fMRI Investigations of Flexible Behavioural Control Based on Eye Movement Models

(Spine title: fMRI of Flexible Eye Movement Control)

(Thesis format: Integrated-Article)

by

Matthew Robert Graham <u>Brown</u>

Graduate Program in Neuroscience

Submitted as partial fulfillment
of the degree of Doctor of Philosophy

Faculty of Graduate Studies,
The University of Western Ontario,
London, Ontario, Canada

THE UNIVERSITY OF WESTERN ONTARIO
FACULTY OF GRADUATE STUDIES

CERTIFICATE OF EXAMINATION

Supervisor

_____
Dr. Stefan Everling

Co-supervisor

_____
Dr. Tutis Vilis

Supervisory Committee

_____
Dr. Jody Culham

_____
Dr. Paul Gribble

Examiners

_____
Dr. Blaine Chronik

_____
Dr. Brian Corneil

_____
Dr. Clayton Curtis

_____
Dr. Bruce Morton

The thesis by

Matthew Robert Graham <u>Brown</u>

Entitled:

**fMRI Investigations of Flexible Behavioural Control Based on Eye
Movement Models**

Is accepted in partial fulfillment of the
requirements for the degree of
Doctor of Philosophy

Date _____   _____
                                  Chair of the Thesis Examination Board

Abstract

The antisaccade task is an important model of flexible behavioural control. This task requires subjects to inhibit the automatic saccade toward a flashed peripheral visual stimulus and to generate a voluntary antisaccade toward the stimulus' mirror location in the opposite visual hemifield. Previous functional magnetic resonance imaging (fMRI) studies support the involvement of frontoparietal regions in antisaccade performance. Experiment 1 was designed to dissociate saccade inhibition from saccade generation processes by comparing prosaccades, antisaccades, and nogo trials in a rapid fMRI design. Trials included a task instruction followed by peripheral stimulus presentation and response. Prosaccade, antisaccade, and nogo trial responses were, respectively, to look at the stimulus, look away from it, and inhibit the automatic saccade while maintaining central fixation. Frontal eye field (FEF), supplementary eye field (SEF), anterior cingulate cortex (ACC), intraparietal sulcus (IPS), and precuneus exhibited surprisingly similar activations for prosaccade and nogo responses, suggesting that their fMRI signals might reflect visual detection and attention processes rather than saccade generation or inhibition. Inconsistently with previous studies, Experiment 1 revealed few instruction-related differences. In Experiment 2, we compared prosaccades and antisaccades using half trials to separate instruction- and response-related signals, rather than jittered instruction intervals as in Experiment 1. FEF, SEF, IPS, precuneus, ACC, and left dorsolateral prefrontal cortex (DLPFC) exhibited greater instruction-related activation for antisaccades, demonstrating that a rapid fMRI design can detect instruction-related differences. The first four regions also exhibited greater antisaccade response activation, unlike DLPFC and ACC, which might therefore be involved more in antisaccade preparation and task set rather than execution. In Experiment 3, we looked

for the saccade inhibition signature not seen in Experiment 1 by comparing frequent prosaccades and rare nogo trials (2:1 ratio). Nogo instruction-related activation was greater in right FEF, DLPFC, IPS, and precuneus, probably due to preparatory and task switching processes. Nogo response-related activation was greater in SEF, ACC, inferior frontal gyrus, and right supramarginal gyrus, probably due to saccade inhibition in nogo trials. Together, these experiments suggest that DLPFC is involved in task set while more posterior regions support a mixture of visual detection, attention, and saccade inhibition processes.

Co-Authorship

Matthew R.G. Brown[1,2,4], Tutis Vilis[1,2,4], Herbert C. Goltz[2,3], Kristen Ford[1,2,3], and Stefan Everling[1,2,3,4].

1. Dept. of Physiology and Pharmacology, University of Western Ontario, London, Ontario, Canada, N6A5C1.
2. Dept. of Psychology, University of Western Ontario, London, Ontario, Canada, N6A5C2.
3. Robarts Research Institute, London, Ontario, Canada, N6A5K8.
4. Graduate Program in Neuroscience, University of Western Ontario, London, Ontario, Canada, N6G2V4.

I, Matthew Brown, as the author of this thesis was responsible for every stage of the research program documented here, from initial design to writing up of results. My supervisors, Drs. Stefan Everling and Tutis Vilis, gave expert advice and supervision pertaining to experimental design and analysis, and they assisted in the preparation of manuscripts for publication. Dr. Herbert Goltz and Kristen Ford assisted with data collection for the first experiment described here (Chapter 3), and they provided comments during manuscript preparation.

The first experiment described here (Chapter 2) has been published in the journal NeuroImage as MR Brown, HC Goltz, T Vilis, KA Ford, and S Everling (2006) Inhibition and generation of saccades: Rapid event-related fMRI of prosaccades, antisaccades, and nogo trials. *NeuroImage* 33: 644-659. The second experiment (Chapter 3) has been accepted for publication with the Journal of Neurophysiology as MR Brown, T Vilis, and S Everling. 2007. Frontoparietal activation with preparation for antisaccades. *J Neurophys*. (in press). The third experiment (Chapter 4) is under review with NeuroImage, at the time of this writing.

First, the path is narrow.

Then the path is broad.

Then the path disappears.

For my wife, Lisa,

my daughters, Rheya and Arianna,

and our soon-to-arrive new baby.

Acknowledgements

Table of Contents

Preface

Chapter 1 – Review of Literature on Eye Movements, Executive Control, and fMRI

Chapter 4 – Isolation of Saccade Inhibition Processes: Rapid Event-related fMRI of Saccades and Nogo Trials

Chapter 5 – Summary and Conclusions

Tables

Figures

Appendices

Abbreviations

3n   – third cranial nerve nucleus (oculomotor nucleus)
4n   – fourth cranial nerve nucleus (trochlear nucleus)
6n   – sixth cranial nerve nucleus (abducens nucleus)
BOLD – blood oxygenation level dependent
df   – degrees of freedom
DLPFC – dorsolateral prefrontal cortex
EPI   – echo planar imaging
FEF   – frontal eye field
FST   – fundus of the superior temporal sulcus
fMRI   – functional magnetic resonance imaging
FOV   – field of view
FWHM – function width at half maximum
GLM   – general linear model
IFG   – inferior frontal gyrus
IPS   – intraparietal sulcus
LIP   – lateral intraparietal area
MFG   – middle frontal gyrus
nIC   – interstitial nucleus of Cajal
nPH   – nucleus prepositus hypoglossi
PEF   – parietal eye field
PET   – positron emission tomography
PPC   – posterior parietal cortex
PPRF   – paramedian pontine reticular formation
PO   – parieto-occipital area
riMLF   – rostral interstitial nucleus of the medial longitudinal fasciculus
rIPS   – rostral intraparietal sulcus
SEF   – supplementary eye field
SFG   – superior frontal gyrus
SNpr   – substantia nigra pars reticulata
TE   – echo time
TMS   – transcranial magnetic stimulation

Symbols

°   – degrees
Hz   – Hertz (cycles/second)
ms   – millisecond(s)
µA   – microamp(s)
s   – second(s)

Preface

The Question

In the study of the mind and brain, we are motivated by weighty questions. What is the mind? How does it work? What is consciousness? What is the relationship between the physical brain and subjective thought? The answers to these will have immeasurable social, ethical, and personal implications. Despite the importance of these issues, we do not yet know enough to address them fully, or even how to ask them properly. The answer to any question resides almost entirely in the way it is asked, in the implicit background knowledge tacitly evoked by the question. This is particularly true when dealing with complex subjects, and the mind and brain are among the most complex of subjects. We are still in the process of acquiring enough background knowledge to ask questions, that are simultaneously simple, broad, and meaningful, about what the mind and brain are and how they work.

In acquiring this background knowledge, we can begin by asking questions about specific facets of behaviour and the neuronal subsystems that support them. One example would be to ask how do we exert flexible and adaptive control over our own behaviour? This question should be somewhat easier to answer, and it is important because flexible behavioural control is something by which we define and identify ourselves as a species. However, flexible behavioural control is still too broad a topic, so researchers limit themselves to model systems and tasks designed to capture the essential patterns of behavioural control in an effort to make the research more tractable.

My own research uses the control of eye movements as a model of general behavioural control. Primates, including humans, tend to look automatically at abrupt, peripheral visual stimuli. The evolutionary advantages of this are obvious. Reflexive, visually guided saccades are rapid, involuntary eye movements toward visual stimuli, and they provide a model of automatic behaviour. The antisaccade task provides a model of flexible behavioural control. This task requires the subject to inhibit the automatic eye movement toward a peripheral stimulus and then to generate a voluntary eye movement to the blank space diametrically opposite the stimulus. Though perhaps superficially artificial and narrow, the antisaccade task requires processes like response inhibition,

attention, task set, and voluntary response generation thought to underlie diverse forms of behaviour. The antisaccade task is also very similar to some common situations in everyday life. When receiving a flu shot, one must inhibit the automatic tendency to flinch away from the painful prick of the needle so that one can voluntarily remain still. If one has learned to drive on the right-hand side of the road and one visits a country that drives on the left, one must inhibit the ingrained habit to position the automobile on the right in favour of the left.

Now I can specify the question underlying my research. How do we exert flexible control over our behaviour generally, and how do we perform the antisaccade task specifically?

Overview of Thesis

In *Chapter 1*, I will review those subsets of the literature directly relevant to my thesis work. The first part of *Chapter 1* will cover the neurophysiology of eye movement control, including what is known of how antisaccades are performed. The second part of *Chapter 1* will discuss functional magnetic resonance imaging (fMRI), and in particular rapid event-related fMRI techniques, which I used throughout my thesis work. It seems to be the unfortunate fate of most graduate theses to be read by those directly involved in the defense and then to gather dust on a library shelf. The scientific community understandably opts for the published papers that form the body of the thesis, instead. In *Chapter 1*, I have deliberately taken the approach of trying to explain the background material a little more simply and less tersely than I could have. It is my hope that this thesis might enjoy a more productive future than shelf occupation, perhaps as an introduction for new students to the neuroscience of eye movements and fMRI.

In *Chapters 2*, *3*, and *4*, I will describe the three experiments that comprise my thesis work. These experiments address the following questions: 1) What is the role of various cortical regions in generating voluntary eye movement responses versus inhibiting unwanted, automatic eye movements? 2) Can we separate instruction- and response-related activation using rapid event-related fMRI? 3) Which brain regions are involved in actively inhibiting automatic eye movements that are made particularly

automatic by being presented more frequently than other trial types? At the time of writing, the first experiment has been published in *NeuroImage* (MR Brown, HC Goltz, T Vilis, KA Ford, and S Everling, 2006. Inhibition and generation of saccades: Rapid event-related fMRI of prosaccades, antisaccades, and nogo trials. *NeuroImage* 33(2): 644-659). The second experiment has been published in the *Journal of Neurophysiology* (MR Brown, T Vilis, and S Everling, 2007. Frontoparietal activation with preparation for antisaccades. *Journal of Neurophysiology* 98(3): 1751-1762). The third experiment has been accepted for publication in *NeuroImage* (MR Brown, T Vilis, and S Everling, 2007. Isolation of saccade inhibition processes: Rapid event-related fMRI of saccades and nogo trials. *NeuroImage* [in press]). I have included the corresponding manuscripts in *Chapters 2*, *3*, and *4*. Finally, in *Chapter 5*, I will provide a summary and discussion of the three experiments, as well as some concluding thoughts on future research directions.

Chapter 1
Review of Literature on Eye Movements,
Executive Control, and fMRI

1.1 – Neurophysiology of Eye Movements and Executive Control

The first part of *Chapter 1* will review the literature on eye movement control. *Section 1.1.1* will motivate the use of the eye movement system to model general behavioural control. *Sections 1.1.2* to *1.1.13* will explore the hierarchical organization of the system, starting with the oculomotor plant and ending with cortical saccade regions. *Sections 1.1.14* through *1.1.18* will discuss the antisaccade task as a model of executive control, while *Section 1.1.19* will summarize of the preceding sections. Finally, *Section 1.1.20* will describe my research questions and hypotheses.

1.1.1 – The Eye Movement System as Model

Humans, like all primates, depend heavily on vision. The primate eye has greatest spatial acuity at its fovea, or centre. Exploring the visual world in high definition involves reorienting the fovea sequentially toward objects appearing in one's peripheral vision. Saccades are the fast and accurate eye movements that perform the necessary reorientations. The retinal slip, or movement of the image on the retina, that occurs during a saccade blurs the image, making vision impossible while the saccade is in flight. As a result, visual input is temporarily suspended during a saccade, a phenomenon which has been termed saccadic suppression (Erdmann and Dodge 1898). Blindness during saccade execution explains why saccades must be fast. In addition to generating reorienting movements, the eye movement system is capable of fixating, or maintaining stable gaze, on an object to facilitate its viewing.

The eye movement system has been used extensively as a model of general behaviour. It is assumed that understanding the intricacies of eye movement control will lead directly to understanding some components of general behavioural control and to a preliminary understanding of other components. The eye movement system constitutes a "microcosm of the brain" (Carpenter 1994, p. 341). That is, the eye movement system

must take sensory input and transform it into an eye movement command. The mapping of sensory input to motor output is, very generally, the function of the entire brain, which evolved first and foremost to direct the body's behaviour.

Like any model system, the eye movement system has advantages and disadvantages. I will start with the advantages. The eye movement system is one of the best understood of all primate neuronal systems at a basic, neurophysiological level (see Munoz and Everling 2004). The neuronal circuitry of the eye movement system has been mapped quite thoroughly, and we have extensive data on the discharge properties of neurons throughout the system. We understand particularly well the "low-level" eye movement centres that lie functionally close to the actual generation of the eye movements, themselves. Eye movements have relatively simple kinematics and dynamics because the eye has negligible inertia and only moves with three degrees of freedom. This makes the analysis of eye movements much more tractable than arm movements, for example, which involve substantial inertias and many degrees of freedom. Eye movement tasks are easy for human subjects to perform, and it is even possible to train monkeys to perform them. From a technological point of view, eye movements are straightforward for experimentalists to measure. They can also be performed in constrained spaces and do not involve large body movements, in contrast to reaching motions for example. These last two points are important particularly for functional magnetic resonance imaging (fMRI) studies due to the small bore space of fMRI scanners and fMRI's susceptibility to magnetic distortions caused by subject movement.

Despite its advantages, the eye control movement system is certainly not a complete model for general behavioural control. Eye movements are primarily orienting movements that, for the most part, do not cause any direct impact on the world. One important exception is that in social situations, gaze direction has significant consequences for what others do and how one interacts with them. In most cases though, eye movements do not exert the kind of direct impact on the environment that we can achieve with our limbs. This likely means that there are differences between eye movement and limb control circuitry, for example in how error signals are detected, processed, and learned from. However, in terms of using eye movement control as a model of general behavioural control, it is not necessary that the former should capture

completely all the properties of the latter. Models are just that, simplified versions of the thing being modeled that capture some of its aspects. Particularly in biology and neuroscience, the only perfect model of a system is the system, itself.

1.1.2 – Characteristics of Eye Movements and the Oculomotor Plant

In my thesis dissertation, I will focus on two types of eye movements, saccades and fixation. As mentioned above, saccades are the rapid, accurate eye movements we use to reorient the visually-acute retinal fovea onto different objects in the visual environment. Fixation is the maintenance of stable gaze on an object. Because of saccadic suppression, vision occurs only under the stable gaze conditions that exist during fixation periods (Erdmann and Dodge 1898). We humans typically make about 3 saccades per second interleaved with brief periods of fixation lasting 85 to 500 ms as we go about the business of seeing the world (Schiller and Tehovnik 2005; Schall 2002).

To understand the control of eye movements, we must understand the eye movements, themselves, as well as the oculomotor plant, which is the system comprised of the eyeball and the muscles, connective tissues, and bony orbit that support it. Saccadic eye movements have been studied psychophysically using the visually guided saccade task (Saslow 1967; Fischer and Boch 1983; Fischer and Ramsperger 1984; Abrams et al. 1989). This task simply involves looking at a central fixation point and then looking at a peripheral stimulus flashed in the periphery. Primate saccades take 15-100 ms to complete, depending on their amplitude, and can reach peak angular velocities of 500 °/s (Fuchs et al. 1985). Even large saccades are accurate to within about 1°, with smaller saccades being more accurate. Latency is the time interval between peripheral stimulus onset and the start of a saccade. Latencies of visually guided saccades vary depending on when the fixation point is extinguished. If fixation is extinguished simultaneously with or after peripheral stimulus onset, latencies of 200 ms are typical in monkeys and humans. However, if the fixation point disappears 150-250 ms before peripheral stimulus onset, latencies become much shorter, with 140-160 ms being typical values (Saslow 1967). This is called the gap effect. In addition, the presence of a gap between fixation offset and peripheral stimulus onset causes saccade latencies to become bimodal. The extra mode is comprised of express saccades with very short latencies of

70-80 ms in monkeys (Fischer and Boch 1983) and about 100 ms in humans (Fischer and Ramsperger 1984). With gap duration of 200-240 ms, express saccades can constitute up to 80% of saccades in monkeys. As the gap becomes shorter than 200 ms, regular saccades with normal latencies around 140-160 ms become more common, and express saccades are absent with gap durations shorter than about 140 ms. Regular and express saccades are thought to have different neurophysiological underpinnings, as discussed below in *Section 1.1.4.*

Eye movements are physically effected by six extraocular muscles that are arranged in pairs oriented in three approximately orthogonal planes including the horizontal or axial plane and the two vertical planes created by rotating the sagittal plane about 45° either clockwise or counter-clockwise when viewed from above. The extraocular muscles include the medial and lateral recti, the superior and inferior recti, and the superior and inferior oblique muscles. The medial and lateral recti rotate the eye horizontally, within the axial plane. The other muscles cause various combinations of vertical and torsional rotations. The six extraocular muscles work in concert to reorient the eye toward different objects in the environment.

The muscles and connective tissues that support the eyeball exert viscous and elastic forces on it when it deviates from the canonical, forward-looking orientation in the orbit (Robinson 1964). This fact has several implications for eye movement control. First, saccadic eye movements do not need to be braked actively by the extraocular muscles; the visco-elastic properties of the plant stop the eye automatically. Second, to move the eyeball with the high velocities characteristic of saccades, the extraocular muscles must exert considerable force on it to overcome this visco-elastic braking. Third, if the eyeball is deviated artificially away from the centre gaze orientation, for example by applying a tangential force to it with a finger, the muscles and connective tissues will passively pull the eyeball back to the canonical orientation. This means that to maintain an eccentric gaze orientation deviated away from the central orientation, the extraocular muscles must exert a tonic force on the eyeball to counter the visco-elastic returning force. In the next section, we will see that the brainstem saccade generator circuits, which control the extraocular muscles, have evolved to compensate for these mechanical properties of the oculomotor plant.

Eye movements can be classified into involuntary and voluntary groups. Involuntary saccades are generated by the "visual grasp reflex" (Hess et al. 1946), which is the automatic tendency to look toward novel or salient visual stimuli. The evolutionary advantage of this reflex is obvious – novel, salient stimuli might belong to a threatening source like a bear or hostile conspecific. The visual grasp reflex plays an important role in brain and behaviour research because it provides a model of automatic behaviour. In the antisaccade task (Hallett 1978), the subject is presented with a peripheral visual stimulus and must inhibit the automatic saccade response toward the stimulus. The subject must then generate a voluntary saccade toward the stimulus' mirror location in the opposite visual hemifield. The antisaccade task has been used extensively as a model of general behavioural control (Everling and Fischer 1998; Munoz and Everling 2004). Voluntary saccades, of which the antisaccade is a good example, are just as their name implies.

1.1.3 – Brainstem Saccade Generators

The brainstem saccade generators send saccade burst signals to the oculomotor neurons that control the extraocular muscles. The anatomy and physiology of these systems have been studied extensively and are quite well understood (for review see Robinson 1972; Robinson and Keller 1972; Robinson 1981; Fuchs et al. 1985; Munoz et al. 2000; Scudder et al. 2002; Schiller and Tehovnik 2005). The extraocular muscles are innervated by the third, fourth, and sixth cranial nerves. These nerves are comprised of axons from oculomotor neurons in the third (3n), fourth (4n), and sixth (6n) cranial nerve nuclei, also called the oculomotor, trochlear, and abducens nuclei, respectively. The six eye movement motor nuclei (n3, n4, and n6 on each side of the brainstem) send projections to each other in the medial longitudinal fasciculus. These projections allow the nuclei to coordinate their signals to the extraocular muscles so that the eyes move together (conjugate gaze control). The oculomotor neurons are also innervated by the brainstem saccade generators, which can be divided into two sets of nuclei. One set, located mostly in the pons and including 6n and part of 3n, is responsible for horizontal saccadic eye movements, while the other set, located in the mesencephalon or midbrain

and containing 4n and the rest of 3n, is responsible for vertical and torsional saccadic eye movements.

I mentioned in the previous section that considerable force is required to drive a rapid, saccadic eye movement against the visco-elastic properties of the oculomotor plant. Accordingly, during a saccade, the oculomotor neurons send high frequency volleys of action potentials to the extraocular muscles, causing them to contract forcefully. This is the phasic component of the saccadic oculomotor signal. There exists an essentially linear relationship between the frequency of phasic signal action potentials sent to the muscles and the angular velocity of the eye (see Fuchs et al. 1985). Once the eye has reached its target, the phasic signal ceases, and the eye movement is stopped by visco-elastic braking (Robinson 1964). Unless a saccade's end point is the canonical, forward-looking orientation, the eye will be deviated from centre at the end of the saccade. In this case, the phasic signal is followed by a tonic signal from the oculomotor neurons. The tonic signal is a low frequency, continuous sequence of action potentials that serves to contract the appropriate extraocular muscles against the visco-elastic force that would otherwise rotate the eyeball back to its central orientation.

Specific brainstem nuclei are responsible for the phasic and tonic components of the saccadic oculomotor signal. The paramedian pontine reticular formation (PPRF) and the rostral interstitial nucleus of the medial longitudinal fasciculus (riMLF) contain medium lead burst neurons that innervate the oculomotor neurons and discharge high frequency bursts of action potentials during execution of a saccade. PPRF projects to n6, while, riMLF projects to n3 and n4. Medium lead burst neurons, themselves, receive input from long lead burst neurons, which exhibit saccade-related bursts of action potentials as well as gradual increases or build ups of discharge activity preceding saccade onset. The medium and long lead burst neurons are quiescent during the fixation periods between saccades due to an inhibitory input from omnipause neurons in the medial raphé nuclei. The omnipause neurons are named for their continuous, tonic discharge behaviour, which ceases, or pauses, only during the execution of a saccade to allow the burst neurons to fire. In addition, burst neurons in PPRF project to the nucleus prepositus hypoglossi (nPH), while those in riMLF project to the interstitial nucleus of Cajal (nIC). Neurons in nPH and nIC integrate the burst neurons' signals, which

constitute a velocity signal for the eye, to derive the change in eye position due to the saccade (see Moschovakis 1997). Based on the computed eye position, nPH and nIC send appropriate tonic signals to the oculomotor neurons to maintain the eye in a deviated orientation. nPH innervates n6 and indirectly n3, and nIC innervates n3 and n4.

In summary, the brainstem saccade generators are responsible for the oculomotor signal that controls the extraocular muscles. The saccade generators also take care of "low level" aspects of eye movement control including the high frequency phasic burst signal responsible for high velocity saccades and the tonic signal that maintains stable, eccentric gaze positions. The saccade generators are innervated by the superior colliculus (SC) (Grantyn and Grantyn 1982; Moschovakis and Karabelas 1985; and see Munoz et al. 2000), a small, layered structure on the superior aspect of the mesencephalon, as well as by numerous cortical regions including the frontal eye field (FEF) (Leichnetz 1981), supplementary eye field (SEF) (Huerta and Kaas 1990), the lateral intarparietal area (LIP) (May and Andersen 1986), and dorsolateral prefrontal cortex (DLPFC) (Goldman and Nauta 1976; Leichnetz et al. 1981). These cortical structures constitute the "higher-level" eye movement control system, described in later sections.

1.1.4 – Superior Colliculus (SC)

As mentioned in the last section, the superior colliculus (SC) sends important projections to the brainstem saccade generators. The SC, itself, receives input directly from the retina (Schiller and Malpeli 1977) as well as from V1, FEF, SEF, LIP, and DLPFC (Finlay et al. 1976; Leichnetz et al. 1981; Lynch et al. 1985). SC is important for reflexive visually guided saccades and as a "gateway" between the cortical and brainstem saccade systems (for review see Munoz et al. 2000).

The SC contains a variety of neurons arranged in a retinotopic map (Wurtz and Albano 1980; Sparks 1986; Sparks and Mays 1990; Munoz et al. 1991; and see Munoz et al. 2000). The rostal end of SC represents the fovea, while more posterior parts of SC represent increasingly eccentric locations in the visual field. Each half of SC represents the contralateral visual hemifield, and the left-to-right axis of SC maps onto the direction within the visual field (in the polar coordinate system). The superficial layers of SC contain visually-responsive neurons that receive direct projections from retinal ganglion

cells as well as other visual areas (Robinson and McClurkin 1989a; Robinson and McClurkin 1989b). The intermediate layers of SC contain several types of visuo-motor neurons. Superior colliculus burst neurons (SCBNs), like medium lead burst neurons in the brainstem, exhibit high frequency bursts of action potentials slightly before and during saccades into their retinotopic movement fields (Munoz and Wurtz 1995; and see Munoz et al. 2000). Superior colliculus build-up neurons (SCBUNs), like brainstem long lead burst neurons, increase their action potential firing rates before saccades and discharge saccade-related bursts of action potentials (Munoz and Wurtz 1995). SCBNs and SCBUNs are thought to innervate, respectively, the medium and long lead burst neurons in the PPRF and riMLF in the brainstem (Munoz et al. 2000). Finally, the rostral pole of each half of SC contains fixation neurons that fire tonically during periods of fixation and cease discharging during saccade execution (Munoz and Wurtz 1993a; and see Munoz et al. 2000). The rostral fixation cells inhibit the motor neurons in more caudal parts of SC via inhibitory interneurons, and they also project to the omnipause neurons in the brainstem (Munoz et al. 2000). SC fixation neurons have been shown to play an important functional role in saccade control (Munoz and Wurtz 1993b). Increasing SC fixation neuron activity by means of microstimulation or introduction of the GABA antagonist bicuculline increases latencies of saccades. Bilateral microstimulation in rostral SC completely abolishes saccades, and stimulation while a saccade is in flight halts the eye movement. In contrast, injecting the GABA agonist muscimol into the SC fixation region decreases saccade latencies and interferes with the monkey's ability to maintain fixation and suppress unwanted saccades. SC neurons are thought to implement a population vector code of the desired saccade vector in retinotopic space (Munoz et al. 2000; Schiller and Tehovnik 2001).

The retina-SC-brainstem pathway is the shortest saccade pathway in the brain, and it is thought to be partially responsible for reflexive visually-guided saccades (see Munoz et al. 2000). As described above (*Section 1.1.2*), express saccades are visually guided saccades with unusually short latencies of around 70-100 ms (Fischer and Boch 1983; Fischer and Ramsperger 1984; Fischer and Weber 1993). Express saccades are potentiated with the gap paradigm, in which the central fixation point is removed about 200 ms before the onset of a peripheral stimulus, which evokes a visually guided saccade.

Over a series of gap saccade trials, one often observes a bimodal reaction time distribution, with the earlier mode corresponding to the express saccades and the later mode corresponding to the regular saccades. The bursting activity found in SCBNs and SCBUNs is actually composed of two separate bursts for visually-guided saccades with regular latencies (150-200 ms). The two bursts are thought to represent visual and motor components of visuo-motor transformation that generates a saccade. In contrast, only a single burst is observed for express saccades. SC is thought to use a threshold mechanism for saccade triggering; saccades are generated when discharge rates in SCBNs and SCBUNs reach threshold (Carpenter and Williams 1995; Hanes and Schall 1996; Dorris et al. 1997; Ratcliff et al. 1999). In regular saccades, the visual burst has a lower frequency than the motor burst. Only the motor burst is high enough to surpass the threshold to trigger a saccade. It has been suggested that express saccades occur when the initial visual burst has a high enough frequency to surpass threshold, causing the visual and motor bursts to merge into a single visuo-motor burst (Dorris et al. 1997; Edelman and Keller 1996). Because of noise in the neuronal circuitry, the precise height of the visual burst in SCBNs and SCBUNs fluctuates randomly from saccade to saccade. When a fixation point is present, fixation neurons in rostral SC are stimulated, and they send inhibitory input to SCBNs and SCBUNs via inhibitory interneurons. This inhibition prevents random fluctuations in SCBN and SCBUN firing rates from surpassing the saccade threshold. Removal of the fixation point about 200 ms before stimulus onset in the gap task causes decreased discharge rates in the fixation neurons in rostral SC. This in turn disinhibits the SCBUNs, which exhibit buildup activity during the gap. This increased preparatory activity is thought to predispose the system to generate express saccades (Munoz and Everling 2004). On a certain proportion of gap saccade trials, the background fluctuations in SC circuitry contrive to put the visual burst over threshold, leading to express saccades. Intact SC function is necessary for express saccade generation as destruction of one half of SC abolishes contralateral express saccades (Schiller et al. 1987).

In the antisaccade task (Hallett 1978), subjects must inhibit the automatic saccade toward a stimulus in favour of a voluntary saccade to the stimulus' mirror location in the opposite visual hemifield. Everling and Munoz (Everling et al. 1998; Everling et al.

1999) recorded from SC neurons while monkeys performed prosaccades and antisaccades. A prosaccade trial simply involves making a visually guided saccade to a peripheral stimulus. On each trial, a coloured fixation point instructed the animal on which task to perform. A peripheral visual stimulus then appeared, and the monkey was supposed to make a prosaccade toward or an antisaccade away from it, as previously instructed. During the instruction interval between instruction onset and peripheral stimulus onset, they found the following results. Fixation neurons in rostral SC discharged at higher rates for correct antisaccades than for prosaccades, and SCBUNs whose response fields encompassed the visual stimulus' location discharged at lower rates for antisaccades versus prosaccades. In addition, antisaccade error trials, in which the monkey failed to inhibit the automatic saccade and in fact looked at the visual stimulus, were associated with elevated discharges during the instruction period in SCBUNs coding for the visual stimulus' location. It appears then that successful performance in the antisaccade task involves enhancement of fixation-related discharge activity and suppression of pre-saccadic discharge activity in SC, as both of these measures would predispose the system against making erroneous automatic prosaccades (Everling et al. 1998; Everling et al. 1999).

1.1.5 – Basal Ganglia and Cerebellum

I will just briefly mention two other subcortical circuits important in eye movement control. In the basal ganglia, the caudate nucleus receives cortical projections and sends inhibitory projections to substantia nigra pars reticulata (SNpr) (see Hikosaka et al. 2000; Munoz and Everling 2004). SNpr in turn sends inhibitory projections to saccade neurons in the intermediate layers of SC. Thus, cortical input into this direct pathway through the basal ganglia serves to disinhibit SC saccade neurons. There is also an indirect pathway through caudate nucleus, globus pallidus externus, the subthalamic nucleus, and substantia nigra pars reticulata. The cerebellum is involved in fine-tuning saccadic eye movements (Robinson and Fuchs 2001).

1.1.6 – Introduction to Cortical Saccade System

The cortical saccade system is comprised of several regions distributed throughout the cortex. These regions include the frontal eye field (FEF), supplementary eye field (SEF), lateral intraparietal area (LIP), dorsolateral prefrontal cortex (DLPFC), and anterior cingulate cortex (ACC). These regions are heavily connected with each other and with other cortical and subcortical areas. FEF, SEF, LIP, and DLPFC all send reciprocal connections to each other (Andersen et al. 1985a; Huerta et al. 1987; Schall et al. 1993; Stanton et al. 1993; also see Colby and Duhamel 1991; Schall 1997). ACC is well connected with DLPFC and SEF (Bates and Goldman-Rakic 1993; Paus et al. 2001; Wang et al. 2004). As mentioned above, FEF, SEF, LIP, and DLPFC send projections to SC (Goldman and Nauta 1976; Leichnetz et al. 1981; Lynch et al. 1985). FEF, SEF, and LIP also project directly to brainstem saccade generator nuclei (Leichnetz 1981; Huerta and Kaas 1990; May and Andersen 1986). LIP receives visual input from multiple areas including V2, V3, MT, PO, FST, IT, and other posterior parietal regions (Asanuma et al. 1985; Blatt et al. 1990; also see Colby and Duhamel 1991). Frontal areas including FEF, SEF, and DLPFC receive visual input from posterior visual areas such as V4, PO, FST, and MT (Schall et al. 1995; Schall 1997). These visual areas are specialized for various aspects of visual processing, such as the extraction of features for object identification in IT and motion detection in MT. These "higher" visual areas receive input from "lower" visual areas including V1, V2, V3 (see Colby and Duhamel 1991).

Much work has gone into understanding the functional roles played by the various cortical saccade regions in eye movement control. Early efforts attempted to explain each cortical region in terms of a single, discrete function. LIP was hypothesized to code the locus of visuo-spatial attention by one group and the end point of an intended saccade by another (see *Section 1.1.12*). It is now appreciated that the organization of the mind and brain is highly distributed across different brain regions, each of which seems to perform a set of broad and overlapping functions. LIP seems to have roles in both visuo-spatial attention and saccade generation. Our understanding of what neurons do has also evolved, as illustrated in the following quote from Carpenter (p. 416 Carpenter 2004):

Cortical neurons are now seen to be not so much efficient little clerical *fonctionnaires*, tidily processing their incoming paper-work, as a turbulent, chaotic mob of self-opinionated demagogues, fighting with one another through lateral inhibition to make their voice heard above the crowd. Decisions come about through the growth of tides of opinion — 'rise to threshold' — resulting in the 'population vectors' (Georgopoulos 1994) that determine actions rather as a lynch-mob might suddenly form in a hostile crowd.

The theme of distributed computation is evident throughout the eye movement control literature. In *Sections 1.1.7-1.1.13* will cover background on the cortical saccade system, including regions FEF, SEF, and LIP. I will discuss the cortical saccade system in relation to the antisaccade task and executive control in *Sections 1.1.14-1.1.18*.

1.1.7 – Frontal Eye Field (FEF)

The frontal eye field (FEF) is important in the visuo-motor processing underlying saccade control. FEF is located at the junction of the superior frontal sulcus and precentral sulcus in humans (Paus 1996) and on the anterior bank of the arcuate sulcus in macaques (Bruce et al. 1985). Bruce and Goldberg (1985) provided the classical 'catalogue' of neuronal response patterns in FEF based on single unit recordings. Visual neurons comprise 40% of FEF neurons and respond to presentation of visual stimuli in their response fields with elevated discharge rates regardless of whether the monkey makes a saccade to the stimulus or not. Movement neurons comprise 20% of FEF neurons and discharge at higher frequencies before and during execution of saccades into their response fields, with or without a visual stimulus as target. Movement cells display weak or no response to visual stimulus presentation. Visuomovement neurons, comprising the remaining 40% of FEF neurons, have both visual and movement response components to their behaviour. As Bruce and Goldberg (1985) point out, this tripartite classification might be artificial in that the various response types could represent points on a continuum. Some FEF neurons also discharge at elevated rates during the delay period of the memory-guided saccade task (Bruce and Goldberg 1985).

FEF neurons code in retinotopic or eye-centred coordinates. Their response fields are 'locked' to the eye, always coding a particular location in visual space relative to the eye itself. Some other brain regions use different coordinate systems; we will see below that some SEF neurons code in head-centred coordinates. Like SC, FEF is also organized retinotopically. Neurons in the ventrolateral part of FEF have response fields representing locations near the centre of the contralateral visual field, with response field eccentricity increasing as one moves dorsomedially within FEF.

Electrical microstimulation in FEF evokes contralateral saccades (Robinson and Fuchs 1969; Bruce et al. 1985), and the magnitude of the saccade varies depending upon which part of the retinotopic map is stimulated (although the retinotopy in FEF is somewhat chaotic and not as well-defined as in SC). That is, the magnitude of an evoked saccade increases along the ventrolateral to dorsomedial axis (Bruce et al. 1985). The direction (the angle coordinate in the polar coordinate system) of response fields and evoked saccades in FEF does not change in a globally systematic way like the magnitude coordinate, but direction is organized in locally systematic patterns. As one advances an electrode through the bank of the arcuate sulcus, response direction changes incrementally. The amount of current necessary to evoke saccades in FEF is very low, as small as 10 µA. Bruce et al. (1985) provided one definition of FEF as the region in and around the arcuate sulcus from which saccades can be evoked with current less than 50 µA. Microstimulation sites in FEF with lower thresholds (below 50 µA) tend to contain movement and visuomovement neurons, whereas sites around FEF with higher thresholds (above 100 µA) tend to contain visual neurons.

Ablation of FEF partially or completely abolishes saccadic eye movements (Schiller et al. 1979, 1980; Keating and Gooley 1988b). Unilateral FEF lesions prevent only contralateral saccades. These deficits are temporary, however, with essentially complete recovery occurring within 2-4 weeks. If SC is also ablated along with FEF, saccades are permanently abolished. Inactivation of FEF by muscimol or lidocaine disrupts performance of visually guided and memory-guided saccades with endpoints corresponding to the inactivated FEF region's retinotopic representation (Sommer and Tehovnik 1997; Dias and Segraves 1999). This disruption takes the form of increased saccade latencies and decreased saccade velocities and accuracies.

Antidromic methods have established that FEF sends functional connections to SC (Segraves and Goldberg 1987; Sommer and Wurtz 2000, 2001; Everling and Munoz 2000; Wurtz et al. 2001) and to the brainstem saccade generators (Segraves 1992). FEF neurons projecting to SC include the entire range of FEF neuronal response profiles (Sommer and Wurtz 2000). Stimulation in FEF evokes orthodromic action potentials in SC neurons with movement and visual activity located in the intermediate and deep layers of SC but not in purely visual cells in the superficial layers of SC (Helminski and Segraves 2003). SC also sends reciprocal connections back to FEF (Sommer and Wurtz 1998). The precise anatomy of this pathway is not known but Lynch et al. (1994) have suggested that SC projects to FEF indirectly by way of the thalamus.

Functional imaging studies based on positron emission tomography (PET) and functional magnetic resonance imaging (fMRI) (see *Sections 1.2.1-1.2.6* and *Appendix 3*) have found that FEF as a whole exhibits more regional cerebral blood flow and higher blood oxygenation level dependent (BOLD) signal levels during visually guided saccade performance compared to fixation and during memory-guided saccade or antisaccade performance compared to visually guided saccade performance (Sweeney et al. 1996; Connolly et al. 2002; Connolly et al. 2005; Curtis and D'Esposito 2003; Curtis et al. 2004; Curtis et al. 2005; DeSouza et al. 2003; Brown et al. 2004; Ford et al. 2005).

1.1.8 – Lateral Intraparietal Area (LIP)

The lateral intraparietal area (LIP) has been implicated in allocating visuospatial attention and in saccade generation. In macaques, LIP is located, as the name suggests, on the lateral bank of the intraparietal sulcus (Colby and Duhamel 1991). It is unclear how to define a human homologue for LIP. Functional imaging studies often find two different saccade-related activation foci in the posterior part of human intraparietal sulcus (Brown et al. 2004; Schluppeck et al. 2006). Muri et al. (1996) defined a human parietal eye field (PEF) as the human homologue of LIP. PEF occupies the lateral bank of the posterior part of the intraparietal sulcus.

LIP shares many similarities with FEF. Neurons in LIP discharge at elevated rates in response to visual stimulus presentation in their response fields and before and during saccades into their response fields (Barash et al. 1991b; Colby et al. 1996). Some LIP

neurons also fire at higher rates during the delay period of the memory-guided saccade task. LIP neurons are tuned for both the direction and amplitude of a visual stimulus location or saccade end point (Platt and Glimcher 1998). More specifically, Platt and Glimcher tested LIP neurons with response fields within 40° of the centre of visual space and found that their response fields could be modeled as 2D Gaussian functions in retinotopic space. The mean and mode of these Gausian response fields' radii was 10° and 5°, respectively. Two thirds of LIP neuron response fields are contralateral, though neuronal recording has found them not to be arranged retinotopically as in SC (Andersen et al. 1985b; Andersen et al. 1990; Barash et al. 1991a; for review see Snyder et al. 2000; Colby and Duhamel 1991). In contrast, Sereno et al. (2001) found retinotopic fMRI signal patterns in an intraparietal sulcus region, which they presented as a potential human homologue of LIP. Further research will be necessary to resolve the issue of retinotopy in LIP. Visual, saccadic, and delay period-related discharge patterns can be present singly or more typically in combination in a given LIP neuron. LIP neurons also exhibit discharge patterns related to visual attention (see Goldberg et al. 2002) (and also see *Section 1.1.12* below). Using antidromic methods, Ferraina et al. (2002) found that LIP projections to FEF and SC carry visual, saccadic, and delay period-related signals. They also found that a larger proportion of LIP neurons projecting to FEF compared to SC had visual responses without saccade responses (78% versus 55%, respectively, for neurons projecting to FEF and SC), while a greater proportion of neurons projecting to SC had saccade responses (17% versus 42% for projections to FEF and SC).

Microstimulation in LIP evokes mainly contralateral saccades at low current thresholds (Thier and Andersen 1996, 1998). Lesions to posterior parietal cortex in humans cause increased latencies in visually and memory-guided saccades (Pierrot-Deseilligny et al. 1991b, a) and decreased target accuracy for memory-guided saccades (Pierrot-Deseilligny et al. 1991a). Inactivation of LIP by muscimol injection produces deficits in saccade behaviour including increases in the mean and variance of latencies for visually and memory-guided saccades contralateral to the injection (Li et al. 1999). Contralateral memory-guided saccades are also hypometric, have lower peak velocity, and, in some cases, show increased end point scattering.

PET and fMRI studies in humans have found that the intraparietal sulcus (IPS) exhibits more regional cerebral blood flow and higher blood oxygenation level dependent (BOLD) signal levels during visually guided saccade performance compared to fixation and during memory-guided saccade or antisaccade performance compared to visually guided saccade performance (Sweeney et al. 1996; Connolly et al. 2002; Connolly et al. 2005; Curtis and D'Esposito 2003; Curtis et al. 2004; Curtis et al. 2005; DeSouza et al. 2003; Brown et al. 2004; Ford et al. 2005).

LIP is discussed further in *Section 1.1.12*.

1.1.9 – Supplementary Eye Field (SEF)

The supplementary eye field (SEF) is located medial to the frontal eye field on the medial frontal gyrus. Specifically, SEF circumscribes the paracentral sulcus in humans (Grosbras et al. 1999), and in macaques, its position is several millimeters medial to the superior limb of the arcuate sulcus, dorsolateral to the pre-supplementary motor area (Huerta and Kaas 1990; Schlag and Schlag-Rey 1987). Many neurons in SEF, like those in FEF and LIP, change their firing rates in response to various combinations of visual stimulus presentation and saccade generation into their response fields (Schlag and Schlag-Rey 1987; Schall 1991). Microstimulation in SEF evokes saccades (Schlag and Schlag-Rey 1987; Tehovnik and Lee 1993; also see Schiller and Tehovnik 2001). Unlike FEF or LIP, SEF does not seem to code in retinotopic coordinates. Rather, many SEF neurons code in craniocentric coordinates, such that their response fields always occupy specific locations relative to the head (Schlag and Schlag-Rey 1987; Schall 1991). Microstimulation in a given SEF location also tends to evoke saccades with a common craniocentric endpoint irrespective of the eyes' starting position (Schlag and Schlag-Rey 1987; Tehovnik and Lee 1993; also see Schiller and Tehovnik 2001). Some SEF neurons also seem to use object-centred coordinates, such that they respond, for example, when the subject makes a saccade to a specific part of an object, say it's left side (Olson and Gettner 1995). SEF has been implicated in a variety of "higher order" functions. Evidence exists that SEF stores and manipulates sequence information in working memory. Patients with SEF lesions do poorly on tasks requiring them to perform a series of saccades based on remembered target locations (Gaymard et al. 1990). Isoda and Tanji

(2002) also found neuronal activity patterns in SEF coding for various aspects of saccade sequences. SEF has been implicated in coding arbitrary stimulus-response associations. For example, many SEF neurons code for eye movement responses that are inconsistent with the locations of visual stimuli, as in the antisaccade task (Schlag-Rey et al. 1997; Amador et al. 1998, 2004). (I will discuss the antisaccade task in more detail in *Section 1.1.14* below.) Using the countermanding saccade task, Stuphorn et al. (2000) found three groups of SEF neurons that were active after correctly performed trials, error trials, and reward delivery. PET and fMRI studies have found similar results in SEF, FEF, and IPS, consistent with the idea that these regions are components of a single, interconnected system for saccade control. SEF exhibits greater regional cerebral blood flow and higher BOLD signal levels for visually guided saccades compared to fixation and for memory-guided saccade or antisaccade performance compared to generation of visually guided saccades (Sweeney et al. 1996; Connolly et al. 2002; Connolly et al. 2005; Curtis and D'Esposito 2003; Curtis et al. 2004; Curtis et al. 2005; DeSouza et al. 2003; Brown et al. 2004; Ford et al. 2005).

1.1.10 – Automatic versus Voluntary Saccades

The saccade control system is capable of generating saccades along a spectrum from completely involuntary, reflexive saccades to voluntary, deliberate saccades. Schiller and colleagues have proposed that the saccade circuitry is organized into two pathways (Schiller 1998; Schiller and Tehovnik 2001). In their model, an anterior pathway, which passes from FEF and SEF directly to the brainstem saccade generators, bypassing SC, is responsible for voluntary saccades. A posterior pathway, which runs from LIP, through SC, and then to the brainstem, is responsible for automatic saccades. Early neuronal recording results seemed to indicate that FEF neurons might respond to visual stimuli in their response fields only when those stimuli are candidate targets for deliberate saccades (Bizzi 1968; Bizzi and Schiller 1970) (but see *Section 1.1.11*). FEF lesions also do not abolish express saccade generation (Schiller et al. 1987), whereas SC lesions do (Schiller et al. 1987). In addition, after SC has been lesioned, it is still possible to evoke saccades with microstimulation in FEF and SEF but not in LIP (Schiller and Malpeli 1977; Keating et al. 1983; Keating and Gooley 1988a; Tehovnik et al. 1994).

Hanes and Wurtz (2001) have provided a counterargument to the above hypothesis. They attempted to evoke saccades with microstimulation in FEF while simultaneously inactivating specific parts of SC using muscimol. Hanes and Wurtz found that when FEF stimulation and SC inactivation were in retinotopically correspondent locations, no saccades were evoked. Using muscimol, rather than ablative lesioning, to inactivate SC was important because it did not allow time for plastic reorganization and recovery to occur. Hanes and Wurtz proposed that evoking saccades with FEF stimulation after SC ablation is possible only due to recovery subsequent to the lesion. They suggested that direct projections from FEF to the brainstem saccade generators are normally weak and incapable of conveying a full saccade command. Only after ablation of SC would the direct FEF-brainstem projects undergo plastic changes resulting in the strengthening of the pathway.

1.1.11 – Visual Stimulus Selection and Saccade Generation in FEF

Making a saccade theoretically involves at least two processes. First a saccade target must be selected from among the array of stimuli in the visual world. A saccadic motor command must then be issued to generate a saccade to the target. Stimulus selection and saccade generation may or may not be discrete processes in different regions. We saw above (*Section 1.1.4*) that SC contains neurons that code both visual stimulus location and saccade endpoint location in retinotopic coordinates. That is, a retinotopic representation of stimulus location can also serve as a saccade motor command vector. It is arguable that there is not much "selection" going on here, though. Several studies have investigated stimulus selection and saccade generation in FEF, in which it turns out that stimulus selection and saccade generation are discrete processes that interact with each other.

Thompson et al. (1997) recorded from FEF neurons while monkeys discriminated the oddball among an array of distractor stimuli and then either made a saccade to the oddball in saccade trials or maintained central fixation without looking toward the oddball in nogo trials. Discriminating the oddball stimulus was actually unnecessary in the nogo trials, but FEF neurons that discriminated between oddball and distractor stimuli in the go task continued to exhibit the discrimination in the nogo task. These neurons

seem to be involved in stimulus selection rather than saccade generation. Thompson et al. (2005) showed that stimulus selection in FEF neurons is not contingent on making saccades. In their task, monkeys pressed a lever in the direction of the oddball stimulus in a distractor array without making any saccades. FEF neurons that exhibited visual or visuo-motor tuning on memory-guided saccades also discriminated the oddball in the lever task, while purely motor neurons were suppressed during the lever task.

Sato and Schall (2003) and Schall (2004) recorded from FEF neurons while monkeys performed a task similar to that used by Thompson et al. (1997). Monkeys discriminated the oddball stimulus against an array of distractor stimuli and then looked toward it on prosaccade trials, away from it on antisaccade trials, or maintained central fixation on nogo trials. The oddball stimulus' shape indicated which task to perform. In two thirds of FEF neurons, which they called Type I, discharge patterns were selective for the oddball stimulus' location. In the antisaccade task, the majority of Type I neurons subsequently changed their activity to be selective for the antisaccade end point. Most Type I neurons also exhibited oddball stimulus selectivity in the nogo task, which did not require any saccade execution. Type II neurons comprised the remaining third of FEF neurons. These always selected the saccade end point, regardless of whether this corresponded to the stimulus location. On nogo trials, Type II neurons selected what would be the antisaccade end point if the trial were an antisaccade trial. The authors attributed this result to an effect of how the monkeys were trained.

In summary, visual stimulus selection and saccade command generation are separate processes supported by distinct neuron populations in FEF. These processes must also interact extensively. Schall (2004) and Thompson et al. (2005) suggested that overt saccade generation is based upon visual stimulus selection followed by saccade endpoint selection, while the former process but not the latter is involved in covert attention shifting. Covert attention shifts involve reorienting the locus of visuo-spatial attention without making actual saccades. The above suggestion is a variant of the premotor theory of attention (Rizzolatti et al. 1987), which holds that the computations underling visuo-spatial attention shifting and saccade execution are shared up to some point before the issuing of a saccade command.

1.1.12 – Sensory versus Motor Processing in LIP

Early studies attempted to ascribe fairly narrow functional roles to cortical saccade regions. One group ascribed to LIP the role of coding the location of visuo-spatial attention (see Colby and Goldberg 1999; Goldberg et al. 2002; Goldberg et al. 2006), while another group held that LIP coded an intended saccade's endpoint (see Snyder et al. 2000; Andersen and Buneo 2002). The last 20 plus years have witnessed a long and occasionally rancorous debate between these two camps. Substantial evidence now suggests that saccadic motor processes and visuo-spatial attention processes are co-localized in cortical saccade regions with extensive interactions occurring between the two systems. I should point out that not everyone agrees with this view and that the debate is ongoing.

Snyder, Batista, and Andersen (1997) recorded from neurons in LIP and in what they call the parietal reach region (PRR), which lies immediately caudal and dorsal to LIP. Their monkey subjects performed delayed response tasks requiring either a saccade or reaching movement to the remembered location of a flashed peripheral stimulus. Importantly, when the monkeys performed one type of action, they were required not to perform the other. Throughout reach trials, they had to maintain central fixation and vice-versa. A majority of LIP neurons fired at higher rates during the delay period in saccade trials compared to reach trials, while a majority of PRR neurons displayed the reverse pattern. Snyder et al. argue that this double dissociation evinces that LIP and PRR neurons carry a motor intention or motor plan signal for eye and hand movements, respectively. They argue that visual attention effects cannot explain these results because if the LIP and PRR neurons were coding only the locus of visual attention, their activity should be independent of the type of movement planned. Furthermore, one half of those neurons that exhibited elevated delay activity in both saccade and reach trials showed a decisive preference for one movement type on dissociation trials in which the monkey made an eye movement into / out of the neuron's response field and an opposing reach movement out of / into the response field. In these trials, the neurons in question discharged at elevated rates during the delay period only when the monkey planned the preferred movement type into the response field. Finally, Snyder, Batista, and Andersen (1998) used a task in which the movement modality, saccade versus reach, was specified

twice. That is, a coloured peripheral stimulus was presented twice, at the same location, with the colour indicating the modality. Modality could either remain the same across both presentations, or it could switch. During the first and second delay periods, after the first and second peripheral stimulus presentations, respectively, LIP neurons discharged at higher rates for a saccade instruction compared to a reach instruction. In addition, the largest firing rates were evoked by the second stimulus presentation when the instruction switched from reach to saccade, whereas a repeated saccade instruction stimulus evoked lower discharge rates, though these were still higher than those for reach instructions. PRR displayed the same pattern of results, with the exception that PRR cells preferred reach over saccade instructions. Snyder et al. argued that visual attention is more or less constant across a given trial and that saccade versus reach instruction-based modulations of neuronal firing rates must therefore reflect motor planning processes. One possibility that none of the above studies addresses is that visuo-spatial attention might consist of multiple subprocesses specific to and partially identical to the eye movement and reaching command processes. Recall the premotor theory of attention, which we encountered in the *Section 1.1.11* and which holds that a visuo-spatial attention signal is a precursor to a saccade command (Rizzolatti et al. 1987). This hypothesis suggests that it is meaningless to attempt to separate neuronal signals into attention and saccadic motor categories because at an initial stage of processing, they are the same.

Goldberg and colleagues have argued that LIP's role is primarily in attention rather than intention (Goldberg et al. 2002; Goldberg et al. 2006), in opposition to Andersen and company. The antisaccade task dissociates visual stimulus and saccade end point locations by requiring subjects to generate a saccade away from a peripheral stimulus to its mirror location in the opposite visual hemifield. Gottlieb and Goldberg (1999) recorded from LIP neurons in monkeys performing the prosaccade and antisaccade tasks. They found that the vast majority of neurons responded reliably when monkeys made prosaccades, which are simply visually guided saccades, into their response fields whereas these neurons responded only weakly or not at all when an antisaccade was made into the response field. A minority of LIP neurons (12%) did exhibit robust activity for either prosaccades or antisaccades into their response fields, suggesting that these neurons might code saccade endpoint. Gottlieb and Goldberg's

results suggest that most LIP neurons do not code arbitrary saccade endpoint locations. This study did not rule out whether LIP might specifically code visually guided saccade endpoints.

To address visually guided saccade coding in LIP, Powell and Goldberg (2000) recorded from LIP neurons while monkeys performed memory-guided saccades, in which the monkeys had to remember the location of a flashed peripheral stimulus and then generate a memory-guided saccade to that location at the end of a delay period. In some trials, a distractor stimulus was presented during the delay period. LIP neurons responded with higher firing rates to the distractor than when the same stimulus was the target of a memory-guided saccade. Powell and Goldberg suggested that these findings are difficult to reconcile with the hypothesis that LIP codes a saccade generation signal whereas a role in coding the locus of visuo-spatial attention does fit the results. It is possible that increases in neuronal firing rates in response to the distractor might provide a signal not only for attention but also for a new visually guided saccade plan. This brings up the premotor theory of attention again (Rizzolatti et al. 1987).

Bisley and Goldberg (2003) used a sophisticated task to probe LIP neuronal activity. They had monkeys plan to make a saccade to a remembered target location. They then flashed a probe stimulus either at the target location or in a different location. The probe served as a go/nogo cue instructing the monkey whether the execute the planned saccade or not. In addition, they varied the contrast of the go/nogo cue. When the cue was presented at the location of the planned saccade, they saw an attentional effect, namely that the monkeys could correctly discriminate cue with lower contrast. In addition, some trials included a distractor stimulus flashed in a location different from the planned saccade end point. Bisley and Goldberg found that 200 ms after distractor presentation, the attentional effect was shifted away from the planned saccade end point to the distractor's prior location. By 700 ms after distractor offset, the attentional effect was shifted back to the planned saccade end point. LIP neurons were observed to respond either to the saccade target presentation or to the distractor presentation when their receptive fields encompassed those locations. That is, even though the monkey was planning a saccade toward the remembered target location, LIP neurons still responded to distractor presentation in their receptive fields. This distractor-related response was also

larger than the sustained, tonic activation coding for a remembered target location. Bisley and Goldberg suggest the LIP maintains a map of visuospatial attention, with attention being allocated to the location coded by neurons with the largest relative firing rates. One possible alternative interpretation of Bisley and Goldberg's results is that LIP neurons responding to the distractor might be coding both a shift in attention location as well as a potential, new saccade endpoint. In the absence of other task requirements, that location might become the end point of a visually-guided saccade, but since the monkey is required to look toward a remembered saccade target, the potential visually guided saccade is not executed, and only the location of attention is briefly shifted.

The above-cited studies on LIP emphasized its role either in saccade generation or visuo-spatial attention. The premotor theory of attention (Rizzolatti et al. 1987) suggests that it is impossible to separate LIP neuronal activity cleanly along the line of visuo-spatial attention versus saccadic motor intention. Perhaps this is not surprising given that visual attention and saccades play related roles in determining the input to the (resource limited) visual system (Rizzolatti et al. 1987). Corbetta et al. (1998; and see Corbetta and Shulman 2002) showed that visual attention and saccade tasks evoked similar patterns of fMRI activation in cortical areas, including FEF, SEF, and intraparietal sulcus. Current consensus seems to support a synthesis of the attention and intention hypotheses on LIP function. LIP maintains a saliency map of the visual world, and this is involved in directing both visual attention and visually guided saccades.

1.1.13 – Summary of Cortical Saccade System

We have seen that control of saccadic eye movements involves a distributed system of cortical regions including FEF, SEF, and IPS, as well as subcortical regions, including SC and the brainstem saccade generators. The brainstem saccade generators receives saccade command input from SC and cortical areas including FEF, SEF, and LIP. The brainstem circuitry performs processing related to the physics and geometry of eye movements and issues motor commands to the extraocular muscles, thereby relieving "higher" regions of these "low level" responsibilities. SC is important in triggering reflexive visually-guided saccades and in integrating input from cortical regions. Cortical saccade areas perform the sensory-motor transformations that underlie saccades,

particularly deliberate saccades. Many aspects of sensory-motor processing are organized in a distributed fashion across multiple cortical regions. For example, neurons in all three cortical areas exhibit visual or motor responses or most often a combination of the two. The premotor theory of attention holds that the allocation of visuo-spatial attention is an early component of the visuo-motor transformation and that attention shares neuronal circuitry with the saccade control system. Despite the similarities among FEF, SEF, and LIP, these regions do not have precisely the same roles. LIP seems to be somewhat more specialized for visual processing and attention, while FEF and SEF seem to play a relatively larger role in saccadic motor processing. SEF, in particular, is important for "higher level" computations underlying the storage and manipulation of sequence information as well as performance-related processing. In the next sections, I will discuss the antisaccade task as a model of flexible behavioural control involving high-level executive functions. We will see that DLPFC is an important player in executive function as well as in eye movement control.

1.1.14 – The Antisaccade Task

The antisaccade task requires the subject to inhibit the automatic saccade toward a flashed peripheral stimulus and generate a voluntary antisaccade to the stimulus' mirror location in the opposite visual hemifield. Compared to prosaccades (visually guided saccades), antisaccades are less accurate, have lower peak velocities, and have latencies that are longer by about 15-90 ms depending on the subject and experimental parameters such as the presence or absence of a gap, gap duration, and peripheral stimulus eccentricity (Hallett 1978; Fischer and Weber 1992, 1997; Bell et al. 2000). The longer latencies of antisaccade versus prosaccade trials are thought to be caused by additional processing underlying some combination of saccade inhibition, visuo-spatial remapping, and deliberate saccade generation in the antisaccade task (Hallett 1978; and see Munoz and Everling 2004). On antisaccade error trials, the subject fails to inhibit the automatic saccade toward the peripheral stimulus and generates a saccade toward it. Antisaccade error trials have shorter latencies than correct antisaccades, which reflects the uninhibited triggering of an automatic visually guided saccade command.

Everling and Munoz (2000) have recorded from FEF when monkeys performed antisaccades and prosaccades. They also used antidromic methods to identify FEF neurons projecting to SC. During the preparatory period immediately before peripheral stimulus presentation, saccade-related neurons discharged at lower rates for antisaccades versus prosaccades. Levels of prestimulus activity also correlated with saccadic reaction times, the frequency of express saccades, and antisaccade errors, in which the monkeys failed to inhibit the automatic saccade and looked toward the peripheral stimulus. Everling and Munoz interpreted these results as representing different preparatory set states preceding antisaccade and prosaccade performance, with reduced presaccadic activity in the antisaccade task necessary to allow for inhibition of the automatic prosaccade response. Stimulus- and saccade-related firing rates were also lower for antisaccades versus prosaccades, which accords with the preparatory set hypothesis. Lesions to FEF do not cause subjects to make antisaccade errors at an increased rate, suggesting that FEF is not involved directly in inhibiting the automatic visually guided saccade (Pierrot-Deseilligny et al. 1991b; Gaymard et al. 1999). FEF lesions do increase latencies and decrease velocities of antisaccades and prosaccades.

If FEF neurons discharge at lower rates for antisaccades into their response fields than for prosaccades, it would seem that this might create a deficiency in excitatory drive in SC and the brainstem saccade generators necessary to evoke a saccade. Indeed, this is the point; without this deficiency, the visual stimulus might evoke an unwanted reflex saccade. The deficiency in FEF saccade drive might conceivably be made up for by SEF neurons, which exhibit the opposite behaviour to FEF neurons. A large majority of saccade-related SEF neurons have been shown to fire at higher rates before and during antisaccades compared to prosaccades (Schlag-Rey et al. 1997; Amador et al. 2004). In addition, SEF firing rates 40 ms before saccade onset were strongly predictive of whether monkeys generated correct antisaccades or erroneously made prosaccade responses on antisaccade trials. These results were interpreted to suggest that SEF is involved in preparation to perform the antisaccade task as well as in voluntary generation of an antisaccade (Schlag-Rey et al. 1997; Amador et al. 2004). However, lesions to SEF do not increase the rate at which subjects make errors in the antisaccade task, which has

been suggested to preclude a role for SEF in saccade inhibition (Pierrot-Deseilligny et al. 1991b).

Posterior parietal cortex, including LIP in the intraparietal sulcus, has been studied with the antisaccade task. Gottlieb and Goldberg (1999) compared prosaccades and antisaccades in LIP and found that the majority of LIP neurons' discharge activity at the time of peripheral stimulus presentation and saccade generation reflected the location of the stimulus, rather than the saccade endpoint in the antisaccade task. A minority of neurons did code for the saccade endpoint, however, irrespective of task (prosaccade versus antisaccade). Zhang and Barash (2000; 2004) compared memory-guided prosaccades and antisaccades in LIP. They found a subset of LIP neurons whose activity seemed to change from coding the peripheral stimulus location to coding the antisaccade endpoint, and they suggested that LIP might be involved in the visuo-spatial remapping of the stimulus into its mirror position. No single unit study of LIP has investigated preparatory activity differences between prosaccades and antisaccades. That is no study has examined the time period following instruction presentation but preceding stimulus onset. Lesion studies in humans have found that damage to posterior parietal cortex increases prosaccade and antisaccade latencies but not the error rate in the antisaccade task, which has been interpreted as evidence against LIP's playing a direct role in saccade inhibition (Pierrot-Deseilligny et al. 1991b). These lesions did increase latencies and decrease velocities of antisaccades and prosaccades.

One important region about which I have said little thus far is the dorsolateral prefrontal cortex (DLPFC) in monkeys. The human homologue is thought to be Brodmann's area 46 (Pierrot-Deseilligny et al. 2005), also called posterior middle frontal gyrus (MFG). It is widely accepted that DLPFC plays an important role in inhibiting unwanted responses. Lesions to this region cause marked increases in the error rate for the antisaccade task such that patients with DLPFC lesions have difficulty suppressing the automatic saccade and tend to look at the peripheral stimulus on antisaccade trials (Guitton et al. 1985; Pierrot-Deseilligny et al. 1991b; Ploner et al. 2005; also see Pierrot-Deseilligny et al. 2005). Neuronal recording studies have found that many DLPFC neurons' discharge patterns are selective for whether the subject is performing the prosaccade or antisaccade task (Everling and DeSouza 2005; Johnston and Everling

2006b, a; Johnston et al. 2007). This activity also discriminates between correct and error performance on the antisaccade task. One hypothesis is that DLPFC supports task set and preparatory processing in combination with sending bias signals that infuence computations in other regions. I will discuss DLPFC and its role in executive control in *Section 1.1.17.*

1.1.15 – Functional Imaging of the Antisaccade Task

Functional imaging based on positron emission tomography (PET) and functional magnetic resonance imaging (fMRI) has been used to compare prosaccades and antisaccades in human subjects (Sweeney et al. 1996; Connolly et al. 2002; Curtis and D'Esposito 2003; DeSouza et al. 2003; Ford et al. 2005). Jointly, these studies have shown that cortical saccade regions including FEF, SEF, IPS, DLPFC, and ACC all exhibit increases in various proxies for population level neuronal activation. PET typically measures regional cerebral blood flow, and fMRI usually measures the blood oxygenation level dependent (BOLD) signal (see *Appendix 3.2* and *3.3*). The cortical saccade system is also involved in many other neuronal activities including performance of memory-guided saccades (Sweeney et al. 1996; Brown et al. 2004; Curtis et al. 2004) and reorientation of visuospatial attention (Corbetta et al. 1998; Corbetta and Shulman 2002), to name two. Returning to prosaccades and antisaccades, FEF, SEF, IPS, DLPFC, and ACC also display greater activation on antisaccade trials compared to prosaccade trials (Sweeney et al. 1996; Connolly et al. 2002; Curtis and D'Esposito 2003; DeSouza et al. 2003; Ford et al. 2005). Sweeney et al. (1996) used PET, which limited them to blocked trial design (see *Section 1.2.2*). The other studies used widely-spaced event-related fMRI designs (see *Section 1.2.3*) to separate instruction- and response-related signals. In what is called the instruction period of each trial, the subject was presented with a coloured fixation point, which indicated whether to perform the prosaccade or antisaccade task. The instruction period ended with peripheral stimulus onset, which elicited a prosaccade or antisaccade response from the subject, as previously instructed. Instruction-related BOLD signal was higher for antisaccades than prosaccades in cortical saccade regions. This finding has been interpreted as the signature of preparatory set processes, that is, preparation to perform an antisaccade versus a prosaccade. In some

cases, response-related activation was also greater for antisaccades compared to prosaccades (Curtis and D'Esposito 2003; DeSouza et al. 2003), though DeSouza et al. (2003) attributed this to carryover from the instruction period. Greater antisaccade response activation could reflect a combination of voluntary saccade generation, visuo-spatial remapping, and attention processes. These results are not entirely consistent across the above-cited studies in that certain brain regions exhibit significance in some experiments but not in others. The discrepancies seem to be primarily a matter of differences in statistical sensitivity. Previous studies have not actually reported conflicting results.

Antisaccade error trials occur when the subject fails to inhibit the automatic prosaccade during an antisaccade trial and instead looks at the peripheral stimulus. Ford et al. (2005) found greater activation during the late instruction period for correct antisaccades compared to error antisaccades in DLPFC and ACC, suggesting that insufficient engagement of preparatory processes in these regions leads to errors in the antisaccade task. Curtis et al. (2003) also found greater response-related activation for correct versus error antisaccades in FEF, SEF, and IPS, which could be attributed to insufficient recruitment of saccade inhibition processes in the error trials at the time of peripheral stimulus onset.

1.1.16 – fMRI versus Neuronal Recording Results in FEF

A discrepancy exists between single neuron recording and fMRI results with the antisaccade task in FEF, as has been discussed by Ford et al. (2005). FEF neurons discharge at higher frequencies before and during prosaccades compared to antisaccades (Everling and Munoz 2000) whereas instruction- and response-related fMRI activation evoked in FEF is higher for antisaccades compared to prosaccades (Connolly et al. 2002; Curtis and D'Esposito 2003; DeSouza et al. 2003; Ford et al. 2005). There are several potential explanations for this. Neuronal recording results were made in macaque monkeys whereas the fMRI results came from humans, and there might exist an intrinsic difference between the species. Monkeys are over-trained, with months of practice on the prosaccade and antisaccade tasks preceding actual recording. Human fMRI subjects are simply instructed on task performance on the day of the experiment and given perhaps 15

minutes of practice before data collection. Practice on a task is known to increase the involvement of basal ganglia circuits during task execution as the task becomes more automatic (see Graybiel 1995). Neuronal recording and fMRI also measure different aspects of neuronal computation. Finally, previous event-related fMRI comparisons of prosaccades and antisaccades used long fixation intervals (10-14s) between trials and long instruction periods (6-14s) to prevent overlap of the BOLD signal (see *Section 1.2.1*). Neuronal recording studies use much more rapid trial presentations with individual trials lasting on the order of 1 s and inter-trial intervals lasting 1-3 s. The difference in timing regimens might cause differences in task execution. My own research using rapid event-related designs allows us to discard this last hypothesis, as I will discuss in *Section 5.4*.

1.1.17 – Dorsolateral Prefrontal Cortex (DLPFC) and Saccade Inhibition

Dorsolateral prefrontal cortex (DLPFC) plays an important role in inhibiting the automatic prosaccade response in the antisaccade task. In macaque monkeys, DLPFC occupies the cortical convexity between the superior branch of the arcuate sulcus and the principle sulcus in the frontal lobe. How to define a human homologue of monkey DLPFC is debated. Human frontal cortex is more highly developed and constitutes a much larger proportion of the total cortical volume than in monkeys. Pierrot-Deseilligny et al. (2005) have suggested that a human homologue of DLPFC might be Brodmann's area 46, in the posterior part of the middle frontal gyrus (MFG). Of course, it is entirely possible that multiple prefrontal regions in humans could serve functions related to that of monkey DLPFC (for example see Courtney et al. 1998).

Lesion evidence has long pointed to prefrontal cortex's involvement in executive control. Frontal lobe syndrome, caused by prefrontal cortical damage, is characterized by a combination of impulsive, uncontrolled behaviour and a lack of initiative or impetus in various situations (see Luria 1966). At present, the phrase "executive control" cannot be rigorously defined, but Miller and Cohen (2001) have explained it as follows. Animals with more highly developed brains can express a greater range of behaviours. This has the secondary effect of making it more difficult to choose which among many potential actions to take in a given circumstance. Executive control is the set of functions that

prevent greater behavioural flexibility from degenerating into chaos and confusion. Examples of situations that engage executive control are those requiring inhibition of an automatic response in favour of a more appropriate voluntary response and those involving rapid changes in the stimulus-response mappings that determine appropriate behaviour. Miller and Cohen (2001) put forward a theory of executive control based on earlier work by Desimone and Duncan (1995), who described a model of visual attention. First, visual cortical neurons specialized to process different aspects of the visual scene, such as colour, shape, orientation, and location, compete with each other via inhibitory interneurons. In bottom-up attention, one set of neurons wins the competition because its favoured stimulus attribute is particularly salient, and these neurons then sustain higher firing rates than the other neurons. Top-down attention can be imposed by signals from regions like prefrontal cortex. These signals bias the competition in favour of one group of neurons or another. Miller and Cohen (2001) extended this idea to include general decision making, sensory-motor transformation, maintenance of working memories, and retrieval of long term memories. They hold that these processes can all be viewed as competitive computations. For example, sensory input gives rise to multiple motor plans for different potential responses, and the neuronal representations of these motor plans compete, with the winner expressed as behaviour. Miller and Cohen (2001) suggested that prefrontal cortex maintains representations of pertinent rules and of the desired goals in a situation, and that these goal representations then exert executive control by biasing the competitive computations in other brain regions. This view is supported by an extensive literature on the prefrontal cortex (for review, see Miller and Cohen 2001). Area DLPFC is currently the best-understood prefrontal subregion in terms of eye movement control, and I will now focus on DLPFC.

Anatomically, DLPFC is well-connected to other cortical eye movement areas including FEF, SEF, LIP, and SC (see *Section 1.1.6*). Lesions to DLPFC do not alter the ability to generate saccades, but they do impair the ability to suppress the visual grasp reflex, thereby increasing error rates in the antisaccade task (Guitton et al. 1985; Pierrot-Deseilligny et al. 1991b; Pierrot-Deseilligny et al. 2005; Ploner et al. 2005). Neuronal recording results in DLPFC are consistent with a role in top down control of eye movements. Everling and Desouza (2005) recorded from neurons in and around DLPFC

while monkeys performed a version of the antisaccade task in which explicit task cues were not given. Rather, prosaccades and antisaccades were presented in alternating blocks, and the monkeys had to figure out which task to perform based on which response yielded a reward. This approach had the advantage that no changes in fixation colour were present as potential confounds. One quarter of DLPFC neurons exhibited selectivity either for prosaccades or for antisaccades during the fixation period before peripheral stimulus presentation. About two thirds of task selective neurons preferred prosaccades, with the remainder preferring antisaccades. Importantly, task selectivity was not present before antisaccade error trials, suggesting that the selectivity before correct trials represents preparatory set to perform the antisaccade task. Everling and Desouza (2005) also investigated presaccadic activity during the 250ms preceding saccade initiation. Only very small (not statistically significant) proportions of DLPFC neurons displayed simple (contralateral versus ipsilateral) selectivity either for stimulus location or for saccade end point. In contrast, about one fifth of DLPFC neurons exhibited discharge patterns based on an interaction between stimulus location and saccade endpoint. Using antidromic methods, Johnston and Everling (2006b) isolated and studied DLPFC neurons that projected to SC. A third of DLPFC cortico-tectal neurons exhibited task selectivity during the fixation period preceding peripheral stimulus presentation. Unlike the DLPFC population as a whole, four fifths of task-selective cortico-tectal neurons preferred antisaccades. The authors suggested that these antisaccade-selective signals to SC could underlie the presaccadic suppression of firing rates seen in SC burst neurons in the antisaccade task (see *Section 1.1.4*). The obvious supposition is that DLPFC input to FEF probably underlies the similar presaccadic decrease in FEF saccade-related neurons' firing rates in the antisaccade task (see *Section 1.1.14*), though this has not been directly tested. So far, I have discussed how DLPFC might exert executive control on other brain regions but not how and when the brain might determine that executive control is necessary in the first place. This will be the topic of the next section.

1.1.18 – Anterior Cingulate Cortex (ACC) and Conflict Monitoring

In monkeys and humans, anterior cingulate cortex (ACC) is located immediately ventral to SEF in the cingulate sulcus and gyrus. ACC has been implicated in conflict

monitoring. The conflict monitoring hypothesis holds that computational conflict can serve as a measure of when executive control is necessary and that ACC, in particular, monitors conflict and then signals regions like DLPFC and SEF to exert executive control on other regions (see Botvinick et al. 2004). Firstly, ACC is anatomically well-connected with both DLPFC and SEF (Bates and Goldman-Rakic 1993; Paus et al. 2001; Wang et al. 2004). Lesions to ACC also increase the rate of error commission in the antisaccade task (Gaymard et al. 1998). ACC has been studied using fMRI with a range of tasks thought to elicit computational conflict. Response inhibition tasks require suppression of a prepotent response in favour of an alternative. ACC activation is greater for high conflict versus low conflict trials in several such tasks including the antisaccade task (Ford et al. 2005), the Stroop task (see MacLeod and MacDonald 2000; Barch et al. 2001), the flanker task (Botvinick et al. 1999; Casey et al. 2000; Hazeltine et al. 2000; van Veen et al. 2001; Bunge et al. 2002; Durston et al. 2003), and the go/nogo task (de Zubicaray et al. 2000; Braver et al. 2001; Durston et al. 2002). Underdetermined response tasks involve competition among multiple potential responses, all of which are equally valid. ACC exhibits increased activation in several underdetermined response tasks including the stem-completion task (Palmer et al. 2001) and the verb generation task (Thompson-Schill et al. 1997; Barch et al. 2000). ACC also displays greater activation on error trials versus correct trials for a variety of tasks (Carter et al. 1998; Kiehl et al. 2000; Braver et al. 2001; Menon et al. 2001). It has been argued that error processing can be subsumed under the conflict monitoring scheme since competing motor plans for both the correct and incorrect responses are thought to be co-active around the time of error commission (see Botvinick et al. 2004).

As an alternative to the conflict monitoring hypothesis, Posner and DiGirolamo (1998) have proposed that ACC plays a role in exerting top down executive control, along the lines of what DLPFC is thought to do (see *Section 1.1.17* above). Botvinick et al. (2004) claim that results from several fMRI studies refute this hypothesis. These studies found that ACC exhibited greater activation for trials with high conflict and low top down control requirements compared to trials with low conflict and high top down control, using the Stroop task (Carter et al. 2000; MacDonald et al. 2000) and the flanker task (Botvinick et al. 1999; Casey et al. 2000; Durston et al. 2003). However, monkey

neuronal recording studies did not find neurons carrying conflict-related signals in ACC using the saccade countermanding task (Ito et al. 2003) or a stimulus-response task with various conflict-inducing permutations (Nakamura et al. 2005). These studies did find ACC neurons that exhibited different discharge rates for correct versus incorrect responses and for rewarded versus unrewarded behaviour. Johnston et al. (2007) also found that many ACC neurons exhibited task type selectivity (prosaccade versus antisaccade) during the preparatory period preceding stimulus presentation and response execution, and they suggested a role for ACC in executive control rather than conflict monitoring. Brown and Braver (2005) proposed the error likelihood hypothesis, in which ACC predicts the likelihood of error on a given trial. Conflict monitoring and response error detection can be formulated as special cases of error likelihood monitoring.

More research will be required to resolve the discrepancies among the above-presented results and hypotheses. In *Section 1.1.12*, I discussed the difficulty of differentiating between visuo-spatial attention and intentional saccade coding in LIP, based on the tight relationship between attention and intention. The case seems to be similar in ACC. Conflict monitoring, error detection, error likelihood, and exertion of top down executive control are all intertwined, and it will likely turn out that all of these processes are submanifestations of ACC's true function.

1.1.19 – Synthesis and Summary of Executive Saccade Control

I started this discussion with the question how does the brain perform the antisaccade task? We are now in the position to provide a preliminary answer. The task recruits processing at all levels of the saccade system's hierarchy, from automatic response triggering to exertion of executive control. At the lower levels, the brainstem saccade generators, which control the extraocular muscles, receive input from the superior colliculus (SC) and lateral intraparietal area (LIP). SC and LIP perform the sensory-motor transformations that generate automatic saccades toward novel, salient visual stimuli, including the peripheral stimulus in the prosaccade task. Inhibition of the automatic saccade is the job of higher areas in the prefrontal cortex, specifically the dorsolateral prefrontal cortex (DLPFC) region. The need to exert top down inhibition of the automatic response is signaled by the anterior cingulate cortex (ACC) based on

response conflict detection. Remapping of the stimulus into its mirror position is not currently understood, but it might involve LIP and prefrontal cortex. Triggering of the voluntary antisaccade response originates in the frontal and supplementary eye fields (FEF and SEF).

The above description of how the brain might perform the antisaccade task is a good starting point, but many fundamental questions remain unanswered. Different experimental techniques paint somewhat different overall pictures of the eye movement control system. Attempts have been made to assign specific, well-defined functions to specific brain regions. The human lesion literature has been most successful with this approach, which is epitomized by the work of Pierrot-Deseilligny and colleagues (see Pierrot-Deseilligny et al. 2003). The description given above of antisaccade performance is based on their model. On the other hand, evidence from neuronal recording studies and particularly fMRI studies does not seem to favour this approach. Neuronal recording studies initially described LIP as a visuo-spatial attention area or as an area coding for saccade commands, but the evidence suggests increasingly that LIP performs both functions (see *Sections 1.1.8* and *1.1.12*). FEF was once thought only to trigger deliberate saccade responses but is now known also to be involved in stimulus discrimination and selection as well as in attention (see *Sections 1.1.7* and *1.1.11*). Evidence exists that SEF is involved in a range of computations, including remembering and manipulating sequence information, coding arbitrary stimulus-response associations, triggering antisaccades, and processing information related to correct and error performance and to reward (see *Section 1.1.9*). ACC is suggested to signal executive control regions like DLPFC when to exert control. It is actively debated what ACC neurons might code, whether it be response conflict, reward receipt, error performance, error likelihood or some combination of these (see *Section 1.1.18*). ACC might also be involved in exerting executive control, itself. In brief, cortical regions do not seem to perform just one function, and many functions seem to be distributed across multiple regions. This is highlighted by functional imaging studies, which have found that the entire cortical saccade system, including LIP, FEF, and SEF as well as DLPFC and ACC, exhibits activation changes across a range of eye movement tasks and task subcomponents (see

Section 1.1.15). Discussions of cortical eye movement control have come increasingly to emphasize distributed processing and distributed functional organization.

I would like to make one final point on the topic of distributed organization. One criticism of this idea is that using multiple brain regions to perform very similar functions would be wasteful. Given that neural tissue is expensive to build and maintain, we would not expect large parts of the brain to be functionally homogeneous. This point is valid, but it addresses only a very narrow version of distributed organization. The hypothesis that such functions as attention or sensory-motor transformation are distributed across multiple cortical regions requires only that these regions contribute jointly in some way to those functions. It does not necessitate that the regions contribute identically.

A good analogy is a nation's provision for the physical defense of its citizens against attack. Many institutions contribute to defense including not only the various branches of the military, such as the army, navy, and air force, but also the intelligence service, national and local police forces, and many others. We could describe the function of defense as being distributed across all of these organizations. It is interesting to note that the different institutions must often work together very closely, mirroring the tightly-coupled interactions among various cortical regions. The key point is that the differences among the organizations listed above are based on many factors, such as where they operate in terms of national borders and in terms of terrain, the type of aggression they guard against, the human and material resources they employ, and so on. These factors must be understood if we are to understand the differences among the organizations. It is probably because we do in fact understand the relevant factors that we don't typically lump all the organizations I mentioned under a single heading like defense. In conversation, we tend to discuss the police in the context of preventing crime, the intelligence service in terms of spying, and the military in the context of fighting battles. Perhaps then, the fact that functions such as attention and saccade control seem to be distributed across multiple brain regions reflects more our ignorance of what those processes involve than a deep insight into the brain's functional organization.

1.1.20 – Research Questions and Hypotheses

Systems neuroscience is a young field, only just beginning to grapple with the brain's complex and subtle organization. It behooves us, then, to conduct our research by starting at the beginning, with simple questions. Even incorrect simple hypotheses can contain elements of the truth, and they stimulate subsequent investigation. My thesis research is grounded in a very simple conception of executive processing in the antisaccade task. I took the view that successful antisaccade performance requires two separate processes, inhibition of the prepotent automatic saccade and generation of a voluntary antisaccade response. In my first experiment, which is described in *Chapter 2*, I used fMRI to investigate the relative contributions to these two putative functions by various cortical saccade regions. I should note that automatic saccade inhibition and voluntary saccade generation could consist of multiple subfunctions. More importantly, the two functions need not be implemented as completely distinct processes in the brain. As a starting point for my research, I assumed these processes to be at least partially separate and differentiable with fMRI. My expectation was that either success or failure to find such a separation would open interesting, new questions.

Experiment 1 compared prosaccades, antisaccades, and nogo trials using rapid event-related fMRI (see *Section 1.2.4*). Each trial included an instruction event (coloured fixation point) that indicated which task to perform, followed by peripheral stimulus presentation and an appropriate response by the subject. Required responses were to look toward the stimulus on prosaccade trials, to look away from it on antisaccade trials, and to maintain central fixation on nogo trials. Where antisaccade trials involved both saccade inhibition and generation processes, nogo trials required only saccade inhibition and prosaccade trials required only saccade generation. Our design also allowed us to separate fMRI activation patterns related to the instruction and response events. Our hypothesis was that cortical regions involved in antisaccade performance are involved preferentially either in saccade generation or saccade inhibition. We expected these regions to exhibit greater activation on prosaccades versus nogo trials or vice versa during the instruction event and/or the response event.

The first experiment yielded two surprising results. First, we did not find instruction-related differences between prosaccades and antisaccades. This was surprising

as instruction differences had been demonstrated previously with widely-spaced fMRI experiments (Curtis and D'Esposito 2003; DeSouza et al. 2003; Ford et al. 2005). Another surprising result from Experiment 1 was that prosaccades and nogo trials evoked almost identical response-related activation levels in cortical saccade regions, despite the fact that prosaccade trials required saccade generation whereas nogo trials required inhibition of saccades.

Experiment 2 investigated the lack of instruction-related differences between prosaccades and antisaccades in Experiment 1. One important difference between Experiment 1 and previous fMRI studies was that it used a rapid event-related design. We also had technical concerns regarding the capability of the particular design used in Experiment 1 to completely separate instruction- from response-related differences (see *Section 1.2.5* and *Chapter 3*). Experiment 2 tested whether a rapid event-related design is in fact capable of detecting instruction-related differences. We compared prosaccades and antisaccades using a rapid design that incorporated a different method of separating instruction- and response-related signals from that used in Experiment 1. This alternate method allowed us to resolve differences in instruction-related activation between prosaccades and antisaccades.

Experiment 3 sought to address the lack of activation differences between prosaccades and nogo trials in Experiment 1. We compared prosaccade and nogo trials using a rapid fMRI design. As in Experiments 1 and 2, we separated instruction- and response-related fMRI signals. In this paradigm, we had subjects perform twice as many prosaccade trials as nogo trials. This manipulation was meant to make the prosaccade response more automatic, thereby requiring the subject to recruit saccade inhibition processes at an increased level to perform the nogo task successfully. We hypothesized that regions involved in saccade inhibition should exhibit greater activation for the rare nogo trials compared to the frequent prosaccade trials because of the increased requirement for saccade inhibition in the nogo trials. We found just this pattern of results, with various cortical regions exhibiting either instruction- or response-related differences.

1.2 – BOLD fMRI and Rapid Event-related Designs

Before continuing with my own research in *Chapters 2-4*, I will devote the second half of *Chapter 1* to an overview of fMRI and specifically of rapid event-related methodology, which I used throughout my thesis work.

Most functional magnetic resonance imaging (fMRI) protocols measure the blood oxygenation level dependent (BOLD) signal, which is a physiological correlate of population-level neuronal activity. Changes in neuronal activity lead to changes in the blood perfusing the neuronal tissue. Blood flow rate, blood volume, and the ratio of oxyhaemoglobin to deoxyhaemoglobin are altered. These changes alter the signal measured with T2*-weighted MRI (Bandettini et al. 1992; Ogawa et al. 1992; Kwong et al. 1992; also see Logothetis and Wandell 2004). Functional MRI scanning consists of taking a series of T2*-weighted image volumes while the subject performs a psychological task. The researcher then looks for task-related BOLD signal changes in various brain regions. The fact that fMRI measures the BOLD signal, rather than action potentials or dendritic potentials, has implications for how we interpret fMRI data and for how we design fMRI experiments.

Section 1.2.1 will cover limitations in the design of fMRI experiments imposed by the BOLD signal. *Sections 1.2.2* and *1.2.3* will discuss traditional blocked and widely-spaced event-related experimental design methods. Rapid event-related fMRI, which allows one to circumvent some of the constraints of earlier fMRI designs, is discussed in *Section 1.2.4*. I will also cover more sophisticated rapid fMRI designs using compound (multi-component) trials in *Section 1.2.5*. I will conclude in *Section 1.2.6* with a discussion of non-linearity in the BOLD signal and how it affects rapid fMRI designs. *Appendix 3* discusses the analysis of fMRI signals based on the general linear model (GLM) including the issue of non-orthogonality ("non-separateness") of the predictions one builds into the GLM.

1.2.1 – Implications of the BOLD Response for fMRI Experimental Design

The haemodynamic response function (HRF), which couples neuronal activity to the BOLD signal, acts as a low-pass filter, attenuating the high-frequency details of the underlying neuronal activity that gives rise to the BOLD signal. A brief

neurocomputational event, such as the neuronal response in primary visual cortex to a flashed visual stimulus, evokes a BOLD response that is first lagged by 2-3 s, takes 4-6 s to reach peak amplitude, and requires a total of 10-14 s to return to baseline from beginning to end (see *Figure 1.1*) (Boynton et al. 1996; Dale and Buckner 1997; Friston et al. 1998). Adjacent neuronal events that occur more closely together than the 10-14s evolution time of the haemodynamic response evoke a succession of BOLD signatures that overlap and summate with each other (see Buxton and Frank 1997). This phenomenon must be accounted for in the design of fMRI experiments. Suppose we wanted to compare the activation associated with two behavioural tasks using fMRI. If we simply presented the two tasks to the subject randomly interleaved at a frequency of one trial per second, the resultant BOLD signal would be a summation of signals from the two tasks with no obvious way to separate the two. Several strategies have been developed to accommodate the long evolution time and summation of the BOLD signal. Blocked and widely-spaced event-related designs space task events widely enough that their individual BOLD signal components do not overlap. Rapid event-related designs place task events close together and incorporate careful timing arrangements to allow for deconvolution of event-specific signals. I will describe each of these methods in the next three sections.

1.2.2 – Blocked fMRI Designs

In blocked fMRI designs, trials of a given type are presented close together in long batches, or blocks, often with periods of rest in between the blocks (*Figure 1.2A*) (Bandettini et al. 1992; Kwong et al. 1992; Ogawa et al. 1992; also see Aguirre and D'Esposito 1999; Donaldson and Buckner 2001; Huettel et al. 2004). The BOLD response is allowed to summate across trials within a block (*Figure 1.2B*), and the summated BOLD signal levels from different block types are then compared. For example, to localize saccade-related brain regions, one could compare blocks of visually-guided trials with blocks of fixation. Saccade blocks would evoke sustained, elevated BOLD signal levels in saccade-related regions, while the BOLD activation would decay back to baseline during the fixation periods. In typical blocked fMRI designs, blocks last about 10 to 20 s and might include 10 to 20 trials. To achieve optimal statistical

Haemodynamic Response Function (difference of gammas)
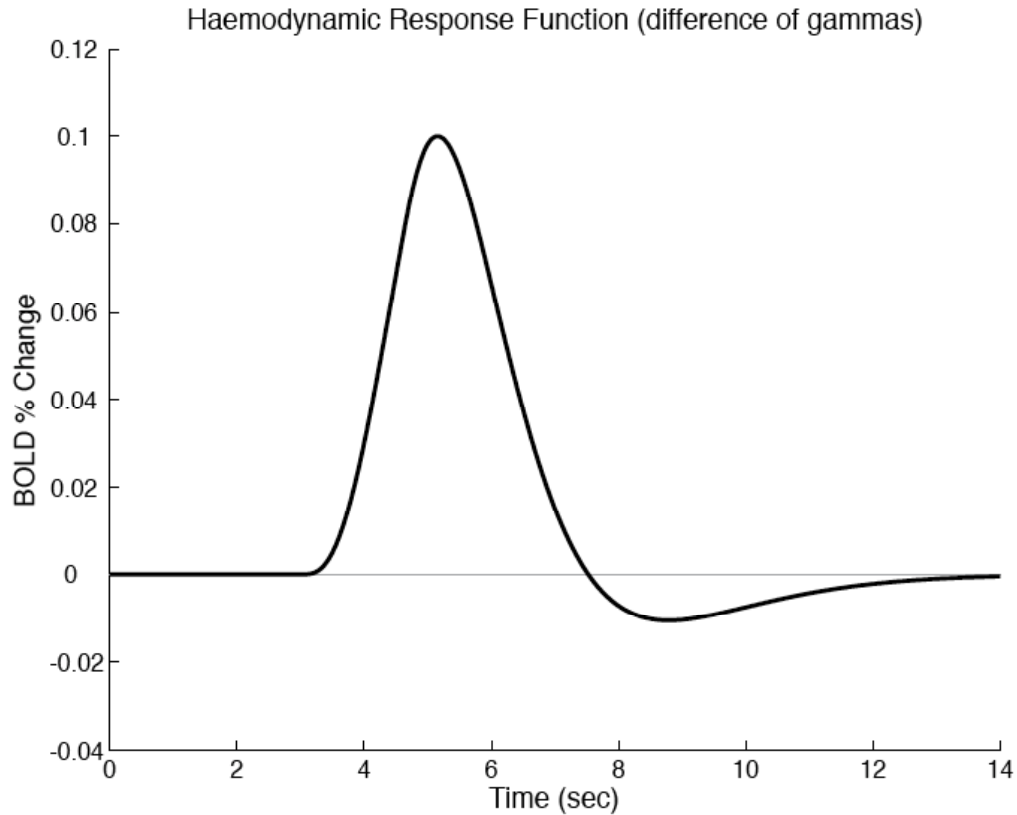
Figure 1.1 – Haemodynamic Response Function

This figure shows a difference of gammas curve, which is one widely-used model of the haemodynamic response function. Here the BOLD signal would be evoked by a brief neuronal event occurring at time 0 s. Note the lag of about 3 s between time 0, corresponding to the neuronal activity, and the deflection of the BOLD signal above baseline. Also note the long time taken to return to baseline, 14 s in this case. (Also see *Appendix 3.2.*)
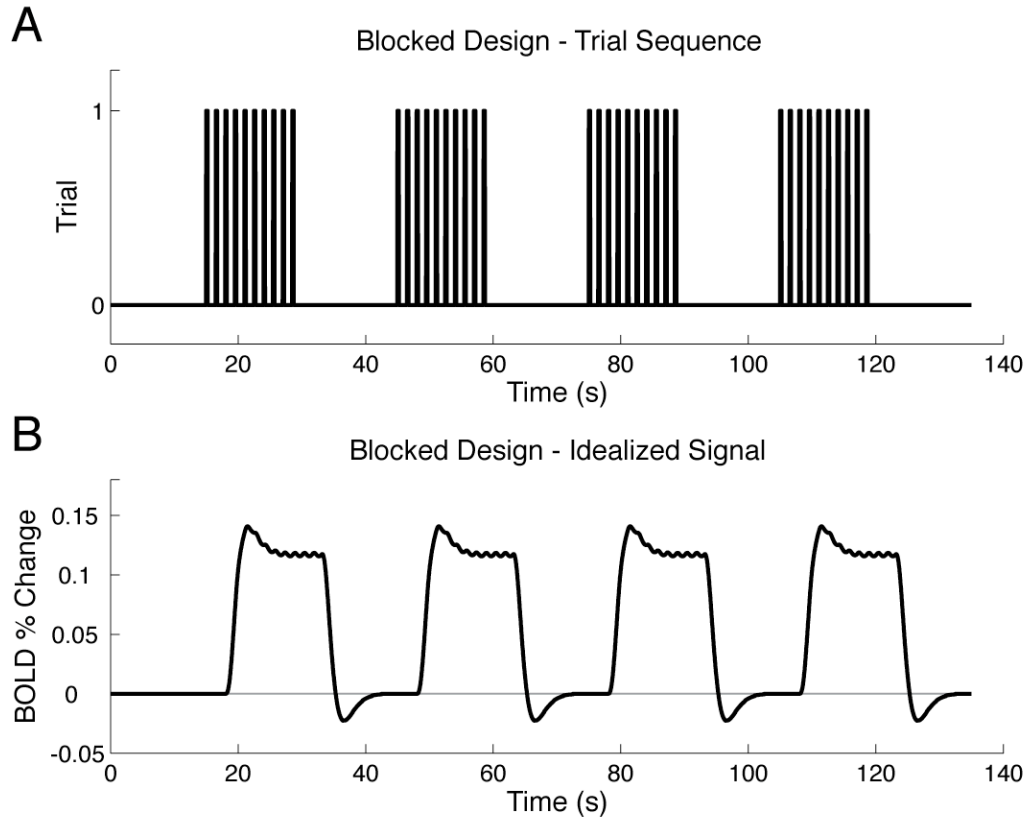
Figure 1.2 – Blocked fMRI Design with One Task

This figure illustrates a blocked fMRI design with a single trial type.

A: An impulse sequence where each impulse defines the onset of a single trial. Notice that the trials are presented in batches or blocks, interspersed with long rest periods.

B: BOLD signal evoked by the trial sequence shown in *A*. The BOLD signal summates across the closely-spaced trials in each block. The standard statistical comparison with a single trial blocked design would be between the summated block activation and the rest periods. This signal is "idealized" because it does not contain noise as a real fMRI signal would. (Also see *Section 1.2.2.*)

sensitivity in a pair-wise comparison of this type, one should collect an equal number of data points for each of the two conditions (see Aguirre and D'Esposito 1999; Donaldson and Buckner 2001; Huettel et al. 2004). In the visually guided saccade example, the fixation blocks should be equal in number and duration to the saccade blocks.

To compare two tasks, such as the prosaccade and antisaccade tasks, one could present subjects with blocks of prosaccades and blocks of antisaccades, perhaps separated by fixation rest periods (*Figure 1.3*). If one were interested only in the difference between prosaccades and antisaccades and not in localizing saccade-related activity against a baseline rest condition, the fixation periods could be excluded all-together. In this case, the alternating prosaccade and antisaccade blocks would cause the BOLD signal to stabilize at steady-state levels specific to each task type, again allowing for straight-forward comparison between block types.

The typical 10-20 s block length mentioned above results from a trade-off between the BOLD signal's low-pass filter property and fMRI signal noise. The haemodynamic response filters out high frequency signal components more than low frequency components. The MRI scanner also introduces noise that obeys an inverse frequency relationship, such that the lower frequencies of an fMRI time series are more noisy (Zarahn et al. 1997b; Aguirre and D'Esposito 1999). This effect becomes especially pronounced for frequencies under 0.017 Hz or 1 cycle / minute. Therefore, using a task with high signal power at very high or very low frequencies is inadvisable. In addition, human physiology introduces noise at specific higher frequencies, because of breathing and the heart beat. The heart beat, at a frequency of approximately 1-1.5 Hz can be particularly problematic because it can alias with the fMRI sampling rate, or volume collection rate, which is typically one sample every 1-4 s (or 0.25-1 Hz). Taken together, these factors cause the optimal task frequency for blocked designs to lie in the 0.025-0.05 Hz range, or about one cycle every 20 to 40 seconds. Therefore, when comparing two different tasks in a blocked design, setting the block length in the range of 10 to 20 s will put the evoked BOLD signal's maximal power in the ideal 0.025-0.05 Hz band.
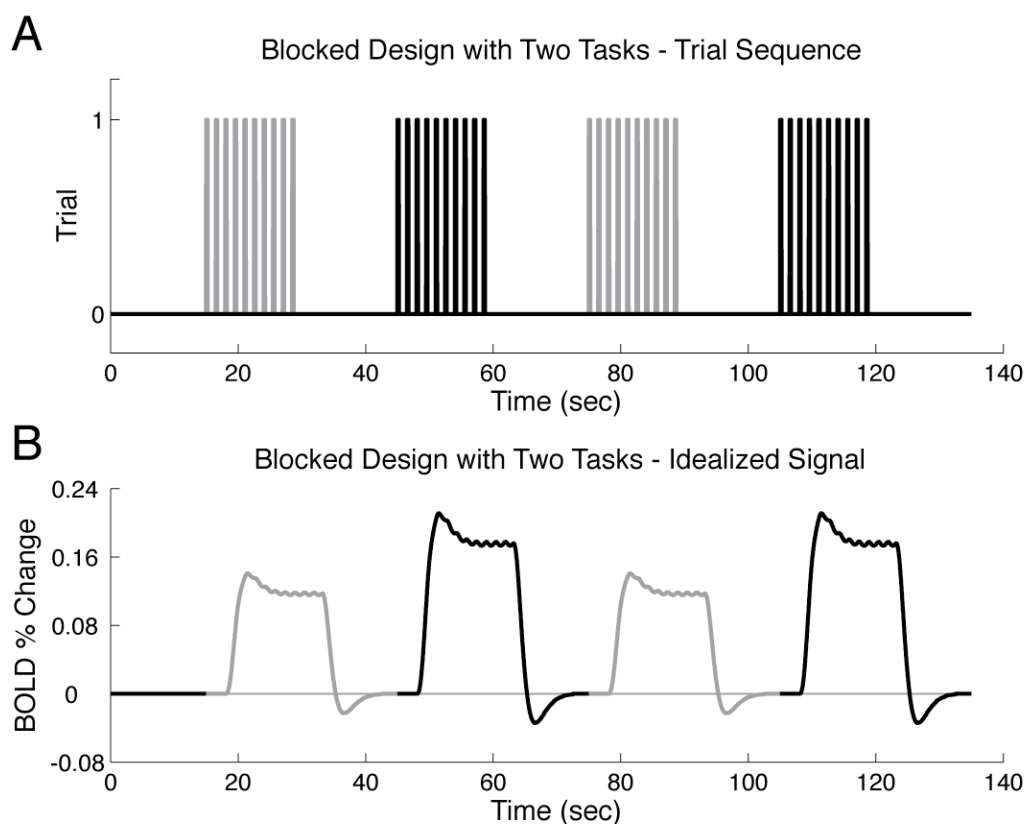
Figure 1.3 – Blocked fMRI Design with Two Tasks

This figure illustrates a blocked fMRI design comparing two tasks.

A: Impulse sequence showing trial onsets for two tasks, coloured gray and black.

B: BOLD signal evoked by impulse sequence in *A*. Notice that BOLD activation summated across the gray trial blocks is lower than that from the black trial blocks. Typical comparisons in such a design would be between the black and gray trial blocks and between the trial blocks and the rest periods. (Also see *Section 1.2.2*.)

A disadvantage of the blocked approach is that one cannot separate signals from individual trials or from subcomponents within tasks like the memory-guided saccade task, which includes memory encoding, maintenance, and retrieval stages. Error trials cannot be discarded or investigated separately. In addition, one is obviously prevented from interleaving trials of different types together within a block. Interleaving prevents subjects from anticipating the upcoming task type and also is necessary when studying the effects of repeating versus switching task set (Monsell 2003).

Despite these disadvantages, blocked designs have several strengths that have contributed to their wide-spread use, particularly in the early years of fMRI. Because they summate across trials, blocked designs evoke very strong BOLD responses, making them ideal for detecting small effects in the presence of background noise. Designing the timing and sequence structures for blocked experiments is fairly straight-forward, as is processing and analysing blocked fMRI data. Of course, devising informative tasks and interpreting statistical results from blocked experiments takes considerable thought and creativity, as is the case with all fMRI experiments.

1.2.3 – Widely-spaced Event-related fMRI Designs

Unlike blocked designs, event-related experimental designs allow investigation of signals specific to individual trials or trial subcomponents. The widely-spaced event-related technique interleaves long rest periods between trial events such that the BOLD signal can return substantially or fully to baseline between the events (*Figure 1.4*) (Buckner et al. 1996; Dale and Buckner 1997; Josephs et al. 1997; McCarthy et al. 1997; Zarahn et al. 1997a; also see Aguirre and D'Esposito 1999; Donaldson and Buckner 2001; Huettel et al. 2004). An event, in this context, can refer to a whole trial or to a subcomponent of a compound trial. Inter-event spacing in widely-spaced designs is typically at least 7 to 8 s, though 10 to 14 s gives better separation between signals. For example, several studies have compared antisaccades and prosaccades using widely-spaced event-related fMRI (Connolly et al. 2002; Curtis and D'Esposito 2003; DeSouza et al. 2003; Ford et al. 2005) (also see *Section 1.1.15*). Each trial began with a trial type instruction, which was followed by a long fixation interval, at the end of which the subject made a prosaccade or antisaccade as previously instructed in response to a

A: Widely-spaced Event-related Design - Trial Sequence

B: Widely-spaced Event-related Design - Idealized Signal
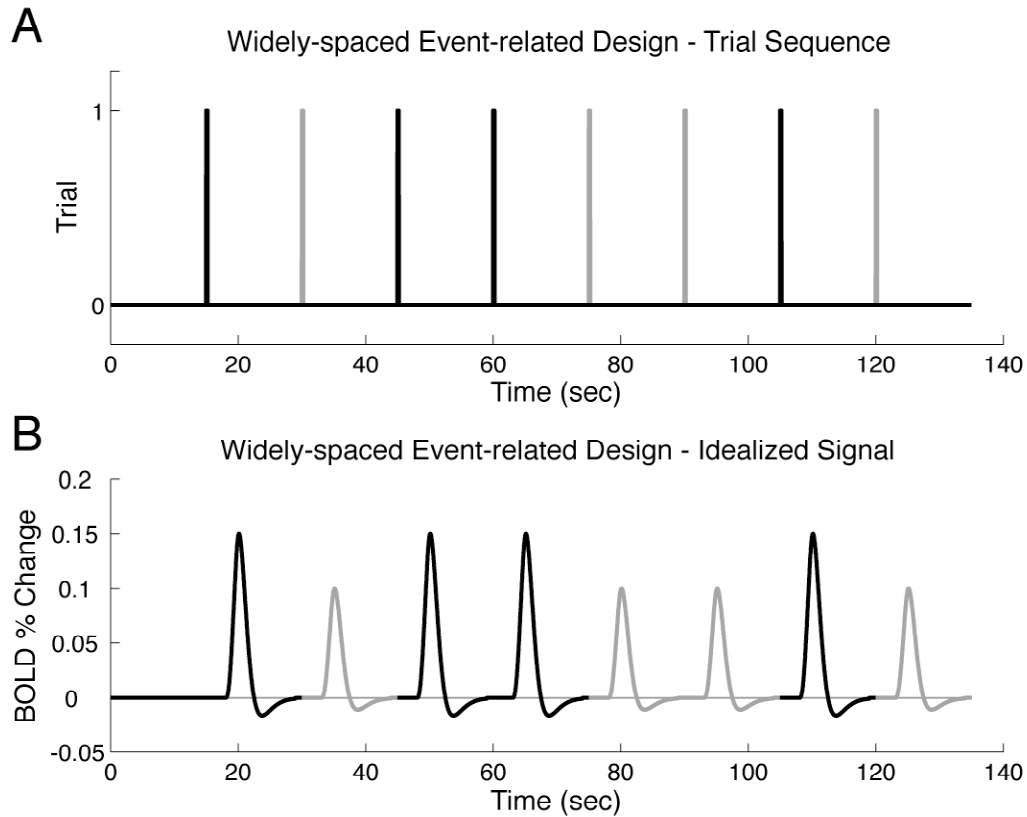
Figure 1.4 – Widely-spaced Event-related Design

This figure illustrated a widely-spaced event-related fMRI design with two tasks.

A: Impulse sequence showing trial onsets for two tasks, coloured gray and black. Note that individual trials are widely-spaced to allow the BOLD signal to return to baseline between trials.

B: BOLD signal evoked by impulse sequence in *A*. Notice that the widely-spaced individual trials evoke distinct BOLD response signatures that are easily isolated because they do not overlap and summate. (Also see *Section 1.2.3*.)

peripheral stimulus. Long fixation intervals also separated adjacent trials. These studies were able to compare activation related specifically to the appearance of the instruction cue, to maintenance and preparation of instruction set at the end of the instruction interval, and to response execution. Curtis et al. (2003) and Ford et al. (2005) were also able to identify error trials and examine them separately, something that is clearly not possible with a blocked design.

In brief, the advantages offered by the widely-spaced event-related framework are as follows. One can examine individual trials. Error trials can be discarded or examined separately if interesting in their own right. One can interleave task types and randomize the order of trial presentation within event-related trial sequences. The long inter-event intervals required in widely-spaced event-related designs impose the method's main disadvantages. Long rest periods are tedious for subjects and promote subject fatigue, both of which tend to reduce BOLD activation levels evoked by experimental tasks. The long rest intervals are also unrealistic in that human behaviour typically does not include arbitrary, long pauses between successive actions. Turning a door knob and opening the door proceed continuously and fluidly, without a long gap in between. Most psychophysical and electrophysiological experiments use fairly rapid experimental sequences, which makes their comparison with widely-spaced fMRI results more problematic. Lastly, the long rest intervals between trials use up much of the time available in an experimental session, preventing one from including very many trials. For each given subject, the antisaccade studies cited above collected at most 30 (DeSouza et al. 2003), 80 (Curtis and D'Esposito 2003), and 56 (Ford et al. 2005) antisaccade trials, in comparison to the hundreds of trials collected with other experimental methodologies (for example, see Hallett 1978; Everling and Munoz 2000). Low trial numbers can limit averaging and statistical sensitivity in some widely-spaced event-related fMRI studies. Regardless of these drawbacks, the advantages of widely-spaced event-related designs, particularly their ability to interrogate subprocesses within trials, allowed these designs to open up substantial new areas of investigation that were impossible with blocked fMRI.

Rapid event-related designs allow even more freedom in specifying trial structure than widely-spaced designs, in many cases. Rapid designs use clever timing structures rather than long fixation intervals to separate trials and trial subcomponents. To

understand rapid designs, we must understand something about how fMRI data are analyzed, which is the topic of the next section.

1.2.4 – Rapid Event-related fMRI Designs

Rapid event-related fMRI addresses the limitations of widely-spaced designs, which were discussed in *Section 1.2.3* above, by placing task events close together. This avoids the need for long fixation periods that cause fatigue and wandering attention in subjects, and valuable scanning time is not spent imaging the BOLD signal's return to baseline. Trial timing in rapid fMRI designs is similar to what is used with other techniques like psychophysics and electrophysiology, facilitating comparison of rapid fMRI results with these other methodologies. Like widely-spaced event-related designs, rapid designs allow one to examine activation patterns from individual trial types or trial subcomponents. Rapid designs also have the potential to make participation in fMRI experiments easier for certain subject groups, such as neurological patients, who can exhibit poor performance in the demanding and sometimes uncomfortable fMRI scanner environment.

The close spacing of trials in rapid fMRI designs causes the BOLD signals evoked by adjacent task events to overlap and summate, necessitating the use of deconvolution techniques to resolve event-specific activation profiles. The deconvolution approach I use is based on the general linear model (GLM), which is described in *Appendix 3.1*. The set of predictor functions used in the GLM must be linearly dependent so that we can compute a unique solution to the problem of fitting the GLM to the data. Furthermore, the predictor vectors must be chosen carefully to avoid poor conditioning (see *Appendix 3.2*) if we want our solution to be at all robust against fMRI signal noise. Rapid fMRI designs jitter, or randomly vary, the time interval between adjacent task events so as to maximize the linear separation between predictor functions.

To understand why jittering is effective, consider an fMRI design without jittering. Suppose we ran an fMRI experiment containing a single trial type presented every three seconds (*Figure 1.5A*). Suppose also that we collected one functional volume every second (volume collection time of 1s). For now, ignore fMRI signal noise.

Figure 1.5 – BOLD Summation and Rapid fMRI

A: Impulse sequence specifying the onset times of trials in a rapid event-related fMRI design with a constant 3 s inter-trial interval (ITI).

B: The black line is the BOLD signal evoked by the impulse sequence in *A*. The gray lines are the BOLD signal components from the individual trials (impulses) in *A*. The fMRI scanner would record the black line, but one is actually interested in the signal profile of the gray lines. Notice that the summated signal rapidly settles into an equilibrium consisting of the same 3 repeating values. This is a result of the constant, small inter-trial interval.

C: The reason that the summated BOLD signal (black line in *B*) settles into an equilibrium is that the constant 3 s inter-trial interval creates a linear dependence relationship among the BOLD response profiles from the individual trials (gray lines in *B*). The left panel shows the haemodynamic response function (HRF) consisting of 15 points labeled *h0* through *h14*. The right panel shows a short segment of BOLD signal (black line) that is summated over six individual trial activation profiles (gray curves). Notice that the black curve settles into an equilibrium consisting of the same three repeated values, labeled *x1*, *x2*, and *x3*. Because of the 3 s constant inter-trial interval, each of the *x1*, *x2*, and *x3* values is just the sum of five of the haemodynamic response function values separated by 3 s intervals. So, *x1=h0+h3+h6+h9+h12*, and similarly for *x2* and *x3*. This creates a problem because it is impossible to recover the 15 values of the haemodynamic response function from the three repeated *x* values in the summated signal. (It is still technically possible to recover the haemodynamic response function from the entire summated BOLD signal because the linear dependence relationship breaks down for the few points at the beginning and end. However, any such solution would be entirely dependent on those few data points and highly susceptible to signal noise.) (Also see *Section 1.2.4*.)

A Rapid Design with 3 sec Constant ITI - Trial Sequence

B Rapid Design with 3 sec Constant ITI - Idealized Signal

C HRF   Linear Dependence

$x1 = h0 + h3 + h6 + h9 + h12$
$x2 = h1 + h4 + h7 + h10 + h13$
$x3 = h2 + h5 + h8 + h11 + h14$
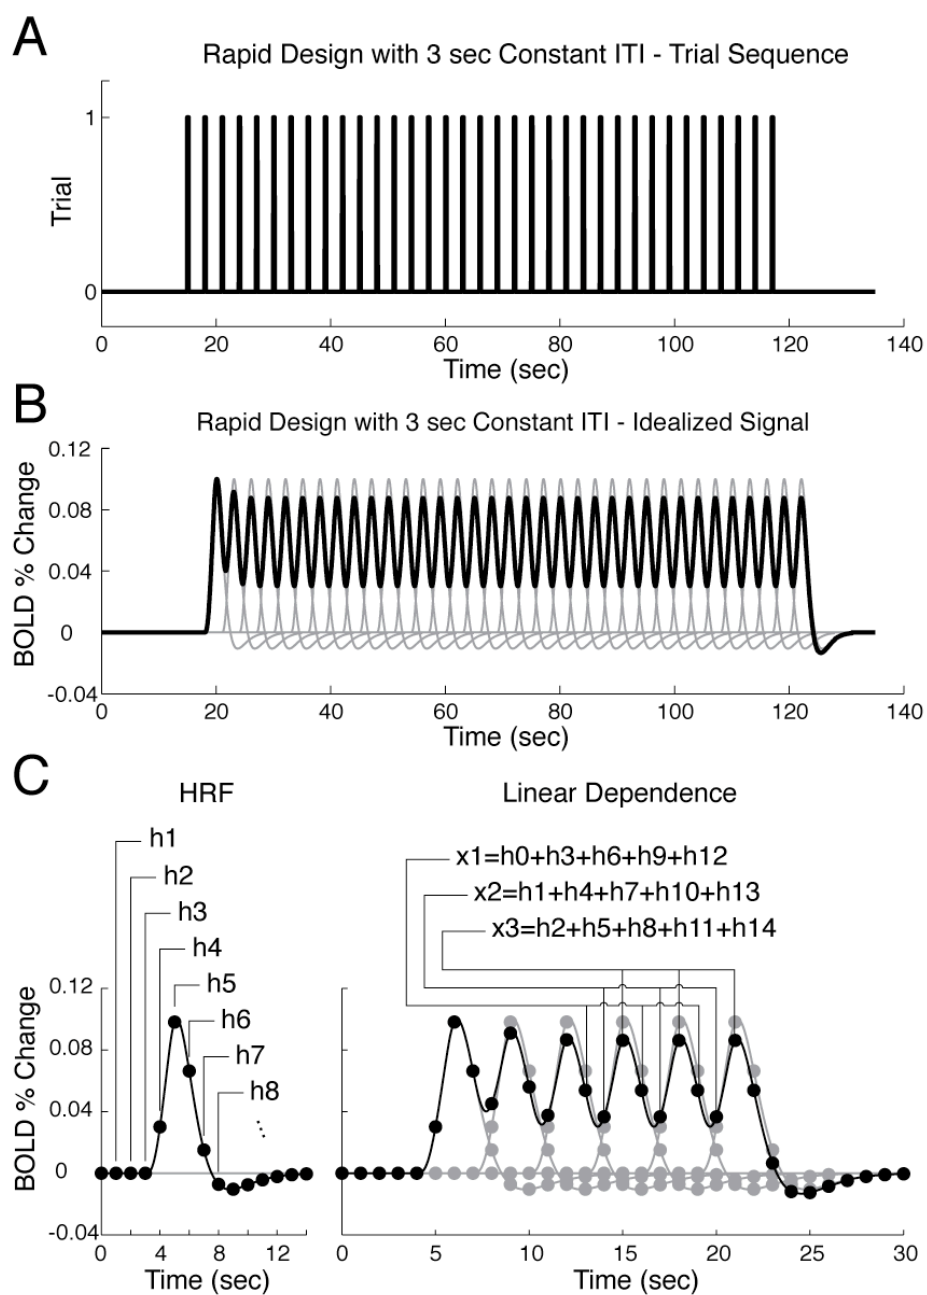
As can be seen in *Figure 1.5B* and *C*, the BOLD time course from a single voxel in this experiment would rapidly stabilize into a steady state consisting only of three values repeated in sequence. This would happen because the constant inter-trial interval would cause the BOLD responses evoked by the individual trials to overlap in a repeating pattern. The problem would get worse if we shortened the inter-trial interval to 2 s because the steady state time course would then consist only of two unique values (*Figure 1.6*, black line). If we were to shorten the inter-trial interval further to 1 s, the steady state time course would be comprised only of a single repeated value (*Figure 1.6*, gray line). The scenarios described here, with 1, 2, or 3 s constant inter-trial intervals, would not actually prevent us from fitting a GLM to the data because the pattern of repetition would not hold for those few data points at the very beginning and end of the time course. Because of the beginning and end points, it would be possible to define a set of linearly independent predictor functions, and it would then be possible to fit the GLM to the data. However, the GLM solution would be extremely fragile as it would depend entirely on those few points at the beginning and end, and those few points would be influenced by fMRI signal noise in a real world situation. Put another way, all possible predictor sets would be poorly-conditioned.

Now consider an adaptation of the above example to make it a rapid event-related design, with the inter-trial interval jittered over a set including 1, 2, 3, 4, and 5 s (*Figure 1.7A*). Notice that with jittering the fMRI time course never stabilizes into an equilibrium (*Figure 1.7B*) because its data points represent many different combinations of points from the preceding haemodynamic responses (*Figure 1.7C*). The jittering could be said to preserve the variance in the BOLD signal. The jittered trial structure can also be described with a set of predictor functions that point in fairly different directions, meaning that they are well-conditioned and not close to being linearly dependent. Note that the predictor vectors in a rapid design do not have to be (and seldom are) perfectly orthogonal. Good conditioning is sufficient for a robust, noise-resistant analysis. Thus, jittering the time interval between task events can allow one to space the events close together in a rapid event-related design. Note that 'task events' could refer to individual trials, in which case it would be the inter-trial interval that was jittered. It is also possible
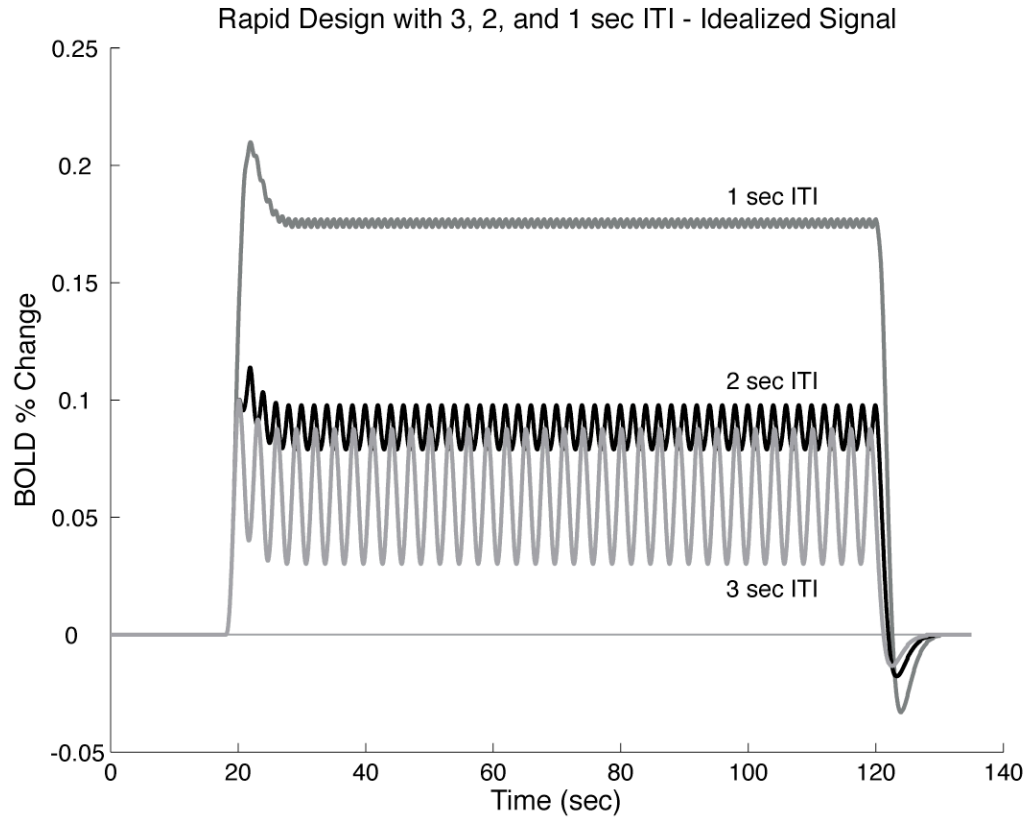
Figure 1.6 – Rapid fMRI Design with 1, 2, and 3 s ITI

As we saw previously, *Figure 1.5* illustrated the problem of BOLD summation in rapid event-related designs with constant inter-trial interval (ITI). *Figure 1.6* shows that with a 3 s inter-trial interval, the BOLD signal settles into an equilibrium consisting of 3 distinct, repeated values (light gray line). With 2 s and 1 s constant inter-trial interval (black line and dark gray line, respectively), the BOLD signal's equilibrium state consists of 2 and 1 unique, repeated values, in that order. In other words, decreasing the inter-trial interval in a constant inter-trial interval rapid design decreases the amount of information in the summated BOLD signal recorded by the MRI scanner. (Also see *Section 1.2.4*.)

Figure 1.7 – Rapid fMRI Design with Jittered ITI

A: Impulse sequence specifying the onset times of trials in a rapid event-related fMRI design with inter-trial interval (ITI) jittered, or varied, over the values 1, 2, 3, 4, and 5 s.

B: Summated BOLD signal (black line) evoked by impulse sequence in *A* and activation profiles (gray lines) from individual trials. Notice that the summated signal does not settle into an equilibrium as occurs with a constant inter-trial interval (see *Figures 1.5* and *1.6*). Jittering the inter-trial interval could be said to preserve the variance in the summated signal.

C: Jittering the inter-trial interval preserves linear independence among the BOLD activation profiles from the individual trials. The left panel shows the haemodynamic response function (HRF) consisting of 15 points labeled *h0* through *h14*. The right panel shows a short segment of BOLD signal (black line) that is summated over six individual trial activation profiles (gray curves) with progressively longer inter-trial intervals (1, 2, 3, 4, 5 s). Three arbitrary points from the summated BOLD signal are labeled *y1*, *y2*, and *y3*. The equations specifying these values are shown for illustration. Note that points *y1* and *y3*, which occupy successive peaks in the BOLD signal curve, have different values because they are the sums of different combinations of points from the haemodynamic response function. Though equations are given only for the points *y1*, *y2*, and *y3*, note that the summated BOLD signal consists of many additional distinct values, whose equations are not shown because of space constraints. Because the summated signal does include many points with distinct values, it is possible to recover the haemodynamic response function using deconvolution. (Also see *Section 1.2.4.*)

A

Rapid Design with Jittered ITI - Trial Sequence

B

Rapid Design with Jittered ITI - Idealized Signal

C

HRF

Linear Independence

$y1=h1+h5+h8+h10+h11$
$y2=h3+h7+h10+h12+h13$
$y3=h0+h5+h9+h12+h14$
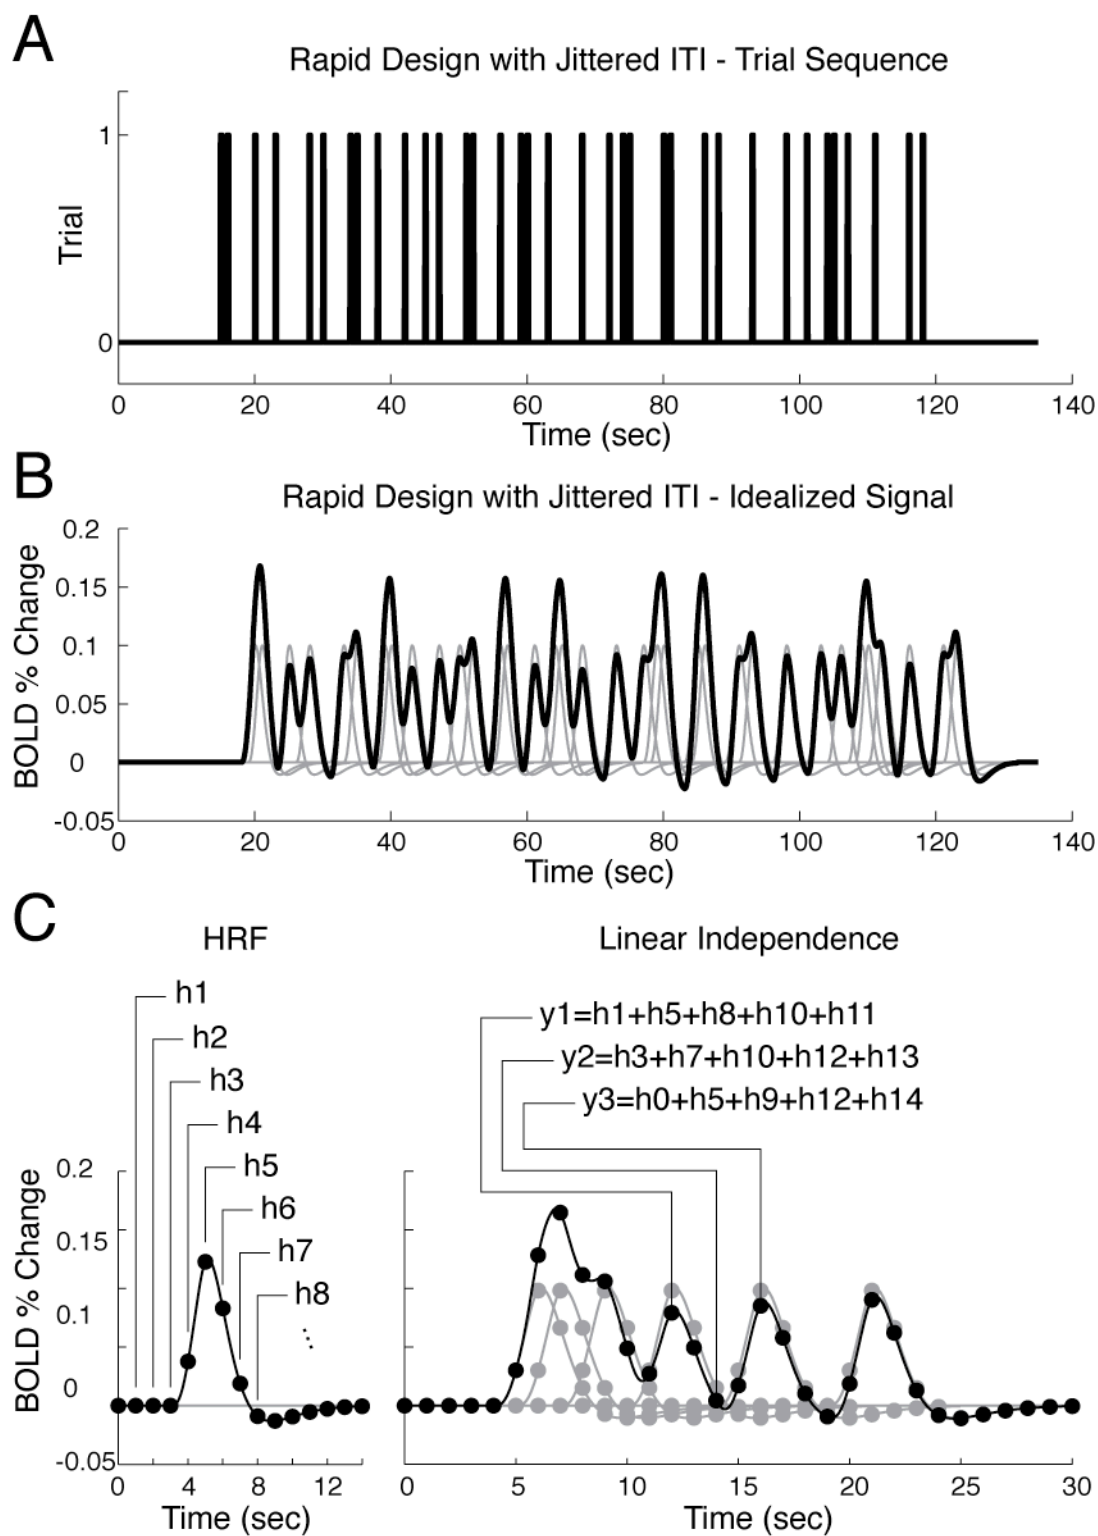
to jitter intervals within trials to allow for separation of BOLD signals evoked by subcomponents within compound trials. For example, in the memory-guided saccade task, the subject must first remember the location of a flashed visual stimulus and then generate a saccade to the remembered location after a delay period. By jittering the length of the delay period, it is possible to separate stimulus- and saccade-evoked BOLD signals in this task. I will discuss the use of compound trials in rapid fMRI designs in *Section 1.2.5*.

Once fMRI data have been collected using a proper rapid event-related trial design, with jittered inter-event intervals, the actual deconvolution process is straight-forward. As is widely done in GLM-based fMRI analysis, one could simply convolve a model of the BOLD response, for example a gamma function, with impulse sequences that specify the start times of the events of interest. This would provide predictor curves that would constitute the columns of the GLM design matrix (see *Appendix 3.1*). The process of fitting a GLM to rapid event-related fMRI data can effect a deconvolution because it finds the best scaling of predictor curves which model overlapping, summated BOLD signal components.

The approach described above assumes a specific shape for the haemodynamic response, whether it be gamma-shaped, difference of gammas-shaped, and so on. This assumption can yield good results if it is accurate, but mis-specifying the shape of the predictors in the GLM can cause the GLM not to model important aspects of the data (see *Appendix 3.1*). Finite impulse response predictors provide another approach to modeling the BOLD response (Serences 2004). Unlike the models described in the previous paragraph, which involved convolving a BOLD response model with impulse sequences defining task event onset times, a finite impulse response model consists simply of the impulse sequences, themselves, as well as time-shifted copies of them (*Figure 1.8A*). Each impulse sequence models one data point in the evolution of the BOLD response curve (*Figure 1.8B*). For example, suppose we were imaging a subject performing visually guided saccades in a jittered rapid design. Suppose also that we used a volume collection time of 1 s and wanted to model saccade-evoked activation using a finite impulse response predictor set over 8 s, which would be 8 functional data points.

A: Finite Impulse Response Predictors

B: Estimated BOLD Response
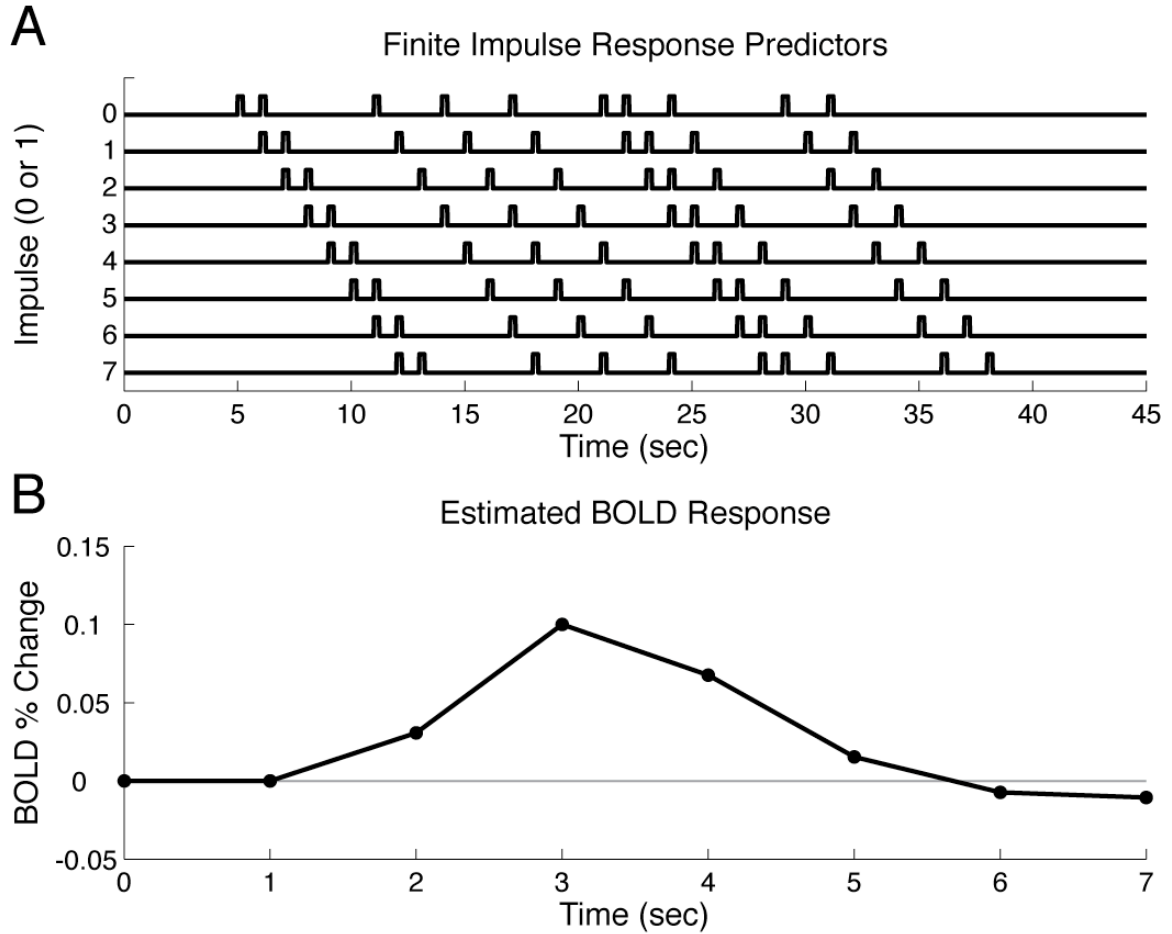
Figure 1.8 – Finite Impulse Response Predictors

A: Illustration of eight impulse sequences, which are time-shifted copies of each other. This set of impulse sequences constitutes a finite impulse response predictor set that can be used to deconvolve the BOLD response profile in a single task rapid event-related design. The sequence labeled *0* indicates the onset times of the trials, and it models the first data point in the BOLD response profile. The sequences labeled *1* through *7* model subsequent time points in the BOLD response. The individual impulses have been widened for visual clarity; each impulse would actually have a duration of one sample point.

B: Example BOLD response curve that might be deconvolved using the predictor set in *A*. Time point *x* in the curve is derived from impulse sequence *x* in *A*, where *x* is *0, 1, 2*, and so on up to *7*. (Also see *Section 1.2.4*.)

A finite impulse response model would include the impulse sequence specifying the onset of each trial and 7 copies of this impulse sequence, with each successive copy shifted one more second in time later than the last. Thus, we would have a total of 8 copies of the impulse sequence, with one copy per second. After fitting the resulting GLM to the data, we would have 8 coefficients corresponding to the 8 impulse sequences in the finite impulse response predictor set. These 8 coefficients would comprise a deconvolved time course of the BOLD response for the visually guided saccade task in our experiment. The advantage to using a finite impulse response model is that it involves no assumptions about the shape of the haemodynamic response other than the time it takes to return to baseline. In addition, using the GLM to begin with assumes that the haemodynamic response is linear (see *Appendix 3.1*).

The last two paragraphs illustrate two different approaches to analyzing fMRI data. One could either model the haemodynamic function explicitly, for example with a gamma function, or one could use a finite impulse response predictor set to allow the GLM to determine the haemdynamic response shape with more freedom. Both approaches are valid. Their relative strengths and weaknesses have to do with what is called the bias-variance tradeoff (see p. 87 of Haykin 1999). Any estimation procedure can incorporate more or less prior knowledge about the signal being estimated. Building in prior expectations can make the estimation more robust against random noise in the data, thereby reducing the variance of the estimate. The risk is that incorrect expectations can bias the estimate away from its target value. The more explicit fMRI analysis deliberately biases the results of a GLM analysis by assuming a specific shape for the haemdynamic response function. The potential disadvantage is that using the wrong shape for the haemdynamic response would cause the GLM not to capture some aspects of the data. However, there are instances when forcing the GLM to assume a certain haemodynamic shape can be useful. When dealing with weak fMRI signal changes in the face of noise, this approach causes the GLM to ignore aspects of the noise incompatible with the assumed BOLD response shape, thereby reducing the variance of the estimated coefficients. This reduction in variance can very properly increase the sensitivity of a statistical analysis when used correctly. The finite impulse response predictor set takes the opposite approach; it reduces the bias in the estimated coefficients at the cost of

increased variance. This approach will, therefore, be more susceptible to noise in the data, but it will be less liable to generate misleading results because of incorrect assumptions about the shape of the haemodynamic response function.

1.2.5 – Compound Trials and Rapid Event-related fMRI

The technique of jittering time intervals within compound trials can be used to separate activation profiles for task subcomponents. This is the approach I took in Experiment 1, in which I compared prosaccades, antisaccades, and nogo trials (see *Chapter 2*). Each trial started with a coloured fixation point to specify the trial type, followed by a flashed peripheral stimulus. Subjects had to respond to the stimulus by looking toward it on prosaccade trials, by looking away from it on antisaccade trials, and by maintaining central fixation on nogo trials. I jittered the interval between adjacent trials, and I also jittered the length of the instruction interval between instruction onset and peripheral stimulus presentation to allow for separation of instruction- and response-related activation patterns.

The drawback of separating compound trial components by jittering intervals within the trials is that complete counterbalancing is impossible. In my design for Experiment 1, response events for a given task always followed instruction events for the same task. Antisaccade responses always came after antisaccade task instructions and so on. This was logically necessary, but it created a sequence bias. Deconvolution based on finite impulse response predictor sets has been reported to be robust against sequence biasing (Serences 2004). In Experiment 1, I found that using finite impulse response predictor sets generated very noisy deconvolved time courses for events with weak BOLD activation signatures. The approach did work well for the actual saccade responses, which evoked strong signals in cortical saccade regions. To address the noise susceptibility with weaker signals, I opted to use a difference of gammas BOLD response model. By biasing the GLM to conform to an expected haemodynamic shape, I was able to reduce the impact of noise on the analysis.

Ollinger et al. (2001a; 2001b) have devised an alternative to jittering intervals within trials called the catch method. This method includes a subset of catch trials, which present only half of the task, in the trial sequence. Catch trials are also referred to as

partial trials or half trials. I used this technique in Experiments 2 and 3 (see *Chapters 3 and 4*). For example, in Experiment 2, I compared prosaccades and antisaccades using a rapid design. I used whole and half trial versions of these tasks. On whole trials, subjects were shown a task instruction and then a peripheral stimulus, to which they responded with a prosaccade or antisaccade as instructed. On half trials, subjects were shown only the trial type instruction and not the peripheral stimulus. Subjects did not make any eye movements on half trials. Thus, the half trials provided a measure of instruction-related activation separate from response-related activation. Subtracting half trial from whole trial activation patterns then yielded measures of response-related activation separate from instruction-related activation. The half trial method worked well in Experiments 2 and 3. It allowed me to resolve instruction- and response-related signal patterns separately from each other using finite impulse response predictor sets. That is, I was not forced to use the biasing solution as I was in Experiment 1. In my experience, the half trial method seems to work better than jittering within trials for separating trial components. The primary disadvantage of the half trial method is that one must include a much smaller proportion of half trials compared to whole trials to prevent subjects from expecting individual trials to be half trials. This expectation would alter the psychological nature of the tasks. One must also be careful not to include too few half trials for a valid statistical analysis. I was successful using a 1:2 ratio of half to whole trials.

Compound trials often have maintained or tonic activation components. In the memory-guided saccade task, the subject must maintain a stimulus location in memory. Using a widely-spaced fMRI design, I have shown that the delay period in the memory-guided saccade task evokes sustained, tonic BOLD activation in cortical saccade regions (Brown et al. 2004). Many experiments, including my own, also use trials consisting of a task instruction event followed by stimulus presentation and response. During the instruction period between instruction onset and stimulus onset, the subject prepares to execute the instructed task, and sustained activation has been observed during the instruction period in tasks like the prosaccade and antisaccade tasks using widely-spaced fMRI (Curtis and D'Esposito 2003; DeSouza et al. 2003; Ford et al. 2005). Deconvolving this sort of tonic activation presents some difficulty for rapid event-related fMRI. Though it is possible to jitter the length of a tonic activation period, it is impossible to jitter a

tonic period's onset and offset times. For example, an instruction period begins with instruction onset and ends with peripheral stimulus onset, and it makes no sense to contemplate jittering the instruction onset relative to the start of the instruction period. Likewise, it is meaningless to discuss a partial trial containing only the instruction onset and not the subsequent sustained activity evoked by the instruction. These considerations make it impossible to separate sustained activation at the beginning of an instruction period from the task event that defines the start of the instruction period. Likewise, it is impossible to separate sustained activation at the end of an instruction period from the task event defining the end using the rapid event-related techniques I have covered here. This limitation ultimately stems from the relatively poor temporal resolution of fMRI. The BOLD signal takes about 2-3 s to respond to a neurocomputational event, 5-6 s to peak, and 12-14 s to evolve back to baseline (see *Section 1.2.1* and *Appendix 3.2*). In comparison, neuronal events happen on the scale of milliseconds. To my knowledge, it is still an open question whether rapid event-related fMRI technique can be devised to separate the phasic and tonic activation components evoked by such task events as an instruction onset.

1.2.6 – BOLD Nonlinearity and Rapid Event-related fMRI

Analyses of fMRI data based on the general linear model (GLM) assume that the haemodynamic response is linear. A linear system obeys superposition and scaling laws. In the case of the BOLD response, linear scaling would necessitate that a 'doubling' of neuronal activity would double the magnitude of the evoked BOLD response without changing its shape. Linear superposition would require that 'adding together' two neuronal activities would generate a BOLD response that was the sum of the BOLD responses evoked by the two neuronal activity patterns separately. Put more colloquially, linear systems are 'well-behaved' because adding or multiplying the inputs to linear systems results in a similar adding or multiplying of the outputs. These properties do not hold in non-linear systems, which can exhibit small changes in output over a large part of the input range while exhibiting large changes in ouput in response to small changes in input over certain parts of the input domain.

Early studies by Boynton et al. (1996) and Dale and Buckner (1997) concluded that the BOLD response was essentially linear. It was later found that linearity holds only with long inter-trial intervals (longer than 6 s). The BOLD response exhibits non-linear refractory effects with shorter intervals (Friston et al. 1998; Huettel and McCarthy 2000). When a task event follows the preceding event by less than 6 s, its BOLD signature is reduced in amplitude. Refractory effects become more pronounced with shorter inter-event intervals, and they are particularly severe with intervals under 2 s (Buckner 1998; Burock et al. 1998; Rosen et al. 1998). One might also expect refractory effects to become larger with increasing magnitude of the preceding BOLD response, though I am not aware of any studies actually testing this possibility. Refractory effects are an example of the failure of linear superposition. Vazquez and Noll (1998) have demonstrated that linear scaling can also fail in V1 for visual stimuli with contrast levels below 40%, implying that BOLD non-linearity is based on more than just event timing.

Non-linearities in the BOLD signal can skew results of improperly designed rapid fMRI studies. If the average time interval preceding task event A were substantially shorter than that preceding task event B, for example, refractory effects would tend to suppress the magnitude of deconvolved activation time courses for event A more so than for event B. My approach in designing rapid event-related experiments has been to counterbalance all event types in terms of preceding and following time intervals and in terms of preceding and following event types, as much as possible. This measure should minimize first-order refractory effects. As discussed in *Section 1.2.5*, some compound trial designs involve unavoidable sequence effects that preclude perfect counterbalancing. Given the large number of trials in rapid designs (750 trials per subject is common), it is necessary to use computer search algorithms to devise appropriately counterbalanced trial sequences when designing rapid fMRI experiments.

1.3 – Bibliography

Abrams RA, Meyer DE, and Kornblum S. 1989. Speed and accuracy of saccadic eye movements: characteristics of impulse variability in the oculomotor system. *J Exp Psychol Hum Percept Perform* 15(3): 529-543.

Aguirre GK and D'Esposito M. 1999. Experimental design for brain fMRI. In: *Functional MRI*, edited by Moonen CTW and Bandettini PA. Berlin: Springer Verlag, p. 369-380.

Amador N, Schlag-Rey M, and Schlag J. 1998. Primate antisaccades. I. Behavioral characteristics. *J Neurophysiol* 80(4): 1775-1786.

Amador N, Schlag-Rey M, and Schlag J. 2004. Primate antisaccade. II. Supplementary eye field neuronal activity predicts correct performance. *J Neurophysiol* 91(4): 1672-1689.

Andersen RA, Asanuma C, and Cowan WM. 1985a. Callosal and prefrontal associational projecting cell populations in area 7A of the macaque monkey: a study using retrogradely transported fluorescent dyes. *J Comp Neurol* 232(4): 443-455.

Andersen RA, Essick GK, and Siegel RM. 1985b. Encoding of spatial location by posterior parietal neurons. *Science* 230(4724): 456-458.

Andersen RA, Bracewell RM, Barash S, Gnadt JW, and Fogassi L. 1990. Eye position effects on visual, memory, and saccade-related activity in areas LIP and 7a of macaque. *J Neurosci* 10(4): 1176-1196.

Andersen RA and Buneo CA. 2002. Intentional maps in posterior parietal cortex. *Annu Rev Neurosci* 25: 189-220.

Asanuma C, Andersen RA, and Cowan WM. 1985. Form of the divergent thalamocortical projections from the medial pulvinar to the caudal inferior parietal lobule and prefrontal cortex: a double-label retrograde fluorescent tracer study in macaque monkey. *J Comp Neurol* 241: 357-381.

Bandettini PA, Wong EC, Hinks RS, Tikofsky RS, and Hyde JS. 1992. Time course EPI of human brain function during task activation. *Magn Reson Med* 25(2): 390-397.

Barash S, Bracewell RM, Fogassi L, Gnadt JW, and Andersen RA. 1991a. Saccade-related activity in the lateral intraparietal area. II. Spatial properties. *J Neurophysiol* 66(3): 1109-1124.

Barash S, Bracewell RM, Fogassi L, Gnadt JW, and Andersen RA. 1991b. Saccade-related activity in the lateral intraparietal area. I. Temporal properties; comparison with area 7a. *J Neurophysiol* 66(3): 1095-1108.

Barch DM, Braver TS, Sabb FW, and Noll DC. 2000. Anterior cingulate and the monitoriing of response conflict: evidence from an fMRI study of overt verb generation. *J Cogn Neurosci* 12(2): 298-309.

Barch DM, Braver TS, Akbudak E, Conturo T, Ollinger J, and Snyder A. 2001. Anterior cingulate cortex and response conflict: effects of response modality and processing domain. *Cereb Cortex* 11(9): 837-848.

Bates JF and Goldman-Rakic PS. 1993. Prefrontal connections of medial motor areas in the rhesus monkey. *J Comp Neurol* 336(2): 211-228.

Bell AH, Everling S, and Munoz DP. 2000. Influence of stimulus eccentricity and direction on characteristics of pro- and antisaccades in non-human primates. *J Neurophysiol* 84(5): 2595-2604.

Bisley JW and Goldberg ME. 2003. Neuronal activity in the lateral intraparietal area and spatial attention. *Science* 299(5603): 81-86.

Bizzi E. 1968. Discharge of frontal eye field neurons during saccadic and following eye movements in unanesthetized monkeys. *Exp Brain Res* 6(1): 69-80.

Bizzi E and Schiller PH. 1970. Single unit activity in the frontal eye fields of unanesthetized monkeys during eye and head movement. *Exp Brain Res* 10(2): 150-158.

Blatt GJ, Andersen RA, and Stoner GR. 1990. Visual receptive field organization and cortico-cortical connections of the lateral intraparietal area (area LIP) in the macaque. *J Comp Neurol* 299(4): 421-445.

Botvinick M, Nystrom LE, Fissell K, Carter CS, and Cohen JD. 1999. Conflict monitoring versus selection-for-action in anterior cingulate cortex. *Nature* 402(6758): 179-181.

Botvinick MM, Cohen JD, and Carter CS. 2004. Conflict monitoring and anterior cingulate cortex: an update. *Trends Cogn Sci* 8(12): 539-546.

Boynton GM, Engel SA, Glover GH, and Heeger DJ. 1996. Linear systems analysis of functional magnetic resonance imaging in human V1. *J Neurosci* 16(13): 4207-4221.

Braver TS, Barch DM, Gray JR, Molfese DL, and Snyder A. 2001. Anterior cingulate cortex and response conflict: effects of frequency, inhibition and errors. *Cereb Cortex* 11(9): 825-836.

Brown JW and Braver TS. 2005. Learned predictions of error likelihood in the anterior cingulate cortex. *Science* 307(5712): 1118-1121.

Brown MRG, DeSouza JFX, Goltz HC, Ford K, Menon RS, Goodale MA, and Everling S. 2004. Comparison of memory- and visually guided saccades using event-related fMRI. *J Neurophysiol* 91(2): 873-889.

Bruce CJ and Goldberg ME. 1985. Primate frontal eye fields. I. Single neurons discharging before saccades. *J Neurophysiol* 53(3): 603-635.

Bruce CJ, Goldberg ME, Bushnell MC, and Stanton GB. 1985. Primate frontal eye fields. II. Physiological and anatomical correlates of electrically evoked eye movements. *J Neurophysiol* 54(3): 714-734.

Buckner RL, Bandettini PA, O'Craven KM, Savoy RL, Petersen SE, Raichle ME, and Rosen BR. 1996. Detection of cortical activation during averaged single trials of a cognitive task using functional magnetic resonance imaging. *Proc Natl Acad Sci U S A* 93(25): 14878-14883.

Buckner RL. 1998. Event-related fMRI and the hemodynamic response. *Hum Brain Mapp* 6(5-6): 373-377.

Bunge SA, Hazeltine E, Scanlon MD, Rosen AC, and Gabrieli JD. 2002. Dissociable contributions of prefrontal and parietal cortices to response selection. *NeuroImage* 17(3): 1562-1571.

Burock MA, Buckner RL, Woldorff MG, Rosen BR, and Dale AM. 1998. Randomized event-related experimental designs allow for extremely rapid presentation rates using functional MRI. *Neuroreport* 9(16): 3735-3739.

Buxton RB and Frank LR. 1997. A model for the coupling between cerebral blood flow and oxygen metabolism during neural stimulation. *J Cereb Blood Flow Metab* 17(1): 64-72.

Carpenter RH. 1994. Frontal cortex: choosing where to look. *Curr Biol* 4: 341-343.

Carpenter RH and Williams ML. 1995. Neural computation of log likelihood in control of saccadic eye movements. *Nature* 377: 59-62.

Carpenter RH. 2004. Supplementary eye field: keeping an eye on eye movement. *Curr Biol* 14(11): R416-418.

Carter CS, Braver TS, Barch DM, Botvinick MM, Noll D, and Cohen JD. 1998. Anterior cingulate cortex, error detection, and the online monitoring of performance. *Science* 280(5364): 747-749.

Carter CS, Macdonald AM, Botvinick M, Ross LL, Stenger VA, Noll D, and Cohen JD. 2000. Parsing executive processes: strategic vs. evaluative functions of the anterior cingulate cortex. *Proc Natl Acad Sci U S A* 97(4): 1944-1948.

Casey BJ, Thomas KM, Welsh TF, Badgaiyan RD, Eccard CH, Jennings JR, and Crone EA. 2000. Dissociation of response conflict, attentional selection, and expectancy with functional magnetic resonance imaging. *Proc Natl Acad Sci U S A* 97(15): 8728-8733.

Colby CL and Duhamel JR. 1991. Heterogeneity of extrastriate visual areas and multiple parietal areas in the macaque monkey. *Neuropsychologia* 29(6): 517-537.

Colby CL, Duhamel JR, and Goldberg ME. 1996. Visual, presaccadic, and cognitive activation of single neurons in monkey lateral intraparietal area. *J Neurophysiol* 76(5): 2841-2852.

Colby CL and Goldberg ME. 1999. Space and attention in parietal cortex. *Annu Rev Neurosci* 22: 319-349.

Connolly JD, Goodale MA, Menon RS, and Munoz DP. 2002. Human fMRI evidence for the neural correlates of preparatory set. *Nat Neurosci* 5(12): 1345-1352.

Connolly JD, Goodale MA, Goltz HC, and Munoz DP. 2005. fMRI activation in the human frontal eye field is correlated with saccadic reaction time. *J Neurophysiol* 94(1): 605-611.

Corbetta M, Akbudak E, Conturo TE, Snyder AZ, Ollinger JM, Drury HA, Linenweber MR, Petersen SE, Raichle ME, Van Essen DC, and Shulman GL. 1998. A common network of functional areas for attention and eye movements. *Neuron* 21(4): 761-773.

Corbetta M and Shulman GL. 2002. Control of goal-directed and stimulus-driven attention in the brain. *Nat Rev Neurosci* 3(3): 201-215.

Courtney SM, Petit L, Maisog JM, Ungerleider LG, and Haxby JV. 1998. An area specialized for spatial working memory in human frontal cortex. *Science* 279(5355): 1347-1351.

Curtis CE and D'Esposito M. 2003. Success and failure suppressing reflexive behavior. *J Cogn Neurosci* 15(3): 409-418.

Curtis CE, Rao VY, and D'Esposito M. 2004. Maintenance of spatial and motor codes during oculomotor delayed response tasks. *J Neurosci* 24(16): 3944-3952.

Curtis CE, Cole MW, Rao VY, and D'Esposito M. 2005. Canceling planned action: an FMRI study of countermanding saccades. *Cereb Cortex* 15(9): 1281-1289.

Dale AM and Buckner RL. 1997. Selective averaging of rapidly-presented individual trials using fMRI. *Hum Brain Mapp* 5: 329-340.

de Zubicaray GI, Andrew C, Zelaya FO, Williams SC, and Dumanoir C. 2000. Motor response suppression and the prepotent tendency to respond: a parametric fMRI study. *Neuropsychologia* 38(9): 1280-1291.

Desimone R and Duncan J. 1995. Neural mechanisms of selective visual attention. *Annu Rev Neurosci* 18: 193-222.

DeSouza JFX, Menon RS, and Everling S. 2003. Preparatory set associated with pro-saccades and anti-saccades in humans investigated with event-related FMRI. *J Neurophysiol* 89(2): 1016-1023.

Dias EC and Segraves MA. 1999. Muscimol-induced inactivation of monkey frontal eye field: effects on visually and memory-guided saccades. *J Neurophysiol* 81(5): 2191-2214.

Donaldson DI and Buckner RL. 2001. Effective paradigm design. In: *Functional magnetic resonance imaging: an introduction to methods*, edited by Jezzard P, Matthews PM, and Smith SM. Oxford: Oxford University Press, p. 177-196.

Dorris MC, Pare M, and Munoz DP. 1997. Neuronal activity in monkey superior colliculus related to the initiation of saccadic eye movements. *J Neurosci* 17(21): 8566-8579.

Durston S, Thomas KM, Worden MS, Yang Y, and Casey BJ. 2002. The effect of preceding context on inhibition: an event-related fMRI study. *NeuroImage* 16(2): 449-453.

Durston S, Davidson MC, Thomas KM, Worden MS, Tottenham N, Martinez A, Watts R, Ulug AM, and Casey BJ. 2003. Parametric manipulation of conflict and response competition using rapid mixed-trial event-related fMRI. *NeuroImage* 20(4): 2135-2141.

Edelman JA and Keller EL. 1996. Activity of visuomotor burst neurons in the superior colliculus accompanying express saccades. *J Neurophysiol* 76(2): 908-926.

Erdmann B and Dodge R. 1898. *Psychologische Untersuchungen uber das Lesen auf experimenteller Grundlage*.

Everling S, Dorris MC, and Munoz DP. 1998. Reflex suppression in the anti-saccade task is dependent on prestimulus neural processes. *J Neurophysiol* 80(3): 1584-1589.

Everling S and Fischer B. 1998. The antisaccade: a review of basic research and clinical studies. *Neuropsychologia* 36(9): 885-899.

Everling S, Dorris MC, Klein RM, and Munoz DP. 1999. Role of primate superior colliculus in preparation and execution of anti-saccades and pro-saccades. *J Neurosci* 19(7): 2740-2754.

Everling S and Munoz DP. 2000. Neuronal correlates for preparatory set associated with pro-saccades and anti-saccades in the primate frontal eye field. *J Neurosci* 20(1): 387-400.

Everling S and DeSouza JF. 2005. Rule-dependent Activity for Prosaccades and Antisaccades in the Primate Prefrontal Cortex. *J Cogn Neurosci* 17(9): 1483-1496.

Ferraina S, Pare M, and Wurtz RH. 2002. Comparison of cortico-cortical and cortico-collicular signals for the generation of saccadic eye movements. *J Neurophysiol* 87(2): 845-858.

Finlay BL, Schiller PH, and Volman SF. 1976. Quantitative studies of single-cell properties in monkey striate cortex. IV. Corticotectal cells. *J Neurophysiol* 39(6): 1352-1361.

Fischer B and Boch R. 1983. Saccadic eye movements after extremely short reaction times in the monkey. *Brain Res* 260(1): 21-26.

Fischer B and Ramsperger E. 1984. Human express saccades: extremely short reaction times of goal directed eye movements. *Exp Brain Res* 57(1): 191-195.

Fischer B and Weber H. 1992. Characteristics of "anti" saccades in man. *Exp Brain Res* 89(2): 415-424.

Fischer B and Weber H. 1993. Express saccades and visual attention. *Behavioral Brain Science* 16: 533-610.

Fischer B and Weber H. 1997. Effects of stimulus conditions on the performance of antisaccades in man. *Exp Brain Res* 116(2): 191-200.

Ford KA, Goltz HC, Brown MRG, and Everling S. 2005. Neural processes associated with antisaccade task performance investigated with event-related FMRI. *J Neurophysiol* 94(1): 429-440.

Friston KJ, Josephs O, Rees G, and Turner R. 1998. Nonlinear event-related responses in fMRI. *Magn Reson Med* 39(1): 41-52.

Fuchs AF, Kaneko CR, and Scudder CA. 1985. Brainstem control of saccadic eye movements. *Annu Rev Neurosci* 8: 307-337.

Gaymard B, Pierrot-Deseilligny C, and Rivaud S. 1990. Impairment of sequences of memory-guided saccades after supplementary motor area lesions. *Ann Neurol* 28(5): 622-626.

Gaymard B, Rivaud S, Cassarini JF, Dubard T, Rancurel G, Agid Y, and Pierrot-Deseilligny C. 1998. Effects of anterior cingulate cortex lesions on ocular saccades in humans. *Exp Brain Res* 120(2): 173-183.

Gaymard B, Ploner CJ, Rivaud-Pechoux S, and Pierrot-Deseilligny C. 1999. The frontal eye field is involved in spatial short-term memory but not in reflexive saccade inhibition. *Exp Brain Res* 129(2): 288-301.

Georgopoulos AP. 1994. New concepts in generation of movement. *Neuron* 13(2): 257-268.

Goldberg ME, Bisley J, Powell KD, Gottlieb J, and Kusunoki M. 2002. The role of the lateral intraparietal area of the monkey in the generation of saccades and visuospatial attention. *Ann N Y Acad Sci* 956: 205-215.

Goldberg ME, Bisley JW, Powell KD, and Gottlieb J. 2006. Saccades, salience and attention: the role of the lateral intraparietal area in visual behavior. *Prog Brain Res* 155: 157-175.

Goldman PS and Nauta WJ. 1976. Autoradiographic demonstration of a projection from prefrontal association cortex to the superior colliculus in the rhesus monkey. *Brain Res* 116(1): 145-149.

Gottlieb J and Goldberg ME. 1999. Activity of neurons in the lateral intraparietal area of the monkey during an antisaccade task. *Nat Neurosci* 2(10): 906-912.

Grantyn A and Grantyn R. 1982. Axonal patterns and sites of termination of cat superior colliculus neurons projecting in the tecto-bulbo-spinal tract. *Exp Brain Res* 46(2): 243-256.

Graybiel AM. 1995. Building action repertoires: memory and learning functions of the basal ganglia. *Curr Opin Neurobiol* 5(6): 733-741.

Grosbras MH, Lobel E, Van de Moortele PF, LeBihan D, and Berthoz A. 1999. An anatomical landmark for the supplementary eye fields in human revealed with functional magnetic resonance imaging. *Cereb Cortex* 9(7): 705-711.

Guitton D, Buchtel HA, and Douglas RM. 1985. Frontal lobe lesions in man cause difficulties in suppressing reflexive glances and in generating goal-directed saccades. *Exp Brain Res* 58(3): 455-472.

Hallett PE. 1978. Primary and secondary saccades to goals defined by instructions. *Vision Res* 18(10): 1279-1296.

Hanes DP and Schall JD. 1996. Neural control of voluntary movement initiation. *Science* 274: 427-430.

Hanes DP and Wurtz RH. 2001. Interaction of the frontal eye field and superior colliculus for saccade generation. *J Neurophysiol* 85(2): 804-815.

Haykin S. 1999. *Neural Networks A Comprehensive Foundation*. Upper Saddle River, New Jersey: Prentice Hall.

Hazeltine E, Poldrack R, and Gabrieli JD. 2000. Neural activation during response competition. *J Cogn Neurosci* 12 Suppl 2: 118-129.

Helminski JO and Segraves MA. 2003. Macaque frontal eye field input to saccade-related neurons in the superior colliculus. *J Neurophysiol* 90(2): 1046-1062.

Hess WR, Burgi S, and Bucher V. 1946. Motor function of tectal and tegmental area. *Monatsschrift fuer Psychiatrie und Neurologie* 112: 1-52.

Hikosaka O, Takikawa Y, and Kawagoe R. 2000. Role of the basal ganglia in the control of purposive saccadic eye movements. *Physiol Rev* 80(3): 953-978.

Huerta MF, Krubitzer LA, and Kaas JH. 1987. Frontal eye field as defined by intracortical microstimulation in squirrel monkeys, owl monkeys, and macaque monkeys. II. Cortical connections. *J Comp Neurol* 265(3): 332-361.

Huerta MF and Kaas JH. 1990. Supplementary eye field as defined by intracortical microstimulation: connections in macaques. *J Comp Neurol* 293(2): 299-330.

Huettel SA and McCarthy G. 2000. Evidence for a refractory period in the hemodynamic response to visual stimuli as measured by MRI. *NeuroImage* 11(5 Pt 1): 547-553.

Huettel SA, Song AW, and McCarthy G. 2004. *Functional Magnetic Resonance Imaging*. Sunderland, MA: Sinauer Associates.

Isoda M and Tanji J. 2002. Cellular activity in the supplementary eye field during sequential performance of multiple saccades. *J Neurophysiol* 88(6): 3541-3545.

Ito S, Stuphorn V, Brown JW, and Schall JD. 2003. Performance monitoring by the anterior cingulate cortex during saccade countermanding. *Science* 302(5642): 120-122.

Johnston K and Everling S. 2006a. Neural activity in monkey prefrontal cortex is modulated by task context and behavioral instruction during delayed-match-to-sample and conditional prosaccade-antisaccade tasks. *J Cogn Neurosci* 18(5): 749-765.

Johnston K and Everling S. 2006b. Monkey dorsolateral prefrontal cortex sends task-selective signals directly to the superior colliculus. *J Neurosci* 26(48): 12471-12478.

Johnston K, Levin HM, Koval MJ, and Everling S. 2007. Top-Down Control-Signal Dynamics in Anterior Cingulate and Prefrontal Cortex Neurons following Task Switching. *Neuron* 53(3): 453-462.

Josephs O, Turner R, and Friston KJ. 1997. Event-related fMRI. *Hum Brain Mapp* 5: 243-248.

Keating EG, Gooley SG, Pratt SE, and Kelsey JE. 1983. Removing the superior colliculus silences eye movements normally evoked from stimulation of the parietal and occipital eye fields. *Brain Res* 269(1): 145-148.

Keating EG and Gooley SG. 1988a. Disconnection of parietal and occipital access to the saccadic oculomotor system. *Exp Brain Res* 70(2): 385-398.

Keating EG and Gooley SG. 1988b. Saccadic disorders caused by cooling the superior colliculus or the frontal eye field, or from combined lesions of both structures. *Brain Res* 438(1-2): 247-255.

Kiehl KA, Liddle PF, and Hopfinger JB. 2000. Error processing and the rostral anterior cingulate: an event-related fMRI study. *Psychophysiology* 37(2): 216-223.

Kwong KK, Belliveau JW, Chesler DA, Goldberg IE, Weisskoff RM, Poncelet BP, Kennedy DN, Hoppel BE, Cohen MS, Turner R, and et al. 1992. Dynamic magnetic resonance imaging of human brain activity during primary sensory stimulation. *Proc Natl Acad Sci U S A* 89(12): 5675-5679.

Leichnetz GR. 1981. The prefrontal cortico-oculomotor trajectories in the monkey. *J Neurol Sci* 49(3): 387-396.

Leichnetz GR, Spencer RF, Hardy SG, and Astruc J. 1981. The prefrontal corticotectal projection in the monkey; an anterograde and retrograde horseradish peroxidase study. *Neuroscience* 6(6): 1023-1041.

Li CS, Mazzoni P, and Andersen RA. 1999. Effect of reversible inactivation of macaque lateral intraparietal area on visual and memory saccades. *J Neurophysiol* 81(4): 1827-1838.

Logothetis NK and Wandell BA. 2004. Interpreting the BOLD signal. *Annu Rev Psychol* 66: 735-769.

Luria AR. 1966. *Higher Cortical Functions in Man*. London: Tavistock.

Lynch JC, Graybiel AM, and Lobeck LJ. 1985. The differential projection of two cytoarchitectonic subregions of the inferior parietal lobule of macaque upon the deep layers of the superior colliculus. *J Comp Neurol* 235(2): 241-254.

Lynch JC, Hoover JE, and Strick PL. 1994. Input to the primate frontal eye field from the substantia nigra, superior colliculus, and dentate nucleus demonstrated by transneuronal transport. *Exp Brain Res* 100: 181-186.

MacDonald AW, 3rd, Cohen JD, Stenger VA, and Carter CS. 2000. Dissociating the role of the dorsolateral prefrontal and anterior cingulate cortex in cognitive control. *Science* 288(5472): 1835-1838.

MacLeod CM and MacDonald PA. 2000. Interdimensional interference in the Stroop effect: uncovering the cognitive and neural anatomy of attention. *Trends Cogn Sci* 4(10): 383-391.

May JG and Andersen RA. 1986. Different patterns of corticopontine projections from separate cortical fields within the inferior parietal lobule and dorsal prelunate gyrus of the macaque. *Exp Brain Res* 63(2): 265-278.

McCarthy G, Luby M, Gore J, and Goldman-Rakic P. 1997. Infrequent events transiently activate human prefrontal and parietal cortex as measured by functional MRI. *J Neurophysiol* 77(3): 1630-1634.

Menon V, Adleman NE, White CD, Glover GH, and Reiss AL. 2001. Error-related brain activation during a Go/NoGo response inhibition task. *Hum Brain Mapp* 12(3): 131-143.

Miller EK and Cohen JD. 2001. An integrative theory of prefrontal cortex function. *Annu Rev Neurosci* 24: 167-202.

Monsell S. 2003. Task switching. *Trends Cogn Sci* 7(3): 134-140.

Moschovakis AK and Karabelas AB. 1985. Observations on the somatodendritic morphology and axonal trajectory of intracellularly HRP-labeled efferent neurons located in the deeper layers of the superior colliculus of the cat. *J Comp Neurol* 239(3): 276-308.

Moschovakis AK. 1997. The neural integrators of the mammalian saccadic system. *Front Biosci* 2: 552-577.

Munoz DP, Pelisson D, and Guitton D. 1991. Movement of neural activity on the superior colliculus motor map during gaze shifts. *Science* 251: 1358-1360.

Munoz DP and Wurtz RH. 1993a. Fixation cells in monkey superior colliculus. I. Characteristics of cell discharge. *J Neurophysiol* 70(2): 559-575.

Munoz DP and Wurtz RH. 1993b. Fixation cells in monkey superior colliculus. II. Reversible activation and deactivation. *J Neurophysiol* 70(2): 576-589.

Munoz DP and Wurtz RH. 1995. Saccade-related activity in monkey superior colliculus. I. Characteristics of burst and buildup cells. *J Neurophysiol* 73(6): 2313-2333.

Munoz DP, Dorris MC, Pare M, and Everling S. 2000. On your mark, get set: brainstem circuitry underlying saccadic initiation. *Can J Physiol Pharmacol* 78(11): 934-944.

Munoz DP and Everling S. 2004. Look away: the anti-saccade task and the voluntary control of eye movement. *Nat Rev Neurosci* 5(3): 218-228.

Muri RM, Iba-Zizen MT, Derosier C, Cabanis EA, and Pierrot-Deseilligny C. 1996. Location of the human posterior eye field with functional magnetic resonance imaging. *J Neurol Neurosurg Psychiatry* 60(4): 445-448.

Nakamura K, Roesch MR, and Olson CR. 2005. Neuronal activity in macaque SEF and ACC during performance of tasks involving conflict. *J Neurophysiol* 93(2): 884-908.

Ogawa S, Tank DW, Menon R, Ellermann JM, Kim SG, Merkle H, and Ugurbil K. 1992. Intrinsic signal changes accompanying sensory stimulation: functional brain mapping with magnetic resonance imaging. *Proc Natl Acad Sci U S A* 89(13): 5951-5955.

Ollinger JM, Corbetta M, and Shulman GL. 2001a. Separating processes within a trial in event-related functional MRI, II. Analysis. *NeuroImage* 13(1): 218-229.

Ollinger JM, Shulman GL, and Corbetta M. 2001b. Separating processes within a trial in event-related functional MRI, I. The Method. *NeuroImage* 13(1): 210-217.

Olson CR and Gettner SN. 1995. Object-centered direction selectivity in the macaque supplementary eye field. *Science* 269(5226): 985-988.

Palmer ED, Rosen HJ, Ojemann JG, Buckner RL, Kelley WM, and Petersen SE. 2001. An event-related fMRI study of overt and covert word stem completion. *NeuroImage* 14(1 Pt 1): 182-193.

Paus T. 1996. Location and function of the human frontal eye-field: a selective review. *Neuropsychologia* 34(6): 475-483.

Paus T, Castro-Alamancos MA, and Petrides M. 2001. Cortico-cortical connectivity of the human mid-dorsolateral frontal cortex and its modulation by repetitive transcranial magnetic stimulation. *J Neurosci* 14(8): 1405-1411.

Pierrot-Deseilligny C, Rivaud S, Gaymard B, and Agid Y. 1991a. Cortical control of memory-guided saccades in man. *Exp Brain Res* 83(3): 607-617.

Pierrot-Deseilligny C, Rivaud S, Gaymard B, and Agid Y. 1991b. Cortical control of reflexive visually-guided saccades. *Brain* 114 (Pt 3): 1473-1485.

Pierrot-Deseilligny C, Muri RM, Ploner CJ, Gaymard B, and Rivaud-Pechoux S. 2003. Cortical control of ocular saccades in humans: a model for motricity. *Prog Brain Res* 142: 3-17.

Pierrot-Deseilligny C, Muri RM, Nyffeler T, and Milea D. 2005. The role of the human dorsolateral prefrontal cortex in ocular motor behavior. *Ann N Y Acad Sci* 1039: 239-251.

Platt ML and Glimcher PW. 1998. Response fields of intraparietal neurons quantified with multiple saccadic targets. *Exp Brain Res* 121(1): 65-75.

Ploner CJ, Gaymard BM, Rivaud-Pechoux S, and Pierrot-Deseilligny C. 2005. The prefrontal substrate of reflexive saccade inhibition in humans. *Biol Psychiatry* 57(10): 1159-1165.

Posner MI and DiGirolamo GJ. 1998. Executive attention: conflict, target detection, and cognitive control. In: *The Attentive Brain*, edited by Parasuraman RMIT Press, p. 401-423.

Powell KD and Goldberg ME. 2000. Response of neurons in the lateral intraparietal area to a distractor flashed during the delay period of a memory-guided saccade. *J Neurophysiol* 84(1): 301-310.

Ratcliff R, Van Zandt T, and McKoon G. 1999. Connectionist and diffusion models of reaction time. *Psychol Rev* 106: 261-300.

Rizzolatti G, Riggio L, Dascola I, and Umilta C. 1987. Reorienting attention across the horizontal and vertical meridians: evidence in favor of a premotor theory of attention. *Neuropsychologia* 25(1A): 31-40.

Robinson DA. 1964. The Mechanics of Human Saccadic Eye Movement. *J Physiol* 174: 245-264.

Robinson DA and Fuchs AF. 1969. Eye movements evoked by stimulation of frontal eye fields. *J Neurophysiol* 32(5): 637-648.

Robinson DA. 1972. On the nature of visual-oculomotor connections. *Invest Ophthalmol* 11(6): 497-503.

Robinson DA and Keller EL. 1972. The behavior of eye movement motoneurons in the alert monkey. *Bibl Ophthalmol* 82: 7-16.

Robinson DA. 1981. The use of control systems analysis in the neurophysiology of eye movements. *Annu Rev Neurosci* 4: 463-503.

Robinson DA and McClurkin JW. 1989a. The visual superior colliculus and pulvinar. In: *The Neurobiology of Saccadic Eye Movements*, edited by Wurtz RH and Goldberg ME. Amsterdam: Elsevier.

Robinson DL and McClurkin JW. 1989b. The visual superior colliculus and pulvinar. *Rev Oculomot Res* 3: 337-360.

Robinson FR and Fuchs AF. 2001. The role of the cerebellum in voluntary eye movements. *Annu Rev Neurosci* 24: 981-1004.

Rosen BR, Buckner RL, and Dale AM. 1998. Event-related functional MRI: past, present, and future. *Proc Natl Acad Sci U S A* 95(3): 773-780.

Saslow MG. 1967. Effects of components of displacement-step stimuli upon latency for saccadic eye movement. *J Opt Soc Am* 57(8): 1024-1029.

Sato TR and Schall JD. 2003. Effects of stimulus-response compatibility on neural selection in frontal eye field. *Neuron* 38(4): 637-648.

Schall JD. 1991. Neuronal activity related to visually guided saccadic eye movements in the supplementary motor area of rhesus monkeys. *J Neurophysiol* 66(2): 530-558.

Schall JD, Morel A, and Kaas JH. 1993. Topography of supplementary eye field afferents to frontal eye field in macaque: implications for mapping between saccade coordinate systems. *Vis Neurosci* 10(2): 385-393.

Schall JD, Morel A, King DJ, and Bullier J. 1995. Topography of visual cortex connections with frontal eye field in macaque: convergence and segregation of processing streams. *J Neurosci* 15(6): 4464-4487.

Schall JD. 1997. Visuomotor areas of the frontal lobe. In: *Cerebral Cortex*, edited by Rockland K, Peters A, and Kaas JH. New York: Plenum, p. 527-638.

Schall JD. 2002. The neural selection and control of saccades by the frontal eye field. *Philos Trans R Soc Lond B Biol Sci* 357(1424): 1073-1082.

Schall JD. 2004. On the role of frontal eye field in guiding attention and saccades. *Vision Res* 44(12): 1453-1467.

Schiller PH and Malpeli JG. 1977. Properties and tectal projections of monkey retinal ganglion cells. *J Neurophysiol* 40(2): 428-445.

Schiller PH, True SD, and Conway JL. 1979. Effects of frontal eye field and superior colliculus ablations on eye movements. *Science* 206(4418): 590-592.

Schiller PH, True SD, and Conway JL. 1980. Deficits in eye movements following frontal eye-field and superior colliculus ablations. *J Neurophysiol* 44(6): 1175-1189.

Schiller PH, Sandell JH, and Maunsell JH. 1987. The effect of frontal eye field and superior colliculus lesions on saccadic latencies in the rhesus monkey. *J Neurophysiol* 57(4): 1033-1049.

Schiller PH. 1998. The neural control of visually guided eye movements. In: *Cognitive Neuroscience of Attention*, edited by Richards J. London: Lawrence Erlbaum.

Schiller PH and Tehovnik EJ. 2001. Look and see: how the brain moves your eyes about. *Prog Brain Res* 134: 127-142.

Schiller PH and Tehovnik EJ. 2005. Neural mechanisms underlying target selection with saccadic eye movements. *Prog Brain Res* 149: 157-171.

Schlag J and Schlag-Rey M. 1987. Evidence for a supplementary eye field. *J Neurophysiol* 57(1): 179-200.

Schlag-Rey M, Amador N, Sanchez H, and Schlag J. 1997. Antisaccade performance predicted by neuronal activity in the supplementary eye field. *Nature* 390(6658): 398-401.

Schluppeck D, Curtis CE, Glimcher PW, and Heeger DJ. 2006. Sustained activity in topographic areas of human posterior parietal cortex during memory-guided saccades. *J Neurosci* 26(19): 5098-5108.

Scudder CA, Kaneko CS, and Fuchs AF. 2002. The brainstem burst generator for saccadic eye movements: a modern synthesis. *Exp Brain Res* 142(4): 439-462.

Segraves MA and Goldberg ME. 1987. Functional properties of corticotectal neurons in the monkey's frontal eye field. *J Neurophysiol* 58(6): 1387-1419.

Segraves MA. 1992. Activity of monkey frontal eye field neurons projecting to oculomotor regions of the pons. *J Neurophysiol* 68(6): 1967-1985.

Serences JT. 2004. A comparison of methods for characterizing the event-related BOLD timeseries in rapid fMRI. *NeuroImage* 21(4): 1690-1700.

Sereno MI, Pitzalis S, and Martinez A. 2001. Mapping of contralateral space in retinotopic coordinates by a parietal cortical area in humans. *Science* 294(5545): 1350-1354.

Snyder LH, Batista AP, and Andersen RA. 1997. Coding of intention in the posterior parietal cortex. *Nature* 386(6621): 167-170.

Snyder LH, Batista AP, and Andersen RA. 1998. Change in motor plan, without a change in the spatial locus of attention, modulates activity in posterior parietal cortex. *J Neurophysiol* 79(5): 2814-2819.

Snyder LH, Batista AP, and Andersen RA. 2000. Intention-related activity in the posterior parietal cortex: a review. *Vision Res* 40(10-12): 1433-1441.

Sommer MA and Tehovnik EJ. 1997. Reversible inactivation of macaque frontal eye field. *Exp Brain Res* 116(2): 229-249.

Sommer MA and Wurtz RH. 1998. Frontal eye field neurons orthodromically activated from the superior colliculus. *J Neurophysiol* 80(6): 3331-3335.

Sommer MA and Wurtz RH. 2000. Composition and topographic organization of signals sent from the frontal eye field to the superior colliculus. *J Neurophysiol* 83(4): 1979-2001.

Sommer MA and Wurtz RH. 2001. Frontal eye field sends delay activity related to movement, memory, and vision to the superior colliculus. *J Neurophysiol* 85(4): 1673-1685.

Sparks DL. 1986. Translation of sensory signals into commands for control of saccadic eye movements: role of primate superior colliculus. *Physiol Rev* 66(1): 118-171.

Sparks DL and Mays LE. 1990. Signal transformations required for the generation of saccadic eye movements. *Annu Rev Neurosci* 13: 309-336.

Stanton GB, Bruce CJ, and Goldberg ME. 1993. Topography of projections to the frontal lobe from the macaque frontal eye fields. *J Comp Neurol* 330(2): 286-301.

Stuphorn V, Taylor TL, and Schall JD. 2000. Performance monitoring by the supplementary eye field. *Nature* 408(6814): 857-860.

Sweeney JA, Mintun MA, Kwee S, Wiseman MB, Brown DL, Rosenberg DR, and Carl JR. 1996. Positron emission tomography study of voluntary saccadic eye movements and spatial working memory. *J Neurophysiol* 75(1): 454-468.

Tehovnik EJ and Lee K. 1993. The dorsomedial frontal cortex of the rhesus monkey: topographic representation of saccades evoked by electrical stimulation. *Exp Brain Res* 96(3): 430-442.

Tehovnik EJ, Lee K, and Schiller PH. 1994. Stimulation-evoked saccades from the dorsomedial frontal cortex of the rhesus monkey following lesions of the frontal eye fields and superior colliculus. *Exp Brain Res* 98(2): 179-190.

Thier P and Andersen RA. 1996. Electrical microstimulation suggests two different forms of representation of head-centered space in the intraparietal sulcus of rhesus monkeys. *Proc Natl Acad Sci U S A* 93(10): 4962-4967.

Thier P and Andersen RA. 1998. Electrical microstimulation distinguishes distinct saccade-related areas in the posterior parietal cortex. *J Neurophysiol* 80(4): 1713-1735.

Thompson KG, Bichot NP, and Schall JD. 1997. Dissociation of visual discrimination from saccade programming in macaque frontal eye field. *J Neurophysiol* 77(2): 1046-1050.

Thompson KG, Bichot NP, and Sato TR. 2005. Frontal eye field activity before visual search errors reveals the integration of bottom-up and top-down salience. *J Neurophysiol* 93(1): 337-351.

Thompson-Schill SL, D'Esposito M, Aguirre GK, and Farah MJ. 1997. Role of left inferior prefrontal cortex in retrieval of semantic knowledge: a reevaluation. *Proc Natl Acad Sci U S A* 94(26): 14792-14797.

van Veen V, Cohen JD, Botvinick MM, Stenger VA, and Carter CS. 2001. Anterior cingulate cortex, conflict monitoring, and levels of processing. *NeuroImage* 14(6): 1302-1308.

Vazquez AL and Noll DC. 1998. Nonlinear aspects of the BOLD response in functional MRI. *NeuroImage* 7(2): 108-118.

Wang Y, Matsuzaka Y, Shima K, and Tanji J. 2004. Cingulate cortical cells projecting to monkey frontal eye field and primary motor cortex. *Neuroreport* 15(10): 1559-1563.

Wurtz RH and Albano JE. 1980. Visual-motor function of the primate superior colliculus. *Annu Rev Neurosci* 3: 189-226.

Wurtz RH, Sommer MA, Pare M, and Ferraina S. 2001. Signal transformations from cerebral cortex to superior colliculus for the generation of saccades. *Vision Res* 41(25-26): 3399-3412.

Zarahn E, Aguirre G, and D'Esposito M. 1997a. A trial-based experimental design for fMRI. *NeuroImage* 6(2): 122-138.

Zarahn E, Aguirre GK, and D'Esposito M. 1997b. Empirical analyses of BOLD fMRI statistics. I. Spatially unsmoothed data collected under null-hypothesis conditions. *NeuroImage* 5(3): 179-197.

Zhang M and Barash S. 2000. Neuronal switching of sensorimotor transformations for antisaccades. *Nature* 408(6815): 971-975.

Zhang M and Barash S. 2004. Persistent LIP activity in memory-antisaccades: working memory for a sensorimotor transformation. *J Neurophysiol* 91(3): 1424-1441.

Chapter 2
Inhibition and Generation of Saccades: Rapid Event-related fMRI of Prosaccades, Antisaccades, and Nogo Trials

The material in *Chapter 2* was published previously as MR Brown, HC Goltz, T Vilis, KA Ford, and S Everling, 2006. Inhibition and generation of saccades: Rapid event-related fMRI of prosaccades, antisaccades, and nogo trials. *NeuroImage* 33(2): 644-659. The material therein is used here with permission (see *Appendix 2*).

2.1 – Introduction

Inhibition of automatic responses and generation of voluntary responses are complementary components of flexible, adaptable behaviour given that unusual circumstances frequently require inhibition of inappropriate automatic responses followed by generation of a suitable voluntary response. The importance of inhibitory and generative processes is illustrated by the behavioural disturbances characteristic of schizophrenia and frontal lobe syndrome. These conditions produce deficits in behavioural inhibition, manifesting as distractibility and impulsive, disorganized behaviour, as well as deficits in the generation of voluntary activity, resulting in flattened cognitive and emotional affect and lack of initiative (Luria 1966; Fuster 1989; Frith 1992).

Behavioural inhibition and voluntary response generation have been studied using the antisaccade task, which requires subjects to inhibit the automatic saccadic eye movement toward an abruptly presented peripheral stimulus and to generate a voluntary saccade to the blank visual field location diametrically opposite the stimulus. Previous human fMRI studies (Connolly et al. 2002; Curtis and D'Esposito 2003; DeSouza et al. 2003; Ford et al. 2005) compared the antisaccade task with the prosaccade task, also called the visually guided saccade task. These studies used compound trial designs consisting of an instruction period, during which subjects viewed a coloured fixation point indicating whether the current trial was a prosaccade or antisaccade trial, followed by a response period, during which subjects were presented with a peripheral stimulus and had to respond appropriately. These studies found greater fMRI activation for antisaccades than prosaccades during the instruction period in frontal eye field,

supplementary eye field, intraparietal sulcus, middle frontal gyrus (MFG), and pre-supplementary motor area (pre-SMA). They also found activation differences during the response period, though some attributed this to carryover from the instruction period. Two studies found greater activation for correct antisaccades compared to error antisaccades, in which subjects failed to inhibit the automatic prosaccade, during the late instruction period in right MFG, anterior cingulate cortex, and pre-SMA (Curtis and D'Esposito 2003; Ford et al. 2005).

These results can be explained by the necessity to inhibit automatic saccades and generate voluntary saccades to endogenously defined targets in the antisaccade but not the prosaccade task. Previous studies did not examine the relative contributions of saccadic inhibition and saccade generation to activation differences between prosaccades and antisaccades. To ascertain whether just one of these processes or some combination of the two is responsible, we used rapid event-related fMRI with human subjects to compare prosaccades, antisaccades, and nogo trials. On nogo trials, subjects had to inhibit the automatic saccade toward a sudden peripheral stimulus while maintaining central fixation. Thus, the antisaccade task recruited saccadic inhibition, visuospatial remapping, and voluntary saccade generation, while the prosaccade and nogo tasks recruited only saccade generation and saccadic inhibition, respectively. Cortical saccade regions specialized for saccadic inhibition should exhibit high activation levels for the nogo and antisaccade tasks and low activation for prosaccades, while regions specialized for voluntary saccade generation should exhibit high activation levels for the prosaccade and antisaccade tasks and low activation for nogo trials. Unlike previous antisaccade studies, we used a rapid event-related design because it avoided the need for long fixation intervals lasting 10 or more seconds between trials. Such long fixation intervals are tedious and have the undesireable potential to alter the neuronal processes under investigation by promoting subject fatigue and blunting subject alertness and, most importantly, attention to the paradigm.

2.2 - Methods

All procedures were approved by the University Research Ethics Board for Health Sciences Research at the University of Western Ontario, London, Ontario, Canada and are in accordance with the 1964 Declaration of Helsinki.

2.2.1 – Subjects

Ten human subjects (6 female) participated in this study after giving informed written consent. Subject age ranged from 22 to 33 years with a mean of 26 years. Subjects reported no history of neurological or psychiatric disorder, and they had normal or corrected to normal vision. One subject described herself as left-handed while the rest described themselves as right-handed.

2.2.2 – fMRI Data Acquisition Procedure

All fMRI data were acquired on a whole body 4 Tesla MRI system (Varian, Palo Alto, CA; Siemens, Erlangen, Germany) operating at a slew rate of 120 T/m/s with 40 mT/m gradients. A transmit only receive only (TORO) cylindrical hybrid birdcage radio frequency (RF) head coil (Barberi et al. 2000) was used for transmission and detection of the signal.

Imaging planes for the functional scans were prescribed from a series of sagittal anatomical images acquired using T1-weighting. Eleven contiguous functional planes, 5 mm thick, were prescribed axially from the top of the brain to the level of the dorsal caudate and thalamus. A constrained, three-dimensional phase shimming procedure (Klassen and Menon 2004) was employed to optimize the magnetic field homogeneity over the prescribed functional volume. During each functional task, blood oxygenation level-dependent (BOLD) images (T_2*-weighted) were acquired continuously using an interleaved, two segment, optimized spiral imaging protocol (volume collection time = 1.0 s, repeat time (TR) = 500 ms, echo time (TE) = 15 ms, flip angle = 30º, 64 x 64 matrix size, 22.0 x 22.0 cm field of view (FOV), 3.44 x 3.44 x 5.00 mm voxel resolution). Each image was corrected for physiologic fluctuations using navigator echo correction. A corresponding high-resolution T1-weighted structural volume was acquired during the same scanning session using a 3D spiral acquisition protocol (TE = 3.0 ms,

inversion time (TI) = 1300 ms, TR = 50 ms) with a voxel resolution of 0.9 x 0.9 x 1.25 mm. Each subject was immobilized during the experimental session within a head cradle, which was packed with foam padding.

2.2.3 – Eye Tracking

Visual stimuli were presented during fMRI scanning using a Silent Vision SV-4021 projection system from Avotec (Stuart, Fl, USA). This system includes an MEyeTrack-SV (Silent Vision) eye tracker from SensoMotoric Instruments GmbH (Teltow, Germany). This setup employs fibre optics housed in dual stalks positioned over the subject's eyes to allow simultaneous presentation of visual stimuli and CCD video-based infrared eye tracking. The visual display subtends 30° horizontally by 23° vertically with a resolution of 800 x 600 pixels and a refresh rate of 60 Hz. Eye tracking was also performed at a 60 Hz sampling rate with an accuracy of approximately 1 degree. Before scanning, we calibrated the system with a 5 point calibration, with one point in each corner of the visual display as well as a central point. All experimental targets were within the range of the calibration. Analysis of the eye movement signals was performed offline.

2.2.4 – Experimental Design

This experiment employed three tasks, a prosaccade task, an antisaccade task, and a nogo task. Trials were compound, consisting of a trial type instruction period and a response period (*Figure 2.1*), and they were presented randomly interleaved in a jittered rapid event-related design. When not performing a task response, the subject fixated a central white dot (0.6° diameter). An individual trial started when the central dot changed from white to red, green, or blue, where the colour indicated the trial type. There exist 6 possible colour-trial assignment patterns using this scheme. Colour-trial assignment patterns were distributed among the ten subjects as follows. We used 4 of the colour-trial assignments twice, that is each was used for two different subjects, and we used the remaining 2 colour-trial assignments once, each for a single subject. The instruction period lasted from 2 to 7 s (mean 3 s, exponential distribution favouring 2 s) at the end of
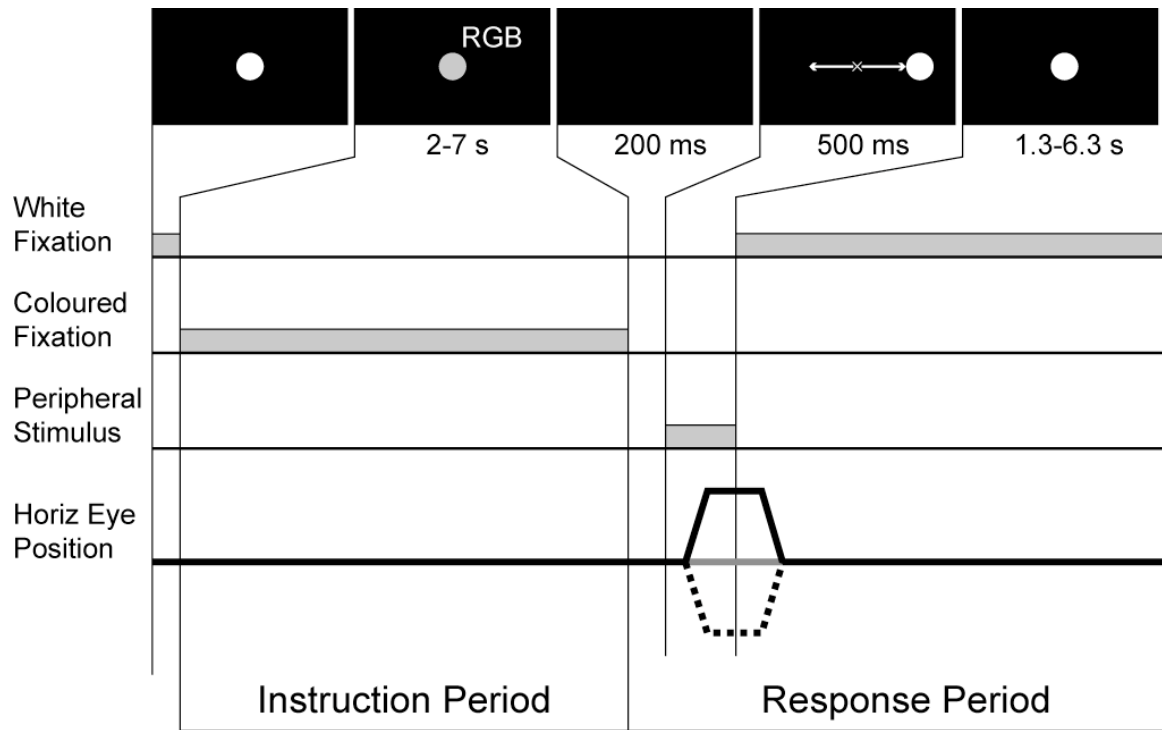
Figure 2.1 – Experimental Design

Timing and visual stimulus presentation for prosaccade, antisaccade, and nogo tasks are illustrated. Topmost images depict sequential visual arrays. The gray fixation point labeled RGB in the second stimulus panel denotes the use of a coloured fixation point to cue task type. The fourth panel indicates a rightward prosaccade, a leftward antisaccade, and a nogo response with a rightward arrow, leftward arrow, and central x, respectively. Gray horizontal bars in lower portion of figure illustrate timings and durations of various visual stimulus components. Horizontal eye position trace at bottom depicts a rightward prosaccade (black solid line), a leftward antisaccade (black dotted line), and a nogo response (gray solid line). See *Section 2.2 – Methods* for details.

which the central dot disappeared. Following a 200 ms gap, during which the subject was presented with a black screen, a peripheral white dot (0.6° diameter) appeared at an eccentricity of 4°, 5°, 6°, or 7° located along the horizontal meridian to the left or right of fixation. The peripheral stimulus was presented for 500 ms, after which it disappeared and the central white fixation dot reappeared simultaneously. Before the peripheral

stimulus disappeared, the subject had to (1) look at it on prosaccade trials, (2) look away from it, toward its mirror location in the opposite visual hemifield on antisaccade trials, or (3) simply maintain central fixation without generating a saccade on nogo trials. Following the task response event, the subject had to return gaze to centre (on prosaccade and antisaccade trials) and maintain central fixation on the newly reappeared central fixation dot. The response period, defined as the interval from central fixation offset at the end of the preceding instruction period until the changing of the fixation cross's colour from white to red, green, or blue marking the start of the next trial's instruction period, lasted from 2 to 7 s (mean 3 s, exponential distribution favouring 2 s).

Data collection runs consisted of 48 trials (16 of each trial type), and each subject performed from 7 to 12 runs (median 11 runs and 176 trials of each type per subject ; total 101 runs and 1616 trials of each type across all 10 subjects). An individual run lasted 320 s. Prosaccade, antisaccade, and nogo trials were interleaved in a random sequence, which was determined using Monte Carlo search to yield high statistical efficiency and low covariance among task components. Each of the eight possible peripheral stimulus locations (4°, 5°, 6°, or 7° to the left or right of fixation) was used twice for each of the three trials types in a given run.

2.2.5 – Behavioural Analysis

We used video eye tracking data collected concurrently with fMRI scanning to determine saccadic response latencies, where latency was defined as the time from peripheral stimulus appearance to the start of the saccade. We also discarded trials with several types of mistakes including inappropriate response type, failure to respond when required, excessively short or long response latency (latency < 100 ms or latency > 500 ms, respectively), and failure to maintain fixation. Inappropriate responses involved making a response based on instructions other than those of the given trial, namely, failing to inhibit the automatic prosaccade on an antisaccade trial, making an antisaccade on a prosaccade trial, or making either a prosaccade or an antisaccade on a nogo trial. It would have been interesting to analyze fMRI data from inappropriate response trials, but our subjects did not consistently make enough inappropriate responses to afford sufficient power for such an analysis. A fixation break was defined as a deviation of eye position

from central fixation with magnitude at least equal to the mean saccadic response magnitude across all trials minus one standard deviation. This definition ensured that fixation break magnitudes fell within the range of the largest 84.13% of all saccadic responses.

We computed percent correct values as the proportion of correct responses relative to the set of correct responses and error responses, where errors included inappropriate responses as defined above and failures to respond when required on prosaccade and antisaccade trials. This computation ignored trials with fixation breaks or late or early responses.

2.2.6 – Preprocessing

Preprocessing was done using Matlab 6 (The Mathworks, Inc., Natick, Ma, USA) and BrainVoyager 2000 (Brain Innovation BV, Maastricht, The Netherlands). All functional data were superimposed on 3D anatomical images, resampled into 3x3x3 mm cubic voxels, aligned onto the anterior commissure – posterior commissure axis, and scaled to Talairach space with BrainVoyager 2000. Functional data then underwent 3D motion correction by trilinear interpolation of each volume to the first volume and spatial low-pass filtering, performed with a low-pass cutoff of 0.069 cycles / mm in the Fourier domain.

We used Matlab 6 to remove the first four volumes from each voxel's time course to allow for complete spin saturation. We performed nonlinear trend removal on the data using a running lines smoother programmed in Matlab 6. This algorithm corrected time courses for low frequency trends as described in Marchini and Ripley (2000) and has the advantage of being robust against edge effects (see *Appendix 4.1*). Finally, voxel time courses were scaled into percent signal change values on a run-by-run basis using the mean time course across the run as baseline.

2.2.7 – Statistical Analysis and Linear Deconvolution

Because we used a rapid event-related design, with task event spacing randomized from 2 to 7 s with a mean of 3 s, the BOLD signal we recorded from any given point was the sum of the BOLD impulse responses from several preceding task events. This

necessitated the use of deconvolution methods to extract BOLD responses from individual task components. We used a general linear model (GLM) method to perform both the statistical analysis and the linear deconvolution. The process of fitting a GLM to rapid event-related fMRI data can effect a linear deconvolution because it finds the best scaling of predictor curves which model overlapping, summated BOLD signal components.

The GLM method we used performs a mixed effects analysis and is based on Worsley and colleagues (2002). Briefly, a GLM was fit to each subject's data after pre-whitening to correct for coloured noise covariance (see *Appendix 4.2*). Pre-whitening avoids the inflation of type I error that occurs with GLM statistical tests which fail to account for coloured noise in the data (Bullmore et al. 2001). The haemodynamic response model incorporated into the GLM included a difference of gammas curve and its temporal derivative (for details see *Appendix 4.3*). The temporal derivative allowed the model to account for differences in the BOLD response latency between regions and subjects (Friston et al. 1998). We chose this model because it achieved a good balance between giving the GLM enough flexibility to fit the data while constraining it so as to fit meaningful signal and not signal noise. For comparison purposes, we also analyzed the data with a finite impulse response model (see Serences 2004). This worked well for regions displaying robust BOLD activation; however, for regions with significant but low amplitude activation patterns, the finite impulse response model's estimates of the BOLD signal were unacceptably sensitive to noise in the data.

The prosaccade, antisaccade, and nogo tasks each consisted of an instruction and a response period, and we defined six task epochs: *pro_instruction*, *anti_instruction*, *nogo_instruction*, *pro_response*, *anti_response*, and *nogo_response*. For each of the ten subjects, we computed twelve statistical parametric maps using massively univariate T-tests. Six of the maps served as activation localizers and were built by contrasting activation from each of the six task epochs to a baseline value of zero. The other six maps compared pairs of tasks during the instruction or the response period, and these were computed from the following contrasts: *anti_instruction − pro_instruction*, *anti_instruction − nogo_instruction*, and *pro_instruction − nogo_instruction*, as well as

anti_response – pro_response, anti_response – nogo_response, and *pro_response – nogo_response*.

For each of the twelve statistical contrasts, the T-statistic maps were combined across subjects using a mixed effects analysis (Worsley et al. 2002), which addressed variability both within subjects and between subjects. Mixed effects contrast maps were thresholded at a T-value of 2.575 ($p < 0.01$ uncorrected, df = 15000) and then cluster size thresholded at a minimum size of 1096 cubic mm to correct for multiple comparisons across the voxel population at $p < 0.05$, based on Gaussian random field theory (see Worsley et al. 2002). These analysis steps are described in detail in *Appendix 4.3*.

2.2.8 – Definitions of Activated Regions

We defined activated regions from the various contrast maps using the patterns of activated voxels from the maps, themselves, combined with structural anatomical scans upon which the activation maps were overlain.

2.3 – Results

We compared prosaccades, antisaccades, and nogo trials in a rapid event-related fMRI design to examine processes underlying saccadic inhibition and the generation of voluntary saccadic eye movements.

2.3.1 – Behaviour

On the basis of video eye position recordings taken during fMRI data acquisition (*Figure 2.2*), we discarded trials with several types of mistakes (see *Section 2.2.5 – Behavioural Analysis*) and verified that subjects exhibited normal task behaviour in the fMRI scanner environment. Mean saccadic latencies over all subjects were 272 ± 51 (standard deviation) ms for prosaccade trials and 335 ± 54 ms for antisaccade trials. This difference was significant ($p = 0.00019$, T = 6.07, df = 9), and it matches the results of other antisaccade studies, which typically find longer latencies for antisaccades compared to prosaccades (Hallett 1978; Everling and Fischer 1998).
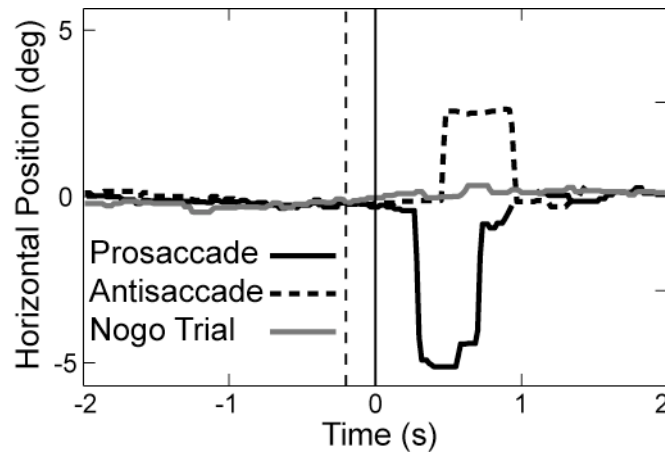
Figure 2.2 – Example Eye Traces

Horizontal component of three example eye traces are shown as a function of time, including a leftward prosaccade (solid black trace) to a peripheral stimulus 5 degrees left of fixation, a rightward antisaccade (dotted black trace) in response to the same stimulus, and a nogo response (solid gray trace). Solid vertical black line denotes peripheral stimulus onset. Dashed vertical black line denotes fixation point offset 200 ms before peripheral stimulus onset.

Mean percent correct values across subjects (see *Section 2.2.5 – Behavioural Analysis*) were 96.7 ± 2.9% , 86.7 ± 9.8%, and 95.6 ±5.6% for prosaccade, antisaccade, and nogo trials, respectively, and these were significantly different (1-way ANOVA, F = 6.56, p = 0.0048, df = 2 / 27). The antisaccade percent correct value was significantly less than the prosaccade and nogo trial values (for prosaccades vs. antisaccades and nogo trials vs. antisaccades respectively, p = 0.0020 and p = 0.00059, T = 4.31 and T = 5.17, df = 9), while the prosaccade and nogo trial values were not significantly different (p = 0.26, T = 1.20, df = 9). Lesser percent correct performance for antisaccades compared to prosaccades also agrees with the antisaccade literature (Hallett 1978; Everling and Fischer 1998; Barton et al. 2006), and lesser percent correct performance for antisaccades compared to nogo trials agrees with Barton and colleagues (2006).

2.3.2 – Response Period Activation Localizer Maps

The prosaccade, antisaccade, and nogo tasks each consisted of an instruction and a response period (see *Section 2.2.4 – Experimental Design*), and we defined six task epochs: *pro_instruction*, *anti_instruction*, *nogo_instruction*, *pro_response*, *anti_response*, and *nogo_response*. Activation levels for these various epochs were contrasted using mixed-effects T-tests as described in *Section 2.2 – Methods*. We found our most important results in the response period, and those are discussed first.

To visualize brain regions recruited by task response execution, we computed activation localizer maps that contrasted the main predictor beta weight for the *pro_response*, *anti_response*, and *nogo_response* epochs against fixation baseline during the inter-trial interval (*Figure 2.3*, *Table 2.1*). We found response-related activation for all three tasks bilaterally in the middle frontal gyrus (MFG), frontal eye field (FEF), supplementary eye field (SEF), anterior and posterior cingulate cortex (ACC and PCC), intraparietal sulcus (IPS), dorsal occipital lobe, and precuneus (*pro_response* activation was limited to right precuneus). That is, all three tasks in our rapid event-related paradigm evoked activation in cortical regions typically recruited by saccade tasks (for examples see Anderson et al. 1994; Sweeney et al. 1996; Berman et al. 1999; Luna and Sweeney 1999; Kimmig et al. 2001; Connolly et al. 2002; Curtis and D'Esposito 2003; DeSouza et al. 2003; Brown et al. 2004; Ford et al. 2005). This provides important verification that a rapid event-related saccade paradigm recruits activation in the same areas found to be activated by saccade tasks in the fMRI literature.

All three response activation localizer maps also included activated regions in left and right anterior insula. The *anti_response* localizer map included activated voxels in bilateral dorsal caudate body, while the *nogo_response* localizer included activated voxels in left dorsal caudate body. In addition, all three response period localizer maps included deactivated voxel clusters in the medial frontal gyrus, and the *pro_response* and *anti_response* localizers contained deactivated regions in posterior cingulate cortex / ventral precuneus and bilateral supramarginal gyrus / posterior superior temporal sulcus / angular gyrus.
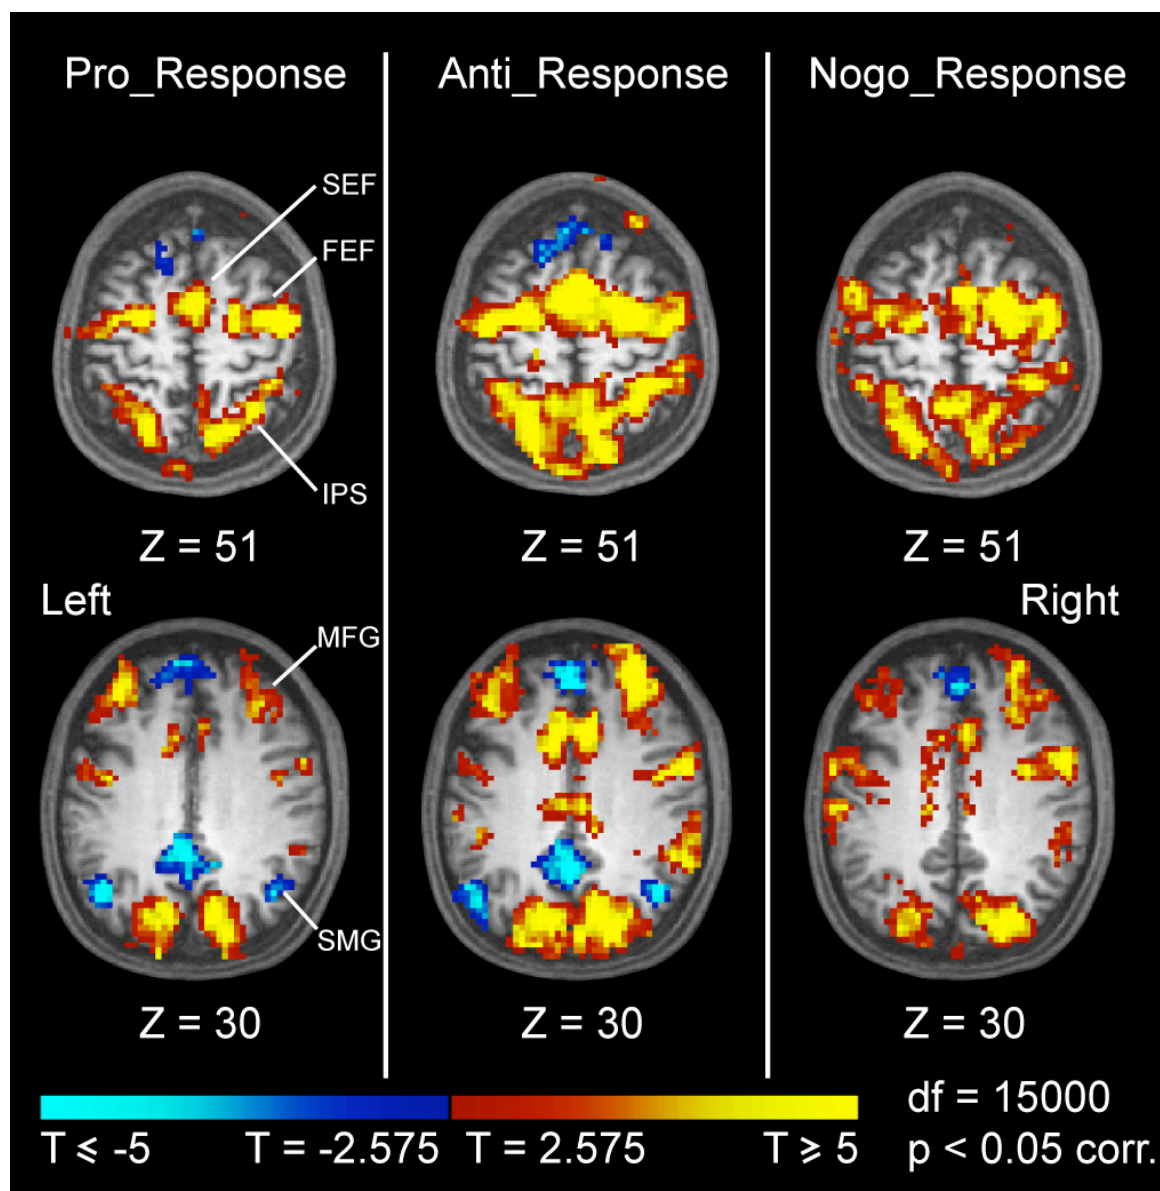
Figure 2.3 – Response Period Localizer Maps

Figure shows activation localizer contrast maps comparing *pro_response*, *anti_response*, and *nogo_response* against baseline. Typical saccade-related areas in frontal cortex and intraparietal sulcus are active. Data average of 10 subjects. Images are in neurological coordinates; the left side of an axial image corresponds to the left side of the brain. FEF – frontal eye field, SEF – supplementary eye field, MFG – middle frontal gyrus, IPS – intraparietal sulcus, SMG – supramarginal gyrus.

Table 2.1 – Response Period Localizer Data

Region	Pro_Response					Anti_Response					Nogo_Response				
	Dir	x	y	z	Vol	Dir	x	y	z	Vol	Dir	x	y	z	Vol
Right FEF	+	36	-5	50	21,330	+	35	-4	50	36,774	+	33	-1	49	39,528
Left FEF	+	-38	-6	47	17,685	+	-34	-5	49	25,677	+	-39	-2	46	28,836
Right MFG	+	34	42	33	6,804	+	31	47	31	19,278	+	32	45	29	13,905
Left MFG	+	-35	44	31	7,317	+	-34	45	29	12,906	+	-34	44	25	9,720
Right Insula	+	49	3.3	17	2,565	+	52	5.3	17	2,268	+	42	1.4	17	4,536
Left Insula	+	-36	4.3	13	3,078	+	-43	7	14	3,078	+	-49	4.8	16	4,482
SEF	+	0.2	-4	56	6,156	+	0.6	0.2	57	15,444	+	3.1	-1	58	10,692
ACC	+	-0	9.1	38	8,019	+	0.1	13	36	20,250	+	0.6	8	38	14,256
Medial PFC	-	-4	44	33	10,557	-	-3	44	37	11,043	-	0.9	47	27	3,942
Perigenu PFC	-	-1	25	15	1,296										
Right Rostral IPS	+	36	-46	47	6,426	+	37	-43	45	15,201	+	38	-42	44	12,690
Left Rostral IPS	+	-36	-45	45	5,022	+	-37	-44	44	10,098	+	-37	-45	45	13,014
Right Caudal IPS	+	19	-66	44	8,073	+	19	-67	44	11,961	+	21	-67	45	11,691
Left Caudal IPS	+	-21	-66	44	8,235	+	-21	-66	46	12,933	+	-23	-66	43	12,555
Right IPL	+	57	-39	40	2,700	+									
Right Parietal Surface						+	58	-37	35	11,151	+	56	-36	37	4,860
Left Parietal Surface						+	-56	-20	22	1,107	+	-58	-23	34	8,802
Right SMG	-	44	-61	27	1,485	-	41	-62	28	1,701					
Left SMG	-	-44	-62	27	2,673	-	-48	-67	29	4,374					
PCC	+	12	-32	42	621	+	-0	-26	29	12,042	+	4.5	-26	33	6,912
Precuneus	+	2.4	-67	40	7,236	+	-0	-62	45	23,436	+	0.4	-60	45	20,466
Dorsal Precuneus	-	6.2	-36	65	1,728										
PCC / vPrecuneus	-	-4	-44	28	6,966	-	-3	-49	26	5,238					
Right dOccipital	+	16	-82	25	8,424	+	17	-82	26	12,393	+	25	-80	26	6,102
Left dOccipital	+	-18	-81	25	7,479	+	-16	-82	26	12,258	+	-22	-81	28	6,156
Right Caudate / Thalamus						+	18	-1	16	1,998					
Left Caudate / Thalamus						+	-19	-4	15	4,671	+	-14	0.3	22	648

Legend: Dir refers to direction of contrast: + above fixation baseline, - below fixation baseline. X, Y, and Z refer to Talairach coordinates of statistical centre of mass. Vol is volume in cubic millimeters. ACC - anterior cingulate cortex, dOccipital - dorsal occipital, FEF - frontal eye field, IPL - inferior parietal lobule, IPS - intraparietal sulcus, MFG - middle frontal gyrus, PCC - posterior cingulate cortex, PFC - prefrontal cortex, SEF - supplementary eye field, SMG - supramarginal gyrus, vPrecuneus - ventral precuneus.

2.3.3 – Response Period Pair Wise Contrasts

To compare response-related activation levels among the three task types, we computed three pair wise contrasts using mixed effects T-tests: *anti_response – pro_response*, *anti_response – nogo_response*, and *pro_response – nogo_response* (*Figures 2.4* and *2.5*; *Table 2.2*) (for details see *Section 2.2 – Methods*).

The *pro_response – nogo_response* contrast revealed no activation differences in FEF, SEF, ACC, and IPS, which was surprising given that the nogo task required saccadic inhibition while the prosaccade required saccade generation. The similarity of the *pro_response* and *nogo_response* activation patterns in these areas is illustrated by

their BOLD time courses (*Figure 2.6*). We also found large voxel clusters with greater *anti_response* activation compared to either *pro_response* or *nogo_response* activation bilaterally in FEF, SEF, ACC, IPS, PCC, and precuneus.

In bilateral middle frontal gyrus (MFG, also called Brodmann's area 46 and dorsolateral prefrontal cortex), we found greater activation in the *anti_response − nogo_response* comparison but not in the *anti_response − pro_response* or *pro_response − nogo_response* comparisons when correcting for multiple comparisons across the voxel population. However, about half of the voxels within both left and right MFG, as defined from the *anti_response − nogo_response* contrast, exhibited significantly greater *anti_response* activation in the *anti_response − pro_response* comparison at an uncorrected $p < 0.05$ (T > 1.96, df = 15000, for 39 / 87 voxels in left MFG and 42 / 69 voxels in right MFG, where each voxel was 27 mm^3 in volume). No voxels within left or right MFG exhibited significant differences in the *pro_response − nogo_response* comparison, even without correction for multiple comparisons (left MFG: maximum single voxel |T| = 1.78, p = 0.076, df = 15000, and right MFG: maximum single voxel |T| = 1.38, p = 0.17, df = 15000).

In dorsal left caudate body, *anti_response* activation was larger than both *pro_response* and *nogo_response* activation. Voxel clusters in superficial right postcentral sulcus and deep left central sulcus exhibited greater *anti_response* versus *pro_response* activation, and bilateral parieto-occipital sulcus exhibited greater *pro_response* versus *nogo_response* activation.

Outside the standard cortical saccade system, several regions exhibited greater *nogo_response* activation compared to *anti_response* and/or *pro_response* activation (blue regions in *Figures 2.4* and *2.5*; see also *Table 2.2* and *Figure 2.6*). We found significantly greater *nogo_response* versus *pro_response* activation in right superior frontal sulcus (SFS), right supramarginal gyrus (SMG), and caudal posterior cingulate sulcus (PCS). In right SFS and right SMG, *anti_response* activation fell in between *nogo_response* and *pro_response* activation (*Figure 2.6*), though *anti_response* activation was not significantly different from either *nogo_response* or *pro_response* activation. In PCS, *anti_response* activation was greater than both *nogo_response* and

Figure 2.4 – Response Period Comparison Maps (Axial)

Figure shows axial sections of statistical contrast maps for the *anti_response –
pro_response, anti_response – nogo_response* and *pro_response – nogo_response*
comparisons. The first two contrasts reveal greater *anti_response* versus either
pro_response or *nogo_response* activation in FEF, SEF, ACC, and intraparietal sulcus,
while the third contrast indicates an absence of significant differences between
pro_response and *nogo_response* signal in these areas. "Hot" colours denote voxels with
greater BOLD activation on the *anti_response*, *anti_response*, and *pro_response* epochs,
respectively for the three comparisons. "Cold" colours denote the opposite. Data average
of 10 subjects. Images are in neurological coordinates. SFS – superior frontal sulcus.
Other abbreviations as in *Figure 2.3*.
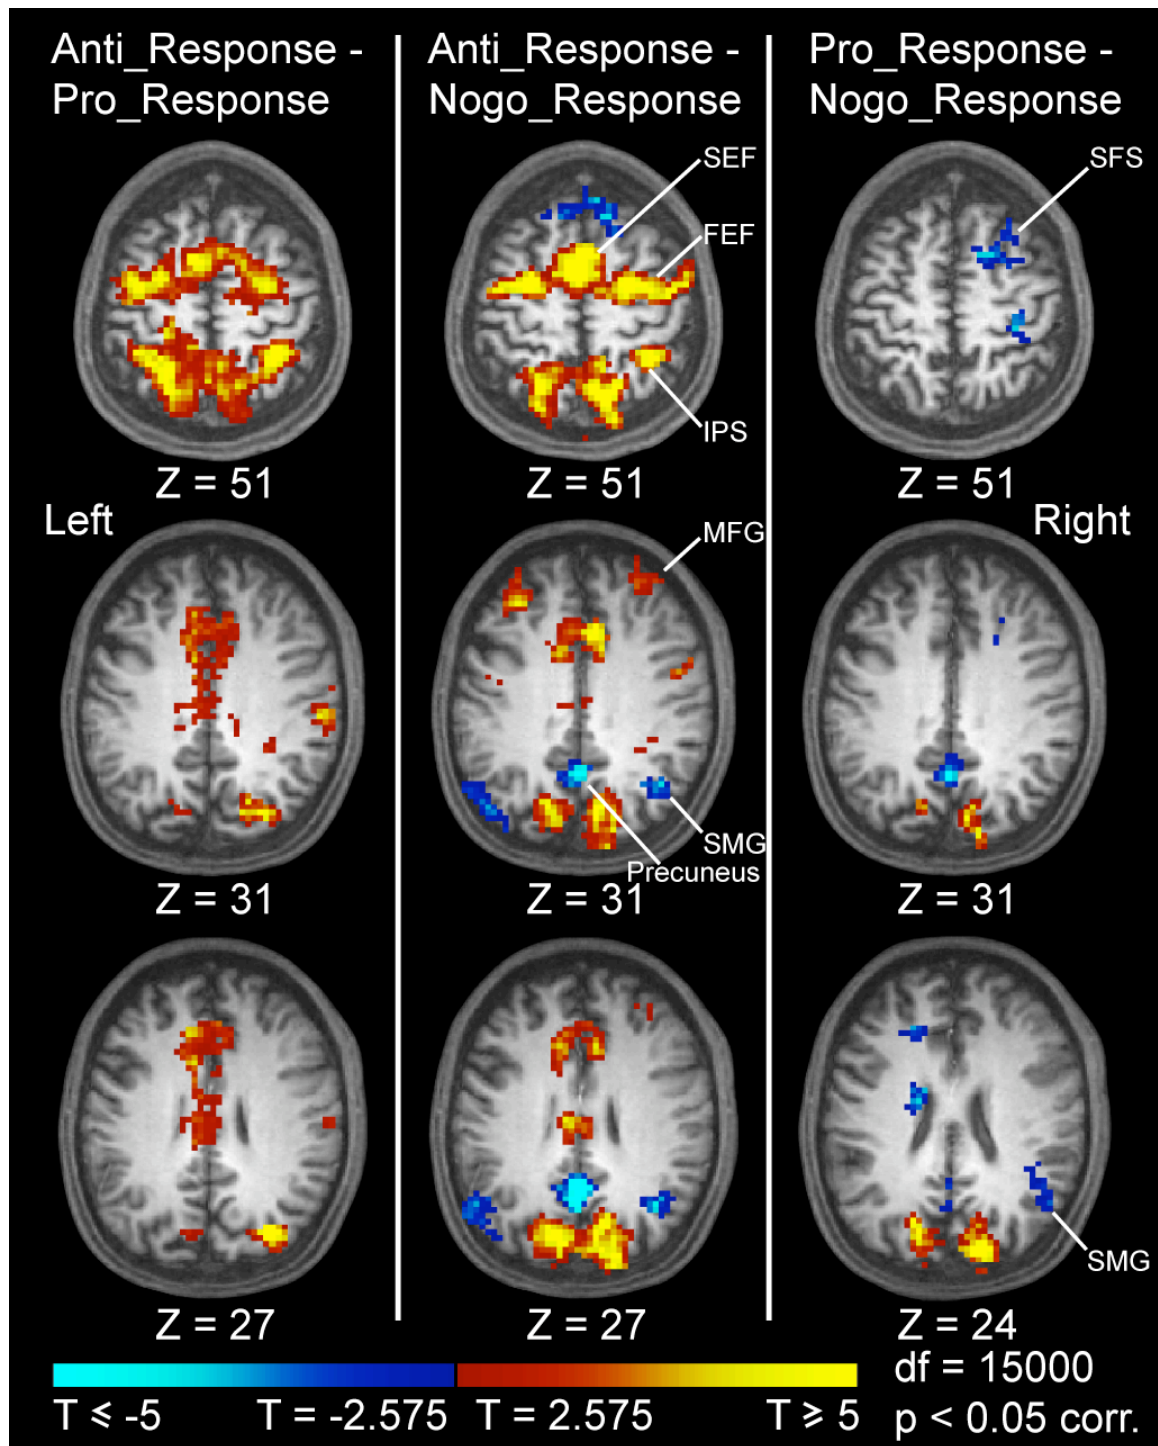
Figure 2.5 – Response Period Comparison Maps (Sagittal)

Figure shows sagittal sections of statistical contrast maps for the *anti_response – pro_response, anti_response – nogo_response* and *pro_response – nogo_response* comparisons. The first two contrasts reveal greater *anti_response* versus either *pro_response* or *nogo_response* activation in SEF, ACC, and precuneus. The second and third contrasts show greater *nogo_response* versus either *anti_response* or *pro_response* activation in the junction of ventral precuneus and posterior cingulate cortex. "Hot" colours denote voxels with greater BOLD activation on the *anti_response*, *anti_response*, and *pro_response* epochs, respectively for the three comparisons. "Cold" colours denote the opposite. Data average of 10 subjects. ACC – anterior cingulate cortex, SEF – supplementary eye field, PCS – posterior cingulate sulcus.
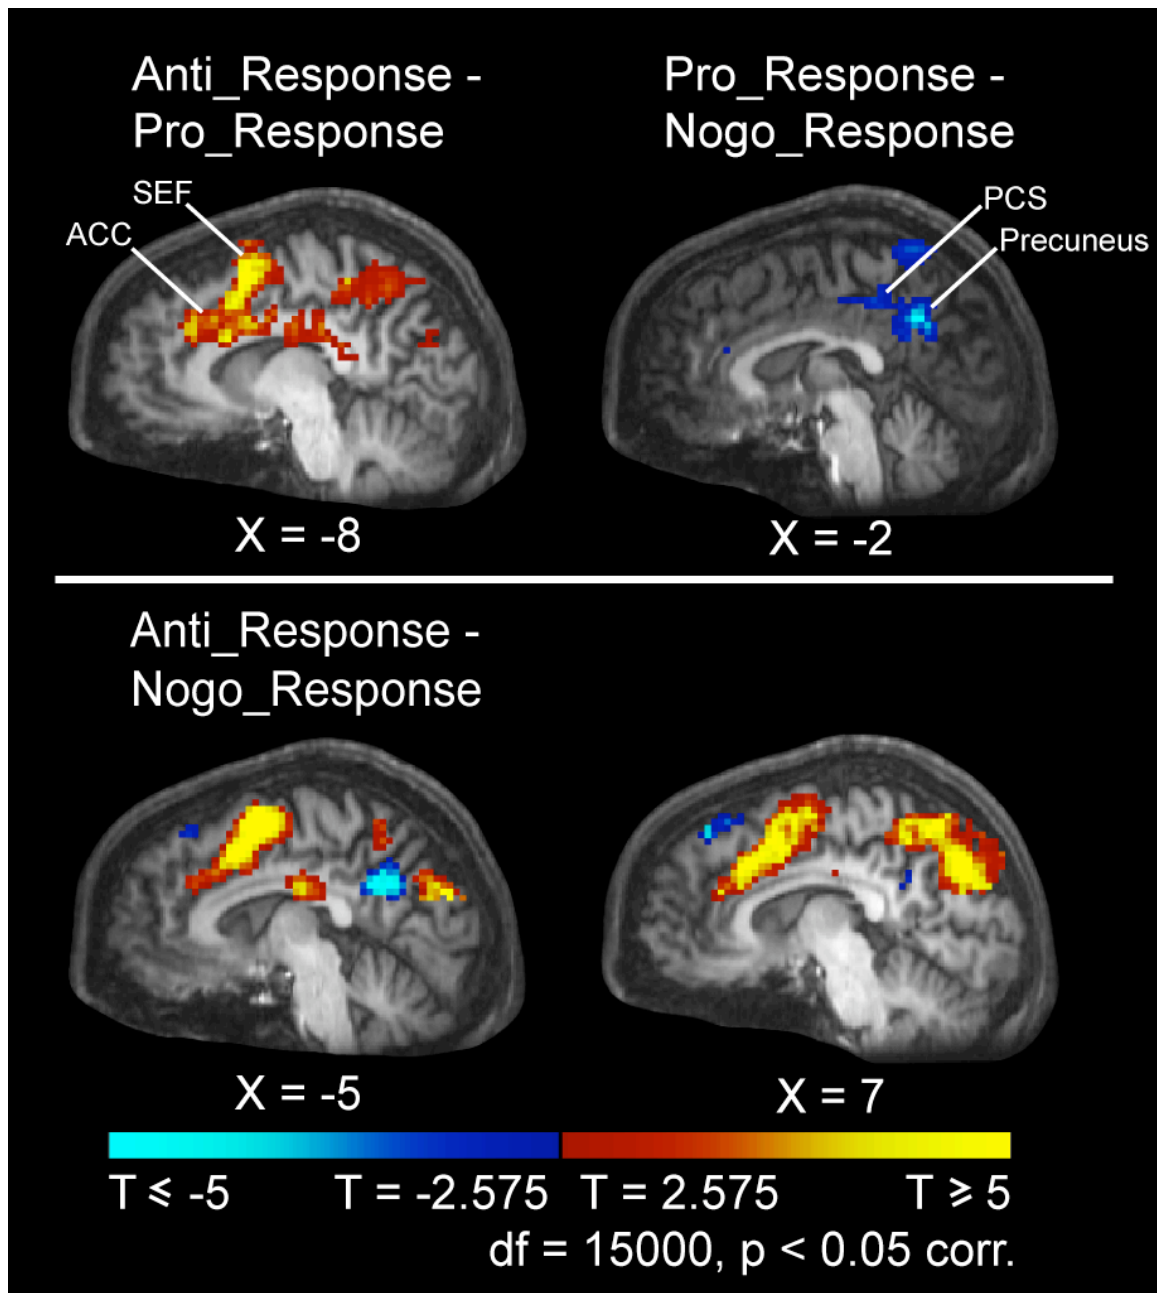
Table 2.2 – Response Period Pair Wise Contrast Data

Region	Anti - Pro_Response					Anti - Nogo_Response					Pro - Nogo_Response				
	Dir	x	y	z	Vol	Dir	x	y	z	Vol	Dir	x	y	z	Vol
Right FEF	+	23	-5	54	10,854	+	33	-6	53	18,387					
Left FEF	+	-29	-5	51	13,770	+	-28	-5	52	12,312					
Right MFG		(see Results)				+	32	48	33	1,863					
Left MFG		(see Results)				+	-32	42	35	2,349					
Right SFS											-	24	18	45	4,806
SEF	+	-4	3.6	56	8,127	+	-1	-0	56	10,611					
ACC	+	-2	14	34	17,226	+	0.3	14	36	13,770					
Bilateral SFG						-	8.1	30	51	4,455					
Medial PFC											-	-14	29	18	1,512
Right Rostral IPS	+	31	-44	46	7,506	+	34	-43	45	4,347					
Left Rostral IPS	+	-29	-45	51	4,995	+	-28	-52	49	351					
Right Caudal IPS	+	19	-66	44	6,750	+	16	-68	45	5,535					
Left Caudal IPS	+	-20	-61	49	6,993	+	-20	-64	49	6,696					
Right Descending IPS	+	30	-77	26	2,187	+	18	-83	27	4,104					
Left Descending IPS	+	-8	-76	25	513	+	-19	-79	25	2,079					
Right SMG						-	39	-61	28	1,566	-	46	-52	22	1,404
Left SMG						-	-48	-67	28	5,238					
Right Central Sulcus	+	59	-23	33	1,917						-	33	-28	49	1,215
Left Central Sulcus	+	-21	-31	49	1,566										
PCC	+	-2	-25	32	7,587	+	-4	-22	25	2,565	-	-1	-36	41	1,107
Precuneus	+	-2	-56	49	14,202	+	2.5	-63	45	12,339					
Dorsal Precuneus											-	-1	-52	63	1,485
Ventral Precuneus						-	-2	-54	27	3,861	-	-1	-54	32	2,160
Right dOccipital											+	13	-82	27	5,184
Left dOccipital											+	-14	-77	26	2,592
Medial Occipital						+	0.1	-81	27	6,939					
Left Caudate	+	-16	1.2	16	1,215	+	-16	-6	15	2,862	-	-17	-4	25	1,458

Legend: Dir refers to direction of contrast: + above fixation baseline, - below fixation baseline. X, Y, and Z refer to Talairach coordinates of statistical centre of mass. Vol is volume in cubic millimeters. ACC - anterior cingulate cortex, dOccipital - dorsal occipital, FEF - frontal eye field, IPS - intraparietal sulcus, MFG - middle frontal gyrus, PCC - posterior cingulate cortex, PFC - prefrontal cortex, SEF - supplementary eye field, SFG - superior frontal gyrus, SFS - superior frontal sulcus, SMG - supramarginal gyrus.

Figure 2.6 – Modeled Time Courses

Figure shows general linear model (GLM) derived instruction and response epoch activation time courses for select regions defined from various response period contrast maps. Prosaccade, antisaccade, and nogo trial curves are depicted in green, red, and blue, respectively. Time courses for right frontal eye field (FEF), supplementary eye field (SEF), anterior cingulate cortex (ACC), rostral intraparietal sulcus (IPS), and caudal IPS, all of which are shown in the two leftmost columns, were derived from voxels that showed significant differences in both the *anti_response* − *pro_response* and *anti_response* − *nogo_response* contrasts. Right middle frontal gyrus (MFG) was taken from the *anti_response* − *nogo_response* contrast map. Right superior frontal sulcus (SFS), right supramarginal gyrus (SMG), posterior cingulate sulcus (PCS), and precuneus were taken from the *pro_response* − *nogo_response* contrast map. Time courses are linear combinations of the canonical predictor set (*Appendix Figure 5.1*) with mean GLM derived beta weights across voxels and across subjects (see *Section 2.3.4 – GLM Derived Bold Time Courses*).
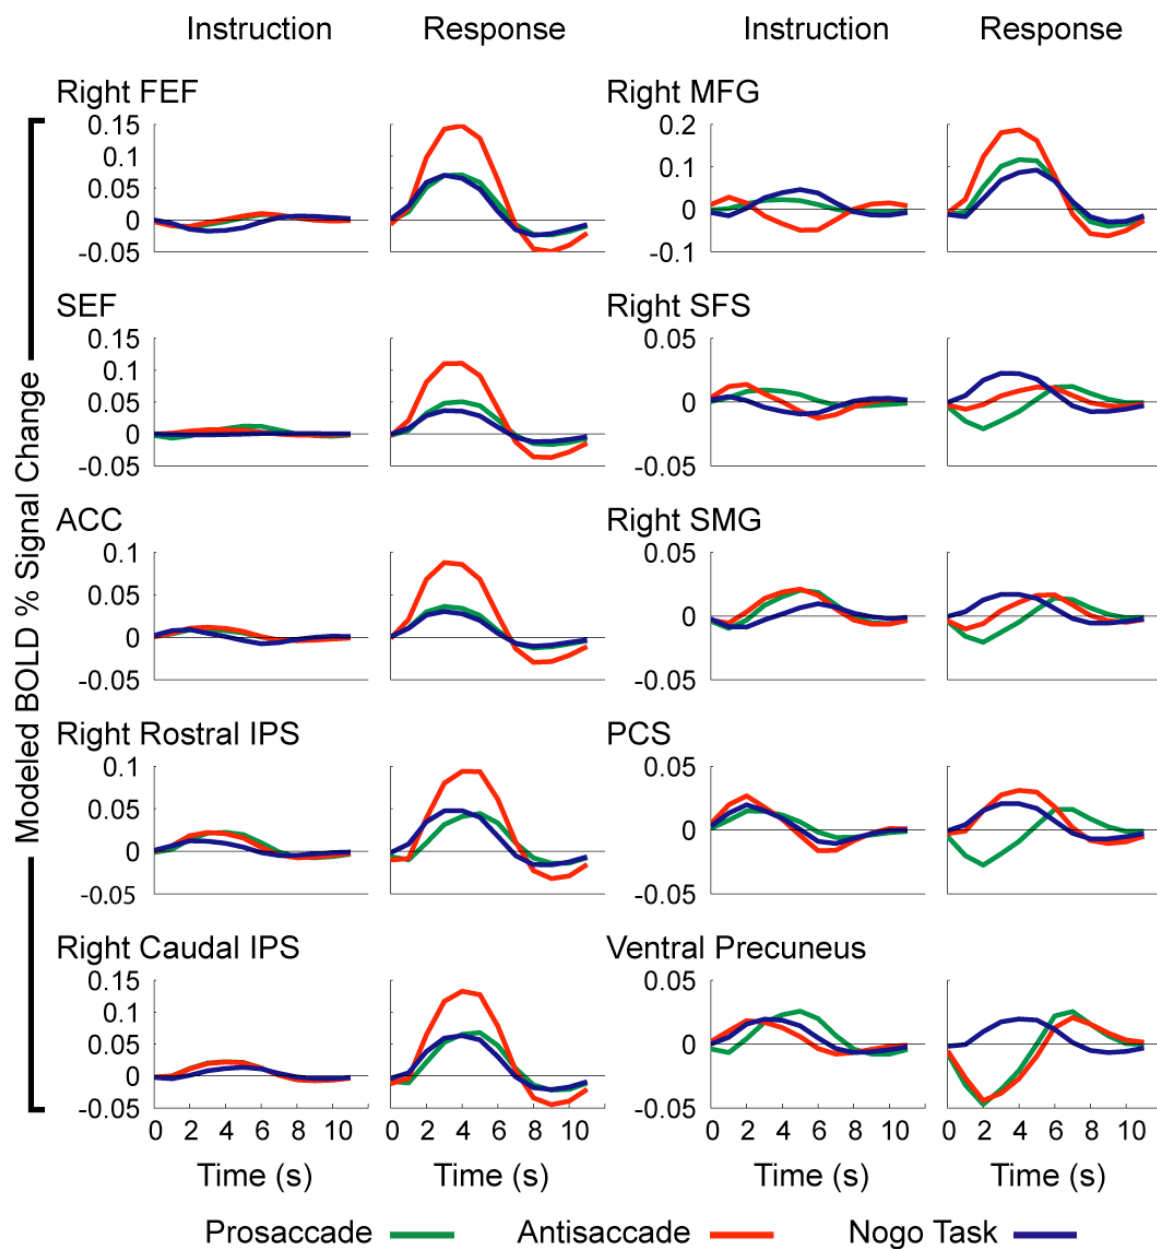
pro_response activation, though these differences were not significant. The activation patterns displayed by SFS, right SMG, and PCS are consistent with their playing a role in active inhibition of saccadic eye movements, as outlined in *Section 2.4 – Discussion*.

On the medial aspect, a bilateral region in ventral precuneus exhibited greater *nogo_response* versus *pro_response* activation. Part of this region (53 of 80 voxels) overlapped with a bilateral region straddling the junction between ventral precuneus and caudal PCC that exhibited greater *nogo_response* versus *anti_response* activation. We found voxel clusters with greater *nogo_response* versus *pro_response* activity in left medial frontal gyrus, right central sulcus, dorsal precuneus, and left caudate body. We also found greater *nogo_response* versus *anti_response* activation in bilateral rostral superior frontal gyrus and in a bilateral region encompassing caudal supramarginal gyrus and anterior angular gyrus. The right supramarginal gyrus regions from the *anti_response – nogo_response* and *pro_response – nogo_response* contrasts were almost completely distinct, with only two voxels overlapping between them.

2.3.4 – GLM Derived BOLD Time Courses

Note that the time courses depicted in *Figure 2.6* were generated by scaling the two curves comprising the canonical BOLD response model (*Appendix Figure 5.1*) by the GLM derived beta weights from a given region and then adding the results. For comparison purposes, we computed time courses (not shown) using a deconvolution model incorporating a finite impulse response (FIR) predictor set (see Serences 2004). For FEF, SEF, ACC, IPS, and MFG, which exhibited robust response epoch fMRI signals, the response period activation time courses computed with the FIR predictor set were similar to the time courses displayed in *Figure 2.6*. However, in instances where the BOLD signal was less robust, the FIR predictor model fit noise rather than meaningful signal, resulting in choppy FIR derived time courses that did not resemble the time courses depicted in *Figure 2.6*. This was the case for the instruction period in FEF, SEF, ACC, IPS, and MFG and for both the instruction and response periods in SFS, SMG, PCS, and precuneus. In contrast to the FIR predictor set, the canonical BOLD model (*Appendix Figure 5.1*) that we incorporated into our GLM analysis was robust against

signal noise because it made assumptions about the haemodynamic response shape which constrained the GLM such that it fit signal rather than noise.

We did not compute event-related average time courses because these would have been severely skewed by our compound trial design. That is, each response epoch for a given trial type was necessarily preceded by that trial type's instruction epoch, creating a sequence effect which would have skewed event-related average time courses locked either to the instruction or to the response period (see Serences 2004).

2.3.5 – Instruction Period Results

We generated localizer maps comparing activation for the *pro_instruction*, *anti_instruction*, and *nogo_instruction* periods to baseline using mixed effects T-tests (*Figure 2.7, Table 2.3*; see *Section 2.2 – Methods* for details). These maps revealed relatively little activity compared to the response period activation localizers discussed above. The *pro_instruction* localizer contained activated regions in the right rostral frontal pole, the left prefrontal cortex anterior to the insula, the left rostral end of IPS extending into the adjacent superficial part of postcentral sulcus, bilateral IPS, and precuneus. The *anti_instruction* localizer contained activated regions in left prefrontal cortex antero-medial to the left insula extending into the medial frontal gyrus, left FEF, bilateral IPS, PCC, and precuneus. The *nogo_instruction* localizer contained activated regions in PCC, precuneus and right IPS as well as deactivated regions in left and right FEF and left posterior insula.

We also found few instruction-related differences in pair wise comparisons of the tasks based on mixed effects T-tests. The *anti_instruction – pro_instruction* comparison revealed no voxel clusters with significant differences. The *anti_instruction – nogo_instruction* comparison (*Figure 2.8, Table 2.3*) revealed three small cortical regions with greater *anti_instruction* activation in left and right FEF and left caudal precuneus. The *pro_instruction – nogo_instruction* comparison (*Figure 2.8, Table 2.3*) included one region with greater *pro_instruction* activation in the left precuneus.
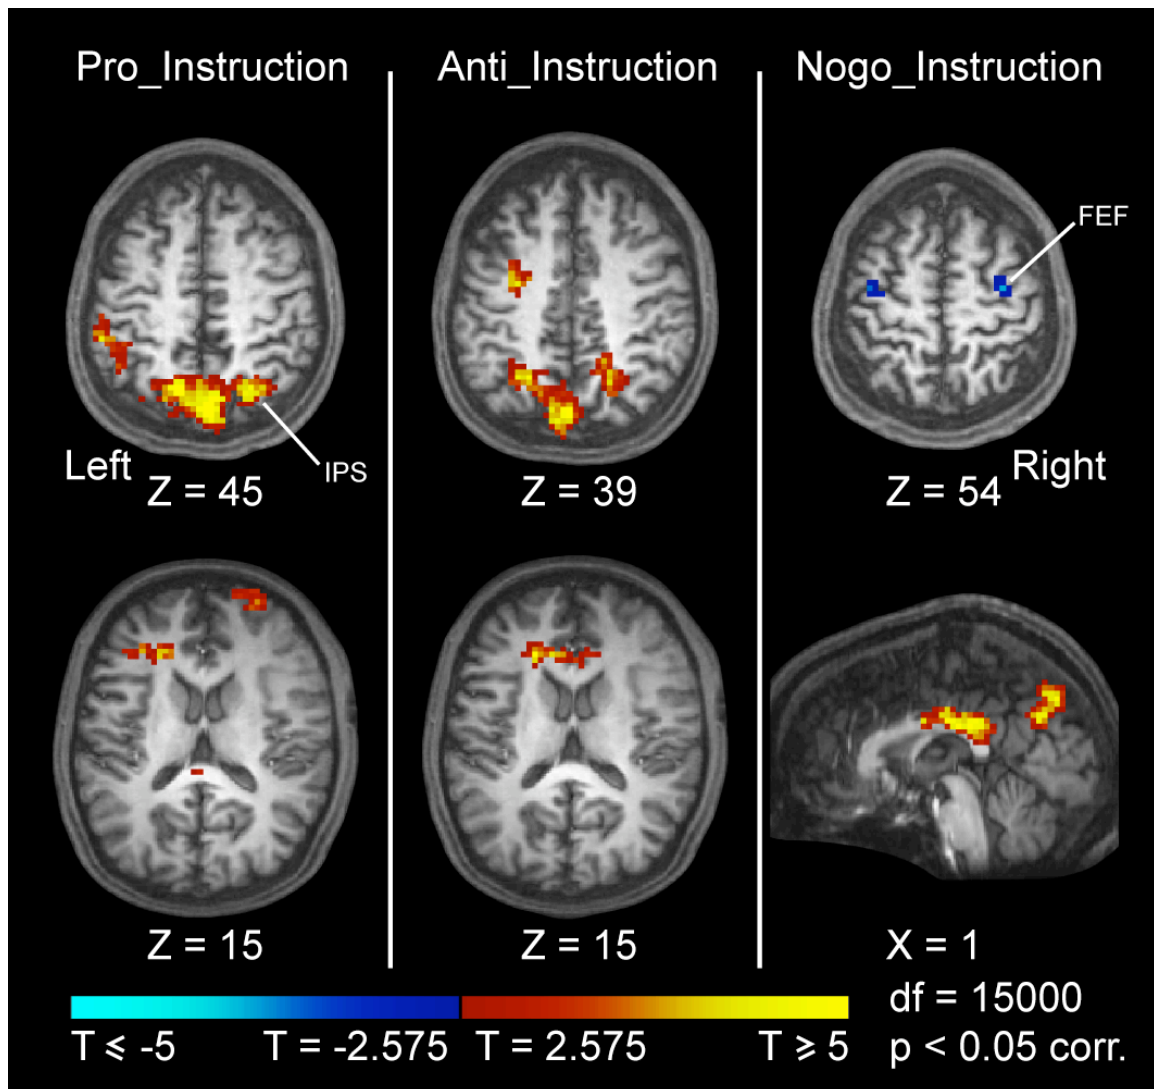
Figure 2.7 – Instruction Period Localizer Maps

Figure shows activation localizer contrast maps comparing *pro_instruction*, *anti_instruction*, and *nogo_instruction* against baseline. "Hot" colours indicate activation above baseline, while "cold" colours denote the opposite. Data average of 10 subjects. Images are in neurological coordinates. FEF – frontal eye field, IPS – intraparietal sulcus.

Table 2.3 – Instruction Period Results

Region	Pro_Instruction					Anti_Instruction					Nogo_Instruction				
	Instruction Period Localizer Data														
	Dir	x	y	z	Vol	Dir	x	y	z	Vol	Dir	x	y	z	Vol
Right FEF											-	30	-9	56	1,431
Left FEF						+	-27	-0	35	2,025	-	-30	-9	61	1,539
Left Insula											-	-50	-26	16	1,269
Left PFC	+	-27	35	18	3,105	+	-15	32	16	2,808					
						+	-22	19	24	972					
Left Frontal Pole	+	25	62	15	1,377										
Medial PFC						+	6.2	29	14	621					
Right Rostral IPS	+	29	-53	30	2,187	+	22	-47	32	1,323	+	28	-48	39	1,107
Left Rostral IPS	+	-30	-53	27	594	+	-31	-52	37	1,620					
Right Caudal IPS	+	23	-63	36	4,968	+	20	-59	36	4,428	+	21	-63	38	2,349
Left Caudal IPS	+	-23	-63	35	4,833	+	-22	-60	39	3,942	+	-16	-65	28	81
Right Descending IPS	+	26	-70	22	243										
Left Descending IPS						+	-29	-73	25	297					
Left Central Sulcus	+	-47	-39	47	3,051										
PCC	+	2.8	-29	23	2,592	+	1.8	-25	25	2,214	+	-0	-26	24	5,940
Precuneus	+	-1	-69	43	13,284	+	-5	-68	47	10,962	+	2.4	-69	33	4,347
Superficial POS						+	-4	-85	33	1,512					

Region	Anti - Pro_Instruction					Anti - Nogo_Instruction					Pro - Nogo_Instruction				
	Instruction Period Pair Wise Contrast Data														
	Dir	x	y	z	Vol	Dir	x	y	z	Vol	Dir	x	y	z	Vol
Right FEF						+	25	-1	43	1,917					
Left FEF						+	-28	-1	41	2,160					
Precuneus						+	-11	-70	45	1,890	+	-6	-67	48	2,376

Legend: Dir refers to direction of contrast: + above fixation baseline, - below fixation baseline. X, Y, and Z refer to Talairach coordinates of statistical centre of mass. Vol is volume in cubic millimeters. FEF - frontal eye field, IPS - intraparietal sulcus, PCC - posterior cingulate cortex, PFC - prefrontal cortex, POS - parieto-occipital sulcus.
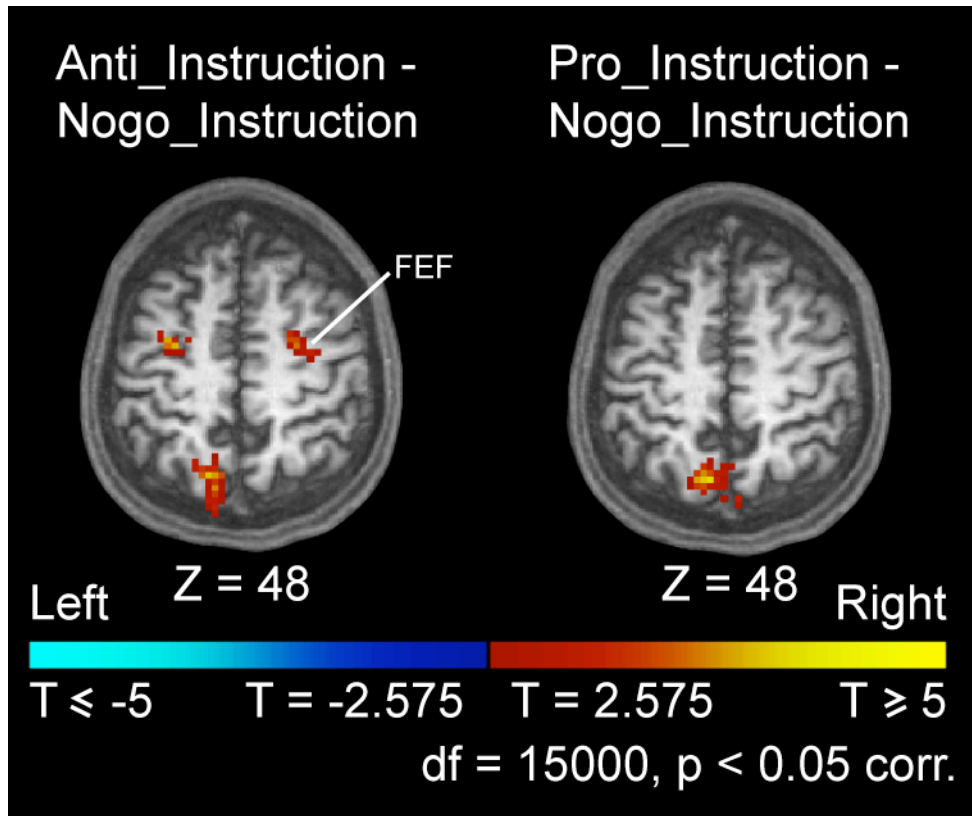
Figure 2.8 – Instruction Period Comparison Maps

Figure shows statistical contrast maps for the *anti_instruction – nogo_instruction* and *pro_instruction – nogo_instruction* comparisons. Maps reveal few statistically significant voxel clusters. The *anti_instruction – pro_instruction* comparison is not shown because it did not reveal any voxel clusters with significant differences. "Hot" colours denote voxels with greater BOLD activation on the *anti_instruction* or *pro_instruction* epoch compared to the *nogo_instruction* epoch, and "cold" colours denote the opposite. Data average of 10 subjects. Images are in neurological coordinates. FEF – frontal eye field.

2.4 – Discussion

We compared prosaccades, antisaccades, and nogo trials using rapid event-related fMRI in humans. These tasks, respectively, required subjects to look at a peripheral stimulus, to look away from it, or to inhibit the automatic saccade and maintain central fixation. Our objective was to ascertain the relative contributions of saccadic inhibition and voluntary saccadic response generation to cortical fMRI signal patterns.

2.4.1 – Visual Detection and Attention in Cortical Saccade Areas

We observed greater BOLD activation for antisaccade responses (*anti_response*) than prosaccade responses (*pro_response*) in FEF, SEF, IPS, and ACC as in previous studies (Connolly et al. 2002; Curtis and D'Esposito 2003; DeSouza et al. 2003; Ford et al. 2005). However, these areas exhibited very similar activation patterns for prosaccade and nogo trial responses (*nogo_response*), which was surprising given these tasks' very different performance requirements, namely saccade generation versus saccadic inhibition, respectively. The absence of greater *nogo_response* activation fails to provide conclusive evidence for saccadic inhibition processes in these areas. Likewise, the lack of greater *pro_response* activation fails to provide evidence for saccade generation processes. This result suggests, then, that neither saccadic inhibition nor saccade generation are responsible for BOLD activation patterns in these areas.

One possibility is that BOLD signals in cortical saccade regions might predominantly reflect processes underlying visual stimulus detection rather than saccade inhibition or generation. This idea is supported by neuronal recording studies indicating that monkey FEF and lateral intraparietal area (LIP) neurons reflect visual detection, attention, and discrimination in instances when monkeys do not make saccades. Two experiments using a go-nogo task found elevated discharge levels in FEF and LIP neurons when a peripheral visual stimulus fell in their receptive fields during nogo trials, which required central fixation maintenance rather than saccade generation (Pare and Wurtz 2001; Sommer and Wurtz 2001). Thompson and colleagues (1997) combined an oddball discrimination task with a go-nogo task, requiring monkeys to make a saccade to an oddball stimulus in a distractor array on go trial blocks, whereas they had to maintain central fixation when presented with the oddball and distractor array on nogo trial blocks.

Although the oddball stimulus was irrelevant to correct performance on the nogo trials, FEF neurons whose discharge levels discriminated between oddball and distractor stimuli in the go task continued to exhibit the discrimination in the nogo task.

Covert shifts of attention during nogo trials might also explain the similar *pro_response* and *nogo_response* data. There exists evidence that attention is briefly diverted to the location of a suddenly appearing visual stimulus (Jonides 1981; Yantis and Jonides 1984), and overt eye movements and covert shifts in attention are known to recruit similar BOLD activation patterns in cortical saccade regions (Corbetta et al. 1998; Corbetta and Shulman 2002).

The hypotheses that BOLD activation patterns in cortical saccade regions are determined mainly by visual stimulus detection or covert attention shifts as discussed in the last two paragraphs are not mutually exclusive. Rizzolatti and colleagues (1987) proposed the premotor theory of attention, which states that covert attention shifting and overt saccade generation share neuronal circuitry in the initial stages of processing. Our results are consistent with this idea.

The premotor theory of attention is also supported by neuronal recording studies in FEF which provided evidence for two processes, one devoted to selection of visual stimuli and one devoted to selection of saccade end points (Sato and Schall 2003; Schall 2004; Thompson et al. 2005). The first two studies (Sato and Schall 2003; Schall 2004) required monkey subjects to discriminate the oddball among several distractors based on colour, and then make prosaccade, antisaccade, or nogo responses to the oddball as indicated by its shape. Two thirds of FEF neurons, which they called Type I, were selective for the location of the oddball stimulus. In the antisaccade task, the majority of Type I neurons subsequently changed their activity to be selective for the antisaccade end point. Most Type I neurons also exhibited oddball stimulus selectivity in the nogo task, which did not require any saccade execution. The remaining third of FEF neurons were of Type II, and these selected the saccade end point and not the oddball stimulus location on the antisaccade task. On the nogo task, Type II neurons selected what would be the antisaccade end point if the trial were an antisaccade trial. Schall (2004) attributed this to an effect of how the monkeys were trained on the tasks. Thompson and colleagues (2005) had monkeys press a lever in the direction of the oddball stimulus in a distractor array

without making any saccades. They found that FEF neurons exibiting visual or visuo-motor tuning on memory-guided saccades also discriminated the oddball in the lever task, while purely motor neurons were suppressed during the lever task. Schall (2004) and Thompson and colleagues (2005) suggest that the visual stimulus selection process is common to covert attention shifting and overt saccade generation, while the saccade end point selection process is unique to saccade generation. That two thirds of FEF neurons were Type I, or visual stimulus selective, is important as it is reasonable to expect these neurons, being the majority, to exert the largest influence on FEF BOLD signals, which fits our results and our suggestion that BOLD activation in FEF is driven by visual stimulus detection and covert attention orientation.

That *pro_response* and *nogo_response* BOLD signals were significantly above baseline eliminates the possibility that their similarity was caused by neither task's recruiting substantial neuronal processing.

To summarize, similar *pro_response* and *nogo_response* activation patterns in FEF, SEF, IPS, and ACC suggest that BOLD signal in these areas is determined by processes underlying visual detection and attention rather than saccade generation or inhibition.

2.4.2 – Saccadic Motor Discharges revealed by Single Neuron Electrophysiology

Our suggestion that BOLD activation in cortical saccade regions might be governed predominantly by visual detection and attention rather than saccade generation or inhibition seems superficially at odds with some results from previous electrophysiological studies that found neuronal discharge patterns related specifically to saccade generation in cortical saccade regions (see Wurtz and Goldberg 1989; Schall 2002; Munoz and Everling 2004). Sato and Schall (2003) also found evidence for a saccade selection process separate from visual stimulus selection in FEF, and neuronal discharge patterns coding for antisaccade performance have been found in SEF (Schlag-Rey et al. 1997; Amador et al. 2004). As discussed above, Sato and Schall (2003) and Schall (2004) found that two thirds of FEF neurons were selective for visual stimuli rather than saccade end points, and one might therefore expect the visual stimulus selection process rather than the saccade end point selection process to exert the strongest

influence on FEF BOLD signal patterns, as we suggest. Furthermore, neuronal recording and fMRI are sensitive to different and complementary aspects of neuronal computation. Neuronal recording measures action potentials from at most a small number of neurons near the electrode tip and is most sensitive to robust changes in the discharge patterns in these few neurons. In contrast, Logothetis et al. (2001) have suggested that dendritic potentials might contribute more significantly to BOLD signal changes than action potentials, and Scannell and Young (1999) contend that small changes in firing rate across a large number of neurons are likely to affect the BOLD signal much more strongly than large firing rate changes in a small number of neurons.

2.4.3 – Inhibitory Processes Outside the 'Classical' Cortical Saccade System

We found three regions potentially involved in actively inhibiting automatic saccades. Right superior frontal sulcus (SFS), right supramarginal gyrus (SMG), and posterior cingulate sulcus (PCS) exhibited significantly greater *nogo_response* versus *pro_response* activation and (non-significantly) greater *anti_response* versus *pro_response* activation. *Anti_response* activation in SFS and SMG was between *pro_response* and *nogo_response* activity, while *anti_response* activation in PCS was slightly larger than *nogo_response* activation.

Condy and colleagues (2007) found that reversible inactivation of a focal region in the inferior bank of the posterior principal sulcus in monkeys impaired inhibition of the automatic saccade in the antisaccade task. We suggest that this posterior principal sulcus region might be the monkey homologue of the human SFS region in which we found greater *nogo_response* versus *pro_response* activation. This suggestion is consistent with the case made by Courtney and colleagues (1998) that monkey principal sulcus is homologous to human superior frontal sulcus.

Picton and colleagues (2007) found that 11 human patients with superior medial frontal lesions made more erroneous go responses on nogo trials in a go/nogo task than controls. The lesions of 5 of the 11 patients (see Alexander et al. 2005) included right SFS, in agreement with our suggestion that right SFS is involved in behavioural inhibition. In addition, one patient's lesion included left SFS.

It is possible that correct antisaccade performance might recruit inhibitory processes at a submaximal level so as to inhibit the automatic saccade without preventing voluntary antisaccade generation. This would explain activation patterns in right SFS and right SMG. Alternately, the antisaccade task might put greater demands on inhibitory processes than the nogo task, if it is, in fact, more difficult to inhibit an automatic saccade while preparing to generate an antisaccade, compared to inhibiting a saccade while maintaining fixation. Antisaccade responses would then elicit more fMRI activation in inhibitory regions than nogo responses, which would explain PCS activation patterns.

We can exclude the possibility that these three regions might be part of a resting state network (Raichle et al. 2001). Given that the antisaccade task was the most difficult of the three tasks, based on performance data, a resting network should exhibit less antisaccade compared to prosaccade activation, which was not the case here.

2.4.4 – Comparison with Previous fMRI Studies of Behavioural Inhibition

Previous fMRI studies found that tasks designed to recruit behavioural inhibition evoked activation in diverse frontal and parietal regions, including middle and inferior frontal gyri, intraparietal sulcus, supramarginal gyrus, and anterior cingulate (Asahi et al. 2004; Braver et al. 2001; Buchsbaum et al. 2005; de Zubicaray et al. 2000; Fassbender et al. 2004; Garavan et al. 1999; Garavan et al. 2002; Horn et al. 2003; Kelly et al. 2004; Maguire et al. 2003; Menon et al. 2001; Mostofsky et al. 2003; Rubia et al. 2001; Rubia et al. 2005; Sylvester et al. 2003; Wager et al. 2005). With the exception of right supramarginal gyrus, we did not find greater *nogo_response* activation in these regions in our *pro_response – nogo_response* comparison. One possible explanation for this is that in many of the behavioural inhibition studies, it was difficult to separate processes underlying behavioural inhibition from those underlying task switching. Four studies compared blocks of go trials with blocks of interleaved go and nogo trials (Asahi et al. 2004; Maguire et al. 2003; Menon et al. 2001; Rubia et al. 2001), in which case only the interleaved trial blocks involved task switching. Seven studies embedded infrequent nogo trials in sequences of rapidly-presented, frequent go trials (de Zubicaray et al. 2000; Fassbender et al. 2004; Garavan et al. 1999; Garavan et al. 2002; Horn et al. 2003; Kelly et al. 2004; Rubia et al. 2005). This design was intended to increase the prepotency of the

go response, and thereby increase the requirement to inhibit the go response in the nogo trials, but it also introduced an asymmetry in the switch requirements between the rare nogo trials, most of which would have involved switching from previous go trials, and the frequent go trials, most of which would have been preceded by go trials and would not have required task switching. Importantly, a meta-analysis by Buchsbaum and colleagues (2005) and a study by Sylvester and colleagues (2003) found overlap between activation related to behavioural inhibition and task switching in intraparietal sulcus, parts of the middle and inferior frontal gyri, and ventral supplementary motor area and dorsal anterior cingulate gyrus. In contrast, our design included prosaccade, antisaccade, and nogo trials with equal frequencies, and trials were compound, with distinct task instruction and response phases, such that response-related results for our *pro_response – nogo_response* comparison should be independent of activation evoked by task switching during the trial type instruction phase.

Wager and colleagues (2005) compared rapid, interleaved go and nogo trials under a rare nogo condition (20% nogo trials) and an equiprobable nogo condition (50% nogo trials). However, only results collapsed across the rare and equiprobable conditions are presented, introducing the issue of task switching as discussed above. Wager and colleagues (2005) also used two blocked design experiments with different inhibitory requirements, a two alternative forced choice response task with compatible and incompatible flanker stimuli and a stimulus-response compatibility task. They found that all three tasks evoked significantly greater activation in their respective inhibition conditions in anterior cingulate and bilateral caudate. When correction for multiple comparisons was not performed, they also found greater inhibition-related activity in bilateral anterior insula, right middle frontal gyrus, and several frontal and parietal regions consisting of single functional voxels.

Braver and colleagues (2001) compared go and nogo trials interleaved in rapid sequences under three conditions, rare nogo (17% nogo), equiprobable (50% nogo), and rare go (17% go). They found regions in anterior cingulate, bilateral anterior insula, and right middle frontal gyrus exhibiting greater activation on both the rare nogo and rare go trials, and they attributed these results to response conflict as opposed to response inhibition. However, as discussed above, the rare nogo and rare go trials would have been

predominantly switch trials while the frequent go and nogo trials would have been predominantly repeat (ie. non-switch) trials, and these activation patterns might be related to task switching. Braver and colleagues (2001) also found regions in right middle frontal gyrus, right ventrolateral prefrontal cortex, right posterior prefrontal cortex, right anterior cingulate, and right supplementary motor area, as well as several predominantly right parietal regions all exhibiting activation specific to behavioural inhibition, that is, greater activation on rare nogo trials versus frequent go trials. Rare go trials did not evoke more activation than frequent nogo trials, indicating that these regions were not concerned simply with task switching independent of task type. In addition, comparison of nogo versus go trials in the equiprobable go/nogo task did not reveal activation differences in these regions. One possibility is that in the equiprobable condition, the go response was not prepotent (see discussion below) and that the nogo trials did not require much behavioural inhibition.

We did not find evidence for saccadic inhibition processes in those regions implicated in behavioural inhibition by the two studies discussed above (Braver et al. 2001; Wager et al. 2005). One possibility is that the mechanism of behavioural inhibition might be specific to the type of process being inhibited. Those two studies used a button press task in which subjects were presented letters and pressed a button for non X stimuli and withheld the button press for X stimuli. Pressing a button in response to non X letters is not intrinsically prepotent, given that most of us read thousands of words each day without pressing buttons in response to each letter read. The go response in rare nogo versions of the button press task is made prepotent by virtue of its high frequency in comparison with the nogo task's low frequency. In contrast, our nogo task involved inhibiting the intrinsically prepotent automatic saccade to a peripheral visual stimulus. It is possible that inhibiting an intrinsically prepotent automatic saccade is performed by a different mechanism than inhibiting a button press response made prepotent by virtue of that response type's recent, frequent repetition. This suggestion is of course very speculative, and more research is necessary to explore the issue fully.

2.4.5 – Antisaccade Task Activation

In FEF, SEF, ACC, and IPS, *anti_response* activation was greater than either *pro_response* or *nogo_response* activation. In middle frontal gyrus (MFG), also called dorsolateral prefrontal cortex, *anti_response* activation was significantly greater than *nogo_response* activation. In addition, about half the voxels that exhibited greater *anti_response* versus *nogo_response* activation in MFG exhibited greater *anti_response* versus *pro_response* activation at a significance of $p < 0.05$ not corrected for multiple comparisons (see *Section 2.3.3 – Response Period Pair Wise Contrasts* for details). This difference is consistent with two previous fMRI studies that showed greater activation for antisaccades compared to prosaccades in MFG (DeSouza et al. 2003; Ford et al. 2005).

How might we explain greater BOLD signal for antisaccade responses given the suggestion discussed above that saccade generation and inhibition do not determine BOLD signal levels in these regions? The antisaccade task required visuospatial remapping of the peripheral stimulus into the opposite visual hemifield, unlike the prosaccade and nogo tasks, and this additional process might account for greater *anti_response* activation. The antisaccade task was also more difficult than either the prosaccade or nogo tasks given the antisaccade task's conflicting requirement to inhibit the automatic saccade to the stimulus while preparing a voluntary antisaccade. In contrast, subjects simply generated visually-guided prosaccades or inhibited automatic saccades in the nogo task. We also observed a higher proportion of erroneous prosaccade responses on antisaccade trials compared to nogo trials, indicating that the antisaccade task was more difficult and might have required more attention to perform correctly, resulting in greater *anti_response* activation.

2.4.6 – Rapid versus Widely-spaced Event-related Designs

Previous fMRI studies comparing prosaccades and antisaccades (Connolly et al. 2002; Curtis and D'Esposito 2003; DeSouza et al. 2003; Ford et al. 2005) used widely-spaced event-related designs with inter-trial intervals lasting approximately 12 s and long instruction intervals within the trials. This was done to compensate for haemodynamic lag but introduced the possibility that neuronal processes underlying prosaccades and antisaccades were altered by the long timing regimen. We can exclude this possibility

because our results, acquired using a rapid event-related design with a 3 s mean inter-event interval, agreed substantially with previous findings.

2.4.7 – 'Amount' of Inhibition Recruited by the Nogo Task

The nogo task included a 200 ms gap between fixation point offset and peripheral stimulus onset. A 200 to 300 ms gap in visually guided saccade trials causes fixation and saccade neurons in the superior colliculus to decrease and increase their firing rates, respectively (Dorris and Munoz 1995; Dorris et al. 1997; Everling et al. 1998). The gap also increases the proportion of erroneous prosaccades made on the antisaccade task (Forbes and Klein 1996; Fischer and Weber 1997; Bell et al. 2000). The 200 ms gap in our nogo task made it more difficult for subjects to inhibit the automatic saccadic response, presumably forcing them to employ active saccadic inhibition to perform the task correctly. It is unlikely that subjects, once informed that a given trial was a nogo trial, could simply "coast" until the next trial, while still performing the nogo trial correctly. That nogo responses evoked BOLD activation significantly above baseline in cortical saccade regions further supports this argument. Finally, if subjects did in fact coast through the nogo trials, we would expect nogo trial activation to be less than that evoked by prosaccade trials, which was not the case here.

2.4.8 – Instruction-related Activation

We found instruction-related activations in intraparietal sulcus and several prefrontal regions (see *Section 2.3 – Results*, *Figure 2.7*, *Table 2.3*). This activity represents some combination of processes underlying interpretation of the trial type instruction (coloured fixation point) based on the visuo-motor conditional associations relating colour to task type as well as processes underlying loading of relevant task instructions and preparation to perform the upcoming task. Previous studies have found evidence for both sets of processes in similar regions throughout the cortical saccade system (Amiez et al. 2006; Deiber et al. 1997; Grafton et al. 1998; Grosbras et al. 2005; Toni and Passingham 1999).

We did not find instruction-related differences between prosaccades and antisaccades unlike some previous experiments. Three studies used long time intervals

between trial type instruction presentation and peripheral stimulus onset (Curtis and D'Esposito 2003 - 7 s; DeSouza et al. 2003 - 6, 10, and 14 s; Ford et al. 2005 - 10 s). They observed bimodal activation time courses with the first mode evoked by instruction onset and the second mode evoked by peripheral stimulus onset and response execution. Tonic activation persisted between the modes, and it was this tonic activity that was larger for antisaccades compared to prosaccades. The initial phasic response to trial instruction onset did not differentiate between trial types, which was also true for the phasic BOLD responses we observed to instruction onset. Our design employed an instruction period of 2, 3, or 4 s for the vast majority of trials, with a maximum of 7 s in a small minority of trials. Thus, we cannot compare our results directly with previous tonic activation data over long instruction periods. Connolly and colleagues (2002), using instruction intervals 0, 2, and 4 s long, did find greater activation in FEF for antisaccades compared to prosaccades during the rising component of the phasic BOLD response to instruction onset. They did not present a fixation point during their instruction interval, unlike other fMRI antisaccades studies, and this might explain this result.

2.4.9 – Summary

We compared fMRI activation patterns for prosaccades, antisaccades, and nogo trials in humans. BOLD activation patterns in cortical saccade regions might predominantly reflect visual detection and attention rather than saccadic motor processes or saccadic inhibition. However, right SFS, right SMG, and PCS might play a role in active saccadic inhibition. Antisaccade BOLD activation in cortical saccade regions could reflect visuospatial remapping or attention related to antisaccade task difficulty. The findings from our rapid event-related comparison of prosaccades and antisaccades also agreed substantially with previous widely-spaced event-related results.

2.5 – Bibliography

Alexander MP, Stuss DT, Shallice T, Picton TW, and Gillingham S. 2005. Impaired concentration due to frontal lobe damage from two distinct lesion sites. *Neurology* 65(4): 572-579.

Amador N, Schlag-Rey M, and Schlag J. 2004. Primate antisaccade. II. Supplementary eye field neuronal activity predicts correct performance. *J Neurophysiol* 91(4): 1672-1689.

Amiez C, Kostopoulos P, Champod AS, and Petrides M. 2006. Local morphology predicts functional organization of the dorsal premotor region in the human brain. *J Neurosci* 26(10): 2724-2731.

Anderson TJ, Jenkins IH, Brooks DJ, Hawken MB, Frackowiak RS, and Kennard C. 1994. Cortical control of saccades and fixation in man. A PET study. *Brain* 117 (Pt 5): 1073-1084.

Asahi S, Okamoto Y, Okada G, Yamawaki S, and Yokota N. 2004. Negative correlation between right prefrontal activity during response inhibition and impulsiveness: a fMRI study. *Eur Arch Psychiatry Clin Neurosci* 254(4): 245-251.

Barberi EA, Gati JS, Rutt BK, and Menon RS. 2000. A transmit-only/receive-only (TORO) RF system for high-field MRI/MRS applications. *Magn Reson Med* 43(2): 284-289.

Barton JJ, Raoof M, Jameel O, and Manoach DS. 2006. Task-switching with antisaccades versus no-go trials: a comparison of inter-trial effects. *Exp Brain Res* 172(1): 114-119.

Bell AH, Everling S, and Munoz DP. 2000. Influence of stimulus eccentricity and direction on characteristics of pro- and antisaccades in non-human primates. *J Neurophysiol* 84(5): 2595-2604.

Berman RA, Colby CL, Genovese CR, Voyvodic JT, Luna B, Thulborn KR, and Sweeney JA. 1999. Cortical networks subserving pursuit and saccadic eye movements in humans: an FMRI study. *Hum Brain Mapp* 8(4): 209-225.

Braver TS, Barch DM, Gray JR, Molfese DL, and Snyder A. 2001. Anterior cingulate cortex and response conflict: effects of frequency, inhibition and errors. *Cereb Cortex* 11(9): 825-836.

Brown MRG, DeSouza JFX, Goltz HC, Ford K, Menon RS, Goodale MA, and Everling S. 2004. Comparison of memory- and visually guided saccades using event-related fMRI. *J Neurophysiol* 91(2): 873-889.

Buchsbaum BR, Greer S, Chang WL, and Berman KF. 2005. Meta-analysis of neuroimaging studies of the Wisconsin card-sorting task and component processes. *Hum Brain Mapp* 25(1): 35-45.

Bullmore E, Long C, Suckling J, Fadili J, Calvert G, Zelaya F, Carpenter TA, and Brammer M. 2001. Colored noise and computational inference in neurophysiological (fMRI) time series analysis: resampling methods in time and wavelet domains. *Hum Brain Mapp* 12(2): 61-78.

Condy C, Wattiez N, Rivaud-Pechoux S, Tremblay L, and Gaymard B. 2007. Antisaccade deficit after inactivation of the principal sulcus in monkeys. *Cereb Cortex* 17(1): 221-229.

Connolly JD, Goodale MA, Menon RS, and Munoz DP. 2002. Human fMRI evidence for the neural correlates of preparatory set. *Nat Neurosci* 5(12): 1345-1352.

Corbetta M, Akbudak E, Conturo TE, Snyder AZ, Ollinger JM, Drury HA, Linenweber MR, Petersen SE, Raichle ME, Van Essen DC, and Shulman GL. 1998. A common network of functional areas for attention and eye movements. *Neuron* 21(4): 761-773.

Corbetta M and Shulman GL. 2002. Control of goal-directed and stimulus-driven attention in the brain. *Nat Rev Neurosci* 3(3): 201-215.

Courtney SM, Petit L, Maisog JM, Ungerleider LG, and Haxby JV. 1998. An area specialized for spatial working memory in human frontal cortex. *Science* 279(5355): 1347-1351.

Curtis CE and D'Esposito M. 2003. Success and failure suppressing reflexive behavior. *J Cogn Neurosci* 15(3): 409-418.

de Zubicaray GI, Andrew C, Zelaya FO, Williams SC, and Dumanoir C. 2000. Motor response suppression and the prepotent tendency to respond: a parametric fMRI study. *Neuropsychologia* 38(9): 1280-1291.

Deiber MP, Wise SP, Honda M, Catalan MJ, Grafman J, and Hallett M. 1997. Frontal and parietal networks for conditional motor learning: a positron emission tomography study. *J Neurophysiol* 78(2): 977-991.

DeSouza JFX, Menon RS, and Everling S. 2003. Preparatory set associated with pro-saccades and anti-saccades in humans investigated with event-related FMRI. *J Neurophysiol* 89(2): 1016-1023.

Dorris MC and Munoz DP. 1995. A neural correlate for the gap effect on saccadic reaction times in monkey. *J Neurophysiol* 73(6): 2558-2562.

Dorris MC, Pare M, and Munoz DP. 1997. Neuronal activity in monkey superior colliculus related to the initiation of saccadic eye movements. *J Neurosci* 17(21): 8566-8579.

Everling S and Fischer B. 1998. The antisaccade: a review of basic research and clinical studies. *Neuropsychologia* 36(9): 885-899.

Everling S, Pare M, Dorris MC, and Munoz DP. 1998. Comparison of the discharge characteristics of brain stem omnipause neurons and superior colliculus fixation neurons in monkey: implications for control of fixation and saccade behavior. *J Neurophysiol* 79(2): 511-528.

Fassbender C, Murphy K, Foxe JJ, Wylie GR, Javitt DC, Robertson IH, and Garavan H. 2004. A topography of executive functions and their interactions revealed by functional magnetic resonance imaging. *Brain Res Cogn Brain Res* 20(2): 132-143.

Fischer B and Weber H. 1997. Effects of stimulus conditions on the performance of antisaccades in man. *Exp Brain Res* 116(2): 191-200.

Forbes K and Klein RM. 1996. The magnitude of the fixation offset effect with endogenously and exogenously controlled saccades. *J Cogn Neurosci* 8: 344-352.

Ford KA, Goltz HC, Brown MRG, and Everling S. 2005. Neural processes associated with antisaccade task performance investigated with event-related FMRI. *J Neurophysiol* 94(1): 429-440.

Friston KJ, Fletcher P, Josephs O, Holmes A, Rugg MD, and Turner R. 1998. Event-related fMRI: characterizing differential responses. *NeuroImage* 7(1): 30-40.

Frith CD. 1992. *The cognitive neuropsychology of schizophrenia*. Hove, UK: Lawrence Erlbaum Associates.

Fuster JM. 1989. *The Prefrontal Cortex: Anatomy, Physiology, and Neuropsychology of the Frontal Lobe*. New York: Raven.

Garavan H, Ross TJ, Murphy K, Roche RA, and Stein EA. 2002. Dissociable executive functions in the dynamic control of behavior: inhibition, error detection, and correction. *NeuroImage* 17(4): 1820-1829.

Garavan H, Ross TJ, and Stein EA. 1999. Right hemispheric dominance of inhibitory control: an event-related functional MRI study. *Proc Natl Acad Sci U S A* 96(14): 8301-8306.

Grafton ST, Fagg AH, and Arbib MA. 1998. Dorsal premotor cortex and conditional movement selection: A PET functional mapping study. *J Neurophysiol* 79(2): 1092-1097.

Grosbras MH, Laird AR, and Paus T. 2005. Cortical regions involved in eye movements, shifts of attention, and gaze perception. *Hum Brain Mapp* 25(1): 140-154.

Hallett PE. 1978. Primary and secondary saccades to goals defined by instructions. *Vision Res* 18(10): 1279-1296.

Horn NR, Dolan M, Elliott R, Deakin JF, and Woodruff PW. 2003. Response inhibition and impulsivity: an fMRI study. *Neuropsychologia* 41(14): 1959-1966.

Jonides J. 1981. Voluntary vs. automatic control over the mind's eye's movement. In: *Attention and Performance IX*, edited by Long JB and Baddeley AD. Hillsdale, NJ: Lawrence Erlbaum Associates, p.187-203.

Kelly AM, Hester R, Murphy K, Javitt DC, Foxe JJ, and Garavan H. 2004. Prefrontal-subcortical dissociations underlying inhibitory control revealed by event-related fMRI. *Eur J Neurosci* 19(11): 3105-3112.

Kimmig H, Greenlee MW, Gondan M, Schira M, Kassubek J, and Mergner T. 2001. Relationship between saccadic eye movements and cortical activity as measured by fMRI: quantitative and qualitative aspects. *Exp Brain Res* 141(2): 184-194.

Klassen LM and Menon RS. 2004. Robust automated shimming technique using arbitrary mapping acquisition parameters (RASTAMAP). *Magn Reson Med* 51(5): 881-887.

Logothetis NK, Pauls J, Augath M, Trinath T, and Oeltermann A. 2001. Neurophysiological investigation of the basis of the fMRI signal. *Nature* 412(6843): 150-157.

Luna B and Sweeney JA. 1999. Cognitive functional magnetic resonance imaging at very-high-field: eye movement control. *Top Magn Reson Imaging* 10(1): 3-15.

Luria AR. 1966. *Higher Cortical Functions in Man*. London: Tavistock.

Maguire RP, Broerse A, de Jong BM, Cornelissen FW, Meiners LC, Leenders KL, and den Boer JA. 2003. Evidence of enhancement of spatial attention during inhibition of a visuo-motor response. *NeuroImage* 20(2): 1339-1345.

Marchini JL and Ripley BD. 2000. A new statistical approach to detecting significant activation in functional MRI. *NeuroImage* 12(4): 366-380.

Menon V, Adleman NE, White CD, Glover GH, and Reiss AL. 2001. Error-related brain activation during a Go/NoGo response inhibition task. *Hum Brain Mapp* 12(3): 131-143.

Mostofsky SH, Schafer JG, Abrams MT, Goldberg MC, Flower AA, Boyce A, Courtney SM, Calhoun VD, Kraut MA, Denckla MB, and Pekar JJ. 2003. fMRI evidence that the neural basis of response inhibition is task-dependent. *Brain Res Cogn Brain Res* 17(2): 419-430.

Munoz DP and Everling S. 2004. Look away: the anti-saccade task and the voluntary control of eye movement. *Nat Rev Neurosci* 5(3): 218-228.

Pare M and Wurtz RH. 2001. Progression in neuronal processing for saccadic eye movements from parietal cortex area lip to superior colliculus. *J Neurophysiol* 85(6): 2545-2562.

Picton TW, Stuss DT, Alexander MP, Shallice T, Binns MA, and Gillingham S. 2007. Effects of focal frontal lesions on response inhibition. *Cereb Cortex* 17(4): 826-838.

Raichle ME, MacLeod AM, Snyder AZ, Powers WJ, Gusnard DA, and Shulman GL. 2001. A default mode of brain function. *Proc Natl Acad Sci* 98(2): 676-682.

Rizzolatti G, Riggio L, Dascola I, and Umilta C. 1987. Reorienting attention across the horizontal and vertical meridians: evidence in favor of a premotor theory of attention. *Neuropsychologia* 25(1A): 31-40.

Rubia K, Lee F, Cleare AJ, Tunstall N, Fu CH, Brammer M, and McGuire P. 2005. Tryptophan depletion reduces right inferior prefrontal activation during response inhibition in fast, event-related fMRI. *Psychopharmacology (Berl)* 179(4): 791-803.

Rubia K, Russell T, Overmeyer S, Brammer MJ, Bullmore ET, Sharma T, Simmons A, Williams SC, Giampietro V, Andrew CM, and Taylor E. 2001. Mapping motor inhibition: conjunctive brain activations across different versions of go/no-go and stop tasks. *NeuroImage* 13(2): 250-261.

Sato TR and Schall JD. 2003. Effects of stimulus-response compatibility on neural selection in frontal eye field. *Neuron* 38(4): 637-648.

Scannell JW and Young MP. 1999. Neuronal population activity and functional imaging. *Proc Biol Sci* 266(1422): 875-881.

Schall JD. 2002. The neural selection and control of saccades by the frontal eye field. *Philos Trans R Soc Lond B Biol Sci* 357(1424): 1073-1082.

Schall JD. 2004. On the role of frontal eye field in guiding attention and saccades. *Vision Res* 44(12): 1453-1467.

Schlag-Rey M, Amador N, Sanchez H, and Schlag J. 1997. Antisaccade performance predicted by neuronal activity in the supplementary eye field. *Nature* 390(6658): 398-401.

Serences JT. 2004. A comparison of methods for characterizing the event-related BOLD timeseries in rapid fMRI. *NeuroImage* 21(4): 1690-1700.

Sommer MA and Wurtz RH. 2001. Frontal eye field sends delay activity related to movement, memory, and vision to the superior colliculus. *J Neurophysiol* 85(4): 1673-1685.

Sweeney JA, Mintun MA, Kwee S, Wiseman MB, Brown DL, Rosenberg DR, and Carl JR. 1996. Positron emission tomography study of voluntary saccadic eye movements and spatial working memory. *J Neurophysiol* 75(1): 454-468.

Sylvester CY, Wager TD, Lacey SC, Hernandez L, Nichols TE, Smith EE, and Jonides J. 2003. Switching attention and resolving interference: fMRI measures of executive functions. *Neuropsychologia* 41(3): 357-370.

Thompson KG, Bichot NP, and Schall JD. 1997. Dissociation of visual discrimination from saccade programming in macaque frontal eye field. *J Neurophysiol* 77(2): 1046-1050.

Thompson KG, Biscoe KL, and Sato TR. 2005. Neuronal basis of covert spatial attention in the frontal eye field. *J Neurosci* 25(41): 9479-9487.

Toni I and Passingham RE. 1999. Prefrontal-basal ganglia pathways are involved in the learning of arbitrary visuomotor associations: a PET study. *Exp Brain Res* 127(1): 19-32.

Wager TD, Sylvester CY, Lacey SC, Nee DE, Franklin M, and Jonides J. 2005. Common and unique components of response inhibition revealed by fMRI. *NeuroImage* 27(2): 323-340.

Worsley KJ, Liao CH, Aston J, Petre V, Duncan GH, Morales F, and Evans AC. 2002. A general statistical analysis for fMRI data. *NeuroImage* 15(1): 1-15.

Wurtz RH and Goldberg ME. 1989. *The Neurobiology of Saccadic Eye Movements*. Amsterdam: Elsevier.

Yantis S and Jonides J. 1984. Abrupt visual onsets and selective attention: evidence from visual search. *J Exp Psychol Hum Percept Perform* 10(5): 601-621.

Chapter 3
Frontoparietal Activation with Preparation for Antisaccades

The material in Chapter 3 has been published as MR Brown, T Vilis, and S Everling, 2007. Frontoparietal activation with preparation for antisaccades. *Journal of Neurophysiology* 98(3): 1751-1762. The material therein is used here with permission (see *Appendix 2*).

3.1 – Introduction

Primates are not constrained to react to sensory stimuli with reflexive movements, but they rather can acquire almost arbitrary stimulus-response associations. The ability to execute different responses to identical stimuli has been attributed to differences in preparatory set, i.e. the intention and readiness to perform a certain task (Evarts et al. 1984; Hebb 1972, p. 77-93). Examples of two tasks that require very different preparatory sets are the prosaccade and antisaccade tasks (Hallett 1978; Hallett and Adams 1980; Everling and Fischer 1998). The prosaccade task requires subjects to look towards a flashed peripheral stimulus, whereas the antisaccade task requires subjects to suppress the automatic saccade to the stimulus and instead to look away from the stimulus to its mirror location in the opposite visual hemifield. To correctly perform this task, the neural process that triggers the automatic prepotent response to look towards the stimulus must be suppressed so that the vector inversion for the antisaccade can be computed and the saccade can be executed. Single neuron recordings in nonhuman primates have demonstrated that the brain accomplishes this task by reducing the level of preparatory saccade-related activity in the superior colliculus prior to stimulus presentation on antisaccade trials (Munoz and Everling 2004). Neural correlates for different preparatory sets associated with prosaccades and antisaccades have been identified in several frontal cortical brain areas in nonhuman primates, including the frontal eye field (FEF) (Everling and Munoz 2000), supplementary eye field (SEF) (Amador et al. 2003), dorsolateral prefrontal cortex (DLPFC) (Everling and DeSouza 2005; Johnston and Everling 2006a, b), and anterior cingulate cortex (ACC) (Johnston et al. 2007). Single neuron recording has also been used to compare prosaccades and antisaccades in the lateral intraparietal area (LIP) (Gottlieb and Goldberg 1999; Zhang

and Barash 2000, 2004), but these studies did not investigate preparatory differences between the two tasks. Specifically, they did not compare neuronal discharge patterns between prosaccades and antisaccades during the preparatory period in which monkey subjects had been given the trial type instruction cue but had not seen a peripheral stimulus or responded to it with an eye movement.

Studies recording event-related potentials (ERP) in human subjects have found a higher negative potential prior to antisaccades than prosaccades over a wide range of recording sites, supporting the idea that cortical areas in humans exhibit neural correlates for different preparatory sets associated with prosaccades and antisaccades (Everling et al. 1997; Klein et al. 2000). Connolly and colleagues directly compared the activation in the frontal eye field and intraparietal sulcus between prosaccades and antisaccades by using an event-related functional magnetic resonance imaging (fMRI) paradigm (Connolly et al. 2002). Their paradigm included a gap of 0, 2, or 4 s, during which the subject viewed a black screen, between the disappearance of the central fixation point and peripheral stimulus onset. The authors found greater preparatory activation for antisaccades compared to prosaccades in FEF but not in intraparietal sulcus (IPS). Based on this finding they concluded that the human FEF, but not the IPS, is critically involved in preparatory set. The absence of any parietal activation during the preparatory period for antisaccades is somewhat surprising given the activation of a frontoparietal network with many different kinds of tasks (Cabeza and Nyberg 2000; Duncan and Owen 2000).

A potential problem with the analysis carried out by Connolly and colleagues is that they shifted the time courses upon which they performed their analysis by 3 s to account for haemodynamic lag. Therefore, the generation of the saccade, itself, might have contributed to apparent preparatory differences. In order to avoid this problem, several studies have employed event-related fMRI design with compound trials consisting of a long preparatory period (6-14 s) and a response period (Curtis and D'Esposito 2003; Desouza et al. 2003; Ford et al. 2005). Although not consistent, together these studies have reported a higher activation for antisaccades than prosaccades in a variety of cortical areas during the preparatory period, including DLPFC, FEF, ACC, pre-supplementary motor area (SMA), IPS, and parieto-occipital sulcus.

A direct comparison between the results from these designs with very long preparatory periods (>6 s) and the results from electrophysiological studies in nonhuman primates and ERP studies in humans that usually have short preparatory periods (<2 s) is potentially difficult due to additional processes that might be initiated during long delays. In addition, the long intertrial intervals in these event-related fMRI studies (12 or 14 s) likely lead to fatigue and might reduce task-related preparatory activations.

In a rapid event-related fMRI comparison of saccades and countermanded saccades, (Curtis et al. 2005) recently showed that catch trials, in which subjects neither saw peripheral stimuli nor made saccade responses, evoked significant frontoparietal fMRI activation. The authors attributed this activation to covert attention. It is not yet known whether catch trial-evoked activation might be modulated by task type. Here, we compared prosaccades and antisaccades using a rapid event-related fMRI design that included both whole trials and catch trials, which we will call half trials. Whole prosaccade and antisaccade trials included a 1 s preparatory period, in which a coloured fixation point indicated the trial type, followed by a response period, in which subjects made a prosaccade toward or an antisaccade away from a flashed peripheral stimulus. Half trials included only the preparatory period (Ollinger et al. 2001a; Ollinger et al. 2001b). Subjects did not know whether a given trial would be a whole or half trial until the peripheral stimulus appeared or failed to appear, respectively. Our results show that the preparatory activation for an antisaccade, even with no executed saccade, is greater than the preparatory activation for a prosaccades in frontoparietal regions of the human brain.

3.2 – Methods

All procedures were approved by the University Research Ethics Board for Health Sciences Research at the University of Western Ontario, London, Ontario, Canada and are in accordance with the 1964 Declaration of Helsinki.

3.2.1 – Subjects

Eleven human subjects (9 female, 2 male) participated in this study after giving informed written consent. Subject age ranged from 20 to 28 years with a mean of 25 years. Subjects reported no history of neurological or psychiatric disorder, and they had normal or corrected to normal vision. Two subjects described themselves as left-handed while the rest described themselves as right-handed.

3.2.2 – fMRI Data Acquisition Procedure

All fMRI data were acquired on a whole body 4 Tesla MRI system (Varian, Palo Alto, CA; Siemens, Erlangen, Germany) operating at a slew rate of 120 T/m/s and 40 mT/m gradients. A transmit only receive only (TORO) cylindrical hybrid birdcage radio frequency (RF) head coil (Barberi et al. 2000) was used for transmission and detection of the signal.

Imaging planes for the functional scans were defined based on a series of sagittal anatomical images acquired using T1-weighting. Eleven contiguous functional planes, 5 mm thick, were prescribed axially from the top of the brain to the level of the dorsal caudate and thalamus. A constrained, three-dimensional phase shimming procedure (Klassen and Menon 2004) was employed to optimize the magnetic field homogeneity over the functional volume. During each functional task, blood oxygenation level-dependent (BOLD) images (T_2*-weighted) were acquired continuously using an interleaved, two segment, optimized spiral imaging protocol (volume collection time = 1.0 s, repeat time (TR) = 500 ms, echo time (TE) = 15 ms, flip angle = 30º, 64 x 64 matrix size, 22.0 x 22.0 cm field of view (FOV), 3.44 x 3.44 x 5.00 mm voxel resolution). Each image was corrected for physiological fluctuations using navigator echo correction. A corresponding high-resolution T1-weighted structural volume was acquired during the same scanning session using a 3D spiral acquisition protocol (TE = 3.0 ms, inversion time (TI) = 1300 ms, TR = 50 ms) with a voxel resolution of 0.9 x 0.9 x 1.25 mm. Each subject was immobilized during the experimental session within a head cradle, which was packed with foam padding.

3.2.3 – Eye Tracking

Visual stimuli were presented during fMRI scanning using a Silent Vision SV-4021 projection system from Avotec (Stuart, Fl, USA). The system includes an MEyeTrack-SV (Silent Vision) eye tracker from SensoMotoric Instruments GmbH (Teltow, Germany). This setup employs fibre optics housed in dual stalks positioned over the subject's eyes to allow simultaneous presentation of visual stimuli and CCD video-based infrared eye tracking. The visual display subtends 30° horizontally by 23° vertically with a resolution of 800 x 600 pixels and a refresh rate of 60 Hz. Eye tracking was also performed at a 60 Hz sampling rate with an accuracy of approximately 1 degree. Before scanning, we calibrated the system with a 9 point calibration, with four points in the corners, four points at the midpoints of the sides, and one point in the centre of the visual display. All experimental targets were within the range of the calibration. Analysis of the eye movement signals was performed offline.

3.2.4 – Experimental Design

This experiment compared the prosaccade and antisaccade tasks in a rapid event-related fMRI design (*Figure 3.1*). Two thirds of prosaccade and antisaccade trials were whole trials, while the remaining third were half trials. When not performing a task response, the subject fixated a central white dot (0.6° diameter). An individual trial started when the central dot changed from white to magenta or cyan, where colour indicated whether the current trial was a prosaccade or antisaccade trial. For six of the subjects, cyan and magenta mapped, respectively, onto the antisaccade and prosaccade tasks, while for the other five subjects, the mapping was reversed. The coloured fixation point was presented for 1 sec, followed by a 200 ms gap, during which the subject was presented with a black screen. The 200 ms gap was included because it is known to decrease the fixation-related activity and increase the saccade-related activity in the superior colliculus (Dorris et al. 1997; Everling et al. 1999), resulting in a larger proportion of erroneous prosaccade responses on antisaccade trials (Fischer and Weber 1992; Forbes and Klein 1996; Fischer and Weber 1997; Bell et al. 2000) and an increased demand for cognitive control prior to stimulus onset (Munoz and Everling 2004).

Whole Trial
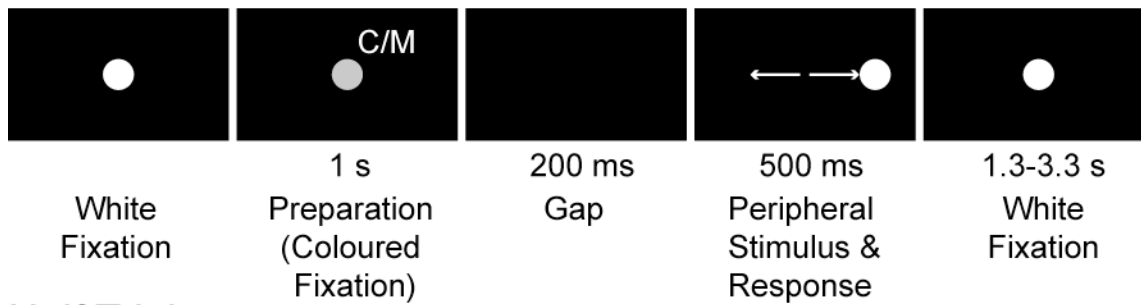
Half Trial
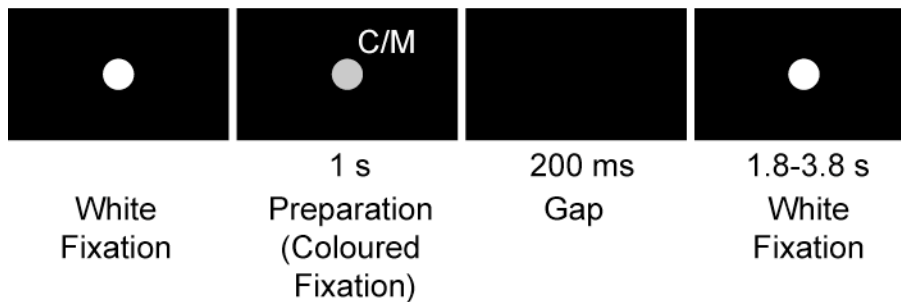
Figure 3.1 – Experimental Design

Schematic respresentation of timing and visual stimulus presentation for prosaccade and antisaccade whole and half trials. The gray fixation point and "C/M" in the preparation panel (second from left) denote that the white fixation point changed to magenta or cyan to instruct the subject whether to make a prosaccade or antisaccade upon stimulus presentation. Colour-trial type assignment depended on the subject. The small white arrows in the fourth panel for the whole trial indicate a rightward prosaccade response and a leftward antisaccade response. Whereas whole trials contained both the preparation and response events, half trials contained only the trial type preparation. On half trials, subjects never saw a peripheral stimulus nor made an eye movement. The half trials allowed us to measure preparatory activation separately from response-related activation. See *Section 3.2 – Methods* for details.

On whole trials, a peripheral stimulus (0.6° diameter white dot) was presented at the end of the 200 ms gap. The peripheral stimulus could appear to the left or right of fixation, at eccentricities of 6° or 10°, either on the horizontal meridian or 20° clockwise or counterclockwise relative it. The peripheral stimulus was presented for 500 ms, after

which it disappeared simultaneously with the reappearance of the central white fixation dot. Before the peripheral stimulus disappeared, the subject had to look at it on prosaccade trials or look away from it, toward its mirror location in the opposite visual quadrant, on antisaccade trials. The subject had then to return gaze to centre and maintain central fixation on the newly reappeared central white fixation dot. On half trials, the central fixation dot reappeared at the end of the 200 ms gap; no peripheral stimulus was presented and the subject was required simply to maintain central fixation without generating either a prosaccade or antisaccade response. Half trials provided a measure of preparatory activation patterns separate from response-related signals. The subject did not know whether a given trial would be a whole or half trial until the peripheral stimulus appeared or failed to appear, respectively. The time interval between the starts of adjacent trials was 3, 4, or 5 s, randomized with a mean of 4 s (uniform distribution). This jittering was done to decorrelate the fMRI signal components evoked by individual trial types allowing for subsequent deconvolution as described below (Dale 1999).

Each subject performed 8 to 15 data collection runs (median 12 runs; total 122 runs across all 11 subjects). An individual run lasted 316 s. One run consisted of 73 trials (24 whole prosaccade trials, 12 half prosaccade trials, 24 whole antisaccade trials, 12 half antisaccade trials, and the first trial of the run). The first trial of a run, which was a whole prosaccade trial on half of the runs and a whole antisaccade trial on the other half, was modeled separately because it could not be counterbalanced in terms of the preceding trial type, given that there was no preceding trial. Each of the 12 possible peripheral stimulus locations (see above) was used twice for the whole prosaccade and whole antisaccade trials in a given run. In addition, the first trial of the run always used a peripheral stimulus with 6° eccentricity and 20° clockwise relative to the horizontal meridian on the right side. The order and timing structure of the trial presentation sequence within each run was determined using computerized random search to minimize correlation among different trial types and to ensure first order counterbalancing of trial parameters (prosaccade vs. antisaccade, half vs. whole trials, left vs. right peripheral stimulus) in the forward and backward directions.

We chose to present whole and half trials at a 2:1 ratio (see above) based on a compromise between two criteria. We presented half trials less frequently than whole

trials so that subjects would not come to expect a large proportion of trials to be half trials with no response requirement. This expectation might have attenuated the recruitment of preparatory processes evoked by the instruction cue. On the other hand, we had to include enough half trials to allow for decent statistical precision when deconvolving half trial activation profiles (see *Section 3.2.7 – Statistical Analysis and Linear Deconvolution* below for details of deconvolution). We decided that a 2:1 ratio of whole to half trials was a good compromise between these two competing concerns, and pilot studies confirmed this.

On the day of an experiment, subjects were instructed on task performance and given 5-15 minutes of practice until they were comfortable with the tasks. Within 15 minutes, all subjects were able to perform the tasks competently as assessed by the authors' visual inspection. We also recorded subject behaviour to verify performance, as described in the next section.

3.2.5 – Behavioural Analysis

We used video eye tracking data collected concurrently with fMRI scanning to determine saccadic response latencies and saccade end points. We discarded trials with several types of mistakes including prosaccade and antisaccade error trials, failure to respond, excessively short or long response latency (latency < 100 ms or latency > 500 ms, respectively), and failure to maintain fixation. The average proportion of trials discarded for any of these reasons was 10.6% ± 5.2% (standard deviation) with a range of 3.2% to 19.1%, across all 11 subjects. Prosaccade error trials, in which subjects incorrectly made antisaccade rather than prosaccade responses, constituted 2.8 ± 2.1% of prosaccade whole trials on average. Error antisaccade trials, in which subjects failed to inhibit the automatic saccade and generated prosaccade responses, constituted 14.7 ± 8.9% of antisaccade whole trials on average. We did not analyze error trials because half trials, which contained no response events, could not be classified in terms of correct and error performance, preventing us from measuring preparatory activation separately from response-related activation on error trials. Percentages of trials discarded due to failure to maintain fixation during fixation periods were as follows: prosaccade half trials 0.9 ±

0.9% (mean and standard deviation), antisaccade half trials 1.9 ± 2.5%, prosaccade whole trials 0.4 ± 0.4%, and antisaccade whole trials 0.7 ± 0.9%.

3.2.6 – Preprocessing

Preprocessing was done using Matlab 6 (The Mathworks, Inc., Natick, Ma, USA) and BrainVoyager 2000 (Brain Innovation BV, Maastricht, The Netherlands). We removed the first four volumes from each voxel's time course to allow for complete spin saturation. We performed nonlinear trend removal on the data using a running lines smoother programmed in Matlab 6. This algorithm corrected time courses for low frequency trends as described in *Appendix 4.1*, and has the advantage of being robust against edge effects. The trend removal algorithm also removed spikes in the signal, defined as any point greater than three standard deviations from the mean, trend-removed signal. All functional data were superimposed onto 3D anatomical images, resampled into 3x3x3 mm cubic voxels, aligned onto the anterior commissure - posterior commissure axis, and scaled to Talairach space with BrainVoyager 2000. Functional data underwent 3D motion correction by trilinear interpolation of each volume to a single reference functional volume. Mean subject head translation was 0.07 ± 0.04 mm (standard deviation), and mean subject head rotation was 0.13 ± 0.06 °. Maximum motion across all subjects was 0.14 mm translation and 0.25 ° rotation. We did not exclude any subjects on the basis of motion. Functional data were smoothed with a spatial low-pass filter using a low-pass cutoff of 0.063 cycles / mm in the Fourier domain. Finally, voxel time courses were scaled into percent signal change values on a run-by-run basis using the mean time course across the run as baseline.

3.2.7 – Statistical Analysis and Linear Deconvolution

Because we used a rapid event-related design, with trial onset spacing randomized from 3-5 s with a mean of 4 s (uniform distribution), the BOLD signal we recorded from any given point was the sum of the BOLD impulse responses from several preceding task events. This necessitated the use of deconvolution methods to extract BOLD responses from individual task components. We used linear regression based on the general linear model (GLM) to perform both the statistical analysis and the linear deconvolution. The

process of fitting a GLM to rapid event-related fMRI data can effect a linear deconvolution because it finds the best scaling of predictor curves which model overlapping, summated BOLD signal components.

We used a mixed effects analysis based on Worsley et al. (2002) (see *Appendix 4.2*). A GLM was fit to each subject's data after pre-whitening to correct for coloured noise covariance. Pre-whitening avoids the inflation of type I error that occurs with GLM statistical tests which fail to account for coloured noise in the data (Bullmore et al. 2001). Our GLM design matrix incorporated sets of finite impulse response predictor functions (Serences 2004) for eight groups of trials including 1) correct half prosaccade trials; 2) correct whole prosaccade trials; 3) correct half antisaccade trials; 4) correct whole antisaccade trials; 5) error whole antisaccade trials (when subjects made prosaccades rather than antisaccades); 6) the first trials of each run, which were modeled separately because they could not be counterbalanced for the preceding trial type; 7) any prosaccade or antisaccade trials with responses greater than 500 ms in latency; and 8) other discarded trials with latencies within the 100 to 500 ms range. That is, for each of these groups of trials, a set of 12 impulse responses (one impulse per 1 sec volume) was locked to the start of each trial in the group. Together, the 12 impulse predictors for a given group of trials, for example antisaccade whole trials, modeled the activation time course for that group of trials with one point per second over 12 seconds. Viewed another way, for each of the 8 groups of trials, a subject's GLM design matrix contained 12 columns, with ones at the appropriate locations, to model the 12 impulse functions for that trial group (Serences 2004). The design matrix thus contained a total of 97 columns (8 groups of trials times 12 columns per trial group, plus one column for the offset predictor). The constant, offset predictor, which was composed all of ones, was included to model the baseline across all (percent scaled) functional runs in a given subject. We chose the finite impulse response model because it makes no assumptions about the shape of the haemodynamic response function other than its length (12 sec in this case). Note that this is unlike more constrained approaches, such as incorporating a gamma function into the model, which assumes that the haemodynamic response is gamma-shaped. The use of linear regression also assumes of course that the haemodynamic response is linear (Boynton et al. 1996). It was necessary to model discarded trials explicitly in the design

matrix because the fMRI signals they evoked would otherwise contaminate the deconvolution of signals from adjacent correct trials. The prosaccade and antisaccade whole trials each consisted of a preparation and a response event, whereas the prosaccade and antisaccade half trials contained only the preparation event. Half trials then provided a measure of preparation-related fMRI activation, which was subtracted from whole trial activation patterns to yield estimates of purely response-related fMRI signals.

For each of the eleven subjects, we computed two statistical parametric maps by contrasting four pairs of activation profiles. The first contrast was antisaccade preparatory activation minus prosaccade preparatory activation over a time interval from 4-6 s after the instruction onset. This interval was chosen a priori because it straddles the peak of the haemodynamic impulse response curve as described previously (Boynton et al. 1996). The second contrast was antisaccade response activation minus prosaccade response activation over the time interval from 5-7 s after instruction onset. The response contrast time interval was shifted 1 s later than the preparation contrast interval because the onset of the peripheral stimulus which evoked a prosaccade or antisaccade response in the whole trials occurred 1.2 seconds after the instruction stimulus onset.

For each of the statistical contrasts, individual subject maps were combined using a mixed effects analysis (Worsley et al. 2002), which addressed variability both within subjects and between subjects. Mixed effects contrast maps were thresholded at a T-value of 2.0 ($p < 0.05$ uncorrected, df = 15000) and then cluster size thresholded at a minimum size of 1742 cubic mm to correct for multiple comparisons across the voxel population at $p < 0.05$, based on Gaussian random field theory (Worsley et al. 2002). For details of the mixed effects analysis, see *Appendix 4.3* and (Worsley et al. 2002).

3.2.8 – Presentation of Statistical Analysis

We present the results of the above-described analysis in three formats. *Figure 3.2* shows statistical parametric maps overlaid on brain slices in the typical style of the fMRI literature. The preparation and response contrast maps, computed as described earlier, are shown on the left and right of *Figure 3.2A* and *B*, respectively. Orange and yellow voxels exhibited greater activation on antisaccades, while blue voxels exhibited greater

Figure 3.2 – Preparation and Response Contrast Maps

A: Preparation difference map built by contrasting antisaccade half trials with prosaccade half trials. Yellow/red regions exhibited greater antisaccade activation. Note that preparatory activation was greater for antisaccades in cortical saccade regions including the frontal eye field (FEF), supplementary eye field (SEF), and intraparietal sulcus (IPS) as well in left dorsolateral prefrontal cortex (DLPFC), anterior cingulate cortex (ACC), and precuneus. Maps are mixed effects maps across 11 subjects. Images obey neurological convention.

B: Response difference map built by contrasting antisaccade and prosaccade response-related activation. Response activation for both trial types was computed by subtracting half trial activation from whole trial activation (see *Section 3.2.7 – Statistical Analysis and Linear Deconvolution* and *Figure 3.1*). Conventions as in *A*. Note that response-related activation was greater for antisaccades in FEF, SEF, IPS, and precuneus but not in left DLPFC or ACC. Response activation was greater for prosaccades in angular gyrus, which is shown in blue.
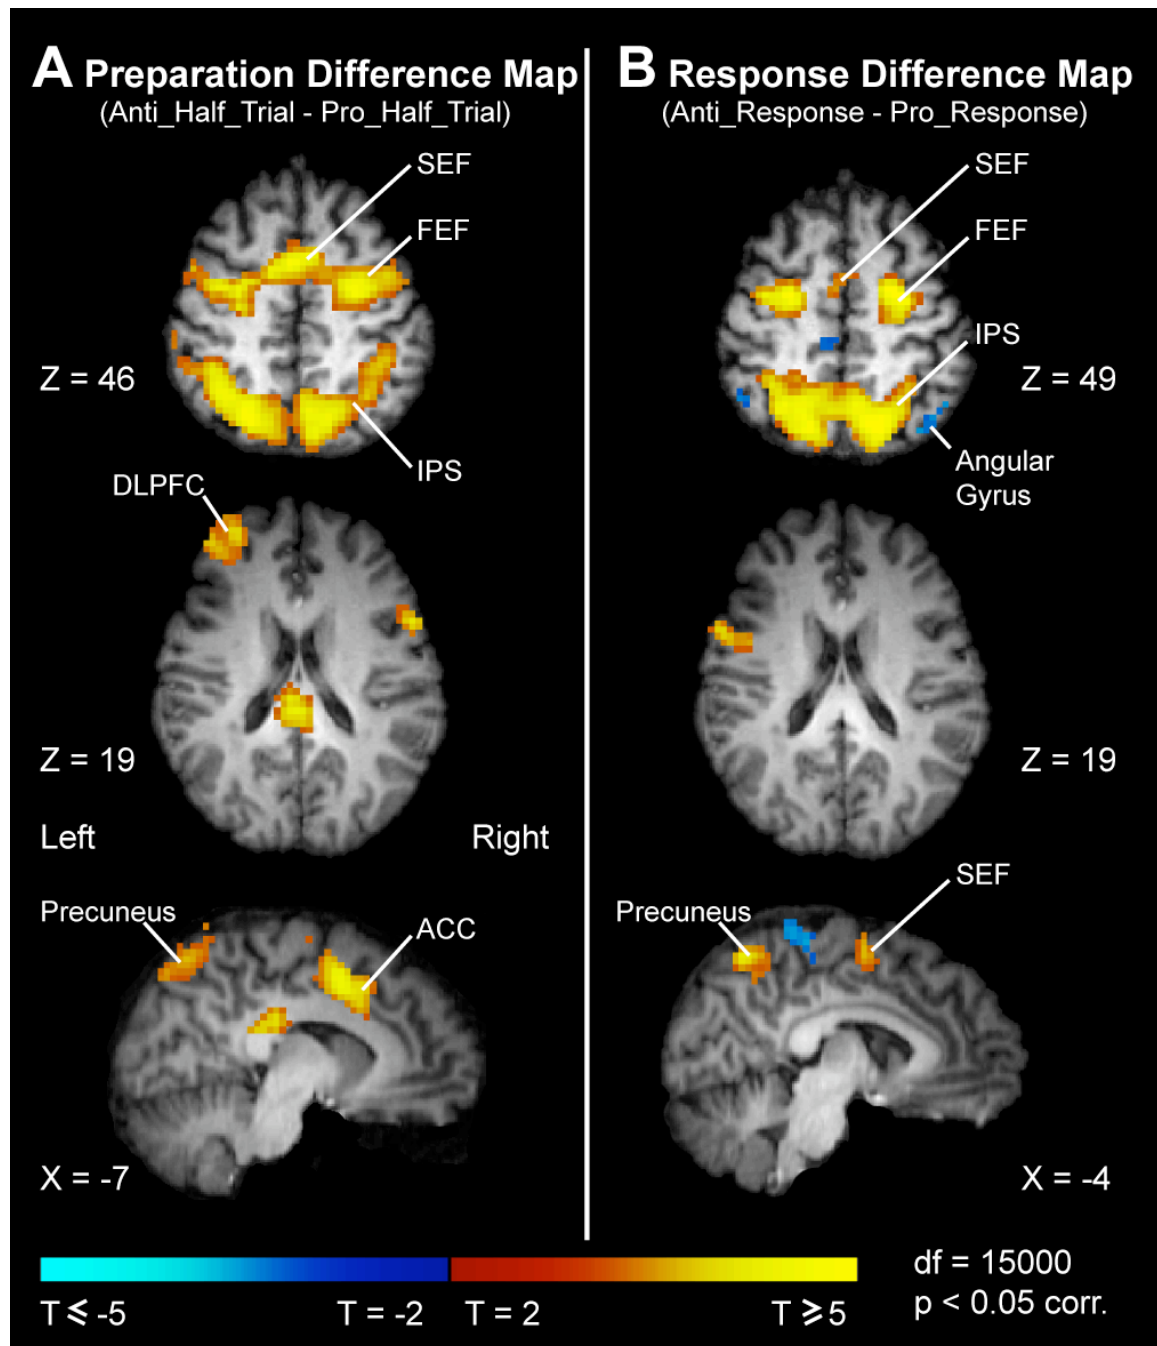
activation on prosaccades. Axial slices are in neurological coordinates, with the left side of the image corresponding to the left side of the brain.

Figure 3.3 displays activation time courses for select regions of interest. These time courses were computed as follows. First consider the time course one might extract from a single voxel in a single subject. In our GLM, we used sets of finite impulse response predictors for various groups of trials, as described above. The beta weights computed by the GLM for a specific set of finite impulse responses in fact constitute a time course for the corresponding group of trials (Serences 2004). For example, the first 12 columns of our GLM design matrix included impulses modeling the prosaccade half trials, so the first 12 beta weights in the beta weight vector computed for a given voxel would constitute that voxel's activation time course for prosaccade half trials, with one activation data point per second over 12 seconds. In this way, one could derive voxel-specific time courses for prosaccade and antisaccade half and whole trials. Note that in this case no special baseline adjustment would have occurred beyond the very simple baseline that is part of the GLM design matrix. By this, we refer to the last column of a subject's design matrix, which was all ones and modeled the mean activation across that subject's functional runs. Recall that all functional runs were percent scaled during preprocessing, avoiding the necessity to model each run with its own constant offset column in the design matrix. *Figure 3.3* shows averaged time courses, and averaging was performed as follows. Time courses for each voxel in a region of interest were first derived from the beta weights on a subject-by-subject basis. For each time point in the time course, activation data were then averaged across all voxels in the region of interest and across all subjects. Regions of interest were defined from clusters of voxels exhibiting significance in either or both of the preparation or response contrasts. In right frontal eye field (FEF), supplementary eye field (SEF), right rostral intraparietal sulcus (IPS), right caudal IPS, and precuneus, we found clusters of statistically significant voxels that overlapped in both the preparation and response contrasts. The time courses for these regions were computed across the intersection of significant voxels from both contrasts. In anterior cingulate cortex (ACC) and left dorsolateral prefrontal cortex (DLPFC), we found significant voxel clusters only in the preparation contrast, and time courses for these regions were computed from the preparation contrast-derived clusters.

Figure 3.3 – Deconvolved Time Courses

Deconvolved time courses for *A*: half trials (preparation), *B*: whole trials (preparation and response), and *C*: whole trials minus half trials (response). Antisaccade and prosaccade time courses are depicted with solid and dashed lines, respectively. Time courses for left and right frontal eye field (FEF), supplementary eye field (SEF), left and right rostral intraparietal sulcus (IPS), left and right caudal IPS, and precuneus were derived from voxels exhibiting significance in both the preparation and response contrast maps (see Figure *2A* and *2B*). Anterior cingulate cortex (ACC) and left dorsolateral prefrontal cortex (DLPFC) time courses were taken from voxels exhibiting significance in the preparation contrast map only, as these regions showed no significant differences in the response contrast map. Gray rectangles in the left and right columns denote the 4-6 s interval and the 5-7 s interval used, respectively, to build the preparation and response contrast maps in *Figure 3.2A* and *3.2B* (see *Section 3.2 – Methods*). Asterisks denote statistically significant differences between antisaccades and prosaccades during the appropriate time intervals ($p < 0.05$ corrected). In columns *B* and *C*, the vertical gray lines positioned at 1.2 s denote the peripheral stimulus onset in whole trials. The time courses shown here were computed first by deconvolving time courses for individual voxels (see Methods) and then averaging across voxels in a given region of interest. Ordinate units (BOLD percent signal change) were derived by percent scaling individual functional runs on a voxel-by-voxel basis using each voxel's mean activation level as baseline. Further details of deconvolution and derivation of time courses are given in *Section 3.2.8 – Presentation of Statistical Analysis*.
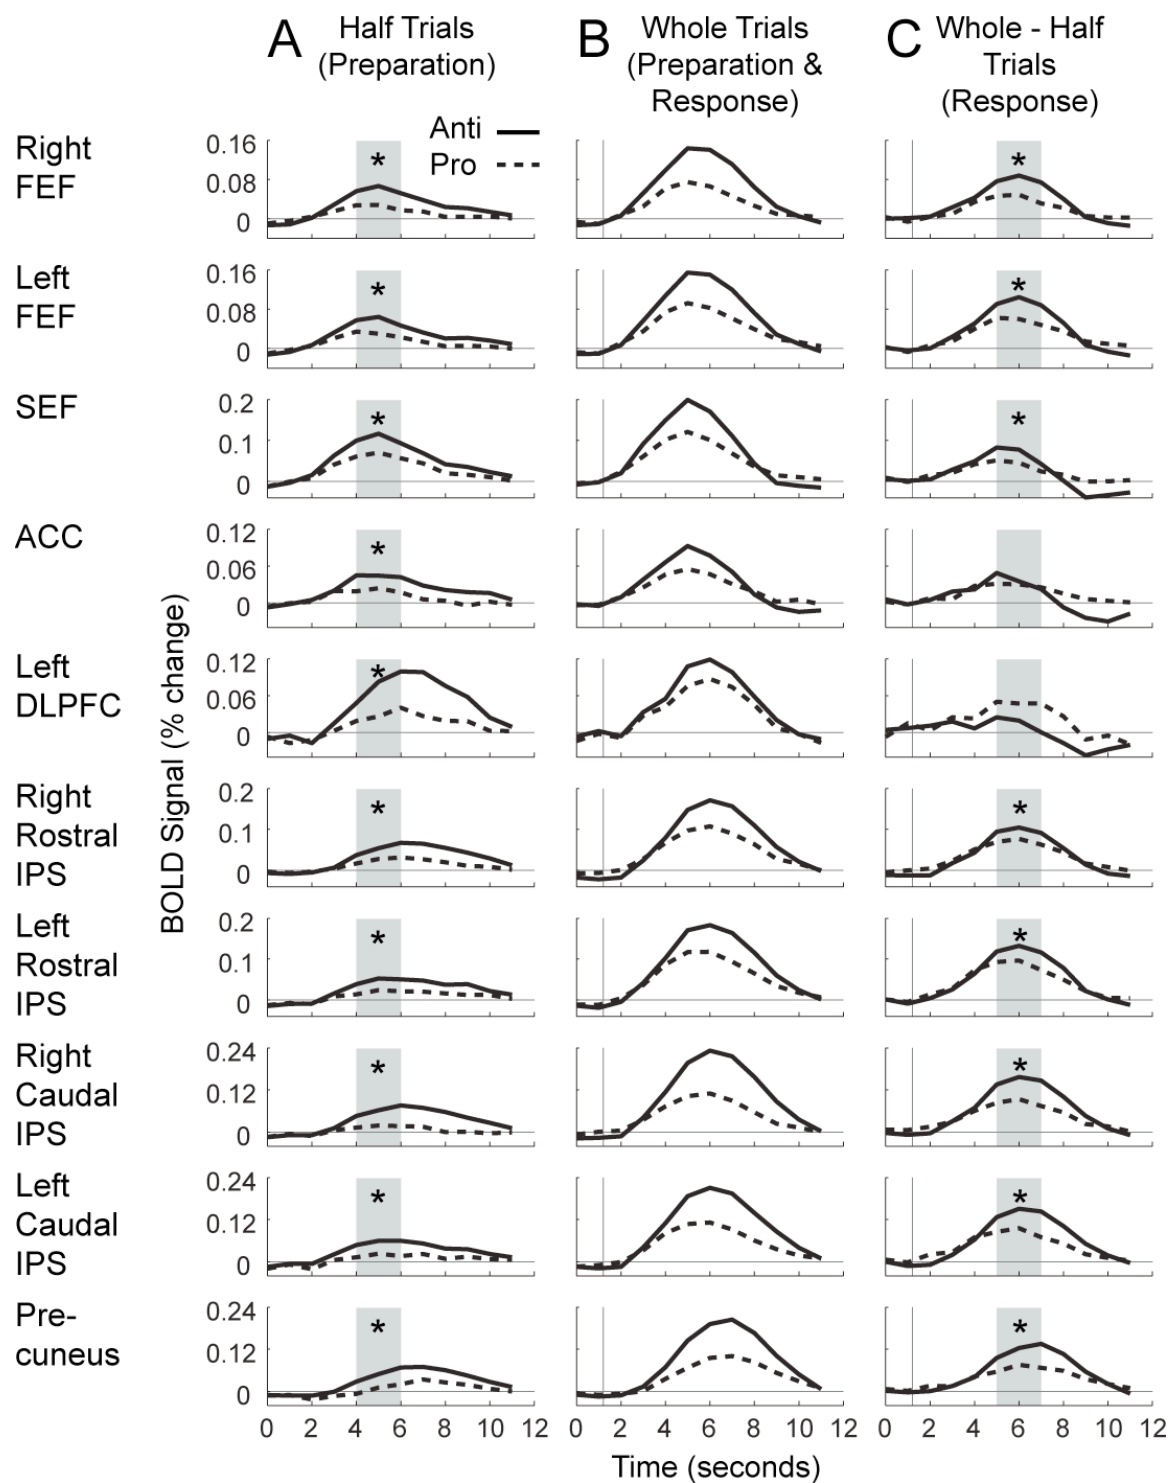
Figure 3.4 shows the activation differences values that were computed in the course of building the preparation and response contrast maps. More specifically, in building the preparation and response contrast maps, we computed preparation- and response-related activation difference values (antisaccade – prosaccade) for each voxel. Recall that preparation-related activation was obtained from half trials, whereas response-related activation was obtained by subtracting half trial from whole trial activation. As described above, time windows of 4-6 s and 5-7 s were used, respectively, for the preparation and response contrasts. These voxel-specific activation difference values were combined across subjects using the expectation maximization algorithm rather than averaging, again as described above. The expectation maximization algorithm also computed between-subjects variance for each difference value, and we then combined within- and between-subjects variances to derive the total variance for each voxel-specific difference value, in both the preparation and response contasts. To derive the bar graphs shown in *Figure 3.4*, activation difference values were averaged across all voxels in a given region of interest, where regions of interest were defined as statistically significant clusters of voxels in either the preparation or response contrast. Error bars shown in *Figure 3.4* were derived by averaging across voxel-specific standard deviations within a region of interest. *Figure 3.4* shows activation difference values computed from the preparation contrast (A) and the response contrast (B) for all regions of interest found to be significant from both contrasts. Asterisks indicate that the activation differences are significantly greater or less than zero (p < 0.05, corrected for multiple comparisons), and absence of an asterisk denotes lack of significance. If a given anatomical location, for example right frontal eye field (FEF), contained statistically significant voxel clusters in both the preparation and response contrasts, then the corresponding preparation- and response-related activation difference values shown in *Figure 3.4A* and *B* were computed based on the clusters of voxels derived separately from the preparation and response contrasts. If an anatomical location exhibited a significant cluster of voxels in only one contrast, that single cluster of voxels was used to compute the activation difference values for both *Figure 3.4A* and *B*. For example, we found a significant voxel cluster in anterior cingulate cortex (ACC) in the preparation contrast but not in the response

Figure 3.4 – Mean Activation Differences

Bar graphs showing mean *A*: preparation- and *B*: response-related activation differences (antisaccade – prosaccade) used to compute the preparation and response contrasts, respectively, that are displayed in Figure *2*. Differences are shown for all statistically significant regions in the preparation and response contrasts. Error bars are (mean) standard deviations of mean differences. Activation differences were computed over the 4-6 s time interval for preparation differences (*A*) and the 5-7 s interval for response differences (*B*). These time windows are also shown in *Figure 3.3* as light gray rectangles. Asterisks denote statistical significance (p < 0.05, corrected for multiple comparisons), while lack of an asterisk denotes lack of significance. DLPFC stands for dorsolateral prefrontal cortex. IPS stands for intraparietal sulcus. For further details of bar graph computation, see *Section 3.2.8 – Presentation of Statistical Analysis*. For further details of preparation and response contrasts, see *Section 3.2.7 – Statistical Analysis and Linear Deconvolution*.

A. Preparation Diffference
(Anti - Pro)

B. Response Difference
(Anti - Pro)

Right Frontal Eye Field (FEF)
Left Frontal Eye Field (FEF)
Supplementary Eye Field (SEF)
Anterior Cingulate Cortex (ACC)
Left Dorsolateral PFC (DLPFC)
Right Precentral Gyrus
Left Precentral Gyrus
Right Inferior Frontal Gyrus
Right Postcentral Sulcus
Left Postcentral Sulcus
Right Rostral IPS
Left Rostral IPS
Right Caudal IPS
Left Caudal IPS
Precuneus
Ventral Posterior Cingulate Cortex
Right Angular Gyrus
Left Angular Gyrus
Dorsal Precuneus
Right Dorsal Postcentral Sulcus
Left Dorsal Postcentral Sulcus

BOLD % Signal Change

Table 3.1 – Preparation and Response Contrasts

Region	A. Preparation Contrast					B. Response Contrast				
	Dir	X	Y	Z	Volume	Dir	X	Y	Z	Volume
Right Frontal Eye Field (FEF)	+	27	-5	53	14,769	+	23	-8	54	7,830
Left Frontal Eye Field (FEF)	+	-28	-6	54	11,529	+	-26	-8	52	8,262
Supplementary Eye Field (SEF)	+	0	2	50	6,048	+	0	-1	55	1,917
Anterior Cingulate Cortex (ACC)	+	-1	9	39	6,453					
Left Dorsolateral Prefrontal Cortex (DLPFC)	+	-32	49	21	4,212					
Right Precentral Gyrus	+	49	5	35	3,213					
Left Precentral Gyrus	+	-52	4	37	1,296	+	-50	4	25	2,754
Right Inferior Frontal Gyrus	+	46	21	26	3,429					
Right Postcentral Sulcus						+	52	-24	34	2,619
Left Postcentral Sulcus						+	-36	-33	35	2,079
Right Rostral Intraparietal Sulcus (IPS)	+	35	-47	46	6,075	+	30	-48	49	2,160
Left Rostral Intraparietal Sulcus (IPS)	+	-37	-47	47	8,370	+	-29	-47	48	5,103
Right Caudal Intraparietal Sulcus (IPS)	+	16	-67	48	9,342	+	17	-65	51	6,669
Left Caudal Intraparietal Sulcus (IPS)	+	-20	-66	49	9,369	+	-19	-63	49	8,478
Precuneus	+	2	-60	54	6,669	+	0	-56	51	6,318
Ventral Posterior Cingulate Cortex	+	0	-27	23	2,943					
Right Angular Gyrus						-	44	-63	40	4,347
Left Angular Gyrus	-	-47	-65	32	4,023	-	-46	-65	40	1,998
Dorsal Precuneus						-	-5	-32	60	1,836
Right Dorsal Postcentral Sulcus						-	15	-36	68	243
Left Dorsal Postcentral Sulcus						-	-32	-30	65	1,323

Legend: Region of interest data for A) preparation and B) response contrast maps. 'Dir' denotes direction of contrast: greater activation for antisaccades (+) or for prosaccades (-). X, Y, and Z are the Talairach coordinates of the statistical centres of mass in mm. Volume is in cubic mm. Regions not explicitly listed as 'left' or 'right' straddled the midline and included voxels in both hemispheres. See *Section 3.2.8 – Presentation of Statistical Analysis* for details.

contrast. We took this single cluster of voxels and computed the average preparation-related activation difference across the cluster (shown in *Figure 3.4A*), and we computed the average response-related activation difference across the same cluster (shown in *Figure 3.4B*).

Table 3.1 lists direction of contrast, Talairach coordinates for statistical centre of mass, and volume in cubic mm for all regions of interest exhibiting statistical significance in A) the preparation contrast and B) the response contrast. Direction of contrast, denoted 'Dir' in *Table 3.1*, is given as "+" to indicate greater activation on antisaccades versus prosaccades or "−" for the converse. The three Talairach coordinates for the statistical centres of mass are given in the X, Y, and Z columns. The X coordinate for the statistical centre of mass (x_{CM}) of a region of interest was calculated with the following formula:

$$x_{CM} = \frac{\sum_x x \sum_y \sum_z T(x,y,z)}{\sum_x \sum_y \sum_z T(x,y,z)}$$, where $T(x,y,z)$ denotes the T-

statistic at a given position and x, y, and z are taken across all voxels within the region of

interest. The Y and Z coordinates for the centre of mass were calculated with equivalent formulae.

3.3 – Results

3.3.1 – Behaviour

Unlike widely-spaced fMRI studies, our rapid fMRI trial design used a similar timing regimen to psychophysical and electrophysiological experiments, facilitating comparison of our results with those from these other techniques. However, despite the similarities in timing, it is still possible that the fMRI scanning environment might have altered subject performance characteristics in some way. To confirm that subject performance was similar to what has been seen in psychophysical and electrophysiological settings, we took video eye position recordings during fMRI data acquisition. We also used these recordings to exclude the few trials that contained one of several types of mistakes (see *Section 3.2.5 – Behavioural Analysis*). Across all 11 subjects, mean correct performance values were high, $97.2 \pm 2.1\%$ (standard deviation) and $85.3 \pm 8.9\%$ for prosaccade and antisaccade trials, respectively. Antisaccade correct performance rates were significantly below those of the prosaccade task ($p < 0.001$, $T = 5.1$, $df = 10$), which is consistent with previous antisaccade studies (Everling and Fischer 1998). Mean saccadic latencies over all subjects were 260 ± 28 ms for correct prosaccade trials, 355 ± 36 ms for correct antisaccade trials, and 276 ± 40 ms for error antisaccade trials. Differences among these values were significant ($p < 10^{-6}$, $F(2, 30) = 23.04$, one-way ANOVA). Correct antisaccade latencies were significantly greater than correct prosaccade latencies ($p < 10^{-5}$, $T = 8.8$, $df = 10$) and significantly greater than error antisaccades latencies ($p < 10^{-5}$, $T = 8.5$, $df = 10$). Correct prosaccade latencies were significantly shorter than error antisaccade latencies ($p < 0.05$, $T = 2.4$, $df = 10$). In summary, the antisaccades in this study closely match those of other antisaccade studies, which typically find longer latencies for antisaccades compared to prosaccades and longer latencies for correct compared to error antisaccades (Everling and Fischer 1998).

3.3.2 – Preparatory Activation

We compared prosaccades and antisaccades using rapid event-related fMRI to test whether a rapid fMRI design can resolve the preparatory activation differences previously observed between prosaccades and antisaccades that used widely-spaced event-related fMRI. Our design included prosaccade and antisaccade half trials, comprised only of the preparation event, as a means of separating preparatory and response-related activation (see *Section 3.2.4 – Experimental Design*). To investigate preparatory differences between the two tasks, we built a statistical difference map by contrasting activation from antisaccade half trials with that from prosaccade half trials over a time interval from 4-6 seconds. *Figure 3.2A* displays the most relevant slices of this map, and *Table 3.1A* lists the positions and sizes of regions exhibiting significant differences in the map. *Figure 3.3* includes activation time courses for select regions, and *Figure 3.4A* shows mean preparatory differences between prosaccades and antisaccades for all regions.

We found that cortical saccade regions typically associated with antisaccade task performance exhibited greater antisaccade versus prosaccade activation in the preparatory period. These regions included left and right frontal eye field (FEF), bilateral supplementary eye field (SEF), left dorsolateral prefrontal cortex (DLPFC) in the left middle frontal gyrus, bilateral anterior cingulate cortex (ACC), left and right intraparietal sulcus (IPS), which contained statistically significant clusters both in the rostral and caudal portions, bilateral posterior cingulate cortex, and bilateral precuneus. In addition, the statistically "activated" voxel clusters constituting left and right FEF were continuous with activated clusters located ventrolaterally to FEF in the precentral gyrus and with a cluster in the inferior frontal gyrus on the right side. One region, DLPFC, exhibited strong hemispheric asymmetry with a large cluster of voxels (4212 mm^3) in left DLPFC displaying greater preparatory activation for antisaccades. In right DLPFC, a much smaller cluster (837 mm^3) also displayed greater preparatory signal for antisaccades, but this cluster's size was well below the 1742 mm^3 threshold for significance (see *Section 3.2.7 – Statistical Analysis and Linear Deconvolution*).

Greater antisaccade preparatory activation is also evident in the activation time courses for the prosaccade and antisaccade half trials displayed for selected regions in

Figure 3.3A. Note that these preparatory differences manifested as soon as the haemodynamic lag would allow, with the antisaccade and prosaccade half trial activation curves separating at the 3 second mark.

We also found voxel clusters exhibiting significantly greater prosaccade preparatory versus antisaccade preparatory activation in left angular gyrus (*Table 3.1A*).

3.3.3 – Response-related Activation

The whole prosaccade and antisaccade trials contained both a preparation and a response event. To isolate response-related activation, we subtracted deconvolved time courses for the half trials, which contained only the preparation and not the response event, from the whole trial time courses. *Figure 3.3* shows examples of activation time courses for *A*) half trials, *B*) whole trials, and *C*) the response-related activation curves generated by subtracting half trials from whole trials. *Figure 3.4B* shows mean response-related differences between prosaccades and antisaccades for all regions.

We contrasted antisaccade versus prosaccade response-related activation over a time interval from 5-7 seconds. The response contrast's statistical window was one second later than the preparation contrast's because the peripheral stimulus onset that evoked prosaccade or antisaccade responses in the whole trials occurred 1.2 s after the instruction onset (see *Section 3.2.4 – Experimental Design* and *Section 3.2.7 – Statistical Analysis and Linear Deconvolution*). Slices of the statistical map built from the response contrast are shown in *Figure 3.2B*, and the sizes and locations of the statistical centres of mass for significant regions are listed in *Table 3.1B*.

Many of the areas that exhibited significant differences in the preparation contrast described above also exhibited significant differences in the response contrast. We found voxel clusters with greater antisaccade versus prosaccade response-related activation in left and right FEF, bilateral SEF, left and right rostral IPS, left and right caudal IPS, and bilateral precuneus. There were also voxel clusters with greater antisaccade response-related activation in left precentral gyrus and left and right postcentral sulcus. *Figure 3.3C* shows response-related activation profiles with greater activation for antisaccade responses in left and right FEF, SEF, and left and right rostral and caudal IPS. Left DLPFC and ACC, which are also shown in *Figure 3.3C*, did not display response-related

differences. It is unlikely that this result was due to insufficient statistical sensitivity. Of all the voxels in the left DLPFC cluster that was significant in the preparation contrast (see above), none were individually significant in the response contrast ($p > 0.05$, without correction for multiple comparisons). Furthermore, 98% of left DLPFC voxels were not significant at $p > 0.1$ ($|T|<1.645$, df=15000), and 93% were not significant at $p > 0.2$ ($|T|<1.2815$, df=15000). Similarly, for voxels in the ACC cluster that exhibited significance in the preparation contrast, 92% were not significant at $p > 0.1$ ($|T|<1.645$, df=15000), and 74% were not significant at $p > 0.2$ ($|T|<1.2815$, df=15000). Only 2% of ACC voxels were individually significant in the response contrast ($p<0.05$, $T>1.96$, df=15000, uncorrected).

Voxel clusters with greater prosaccade compared to antisaccade response-related activation (blue regions in *Figure 3.2B*) were located in left and right angular gyrus, bilateral dorsal precuneus, and left and right dorsal postcentral sulcus.

As can be seen in *Figure 3.3*, particularly in *Figure 3.3B*, BOLD time courses tended to peak at 5 seconds in FEF, SEF, and ACC, whereas they peaked at 6 seconds in left DLPFC and IPS and at 7 seconds in precuneus. These differences might reflect timing differences with respect to neurocomputational processing in the various regions, but it is also possible that the differences in time to peak might reflect vascular differences among the regions. For this reason, we hesitate to draw conclusions on brain function from these timing differences.

3.4 – Discussion

Our results show strong frontoparietal activation associated with the preparation for an antisaccade. This finding is in line with single unit recording studies in nonhuman primates and event-related potential and event-related fMRI studies in human subjects. The data support the influential model that frontoparietal regions serve as an important source for cognitive control by biasing task-relevant processes in other brain regions (Desimone and Duncan 1995; Miller and Cohen 2001). We found that even a brief preparatory period of one second with no stimulus presentation and no executed saccade was sufficient to elicit a stronger BOLD fMRI activation on antisaccade trials than on

prosaccade trials in frontal eye field (FEF), supplementary eye field (SEF), anterior cingulate cortex (ACC), dorsolateral prefrontal cortex (DLPFC), and intraparietal sulcus (IPS). By subtracting preparatory activation as measured in the half trials from the combination of preparation- and response-related activation measured in the whole trials, we were able to resolve response-related activation profiles. While FEF, SEF, and IPS were also more active for the antisaccade response than the prosaccade response, the ACC and DLPFC where only more active for antisaccade than prosaccade trials during the preparatory period.

3.4.1 – Preparatory Activation

To our knowledge, this is the first fMRI study that found significantly higher preparatory activations in frontoparietal regions for antisaccades compared with prosaccades for very short preparatory periods. Our results extend those of (Curtis et al. 2005) who demonstrated that short catch trials, in which subjects fixated a central fixation point and expected a peripheral stimulus to appear, evoked significant activation in frontoparietal regions. Here we have shown that this frontoparietal activation in specially designed catch trials, namely the half trials we used in this experiment, is task-dependent. Recently, we compared prosaccade, antisaccade, and nogo trials using a rapid event-related compound fMRI design and did not find differences during the preparatory period between these tasks (Brown et al. 2006). In that study, the preparation interval length was jittered to decorrelate preparation- and response-related fMRI signal components. One potential limitation of the previous study's design was that the instruction for a given trial type was always followed by the response for that trial type, introducing a sequence bias. This sequence effect might have caused any preparatory differences to be obscured by the subsequent response-related signals, which did exhibit robust differences between prosaccades and antisaccades in frontoparietal regions. We therefore suggest that a half-trial design is better-suited to detect task-dependent preparatory processes than jittered compound trials in rapid fMRI.

Several previous event-related fMRI studies have found higher activations in some of these frontoparietal areas for antisaccades compared with prosaccades (Connolly et al. 2002; Desouza et al. 2003; Ford et al. 2005; Curtis and D'Esposito 2003). Connolly

and colleagues found an increased activation for antisaccades compared with prosaccades in FEF but not in IPS. The authors imaged only one brain slice, and they used compound trials with very short preparatory periods (0, 2, and 4 s). Because of the haemodynamic lag, they compared levels of preparatory activation 3 s after instruction onset, and therefore response processes might have influenced this activation. The study by Curtis and D'Esposito (2003) found differences only at the end of a 7 s preparatory period in the pre-SMA. Subjects in that study were required to hold the task instruction in working memory, a process that is known to activate frontoparietal regions. Such a working memory-related activation in both tasks might have masked any further potential differences between the preparation of prosaccades and antisaccades. As in the present study, Ford and colleagues found differences in similar frontoparietal regions, although considerably weaker (Ford et al. 2005). These differences, however, only appeared at the end of their 10 s preparatory period, throughout which the task instruction cue (coloured fixation point) remained visible. The findings of Curtis and D'Esposito (2003) and Ford et al. (2005) suggest that subjects may postpone task-specific preparation until the end of a fixed-length instruction interval. Future studies making a direct comparison of short and long instruction intervals could test this prediction.

Preparatory differences in DLPFC were highly asymmetric in our study, with only left DLPFC exhibiting a large, statistically significant cluster of voxels with higher preparatory activation for antisaccade trials. The antisaccade fMRI literature is somewhat inconsistent on DLPFC and hemispheric asymmetry. Desouza et al. (2003) found greater preparatory activity for antisaccades versus prosaccades in right DLPFC only. Ford et al. (2005) found a similar result but in left rather than right DLPFC. Ford et al. also found greater preparatory activity for correct antisaccades compared to error antisaccades in right DLPFC only. Curtis and D'Esposito (2003) and Connolly et al. (2002) did not investigate DLPFC for technical reasons. Given these inconsistencies, it would be premature to make strong statements on hemispheric asymmetry with regard to DLPFC's role in antisaccade performance.

3.4.2 – Task Preparation Versus Response Generation

Response-related activations were greater for antisaccades versus prosaccades in several of the cortical saccade regions that exhibited preparatory differences (see above), consistent with our previous rapid event-related results (Brown et al. 2006). As we suggested in that study, greater response-related activation for antisaccades compared to prosaccades could arise from a variety of processes, including the visuospatial remapping required to perform the antisaccade task and a more voluntary motor command with increased attention on antisaccade trials. Interestingly, an event-related fMRI study that used long instruction periods of 6 or more seconds did not find any differences between prosaccade and antisaccades during the response period (Desouza et al. 2003). We hypothesize that the long instruction periods in that study made the prosaccade task less automatic and therefore diminished activation differences.

While the areas that make up the "classical" cortical saccade network, namely FEF, SEF, and IPS, showed differences during the preparatory period and during the response period, the ACC and the DLPFC did not show any differences in their response-related activation between prosaccades and antisaccades. This suggests that these two classical "control" areas are involved in the preparatory set for an antisaccade but not in the generation of the actual response. This hypothesis is also supported by studies in human subjects who suffered brain damage. Lesions of the dorsolateral prefrontal cortex in humans lead to increased error rates in the antisaccade task, but they do not impair the ability to generate an antisaccade or prosaccade (Guitton et al. 1985; Pierrot-Deseilligny et al. 2003; Ploner et al. 2005; Pierrot-Deseilligny et al. 2002; Pierrot-Deseilligny et al. 1991). The same finding has been reported for lesions of the anterior cingulate cortex (Gaymard et al. 1998). In contrast, damage to FEF (Gaymard et al. 1999) and IPS (Pierrot-Deseilligny et al. 1991) does not increase error rates in the antisaccade task, but it decreases the gain and increases the reaction times of contralateral saccades.

3.4.3 – Processes Related to Frontoparietal Activation

The frontoparietal activation evoked by the onset of an instruction stimulus will certainly reflect a variety of different processes. Indeed, a similar activation pattern is produced by increases in task difficulty in a multitude of cognitive tasks (Desimone and

Duncan 1995; Miller and Cohen 2001). Compared to the performance of a prosaccade, the correct performance of an antisaccade requires at least two additional processes. The subject must first suppress the automatic response to look towards the flashed visual stimulus and then invert the stimulus location into a voluntary motor-command to look away from the stimulus. Single neuron recordings in monkeys have provided evidence that the brain accomplishes the first process by reducing the level of preparatory activity in the saccadic eye movement system before the stimulus appears (Munoz and Everling 2004). Saccade-related neurons in the superior colliculus (SC) and FEF have a lower level of activity during the preparatory period on antisaccade trials compared with prosaccade trials. Fixation-related neurons in these areas show the reverse pattern of activity. Higher levels of preparatory activity for antisaccades compared with prosaccades have been found in the SEF (Amador et al. 2003). A large percentage of neurons in the DLPFC also exhibit differential preparatory activity between pro- and antisaccade trials (Everling and DeSouza 2005). The DLPFC is connected to a large number of cortical and subcortical areas, and Miller and Cohen (2001) have proposed that it provides bias signals for task-relevant processing to other areas. Indeed, it has recently been demonstrated that neurons in the DLPFC that project directly to the SC exhibit higher activity during the preparatory period on antisaccade trials than on prosaccade trials (Johnston and Everling 2006b). This finding is consistent with the idea that the DLPFC provides saccade suppression signals to the SC (Gaymard et al. 2003). Many neurons in the ACC also exhibit different levels of preparatory activity between prosaccade and antisaccade trials, and this difference actually appears earlier than those in the DLPFC following a task switch (Johnston et al. 2007). Johnston and colleagues (2007) therefore suggested that the ACC also provides preparatory bias signals to other brain areas. These areas include both the SC (Leichnetz et al. 1981) and FEF (Wang et al. 2004). Together these data suggest that the frontal activation reflects at least in part this presetting of the saccade circuitry for the antisaccade task.

It is not yet known, whether neurons in the lateral intraparietal (LIP) area, the putative monkey homologue of the human IPS, also exhibit task-related differences in prestimulus excitation. Gottlieb and Goldberg (1999) recorded single unit activity in area LIP while monkeys performed prosaccade and antisaccade trials. Similar to our

paradigm, monkeys were instructed at the onset of each trial by the colour of the central fixation point to generate either a prosaccade or antisaccade upon stimulus presentation. However, half of the trials were randomly interleaved memory-guided prosaccade and antisaccade trials, in which the animals had initially to suppress a saccade. Therefore, monkeys had to suppress a saccade also on 50% of the prosaccade trials. Although Gottlieb and Goldberg (1999) did not compare preparatory activity between pro- and antisaccade trials in their study, it would be surprising to find any differences in this particular design, because the animals likely adopted a strategy in which they prepared initially to suppress a saccade on both antisaccade and prosaccade trials. The same argument applies to another study that employed only memory-guided prosaccade and antisaccade trials (Zhang and Barash 2000).

While Gottlieb and Goldberg (1999) did not report any differences in preparatory activity between pro- and antisaccades, they observed a modest, but significant increase in stimulus-related activity for antisaccades compared with prosaccades in the population of LIP neurons. This higher activity for antisaccades is likely critical for task performance because LIP neurons showed a reduced stimulus-related response on error antisaccade trials when monkeys looked towards the stimulus. Moreover, a small percentage of LIP neurons (12%) encoded reliable saccade direction signals. The higher stimulus-related activity of LIP neurons for antisaccades could be the basis of the enhanced BOLD activation that we found for antisaccade trials compared with prosaccade trials during the response period.

The parietal cortex has also been implicated in the process of vector inversion required for antisaccades by single unit recordings in monkeys (Zhang and Barash 2000, 2004), human ERP (Everling et al. 1998), and human fMRI studies (Medendorp et al. 2005). The vector inversion, however, can only start once the peripheral stimulus has been presented. Although the parietal activation might reflect the preparation of neurons in these areas for this process, we would regard this explanation as unlikely.

3.4.4 – Frontoparietal Cortex and Visuospatial Attention

Curtis and colleagues interpreted the activation in FEF and IPS on catch trials as evidence for a role of these areas in the covert direction of attention to the space where

the targets were expected to appear (Curtis et al. 2005). This interpretation is in line with many experiments that have demonstrated that reorientation of attention is a core function of frontoparietal regions (Posner and Petersen 1990; Corbetta et al. 1998; Corbetta et al. 2000). The shift of attention to peripheral locations certainly contributes to the frontoparietal activation that we observed on half trials. Such shifts of attention, however, should occur on both prosaccade and antisaccade trials as subjects have to localise the stimulus in both tasks.

Hon et al. (2006) demonstrated that attended stimulus changes, even when they require no behavioural responses, are sufficient to evoke frontoparietal activation. The authors interpreted this finding as evidence that these regions are activated by the simple updating of attended information. According to their model, the change of the colour of the attended central fixation point might have been sufficient to evoke frontoparietal activation in our paradigm. However, this process cannot account for the higher activation for antisaccades compared with prosaccades that we found in frontoparietal regions in this study.

3.4.5 – Retinotopy

Although not shown here, we did not find differences in activation between contralateral and ipsilateral saccades. A contralateral bias has recently been reported for the IPS and FEF in human fMRI studies (Medendorp et al. 2005; Curtis and D'Esposito 2006; Schluppeck et al. 2006), but this bias originated during delay periods when subjects prepared to generate a saccade towards a previously cued location. Sereno et al. (2001) and Kastner et al. (2007) also found contralateral activation biases in IPS and FEF, respectively, using memory-guided saccade tasks, but they did not separate activation evoked by visual stimulus processing, working memory, and saccade generation. All of these could have contributed to the contralateral biases observed. To our knowledge, no fMRI study that did not contain a spatial working memory / saccade preparation component has reported a contralateral saccade activation bias in humans. This might be surprising given the strong contralateral bias in cortical neural activity found in monkey electrophysiology studies, especially in the FEF (Bruce and Goldberg 1985; Schall 1997). This difference might be related to the poor temporal resolution of fMRI, which makes it

impossible to distinguish between activation caused by the saccade response and activation evoked by the subsequent return saccade. It is also known that the activity of FEF neurons in monkeys is suppressed prior to saccade onset if the saccade is directed outside of their response fields (Schall et al. 1995; Everling and Munoz 2000; Seidemann et al. 2002). Thus, the metabolic activity associated with the suppression of large populations of neurons in the ipsilateral and contralateral FEF might prevent the detection of a contralateral saccade bias with fMRI.

3.4.6 – Conclusions

Taken together, we suggest that higher activation on antisaccade trials in the frontoparietal network reflects mainly the presetting of the oculomotor circuitry for the antisaccade task. This presetting ultimately requires a reduction in preparatory activity in the superior colliculus prior to stimulus presentation so that the process initiated by the incoming visual signal does not reach the saccade threshold before the vector inversion process for the voluntary antisaccade is completed (Munoz and Everling 2004). We propose that the DLPFC and ACC, the only areas that were more active for antisaccades than prosaccades during the preparatory period but not during the response period, bias activation in FEF, SEF, IPS, and SC prior to stimulus onset.

3.5 – Bibliography

Amador N, Schlag-Rey M, and Schlag J. 2003. Primate Antisaccade II. Supplementary Eye Field Neuronal Activity Predicts Correct Performance. *J Neurophysiol* 91: 1672-1689.

Barberi EA, Gati JS, Rutt BK, and Menon RS. 2000. A transmit-only/receive-only (TORO) RF system for high-field MRI/MRS applications. *Magn Reson Med* 43(2): 284-289.

Bell AH, Everling S, and Munoz DP. 2000. Influence of stimulus eccentricity and direction on characteristics of pro- and antisaccades in non-human primates. *J Neurophysiol* 84(5): 2595-2604.

Boynton GM, Engel SA, Glover GH, and Heeger DJ. 1996. Linear systems analysis of functional magnetic resonance imaging in human V1. *J Neurosci* 16(13): 4207-4221.

Brown MR, Goltz HC, Vilis T, Ford KA, and Everling S. 2006. Inhibition and generation of saccades: Rapid event-related fMRI of prosaccades, antisaccades, and nogo trials. *NeuroImage* 33: 644-659.

Bruce CJ and Goldberg ME. 1985. Primate frontal eye fields. I. Single neurons discharging before saccades. *J Neurophysiol* 53(3): 603-635.

Bullmore E, Long C, Suckling J, Fadili J, Calvert G, Zelaya F, Carpenter TA, and Brammer M. 2001. Colored noise and computational inference in neurophysiological (fMRI) time series analysis: resampling methods in time and wavelet domains. *Hum Brain Mapp* 12(2): 61-78.

Cabeza R and Nyberg L. 2000. Imaging cognition II: An empirical review of 275 PET and fMRI studies. *J Cogn Neurosci* 12(1): 1-47.

Connolly JD, Goodale MA, Menon RS, and Munoz DP. 2002. Human fMRI evidence for the neural correlates of preparatory set. *Nat Neurosci* 5(12): 1345-1352.

Corbetta M, Akbudak E, Conturo TE, Snyder AZ, Ollinger JM, Drury HA, Linenweber MR, Petersen SE, Raichle ME, Van Essen DC, and Shulman GL. 1998. A common network of functional areas for attention and eye movements. *Neuron* 21(4): 761-773.

Corbetta M, Kincade JM, Ollinger JM, McAvoy MP, and Shulman GL. 2000. Voluntary orienting is dissociated from target detection in human posterior parietal cortex. *Nat Neurosci* 3(3): 292-297.

Curtis CE and D'Esposito M. 2003. Success and failure suppressing reflexive behavior. *J Cogn Neurosci* 15(3): 409-418.

Curtis CE, Cole MW, Rao VY, and D'Esposito M. 2005. Canceling planned action: an FMRI study of countermanding saccades. *Cereb Cortex* 15(9): 1281-1289.

Curtis CE and D'Esposito M. 2006. Selection and maintenance of saccade goals in the human frontal eye fields. *J Neurophysiol* 95(6): 3923-3927.

Dale AM. 1999. Optimal experimental design for event-related fMRI. *Hum Brain Mapp* 8(2-3): 109-114.

Desimone R and Duncan J. 1995. Neural mechanisms of selective visual attention. *Annu Rev Neurosci* 18: 193-222.

Desouza JF, Menon RS, and Everling S. 2003. Preparatory set associated with pro-saccades and anti-saccades in humans investigated with event-related fMRI. *J Neurophysiol* 89(2): 1016-1023.

Dorris MC, Pare M, and Munoz DP. 1997. Neuronal activity in monkey superior colliculus related to the initiation of saccadic eye movements. *J Neurosci* 17(21): 8566-8579.

Duncan J and Owen AM. 2000. Common regions of the human frontal lobe recruited by diverse cognitive demands. *Trends Neurosci* 23(10): 475-483.

Evarts EV, Shinoda Y, and Wise SP. 1984. *Neurophysiological approaches to higher brain function*. New York: Wiley.

Everling S, Krappmann P, and Flohr H. 1997. Cortical potentials preceding pro- and antisaccades in man. *Electroencephalogr Clin Neurophysiol* 102(4): 356-362.

Everling S and Fischer B. 1998. The antisaccade: a review of basic research and clinical studies. *Neuropsychologia* 36(9): 885-899.

Everling S, Spantekow A, Krappmann P, and Flohr H. 1998. Event-related potentials associated with correct and incorrect responses in a cued antisaccade task. *Exp Brain Res* 118(1): 27-34.

Everling S, Dorris MC, Klein RM, and Munoz DP. 1999. Role of primate superior colliculus in preparation and execution of anti-saccades and pro-saccades. *J Neurosci* 19(7): 2740-2754.

Everling S and Munoz DP. 2000. Neuronal correlates for preparatory set associated with pro-saccades and anti-saccades in the primate frontal eye field. *J Neurosci* 20(1): 387-400.

Everling S and DeSouza JF. 2005. Rule-dependent activity for prosaccades and antisaccades in the primate prefrontal cortex. *J Cogn Neurosci* 17(9): 1483-1496.

Fischer B and Weber H. 1992. Characteristics of "anti" saccades in man. *Exp Brain Res* 89(2): 415-424.

Fischer B and Weber H. 1997. Effects of stimulus conditions on the performance of antisaccades in man. *Exp Brain Res* 116(2): 191-200.

Forbes K and Klein RM. 1996. The magnitude of the fixation offset effect with endogenously and exogenously controlled saccades. *J Cogn Neurosci* 8: 344-352.

Ford KA, Goltz HC, Brown MR, and Everling S. 2005. Neural processes associated with antisaccade task performance investigated with event-related FMRI. *J Neurophysiol* 94(1): 429-440.

Gaymard B, Rivaud S, Cassarini JF, Dubard T, Rancurel G, Agid Y, and Pierrot-Deseilligny C. 1998. Effects of anterior cingulate cortex lesions on ocular saccades in humans. *Exp Brain Res* 120(2): 173-183.

Gaymard B, Ploner CJ, Rivaud-Pechoux S, and Pierrot-Deseilligny C. 1999. The frontal eye field is involved in spatial short-term memory but not in reflexive saccade inhibition. *Exp Brain Res* 129(2): 288-301.

Gaymard B, Francois C, Ploner CJ, Condy C, and Rivaud-Pechoux S. 2003. A direct prefrontotectal tract against distractibility in the human brain. *Ann Neurol* 53(4): 542-545.

Gottlieb J and Goldberg ME. 1999. Activity of neurons in the lateral intraparietal area of the monkey during an antisaccade task. *Nat Neurosci* 2(10): 906-912.

Guitton D, Buchtel HA, and Douglas RM. 1985. Frontal lobe lesions in man cause difficulties in suppressing reflexive glances and in generating goal-directed saccades. *Exp Brain Res* 58(3): 455-472.

Hallett PE. 1978. Primary and secondary saccades to goals defined by instructions. *Vision Res* 18(10): 1279-1296.

Hallett PE and Adams BD. 1980. The predictability of saccadic latency in a novel voluntary oculomotor task. *Vision Res* 20(4): 329-339.

Hebb DO. 1972. *Textbook of physiology*. Philadelphia: Saunders.

Hon N, Epstein RA, Owen AM, and Duncan J. 2006. Frontoparietal activity with minimal decision and control. *J Neurosci* 26(38): 9805-9809.

Johnston K and Everling S. 2006a. Neural activity in monkey prefrontal cortex is modulated by task context and behavioral instruction during delayed-match-to-sample and conditional prosaccade-antisaccade tasks. *J Cogn Neurosci* 18(5): 749-765.

Johnston K and Everling S. 2006b. Monkey dorsolateral prefrontal cortex sends task-selective signals directly to the superior colliculus. *J Neurosci* 26(48): 12471-12478.

Johnston K, Levin HM, Koval MJ, and Everling S. 2007. Top-Down Control-Signal Dynamics in Anterior Cingulate and Prefrontal Cortex Neurons following Task Switching. *Neuron* 53(3): 453-462.

Kastner S, Desimone K, Konen CS, Szczepanski SM, Weiner KS, and Schneider KA. 2007. Topographic maps in human frontal cortex revealed in memory-guided saccade and spatial working-memory tasks. *J Neurophysiol* 97(5): 3494-3507.

Klassen LM and Menon RS. 2004. Robust automated shimming technique using arbitrary mapping acquisition parameters (RASTAMAP). *Magn Reson Med* 51(5): 881-887.

Klein C, Heinks T, Andresen B, Berg P, and Moritz S. 2000. Impaired modulation of the saccadic contingent negative variation preceding antisaccades in schizophrenia. *Biol Psychiatry* 47(11): 978-990.

Leichnetz GR, Spencer RF, Hardy SG, and Astruc J. 1981. The prefrontal corticotectal projection in the monkey; an anterograde and retrograde horseradish peroxidase study. *Neuroscience* 6(6): 1023-1041.

Medendorp WP, Goltz HC, and Vilis T. 2005. Remapping the remembered target location for anti-saccades in human posterior parietal cortex. *J Neurophysiol* 94(1): 734-740.

Miller EK and Cohen JD. 2001. An integrative theory of prefrontal cortex function. *Annu Rev Neurosci* 24: 167-202.

Munoz DP and Everling S. 2004. Look away: the anti-saccade task and the voluntary control of eye movement. *Nat Rev Neurosci* 5(3): 218-228.

Ollinger JM, Corbetta M, and Shulman GL. 2001a. Separating processes within a trial in event-related functional MRI. *NeuroImage* 13(1): 218-229.

Ollinger JM, Shulman GL, and Corbetta M. 2001b. Separating processes within a trial in event-related functional MRI. *NeuroImage* 13(1): 210-217.

Pierrot-Deseilligny C, Rivaud S, Gaymard B, and Agid Y. 1991. Cortical control of reflexive visually-guided saccades. *Brain* 114: 1473-1485.

Pierrot-Deseilligny C, Ploner CJ, Muri RM, Gaymard B, and Rivaud-Pechoux S. 2002. Effects of cortical lesions on saccadic: eye movements in humans. *Ann N Y Acad Sci* 956: 216-229.

Pierrot-Deseilligny C, Muri RM, Ploner CJ, Gaymard B, Demeret S, and Rivaud-Pechoux S. 2003. Decisional role of the dorsolateral prefrontal cortex in ocular motor behaviour. *Brain* 126(Pt 6): 1460-1473.

Ploner CJ, Gaymard BM, Rivaud-Pechoux S, and Pierrot-Deseilligny C. 2005. The prefrontal substrate of reflexive saccade inhibition in humans. *Biol Psychiatry* 57(10): 1159-1165.

Posner MI and Petersen SE. 1990. The attention system of the human brain. *Annu Rev Neurosci* 13: 25-42.

Schall JD, Hanes DP, Thompson KG, and King DJ. 1995. Saccade target selection in frontal eye field of macaque. I. Visual and premovement activation. *J Neurosci* 15(10): 6905-6918.

Schall JD. 1997. Visuomotor areas of the frontal lobe. In: *Cerebral Cortex*, edited by Rockland K, Peters A, and Kaas JH. New York: Plenum, p. 527-638.

Schluppeck D, Curtis CE, Glimcher PW, and Heeger DJ. 2006. Sustained activity in topographic areas of human posterior parietal cortex during memory-guided saccades. *J Neurosci* 26(19): 5098-5108.

Seidemann E, Arieli A, Grinvald A, and Slovin H. 2002. Dynamics of depolarization and hyperpolarization in the frontal cortex and saccade goal. *Science* 295(5556): 862-865.

Serences JT. 2004. A comparison of methods for characterizing the event-related BOLD timeseries in rapid fMRI. *NeuroImage* 21(4): 1690-1700.

Sereno MI, Pitzalis S, and Martinez A. 2001. Mapping of contralateral space in retinotopic coordinates by a parietal cortical area in humans. *Science* 294(5545): 1350-1354.

Wang Y, Matsuzaka Y, Shima K, and Tanji J. 2004. Cingulate cortical cells projecting to monkey frontal eye field and primary motor cortex. *Neuroreport* 15(10): 1559-1563.

Worsley KJ, Liao CH, Aston J, Petre V, Duncan GH, Morales F, and Evans AC. 2002. A general statistical analysis for fMRI data. *NeuroImage* 15(1): 1-15.

Zhang M and Barash S. 2000. Neuronal switching of sensorimotor transformations for antisaccades. *Nature* 408(6815): 971-975.

Zhang M and Barash S. 2004. Persistent LIP activity in memory-antisaccades: working memory for a sensorimotor transformation. *J Neurophysiol* 91(3): 1424-1441.

Chapter 4
Isolation of Saccade Inhibition Processes: Rapid Event-related fMRI of Saccades and Nogo Trials

4.1 – Introduction

The ability to inhibit automatic behaviours that are inappropriate for a given situation in favour of more suitable voluntary responses is an important component of adaptive behavioural control. This topic has been studied using the eye movement system extensively as a model. For example, the antisaccade task requires subjects to inhibit the automatic saccade toward a flashed peripheral stimulus and then to generate a voluntary saccade to the stimulus' mirror location in the opposite visual hemifield (Hallett 1978; Munoz and Everling 2004). According to popular models of behavioural control, fronto-parietal areas provide bias signals to other brain regions to support task-processing (Desimone and Duncan 1995; Duncan 2001; Miller and Cohen 2001).

Several previous studies have compared the saccade and nogo tasks to investigate inhibition of automatic saccade responses. The saccade task, which requires the subject to look at a flashed peripheral visual stimulus, engages the visual grasp reflex, the tendency of primates to look automatically at novel visual stimuli (Hess et al. 1946). The nogo task requires the subject to inhibit this automatic tendency in favour of maintaining central visual fixation (Machado and Rafal 2000; Sommer and Wurtz 2001). Like the antisaccade task, the nogo task recruits response inhibition processes, but unlike the antisaccade task, the nogo task does not involve an explicit motor component, making it easier to isolate response inhibition-related signals.

Functional magnetic resonance imaging (fMRI) studies have also compared saccade and nogo trials (Brown et al. 2006; Neggers et al. 2005). In a previous study (Brown et al. 2006), we found that these two tasks evoked essentially identical fMRI signal patterns in cortical saccade regions, despite their very different response requirements. To explain this result, we suggested that fMRI activation in cortical

saccade regions might reflect stimulus-related processes, such as visuospatial attention and stimulus selection, rather than processes related to response execution, such as generation or inhibition of the eye movements, themselves. We are aware of only one other fMRI study whose design included saccades and nogo trials (Neggers et al. 2005), and this study, which focused on the gap effect rather than saccade inhibition, did not include a direct comparison of saccades and nogo trials.

Curtis et al. (2005) also studied response inhibition with fMRI by comparing saccades and countermanded saccades. The countermanding task presents the subject with a visual stimulus toward which he or she must make a saccade unless a countermanding signal (in Curtis et al.'s case an auditory stimulus) is presented after peripheral stimulus onset, instructing the subject to cancel the developing saccade command. Curtis et al. (2005) found greater activation for countermanded saccades versus non-countermanded saccades in cortical saccade areas, in contrast to our previous results from Brown et al. (2006). This difference might be explained by differences in when the response inhibition instruction was presented. Subjects were instructed to withhold the saccade after stimulus onset in Curtis et al. (2005) and before stimulus onset in Brown et al. (2006). Therefore, it was probably harder for subjects to cancel a partially-developed saccade command in Curtis and colleagues' countermanding task than to pre-emptively inhibit the initiation of the automatic saccade in our nogo task. This raises the possibility that we might have found activation differences between saccades and nogo trials in Brown et al. (2006) if the nogo task had had a stronger response inhibition component.

In the current study, we tested the hypothesis from Brown et al. (2006) that fMRI activation in cortical saccade regions primarily reflects stimulus-related processes, such as stimulus detection and attention, rather than response-related processes, such as saccade inhibition or generation. We compared saccades and nogo trials using a rapid event-related fMRI design which included twice as many saccade trials as nogo trials. The higher frequency of saccade trials was intended to increase the prepotency of the saccade response, thereby also increasing the need for active inhibition of the automatic saccade in the rare nogo trials. Comparing a frequent and therefore prepotent go response with a rare nogo type response is a standard procedure found in many response inhibition

studies using button press tasks (see for example Braver et al. 2001; de Zubicaray et al. 2000; Fassbender et al. 2004; Garavan et al. 1999; Garavan et al. 2002; Horn et al. 2003; Kelly et al. 2004; Rubia et al. 2005; Wager et al. 2005). Decreasing the frequency of the nogo trial type also increases the rate of commission errors (Wager et al. 2005). In this way, we intended to increase the recruitment of inhibitory processes in the nogo trials in the current experiment compared to the nogo trials used in our previous study (Brown et al. 2006). We expected that increasing the demand for response-related processes associated with saccade inhibition in the nogo task would lead to greater fMRI signal levels for nogo trials compared to saccade trials, given that the two trial types otherwise seem to evoke attention and detection processes with similar intensity (Brown et al. 2006).

4.2 – Methods

All procedures were approved by the University Research Ethics Board for Health Sciences Research at the University of Western Ontario, London, Ontario, Canada and are in accordance with the 1964 Declaration of Helsinki.

4.2.1 – Subjects

This study included eleven human subjects (5 female, 6 male), all of whom gave informed written consent. Subjects ranged in age from 21 to 35 years with a mean of 27 years. Subjects reported no history of neurological or psychiatric disorder, and they had normal or corrected to normal vision. Three of the subjects described themselves as left-handed, and the rest described themselves as right-handed.

4.2.2 – fMRI Data Acquisition Procedure

A whole body 4 Tesla MRI system (Varian, Palo Alto, CA; Siemens, Erlangen, Germany) was used to collect fMRI data. The system operated at a slew rate of 120 T/m/s and 40 mT/m gradients. A transmit only receive only (TORO) cylindrical hybrid birdcage radio frequency (RF) head coil (Barberi et al. 2000) was used for transmission and detection of the signal.

A series of sagittal anatomical images acquired using T1-weighting was used to define imaging planes for the functional scans. Eleven contiguous functional planes, each 5 mm thick, were prescribed axially from the top of the brain to the level of the dorsal caudate and thalamus. A constrained, three-dimensional phase shimming procedure (Klassen and Menon 2004) was employed to optimize the magnetic field homogeneity over the functional volume. During each functional task, blood oxygenation level-dependent (BOLD) images (T_2*-weighted) were acquired continuously using an interleaved, two segment, optimized spiral imaging protocol (volume collection time = 1.0 s, repeat time (TR) = 500 ms, echo time (TE) = 15 ms, flip angle = 30º, 64 x 64 matrix size, 22.0 x 22.0 cm field of view (FOV), 3.44 x 3.44 x 5.00 mm voxel resolution). Navigator echo correction was used to correct each image for physiological fluctuations. A corresponding high-resolution T1-weighted structural volume was acquired during the same scanning session using a 3D spiral acquisition protocol (TE = 3.0 ms, inversion time (TI) = 1300 ms, TR = 50 ms) with a voxel resolution of 0.9 x 0.9 x 1.25 mm. Each subject was immobilized during the experimental session within a head cradle packed with foam padding.
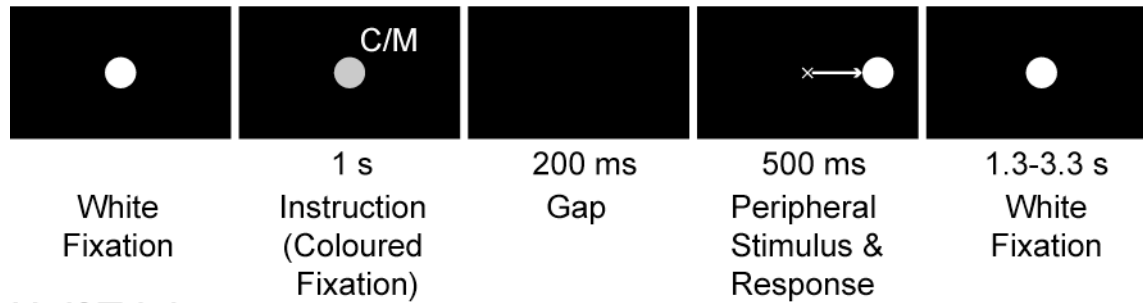
4.2.3 – Eye Tracking

We presented visual stimuli during fMRI scanning using a Silent Vision SV-4021 fibre optic projection system from Avotec (Stuart, Fl, USA). The system also includes an MEyeTrack-SV (Silent Vision) eye tracker from SensoMotoric Instruments GmbH (Teltow, Germany). The presentation and eye tracking equipment is housed in dual stalks positioned over the subject's eyes to allow simultaneous presentation of visual stimuli and CCD video-based infrared eye tracking. The visual display subtends 30° horizontally by 23° vertically with a resolution of 800 x 600 pixels and a refresh rate of 60 Hz. Eye tracking was also performed at a 60 Hz sampling rate with an accuracy of approximately 1 degree. Before scanning, we calibrated the system with a 9 point calibration, with four points in the corners, four points at the midpoints of the sides, and one point in the centre of the visual display. We verified that all experimental targets were within the range of calibration. Analysis of eye movement signals was performed offline using custom software built in Matlab 6 (The Mathworks, Inc., Natick, Ma, USA).
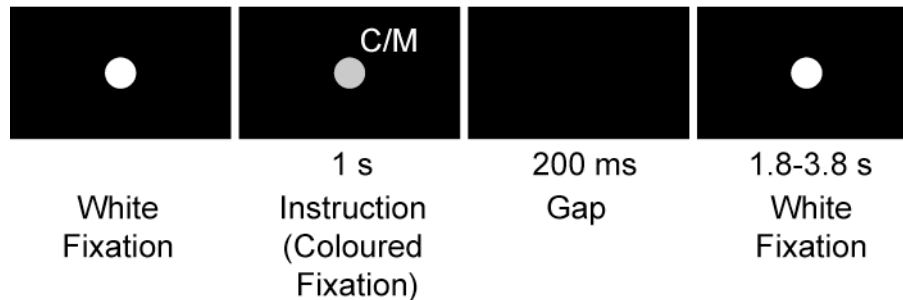
4.2.4 – Experimental Design

We compared the saccade and nogo tasks in a rapid event-related fMRI design. The saccade task required the subjects to look toward a peripheral visual stimulus, while the nogo task required him or her to inhibit the automatic saccade toward the stimulus. Saccade and nogo trials were presented in whole and half trial versions (*Figure 4.1*), as explained below. Between trials, the subject fixated a central white dot (0.6° diameter). An individual trial began when the fixation point changed from white to magenta or cyan, with the colour indicating the current trial's type: saccade or nogo trial. Magenta and cyan fixation points instructed six of the subjects to perform the nogo and saccade tasks, respectively, and the colour-task assignment was reversed for the other five subjects. The instruction period, which occupied the interval from the onset of the coloured fixation point until its offset, lasted for 1 sec. On saccade and nogo whole trials, a peripheral visual stimulus (0.6° diameter white dot) was presented 200 ms after the coloured fixation offset. During the intervening 200 ms gap, the subject was shown a black screen. The 200 msec gap was included because it is known to increase the difficulty of inhibiting the automatic saccade response toward an abrupt peripheral stimulus (Fischer and Weber 1992; Forbes and Klein 1996; Fischer and Weber 1997; Bell et al. 2000). Presumably, the presence of the gap increases the requirement for inhibitory control in the nogo task. On whole trials, the peripheral stimulus was presented for 500 ms to the left or right of fixation, at peripheral eccentricities of 6°, 8°, or 10°, either on the horizontal meridian or 20° clockwise or counterclockwise relative it. After 500 ms, the peripheral stimulus disappeared, and the central white fixation point reappeared. On whole saccade trials, the subject had to look at the peripheral stimulus before it disappeared and then return gaze to centre. On whole nogo trials, the subject had to inhibit the automatic saccade toward the peripheral stimulus while maintaining central fixation. Two thirds of saccade and nogo trials were whole trials, while the remaining third were half trials. On half trials, the central white fixation dot reappeared at the end of the 200 msec gap. The subject presumably prepared to perform the saccade or nogo task based on the fixation point's colour in the preceding instruction interval.

Whole Trial

| White Fixation | 1 s Instruction (Coloured Fixation) | 200 ms Gap | 500 ms Peripheral Stimulus & Response | 1.3-3.3 s White Fixation |

Half Trial

| White Fixation | 1 s Instruction (Coloured Fixation) | 200 ms Gap | 1.8-3.8 s White Fixation |

Figure 4.1 – Experimental Design

Timing and visual stimulus presentation for saccade and nogo whole and half trials. The gray fixation point and "C/M" in the instruction panel (second from left) denote that the white fixation point changed to magenta or cyan to instruct the subject whether to make a saccade or nogo response. The assignment of fixation colour to trial type varied from subject to subject. The small white arrow and white X in the fourth panel for the whole trial indicate a rightward saccade response and a nogo response, namely central fixation maintenance. Whole trials contained both an instruction and response event, whereas half trials contained only the instruction. On half trials, subjects never saw a peripheral stimulus, nor did they make a saccade or active nogo fixation response. The half trials allowed us to measure instruction-related fMRI signals separately from response-related signals. See *Section 4.2 – Methods* for details.

However, no peripheral stimulus was presented, and the subject was not required to make an actual saccade or nogo response on half trials. The half trials provided a measure of instruction-related fMRI activation that was separate from response-evoked activation.

Half and whole trials were presented at a 1:2 ratio to prevent subjects from anticipating that a given trial was a half trial, which would have reduced any preparatory activity during the instruction period. The time interval between the starts of adjacent trials was 3, 4, or 5 s, randomized with a mean of 4 s (uniform distribution). This jittering of the trial-to-trial interval was effected by varying the duration of central white fixation at the end of the trial (*Figure 4.1*, right-most panel for both whole and half trials). At the end of whole trials, central white fixation was presented for 1.3, 2.3, or 3.3 s, and at the end of half trials, which did not include the 500 ms peripheral stimulus, central white fixation was presented for 1.8, 2.8, or 3.8 s. Thus, the time interval from the start of any given trial to the start of the next trial was 3, 4, or 5 s. Jittering the trial-to-trial interval was done to decorrelate the fMRI signal components evoked by individual trial types allowing for subsequent deconvolution as described below (Dale 1999).

Each subject completed 7 to 12 data collection runs (median 10 runs; total 106 runs across all 11 subjects). An individual run lasted 352 s and consisted of 82 trials (36 whole saccade trials, 18 half saccade trials, 18 whole nogo trials, 9 half nogo trials, and the first trial of the run). As was discussed in the Introduction, we included twice as many saccade trials as nogo trials to increase the prepotency of the saccade response and to increase the requirement for inhibitory control of the automatic saccade in the rare nogo trials. The first trial of a run, which was always a whole saccade trial, was modeled separately in the analysis because it was not preceded by a trial and, therefore, was not part of the set of counterbalanced trials in which each trial type (saccade vs. nogo, half trial vs. whole trial) was preceded proportionately by every other trial type. Each of the 18 possible peripheral stimulus locations (see above) was used twice for the whole saccade trials and once for the whole nogo trials in a given run. In addition, the first trial of the run always used a peripheral stimulus with 6° eccentricity and 20° clockwise relative to the horizontal meridian on the right side. The order and timing structure of the trial presentation sequence within each run was determined using computerized random search to minimize correlation among different trial types and to ensure first order counterbalancing of trial parameters (saccade vs. nogo trials, half vs. whole trials, left vs. right peripheral stimulus) in the forward and backward directions.

4.2.5 – Behavioural Analysis

We used video eye tracking data collected concurrently with fMRI scanning to mark subject performance and to determine saccadic response latencies. We discarded trials with several types of mistakes including error responses, failure to respond when required, excessively short or long response latency (latency < 100 ms or latency > 500 ms, respectively), and failure to maintain fixation when required. The average proportion of discarded trials across all 11 subjects was 3.7% ± 3.2% (standard deviation), with a range of 0.3% to 11.0%. It would have been interesting to analyze fMRI activation patterns for error trials, in which subjects failed to generate saccades on saccade trials or incorrectly generated saccades on nogo trials. Unfortunately, subjects made too few errors for a statistically sound examination (see *Section 4.3.1 – Behaviour* for performance data). We did attempt to deconvolve BOLD time courses for nogo error trials (see below for discussion of deconvolution), but these were very noisy because of the small samples sizes, particularly in 7 of the subjects from whom we collected fewer than 10 nogo error trials each.

4.2.6 – Preprocessing

Preprocessing was performed using Matlab 6 (The Mathworks, Inc., Natick, Ma, USA) and BrainVoyager 2000 (Brain Innovation BV, Maastricht, The Netherlands). The first four volumes from each voxel's time course were removed using Matlab 6 to allow for complete spin saturation. We performed nonlinear trend removal on the data using a running lines smoother programmed in Matlab 6. This algorithm corrected time courses for low frequency trends as described in *Appendix 4.1* and has the advantage of being robust against edge effects. All functional data were superimposed onto 3D anatomical images, resampled into 3x3x3 mm cubic voxels, aligned onto the anterior commissure – posterior commissure axis, and scaled into Talairach space with BrainVoyager 2000. Functional data then underwent 3D motion correction by trilinear interpolation of each volume to a single reference functional volume and spatial low-pass filtering, performed with a low-pass cutoff of 0.063 cycles / mm in the Fourier domain. Finally, voxel time courses were scaled into percent signal change values on a run-by-run basis using the mean time course across the run as baseline.

4.2.7 – Statistical Analysis and Linear Deconvolution

The close spacing of task events in our rapid event-related fMRI design necessitated the use of deconvolution to separate activation profiles for individual event types. That is, because the trial onset spacing (3-5 s) was shorter than the 12 to 14 second evolution time of the haemodynamic response (Boynton et al. 1996), the BOLD signal we recorded at any given time point was the sum of the BOLD impulse responses from several preceding task events. We used linear regression based on the general linear model (GLM) to perform both the statistical analysis and the linear deconvolution. A GLM is fit to a dataset by scaling its component predictor curves so as to minimize mean squared error between the estimated and actual data curves. This process automatically deconvolves modeled task components because it finds the best scaling of predictor curves which model overlapping, summated BOLD signal components.

We used a mixed effects analysis based on Worsley and colleagues (2002). A GLM was fit to each subject's data after pre-whitening to correct for coloured noise covariance (see *Appendix 4.2*). Pre-whitening avoids the inflation of type I error that occurs with GLM statistical tests that fail to account for coloured noise in the data (Bullmore et al. 2001). Our GLM design matrix incorporated sets of finite impulse response predictor functions (see Serences 2004) for 1) correct half saccade trials, 2) correct whole saccade trials, 3) correct half nogo trials, 4) correct whole nogo trials, 5) error whole nogo trials (when subjects made saccades rather than nogo responses), 6) the first trial of the run, which was always a whole saccade trial, 7) any saccade or nogo trials with responses greater than 500 msec in latency, and 8) other discarded trials with latencies within the 100 to 500 msec range. That is, for each of these groups of trials, a set of 12 impulse responses (one impulse per 1 sec volume) was locked to the start of each trial in the group. The finite impulse response model has the advantage that it does not incorporate assumptions about the shape of the haemodynamic response function other than its length (12 sec in this case). The use of linear regression also assumes that the haemodynamic response is linear. It was necessary to model discarded trials explicitly in the design matrix because the fMRI signals they evoked would otherwise contaminate the deconvolution of signals from adjacent correct trials.

For each of the eleven subjects, we computed two statistical parametric maps, one based on instruction differences and one based on response differences. The instruction contrast was built by subtracting saccade half trial activation from nogo half trial activation over a time interval from 4-6 s after the instruction onset. This interval was chosen beforehand because it occupies the peak of the haemodynamic impulse response curve as described previously (Boynton et al. 1996). The response contrast was computed in two steps. For the saccade and nogo tasks separately, half trial activation was first subtracted from whole trial activation to produce measures of response-related activation. Saccade response activation was then subtracted from nogo response activation over the time interval from 5-7 s after instruction onset. The response contrast time interval was 1 s later than the instruction contrast interval because the onset of the peripheral stimulus which evoked a saccade or nogo response in the whole trials occurred 1.2 s after the instruction stimulus onset.

For each of the statistical contrasts, individual subject maps were combined using a mixed effects analysis (Worsley et al. 2002), which addressed variability both within subjects and between subjects. Mixed effects contrast maps were thresholded at a T-value of 2.0 ($p < 0.05$ uncorrected, df = 15000) and then cluster size thresholded at a minimum size of 1184 cubic mm to correct for multiple comparisons across the voxel population at $p < 0.05$, based on Gaussian random field theory (see Worsley et al. 2002). For details of how the mixed effects analysis works, see *Appendix 4.3* and Worsley et al. (2002).

4.3 – Results

We compared frequent saccades and rare nogo trials (2:1 ratio saccades : nogo trials) using rapid event-related fMRI to investigate processes underlying inhibition of the automatic saccade response in the nogo trials.

4.3.1 – Behaviour

We recorded subjects' eye positions using video eye tracking during fMRI data acquisition to verify that subjects performed correctly in the experiment. Compared to widely-spaced fMRI designs, the timing structure of our rapid trial design was much

more similar to typical psychophysical and electrophysiological designs, which should facilitate comparison between results from these modalities and those from this study. However, it was still possible that the fMRI scanning environment might have altered subject performance in some way, and the eye position recordings allowed us to eliminate this possibility. Correct performance in the saccade and nogo tasks, respectively, involved making or withholding saccades toward peripheral stimuli. In saccade error trials, subjects failed to initiate a saccade, and in nogo error trials, they failed to inhibit the automatic saccade. Of course, it is impossible to distinguish whether a given saccade error trial occurred because of a simple lapse in attention or because the subject accidentally made a voluntary nogo response. Mean percent correct values across all 11 subjects were $97.6 \pm 2.7\%$ (standard deviation) for saccades and $92.9 \pm 6.1\%$ for nogo trials, and these values were significantly different ($p = 0.012$, $T = 3.05$, $df = 10$). Lesser percent correct performance values for rare nogo trials compared to frequent go trials (that is, saccade trials in this study) are consistent with the response inhibition literature (Braver et al. 2001; de Zubicaray et al. 2000; Fassbender et al. 2004; Garavan et al. 1999; Garavan et al. 2002; Horn et al. 2003; Kelly et al. 2004; Maguire et al. 2003; Menon et al. 2001; Mostofsky et al. 2003; Rubia et al. 2001; Rubia et al. 2005; Wager et al. 2005). Mean saccadic latencies over all subjects were 257 ± 30 ms for correct saccade trials and 274 ± 80 ms for error nogo trials. These values were not significantly different ($p > 0.5$, $T = 0.68$, $df = 10$). Correct nogo trials, of course, did not have responses on which to measure latencies.

4.3.2 – Instruction-related Results

We included saccade and nogo half trials (Ollinger et al. 2001a; Ollinger et al. 2001b), which were comprised only of the instruction event, as a means of separating instruction- from response-related activation profiles (see *Figure 4.1* and *Section 4.2.4 – Experimental Design*). To investigate instruction-related differences between the two tasks, we statistically contrasted activation from nogo half trials with that from saccade half trials over a time interval from 4-6 seconds (see *Section 4.2.7 – Statistical Analysis and Linear Deconvolution*). A subset of slices from the resulting statistical map is shown in *Figure 4.2A*. Size and location data for statistically significant voxel clusters are shown

in *Table 4.1A*. We found greater instruction-related activation for nogo versus saccade trials in right frontal eye field (FEF), left and right middle frontal gyrus (MFG), left rostral intraparietal sulcus (IPS), left and right caudal IPS, and precuneus. Greater nogo trial instruction-related activation is also evident in the half trial time courses for right FEF and right MFG, which are displayed in *Figure 4.3A*.

Note that instruction differences in frontal cortex were partially lateralized. The right MFG voxel cluster was much larger than the left MFG equivalent (10,661 and 2,862 cubic mm for right and left MFG, respectively, see *Table 4.1A*), and right FEF but not left FEF exhibited statistically significant instruction differences.

4.3.3 – Response-related Results

Whole saccade and nogo trials each contained an instruction and a response event. To isolate response-related activation patterns, we subtracted deconvolved time courses for the half trials, which contained only the instruction and not the response event, from the whole trial time courses. Activation time courses for half trials, whole trials, and the response-related activation curves generated by subtracting half trials from whole trials are shown, respectively, in *Figure 4.3A*, *B*, and *C* for select regions.

We contrasted nogo trial versus saccade response-related activation over a time interval from 5-7 seconds (see *Section 4.2.7 – Statistical Analysis and Linear Deconvolution*). Certain slices of the response contrast map are shown in *Figure 4.2B*, and size and location data for all statistically significant voxel clusters are listed in *Table 4.1B*. We found large voxel clusters with greater nogo trial versus saccade response-related activation (red regions in *Figure 4.2B*) in left and right inferior frontal gyrus (IFG), supplementary eye field (SEF), anterior cingulate cortex (ACC), right postcentral gyrus, right supramarginal gyrus (SMG), and right posterior inferior parietal lobule (IPL). *Figure 4.3C* shows response-related activation profiles with greater activation for nogo responses in right IFG, SEF, and right SMG. ACC time courses (not shown) were similar to SEF time courses. There was also a region with greater nogo response activation extending from the right posterior superior frontal gyrus (SFG) into the right dorsal precentral gyrus.

Figure 4.2 – Instruction and Response Contrast Maps

A: Instruction difference map built by contrasting nogo half trials with saccade half trial. Yellow/red regions exhibited greater nogo activation, and blue regions exhibited greater saccade activation. Nogo response activation was greater in right frontal eye field (FEF), bilateral middle frontal gyrus (MFG), left rostral intraparietal sulcus (IPS), bilateral caudal IPS, and precuneus. Maps are mixed effects maps across 11 subjects. Images obey neurological convention. The yellow lines in the sagittal section demark the edges of the brain volume included in our scanning profile (see *Section 4.2.2 – fMRI Data Acquisition Procedure* for details).

B: Response difference map built by contrasting nogo and saccade response activation. Response activation for both trial types was computed by subtracting half trial activation from whole trial activation (see Methods and Figure 1). Nogo response activation was greater in bilateral inferior frontal gyrus (IFG), supplementary eye field (SEF), anterior cingulate cortex (ACC), right supramarginal gyrus (SMG), and right posterior inferior parietal lobule (IPL). Other conventions as in *Figure 4.2A*.
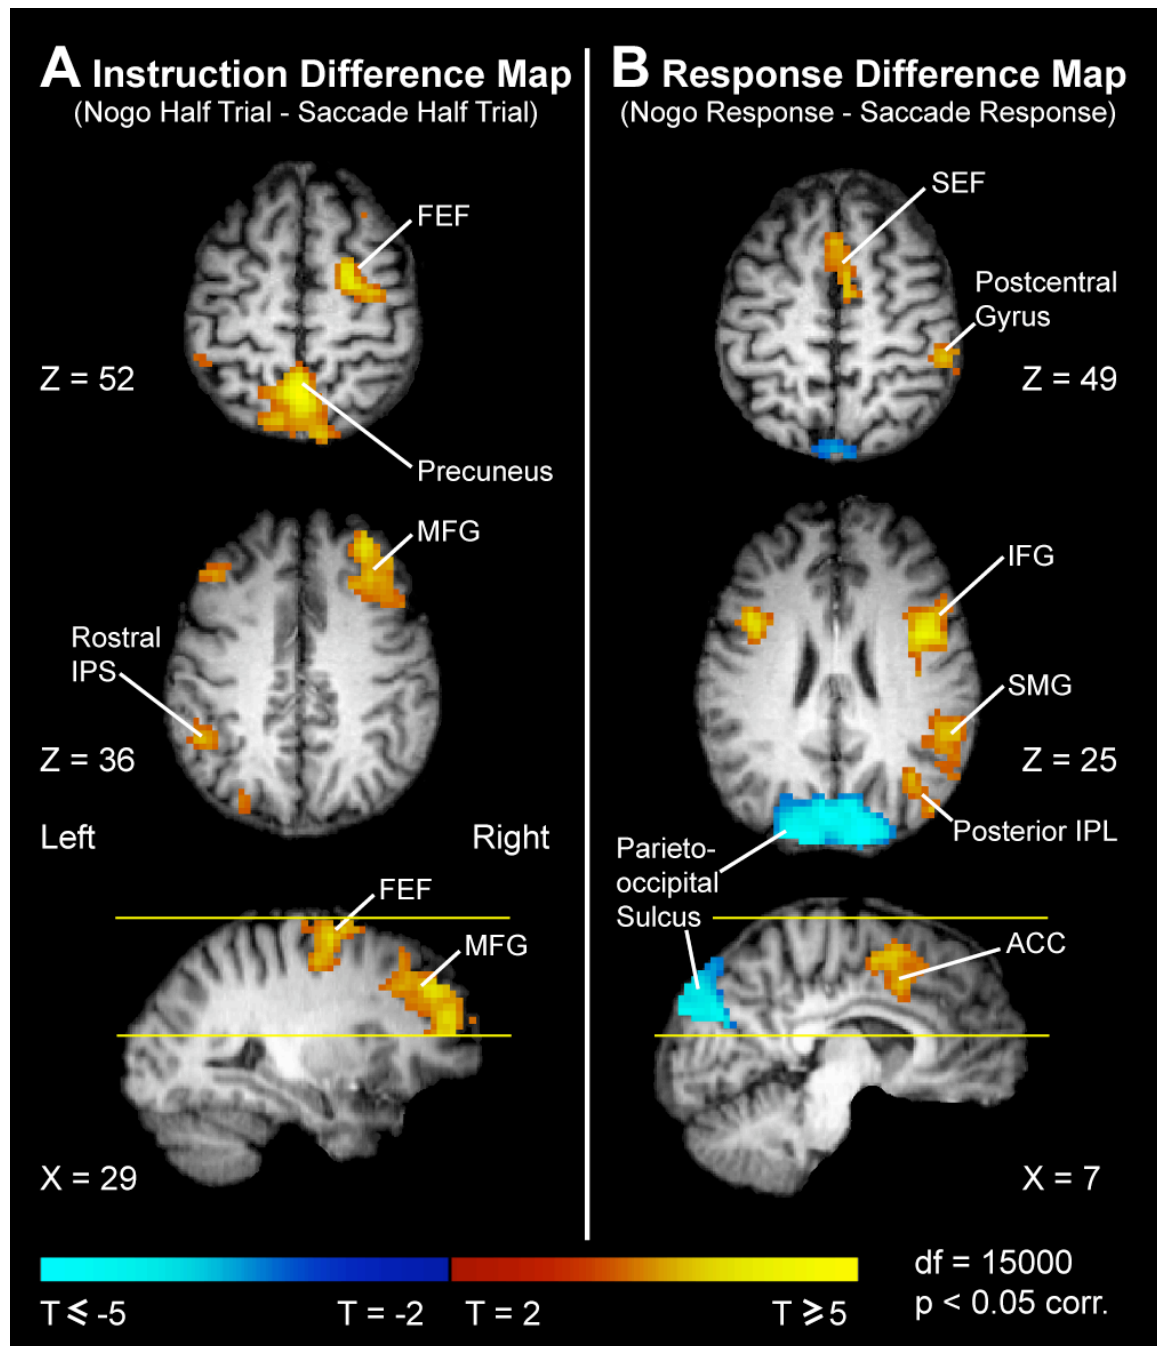
Table 4.1 – Instruction and Response Contrasts

Region	A. Instruction Contrast					B. Response Contrast				
	Dir	X	Y	Z	Volume	Dir	X	Y	Z	Volume
Right Frontal Eye Field (FEF)	+	25	-1	56	8,478					
Right Middle Frontal Gyrus (MFG)	+	32	39	32	10,611					
Left Middle Frontal Gyrus (MFG)	+	-36	44	29	2,862					
Left Rostral Intraparietal Sulcus (IPS)	+	-45	-40	38	2,133					
Right Caudal Intraparietal Sulcus (IPS)	+	16	-68	49	2,106					
Left Caudal Intraparietal Sulcus (IPS)	+	-24	-76	34	2,484					
Precuneus	+	1	-56	52	11,745					
Supplementary Eye Field (SEF)						+	5	9	52	1,809
Anterior Cingulate Cortex (ACC)						+	4	16	44	2,538
Right Superior Frontal Gyrus (SFG) / Precentral Gyrus						+	15	-2	62	1,350
Right Inferior Frontal Gyrus (IFG)						+	42	6	30	7,209
Left Inferior Frontal Gyrus (IFG)						+	-38	11	26	2,349
Right Postcentral Gyrus						+	53	-28	45	3,024
Right Supramarginal Gyrus (SMG)						+	54	-43	29	5,076
Right Posterior Inferior Parietal Lobule (IPL)						+	39	-66	27	1,593
Parieto-Occipital Sulcus						-	-1	-80	31	18,846

Legend: Region of interest data for A) instruction and B) response contrast maps. 'Dir' denotes direction of contrast: greater activation for nogo trials (+) or for saccade trials (-). X, Y, and Z are the Talairach coordinates of the statistical centres of mass in mm. Volume is in cubic mm. Regions not explicitly listed as 'left' or 'right' straddled the midline and included voxels in both hemispheres.

Figure 4.3 – Deconvolved Time Courses

Deconvolved time courses for *A*: half trials (instruction), *B*: whole trials (instruction and response), and *C*: whole trials minus half trials (response). Nogo trial and saccade time courses are depicted with solid and dashed lines, respectively. Time courses for right frontal eye field (FEF) and right middle frontal gyrus (MFG) were derived from voxels exhibiting significance in the instruction contrast map (see *Figure 4.2A*), while supplementary eye field (SEF), right inferior frontal gyrus (IFG), right supramarginal gyrus (SMG), and parieto-occipital sulcus time courses were taken from voxels exhibiting significance in the response contrast map (*Figure 4.2B*). Gray rectangles in the left and right columns denote the 4-6 s interval and the 5-7 s interval used, respectively, to build the instruction and response contrast maps in *Figure 4.2A* and *4.2B*. Asterisks denote statistically significant differences between nogo trials and saccades during the appropriate time intervals ($p < 0.05$ corrected). In columns *B* and *C*, the vertical gray lines positioned at 1.2 s denote the peripheral stimulus onset in whole trials.
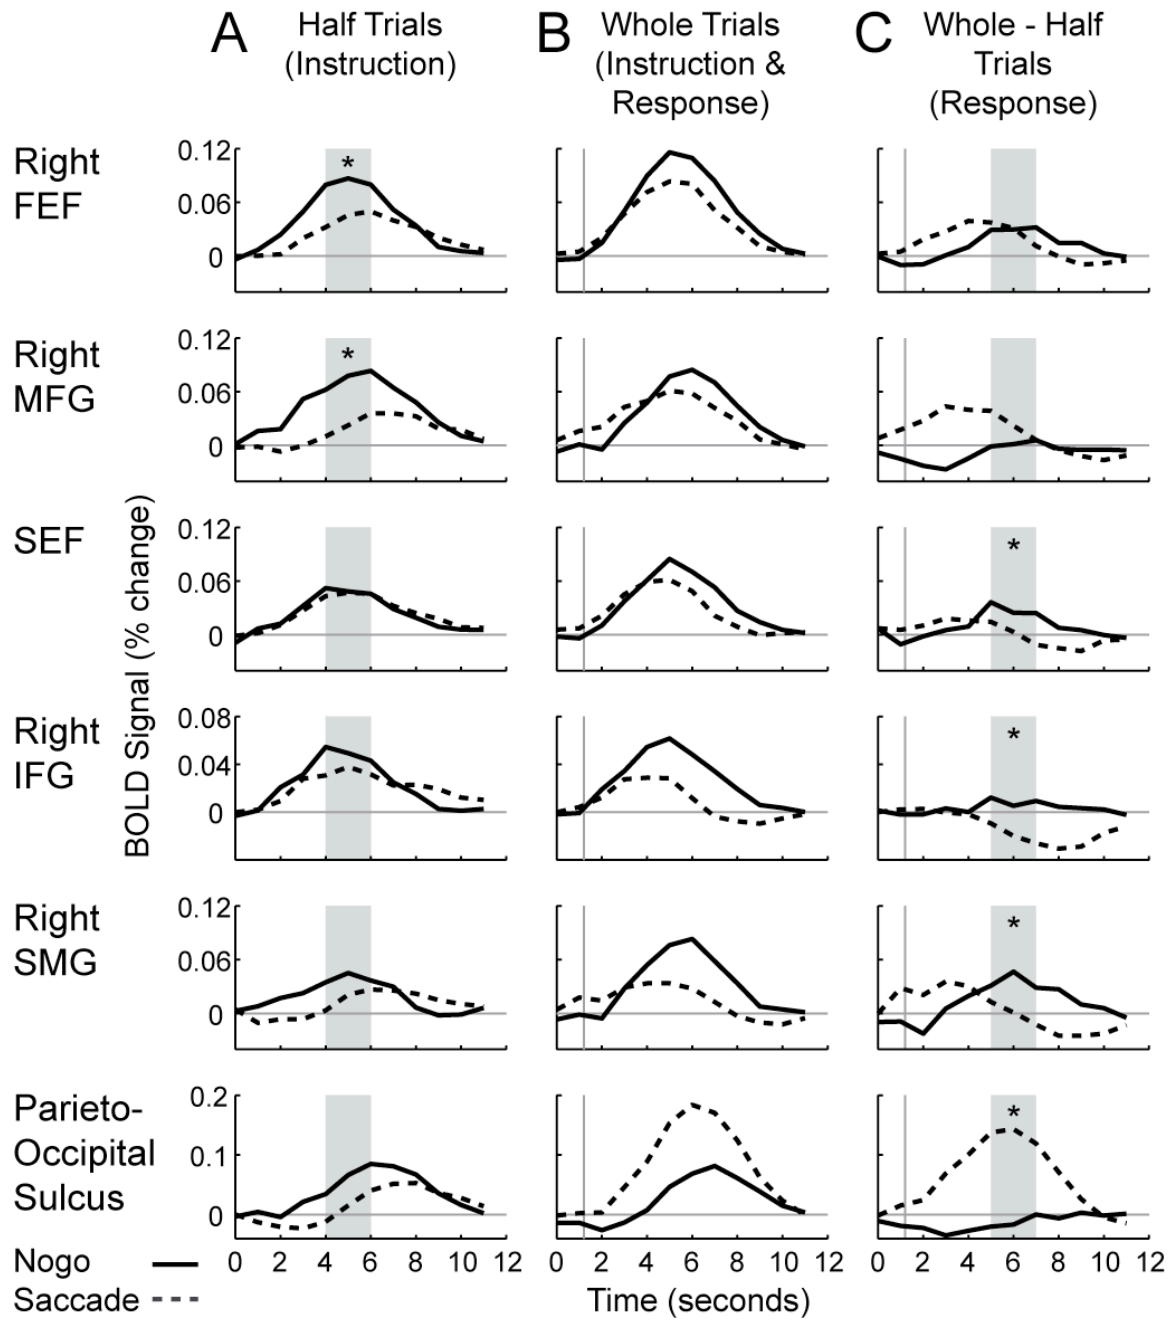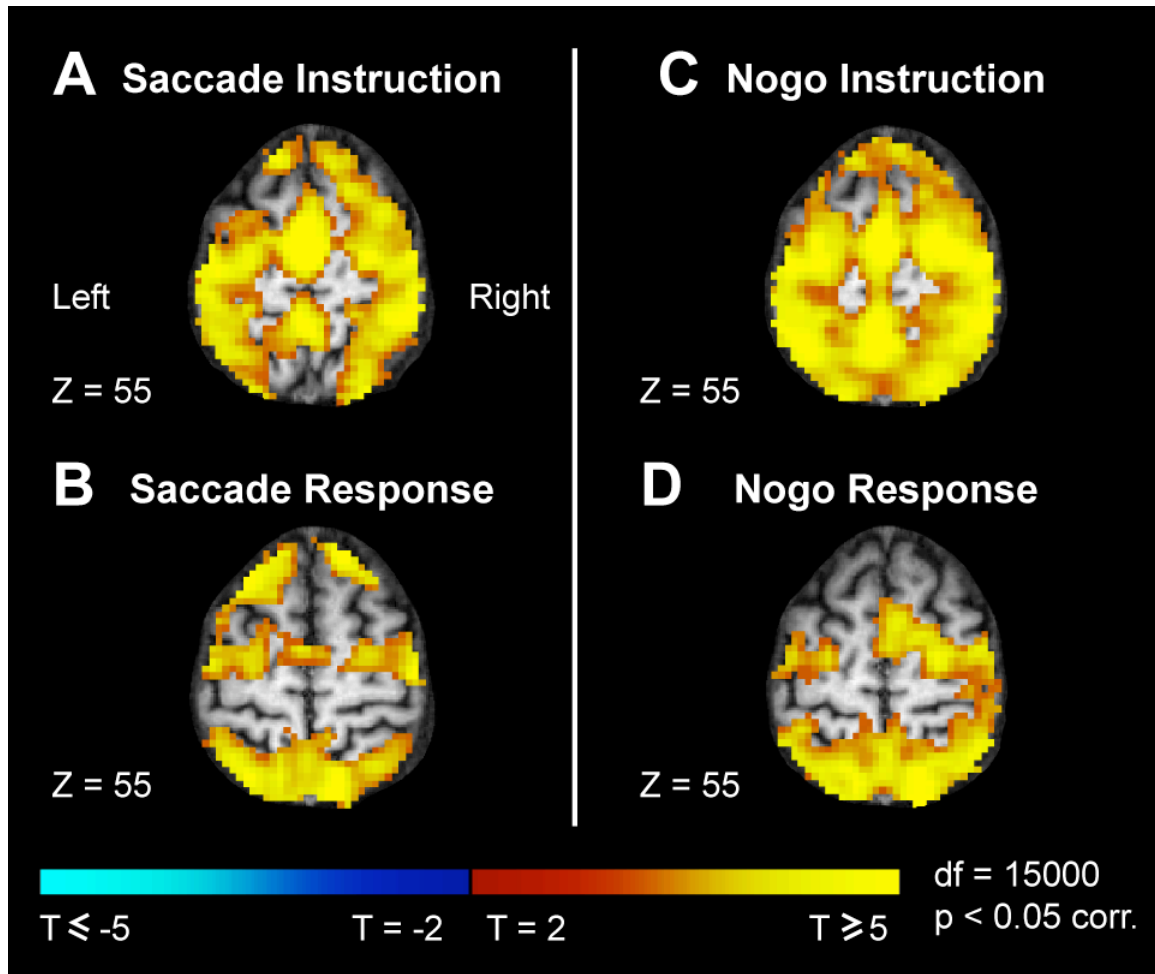
Figure 4.4 – Instruction and Response Localizer Maps

This figure shows localizer maps including *A*: saccade instruction (half trial) activation compared to fixation baseline, *B*: saccade response (whole – half trial) activation compared to fixation baseline, *C*: nogo instruction (half trial) activation compared to fixation baseline, and *D*: nogo response (whole – half trial) activation compared to fixation baseline. Yellow/red regions exhibited activation above fixation baseline. Maps are mixed effects maps across 11 subjects. Instruction localizer maps were computed across the 4-6 s interval, and response localizer maps were computed across the 5-7 s interval (see *Figure 4.3*). Images obey neurological convention.

Regions exhibiting greater nogo response activation were partially lateralized to the right side. The right IFG voxel cluster was larger than the left IFG cluster, at 7,209 and 2,349 cubic mm, respectively (see *Table 4.1B*), and statistically significant voxel clusters were present in right but not left SMG and IPL.

Parieto-occipital sulcus exhibited greater saccade compared to nogo response activation (blue region in *Figure 4.2B*, also see *Figure 4.3C*). This is consistent with single neuron recording studies which have found visual and saccade-related neuronal discharges in monkey area V6A (Nakamura et al. 1999; Kutz et al. 2003).

4.3.4 – Localizer Results

We also compared instruction- and response-related activation for saccade and nogo trials (separately) against fixation baseline. This comparison generated the localizer maps shown in *Figure 4.4*. All regions that exhibited differences between saccade and nogo trials in the instruction and response contrasts also exhibited saccade and nogo activation levels above baseline. The single exception was right IFG (not shown), which exhibited saccade response activation below baseline and nogo response activation above baseline. One possibility is that this region might be involved in inhibiting the generation of saccades during nogo trials and during the fixation intervals between trials.

4.4 – Discussion

We compared frequent saccade trials with rare nogo trials in a rapid event-related fMRI design. Saccade trials were presented twice as frequently as nogo trials as a means of making the saccade response more automatic, thereby increasing the need for active inhibition of the automatic saccade response in the nogo trials. We included half trials as a means of separating instruction- and response-related signals (see *Section 4.2.4 – Experimental Design*). Instruction-related activation was higher for nogo trials than saccade trials in right frontal eye field (FEF), bilateral middle frontal gyrus (MFG), intraparietal sulcus (IPS), and precuneus. Response-related activation was greater for nogo trials than saccade trials in bilateral inferior frontal gyrus (IFG), supplementary eye field (SEF), anterior cingulate cortex (ACC), and right supramarginal gyrus (SMG).

4.4.1 – Instruction-related Activation Differences

Greater instruction-related activation levels for nogo trials versus saccade trials in right FEF, bilateral MFG, IPS, and precuneus likely reflect preparation for active inhibition of the automatic saccade in the nogo task. The instruction differences we found could also reflect task switching processes (for review see Monsell 2003). Because saccade trials outnumbered nogo trials by two to one, a given trial of either type was most likely to be preceded by a saccade trial. This would make most saccade trials repeat trials and most nogo trials switch trials, and the above-mentioned instruction differences might reflect the need to switch task set more often in nogo trials. Of course, switching task set from saccade to nogo task performance might, itself, recruit inhibitory processes to inhibit the previously resident saccade task set, and in this sense, explanations of instruction-related differences based on preparation for saccadic inhibition versus task switching need not be mutually exclusive. We discuss the task-switching implications of our results further below.

4.4.2 – Response-related Activation Differences

Nogo trials evoked greater response-related activation in bilateral IFG, SEF, ACC, and right SMG. We attribute these results to the recruitment of processes related to inhibiting the generation of an automatic saccade in the nogo task. In our previous fMRI comparison of saccades and nogo trials (Brown et al. 2006), we did not find activation differences between the two tasks in cortical saccade regions including FEF, SEF, IPS, and ACC. To explain this, we suggested that fMRI activation in these regions might reflect stimulus-related processes, such as visual attention and stimulus selection, rather than response-related processes, such as saccade generation or inhibition. In the previous study, saccades and nogo trials were presented with equal frequency, whereas saccade trials were presented twice as frequently as nogo trials in the current experiment, with the intention of increasing the need for inhibitory control in the nogo task. Given the greater response-related activation for nogo trials in the current study, we must modify our previous suggestion. Activation patterns in FEF, SEF, IPS, and ACC might be determined predominantly by stimulus-related processing, but the current study's results

suggest that saccade inhibition and other processes associated with an increased requirement for inhibition of saccade responses also contribute to fMRI signal patterns in SEF and ACC.

The saccades used in this study accrued automaticity or prepotency from at least two, potentially distinct sources: the visual grasp reflex and the high frequency of saccade trials in the trial presentation sequence. It is currently an open question whether a single neuronal mechanism underlies these two sources of prepotency or whether they have distinct neuronal underpinnings. Accordingly, the response-related differences cited above could reflect a single inhibitory mechanism or some subset or combination of inhibitory systems specialized to address different aspects of prepotency. One possible avenue for future research would be to address this issue.

One potential criticism of our attributing the response-related results to saccade inhibition is that the rare nogo trials might have been more difficult than the frequent saccade trials. A difference in difficulty between the tasks does not provide a good explanation for our results. If the more difficult nogo task were to evoke higher levels of a non-specific attention- or arousal-related BOLD signal, we would expect the effect to occur throughout the fronto-parietal attention / eye movement system, including FEF, SEF, and IPS, based on the work of Raichle and colleagues (Fox et al. 2005; Vincent et al. 2007). On the contrary, we observed response-related differences only in SEF, ACC, IFG, and right SMG and instruction-related differences only in right FEF, MFG, IPS, and precuneus. It is unlikely that a non-specific, difficulty-related signal would manifest as a region-selective pattern of instruction- versus response-related differences.

The instruction- and response-related differences we observed were partially lateralized to the right hemisphere. Left MFG and IFG clusters were smaller than their right hemisphere counterparts, and FEF and SMG displayed significant clusters in the right hemisphere only. This right lateralization of inhibition-related differences agrees with the inhibition literature (see Aron et al. 2004; Buchsbaum et al. 2005).

4.4.3 – Response Inhibition and Task Switching

Previous fMRI studies using button press go/nogo tasks designed to recruit response inhibition found greater nogo activation in diverse frontal and parietal regions,

including middle and inferior frontal gyri, intraparietal sulcus, supramarginal gyrus, and anterior cingulate (Asahi et al. 2004; Braver et al. 2001; de Zubicaray et al. 2000; Fassbender et al. 2004; Garavan et al. 1999; Garavan et al. 2002; Horn et al. 2003; Kelly et al. 2004; Maguire et al. 2003; Menon et al. 2001; Rubia et al. 2001; Rubia et al. 2005; Wager et al. 2005). As discussed in Brown et al. (2006), these previous studies were unable to separate activation related to task switching from that related to active response inhibition processes. In four of the studies (Asahi et al. 2004; Maguire et al. 2003; Menon et al. 2001; Rubia et al. 2001), blocks of pure go trials were compared with mixed go / nogo blocks, and the mixed blocks involved switching between the two tasks, whereas the pure blocks did not. The other nine studies (Braver et al. 2001; de Zubicaray et al. 2000; Fassbender et al. 2004; Garavan et al. 1999; Garavan et al. 2002; Horn et al. 2003; Kelly et al. 2004; Rubia et al. 2005; Wager et al. 2005) included comparisons of go and nogo trials in rapidly-presented sequences, with nogo trials presented less frequently than go trials. In this case, the majority of trials of either type were preceded by go trials, making most go trials repeat trials and most nogo trials switch trials (Monsell 2003). This raises the possibility that task switching, in addition to or instead of response inhibition, might account for the differences cited above.

Our results can shed some light on this issue. We included half and whole trials in our design to separate instruction- and response-related signals (Ollinger et al. 2001a; Ollinger et al. 2001b). The greater nogo response-related activations we observed in IFG, SEF, ACC, and right SMG agree with the response inhibition studies listed above. Furthermore, we can eliminate the possibility that these differences were caused by preemptive task-switching processes during the instruction period because we subtracted instruction-related signals out of the whole trial signals to resolve response-related activation patterns. It is tempting, then, to attribute these differences to response inhibition processes recruited by the nogo task, but there still exists the possibility that task switching in the response-period, itself, might account for these results. The switch cost, or latency increase after a switch trial, is reduced only partially by prior warning of the switch (see Monsell 2003). To explain this 'residual switch cost', Rogers and Monsell (1995) have suggested that an exogenous component of task switching must be performed at the time of the response, irrespective of any prior preparation (but see de

Jong 2000). For example, most task switching studies involve switching between two tasks that prescribe different responses to the same stimulus array. It is possible then that, in switch trials, response inhibition processes might be evoked to inhibit left-over stimulus-response mappings that would tend to trigger the previous task's response upon appearance of the stimulus array. This hypothesis is consistent with fMRI results based on task switching paradigms. Response inhibition-related fMRI activation patterns and task switching-related activation patterns have been found to overlap in MFG, IFG, the junction of ACC and supplementary motor area (SMA), and IPS (see Buchsbaum et al. 2005; see also Brass et al. 2003; Dove et al. 2000; Dreher et al. 2002; Smith et al. 2004; Sohn et al. 2000). Based on our results, we suggest that the overlap in IFG, SEF / SMA, and ACC represents recruitment of response inhibition in switch trials at the time of stimulus presentation and response execution. In contrast, we found greater instruction-related activation for nogo trials but no response-related differences in MFG and IPS. These results imply that the overlap between inhibition- and task switching-related activations in MFG and IPS might represent aspects of task switching not based on response inhibition, for example reconfiguration of task set as well as preparatory processes in studies that warned of an imminent switch before the appearance of the stimulus array.

4.4.4 – Neuronal Recording with Nogo Saccade Task

Sommer and Wurtz (2001) and Paré and Wurtz (2001) recorded in FEF and monkey lateral intraparietal (LIP) area, respectively, while monkeys performed versions of the saccade and nogo tasks that included a go-nogo instruction cue (coloured fixation point), followed by peripheral stimulus presentation, followed in turn by a cue to respond. It is important to note that the monkeys waited for a specific cue before responding to the peripheral stimulus; the stimulus itself did not evoke the response. Therefore, the go and nogo tasks involved identical response inhibition requirements. Both studies found a variety of neuronal response profiles, including a subset of neurons with selectivity for go trials and an even smaller subset with selectivity for nogo trials. Paré and Wurtz (2001) also found that the instruction cue, itself, did not evoke elevated neuronal discharge rates in LIP. Only a small minority of neurons exhibited differences in peripheral stimulus-

evoked discharge activity in LIP with 18% and 3% of LIP neurons showing greater stimulus-related discharges for go and nogo trials, in that order. Sommer and Wurtz (2001) did not include a systematic analysis of neuronal discharges evoked by task instructions that were not preceded by peripheral stimulus presentation. They also found essentially no differences in peripheral stimulus-evoked discharge rates between go and nogo trials in FEF. We ascribe the differences between these results and those of the current study to the fact that our nogo task imposed the requirement to inhibit the automatic saccade response, unlike our saccade task, whereas the tasks in Sommer and Wurtz (2001) and Paré and Wurtz (2001) required monkeys to inhibit the automatic saccade to the peripheral stimulus in both the go and nogo tasks, while then waiting for a specific cue to respond.

Thompson et al. (1997) and Schall (2004) also compared saccade and nogo trials in the context of a visual search task. Their analysis focused on comparing neuronal discharges in response to the target versus distractor stimuli, and they did not include a direct comparison of saccade and nogo trial activity.

4.4.5 – Positron Emission Tomography with Nogo Task

Petit et al. (1999) used positron emission tomography (PET) to compare the nogo task as well as fixation with a rest condition involving fixation without visual stimulation. Both the nogo and fixation tasks evoked more activation than the rest condition in bilateral FEF and IPS as well as in right MFG, right IFG, and right inferior precentral gyrus extending ventro-laterally from the right FEF activation focus. The nogo task also evoked more activation than the fixation task in right FEF, right IPS, and bilateral area V5/MT. These results broadly agree with those from the current study. We did not image inferior temporal cortex, so we have no data on area V5 for comparison. We also found greater response-related activation for nogo versus saccade trials in SEF and ACC. That Petit et al. (1999) did not find differences in SEF or ACC might be explained by their not including saccade trials in their design. Subjects were never actually required to make eye movement responses, which might have lessened the recruitment of response inhibition processes in their nogo task. Another potential explanation is that Petit et al. (1999) used a blocked trial design in contrast to our rapid event-related design.

4.4.6 – Antisaccade Literature

The antisaccade task (Hallett 1978) requires subjects to inhibit the automatic saccade to a peripheral stimulus in favour of a voluntary saccade to the stimulus' mirror location. The antisaccade task has been used extensively as a model of response inhibition coupled with voluntary response generation. Previous fMRI studies that compared antisaccades and saccades (Connolly et al. 2002; Curtis and D'Esposito 2003; DeSouza et al. 2003; Ford et al. 2005; Brown et al. 2006; Brown et al. 2007) found greater instruction-related fMRI activation for antisaccades in FEF, SEF, IPS, MFG, and pre-supplementary motor area and greater response-related activation for antisaccades in FEF, SEF, and IPS, though DeSouza et al. (2003) attributed this to carryover from the instruction period. Several of these cortical saccade regions are thought to provide inhibitory control over the automatic saccade in the antisaccade task including SEF and MFG.

A role for SEF in inhibition is supported not only by the fMRI results cited above but also by single neuron recording results. Several studies have found greater neuronal discharges in SEF before and during the generation of antisaccades versus saccades (Schlag-Rey et al. 1997; Amador et al. 2004). That we found greater fMRI activation for nogo responses versus saccade responses in SEF supports this claim.

We found greater instruction-related activation for nogo trials in MFG, but we did not find response-related differences in MFG. We found a similar result in a previous rapid event-related comparison of prosaccades and antisaccades (Brown et al. 2007). In that study, left MFG exhibited greater instruction-related, but not response-related, activation for antisaccades versus prosaccades. These results suggest that MFG might be more involved in presetting the system to perform one task or the other rather than in generating task responses. This suggestion is consistent with neuronal recording studies. In monkey dorsolateral prefrontal cortex (dlPFC), a putative homologue of human MFG, Everling and DeSouza (2005) and Johnston and Everling (2006b) found neuronal discharge patterns that coded components of task set, including whether the monkey was required to perform the saccade or antisaccade task. Johnston and Everling (2006a) found

that dlPFC sends task-selective signals to the superior colliculus during the preparatory period before peripheral stimulus presentation.

Previous fMRI studies of antisaccades and saccades (Connolly et al. 2002; Curtis and D'Esposito 2003; DeSouza et al. 2003; Ford et al. 2005; Brown et al. 2006; Brown et al. 2007) did not find differences in IFG, unlike the current study, which found greater nogo compared to saccade response activation in IFG. None of the studies cited above presented saccades and antisaccades in a rapid design with frequent saccades and rare antisaccades in the way that the current study's design included frequent saccades and rare nogo trials. It is possible that the current study evoked greater nogo response activation in IFG because of the need to inhibit a frequent response type during execution of a rare response type. That is, IFG activation might not be related specifically to inhibiting the automatic saccade in nogo trials. One way to test this hypothesis would be to compare rare saccade trials to frequent nogo trials with fMRI. If IFG is specialized for inhibiting a frequent response type, we would expect to see greater activation for rare saccade responses compared to frequent nogo responses.

4.4.7 – Summary

Instruction-related processes, including preparation and task switching, evoked greater activation in nogo trials versus saccades in right FEF, MFG, IPS, and precuneus. In IFG, SEF, ACC, and right SMG, saccade inhibition and other processes associated with an increased requirement to inhibit saccade responses evoked greater response-related activation in nogo trials compared to saccade trials.

4.5 – Bibliography

Amador N, Schlag-Rey M, and Schlag J. 2004. Primate antisaccade. II. Supplementary eye field neuronal activity predicts correct performance. *J Neurophysiol* 91(4): 1672-1689.

Aron AR, Robbins TW, and Poldrack RA. 2004. Inhibition and the right inferior frontal cortex. *Trends Cogn Sci* 8(4): 170-177.

Asahi S, Okamoto Y, Okada G, Yamawaki S, and Yokota N. 2004. Negative correlation between right prefrontal activity during response inhibition and impulsiveness: a fMRI study. *Eur Arch Psychiatry Clin Neurosci* 254(4): 245-251.

Barberi EA, Gati JS, Rutt BK, and Menon RS. 2000. A transmit-only/receive-only (TORO) RF system for high-field MRI/MRS applications. *Magn Reson Med* 43(2): 284-289.

Bell AH, Everling S, and Munoz DP. 2000. Influence of stimulus eccentricity and direction on characteristics of pro- and antisaccades in non-human primates. *J Neurophysiol* 84(5): 2595-2604.

Boynton GM, Engel SA, Glover GH, and Heeger DJ. 1996. Linear systems analysis of functional magnetic resonance imaging in human V1. *J Neurosci* 16(13): 4207-4221.

Brass M, Ruge H, Meiran N, Rubin O, Koch I, Zysset S, Prinz W, and von Cramon DY. 2003. When the same response has different meanings: recoding the response meaning in the lateral prefrontal cortex. *Neuroimage* 20(2): 1026-1031.

Braver TS, Barch DM, Gray JR, Molfese DL, and Snyder A. 2001. Anterior cingulate cortex and response conflict: effects of frequency, inhibition and errors. *Cereb Cortex* 11(9): 825-836.

Brown MR, Goltz HC, Vilis T, Ford KA, and Everling S. 2006. Inhibition and generation of saccades: rapid event-related fMRI of prosaccades, antisaccades, and nogo trials. *Neuroimage* 33(2): 644-659.

Brown MR, Vilis T, and Everling S. 2007. Frontoparietal activation with preparation for antisaccades. *J Neurophysiol* 98(3): 1751-1762.

Buchsbaum BR, Greer S, Chang WL, and Berman KF. 2005. Meta-analysis of neuroimaging studies of the Wisconsin card-sorting task and component processes. *Hum Brain Mapp* 25(1): 35-45.

Bullmore E, Long C, Suckling J, Fadili J, Calvert G, Zelaya F, Carpenter TA, and Brammer M. 2001. Colored noise and computational inference in neurophysiological (fMRI) time series analysis: resampling methods in time and wavelet domains. *Hum Brain Mapp* 12(2): 61-78.

Connolly JD, Goodale MA, Menon RS, and Munoz DP. 2002. Human fMRI evidence for the neural correlates of preparatory set. *Nat Neurosci* 5(12): 1345-1352.

Curtis CE, Cole MW, Rao VY, and D'Esposito M. 2005. Canceling planned action: an FMRI study of countermanding saccades. *Cereb Cortex* 15(9): 1281-1289.

Curtis CE and D'Esposito M. 2003. Success and failure suppressing reflexive behavior. *J Cogn Neurosci* 15(3): 409-418.

Dale AM. 1999. Optimal experimental design for event-related fMRI. *Hum Brain Mapp* 8(2-3): 109-114.

de Jong R. 2000. An intention-activation account of residual switch costs. In: *Control of Cognitive Processes: Attention and Performance XVIII*, edited by Monsell S and Driver J. Cambridge, MA: MIT Press.

de Zubicaray GI, Andrew C, Zelaya FO, Williams SC, and Dumanoir C. 2000. Motor response suppression and the prepotent tendency to respond: a parametric fMRI study. *Neuropsychologia* 38(9): 1280-1291.

Desimone R and Duncan J. 1995. Neural mechanisms of selective visual attention. *Annu Rev Neurosci* 18: 193-222.

DeSouza JFX, Menon RS, and Everling S. 2003. Preparatory set associated with pro-saccades and anti-saccades in humans investigated with event-related FMRI. *J Neurophysiol* 89(2): 1016-1023.

Dove A, Pollmann S, Schubert T, Wiggins CJ, and von Cramon DY. 2000. Prefrontal cortex activation in task switching: an event-related fMRI study. *Brain Res Cogn Brain Res* 9(1): 103-109.

Dreher JC, Koechlin E, Ali SO, and Grafman J. 2002. The roles of timing and task order during task switching. *Neuroimage* 17(1): 95-109.

Duncan J. 2001. An adaptive coding model of neural function in prefrontal cortex. *Nat Rev Neurosci* 2(11): 820-829.

Everling S and DeSouza JF. 2005. Rule-dependent Activity for Prosaccades and Antisaccades in the Primate Prefrontal Cortex. *J Cogn Neurosci* 17(9): 1483-1496.

Fassbender C, Murphy K, Foxe JJ, Wylie GR, Javitt DC, Robertson IH, and Garavan H. 2004. A topography of executive functions and their interactions revealed by functional magnetic resonance imaging. *Brain Res Cogn Brain Res* 20(2): 132-143.

Fischer B and Weber H. 1992. Characteristics of "anti" saccades in man. *Exp Brain Res* 89(2): 415-424.

Fischer B and Weber H. 1997. Effects of stimulus conditions on the performance of antisaccades in man. *Exp Brain Res* 116(2): 191-200.

Forbes K and Klein RM. 1996. The magnitude of the fixation offset effect with endogenously and exogenously controlled saccades. *J Cogn Neurosci* 8: 344-352.

Ford KA, Goltz HC, Brown MRG, and Everling S. 2005. Neural processes associated with antisaccade task performance investigated with event-related FMRI. *J Neurophysiol* 94(1): 429-440.

Fox MD, Snyder AZ, Vincent JL, Corbetta M, Van Essen DC, and Raichle ME. 2005. The human brain is intrinsically organized into dynamic, anticorrelated functional networks. *Proc Natl Acad Sci U S A* 102(27): 9673-9678.

Garavan H, Ross TJ, Murphy K, Roche RA, and Stein EA. 2002. Dissociable executive functions in the dynamic control of behavior: inhibition, error detection, and correction. *Neuroimage* 17(4): 1820-1829.

Garavan H, Ross TJ, and Stein EA. 1999. Right hemispheric dominance of inhibitory control: an event-related functional MRI study. *Proc Natl Acad Sci U S A* 96(14): 8301-8306.

Hallett PE. 1978. Primary and secondary saccades to goals defined by instructions. *Vision Res* 18(10): 1279-1296.

Hess WR, Burgi S, and Bucher V. 1946. Motor function of tectal and tegmental area. *Monatsschr Psychiatr Neurol* 112: 1-52.

Horn NR, Dolan M, Elliott R, Deakin JF, and Woodruff PW. 2003. Response inhibition and impulsivity: an fMRI study. *Neuropsychologia* 41(14): 1959-1966.

Johnston K and Everling S. 2006a. Monkey dorsolateral prefrontal cortex sends task-selective signals directly to the superior colliculus. *J Neurosci* 26(48): 12471-12478.

Johnston K and Everling S. 2006b. Neural activity in monkey prefrontal cortex is modulated by task context and behavioral instruction during delayed-match-to-sample and conditional prosaccade-antisaccade tasks. *J Cogn Neurosci* 18(5): 749-765.

Kelly AM, Hester R, Murphy K, Javitt DC, Foxe JJ, and Garavan H. 2004. Prefrontal-subcortical dissociations underlying inhibitory control revealed by event-related fMRI. *Eur J Neurosci* 19(11): 3105-3112.

Klassen LM and Menon RS. 2004. Robust automated shimming technique using arbitrary mapping acquisition parameters (RASTAMAP). *Magn Reson Med* 51(5): 881-887.

Kutz DF, Fattori P, Gamberini M, Breveglieri R, and Galletti C. 2003. Early- and late-responding cells to saccadic eye movements in the cortical area V6A of macaque monkey. *Exp Brain Res* 149(1): 83-95.

Machado L and Rafal R. 2000. Control of eye movement reflexes. *Exp Brain Res* 135(1): 73-80.

Maguire RP, Broerse A, de Jong BM, Cornelissen FW, Meiners LC, Leenders KL, and den Boer JA. 2003. Evidence of enhancement of spatial attention during inhibition of a visuo-motor response. *Neuroimage* 20(2): 1339-1345.

Menon V, Adleman NE, White CD, Glover GH, and Reiss AL. 2001. Error-related brain activation during a Go/NoGo response inhibition task. *Hum Brain Mapp* 12(3): 131-143.

Miller EK and Cohen JD. 2001. An integrative theory of prefrontal cortex function. *Annu Rev Neurosci* 24: 167-202.

Monsell S. 2003. Task switching. *Trends Cogn Sci* 7(3): 134-140.

Mostofsky SH, Schafer JG, Abrams MT, Goldberg MC, Flower AA, Boyce A, Courtney SM, Calhoun VD, Kraut MA, Denckla MB, and Pekar JJ. 2003. fMRI evidence that the neural basis of response inhibition is task-dependent. *Brain Res Cogn Brain Res* 17(2): 419-430.

Munoz DP and Everling S. 2004. Look away: the anti-saccade task and the voluntary control of eye movement. *Nat Rev Neurosci* 5(3): 218-228.

Nakamura K, Chung HH, Graziano MS, and Gross CG. 1999. Dynamic representation of eye position in the parieto-occipital sulcus. *J Neurophysiol* 81(5): 2374-2385.

Neggers SF, Raemaekers MA, Lampmann EE, Postma A, and Ramsey NF. 2005. Cortical and subcortical contributions to saccade latency in the human brain. *Eur J Neurosci* 21(10): 2853-2863.

Ollinger JM, Corbetta M, and Shulman GL. 2001a. Separating processes within a trial in event-related functional MRI, II. Analysis. *Neuroimage* 13(1): 218-229.

Ollinger JM, Shulman GL, and Corbetta M. 2001b. Separating processes within a trial in event-related functional MRI, I. The Method. *Neuroimage* 13(1): 210-217.

Pare M and Wurtz RH. 2001. Progression in neuronal processing for saccadic eye movements from parietal cortex area lip to superior colliculus. *J Neurophysiol* 85(6): 2545-2562.

Petit L, Dubois S, Tzourio N, Dejardin S, Crivello F, Michel C, Etard O, Denise P, Roucoux A, and Mazoyer B. 1999. PET study of the human foveal fixation system. *Hum Brain Mapp* 8(1): 28-43.

Rogers RD and Monsell S. 1995. The costs of a simple switch between predictable cognitive tasks. *J Exp Psychol Gen* 124(2): 207-231.

Rubia K, Lee F, Cleare AJ, Tunstall N, Fu CH, Brammer M, and McGuire P. 2005. Tryptophan depletion reduces right inferior prefrontal activation during response inhibition in fast, event-related fMRI. *Psychopharmacology (Berl)* 179(4): 791-803.

Rubia K, Russell T, Overmeyer S, Brammer MJ, Bullmore ET, Sharma T, Simmons A, Williams SC, Giampietro V, Andrew CM, and Taylor E. 2001. Mapping motor inhibition: conjunctive brain activations across different versions of go/no-go and stop tasks. *Neuroimage* 13(2): 250-261.

Schall JD. 2004. On the role of frontal eye field in guiding attention and saccades. *Vision Res* 44(12): 1453-1467.

Schlag-Rey M, Amador N, Sanchez H, and Schlag J. 1997. Antisaccade performance predicted by neuronal activity in the supplementary eye field. *Nature* 390(6658): 398-401.

Serences JT. 2004. A comparison of methods for characterizing the event-related BOLD timeseries in rapid fMRI. *Neuroimage* 21(4): 1690-1700.

Smith AB, Taylor E, Brammer M, and Rubia K. 2004. Neural correlates of switching set as measured in fast, event-related functional magnetic resonance imaging. *Hum Brain Mapp* 21(4): 247-256.

Sohn MH, Ursu S, Anderson JR, Stenger VA, and Carter CS. 2000. Inaugural article: the role of prefrontal cortex and posterior parietal cortex in task switching. *Proc Natl Acad Sci U S A* 97(24): 13448-13453.

Sommer MA and Wurtz RH. 2001. Frontal eye field sends delay activity related to movement, memory, and vision to the superior colliculus. *J Neurophysiol* 85(4): 1673-1685.

Thompson KG, Bichot NP, and Schall JD. 1997. Dissociation of visual discrimination from saccade programming in macaque frontal eye field. *J Neurophysiol* 77(2): 1046-1050.

Vincent JL, Patel GH, Fox MD, Snyder AZ, Baker JT, Van Essen DC, Zempel JM, Snyder LH, Corbetta M, and Raichle ME. 2007. Intrinsic functional architecture in the anaesthetized monkey brain. *Nature* 447(7140): 83-86.

Wager TD, Sylvester CY, Lacey SC, Nee DE, Franklin M, and Jonides J. 2005. Common and unique components of response inhibition revealed by fMRI. *Neuroimage* 27(2): 323-340.

Worsley KJ, Liao CH, Aston J, Petre V, Duncan GH, Morales F, and Evans AC. 2002. A general statistical analysis for fMRI data. *Neuroimage* 15(1): 1-15.

Chapter 5
Summary and Conclusions

5.1 – Overview of Thesis Work

The first theme of my thesis research has been understanding performance of the antisaccade task, which is an important model of flexible behavioural control. The other theme has been the use of rapid event-related functional magnetic resonance imaging (fMRI) techniques, which I will discuss in the next section. In Experiment 1, I looked for separate signatures for processes underlying active inhibition of saccades and voluntary generation of saccades. We compared prosaccades, antisaccades, and nogo trials with a rapid event-related fMRI design. The prosaccade task was meant to recruit saccade generation processes, the nogo task was meant to recruit saccade inhibition processes, and the antisaccade task was meant to recruit processes underlying both saccade generation and saccade inhibition. Comparisons among the three tasks were then intended to reveal the involvement of different brain regions in the saccade inhibition and saccade generation components of the antisaccade task. The results of Experiment 1 contained two surprises. First, we did not observe significant activation differences between the instruction events for prosaccades and antisaccades even though these had been demonstrated previously using widely-spaced fMRI (see *Section 1.1.15*). Second, there were no activation differences between prosaccades and nogo trials in the cortical saccade-related regions including frontal eye field (FEF), supplementary eye field (SEF), anterior cingulate cortex (ACC), intraparietal sulcus (IPS), and precuneus. This was despite the fact that the tasks had opposite response requirements. Experiments 2 and 3 explored these two issues further. We suggested based on these results that the fMRI BOLD signal in these regions might reflect stimulus detection and attention processes more than saccade generation or inhibition. We also found greater response-related activation for antisaccades compared to either prosaccades or nogo trials in FEF, SEF, ACC, IPS, and precuneus, which we attributed to visuo-spatial remapping and heightened attention levels in the antisaccade task.

In Experiment 2, we tested whether a rapid event-related paradigm could detect instruction-related differences between prosaccades and antisaccades. To separate

instruction- and response-related signals in Experiment 1, we varied the length of the instruction interval, but this method did not address the sequence bias inherent in compound trial designs (see *Section 5.3* below for details). To separate instruction and response activation while avoiding problems with sequence bias in Experiment 2, we used the half trial method. We found greater instruction-related activation for antisaccades versus prosaccades in FEF, SEF, IPS, precuneus, ACC, and left dorsolateral prefrontal cortex (DLPFC). This result demonstrated that a rapid event-related design could detect instruction-related differences observed previously in widely-spaced fMRI studies on antisaccades and prosaccades. We also found greater response-related activation for antisaccades versus prosaccades in FEF, SEF, IPS, and precuneus but not in ACC or DLPFC. We suggested that ACC and DLPFC are more involved in presetting the saccade system to perform the antisaccade task than in generating the antisaccade response. ACC exhibited response-related differences in Experiment 1 but not Experiment 2. ACC has been implicated in conflict monitoring (see *Section 1.1.19*). Experiment 1 forced the subject to keep track of three different tasks, whereas Experiment 2 involved only two tasks. This might have heightened conflict at the time of the response in Experiment 1, accounting for the response-related differences we observed in ACC.

Recall that we did not find differences between prosaccades and nogo trials in Experiment 1. In Experiment 3, we attempted again to isolate a signature for saccade inhibition processes by comparing saccades and nogo trials with a rapid fMRI design. We increased the difficulty of the nogo task in comparison with Experiment 1 by presenting nogo trials only half as frequently as saccade trials. This manipulation was intended to increase both the prepotency of the saccade response and the recruitment of response inhibition processes in the nogo task to suppress the saccade response. We reasoned that increased recruitment of response inhibition processes should manifest as greater BOLD activation levels on nogo trials. We observed greater instruction-related activation for nogo trials versus saccade trials in right FEF, DLPFC, IPS, and precuneus, and we suggested that these results reflected preparatory processes and task switching. We also observed greater response-related activation for nogo versus saccade trials in SEF, ACC, inferior frontal gyrus, and right supramarginal gyrus, and we attributed these differences

to inhibition of the automatic saccade response in the nogo task as the most straight-forward interpretation. It seems then that that the infrequent nogo trials did indeed recruit saccade inhibition processes more strongly than those in Experiment 1, hence the lack of results in Experiment 1. I must also mention some other possibilities, which I discuss in detail in *Chapter 4*. The rare nogo trials were mostly switch trials, whereas the frequent saccade trials were mostly non-switch trials, and task switching processes could have contributed to the instruction- and response-related results. Two varieties of prepotency might also have been involved in Experiment 3. Automatic saccades might have a hard-wired prepotency because of their particular importance, though it is also true that we frequently make automatic visually-guided saccades so they are well practiced. It is possible that presenting saccade trials more frequently than nogo trials might increase the prepotency of the saccade response by a mechanism not directly related to the intrinsic saccade prepotency. That could have contributed to our results in Experiment 3.

Taken together, the three experiment that make up my thesis work suggest the following picture of cortical saccade control. Visual stimulus detection and attention have a large impact on the BOLD signal in FEF, SEF, IPS, ACC, and precuneus, though inhibition of the automatic saccade response also plays a role. Inferior frontal gyrus and supramarginal gyrus are implicated in saccade response inhibition. DLPFC and perhaps ACC seem to be more involved in preparatory processes, task switching, and establishing general task set than in generating antisaccade, prosaccade, or nogo responses. ACC also seems to exhibit high levels of BOLD activation for high processing conflict situations around the time of response.

5.2 – Big Picture View

In the preceding chapters, I have limited myself to the typical kind of science writing, in which we restrict ourselves to theorizing that is close to the data and reasonably well supported by it. I will now take a more speculative stance to describe my own intuitions about how my research relates to the big picture. I am partial to the view that neuronal decision-making involves winner-take-all competition among mutually inhibitory activation patterns that are resident simultaneously in a neural network (for

example, see Coultrip et al. 1992; Chen and Yang 2000). These activity patterns "churn around" in the network, eventually settling into an equilibrium state. The pattern that "wins" the competition determines the final equilibrium state, and the final equilibrium state represents the network's decision. This type of computation can encompass the functions for which my thesis work attempted to detect isolated fMRI signatures, including deliberate saccade response generation and inhibition. For example, when a subject performs the antisaccade task, his or her eye movement control network has to decide whether to generate a prosaccade or antisaccade response. I think it probable that an important component of this decision involves mutually inhibitory interaction between prosaccade and antisaccade "programs". In this context, it does not make sense to conceptualize the computation in terms of separate voluntary saccade generation or inhibition processes because either a prosaccade or antisaccade program would encompass both inhibition of competing programs as well as the generation of saccade commands in its activity pattern. Since my own experiments are built explicitly on a conceptual dichotomy between saccade generation and inhibition, they are not designed to probe this "competing network states" aspect of neurocomputation. However, it also seems reasonable that this competition-based processing should take place at multiple scales ranging from very local interactions among physically adjacent neurons to long-range interactions among cortical regions. It also seems likely that individual activity patterns, such as prosaccade or antisaccade programs, are supported to differing degrees by different cortical regions. For example, prosaccade execution seems dependent on LIP and FEF but not on DLPFC. It seems that DLPFC is important in creating and/or maintaining the distributed pattern of activity that constitutes an antisaccade program, though. DLPFC also seems important for nogo trial programs. I suspect that it is these differences among cortical regions in their contributions to different task programs that are detected by fMRI in typical experiments. Furthermore, the idea that saccade task performance is the outcome of competitive, "winner takes all" interactions among different task programs has implications for how we interpret fMRI results. When we observe BOLD activation patterns that seem to be attributable to a process like voluntary saccade generation or inhibition of automatic saccades, this does not mean that the region exhibiting that activation pattern is necessarily devoted specifically to the given function.

It seems much more likely to me that those activation patterns simply reflect the operation of a relevant task program. In Experiment 3, inferior frontal gyrus exhibited more activation on nogo trial responses than prosaccade responses. Rather than labeling the region as an "automatic saccade inhibition module", I suggest that the fMRI results reflect the participation of inferior frontal gyrus in a nogo response program, which in this case competes with and inhibits the automatic saccade program that is evoked by peripheral stimulus presentation. This section has been quite speculative and more research will be necessary to firmly ground or else refute many of the ideas discussed here. At the same time, it is important to think about how individual results might eventually fit into a global view of ones research topic.

In the next two sections, I will discuss more technical contributions from my thesis work.

5.3 – Rapid Event-related fMRI

When we started the first of the three experiments described here, only a few proof of concept papers on rapid event-related fMRI had been published (for example Burock et al. 1998; Miezin et al. 2000). Now, the technique is fairly widely used among fMRI practitioners. This is not surprising given the freedom that rapid designs afford in terms of trial spacing. I found rapid designs useful in particular because they allowed us to present meaningful nogo trials in an event-related setting. Incorporating nogo trials into a widely-spaced event-related paradigm would have been problematic because the fatigue and dulling of attention that can occur during the long fixation intervals in widely-spaced designs would likely have blunted the prepotency of the automatic saccade. This would have raised the possibility that subjects might have "correctly" performed the nogo task because they were not engaged in the paradigm rather than because they deliberately inhibited the automatic saccade response.

The experiments described here were some of the first to apply rapid event-related methodology to the study of eye movement control. In addition, our studies contributed to the initial wave of rapid fMRI studies, which were important for demonstrating that rapid designs do in fact work. In this context, I should also mention Curtis et al. (2005),

who deserve credit for applying rapid fMRI methodology to the study of eye movements with their investigation of the countermanding task.

Rapid event-related designs rely on randomizing or jittering the inter-trial interval to preserve linear independence among the BOLD signal components evoked by individual tasks (see *Section 1.2.4*). Linear analyses, such as those based on averaging or on the general linear model, can then be used to tease apart the respective BOLD signatures of different trial types. Rapid designs can also be used to separate signal components for events within compound trials. This is more challenging than just separating trial type activation profiles because the events in a compound trial must almost certainly have a fixed order, which introduces sequence bias. We encountered this problem in Experiment 1 (also see *Chapter 2* for details). To decorrelate instruction- and response-evoked fMRI signals, we randomly varied the interval between instruction onset and peripheral stimulus onset. This manipulation did allow us to deconvolve activation time courses for the instruction and response events in the three trial types, but we found almost no statistically significant differences between the instruction curves. This was particularly surprising for the antisaccade versus prosaccade instruction comparison because previous studies using widely-spaced fMRI did find instruction differences with these tasks (see *Section 1.1.15*). It is possible that jittering the instruction interval was not entirely successful in separating the instruction- and response-related signals and that the very large response-related activation profiles therefore obscured instruction-related differences. This suspicion motivated Experiment 2, which compared prosaccades and antisaccades using a rapid design but incorporated half trials as a means of separating instruction- and response-related signals. Recall that half trials contain only the first of two events in a whole compound trial. The half trial approach worked well, and we found strong statistical differences between prosaccade and antisaccade instruction signals, in line with previous studies. These results suggest that the half trial method is better than jittering inter-event intervals for separating activation patterns evoked by events within compound trials, probably because using half trials goes a long way to remove the sequence bias inherent in compound trial designs. For example, including half trials in Experiment 2 allowed us to precede an instruction event not just with a response event, as was the case in Experiment 1, but also with another instruction event. The sequence bias

was not entirely removed in Experiment 2 (or Experiment 3, which also used half trials), however. An antisaccade response had necessarily to be preceded by an antisaccade instruction. The various events in Experiment 2's design were still less correlated with each other than were those in Experiment 1.

One other argument for the superiority of the half trial method over jittering intervals within trials was that the half trials in Experiments 2 and 3 allowed us to use a much less constrained model of the haemodynamic function compared to Experiment 1. In Experiment 1, we specified a difference of gamma functions as an explicit model of the haemodynamic function, whereas in Experiments 2 and 3, we were able to use finite impulse response predictors, which do not incorporate specific assumptions about the haemodynamic response function's shape (see *Appendix 3.1.2*). In Experiment 1, attempting to deconvolve instruction activation curves using finite impulse response predictors produced very noisy results. Closer examination revealed that the finite impulse predictors were particularly vulnerable to the sequence bias described in the last paragraph.

5.4 – Neuronal Recording vs. fMRI with the Antisaccade Task in FEF

As Ford et al. (2005) have pointed out, there exists a discrepancy between neuronal recording results and fMRI results in FEF (see also *Section 1.1.16*). Saccade-related FEF neurons discharge at higher rates before and during prosaccades versus antisaccades, whereas fMRI activation is higher for antisaccades than prosaccades in FEF. Several factors might explain this difference. The techniques of neuronal recording and fMRI might be measuring different aspects of prosaccade and antisaccade performance. Neuronal recording is done in Rhesus macaque monkeys, whereas fMRI experiments on prosaccades and antisaccades have all been done on humans. The two species might perform the task differently. As I write this, Kristen Ford in Dr. Stefan Everling's laboratory is performing the equivalent fMRI experiment with monkeys to test this hypothesis. Monkeys are also over-trained on the tasks, with training occurring for many months before cellular discharge data are collected. Humans in fMRI experiments are usually told how to perform the tasks and given only a few minutes of training.

Finally, previous event-related fMRI experiments on prosaccades and antisaccades used widely-spaced designs, which might have changed the psychological nature of the tasks in comparison to the rapid pace of trial presentation typical of monkey electrophysiology experiments. In Experiments 1 and 2, I compared prosaccades and antisaccades using rapid event-related fMRI with a quick pace of trial presentation fairly similar to what is used in monkey experiments. Because I also found greater activation for antisaccades versus prosaccades in cortical saccade regions, in agreement with previous human fMRI studies, we can rule out the third explanation listed above. Further research is needed to test the other hypotheses.

5.5 – Future Directions

In Experiment 3, we compared frequent saccade trials with rare nogo trials. The manipulation of task frequencies was intended to increase the prepotency of the saccade response, forcing subjects to recruit saccade inhibition processes to a greater degree in the nogo task. One potential criticism of this approach is that inhibition of responses made prepotent through high frequency might be different from inhibition of intrinsically automatic saccade responses (also see *Section 4.4.2*). If so, our results from Experiment 3 might reflect processes related to inhibiting a frequent response rather than inhibiting automatic saccades, per se. Two different approaches could be taken to address this issue. One test would be to run a version of Experiment 3 in which the saccade trials were rare and the nogo trials frequent, as I mentioned in *Section 4.4.6*. In cortical regions involved predominantly in inhibiting a frequent response type as opposed to automatic saccade responses, I would expect to see greater activation for rare saccade responses compared to frequent nogo responses.

A second test of the results from Experiment 3 would manipulate visual stimulus saliency to make the nogo trials more or less difficult. One could compare saccade and nogo trials in a rapid fMRI design and include dim and bright peripheral stimulus conditions for both trial types. One would then compare saccade and nogo trial activation for the dim stimulus and for the bright stimulus conditions, separately. (It would be difficult to interpret a comparison across saliency conditions, for example bright nogo

trials to dim nogo trials, because the difference in stimulus characteristics would introduce a confound.) I would expect regions involved in automatic saccade inhibition to show greater activation for nogo versus saccade trials, and I would expect this difference to be greater in the high saliency condition.

5.6 – Thoughts on the Field

This final section includes some more general thoughts on the state of neuroscience and where it is going. The neuroscientific community devotes itself to studying a very difficult topic, how the mind and brain work. Science generally takes the approach of attempting to reduce the world into components that can be described with simple yet profound rules. Certain domains are intrinsically amenable to this approach with physics providing the quintessential example. It has proven possible to isolate physical systems well enough that they can be studied separately (though difficult technical and technological manipulations are often necessary to do so). For example, the solar system provided Kepler, Newton, and others with a ready-made example of a set of objects interacting based on classical gravitational rules, essentially to the exclusion of other influences. Even relativity, which is still a gravitational phenomenon though not a classical one, only comes into play when considering the orbit of a single planet, Mercury. Neuroscience has enjoyed some similar successes. Hodgkin and Huxley's model of action potential firing was built on experiments in the squid giant axon, which constituted a sufficiently well isolated system (Hodgkin and Huxley 1952). In general, however, the mind and brain resist decomposition into discrete components that can be individually inspected and manipulated. Put colloquially, the neurocomputational processes we investigate do not seem to "sit still". When we change one part of a task to manipulate a specific aspect of neuronal processing, other processes invariably change in ways that are not desired or not necessarily even anticipated. It is of course possible to run more experiments to investigate confounds and complications, but then one runs into the same sort of trouble. Though some loose ends are tied up, each step of the investigative process produces many more unresolved questions, which multiply

exponentially as a consequence. All this makes it very difficult to get any traction when trying to build theories on brain function.

How are we to address the difficulty posed by the brain's "resistance to reduction". My own suspicion is that the root of the problem lies in the interaction between theoretical and empirical neuroscience. On the one hand, theoretical neuroscience has developed the mathematical and computational tools to build, test, and think about complex models of neuronal systems. These models are capable of expressing the kinds of dynamic, recursive interactions, with their counterintuitive implications, that biological neural networks seem to exhibit. On the other hand, experimental neuroscience has developed a plethora of sophisticated techniques for exploring brain function at all levels of its organization. The shortcoming, I believe, is that experimental and theoretical neuroscience do not connect with each other well enough or often enough. Many, I would almost say most, theoretical models are inspired by empirical results but are not necessarily designed to be firmly grounded in them. Large neural network models with many parameters can illustrate principles of organization or computation that might be implemented in the brain, but it is usually impossible to establish whether those principles actually are or are not so implemented, given the many degrees of freedom in these models. Conversely, most empirical studies use designs that can only relate to simple hypotheses. These hypotheses often involve dichotomies between putative neurocomputational functions that are described verbally. Of course, a good hypothesis should be simple, and one should obviously start with the simplest hypotheses when broaching an unknown topic area. However, given the difficulties with trying to apply simple, qualitative verbal concepts to complicated neural networks, I propose that neuroscience must now move beyond this approach.

First, let me point out that technological advancements will contribute to closing the gap between empirical and theoretical neuroscience. I suspect that the new multi-electrode recording techniques being developed for brain-machine interfaces (Lebedev and Nicolelis 2006) will be particularly important. Recording from hundreds or thousands of neurons simultaneously would make it much easier to relate neuronal discharge data to neural network models. More important, however, is a deliberate effort to base research on reasonably sophisticated, quantitative theoretical models that are designed explicitly to

interface with empirical studies. I will briefly mention three examples of this kind of research. Gold and Shadlen (2002) relate an information-theoretic model of decision making to neuronal discharge patterns recorded in monkeys performing a forced choice response task. The application of information theory is appropriate here not only because it is well understood but also because the authors could shape it to make predictions about the discharge activities of individual neurons. The underlying neurophysiology is also well understood. Brown and Braver (2005) built a model of error likelihood learning in anterior cingulate cortex (ACC) inspired by general electrophysiological considerations and tested it using fMRI. A specific effort was made in building the model to relate its output meaningfully to fMRI measurements, and those experiments were designed explicitly to test hypotheses generated by considering the model. Finally, Trappenberg et al. (2001) built a model of how superior colliculus (SC) performs computations underlying the selection of eye movement targets and saccade generation. This model incorporated extensive electrophysiological data from neuronal recordings in SC. They tested different configurations of target stimuli and related the model's output to the behavioural performance of monkeys doing the same tasks. The common thread in all three approaches is a deliberate attempt to relate theory and empirical results. I believe that this theme will become increasingly important in neuroscientific research in the near future.

5.7 – Bibliography

Brown JW and Braver TS. 2005. Learned predictions of error likelihood in the anterior cingulate cortex. *Science* 307(5712): 1118-1121.

Burock MA, Buckner RL, Woldorff MG, Rosen BR, and Dale AM. 1998. Randomized event-related experimental designs allow for extremely rapid presentation rates using functional MRI. *Neuroreport* 9(16): 3735-3739.

Chen C and Yang J. 2000. Layer winner-take-all nerual networks based on existing competitive structures. *IEEE Transactions on Systems, Man and Cybernetics B* 30(1): 25-30.

Coultrip R, Granger R, and Lynch G. 1992. A cortical model of winner-take-all competition via lateral inhibition. *Neural Networks* 5: 47-54.

Curtis CE, Cole MW, Rao VY, and D'Esposito M. 2005. Canceling planned action: an FMRI study of countermanding saccades. *Cereb Cortex* 15(9): 1281-1289.

Ford KA, Goltz HC, Brown MRG, and Everling S. 2005. Neural processes associated with antisaccade task performance investigated with event-related FMRI. *J Neurophysiol* 94(1): 429-440.

Gold JI and Shadlen MN. 2002. Banburismus and the brain: decoding the relationship between sensory stimuli, decisions, and reward. *Neuron* 36(2): 299-308.

Hodgkin AL and Huxley AF. 1952. A quantitative description of membrane current and its application to conduction and excitation in nerve. *J Physiol* 117(4): 500-544.

Lebedev MA and Nicolelis MA. 2006. Brain-machine interfaces: past, present and future. *Trends Neurosci* 29(9): 536-546.

Miezin FM, Maccotta L, Ollinger JM, Petersen SE, and Buckner RL. 2000. Characterizing the hemodynamic response: effects of presentation rate, sampling procedure, and the possibility of ordering brain activity based on relative timing. *NeuroImage* 11(6 Pt 1): 735-759.

Trappenberg TP, Dorris MC, Munoz DP, and Klein RM. 2001. A model of saccade initiation based on the competitive integration of exogenous and endogenous signals in the superior colliculus. *J Cogn Neurosci* 13(2): 256-271.

Appendix 1
Ethics Approval Forms

Office of Research Ethics

The University of Western Ontario
Room 00045 Dental Sciences Building, London, ON, Canada N6A 5C1
Telephone: (519) 661-3036 Fax: (519) 850-2466 Email: ethics@uwo.c
Website www.uwo.ca/research/ethics

Use of Human Subjects - Ethics Approval Notice

Principal Investigator: Dr. S. Everling
Review Number: 08465
Revision Number: 1
Protocol Title: Neural control of saccadic eye movements and visual attention.
Department and Institution: Physiology, University of Western Ontario
Sponsor: MRC
Approval Date: 01-Nov-04
End Date: 31-Oct-05
Documents Reviewed and Approved: Revised Study End Date

Documents Received for Information:

This is to notify you that the University of Western Ontario Research Ethics Board for Health Sciences Research Involving Human Subjects (HSREB) which is organized and operates according to the Tri-Council Policy Statement and the Health Canada/ICH Good Clinical Practice Practices: Consolidated Guidelines; and the applicable laws and regulations of Ontario has received and granted full board approval to the above named research study on the date noted above. The membership of this REB also complies with the membership requirements for REB's as defined in Division 5 of the Food and Drug

This approval shall remain valid until end date noted above assuming timely and acceptable responses to the HSREB's periodic requests for surveillance and monitoring information. If you require an updated approval notice prior to that time you must request it using the UWO Updated Approval Request Form.

During the course of the research, no deviations from, or changes to, the protocol or consent form may be initiated without prior written approval from the HSREB except when necessary to eliminate immediate hazards to the subject or when the change(s) involve only logistical or administrative aspects of the study (e.g. change of monitor, telephone number). Expedited review of minor change(s) in ongoing studies will be considered. Subjects must receive a copy of the signed information/consent documentation.

Investigators must promptly also report to the HSREB:
a) changes increasing the risk to the participant(s) and/or affecting significantly the conduct of the study;
b) all adverse and unexpected experiences or events that are both serious and unexpected
c) new information that may adversely affect the safety of the subjects or the conduct of the study

If these changes/adverse events require a change to the information/consent documentation, and/or recruitment advertisement, the newly revised information/consent documentation, and/or advertisement, must be submitted to this office for approval.

Members of the HSREB who are named as investigators in research studies, or declare a conflict of interest, do not participate in discussion related to, nor vote on, such studies when they are presented to the HSREB.

Chair of HSREB: Dr. Paul Harding

Janice Sutherland, M.A., Ethics Officer HSREB

Faxed: Y/N
Date:

This is an official document. Please retain the original in your files.
UWO HSREB Ethics Approval   08465   Page 1 of 1

Office of Research Ethics

The University of Western Ontario
Room 00045 Dental Sciences Building, London, ON Canada N6A 5C1
Telephone: (519) 661-3036 Fax: (519) 850-2466 Email: ethics@uwo.ca
Website: www.uwo.ca/research/ethics

Use of Human Subjects - Ethics Approval Notice

Principal Investigator:	Dr. S. Everling
Review Number: 10254	Revision Number:
Protocol Title:	Cognitive Processes Underlying Eye Movement Performance Investigated With fMRI
Department and Institution:	Physiology & Pharmacology, Robarts Research Institute
Sponsor:	NSERC
Ethics Approval Date: April 15, 2005	Expiry Date: August 31, 2007
Documents Reviewed and Approved:	UWO Protocol, Letter of Information

Documents Received for Information:

This is to notify you that The University of Western Ontario Research Ethics Board for Health Sciences Research Involving Human Subjects (HSREB) which is organized and operates according to the Tri-Council Policy Statement: and the Health Canada/ICH Good Clinical Practice Practices' Consolidated Guidelines, and the applicable laws and regulations of Ontario has reviewed and granted full board approval to the above named research study on the approval date noted above. The membership of this REB also complies with the membership requirements for REB's as defined in Division 5 of the Food and Drug Regulations.

This approval shall remain valid until the expiry date noted above assuming timely and acceptable responses to the HSREB's periodic requests for surveillance and monitoring information. If you require an updated approval notice prior to that time you must request it using the UWO Updated Approval Request Form.

During the course of the research, no deviations from, or changes to, the protocol or consent form may be initiated without prior written approval from the HSREB except when necessary to eliminate immediate hazards to the subject or when the change(s) involve only logistical or administrative aspects of the study (e.g. change of monitor, telephone number). Expedited review of minor change(s) in ongoing studies will be considered. Subjects must receive a copy of the signed information/consent documentation.

Investigators must promptly also report to the HSREB:
a) changes increasing the risk to the participant(s) and/or effecting significantly the conduct of the study;
b) all adverse and unexpected experiences or events that are both serious and unexpected;
c) new information that may adversely affect the safety of the subjects or the conduct of the study.

If these changes/adverse events require a change to the information/consent documentation, and/or recruitment advertisement, the newly revised information/consent documentation, and/or advertisement, must be submitted to this office for approval.

Members of the HSREB who are named as investigators in research studies, or declare a conflict of interest, do not participate in discussion related to, nor vote on, such studies when they are presented to the HSREB.

Chair of HSREB: Dr. Paul Harding
Deputy Chair: Susan Hodgins

Ethics Officer to Contact for Further Information			
☐ Janice Sutherland	☐ Karen Kueneman	☐ Susan Underhill	☑ Jennifer McEwen

Faxed: Y/N
Date:

This is an official document. Please retain the original in your files.

UWO HSREB Ethics Approval   10254   Page 1 of 1

Appendix 2
Copyright Permissions

With regard to the material in Chapters 2 and 4:

The material in Chapter 2 was published as MR Brown, HC Goltz, T Vilis, KA Ford, and S Everling, 2006. Inhibition and generation of saccades: Rapid event-related fMRI of prosaccades, antisaccades, and nogo trials. *NeuroImage* 33(2): 644-659.

The material in Chapter 4 has been accepted for publication as MR Brown, T Vilis, and S Everling, 2007. Isolation of saccade inhibition processes: Rapid event-related fMRI of saccades and nogo trials. *NeuroImage* (in press).

Quoted from the website of Elsevier, which publishes *NeuroImage*:
(http://www.elsevier.com/wps/find/authorsview.authors/copyright#whatrights)

"As a journal author, you retain rights for large number of author uses, including use by your employing institute or company. These rights are retained and permitted without the need to obtain specific permission from Elsevier. These include:
[…]
  • the right to include the journal article, in full or in part, in a thesis or dissertation"

The American Physiological Society
9650 Rockville Pike, Bethesda, MD 20814-3991, USA
Phone: (301) 634-7070
Fax: (301) 634-7243

August 28, 2007

Mr. Matthew Brown
Medical Sciences Building, Room 216
University of Western Ontario
London, Ontario, Canada N6A 5C1

Dear Mr. Brown:

The American Physiological Society grants you permission to use the following article for your
PhD thesis.:

Matthew R G Brown, Tutis Vilis, and Stefan Everling
   Frontoparietal Activation with Preparation for Antisaccades
   J Neurophysiol (June 27, 2007). doi:10.1152/jn.00460.2007

The American Physiological Society publication must be credited as the source with the words
"used with permission" added.

The permission includes the National Library's use of the thesis.

Sincerely,

Ms. Margaret Reich
Director of Publications
The American Physiological Society

MR/pr

Appendix 3
Functional Magnetic Resonance Imaging (fMRI) Signal Analysis

Appendix 3 reviews the analysis of fMRI signal data based on the general linear model (GLM) and issues associated with non-orthogonality ("non-separateness") of the predictions one builds into the model.

Appendix 3.1 – General Linear Model (GLM) Analysis of fMRI Time Series

Rapid event-related designs involve close spacing of task events, which causes their blood oxygenation level dependent (BOLD) signal components to overlap and summate. This necessitates the use of deconvolution methods to separate individual fMRI activation components from the raw signal. The most widely-used deconvolution is based on the general linear model (GLM). In this section, I will review the GLM, and in the next section, I will show how it can be adapted for deconvolution of rapid fMRI data.

GLM-based analysis of fMRI time series is usually done on a voxel-by-voxel basis. That is, the BOLD time series from each voxel is analyzed separately, and the results are combined across voxels at the end. In this approach, the voxel's time series is represented as a vector or sequence of BOLD signal values with length equal to the number of functional volumes (time points) collected in the experiment. A vector is simply an ordered list of numbers that obey certain algebraic properties that I will not discuss here (but see Strang 1988). What is important is that the head of an fMRI data vector can be represented as a point in a high-dimensional space. The vector can then be represented as an arrow (a directed line) starting at the origin (centre of the coordinate system) and ending at the point. If the experimenter collected 1800 time points, that is 1800 functional volumes, the vector would be a line in an 1800-dimensional space.

GLM analysis consists of two steps, defining the GLM design matrix and fitting it to the data. The GLM design matrix is built by defining a set of expected fMRI signal profiles, called predictors, which are simply vectors (with length equal to that of the fMRI data vector) that each express some expected structural aspect of the data. The individual predictor vectors are then assembled side-by-side into the GLM design matrix. Put another way, the predictors form the columns of the GLM matrix. We call these

vectors columns vectors. Predictors are typically built using the finite impulse response framework.

The finite impulse response framework (also see Rabiner and Gold 1975) models a system's output by combining its input with its impulse response function. The impulse response function is the response of the system to an input which is a single impulse. (A finite impulse response function eventually returns to zero after some finite time.) Informally, an impulse is a pulse of signal with "infinite" height whose duration is "infinitely" small, resulting in an integral value or "area" of 1. The concept of Dirac's delta function formalizes this idea. When expressed in vector form, an impulse must be discretized, such that it takes the value 1 at one time point (the time of the impulse) and 0 everywhere else. (Actually, the non-zero value of the impulse equals the sampling rate, so the impulse will only have a value of 1 if the sampling rate is 1Hz.) In fMRI, the haemodynamic response function can be thought of as an impulse response function (Dale 1999; see also Boynton et al. 1996; Dale and Buckner 1997; Friston et al. 1994). The haemodynamic response describes the coupling between neuronal activity and the BOLD signal. To build predictor vectors, one combines a model of the haemodynamic response function with a model of the onset times of each task event that is expected to evoke a BOLD response. The haemodynamic response is modeled as a curve such as a gamma function or difference of gamma functions (Figure *1.1* in *Chapter 1*) (Boynton et al. 1996; Dale and Buckner 1997; Friston et al. 1998). The onset times of each task or task subcomponent of interest are modeled as a separate sequence of impulses (see *Appendix Figures 3.2A*, *3.3A*, and *3.4A* below). For example, in a simple localizer experiment using visually guided saccades to identify the saccade system, one might use a single impulse sequence specifying the onsets of all the trials in the experiment. In a blocked design with three different tasks, one might include three impulse sequences, one to model trials in each of the three block types. When comparing prosaccades and antisaccades in a widely-spaced design, one could use four different impulse sequences to model the onsets of the instruction periods and the response periods for both the prosaccade the antisaccade trials. Once one has specified the haemodynamic response function and the impulse sequences, one uses convolution, a simple mathematical operation (see *Appendix 3.3*), to combine each impulse sequence with the haemodynamic

response and generate an expected signal shape for the BOLD activation evoked by that impulse sequence. These expected signal functions are the predictors that comprise the GLM matrix. One can also include extra columns in the GLM matrix for so-called nuisance predictors that model, for example, changes in baseline from run to run or low frequency noise components in the signal (for more details, see Chapter 12 of Huettel et al. 2004). Incidentally, a finite impulse response system is a linear system, and the approach described in the last paragraph is formally valid only if one assumes that the haemodynamic response function and the BOLD signal are linear. This is true for long inter-event intervals (greater than 6 s), but the haemodynamic response becomes non-linear with shorter intervals (see Section *1.2.6*).

Once the GLM design matrix has been built, it is fit to the data as follows. By multiplying the set of predictor vectors by beta weight coefficients (which are just numbers) and then adding the products together, one can create an estimate of the measured data time series. The operation of scaling and adding vectors in this way is called linear combination. In general linear modeling, the aim is to find a linear combination of the predictors that is the "best" estimate of the actual data time series. The best estimate is the one for which the sum of squared errors between the estimated and actual time series is minimized. The computations underlying the fitting process are straight-forward (see chapter 3.3 of Strang 1988), and they return a unique solution provided the predictors are chosen to be linearly independent. The meaning of linear independence will be provided shortly.

There is actually a very elegant geometrical conception of GLM fitting process. As I mentioned earlier, a vector of a given length, call it *n*, can be envisaged as existing in an *n*-dimensional space. For example, a vector containing two elements exists in a 2-dimensional plane. The vector [2, 3] can be represented as an arrow in the plane with its tail at the origin (coordinate [0, 0]) and its head at position [2, 3]. A vector with three elements exists in 3-dimensional space. By extension, a vector with *n* elements exists in an *n*-dimensional space. We cannot actually visualize spaces with dimension greater than 3, but we can describe and manipulate them mathematically. Suppose we ran an fMRI experiment and collected one volume per second for 30 minutes, or 1800 seconds. This would give us a BOLD time series with 1800 values, which we would represent as a

vector with 1800 elements existing in an 1800-dimensional space. In addition, the predictor vectors in our GLM would be vectors in an 1800-dimensional space. Suppose in what follows that we used 10 predictor vectors in our GLM. This is realistic as the number of predictors in fMRI analyses is typically no more than several dozen, which is much smaller than the number of data points in most experiments.

It is now necessary to introduce the idea of a subspace. A subspace is a space that is embedded in a bigger space, for example a plane embedded in 3-dimensional space or a line embedded in a plane. We can also embed a 10-dimensional subspace in an 1800-dimensional space, though these spaces are impossible to visualize. One way to describe a space (or subspace) is to define a coordinate system for it. A coordinate system requires one axis for each of its dimensions. The standard Cartesian coordinate system has orthogonal axes oriented in the familiar up-down, left-right configuration (*Appendix Figure 3.1A*). Coordinate systems specify the positions of points by constraining those positions along the coordinate axes. The point [1, 3] in *Appendix Figure 3.1A* is defined as the intersection of the lines X=1 and Y=3. This is an implicit way of describing locations. It turns out to be easier to do general linear modeling using a more explicit approach. Instead of specifying locations by constraining them, we will build the locations by combining a set of canonical vectors which constitute a basis set. In the equivalent to the Cartesian system, the basis set includes the vectors [1, 0] and [0, 1], which are the vectors of length one pointing rightward and upward, respectively (*Appendix Figure 3.1B*). Using linear combination, which we saw above involves scaling and adding vectors, we can combine the Cartesian basis vectors to create any point in the plane. In *Appendix Figure 3.1B*, the point [1, 3] is created by multiplying the [1, 0] basis vector by 1, multiplying the [0, 1] basis vector by 3, and then adding the products. In general, we can label any point with the coefficients that create that point through linear combination of the basis vectors.

The basis vector approach is not limited to the standard 2- or 3-dimensional space. We can describe a space or subspace of arbitrarily large dimension. A basis set for a given space can be created from any set of vectors which span the space and which satisfy the property of linear independence. Spanning the space simply means that the

Appendix Figure 3.1 – Descriptions of Point Locations

A: The standard Cartesian coordinate system. Point [1, 3] is defined as the only point satisfying the constraint that its x-coordinate is 1 and its y-coordinate is 3.

B: Basis set consisting of the standard unit vectors pointing along the Cartesian axes. The *x*- and *y*-vectors have length 1 and point along the *x*- and *y*-axes, respectively. Locations of points are specified by multiplying the *x*- and *y*-vectors by some coefficients and then combining the products with vector addition. To derive point [1, 3], one would multiply the *x*-vector by 1 and the *y*-vector by 3 and then perform vector addition on the results. In this case, each point is defined by its coefficients rather than by coordinates. Notice that under this particular basis set the coefficients are identical to the Cartesian coordinates.

C: A basis set consisting of an *x*-vector pointing along the Cartesian *x*-axis, as in *B*, but with a new *y*-vector that has been rotated 45° clockwise compared to the *y*-vector in *B*. Under this basis set, the point [1, 3] in *B* would have coefficients [-2, 4.2426]. That is, the coefficients no longer correspond to the Cartesian coordinates. Though possibly unfamiliar, this new basis set is equally valid to that used in *B*. (Also see *Appendix 3.2.1.*)
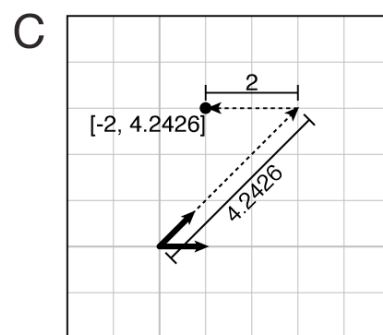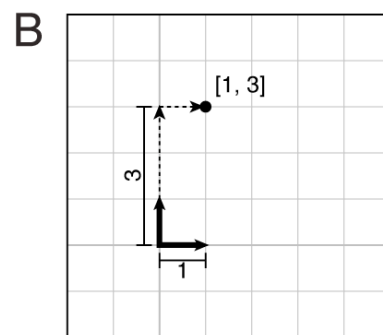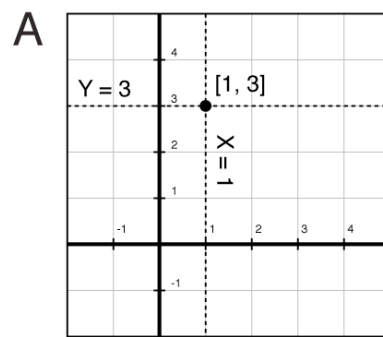
basis set is capable of describing all the points in the space. Linear independence obtains when no single one of the vectors can be created as a linear combination of the other vectors. For example, if we wish to define a 10-dimensional space, the 10 basis vectors would be linearly independent if it were impossible to scale and add together 9 of them to derive the tenth. Linear independence would fail if two or more of the vectors were to lie within the same straight line, that is, if they were collinear. Linear independence would also fail if three or more of the vectors were to lie in a common plane, and so on. Colloquially, linear independence ensures that none of the basis vectors "interferes" with the others. The basis set formalism is implicit in the standard Cartesian system. This fact can be safely ignored because the Cartesian basis vectors are orthogonal and of unit length, which causes the coefficients used to combine the basis vectors to be identical to the Cartesian coordinates, themselves (compare *Appendix Figure 3.1A* and *B*). With more unusual basis vectors, we cannot ignore the formalism. *Appendix Figure 3.1C* shows a non-Cartesian basis set consisting of the vectors [1, 0] and [0.7071, 0.7071], which have unit length and point right and up-right, respectively. In this example, the point with Cartesian coordinates [1, 3] has coefficients for this basis set of [-2, 4.2426].

The approach in GLM analysis is to treat the predictor vectors as a basis set defining a low-dimensional subspace embedded within the high-dimensional space that describes the fMRI signal data. Recalling our previous example of an fMRI experiment, the 10 predictor vectors would define a 10-dimensional subspace embedded in the 1800-dimensional space that describes the data vector for each voxel. One then finds the best representation of the data in terms of the basis set comprised of the predictor vectors by fitting the predictor vectors to the data vector. If the predictor vectors are linearly independent, the representation of the data vector will be unique; in other words, there will be a unique solution to the problem of fitting the predictor vectors to the data. On the other hand, if the predictors are not linearly independent, there will be multiple (in fact, infinitely many) solutions to the fitting problem, which is obviously undesirable. I mentioned above that there are typically far fewer predictor vectors than there are data points in the time series. In our example, we have 1800 data points but only 10 predictor vectors. This means that in trying to represent the data vector as a coordinate in the subspace defined by the predictor vectors, we are trying to represent an 1800-dimensional

vector as a 10-dimensional vector. Most often, this is impossible without losing some information. The best we can do is represent the 1800-dimensional data vector as that 10-dimensional vector (in the subspace defined by the predictors) whose position lies closest to that of the 1800-dimensional data vector (in the original 1800-dimensional space, in which is embedded the 10-dimensional predictor-based subspace). This process of finding a vector's nearest point within a subspace is performed by projection, which I describe in the next paragraph. Incidentally, projection minimizes the sum of squared errors between the estimated and actual time series, which I mentioned earlier is the criterion that the GLM fitting procedure satisfies.

Projection is easy to understand intuitively. Consider a flagpole stuck in the ground. Suppose that the pole is bent 45° toward the north-west. Let us define a 3-dimensional basis set centred on the pole's base, with the *x*-vector pointing due east, the *y*-vector pointing due north, and the *z*-vector pointing straight up (*Appendix Figure 3.2A*). Now, suppose that the sun were directly overhead the pole. In *Appendix Figure 3.2A*, the shadow cast by the pole upon the ground would be the projection of the pole, viewed as a 3-dimensional vector in our *x*-*y*-*z* system, onto the 2-dimensional *x*-*y* planar subspace defined by the flat ground. From this example, we can see that projection of a vector onto a subspace essentially separates the vector into two parts, one part that lies within the subspace and one part that is orthogonal to it. If the pole were straight and not bent, its shadow would fall right at its base, meaning that its projection onto the ground would have coordinate [0, 0] (*Appendix Figure 3.2B*). In this case, the pole would be completely orthogonal to the ground, so the *x*-*y* plane would be incapable of representing the flagpole vector at all! On the other hand, suppose the pole were bent so far over that it lay on its side. Then the pole's shadow on the ground, that is, the projection of the flagpole vector onto the *x*-*y* plane, would stretch from the pole's base to its top (*Appendix Figure 3.2C*). In this case, the flagpole vector would actually lie within the *x*-*y* planar subspace, and its projection into that subspace would simply be itself. When a vector lies within a subspace, that subspace can completely describe the vector. As the vector rotates closer to orthogonality with the subspace, the subspace's ability to represent the vector decreases.

Appendix Figure 3.2 – Projection onto a Subspace

A: A flagpole that is bent 45° toward the north-west is used to represent a 3-dimensional vector. If the sun were directly overhead, the shadow cast by the pole on the ground would represent the projection of the vector onto the 2D planar subspace that is the ground. Also shown are the standard unit x-, y-, and z-basis vectors pointing along the Cartesian axes.

B: The flagpole is now straight up and down, rather than bent. Its shadow falls at its base, meaning that the projection of the "flagpole vector" onto the ground has length zero. Because the flagpole vector is orthogonal to the 2D ground subspace, that subspace cannot capture any of the vector's structure.

C: The flagpole is tipped on its side and lies on the ground. The shadow of the tipped over flagpole runs its entire length. Mathematically, the flagpole vector is within the 2D ground subspace, and the projection of that vector onto the subspace is simply the vector, itself. When a vector lies within a subspace, that subspace can completely capture the vector's structure. (Also see *Appendix 3.2.1*.)
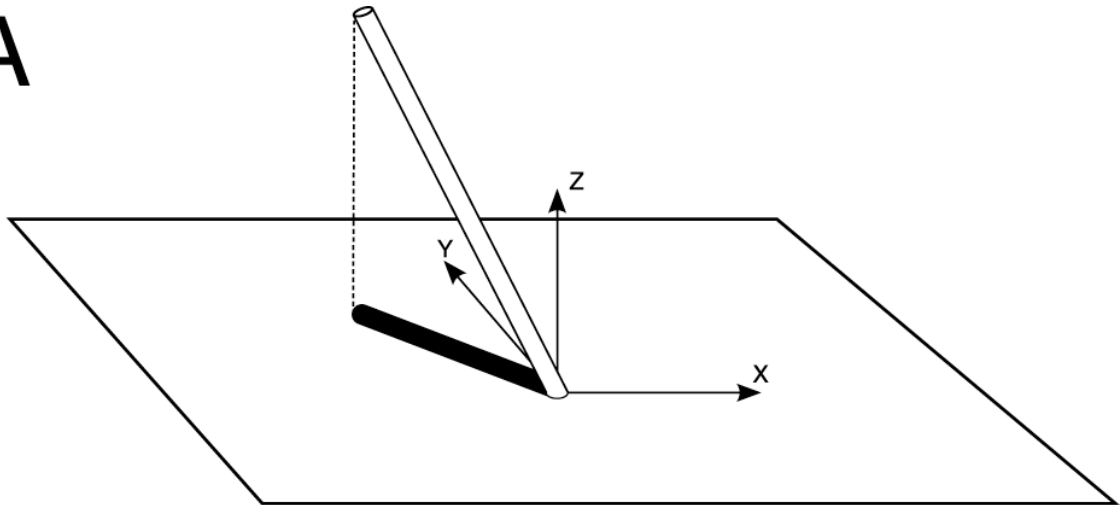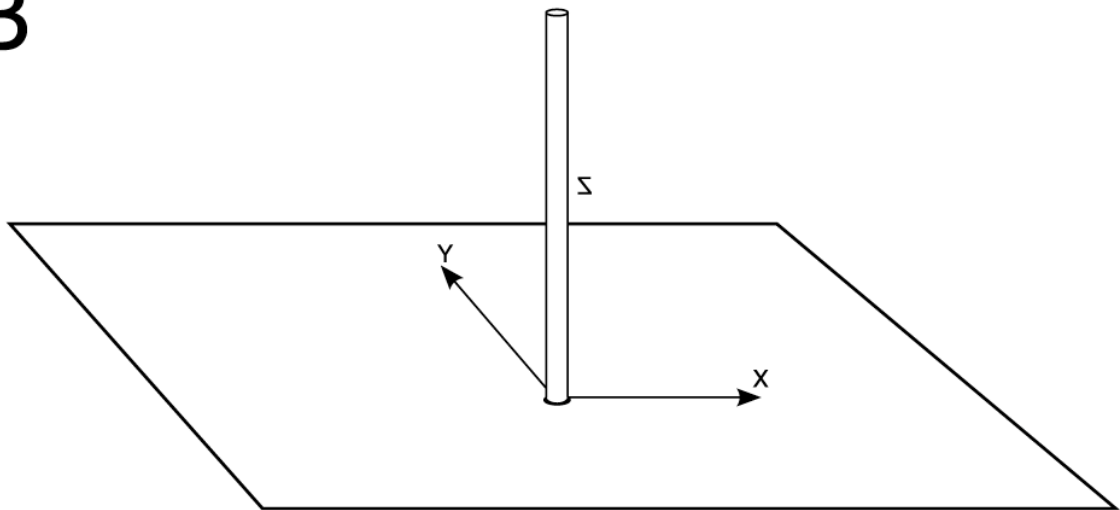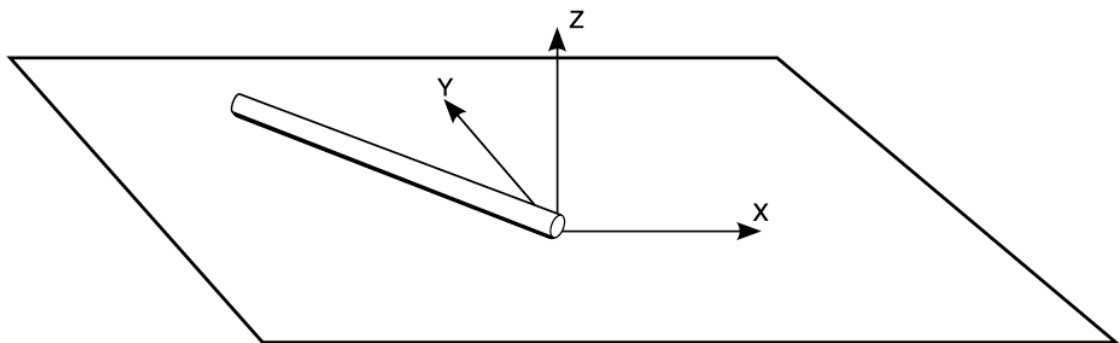
The ideas described here in terms of projecting a 3-dimensional vector onto a 2-dimensional plane can be extended to higher dimensions. In the fMRI example from before, we would project an 1800-dimensional data vector onto a 10-dimensional subspace defined by 10 predictor vectors. The reason that it is possible to represent an 1800-dimensional data vector fairly well with a mere 10 predictor vectors is that fMRI data vectors are not "allowed" to occupy arbitrary positions within their space. Whereas an 1800-dimensional vector with random element values can occupy any position in the 1800-dimensional space, an 1800-dimensional fMRI data vector is generated by physical processes that constrain its shape, that is its location in 1800-dimensional space. This makes it possible to define a set of predictor functions capable of capturing a substantial 'amount' of a BOLD signal time series' shape. It is worth highlighting, though, that a poor choice of predictor vectors on the experimenter's part can render the subspace defined by the predictors incapable of representing a given data vector well. The GLM would then fail to capture important aspects of the experiment's time series. Even if the set of predictor vectors is capable of representing the major components of a time series, the exact form of that representation can be unstable and/or uninterpretable if the predictor vectors are very non-orthogonal, as discussed in the next section.

**Appendix 3.2 – Non-orthogonality and Poor Conditioning in the GLM**

I have two points to make on the practical implementation of the general linear model (GLM), which was described in the last section. The first is that many possible sets of predictor vectors can define the same subspace. It is possible to change the predictors, and thus the coefficients that specify a given point's location, without changing the actual subspace described by the predictors. Consider again the bent flagpole example (from *Appendix 3.1*), in which we projected the flagpole onto the ground by means of its shadow. The original example uses a Cartesian *x-y-z* basis set with the *x*-vector pointing east, the *y*-vector pointing north, and the *z*-vector pointing up (*Appendix Figure 3.3A*). Here, the *x*- and *y*-vectors also define a planar subspace, namely the ground. We could define the *x*-vector as pointing south-east and the *y*-vector as pointing west, and we could also make the *x*-vector longer and the *y*-vector shorter (*Appendix Figure 3.3B*). This new *x-y* basis set would define precisely the same planar

Appendix Figure 3.3 – Subspaces and Basis Sets

A: This is identical to *Appendix Figure 3.2A*. Note the standard unit *x*-, *y*-, and *z*-basis vectors pointing along the Cartesian axes.

B: This is identical to *A* except that the *x*- and *y*-basis vectors have undergone changes of length as well as rotation within the 2D planar ground subspace. Importantly, these new *x*- and *y*-basis vectors define the same 2D ground subspace as in *A*, though the coefficients that define specific points would be different in *B* compared to *A*.

C: The *x*- and *y*-basis vectors have undergone rotations in 3 dimensions, and are no longer constrained to the 2D ground subspace in *B*. These new *x*- and *y*-basis vectors define a different 2D subspace from that in *A* or *B*. Notice that the projection (shadow) of the flagpole onto this new subspace is also different compared to *A* and *B* because the new subspace is closer to being parallel with the flagpole. (Also see *Appendix 3.2.1*.)
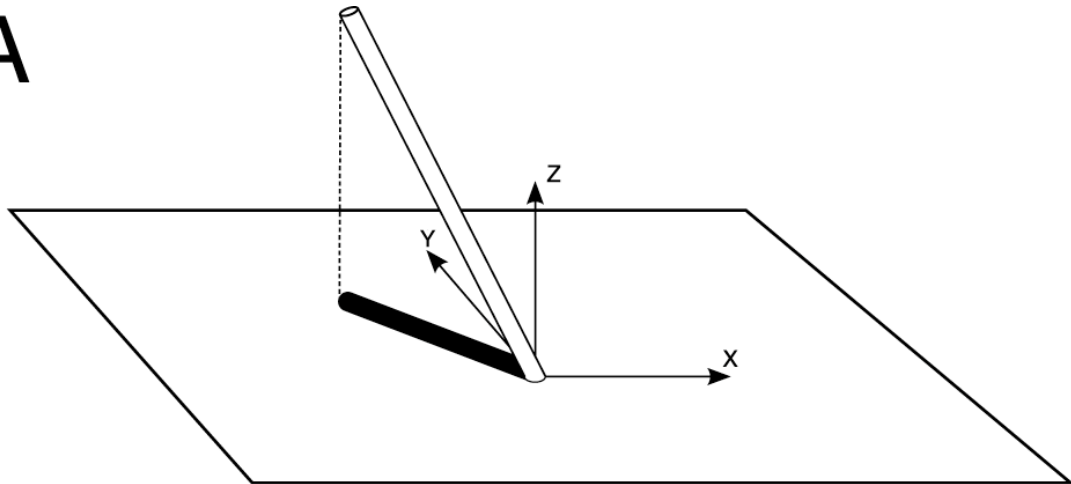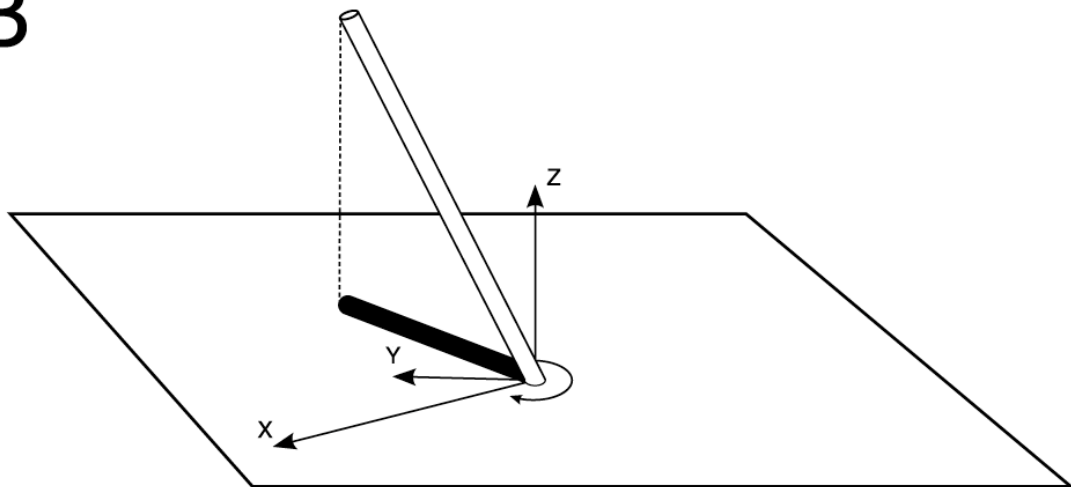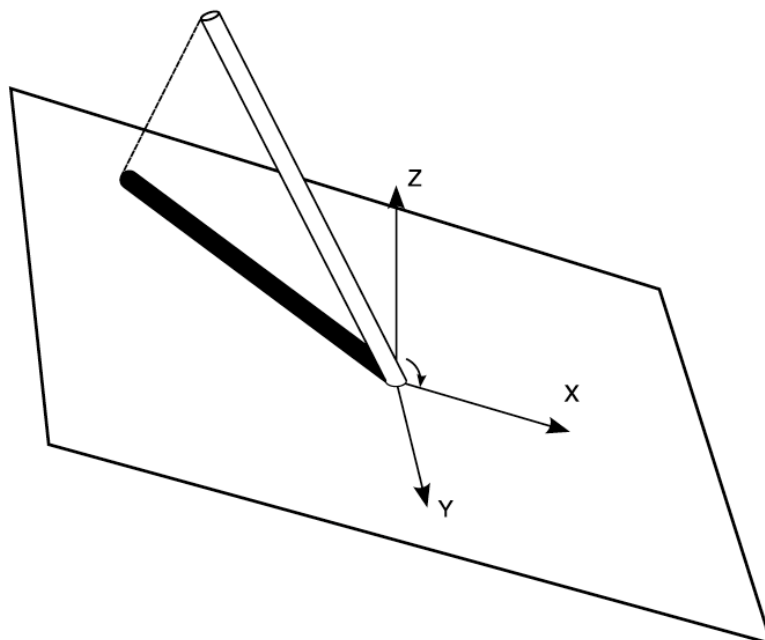
subspace (that is, the ground) as in the original system. The new basis set would also be as capable of describing positions in the plane as the original, though a point's coefficients would change depending upon which basis set was used. Notice that the proposed manipulation of the *x*- and *y*-vectors would only involve rotating them within the *x-y* plane (and changing their lengths). If we redefined the *x*- and *y*-vectors to include a non-zero height component (along the *z*-vector), we would actually change the subspace defined by those vectors such that it would no longer correspond to the ground (*Appendix Figure 3.3C*). To some extent, the choice of basis set is a matter of convenience and of what one wants to do with it. In fMRI experiments, it is typical to define predictor vectors that each represent one task event. When comparing prosaccades and antisaccades, one might include a predictor for prosaccade trials and one for antisaccade trials. This would allow for localization of activation changes specific to prosaccade performance and specific to antisaccade performance. However, if one were interested only in the difference between prosaccades and antisaccades, one might include a predictor for prosaccade trials and a predictor for the difference between antisaccade and prosaccade trials. Theoretically, one could use strange predictors such as one modeling 30% of prosaccade trial activation and 70% of antisaccade trial activation and a second predictor modeling 70% of prosaccade activation and -30% of antisaccade activation. This predictor set would be mathematically sound, but it would be unnatural to interpret.

The basis vectors in the standard Cartesian basis set are orthogonal, or at right angles, to each other. Orthogonal systems have the desirable property that the vectors do not 'interfere' with each other in the task of describing a position coordinate. Movement along the *y*-direction in the Cartesian system affects only the *y*-coefficient, and the *x*-coefficient remains unchanged (*Appendix Figure 3.4A*). Non-orthogonal systems do not have this property. Suppose we were to create a new system by rotating the *y*-vector 45° clockwise and keeping the *x*-vector unchanged (*Appendix Figure 3.4B*). In the new system, upward motion along the up-down direction, would increase the *y*-coordinate and decrease the *x*-coordinate to offset the left-right aspect of the new *y*-axis. This phenomenon in which a change along one basis vector must be offset by changes

Appendix Figure 3.4 – Basis Sets and Orthogonality

A: Basis set consisting of the standard *x*- and *y*-vectors with unit length pointing along the Cartesian axes. Notice that in moving from point [1, 2] to point [1, 3], only the *y*-coefficient changes. This is due to the orthogonality of the basis vectors. Solid black lines with arrowheads are the basis vectors. Dotted lines with arrowheads illustrate the linear combination (scaling and addition) of the basis vectors to derive the points' locations.

B: Basis set consisting of the *x*-vector from *A* and a *y*-vector that is rotated 45° clockwise relative to that in *A*. The vectors in this basis set are now 45° apart, rather than orthogonal. Now, the points [1, 2] and [1, 3] from *A* have coefficients [-1, 2.8284] and [-2, 4.2426]. Notice that moving from the first to the second point involves a change in the *x*-coefficient as well as the *y*-coefficient, which results from the non-orthogonality of the basis-vectors. Other conventions as in *A*.

C: Extreme example of non-orthogonality in which the basis set includes the *x*-vector from *A* and a *y*-vector that is only 1° away from the *x*-vector. In the illustration, the *x*- and *y*-vectors cannot be resolved because they almost completely overlap each other. Points [1, 2] and [1, 3] from *A* now have coefficients [-113.6, 114.6] and [-170.9, 171.9]. The large absolute values of the coefficients and the large differences in coefficients for points that are only 1 unit apart result from the basis vectors' being almost co-linear. Other conventions as in *A*. The dotted lines with arrowheads are only partly shown as they would extend far past of the edge of the page. (Also see *Appendix 3.2.2*.)

A

[1, 3]

[1, 2]

B

[-2, 4.2426]

[-1, 2.8284]

C

[-170.9, 171.9]

[-113.6, 114.6]
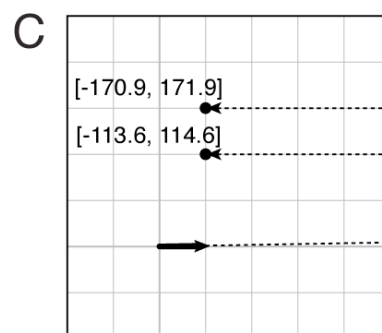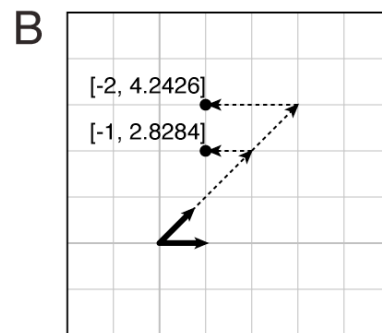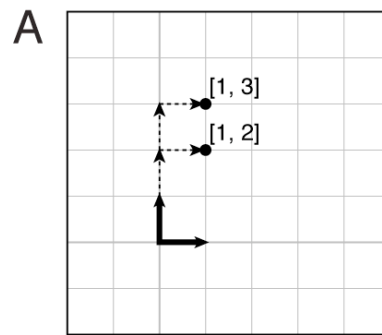
along the other basis vectors becomes more pronounced as the vectors rotate closer together, that is, as the vectors deviate further from orthogonality. Imagine a basis set consisting of the *x*-vector as before but with a *y*-vector pointing only 1° counter-clockwise from the *x*-vector (*Appendix Figure 3.4C*). In this example, the *x*- and *y*-vectors would be almost co-linear. This system would still be capable of describing any position in the plane, but the coordinates would tend toward extremely large values. For example, consider location [1, 2] in Cartesian coordinates. In the new, strange *x*-*y* system this point would have coefficients [-113.5799, 114.5974]. The coefficients in the nearly co-linear *x*-*y* system would also change drastically with small changes in position along certain directions (*Appendix Figure 3.4C*). Moving position 1 unit upward, to location [1, 3] in Cartesian coordinates, the new coefficients in the nearly co-linear *x*-*y* system would be [-170.8699, 171.8961]. The nearly co-linear basis set is clearly very sensitive to small movements along the up-down axis, and this is caused by the *x*- and *y*-vectors' being almost co-linear.

This brings me to one last point on the GLM, which is that when fitting a GLM to data, the solution becomes more delicate and more susceptible to noise as the predictor vectors deviate from orthogonality. This is not a large problem as long as the predictor vectors are at least 45° apart (see *Appendix 3.4*), which corresponds to the predictor vectors' not having excessively high covariance. In extreme cases, such as when two or more predictor vectors point in very similar directions, the GLM fitting solution can be determined almost entirely by noise, giving rise to 'garbage' results. The example of the nearly co-linear *x*-*y* basis set illustrates this phenomenon. This is called the problem of poor conditioning. Note that the predictor vectors in a poorly conditioned GLM design matrix are still linearly independent (though they are not far from linear dependence). It is still (barely) possible to compute a solution to fit the GLM matrix to the data. The problem is that the solution will be extremely vulnerable to noise in the data. So, in building GLM design matrices, we must avoid not only the problem of linear dependence (discussed in *Appendix 3.1* above), which will prevent any GLM fitting solution what-so-ever, but also the problem of poor conditioning, which will make the solution overly susceptible to noise in the data. In the *Section 1.2.4*, I discuss how randomizing, or

jittering, the time interval between adjacent task events in a rapid event-related design can be used to avoid poorly-conditioned GLM predictor sets.

Appendix 3.3 – Convolution

In discrete form, given two signals $x(\tau)$ and $y(\tau)$, where τ is a whole number indexing time, the convolution of x and y at a given time t is

$$x \otimes y(t) = \sum_{\tau \in [0,t]} x(\tau)y(t-\tau)$$

The continuous form of convolution is defined analogously using integration rather than summation.

Appendix 3.4 – Angle Between Two Vectors

Given column vectors x and y, the angle between them, a, is

$a = \cos^{-1}\left(\frac{x^T \cdot y}{\|x\|\|y\|}\right)$, where superscript T denotes the transpose and $\|x\|$ is the

norm or length of vector x.

Appendix 3.5 – Bibliography

Abeles M. 1991. *Corticonics: Neural Circuits of the Cerebral Cortex*. Cambridge: Cambridge University Press.

Astrup J, Heuser D, Lassen NA, Nilsson B, Norberg K, and Siesjo BK. 1978. Evidence against H+ and K+ as main factors for the control of cerebral blood flow: a microelectrode study. *Ciba Found Symp* (56): 313-337.

Attwell D and Laughlin SB. 2001. An energy budget for signaling in the grey matter of the brain. *J Cereb Blood Flow Metab* 21(10): 1133-1145.

Attwell D and Iadecola C. 2002. The neural basis of functional brain imaging signals. *Trends Neurosci* 25(12): 621-625.

Bandettini PA, Wong EC, Hinks RS, Tikofsky RS, and Hyde JS. 1992. Time course EPI of human brain function during task activation. *Magn Reson Med* 25(2): 390-397.

Boynton GM, Engel SA, Glover GH, and Heeger DJ. 1996. Linear systems analysis of functional magnetic resonance imaging in human V1. *J Neurosci* 16(13): 4207-4221.

Braitenberg V and Schuz A. 1998. *Cortex: Statistics and Geometry of Neuronal Connectivity*. Springer Verlag.

Buxton RB and Frank LR. 1997. A model for the coupling between cerebral blood flow and oxygen metabolism during neural stimulation. *J Cereb Blood Flow Metab* 17(1): 64-72.

Dale AM and Buckner RL. 1997. Selective averaging of rapidly-presented individual trials using fMRI. *Hum Brain Mapp* 5: 329-340.

Dale AM. 1999. Optimal experimental design for event-related fMRI. *Hum Brain Mapp* 8(2-3): 109-114.

Fergus A and Lee KS. 1997. GABAergic regulation of cerebral microvascular tone in the rat. *J Cereb Blood Flow Metab* 17(9): 992-1003.

Fox PT and Raichle ME. 1986. Focal physiological uncoupling of cerebral blood flow and oxidative metabolism during somatosensory stimulation in human subjects. *Proc Natl Acad Sci U S A* 83(4): 1140-1144.

Fox PT, Raichle ME, Mintun MA, and Dence C. 1988. Nonoxidative glucose consumption during focal physiologic neural activity. *Science* 241(4864): 462-464.

Friston KJ, Jezzard P, and Turner R. 1994. Analysis of functional MRI time series. *Hum Brain Mapp* 1: 153-171.

Friston KJ, Josephs O, Rees G, and Turner R. 1998. Nonlinear event-related responses in fMRI. *Magn Reson Med* 39(1): 41-52.

Huettel SA, Song AW, and McCarthy G. 2004. *Functional Magnetic Resonance Imaging*. Sunderland, MA: Sinauer Associates.

Kwong KK, Belliveau JW, Chesler DA, Goldberg IE, Weisskoff RM, Poncelet BP, Kennedy DN, Hoppel BE, Cohen MS, Turner R, and et al. 1992. Dynamic magnetic resonance imaging of human brain activity during primary sensory stimulation. *Proc Natl Acad Sci U S A* 89(12): 5675-5679.

Logothetis NK, Pauls J, Augath M, Trinath T, and Oeltermann A. 2001. Neurophysiological investigation of the basis of the fMRI signal. *Nature* 412(6843): 150-157.

Logothetis NK. 2002. The neural basis of the blood-oxygen-level-dependent functional magnetic resonance imaging signal. *Philos Trans R Soc Lond B Biol Sci* 357(1424): 1003-1037.

Logothetis NK and Wandell BA. 2004. Interpreting the BOLD signal. *Annu Rev Psychol* 66: 735-769.

Mintun MA, Lundstrom BN, Snyder AZ, Vlassenko AG, Shulman GL, and Raichle ME. 2001. Blood flow and oxygen delivery to human brain during functional activity: theoretical modeling and experimental data. *Proc Natl Acad Sci U S A* 98(12): 6859-6864.

Miyashita Y. 1988. Neuronal correlate of visual associative long-term memory in the primate temporal cortex. *Nature* 335(6193): 817-820.

Ogawa S, Tank DW, Menon R, Ellermann JM, Kim SG, Merkle H, and Ugurbil K. 1992. Intrinsic signal changes accompanying sensory stimulation: functional brain mapping with magnetic resonance imaging. *Proc Natl Acad Sci U S A* 89(13): 5951-5955.

Pinard E, Tremblay E, Ben-Ari Y, and Seylaz J. 1984. Blood flow compensates oxygen demand in the vulnerable CA3 region of the hippocampus during kainate-induced seizures. *Neuroscience* 13(4): 1039-1049.

Powers WJ, Hirsch IB, and Cryer PE. 1996. Effect of stepped hypoglycemia on regional cerebral blood flow response to physiological brain activation. *Am J Physiol* 270(2 Pt 2): H554-559.

Rabiner R and Gold B. 1975. *Theory and Application of Digital Signal Processing*. Englewood Cliffs, New Jersey: Prentice-Hall.

Sakai K and Miyashita Y. 1991. Neural organization for the long-term memory of paired associates. *Nature* 354(6349): 152-155.

Scannell JW and Young MP. 1999. Neuronal population activity and functional imaging. *Proc Biol Sci* 266(1422): 875-881.

Strang G. 1988. *Linear Algebra and its Applications*. Brooks Cole.

Appendix 4
Details of Analysis for Experiment 1

Appendix 4.1 – Running Lines Smoother

We used a running lines smoother, which is a type of non-linear temporal filter, to remove low-frequency trends from the data. This algorithm was based on that described in Marchini and Ripley (2000). The algorithm operated on a voxel-by-voxel basis, removing low frequency trends from each functional run separately. For each time point in a time course, a line was fit by least squares regression to the 22 nearest neighbours on either side of the point (ie. we used a computational window 45 points wide), and the centre point of the fitted line was taken as an estimate of the centre point of the data window. For time points less than 22 points from the beginning or end of the time course, the window was truncated so as not to extend beyond the edge of the time course. The values thus computed across all points in the time course constituted an estimate of the low frequency trends, which we then subtracted from the original time course. The running lines smoother has the advantage of being robust against edge effects (Marchini and Ripley 2000).

Because a running lines smoother, like many other temporal filters, is sensitive to large isolated spikes in the signal, we removed spikes from each time course as follows. We defined a spike as a point greater than 3 standard deviations from the mean of the trend-removed signal. In the presence of severe low frequency trends, this 3 standard deviation criterion does not distinguish between spikes and time points that lie far from the mean by virtue of the trend. Therefore, it is necessary to compute an interim trend-removed time course upon which to perform spike removal. We computed an interim estimate of the low frequency trend function using the running lines smoother and subtracted this estimate from the original data time course. Any point more than 3 standard deviations from the trend-removed mean signal was classified as a spike and replaced with a linear estimate based on the point's two nearest neighbours. The interim low frequency trend function was then added back into the time course, and the running lines smoother was used again, this time on the spike-removed signal. The reason for this was that any spikes would have perturbed the running lines interim estimate of the low

frequency trends, making it was necessary to re-estimate the trend function based on the spike-removed time course.

Appendix 4.2 – Autocovariance Computation

General linear model (GLM) statistical methodology, which is widely used to analyze functional magnetic resonance imaging (fMRI) data, assumes that additive noise in the fMRI signal is white, or uncorrelated with itself. Violation of this assumption, which is typical with fMRI data sets, increases the false positive rate of GLM statistics, and techniques have been developed to address autocovariance, or colouring, in fMRI noise to avoid this problem. Worsley and colleagues (2002, Appendix A) describe a fast, elegant method for computing the parameters of a noise autocovariance model for fMRI time series. This model can then be used to prewhiten the fMRI data and GLM design matrix to avoid problems with inflated false positive rate. Worsley and colleagues' method has the advantage of producing less biased autocovariance parameter estimates compared to other techniques that estimate autocovariance from residuals without accounting for the effects of projecting the data onto the GLM design matrix subspace (see Worsley et al. 2002, section 3.3).

In Worsley and colleagues' (2002) scheme, a general linear model is fit voxel by voxel to the activation time series from each run within an experimental session, and autocovariance parameters are estimated from the residuals. One potential disadvantage of fitting a GLM to individual runs is that infrequent events, with few instances occurring in each run, cannot be modeled accurately. For example, in an fMRI comparison of two trials types, A and B, with 20 A trials and 20 B trials in each run and 10 runs in an experimental session, suppose that subjects make interesting errors on 10% of trials and that the experimenter wants to model these in the GLM. Models fit to individual runs would include only 2 A trial errors and 2 B trial errors, on average, leading to high variance in error trial parameter estimates. A model fit to all 10 runs from a session would include 20 error trials for tasks A and B, yielding error trial parameter estimates with lower variance. If one also wanted to use Worsley and colleagues' (2002) method to estimate autocovariance parameters in this experiment, one would have to apply it to the concatenated time series constructed from all 10 runs within a session, thereby making

the tacit assumption that covariance parameters remain constant across runs. Here, we describe a simple extension of Worsley and colleagues' (2002) method for autocovariance estimation that avoids this assumption. The modified procedure computes run-specific autocovariance parameters based on the residuals from a GLM that estimates task activation parameters from all the runs in a single session. This approach assumes that activation parameters, but not autocovariance parameters, are constant across runs in the experimental session.

Worsley and colleagues compute unbiased estimates of the noise covariance parameters in a single functional run by accounting for the fact that removing the effects of interest (as modeled by the GLM) from the fMRI time series changes the covariance structure of the residuals compared to that of the original signal noise (see Worsley et al. 2002, section 3.3). We use a similar approach to compute run-specific noise covariance parameters based on the residuals derived from a GLM incorporating all runs in an experimental session. That is, we assume that parameters governing effects of interest are constant across the session, but we allow noise autocovariance parameters to vary from run to run. The method is massively univariate, or applicable to each fMRI voxel separately.

Let y_i be a column vector containing the fMRI time series for run i. Index i is an integer from 1 to n, where n is the number runs within the experimental session. Let y be the concatenated time series from all runs 1 through n: $y' = [y_1', y_2', \ldots, y_n']$. Without loss of generality, we assume that the length of all runs y_i is a constant l. Let $\mathbf{X}$ be the nl by m GLM design matrix, incorporating m predictor curves represented as column vectors of length nl. Then the ordinary least squares solution for the beta weights is given by

$$\hat{\beta} = \left(\mathbf{X'X}\right)^{-1}\mathbf{X'}y,$$

and the residual vector r is given by

$$r = y - \mathbf{X}\hat{\beta} = \left(\mathbf{I} - \mathbf{X}(\mathbf{X'X})^{-1}\mathbf{X'}\right)y,$$

where $\left(\mathbf{I} - \mathbf{X}(\mathbf{X'X})^{-1}\mathbf{X'}\right)$ is the nl by nl residual forming matrix, which we will call $\mathbf{R}$.

Now, let r_i be the length l vector of residuals from run i. And let $\mathbf{R}_i$ be the l by nl residual forming matrix for run i. That is, $\mathbf{R}_i$ is the l by nl submatrix of $\mathbf{R}$ corresponding to run i. Then we have

$$r_i = \mathbf{R}_i y.$$

Recall that y is the time series built by concatenating all n runs. Let $\mathbf{V}$ be the nl by nl variance-covariance matrix of y. Then, the variance-covariance matrix of the vector r_i is

$$\mathbf{R}_i \mathbf{V} \mathbf{R}_i'.$$

Consider the quadratic form $r_i' \mathbf{D}_\lambda r_i$, where $\mathbf{D}_\lambda$ is an l by l Toeplitz matrix with ones in the λ^{th} upper diagonal and λ is a lag parameter ranging from 0 to p, where p is the order of the covariance model as specified by the experimenter. $\mathbf{D}_\lambda$ can be thought of as a "shifting matrix", and $r_i' \mathbf{D}_\lambda r_i$ is the sum of the product of residuals separated by lag λ. Note that Worsley and colleagues (2002) use $\mathbf{D}_l$ instead of $\mathbf{D}_\lambda$; we use λ here because we used l to denote run length above.

The expectation of $r_i' \mathbf{D}_\lambda r_i$ is given by

$$E\left(r_i' \mathbf{D}_\lambda r_i\right) = Tr(\mathbf{D}_\lambda \mathbf{R}_i \mathbf{V} \mathbf{R}_i') = Tr(\mathbf{R}_i' \mathbf{D}_\lambda \mathbf{R}_i \mathbf{V}),$$

where Tr denotes the trace of the matrix. The second equality is based on the fact that $Tr(\mathbf{AB}) = Tr(\mathbf{BA})$. We now approximate the variance-covariance matrix $\mathbf{V}$ as a block diagonal matrix with n submatrices $\mathbf{V}_k$, each l by l in size, along the diagonal. $\mathbf{V}_k$ models the variance-covariance matrix for run k and is a Toeplitz matrix with elements $v_{k,j}$ along the main diagonal ($j = 0$) or the j^{th} off-diagonal ($j \in [1, p]$). Approximating $\mathbf{V}$ with the submatrices $\mathbf{V}_k$:

$$\mathbf{V} = \sum_{k=1}^{n} \mathbf{V}_k,$$

Then, $E\left(r_i' \mathbf{D}_\lambda r_i\right) = Tr(\mathbf{R}_i' \mathbf{D}_\lambda \mathbf{R}_i \mathbf{V}) \approx \sum_{k=1}^{n} Tr(\mathbf{R}_i' \mathbf{D}_\lambda \mathbf{R}_i \mathbf{V}_k)$

Now, $\mathbf{V}_k$ is all zeroes except for a $2p+1$ by $2p+1$ block of parameters $v_{k,j}$ in the k^{th} block position along the main diagonal (remember, we model $\mathbf{V}$ as a block diagonal matrix). In this case, $\mathbf{R}_i' \mathbf{D}_\lambda \mathbf{R}_i \mathbf{V}_k = \mathbf{R}_{i,k}' \mathbf{D}_\lambda \mathbf{R}_{i,k} \mathbf{V}_k$, where $\mathbf{R}_{i,k}$ is a submatrix of $\mathbf{R}_i$

corresponding to run k. For example, suppose we had $n=5$ runs of length $l=100$ time points. Then the residual-forming matrix $\mathbf{R}$ for the whole data set would be of size 500 by 500, and each of the 5 run-specific residual-forming submatrices $\mathbf{R}_i$ would be of size 100 by 500. $\mathbf{R}$ would be composed of the $\mathbf{R}_i$ submatrices stacked vertically, one of top of the other. Each of the $\mathbf{R}_i$ submatrices would in turn be composed of 5 submatrices $\mathbf{R}_{i,k}$, which would have size 100 by 100, arranged horizontally side-by-side. Then we have

$$E\left(r_i'\mathbf{D}_\lambda r_i\right) = Tr(\mathbf{R}_i'\mathbf{D}_\lambda \mathbf{R}_i \mathbf{V}) \approx \sum_{k=1}^{n} Tr(\mathbf{R}_i'\mathbf{D}_\lambda \mathbf{R}_i \mathbf{V}_k) = \sum_{k=1}^{n} Tr(\mathbf{R}_{i,k}'\mathbf{D}_\lambda \mathbf{R}_{i,k} \mathbf{V}_k) =$$

$$\sum_{k=1}^{n}\left(Tr(\mathbf{R}_{i,k}'\mathbf{D}_\lambda \mathbf{R}_{i,k})v_{k,0} + \sum_{j=1}^{p} Tr(\mathbf{R}_{i,k}'\mathbf{D}_\lambda \mathbf{R}_{i,k}(\mathbf{D}_j + \mathbf{D}_j'))v_{k,j} \right)$$

Note that k and i both index run number and range from 1 to n. We had to use two run indices to keep track of independently indexed run specific residual vectors r_i and run specific covariance matrices $\mathbf{V}_k$ and covariance parameters $v_{k,j}$.

By equating the product $r_i'\mathbf{D}_\lambda r_i$ with its expectation and using the above equation, we create a system of $(p+1)n$ equations in $(p+1)n$ unknowns, where the unknowns are the $(p+1)n$ variance-covariance parameters $v_{k,j}$.

$$a_{i,\lambda} = r_i'\mathbf{D}_\lambda r_i$$

$$m_{i,\lambda,k,j} = \begin{cases} Tr(\mathbf{R}_{i,k}'\mathbf{D}_\lambda \mathbf{R}_{i,k}) & j = 0 \\ Tr(\mathbf{R}_{i,k}'\mathbf{D}_\lambda \mathbf{R}_{i,k}(\mathbf{D}_j + \mathbf{D}_j')) & 1 \le j \le p \end{cases}$$

where $i \in [1, n]$, $\lambda \in [0, p]$, $k \in [1, n]$, and $j \in [0, p]$ in the integers.

$$\mathbf{a}' = \left[a_{1,0}, \ldots, a_{1,p}, a_{2,0}, \ldots, a_{2,p}, \ldots, a_{n,0}, \ldots, a_{n,p} \right]$$

$$\mathbf{M} = \begin{bmatrix} m_{1,0,1,0},\ldots,m_{1,0,1,p},m_{1,0,2,0},\ldots,m_{1,0,2,p},\ldots,m_{1,0,n,0},\ldots,m_{1,0,n,p} \\ \vdots \\ m_{1,p,1,0},\ldots,m_{1,p,1,p},m_{1,p,2,0},\ldots,m_{1,p,2,p},\ldots,m_{1,p,n,0},\ldots,m_{1,p,n,p} \\ \vdots \\ m_{n,0,1,0},\ldots,m_{n,0,1,p},m_{n,0,2,0},\ldots,m_{n,0,2,p},\ldots,m_{n,0,n,0},\ldots,m_{n,0,n,p} \\ \vdots \\ m_{n,p,1,0},\ldots,m_{n,p,1,p},m_{n,p,2,0},\ldots,m_{n,p,2,p},\ldots,m_{n,p,n,0},\ldots,m_{n,p,n,p} \end{bmatrix}$$

$$\hat{\mathbf{v}}' = \left[\hat{v}_{1,0},\ldots,\hat{v}_{1,p},\hat{v}_{2,0},\ldots,\hat{v}_{2,p},\ldots,\hat{v}_{n,0},\ldots,\hat{v}_{n,p} \right]$$

Then, to estimate the variance-covariance parameters, we compute

$$\hat{\mathbf{v}} = \mathbf{M}^{-1}\mathbf{a} .$$

To estimate the correlation coefficient at lag λ for run k, we use

$$\hat{\rho}_{k\lambda} = \frac{\hat{v}_{k\lambda}}{\hat{v}_{k0}} .$$

The correlation coefficients having been estimated, pre-whitening matrices can be computed for each functional run using the method described in Worsley and colleagues (2002, Appendix A.3). Each functional run's length l time series is then pre-whitened by pre-multiplication with the run-specific pre-whitening matrix. The nl by m GLM design matrix $\mathbf{X}$ is pre-whitened by pre-multiplication with the matrix formed by vertically concatenating the run-specific pre-whitening matrices. Ordinary least squares regression of the pre-whitened data against the pre-whitened design matrix then provides estimates of the beta weights, with noise autocovariance taken into account.

Appendix 4.3 – Details of Statistical Analysis

Our analysis was based on the methods described in Worsley and colleagues (2002). This technique uses a hierarchical model of the data, which estimates effect size and covariance parameters at the single functional run level, the single subject level, and the between subject level. We modified this technique slightly and modeled effects parameters at the within and between subjects levels only.
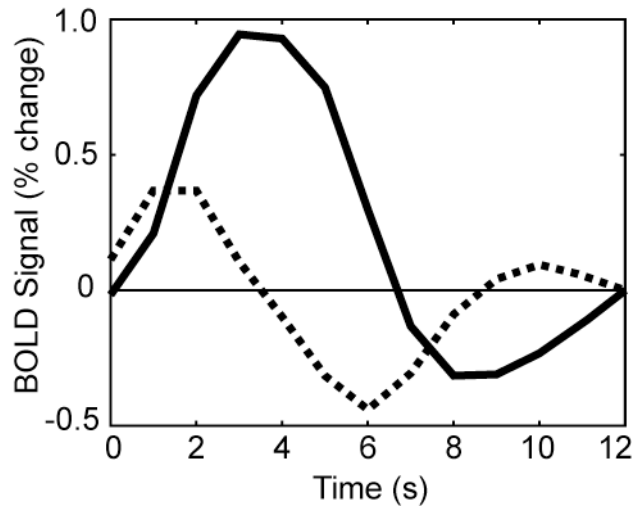
Appendix Figure 4.1 – Haemodynamic Response Model

The model of the haemodynamic response function in Experiment 1 (*Chapter 2*) consisted of two curves, a difference of gamma functions (solid line) and its temporal derivative (dotted line). This model was convolved with trial-locked impulse sequences to derive the predictor functions used in the statistical analysis for Experiment 1. See *Section 2.2 – Methods* and *Appendix 4.3* for details.

This was done to allow the use of nuisance predictors (see below) in the single subject general linear models (GLMs) since the nuisance events, namely late response trials and other discarded trials, were too rare in many subjects to yield good estimates of their effect sizes had they been modeled for each functional run individually. This modification assumed that effects parameters were fairly constant across all runs from a given subject.

A GLM was computed for each subject as follows. All computations described here were coded in the Matlab 6 numerical analysis environment. We modeled the haemodynamic response function (HRF) with a 'basis' set comprised of the difference of two gamma functions ($\Gamma_1(t) - \Gamma_2(t)$) and its temporal derivative (*Appendix Figure 4.1*). Here,

$$\Gamma_i(t) = s_i \left((t - \delta_i)/\tau_i \right)^{n_i - 1} \exp\left(-(t - \delta_i)/\tau_i \right) / \tau_i (n_i - 1)!, \ t \in [0,11],$$

$$s_1 = 5.3, \ \tau_1 = 1.5, \ \delta_1 = 0.35, \ n_1 = 3$$

$$s_2 = 2.9, \ \tau_2 = 1.32, \ \delta_2 = 5, \ n_2 = 3.$$

The temporal derivative was included to allow the model to address differences in BOLD response latency between regions and between subjects (Friston et al. 1998). BOLD signal time courses were modeled as the convolution of the above HRF model with a set of 10 impulse response sequences (Dale 1999). The impulse sequences were zero-baseline sequences containing a one for each impulse. Three of the sequences contained impulses locked to the first time point of each instruction period for correct prosaccades, correct antisaccades, or correct nogo trials, respectively. The three sets of (two) predictor curves generated by convolving these impulse sequences with the HRF model we called *pro_instruction*, *anti_instruction*, and *nogo_instruction*. Three of the impulse sequences contained impulses during the first time point of the response period for correct prosaccades, correct antisaccades, or correct nogo trials, and the corresponding predictor curve sets were called *pro_response*, *anti_response*, and *nogo_response*. Two impulse sequences contained impulses during the instruction period and response period, respectively, of any trial with a late response (> 500 ms latency). Two impulse sequences contained impulses during the instruction period and response period, respectively, for discarded trials containing various types of mistake including incorrect response, lack of response when required, or inappropriate fixation break (see *Section 2.2.5 – Behavioural Analysis*). In total, the statistical model included a 'basis' set with two functions, two epochs (instruction period and response period) for each group of trials, and five groups of trials (correct prosaccades, correct antisaccade, correct nogo trials, late trials of any kind, and discarded trials of any kind) for a total of 20 predictor curves. The two sets of four predictor curves for late response trials and discarded trials were the nuisance

predictors alluded to above. We also included a constant offset predictor composed entirely of ones. Therefore, each subject's design matrix consisted of 21 columns, each containing one predictor curve, and 316 x *nr* rows, where *nr* was the number of functional runs recorded for the subject. We analyzed statistically only the predictor weights for correct trials but included predictors for late and discarded trials to reduce residual variance of the model and to reduce contamination of the correct trial predictors by partially overlapping error trial activation.

For each subject, we computed minimum variance unbiased estimates of the beta weights for each curve after pre-whitening the data and design matrix. (The beta weights are the appropriate scaling or weighting that one must apply to the predictor curves in the design matrix to yield the best model of the data. In this case, the 'best' model is that which minimizes residual squared error after pre-whitening.) Pre-whitening the data and design matrix compensates for coloured noise in the data, which is important to avoid the inflation of type I error that coloured noise introduces into standard (unweighted) least squares GLM statistics (Bullmore et al. 2001). Pre-whitening followed by least squares regression is a standard technique in statistical packages like SPM2 and AFNI. We first fit the design matrix to the data by (unweighted) least squares regression. From the unweighted least squares model, we used an extension of the method described in Worsley and colleagues (2002) to compute, on a voxel-by-voxel basis, the parameters of a fifth order noise autocovariance model which included individual autocovariance parameters for each run recorded from the given subject (see Worsley et al. 2002, Appendix A.1-3). We chose a fifth order model because the data contained appreciable autocovariance structure at fifth and lower order but very little sixth and higher order autocovariance structure. The covariance parameters were smoothed across the scanned volume using a 4 mm FWHM Gaussian filter to reduce noise in the parameter estimates. For each voxel, we whitened both the voxel time course and the design matrix of the given subject using the voxel-specific covariance model in accordance with Worsley and colleagues (2002, Appendix A.3). We then fit the whitened design matrix to the whitened data using least squares.

Having computed estimates of the beta weights for each voxel for each subject, we derived 3D massively univariate T-statistic maps for twelve different contrasts. Six of

the contrasts served as activation localizers and compared the main predictor curve (difference of gammas only, not its temporal derivative) for *pro_instruction*, *anti_instruction*, *nogo_instruction*, *pro_response*, *anti_response*, and *nogo_response* to a null value of zero. The other six of contrasts compared the main predictor curve (again, difference of gammas only, not its temporal derivative) from two of the three trial types during either the instruction period or the response period. The comparison contrasts were *anti_instruction − pro_instruction*, *anti_instruction − nogo_instruction*, and *pro_instruction − nogo_instruction*, as well as *anti_response − pro_response*, *anti_response − nogo_response*, and *pro_response − nogo_response*.

Single subject maps were combined into a mixed effects analysis based on the methods described in Worsley and colleagues (2002). Expectation maximization (EM) was used to estimate between subjects variance for each of the contrasts, and T-statistic maps were built by scaling expected contrast size across subjects by the square root of the combined within subjects variance and between subjects variance. T-maps were regularized using the method described in Worsley and colleagues (2002). Briefly, this technique allows one to use within subject statistical estimates to enhance the degrees of freedom of between subjects estimates. In fMRI experiments, each subject's data typically includes many time points per voxel resulting in high within subject degrees of freedom and low within subject variance for any statistical estimate of interest. On the other hand, fMRI experiments usually include somewhere from 6 to 20 subjects, which imparts low between subject degrees of freedom and correspondingly high variance for between subject statistical estimates. The variance regularization method described in Worsley and colleagues (2002) high degrees of freedom of the within subject estimates of variance of a T-statistic as a template to enhance the degrees of freedom, and reduce the variance, of between subject T-statistics. In our application of this regularizing method, we used a Gaussian smoothing kernel with 3.2719 mm FWHM, resulting in a mixed-effects degrees of freedom for the T-statistic of 15000 (see Worsley et al. 2002 for details). All statistical maps were thresholded twice: first at a single voxel T-threshold of 2.575, corresponding to an uncorrected p-value of 0.01 (df = 15000), and second at a cluster size threshold of 1096 cubic mm, yielding a p-value of 0.05 corrected for multiple comparisons across the voxel population. The cluster size threshold value of 1096 cubic mm was computed using

fmristat with the FWHM parameter for the T-maps taken as 6 mm. Fmristat is a Matlab program, written by Worsley and colleagues, that uses random field theory to compute cluster size thresholds (Worsley et al. 2002).

Appendix 4.4 – Bibliography

Bullmore E, Long C, Suckling J, Fadili J, Calvert G, Zelaya F, Carpenter TA, and Brammer M. 2001. Colored noise and computational inference in neurophysiological (fMRI) time series analysis: resampling methods in time and wavelet domains. *Hum Brain Mapp* 12(2): 61-78.

Dale AM. 1999. Optimal experimental design for event-related fMRI. *Hum Brain Mapp* 8(2-3): 109-114.

Friston KJ, Fletcher P, Josephs O, Holmes A, Rugg MD, and Turner R. 1998. Event-related fMRI: characterizing differential responses. *NeuroImage* 7(1): 30-40.

Marchini JL and Ripley BD. 2000. A new statistical approach to detecting significant activation in functional MRI. *NeuroImage* 12(4): 366-380.

Worsley KJ, Liao CH, Aston J, Petre V, Duncan GH, Morales F, and Evans AC. 2002. A general statistical analysis for fMRI data. *NeuroImage* 15(1): 1-15.

Matthew R. G. Brown – Curriculum Vitae

Education

Post-Graduate

Ph.D. Neuroscience – University of Western Ontario, London, Ontario
                    – September 2003 - November 2007 (expected completion date)

M.Sc. Neuroscience – University of Western Ontario, London, Ontario
                    – August 2001 - August 2003

Undergraduate

B.Sc. Honours Neuroscience – University of Alberta, Edmonton, Alberta
                    – September 1997 - May 2001

Scholarly Publications

Journal Publications

Brown MRG, Vilis T, Everling S. (2007) Isolation of saccade inhibition processes: Rapid event-related fMRI of saccades and nogo trials. *NeuroImage* (in press).

Brown MR, Vilis T, Everling S. (2007) Frontoparietal activation with preparation for antisaccades. *Journal of Neurophysiology* 2007; 98(3): 1751-1762.

Brown MRG, Goltz HC, Vilis T, Ford KA, Everling S. (2006) Inhibition and generation of saccades: Rapid event-related fMRI of prosaccades, antisaccades, and nogo trials. *NeuroImage* 2006; 33(2): 644-659.

Ford KA, Goltz HC, Brown MRG, Everling S. (2005) Neural processes associated with antisaccade task performance investigated with event-related fMRI. *Journal of Neurophysiology* 2005; 94(1): 429-440.

Brown MRG, Desouza JFX, Goltz HC, Ford K, Menon RS, Goodale MA, Everling S. (2004) Comparison of memory- and visually guided saccades using event-related fMRI. *Journal of Neurophysiology* 2004; 91(2): 873-889.

Wylie DRW, Brown MR, Winship IR, Crowder NA, Todd KG. (2003) Zonal organization of the vestibulocerebellum in pigeons *(Columba livia)*: III. projections of the translation zones of the ventral uvula and nodulus. *Journal of Comparative Neurology* 2003 Oct 13; 465(2): 179-94.

Wylie DRW, Brown MR, Barclay RR, Winship IR, Crowder NA, Todd KG. (2003) Zonal organization of the vestibulocerebellum in pigeons (*Columba livia*): II. projections of the rotation zones of the flocculus. *Journal of Comparative Neurology* 2003 Feb 3; 456(2): 140-53.

Refereed Conference Presentations

Brown MR, Vilis T, Everling S. (2007) Saccade inhibition: Rapid event-related fMRI of saccades and nogo trials. *Society for Neuroscience Annual Meeting*, San Diego, CA.

Brown MR, Vilis T, Everling S. (2006) Prosaccade and antisaccade instruction differences investigated with rapid event-related fMRI. 48.21 *Society for Neuroscience Annual Meeting*, Atlanta, GA.

Brown MRG, Goltz HC, Vilis T, Ford KA, Everling S. (2005) Rapid event-related fMRI of prosaccades, antisaccades, and nogo trials. 167.10 *Society for Neuroscience Annual Meeting*, Washington, D.C.

Ford KA, Levin HM, Brown MR, Everling S. (2005) Neural circuitry underlying anti-saccade task performance in humans and monkeys investigated with fMRI. 166.10 *Society for Neuroscience Annual Meeting*, Washington, D.C.

Brown MRG, Goltz HC, Vilis T, Ford KA, Everling S. (2005) Fast event-related fMRI of prosaccades, antisaccades, and nogo trials. TK-128 *13th European Conference on Eye Movements*, Bern, Switzerland.

Brown MRG, Desouza JFX, Ford K, Goltz HC, Goodale MA, Everling S. (2004) Cortical connectivity for memory- versus visually guided saccades. 991.8 *Society for Neuroscience Annual Meeting*, San Diego, California.

Ford KA, Goltz HC, Brown MR, Everling S. (2004) A distributed frontal cortical network for saccade suppression. 313.14 *Society for Neuroscience Annual Meeting*, San Diego, California.

Cavina Pratesi C, Valyear KF, Obhi SS, Brown MR, Marzi C, Goodale MA. (2004) Neural correlates of preparatory set: response selection versus movement planning. 202.12 *Society for Neuroscience Annual Meeting*, San Diego, California.

Brown MRG, DeSouza JFX, Ford K, Goltz HC, Menon R, Goodale MA, Everling S. (2003) Neural correlates for memory- and visually guided saccades investigated with event-related fMRI. *12th European Conference on Eye Movements*, Dundee, Scotland, U.K.

Brown MRG, Desouza JFX, Ford K, Goltz HC, Goodale MA, Everling S. (2002) Comparison of memory- and visually guided saccades using event-related fMRI. *Society for Neuroscience Annual Meeting*, Orlando, Florida.

Non-refereed Conference Presentations

Brown MRG. (2004) Fast event-related fMRI of saccadic eye movements. *CIHR Group for Action and Perception Annual Retreat*, London, On.

Brown MRG. (2001) fMRI study of delayed saccades: Circuitry for visual vs. memory-guided saccades. *CIHR Group for Action and Perception Annual Retreat*, London, On.

Brown MRG, Winship I, Wong-Wylie DR. (2001) Zonal innervation of vestibular and cerebellar nuclei by translation-sensitive Purkinje cells in the vestibulocerebellum of the pigeon (*Columba livia*). *Tenth Canadian Spring Conference on Behaviour and Brain*, Fernie, B.C.

Scholarships and Awards

Awards	Location*	Period Held
The G. Keith Humphrey Memorial Award	UWO	Awarded 06/2006
NSERC Canada Graduate Scholarship D	UWO	09/2003 - 08/2005
Graduate Tuition Scholarship	UWO	09/2003 - 08/2005
Neuroscience M.Sc. Thesis Award	UWO	Awarded 01/2004
NSERC Postgraduate Scholarship A	UWO	09/2001 - 08/2003
Graduate Tuition Scholarship	UWO	09/2001 - 08/2003
NSERC Summer Studentship	UAlberta	05/2001 - 07/2001
AHFMR Summer Studentship	UAlberta	05/2001 - 07/2001
Dean's Silver Medal in Science	UAlberta	Awarded 05/2001
William Irvine Memorial Prize in Ethics	UAlberta	09/2000 - 04/2001
Louise McKinney Scholarship	UAlberta	09/2000 - 04/2001
4th Year Science Scholarship	UAlberta	09/2000 - 04/2001
NSERC Summer Studentship	UAlberta	05/2000 - 08/2000
AHFMR Summer Studentship	UAlberta	05/2000 - 08/2000
Louise McKinney Scholarship	UAlberta	09/1999 - 04/2000
3rd Year Science Scholarship	UAlberta	09/1999 - 04/2000
NSERC Summer Studentship	UAlberta	05/1999 - 08/1999
AHFMR Summer Studentship	UAlberta	05/1999 - 08/1999
2nd Year Science Scholarship	UAlberta	09/1998 - 04/1999
Academic Excellence Scholarship	UAlberta	09/1997 - 04/1998
Science Entrance Scholarship	UAlberta	09/1997 - 04/1998
Alexander Rutherford Scholarship	UAlberta	09/1997 - 04/1998

Awards (Declined)	Location*	Year Awarded
Ontario Graduate Scholarship	declined	2003
Ontario Graduate Scholarship	declined	2001

Teaching Assistantships	Location*	Period Held
Teaching Assistantship	UWO	09/2005 - 04/2006
Teaching Assistantship	UWO	09/2004 - 12/2004
Teaching Assistantship	UWO	09/2003 - 12/2003
Teaching Assistantship	UWO	09/2002 - 12/2002
Teaching Assistantship	UWO	09/2001 - 04/2002

* UWO = University of Western Ontario, London, Ontario, Canada

UAlberta = University of Alberta, Edmonton, Alberta, Canada