Data from the estimation task for one subject, four coherence levels. Density plots in row one and three display the distribution of estimated directions of motion as a function of stimulus direction. Red dots indicate peaks of column-wise distributions for each direction presented. Histograms in the second and fourth rows represent the column-wise distribution highlighted by the red rectangle in the density plots of row one and three, respectively.

From the distribution of the red circles (peaks of distributions in Figure 3.7), it is possible to see that responses were, on average, unbiased: across all conditions they are close to the identity line (dashed white line). It is also apparent in the data that the distributions of estimates show less evidence of bimodality. This pattern in the direction estimates can be appreciated by observing the histograms in row 2, for a specific direction of motion. The histograms here show estimates for trials in which the stimulus direction was 0°, i.e. when the signal dots moved in the same direction as the orientation of the reference line.

To quantify the discrepancy between biased responses in Experiment 1 and unbiased estimates in Experiment 2 we used the *dip*-statistic (Hartigan & Hartigan, 1985), a descriptive statistic quantifying deviations from unimodality. A representation of this analysis is shown in Figure 3.8 for one example subject in the dataset. The scatter plots show estimated directions as a function of stimulus direction with respect to the reference; the four panels show data for different stimulus coherence levels. Data for the 'reference present' condition (Experiment 1) are shown in red and those for the 'reference absent' condition (Experiment 2) are shown in grey. The repulsive bias from the reference at stimulus direction around 0° can clearly be appreciated for the 'reference present' data. On the other hand, for the data shown in grey the bias is not apparent. Using the marginal histograms, Hartigan's *dip*-statistic was computed (see insertion in each panel of Figure 3.8). A large dip-statistic provides evidence for deviations from a unimodal

distribution: this is the case for data from the 'reference present' condition, but not for those from the 'reference absent' condition.

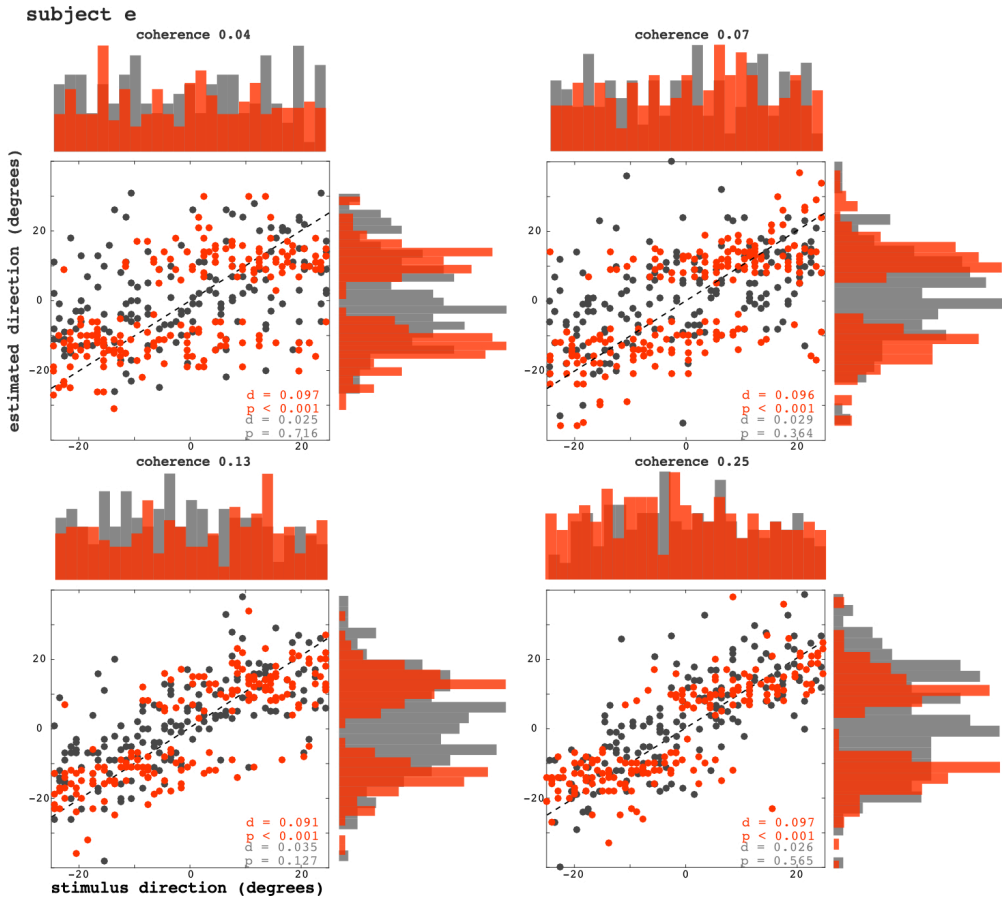


Figure 3.8 Data from one subject, four coherence levels (different panels), two conditions (different colours). Estimated directions of motion are plotted as a function of stimulus direction for the reference present (red) and reference absent (grey) conditions. Marginal histograms at the sides of the scatter plots indicate the marginal distributions of the estimates (along the y-axis) and the stimulus directions (along the x-axis). Dip-statistic values (d) and corresponding p-values (p) are provided in each panel (colour of text relates to the condition tested).

Table 3.1 in Appendix provides computed *dip*-statistic values for each subject, for each coherence level tested, and the two conditions (i.e., 'reference present' and 'reference absent').

Finally, Figure 3.9 represents a summary scatter plot with dip-statistic values for the 'reference absent' (abscissa) and 'reference present' (ordinate) data, for all subjects (N=5), all conditions (4 coherence levels, 2 reference states). The *dip*-statistic values for the whole dataset show that the same pattern is evident across all subjects. *Dip*-statistic values cluster in

the top-left quadrant of the figure: low values for data from the 'reference absent' condition and high values for the 'reference present' condition.

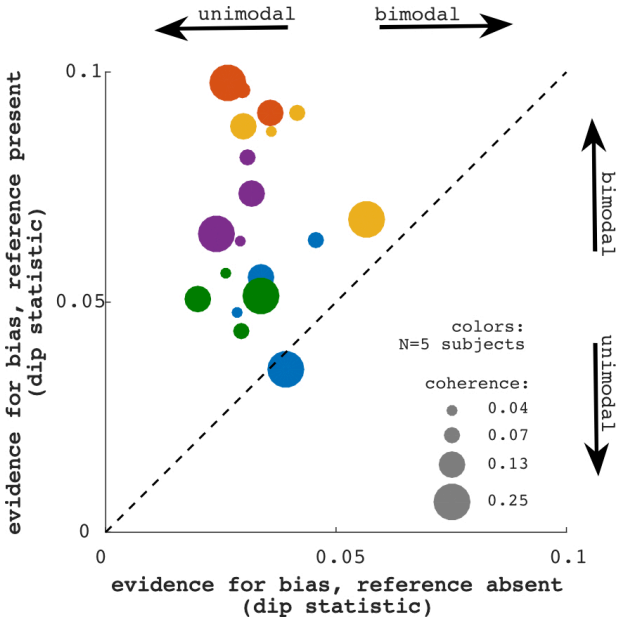


Figure 3.9 Summary plot of dip-statistic values for each subject (different colours), for each coherence tested (different symbol sizes). D-values for reference present condition are plotted against d-values for reference absent condition. Source Zamboni et al., 2016

The data across all subjects and conditions (with one exception) show that the distribution of direction estimates were unimodal when the reference was absent, bimodal when the reference was present (all individually, statistically significant at $p < 0.05$ – see Table 3.1 in Appendix). The results presented here strongly challenge the findings reported in Section 3.1.2. The reasoning for this is as follows: given that subjects are asked to estimate the direction in which they perceived the coherent motion early on during stimulus presentation, removing the oriented reference during the estimation task was not expected to affect the perceptual bias. However, the repulsive effect observed when the reference was present, was not apparent when the reference was removed at the end of the trial. Therefore, the presence of the oriented reference line

during the estimation task plays an important role for setting up the bias. Moreover, the fact that changes relatively late in the task sequence affect the repulsive effect also suggest that we are looking at a phenomenon manifesting itself at a processing stage well after the encoding of the stimulus and possibly the decoding of such information. This last aspect is addressed directly in Experiment 3.

3.3 Experiment 3: Is the Fine Discrimination Task Sufficient for the Bias to Arise?

Results from Experiment 2 showed that making a binary judgement on a stimulus feature – performing the fine-discrimination task at the beginning of a trial – is not sufficient for the repulsive bias to arise. Despite having subjects discriminating between CW and CCW directions of motion, it was the absence of the oriented reference line to impact the resulting biases. The experiment outlined here addressed both the role of the fine-discrimination task, as well as shedding further light on the role that the oriented reference line during the estimation task has on the overall repulsive bias effect.

3.3.1 Materials and Methods

The methods used here are similar to those described for Experiment 1 and any differences are detailed in the appropriate sections below.

Participants performed a minimum of 2500 trials and those that did not take part in the previous studies performed 400 practice trials on average before engaging in the experimental ones.

3.3.1.1 Observers

Five subjects (four females, age: 23.4 ± 2.8 years, mean ± s.d.) took part in the experiment: three of them participated in Experiments 1 and 2 as well, two were recruited anew. All had normal or corrected-to-normal acuity.

3.3.1.2 Apparatus and Stimuli

Stimuli were defined and presented as described in Experiment 1 – Apparatus and Stimuli.

3.3.1.3 Procedure

Subjects were presented with the same stimuli as in Experiments 1 and 2 with the same timing and procedure, with the following modifications: (i) subjects were instructed to passively view the stimuli and *not* to perform the fine-discrimination task, (ii) subjects manually estimated the direction of motion in *all* the trials and *crucially* (iii) unbeknownst to the subjects the orientation of the reference line was manipulated just before the onset of the estimation task. In particular, the reference could either remain consistent with that at the start of the trial, or it could vary by ±6° from the initial orientation (see Figure 3.10). Thus, for this experiment, the task

consisted purely in *estimating* the direction of motion of the coherently moving dots.

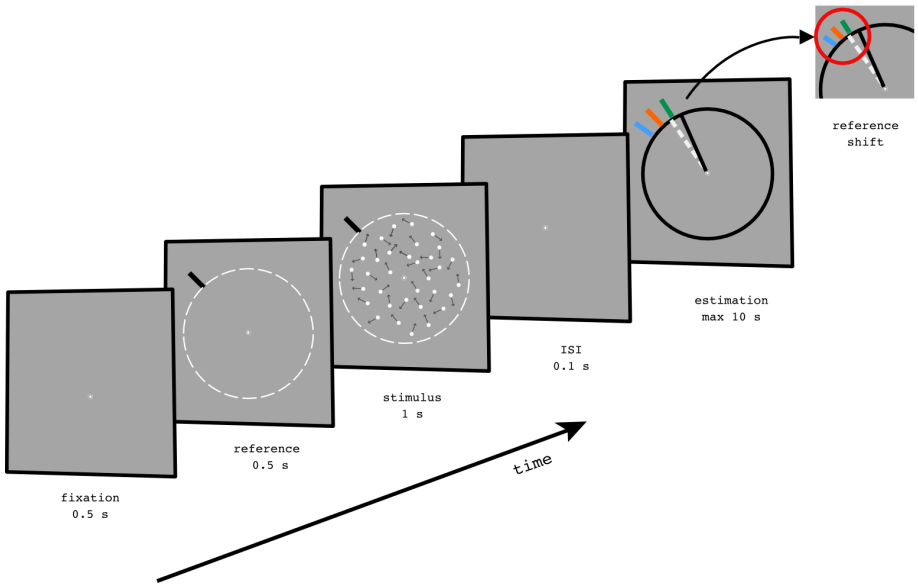


Figure 3.10 Schematic representation of a trial sequence: the oriented reference line is present throughout the whole duration of the trial but its angular position is changed during the estimation task. This is highlighted in the magnification on the top right corner. Different colours of the reference line correspond to different conditions tested (blue: reference moved CCW, orange: reference did not move; green: reference moved CW).

3.3.2 Results

Data from the estimation task were sorted according to the shift in the oriented reference line (i.e., -6°, 0°, and +6° with respect to the angular position of the reference at the beginning of the trial; see Figure 3.11). As observed in Experiment 1, when the shift in angular position in the reference orientation between time at stimulus presentation and time of estimation response remained unchanged, there was a repulsive bias. Interestingly, this was despite the subject not performing the fine-discrimination task, suggesting that an overt / explicit binary judgement is not necessary in generating a biased response. For the conditions in which the reference line

was shifted by 6° CW from its location during stimulus presentation, the distribution of estimates shifted with it, revealing the same repulsive bias from the reference line at time of estimation. In a similar fashion, the distribution of estimates shifted in a consistent direction for the CCW manipulation.

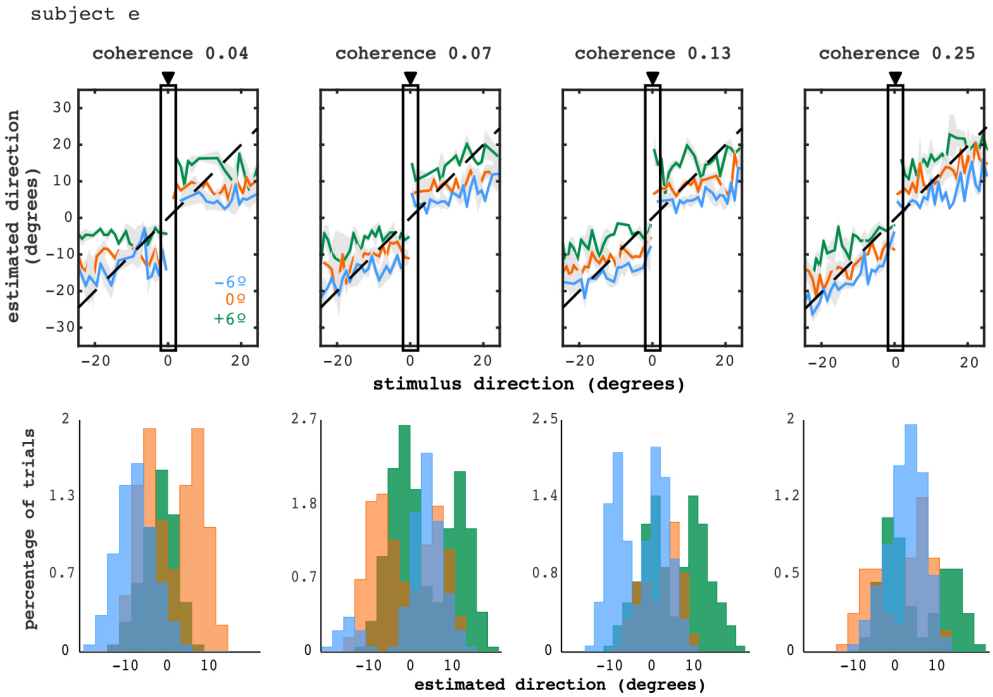


Figure 3.11 Data from one subject, four coherence levels tested. Each colour represents a reference condition (see legend in panel 1, row 1). Row one: mean direction estimates plotted as a function of stimulus direction. Shaded grey areas represent ±1 s.e.m. Row two: histogram of estimates distribution for the column-wise selection highlighted by the black rectangle on the plots in row one.

Interestingly, when inspecting the centres of the distributions of estimates, these resulted to overall agree with the angular position of the reference line at time of performing the estimation task. The repulsive effect was therefore yoked to the oriented reference available when performing the task, whereas information from the location of the oriented reference during stimulus presentation was discarded (see Figure 3.11, second row).

Data from all subjects (one for each row), all coherence levels (one for each column) tested are presented in Figure 3.12. Mean direction estimates (± 1 s.e.m., shaded area) are shown as a function of stimulus direction, and colours indicate data for each reference shift condition (-6° blue, 0° orange, and +6° green). A clear gap around the 0° condition (i.e., when signal dots moved in the same direction as the orientation of the reference line) is evident, and this repulsive effect is shifted CW / CCW in agreement with the corresponding shift in the boundary during the trial.

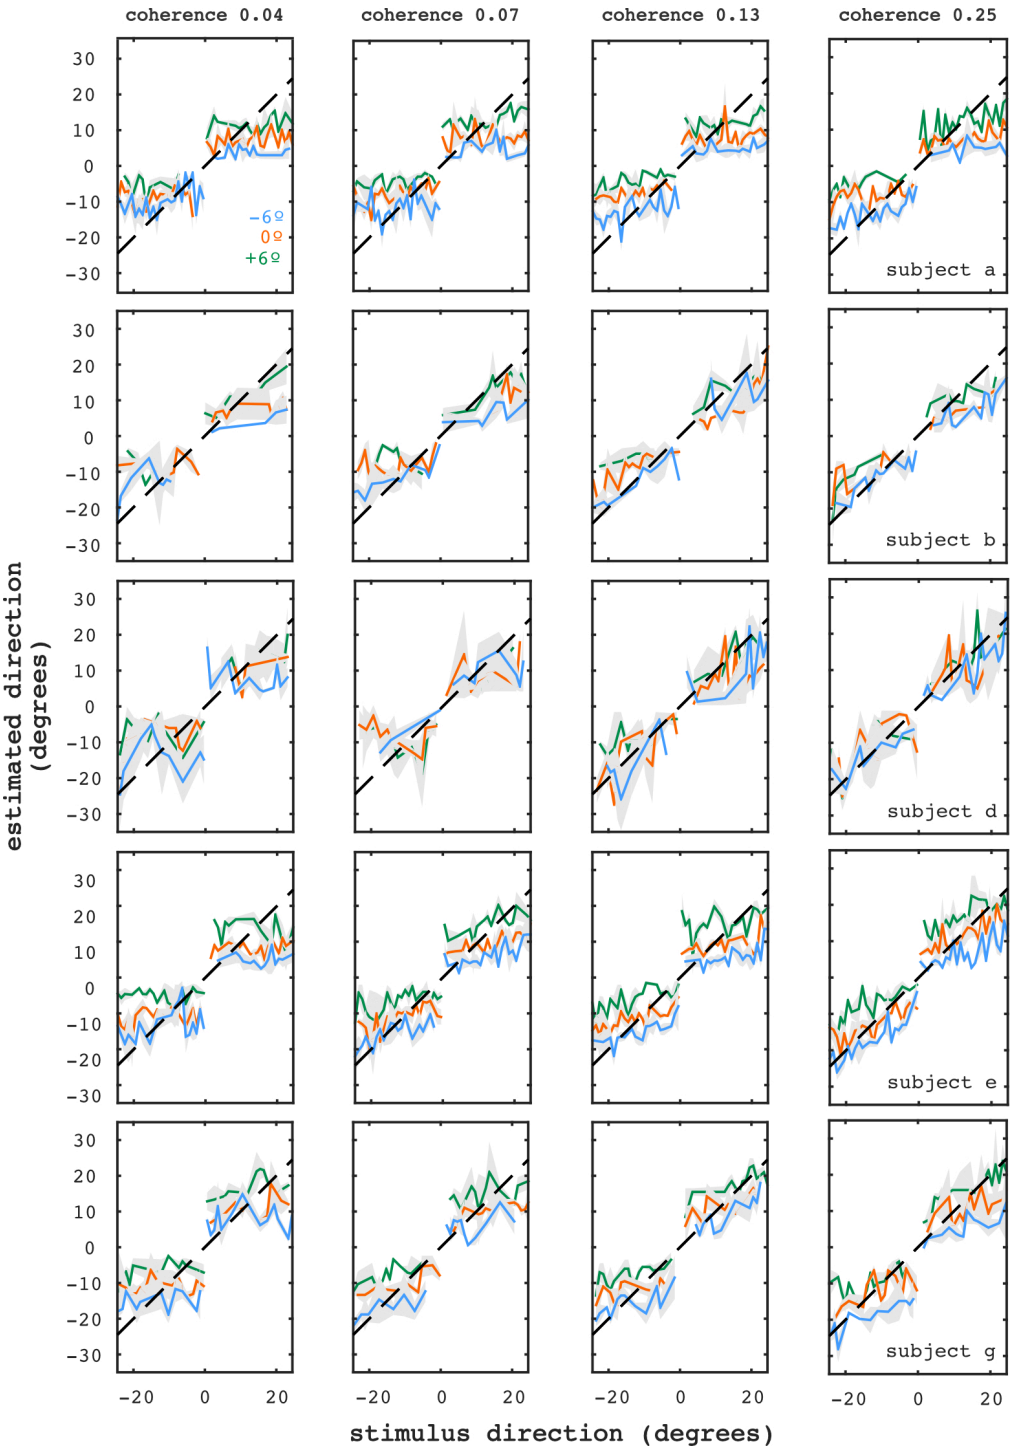


Figure 3.12 Same representation of data as in row one of Figure 3.11. Each row here presents data from a different subject. Each column represents a different coherence level tested.

Thus, not only do these findings support those in Experiment 1, showing that subjects are biased when estimating the perceived direction of motion of an RDK in the presence of an oriented reference line, but they further extend their impact. Here, subjects passively viewed the stimuli, but were

not required to make a forced choice. Nonetheless, the repulsive effect in the estimates was still evident and yoked to the reference during the task, regardless of its angular location at stimulus presentation. Moreover, these results highlight the importance of the reference line during the estimation task: its availability is fundamental for the biased responses to arise (findings from Experiment 2 indeed suggest that, when the boundary is absent, subjects estimate perceived direction of motion in a veridical, although noisy, way).

3.4 Discussion

The experiments presented in this chapter provide new insights into the nature of a well-known perceptual bias: *reference repulsion*. The point, during sensory information processing, at which such bias emerges was revealed by applying a series of small and systematic manipulations to a task sequence used in previous studies. In particular, the aim was to disentangle the contribution of making a forced judgement to that of contextual cues in biased estimates of perceived directions of motion.

In Experiment 1 we assessed that we could actually replicate the previous finding of a repulsive effect when estimating the perceived direction of motion of a field of signal dots after having performed a fine-discrimination task.

We then proceeded to check whether the simple act of performing a forced choice discrimination task results in biased estimates. Indeed, if the bias is a consequence / result of the strategy used in pooling neural

population responses to decode sensory information and depends on the task to be performed, changes in the trial that follow the forced choice should not affect the overall distribution of estimates. In Experiment 2 subjects, therefore, were asked to perform the same fine-discrimination task as in Experiment 1, but when prompted to perform the estimation task, the oriented reference line was removed, so that they were presented with a blank context.

Another way to test, not only the effect of contextual information present during the estimation phase, but also to address the impact of an explicit binary judgements to the overall bias was presented in Experiment 3. Here subjects were instructed to passively view the moving stimulus, together with the reference line, and only perform the estimation task in the presence of either the very same oriented line, or of a shifted version of it.

Experiment 1 supported the findings of Jazayeri & Movshon (2007a) that subjects are biased when estimating the perceived direction of motion of a field of moving dots in the presence of an oriented reference line. This was true for all coherence levels tested (see Figure 10.1 in Appendix), although the strength and reliability of the effect varied with the performance of subjects in the fine-discrimination task preceding the estimation (see Figure 10.2, subject c and subject f in Appendix).

Experiment 2 showed that, when the reference line is removed during the estimation task, the bias disappears (see Figure 10.4 in Appendix). This suggests that subjects have a veridical representation of motion direction

that can be used during the estimation task. Further, this indicates that: (i) making a binary judgement on a stimulus feature is not sufficient for biasing its representation, and (ii) the presence of the reference line during the estimation task is necessary for the bias to occur.

Experiment 3, further strengthens the major role played by the oriented reference line during the estimation task in the emergence of the repulsive bias: directions of motion estimated after passively viewing the same stimulus configurations as in Experiments 1 and 2 result in biased distributions according to the angular location of the reference at time of task execution (i.e., regardless of reference location during stimulus presentation). Furthermore, given that the bias is still present, it is evident that an explicit binary judgement is not necessary to generate a biased response.

The findings presented in this chapter further inform a perceptual phenomenon widely investigated in the field of visual neuroscience. These same data were also used to test and modify a model previously proposed by Jazayeri & Movshon (2007a). The model is described in Chapter 4: a *late-decision related bias model,* integrating information about motion with that of an explicit reference cue can account for the perceptual bias described in this chapter.

Another important aspect highlighted by the findings presented here is related to the role of the modality with which subjects responded. A particularly interesting question is whether the biases also arise when subjects report their percepts with different means. This specific question is

addressed in Chapter 5, where the logic underlying the psychophysical experiments outlined in this chapter is translated to an eye-movement study.

CHAPTER 4[2]

4 Modelling Biases in Estimates of Motion Direction

Visual motion perception is tightly linked to perceptual decision making for action. Tasks used to investigate how perceptual decisions are formed often exploit the flexibility with which stimuli conveying motion information can be manipulated. Thus, assumptions and models aimed at characterising the processes involved in forming perceptual judgements are often presented within a framework of vision. In what follows, I will therefore first give a brief overview of some of the most well-known computational models describing mechanisms involved in visual motion perception. I will then provide the link to perceptual decision making and focus on decoding strategies that may be deployed in neuronal populations of the visual system to optimally perform a specific set of tasks. This leads to the description of a *decoding model* proposed by Jazayeri & Movshon (2007a). I will also describe its implementation, together with that of a modified version – the *late, decision-related bias model* – I propose to account for the data presented in Chapter 3.

---

[2] Some of the work presented in this chapter has been published in the following article:
Zamboni E, Ledgeway T, McGraw PV, Schluppeck D (2016). Do perceptual biases emerge early or late in visual processing? Decision-biases in motion perception. *Proc. R. Soc. B* **283:** 2016.0263 http://dx.doi.org/10.1098/rspb.2016.0263

Computational models of visual *motion* perception generally assume two separate stages of processing: first low-level information is extracted from the stimulus signal, then the local information is combined together (pooled, integrated or compared in some manner) following a set of rules to generate a global percept. These two processes are often mapped onto areas V1 and MT, respectively. Models that include two such stages capture well some of the observed properties of MT neurons such as their direction selectivity and the spatial layout of their receptive fields (Rust et al., 2006; Simoncelli & Heeger, 1998). The fundamental goal of these types of models is to describe the *bottom-up* encoding process of sensory information and do so by focussing on integration of low-level properties inferred from the moving stimulus. On the other hand, a slightly different approach proposes to take into account strategies used in *decoding* visual information by the brain. Here, a model is generally based on top-down assumptions regarding the way in which certain visual features are represented in the brain and it is this specific strategy that is modelled. The most-known algorithms used to capture how visual information may be decoded from populations of neurons are: *vector averaging* (VA), *winner-takes-all* (WTA), and *maximum likelihood* (ML) models (Webb et al., 2007).

The VA algorithm assumes that direction selectivity in visual area MT is defined as an oriented vector. Imagine representing the preferred direction of motion of each neuron as a tuning function / vector. The contribution of each such neuron to the perceived direction of motion is then given by weighting the neuron's tuning by the response strength observed for a

certain stimulus motion. The final output, the information decoded from a population of neurons, is obtain by averaging each such vectors (Groh et al., 1997). In particular, the weight to be applied to each vector is given by the strength (firing rates) with which the corresponding neuron responds to the presented moving stimulus. This process is illustrated in Figure 4.1 presenting vector responses (black segments radiating around axis of motion) for a population of neurons to eight different directions of motion. The resulting average population vector is indicated by the arrow emerging from the eight populations of vectors (Dayan & Abbott, 2005).

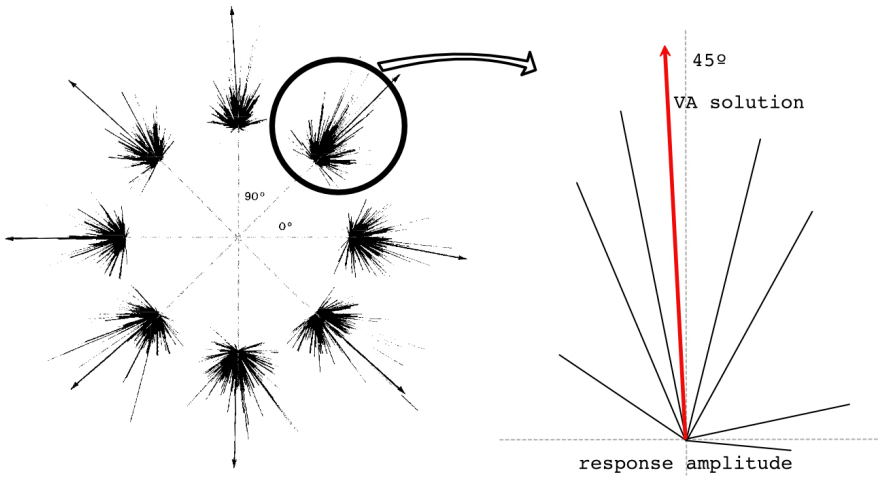


Figure 4.1 Population vector representation: From the left, visualisation of eight stimulus directions and corresponding response directions for each movement. The arrows in each group of line represent the vector sums of responses for each direction. On the right a magnification of the responses for stimulus direction of movement of 45° and corresponding VA solution is shown. Adapted from Dayan & Abbott, 2005

An alternative view proposes to discard weak responses to a stimulus and rather use the information provided by the most active neuron as the decoded feature. This, therefore, is less of a *pooling* mechanism: contributions from individual neurons in response to a stimulus are not necessarily combined (e.g. through a weighted average) to generate the final output. Instead, signals are kept segregated and the perceived, global

direction of motion strictly depends on the neuron with highest response to the stimulus. In this sense the strategy is called *winner-takes-all* (Salzman & Newsome, 1994). This mechanism is schematically represented in Figure 4.2: stimulus direction of motion is given by the dashed line, solid lines radiating from it represent individual neural responses to the stimulus, and the decoded direction of motion is given by the arrow superimposed to the highest neural response within the population.

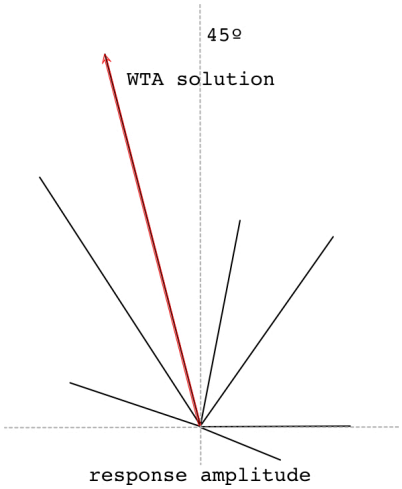


Figure 4.2 Winner takes all representation: black lines represent neuronal responses to a direction of motion of 45°. The neuron showing the highest response to the stimulus is taken as the most informative (red arrow). Adapted from Rocchi, 2013

The validity of this model was tested in a study by Salzman and Newsome (1994): direction-selective columns of monkeys' MT areas were microstimulated while the animals performed an eight-alternative direction discrimination task. After having assessed the preferred direction of motion of the recorded / stimulated MT column, monkeys were presented with RDKs, computed at different coherence levels, moving in one of eight possible directions. When stimulus and fixation point were extinguished, the animals had to indicate their choice – the perceived direction of motion

– by making a saccadic eye movement to one of the eight targets displayed on the screen. According to a VA framework, stimulating neurons whose preferred direction of motion is at least 90° apart from the stimulus direction, should result in choices lying in between stimulus and preferred, stimulated directions. On the other hand, a WTA approach predicts that either the direction preferred by the stimulated neurons, or that of the presented stimulus should result in the monkeys' choice (these two options should be those with highest neural response intensity). The data showed a larger number of choices for the direction preferred by the stimulated neurons compared to the number of responses to directions in between that of the stimulus and the one preferred by the stimulated neurons. The results, therefore, were consistent with the WTA assumption. A more flexible view of read-out strategies used in visual area MT when reporting perceived directions of motion is provided by Nichols & Newsome (2002). Applying the same experimental paradigm described above, monkeys had to indicate the direction of a moving stimulus by means of either a saccadic eye movement or a manual response to one of the targets displayed on the screen. Interestingly, results showed that both VA and WTA can be used to capture specific neuronal properties. While vector averaging captured the behaviour of neurons representing direction selectivity within a range of up to 140° and contributing to the decoding process of the moving stimulus (see also Lisberger & Ferrera, 1997), for neurons with tuning properties outside such range, a WTA strategy worked best in predicting direction read-out (in agreement with Salzman & Newsome, 1994). Therefore, rather

than providing one exclusive mechanisms, it is possible that activity in area MT may be read out through different decoding strategies.

Despite evidence supporting both VA and WTA algorithms, a common drawback is that neither accounts for the stochastic nature of neuronal responses – they do not take into account the inherent noise characterising the response, but rather rely on estimating a single best value representing the sensory information. An approach taking into account neuronal response variability adopts *likelihood functions* to describe the sensory information encoded in the brain (Rocchi, Ledgeway, & Webb, 2013). By considering the task of perceptual decision making in terms of statistical inference, it is assumed that the incoming sensory input evokes noisy signals in a population of (motion selective) neurons and the direction selectivity of each cell is described by *distribution of possible internal responses*. The population response is computed as a combination of these individual tuning functions: an overall likelihood describing the encoded stimulus (Jazayeri & Movshon, 2006; Webb, Ledgeway, & McGraw, 2007). The decoded direction of motion is then assumed to be represented by the peak (maximum value) of the likelihood distribution. This strategy is well described in Figure 4.3: each neuron contributes to a certain extent in encoding a stimulus property. The noisy response of such neurons is described as a weighted likelihood function: a product of the spiking response and the tuning curve of the cell. The overall likelihood function is obtained by adding the individual contributions from each cell and the peak

of the resulting distribution can be thought of representing the most likely stimulus feature (Jazayeri & Movshon, 2006).

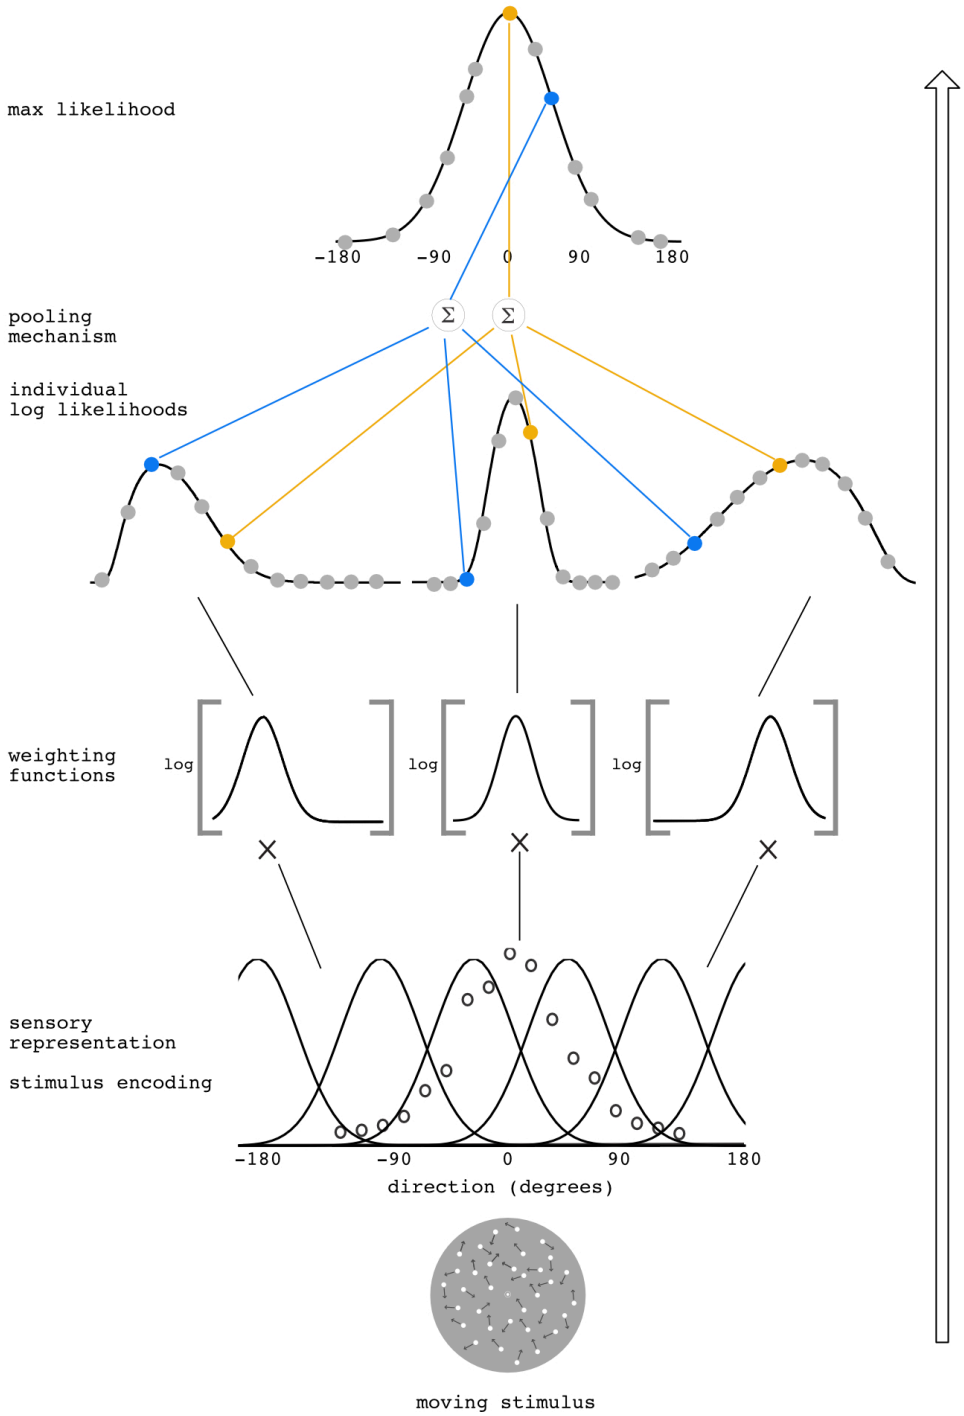


Figure 4.3 Maximum likelihood pooling representation. From the bottom to the top (as indicated by the arrow on the right): sensory information (moving stimulus) is encoded by a population of neurons (individual distributions). Each response is multiplied by the logarithm of that neuron's tuning function. The resulting likelihood function is obtained by summing together the individual log likelihoods. The decoded feature (direction) is represented by the maximum of the overall likelihood. Adapted from Jazayeri & Movshon, 2006

The maximum likelihood algorithm is a key concept in understanding the link between responses of visual neurons selective for sensory features and responses related to perceptual decisions. Indeed, as described in Chapter

1, Section 1.1, discrimination of visual motion direction has been extensively used as a proxy to study processes underlying perceptual decision making mechanisms. Some of these studies show that it is possible to measure the outcome of a perceptual decision based on responses from sensory neurons encoding a stimulus.

For the specific question addressed in this thesis, we need to consider the mechanisms involved in discriminating directions of motion. When stimuli present highly dissimilar features (e.g., directions of motion far apart from each other), the best strategy to encode the sensory information is by preferentially weighting the responses of neurons that are selectively tuned to such direction (Hol & Treue, 2001). This can be described in terms of maximum likelihood functions as mentioned above and represented in Figure 4.4: given that the two alternatives to be discriminated are highly dissimilar, neurons tuned to each alternative are weighted the most (Jazayeri & Movshon, 2006). Variability is associated with the contribution of responses from neurons tuned to close directions that will still provide a considerable amount of information, although it will be less accurate (Jazayeri & Movshon, 2007b).

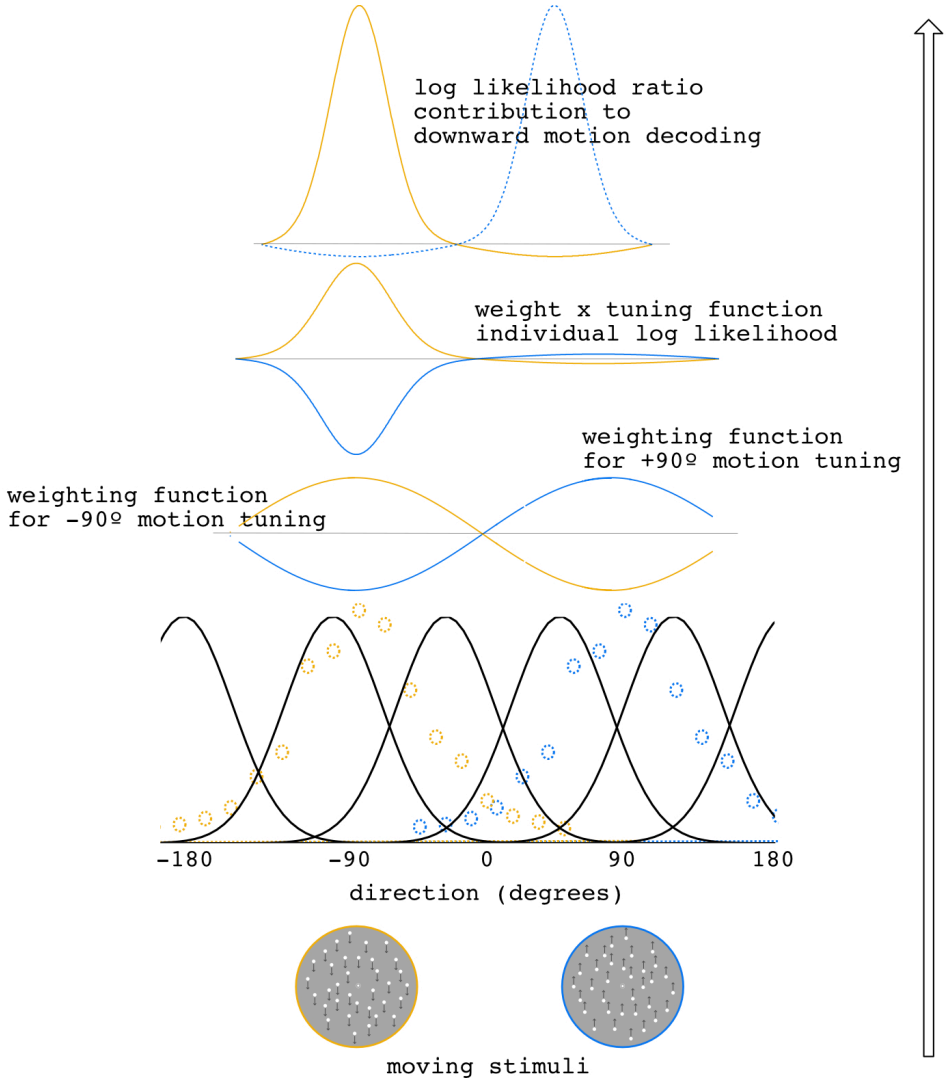


Figure 4.4 Likelihood computation. From bottom to top (as indicated by the arrow on the right): random dot stimuli evoke responses in neurons sensitive to motion. Neuronal responses are represented by each distribution (yellow and blue circles show noise perturbing the responses). By multiplying the encoded signal with an appropriate weight, individual log likelihood functions are obtained. The difference between the contributions of one neuron to the two sensory signals represents the log likelihood ratio. Adapted from Jazayeri & Movshon, 2006

On the other hand, when discriminating between two highly similar alternatives, information coming from neurons that are each preferentially tuned to one of the two directions of motion will also be very similar, making it difficult to distinguish between options. Analogously, neurons whose preferred direction of motion is far removed from the target ones will have little to no contribution in the discrimination task. The most informative neurons for this specific task are therefore those with tuning functions that

peak at directions slightly away from those of the targets (Regan & Beverley, 1985; Hol & Treue, 2001). The mechanisms involved in applying this task-specific pooling strategy are well described through the maximum likelihood algorithm mentioned previously. Interestingly, while the models presented so far rely on the assumption of having an ideal observer (i.e., an observer whose estimates / responses are veridical and not biased), perceptual biases have been reported for several visual motion tasks (Putzeys et al., 2012).

In particular, the goal of this chapter is to explain the perceptual bias of *reference repulsion* (Rauber & Treue, 1998). Thus, starting from the assumptions of Jazayeri and Movshon's decoding model (2007a), together with the rationale at the core of the experiments described previously, the *Late, Decision-Related Bias Model* is developed. Recall from Chapter 3 that the presence or absence of a reference line has a marked effect on how perceptual biases present. In what follows, a description of implementation and validation of both decoding and late, decision-related bias models is provided.

4.1 Decoding Model

A version of the Jazayeri and Movshon (2007a) *decoding model* – comprising of two main components – was implemented and fitted to the data. The first component reflects the *sensory representation* of the moving stimulus, modelled as a Gaussian probability density function, centred on

the stimulus direction of motion and whose variance is derived from the discrimination psychometric functions. In particular, the amount of spread of the representation changes with the coherence level: as the proportion of dots moving coherently (evidence for a particular direction) decreases, the variance of the sensory representation increases. Mathematically speaking, the sensory representation is described as:

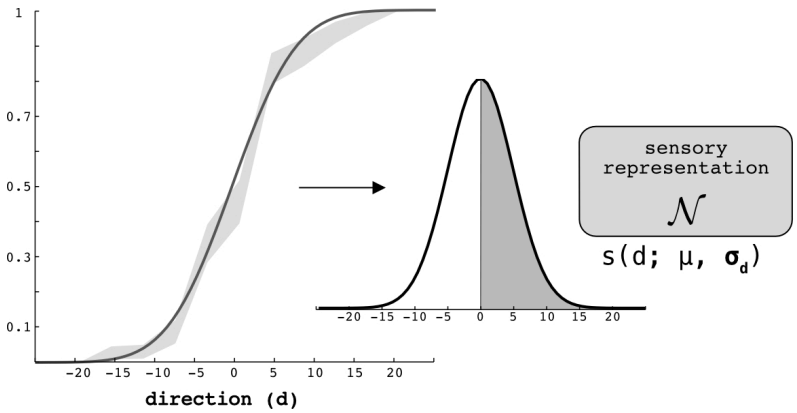


Figure 4.5 Sensory representation: relationship between cumulative distribution characterising behavioural performance and probability density function characterising the internal representation of a certain stimulus feature. The insertion on the right highlights the description of the representation as a function of d (direction) with parameters μ (mean) and σ (variance).

As described in Chapter 2, Section 2.1.1.5, the proportion of CW-responses is plotted as a function of the angular difference (in degrees) between the orientation of the reference line and the direction of motion of the signal dots in the moving stimulus. This is done for each coherence level tested, for each subject. A *non-linear least squares* method is used to estimate the two parameters (μ and σ) of the best fitting cumulative Gaussian distribution. The σ is then used to compute the sensory representation of the moving stimulus: the angular differences (*deltas*) between reference line and signal dots are iteratively used as mean values to obtain Gaussian probability density functions. Each distribution is then

evaluated on a finely-spaced set of directions of motion (ranging from the largest anti-clockwise delta to the largest clockwise one), and its corresponding variance is that derived from fitting a Gaussian cumulative density function to the discrimination data. The process is described below (Figure 4.6) and results in a normally distributed sensory representation centred at each tested delta, for each coherence level and each subject. Figure 4.6 shows the sensory representation for stimuli that move in direction 0° from the reference, at different coherence levels.

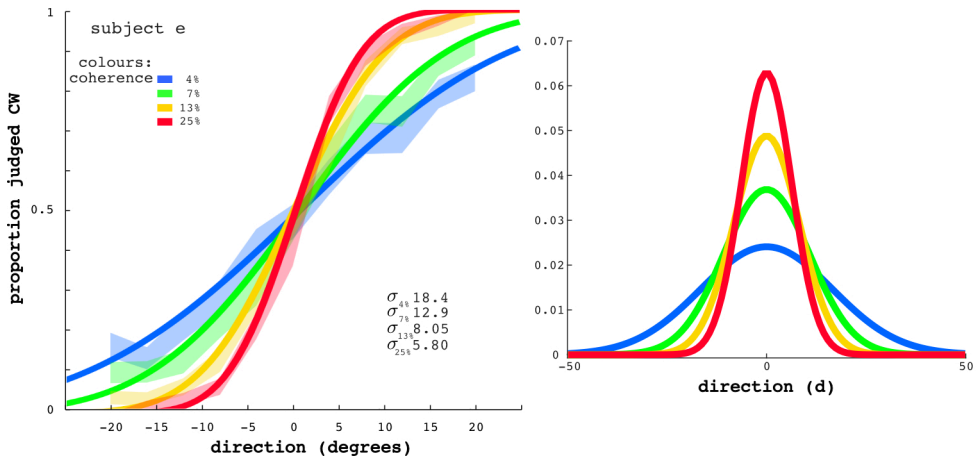


Figure 4.6 Derived sensory representation from psychometric functions. On the left, psychometric functions obtained from one subject tested at four levels of motion coherence. On the right, probability density functions derived from the behavioural data indicating sensory representations for the different stimulus conditions.

The second component of the model is a *weighting profile*, whose aim is to boost the responses of directionally tuned mechanisms that are optimal for performing a fine-discrimination task: neurons with tuning functions peaking slightly to the sides of the decision boundary are more informative for such tasks than those tuned to directions very close to (or far away from) the boundary (see Regan & Beverley ,1985; Hol & Treue ,2001; Figure 1.6).

Different functional forms can capture this pattern, but to stay in line with Jazayeri & Movson's (2007a) model implementation, the weighting of information for the decoding model was implemented as a gamma probability density function. The gamma function is controlled by two parameters: *shape* and *scale*, and is mathematically described as:

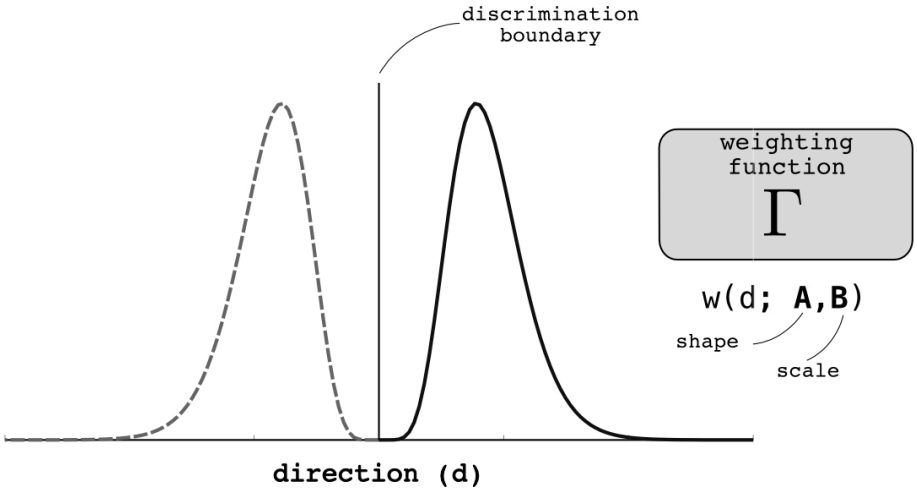


Figure 4.7 Weighting profile. Gamma probability density function used to systematically weight the sensory representation. The insertion on the right highlights the description of the weighing profile as a function of d (direction) with parameters A (shape) and B (scale).

The specific *shape* of the function can capture not only a *repulsive* effect around the reference line (as a result of the response pattern of directionally tuned mechanisms involved in the discrimination task), but also the influence of a boundary imposed by the stimulus range. In fact, by performing sequences of trials, subjects experience a range of stimuli, each related to the orientation of the reference line. Despite the randomised order in which different trials are presented, it is possible that subjects build an internal representation of the range of possible directions that the moving stimulus can assume. This particular boundary, however, is unknown to the subjects, and therefore noisier than the boundary defined by the reference line. The asymmetric nature of the gamma function,

steeper close to the oriented reference line and shallower at the other end, captures the stimulus range and its uncertainty.

A key aspect here is that the weighting profile is yoked to the angular location at which the reference line was presented during stimulus presentation for that specific trial. The *late, decision-related bias model* enhances the detail of information provided by the weighting profile (see below).

In order to predict the estimated direction of motion, the weighting profile is applied to the sensory representation for that trial. Mathematically, this is obtained by multiplying the Gaussian probability density function with the gamma distribution:

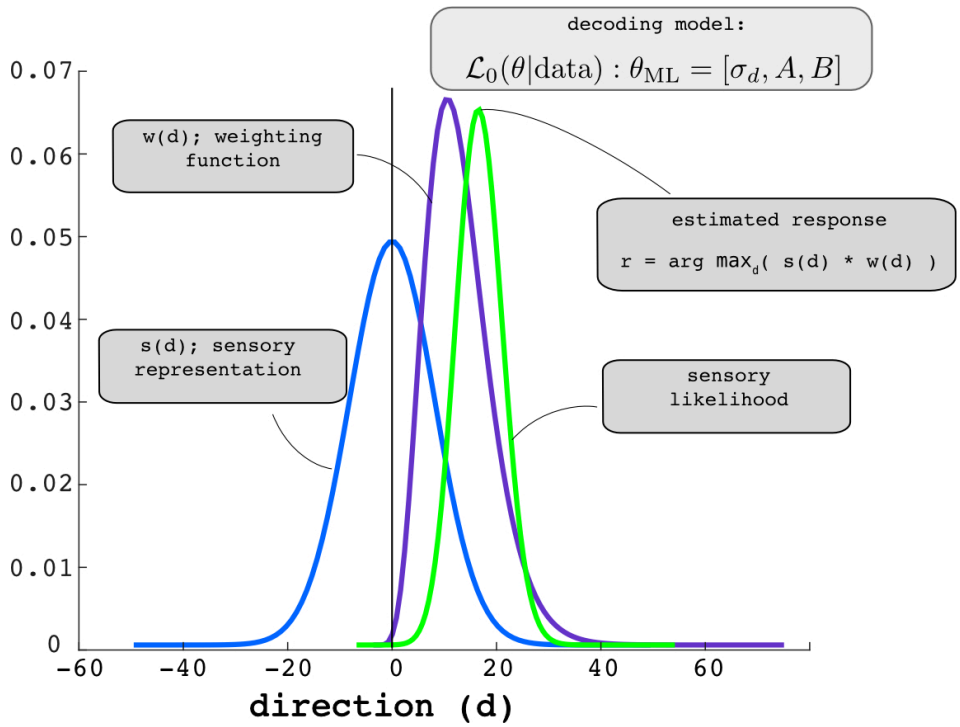


Figure 4.8 Decoding model. From left to right: the sensory representation is centred at the direction in which the stimulus was moving. The distribution is weighted by a gamma function estimated for each individual subject. The result is a sensory likelihood distribution whose peak, shifted from the centre of the sensory representation, is taken as the estimated direction for a certain trial.

The resulting distribution, called the *sensory likelihood distribution*, is shifted from the derived internal representation of the moving stimulus and its peak is taken as the best prediction of the estimated direction of motion for that subject, for that trial. A detailed description of the computational steps involved in estimating sensory representation, weighting profile, and likelihood distribution, for each subject, each coherence level, and experiment condition (i.e., *reference present*, *reference shift*) are given in the *Late, Decision-Related Bias Model*, Section 4.2. Here, I will first give a description of the model, followed by the assessment of the data fitting performance for each testing condition and the resulting implications with regard to interpreting the reference repulsion effect.

4.2 Late, Decision-Related Bias Model

The goal of the studies presented in Chapter 3 is to address the question of how the presence of a reference line leads to large perceptual biases? To do this, information available from the line reference was manipulated, which required adjustments to the *decoding model* described above. It is worth noting that because both the decoding model and the late, decision-related bias model share the final aim of capturing and explaining perceptual biases, components playing a role in these phenomena are also shared between the two models. However, the detailed timeline / order in which specific processes take place result in a substantial difference in the logic (and in the fitting performance) of the models.

The first step of the models (both, decoding and late, decision-related bias) is to estimate the internal / sensory representation of the moving stimulus that is presented during a certain trial. As mentioned in the *Decoding Model* section, this is obtained by fitting a Gaussian cumulative distribution function to the data from the discrimination task, estimating the variance of the curve and use this variability as an indicator of the internal representation reliability.

This procedure can be applied to data from Experiment 1. However, in Experiment 3, subjects viewed the oriented reference line and the moving stimulus without actively performing a binary judgement. Therefore, the data from the discrimination task were not available, requiring the inclusion of an extra free parameter to the late, decision-related bias model capturing the spread of the sensory representation of the moving stimulus.

Subsequently, a weighting profile needs to be estimated in order to be able to capture the observed repulsive effect in the estimated directions of motion. We saw that the function used in this process is a gamma distribution with shape and scale as parameters, yoked to the angular position of the reference line. However, in Experiment 3, the position of such a reference could vary during the trial, shifting by 6° clockwise /anti-clockwise between the time of stimulus presentation and that of response. In order to account for this manipulation, the weighting function needs to shift along its abscissa (providing a free parameter in direction space). From a theoretical point of view, this means that the profile captures the state of the oriented reference line throughout the whole trial, allowing the

weighting of information to happen at a stage later than the decoding one, that is during the response. Computationally, this is obtained simply by allowing an extra free parameter in the model:

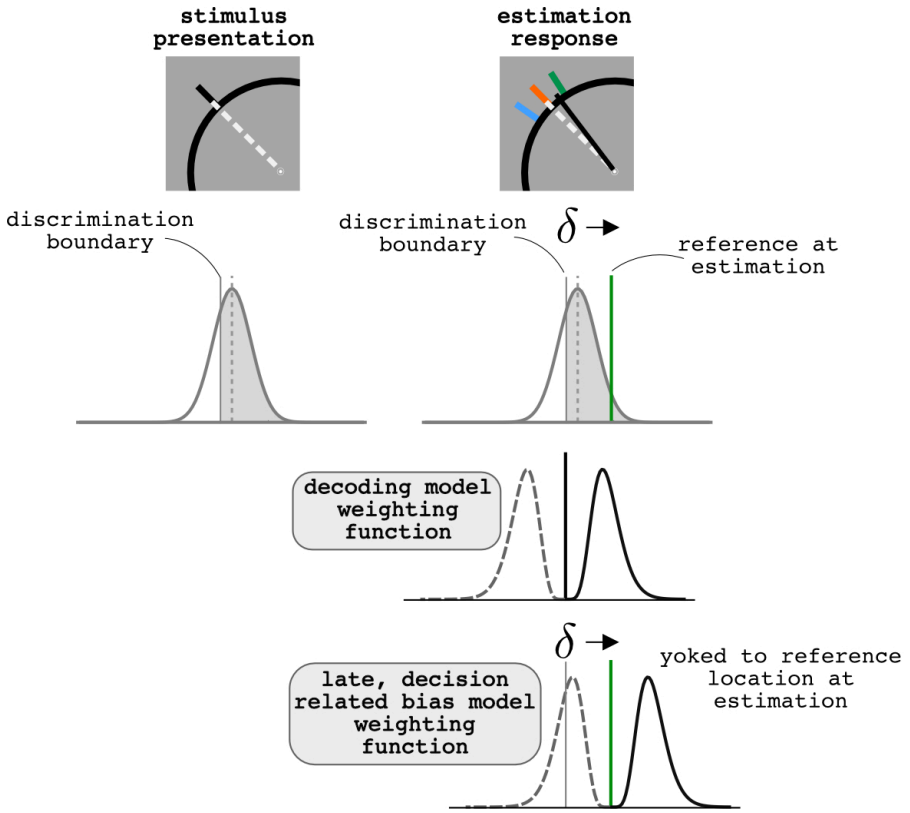


Figure 4.9 Late, decision related bias model applied to 'reference shift' condition. The distinction between the angular location of the reference line at time of stimulus presentation and of response is shown by the difference between the solid grey and green lines on top of the sensory representation. The distinction between the decoding and late, decision related bias model is highlighted in the second and third rows: the weight in the decoding model is yoked to the angular location of the reference at the time of stimulus presentation. In the late, decision related bias model, on the other hand, the weight is yoked to the angular position of the reference at the time of response.

The full description of the late, decision-related bias model is given in Figure 4.10. Depending on the task condition performed (reference *present* / reference *shift*), the parameters of the weighting function are estimated using all data from all coherence levels tested (reference *present*) or one set of data at a time (reference *shift*). In a nested fashion, these parameters are also iteratively applied to obtain the best fit to the data, for each coherence level, for each condition (reference *shift*). Thus, on one hand there is a

distribution of estimated directions of motion (one for each trial, for each stimulus direction). On the other hand, there is the maximum value of the distribution resulting from the product between the sensory representation and the weighting profile. This maximum value is used to minimise (in the least squares sense) its distance from the centre of the distribution of estimated directions. Once the optimal solution is found, the best parameters describing the weighting profile and the best fit to the data are returned.

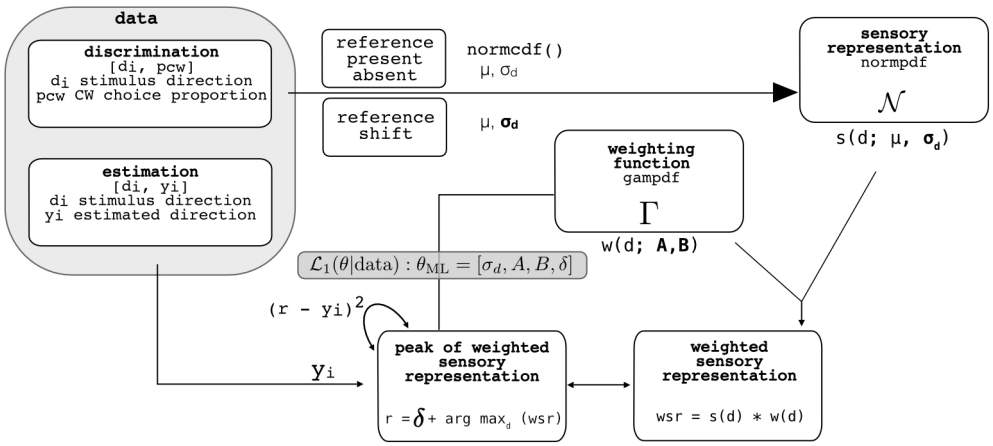


Figure 4.10 Graphical representation of the late, decision related bias model. From the left: when data from the discrimination task are available (reference present / absent conditions), these are used to derive the sensory representation of the moving stimulus. When discrimination data are not available (reference shift condition), the variability of the sensory representation is estimated. Data from the estimation task are used to derive the optimal weighting function for each subject such that, when multiplied with the sensory representation, the peak of the resulting distribution best fits (in a least squares sense) the data. Parameters that are estimated by the model are represented in bold.

By applying this model, therefore, it is possible not only to fit the data, but also to estimate the angular shift of the oriented reference during the trial and the bias compared with the angular position of the reference at the start of the trial. The results obtained from fitting the model to the data collected in Experiment 1 and Experiment 3 (Chapter 3) are presented next.

4.2.1 *Modelling Performance: Reference Present*

Data for the reference present condition were fitted using the late, decision-related bias model. Figure 4.11 shows data from one subject, four coherence levels tested: grey scale joint histograms represent the data from the estimation task (see detailed description in Chapter 3 – Experiment 1), the model fit is displayed in green on top of the data. Graphically, it is possible to appreciate the reliability with which the fit describes the data. Moreover, the reconstructed weighting profile for this subject is presented in Figure 4.11.

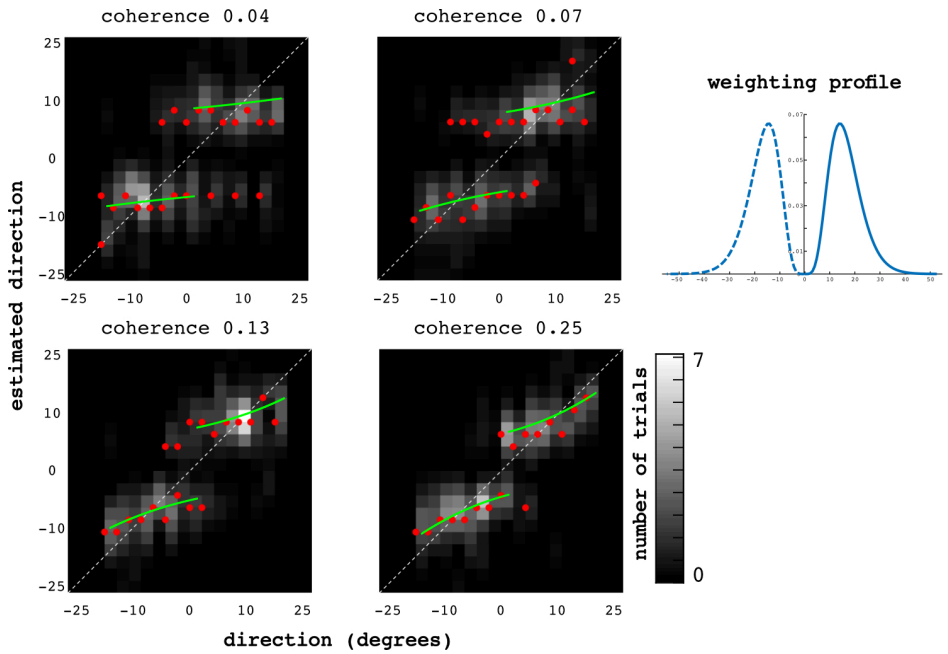


Figure 4.11 Model fit of data from the reference present condition. Data from one subject, four coherence levels tested. Model fit is displayed in green, red dots represent the peaks of each distribution of estimated directions for a presented stimulus direction (column-wise distributions). White dashed line represents veridicality of responses. The light-scale shows the scaled number of trials constituting each distribution. On the right, the graphical description of the weighting profile estimated for this subject.

Table 4.1 presents the computed coefficient of determination (r^2) for each coherence level, for each subject. For all five subjects tested, the model captured the distribution of direction of motion estimates with high reliability (capturing on average 84% of the variance in the data).

r^2	4% coherence	7% coherence	13% coherence	25% coherence
Subject a	0.705	0.71	0.743	0.765
Subject c	0.830	0.836	0.839	0.868
Subject d	0.869	0.900	0.914	0.911
Subject e	0.861	0.859	0.893	0.868
Subject f	0.816	0.813	0.814	0.833

Table 4.1 Coefficient of Determination (r^2) for each coherence level (columns), each subject (rows).

The results show that the model implemented is consistent with the repulsion effect described in Chapter 3 and also reported by Jazayeri and Movshon (2007a). Therefore, the late, decision-related bias model well predicts the repulsive bias under investigation, as was also the case for the decoding model. In fact, in this specific experimental setting (i.e., subjects first perform a fine-direction discrimination task and subsequently estimate the perceived direction of motion in the presence of an oriented and *unchanged* reference line), the model predictions are indistinguishable.

4.2.2 Modelling Performance: Reference Shift

The manipulation of the task sequence in Experiment 3 required more flexibility in designing the model, so that it could account for the information that is provided by the oriented reference line at the time of response. According to the findings described in Chapter 3, perceptual biases arise at a late stage and are strictly related to a boundary presented during task execution, even though this should provide no relevant

information for the actual performance. The late, decision-related bias model provided a good fit of the data as shown in Figure 4.12. Thin line with shaded grey areas correspond to mean ± 1 s.e.m. of the estimated directions of motion, whereas the thick line shows the model fit. Three different colours indicate the shift-condition of the reference line: -6° in blue, 0° in orange, and +6° in green and each panel corresponds to a tested coherence level.

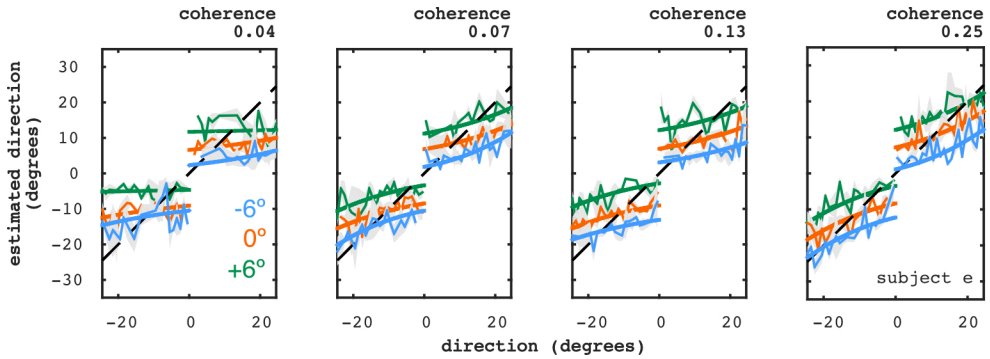


Figure 4.12 Model fit for the reference shift condition. Data shown for one subject, four coherence levels tested, three shift conditions. The thick, solid line shows the model fit to the data, whereas the thinner, jittering lines represent the running average for the estimated directions of motion. Black dashed line indicates veridicality (estimated and stimulus directions are equal).

In order to quantify the goodness of fit of the model, the coefficient of determination (r^2) was computed, and this is reported in Table 4.2 in Appendix, for each subject, each condition (4 coherence levels x 3 shift conditions).

Another important characteristic of the model that we propose, is that it allows the shift of the reference line for each trial type to be estimated. It is therefore possible to account for the bias tied to the change in angular position of the oriented line between the time of stimulus presentation and that of the response. Such values are reported in Table 4.3 in Appendix, for each subject, each coherence level tested, and each reference line

condition. From these data it is also possible to compute the estimated mean (± 1 s.e.m.) directions for the -6°, 0°, and +6° condition, respectively: -3.48° (±0.054°), -0.37° (±0.045°), and 3.49° (±0.069°). A repeated-measures two-way ANOVA with these values showed a significant main effect of shift condition ($F_{2,8}$ = 58.731, p < 0.0001), but no significant main effect of coherence level ($F_{3,12}$ = 0.776, p = 0.5294). The interaction between shift condition and coherence level was not significant ($F_{6,24}$, p = 0.8154). Thus, the bias observed is driven by the position of the reference line during the estimation task, but is independent of stimulus uncertainty (coherence level).

Because the main difference between the decoding model and the late, decision-related bias model, from a computational point of view, lays in an additional parameter accounting for possible changes in the reference line throughout the task sequence, a likelihood ratio test is computed in order to compare fitting performance for this set of data from the two models. In order to perform this test, each set of data (4 coherence levels x 3 shift conditions) is taken independently and passed through both models. The ratio between the likelihood of the data being under the decoding model over the late, decision-related bias model is then taken. The statistical test provides an objective means to assess whether a set of data is better described by either the decoding model (i.e., here the *null model*) or the late, decision-related bias model (*alternative model*). Results from the likelihood ratio test are given in Table 4.4 in the Appendix, for each subject, each coherence level, and each shift condition.

With only few exceptions, non-significant tests where the late, decision-related bias model did *not* outperform the decoding model, are mostly restricted to the condition in which the reference did not move during the task sequence. Indeed, in this specific circumstance, the two models can be said to be the same, thus it is not surprising that the fitting performance works equally well in both scenarios. The fundamental gain provided by the new model proposed in this chapter can be appreciated for the trials where the reference line shifted between time of stimulus presentation and that of response. Taking into account this change implies a fundamentally different process from that described by the original model: it ties the repulsive biases to the position of the reference at the time of estimation and, crucially, not at the time of stimulus presentation.

4.3 Discussion

In this Chapter I presented the rationale and implementation of the *late, decision-related bias model*. The purpose of this model is to help understanding the role of a reference line in generating perceptual biases. A property of the boundary was systematically manipulated in the experiments described in Chapter 3. Data from Experiment 1 and Experiment 3 guided the implementation of the model. Here the oriented reference line was consistently present throughout the trial, with the caveat that, in Experiment 3, its angular position could vary by ±6° from the time of the stimulus presentation to that of the response. This fundamental change

represents the key to understanding the model and timeline of processing suggested to account for the perceptual biases observed empirically. The weighting function, a component of the model representing a specific tuning mechanism involved in processing directions of motion relative to a certain reference, only applies to the internal, sensory representation of the moving stimulus, when an external reference is present. In particular, we suggest that the reference has to be present at time of estimation / response: the weighting function not only describes a tuning mechanism, but also conveys specific information about the location of the reference when the subject has to give a response.

In order to quantitatively assess the performance of both models in fitting the data for both experiments, coefficients of determination were computed. Although both models presented a good description of the estimated directions of motion for both reference *present* and reference *shift* conditions, a likelihood ratio test showed that the late, decision-related bias model outperformed the decoding model in almost all circumstances (it is worth noting that most of the instances where this was not the case, were for the 0° shift condition, i.e. estimation task in Experiment 1 and Experiment 3 were indistinguishable).

The evidence presented in this Chapter captures the fundamental importance of the reference line during the estimation task and suggests a possible interpretation of its role in biased responses. However, another factor that could potentially have an impact on the repulsion effects reported in the literature, and those in the current thesis, is the response

modality. Participants consistently estimated the perceived directions of motion by manually adjusting an oriented line presented on the screen in front of them. It would be interesting to see if perceptual reports via an alternative response modality show a similar pattern. Indeed, manual responses involve a substantial amount of voluntary control, and this could result in variability in the responses: the repulsion we measured could be the consequence of inaccuracy in motor response. Thus, to test the generality of the bias, a motor response that can still be voluntarily initiated, but requires less top-down control, would be ideal. One such example is provided by saccadic eye movements (Abrams, Meyer, & Kornblum, 1989): observers can initiate the movement, but have less direct control on properties such as speed and duration (Becker, 1989). If the same repulsive effect was observed for both manual and eye movement estimates of motion direction, then the generality of the effect would be supported. On the other hand, if this were not the case, then a motor component to the effect would have to be considered, affecting the conceptualisation of the repulsion effect as a perceptual/cognitive phenomenon. This is addressed in Chapter 5, where subjects were presented with stimuli described previously, but performed the estimation task by means of a saccadic eye movement.

CHAPTER 5

5 Oculomotor Responses and Perceptual Decision-Making

Many processes underlying perceptual decision making were originally investigated with electrophysiological methods involving animal subjects (for example Britten et al., 1992; Britten, Shadlen, Newsome, & Movshon, 1993; Newsome, 1989, 1995). The background literature on this has already been discussed in more detail in Chapter 1. In a series of studies, Newsome and colleagues (e.g., 1989) implemented a visual motion discrimination task whereby monkeys were presented with a field of randomly moving dots and had to make a *saccadic eye movement* in the direction in which they perceived a coherent motion. The specific choice of using a visual motion stimulus paired with an eye movement response was motivated by the fact that the networks involved in both processes are generally well understood (Blum & Price, 2014; Johnston & Everling, 2008; Maunsell & Newsome, 1987). Moreover, the neural circuits governing eye movements have been shown to be located in the *posterior parietal cortex* (Andersen, 1989), a wide region comprising amongst others the *lateral intraparietal area* (LIP). Visual areas MT, MST in the temporal lobe also form important components of this circuit.

Because of the proximity of areas involved in visual motion perception and saccadic eye movements, it has been suggested that perceptual decisions in this specific context could *arise* in the posterior parietal cortex.

Newsome and colleagues initially focussed on the response properties of neurons in area MT, linking them to perceptual judgements (1989, Britten et al., 1992). Trained monkeys performed a visual motion discrimination task while neuronal activity from area MT was recorded. Stimuli were adjusted so that they reflected the tuning properties of the neuron under study (i.e., direction of motion, speed, stimulus size and location were optimal for the recorded neuron). Monkeys had to maintain fixation on a circular spot while stimuli were displayed on the screen and once it was extinguished, they had to make a saccade to either of two presented targets: a *preferred target* indicated the preferred direction for the neuron being recorded, a *null target* located 180° apart from the preferred one. Data resulting from the overt discrimination task were then compared to activity recorded from MT neurons: firing rates were correlated with the coherence level of the stimulus – the proportion of dots in the stimulus moving in the same direction. Strikingly, there was similar sensitivity in the neuronal responses and psychophysical performance, as captured in neurometric and psychometric curves respectively. These observations led to the hypothesis that visual motion discrimination judgements result from activity in an ensemble of MT neurons (Britten et al., 1996). In a similar fashion and using the paradigm described above, Shadlen and colleagues (1996) investigated the link between sensory evidence and behaviour while recording activity from monkeys' area LIP. While Salzman and colleagues (1990) showed that firing activity of neurons in visual areas MT and MST of monkeys is related to the sensory processing of visual information (electrical stimulation of

directionally selective columns in area MT resulted in judgements biased towards the direction preferred by the stimulated neurons compared to judgements given after no stimulation was applied), Shadlen et al. (1996) showed that motor responses (here *saccadic eye movements*) are processed in more frontal areas known as *Frontal Eye Fields* and the *Superior Colliculus* (FEF and SC, respectively –Glimcher, 2001). Thus, area LIP, located between and anatomically connected to sensory and motor planning areas (Andersen, Brotchie, & Mazzoni, 1992), seems like a good candidate in the decision process, linking sensory evidence to behavioural response through a choice-dependent action planning. Interestingly, it was observed that the neuronal activity recorded in area LIP could predict the saccadic response to be performed by the monkey.

The relationship between choice and LIP firing rate also emerged from trials in which the visual stimulus presented incoherent motion: despite being presented with the same physical stimulus and responding at chance level, the neuronal activity recorded in LIP predicted the saccadic response of the monkey. More specifically, the firing rate of LIP neurons gradually increased during the display of the moving stimulus, the rate of increment being correlated to the proportion of direction signal present in the stimulus, and once a threshold is reached, the eye movement is elicited (Shadlen & Newsome, 1996). This process agrees well with the accumulation of evidence model described in Chapter 1, Section 1.1 and is also consistent with the results of a study by Hanks, Ditterich, & Shadlen (2006). Here, monkeys performed the same visual motion direction

discrimination task as described in the previous studies, but were allowed to respond via a saccadic eye movement to the chosen visual target at any point following stimulus onset. Microstimulation was applied to an ensemble of LIP neurons in half of the trials during stimulus presentation. This did not directly elicit a saccadic response, but affected both choice and reaction time for that trial: an increased choice for the target located within the stimulated response field for weak signal stimuli was observed.

Similarly, reaction times for the target inside the stimulated response field were faster than those toward the other target. Moreover, stimulation of LIP neurons with a RF tuned to one target resulted in slower reaction times whenever the opposite target was chosen. Thus, the authors concluded that neurons in area LIP are involved in accumulating evidence during a perceptual decision task: stimulating a preferred target location reduces the time needed to reach a threshold that will then elicit the appropriate motor response.

The relationship between sensory representation and perceptual decision expressed through a saccadic eye movement has been investigated in human subjects as well. For example, McSorley & McCloy (2009) used a visual motion discrimination task similar to that by Newsome and colleagues and studied the effect of signal strength (i.e., proportion of dots moving coherently across the stimulus aperture) upon saccadic response times. They measured saccade *latency*, *trajectory*, and *landing position* and correlated them with the target chosen for the saccadic response. Overall, trajectory and landing position were more accurate the

higher the coherence level of the moving stimulus. Saccade latency, the time between stimulus onset and saccade initiation, decreased as a function of stimulus coherence: the stronger the sensory evidence, the faster the response. These results are in broad agreement with those found with monkey electrophysiology and, importantly, they also hint at the precision of this type of measurement: saccades are fast, voluntary eye movements, with low levels of variability, mostly explained by uncertainty in external factors (e.g., target location) rather than in the motor response per se (see van Beers, 2007).

Given the high precision of this motor response and the extensive use of the paradigm in the perceptual decision making literature, an adapted version of the behavioural task was implemented in the current study. The behavioural experiments that form the basis of this chapter are described in Chapter 3, Section 3.1.1.3. The aim here was to investigate the role of response modality in the perceptual bias observed when subjects have to manually estimate the perceived direction of motion of a visual stimulus. Subjects were presented with an oriented reference line and an RDK with variable coherence levels and had to perform a saccade indicating the perceived direction of motion once the fixation point at the centre of the screen disappeared.

5.1.1 Materials and Methods

5.1.1.1 Observers

Four subjects participated in this study (three females; age 25.25 ± 1.89 years, mean ± s.d.). Two out of the four took also part in Experiment 3 presented in Chapter 3 (section 3.3). All had normal or corrected-to-normal visual acuity.

5.1.1.2 Apparatus and Stimuli

A detailed description is provided in Section 2.2.3 and 2.2.3.1 of Chapter 2. It is important to notice that the stimuli for this experiment were computed at four coherence levels: 7, 15, 25, and 50%. Two of the four values differ from those used in the experiments described in Chapter 3 (i.e., 4, 7, 13, 25% coherence levels were used). The choice of different values was guided by the results obtained by both Jazayeri & Movshon (2007a) and Experiment 1 of Chapter 3. In particular, Jazayeri & Movshon relate the repulsion effect to the uncertainty of the motion signal: the lower the strength, the stronger the effect. However, in Experiment 1 (Section 3.1) the repulsion effect was still evident in stimuli with a coherence level of 25% (twice the highest coherence used in the study by Jazayeri & Movshon). For the eye-tracking version presented here, we therefore decided to test an even higher level (i.e., 50%) to investigate the impact of the sensory evidence reliability, together with levels comparable to those used in the manual response task.

5.1.1.3 Procedure

The experimental procedure is described in Chapter 2, Section 2.2.3.2. Briefly, subjects underwent extensive practice prior to performing the experiment. Training consisted of a minimum of 600 trials at a suprathreshold coherence level (coherence, 80%), where observers were asked to estimate the perceived direction of coherent motion in the RDK stimulus with a saccadic eye movement. Once the observer felt comfortable with the task, a minimum of 1260 trials were performed, spread over 30 experimental runs.

In a particular trial, participants had to fixate the centre of the screen while presented with an oriented reference line, followed by an RDK. The latter was presented at a specific motion coherence level (blocked within run) of either 7, 15, 25, or 50%, and the signal direction was randomly selected between ±20° centred on the reference line. The presentation of a black circular aperture, together with the extinction of the fixation point, cued the subjects to make a response: saccade in the direction of perceived coherent motion. The reference line was displayed throughout the trial. This particular condition was chosen, because it led to *biased* manual responses in the previous psychophysical experiment. It is important to address the different estimation times allowed for the two conditions (i.e., manual adjustment described in Chapter 3 and saccadic eye movement). In particular, when subjects were asked to manually estimate the perceived direction of motion by aligning a segment on the screen using a cursor, they were allowed up to 10s. Here, the next trial started 750ms after response

cue was presented. The difference in timings is remarkable, however, an analysis of response times across the three experiments described in Chapter 3 revealed that subjects gave their response within 1.53s on average. Figures 10.5 and 10.6 (data from Experiment 3, Chapter 3, only) in the Appendix show the frequency distributions of time elapsed between presentation of estimation cue (i.e., black circular ring displayed on the screen) and the moment when the response was given (i.e., the segment line adjusted by the subject hit the circular ring). The distributions are skewed to the left, with median response time of 1.39s and mean 1.53s (data from all experiments), resulting in 70.15% of trials with a response time below the average (median response time of 1.40s, mean of 1.52s for data from Experiment 3, only, with 69.33% of trials presenting faster responses than the average). This shows that, despite allowing 10s for an estimate to be provided in the manual task, subjects responded quicker, sensibly reducing the difference in response times highlighted in the paradigm between the manual and saccadic movement.

5.1.2 Results

In Section 5.1.2.1, general statistics of saccadic eye movements are presented. Section 5.1.2.2 provides results related to the research question: is the repulsive effect in estimates of motion direction evident when responses are provided via a voluntary eye movement?

5.1.2.1 Dynamics of Saccadic Eye Movements

Parameters of the main sequence (Bahill, Clark, & Stark, 1975) were assessed, for each participant in two ways. First, eye movement data were pooled across all coherence levels tested for that observer. Second, eye movement traces were analysed separately for each coherence level tested. Figure 5.1A shows mean amplitude of saccadic eye movements (in degrees) as a function of stimulus direction (in degrees). Despite the fact that accuracy in saccadic eye movement direction, rather than amplitude (i.e., landing position on the border of the circular aperture), was emphasised when performing the task, a significant difference in saccade amplitudes for different stimulus directions was observed for all participants. Interestingly, amplitudes for eye movements to the right of vertical were consistently smaller than those for eye movements directed to the left of vertical. Another interesting pattern emerged when taking into account the coherence level of the stimulus: for subjects tested with a 15% coherent stimulus (subjects b and g), saccade amplitudes in this condition were bigger than for any other level (Figure 5.1B). No effect was found for subjects tested with signal strengths other than the 15%.

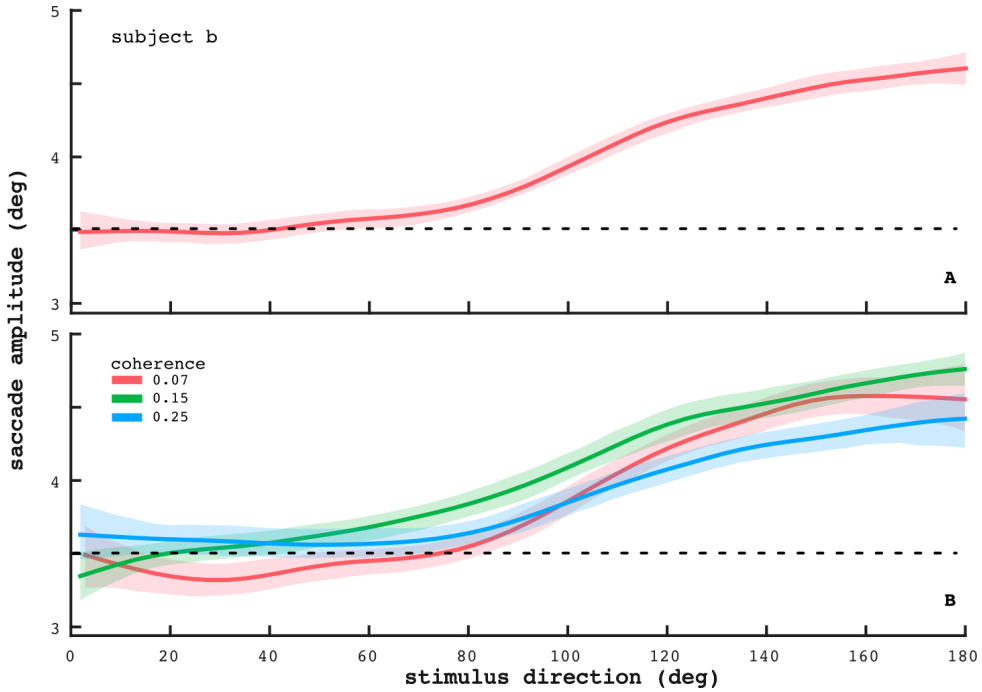


Figure 5.1 Saccadic amplitude as a function of stimulus direction. Data shown for one subject. A) Mean saccadic amplitude (red line) and 95% CI (shaded area), black dashed line indicates the location of the circular aperture (target landing location). B) Mean saccadic amplitude (solid lines, colours indicate different coherence levels of the stimulus), shaded areas and dashed line as described in A). Amplitudes to the right of vertical (between 0 and 90° along the x axis) are smaller than those to the left of vertical (between 90 and 180° along the x axis).

Figure 5.2A shows mean saccade duration (ms) as a function of stimulus direction: two of the observers (subjects b and h) tested showed shorter saccade durations for movements directed to the left along the horizontal axis. This effect was reversed for one of the subjects (subject i) and absent in the other subject tested (subject g) so was not a consistent characteristic of the data set. When looking at the effect of stimulus strength on saccade duration, two participants (subjects h and i) performed faster saccadic eye movements when presented with a high signal strength stimulus (50% coherence), and the same pattern was true for observer b when tested with a stimulus of 15% coherence.

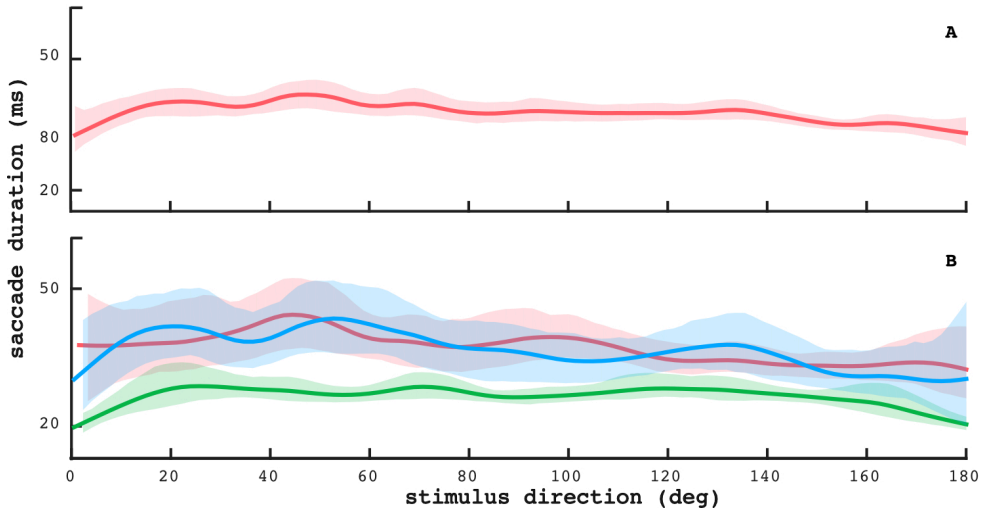


Figure 5.2 Saccadic duration as a function of stimulus direction. Data from one subject. A) Mean duration (red line) and 95% CI (shaded area): the distribution is relatively flat. B) Mean duration (solid lines, colours indicate different coherence levels of the stimulus), shaded area as described in A): distributions are not affected by stimulus direction. Durations of saccades for stimuli computed at 15% coherence are significantly faster than when these are computed at either 7 or 25% coherence.

Peak velocity of saccadic eye movements showed a pattern consistent with the results provided by the analysis of amplitude and duration: higher velocity for saccades towards the left of vertical compared to saccades towards the right (0 to 90° on the x axis, Figure 5.3A). This was indeed expected: larger amplitudes over the same time (duration) of smaller ones require faster velocity. The same pattern was apparent when the data were split according to the stimulus strength (Figure 5.3B).

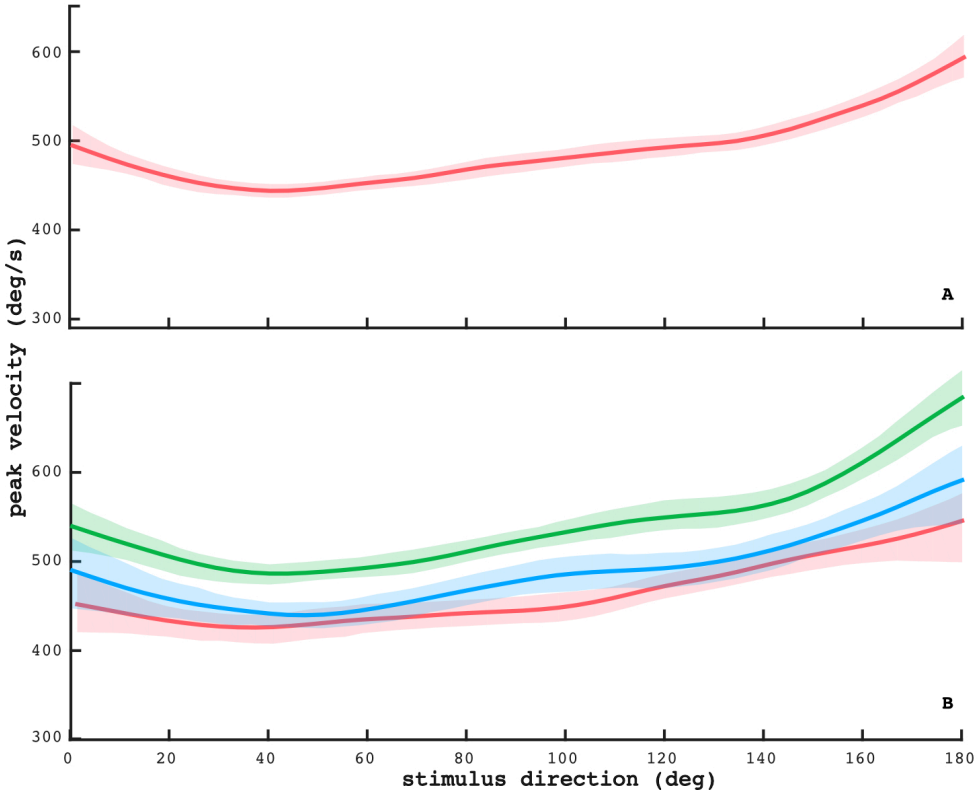


Figure 5.3 Saccadic velocity as a function of stimulus direction. Data from one subject. A) Mean velocity (red line) and 95% CI (shaded area): velocity is faster for eye movements to the left of vertical (90 to 180° along the x axis), consistent with bigger amplitudes over a constant amount of time. B) Mean velocity (solid lines, colours indicate different coherence levels of the stimulus) and shaded areas as described in A). The same pattern observed in A) can be appreciated.

Variability of saccadic endpoints was investigated by fitting confidence ellipses (see Section 2.2.3.3, Chapter 2 for more details) around the data. First, given that the directions of motion of the RDKs spanned the range between 0 and 180 degrees uniformly, six bins with a width of 40° (thus including also directions ±20° from the horizontal meridian), were computed. Within each bin, the angular distance between the saccadic response and the actual stimulus direction, for each individual trial, was calculated and plotted as a function of true direction (dot symbols in Figure 5.4). The 95% confidence ellipses for each direction bin provided a picture of the variance of each distribution. (Note the major and minor axes of the ellipse) This was used to test whether variability in saccade endpoints is

affected by the direction in which the eye movement is performed. Although the Levene's test used to check for homoscedasticity in the distributions of estimated directions in each direction bin did not result in a clear pattern across individuals, two subjects showed less variability for saccades performed along the horizontal axis compared to any other direction (subject b: $F(5,1606) = 9.944$, $p < 0.0001$; subject g: $F(5,999) = 3.12$, $p = 0.015$). Subject i showed the reverse pattern, with less variability for vertical saccades compared to any other direction ($F(5,870) = 2.635$, $p = 0.022$) and for subject h distributions of estimates resulted to have equal variances ($F(5,993) = 0.8569$, $p = 0.51$).

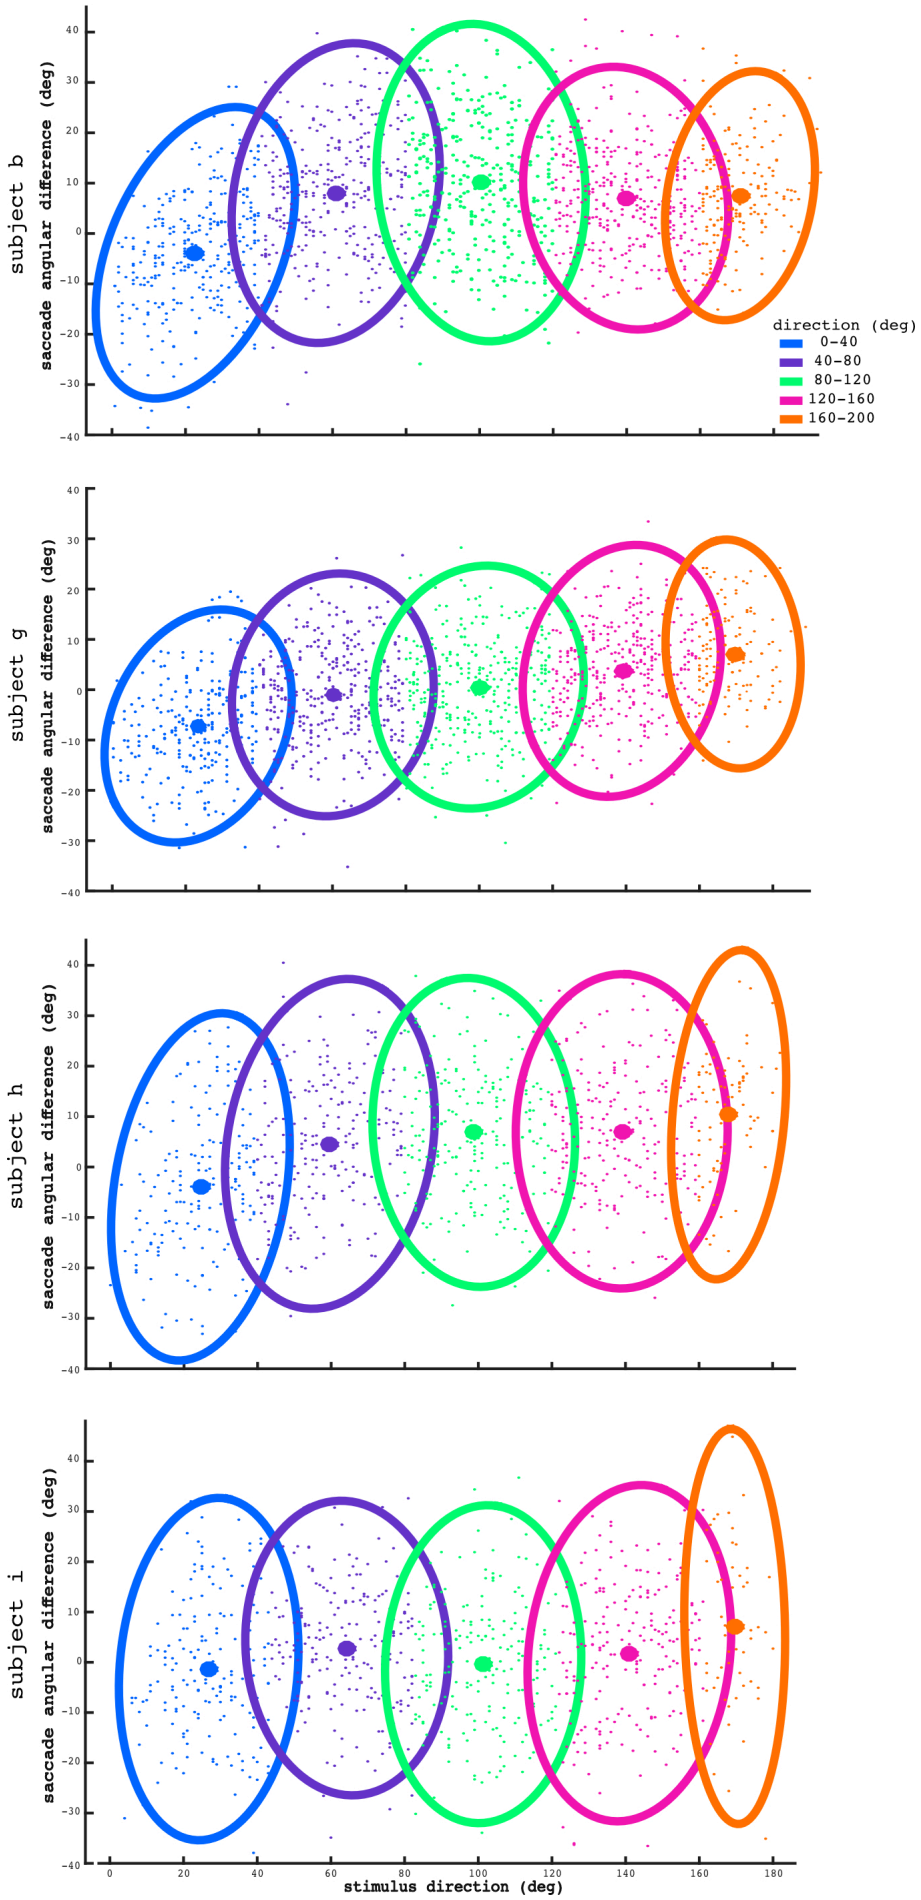

subject b
subject g
subject h
subject i
saccade angular difference (deg)
stimulus direction (deg)
direction (deg)
0–40
40–80
80–120
120–160
160–200

Figure 5.4 Variability of saccadic endpoints. Each row shows data from one subject. Colours indicate different bins of stimulus direction, with small points showing individual trials, large dot symbols represent the centre of the 95% confidence ellipse delimited by solid line. 0° along the x axis indicates rightward motion, 90° indicates upward motion, 180° leftward motion. Along the y-axis: 0° indicates veridicality of saccadic response (estimate of motion direction equals stimulus motion direction), positive values represent saccades more CCW than stimulus direction.

5.1.2.2 *Saccadic Estimates of Motion Direction*

Following the analysis adopted in Chapter 3 (see Section 2.1.1.5 for details), we looked at estimates of motion direction as a function of the angular difference between the reference line and the stimulus. Figure 5.5 shows data from one subject: each symbol represents the relative error in saccadic response as a function of relative direction of motion of the stimulus. The solid line and shaded area in each scatterplot shows the mean and 95% confidence interval, respectively, while the histogram in the top right corner shows the distribution of estimated directions of motion in the stimulus range between -6 and 6 degrees (relative to the oriented reference line).

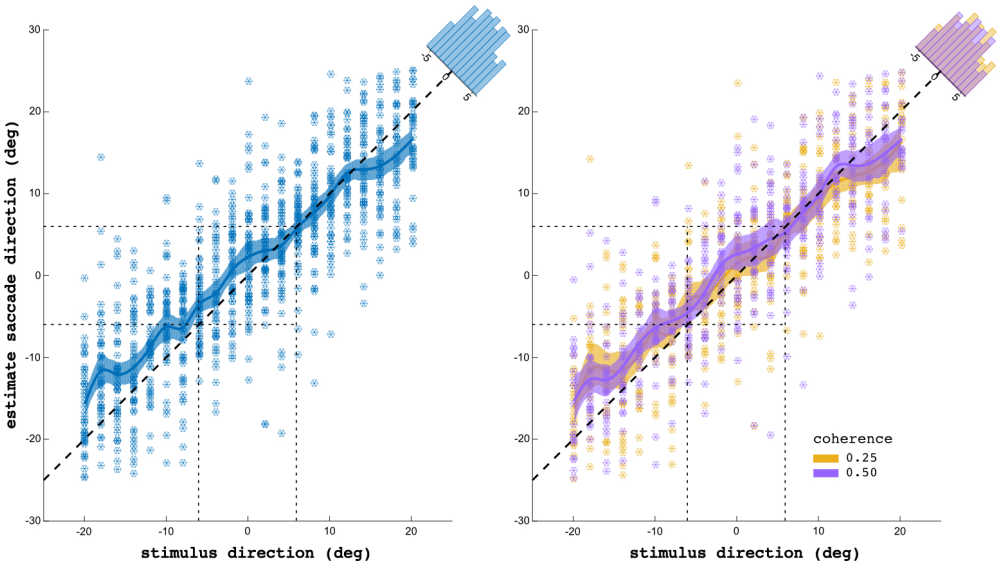


Figure 5.5 Estimates as a function of stimulus direction. On the left: all data from one subject are pooled together and presented as a scatter plot. Thick solid line indicates mean estimate in a column-wise manner and shaded area shows 95% CI. On the right: data from the same subject are split according to the coherence level of the stimulus (indicated by different colours). Dash black line indicates identity (veridical responses lay along this diagonal). Selection indicated by dash lines on top of each plot aid visualisation of the interval across which the corner histograms on the top right of each figure are computed. Overall, no biases (deviations from identity line) are observed.

The corner histograms, shown in the top right corner, clearly illustrate that saccadic responses did not show a repulsive effect. A repulsive effect as seen in the behavioural task would have led to deviations from the unimodal pattern seen here. This is even more evident when looking at Figure 5.6: the joint histogram (described in Section 2.1.1.5, Chapter 2) represents the same data as in Figure 5.5 but this time each column is displayed as a density distribution and the lightness scale indicates the percentage of trials for each distribution (high intensity for high values). Red symbols represent the peaks of each distribution (procedure also described in Section 2.1.1.5, Chapter 2). Responses, therefore, were unbiased across different signal strengths and the distribution of saccadic responses were centred around the identity line (dashed white line representing veridical responses).

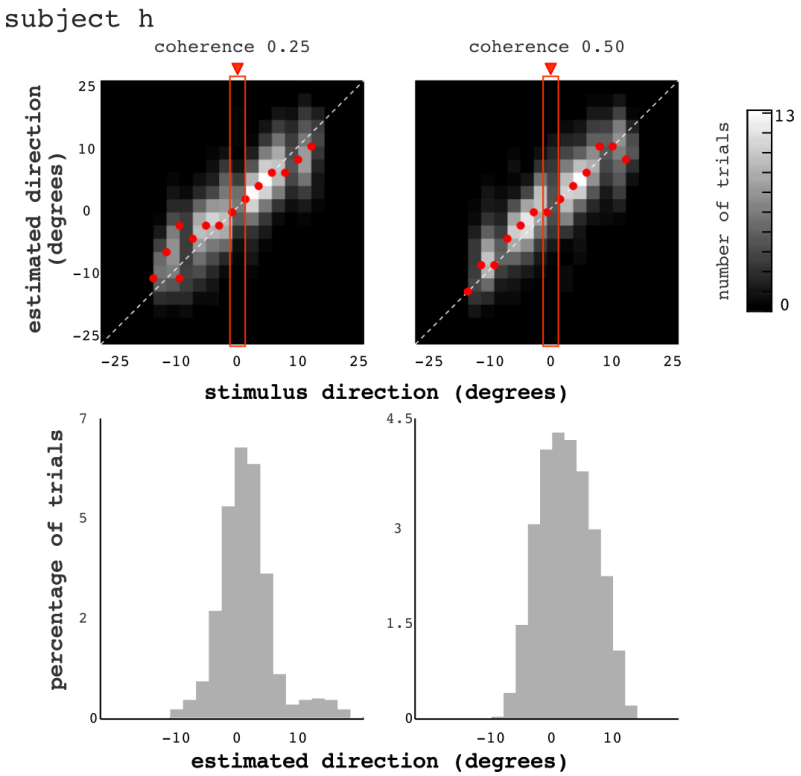
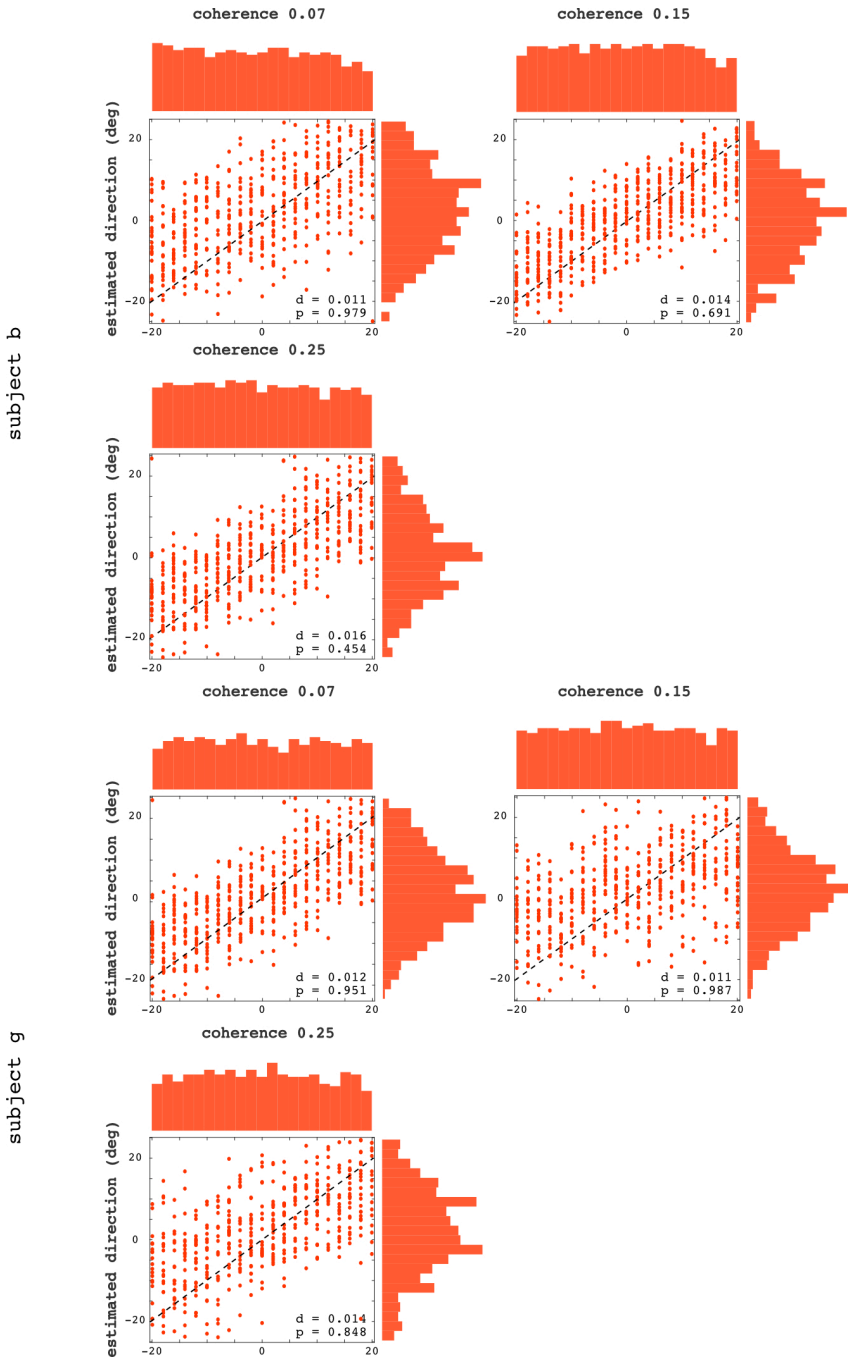


Figure 5.6 Data for one subject, two coherence levels. Density plots in row one display the distribution of saccadic estimates as a function of stimulus direction. Red dots indicate peaks of column-wise distributions for each direction presented. Histograms in the second row represent the column-wise distribution highlighted by the red rectangle in the density plots of row one.

Dip-statistic values were computed, for each subject and each coherence level adopted, in order to test for deviations from unimodality (see insertions of each panel in Figure 5.7). This confirmed the absence of biases (in terms of repulsive effects) in estimates of direction of motion provided via a saccadic eye movement (Figure 5.7).


subject b
coherence 0.07
coherence 0.15
coherence 0.25
estimated direction (deg)
d = 0.011
p = 0.979
d = 0.014
p = 0.691
d = 0.016
p = 0.454
subject g
coherence 0.07
coherence 0.15
coherence 0.25
estimated direction (deg)
d = 0.012
p = 0.951
d = 0.011
p = 0.987
d = 0.014
p = 0.848

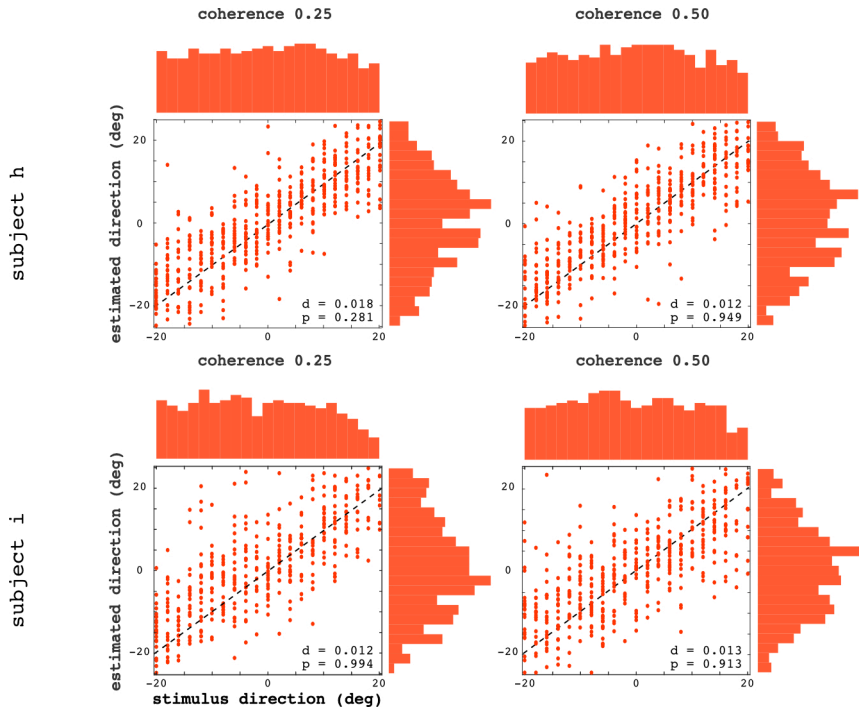


Figure 5.7 Distribution of estimated directions as a function of stimulus direction. Rows one and two present data for subject b, rows three and four data for subject g, the following rows present data for individual subjects. Each column presents data from a different coherence level. Marginal histograms along the y-axis show unimodality in the distribution of estimates, for all subjects, all coherence levels tested. Dip-statistic values are indicated at the bottom-right corner of each panel.

5.1.3 Discussion

The study presented in this chapter aimed at investigating whether the repulsive effect in estimates of motion direction observed when subjects respond to a moving stimulus is also present when such estimates are given via an eye movement, rather than via a manual adjustment of a dial. This was assessed by asking observers to indicate, via a saccadic eye movement, the coherent direction of motion present within an RDK.

Results from evaluating standard quantitative metrics of eye movements generally agreed with findings reported in the literature: saccade amplitude generally agreed with the amplitude of the circular aperture displayed at the time of response, however a marked difference

between saccade amplitudes to the right and left of vertical was observed. Amplitudes were bigger, and less accurate, for eye movements directed to the left compared to those directed to the right of the vertical axis. A possible explanation of this finding relates to the way in which data are treated during the analysis: despite tracking both eyes during task performance, only gaze positions from the left eye are then used in the analysis. Thus, points on the left visual field will fall onto the nasal retinal of the left eye (and temporal retina of the right eye). It is known that the density of photoreceptors across the retina varies, with higher density for nasal, compared to temporal portions (Curcio et al., 1990; Østerberg, 1935). Advantages dependent on this asymmetric distribution of receptors have been observed for stimuli presented far in the periphery of the visual field for visual tasks such as the Vernier-acuity test (Harris & Fahle, 1996), contrast discrimination and lexical decision judgements (Lavidor, Alexander, & McGraw, 2009). Because our stimuli were not presented in the periphery of the visual field, the difference in saccade amplitudes is unlikely to derive from this asymmetry. However, Lavidor and colleagues (2009) reported a naso-temporal asymmetry (NTA) effect in a visual word recognition task: words presented to the temporal hemifield of the left eye, and therefore encoded in the nasal portion of the left eye's retina, were identified more easily (lower threshold) than words presented to the right hemifield of the same eye. This effect, however, was found for subjects whose native language was Hebrew (reading direction right to left), but not for English readers (reading direction left to right). Although the stimuli

were still presented in the periphery of the visual field, the fact that the effect was present for Hebrew readers and not for English readers, suggests that the reading habits (i.e., reading direction) may affect word recognition. Thus, a possible explanation for the left-right anisotropy observed in saccadic amplitudes could be related to the fact that during reading most western languages, observers make a series of saccades toward the right and the accuracy of the landing points is important to effectively encode the next word (these eye movements are usually directed towards the centre of a word, see for example McConkie et al., 1988). This 'reading habits' dependent effect could be tested in two ways: by recruiting subjects whose native language is encoded from right to left and test them with the paradigm used in this Chapter, expecting an inversion in the amplitude asymmetry reported for English readers that participated in this study. Alternatively, one could adapt the current paradigm to test children who have yet to learn to read. In this scenario we would not expect to see a left/right anisotropy for saccadic amplitudes, given that at this early stage eye movement patterns could potentially not be yet biased towards one specific direction (however, see Simion et al., 1998).

Moreover, a signal strength in the stimulus of 15% produced bigger saccade amplitudes than any other strengths tested. Saccade duration changed with the direction of the movement: saccades directed to the left along the horizontal axis were briefer than saccades in any other directions. Moreover, saccade duration was shorter when the signal strength of the moving stimulus was high (e.g., 50%). Saccade velocity conformed with the

expected pattern given the findings from eye movement amplitudes and duration: saccades were faster for bigger amplitudes. The analysis of the variability of saccadic endpoints did not result in a systematic pattern across subjects. While two observers presented more accurate eye movements along the horizontal axis, another observer showed more accurate saccade endpoints for vertical directions, and the responses of the fourth subject were unaffected by the direction towards which the saccade was performed. Despite the lack of consensus in the findings, the patterns observed generally agree with previous studies showing that endpoints of saccadic eye movements are less variable for cardinal directions (horizontal and vertical axes of motion), compared to oblique directions (see van Beers, 2007). Given that horizontal, vertical, and oblique eye movements are governed by a different combination of six muscles (medial and lateral rectus muscles for horizontal eye movements with the addition of superior / inferior rectus and oblique muscles for movements that deviate from horizontal), and that different neuronal circuitries process information for horizontal and vertical gaze (see for example Sparks, 2002), one could speculate in favour of a physiological or anatomical explanation to this discrepancy in accuracy of saccadic eye movements between cardinal and oblique axes. In particular, while the commands generating horizontal voluntary eye movements can be traced from regions of the FEF (Frontal Eye Fields) to a portion of the brainstem known as *paramedian pontine reticular formation* (PPRF), either directly or indirectly (via the superior colliculus, regions of the pons and cerebellum), the pathway for vertical and

oblique eye movements is more diffuse (Leigh & Zee, 2013). The systematic control of a larger set of muscles, together with an increased sparsity in the representation / processing of the motor command for voluntary eye movements that deviate from the horizontal / vertical axes may contribute towards the increased variability in endpoints observed in our analysis.

The results of the current study are important in that they show no evidence of a reference repulsion effect: the angular direction of saccade endpoints overall agrees with that of the presented stimulus. This suggests that the response modality plays a key role in the repulsion effect observed in the studies presented in Chapter 3. The findings, therefore, support the idea that the repulsion is due, in part, to a motor component: a manual response involves a high degree of voluntary control (i.e., characteristics such as speed and duration of movement can be overtly controlled by the observer), whereas this is much reduced in the case of eye movements (saccade speed and durations cannot be adjusted when performing the action, see Becker, 1989).

However, it is important to note that for the study presented in this Chapter, observers were allowed a maximum response time of 750ms. This is a relatively short time interval and was chosen in order to control for the amount of corrective saccades that could follow the first response. Experiments that required a manual response, on the other hand, allowed for a maximum of 10s for an estimate to be given. Therefore, it is possible that the repulsion effect (bias) is not instantaneous but evolves over the relatively long time course of the manual adjustment procedure. A way to

address this possibility would be to continuously track the adjustment of the estimate from response onset until the end of the trial. On the other hand, looking at saccades that shortly followed the one indicating the response could provide information on successive adjustments in the judgement.

Also, relating to the variability in motor response and amount of voluntary control over a particular action, one could speculate that a component of error might derive from translating visual information to a different space where a hand motor plan has to be formulated and accordingly acted upon. For eye movements specifically, studies have shown that a motor command is encoded by ocular motoneurons of the cranial nerves in a coordinate frame that is relative to the position of the head (i.e., *craniotopic coordinates*), and determined by the temporal dynamics (duration and frequency) of discharged spikes (Leigh & Zee, 2013). However, a visual stimulus (e.g., the saccadic target) is encoded by the visual cortex in terms of neuronal activity patterns that are retinotopically organised. Therefore, for a voluntary eye movement towards a visual stimulus to be performed, a transformation of the sensory information from retinotopic to craniotopic coordinates is necessary: the space-encoded stimulus needs to be mapped onto the temporally coded one. Areas such as the frontal and posterior eye fields have been appointed as involved in this transformation (see for example Lynch, 1992), whereas in the superior colliculus, motor components encoded in the ventral layers still present a retinotopic coordinate framework (Optican, 2005). It is therefore

reasonable to think that a similar remapping of *place-coded* (retinotopic) to *motor-coded* information happens when performing a manual response.

The difference between the two types of transformations could arise from the fact that, in the voluntary eye movement case, the sensory system is shared between perception and response. The stimulus is encoded in the visual cortex and the estimate requires to move the eyes in a sort of 'visually guided response', whereas in the experiments presented in Chapter 3, a transformation across sensory systems was required. In this case, the manual response was not 'visually guided': subjects could monitor the alignment of the bar on the screen while estimating the direction of motion, but could not look at the changes in the manual movements causing the displacements of the line on the screen. It seems that manual responses require extra layers of complexity and transformations across coordinate frames that could affect the bias we measured.

Estimates of motion direction via saccadic eye movements were essentially veridical for all coherence levels tested. It is true that the signal strengths used in the present study were slightly higher than for the previous experiment. For the saccade response experiment the maximum coherence was 50% compared to a maximum of 25% for manual version presented in Chapter 3. However, for coherence levels of 15 (close to 13% used in Chapter 3) and 7% the repulsive effect was still absent. This suggests that the uncertainty of the stimulus does not affect the final response, at least when this is given via a fast, voluntarily-initiated, eye movement.

These results challenge the idea that the repulsive effect is independent of response modality. A better understanding of the elements at play here would be given by testing a wider and more finely sampled range of coherence levels, as well as investigating corrective saccades when allowing longer response intervals in the trial.

CHAPTER 6

6 *Imaging Mechanisms Involved in Fine and Coarse Discrimination Tasks*

In Chapters 3 to 5 I have shown that a perceptual bias known as reference repulsion is a phenomenon manifesting itself at the time of indicating a decision based on sensory evidence and that the modality with which the response is provided also affects the resulting error. However, the findings presented do not rule out the fact that different sensory mechanisms are implicated in discriminating between two highly similar versus two highly dissimilar stimuli. The work presented in this and the following Chapter (7) addresses this point: is it possible to model different patterns of responses from early visual cortex when engaging in a fine versus coarse discrimination task?

As discussed in Chapters 1 and 3, when discriminating between two highly dissimilar stimuli (e.g., motion upwards versus downwards), the optimal strategy to maximise performance is to base the decision on those neurons tuned to the feature under investigation. On the other hand, a better performance in discriminating between stimuli closely matching each other is driven by activity of neurons presenting a higher differential response between the two alternatives, that is neurons tuned away from the specific value of the stimulus feature under study (Beck et al., 2008; Hol & Treue, 2001; Jazayeri & Movshon, 2007a; Purushothaman & Bradley, 2005; Regan

& Beverley, 1985). Scolari & Serences (2010) investigated the relationship between the activity of so-called *off-channels* (i.e., neurons tuned away from a critical feature value) in early visual cortex and discrimination performance using an orientation discrimination task. Subjects were presented with a *sample* (reference) oriented grating while fixating at the centre of the screen and had to judge whether the *test* (comparison) grating was rotated CW / CCW from the reference while brain activity was measured using functional magnetic resonance imaging (fMRI). Once the authors identified area V1, responses in each voxel of this region of interest (ROI) were used to compute a tuning function: a voxel was tuned to the orientation that evoked the greatest response. By re-computing the tuning functions of each voxel in terms of offset from the orientations of the sample stimulus (i.e., the preferred orientation of a voxel is now centred at 0° and the other possible orientations of the sample stimulus are defined relative to this 0 point), Scolari and colleague predicted that a higher response should be expected in voxels tuned away from the orientation of the sample stimulus for correct trials. This is equivalent to the idea of observing a gain in off-channel neurons responses to a stimulus orientation different from the neuron's preferred one. Indeed, they reported higher responses on correct trials in voxels that were tuned 36° away from the orientation of the sample stimulus. These results seem therefore to support the idea of an off-channel mechanism underlying perceptual decisions based on discrimination of similar stimuli. It is important to note here that the authors focussed their investigation on responses in area V1. This is because primary visual cortex

presents a high number of orientation selective neurons (Hubel & Wiesel, 1962), thus making it the most likely region to observe this tuning mechanism.

The aim of the study presented in this Chapter is therefore to test for a similar mechanism as described by Scolari & Serences (2010) using a different visual feature – directions of motion. The advantage of this lays in the possibility of also looking for differences in feature processing between area V1 and MT, both presenting tuning for visual motion. Moreover, two different approaches were taken in addressing this research question: first, and described in this Chapter, differences between fine and coarse discrimination processes were investigated at a univariate level. By applying a version of the Generalized Linear Model (see Chapter 2, Section 2.3.6.1) we looked at modulations in voxels' responses due to performing either of the two discriminations across the whole brain. This is because, given that the two tasks are different, it is possible that top-down effects (e.g., different levels of attentional demands between the two tasks) influence the gain of neurons in early sensory areas (Scolari, Byers & Serences, 2012). The second approach, described in Chapter 7, considers the question in a multivariate framework, where simultaneous changes across multiple voxels are taken into account.

6.1.1 Materials and Methods

6.1.1.1 Observers

Five subjects participated in this study (three females; age 25 ± 3.81 years, mean ± s.d.). Three of them took part in other experiments described in this thesis and the remaining two were naïve to the purpose of the study. All had normal or corrected-to-normal visual acuity, gave written informed consent previous to scanning and procedures were approved by the Medical School Research Ethics Committee at the University of Nottingham. Data from one of the subjects was discarded due to motion during scanning procedure.

6.1.1.2 Apparatus and Stimuli

A detailed description of Apparatus and Stimuli is given in Section 2.3.7 (subsection 2.3.7.2) in Chapter 2. Only data from the discrimination task are considered here.

6.1.1.3 Procedure

A schematic representation of the experimental procedure is presented in Figure 6.1. Briefly, prior to starting the scanning session, observers were informed as to which task to perform (fine / coarse discrimination) and *reference* and *comparison* stimuli were identified (left/right reference, right/left comparison). A trial started with the presentation of a white dot that subjects had to fixate for the whole duration of the scan. The two circular apertures representing the reference and comparison stimuli were

presented simultaneously, one on each side of the fixation point. Once the stimuli were extinguished, participants had to indicate whether the comparison stimulus moved CW/CCW (fine discrimination task) or in the same/opposite (coarse discrimination task) direction as that of the reference stimulus. Visual feedback was provided after each trial by changing the colour of the fixation point. Each participant performed two scans for each task during the same session.

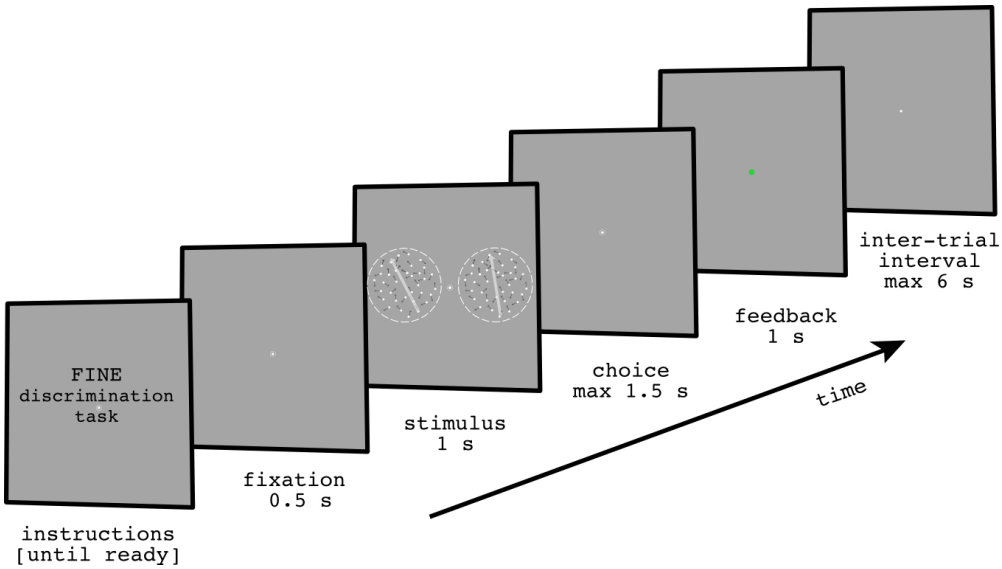


Figure 6.1 Schematic representation of a trial sequence. Subjects are presented with two circular apertures, one (left) contains the reference stimulus (light grey arrows on top of RDKs indicate overall direction – for display purposes only). Subjects indicated whether the stimulus on the right moved CW/CCW (fine) or in the same/opposite (coarse) direction of the reference stimulus. Feedback was provided by changing colour of the fixation point (green for correct, red for incorrect).

6.1.1.4 Data Preprocessing and Analysis

In order to minimise artefacts induced by head motion during the scanning session, data underwent motion correction (see Chapter 2, Section 2.3.5.1). By examining the resulting matrix indicating the amount of translation, rotation, and shear during each scan, subject l was excluded: clear big deviations in the traces in Figure 6.2 indicate that the participant

moved to a degree that could not be accounted for by the preprocessing stage. Motion correction matrices for the remaining four observers were reasonable and data were therefore used in the following analyses.

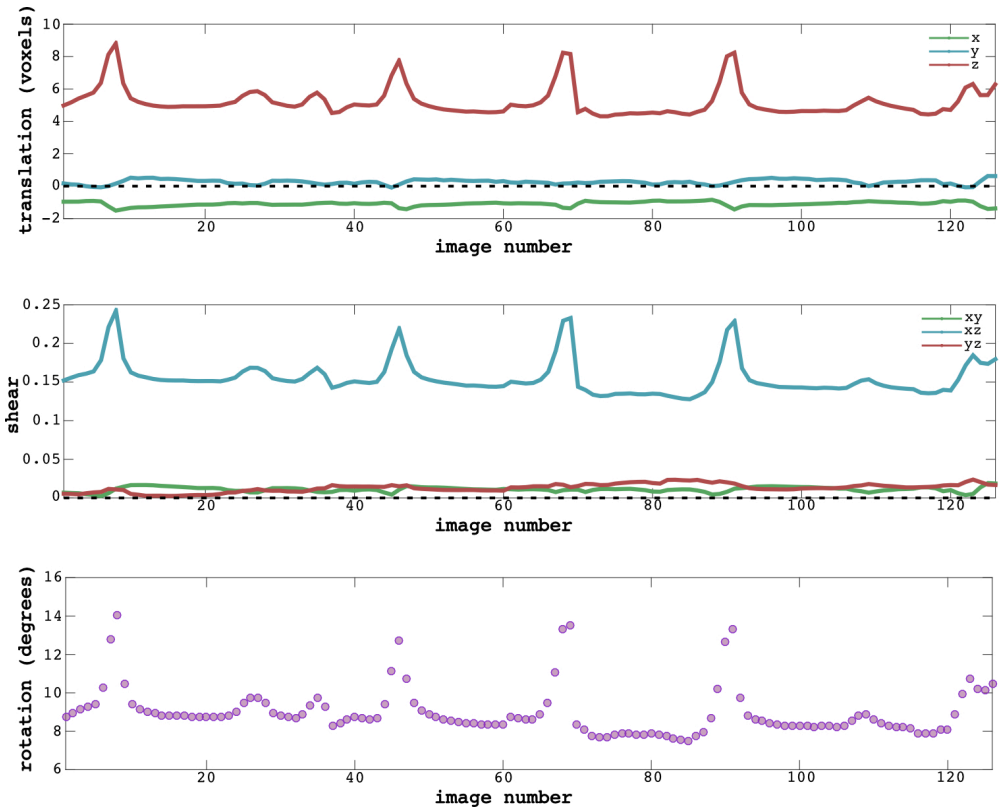


Figure 6.2 Motion correction parameters for the excluded subject. From the top: translation along the main axes, shear motion, and rotation parameters plotted as a function of time.

Generalised Linear Model was applied, performing four different analyses with the following factors: key-press, investigating modulations of voxel responses due to motor response; coherence level, looking at modulations of voxels in the visual cortex due to the different strengths of the motion signal in the stimuli; task type, looking at differences in task performance (i.e., ensuring that the difficulty of the tasks was balanced).

6.1.2 Results

This section is divided into two parts. First, results from two control analyses are presented: did subjects perform the tasks (Section 6.1.2.1) and can we measure a response modulation from voxels in the visual cortex due to the coherence level of the stimuli (Section 6.1.2.2). The second part looks at differences in voxel responses due to the type of task performed (fine / coarse) across the whole brain, as well as in striate and extrastriate visual areas (Sections 6.1.2.3 and Section 6.1.2.4, respectively).

6.1.2.1 Control Analysis: Key-Press Related Activity

As a simple reality check of the data acquired, a GLM analysis was performed, for each subject separately, investigating modulations of the timeseries due to the key-press event (i.e., the action indicating the choice made by the participant during the response interval of the trial). Figure 6.3 shows the results of this analysis for one subject: an increase in activity (Figure 6.3A) in a region anterior to the central sulcus (motor cortex, M1, Figure 6.3B) indicates activity related to finger-pressing. This was observed in all subjects, located in the hemisphere contralateral to the hand used to give the response (this was counterbalanced across participants).

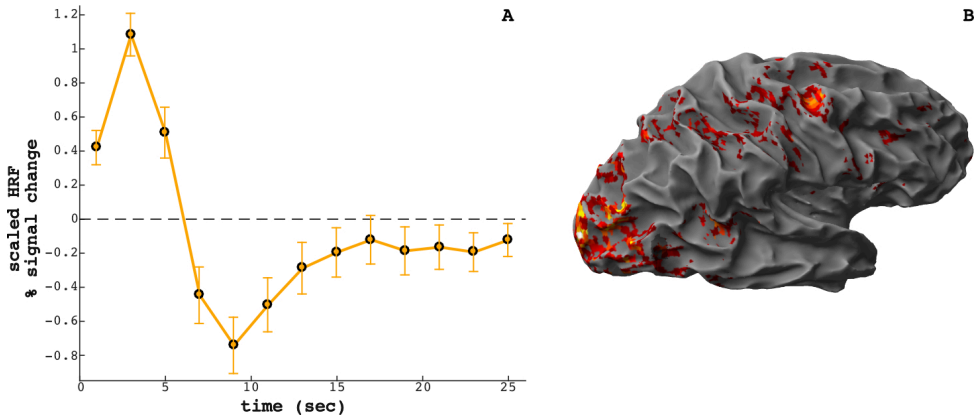


Figure 6.3 Motor response activity for one example subject. On the left: estimated HRF for a voxel in the motor cortex indicating the key press when the subject made his choice. On the right: localised activity for this ROI.

It is worth noticing that increased activity can also be observed in visual areas. This is likely due to the timing of the trial: a response immediately followed the presentation of the stimuli, making the two events too close in time for the haemodynamic response in visual cortex to go back to baseline.

6.1.2.2 Control Analysis: Coherence Response Modulation

Given that previous studies have shown a systematic modulation of BOLD response in area MT dependent on the strength of coherent motion in an RDK (see for example Rees, Friston, & Koch, 2000), we tried to measure the response curve of voxels in area MT and see if a difference between the different levels adopted in the coarse discrimination scans could be observed. Figure 6.4 shows the results of this analysis for subject j: a strong response to the moving stimuli is observed in area MT (Figure 6.4A). However, there is no significant difference in the response curves depending on the strength of the signal (Figure 6.4B). This was true for all four subjects.

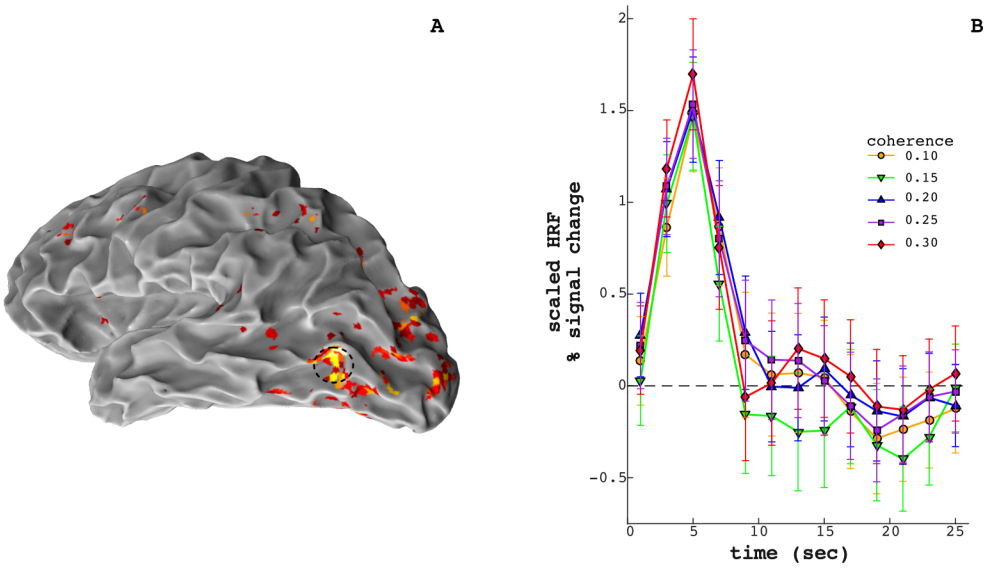


Figure 6.4 Coherence modulation in left MT for one subject. A) localised response in area MT of the left hemisphere. B) HRF for different coherence levels used in the coarse discrimination task.

An explanation for this lack of differences could be that the changes in coherence in the stimuli were too subtle to evoke a significant differential response across conditions. Indeed, in studies where this effect was observed, the levels contrasted were either random noise (0%) versus fully coherent motion (100%), whereas intermediate values show response amplitudes in between these two range limits. Despite this fact, we were able to show a modulation of response amplitudes for coherent motion in areas MT, and this suggested that the features of the stimuli and the data acquisition process worked as expected.

6.1.2.3 Task-Related Activity Changes

Another important aspect to check was to ensure that the demands in performing both fine and coarse discrimination tasks were balanced. This means that, when contrasting trials in which participants judged two highly similar stimuli with trials where the stimuli differed in a more substantial way

from each other, no significant differences in BOLD response should be observed. This was tested, for each participant, by contrasting the task condition on the whole volume. The results (Figure 6.5 – data from subject g for areas V1 (A) and MT (B), left hemisphere) show no significant differences across the two tasks.

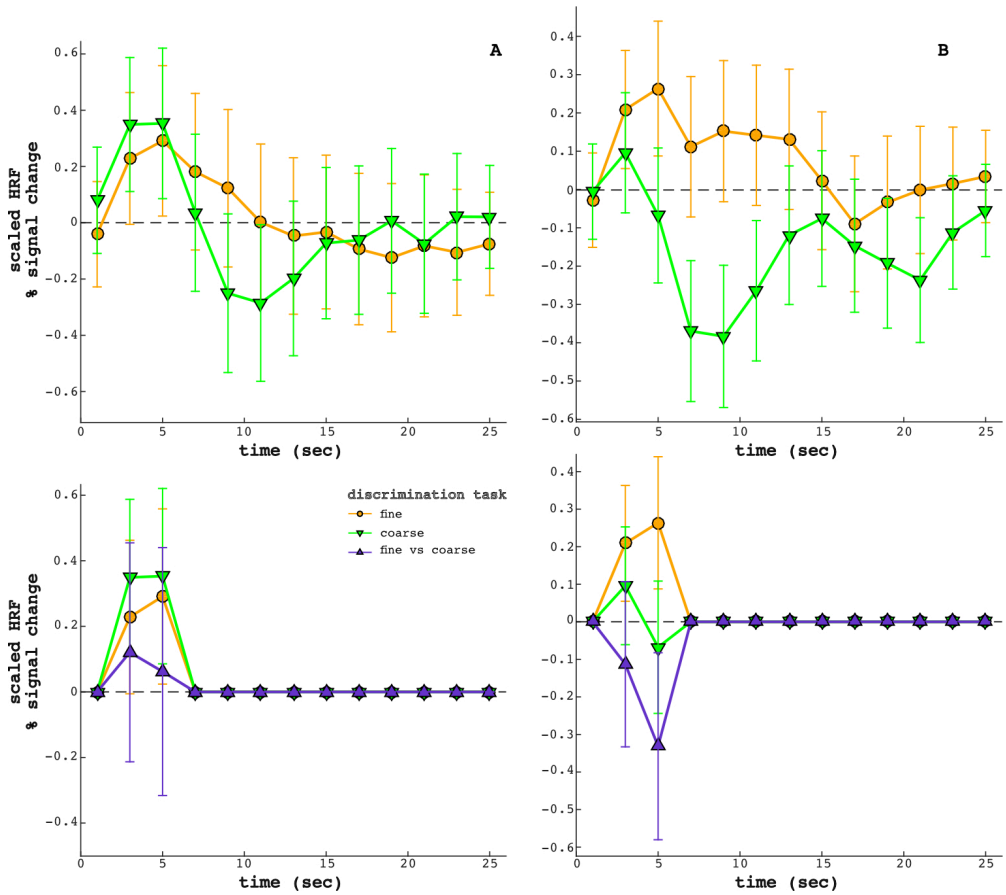


Figure 6.5 Comparison of Fine and Coarse tasks for one subject across the whole volume. A HRF curve for fine and coarse discrimination (contrast displayed in purple) for area V1. B) HRF curve for same comparison in area MT.

This was true for three participants, whereas subject j showed an effect of task in area MT lateralised to the left hemisphere (Figure 6.6). BOLD response to trials in the coarse discrimination presented a larger amplitude than responses to trials in the fine discrimination task.

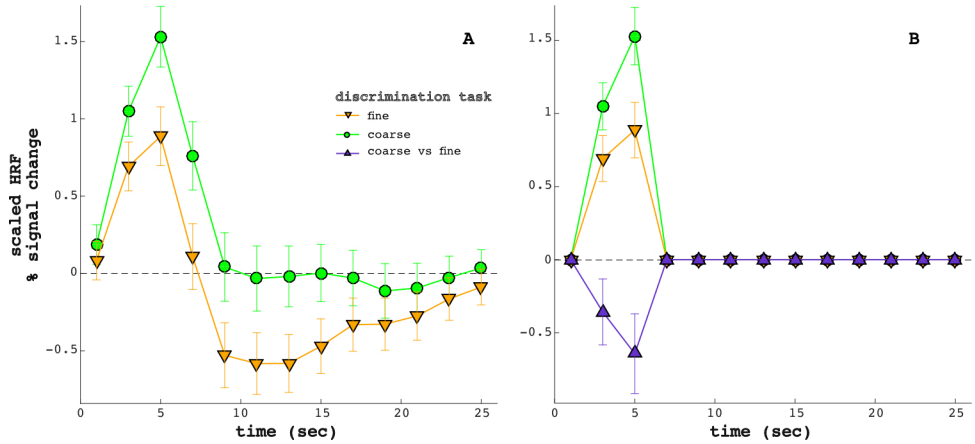


Figure 6.6 Fine versus Coarse discrimination in area MT for subject j. A) HRF for fine and coarse tasks. B) HRF for contrast between coarse and fine tasks (in purple) for area MT (left hemisphere).

The discrepancy in results across subjects for this analysis led to question whether a difference in task demand could be obscured for the remaining three participants by the fact that stimuli, in the two conditions, were different for half of the trials. In order to account for this difference in sensory stimulation, a second, more restrictive, GLM analysis was performed. This time, only trials where the same stimuli were presented, but different tasks performed, were selected. Subject j presented the same significant difference in response amplitude as described in Figure 6.6. The other three observers showed no effect as from the previous analysis.

6.1.3 Discussion

The analyses presented in this chapter aimed at ensuring that standard levels of data quality were met during acquisition. In particular, as a first control, response amplitude changes modulated by key-press (i.e., indicating the action performed by subjects when reporting their choices) were identified for all observers in localised areas of the motor cortex, contralateral to the hand used to give the response. As highlighted in

Section 6.1.2.1, large BOLD responses visible in areas of the visual cortex are likely to be related to the timing of the experiment: action performance immediately followed the presentation of the visual stimuli, thus the haemodynamic response evoked by the sensory stimulation did not have enough time to return to baseline before the choice-related key was pressed. A second reality check of the data, as well as of the adequacy of the stimuli used in the study, consisted in deriving haemodynamic response functions modulated by the moving stimuli in visual areas known to present motion-selective neurons (i.e., MT). Previous studies reported a strong modulation of BOLD response in voxels of area MT when comparing signals evoked by static stimuli (or blank epochs) with signals evoked by fully coherent stimuli (100% coherence) (see, for example, Watson et al., 1993). In the experiment presented here, changes in coherence level when performing a coarse discrimination task were implemented in order to ensure that the difficulty of this task was in line with that of the fine discrimination task. For this reason, the actual levels adopted differed by a small amount from each other (from 10 to 30% in steps of 5%), resulting in a change in response amplitude too small to be identified. Despite this detail, all subjects showed an overall, consistent modulation of the haemodynamic response in area MT to the stimuli, which confirmed the adequacy of the RDKs presented during the scanning sessions.

Given that the overall aim of this study was to investigate sensory mechanisms processing information from a fine versus coarse visual motion discrimination task, it was important to ensure that differences in activity

patterns evoked during these two, blocked, conditions were not related to higher order levels of processing (e.g., different attentional demand due to unbalanced task difficulty levels). This was achieved by contrasting response related activity between the two conditions on the whole volume, for each participant. Results were unsystematic: while one subject presented a significantly larger response amplitude in area MT when performing a coarse discrimination, compared to a fine discrimination task, the remaining three subjects showed no effect of condition. It is important to remember the way in which the stimuli were generated for the two types of tasks. The RDK designated to be the reference stimulus presented coherent motion in five possible directions, and these features were shared by stimuli in both coarse and fine tasks. The comparison stimulus, on the other hand, differed in approximately half of the trials between the two task conditions: while in the fine discrimination the direction of motion of the comparison RDK was always computed within a range of ±20° from that of the reference stimulus, this was true only half of the times in the coarse discrimination. The remaining trials presented a comparison with direction of motion approximately 180° away from that of the reference stimulus. Given that all trials were used in the analysis described above, it could be possible that in the three subjects where no task effect was observed, this was simply obscured by changes in sensory inputs between fine and coarse conditions. To control for this, a second analysis restricted the trials pool to those with physical stimuli presenting the very same characteristics. However, subjects that did not show significant differences in the less

restrictive analysis still presented a comparable BOLD response for fine and coarse conditions. In a follow-up after scanning with subject j it was reported that the coarse discrimination task required a higher level of engagement. It is therefore possible that the difference observed for area MT derive from a top-down, attentional gain mechanism.

Thus, the results presented in this chapter show that the paradigm applied in the imaging study allowed to reliably detect event-related haemodynamic responses. However, the mapping of the question we are addressing with this study happens at a scale much smaller than that of the voxel. This means that a univariate approach is inadequate for distinguishing decoding strategies applied during a fine / coarse discrimination task. A way to bypass this problem is to employ a multivariate approach. Chapter 7 describes this methodology and is divided into two parts. First, a discussion of the population receptive field model and relative results from the data collected during the retinotopy session is presented. The description of the data analysis through a forward encoding model I implemented is the focus of the second part.

CHAPTER 7

7 Encoding Models of Visual Representations

This Chapter describes two types of multivariate approaches applied to the functional imaging data presented in Chapter 6. The goal of the first part (Section 7.1) is to show the feasibility of modelling internal visual representations by measuring activity patterns when systematic changes are made to a visual stimulus along a particular dimension of interest. This is achieved via population receptive field modelling with contrast-defined visual stimuli (Dumoulin & Wandell, 2008). The second part of the Chapter (Section 7.2) presents a multivariate approach (forward encoding model, see as an example Brouwer & Heeger, 2009) for estimating visual representations of highly similar (fine) versus highly dissimilar (coarse) moving dot patterns.

7.1 Population Receptive Field (pRF)

One central idea in the study of visual representations is the measurement of *tuning curves*. Rooted in invasive, electrophysiology studies on animals (e.g., Campbell et al., 1969; Hubel & Wiesel, 1959), tuning curve measurement allows the experimenter to examine functional properties of a neuron or an area of interest in a fine-grained way. PRF modelling can be seen as an fMRI version of such tuning curve measurements and provides several advantages over the univariate

approach described in Chapter 6. The aim of tuning curve and pRF models is to *link* different levels of a stimulus dimension to responses in the unit under investigation, be it a neuron, voxel, region of interest, or even a visual area. In practice, this is achieved by systematically manipulating a predetermined feature of interest of the stimulus while measuring responses in the unit of interest. The estimated pRF / tuning curve model relates the levels of the manipulated stimulus dimension to modulations in the measured activity patterns (Dumoulin & Wandell, 2008; Larsson & Heeger, 2006; Wandell, 1999; Wandell, Dumoulin, & Brewer, 2007).

In what follows, the receptive field properties obtained using the model proposed by Dumoulin & Wandell (2008) and described in Chapter 2, Section 2.3.6.3 are presented.

7.1.1 Material and Methods

7.1.1.1 Observers

Six subjects participated in this study (four females; age 25 ± 3.41 years, mean ± s.d.). One out of the six was excluded due to excessive head motion (see Chapter 2, Section 2.3.5.1). Of the remaining five subjects, four were familiar with the stimuli used in the session, while the remaining subject was new to both stimuli and scanning. All had normal or corrected-to-normal visual acuity and gave written, informed consent prior to taking part in the study. Procedures were approved by the Medical School Research Ethics Committee of the University of Nottingham.

7.1.1.2 Apparatus and Stimuli

A description of the stimuli used in this study is provided in Chapter 2, Section 2.3.6.2. Briefly, contrast-defined expanding and contracting rings, rotating wedges, and sliding bars were used to derive standard eccentricity and polar angle maps (rings and wedges stimuli only, shown in Chapter 2, Section 2.3.6.2) and to estimate population receptive fields of early visual areas (responses to all stimuli were used). Details of the acquisition sequence and parameters are provided in Chapter 2, Section 2.3.7.

7.1.1.3 Procedure

The experimental procedure is described in Chapter 2, Section 2.3.6.2. Briefly, subjects undertook a total of 10 scans: four rotating wedges (two in CW direction, two in CCW direction), two ring stimuli (one contracting, one expanding), and four bar stimuli (two for each motion sequence direction). While presented with these 'checkerboard' patterns, participants fixated at the centre of the screen and performed a two-interval forced-choice contrast discrimination task: the contrast of the cross at fixation changed between two intervals and subjects indicated whether the first or second presentation appeared darker. To ensure a constant attentional level throughout the scanning sequence, the task was controlled by a staircase procedure. A representation of the stimuli used is provided in Figure 2.25.

7.1.1.4 Data pre-processing and Analysis

Motion correction as described in Chapter 2, Section 2.3.5.1, was applied to the data in order to minimise artefacts induced by head motion. This resulted in the exclusion of one subject. Examination of the parameters from the remaining five subjects proved a residual amount of motion that would not interfere with the analysis. Anatomical and functional data were aligned to the high resolution image as described in Chapter 2, Section 2.3.5.3, and surfaces and flat maps computed as in Section 2.3.5.4 and 2.3.5.5. The pRF method used to analyse the data is described in Chapter 2, Section 2.3.6.3 and was implemented in mrTools (available from http://gru.stanford.edu/doku.php/mrTools/tutorialsprf).

7.1.2 Results

An example of the pRF analysis output, for subject j, visual area V2d of the right hemisphere, is provided in Figure 7.1. In A) the surface of the right hemisphere is shown, with areas V1 and V2d indicated by the coloured shaded areas. B) shows the model of the haemodynamic response function (HRF) used in the analysis: the characteristic initial delay, with peak around 5s from stimulus onset can be appreciated. The density plot in C) shows a representation of the receptive field fitted to the region of interest (ROI) V2d. As expected, this is located in the third quadrant of the plot: dorsal area V2 contains the representation of the ventral portion of the contralateral visual field (bottom left). Sections of the timecourse are presented in D), black dotted line (grey, segmented bars indicate cuts in time in displaying

the data). The model fit of the data obtained from the analysis is displayed in red (solid line on top of timeseries). A quick inspection of the fit suggests the model well explains the variability in the data (r^2 = 0.84). Examples of the stimuli used are plotted at the bottom of D): each sequence shows the location of the stimulus, across the time window indicated by the straight grey bar on the timeseries, with respect to the RF (light red shaded area). Plots in E, F, and G present estimated RF properties for each voxel in the ROI. In E) RFs are represented as circles and corresponding centres are displayed in F). RF half width, the parameter indicative of the RF size, is plotted as a function of eccentricity in G): as expected, the size increases linearly as we move from the centre to the periphery of the visual field.

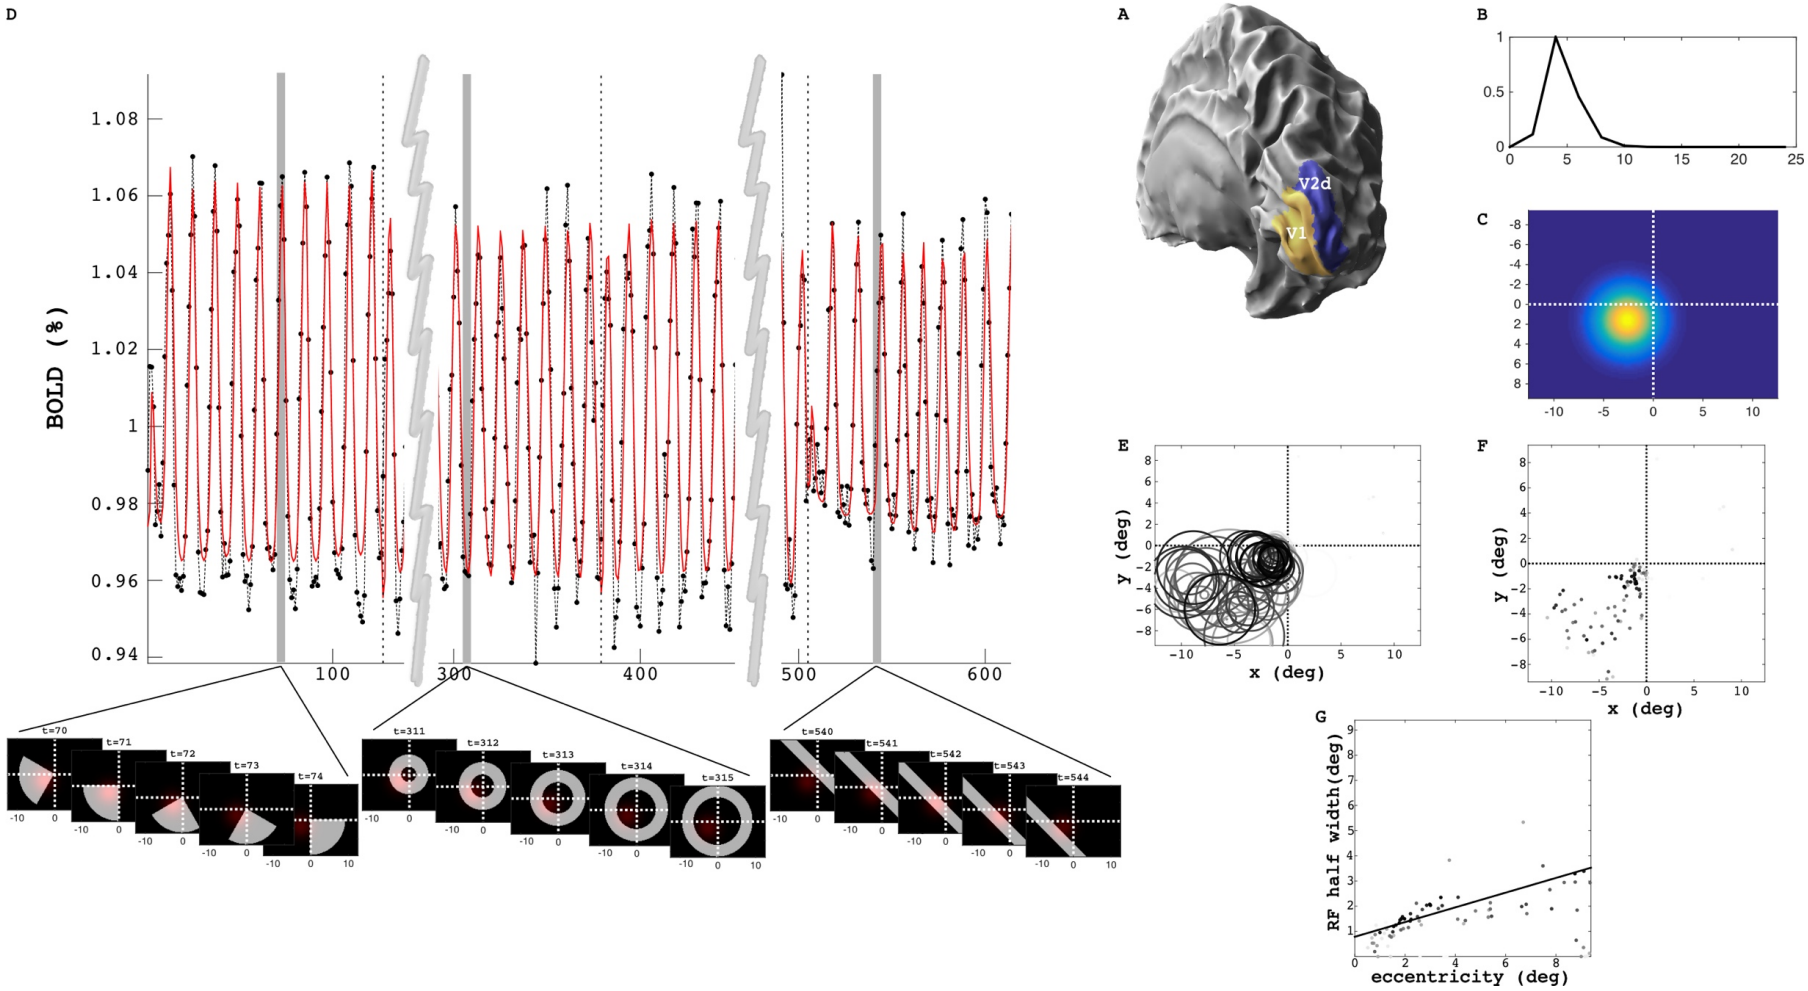

D
BOLD (%)
1.08
1.06
1.04
1.02
1
0.98
0.96
0.94
100
300
400
500
600
t=70
t=71
t=72
t=73
t=74
t=311
t=312
t=313
t=314
t=315
t=540
t=541
t=542
t=543
t=544
A
v2d
v1
B
C
E
y (deg)
x (deg)
F
y (deg)
x (deg)
G
RF half width(deg)
eccentricity (deg)

Figure 7.1 pRF analysis output. Data from one subject, visual area V2d of the right hemisphere. A) Localisation of area V2d. B) Representation of the HRF used in the analysis. C) Density plot locating the estimated RF for the selected visual area within the cortical representation of the visual field. D) Timecourse and model fit. At the bottom the sequence of stimulus presentations for selected time windows across the scanning sequence are shown relative to the estimated RF (light red areas). E) RF for individual voxels in V2d are represented as circles. F) Centres of RF for individual voxels. G) Relationship between RF size of single voxels and eccentricity.

Eccentricity and polar angle maps are also obtained from the model fitting and an example is provided in Figure 7.2. The maps derived from the pRF analysis (on the right) are compared to those computed via classic, phase-encoded retinotopic analysis (as described in Chapter 2, Section 2.3.6.2) for one subject. The validity and reliability of the two methods can be appreciated: borders of visual areas are equally identified through the two methods.

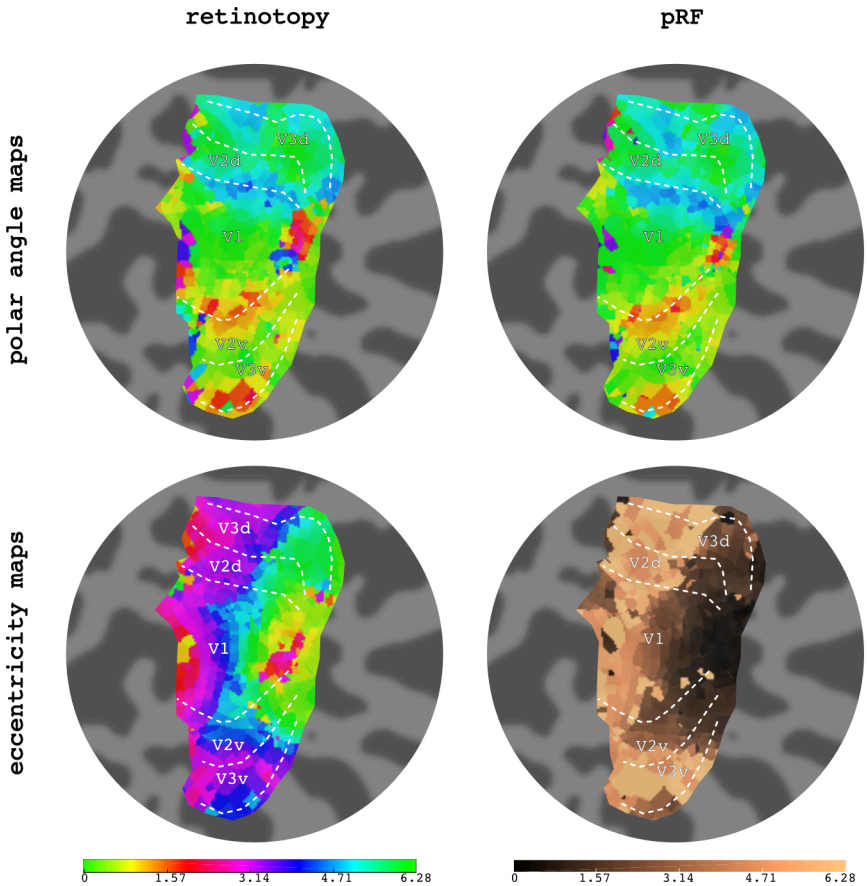


Figure 7.2 pRF and retinotopic maps. On the right: polar angle and eccentricity maps obtained with the pRF analysis for one subject, right hemisphere. On the left: polar angle and eccentricity maps for the same subject, same hemisphere, as shown in Figure 2.27 (subject d).

A summary of the pRF parameters obtained for all subjects is presented in Figure 7.3. Rows show centres and sizes of pRFs for left (column one) and right (column 2) visual areas V1, V2, and V3, in turns. Different colours represent data from different subjects, whereas different symbols (for areas V2 and V3 only) indicate whether data are from the dorsal (circles) or ventral (squares) portions of the visual area. Plots of the pRF centres (rows one, three, and five) confirm that the 'correct' regions of the visual field are encoded by each area: visual areas of the left hemisphere map the right visual field, visual areas of the right hemisphere map the left visual field. Similarly, dorsal regions of areas V2 and V3 map parts of the visual field below the horizontal meridian, and ventral regions of V2 and V3 map portions of the visual field above the horizontal meridian.

Further, plots of pRF sizes as a function of eccentricity (rows two, four, and six) show a linear relationship between the two parameters for all visual areas: as the eccentricity increases (from centre to periphery) the size of the pRF increases. This is in line with previous observations by Dumoulin and colleagues and other groups who have been applying this methodology.

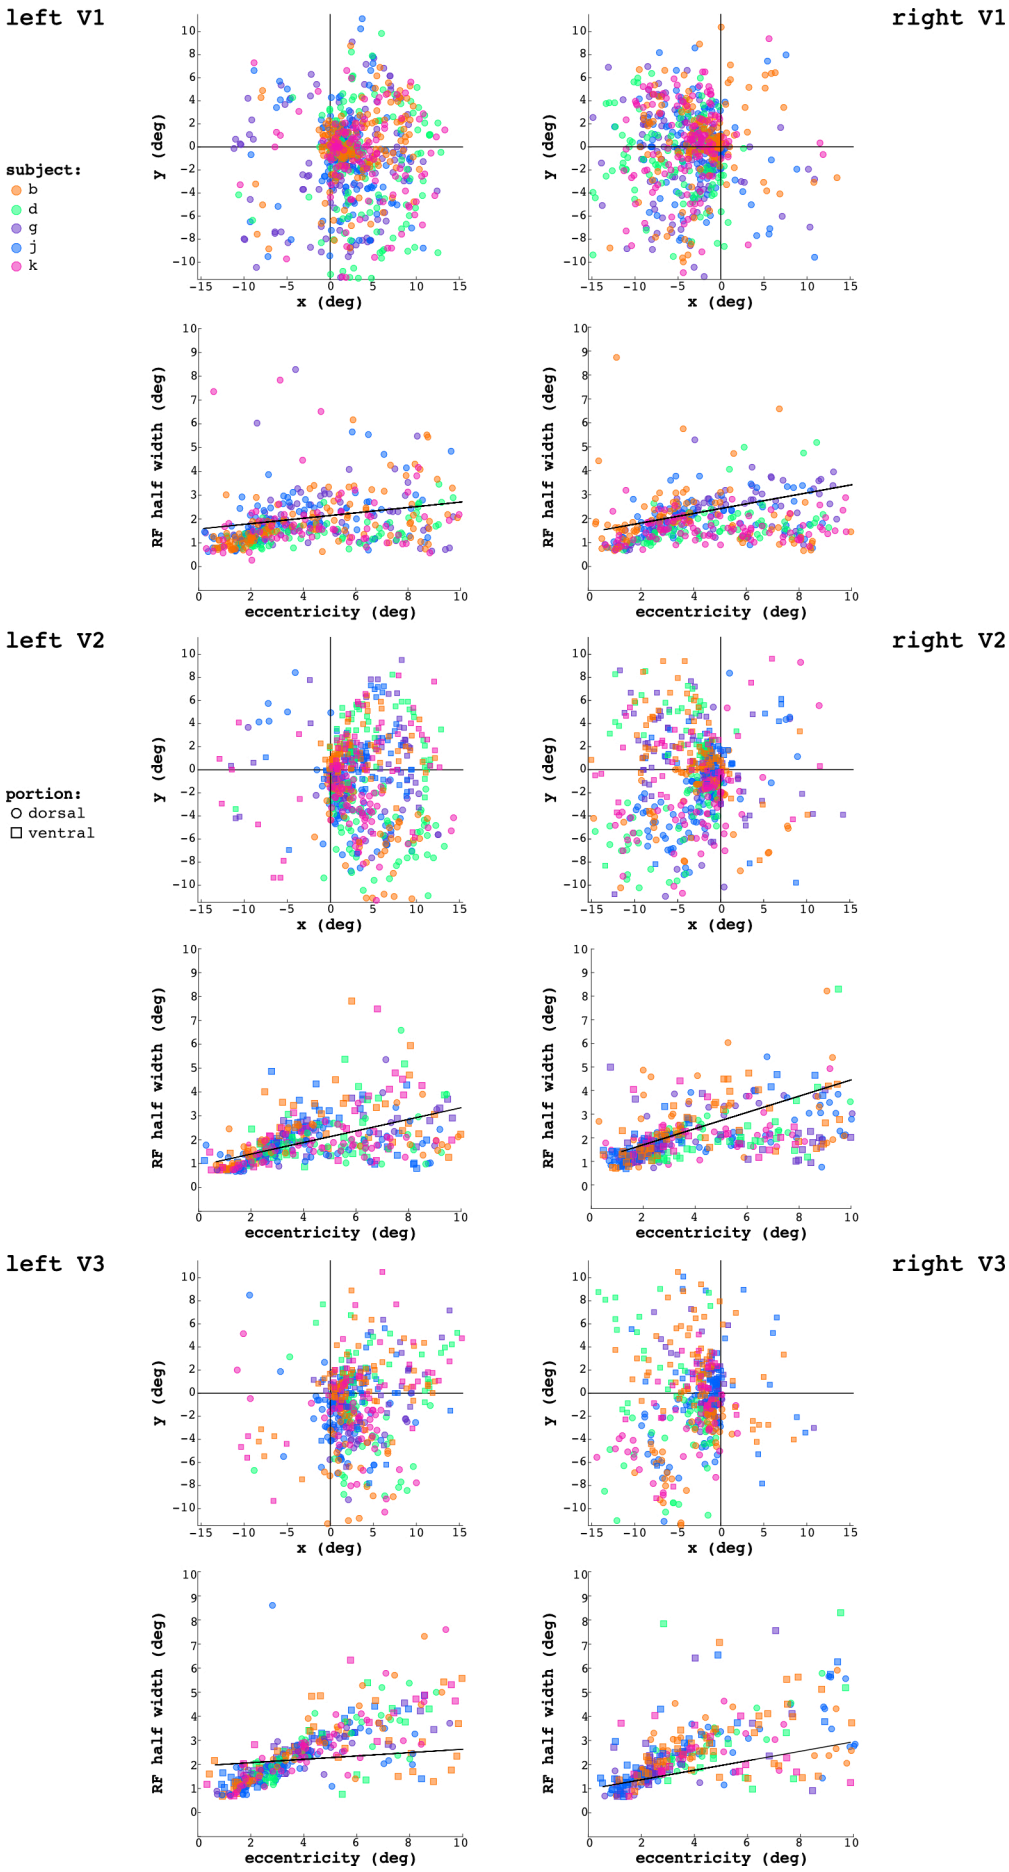

left V1
right V1
subject:
b
d
g
j
k
y (deg)
x (deg)
RF half width (deg)
eccentricity (deg)
left V2
right V2
portion:
dorsal
ventral
left V3
right V3

Figure 7.3 pRF summary across subjects. Rows one, three, and five show pRF centres for visual areas V1, V2, and V3, left (column one) and right (column two) hemispheres. Rows two, four, and six show pRF sizes as a function of eccentricity. Colours indicate data from different subjects, symbols in rows three to six indicate data for dorsal (circles) and ventral (squares) portions of areas V2 and V3. The correct parts of the visual field are mapped by each visual area and sizes of receptive fields increase with increasing eccentricity, for all subjects.

7.1.3 Discussion

The results presented in Section 7.1.2 show that a pRF analysis of the mapping data allows us to examine functional properties of regions of the brain in a more fine-grained way, compared to the type of analyses presented, for example, in Chapter 6. Properties of population receptive fields estimated via this 'tuning curve' approach generally agree with results obtained by invasively recording cell activity in animals (e.g., Felleman & Van Essen, 1987; Gattass et al., 1981; Van Essen, Newsome, & Maunsell, 1984). The validity of this analysis is supported by the fact that centres of receptive fields of voxels contained in specific visual areas were estimated to map the correct regions of the visual field. Similarly, data from all subjects showed the well-established relationship between the size of the population receptive field and eccentricity: voxels representing sensory signals in the periphery have larger receptive field sizes than voxels representing information around the centre of the visual field. Despite these positive results, it is worth noting a few limitations of this approach.

First, because the model aims at identifying borders of visual areas by considering modulations in the measured responses due to changes in stimulus contrast, the results only capture the spatial envelope of the population receptive field. The results cannot easily be generalised to other

visual features. The pRF model maps the levels of the dimension of interest onto the activity measured from voxels of the visual cortex: specifying a different dimension of the visual stimulus would result in different maps of the cortex (Kay, 2012).

Second, despite representing a different approach from the Generalised Linear Model, it still requires to build individual models for each voxel included in the region of interest. In this sense, the pRF approach is still a 'univariate analysis': each unit is considered individually, limiting the statistical power that can be achieved in the analysis. Thus, while the pRF model represents a sophisticated forward encoding approach of linking modulations in measured responses to changes along a dimension of the stimulus and is appropriate for measuring the layout of maps onto the cortex, it lacks the resolution of addressing questions founded at a scale finer than that of the voxels.

An alternative to this univariate approach is provided by a *multivariate* type of forward encoding model: responses from multiple voxels are incorporated in determining a functional link with a more formal description of stimulus features. Here, the high-dimensional problem that follows testing a large number of voxels is simplified by projecting voxel responses onto an abstract, more manageable space: contributions from each voxel are combined – linearly – to result in a modelled response that is 'stimulus-informed'. The forward encoding model approach, suited in addressing the fine / coarse processing difference we are interested in, is presented in the following section.

7.2 Forward Encoding Model

A description of the forward encoding, multivariate approach is provided in Chapter 2, Section 2.3.6.5. The ultimate goal of this type of analysis is to test the validity of information processing models (Kriegeskorte, 2011): to what extent can different formal descriptions of e.g. the encoding of *sensory information* account for measured response patterns in the brain? Because the organisation of the brain can be studied at multiple scales (from single neurons, to ensemble of neurons, to functionally / anatomically identified areas), computational models that define information processing need to also be constructed for different scales (Naselaris et al., 2011). Indeed, in Section 7.1 I have described a type of encoding model that treats information at the level of individual voxels. However, there are questions that can only be addressed at a much finer spatial scale (below the standard resolution limits of fMRI measurements).

One popular approach to overcome the limitations on spatial resolution is to integrate, or combine, the responses across many voxels and therefore leverage small biases in the sampling of the underlying neural substrate. This means taking the analysis into a higher-dimensional space: responses from many different voxels are considered at the same time to obtain a reliable measurement of the representation of, say, a particular sensory input.

An example of this approach is explained very clearly in a recent study by Vo, Sprague, & Serences (2017). The aim of their experiment was to

investigate the effects of spatial selective attention when encoding visual stimuli. While it is known that an attentional gain can be measured at the level of single-cell receptive fields (e.g., David et al., 2008; Moran & Desimone, 1985), as well as at the level of single-voxel receptive fields (e.g., De Haas et al., 2014; Kay, Weiner, & Grill-Spector, 2015), the way in which modulations of individual units are combined to aid encoding of a visual scene is yet unclear (Vo et al., 2017).

Two distinct, and yet related, levels of investigation were put forward: first, the effects of selective spatial attention on parameters defining voxel RFs were evaluated. Second, the impact of these individual modulations on the 'population code', the high-dimensional representation of the stimulus, was derived by aggregating information from multiple voxels into a *spatial* encoding model. Although the fine details of this study go beyond the scope of this chapter, it represents an excellent example of how encoding models can be used at multiple scales to look at effects of one factor at different representational levels. Vo and colleagues observed that attention modulates parameters of voxel RFs (vRFs) that in turn explain the quality of the representation of the visual space at a population (multivariate, ROI) level.

Similarly, while a voxel-wise approach would not allow a strong enough explanatory power to tease apart small voxel-wise differences in the responses during fine / coarse discriminations of visual stimuli, combining responses across many voxels is likely to improve the encoding of the visual representation and allow to address this question. In the following sections,

I describe the application of a forward encoding model to the functional imaging data used in Chapter 6.

7.2.1 Material and Methods

7.2.1.1 Observers

Five subjects participated in this study, although one of these was later excluded due to excessive head motion during scanning. A description of the volunteers is provided in Chapter 6, Section 6.1.1.1.

7.2.1.2 Apparatus and Stimuli

A detailed description of Apparatus and Stimuli is provided in Chapter 2, Section 2.3.7 (subsections 2.3.7.1 and 2.3.7.2). Data from both, localiser and discrimination tasks were used in this modelling approach.

7.2.1.3 Procedure

The experimental procedure for the localiser task is given in Chapter 2 (Section 2.3.7.1). Briefly, subjects were presented with the same circular apertures used for the discrimination task with the aim of identifying visually responsive voxels within these regions of the visual field. The two RDKs displayed inside the apertures moved coherently in a direction randomly selected from eight possible values (0 to 360° range in steps of 45°). The second aim of this procedure was to derive population tuning curves for the eight directions of motion, for each visual area of interest, that will then

serve as base set in the forward encoding model. In order to ensure fixation and a constant attentional level throughout the scan, participants engaged in a two-interval forced-choice contrast discrimination task (see Section 7.1.1.3) at fixation. Two scans of the localiser task were obtained for each participant.

The design of the discrimination task is provided in both Chapters 2 (Section 2.3.6.2) and 6 (Section 6.1.1.3). Subjects performed two scans of the coarse discrimination task, and two scans of the fine discrimination task.

7.2.1.4 Data pre-processing and Analysis

Data pre-processing was carried on as described in Chapter 6 (Section 6.1.1.4). Data analysis was performed at the level of individual subjects and then summarised. This particular analysis path does not lend itself to more traditional approaches to group analysis (e.g. fixed or random effects GLM on group data) and they were therefore not performed. First, the motion-corrected scans for the localiser scans and the discrimination scans – (fine and coarse together) were temporally high-pass filtered and concatenated – this merging resulted in two timeseries, which are more easily handled in further analysis than many individual scans. A version of GLM was applied separately to each of these merged scans.

From responses measured during the localiser task, regression (beta) weights for each individual stimulus direction were obtained via ordinary least squares. Similarly, beta weights for the different stimulus conditions were obtained – one for each direction of motion in the 'reference' stimulus

aperture. The reason for this was that as a test of the forward model, reconstructing the directions of motion of the reference stimuli should be relatively straightforward. Note that the reference stimulus aperture – always in the left visual field – contained RDKs computed at five possible directions. This was common across subjects. Given the lateralisation of the stimulus, we reasoned that it would be possible to reconstruct the direction of the RDK from responses in visual areas of the right hemisphere.

The data from the initial (GLM) regression step were then fed into the forward encoding model as described in Chapter 2, Section 2.3.6.5. Briefly, data from the localiser task were used to estimate a weight matrix indicating the contribution of each voxel of the ROI, e.g., right V1, to each of the hypothetical channels modelling the feature space. We assumed six channels – a weighted combination of the tuning functions for these channels can describe any direction of motion in the 0 – 360° space[3].

Data from the discrimination task, together with the estimated weight matrix, were then used to predict the matrix of channel outputs for the directions of motion of the reference stimulus. This output is the *relative activation* of each channel for each particular reference direction. To decode the direction of the reference stimulus, the channel matrix was multiplied by the tuning functions describing the entire feature space, resulting in a weighted channel output for each direction of motion. At this stage, a grand-average effect emerged from manual inspection of the data: the

[3] The 6 channels form a set of basis vectors.

distributions of the individual directions were revealed by subtracting the mean amplitude of the channel responses. This point is addressed in more detail in the Results section (7.2.2). The 'decoded' direction of motion corresponded to the location of the peak of each direction distribution (arg max).

7.2.2 Results

The reconstruction of the direction of motion presented to each subject in the left circular aperture was attempted for each participant, using data from individual visual areas (V1, V2, V3, MT) of the right hemisphere, as well as using a combined region of interest (ROI), pooling voxels from V1, V2, and V3. Despite careful construction of code that captured the rationale behind the modelling and the precise selection of the visual responsive areas, we were unable to accurately reconstruct the directions of motion of the reference stimulus, for any of the subjects. An example of the weight matrix, estimated channel outputs, and reconstructed directions for one subject is provided in Figure 7.4. The weight matrix obtained via linear regression using the data from the localiser scans and the formal description of the stimulus feature under investigation (direction of motion – channel representation), indicates the contribution of each voxel in the ROI to each hypothetical channel. For ease of visualisation, only ten voxels, randomly selected from the ROI, are displayed (different colours in Figure 7.4A) and their relative weight is plotted as a function of channels in the model. Thus, the response of one voxel can be decomposed into weighted contributions

of six hypothetical channels. Figure 7.4B shows the expected channel outputs given a set of test data (here, discrimination scans) and the weight matrix (Figure 7.4A). Different colours represent different channel outputs, plotted as a function of the five reference directions of motion (used in the fine/coarse discrimination tasks). A projection of these relative responses onto the feature space (directions of motion defined across the 0-360° range) results in the reconstructed directions of the reference stimulus (Figure 7.4C).

Here, each distribution relates to a particular reference direction, and the decoded / reconstructed one is taken to be the peak of such distribution. It is worth noting that, in order to obtain these distributions, a global, mean component was subtracted from each weighted combination of channel responses. One could speculate that this is due to the encoding of an average direction of motion between the reference and target stimuli, on each trial. If we consider that, on each trial, subjects were presented with two moving stimuli, it is possible that, following a vector averaging approach, a representation of the mean direction between that of the reference and of the target stimuli is encoded in the measured signal. This pattern would be shared across trials (and each channel would carry this global mean component in its weighted contribution). The impact of individual channels is revealed once this masking component is subtracted.

Unfortunately, none of these decoded values resembled the ones used in the experiment. Possible explanations are discussed in Section 7.2.3.

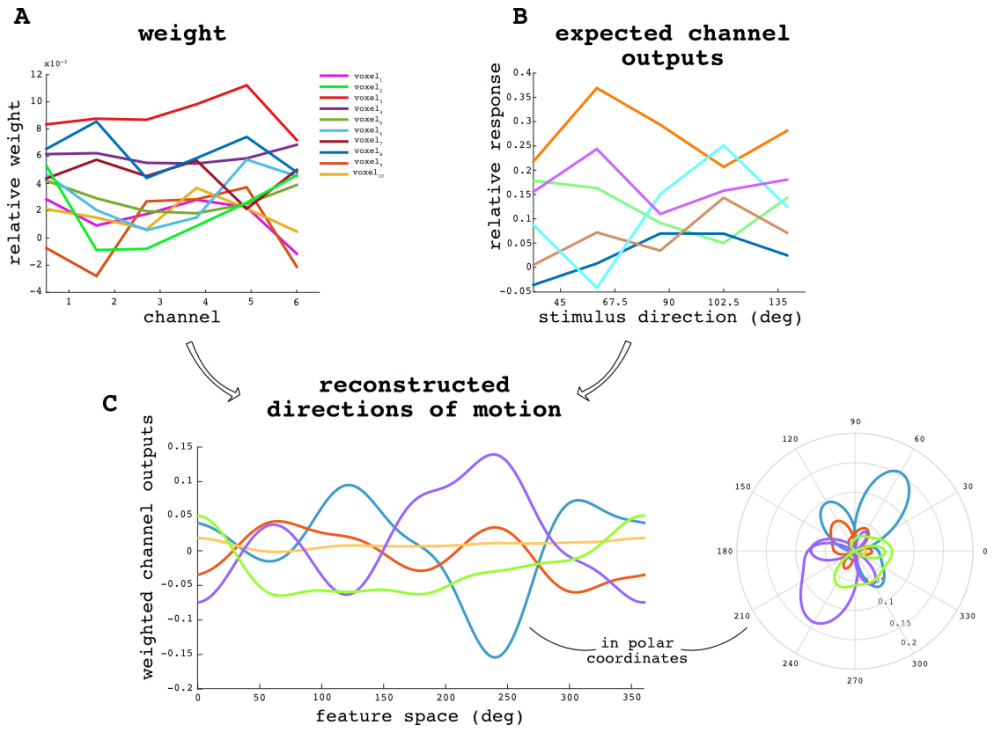


Figure 7.4 Forward Encoding Model, data from one subject, one ROI (combined V1, V2, V3 of right hemisphere). Only responses from ten, randomly selected voxels belonging to the ROI are presented. A) Weight matrix: relative contribution of each channel to the voxel responses (different colours). B) Expected channel outputs: channel contribution (different colours) to each stimulus direction (five directions of motion of the reference stimulus in the experimental / discrimination task). C) Reconstructed direction: distributions of reference stimulus directions of motion (different colours). The decoded direction corresponds to the peak of each individual distribution across the feature space.

Figure 7.5 shows the haemodynamic response function computed for a single voxel (left) and averaged across voxels for visual areas V1, V2, and V3 of the right hemisphere (right) in response to the five reference directions of motion (different colours). The curves show the expected pattern for all conditions: an increase in signal change with peak around 5s after stimulus onset, followed by a decrease, undershoot and return to baseline within 25s from stimulus onset. Again, this is a confirmation that the stimuli were eliciting an appropriate response in the visual areas of interest.

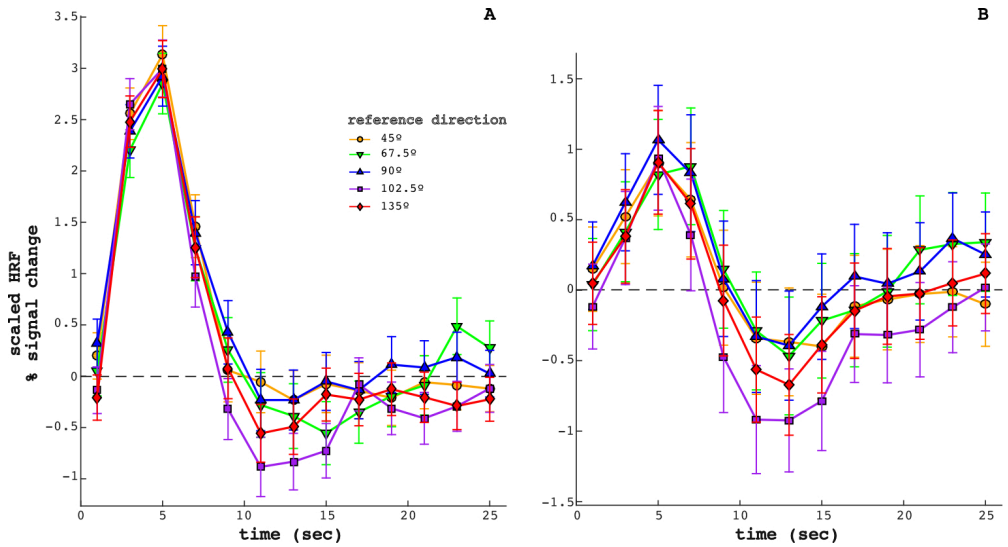


Figure 7.5 HRF in combined ROI for a single subject. On the left: HRF for five directions of motion (colour, indicating directions of the reference stimulus) for one voxel of right V1. On the right: average HRF, across voxels, for reference directions of motion (colours).

In a similar fashion, Figure 7.6 shows the HRFs computed for the directions of motion used in the localiser experiment. The same principle applies here as in Figure 7.5: on the left the computed HRFs for a single voxel in area V1 of the right hemisphere are presented. On the right, the average HRFs for the ROI (combination of right V1, V2, and V3 areas) is shown. As for Figure 7.5, we confirmed that a selective response for the directions of the stimuli used in the experiment could be estimated.

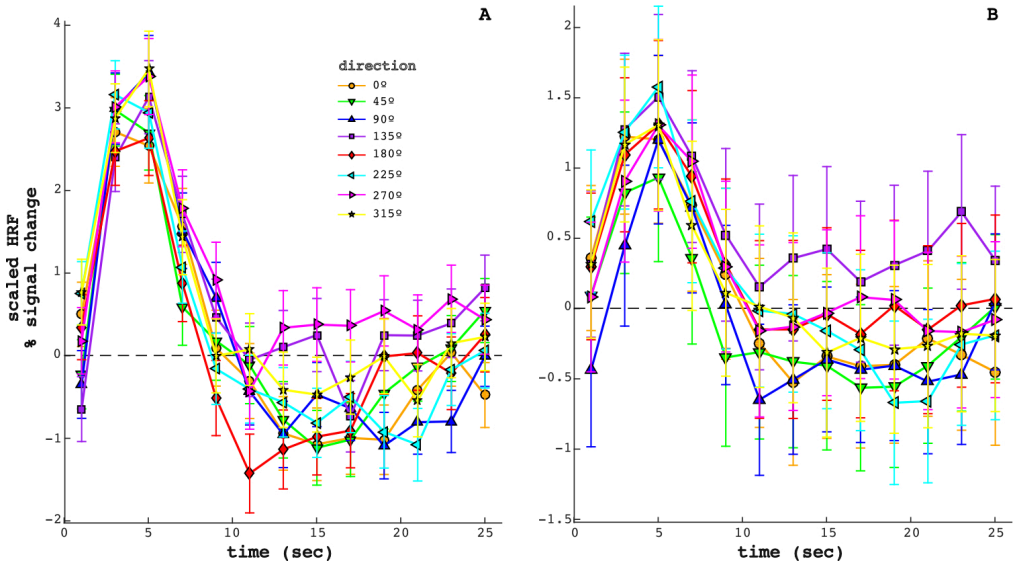


Figure 7.6 HRF for a single subject. On the left: HRFs for individual directions of motion (different colours) used during the localiser experiment, derived from one single voxel. On the right: average HRFs from voxels of the selected ROI (combined V1, V2, V3 of right hemisphere).

7.2.3 Discussion

The results discussed in this section of the chapter showed two main points. First, the stimuli used drove an appropriate response in the visual areas of interest, in both the localiser and the discrimination tasks. Second, despite other groups successfully using a form of forward encoding model to address questions related to, e.g., modulation of attention in encoding visual information (Scolari & Serences, 2012; Vo et al., 2017), our results were inconclusive. Reasons behind this null finding could relate to the choices in the experimental design. In particular, the relatively large number of levels of the experimental variable in both localiser and fine / coarse discrimination tasks might have been too ambitious. In particular, the choice of parameters made aggregating responses across trials without combining different conditions difficult. For example, even though the same reference direction was used, the direction of motion of the target stimulus might have been

different across trials, possibly adding a source of noise that prevented from obtaining a meaningful signal for the model to work. Another aspect that could have influenced the reliability of the data is the coherence level used in the discrimination tasks. While we needed a way to ensure that the difficulty of both coarse and fine discriminations was comparable, when combining data from the two tasks, the fact that target stimuli in the coarse discrimination were presented at different coherence levels would have resulted in slightly noisier signals, depending on the strength of the moving stimulus. Thus, an additional source of noise contributed to the data used in the model, potentially reducing its explanatory power. It is therefore possible that this version of the experiment was not adequate in picking up the effect of interest.

On the other hand, because the model did not allow us to compare representations from fine and coarse trials directly, we can't dismiss the possibility that the tuning of on / off tuned units differed between the two tasks. A more restrictive design, maximising the number of trials per stimulus conditions, and still controlling for task difficulty could be a starting point in improving the study presented here.

Having pointed out these limitations, when looking at studies that also used this type of modelling approach, it must be stressed that different groups often use several pre-processing stages which – while well motivated – are sometimes ad hoc. These include principal component analysis to reduce variability (e.g., Brouwer & Heeger, 2009, 2013; Kok et al., 2013), estimating response functions on a trial by trial basis and averaging

them to increase selectivity of the response (e.g., Brouwer & Heeger, 2011; Kok et al., 2013), collapsing data from visual areas of the two hemispheres to increase the power of the signal (e.g., Vo et al., 2017). Often, these are applied to the data before performing the forward model. Some of these approaches were incompatible with the data we collected in our experiment – for example we would not have been able to collapse data across hemispheres because stimuli were purposely segregated in space to be able to exploit the lateralisation of the representation of the visual field in the visual cortex. Future version of this experiment should allow for such additional data manipulation and dimensionality reduction prior to the application of the model.

CHAPTER 8

8 General Discussion

8.1 Summary of Findings

Perceptual decision making is a pervasive process in our everyday life: confronted with a set of alternatives, we choose one option based on inputs from the environment that have been encoded by our sensory apparatus. On some occasions, the perceptual judgments we make consistently deviate from the sensory evidence we perceive, resulting in biased decisions. The focus of this thesis was on a type of perceptual bias known as *reference repulsion*: estimates of motion direction are systematically biased away from a reference relative to which a stimulus is moving. The goal of the studies I have presented was to exploit a widely used perceptual decision making task in order to shed light on when, in the sequence of events, perceptual biases arise. Moreover, modulations of the reference repulsion effect due to response effector and higher-order factors were addressed.

Exploring when and how visual inputs and motor responses affect the outcome of sensory-based judgements could improve the understanding of how sensory information, action planning and execution, and prior knowledge about the world are integrated to obtain a coherent, though sometimes biased, behaviour. Furthermore, these findings inform a long-

standing debate between sensory versus cognitive accounts of biases such as the reference repulsion effect.

Experiments 1 and 2 (which were presented in Chapter 3[4]) explored the role played by contextual information on perceptual reports: how does the presence / absence of an oriented reference line at different points in a trial affect the well-known reference repulsion bias. Subjects were presented with a static, oriented reference line and had to judge whether the coherent signal within an RDK had an axis of motion CW or CCW relative to the reference. Following the binary choice, they *estimated* the perceived direction of motion of the stimulus. Estimates in the presence of the reference line were systematically biased away from it. On the other hand, estimates in the absence of such a reference were unbiased and generally matched the direction of motion of the stimulus. These results suggest that, despite the irrelevance of the oriented reference in estimating the direction of motion of the stimulus, its presence strongly affects the final outcome: the reference repulsion effect is subject to the presence of the reference when *reproducing* the perceived motion direction of a stimulus.

Experiment 3 further supports these findings: systematically manipulating the orientation of the reference line when subjects made their judgements resulted in biases yoked to the location of the oriented segment at time of response. This not only confirmed the role of the reference line in the

[4] This work has been published as part of the journal article: Zamboni E, Ledgeway T, McGraw PV, & Schluppeck D (2016). Do perceptual biases emerge early or late in visual processing? Decision-biases in motion perception. *Proc. R. Soc. B* **283:** 2016.0263 http://dx.doi.org/10.1098/rspb.2016.0263

reference repulsion effect, but also showed that the bias emerges at a relatively later stage than that of the sensory information encoding one: distributions of motion direction estimates were biased away from the angular location of the reference at time of response, not at time of stimulus presentation. Overall, these results have strong implications on previously proposed accounts of the reference repulsion effect: a strong component of this phenomenon is explained by the perceptual and contextual information that is provided by the reference itself. Thus, events happening relatively late in the task sequence modulate the way in which subjects respond, even though the sensory information provided by the moving stimulus does not change.

In the work presented in Chapter 4[5] I applied a computational modelling approach to the experimental data. In particular, I compared a decoding model proposed by Jazayeri & Movshon (2007a) with a late, decision related bias model that takes into account not only the decoding strategies used to optimally choose between two alternatives (see Jazayeri & Movshon, 2007a), but also the perceptual and contextual information provided by a reference line when estimating a feature of the encoded visual stimulus. The model implemented here can account for the findings from Experiments 1, 2, and 3 of Chapter 3, implying that a richer, quantitative description of the repulsion effect can be obtained when multiple levels of information

[5] Part of this work has been published in the journal article: Zamboni E, Ledgeway T, McGraw PV, & Schluppeck D (2016). Do perceptual biases emerge early or late in visual processing? Decision-biases in motion perception. *Proc. R. Soc. B* **283:** 2016.0263 http://dx.doi.org/10.1098/rspb.2016.0263

processing are considered. The results also highlight that not only does the visual stimulus have to be optimally encoded and decoded by the sensory apparatus, but this representation needs to be constantly updated relative to the context within which a response is expressed. It is the interplay of these 'layers' that results in measured perceptual errors.

Chapter 5 explored the role of *response modality* in estimates of motion direction. Subjects, presented with a static reference line and an RDK with coherent dots moving CW / CCW relative to the line (paradigm as in Chapter 3, Experiment 1), made a saccadic eye movement in the direction they perceived the coherent signal. Voluntary eye movements were chosen as response effector because of their fundamental characteristics: although voluntarily initiated, they are inherently less variable than manual responses. This is because, aspects such as speed and duration of saccades cannot be overtly controlled by the observer (Becker, 1989), thus reducing the impact of spurious variables when investigating perceptual biases. The results obtained were particularly compelling, in that they showed no reference repulsion effect: saccadic estimates of perceived directions of motion were veridical and unaffected by the location of the reference line during task performance. These findings, together with the results from Chapter 3, suggest that the perceptual phenomenon under investigation has a multi-faceted nature and that the integration of evidence from different stages of sensory and motor information processing is key in accounting for the repulsion phenomenon.

Although results from previous Chapters argue in favour of a late emergence of the bias, this claim does not exclude the possibility that different decoding processes are employed when discriminating similar versus dissimilar sensory inputs. Chapters 6 and 7 attempted to address this specific issue using functional magnetic resonance imaging techniques and a multivariate data analysis approach. To this end, participants performed a fine and a coarse discrimination task, in an event-related fMRI experiment, on two moving stimuli: the RDK presented to the left of the central fixation point contained the reference stimulus, the RDK presented to the right of the fixation point contained the target / comparison stimulus. The latter moved either CW / CCW (fine) or in the same / opposite (coarse) direction relative to the reference, and subjects judged the incoming sensory information along this dimension. My starting assumption, based on previous work by other researchers, was that a multivariate approach would provide higher sensitivity than that obtained from analyses at a voxel level. I therefore implemented a forward encoding model (see Brouwer & Heeger, 2009). The goal of this approach consists in reducing a high dimensional problem – analysing responses from hundreds of voxels – to a more manageable space, exploiting not only the selective responses of each voxel within a region of interest, but also the way in which such responses vary when considered as a pattern. This allowed me to take advantage of informative signals that would be lost on a voxel-level framework.

As a first step (Chapter 6), measured imaging data were assessed: response-related activity was recovered from each subject, potential confounds by

differences in task difficulty were excluded (no significant differences in performance between fine and coarse tasks was observed). This allowed me to move onto the second step (Chapter 7): implementing the forward encoding model, trying to assess differences in decoding information from a fine versus coarse discrimination.

The assumption behind the model is that it is possible to describe the response of a region of interest along a particular stimulus dimension as a weighted combination of hypothetical 'channels'. For example, the space determining the directions in which a stimulus can move can be described by a set of basis functions (hypothetical channels, Brouwer & Heeger, 2009). The responses to different motion directions measured in a brain region (many voxels) can be modelled as a weighted combination of these (fewer) channels, thus reducing the dimension of the problem from the number of voxels, which is often in the hundreds of thousdands, contributing to the ROI to the number of channels implemented in the model.

To test this approach, data from a localiser experiment were used. Here subjects engaged in an attentional demanding task while presented with two RDKs, separated in space, but synchronously moving in the same coherent directions. These (training) data were used to derive a weight matrix. This matrix shows the contribution of each channel to the response measured in a voxel / unit. The weight matrix was then used, in combination with the (test) data obtained from the experimental session, where subjects performed blocks of fine and coarse discriminations, to estimate the channel outputs describing the relative contribution of each channel to the

stimulus direction of interest. To minimize confounds in the cortical space, we separated in space the reference and target / comparison stimuli to exploit the lateralised organisation of the visual system. Despite the careful implementation of the model and choice of stimuli display we were unable to accurately decode the directions in which the stimuli moved. The null result deriving from this analysis, however, does not conclusively speak in favour of either hypothesis: different decoding strategies may still be applied when discriminating between similar versus dissimilar moving stimuli. One possibility is that the specific paradigm used in this set of experiments was suboptimal for addressing the question.

8.2 Conclusions and Future Directions

The results presented in this thesis show the importance of investigating perceptual phenomena at multiple processing levels: encoding and decoding of sensory information, influence of contextual information on visual representations, effector of response and its intrinsic properties. The work in this thesis has shown that biases in the perception of visual motion are affected, to different extents, by all these factors.

A psychophysical approach to further address the importance of the oriented reference line, as well as to investigate the way in which motion direction of the RDK is encoded – in absolute or relative terms – would be to invert the order of presentation of the stimuli in the paradigm used throughout this thesis. By presenting subjects with the moving stimulus

first, and providing a reference line only after the visual input to be estimated is encoded, would a bias still emerge? According to the evidence presented in this thesis, the order of presentation should not affect the final outcome as much as the presence / absence of the reference during the estimation would. Another interesting aspect that could explain part of the effect observed is the possibility of having residual repulsion from preceding trials. This phenomenon, also known as *serial dependence*, has been reported in several psychophysical tasks (e.g., Fischer & Whitney, 2014). This effect, although more recently suggested as an explanation for the stability with which sensory information is perceived (see Liberman, Zhang, & Whitney, 2016), has been shown to account for an attraction effect across trials when estimating the orientation of visual stimuli. In this respect, it would be interesting to see if, given a set of trials where the target stimulus is consistently moving in a CW direction relative to the reference, trials following the first CW estimate are reported as moving in a direction even more CW. Also, the interaction of this potential effect with the magnitude of the reference repulsion, dependent on the relative distance between coherent motion of the target stimulus and reference, as well as strength of the coherent signal, would further expand the knowledge of the factors underlying this perceptual bias.

The role of the response modality, highlighted in Chapter 5, could be further investigated in two ways. On the one hand, continuously tracking the manual adjustments leading to the final estimate of perceived direction of motion would inform us on the timing of the effect: is the bias emerging

due to changes in the positioning of the response bar following the initial choice? Similarly, by recording saccades following the first response and looking at these secondary adjustments in the positioning of eye gaze would help understanding whether top-down modulations, resulting in these corrected estimates, are a primary source of the observed bias.

Finally, a careful consideration of the imaging data is required. In particular, in future research it would be fundamental to address whether the paradigm adopted is adequate in studying this fine / coarse decoding distinction. Certain improvements would be obtained by drastically reducing the number of directions used in the experiment presented in this thesis. A restricted amount of variables will guarantee stronger signals, thus an increased signal to noise ratio, as well as an easier way of grouping trials together. I think, ensuring that reference and comparison stimuli are spatially segregated will still be of fundamental importance, as it allows the separation of respective representations from distinct areas of the visual cortex. Implementing a two-interval version of this task would be suboptimal, potentially introducing a memory component that could affect the stability of the stimulus representation. From a more analytical point of view, the analysis might benefit from implementing a few pre-processing steps, as suggested in other published work using this approach. For example, applying a principal component analysis could increase the signal to noise ratio, as well as allowing to further reduce the dimensions treated during the modelling stage. The formal description of the stimulus dimension under investigation (direction of motion) could benefit from a

finer sampling of the feature space. The choice of using only six channels was the outcome of a mathematical convention: a circular space can be described by two channels with sinusoidal shape, however, to approximate firing properties of neurons, each function has to be half rectified and strictly positive, resulting in a least number of six channels (Brouwer & Heeger, 2009). Nevertheless, the width of the basis functions depends on the number of channels covering the feature space: narrower functions result in higher number of channels, which in turn have a higher selectivity for the stimulus dimension under investigation. Hence, finding a good compromise between sparsity of feature space coverage, data pre-processing steps, and number of experimental variables modulated in the paradigm could result in an effective tool for better understanding decoding strategies used in early sensory areas when performing different tasks on the same sensory inputs.

9 References

Abrams, R. A., Meyer, D. E., & Kornblum, S. (1989). Speed and accuracy of saccadic eye movements: characteristics of impulse variability in the oculomotor system. *Journal of Experimental Psychology. Human Perception and Performance*, *15*(3), 529–43. Retrieved from http://www.ncbi.nlm.nih.gov/pubmed/2527960

Acerbi, L., Vijayakumar, S., & Wolpert, D. M. (2014). On the Origins of Suboptimality in Human Probabilistic Inference. *PLoS Computational Biology*, *10*(6). https://doi.org/10.1371/journal.pcbi.1003661

Albright, T. D., Desimone, R., & Gross, C. G. (1984). Columnar organization of directionally selective cells in visual area MT of the macaque. *Journal of Neurophysiology*, *51*(1), 16–31. https://doi.org/6693933

Albright, T. D., Jessell, T. M., Kandel, E. R., & Posner, M. I. (2000). Neural science: a century of progress and the mysteries that remain. *Neuron*, *25 Suppl*, S1--S55. https://doi.org/10.1016/S0896-6273(00)80912-5

Amano, K., Wandell, B. a, & Dumoulin, S. O. (2009). Visual field maps, population receptive field sizes, and visual field coverage in the human MT+ complex. *Journal of Neurophysiology*, *102*(5), 2704–2718. https://doi.org/10.1152/jn.00102.2009

Andersen, R. a. (1989). Visual and eye movement functions of the posterior parietal cortex. *Annual Review of Neuroscience*, *12*, 377–403. https://doi.org/10.1146/annurev.neuro.12.1.377

Andersen, R. A., Brotchie, P. R., & Mazzoni, P. (1992). Evidence for the lateral intraparietal area as the parietal eye field. *Current Opinion in Neurobiology*, *2*(6), 840–846. https://doi.org/10.1016/0959-4388(92)90143-9

Andersen, R. R. A. (1997). Neural mechanisms of visual motion perception in primates. *Neuron*, *18*(6), 865–872. https://doi.org/10.1016/S0896-6273(00)80326-8

Bahill, A. T., Clark, R. C., & Stark, L. (1975). The main sequence, a tool for studying human eye movements. *Mathematical Biosciences*, *24*, 191–204.

Baizer, J. S. (1982). Receptive field properties of V3 neurons in monkey. *Investigative Ophthalmology & Visual Science*, *23*(1), 87–95.

Barlow, H. (1995). The neuron doctrine in perception. *The Cognitive Neurosciences*.

Barlow, H. B., Kaushal, T. P., Hawken, M., & Parker, a J. (1987). Human contrast discrimination and the threshold of cortical neurons. *Journal of the Optical Society of America. A, Optics and Image Science*. https://doi.org/10.1364/JOSAA.4.002366

Bartels, A., & Zeki, S. (2000). The architecture of the colour centre in the human visual brain: new results and a review. *European Journal of*

Neuroscience, *12*(1), 172–193. https://doi.org/10.1046/j.1460-9568.2000.00905.x

Beck, J. M., Ma, W. J., Kiani, R., Hanks, T., Churchland, A. K., Roitman, J., … Pouget, A. (2008). Probabilistic Population Codes for Bayesian Decision Making. *Neuron*, *60*(6), 1142–1152. https://doi.org/10.1016/j.neuron.2008.09.021

Beck, J. M., Ma, W. J., Pitkow, X., Latham, P. E., & Pouget, A. (2012). Not Noisy, Just Wrong: The Role of Suboptimal Inference in Behavioral Variability. *Neuron*, *74*(1), 30–39. https://doi.org/10.1016/j.neuron.2012.03.016

Becker, W. (1989). The neurobiology of saccadic eye movements. Metrics. *Reviews of Oculomotor Research*, *3*, 13–67. https://doi.org/044481017X

Behrens, F., Mackeben, M., & Schröder-Preikschat, W. (2010). An improved algorithm for automatic detection of saccades in eye movement data and for calculating saccade parameters. *Behavior Research Methods*, *42*(3), 701–708. https://doi.org/10.3758/BRM.42.3.701

Bloch, F. (1946). Nuclear induction. *Physical Review*, *70*(7–8), 460–474. https://doi.org/10.1103/PhysRev.70.460

Blum, J., & Price, N. S. C. (2014). Reflexive tracking eye movements and motion perception: one or two neural populations? *Journal of Vision*, *14*(3), 23. https://doi.org/10.1167/14.3.23

Bonhoeffer, T., & Grinvald, A. (1991). Iso-orientation domains in cat visual cortex are arranged in pinwheel-like patterns. *Nature*, *353*(6343), 429–31. https://doi.org/10.1038/353429a0

Born, R., & Bradley, D. (2005). Structure and function of visual area MT. *Annu. Rev. Neurosci.*, *28*, 157–189. https://doi.org/10.1146/annurev.neuro.26.041002.131052

Boynton, G. M. (2005). Imaging orientation selectivity: decoding conscious perception in V1. *Nature Neuroscience*, *8*(5), 541–2. https://doi.org/10.1038/nn0505-541

Boynton, G. M., Engel, S. A., Glover, G. H., & Heeger, D. J. (1996). Linear systems analysis of functional magnetic resonance imaging in human V1. *The Journal of Neuroscience*, *16*(13), 4207–21. Retrieved from http://www.ncbi.nlm.nih.gov/pubmed/8753882

Bradley, A., Skottun, B. C., Ohzawa, I., Sclar, G., & Freeman, R. D. (1987). Visual orientation and spatial frequency discrimination: a comparison of single neurons and behavior. *Journal of Neurophysiology*, *57*(3), 755–72. Retrieved from http://www.ncbi.nlm.nih.gov/pubmed/3559700

Braun, D. I., & Gegenfurtner, K. R. (2016). Dynamics of oculomotor direction discrimination. *Journal of Vision*, *16*(13), 4. https://doi.org/10.1167/16.13.4

Britten, K. H., Newsome, W. T., Shadlen, M. N., Celebrini, S., & Movshon, J. a. (1996). A relationship between behavioral choice and the visual responses of neurons in macaque MT. *Visual Neuroscience*, *13*(1), 87–100. Retrieved from http://www.ncbi.nlm.nih.gov/pubmed/8730992

Britten, K. H., Shadlen, M. N., Newsome, W. T., & Movshon, J. a. (1992). The

analysis of visual motion: a comparison of neuronal and psychophysical performance. *The Journal of Neuroscience*, *12*(12), 4745–4765. https://doi.org/10.1.1.123.9899

Britten, K. H., Shadlen, M. N., Newsome, W. T., & Movshon, J. A. (1993). Responses of neurons in macaque MT to stochastic motion signals. *Visual Neuroscience*, *10*(6), 1157–1169. https://doi.org/10.1017/S0952523800010269

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

Brouwer, G. J., & Heeger, D. J. (2011). Cross-orientation suppression in human visual cortex. *Journal of Neurophysiology*, *106*(5), 2108–2119. https://doi.org/10.1152/jn.00540.2011

Brouwer, G. J., & Heeger, D. J. (2013). Categorical clustering of the neural representation of color. *The Journal of Neuroscience*, *33*(39), 15454–65. https://doi.org/10.1523/JNEUROSCI.2472-13.2013

Brown, S. D., & Heathcote, A. (2005). A ballistic model of choice response time. *Psychological Review*, *112*(1), 117–128. https://doi.org/10.1037/0033-295X.112.1.117

Brown, S. D., & Heathcote, A. (2008). The simplest complete model of choice response time: Linear ballistic accumulation. *Cognitive Psychology*, *57*(3), 153–178. https://doi.org/10.1016/j.cogpsych.2007.12.002

Busemeyer, J. R., & Townsend, J. T. (1993). Decision field theory: a dynamic-cognitive approach to decision making in uncertain environment. *Psychological Review*, *100*(3), 432–459.

Buxton, R. B. (2002). Introduction to Functional Magnetic Resonance Imaging: Principles and Techniques. *Energy*, *24*(2), xi, 523 . https://doi.org/10.1017/CBO9780511605505

Buxton, R. B., Wong, E. C., & Frank, L. R. (1998). Dynamics of blood flow and oxygenation changes during brain activation: The balloon model. *Magnetic Resonance in Medicine*, *39*(6), 855–864. https://doi.org/10.1002/mrm.1910390602

Callaway, E. M. (2005). Structure and function of parallel pathways in the primate early visual system. *The Journal of Physiology*, *566*(Pt 1), 13–9. https://doi.org/10.1113/jphysiol.2005.088047

Campbell, F. W., Cooper, G. F., & Enroth-Cugell, C. (1969). The spatial selectivity of the visual cells of the cat. *The Journal of Physiology*, *203*(1), 223–35. https://doi.org/10.1113/jphysiol.1969.sp008861

Cao, B., Mingolla, E., & Yazdanbakhsh, A. (2015). Tuning Properties of MT and MSTd and Divisive Interactions for Eye-Movement Compensation. *PloS One*, *10*(11), e0142964. https://doi.org/10.1371/journal.pone.0142964

Chalk, M., Seitz, A. R., & Seriès, P. (2010). Rapidly learned stimulus expectations alter perception of motion. *Journal of Vision*, *10*(8), 2. https://doi.org/10.1167/10.8.2

Chen, W., Zhu, X.-H., Thulborn, K. R., & Ugurbil, K. (1999). Retinotopic mapping of lateral geniculate nucleus in humans using functional magnetic resonance imaging. *Proceedings of the National Academy of Sciences* , *96*(5), 2430–2434. https://doi.org/10.1073/pnas.96.5.2430

Churchland, M. M., Priebe, N. J., & Lisberger, S. G. (2005). Comparison of the Spatial Limits on Direction Selectivity in Visual Areas MT and V1 Comparison of the Spatial Limits on Direction Selectivity in Visual Areas MT and V1. *Journal of Neurophysiology*, *93*(October 2004), 1235–1245. https://doi.org/10.1152/jn.00767.2004

Cook, E. P., Guest, J. a, Liang, Y., Masse, N. Y., & Colbert, C. M. (2007). Dendrite-to-Soma Input / Output Function of Continuous Time-Varying Signals in Hippocampal CA1 Pyramidal Neurons. *Journal of Neurophysiology*, (September 2007), 2943–2955. https://doi.org/10.1152/jn.00414.2007.

Curcio, C. A., Sloan, K. R., Kalina, R. E., & Hendrickson, A. E. (1990). Human Photoreceptor Topography. *The Journal of Comparative Neurology*, *523*(292), 497–523. https://doi.org/10.1002/cne.902920402

Dacey, D. M., & Petersen, M. R. (1992). Dendritic field size and morphology of midget and parasol ganglion cells of the human retina. *Proceedings of the National Academy of Sciences of the United States of America*, *89*(20), 9666–70. https://doi.org/10.1073/pnas.89.20.9666

Daniel, B. Y. P. M., & Whitteridge, D. (1961). The Representation of the Visual Field on the Cerebral Cortex in Monkeys. *Journal of Physiology*, *159*, 203–221.

David, S. V., Hayden, B. Y., Mazer, J. A., & Gallant, J. L. (2008). Attention to Stimulus Features Shifts Spectral Tuning of V4 Neurons during Natural Vision. *Neuron*, *59*(3), 509–521. https://doi.org/10.1016/j.neuron.2008.07.001

Dayan, P., & Abbott, L. F. (2005). *Theoretical Neuroscience: Computational And Mathematical Modeling of Neural Systems*. (L. F. Abbott, Ed.) (6th ed.). Massachusetts Institute of Technology Press. Retrieved from https://books.google.co.uk/books?id=hrZYAAAACAAJ

De Haas, B., Schwarzkopf, D. S., Anderson, E. J., & Rees, G. (2014). Perceptual load affects spatial tuning of neuronal populations in human early visual cortex. *Current Biology*. https://doi.org/10.1016/j.cub.2013.11.061

de Lafuente, V., Jazayeri, M., & Shadlen, M. N. (2015). Representation of accumulating evidence for a decision in two parietal areas. *The Journal of Neuroscience : The Official Journal of the Society for Neuroscience*, *35*(10), 4306–18. https://doi.org/10.1523/JNEUROSCI.2451-14.2015

DeAngelis, G. C., & Newsome, W. T. (1999). Organization of disparity-selective neurons in macaque area MT. *J Neurosci*, *19*(4), 1398–1415. Retrieved from http://eutils.ncbi.nlm.nih.gov/entrez/eutils/elink.fcgi?dbfrom=pubmed&id=9952417&retmode=ref&cmd=prlinks

De Valois, R. L., William Yund, E., & Hepler, N. (1982). The orientation and direction selectivity of cells in macaque visual cortex. *Vision Research*,

22(5), 531–544. https://doi.org/10.1016/0042-6989(82)90112-2
Desimone, R., & Gross, C. G. (1979). Visual areas in the temporal cortex of the macaque. *Brain Research*, *178*(2–3), 363–380. https://doi.org/10.1016/0006-8993(79)90699-1
DeYoe, E. A., Carman, G. J., Bandettini, P., Glickman, S., Wieser, J., Cox, R., … Neitz, J. (1996). Mapping striate and extrastriate visual areas in human cerebral cortex. *Proceedings of the National Academy of Sciences of the United States of America*, *93*(6), 2382–6. https://doi.org/10.1073/pnas.93.6.2382
Diogo, A. C. M., Soares, J. G. M., Koulakov, A., Albright, T. D., & Gattass, R. (2003). Electrophysiological imaging of functional architecture in the cortical middle temporal visual area of Cebus apella monkey. *Journal of Neuroscience*, *23*(9), 3881–3898. Retrieved from http://www.ncbi.nlm.nih.gov/pubmed/12736358
Ditchburn, R. W. (1980). The Function of Small Saccades. *Vision Research*, *20*(3), 271–272. https://doi.org/10.1016/0042-6989(80)90112-1
Donner, T. H., Siegel, M., Fries, P., & Engel, A. K. (2009). Buildup of Choice-Predictive Activity in Human Motor Cortex during Perceptual Decision Making. *Current Biology*, *19*(18), 1581–1585. https://doi.org/10.1016/j.cub.2009.07.066
Dorris, M. C., Paré, M., & Munoz, D. P. (1997). Neuronal activity in monkey superior colliculus related to the initiation of saccadic eye movements. *The Journal of Neuroscience*, *17*(21), 8566–8579. Retrieved from http://www.jneurosci.org/cgi/content/full/17/21/8566
Dougherty, R. F., Koch, V. M., Brewer, A. a, Fischer, B., Modersitzki, J., & Wandell, B. a. (2003). Visual field representations and locations of visual areas V1/2/3 in human visual cortex. *Journal of Vision*, *3*(10), 586–598. https://doi.org/10.1167/3.10.1
Draper, D. (1995). Assessment and propagation of model uncertainty. *Journal of the Royal Statistical Society: Series B (Methodological)*. https://doi.org/10.2307/2346087
Dubner, R., & Zeki, S. M. (1971). Response properties and receptive fields of cells in an anatomically defined region of the superior temporal sulcus in the monkey. *Brain Research*, *35*(2), 528–532. https://doi.org/10.1016/0006-8993(71)90494-X
Duda, R. O., Hart, P. E., & Stork, D. G. (2000). Pattern Classification. *New York: John Wiley, Section*. https://doi.org/10.1038/npp.2011.9
Duffy, C. J. (1998). MST neurons respond to optic flow and translational movement. *Journal of Neurophysiology*, *80*(4), 1816–1827.
Dumoulin, S. O., & Wandell, B. a. (2008). Population receptive field estimates in human visual cortex. *NeuroImage*, *39*(2), 647–60. https://doi.org/10.1016/j.neuroimage.2007.09.034
Eagleman, D. M. (2001). Visual illusions and neurobiology. *Nature Reviews. Neuroscience*, *2*(12), 920–926. https://doi.org/10.1038/35104092
Engbert, R., & Mergenthaler, K. (2006). Microsaccades are triggered by low retinal image slip. *Proceedings of the National Academy of Sciences of the United States of America*, *103*(18), 7192–7.

https://doi.org/10.1073/pnas.0509557103
Engel, S. A., Glover, G. G., & Wandell, B. A. (1997). Retintopic organization in human visual cortex and the spatial precision of functional MRI. *Cereb Cortex*, *7*(2), 181–192. https://doi.org/10.1093/cercor/7.2.181
Engel, S. a., Glover, G. H. G., Wandell, B. a., & Glover, G. H. G. (1997). Retintopic organization in human visual cortex and the spatial precision of functional MRI. *Cereb Cortex*, *7*(2), 181–192. https://doi.org/10.1093/cercor/7.2.181
Essen, D. C., & Zeki, S. M. (1978). The topographic organization of rhesus monkey prestriate cortex. *The Journal of Physiology*, *277*, 193–226. https://doi.org/10.1113/jphysiol.1978.sp012269
Fechner, G. T., Adler, H. E., Howes, D. H., & Boring, E. G. (1966). *Elements of psychophysics*. (D. H. Howes & E. G. Boring, Eds.). Holt, Rinehart and Winston. Retrieved from https://books.google.co.uk/books?id=UKW2AAAAIAAJ
Felleman, D. J., & Kaas, J. H. (1984). Receptive-field properties of neurons in middle temporal visual area (MT) of owl monkeys. *Journal of Neurophysiology*, *52*(3), 488–513. Retrieved from http://www.ncbi.nlm.nih.gov/pubmed/6481441
Felleman, D. J., & Van Essen, D. C. (1987). Receptive field properties of neurons in area V3 of macaque monkey extrastriate cortex. *Journal of Neurophysiology*, *57*(4), 889–920. Retrieved from http://www.ncbi.nlm.nih.gov/pubmed/3585463
Felleman, D. J., & Van Essen, D. C. (1991). Distributed hierarchical processing in the primate cerebral cortex. *Cerebral Cortex (New York, N.Y. : 1991)*, *1*(1), 1–47. Retrieved from http://www.ncbi.nlm.nih.gov/pubmed/1822724
Fischer, J., & Whitney, D. (2014). Serial dependence in visual perception. *Nature Neuroscience*, *17*(5), 738–743. https://doi.org/10.1016/j.cub.2014.09.025
Fox, P. T., & Raichle, M. E. (1985). Stimulus rate determines regional brain blood flow in striate cortex. *Annals of Neurology*, *17*(3), 303–305. https://doi.org/10.1002/ana.410170315
Fox, P. T., & Raichle, M. E. (1986). Focal physiological uncoupling of cerebral blood flow and oxidative metabolism during somatosensory stimulation in human subjects. *Proceedings of the National Academy of Sciences* , *83*(4), 1140–1144. https://doi.org/10.1073/pnas.83.4.1140
Freeman, J., Brouwer, G. J., Heeger, D. J., & Merriam, E. P. (2011). Orientation decoding depends on maps, not columns. *The Journal of Neuroscience : The Official Journal of the Society for Neuroscience*, *31*(13), 4792–804. https://doi.org/10.1523/JNEUROSCI.5160-10.2011
Freeman, W. T., & Adelson, E. H. (1991). The Design and Use of Steerable Filters. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, *13*(9), 891–906. https://doi.org/10.1109/34.93808
Friston, K., Chu, C., Mourão-Miranda, J., Hulme, O., Rees, G., Penny, W., & Ashburner, J. (2008). Bayesian decoding of brain images. *NeuroImage*, *39*(1), 181–205. https://doi.org/10.1016/j.neuroimage.2007.08.013

Gallant, J., Nishimoto, S., Naslaris, T., & Wu, M. C. K. (2010). System identification, encoding models, and decoding models, a powerful new approach to fMRI research. *Understanding Visual Population Codes*, 1–26.

Gardner, J. L., Anzai, a, Ohzawa, I., & Freeman, R. D. (1999). Linear and nonlinear contributions to orientation tuning of simple cells in the cat's striate cortex. *Visual Neuroscience*, *16*(6), 1115–1121. https://doi.org/10.1017/S0952523899166112

Garroway, a N., Grannell, P. K., & Mansfield, P. (1974). Image formation in NMR by a selective irradiative process. *Journal of Physics C: Solid State Physics*, *7*(24), L457–L462. https://doi.org/10.1088/0022-3719/7/24/006

Gattass, R., Gross, C. G., & Sandell, J. H. (1981). Visual topography of V2 in the macaque. *The Journal of Comparative Neurology*, *201*(4), 519–539. https://doi.org/10.1002/cne.902010405

Gegenfurtner, K. R., Kiper, D. C., & Levitt, J. B. (1997). Functional Properties of Neurons in Macaque Area V3. *Journal of Neurophysiology*, *77*(4), 1906–1923.

Geisler, W. S. (2003). Ideal Observer Analysis. In L. M. Chalupa & J. S. Werner (Eds.), *The Visual Neurosciences* (Vol. 10, pp. 12–12). The MIT Press. https://doi.org/10.1167/10.7.12

Geisler, W. S., & Albrecht, D. G. (1997). Visual cortex neurons in monkeys and cats: detection, discrimination, and identification. *Visual Neuroscience*, *14*(1997), 897–919. https://doi.org/10.1017/S0952523800011627

Gescheider, G. A. (1997). *Psychophysics: The Fundamentals*. *scan psychology* (Vol. 435).

Glimcher, P. W. (2001). Making choices: The neurophysiology of visual-saccadic decision making. *Trends in Neurosciences*. https://doi.org/10.1016/S0166-2236(00)01932-9

Gold, J. I., & Ding, L. (2013). How mechanisms of perceptual decision-making affect the psychometric function. *Progress in Neurobiology*, *103*, 98–114. https://doi.org/10.1016/j.pneurobio.2012.05.008

Gold, J. I., & Shadlen, M. N. (2007). The neural basis of decision making. *Annual Review of Neuroscience*, *30*, 535–74. https://doi.org/10.1146/annurev.neuro.29.051605.113038

Green, D. M., & Swets, J. A. (1966). Signal detection theory and psychophysics. *Society*, *1*, 521. https://doi.org/10.1901/jeab.1969.12-475

Gregory, R. L. (1972). Cognitive Contours. *Nature*, *238*, 51.

Gregory, R. L. (1973). Eye and brain : the psychology of seeing. *British Medical Journal*, *4*(5893), 682. https://doi.org/10.1136/bmj.4.5893.682-b

Gregory, R. L. (1997). Knowledge in perception and illusion. *Philosophical Transactions of the Royal Society of London. Series B, Biological Sciences*, *352*(1358), 1121–1127. https://doi.org/10.1098/rstb.1997.0095

Gregory, R. L., Gombrich, E. H., & Blakemore, C. (1973). *Illusion in nature and art*. Scribner. Retrieved from https://books.google.co.uk/books?id=zUFRAAAAMAAJ

Groh, J. M., Born, R. T., & Newsome, W. T. (1997). How is a sensory map

read Out? Effects of microstimulation in visual area MT on saccades and smooth pursuit eye movements. *The Journal of Neuroscience : The Official Journal of the Society for Neuroscience*, *17*(11), 4312–4330.
Gur, M., Kagan, I., & Snodderly, D. M. (2005). Orientation and direction selectivity of neurons in V1 of alert monkeys: Functional relationships and laminar distributions. *Cerebral Cortex*, *15*(8), 1207–1221. https://doi.org/10.1093/cercor/bhi003
Hanks, T. D., Ditterich, J., & Shadlen, M. N. (2006). Microstimulation of macaque area LIP affects decision-making in a motion discrimination task. *Nature Neuroscience*, *9*(5), 682–689. https://doi.org/10.1038/nn1683
Hannula, D. E., Simons, D. J., & Cohen, N. J. (2005). Imaging implicit perception: promise and pitfalls. *Nature Reviews. Neuroscience*, *6*(3), 247–55. https://doi.org/10.1038/nrn1630
Harris, J. P., & Fahle, M. (1996). Differences between fovea and periphery in the detection and discrimination of spatial offsets. *Vision Research*, *36*(21), 3469–3477. https://doi.org/10.1016/0042-6989(96)00076-4
Hartigan, J. a., & Hartigan, P. M. (1985). The Dip Test of Unimodality. *The Annals of Statistics*, *13*(1), 70–84. https://doi.org/10.1214/aos/1176346577
Hartigan, P. M., & Hartigan, P. M. (1985). Algorithm AS 217: Computation of the Dip Statistic to Test for Unimodality. *Applied Statistics*, *34*(3), 320–325. https://doi.org/10.2307/2347485
Harvey, B. M., & Dumoulin, S. O. (2011). The relationship between cortical magnification factor and population receptive field size in human visual cortex: constancies in cortical architecture. *J Neurosci*, *31*(38), 13604–13612. https://doi.org/10.1523/JNEUROSCI.2572-11.2011
Haslwanter, T. (1995). Mathematics of three-dimensional eye rotations. *Vision Research*, *35*(12), 1727–1739. https://doi.org/10.1016/0042-6989(94)00257-M
Haynes, J.-D., & Rees, G. (2005). Predicting the orientation of invisible stimuli from activity in human primary visual cortex. *Nature Neuroscience*, *8*(5), 686–91. https://doi.org/10.1038/nn1445
Hecht, S., Shlaer, S., & Pirenne, M. H. (1942). ENERGY, QUANTA, AND VISION. *The Journal of General Physiology*, *25*(6), 819–40. https://doi.org/10.1085/jgp.25.6.819
Heekeren, H. R., Marrett, S., Bandettini, P. a, & Ungerleider, L. G. (2004). A general mechanism for perceptual decision-making in the human brain. *Nature*, *431*(7010), 859–62. https://doi.org/10.1038/nature02966
Hegdé, J., & Van Essen, D. C. (2000). Selectivity for complex shapes in primate visual area V2. *The Journal of Neuroscience : The Official Journal of the Society for Neuroscience*, *20*(5), RC61. Retrieved from http://www.ncbi.nlm.nih.gov/pubmed/10684908
Hendry, S. H. C., & Reid, R. C. (2000). The Koniocellular Pathway in Primate Vision. *Annual Review of Neuroscience*, *23*, 127–153.
Henry, G. H., Dreher, B., & Bishop, P. O. (1974). Orientation specificity of cells in cat striate cortex. *Journal of Neurophysiology*, *37*(6), 1394–1409. Retrieved from papers://27662af5-e24b-497a-8780-

4dcec04ea037/Paper/p4262
Hiris, E., & Blake, R. (1996). Direction repulsion between components in motion transparency. *Visual Neuroscience*, *13*, 187–197. https://doi.org/10.1016/0042-6989(95)00153-0
Hoeting, J. a, Madigan, D., Raftery, a E., & Volinsky, C. T. (1999). Bayesian model averaging: A tutorial. *Statistical Science*, *14*(4), 382–401. https://doi.org/10.2307/2676803
Hogg, R. V, & Craig, A. T. (1995). *Introduction to Mathematical Statistics*. *Technometrics* (Vol. 14). https://doi.org/10.2307/1267313
Hol, K., & Treue, S. (2001). Different populations of neurons contribute to the detection and discrimination of visual motion. *Vision Research*, *41*(6), 685–9. Retrieved from http://www.ncbi.nlm.nih.gov/pubmed/11248258
Holmqvist, K., Nyström, M., Andersson, R., Dewhurst, R., Jarodzka, H., & van de Weijer, J. (2011). *Eye Tracking: A comprehensive guide to methods and measures*. OUP Oxford. Retrieved from https://books.google.co.uk/books?id=5rIDPV1EoLUC
Hubel, D. H., LeVay, S., & Wiesel, T. N. (1975). Mode of termination of retinotectal fibers in macaque monkey: An autoradiographic study. *Brain Research*, *96*(1), 25–40. https://doi.org/10.1016/0006-8993(75)90567-3
Hubel, D. H., & Wiesel, T. N. (1959). Receptive fields of single neurones in the cat's striate cortex. *Journal of Physiology*, *148*, 574–591. https://doi.org/10.1113/jphysiol.2009.174151
Hubel, D. H., & Wiesel, T. N. (1961). Integrative Action in the Cat's Lateral Geniculate Body. *The Journal of Physiology*, *155*, 385–398. https://doi.org/10.1113/jphysiol.1961.sp006635
Hubel, D. H., & Wiesel, T. N. (1962). Receptive fields, binocular interaction and functional architecture in the cat's visual cortex. *The Journal of Physiology*, *160*(1), 106–154.2. https://doi.org/10.1523/JNEUROSCI.1991-09.2009
Hubel, D. H., & Wiesel, T. N. (1963). Shape and arrangement of columns in cat's striate cortex. *The Journal of Physiology*, *165*(3), 559–568. https://doi.org/10.1113/jphysiol.1963.sp007079
Hubel, D. H., & Wiesel, T. N. (1968). Receptive fields and functional architecture of monkey striate cortex. *Journal of Physiology*, 215–243.
Hubel, D. H., & Wiesel, T. N. (1969). Anatomical demonstration of columns in the monkey striate cortex. *Nature*, *221*(5182), 747–750. https://doi.org/10.1038/221747a0
Hubel, D. H., & Wiesel, T. N. (1974). Sequence regularity and geometry of orientation columns in the monkey striate cortex. *Journal of Comparative Neurology*, *158*(3), 267–293. https://doi.org/10.1002/cne.901580304
Hubel, D. H., & Wiesel, T. N. (1977). Ferrier Lecture: Functional Architecture of Macaque Monkey Visual Cortex. *Proceedings of the Royal Society B: Biological Sciences*, *198*(1130), 1–59. https://doi.org/10.1098/rspb.1977.0085

Huettel, S. A., Song, A. W., & McCarthy, G. (2009). *Functional Magnetic Resonance Imaging*. Freeman. Retrieved from https://books.google.co.uk/books?id=BNhMPgAACAAJ

Huk, A. C., Dougherty, R. F., & Heeger, D. J. (2002). Retinotopy and functional subdivision of human areas MT and MST. *The Journal of Neuroscience : The Official Journal of the Society for Neuroscience*, *22*(16), 7195–7205. https://doi.org/20026661

Huk, A. C., & Heeger, D. J. (2002). Pattern-motion responses in human visual cortex. *Nature Neuroscience*, *5*(1), 72–75. https://doi.org/10.1038/nn774

Huk, A. C., Ress, D., & Heeger, D. J. (2001). Neuronal basis of the motion aftereffect reconsidered. *Neuron*, *32*(1), 161–172. https://doi.org/10.1016/S0896-6273(01)00452-4

Huttenlocher, J., Hedges, L. V, & Duncan, S. (1991). Categories and particulars: prototype effects in estimating spatial location. *Psychological Review*, *98*(3), 352–376. https://doi.org/10.1037/0033-295X.98.3.352

Jazayeri, M., & Movshon, J. A. (2006). Optimal representation of sensory information by neural populations. *Nature Neuroscience*, *9*(5), 690–6. https://doi.org/10.1038/nn1691

Jazayeri, M., & Movshon, J. A. (2007a). A new perceptual illusion reveals mechanisms of sensory decoding. *Nature*, *446*(7138), 912–5. https://doi.org/10.1038/nature05739

Jazayeri, M., & Movshon, J. A. (2007b). Integration of sensory evidence in motion discrimination. *J Vis*, *7*(12), 1–7. https://doi.org/10.1167/7.12.7

Jezzard, P., & Clare, S. (2001). Principles of nuclear magnetic resonance and MRI. *Functional MRI: An Introduction to Methods*, 67–93. https://doi.org/10.1002/jmri.10284

Johnston, K., & Everling, S. (2008). Neurophysiology and neuroanatomy of reflexive and voluntary saccades in non-human primates. *Brain and Cognition*, *68*(3), 271–283. https://doi.org/10.1016/j.bandc.2008.08.017

Kandel, E. R., Schwartz, J. H., & Jessell, T. M. (2000). *Principles of Neural Science*. *Neurology* (Vol. 3). https://doi.org/10.1036/0838577016

Kay, K. N. (2012). Understanding Visual Representation by Developing Receptive-Field Models. In *Visual Population Codes : Toward a Common Multivariate Framework for Cell Recording and Functional Imaging* (pp. 133–161). MIT Press.

Kay, K. N., Naselaris, T., Prenger, R. J., & Gallant, J. L. (2008). Identifying natural images from human brain activity. *Nature*, *452*(7185), 352–5. https://doi.org/10.1038/nature06713

Kay, K. N., Weiner, K. S., & Grill-Spector, K. (2015). Attention reduces spatial uncertainty in human ventral temporal cortex. *Current Biology*, *25*(5), 595–600. https://doi.org/10.1016/j.cub.2014.12.050

Kersten, D. (2000). High-Level Vision as Statistical Inference. In Michael S. Gazzaniga & E. Bizzi (Eds.), *The new cognitive neurosciences* (p. 353). MIT Press, 2000. Retrieved from http://books.google.com/books?hl=en&lr=&id=0wx17lC075E

C&oi=fnd&pg=PA353&dq=High-level+Vision+as+Statistical+Inference&ots=RmDOzvL_Rw&sig=vvUwZ0Xhi6DDkH4f3fvDNkWfz3s
Kim, J., & Wilson, H. R. (1996). Direction repulsion between components in motion transparency. *Vision Research*, *36*(8), 1177–1187. https://doi.org/10.1016/0042-6989(95)00153-0
Koch, C., & Segev, I. (2000). The role of single neurons in information processing. *Nature Neuroscience*, *3 Suppl*, 1171–7. https://doi.org/10.1038/81444
Kok, P., Brouwer, G. J., van Gerven, M. a. J., & de Lange, F. P. (2013). Prior Expectations Bias Sensory Representations in Visual Cortex. *Journal of Neuroscience*, *33*(41), 16275–16284. https://doi.org/10.1523/JNEUROSCI.0742-13.2013
Kok, P., Failing, M. F., & de Lange, F. P. (2014). Prior Expectations Evoke Stimulus Templates in the Primary Visual Cortex. *Journal of Cognitive Neuroscience*, *26*(7), 1546–1554. https://doi.org/10.1162/jocn_a_00562
Kolb, H., Linberg, K. a, & Fisher, S. K. (1992). Neurons of the human retina: a Golgi study. *Journal of Comparative Neurology*, *318*, 147–187. https://doi.org/10.1002/cne.903180204
Kolster, H., Peeters, R., & Orban, G. a. (2010). The retinotopic organization of the human middle temporal area MT/V5 and its cortical neighbors. *The Journal of Neuroscience : The Official Journal of the Society for Neuroscience*, *30*(29), 9801–9820. https://doi.org/10.1523/JNEUROSCI.2069-10.2010
Kowler, E. (2011). Eye movements: The past 25years. *Vision Research*, *51*(13), 1457–1483. https://doi.org/10.1016/j.visres.2010.12.014
Kriegeskorte, N. (2011). Pattern-information analysis: From stimulus decoding to computational-model testing. *NeuroImage*, *56*(2), 411–421. https://doi.org/10.1016/j.neuroimage.2011.01.061
Kubo, T., Saika, T., Sakata, Y., Morita, Y., Matsunaga, T., & Kasahara, T. (1991). Analysis of Saccadic Eye Movements Using an Infrared Video System in Human Subjects. *Acta Oto-Laryngologica*, *111*(sup481), 382–387. https://doi.org/10.3109/00016489109131427
Kuffler, S. W. (1953). Discharge Patterns and Functional Organization of Mammalian Retina. *Journal of Neurophysiology*, *16*(1), 37–68.
Larsson, J. (2001). *Imaging Vision: Functional Mapping of Intermediate Visual Processes in Man*. Karolinska institutets bibl. Retrieved from https://books.google.co.uk/books?id=xSMrPQAACAAJ
Larsson, J., & Heeger, D. J. (2006). Two retinotopic visual areas in human lateral occipital cortex. *The Journal of Neuroscience : The Official Journal of the Society for Neuroscience*, *26*(51), 13128–42. https://doi.org/10.1523/JNEUROSCI.1657-06.2006
Lavidor, M., Alexander, T., & McGraw, P. V. (2009). Word recognition processes modulate the naso-temporal asymmetry of the human visual field. *Perception*, *38*(10), 1536–1541. https://doi.org/10.1068/p6078
Lee, C., Rohrer, W., & Sparks, D. (1988). Population coding of saccadic eye

movements by neurons in the superior colliculus. *Nature*, *332*, 357–360. https://doi.org/10.1038/332357a0
Leigh, J., & Zee, D. S. (2013). *The neurology of eye movements*. *Journal of Chemical Information and Modeling* (Vol. 53). https://doi.org/10.1017/CBO9781107415324.004
Lepora, N.F., & Pezzulo, G. (2015). Embodied Choice: How Action Ingluences Perceptual Decision Making. *PLoS Comput Biol 11*(4): e1004110. https://doi.org/10.1371/journal.pcbi.1004110
Leventhal, A. G., Rodieck, R. W., & Dreher, B. (1981). Retinal ganglion cell classes in the Old World monkey: morphology and central projections. *Science*, *213*(4512), 1139–1142. https://doi.org/10.1126/science.7268423
Liberman, A., Zhang, K., & Whitney, D. (2016). Serial dependence promotes object stability during occlusion. *Journal of Vision*, *16*(15), 1–10. https://doi.org/10.1167/16.15.16.doi
Lisberger, S. G., & Ferrera, V. P. (1997). Vector averaging for smooth pursuit eye movements initiated by two moving targets in monkeys. *The Journal of Neuroscience*, *17*(19), 7490–7502.
Lisberger, S. G., Morris, E. J., & Tychsen, L. (1987). Visual motion processing and sensory-motor integration for smooth pursuit eye movements. *Annual Review of Neuroscience*, *10*(1), 97–129. https://doi.org/10.1146/annurev.ne.10.030187.000525
Liu, J., & Newsome, W. T. (2005). Correlation between speed perception and neural activity in the middle temporal visual area. *The Journal of Neuroscience : The Official Journal of the Society for Neuroscience*, *25*(3), 711–22. https://doi.org/10.1523/JNEUROSCI.4034-04.2005
Livingstone, M., & Hubel, D. (1984). Anatomy and physiology of a color system in the primate visual cortex. *The Journal of Neuroscience*, *4*(1), 309–356.
Ljunggren, S. (1983). A simple graphical representation of fourier-based imaging methods. *Journal of Magnetic Resonance*, *54*(2), 338–343. https://doi.org/10.1016/0022-2364(83)90060-4
Lu, H. D., Chen, G., Tanigawa, H., & Roe, A. W. (2010). A Motion Direction Map in Macaque V2. *Neuron*, *68*(5), 1002–1013. https://doi.org/10.1016/j.neuron.2010.11.020
Luce, R. D. (1986). *Response Times: Their Role in Inferring Elementary Mental Organization*. Oxford University Press. Retrieved from https://books.google.co.uk/books?id=h_9i5bSWjHEC
Lynch, J. C. (1992). Saccade initiation and latency deficits after combined lesions of the frontal and posterior eye fields in monkeys. *Journal of Neurophysiology*, *68*(5), 1913–1916.
Lyon, D. C., & Connolly, J. D. (2012). The case for primate V3. *Proceedings. Biological Sciences / The Royal Society*, *279*(1729), 625–33. https://doi.org/10.1098/rspb.2011.2048
Malonek, D., & Grinvald, A. (1996). Interactions between electrical activity and cortical microcirculation revealed by imaging spectroscopy: implications for functional brain mapping. *Science (New York, N.Y.)*, *272*(5261), 551–4. https://doi.org/10.1126/science.272.5261.551

Mansfield, P. (1977). Multi-Planar Image Formation Using NMR Spin Echoes. *Journal of Physics C: Solid State Physics*, *10*, L55–L58. https://doi.org/10.1088/0022-3719/10/3/004
Marmarelis, P. V. Z. (2004). *Nonlinear Dynamic Modeling of Physiological Systems*. John Wiley & Sons. Retrieved from https://books.google.co.uk/books?id=JgW3kL6beBgC
Marmarelis, P. Z., & Marmarelis, V. Z. (1978). *Analysis of Physiological Systems: The White-noise Approach*. Kluwer Academic Pub. Retrieved from https://books.google.co.uk/books?id=esdqAAAAMAAJ
Marshak, W., & Sekuler, R. (1979). Mutual Repulsion Between Moving Targets. *Science*, *205*(28), 1399–1401.
Martinez-Conde, S., Macknik, S. L., & Hubel, D. H. (2004). The role of fixational eye movements in visual perception. *Nature Reviews Neuroscience*, *5*(3), 229–240. https://doi.org/10.1038/nrn1348
Masland, R. H. (2001). The fundamental plan of the retina. *Nature Neuroscience*, *4*(9), 877–886. https://doi.org/10.1038/nn0901-877
Mather, G., & Moulden, B. (1980). A simultaneous shift in apparent direction: further evidence for a "distribution-shift" model of direction coding. *The Quarterly Journal of Experimental Psychology*, *32*(2), 325–333. https://doi.org/10.1080/14640748008401168
Maunsell, J. H., & Newsome, W. T. (1987). Visual processing in monkey extrastriate cortex. *Annual Review of Neuroscience*, *10*, 363–401. https://doi.org/10.1146/annurev.neuro.10.1.363
Maunsell, J. H., & van Essen, D. C. (1983). The connections of the middle temporal visual area (MT) and their relationship to a cortical hierarchy in the macaque monkey. *The Journal of Neuroscience : The Official Journal of the Society for Neuroscience*, *3*(12), 2563–86. Retrieved from http://www.ncbi.nlm.nih.gov/pubmed/6655500
McConkie, G. W., Kerr, P. W., Reddix, M. D., & Zola, D. (1988). Eye movement control during reading: I. The location of initial eye fixations on words. *Vision Research*, *28*(10). https://doi.org/10.1016/0042-6989(88)90137-X
McSorley, E., & McCloy, R. (2009). Saccadic eye movements as an index of perceptual decision-making. *Experimental Brain Research*, *198*(4), 513–520. https://doi.org/10.1007/s00221-009-1952-9
Mikami, a, Newsome, W. T., & Wurtz, R. H. (1986). Motion selectivity in macaque visual cortex. I. Mechanisms of direction and speed selectivity in extrastriate area MT. *Journal of Neurophysiology*, *55*(6), 1308–1327.
Milner, A. D., & Goodale, M. A. (2008). Two visual systems re-viewed. *Neuropsychologia*, *46*(3), 774–785. https://doi.org/10.1016/j.neuropsychologia.2007.10.005
Mohler, C. W., & Wurtz, R. H. (1977). Role of striate cortex and superior colliculus in visual guidance of saccadic eye movements in monkeys. *Journal of Neurophysiology*, *40*(1), 74–94.
Moran, J., & Desimone, R. (1985). Selective attention gates visual processing in the extrastriate cortex. *Science (New York, N.Y.)*,

229(4715), 782–784. https://doi.org/10.1126/science.4023713
Mountcastle, V. B. (1957). Modality and topographic properties of single neurons of cat's somatic sensory cortex. *Journal of Neurophysiology*.
Movshon, J. A., Adelson, E. H., Gizzi, M. S., & Newsome, W. T. (1985). *The analysis of moving visual patterns. Pontificiae Academiae Scientiarum Scripta Varia* (Vol. 54). https://doi.org/10.1098/rstb.1998.0333
Naselaris, T., Kay, K. N., Nishimoto, S., & Gallant, J. L. (2011). Encoding and decoding in fMRI. *NeuroImage*, *56*(2), 400–410. https://doi.org/10.1016/j.neuroimage.2010.07.073
Naselaris, T., Prenger, R. J., Kay, K. N., Oliver, M., & Gallant, J. L. (2009). Bayesian Reconstruction of Natural Images from Human Brain Activity. *Neuron*, *63*(6), 902–915. https://doi.org/10.1016/j.neuron.2009.09.006
Nelder, J. a., Mead, R., Nelder, B. J. a, & Mead, R. (1964). A simplex method for function minimization. *Computer Journal*, *7*(4), 308–313. https://doi.org/10.1093/comjnl/7.4.308
Nestares, O., & Heeger, D. J. (2000). Robust multiresolution alignment of MRI brain volumes. *Magnetic Resonance in Medicine*, *43*(5), 705–715. https://doi.org/10.1002/(SICI)1522-2594(200005)43:5<705::AID-MRM13>3.0.CO;2-R
Newsome, W. (1989). Neuronal correlates of a perceptual decision. *Nature*, *341*, 52–54.
Nichols, M. J., & Newsome, W. T. (2002). Middle Temporal Visual Area Microstimulation Influences Veridical Judgments of Motion Direction. *The Journal of Neuroscience*, *22*(21), 9530–9540. https://doi.org/22/21/9530 [pii]
Nishimoto, S., Ishida, T., & Ohzawa, I. (2006). Receptive field properties of neurons in the early visual cortex revealed by local spectral reverse correlation. *Journal of Neuroscience*, *26*(12), 3269–80. https://doi.org/10.1523/JNEUROSCI.4558-05.2006
Ogawa, S., Lee, T. M., Kay, A. R., & Tank, D. W. (1990). Brain magnetic resonance imaging with contrast dependent on blood oxygenation. *Proceedings of the National Academy of Sciences*, *87*(24), 9868–9872. https://doi.org/10.1073/pnas.87.24.9868
Ogawa, S., Tank, D. W., Menon, R., Ellermann, J. M., Kim, S. G., Merkle, H., & Ugurbil, K. (1992). Intrinsic signal changes accompanying sensory stimulation: Functional brain mapping with magnetic resonance imaging. *Proceedings of the National Academy of Sciences of the United States of America*, *89*(13), 5951–5955. https://doi.org/10.1073/pnas.89.13.5951
Olman, C. A., Ugurbil, K., Schrater, P., & Kersten, D. (2004). BOLD fMRI and psychophysical measurements of contrast response to broadband images. *Vision Research*, *44*(7), 669–683. https://doi.org/10.1016/j.visres.2003.10.022
Oohira, A., Goto, K., & Ozawa, T. (1983). Vertical and oblique saccadic eye movements. *Japanese Journal of Ophthalmology*, *27*(4), 631–646.
Optican, L. M. (2005). Sensorimotor transformation for visually guided

saccades. In *Annals of the New York Academy of Sciences* (Vol. 1039, pp. 132–148). https://doi.org/10.1196/annals.1325.013
Østerberg, G. (1935). *Topography of the Layer of Rods and Cones in the Human Retina*. A. Busck. Retrieved from https://books.google.co.uk/books?id=0NEVQgAACAAJ
Parker, A. J. (2013). A micro-pool model for decision-related signals in visual cortical areas. *Frontiers in Computational Neuroscience*, *7*, 115. https://doi.org/10.3389/fncom.2013.00115
Purcell, E., Torrey, H., & Pound, R. (1946). Resonance Absorption by Nuclear Magnetic Moments in a Solid. *Physical Review*, *69*(1–2), 37–38. https://doi.org/10.1103/PhysRev.69.37
Purushothaman, G., & Bradley, D. C. (2005). Neural population code for fine perceptual decisions in area MT. *Nature Neuroscience*, *8*(1), 99–106. https://doi.org/10.1038/nn1373
Purves, D., Augustine, G. J., Fitzpatrick, D., Hall, W. C., Lamantia, A.-S., Mcnamara, J. O., & Willians, S. M. (2004). *Neuroscience. Sunderland* (Vol. 3). https://doi.org/978-0878937257
Putzeys, T., Bethge, M., Wichmann, F., Wagemans, J., & Goris, R. (2012). A new perceptual bias reveals suboptimal population decoding of sensory responses. *PLoS Computational Biology*, *8*(4). https://doi.org/10.1371/journal.pcbi.1002453
Ratcliff, R., Smith, P. L., Brown, S. D., & McKoon, G. (2016). Diffusion Decision Model: Current Issues and History. *Trends in Cognitive Sciences*, *20*(4), 260–281. https://doi.org/10.1016/j.tics.2016.01.007
Rauber, H.-J., & Treue, S. (1998). Reference repulsion when judging the direction of visual motion. *Perception*, *27*, 393–402.
Rauber, H. J., & Treue, S. (1999). Revisiting motion repulsion: Evidence for a general phenomenon? *Vision Research*, *39*(19), 3187–3196. https://doi.org/10.1016/S0042-6989(99)00025-5
Rees, G., Friston, K., & Koch, C. (2000). A direct quantitative relationship between the functional properties of human and macaque V5. *Nature Neuroscience*, *3*(7), 716–23. https://doi.org/10.1038/76673
Regan, D., & Beverley, K. I. (1985). Postadaptation orientation discrimination. *Journal of the Optical Society of America. A, Optics and Image Science*, *2*(2), 147–55. Retrieved from http://www.ncbi.nlm.nih.gov/pubmed/3973752
Richard, A., Nystrom, M., & Kenneth, H. (2010). Sampling frequency and eye-tracking measures: how speed affects durations, latencies, and more. *Journal of Eye Movement Research*, *3*(3), 1–12. https://doi.org/10.1145/1344471.1344500
Riggs, L. A., Ratliff, F., Cornsweet, J. C., & Cornsweet, T. N. (1953). The disappearance of steadily fixated visual test objects. *Journal of the Optical Society of America*, *43*(6), 495–501. https://doi.org/10.1364/JOSA.43.000495
Rocchi, F., Ledgeway, T., & Webb, B. S. (2013). Visual motion integration is mediated by directional ambiguities in local motion signals. *Frontiers in Computational Neuroscience*, *7*, 167.

https://doi.org/10.3389/fncom.2013.00167
Roe, A. W., & Tso, D. Y. (1995). Visual Topography in Primate V2 - Multiple Representation across Functional Stripes. *Journal of Neuroscience*, *15*(5), 3689–3715.
Romo, R., Hernández, A., Zainos, A., Lemus, L., & Brody, C. D. (2002). Neuronal correlates of decision-making in secondary somatosensory cortex. *Nature Neuroscience*, *5*(11), 1217–25. https://doi.org/10.1038/nn950
Romo, R., & Salinas, E. (2003). Flutter discrimination: neural codes, perception, memory and decision making. *Nature Reviews. Neuroscience*, *4*, 203–218.
Roy, C., & Sherrington, C. (1890). On the Regulation of the Blood-Supply of the Brain. *The Journal of Physiology*, *11*(1–2), 85–158. https://doi.org/10.1186/1471-244X-11-115
Rust, N. C., Mante, V., Simoncelli, E. P., & Movshon, J. A. (2006). How MT cells analyze the motion of visual patterns. *Nature Neuroscience*, *9*(11), 1421–1431. https://doi.org/10.1038/nn1786
Salzman, C. D., & Newsome, W. T. (1994). Neural Mechanisms for Forming a Perceptual Decision. *Science*, *264*(36), 231–237. https://doi.org/10.1126/science.8146653
Salzman, D. (1990). Cortical microstimulation influences perceptual judgments of motion direction. *Nature*, *346*(6280), 174–177.
Salzman, D. (1992). Microstimulation in Visual Area MT: Effects on Direction Discrimination Performance. *The Journal of Neuroscience*, *12*(6), 2331–2355. Retrieved from http://salzmanlab.neuroscience.columbia.edu/salzmanetal1992.pdf
Salzman, D., Britten, K. H., & Newsome, W. T. (1990). Cortical microstimulation influences perceptual judgments of motion direction. *Nature*, *346*, 174–177.
Schein, S. J., & de Monasterio, F. M. (1987). Mapping of retinal and geniculate neurons onto striate cortex of macaque. *The Journal of Neuroscience : The Official Journal of the Society for Neuroscience*, *7*(4), 996–1009.
Schneider, K. A., Richter, M. C., & Kastner, S. (2004). Retinotopic organization and functional subdivisions of the human lateral geniculate nucleus: A high-resolution functional magnetic resonance imaging study. *The Journal of Neuroscience*, *24*(41), 8975–8985. https://doi.org/10.1523/JNEUROSCI.2413-04.2004
Schonhoff, T. A., & Giordano, A. A. (2006). *Detection and Estimation Theory and Its Applications*. Prentice Hall.
Scolari, M., Byers, A., & Serences, J. T. (2012). Optimal Deployment of Attentional Gain during Fine Discriminations. *Journal of Neuroscience*, *32*(22), 7723–7733. https://doi.org/10.1523/JNEUROSCI.5558-11.2012
Scolari, M., & Serences, J. T. (2010). Basing perceptual decisions on the most informative sensory neurons. *Journal of Neurophysiology*, *104*(4), 2266–73. https://doi.org/10.1152/jn.00273.2010
Serences, J. T., & Saproo, S. (2012). Computational advances towards

linking BOLD and behavior. *Neuropsychologia*, *50*(4), 435–446. https://doi.org/10.1016/j.neuropsychologia.2011.07.013
Sereno, M. I., & Allman, J. M. (1991). Cortical visual areas in mammals. *The Neural Basis of Visual Function*.
Sereno, M. I., Dale, A. M., Reppas, J. B., Kwong, K. K., Belliveau, J. W., Brady, T. J., … Tootell, R. B. (1995). Borders of multiple visual areas in humans revealed by functional magnetic resonance imaging. *Science*, *268*(5212), 889–893. https://doi.org/10.1126/science.7754376
Shadlen, M. N., Britten, K. H., Newsome, W. T., & Movshon, J. a. (1996). A computational analysis of the relationship between neuronal and behavioral responses to visual motion. *The Journal of Neuroscience : The Official Journal of the Society for Neuroscience*, *16*(4), 1486–510.
Shadlen, M. N., Britten, K. H., Newsome, W. T., & Movshon, J. A. (1996). A Computational Analysis of the Relationship between Neuronal and Behavioral Responses to Visual Motion. *The Journal of Neuroscience*, *16*(4), 1486–1510.
Shadlen, M. N., & Newsome, W. T. (1996). Motion perception: seeing and deciding. *Proceedings of the National Academy of Sciences of the United States of America*, *93*(2), 628–33. Retrieved from http://www.pubmedcentral.nih.gov/articlerender.fcgi?artid=40102&tool=pmcentrez&rendertype=abstract
Shadlen, M. N., & Newsome, W. T. (2001). Neural basis of a perceptual decision in the parietal cortex (area LIP) of the rhesus monkey. *Journal of Neurophysiology*, *86*(4), 1916–36. Retrieved from http://www.ncbi.nlm.nih.gov/pubmed/11600651
Shipp, S., Watson, J. D., Frackowiak, R. S., & Zeki, S. (1995). Retinotopic maps in human prestriate visual cortex: the demarcation of areas V2 and V3. *NeuroImage*. https://doi.org/10.1006/nimg.1995.1015
Siebert, W. M. (1965). Some implications of the stochastic behavior of primary auditory neurons. *Kybernetik*, *2*(5), 206–215. https://doi.org/10.1007/BF00306416
Siegel, M., Engel, A. K., & Donner, T. H. (2011). Cortical network dynamics of perceptual decision-making in the human brain. *Frontiers in Human Neuroscience*, *5*, 21. https://doi.org/10.3389/fnhum.2011.00021
Simion, F., Valenza, E., Umilta, C., & Dalla Barba, B. (1998). Preferential Orienting to Faces in Newborns: A Temporal-Nasal Asymmetry. *Journal of Experimental Psychology: Human Perception and Performance*, *24*(5), 1399–1405. https://doi.org/10.1037/0096-1523.24.5.1399
Simoncelli, E. P., & Heeger, D. J. (1998). A model of neuronal responses in visual area MT. *Vision Research*, *38*(5), 743–761. https://doi.org/10.1016/S0042-6989(97)00183-1
Singh, K. D., Smith, A. T., & Greenlee, M. W. (2000). Spatiotemporal frequency and direction sensitivities of human visual areas measured using fMRI. *Neuroimage*, *12*(5), 550–564. https://doi.org/10.1006/nimg.2000.0642
Smith, A. T., Singh, K. D., Williams, A. L., & Greenlee, M. W. (2001).

Estimating receptive field size from fMRI data in human striate and extrastriate visual cortex. *Cerebral Cortex*, *11*(12), 1182–90. https://doi.org/10.1093/cercor/11.12.1182
Smith, P. L., & Ratcliff, R. (2004). Psychology and neurobiology of simple decisions. *Trends in Neurosciences*, *27*(3), 161–8. https://doi.org/10.1016/j.tins.2004.01.006
Sparks, D. L. (2002). The brainstem control of saccadic eye movements. *Nature Reviews. Neuroscience*, *3*(12), 952–964. https://doi.org/10.1038/nrn986 [doi]; nrn986 [pii]
Sprague, T. C., Ester, E. F., Serences, J. T., Albers, A. M., Kok, P., Toni, I., … Luck, S. J. (2016). Restoring Latent Visual Working Memory Representations in Human Cortex. *Neuron*, *91*(3), 694–707. https://doi.org/10.1016/j.neuron.2016.07.006
Stocker, A. A., & Simoncelli, E. P. (2008). A Bayesian Model of Conditioned Perception. In *NIPS Proceedings*.
Tarantola, A. (2006). Popper, Bayes and the inverse problem. *Nature Physics*, *2*(8), 492–494. https://doi.org/10.1038/nphys375
Toga, A. W., & Mazziotta, J. C. (2002). *Brain Mapping: The Methods*. Elsevier Science. Retrieved from https://books.google.co.uk/books?id=6FjVzbll7NAC
Tootell, R. B. H., Mendola, J. D., Hadjikhani, N. K., Ledden, P. J., Liu, A. K., Reppas, J. B., … Dale, A. M. (1997). Functional Analysis of V3A and Related Areas in Human Visual Cortex. *The Journal of Neuroscience*, *17*(18), 7060–7078.
Tootell, R. B., & Hadjikhani, N. (2001). Where is "dorsal V4" in human visual cortex? Retinotopic, topographic and functional evidence. *Cerebral Cortex (New York, N.Y. : 1991)*, *11*(4), 298–311. https://doi.org/DOI 10.1093/cercor/11.4.298
Townsend, J. T., & Ashby, F. G. (1983). *Stochastic Modeling of Elementary Psychological Processes*. Cambridge University Press. Retrieved from https://books.google.co.uk/books?id=pYw6AAAAIAAJ
Ungerleider, L. G., & Desimone, R. (1986). Cortical connections of visual area MT in the macaque. *The Journal of Comparative Neurology*, *248*(2), 190–222. https://doi.org/10.1002/cne.902480204
Ungerleider, L. G., & Mishkin, M. (1982). Two cortical visual systems. In *Analysis of Visual Behavior* (pp. 549–586). https://doi.org/10.2139/ssrn.1353746
Usher, M., & McClelland, J. L. (2001). The time course of perceptual choice: the leaky, competing accumulator model. *Psychological Review*. https://doi.org/10.1037/0033-295X.108.3.550
van Beers, R. J. (2007). The sources of variability in saccadic eye movements. *Journal of Neuroscience*, *27*(33), 8757–8770. https://doi.org/10.1523/JNEUROSCI.2311-07.2007
Van Essen, D. C., & Gallant, J. L. (1994). Neural mechanism of form and motion processing in the primate visual system. *Neuron*, *13*(94), 1–10.
Van Essen, D. C., Newsome, W. T., & Maunsell, J. H. R. (1984). The visual field representation in striate cortex of the macaque monkey:

Asymmetries, anisotropies, and individual variability. *Vision Research*, *24*(5), 429–448. https://doi.org/10.1016/0042-6989(84)90041-5
Van Essen, D., Maunsell, J., & Bixby, J. (1981). The middle temporal visual area in the macaque: myeloarchitecture, connections, functional properties and topographic organization. *The Journal of Comparative Neurology*, *199*(3), 293–326. https://doi.org/10.1002/cne.901990302
Vickers, D., Caudrey, D., & Willson, R. (1971). Discriminating between the frequency of occurrence of two alternative events. *Acta Psychologica*, *35*(2), 151–172. https://doi.org/10.1016/0001-6918(71)90018-7
Vincent, B. T. (2015). A tutorial on Bayesian models of perception. *Journal of Mathematical Psychology*, *66*, 103–114. https://doi.org/10.1016/j.jmp.2015.02.001
Vo, V., Sprague, T. C., & Serences, J. T. (2017). Spatial tuning shifts increase the discriminability and fidelity of population codes in visual cortex. *The Journal of Neuroscience*, *37*(12), 3386–3401. https://doi.org/10.1523/JNEUROSCI.3484-16.2017
Wandell, B. a. (1999). Computational neuroimaging of human visual cortex. *Annual Review of Neuroscience*, *22*(FEBRUARY 1999), 145–173. https://doi.org/10.1146/annurev.neuro.22.1.145
Wandell, B. A. (1995). *Foundations of vision*. *Diabetes*. https://doi.org/10.1039/c1pp90008k
Wandell, B. A., Dumoulin, S. O., & Brewer, A. A. (2007). Visual field maps in human cortex. *Neuron*, *56*(2), 366–83. https://doi.org/10.1016/j.neuron.2007.10.012
Wandell, B., Dumoulin, S. O. S. O., & Brewer, A. A. A. a. (2007). Visual Field Maps in Human Cortex. *Neuron*, *56*(2), 366–383. https://doi.org/10.1016/j.neuron.2007.10.012
Watson, J., Myers, R., Frackowiak, R. S., Hajnal, J. V, Woods, R., Mazziotta, J., … Zeki, S. (1993). Area V5 of the human brain: evidence from a combined study using positron emission tomography and magnetic resonance imaging. *Cerebral Cortex*, *3*(2), 79–94.
Webb, B. S., Ledgeway, T., & McGraw, P. V. (2007). Cortical pooling algorithms for judging global motion direction. *Proceedings of the National Academy of Sciences of the United States of America*, *104*(9), 3532–7. https://doi.org/10.1073/pnas.0611288104
Wickens, T. D. (2002). *Elementary signal detection theory*. *New York* (Vol. 1). https://doi.org/10.1093/acprof:oso/9780195092509.001.0001
Wiesel, T. N., & Hubel, D. H. (1966). Spatial and chromatic interactions in the lateral geniculate body of the rhesus monkey. *Journal of Neurophysiology*, 1115–1156.
Wu, M. C.-K., David, S. V, & Gallant, J. L. (2006). Complete functional characterization of sensory neurons by system identification. *Annual Review of Neuroscience*, *29*, 477–505. https://doi.org/10.1146/annurev.neuro.29.051605.113024
Yacoub, E., Harel, N., & Ugurbil, K. (2008). High-field fMRI unveils orientation columns in humans. *Proceedings of the National Academy of Sciences of the United States of America*, *105*(30), 10607–12.

https://doi.org/10.1073/pnas.0804110105
Zamboni, E., Ledgeway, T., McGraw, P. V., & Schluppeck, D. (2016). Do perceptual biases emerge early or late in visual processing? Decision-biases in motion perception. *Proceedings of the Royal Society B: Biological Sciences*, *283*(1833). https://doi.org/10.1098/rspb.2016.0263
Zeki, S. (1974). Funcitonal organisation of a visual area in the posterior bank of the superior temporal sulcus of the rhesus monkey. *Journal of Physiology*, *236*, 549–573.
Zeki, S. (1983). The distribution of wavelength and orientation selective cells in different areas of monkey visual cortex. *Proceedings of the Royal Society of London. Series B, Biological Sciences*, *217*(1209), 449–70. https://doi.org/10.1098/rspb.1983.0020
Zeki, S. M. (1978). The third visual complex of rhesus monkey prestriate cortex. *The Journal of Physiology*. https://doi.org/10.1113/jphysiol.1978.sp012271
Zhang, K., Ginzburg, I., McNaughton, B. L., & Sejnowski, T. J. (1998). Interpreting neuronal population activity by reconstruction: unified framework with application to hippocampal place cells. *Journal of Neurophysiology*, *79*(2), 1017–44. https://doi.org/10.1093/cercor/6.3.406

10 Appendix – Supplementary Material

CHAPTER 3

Table 3.1 – Summary of Dip Statistics Values: Reference Present / Absent

		Subject a	Subject c	Subject d	Subject e	Subject f
Coherence 0.04	Present	$d = 0.056$	$d = 0.048$	$d = 0.087$	$d = 0.097$	$d = 0.063$
		$p < 0.001$	$p = 0.012$	$p < 0.001$	$p < 0.001$	$p < 0.001$
	Absent	$d = 0.026$	$d = 0.029$	$d = 0.036$	$d = 0.025$	$d = 0.029$
		$p = 0.51$	$p = 0.469$	$p = 0.113$	$p = 0.713$	$p = 0.482$
Coherence 0.07	Present	$d = 0.044$	$d = 0.063$	$d = 0.091$	$d = 0.096$	$d = 0.082$
		$p < 0.001$	$p < 0.001$	$p < 0.001$	$p < 0.001$	$p < 0.001$
	Absent	$d = 0.03$	$d = 0.046$	$d = 0.042$	$d = 0.029$	$d = 0.031$
		$p = 0.294$	$p = 0.002$	$p = 0.028$	$p = 0.332$	$p = 0.394$
Coherence 0.13	Present	$d = 0.051$	$d = 0.055$	$d = 0.088$	$d = 0.091$	$d = 0.074$
		$p < 0.001$	$p < 0.001$	$p < 0.001$	$p < 0.001$	$p < 0.001$
	Absent	$d = 0.02$	$d = 0.034$	$d = 0.029$	$d = 0.036$	$d = 0.032$
		$p = 0.861$	$p = 0.161$	$p = 0.399$	$p = 0.128$	$p = 0.454$
Coherence 0.25	Present	$d = 0.051$	$d = 0.035$	$d = 0.068$	$d = 0.098$	$d = 0.065$
		$p < 0.001$	$p = 0.060$	$p < 0.001$	$p < 0.001$	$p < 0.001$
	Absent	$d = 0.034$	$d = 0.039$	$d = 0.057$	$d = 0.026$	$d = 0.024$
		$p = 0.124$	$p = 0.032$	$p < 0.001$	$p = 0.562$	$p = 0.818$

Table 10.1 Summary of Dip Statistics for four coherence levels, two reference conditions (present / absent) five subjects tested. Highlighted cells indicate statistically significant results (for $p < 0.05$).

Figure 10.1 – Summary of Direction Estimates for Reference Present Condition

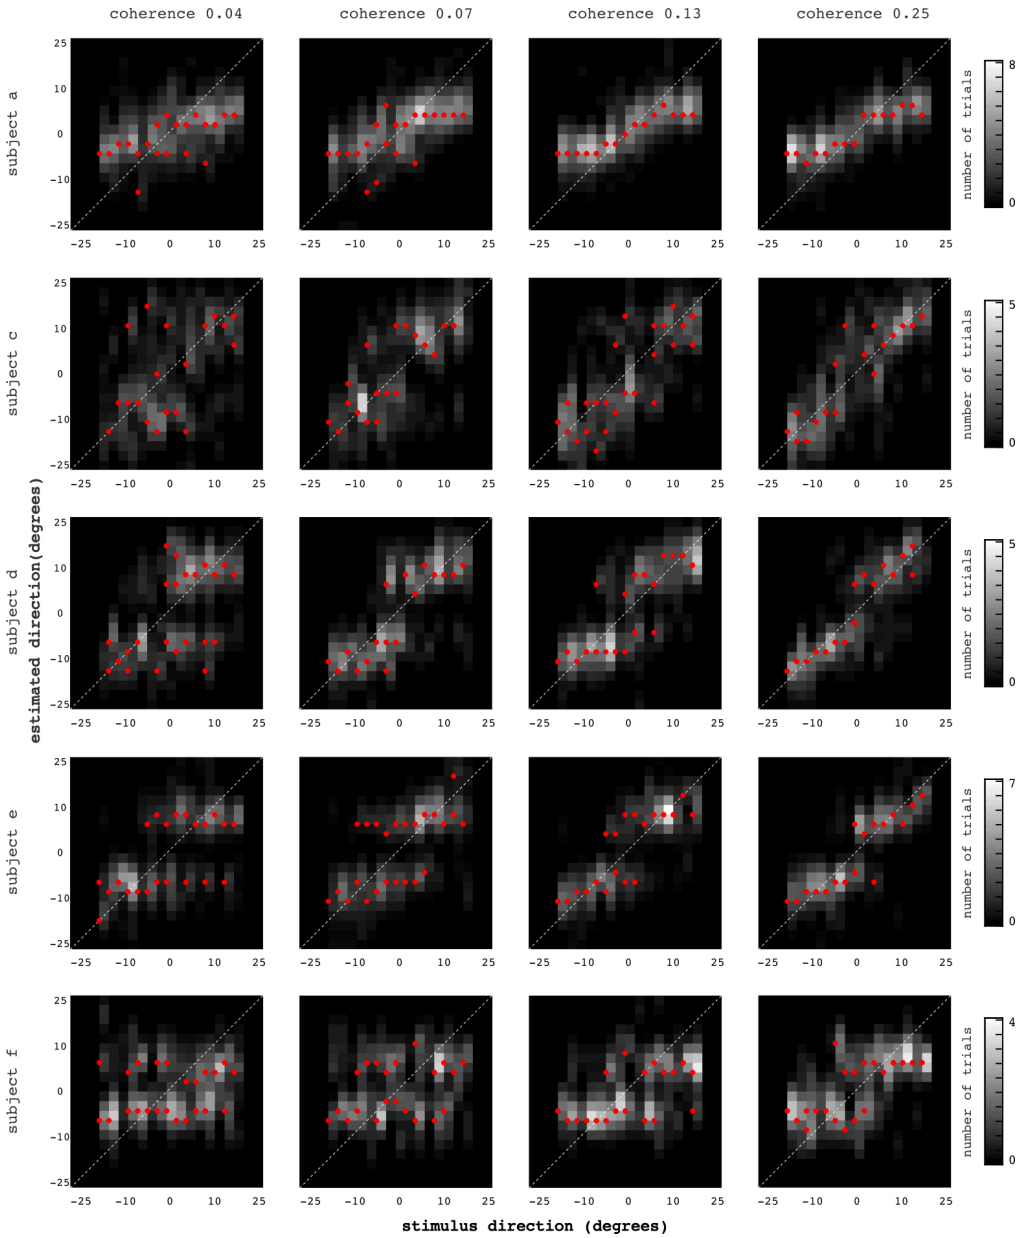


Figure 10.1 Data from the estimation task for each subject (rows), each coherence level tested (columns), reference present condition. Density plots as described in Chapter 3, Figures 3.4 and 3.7. Red dots indicate peaks of column-wise distributions of estimates for each direction of motion presented. White dashed line represents the identity line.

Figure 10.2 – Summary of Fine Discrimination Task Performance

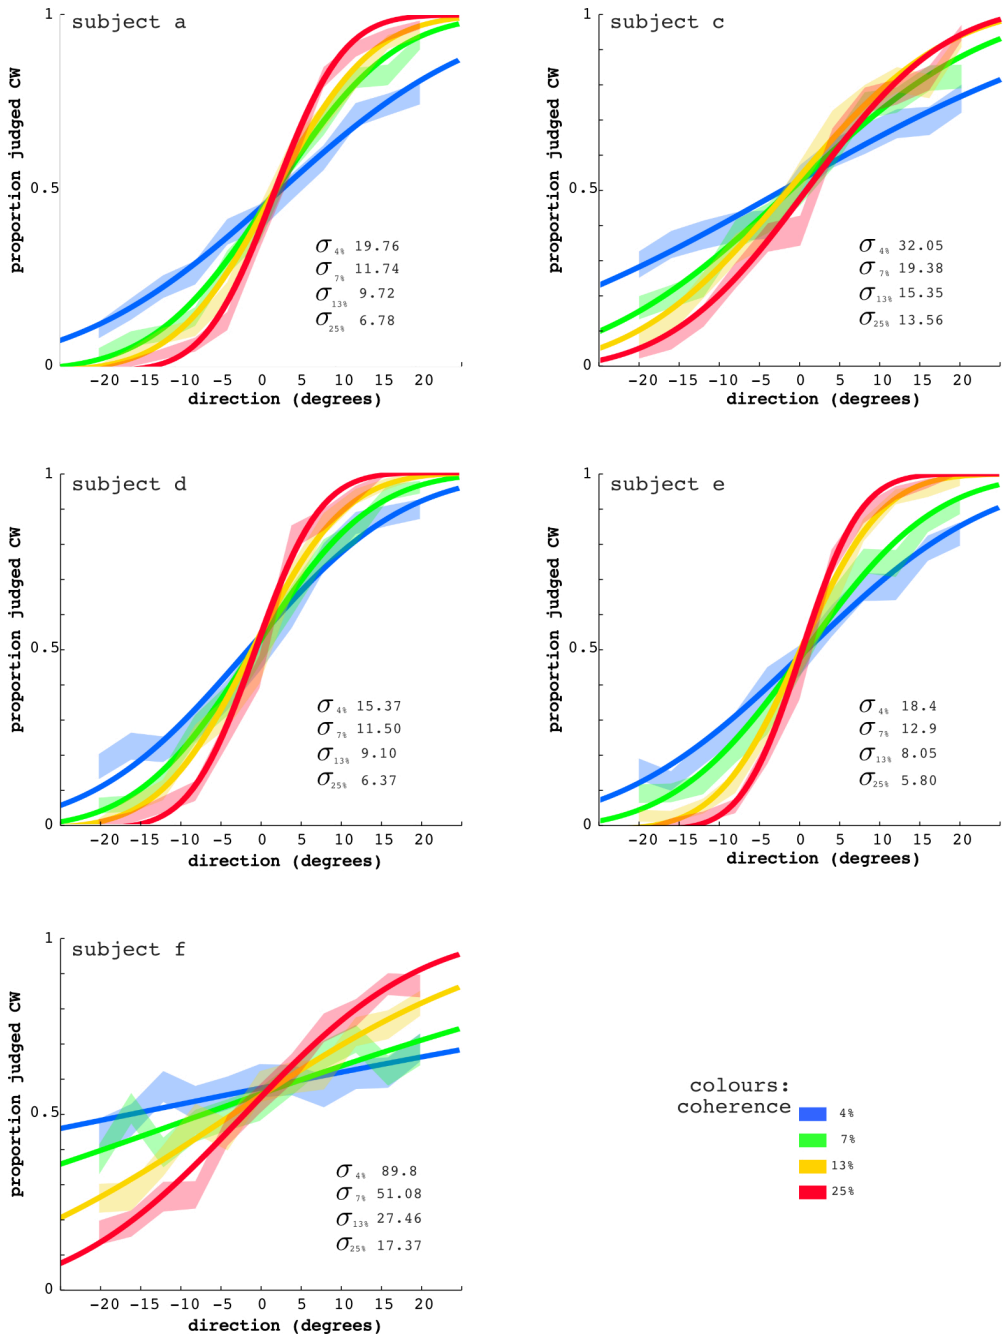


Figure 10.2 Proportion of CW judgements plotted as a function of stimulus direction. Data from five subjects (five panels), four coherence levels (colours), as described in Chapter 3 (Figures 3.2 and 3.6).

Figure 10.3 – Summary of Psychometric Functions Slopes

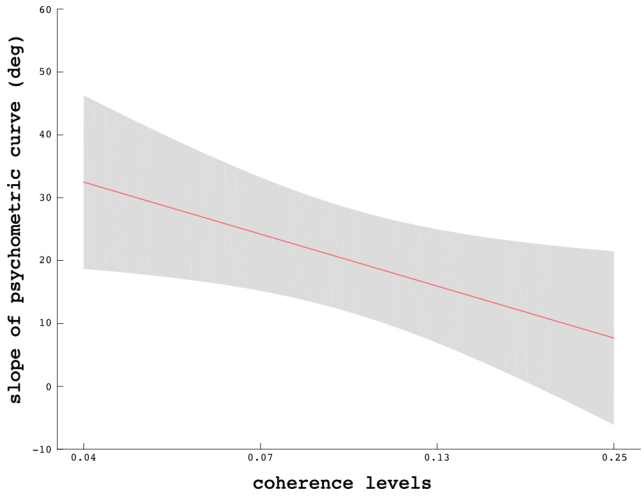


Figure 10.3 Trend of psychometric curves slopes. As coherence (x-axis) increases, the discrimination performance increases (the slope of the curves decreases). Red solid line indicates regression across subjects (N=5), grey shaded area represents s.e.m.

Figure 10.4 – Summary of Direction Estimates for Reference Absent Condition

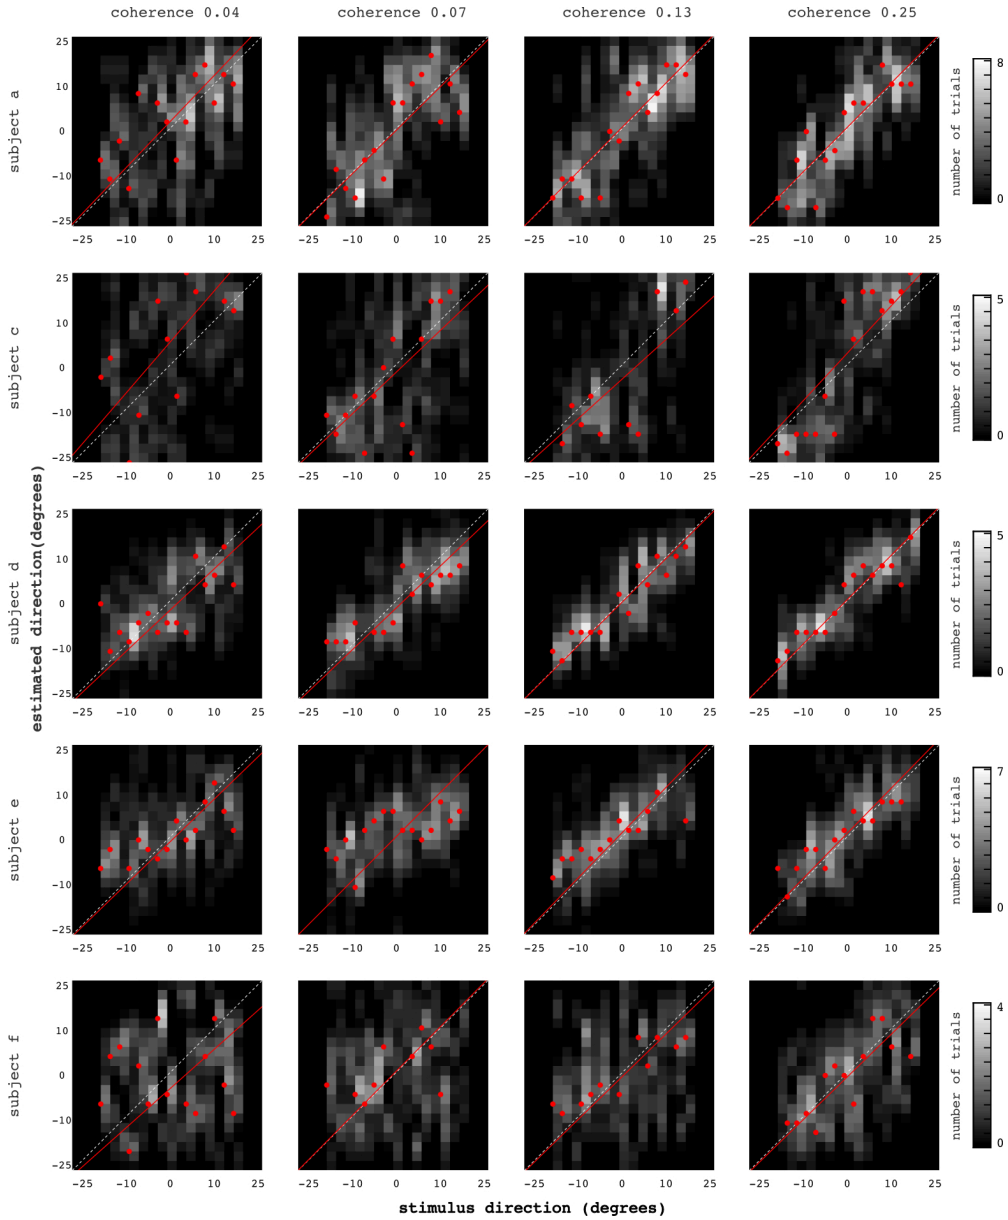


Figure 10.4 Data from the estimation task for each subject (rows), each coherence level (columns), reference absent condition. Density plots as described in Chapter 3 (Figures 3.4 and 3.7). Red dots indicate peaks of column-wise distributions of estimates for each direction of motion presented. Red solid line shows linear regression of these peaks, white dashed line represents identity.

Table 4.2 – Summary of Coefficient of Determination for Reference Shift Condition

		4%	7%	13%	25%
Subject a	-6°	0.74	0.85	0.83	0.83
	0°	0.81	0.78	0.85	0.83
	+6°	0.65	0.75	0.81	0.79
Subject b	-6°	0.83	0.81	0.82	0.88
	0°	0.78	0.77	0.84	0.88
	+6°	0.75	0.78	0.84	0.84
Subject d	-6°	0.90	0.88	0.91	0.91
	0°	0.86	0.91	0.86	0.87
	+6°	0.85	0.88	0.89	0.89
Subject e	-6°	0.82	0.82	0.86	0.87
	0°	0.88	0.87	0.89	0.88
	+6°	0.72	0.83	0.86	0.91
Subject g	-6°	0.64	0.66	0.70	0.78
	0°	0.71	0.67	0.72	0.73
	+6°	0.68	0.69	0.70	0.77

Table 10.2 Summary of Coefficient of Determination. Five subjects (rows), three reference shifts (rows), four coherence levels tested (columns). Colours refer to reference condition as in Chapter 3, Section 3.3.1

Table 4.3 – Summary of Estimated Reference Shifts

		4%	7%	13%	25%
Subject a	-6°	-3.29	-3.95	-3.81	-4.56
	0°	-0.53	-0.47	0.24	-0.36
	+6°	2.98	4.25	3.95	4.22
Subject b	-6°	-4.02	-2.63	-2.30	-3.39
	0°	-0.05	0.47	-0.05	-1.34
	+6°	2.48	3.12	3.16	1.69
Subject d	-6°	-2.89	-4.2	-2.63	-4.5
	0°	0.44	-0.15	1.38	0.66
	+6°	5.53	4.38	5.22	4.80
Subject e	-6°	-4.09	-4.34	-4.99	-5.31
	0°	-1.23	-0.82	-1.04	-0.56
	+6°	3.53	3.49	4.69	4.32
Subject g	-6°	-3.71	-2.16	-1.94	-1.16
	0°	0.02	-0.81	-1.42	0.04
	+6°	2.22	0.64	2.64	1.88

Table 10.3 Estimates of reference shift for each subject (rows), each reference shift condition (rows), and each coherence level tested (columns). Colours indicate the reference condition as in Chapter 3, Section 3.3.1.

Table 4.4 – Summary of Likelihood Ratio Test for Decoding versus Late, Decision Related Bias Model, Reference Shift Condition

		4%	7%	13%	25%
Subject a	-6°	$\chi^2(1) = 3.8415$ $p < 0.0001$	$\chi^2(1) = 3.8415$ $p < 0.0001$	$\chi^2(1) = 3.8415$ $p < 0.0001$	$\chi^2(1) = 3.8415$ $p < 0.0001$
	0°	$\chi^2(1) = 3.8415$ $p < 0.0001$	$\chi^2(1) = 3.8415$ $p = 1$	$\chi^2(1) = 3.8415$ $p = 0.55$	$\chi^2(1) = 3.8415$ $p < 0.0001$
	+6°	$\chi^2(1) = 3.8415$ $p < 0.0001$	$\chi^2(1) = 3.8415$ $p < 0.0001$	$\chi^2(1) = 3.8415$ $p = 1$	$\chi^2(1) = 3.8415$ $p < 0.0001$
Subject b	-6°	$\chi^2(1) = 3.8415$ $p < 0.0001$	$\chi^2(1) = 3.8415$ $p < 0.0001$	$\chi^2(1) = 3.8415$ $p < 0.0001$	$\chi^2(1) = 3.8415$ $p < 0.0001$
	0°	$\chi^2(1) = 3.8415$ $p = 0.31$	$\chi^2(1) = 3.8415$ $p < 0.0001$	$\chi^2(1) = 3.8415$ $p = 0.56$	$\chi^2(1) = 3.8415$ $p < 0.0001$
	+6°	$\chi^2(1) = 3.8415$ $p < 0.0001$	$\chi^2(1) = 3.8415$ $p < 0.0001$	$\chi^2(1) = 3.8415$ $p = 1$	$\chi^2(1) = 3.8415$ $p < 0.0001$
Subject d	-6°	$\chi^2(1) = 3.8415$ $p < 0.0001$	$\chi^2(1) = 3.8415$ $p < 0.0001$	$\chi^2(1) = 3.8415$ $p < 0.0001$	$\chi^2(1) = 3.8415$ $p < 0.0001$
	0°	$\chi^2(1) = 3.8415$ $p = 0.02$	$\chi^2(1) = 3.8415$ $p = 0.41$	$\chi^2(1) = 3.8415$ $p = 1$	$\chi^2(1) = 3.8415$ $p < 0.0001$
	+6°	$\chi^2(1) = 3.8415$ $p < 0.0001$	$\chi^2(1) = 3.8415$ $p < 0.0001$	$\chi^2(1) = 3.8415$ $p < 0.0001$	$\chi^2(1) = 3.8415$ $p < 0.0001$
Subject e	-6°	$\chi^2(1) = 3.8415$ $p < 0.0001$	$\chi^2(1) = 3.8415$ $p < 0.0001$	$\chi^2(1) = 3.8415$ $p < 0.0001$	$\chi^2(1) = 3.8415$ $p < 0.0001$
	0°	$\chi^2(1) = 3.8415$ $p < 0.0001$	$\chi^2(1) = 3.8415$ $p < 0.0001$	$\chi^2(1) = 3.8415$ $p < 0.0001$	$\chi^2(1) = 3.8415$ $p < 0.0001$
	+6°	$\chi^2(1) = 3.8415$ $p < 0.0001$	$\chi^2(1) = 3.8415$ $p < 0.0001$	$\chi^2(1) = 3.8415$ $p < 0.0001$	$\chi^2(1) = 3.8415$ $p < 0.0001$
Subject g	-6°	$\chi^2(1) = 3.8415$ $p < 0.0001$	$\chi^2(1) = 3.8415$ $p = 1$	$\chi^2(1) = 3.8415$ $p < 0.0001$	$\chi^2(1) = 3.8415$ $p < 0.0001$
	0°	$\chi^2(1) = 3.8415$ $p = 1$	$\chi^2(1) = 3.8415$ $p = 1$	$\chi^2(1) = 3.8415$ $p < 0.0001$	$\chi^2(1) = 3.8415$ $p = 0.28$
	+6°	$\chi^2(1) = 3.8415$ $p < 0.0001$	$\chi^2(1) = 3.8415$ $p = 1$	$\chi^2(1) = 3.8415$ $p < 0.0001$	$\chi^2(1) = 3.8415$ $p < 0.0001$

Table 10.4 Likelihood Ratio Test. Performance in fitting estimates of motion direction from the decoding and late, decision related bias models are compared. Cells highlighted in grey show tests where the late, decision-related bias model did *not* outperform the decoding model. Subjects and reference shift conditions in rows (colours as in Chapter 3, Section 3.3.1), coherence levels in columns.

Figure 10.5 – Summary of Estimation Times for Manual Responses

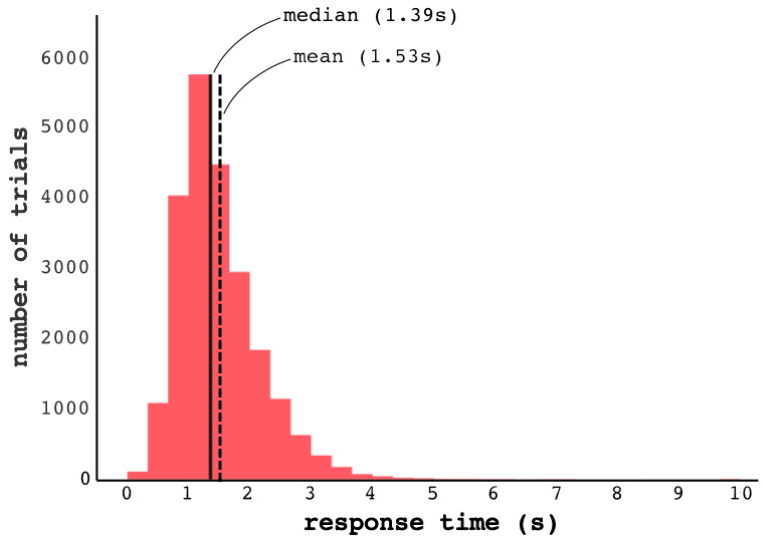


Figure 10.5 Manual response time frequency distribution. Data from all experiments described in Chapter 3. The histogram shows that, although subjects were allowed 10s to provide an estimate for the perceived direction of motion, this took 1.53s on average (black dashed line). Thick solid line indicates median response time (1.39s).

Figure 10.6 – Summary of Estimation Times for Manual Responses in Experiment 3, Chapter 3

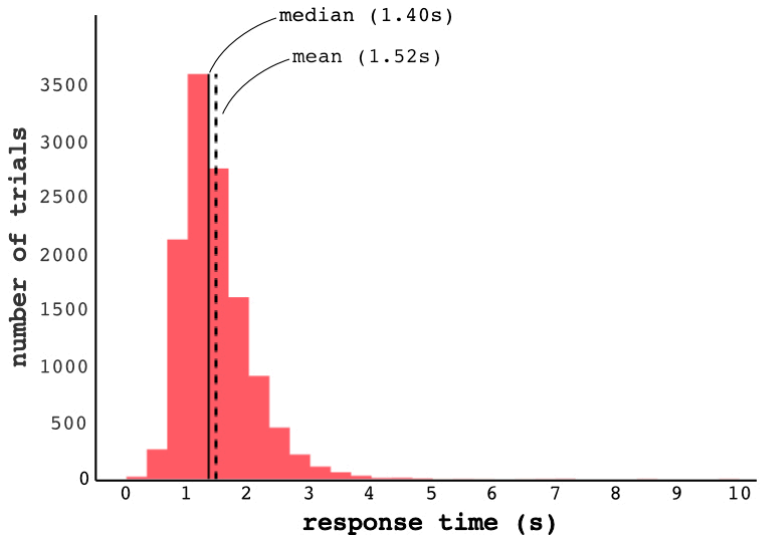


Figure 10.6 Manual response time frequency distribution. Data as described in Figure 10.5 for responses provided in Experiment 3, Chapter 3 only. Again, estimates of perceived direction of motion took 1.52s on average (1.40s median of skewed frequency distribution).