

CIRCUITS UNDERLYING VISUAL ATTENTION IN PRIMATE NEOCORTEX

A DISSERTATION

SUBMITTED TO THE PROGRAM IN NEUROSCIENCE

AND THE COMMITTEE ON GRADUATE STUDIES

OF STANFORD UNIVERSITY

IN PARTIAL FULFILLMENT OF THE REQUIREMENTS

FOR THE DEGREE OF

DOCTOR OF PHILOSOPHY

Nicholas A. Steinmetz

March, 2014

© 2014 by Nicholas Steinmetz. All Rights Reserved.

Re-distributed by Stanford University under license with the author.



This work is licensed under a Creative Commons Attribution-Noncommercial 3.0 United States License.

<http://creativecommons.org/licenses/by-nc/3.0/us/>

This dissertation is online at: <http://purl.stanford.edu/wb273ht0017>

I certify that I have read this dissertation and that, in my opinion, it is fully adequate in scope and quality as a dissertation for the degree of Doctor of Philosophy.

Tirin Moore, Primary Adviser

I certify that I have read this dissertation and that, in my opinion, it is fully adequate in scope and quality as a dissertation for the degree of Doctor of Philosophy.

Kwabena Boahen, Co-Adviser

I certify that I have read this dissertation and that, in my opinion, it is fully adequate in scope and quality as a dissertation for the degree of Doctor of Philosophy.

Shaul Hestrin

I certify that I have read this dissertation and that, in my opinion, it is fully adequate in scope and quality as a dissertation for the degree of Doctor of Philosophy.

Eric Knudsen

I certify that I have read this dissertation and that, in my opinion, it is fully adequate in scope and quality as a dissertation for the degree of Doctor of Philosophy.

William Newsome

Approved for the Stanford University Committee on Graduate Studies.

Patricia J. Gumport, Vice Provost for Graduate Education

This signature page was generated electronically upon submission of this dissertation in electronic format. An original signed hard copy of the signature page is on file in University Archives.

Abstract

Humans and many other species attend to only a small portion of available visual information at any given moment. They enhance perception of the attended stimulus either overtly, making an eye or head movement to orient toward it, or covertly, without any such movements. The neural circuits that underlie these two types of attention behaviors, and the relationship between them, remain unclear. To investigate the interdependence of them we trained monkeys on a task that behaviorally dissociated the location of covert attention from the location of a saccade target. Recordings in extrastriate visual cortical area V4 surprisingly revealed that enhanced firing rates and other modulations of neural activity accompanied both covert attention and saccade preparation. These results suggested a hypothesis about the circuits that could mediate the control of both behaviors. We recorded neurons in the frontal eye field, an area involved in controlling both behaviors, and found evidence contradicting our hypothesis. Separately, we examined the circuit underlying the integration of attention-related feedback signals with visual information in visual cortex by recording from distinct neuron populations, defined by laminar depth, within V4 during the covert attention task. We found that all neuron populations were modulated indistinguishably during attention. Finally, we constructed a large-scale model of FEF and V4 on neuromorphic hardware and used it to investigate a novel hypothesis about the way feedback from FEF influences V4, namely, via NMDA synapses. This model makes predictions for future experiments that could help uncover the mechanism of attention-related modulation of visual cortex. Taken together, these results have helped to elucidate our understanding of the circuits within and between frontal and visual cortical areas underlying attention.

Acknowledgements

I am indebted to a great number of people, both personal friends and family as well as scientific colleagues, many of whom I am happy to now count among my closest friends as well.

Thanks first to Tirin Moore and to Kwabena Boahen, my scientific advisors and mentors who have given generously of their time and expertise in matters scientific, political, practical, and otherwise. I feel tremendously grateful to have had the opportunity to learn from such wonderful scientists and human beings.

Thanks to the Moore and Boahen labs, who provided constant intellectual engagement as well as moral support: Marc Zirnsak, Tatiana Engel, Ryan Squire, Behrad Noudoost, Ben Benjamin, Kelsey Clark, Alireza Soltani, Becca Krock, Peiran Gao, Mindy Chang, Brittany Burrows, Bob Schafer and others have made important contributions to my scientific thinking as well as to my overall happiness. Thanks to Doug Aldrich and Bob Schneeveis for important technical contributions. Thanks also to other scientific collaborators and interlocutors, most notably Sridhar Deverajan, as well as other friends in neuroscience, including James Fitzgerald, Chris Rogers, Forrest Collman, David Kastner, Georgia Panagiotakos, Egle Cekanaviciute, Corbio Bennett, Sergi Arroyo, Jaimie Adelson, and others I don't have space to name.

Thanks to my thesis committee – Eric Knudsen, Shaul Hestrin, and Krishna Shenoy – for their valuable advice and guidance throughout the years. I have also benefitted from the generosity of a number of other Stanford faculty who have taken time to talk with me on various occasions, including Bill Newsome, Jay McClelland, Surya Ganguli, John Huguenard, Dick Tsien, Steve Baccus, and Jennifer Raymond. I would like to thank Jay in particular for his role in running the Mind, Brain, and Computation Center at Stanford, which has been valuable both in providing a stimulating community of like-minded scientists as well as funding two years of my research.

I am happy to have the opportunity to thank John Huguenard again, as well as Ross Colvin and Katie Johnson, for running the Neuroscience Program, which has been a truly incredible environment in which to grow as a scientist. I'm sure I couldn't have picked a better place to learn and live, a fact due in large part to their efforts in building the community we all now enjoy.

I won't try to list all of my friends outside of science who have provided the necessary reprieves from the struggle of lab life, but I am happy to thank two, Steve Connor and Nick Weiler, who I have been fortunate to count as my best friends for the entire duration of my time at Stanford. As a source of humor, companionship, and adventure, I could have asked for none better.

Finally, I would like to thank my family for their unwavering support as I undertook what surely has seemed to them a bizarre journey across the country to do who-knows-what. It was through the love and guidance of my parents, Michael and Kimberly Steinmetz, that I became the curious and critical person I am today, and it was only because of their commitment to allowing me to fulfill my dreams, no matter what they were, that I am here today writing a doctoral dissertation in neuroscience. Thank you, Mom and Dad.

Author contributions

All chapters were written by N. Steinmetz, and describe research designed and executed solely by him, with the following exceptions:

Chapters 2, 3, 4, and 5 describe research designed in collaboration with T. Moore.

Chapters 2 and 3 were written in collaboration with T. Moore.

Chapter 3 includes some analyses performed and some figures made by S. Deverajan (where noted in the text).

Chapter 6 describes research designed in collaboration with K. Boahen and B. Benjamin. Some data was collected and some figures were produced by B. Benjamin (where noted in the text).

Table of Contents

1	Introduction	1
1.1	Overview	1
1.2	Overt and Covert Visual Attention	3
1.3	Area V4	7
1.3.1	Role of area V4 in visual perception	8
1.3.2	Role of V4 in attention and contextually-sensitive visual processing	9
1.3.3	Possible sources and mechanisms of V4 modulation	13
1.3.4	Microcircuitry within area V4	15
1.4	The Frontal Eye Field (FEF)	18
1.4.1	The role of the FEF in saccadic behavior	19
1.4.2	The role of the FEF in covert attention	22
1.4.3	Circuitry within FEF mediating overt and covert attentional functions	24
1.5	Models of Visual Attention	29
2	Changes in the Response Rate and Response Variability of Area V4 Neurons During the Preparation of Saccadic Eye Movements	33
2.1	Abstract	33
2.2	Introduction	34
2.3	Methods	36
2.3.1	Subjects	36
2.3.2	Behavioral task	36
2.3.3	Recording	37
2.3.4	RF stimuli	37
2.3.5	Data analyses	38
2.4	Results	41
2.4.1	Mean-matched control for presaccadic firing rate changes	43
2.4.2	Dependence of FF changes on saccade direction	44
2.4.3	Predicting Saccadic RT	45
2.4.4	Possible influence of microsaccades	48
2.5	Discussion	49
3	Interdependence of saccadic and attentional modulation of visual cortical signals	57
3.1	Abstract	57
3.2	Introduction	58
3.3	Methods	60
3.3.1	Subjects	60
3.3.2	Behavioral task and visual stimuli	60
3.3.3	Behavioral Analysis	62
3.3.4	Neural recordings	63
3.3.5	Quantification of firing rate modulation	66
3.3.6	Quantification of tuning amplitude	66
3.3.7	Quantification of across-trial spiking reliability (Fano factor, FF)	68
3.3.8	Quantification of power in local field potentials (LFP) within a frequency band	68
3.4	Results	69
3.4.1	Behavioral Performance	70
3.4.2	Firing rate modulation	73
3.4.3	Orientation tuning	76
3.4.4	Response reliability	78

3.4.5	Relationship between attention-related and saccadic modulations for individual neurons	79
3.4.6	Local field potential power	80
3.5	Discussion	82
3.5.1	Size of attention effects.....	82
3.5.2	Presaccadic modulation in visual cortex.....	83
3.5.3	Implications for the circuits controlling attention.....	84
3.5.4	What is the source of these modulations, if they share a single source?	84
3.5.5	An alternate interpretation: two separate and independent sources with nearly identical effects.....	85
3.5.6	Behavioral performance.....	86
3.5.7	Noise correlation analyses	86
3.5.8	Dissociation of saccade location from cued location in this experiment.....	86
4	Attentional and saccadic modulation across neocortical layers in area V4	89
4.1	Abstract	89
4.2	Introduction	90
4.3	Methods.....	92
4.3.1	Subjects.....	92
4.3.2	Behavioral task and visual stimuli	92
4.3.3	Neural recordings.....	93
4.3.4	Data Analysis	96
4.4	Results	101
4.4.1	RF alignment.....	101
4.4.2	Depth registration	102
4.4.3	Firing rate.....	103
4.4.4	Stimulus onset latency	104
4.4.5	Orientation tuning.....	105
4.4.6	Firing rate modulation during attention and saccade preparation.....	106
4.5	Discussion	108
4.5.1	Overlapping RFs and recording from a single “column”	108
4.5.2	Correspondence of CSD features with anatomical layers	108
4.5.3	Differences in response properties across depth.....	109
4.5.4	Lack of differences in feedback-driven modulations	111
5	Frontal eye field circuitry underlying attentional and saccadic modulation	113
5.1	Abstract	113
5.2	Introduction	114
5.3	Methods.....	116
5.3.1	Subjects.....	116
5.3.2	Behavioral task and visual stimuli	116
5.3.3	Neural recordings.....	116
5.3.4	Data analysis	117
5.4	Results	118
5.5	Discussion	122
6	Modeling the circuitry of attention in frontal and visual cortices with Neurogrid.....	123
6.1	Abstract	123
6.2	Introduction	124
6.3	Methods.....	126
6.3.1	Neurogrid simulation platform	126
6.3.2	Neurogrid model construction and execution.....	126
6.3.3	Neurogrid model design.....	126

6.3.4	Simplified model	128
6.4	Results	130
6.4.1	Neurogrid model of FEF and V4 accounts for modulation of V4 responses during selective attention.....	130
6.4.2	Simplified model of modulation by NMDA in a single neuron	131
6.4.3	Membrane potential bistability with NMDA currents.....	133
6.5	Discussion	136
6.5.1	Model FEF activity and working memory representations	136
6.5.2	Model feedback with NMDA	137
6.5.3	Normalization models of attention	138
6.5.4	Relationship to other studies of NMDA effects	139
6.6	Appendix: Neurogrid model specification code	141
7	Conclusions	143
8	References	147

Table of Figures

Figure 1-1. The location of area V4 within the macaque brain and the visual network..	7
Figure 1-2. Projections to area V4 from areas that may drive attention-related modulation.....	14
Figure 1-3. The connectivity and anatomical locations of the Frontal Eye Field (FEF) and other structures within the visual and oculomotor systems of the rhesus macaque monkey brain	19
Figure 1-4. Hypothetical circuitry mediating the control of covert attention and saccades.....	26
Figure 2-1. The visually guided delayed saccade task	37
Figure 2-2. Effects of RF stimulation and saccade preparation on the mean firing rate and response variability for the population of area V4 neurons	42
Figure 2-3. Presaccadic changes in FF for “mean-matched” conditions.....	44
Figure 2-4. Relationship of presaccadic mean firing rate and FF to saccadic RT for the population of V4 neurons.....	47
Figure 2-5. Cartoon model of changes in firing rate distributions during saccade preparation.....	53
Figure 3-1. Graphical depiction of the m-ADC model for the two-alternative case.....	63
Figure 3-2. Cued change detection and antisaccade task. A. Task design and trial sequence.....	69
Figure 3-3. Reaction time dependence on cue validity. Mean \pm S.E.M. for correctly performed trials when the cue was valid, invalid indicating the opposite location, and invalid indicating either orthogonal location.....	71
Figure 3-4. Behavioral analysis for monkeys performing the cued change detection task, performed with signal detection theory model	73
Figure 3-5. Responses of three example single neuron in the cued change-detection task	74

Figure 3-6. Effects of cue direction on post-cue period firing rate	75
Figure 3-7. Effects of cue direction on tuning amplitude	76
Figure 3-8. Effects of cue direction on across-trial spiking reliability	78
Figure 3-9. Cue-opposite firing rate modulation for those neurons significantly modulated during cue-RF trials	80
Figure 3-10. Difference in LFP power between cue conditions across frequencies	81
Figure 4-1. Recording setup, electrode array, and sample data.....	93
Figure 4-2. “Well angler”: a device to tilt and position an electrode at an arbitrary vector.....	95
Figure 4-3. Waveform duration distribution	99
Figure 4-4. Example RF alignment using the MRI targeting procedure.....	101
Figure 4-5. Example current source density (CSD) with alignment feature	103
Figure 4-6. Firing rates as a function of depth in cortex	104
Figure 4-7. Stimulus onset latency as a function of depth	105
Figure 4-8. Orientation tuning across depth.....	106
Figure 4-9. Firing rate modulation by cue direction across depth.....	107
Figure 4-10. Two possible alignments between the anatomical layers in area V4 and the CSD as measured in response to the full-field flash stimulus	109
Figure 5-1. The activity of an example FEF neuron	118
Figure 5-2. A second example neuron.....	119
Figure 5-3. A third example neuron	120
Figure 5-4. Visual versus saccade responsiveness for 189 isolated FEF neurons.....	121
Figure 6-1. Multiple cortical neuron types simulated on Neurogrid	127
Figure 6-2. Modeling attention-related modulation of visual cortical responses.....	130

Figure 6-3. Simplified model of NMDA modulation of visually driven responses in a single neuron 132

Figure 6-4. Basic conditions necessary for bistability..... 134

Figure 6-5. Analysis of the region of bistability in two different descriptions of NMDA voltage dependence 135

1 Introduction

1.1 Overview

Our understanding of the brain has progressed rapidly in the just the last few decades simultaneously at both large and small scales. We have come to understand in tremendous detail the operation of individual neurons, the ion channels and synapses that govern their responses, and the ways in which they can combine inputs to generate output. At another scale, we now have access to a flood of information about the connectivity between every part of the brain and the diverse behaviors in which they are involved. We are therefore now entering an era in which a primary objective of neuroscience is to attempt to link these levels: to understand the circuits connecting neurons both within and between brain regions and to understand how these circuits give rise to the dynamics and response properties of the neural activity we observe. Ultimately we hope be able to describe the highest-level features of behavior, such as visual attention, in terms of the brain regions, circuits, neuronal properties, and biophysical features that ultimately underlie them.

In this thesis, I investigate the circuits underlying visual attention in primate neocortex. I explore the relationship between the volitional control of eye movements and the perceptual effects of attention, as well as the neural mechanisms underlying these processes. I have pursued this experimentally by recording neural activity from the Frontal Eye Field (FEF), a region of frontal cortex implicated in the control of both eye movements and attention, and from a visual cortical area (V4) in which visual responses are modulated during overt and covert attention.

In Chapter 2 I describe a novel aspect of V4 response modulation during eye movement preparation – a reduction in across-trial variability – which suggests that visual cortex may be more intimately involved in the circuitry for overt attention than previously believed.

In Chapter 3 I report the unexpected finding that saccade preparation alone, when behaviorally dissociated from covert attention, drives robust modulation of visual cortical responses. This result highlights just how interdependent are the circuits controlling saccades and covert attention, even as they extend to posterior visual cortex.

In Chapter 4 I investigate the circuitry underlying the integration of the modulatory feedback signals related to attention with visual signals within area V4 by examining how different populations of neurons within V4, populations defined by depth in cortex, are modulated during attention and saccade preparation.

In Chapter 5 I study the circuitry underlying the control of overt and covert attention with area FEF by characterizing the responses of different functionally defined classes of neurons during a task that dissociates the two behaviors from each other.

Finally, in Chapter 6 I develop a model of the control and implementation of attention in these two brain regions, FEF and V4, and implement the model on neuromorphic hardware. The model connects observations about the behavior of these areas during attention with specific biophysical mechanisms that may underlie these behaviors, and makes specific predictions for experiments going forward.

Before detailing my experimental and modeling efforts, I will introduce these topics by reviewing what is known about the relationship between gaze control and visual attention behaviors and about their shared underlying neural substrates. In particular, I will focus on the roles of the areas I have studied, FEF and V4, and the anatomical and functional circuitry within and between these areas. I will also discuss existing ideas and models of the circuits linking the neurons within and between these brain regions to account for the observed behavior.

1.2 Overt and Covert Visual Attention

Humans and other primates rely on vision more than any other sense. We primarily perceive the world visually, using our extraordinarily high visual acuity to identify objects from a distance, to comprehend and navigate the three-dimensional structure of our environment, to precisely guide our movements, to infer social cues, and to instill abstract ideas in our minds, as you are experiencing now. Of the torrent of visual information that rushes into our brain each waking moment, however, only a small fraction may be processed in depth. On the basis of our expectations or goals, we choose only a small portion of the visual scene at a time to perceive, remember, and fully understand while discarding and largely failing to perceive the rest. That is, we choose a portion to attend and ignore the rest.

The primate retina contains a region of especially high-density sensors called the fovea. Thus, to acquire the highest quality visual information about the attended part of the scene, we typically reposition our eyes such that light from a peripheral feature of interest now arrives at the fovea. These eye movements are called saccades, and we make them quite frequently, about two or three times per second (Yarbus, 1967). They are ballistic in nature, jumping rapidly from point to point as we examine objects of attention. Performing these saccades accurately requires using visual features of the saccade target, such as its position, form, or motion, for accurate guidance (Moore et al., 1998). Given this need for information about the target to guide saccades, it is perhaps not surprising that psychophysical studies have demonstrated perceptual enhancement at saccade targets in the moments leading up to saccade execution (Hoffman and Subramaniam, 1995). Thus, preparing to saccade to a target draws perceptual resources to that location, that is, it directs attention there. This process of making saccades to attended features, because it involves easily observable movements of the eyes, is referred to as “overt” attention.

By contrast, one may attend an object “covertly,” without making any observable eye movements. In this case, though gaze remains fixed on an object other than the one being attended, the peripheral stimulus at the focus of attention nevertheless becomes preferentially processed, resulting in improved discrimination of stimulus

features and speeded responses to relevant stimulus changes (Carrasco, 2011; Posner, 1980). This ability is known colloquially as “looking at something out of the corner of your eye.” As direct gaze can often be an important social cue, for instance as a sign of aggression or of trust, primates often choose to covertly attend to other members of their species rather than directly fixate on them (Mendelson et al., 1982).

A particularly well-studied form of covert attention, and the one explored in this thesis, is top-down spatial attention, in which attention is volitionally focused on a particular region of visual space. While attention can also be directed to particular stimulus properties no matter where they fall in the visual field (“feature-based attention”; e.g. Bichot et al., 2005) and can be driven by salient stimulus properties rather than by endogenous factors (“bottom-up attention”; e.g. Burrows and Moore, 2009), I will focus here on the top-down, spatially selective type of attention.

Covert and overt attention are clearly unified in that they both involve the selection and detailed inspection of one visual element at the expense of others. But might they be unified in their neural mechanisms as well? As noted above, preparing to make a saccade to a peripheral stimulus brings about many of the same perceptual effects at that target stimulus in the moments just prior to saccade execution as does choosing to covertly attend to it. For instance, discrimination performance is enhanced (Deubel and Schneider, 1996; Hoffman and Subramaniam, 1995; Kowler, 2011). Might our ability to covertly attend peripheral stimuli reflect an adaptation of this process? Might our faculty of covert attention in fact be just the process of preparing to execute a saccade but withholding the actual movement?

The hypothesis that overt and covert spatial attention share the same underlying neural mechanisms is known as the “premotor theory of attention” (Awh et al., 2006; Craighero and Rizzolatti, 2005; Steinmetz and Moore, 2012). This hypothesis states that the neural elements involved in planning and executing saccades to a part of visual space are the same as the ones involved in directing attention covertly to that part of space. By preparing to execute an eye movement but by withholding the actual execution of the movement, we nevertheless invoke the same neural processes that

select a peripheral stimulus and confer perceptual benefits as seen prior to saccades to those stimuli.

A great deal of physiological evidence has supported the idea that the spatial control of saccades and of covert attention are mediated by shared neural mechanisms. Brain regions such as the frontal eye field (FEF), the lateral intraparietal area (LIP), and superior colliculus (SC) seem to be intimately involved in saccade preparation and execution as well as in controlling the locus of covert attention. Several lines of evidence support this view. Recordings from individual neurons as well as neuroimaging studies have shown that neural activity in these regions is modulated by the execution of saccades (FEF: Bruce and Goldberg, 1985; LIP: Barash et al., 1991; SC: Schiller and Stryker, 1972) and by the allocation of covert attention (FEF: Gregoriou et al., 2012; Thompson et al., 2005; LIP: Bisley and Goldberg, 2003; Bushnell et al., 1981; SC: Goldberg and Wurtz, 1972; Ignashchenkova et al., 2004;). Recordings of activity in the brain made with fMRI point to a similar set of regions involved (Corbetta et al., 1998). Permanent or reversible inactivation of these areas degrades saccadic behavior (FEF: Sommer and Tehovnik, 1997; LIP: Li et al., 1999; SC: Schiller et al., 1980;) and attentional performance (FEF: Wardak et al., 2006; LIP: Wardak et al., 2004; SC: Zénon and Krauzlis, 2012). Electrical stimulation of these areas elicits saccades (FEF: Robinson and Fuchs, 1969; LIP: Shibutani et al., 1984; SC: Robinson, 1972), and can also influence performance on attention tasks (FEF: Moore and Fallah, 2004; LIP: Cutrell and Marrocco, 2002; SC: Muller et al., 2005). A detailed comparison of the relative roles of these structures in saccadic and covert attentional behaviors is beyond the scope of this document, though the particular role of the FEF is described in much greater detail in section 1.4. Furthermore, while the greatest amount of evidence supports the roles of these three areas in both functions, other areas may also play a role: other subcortical structures such as the pulvinar nucleus of the thalamus (Petersen et al., 1987) and the cholinergic nucleus basalis (Herrero et al., 2008); other frontal areas such as the dlPFC (Messinger et al., 2009); or other eye fields such as the supplementary eye field (SEF) and area 7m (Lynch and Tian, 2006). Taken together, ample evidence exists supporting the joint control of overt and covert attention by a unified network of cortical and subcortical regions.

Nevertheless, at some level these processes must be dissociable. After all, one involves the overt repositioning of the eyes while the other does not. A key question regarding these foundational visual behaviors, then, is at what stage their neural mechanisms are unified and at what stage they are distinct. Do they both necessarily invoke modulation of perception and of visual cortical representations? Can the particular neurons and circuits where they diverge or interact be identified? These questions form a central topic of this thesis.

I will provide evidence that both overt and covert attention do necessarily modulate visual cortical representations and perception (Chapter 3), indicating that their neural circuits are shared at the level of the neurons at the source of those modulations. I will investigate a novel approach to attempt to determine what neurons constitute that source (Chapter 4). Finally, I will explore the circuits that may be at the nexus of both behaviors, within area FEF, to try to understand how the two processes interact (Chapter 5).

1.3 Area V4

Area V4 is the region of extrastriate visual cortex that we studied in this thesis. Anatomically, V4 has its lower visual field representation on the prelunate gyrus (Figure 1-1). It is primarily a ventral stream area and sits at an intermediate level of the visual hierarchy (at level 6 ± 1 out of 12-20 total levels; Felleman and Van Essen, 1991; Hilgetag et al., 1996). Its prominent and accessible anatomical location, large size, easily identifiable receptive field locations, and wide array of interesting response properties have probably all contributed to making it one of the two or three best-studied extrastriate visual areas out of the 30 or so that exist in the macaque. Here I will briefly review the properties of area V4 neurons, particularly as relevant to the present work.




Figure 1-1. The location of area V4 within the macaque brain and the visual network. A, Anatomical location of area V4 on the prelunate gyrus and approximate locations of some other cortical visual areas mentioned in this document. B, Two-dimensional representation of visual cortex organization, with V4 at

bottom center. Areas with similar anatomical connectivity are grouped together and connected with solid lines. Visual information flows roughly left to right through both the dorsal (top branch) and ventral (bottom) streams. V4 appears to constitute a critical link in the ventral stream. Dotted lines connect nodes to their labels. The plotted position of each node (cortical area) was determined with multi-dimensional scaling applied to the visual cortical connectivity matrix. Connectivity data from (Sporns and Zwi, 2004).

1.3.1 Role of area V4 in visual perception

V4 neurons, like neurons in many other visual areas, respond only to visual stimuli in certain, limited regions of space, called receptive fields (RFs), and respond differently depending on the properties of the stimuli presented there. V4 neurons' RFs have diameters slightly smaller than the eccentricity of the center (Gattass et al., 1988). In terms of visual stimulus form, V4 neurons can be identified that are tuned for the orientation of simple bars and Cartesian gratings, spatial frequency (Desimone and Schein, 1987), 3D orientation (Hinkle and Connor, 2002), binocular disparity (i.e. depth, Hinkle and Connor, 2005), non-Cartesian gratings such as spirals and sinusoidal crosses (Gallant et al., 1996), and the curvature of the edges of small objects (Pasupathy and Connor, 2001). A more recent study has described these complex shape tuning properties as reflecting tuning within the space defined by the Fourier power spectra of the images, i.e. the space defined by the orientation and spatial frequency of image components (David et al., 2006). Another recent study suggested that the boundary curvature tuning can be understood as the composition of tuning for many small localized edge fragments (Nandy et al., 2013). In both of these recent studies, then, shape tuning in V4 can be thought of as arising from a weighted sum over inputs from neurons in earlier visual areas, like area V1 neurons, which represent images by their edge components, or what amounts to a Fourier decomposition.

This basic idea – forming tuning for a new stimulus property by the weighted sum of inputs selective for simpler stimulus properties – has been a foundational idea in visual neuroscience (Hubel and Wiesel, 1962; Riesenhuber and Poggio, 1999) and so seems nicely confirmed by these descriptions of V4 activity. Indeed, specific computational models have been built to account for V4 response properties in this way (Cadiou et al., 2007). However, experimental evidence that these properties in

fact arise at the level of V4 neurons has been inconclusive. For instance, area V2 neurons show many identical or similar tuning properties to these complex stimuli as V4 neurons (Hegd  and Van Essen, 2007). Nevertheless, V4 inactivation does impair shape discrimination (Girard et al., 2002; Merigan and Pham, 1998), so it seems likely that V4 plays some crucial role in the process of object identification.

V4 neurons are also selective for the color of RF stimuli independent of stimulus form. Indeed, it was this property of V4 that was first noted and V4 was initially known as the primary color processing region of extrastriate visual cortex (Zeki, 1983), a characterization bolstered by lesion experiments showing deficits in color discrimination (Walsh et al., 1993). Since then a number of studies have confirmed and refined the observation of color-selective neurons in area V4 (Chang et al., 2013; Heywood and Cowey, 1987; Schein and Desimone, 1990). Recently, the discovery of separate color selective and orientation selective domains within V4 has suggested that these different views of V4 function – color versus form processing – may have arisen from different distributions of recording sites within V4 in the different studies (Kotake et al., 2009; Tanigawa et al., 2010).

In deciding which stimuli to use to drive the V4 neurons from which we recorded, there were many suitable options available. A suitable option would drive a large proportion of V4 neurons and have an easily tunable stimulus dimension along which the stimuli could be continuously varied. We selected 2D, oriented, black-and-white Cartesian grating stimuli largely for simplicity and because the largest number of past studies had described attentional modulation of responses to these stimuli (reviewed below).

1.3.2 Role of V4 in attention and contextually-sensitive visual processing

Contextual cues influence the responses of V4 neurons, both in terms of the relationship between RF stimuli and surrounding stimulus information as well as the learned association of cue stimuli with the identity of rewarding information (i.e. learned orienting of attention). I will first discuss the evidence for the role of V4

during attention, central to the present thesis, and then briefly discuss several other types of contextual modulation in V4 – surround suppression, pop-out, and figure-ground segregation – which have interesting connections with the attentional function.

1.3.2.1 Role during attention

Some evidence implicates area V4 as having a particularly critical role in visual attention. Lesion studies show specific attentional defects after loss of V4 (Buffalo et al., 2005; Schiller and Lee, 1991; De Weerd et al., 1999, 2003). These studies found that lesion of area V4 produced deficits in visual performance that were distracter dependent, a type of effect classically interpreted as being indicative of attentional disruption. Perhaps, these studies suggest, area V4 has a particular role in the selection of visual objects for attention or in the integration of attentional signals with visual representations. However, it is worth noting that distracter dependent deficits may well reflect mere loss of visual representation strength relative to competing representations rather than any deficit in controlling the focus of attention or in applying the concurrent modulatory effects on visual cortex (Squire et al., 2013). Thus, whether V4 plays a particular role in integrating attention-related signals with visual representations or is merely one of many areas to reflect this integrated signal remains unclear.

Nevertheless, V4 neural activity is robustly modulated during attention in a spatially specific manner. When covert attention is directed to RF stimuli, relative to when attention is directed to stimuli outside their RFs, they exhibit increased firing rates (Moran and Desimone, 1985), sensitivity to low-contrast stimuli (Reynolds et al., 2000), and across-trial reliability (Mitchell et al., 2007). Their orientation tuning amplitude is typically multiplicatively enhanced (McAdams and Maunsell, 1999; Williford and Maunsell, 2006), and their tuning in Fourier space shifts toward the attended features (David et al., 2008). They have enhanced high frequency coherence with the V4 population activity (Fries et al., 2001) and with attentionally modulated neurons in FEF (Gregoriou et al., 2009, 2012), while low frequency correlations within V4 decrease (Cohen and Maunsell, 2009; Fries et al., 2001; Mitchell et al., 2009). Their RF locations shift toward the attended target (Connor et al., 1997). All of

these changes in V4 neural activity during attention have been suggested to enhance the amount of information in the population accessible to downstream neurons and/or to enhance the ability of the attended representation to outcompete other stimulus representations for control of behavior in downstream targets.

Interestingly, several of the same effects observed during sustained covert attention tasks have also been observed in visually-guided saccade tasks in the hundred milliseconds or so immediately preceding the execution of the saccades to stimuli in their RFs. During this period, V4 neurons have increased firing rates (Fischer and Boch, 1981; Mazer and Gallant, 2003), enhanced orientation selectivity (Moore and Chang, 2009), enhanced contrast sensitivity (Han et al., 2009), and their RF locations shift toward the saccade target (Tolias et al., 2001). In Chapter 2, we show that another of the signatures of covert attention, decreased across-trial variability, is also true of this presaccadic period (Steinmetz and Moore, 2010). We argue that these striking similarities of visual cortical modulation during saccade preparation to that during covert attention comprise another strong reason, along with the psychophysical similarities and shared network of controlling brain regions, to suspect that covert attention is saccade preparation without execution. Nevertheless, the possibility remains that the similarity of visual cortical modulations merely results from an optional shifting of covert attention to the saccade target. Therefore, in Chapter 3, we test the hypothesis directly by dissociating saccade preparation from covert attention behaviorally and test how these two processes modulate visual cortex.

Altogether, there is strong evidence that area V4 neurons comprise a critical visual representation underlying the enhancements in perception associated with covert attention and saccade preparation (though a causal demonstration of this role is still lacking).

1.3.2.2 Other forms of modulation: pop-out, surround suppression, figure-ground segregation

Area V4 responses are also modulated when certain stimulus configurations occur outside of their RFs. For instance, V4 responses to RF stimuli are attenuated when

large stimuli are displayed in other parts of the visual field, a phenomenon known as surround suppression (Desimone and Schein, 1987; Sundberg et al., 2009). They are often modulated by the identity of the RF stimulus as either part of the “figure” (coherent object in the foreground) or the “ground” (background) of a scene, even though the visual stimulus within the RF is identical in both cases (Zhou et al., 2000). Finally, they may be modulated by the relative salience of the RF stimulus compared to the rest of the scene, as in the pop-out effect (Burrows and Moore, 2009). In the pop-out effect, the response of a V4 neuron to a RF stimulus that is unique among all stimuli on the screen was shown to be larger if the other stimuli all differed from the RF stimulus along some dimension compared to if the other stimuli contained mixtures of the properties of the RF stimulus. For instance, a vertical green bar amongst vertical red bars will elicit a strong response, while a vertical green bar amongst horizontal green, horizontal red, and vertical red bars will elicit only a weak response.

These three types of V4 modulation depend on the stimulus properties rather than on any learned task contingencies, unlike attention-related modulation, and could therefore be plausibly implemented by connections solely within area V4, such as the superficial layer horizontal connections within V4 (Yoshioka et al., 1992). However, evidence suggests that all three forms of modulation may in fact be mediated by feedback rather than horizontal interactions. The timing of surround suppression interactions is such that they are more likely to be implemented by cortico-cortical feedback (Angelucci and Bullier, 2003). Figure-ground modulation, at least in area V1, depends on feedback from higher cortical areas (Hupé et al., 1998). Finally, the pop-out effect was eliminated when saccades were executed away from the RF stimulus, indicating that it depends on the availability of a top-down resource (Burrows and Moore, 2009).

All three of these types of modulation have some relevance for the present work. First, surround-modulation may be particularly influenced during covert attention (Sundberg et al., 2009), an idea that has been influential in understanding the potential mechanisms of the modulation of visual cortical responses during attention (Reynolds

and Heeger, 2009), and will be discussed further below. Second, recent results have revealed that the figure-ground modulation of area V1 may be almost exclusively driven via feedback onto NMDA synapses (Self et al., 2012). This observation corroborates recent experimental work suggesting that NMDA-mediated inputs may have the effect of a gain modulation on visual cortical responses (Smith et al., 2013). Therefore, in Chapter 6, we chose to investigate NMDA-mediated feedback as a mechanism to achieve gain modulation in a large-scale model of visual and frontal cortex. Finally, the observation that pop-out modulation depends on top-down feedback suggests that the work described here in discovering and modeling the circuitry mediating top-down covert attention may have broader relevance to so-called “bottom-up” attention as well.

1.3.3 Possible sources and mechanisms of V4 modulation

What is the source of attention-related modulation of area V4 neurons? Any full account of the circuitry underlying attention should provide an answer to this simple but unresolved question. Addressing this question comprises part of the aim of this thesis.

The particular anatomical pathway that drives effects in area V4 could be any one of the following (see Figure 1-2): direct projections from FEF to V4 (Anderson et al., 2011; Pouget et al., 2009; Stanton et al., 1995); indirect corticocortical feedback projections to V4 via area LIP (Andersen et al., 1990); subcortical projections through the superior colliculus and the pulvinar nucleus of the thalamus (either specific or non-specific) (Jones, 2007; Trojanowski and Jacobson, 1976); feedback down the visual hierarchy, from areas like IT (Felleman and Van Essen, 1991; Ungerleider et al., 2008); or indirect projections through a neuromodulatory system, such as the cholinergic nucleus basalis of Meynert (Tanaka et al., 1990).




Figure 1-2. Projections to area V4 from areas that may drive attention-related modulation. A, Anatomical locations of some areas that may provide attention-related feedback to V4. B, Nissl-stained slice of area V4 with locations of axon terminals from the three feedback pathways shown in A (green arrows, SC/Pulvinar; yellow, LIP; red, FEF). Cortical layers are numbered. Not all attention projections are shown (see text). Anatomical image from www.brainmaps.org.

These different pathways have distinct patterns of termination across layers in area V4. Projections from FEF terminate in all cortical layers. Projections from IT and LIP terminate in superficial and deep layers. The pulvinar has both a specific pathway projecting to deep layer 3 and a non-specific pathway projecting to apical dendrites (Benevento and Rezak, 1976; Jones, 2007). The laminar distribution of cholinergic projections to area V4 is not known; however in macaque area V1 there are nicotinic AchRs on thalamocortical afferent axons in layer 4 (Disney et al., 2007) and muscarinic receptors elsewhere (primarily on GABAergic interneurons, Disney et al., 2006).

All of these laminar projection targets have at least some support from the literature regarding possible roles in attention. Experimental data and models of selective attention variously predict that attention-related modulation is mediated through Ach receptors (Disney et al., 2007; Herrero et al., 2008), through projections onto apical dendrites (Spratling and Johnson, 2004), through projections to layer 6 (Raizada and Grossberg, 2003), and through thalamic gating (Crick, 1984; McAlonan et al., 2008).

These different termination patterns suggest a novel approach for identifying the source area of attention-related signals arriving in area V4: by measuring the latency and strength of attention-related modulation across cortical layers, the input pattern of attention-related signals in depth may be discovered, thereby disambiguating which projection is the source of those signals. In Chapter 4, I will describe an experiment designed to collect the data necessary to address this question by recording across layers in area V4 while monkeys performed a task requiring covert attention.

1.3.4 Microcircuitry within area V4

All regions of neocortex are defined by their characteristic laminar structure consisting usually of six layers each with typical cell types and connectivity, briefly outlined here (for further review, see Douglas and Martin, 2004). Layer 4, the “granular” layer, receives bottom-up sensory input from thalamus in primary sensory cortex or, in the case of V4, L4 receives bottom-up input from hierarchically lower visual cortical areas, primarily V1 and V2. L4 projects strongly to Layers 2 and 3 (typically considered together, though they are visually distinct in area V4), whose excitatory population consists of intracortically-projecting pyramidal neurons. L2/3 projects to L5, also consisting of pyramidal neurons, but which project subcortically. Both L2/3 and L5 pyramidal neurons have apical dendrites that extend into L1 and there receive feedback inputs from higher cortical areas and parts of the thalamus. L6 also receives feedback projections from higher cortical areas, and projects strongly to L4. All layers have inhibitory neurons of many types, though one notable population is the fast-spiking inhibitory interneuron, also called basket cells, found in L4 and L2/3. This is but a brief overview of the most prominent connections; in reality the connectivity is much more diverse and there are many cell types not mentioned here (Binzegger et al., 2004; Cauli et al., 1997; Thomson and Bannister, 2003; Thomson and Lamy, 2007). However, considering only these connections forms a useful approximation to the so-called “canonical” neocortical microcircuit (Douglas and Martin, 2007; Douglas et al., 1989). Finally, the neocortex is organized into repeating columns of the cell types and connections mentioned here (Mountcastle, 1997). Connections between layers occur primarily in the vertical direction, between other

neurons in the same column. This property lends vertically offset cells the property of having coextensive receptive fields (Albright et al., 1984; Hubel and Wiesel, 1968).

Comparison of the functional properties of neurons in different cortical layers within area V1 comprises perhaps the greatest success story in understanding the computation performed by neurons within any cortical area (Hirsch and Martinez, 2006; Hubel and Wiesel, 1968). Nevertheless, experiments to date recording from area V4 (and most other cortical areas) have almost uniformly ignored this complexity and considered neurons indiscriminately from different depths and neuron populations. This oversight has been undoubtedly due in large part to the difficulty of obtaining those identifications due to the lack of suitable tools, aside from the cumbersome, coarse, and ethically-weighty method of post-hoc histology.

Experiments that have attempted to measure differences across layers within V4 have primarily done so with electrodes only capable of recording local field potential responses, not the activity of single neurons. These studies have therefore focused on visually evoked potentials (VEPs) (Givre et al., 1994) or the timing of different signals arriving to area V4 (Chen et al., 2007). One study reported the modulation of VEPs across depth in V4 during an intermodal (visual versus auditory) attention task (Mehta et al., 2000). Authors reported that initial visual-driven responses in layer 4, the ventral stream input layer, were unmodulated by attention. However, because the authors studied VEPs, which are of uncertain origin, rather the spiking activity of neurons, their results are difficult to understand in terms of the underlying microcircuitry. Furthermore, because they did not use a spatially selective attention task, the results are difficult to put into context with other studies of attention.

Another study also measured effects of attention in different layers within area V4 (Buffalo et al., 2011). In this experiment, authors used the method of labeling neurons as “superficial” if they were encountered within 1mm after entering the brain and as “deep” if they were encountered after that. Without a functional register of depth, these labels seem uncertain. Nevertheless, authors reported large differences in spike-field coherence across depth, with strong gamma coherence in superficial layers and

alpha (low-frequency) coherence in deep, suggesting that attention may well differently influence the distinct elements of the V4 anatomical microcircuit, as outlined above. This study did not address whether neurons in superficial or deep layers were modulated more or less strongly during attention.

In Chapter 4, I describe the development of a recording technique capable of recording single neurons at multiple depths within area V4 simultaneously. I describe some of the basic differences observed between layers, including differences in orientation tuning and response latency. I also report measurements of the magnitude of attention-related modulation across depths. These experiments thus represent a step forward for attempts to decipher the functional consequences of the complex neocortical circuitry for visual and cognitive processing in extrastriate cortex.

1.4 The Frontal Eye Field (FEF)

The FEF is a region of primate prefrontal cortex defined as the area in which low-current electrical stimulation evokes saccades. Anatomically, the FEF is located primarily in the prearcuate sulcus, specifically Brodmann areas 45A, 45B, and 8Ac (see Figure 1-3). FEF notably shares cytoarchitectural characteristics of both primary motor cortex (a high concentration of large layer 5 pyramidal neurons) and granular frontal cortex (GFC; a somewhat distinct granular layer 4 as well as large pyramidal neurons in layer 3) (Stanton et al., 1989). This unique duality at the cytoarchitectural level of description links FEF both with purely motor areas, from which other types of motor actions can be elicited with electrical stimulation, and with GFC regions, which are more associated with higher-order cognitive functions. As we will see, this duality has been borne out through detailed study of the functional roles of FEF.

Stimulation, recording, and inactivation experiments have shown both that the FEF appears to play a significant role in the planning and execution of saccadic eye movements and that it participates in the control of visual selective attention. Recent experiments have furthermore begun to yield some understanding of the functional circuitry within FEF linking neurons involved in these two roles. In this section I will discuss the evidence for FEF's roles in overt and covert attention. I will particularly examine evidence about the circuitry both within FEF and between FEF and other structures in visual and oculomotor regions of the brain, circuitry that underlies the FEF's dual roles in overt and covert attention.




Figure 1-3. The connectivity and anatomical locations of the Frontal Eye Field (FEF) and other structures within the visual and oculomotor systems of the rhesus macaque monkey brain. Left panel, diagram of the connectivity of the FEF (orange) with other visual and oculomotor structures. Some connections that do not directly involve the FEF are omitted. Right panel, locations of the brain regions pictured in the left panel shown in the lateral view of the monkey brain. In both panels, surface structures other than the FEF are colored darker blue. Deep structures are illustrated as lighter blue and with a dashed outline. Abbreviations: BG, basal ganglia; BSG, brainstem saccade generator; Cbl, cerebellum; Ex Vis Ctx, extrastriate visual cortex; LIP, lateral intraparietal area; SC, superior colliculus; SEF, supplementary eye field; Thal, thalamus; V1, primary visual cortex.

1.4.1 The role of the FEF in saccadic behavior

[Portions of this section of text has been previously published in: Squire, R. F., Steinmetz, N. A., & Moore, T. (2012). Frontal Eye Field. *Scholarpedia*, 7(10), 5341.]

1.4.1.1 Connectivity with other saccade-related areas

Frontal Eye Field neurons interconnect extensively with other known structures of the primate saccadic system (see Figure 1-3). The FEF has topographic projections directly to the intermediate layers of the ipsilateral superior colliculus (SC) (Leichnetz et al., 1981), particularly to neurons exhibiting saccade related activity (Helminski and Segraves, 2003). The FEF also projects to the ipsilateral caudate and putamen (Kunzle and Akert, 1977), to the cerebellum via the pontine nuclei, and to many oculomotor-associated nuclei in the midbrain and pons including the “brainstem saccade generator” nuclei (Huerta et al., 1986; Leichnetz et al., 1984). In the brainstem, the FEF drives “burst” neurons in the paramedian pontine reticular formation, which control the direction and amplitude of saccades. These burst neurons are gated by “pause” neurons in the nucleus raphe interpositus, and directly drive the motor

neurons when un-gated. The FEF receives subcortical projections from various saccade-related thalamic nuclei including the lateral mediodorsal (MD), medial ventral anterior (VA), and medial pulvinar, among others (Huerta et al., 1986). In addition, the intermediate layers of the SC, the substantia nigra pars reticulata, and the dentate nucleus of the cerebellum all project to the FEF indirectly via the MD and VA thalamic nuclei (Lynch et al., 1994). Along with the extensive connections with subcortical oculomotor structures, the FEF has extensive connections with other saccade-related cortical areas, including the supplementary eye field (SEF), dorsolateral prefrontal cortex (dlPFC), area 7m in the medial parietal lobe, and the lateral intraparietal area (LIP), all cortical regions from which saccades can also be elicited via electrical stimulation, albeit with higher currents (Lynch and Tian, 2006).

1.4.1.2 Stimulation experiments

The saccades evoked by electrical stimulation of the FEF can be used to reveal its functional topography as stimulation of nearby sites usually elicits saccades to nearby locations in retinotopic space (Bruce et al., 1985; Robinson and Fuchs, 1969). However, discontinuities in a series of evoked saccades from nearby cortical sites are frequently encountered, suggesting that the topography is not strictly retinotopic. Nonetheless, small saccades are generally elicited by stimulation of area 45 within the ventrolateral limb of the arcuate sulcus, while large amplitude saccades are elicited by stimulation of area 8A within the dorsomedial limb. In addition, within the dorsomedial portion of the FEF, combined head and eye movements can be elicited (Monteon et al., 2010; Tu and Keating, 2000) consistent with a general role of the FEF in gaze control, as gaze is defined as involving both head and eye components.

1.4.1.3 Lesion and inactivation experiments

Many years of lesion experiments and, more recently, reversible inactivation experiments, have elucidated the particular role of the FEF in saccadic behaviors (Tehovnik et al., 2000). When inactivated, saccades to visual targets slower, less frequent, and less accurate (Dias and Segraves, 1999; Sommer and Tehovnik, 1997). Lesion experiments produce similar deficits though not as pronounced, likely due to recovery or reallocation of function over time (Schiller et al., 1980, 1987). However,

in both inactivation and lesion experiments, deficits are substantially greater for saccades without visual targets, that is, for memory-guided saccades. The FEF may therefore be particularly important for guiding eye movements on the basis of endogenous rather than exogenous factors.

1.4.1.4 Activity of FEF neurons related to saccades

Many studies have characterized the activity of FEF neurons in a variety of tasks requiring saccades. The most commonly used paradigm is the memory-guided saccade task, as it allows for the identification of neurons with saccade-related activity independent of any visual stimuli. Approximately 10% of FEF neurons have purely movement-related activity (“M-type”), responding only around the time of saccades and not when the visual stimulus appears to cue the location of the future saccade (Bruce and Goldberg, 1985). A further 20% have movement-related activity but also respond to the presentation of visual stimuli (“VM-type”). Other FEF neurons in this task respond post-saccadically or else not around the time of saccades at all; this latter class comprises the majority of FEF neurons (Bruce and Goldberg, 1985). In a visually-guided saccade task, some of the neurons which did not appear to respond around the time of saccades nevertheless do respond when saccades are directed to RF stimuli. Importantly, in these visually-guided saccade tasks, saccades appear to be initiated when FEF movement-related activity reaches a threshold, suggesting that the FEF neurons are directly involved in, or perhaps themselves constitute, the trigger for saccade initiation (Hanes and Schall, 1996). Saccade-related activity has also been observed in antisaccade tasks, which require monkeys to implement more complicated correspondences between visual stimuli and eye movements (Everling and Munoz, 2000; Sato and Schall, 2003). In these tasks, some neurons respond when either the informative visual cue or the saccade target falls in their RF, while other neurons respond only to saccade targets. Taken together, these findings provide strong evidence that the FEF is intimately involved in most or all gaze-related behavior.

1.4.2 The role of the FEF in covert attention

1.4.2.1 Connections to visual areas of the brain, including area V4

In addition to its connectivity with structures related to eye movements, the FEF contains extensive reciprocal connectivity with areas containing visual representations. These areas include ventral stream areas such as V2, V3, V4, TEO, and TE as well as dorsal stream areas such as MT, MST, and LIP (Schall et al., 1995; Stanton et al., 1995). Furthermore, these connections appear to be at least coarsely topographic, in that the part of the FEF that represents larger-amplitude saccades (area 8A; see Section 1.4.1.2) interconnects with peripheral regions of visual cortical retinotopic maps, while the part representing smaller-amplitude saccades (area 46) interconnects with more central representations in visual cortex. In general the projections from FEF to these visual areas resemble “feedback” projections based on their laminar pattern of axonal terminations while the reverse projections, from visual areas to FEF, are “feedforward” type (Felleman and Van Essen, 1991). Interestingly, one exception to this rule is the projection from FEF to area V4, which more resembles a feedforward projection (Barone et al., 2000); the significance of this observation remains unclear. These extensive connections with the visual cortex grant FEF the ability to influence many aspects of visual processing and potentially provide the substrate for the FEF’s role in covert attention, as discussed in this section.

1.4.2.2 Lesion and inactivation experiments

Some of the earliest evidence linking FEF with a role in attention came from studies that described contralateral neglect, a condition canonically characterized as an attention-related deficit, after FEF lesions (Lynch and McLaren, 1989; Welch and Stuteville, 1958). However, given the known deficits in eye movement production following FEF inactivation (see Section 1.4.1.3), a more thorough test of the FEF’s role in covert visual attention would involve a task specifically designed to not require eye movements. Indeed, reversible inactivation of FEF in such covert visual search tasks also produced behavioral impairments (Monosov and Thompson, 2009; Wardak et al., 2006).

1.4.2.3 Stimulation experiments

Even more compelling causal demonstrations of the role of the FEF in covert attention have come from experiments employing electrical microstimulation to subtly enhance FEF activity in a spatially specific manner. When stimulated with currents below the threshold for eliciting saccades, monkeys exhibit spatially specific enhancements in their ability to make difficult visual change detections, similar to the perceptual benefits from endogenous selective attention (Moore and Fallah, 2001, 2004). Further experiments also observed spatially specific enhancements of visual cortical representations following the same subthreshold electrical stimulation approach (Armstrong and Moore, 2007; Armstrong et al., 2006; Ekstrom et al., 2009; Moore and Armstrong, 2003).

Since electrical microstimulation may activate both local neurons as well as passing axons from distant structures (Clark et al., 2011), an even more direct test might involve causal manipulations that could more specifically activate only FEF neurons. Two recent experiments have provided such a test. In the first, monkeys controlled the activity of their own FEF neurons via operant conditioning (Schafer and Moore, 2011). When monkeys succeeded in enhancing the activity of FEF neurons, there were concurrent spatially specific perceptual enhancements. In the second, local injections of dopamine D1-receptor antagonists, a manipulation which enhances activity in prefrontal cortical neurons, brought about changes in visual cortical activity similar to those seen with covert attention and with FEF microstimulation (Noudoost and Moore, 2011a).

These experiments provide a strong causal evidence for the FEF's involvement in the control of covert attention and in the production of the visual cortical modulations that presumably underlie this behavioral phenomenon.

1.4.2.4 Activity of FEF neurons related to attention

Given the evidence above, it would be highly surprising if FEF activity were not modulated during covert attention. Indeed, when covert attention is directed to stimuli in FEF neurons' RFs, activity of these neurons is enhanced, even in tasks which do not

require eye movements (Armstrong et al., 2009; Gregoriou et al., 2012; Kodaka et al., 1997; Monosov et al., 2008; Thompson et al., 2005). Interestingly, but perhaps not surprisingly, only certain neurons within FEF exhibit enhanced firing rates with covert attention. Specifically, while neurons with visual but not movement-related activity (“V-type”) as well as those with both visual and movement activity (“VM-type”) are enhanced, neurons with only movement-related activity (“M-type”) are suppressed or unaffected (Gregoriou et al., 2012; Thompson et al., 2005). Furthermore, it is only the V-type neurons that become time-locked to high frequency activity in visual cortical representations (Gregoriou et al., 2012), perhaps suggesting that just these neurons project directly to, or receive projections directly from, visual cortical neurons. The significance of this work is discussed further below.

1.4.3 Circuitry within FEF mediating overt and covert attentional functions

Covert and overt attention are inextricably linked at the psychophysical level and at the level of visual cortical modulations (see Section 1.2, above, and Chapter 3, below). In this section I consider the neural circuitry that could underlie this relationship.

There are three basic hypotheses for the circuits that could result in an obligatory invocation of attention with saccade preparation. First, there could be just one population of neurons mediating both functions. Second, there could be a population of neurons devoted to each function, and the attention-related population, which drives modulation of visual cortex by definition, could be downstream of the saccade preparation neurons, which drive saccades when activated strongly enough. Third, there could be two populations as in the second hypothesis, but the saccade preparation population could be downstream of the attention population.

The first hypothesis, one population mediating both functions, is unlikely for several reasons. Empirical evidence about the nature of the particular instantiations of these populations suggest this is unlikely, since these functions appear to be subserved by distinct populations of neurons in likely source areas such as FEF (Thompson et al.,

2005), SC (Ignashchenkova et al., 2004), and LIP (Gottlieb et al., 1998). Secondly, certain manipulations of cortical activity can influence attention but not saccade preparation (Schafer and Moore, 2011) or saccade preparation but not attention (Noudoost and Moore, 2011a), indicating that the two populations must be at least experimentally dissociable.

The second hypothesis seems likely at a glance, since the obligatory linking of attention to saccade preparation would be easily accounted for if the former were downstream of the latter. However, manipulating FEF activity via dopamine D2-receptor antagonist influenced saccadic but not attentional behavior, suggesting that instead, the population driving saccades is downstream of the attention population. Furthermore, from a high-level perspective, it seems that since visual stimuli guide eye movements in natural behavior, it is only reasonable that the saccade population is downstream from the attention population.

Thus, the third hypothesis, that the population responsible for driving visual cortical modulations drives the population that controls saccades (Figure 1-4), seems the most reasonable. In this case, the ability to covertly attend can be understood as activation of the attention population with, additionally, suppression of the saccade population. This fits nicely with the result that movement-related FEF neurons are suppressed during covert attention (Gregoriou et al., 2012; Thompson et al., 2005), as if being actively inhibited.

Furthermore, it explains how the dopamine D2R manipulation mentioned above (Noudoost and Moore, 2011a, 2011b) could influence saccade preparation without influencing attentional allocation, though this appears to be impossible via endogenous behavior. That is, when the monkey endogenously drives the activity of the saccade population, it can only do so via the attention population, which provides driving input to the saccade population. However, an experimental manipulation (D2R antagonism) may be, and apparently is, capable of intervening in this circuit downstream of the attention population, thereby influencing only the saccade process. We would certainly expect that there must be some point in the circuit to saccade production at

which activity could be unnaturally (i.e. not endogenously; not voluntarily) manipulated to affect saccade production without modulating visual cortex. For instance, a manipulation of brainstem saccade nuclei would modulate propensity to saccade one direction versus another, but there would be no route back to visual cortex for a signal injected at that late stage of the anatomical circuit. Thus, the results of Noudoost et al. importantly provide an “upper bound” on where such an exogenous manipulation can affect saccade behavior but not visual cortical modulation. Specifically, as those authors note, D2Rs are found in deep cortical layers of FEF, so manipulating activity directly in these layers must fail to modulate visual cortical activity (and indeed, deep FEF neurons only rarely project to V4 (Pouget et al., 2009); we do not know specifically whether these sparse, V4-projecting deep FEF neurons express D2Rs). The result from the same experiment that D1R manipulation influenced both visual cortical responses and saccade behavior was, in this understanding, obtained either because the manipulation directly influenced both populations, or because it only influenced the attention population, upstream of the saccade population.




Figure 1-4. Hypothetical circuitry mediating the control of covert attention and saccades. Triangles represent excitatory pyramidal neurons with projections to visual cortex (cyan) and to subcortical saccade

structures (red). Black circles represent a population of inhibitory neurons. Arrows with triangle heads represent excitatory drive and the arrow with a circle head represents inhibition. Dashed arrows represent projections only active in certain contexts (see text).

Thus the suggestion is that it is at the level of different neuron types within FEF where the dissociation between overt and covert attention arises. M-type neurons might normally be activated by V- or VM-type neurons as part of typical visuomotor behaviors in which visual stimuli of interest drive eye movements to foveate them. However, in the case of covert attention, the V- and VM-type neurons will be activated by the visual stimulus of interest even though no eye movement is desired, necessitating the inhibition of saccade-related signals to downstream areas. If V- or VM-type neurons provide feedback to visual areas to mediate covert attention, then by activating these neurons as in overt attention but additionally suppressing M-type neurons, selective attention may be achieved. Thus V- and/or VM- neurons comprise the “attention source neurons” and M-type neurons comprise the “saccade command neurons” of Figure 1-4.

On an anatomical basis, it has been suggested that the superficial layer FEF neurons correspond to those with visually-driven activity (V-type) and the deep neurons correspond to those with motor-related activity (M-type). Specifically, superficial neurons project to visual cortex and deep neurons project to the SC and brainstem saccade nuclei, suggesting their involvement in the functions respectively associated with these areas. Indeed, neurons projecting to the brainstem nuclei are M-type (Segraves, 1992). However, neurons projecting to the SC are mixed, such that some are M-type but others have visual- and memory-related activity (Sommer and Wurtz, 2001). An important outstanding question is what type of FEF neurons project to visual cortex, but we certainly hypothesize that these are primarily neurons representing visual information.

Assuming that the anatomical identification of V-type with superficial and M-type with deep neurons is roughly accurate, then the inference that attention-related (V-type) neurons drive saccade-related (M-type) is consistent with the idea of strong superficial-to-deep connectivity within a cortical area (Douglas and Martin, 2004).

Modulating the deep neurons in such a circuit would require activating the superficial ones en route. Conversely, the superficial could be activated without the deep if there is an additional suppression applied to deep layer neurons.

These same conclusions about the circuitry within area FEF – namely, that superficial neurons represent visual and attention-related information while deep represent movement information, and that the superficial drive the deep – formed the backbone of a computational model of the FEF (Heinzle et al., 2007). This model, with connectivity based primarily on cat visual cortical data (Binzegger et al., 2004), accounted for basic FEF behaviors including memory-guided saccade tasks and antisaccade tasks. Thus this model seems to fit well with the more recent experiments discussed here. Furthermore, it can be used make predictions for how FEF neurons will participate in novel experiments (Chapter 5).

Despite the apparent clarity of these circuit descriptions of activity within area FEF, there is perhaps reason to believe that the true story is not quite so simple. It has always been observed that V-, VM-, and M-type FEF neurons fall on a continuum between purely visual and purely motor activity (Bruce and Goldberg, 1985; Thompson et al., 2005, but see Cohen et al., 2008). Perhaps a clearer way to conceive of these neurons is not that they fall into one of three classes but rather that they each represent some combination of the two (or more) underlying signals. Indeed, a recent study employing a more complicated task design found that FEF neurons represented combinations of stimulus-related information (both color and motion), saccade-related information, and context information (Mante et al., 2013). Perhaps by understanding FEF activity as reflecting combinations of patterns of activity that exist in a distributed way across the population, new or truer features of the visuomotor computation will be revealed.

1.5 Models of Visual Attention

As reviewed in section 1.3.2, the specific effects of spatial attention on the activity of single visual cortical neurons, particularly in area V4, have been well characterized. In particular, those neurons that represent visual information at the focus of attention increase their firing rate while those outside decrease or are unaffected (Moran and Desimone, 1985). The increases in firing rate often appear to be multiplicative, such that responses to preferred stimuli are increased more than responses to non-preferred stimuli (McAdams and Maunsell, 1999; Williford and Maunsell, 2006). However, some reports suggest that the effects of attention look more like an additive input rather than multiplicative (Reynolds et al., 2000). In addition to firing rate changes, the temporal patterns of neural firing change such that neurons representing stimuli under attention tend to oscillate more strongly in the gamma frequency range (30-70Hz) and tend to synchronize their firing as indicated by increased spike-field coherence in the gamma range (Bichot et al., 2005; Fries et al., 2001; Womelsdorf et al., 2006). Finally, there is evidence that surround stimulus interactions are particularly influenced during attention (Sundberg et al., 2009).

Thus, to account for spatial attentional modulation at the cellular level, mechanisms must be specified to implement multiplicative gain, competitive normalization, and oscillatory entrainment. A number of distinct mechanisms are plausible and have been proposed for each of these neuronal level problems. Multiplicative gain could be mediated by nonlinear apical dendritic summation (Larkum et al., 2004; Smith et al., 2013), the activation of acetylcholine receptors (Disney et al., 2007; Herrero et al., 2008), by default upon increased excitatory input (for example from feedback) due to noise properties of synaptic input (Murphy and Miller, 2003), by decreases in the level of background balanced synaptic input (Chance et al., 2002), by inhibitory network oscillations (Tiesinga et al., 2004), by modulation from layer 6 inhibitory neurons (Olsen et al., 2012).

Competitive normalization interactions can be achieved with a population of inhibitory neurons that sums the outputs of a broad range of excitatory neurons and

proceeds to inhibit all of that population or of another downstream population (feedback and feedforward normalization, respectively) (Carandini et al., 1997; Kouh and Poggio, 2008; Swadlow, 2003). These summing interneurons can act divisively through shunting inhibition (Borg-Graham et al., 1998; Mitchell and Silver, 2003, but see Holt and Koch, 1997) or balanced increases in synaptic input (Chance et al., 2002). Many network level models of visual attention include this type of interaction (Buia and Tiesinga, 2008; Hamker and Zirnsak, 2006; Spratling and Johnson, 2004; Wagatsuma et al., 2011). Some have even suggested that competitive normalization may be a mechanism through which various types of attentional effects may be unified (Reynolds and Heeger, 2009). Divisive competition may also be mediated by feedforward synaptic depression (Carandini et al., 2002).

Oscillatory increases can arise naturally out of the interactions of recurrently connected inhibitory populations (Bartos et al., 2002; Börgers et al., 2005), as are known to exist in the neocortex (Beierlein et al., 2003). Recent experiments have provided a causal demonstration of these inhibitory populations as the source of gamma-frequency oscillations (Cardin et al., 2009). Perhaps relatedly, the relative strength of low and high frequency oscillations changes as a function of cortical state (Okun et al., 2010; Poulet and Petersen, 2008), and shifts in cortical state have also been proposed as a mechanism underlying covert attention (Harris and Thiele, 2011).

Given the profusion of possible explanations for each of these phenomena, and particularly for gain modulation, it may be useful to think of each as a “computation” which may have different underlying biological instantiations depending on the particular system (Carandini, 2012; Marr, 1982). Such a framework makes it clear that for each system in which the computation is observed, a different mechanism may underlie it, and experiments will need to be done in that system to determine which. Thus, specific hypotheses for plausible mechanisms cannot be too great in number, even if other mechanisms have been shown conclusively to operate in other systems. In Chapter 6 we describe that a novel mechanism, feedback mediated by voltage-dependent glutamatergic NMDA receptors, can explain some features of attention-related modulation in visual cortex. This mechanism, demonstrated in the context of a

large-scale model of areas FEF and V4, accounts for the multiplicative modulation of firing rates. Future modeling work could investigate the way the NMDA mechanism proposed in Chapter 6 is or is not compatible with ideas about competitive normalization and oscillatory entrainment.

2 Changes in the Response Rate and Response Variability of Area V4 Neurons During the Preparation of Saccadic Eye Movements

This chapter has been previously published, with some modifications, as:

Steinmetz, N. A., & Moore, T. (2010). Changes in the Response Rate and Response Variability of Area V4 Neurons During the Preparation of Saccadic Eye Movements. *Journal of Neurophysiology*, 103(3), 1171–1178.

2.1 Abstract

The visually driven responses of macaque area V4 neurons are modulated during the preparation of saccadic eye movements, but the relationship between presaccadic modulation in area V4 and saccade preparation is poorly understood. Recent neurophysiological studies suggest that the variability across trials of spiking responses provides a more reliable signature of motor preparation than mean firing rate across trials. We compared the dynamics of the response rate and the variability in the rate across trials for area V4 neurons during the preparation of visually guided saccades. As in previous reports, we found that the mean firing rate of V4 neurons was enhanced when saccades were prepared to stimuli within a neuron's receptive field (RF) in comparison with saccades to a non-RF location. Further, we found robust decreases in response variability prior to saccades and found that these decreases predicted saccadic reaction times for saccades both to RF and non-RF stimuli. Importantly, response variability predicted reaction time whether or not there were any accompanying changes in mean firing rate. In addition to predicting saccade direction, the mean firing rate could also predict reaction time, but only for saccades directed to the RF stimuli. These results demonstrate that response variability of area V4 neurons, like mean response rate, provides a signature of saccade preparation. However, the two signatures reflect complementary aspects of that preparation.

2.2 Introduction

Visual perception and oculomotor control are known to interact. In one direction, the features of a visual scene influence the patterns of saccadic eye movements (Vishwanath and Kowler, 2003; Yarbus, 1967). Underlying this influence is presumably the projection of visual cortical representations onto oculomotor structures (Edelman and Keller, 1996; Keller and Edelman, 1994; Moore, 1999). Conversely, psychophysical evidence demonstrates that the preparation of saccadic eye movements informs perception of visual targets, enhancing visual sensitivity at the intended saccade location (Deubel and Schneider, 1996; Hoffman and Subramaniam, 1995). Correspondingly, the mean firing rates of single neurons in some areas of visual cortex have been shown to be modulated during the preparation of saccades to receptive field stimuli, suggesting a direct influence of saccade preparation on these neurons (Chelazzi et al., 1993; Fischer and Boch, 1981; Mazer and Gallant, 2003; Moore et al., 1998; Nakamura and Colby, 2002; Sheinberg and Logothetis, 2001; Tolia et al., 2001). Mimicking endogenous saccade signals by electrically stimulating sites within the frontal eye field (FEF) yields similar modulation of visually driven responses in visual area V4 (Armstrong and Moore, 2007; Moore and Armstrong, 2003), suggesting that the perisaccadic modulation observed during voluntary saccades originates from oculomotor structures (Moore et al., 2003). In spite of the preceding evidence, our understanding of the nature of the oculomotor influence on visual cortex and the contribution of extrastriate areas to saccade preparation remains incomplete.

Thus far, evidence of an influence of saccade preparation on extrastriate neurons has been exclusively examined in terms of perisaccadic modulations in mean firing rate (Fischer and Boch, 1981; Moore and Chang, 2009; Tolia et al., 2001). However, a recent study suggests that the across-trial variability of neuronal firing rate provides a more robust signature of motor preparation (Churchland et al., 2006). This study examined the relationship between the activity of neurons in dorsal premotor cortex and the reaction time of monkeys performing a delayed reach task. Although the mean firing rate of premotor neurons did not predict reaction time, changes in the across-trial variability of firing rate did. This observation suggests that firing rate variability

may be a more sensitive measure of behavioral state than mean firing rate and thus may be a more robust signature of motor preparation. A recent study of extrastriate area V4 observed attention-dependent changes in across-trial variability of neuronal response rates (Mitchell et al., 2007). Given the well established relationship between attention and saccade preparation (Moore, 2006; Schafer and Moore, 2007), across-trial variability of response rates of V4 neurons may also provide an index of motor preparation.

To assess the interaction between saccade preparation and visual cortical representations, we measured the mean firing rate and variability across trials of spike trains recorded from area V4 neurons in monkeys trained to make saccades to visual targets. Response variability was measured by the Fano factor (FF), which was computed by dividing the across-trial variance in spike counts within a small window by the mean count. As expected, the mean firing rate of V4 neurons was enhanced when saccades were prepared to stimuli within a neuron's receptive field (RF) in comparison with saccades to a non-RF location. In contrast, we found robust decreases in FF prior to saccades both to RF and non-RF stimuli, and these decreases predicted saccadic reaction times for saccades to all stimuli. Mean firing rate also predicted reaction time, but only for saccades directed to the RF stimuli. For saccades directed away from the RF, no mean firing rate change was observed yet FF still predicted saccadic reaction time. These results demonstrate that response variability of area V4 neurons, like mean response rate, provides a signature of saccade preparation. However, the two signatures reflect complementary aspects of that preparation.

2.3 Methods

2.3.1 Subjects

Two male monkeys (*Macaca mulatta*, 8–12 kg) were used in these experiments. All experimental procedures were in accordance with National Institutes of Health Guide for the Care and Use of Laboratory Animals, the Society for Neuroscience Guidelines and Policies. General surgical procedures have been described previously (Graziano et al., 1997).

2.3.2 Behavioral task

Monkeys performed a visually guided, delayed saccade task which was initiated by fixation to within 1.0° of the central fixation spot (Figure 2-1). Immediately following fixation, an oriented bar stimulus appeared in the RF of the neuron under study and remained there until the end of the trial. Following the onset of the RF stimulus, the monkey was required to maintain fixation for a fixed delay (0.5–1 s, for a given experiment) while it waited for the appearance of a saccade target (0.25° diam) at one of two locations distant from the RF. In 2/3 of the trials, the target appeared, the fixation spot was extinguished, and the monkey was rewarded for making a saccade to the target. In these conditions, the saccade target could appear either directly upward from the fixation spot (“up” condition) or in the opposite visual hemifield to the RF stimulus (“opposite” condition). In the remaining one-third of trials (“toward” condition), the saccade target did not appear. Instead, when the fixation spot was extinguished, the monkey was rewarded for saccades to the RF stimulus. All conditions were identical until the cue to saccade (disappearance of the fixation spot) and were randomly interleaved. During all behavioral trials, eye position was measured via the scleral search coil method, and digitized at 200 Hz for offline analysis. Trials in which the monkey broke fixation prematurely or made a saccade to an incorrect target were discarded.




Figure 2-1. The visually guided delayed saccade task. In the task, the monkey fixates a central dot while an oriented bar is displayed in the receptive field (RF; dashed circle) of a single V4 neuron. After a delay, the monkey is cued (by fixation spot offset) to make a saccade in 1 of 3 directions. On 2/3 of the trials, a target dot appears in 1 of 2 locations, conditions up (*left*) and opposite (*middle*), and the monkey is rewarded for making a saccade to that dot. If no target appears, the monkey is rewarded for executing a saccade to the RF stimulus (*right*).

2.3.3 Recording

The activity of single V4 neurons was recorded via glass-coated platinum-iridium electrodes lowered into the dorsal surface of the prelunate gyrus. Neural activity was sampled at 32 kHz, digitized and stored. The waveforms of single neurons were isolated by offline clustering (DataWave Technologies).

2.3.4 RF stimuli

RF stimuli were displayed on a 34 x 27 cm Sony video monitor that was driven by a Number Nine graphics board (640 x 480) at a 60 Hz, noninterlaced, refresh rate. The video display was positioned 57 cm in front of the monkey. Visual stimuli consisted of gray-, red-, green-, or blue-colored bars appearing at one of four orientations (0, 45, 90, or 135°), presented at the center of a V4 neuron's RF. The contrast of the oriented bars varied between 5 and 80%, and the sizes varied between 1.0 x 0.1 and 8.0 x 0.8°. In a single block of trials, the RF stimulus varied along only one of the four stimulus dimensions (color, orientation, contrast, or size). The fixation spot was a small (0.25° diameter) circle displayed at the center of the video display. The non-RF saccade

target stimulus used in some behavioral conditions was identical to the fixation spot but located peripherally ($>5.0^\circ$).

2.3.5 Data analyses

We distinguished between tuned and untuned neurons by performing an unpaired t-test between firing rates for trials of each stimulus identity (i.e., “red” or “green”) and trials of each other stimulus along the same dimension (size, contrast, color, or orientation). If any comparison was significant ($P < 0.05$), then the neuron was defined as tuned and the maximal stimulus was taken as “preferred,” whereas the minimal was “nonpreferred.” For trials corresponding to each neuron, each stimulus identity, and each saccade direction [10–20 trials (mean 15.6), hereafter defined as a “neuron-condition”], we used the median reaction time (RT) for that neuron-condition to determine the faster (“short RT”) and slower (“long RT”) trials. Thus the long RT and short RT trials are exactly controlled for the effects of stimulus identity, neuron identity, stimulus preference, and saccade direction. Trials with RTs equal to the median of the neuron-condition were randomly assigned to the short or long RT groups.

Fano factor (FF) was computed by calculating the variance divided by mean of the spike counts across trials for an 80-ms window centered on successive 1-ms time bins for each neuron-condition, i.e., those trials with the same recording site, visual stimulus, and saccade direction. For example, for a time bin centered at -45 ms relative to saccade onset, counts were made within the 80-ms window around that time point (-85 to -5 ms) on each of the 10–20 trials of the neuron-condition, and both the mean and the variance were computed on the resulting set of 10–20 numbers. Finally the variance was divided by the mean to yield the FF, and the population estimate was simply the average of the FF values from all neuron-conditions. Note that the FF measures across-trial variability (Churchland et al., 2006) as opposed to within-trial variability of spike times or interspike intervals (de Ruyter van Steveninck et al., 1997) or variability across neurons (Cohen et al., 2007). Windows with no spikes on any of the trials were excluded from FF calculations. Eighty milliseconds was chosen as a window size, prior to computing any statistics, after trying values between 5 and

150 ms and selecting visually for a window that yielded traces retaining salient features of those generated with shorter windows while smoothing the noise effectively. Mean firing rates were likewise computed with the same 80-ms window.

To determine whether mean firing rate or FF traces at a particular point deviated significantly from a baseline period, we performed Wilcoxon ranked sum tests on the difference between data at the point of interest and data from a set of baseline period time points chosen to include the entire delay period without overlap because each point contains data from an 80-ms window. For saccade aligned data, the delay period was -640 to -320 ms relative to saccade onset, so the selected data points composing the delay period were -600, -520, -440, and -360 ms.

To control for a possible effect of variable firing rates on FF, we employed a “mean-matching” procedure in which the population distribution of mean spike counts was equalized across time (see Fig. 4 of (Churchland et al., 2007)). The algorithm computed the mean spike counts for all neuron-conditions, where each neuron-condition consists of a complete set of trials, 10 –20 total, from a particular neuron, visual stimulus, and saccade direction. Each plotted dot in Figure 2-2G represents the mean and variance across the trials of one neuron-condition. The algorithm determined a common distribution of these mean spike counts that can be found at all time points. It then randomly eliminated neuron-conditions until this common distribution was achieved at each time point. Because individual trials were never deleted from within neuron-conditions, the relationship between the mean and variance of spike counts for any neuron-condition was never altered by this procedure; rather, a different selection of the neuron-conditions (i.e., variance/mean pairs) is taken at each time point to meet the common distribution. The elimination was independent at each time point. The algorithm discarded a minority of the data in each case, keeping 69% for the upward saccade condition, 63% for opposite saccades, and 53% for saccades toward the RF. The FF was then computed only on these remaining data. The process was repeated 10 times, and the results averaged to control for variation due to the randomness of the procedure. We performed this analysis using the “Variance Toolbox” for MATLAB provided by M. M. Churchland.

To assess the possible influence of microsaccades on the mean rate and variability of V4 responses, we performed control analyses in which trials containing microsaccades within relevant time windows were eliminated. Thus for analyses of presaccadic firing rates and FFs, we excluded trials with microsaccades occurring within 200 ms of saccade onset (0.6% of trials). Likewise, for analysis of RT effects around the time of cue onset, we excluded trials with microsaccades occurring within 200 ms of cue onset (2.4% of trials). Microsaccade detection was performed as in (Armstrong et al., 2006). Microsaccades were defined as eye movements that exceeded 0.1° amplitude and had maximum velocity $>10^\circ/\text{s}$ for ≥ 10 ms.

For comparison of two conditions (for example, long RT trials vs. short RT trials), we computed a Wilcoxon signed-rank test on the mean firing rate or FF values for all neurons under the first condition versus those for the second condition at a certain time point. For stimulus aligned responses, we used $t = 100$ ms poststimulus onset, approximately at the peak of responsiveness. For cue aligned, we used $t = 0$ ms (exactly at cue onset) and for saccade aligned, we used $t = -45$ ms (just prior to saccade onset without including any postsaccadic visual responses). Because a window of 80 ms was used for the computation of both mean firing rate and FF, the values at these points include spikes from 40 ms on either side of the point. For comparisons in which many time points were examined to determine the time course of an event, the Simes procedure was used to control the false discovery rate (Benjamini and Hochberg, 1995).

To analyze differences in the magnitude of the presaccadic decline in FF between saccade directions, we performed an analysis of covariance (ANCOVA) on the change in FF over the final 80 ms of saccade preparation with the change in mean firing rate over the same time period as a covariate and the saccade direction as a factor. Thus for each neuron-condition, without mean-matching, we subtracted the FF and mean rate values at -80 ms relative to saccade from the values at the time of saccade onset. These two sets of numbers, ΔFF (dependant variable) and Δ mean firing rate (independent variable), were grouped according to saccade direction and analyzed with the

ANCOVA. The y intercept of the Δ FF versus Δ mean firing rate measures the component of the change in FF that is independent of change in mean firing rate.

2.4 Results

We computed the FF and the mean firing rate for 102 single neurons recorded in area V4 of two macaque monkeys ($n = 28$ neurons from one and $n = 74$ from the other) during the visually guided saccade task. Neurons were visually stimulated with single oriented bars that varied in orientation, color, size, or contrast. Figure 2-2 shows both the mean firing rate and mean FF changes in the population following onset of the RF stimulus, around the time of cue onset, and at the time of saccades to the RF stimulus or to non-RF targets. Stimulus-onset-aligned data from the most effective (preferred) and least effective (nonpreferred) RF stimuli are plotted separately; cue- and saccade-aligned data are divided according to the direction of the saccade. Overall, the sample of V4 neurons was highly selective for the RF stimuli employed, shown by the roughly twofold difference in mean firing rate following stimulus onset between the preferred and nonpreferred stimuli (Figure 2-2A). In contrast, the FF exhibited a marked decrement following stimulus onset (Figure 2-2B), and there was no significant difference in that decrement between the preferred and nonpreferred responses (Wilcoxon signed-rank test, $P = 0.234$; see METHODS for further details on statistical procedures). Instead the dynamics of the stimulus-driven FF changes were similar for the two stimulus divisions during both the initial onset transient and the sustained response in the delay period. The overall decrement in the FF following stimulus presentation is consistent with the stimulus-driven changes in variability reported across many other cortical areas (Churchland et al., 2010). Cue-aligned firing rate and FF are shown only to emphasize that at the time of cue onset the rewarded direction of saccade was unknown to the monkey, and therefore the overall mean firing rate and FF did not differ between the three saccade direction conditions (Figure 2-2, C and D).




Figure 2-2. Effects of RF stimulation and saccade preparation on the mean firing rate and response variability for the population of area V4 neurons. Left: mean firing rate (A) and Fano factor (FF; B) aligned to the time of RF stimulus onset and divided into responses to preferred vs. nonpreferred visual stimuli. These traces, as well as those in C–F, are all smoothed with an 80-ms box filter (see METHODS). C and D: mean firing rate and FF aligned to movement cue onset (i.e., fixation offset) and split by direction of saccade. E and F: the same but aligned to saccade onset. In all traces, means (dark lines) and SE (shading) are shown. In E and F, horizontal bars indicate significant difference from delay period. In E, translucent plots above traces show distributions of cue onset times relative to saccade onset. G: data from individual neuron-conditions for 2 time points: baseline and immediately prior to saccade onset. Each dot represents the mean and variance of spike counts within an 80-ms window for just 1 neuron-condition (those trials corresponding to a particular neuron, stimulus, and saccade direction). Black dots represent variance/mean pairs taken from windows during the baseline period of each saccade condition (1st arrow in F). Colored dots represent variance/mean pairs taken from windows just prior to saccade onset (2nd arrow in F). Thick lines are linear regressions on the data.

The mean firing rate changes we observed prior to saccade onset (Figure 2-2E) confirmed previous findings. Specifically, there was a significant increase in mean firing rate for saccades to the RF stimulus (toward condition; Wilcoxon rank sum test, $P < 0.001$). However, there was no change in mean firing rate for saccades to the opposite hemifield or upward saccade target locations (opposite and up conditions; Wilcoxon rank sum test, $P = 0.155$ and $P = 0.069$, respectively) (Fischer and Boch, 1981; Moore and Chang, 2009; Moore et al., 1998). In contrast to the mean firing rate effects, the FF decreased significantly for all saccade directions when compared with

its value during the delay period (Wilcoxon rank sum test, $P < 0.001$ for each direction; Figure 2-2F). The decrement in the FF was present within the final 100 ms of saccade preparation for each of the three saccade directions, shown by the decreased slope in the presaccadic variance/mean relationship relative to baseline (Figure 2-2G). Saccades to the RF stimulus generally had much longer RTs than saccades to non-RF targets (RTs, mean \pm SD: toward = 224 ± 50 ms; up = 115 ± 28 ms; opposite = 115 ± 24 ms; toward vs. up, $P < 0.001$; toward vs. opposite, $P < 0.001$). The larger RTs of the saccades to the RF stimuli is presumably due to the lack of an abrupt onset of the target (i.e., the RF stimulus) in this condition in contrast to the other two conditions (Yantis and Jonides, 1984). Nonetheless the pattern of presaccadic FF decline was largely similar to the other saccade conditions.

2.4.1 Mean-matched control for presaccadic firing rate changes

Neural firing patterns are commonly approximated as Poisson processes for which the variance of spike counts across trials is equal to the mean and thus FF is unity. However, this assumption may be violated and FF may decrease for extraneous reasons, for example due to an increasing influence of the refractory period at high firing rates. Although average firing rates were low (less than ~ 40 Hz), and thus the refractory period is unlikely to have a large impact on spike train variability (Mitchell et al., 2007), we nonetheless performed an analysis to control for the influence of firing rate dynamics on the FF (Figure 2-3, *A* and *B*). In this analysis, neuron-conditions (sets of trials corresponding to each neuron and stimulus condition; see METHODS) were discarded randomly at each time point to equalize the distribution of mean firing rates across the entire presaccadic period. This was done separately using data for each saccade condition (toward, up, and opposite), and the FF was computed on the remaining data (see METHODS). Thus this procedure eliminated changes in mean firing rate preceding saccades. Nevertheless, the significant decline in FF prior to saccade onset persisted for all three saccade conditions (Wilcoxon signed rank, $P < 0.001$ for toward condition; $P < 0.05$ for up and opposite) even in this mean-matched data set. Thus the observed presaccadic decreases in FF were not due to the changes in mean firing rate.




Figure 2-3. Presaccadic changes in FF for “mean-matched” conditions. The mean-matching algorithm was applied to presaccadic spike trains from the population of recorded V4 neurons to equalize firing rate distributions across time for each of the saccade directions. *A*: mean-matched firing rates for each of the saccade directions (toward, *left*; up, *middle*; opposite, *right*), which no longer vary over time. *B*: FF of the mean-matched data, which still declines presaccadically despite removing variation in firing rate. *C*: the magnitude of FF decline in the final 80-ms period before the saccade for each of the 3 saccade directions. The FF decline plotted corresponds to the component of the FF decline independent of the presaccadic change in mean firing rate, computed in an ANCOVA. Error bars represent 95% confidence intervals from ANCOVA.

2.4.2 Dependence of FF changes on saccade direction

We compared the magnitude of the presaccadic decline in FF between the three saccade directions during the final 80-ms period prior to the saccade onset. We used an ANCOVA to factor out the effect of presaccadic changes in mean firing rate. We found main effects of saccade direction ($P = 0.016$) and mean firing rate ($P = 10^{-4}$) on the magnitude of FF decline (Figure 2-3C). The latter effect demonstrates that firing rate indeed influences the presaccadic change in FF. The main effect of saccade direction, however, demonstrates that the FF declines with different magnitude for the different saccade directions and that this difference is independent of changes in mean firing rate. The overall decline for all saccade directions ($P < 0.01$ for all directions) corroborates the results of the mean-matching analysis in that it confirms a presaccadic decline in FF that is independent of changing firing rates. However, the FF decline was greatest for saccades directed toward the RF compared with the up and

opposite conditions. Thus in addition to a robust overall decline in FF for all saccade directions, we observed a component of that decline that depended on saccade direction.

2.4.3 Predicting Saccadic RT

A recent study found that across trial firing rate variability provides a better predictor of motor preparation than does the mean firing rate (Churchland et al., 2006). We sought to determine whether the variability of V4 responses, measured by FF, might reflect the state of saccade preparation. To do this, we examined the extent to which the FF was predictive of saccadic RT. We divided the trials obtained from all saccade directions and all RF stimuli into two subsets, long and short RT trials, with equal numbers of all conditions in each subset. We then recomputed mean firing rate and FF on these new trial divisions. We reasoned that if either FF or mean firing rate reflects the state of saccade preparation, then we should observe differences in these measures between short and long RT saccades at the time of the movement cue. Because the analysis window was 80 ms in duration, it included spikes occurring from 40 ms prior to the movement cue onset to 40 ms after. V4 neurons have visual onset latencies of ~50 ms (Maunsell, 1987), and in our data closer to 70 ms (Figure 2-2A). Thus the analysis window includes only the activity of neurons prior to any measurable responses to the movement cue (fixation offset) or target onset.

Despite the lack of differential visual stimulation at the time of cue onset, the FF of V4 neurons was significantly different between long and short RT trials, although mean firing rate was not (Figure 2-4). We computed the mean firing rate and FF around the time of movement cue onset separately for trials corresponding to each RT group and saccade direction and depict these data as percent changes from short to long RT trials, plotted for each saccade direction separately (Figure 2-4A). A two-way repeated-measures ANOVA revealed no main effect of either RT or direction on mean firing rate at exactly the time of cue onset, although there was an interaction between the two (Figure 2-4B, $P < 0.001$). Considering only those saccades directed toward the RF, there was a difference in mean firing rate between short and long RT trials at the time of cue onset with 8.6% larger mean firing rate for short RT trials ($P < 0.001$).

Mean firing rate did not differ significantly between RT groups for saccades to other locations, although the trend was toward a suppression of mean firing rate for saccades to the opposite hemifield on short RT trials relative to long (2.4% lower mean firing rate for short RT trials; $P = 0.16$).




Figure 2-4. Relationship of presaccadic mean firing rate and FF to saccadic RT for the population of V4 neurons. *A, left:* traces show percent difference in mean firing rate between short and long reaction time (RT) trials for each saccade condition. *Right:* percent differences in FF. *B:* differences in mean firing rate and FF for short and long RT trials in each saccade condition at the time of the movement cue ($t = 0$). *C:* same data as in *A* but collapsed across the 3 saccade conditions. Horizontal bar indicates a significant difference between long and short RT traces.

We observed a main effect of RT on FF, with short RT saccades having lower FF across directions (Figure 2-4A, *right*; $P = 0.003$). However, there was no main effect of saccade direction ($P = 0.66$) and no interaction between RT and saccade direction ($P = 0.59$). Thus for saccades toward the RF, both mean firing rate and FF predicted RT. However, for saccades directed to the upward and opposite saccade targets, the FF predicted saccadic RT even though there was no change in mean firing rate. Due to the interaction between saccade direction and RT for mean firing rate, collapsing the data across saccade direction largely eliminated the difference between short and long RT trials ($P = 0.15$). In contrast, collapsing the FF data across saccade directions yielded a robust difference between short and long RT trials ($P = 0.006$). FF was significantly lower for short RT trials than long from -35 to 148 ms relative to cue onset ($P < 0.014$; Figure 2-4C). The difference between FF of short and long RT trials (4.5%) is similar in magnitude to the effects reported in the study of premotor cortical neurons in a reaching task (5%) (Churchland et al., 2006). Our results demonstrate that in contrast to the presaccadic decline in FF, a component of which depended on saccade direction, the relationship between FF and RT at the time of cue onset was independent of saccade direction.

2.4.4 Possible influence of microsaccades

Because it is known that fixational saccades (i.e., microsaccades) can affect the firing rates of V4 neurons (Leopold and Logothetis, 1998), we considered their possible influence on the rate and variability of V4 activity in this study. For example, because the rate of microsaccades necessarily (and empirically) decreases in the time leading up to a saccade, this decline in the incidences of microsaccades might have contributed to the decline in FF (Figure 2-2D). To control for any influence of microsaccades, we discarded all of the trials in which a microsaccade occurred within the time window of interest and re-performed the analyses described in the preceding text. There were no differences in the primary effects in this reduced data set compared with the data set in which all trials with microsaccades were included. In particular, the presaccadic decline in FF remained significant for all saccade directions (Wilcoxon rank sum test, $P < 0.001$). The magnitude of the presaccadic decline still

depended on saccade direction ($P < 0.018$). The mean firing rate still predicted saccadic RTs for saccades toward the RF but not for other directions ($P < 0.001$), and the FF predicted saccadic RTs for all directions ($P = 0.05$).

2.5 Discussion

We measured the mean and variability of firing rates across trials of spike trains recorded from area V4 neurons during visually guided saccades. As expected, the mean firing rate of V4 neurons was enhanced when saccades were prepared to stimuli within a neuron's RF in comparison with saccades to a non-RF location. In contrast, we found robust decreases in FF prior to saccades both to RF and non-RF stimuli with only a small influence of saccade direction on the magnitude of the FF decrease. These FF decreases predicted saccadic RTs for all saccade directions. Although mean firing rate also predicted RT, this effect depended on saccades being directed to the RF stimuli. These results demonstrate that mean firing rate and FF exhibit different and complementary signatures of saccade preparation in area V4: while mean firing rate conveys more information about the direction of an imminent saccade, FF primarily reflects the progress of saccade preparation.

The way in which saccades toward the RF were cued differed from that of saccades directed to up or opposite targets. Specifically, in the latter case, the appearance of a saccade target indicated the location of the rewarded saccade, whereas in the former case, the absence of such a target indicated that the rewarded saccade was to the RF stimulus. It could be argued that this unbalanced task design could confound the interpretation. For example, the presaccadic enhancement of mean firing rate for toward saccades might be explained by a difference in the cueing method or by increased RTs in that condition. A previous report has shown that the presaccadic enhancement of mean firing rate is independent of the cue (Moore and Chang, 2009), and thus these mean firing rate changes are not due to the differences in cueing method. The novel changes in FF that we report were also independent of cueing method. Specifically, the ability of FF to predict RT and the presaccadic decline in FF were largely independent of saccade direction and thus cannot be explained by the task design.

FF is a measure of variability normalized to the mean, so it reflects changes in variance relative to concurrent changes in mean firing rate. Nevertheless the FF may vary due to indirect effects caused by changes in mean rate that do not reflect true changes in variability. For example, spike trains may be regularized by an increasing influence of the refractory period at high mean firing rates, such as those observed in our task prior to saccades toward RF stimuli. In some conditions, such indirect effects could not possibly account for the dynamics in the FF. For example, although saccades in different directions were preceded by either enhanced or unchanged mean firing rates, the FF decreased for all directions uniformly (Figure 2-2, *C* and *D*). Thus differences in the dynamics of mean firing rate across conditions do not necessarily result in the same direction of differences in the dynamics of FF. In addition, we controlled for the effect of changes in mean rate on FF by matching the mean across-trial firing rate distributions across time. The effect of this manipulation is to produce a subset of the data that has stable mean firing rate across time. We found that the FF decrease during saccade preparation was still present in the mean-matched data, indicating that the dynamics of the FF response were independent of changes in mean firing rate. The dissociation of responses of these two measures of neural activity demonstrates that they represent different information about the state of the visuosaccadic network.

Firing rates of single neurons predict behavioral RTs in many frontal and parietal cortical regions, such as motor and premotor cortex (Riehle and Requin, 1993), the parietal reach region (Snyder et al., 2006), the frontal eye fields (Hanes and Schall, 1996), and the lateral intraparietal area (Ipata et al., 2006). These correlations may reflect the role of these individual neurons in generating motor behaviors such as arm and eye movements. Recent studies have also shown that several measures of visual cortical activity predict RT, such as LFP in striate and extrastriate cortex (Zhang et al., 2008), spike-field coherence in area V4 (Womelsdorf et al., 2006), and multiunit activity in area V1 (Supèr and Lamme, 2007) as well as single-unit activity in areas MT and VIP (Cook and Maunsell, 2002). Our results are thus consistent with a growing body of evidence that neural activity in visual cortex can predict RT. Our results also demonstrate that the FF of V4 responses provides a reliable prediction of

RT in that it predicted RTs of saccades in all tested directions rather than simply those that target the neuron's RF. Taken together, these findings argue for a more integrated view of the role of visual cortical areas in visually guided behavior, a view that could take advantage of the myriad signatures that predict that behavior.

FF has been interpreted as reflecting the true underlying variability of neuronal firing rate across trials (Churchland et al., 2006). In this view, every spike train recorded from a neuron is a noisy instantiation of some "true" firing rate for that trial. This true firing rate may itself be variable across trials so that the recorded spike trains are in fact noisy realizations of a different true firing rate on each trial. While averaging the firing rate eliminates both sources of variability, FF instead estimates the extent of the underlying true variability with the assumption that spiking noise is invariant. With this context, we can interpret our data in much the same way as did Churchland et al. (2006). The decreased FF for short RT trials relative to long at the time of the cue to move reflects less variability in underlying firing rate, i.e., that more of the trials had the same true firing rate at cue onset for the short RT condition than the long. The precise value of this true firing rate may depend on the particular task, such as in our results in which neurons exhibited higher firing rates for saccades toward the RF and were on average unresponsive to saccades away from the RF. Nevertheless variability decreased in all three conditions, so the FF provides an index of the state of saccade preparation.

Importantly, FF revealed a signature of saccade preparation in the responses of area V4 neurons even when there was no change in mean firing rate. Traditionally, a neuron without changes in mean firing rate would be viewed as nonmodulated and its activity as uninformative during these conditions. Our results indicate that such a view is inaccurate. Even though a neuron may not be modulated in terms of its mean firing rate, a measure of the firing rate distribution may reveal that the activity of such a neuron is indeed modulated. Our results show that such modulation present during saccade preparation occurs to such a degree that the activity predicts saccadic RTs. FF therefore provides a sensitive measure of the influence of saccade preparation on V4 activity that is complimentary with mean firing rate, revealing that neuronal responses

are influenced by saccade preparation even when mean firing rate is neither enhanced nor suppressed.

Nonetheless our results do not in any way undermine the important role that firing rate likely plays in determining how neurons drive behavior. On the contrary, likely it is not the variability per se but rather the particular firing rates on individual trials, as indexed by the FF, which relates to the state of saccade preparation. We assume that the proximity of the firing rate on individual trials to some optimal mean firing rate relates directly to motor preparation (Churchland et al., 2006), an assumption consistent with our result that groups of trials with shorter RTs tend to have firing rates closer to the mean (i.e., lower FF). This view is depicted schematically in Figure 2-5, where the difference between baseline and presaccadic periods, as well as the difference between short and long RT trials, can be understood as a narrowing of the width of the firing rate distribution across trials. Due to a relatively small number of trials for each neuron-condition, and the noisiness of the firing rate measure, these firing rate distributions cannot be visualized directly, but are instead estimated from both the mean and variance of the measured single trial spike counts.




Figure 2-5. Cartoon model of changes in firing rate distributions during saccade preparation. Each firing rate distribution depicted corresponds to the hypothetical probability distribution of firing rates for a given neuron and stimulus. The distributions depicted, which cannot be directly visualized from the data are instead estimates based on the measured means and FFs, with the assumption of a Poisson distribution. A, The baseline firing rate distribution has a comparatively wide spread, reflected in its higher FF. B, Distributions during the presaccadic period move to separate means depending on saccade direction but become narrower, and thus have lower FFs, for all directions. C, At the time of cue onset, firing rates may be within any of the three presaccadic distributions due to advance planning or directed spatial attention (see Discussion). On those trials with the correct plan, firing rates are already in the correct presaccadic distribution and reaction times tend to be short (green). For trials with either of the incorrect plans, firing rates are in either of the other two presaccadic distributions (summed and divided by two) and RTs tend to be longer (red). “Optimal” firing rates are shown as dotted lines corresponding to the presaccadic mean firing rate appropriate for a given condition (note alignment with distribution means in B). Short RT distributions in each case have smaller standard deviations (σ) than long RT distributions while the differences between the means (μ) of the short and long RT distributions in this model correspond qualitatively to our observed results (Figure 2-4A).

There may be some relationship between the presaccadic modulation of variability reported here and modulation of the gamma frequency (30 –70 Hz) spectral power of visual cortical responses. Gamma modulation has been demonstrated to occur presaccadically at least for microsaccades (Bosman et al., 2009), and coherence effects in the gamma range are predictive of behavioral RTs (Womelsdorf et al., 2006) on some tasks, as is the FF reported here. Despite these similarities, it should be re-emphasized that we have measured variability across trials rather than variability in the spike times within trials, which would be most directly related to oscillatory processes. Future studies might address the potential relationship between observed decline in across-trial variability and frequency domain properties of neural activity.

The predictive activity of the mean firing rate for some, and FF for all, saccade directions can also be interpreted in the context of the influence of attention on saccadic RT (Kustov and Robinson, 1996). Because the interval between fixation onset and cue onset during a particular experiment was fixed, monkeys might have anticipated the impending saccade and directed spatial attention accordingly prior to the cue to move. In fact, increased anticipation of a behaviorally relevant stimulus does increase the magnitude of attentional modulation of the firing of area V4 neurons (Ghose and Maunsell, 2002). Thus for example, on some trials the monkey may have anticipated the cue and attended to the RF stimulus, which could have resulted both in reduced FF (Mitchell et al., 2007) and “short” RTs on those trials for which that stimulus became the saccade target (Posner et al., 1980). Likewise, higher FFs and “long” RTs may have resulted from allocation of attention to incorrect target locations or lack of attentional allocation altogether.

Our results do not allow us to determine whether the two measures (mean firing rate and FF) are signatures solely of attentional deployment or saccade preparation. However, given the preponderance of evidence that the effects of attention and saccade preparation on V4 neurons are very similar, if not identical (Moore et al., 2003), it is unclear to what extent such a distinction is possible in this area. However, our results cannot be explained solely by the known influences of covert spatial attention on variability (Mitchell et al., 2007). Because we observed a robust decline

in FF even when the monkey directed saccades, and thus spatial attention (Deubel and Schneider, 1996; Hoffman and Subramaniam, 1995), away from the neuron's RF, any known influence of covert spatial attention on FF must have been combined with some other influence that is independent of the saccade direction. For example, there may be a saccade-direction-independent influence of attention, perhaps merely related to the disengagement of fixation prior to saccades of any direction. On the other hand, such a nonspatial influence need not be directly related to the preparation of the eye movement per se. A number of studies have observed neural correlates of other spatially nonselective factors such as stimulus and reward expectation as well as elapsed time (e.g., Ghose and Maunsell, 2002; Janssen and Shadlen, 2005). Moreover, although the magnitude of presaccadic decline in FF depended on saccade direction, there was a substantial decline for all saccade directions. Thus a more global influence, for example arousal or reward anticipation, could be considered to explain the nonspatial component of the effects. Indeed like attention, these other influences may be associated with saccade preparation but may not require an actual movement to produce the dynamics we observe. Nonetheless the FF predicts saccadic RTs for saccades in all tested directions and thus provides a reliable signature of saccade preparation.

3 Interdependence of saccadic and attentional modulation of visual cortical signals

3.1 Abstract

We examined whether the preparation of saccadic eye movements, when behaviorally dissociated from covert attention, modulates activity within visual cortex. We measured single-neuron and local field potential (LFP) responses to visual stimuli in area V4 while monkeys covertly attended to a stimulus at one location while preparing saccades to a target at another. In spite of the irrelevance of visual information at the saccade target to behavioral performance, visual activity at the saccade target location was modulated at least as much as activity at the attended location. Modulations of activity at the attended and saccade target locations were qualitatively similar, and included increased magnitude, selectivity, and reliability of spiking activity, as well as increased gamma and decreased beta power of LFPs. These results demonstrate the sufficiency of saccade preparation in modulating visual cortical representations and suggest that the interdependence of oculomotor and attentional mechanisms extends to posterior visual cortex.

3.2 Introduction

In order to efficiently interpret the sensory world, many species have evolved powerful orienting systems to select among multiple objects or features for enhanced processing. In primate vision, orienting involves shifting gaze in order to position the foveae on targets of interest, and this behavior requires using the visual parameters of the target (e.g. position, velocity, and shape) to guide gaze shifts. Each orienting movement thus necessarily involves the selection of one stimulus over all others prior to movement onset. Psychophysical studies in human subjects indicate that this selection is accompanied by attention, that is, enhanced detection and discrimination at the location of intended movements (Hoffman and Subramaniam, 1995). Furthermore, this deployment of attention can occur whether or not orienting movements are actually carried out (Posner, 1980); that is, selective attention can be either overt or covert.

Given the co-occurrence of gaze and attentional shifts, investigators have long debated the dissociability of the mechanisms underlying these two functions (e.g. review). For example, while some studies have found that the preparation of saccadic eye movements (saccades) to a particular location is sufficient to improve psychophysical performance at that location, and therefore to direct attention, (Deubel and Schneider, 1996), others have found that saccade preparation can be dissociated from attention (Hunt and Kingstone, 2003). Thus, at present, the degree to which saccade preparation can be accomplished without corresponding changes in perceptual enhancement remains unresolved. In particular, it is unclear whether or not saccade preparation is sufficient to modulate the representations within the visual system, modulation thought to underlie the perceptual enhancements of selective attention (e.g. Reynolds and Chelazzi, 2004).

A number of neurophysiological experiments have provided evidence suggesting that certain brain structures have a role in both overt and covert attention. For example, the frontal eye field (FEF), the lateral intraparietal area (LIP), and the superior colliculus (SC) appear to be involved both in saccade programming and in directing visual spatial attention. Several studies have shown that neural activity in

these regions is modulated prior to saccades (FEF: Bruce and Goldberg, 1985; LIP: Barash et al., 1991; SC: Schiller and Stryker, 1972) as well as during covert spatial attention (FEF: Thompson et al., 2005; LIP: Bushnell et al., 1981; SC: Ignashchenkova et al., 2004). Furthermore, pharmacological inactivation of neurons in these areas affects saccades (FEF: Dias and Segraves, 1999; LIP: Liu et al., 2010; SC: Hikosaka and Wurtz, 1986) and covert attention (FEF: Wardak et al. 2006; LIP: Wardak et al., 2004; SC: Lovejoy and Krauzlis, 2010). However, more recent studies suggest that at the level of single neurons, saccades and attention are nevertheless dissociable. For example, within the FEF, only neurons functional classified as ‘visual’ or ‘visuomovement’ exhibit enhanced sensory responses at attended locations while ‘movement’ neurons do not (Thompson et al., 2005). More importantly, it is only the enhanced responses of ‘visual’ neurons that become more synchronized with activity within posterior visual cortex, suggesting that visual cortex receives only attention-related, but not saccade-related, signals (Gregoriou et al., 2012). However, the prediction that saccade preparation alone is not sufficient to modulate visual cortical responses has not been tested.

We sought to determine whether covert and overt attention are dissociable within visual cortex. We trained monkeys to perform a task in which the target of an upcoming saccade was behaviorally dissociated from the location of relevant visual information. In addition to modulation of the visual responses to covertly attended stimuli, we also found modulation of visual cortical responses to potential targets of saccades. The modulation during saccade preparation was qualitatively similar to modulation by covert attention, including increases of firing rates, stimulus selectivity, across-trial spiking response reliability, and gamma local field potential (LFP) power, as well as decreases in low frequency LFP power. Our results demonstrate that saccade preparation is sufficient to modulate responses in visual cortex.

3.3 Methods

3.3.1 Subjects

Two male monkeys (*Macaca mulatta*, 8–12 kg) were used in these experiments. All experimental procedures were in accordance with National Institutes of Health Guide for the Care and Use of Laboratory Animals, the Society for Neuroscience Guidelines and Policies. General surgical procedures have been described previously (Graziano et al., 1997).

3.3.2 Behavioral task and visual stimuli

We trained two monkeys on a cued change-detection with change-blindness manipulation and antisaccade report task. In brief, the monkey was required to make a difficult visual discrimination at a peripheral location, made easier by a quasi-symbolic cue indicating which location would contain the change, but made more attentionally demanding by the simultaneous disappearance and reappearance of all peripheral stimuli (change-blindness). The monkey was rewarded for reporting a successful detection with a saccade to the diametrically opposite peripheral location (antisaccade response).

The sequence of trial events for most trials was as follows. All time ranges are uniformly distributed and independently chosen unless otherwise stated. A small white dot ($\sim 0.3^\circ$ diameter) appeared on the screen and the monkey initiated a trial by fixating it. Within 100ms, the four peripheral target stimuli appeared (described below). After a brief delay of 300-500ms, the cue appeared: a white line less than half a degree in length and one pixel ($< 0.1^\circ$) in width, originating at the fixation dot and extending in the direction of one of the four stimuli (randomly, independently chosen on each trial with equal probability). The cue indicated with 90 or 93% validity which of the four stimuli would change on this trial (if any). After a post-cue period of 600-2200ms with the display static as described, the four peripheral stimuli synchronously disappeared for a brief (< 270 ms) interval ("blank period"), and then reappeared. Upon reappearance, one of the four stimuli changed its orientation (i.e. was rotated in place) on 50% of trials. On these trials ("change trials"), the monkey could earn a reward by

executing a saccadic eye movement within 800 ms to the stimulus opposite the changed stimulus. On the other 50% of trials (“catch trials”), all four stimuli appeared at identical orientations to those they had before disappearing; in this case, the monkey was rewarded for maintaining fixation on the central dot for 800 ms.

As there were four possible saccade targets as well as the option to make no saccade, chance performance on the task (without making use of the cue or the changing stimulus) was 20% correct. Alternatively, a strategy in which the monkey used the cue to choose the direction of saccade but did not use the visual stimulus information to decide whether to saccade or to continue fixating would result in a chance performance of 50% (on validly cued trials).

The target stimuli were four static gabor patches, i.e. oriented black and white gratings in a circular gaussian aperture. In Monkey G the gratings were square wave; in Monkey B they were sine wave modulated (both types elicited robust responses from the neurons in this study). In both monkeys the gratings were at maximal contrast for the monitor, i.e. the maximum was the brightest white available and the minimum pixels were turned off. The dimensions of the gratings varied somewhat over the course of the experiments but were typically ~ 4 degrees in diameter and ~ 1 cycles/degree in spatial frequency. The location of the gratings varied from session to session depending on the receptive field locations of the neurons being recorded, but the centers were always between 5 and 8 degrees eccentricity. All four gratings had equal eccentricity and were spaced evenly (i.e. at 90 degree intervals around a circle). The screen background was dark gray in Monkey G and middle gray in Monkey B, but in neither monkey were the mean luminances of the gratings matched to the background color. The orientation of the grating took one of 16 possible values, evenly spaced from 0 to 360 degrees in 22.5 degree intervals (note that orientations 180 degrees apart (e.g. 45 and 225 degrees) were identical except for a mirror reflection, and for nearly all neurons drove the neurons identically and have therefore been combined for analysis). The grating orientations were chosen independently for each of the four stimuli and for each trial. The amount of rotation took multiple values to vary the difficulty of the task, but was typically 45, 67, or 90 degrees, and trials

with these different rotation magnitudes were interleaved randomly. For trials with no blank period, the rotation magnitude was typically 2, 5, or 8 degrees. The rotation was clockwise or counterclockwise with equal probability, independently chosen for each trial.

3.3.3 Behavioral Analysis

3.3.3.1 Reaction Times

Reaction times were computed on correctly performed trials as the time between the stimulus re-onset and the initiation of the saccade. This analysis therefore only considered trials when one of the stimuli changed orientation, since correctly performed trials without an orientation change did not result in any saccadic eye movement.

3.3.3.2 Signal detection theory (SDT) model to dissociate sensitivity and bias

We analyzed each monkey's sensitivity and bias to detecting stimulus orientation changes with the 4-alternative detection task (4-ADC) model (Sridharan et al., 2013). In brief, the model assumes that the monkey makes a decision on the basis of "decision variables" which represents the information he has about each stimulus. When there is no change at a certain stimulus, the value of the decision variable is drawn from a Gaussian distribution with zero mean and unit variance. By contrast, when the orientation of stimulus does change, the decision variable is drawn from a Gaussian with unit variance but with a positive mean, called the sensitivity (d'). If the decision variable on any trial exceeds a threshold (called the criterion or bias, c) then the monkey will report a detection at that location. On each trial, the decision variable for each of the four stimuli is chosen independently, and in cases in which multiple decision variables exceed their criteria, the chosen response will be the one that exceeds the criterion by the greatest amount. The model therefore has only eight parameters (sensitivity and bias for each of the four stimulus locations), though there are 25 behavioral measurements with which to fit the model (probability to make each of the five possible responses (saccades to any one of the four stimuli to report

detections or continued fixation to report no detection) on each of the five types of trials (changes at each of the four stimuli or no change); only 20 of these measurements are independent). This conceptually simple model can be visualized graphically (Figure 3-1). It was fit with maximum likelihood estimation (by S. Deverajan), and these best-fit parameters are unique (Sridharan et al., 2013).




Figure 3-1. Graphical depiction of the m-ADC model for the two-alternative case. X- and Y-axes represent the decision variables at each of the two stimulus locations (ψ_1 and ψ_2). The black, red, and blue circles represent the distribution of decision variable values on trials with no changes, changes at location 1, and changes at location 2, respectively. The black, red, and blue regions of the decision variable space represent values of the decision variables for which the subject will report detecting no change ($Y=0$), change at location 1 ($Y=1$), or change at location 2 ($Y=2$), respectively. d_1 and d_2 , and c_1 and c_2 , represent the sensitivity and bias to changes at locations 1 and 2, respectively, respectively.

3.3.4 Neural recordings

3.3.4.1 Linear array recordings

Recordings were made with 16-channel U-Probes (Plexon Inc., Dallas, TX). These electrodes are cylindrical in shape (180 μ m diameter) and have a row of 16 circular platinum/iridium electrical contacts (15 μ m diameter) with 150 μ m spacing (total length of array is 2.25mm).

Data were amplified and recorded using the Omniplex system (Plexon Inc., Dallas, TX). Wide-band data, filtered only in hardware at 0.5Hz highpass and 8kHz lowpass,

were recorded to disk at 40kHz. Spikes were detected from this signal as described below.

3.3.4.2 Spike detection and sorting

We were unable to "isolate" the waveforms of single neurons using the traditional method, by adjusting the position of the electrode carefully throughout the recording to ensure that its recording surface remains as close to the neuron as possible, since any adjustment of the electrode position would alter the isolations all 16 contacts simultaneously. Instead, we set the electrodes in place and left them for the duration of the session, taking whatever units presented themselves there. The spikes we recorded therefore came in a wide range of isolation qualities, and so we took great pains to improve the quality of the data as much as possible in post-processing, using the steps described below, and to quantify the quality of these isolations.

3.3.4.3 Spike detection

The wide-band data was filtered with notch filters at multiples of 60Hz to remove line noise harmonics. "Common average referencing" (CAR) was applied (Ludwig et al., 2009) in order to remove other noise components appearing on all channels. CAR is performed by averaging the signal from all channels together and subtracting this average signal from each individual channel. This filtered and re-referenced signal was used for spike detection using the matched-filter method (Hill et al., 2011). First, the signal is convolved with a waveform representing the average, expected shape of usual cortical neurons. Specifically, the waveform used was biphasic with ~300 ms trough-to-peak duration. Next, a threshold is applied to the new filtered signal and peaks of sections of the signal which cross the threshold are determined. The threshold is chosen such that the rate of crossings is 100Hz. Finally, putative spike waveforms are pulled from the pre-convolution signal at the times of threshold crossing peaks. Spike waveforms were 1ms in duration (40 samples), with 0.2 ms prior to the peak location of the post-convolution signal. Waveforms within 500ms of each other were disallowed. Putative waveforms are selected in order of descending peak height, such that bigger peaks (i.e. waveforms of larger amplitude or more similar to the average

waveform used as in the convolution) have "priority" over smaller ones for cases in which two waveforms are less than the disallowed interval (500ms) from each other.

3.3.4.4 Spike sorting

Spike waveforms were “sorted” in the attempt to classify separately those waveforms originating from one neuron and those from others. Sorting was initially performed manually using Offline Sorter (Plexon) by identifying clusters of waveforms with similar shapes. In many cases, this initial sorting was refined by computing the Fisher Linear Discriminant between the clustered waveforms and all other waveforms on the same channel (Hill et al., 2011), projecting the waveforms along this dimension, and re-classifying waveforms according to their value on this axis. The result of these procedures was the determination of groups of similar waveforms, referred to as “units”. The extent to which the waveforms of any of these units could be confidently reported as originating from a single neuron was determined with further quality metrics (below).

3.3.4.5 Sorting quality quantification

Under the assumption that cortical neurons have a “refractory period,” or minimum time between spikes, we computed an estimation of the false-positive rate for waveforms of each cluster (Hill et al., 2011). This calculation considers the rate of spikes, the duration of the experiment, and the number of waveforms too close together in time to plausibly arise from a single neuron to arrive at a figure estimating what percentage of the total spike count arose from neuron(s) besides the one in question. If greater than 10% of spikes were probably due to contamination from other neurons, the unit was referred to and analyzed as a “multi-unit.” If the false-positive estimate was less than 10%, if the shape of the waveform appeared stable over the duration of the experiment, and if the histogram of waveform amplitudes was approximately symmetric (asymmetric histograms may indicate the failure to detect low-amplitude spikes from that neuron, resulting in only a fraction of the true spikes from the neuron being represented in the unit (Hill et al., 2011)), then the unit was declared a “single-unit,” indicating our confidence in the data from that unit reflecting the responses of just one individual neuron.

3.3.5 Quantification of firing rate modulation

Spikes were counted within the window between 500ms post-cue onset the end of the post-cue period (i.e., start of blank period) and converted to rates for each trial based on the duration of that period. Mean rates in each condition were compared by computing a modulation index, defined as:

$$MI = \frac{A - B}{A + B}$$

where A is the mean firing rate in the modulated condition (either cue-RF or cue-opposite) and B is the mean firing rate in the control condition (cue-orthogonal).

To statistically compare the spike counts on modulated versus unmodulated trials for each individual unit, the Wilcoxon ranksum test was computed between modulated and unmodulated rates. For the population as a whole, the Wilcoxon signrank test was computed between mean rates of modulated and unmodulated conditions on a unit-by-unit basis.

Units with average spike rate over the whole trial of less than 0.1Hz were excluded from this and all other analyses (21 of 717 excluded). Though we included some units which therefore emitted only 1 or 2 spikes per trial on average, visual inspection of these units revealed that many nevertheless showed clear tuning and/or modulation by condition, thanks to the large numbers of trials (>1000) per recording session.

For the purposes of correlating the firing rate modulation on cue-RF trials with modulation on cue-opposite trials, we used independent sets of cue-orthogonal trials as the reference for the computation of modulation index. We selected randomly half of the trials from each cue-orthogonal direction to serve as the reference for the cue-RF modulation and the other half to serve as the reference for the calculation of cue-opposite modulation.

3.3.6 Quantification of tuning amplitude

Spikes were counted within the window between 500ms post-cue and the start of the blank period, each trial's spike count was converted to a firing rate, and rates were

combined across trials that had the same stimulus orientation in the RF. Units were only considered for this analysis if they were significantly modulated by the stimulus orientation during the stimulus onset period ($p < 0.0001$ on Kruskal-Wallis test of spike rates grouped by receptive field stimulus orientation).

The tuning curves were fit to a von Mises distribution (circular gaussian). Just as with a standard gaussian, the von Mises distribution has two parameters, the mean (μ ; preferred direction of the unit) and standard deviation (κ ; tuning width). Two other parameters allow the tuning curves to be fully fit: a baseline offset (b ; added to the tuning curve) and a scaling factor (s ; multiplies the tuning curve). The fit equation for firing rate (r) as a function of stimulus orientation (θ) is given by:

$$r(\theta) = b + s * \frac{e^{\kappa * \cos(\theta - \mu)}}{2\pi I_0(\kappa)}$$

where " I_0 " is the modified zeroth-order Bessel function, computed with the Matlab function "besseli". We fit the equation with constrained least squares curve-fitting, with the mean parameter μ restricted to be in the range $[-\pi, \pi]$, width parameter κ to $[0, 8]$ and all other parameters $[0, \infty]$. The restriction on κ prevented tuning widths being narrower than ~ 45 degrees. Any tuning curves with true widths less than that could not be adequately measured with our sampling of orientations, so with κ unconstrained clearly artifactual fits resulted in some cases. The final value of tuning curve amplitude was computed by subtracting the trough from the peak value of the fit tuning curve, that is:

$$\text{tuning amplitude} = r(\mu) - r(\mu + \pi)$$

To assess whether tuning amplitude was significantly influenced by cue direction for individual units, we computed bootstrapped 95% confidence intervals (CIs) by randomly selecting trials, with replacement, and re-performing the above analyses to determine tuning amplitude for each cue direction on each of 1000 sets of resampled trials. 95% CIs were determined as the 25th and 975th largest tuning amplitudes from this distribution. If the 95% CIs for the modulation condition (either cue-RF or cue-

opposite) were not overlapping with the 95% CI for the cue-orthogonal condition, the difference in tuning amplitude was declared significant at $p=0.05$.

3.3.7 Quantification of across-trial spiking reliability (Fano factor, FF)

Spikes were counted in non-overlapping 50 ms bins (as in Churchland et al., 2010) during the final 400ms of the post-cue period, prior to the blank period. The FF was computed as variance divided by the mean of these spike counts for groups of identical trials (same RF stimulus orientation and cue direction) and was averaged across the groups corresponding to different stimulus orientations.

To assess significance of the difference in FF between conditions for an individual unit, we performed a shuffle test by randomly reassigning cue direction labels for each trial and re-computing FF for each cue direction. The true difference between modulation condition (cue-RF or cue-opposite) and cue-orthogonal was compared to the distribution of 1000 shuffled differences, and declared significant if it was greater or less than 97.5% of this distribution ($p=0.05$ significance level).

3.3.8 Quantification of power in local field potentials (LFP) within a frequency band

The LFP was defined as the continuous voltage signal highpass filtered at 0.5Hz and lowpass filtered at 250Hz, and downsampled to 1kHz. A second-order 60Hz notch filter was also applied. LFP segments were taken from the final 500ms of the post-cue period (prior to the blank period) and the fast fourier transform (FFT) was computed. FFTs were averaged across trials within each cue condition and, due to the large correlations from channel to channel, averaged across channels within each recording.

3.4 Results

Two monkeys (G and B) performed an attention-demanding task that required them to detect orientation changes in one of four peripheral gabor gratings while maintaining central fixation (Figure 3-2A; see Methods). During each trial, the identity of the relevant stimulus was indicated with a central cue. After a variable interval, the complete array of stimuli, the cued stimulus and all three distractors, disappeared for a brief moment and then reappeared. Monkeys were trained to detect changes in orientation of any of the four stimuli upon reappearance. In order to dissociate the locus of attention from that of saccade preparation, monkeys were rewarded for responding to an orientation change with a saccade to the stimulus diametrically opposite of the changed stimulus (antisaccade). The central cue validly indicated the relevant stimulus on a vast majority of trials (90-93%) and orientation changes occurred on only a random half of trials. On trials with no orientation change, monkeys were rewarded for maintaining central fixation.




Figure 3-2. Cued change detection and antisaccade task. A. Task design and trial sequence. In brief, monkeys fixated a white dot while four peripheral oriented-grating stimuli were presented. After a variable delay, stimuli disappeared then reappeared, either with or without one of the four stimuli rotating (change trial or catch trial, respectively). Monkeys could earn a reward by saccading to the diametrically opposite stimulus from the change on change trials, or by maintaining fixation on catch trials. A small, foveal, quasi-symbolic cue (white line) indicated which stimulus, if any, was most likely to

change. Dashed circle indicates area V4 receptive field (RF) locations, yellow spotlight indicates direction of attention, and arrow indicates saccade direction; all three were not visible to the monkey. All graphical elements are not precisely to scale; in particular, the cue is shown much larger than scale for visibility. See Experimental Procedures for further details. B. Task conditions. On cue-RF trials, the relevant visual information was in the RF of recorded neurons (spotlight) while rewarded saccades, on validly cued change trials, were to the opposite stimulus (dashed arrow), due to the antisaccade structure of the task. On cue-opposite trials, conversely, saccades were directed to the RF stimulus while relevant stimulus was opposite. On cue-orthogonal trials, neither relevant visual information nor the saccade target were in the RF.

We recorded activity from 268 single neurons, 428 multi-neuron clusters, and local field potentials (LFPs) at 736 sites in area V4 (see Methods) of the two monkeys while they performed the selective attention task. Monkey G completed 34,803 trials over 25 sessions, and Monkey B completed 33,853 trials over 21 sessions with simultaneous neural recordings. Only neural data from correctly performed trials were analyzed. The task had four conditions with respect to cue direction (Figure 3-2B). In the “cue-RF” condition, the cue directed attention to the receptive field (RF) stimulus. In the “cue-opposite” condition, the cue indicated that the RF stimulus would be the target of rewarded saccades on validly cued change trials. Finally, in the “cue-orthogonal” condition, the RF stimulus was 90° clockwise or 90° counter-clockwise from the cue direction, such that neither attention nor a saccade to the RF stimulus was likely to be required. The two cue-orthogonal conditions were identical in terms of the irrelevance of the RF stimulus to task performance and thus they were combined. Correlates of attention in the neural activity were measured as the difference between cue-RF and cue-orthogonal conditions, while correlates of saccade preparation were measured as the difference between cue-opposite and cue-orthogonal conditions.

3.4.1 Behavioral Performance

The attentionally-demanding “change blindness” manipulation (i.e. the blank period; O’Regan et al., 1999; Simons and Rensink, 2005) as well as smaller than maximal change magnitudes (amount of rotation of changing gratings) were employed to ensure the task did not become easy enough for the monkeys to perform without selective attention. Monkey G correctly responded on 69% of trials on average (77% on change trials, 62% on catch trials). Monkey B correctly responded on 67% of trials

(62% on change trials, 70% on catch trials). Only correctly performed trials were analyzed unless otherwise stated.

If monkeys utilized the cue to deploy covert attention or to prepare a particular saccade, we should expect behavioral evidence of this, in the reaction times and in the choices the monkeys made. Reaction times on the task were strongly influenced by the validity of the cue (Figure 3-3). Mean reaction times were more than 100ms faster when the cue validly indicated the location of the change than when the cue indicated the opposite stimulus. Put another way, the mean reaction times were more than 100ms faster when the prepared saccade was eventually executed than when the monkey instead had to perform a saccade toward the location he had been attending. Reaction times for invalid cues that indicated stimuli orthogonal to the changing stimulus resulted in intermediate reaction times, but still significantly slower than for validly cued changes.




Figure 3-3. Reaction time dependence on cue validity. Mean \pm S.E.M. for correctly performed trials when the cue was valid, invalid indicating the opposite location, and invalid indicating either orthogonal location.

The literature linking macaque visual cortical responses to behavioral performance has been uniformly unclear on the distinction between changes in performance

brought about by alterations in bias or in sensitivity (e.g. Cavanaugh and Wurtz, 2004; Zénon and Krauzlis, 2012). Nevertheless, the phenomenon of covert attention as understood in the psychophysical literature refers to enhancement of sensitivity, not bias. We therefore sought to explain the performance of the monkeys on the task with a behavioral model based on signal detection theory (SDT; See Section 3.3.3.2 and Sridharan et al., 2013). This model effectively dissociates sensitivity to orientation changes from bias to respond to those changes. We found that the largest effect of the cue direction on behavior for both monkeys was a reduction of the criteria, or bias, to respond to changes (Figure 3-4, lower row, cyan bars). That is, monkeys were much more inclined to make the saccade associated with the cued response even when the sensory evidence was weak; they were much more likely to make false alarms in the direction opposite the cue. However, we also observed, most strongly in monkey B, an increase in sensitivity at the cued location, demonstrating that indeed the monkeys used the cue to guide their covert attention. Surprisingly, the sensitivity was also enhanced at the saccade target location. Thus, behaviorally, preparing a saccade to a target is sufficient to bring about perceptual sensitivity enhancements – attention – even when the visual stimulus is not behaviorally any more relevant than the other distractors (orthogonal stimuli) for which no such enhancement was observed. Since the monkeys' criteria for changes at the opposite location was very high, the performance measured by percent correct was in fact substantially worse when changes occurred at that location, emphasizing that it is only by virtue of the behavioral model that we were able to uncover the attentional enhancement in perceptual sensitivity.




Figure 3-4. Behavioral analysis for monkeys performing the cued change detection task, performed with signal detection theory model. Sensitivity (top row) and bias (bottom row) were estimated for trials of each cue direction by fitting the 4-ADC model to the response probabilities of the monkeys. Only trials with changes at the RF (lower left) location were analyzed. Therefore, cue-RF data (cyan bars) were for validly cued changes, cue-opposite (red bars) were for invalidly cued changes when the change happened at the intended saccade location and saccades had to be executed towards the cued location. Analysis and figure by Sridhar Deverajan.

3.4.2 Firing rate modulation

Numerous studies have found that visually driven firing rates of neurons in V4 are enhanced during selective attention to RF stimuli (e.g. Moran and Desimone, 1985; Reynolds et al., 2000). Our task dissociated covert attention and saccade preparation such that we could separately measure the modulations due to both. We computed firing rates during the post-cue period of the task, averaged across trials of the same cue condition, for each single neuron or multi-unit cluster, and compared these average rates between the different cue conditions for the population of neurons (see Methods). We plotted the responses of three example single neurons on cue-RF trials against that on cue-orthogonal trials (Figure 3-5A). Beginning a few hundred milliseconds after cue onset, the responses diverge, with higher firing rates when the cue is directed to the RF stimulus, that is, when the monkey attends the RF stimulus. We also plotted the responses on cue-opposite trials against that on cue-orthogonal trials. Similar to the modulation during cue-RF trials, the responses of these three

neurons are robustly enhanced when the cue is directed to the stimulus opposite to the RF, that is, when saccades for validly cued change trials will be directed to the RF stimulus.




Figure 3-5. Responses of three example single neuron in the cued change-detection task. A, Divergence of spiking responses during the post-cue period for cue-RF (cyan) trials relative to cue-orthogonal (purple) trials. B, As in A but for cue-opposite (red) trials relative to cue-orthogonal.

Across the population, firing rates were significantly larger for cue-RF trials compared to cue-orthogonal ($p < 10^{-10}$, Wilcoxon signrank test; Figure 3-6B). The effect was also significant when considering data from each monkey individually (monkey G, $p < 0.01$; monkey B, $p < 10^{-13}$) and when considering only isolated single neuron responses ($p = 0.03$).




Figure 3-6. Effects of cue direction on post-cue period firing rate. A. Peristimulus time histogram of spiking activity around the time of cue onset averaged across 696 single- and multi-units in two monkeys. Trials were divided by cue direction: cue-RF (cyan); cue-opposite (red); and cue-orthogonal (purple). Individual unit PSTHs were normalized to the average firing rate and smoothed with 150ms box filter prior to computing the mean across units. Shaded regions indicate ± 1 standard error of the mean. B. Histogram of the effect of cueing the RF stimulus on firing rate for each unit. The effect is measured as a modulation index: the difference between mean rates in the cue-RF and the cue-orthogonal conditions divided by the sum. Rates were computed on each trial during the period from 500ms after cue onset until the start of the blank period. p-value shown for Wilcoxon signed rank test. The colored part of the histogram corresponds to units for which firing rates were individually significantly modulated by cue direction (see Experimental Procedures). C. As in B, but for cue-opposite condition compared to cue-orthogonal condition.

In addition to the firing rate increases during the cue-RF condition, we also found that, relative to the cue-orthogonal condition, firing rates increased during the cue-opposite trials, i.e. when the RF stimulus was likely to be the target of saccades on validly cued change trials ($p < 10^{-18}$). Similar to the effect during cue-RF trials, the increase on cue-opposite trials was also significant when considering only isolated single neurons ($p < .001$) and when considering data from each monkey separately (monkey G, $p < 10^{-9}$; monkey B, $p < 10^{-13}$). However, the overall increase in the cue-opposite condition was larger than the increase on cue-RF trials (median 4.8% vs. 2.9%, $p = .002$). Thus saccade preparation, as well as covert attention, correlates with enhanced firing rates of V4 neurons.

3.4.3 Orientation tuning

In addition to the enhancement of firing rates, past studies have described increased tuning amplitudes of V4 neurons during covert attention (McAdams and Maunsell, 1999). We examined the changes of the orientation tuning curves of V4 neurons during both attention and saccade preparation. Since the firing rate effects described above were similar for single neurons and multi-units, we combined the two datasets for this and further analyses. Of the neurons in our sample, 54% were well-tuned for orientation (378 of 696). We fit the responses of these neurons with circular Gaussian (von Mises) functions and quantified the difference between peak and trough (see Methods), referred to as tuning amplitude, separately for trials of each cue condition (Figure 3-7). We found a significant increase in tuning amplitude during the post-cue period in cue-RF trials relative to cue-orthogonal trials (Figure 3-7B; $p < 10^{-5}$). In addition, we found an increase in the same period for cue-opposite compared to cue-orthogonal trials ($p < 10^{-7}$). Thus the amplitudes of V4 orientation tuning functions were increased both during attention and saccade preparation.




Figure 3-7. Effects of cue direction on tuning amplitude. **A.** Tuning curves for four example V4 neurons. Firing rate during the post-cue period is averaged across groups of trials with identical RF stimuli and cue-direction, then plotted against the stimulus orientation for each cue condition. Tuning amplitude modulation indices for cue-RF versus cue-orthogonal (cyan text) and for cue-opposite versus cue-orthogonal (red text) are shown for each neuron. **B.** Histogram of the effect of cueing the RF stimulus on tuning amplitude for each unit. The effect is measured as a modulation index: the difference between tuning amplitudes in the cue-RF and the cue-orthogonal conditions divided by the sum. p -value shown

for Wilcoxon signed rank test. The colored part of histogram corresponds to units for which tuning amplitude was individually significantly modulated by cue direction. C. As in B, but for cue-opposite condition compared to cue-orthogonal condition.

3.4.4 Response reliability

The correlates of covert attention have been interpreted as enhancing the signal-to-noise ratio of neural responses (Noudoost et al., 2010). This can involve both increasing signal by enhancing firing rates or decreasing noise by increasing reliability of spiking responses (Mitchell et al., 2007). To determine whether reliability increases in the post-cue period of our task, we quantified across-trial spiking reliability with the Fano factor, or the variance divided by the mean of spike counts across trials within a sliding 50ms window. We found a significant decrease in Fano factor (i.e. increase in reliability) during cue-RF trials relative to cue-orthogonal (Figure 3-8; median=-0.4%; p=0.002). Similarly, we found a decrease in the same period for cue-opposite compared to cue-orthogonal trials (median=-0.7%; p=0.001). Thus, both attention and saccade preparation correlate with decreased variability, or increased reliability, of area V4 spiking responses.




Figure 3-8. Effects of cue direction on across-trial spiking reliability. A. Histogram of the effect of cueing the RF stimulus on Fano factor (FF) for each unit. The effect is measured as a modulation index: the difference between FF in the cue-RF and the cue-orthogonal conditions divided by the sum. FF was computed with a 50ms bin size and with spike counts in the period from 500ms after cue onset until the start of the blank period. Positive modulation indices indicate larger FF in cue-RF than cue-orthogonal

condition, corresponding to increased variability, i.e. decreased reliability. p-value shown for Wilcoxon signed rank test. The colored part of histogram corresponds to units for which FF was individually significantly modulated by cue direction. B. As in A, but for cue-opposite condition compared to cue-orthogonal condition.

3.4.5 Relationship between attention-related and saccadic modulations for individual neurons

The effects described thus far demonstrate that activity in the population of V4 neurons is robustly modulated during both attention and saccade preparation. However, for individual neurons these effects might be uncorrelated. That is, a given neuron might be enhanced during attention but unaffected during saccade preparation or vice versa, suggesting two distinct underlying mechanisms. Thus we examined the relationship between the firing rate modulations in the two conditions on a neuron-by-neuron basis. First, we measured the correlation between the magnitudes of the two modulations for all neurons, whether or not they showed significant effects. We found that attention-related and saccadic modulation were weakly, but significantly, correlated ($r=0.06$, $p<10^{-10}$). Since this correlation was likely diminished by many unmodulated neurons, we next narrowed our analysis to only those neurons significantly modulated during cue-RF trials (i.e. during attention; $n=143$ neurons, 100 enhanced and 43 suppressed; median +27.8% and -18.5% cue-RF vs cue-orthogonal firing rates, respectively). We asked whether these neurons were likely to be modulated, and in the same direction, during saccade preparation (cue-opposite trials). Of the 143 attentionally modulated neurons, 64 of them (45%) were also significantly modulated during saccade preparation, and of those, 62 (97%) were modulated in the same direction (51 enhanced during both, 11 suppressed during both). Overall, the neurons enhanced during cue-RF trials had on average enhanced responses during cue-opposite trials (median = +17.3%; Figure 3-9). This enhancement significantly exceeded that of the overall population ($p<10^{-5}$). Similarly, neurons suppressed during cue-RF trials were suppressed during cue-opposite trials (median = -4.4%), and this suppression exceeded that of the overall population ($p<10^{-4}$). Moreover, the cue-opposite modulations of these two groups were significantly different from each other (difference = 21.7%; $p<10^{-7}$). Thus, not only did the two behavioral conditions

produce similar effects on the population of neurons, but these effects were also similar on a neuron-by-neuron basis.




Figure 3-9. Cue-opposite firing rate modulation for those neurons significantly modulated during cue-RF trials. Overlaid histograms of cue-opposite firing rate modulation index for neurons significantly enhanced (red) and significantly suppressed (black) during cue-RF trials.

3.4.6 Local field potential power

During covert attention, the frequency spectrum of LFPs in area V4 changes markedly, with increases in power at high frequencies and decreases in low frequencies (Fries et al., 2001), changes that may reflect underlying cortical state dynamics (Harris and Thiele, 2011). We computed the power in the gamma (40-70Hz) and beta (10-20Hz) frequency bands on trials split by cue direction. We found a significant decrease in beta power ($p < 10^{-6}$) and increase in gamma power ($p < 10^{-7}$) during cue-RF trials relative to cue-orthogonal (Figure 3-10A). On cue-opposite trials we found the same changes in LFP power compared to cue-orthogonal trials (beta, $p < 10^{-5}$; gamma, $p < 10^{-8}$). Similar changes in the oscillatory structure of LFPs accompany both attention and saccade preparation.




Figure 3-10. Difference in LFP power between cue conditions across frequencies. Fourier transforms were computed for the final 500ms of the post-cue period for each trial and each channel, then averaged across trials and channels within each cue condition and recording. The mean differences, cue-RF minus cue-orthogonal (cyan) and cue-opposite minus cue-orthogonal (red), across recordings are represented with shaded regions reflecting ± 1 S.E.M. Median percentage differences within two frequency ranges of interest, beta (10-20Hz) and gamma (40-70 Hz), are shown.

3.5 Discussion

We trained monkeys to perform an attentionally demanding cued change-detection task with antisaccade response. While monkeys performed this task we recorded the activity of 696 single- and multi-units from visual area V4.

Critically, the use of the antisaccade response means that the onset of the cue provided two separate but linked pieces of information to the monkey: it identified which of the four peripheral stimuli would contain visual information useful for obtaining rewards (the cued stimulus); and it identified which of the four stimuli should be the target of a saccade if a change was detected (diametrically opposite to the cued stimulus). The other two stimuli, located in orthogonal directions to the cue direction, were irrelevant visually and were unlikely to be the target of saccades. With the onset of the cue, the monkey could both begin covertly attending to the cued visual stimulus as well as preparing the saccade that would be rewarded in the case of an orientation change at the cued stimulus. This task thus behaviorally dissociated attention from saccade preparation by requiring monkeys to attend a stimulus at one location while planning and eventually executing a saccade to another.

Previously described effects on visual cortical responses during covert attention include increases in firing rate (Moran and Desimone, 1985), orientation tuning amplitude (McAdams and Maunsell, 1999), across-trial spiking reliability (Mitchell et al., 2007), gamma-range LFP power (Fries et al., 2001), as well as decreases in low-frequency LFP power (Fries et al., 2001). We observed all of these effects during the cue-RF condition of our task, relative to cue-orthogonal (i.e. attend-to vs. attend-away), confirming that monkeys performing this task were able to adopt a covert attention strategy similar to other tasks, despite the task demand to prepare a saccade away from the location of covert attention.

3.5.1 Size of attention effects

The median size of the attention-related effects on firing rate (2.8%) were somewhat smaller than observed in past studies of attention-related modulations in area V4. Several factors probably contribute to this. First, we recorded with electrode

arrays rather than single electrodes, as most previous studies did, so our sampling of neurons was perhaps somewhat less biased toward neurons with high firing rates than past studies might have been. Indeed, one other study that used electrode arrays to perform recordings also found much smaller effects (8.6%; Cohen and Maunsell, 2010) than previous studies (e.g. 26%, McAdams and Maunsell, 1999). Second, the magnitude of attention-related modulation of firing rates depends strongly on the specific RF stimulus (Figure 3-7), and it was impossible to optimize stimulus properties (e.g. spatial frequency, size, color) for all simultaneously recorded neurons. Modulation is also stronger for low contrast stimuli (Reynolds et al., 2000) whereas we used high contrast stimuli. Finally, and most importantly, the task demands in this study were different in that the saccade target was dissociated from the location of covert attention, a manipulation that induced modulation during the post-cue part of the task at two locations. To the extent that the modulations we observed are smaller than in previous studies, it may reflect the splitting of a single resource involved in directing both covert attention and saccades.

3.5.2 Presaccadic modulation in visual cortex

Firing properties of area V4 neurons are also modulated in advance of saccadic eye movements, similar to the way they are modulated during covert attention. Firing rates increase (Fischer and Boch, 1981), orientation selectivity increases (Moore and Chang, 2009), and across-trial reliability increases (Steinmetz and Moore, 2010). However, these experiments all employed visually-guided saccade tasks, in which the covert attention of the animal was unconstrained by task demands. Therefore, the effects observed in those tasks could have indicated an optional pre-saccadic allocation of attention to the saccade target location, rather than any modulation due directly to the saccade preparation *per se*. By behaviorally dissociating the location of the saccade target from the location of relevant visual stimuli, we have demonstrated that saccade preparation itself is sufficient to induce these modulations.

3.5.3 Implications for the circuits controlling attention

Theoretical accounts of selective attention have suggested that the focus of attention may be driven by a “saliency map,” which reflects the behavioral relevance of visual stimuli based on a combination of their intrinsic visual properties as well as top-down biases related to task demands (Koch and Ullman, 1985). This saliency map would then direct the attentional modulation of visual cortical activity to favor the most salient stimulus. For instance, activity of many neurons in area LIP reflects the salience of visual stimuli, either produced by abrupt onset or by behavioral context, but is independent of saccade planning (Gottlieb et al., 1998). Authors suggested that this population of neurons could drive the modulation of visual cortical responses. However, if so, we would not have seen modulation of visual cortical responses to stimuli that were not task relevant except as saccade targets. Similarly, other work has suggested that attention-related modulations may be driven by neurons that have only visual-related responses and that do not represent information about upcoming saccade plans (Gregoriou et al., 2012). However, this too is not consistent with the modulation we observed during saccade preparation.

3.5.4 What is the source of these modulations, if they share a single source?

A parsimonious interpretation of the strikingly similar modulations of visual cortex brought about by attention and by saccade preparation in our data is that a single pool of neurons drives both modulations and that these source neurons are activated both during covert attention and during saccade preparation. That is, the simplest interpretation of these results is that covert attention and saccade preparation share neural mechanisms at least at the level of those responsible for modulating visual cortex. The source of these modulatory signals may be neurons in any area that is involved in controlling both attention and saccades, such as the LIP, SC, or FEF. The evidence for the FEF’s involvement as a source of attention and visual cortical modulation is more extensive and consistent than the other candidate areas: electrical microstimulation of FEF drives attention-like spatially specific modulation of perception (Moore and Fallah, 2001) and visual cortical activity (Moore and

Armstrong, 2003); subtle alteration of FEF activity via local disruption of neuromodulation (Noudoost and Moore, 2011a) or via voluntary control of firing rate by biofeedback (Schafer and Moore, 2011) induces similar effects; and FEF activity becomes specifically synchronous with V4 firing during covert attention to V4 RF stimuli but not during attention away (Gregoriou et al., 2009). The FEF also plays a critical role in saccade production: fixed-vector, low-latency saccades are produced upon FEF stimulation with low currents (Bruce et al., 1985); some types of saccadic behaviors such as memory-guided saccades are nearly abolished with FEF inactivation (Dias and Segraves, 1999); and monkeys with SC lesions can still produce saccades though monkeys with both SC and FEF lesions cannot (Schiller et al., 1987). That the FEF has such a deep involvement in both functions supports the claim here that both behaviors may be interdependent and subserved by shared circuits.

3.5.5 An alternate interpretation: two separate and independent sources with nearly identical effects

We have made the argument that the interdependence of attention and saccade preparation insofar as they influence visual cortical activity suggests that these visual cortical modulations are driven by the same source, such as the same populations of neurons in an area like the FEF. However, the results presented here are also consistent with another possibility, that attention and saccade preparation both independently induce modulation of visual cortex, modulations which simply have similar effects on neural activity. In fact, several of the effects we have described – increased firing rates, increased across-trial reliability of spiking, increased gamma and decreased beta LFP power – may be manifestations of a single underlying process: changing from a synchronous to desynchronous state (Harris and Thiele, 2011). Such a state change may be brought about by glutamatergic inputs (Zagha et al., 2013), such that glutamatergic drive from different sources may be sufficient to bring about many of the observed effects. It remains to be seen whether the observed effects on orientation tuning could be generated by such a state change, and whether direct evidence of such a state change could be found in macaque visual cortical activity. Given these uncertainties, then, we contend that the parsimonious explanation

presented above, that both modulations arise from a common source, remains the more likely.

3.5.6 Behavioral performance

Monkeys did not perform the task perfectly, averaging 69% and 67% correct. However, this performance was intentionally limited by the change-blindness manipulation (blank period) and the change magnitude (45°, compared to maximum of 90°). With much shorter blank periods and larger changes, monkeys performed much higher accuracy (data not shown). Furthermore, reaction times were substantially shorter for validly cued trials compared to invalidly cued trials, including those invalid cues with changes at the location opposite to the cue. Finally, the sensitivity of the monkeys to detecting changes was enhanced at the cued location, directly demonstrating that they allocated attention there. We conclude that monkeys understood task rules and were limited in their performance only by the difficulty of the visual change detection and by the attentional demands created by the change blindness manipulation.

3.5.7 Noise correlation analyses

In addition to the signatures of attention we examined, it has also been observed that neuronal responses in area V4 become decorrelated during attention (Cohen and Maunsell, 2009; Mitchell et al., 2009). However, as those studies recorded from pairs of neurons separated laterally in cortex whereas we recorded from pairs separated in the vertical dimension, an analysis of this factor in our data would be difficult to interpret.

3.5.8 Dissociation of saccade location from cued location in this experiment

Due to the deterministic link between the cued location and the saccade target (always diametrically opposite), they are technically not dissociated from each other. In other words, neural effects observed during cue-RF trials could have either been due to attention to the RF stimulus or to the saccade plan to the opposite stimulus,

since these events co-occurred on all of those trials. However, V4 neurons have strongly localized spatial sensitivity, so considering the spatial extent of attentional and saccadic modulations, they are dissociated from the perspective of a V4 neuron's RF. Furthermore, since both orthogonal conditions contained saccade preparation to targets closer to the V4 RF than in the cue-RF condition, and modulations were not observed in those cases (likewise for attentional modulation), we conclude it is highly unlikely for either process to have had an influence on V4 responses when directed in the opposite direction.

4 Attentional and saccadic modulation across neocortical layers in area V4

4.1 Abstract

While changes in spiking responses of individual visual cortical neurons during attention have been well-described, the way attention-related signals are processed within neocortical microcircuits has not. That is, are neurons found in distinct cortical layers modulated differently during attention? We recorded the activity of neurons across layers within single cortical columns in area V4. We found that superficial neurons, compared to deep, responded earlier, responded more strongly, and were more likely to be tuned for the orientation of visual stimuli. However, we found that modulation during attention and during saccade preparation did not differ across depth. These results are consistent with a form of modulation that influences all neurons throughout cortical layers equally.

4.2 Introduction

Humans and other primates have the ability to covertly attend peripheral visual stimuli of interest, an ability that consists of a powerful enhancement of the perception of those stimuli relative to others (Carrasco, 2011; Posner, 1980). Changes in the responses of neurons within visual cortex (Desimone and Duncan, 1995) and even in visual thalamus (McAlonan et al., 2008) have been well-studied and may underlie these perceptual changes. Similarly well-investigated are the mechanisms that control the allocation of attention and direct the modulations seen in visual cortex. For instance, the frontal eye field (FEF) has been strongly implicated in the control of attention with a variety of inactivation (Wardak et al., 2006) and recording (Buschman and Miller, 2007) experiments as well as causal demonstrations of its role using electrical microstimulation (Moore and Armstrong, 2003; Moore and Fallah, 2004). Similar evidence implicates the superior colliculus (SC) in the control of attention. Inactivation (Lovejoy and Krauzlis, 2010), recording (Goldberg and Wurtz, 1972), and electrical microstimulation (Cavanaugh and Wurtz, 2004; Muller et al., 2005) experiments all point to a role of this subcortical structure in the control of visual selective attention.

Critically missing is an understanding of the mechanisms that link the control signals in areas like FEF and SC with the modulations observed in visual cortex. How do the signals travel from the source area to the modulated target region? When these attention-related signals arrive there, how do they integrate with feedforward visual signals within the complex microcircuit of neurons in a cortical column representing the attended stimulus? We attempt to address both of these questions in this study.

First, we consider a novel approach to identifying the source of the signals that directly modulate visual cortical responses. Though there are many pathways reaching from frontal, parietal, and subcortical structures back to visual cortical targets, many of these projections have different laminar patterns of termination across layers in visual cortex (see Figure 1-2 and Section 1.3.3 in this thesis). By identifying the pattern across depth of the magnitude and latency of covert attention-related modulation, the particular pathway carrying these signals may be disambiguated.

Second, we asked how distinct neuron populations, defined by their depth in cortex and waveform duration, are differently modulated by attention. Identifying differences in the response properties of distinct populations of neurons within a cortical area, for instance, neurons within different anatomical layers, has historically proven to be a powerful way to understand the computation performed within an area (Hirsch and Martinez, 2006; Hubel and Wiesel, 1968). Characterizing the way that different elements of the microcircuit are modulated during attention may therefore reveal the locus of the integration of attention-related signals with visual processing. It will also provide valuable constraints for detailed microcircuit models of attention (Ardid et al., 2007; Tiesinga and Buia, 2009; Wagatsuma et al., 2011). Furthermore, since the patterns of neuron types and connectivities within and between layers of neocortex are stereotyped (see Section 1.3.4), this approach may generally aid in our understanding of the way cognitive feedback signals interact with sensory processing in the mammalian neocortex.

To address these questions, we recorded from neurons across layers within area V4, a part of extrastriate visual cortex that is robustly modulated during covert attention (e.g. Motter, 1993) and may also play a particularly critical role integrating attention-related feedback with visual representations (Schiller and Lee, 1991; De Weerd et al., 2000). We found that some basic response properties differed across layers: in superficial layers, average firing rates were higher, response latencies were shorter, and neurons were more likely to be tuned for stimulus orientation than in deep layers. Despite these differences, modulation during attention and during saccade preparation did not differ across layers.

4.3 Methods

4.3.1 Subjects

Two male monkeys (*Macaca mulatta*, 8–12 kg) were used in these experiments. All experimental procedures were in accordance with National Institutes of Health Guide for the Care and Use of Laboratory Animals, the Society for Neuroscience Guidelines and Policies. General surgical procedures have been described previously (Graziano et al., 1997).

4.3.2 Behavioral task and visual stimuli

4.3.2.1 Receptive field mapping task

Monkeys fixated a small (~0.3 d.v.a.) white dot against a medium gray background. A horizontally oriented grating was flashed for 50ms at each of six different positions per trial with 150-250ms variable delay between flashes. If the monkey maintained fixation within a 1.8 d.v.a. square window until after the sixth flash, he received a juice reward.

The flashes occurred at a total of 36 locations on a 6x6 grid with 3 d.v.a. spacing (total coverage 15x15 d.v.a.). On each trial the six flash positions were selected from one of the rows of the grid in random order. The upper right position of the grid was at the fovea such that only the lower left visual field was covered by the mapping.

4.3.2.2 Full-field flash task

Monkeys fixated a small (~0.3 d.v.a.) white dot against a black background. The monitor turned maximal white for one frame (~8ms) then back to black. The flash occurred six times per trial with variable delays in the range of 150-250ms. If the monkey maintained fixation within a 1.8 d.v.a. square window until after the sixth flash, he received a juice reward. Approximately 30 trials, or 180 flashes, were completed per day.

4.3.2.3 Cued change-detection task with antisaccade response

Please see section 3.3.2 for details of the behavioral task and visual stimuli used in this experiment, as the data discussed in this chapter are from the same experiment.

4.3.3 Neural recordings

Many aspects of the neural recordings are discussed in detail in section 3.3.4, including descriptions of the electrodes, amplification, digitization, spike detection, and spike sorting. However, the fact that electrodes were linear arrays (Figure 4-1), and the particular advantages conferred by this fact, was not emphasized there, and so is discussed further here.




Figure 4-1. Recording setup, electrode array, and sample data. Recordings were made from area V4, left panel. Electrodes were 16-channel linear arrays with 150 μm contact spacing. These electrodes were capable of recording LFPs, center, from each channel as well as spikes from isolated single neurons on some channels, right.

4.3.3.1 *Electrode targeting: Use of MRI guidance to achieve perpendicularity*

We sought to achieve simultaneous recordings at sites located within a single cortical "column." In particular, the topographic organization of extrastriate visual cortex suggests that vertically separated neurons should have overlapping RFs, so we sought to record from a column by this definition. Since the cortical magnification factor (an estimate of how much physical space in cortical tissue corresponds to an amount of visual space) is approximately 1 deg/mm (Gattass et al., 1988), we could measure the approximate angle with the cortex by the distance between receptive fields measured on the deepest and most superficial recording contacts, and sought to keep this angle at 10 degrees or less, corresponding to a RF shift of ~0.5 degrees, given 2mm thickness of cortex.

In order to achieve these perpendicular penetrations we performed an MRI targeting technique inspired by (Kalwani et al., 2009). We implanted the monkeys with custom built recording chambers made from PEEK-type plastic, rather than from titanium, to avoid "shadows" in the MRI images. While we did not employ ceramic skull screws, we took some care to ensure that the titanium skull screws and plates were not located close to the recording chamber and brain areas of interest. We filled a custom-made plastic cylinder with copper sulfate solution. We anesthetized the monkey and inserted this cylinder into the recording chamber, into which it fit snugly. We performed structural MRI imaging (1.5 Tesla; T-1 weighted image) to visualize the location and orientation of the recording chamber (visible due to the high-contrast copper sulfate solution within it) relative to the position of the prelunate gyrus within the brain. By manually identifying the contours of the prelunate gyrus, we could compute perpendicular vectors to it and project these back to the level of the electrode stage, thus identifying which penetration approach vectors were likely to yield perpendicular penetrations.

4.3.3.2 *Achieving desired approach vectors*

We employed a custom-built targeting device, the "well-angler," to tilt and rotate the electrode into any desired orientation and position in three dimensions, designed

and built in collaboration with Bob Schneeweis, Department of Neurobiology, Stanford University (Figure 4-2). The device consisted of a “double-eccentric” mechanism for positioning the electrode in the x-y plane of the well, a tilting mechanism, and a rotating mechanism. All four coordinates could be set with sub-millimeter precision using notches engraved in the device. The V4 recording chambers on both monkeys projected from the monkeys’ heads at an angle such that there was a unique point on the chamber’s perimeter at the lowest elevation. This point was identified computationally in the MRI images and was identified on the chamber itself by filling the chamber with saline solution until the liquid first contacted the lip of the chamber. With just this one point of alignment between the MRI images and the physical well, the exact X, Y, tilt, and rotation coordinates for an approach vector specified by the MRI images could be geometrically determined.




Figure 4-2. “Well angler”: a device to tilt and position an electrode at an arbitrary vector. A, Schematic of recording chamber (cylinder), electrode path (dotted line), and degrees of freedom. B, Photograph of well angler with mechanisms for adjusting degrees of freedom labeled. The Narishige micromanipulator controls

the z-position of the electrode. The X and Y coordinates are multiplexed by the double-eccentric system at the base of the device, such that these two coordinate settings jointly specify the X, Y position.

4.3.4 Data Analysis

4.3.4.1 Electrode targeting: Assessing perpendicularity with RF overlap

RF positions and extents were estimated by computing the number of multi-unit spikes recorded on each channel in the 200ms period following stimulus onset for each of the 36 stimulus positions in the RF-mapping task (see section 4.3.2.1). This 6x6 matrix of response counts was cubic spline interpolated to produce the full “RF map” and a 75%-of-max contour was determined, defining the RF border. The center of mass of the portion of the RF map within the RF border was defined as the RF center. This analysis was performed after recording RF-mapping task responses but before the change-detection task, so that a stimulus position could be chosen at a location that fell within the RF borders for all channels. If such a position was found, the recording was included in further analyses.

4.3.4.2 Electrode targeting: Depth alignment

We lowered electrodes into the brain rapidly ($\sim 25\mu\text{m}/\text{sec}$) until one channel was in the cortex, based on visual examination of LFP and spiking activity being recorded concurrently. Then we advanced the electrode slowly ($\sim 5\mu\text{m}/\text{sec}$) until the uppermost electrode contact was near the point of entering the brain, being recorded during advancement. We withdrew the electrode $500\mu\text{m}$ to release compression of the brain caused by the electrode. During this brief withdrawal, no apparent change in the LFP or spiking activity was observed, confirming that this served to relax the cortex rather than change the position of the electrode relative to the brain. This manipulation qualitatively improved stability and recording quality. After reaching this position, the full-field flash task (see section 4.3.2.2) was run to assess the depth.

We computed the current source density (CSD) response to the full-field flashes. The CSD reflects the spatial and temporal position of current sources and sinks (i.e. where current flows into and out of the extracellular space, respectively) along the length of the electrode, given certain assumptions likely to be true for our recordings

(Mitzdorf, 1985). The CSD can be computed discretely as the second spatial derivative of the LFP for each point in time, that is:

$$D(z) = \frac{\phi(z + h) - 2\phi(z) + \phi(z - h)}{h^2}$$

where z is the position in depth, h is the distance between potential measurements (in our case, $150\mu\text{m}$), and $\phi(z)$ is the potential measured as a function of depth. We also calculated the CSD according to the inverse estimation method (Pettersen et al., 2006), and display the results of this calculation, which produces smoother, higher resolution plots of CSD, in figures for clarity. However, results were qualitatively indistinguishable with both methods. Borders between current sinks of interest were manually identified and channel depths were computed, in mm, relative to these borders.

4.3.4.3 Data analysis software

Due to the large number of simultaneous recording channels and the overwhelming amounts of data associated with each, novel software solutions were required for accessing and manipulating these datasets. In older datasets, the entire history of spiking activity from all neurons might be loaded into memory simultaneously, for instance in a “struct” variable type in Matlab, and easily accessed and indexed in this way. With data from 717 neurons and LFP data from 736 recording sites, this was no longer possible. Furthermore, if many types of data are recorded and many analyses performed, keeping track of the original data and the results quickly becomes untenable without a well-planned organizational scheme. A new approach was needed to keep data organized and to maintain (or improve) the ease of accessing data without demanding that all the data be simultaneously loaded into memory or that the user remember or employ the hard drive locations of the data.

Therefore, I developed a new standard format for data storage and organization with an object-oriented codebase in Matlab for accessing it. This system, called “dataDB” (I know), allows for data to be accessed from Matlab with simple, intuitive

lines of code, and for analyses to be easily stored in logical locations and accessible at a later date.

The storage format consists of a primary folder containing all the data, subfolders for each date/electrode penetration, subfolders of these for each experiment carried out on that day (e.g. RF mapping, full-field flash, attention task), subfolders of these for each electrode channel that was recorded and a separate subfolder for the behavioral data. Within a channel's folder is a subfolder for each "spikeUnit" and for the LFP. Pieces of data are stored as Matlab-format variables within these folders.

The software works directly with this data structure, parsing it and representing each folder as an object exposed to the user. Data in each folder can be easily viewed and accessed from the command line with index notation, for example, the command:

```
>> st = db.r(1).c(10).s(2).spikeTimes;
```

returns the list of spike times for the second spikeUnit on the tenth channel of the first recording, assuming such data exists. Note the user was not required to know anything about the actual location of these data on the hard drive, nor was the data loaded into the memory of the computer prior to accessing. Similarly, the results of a novel analysis can be saved for later recall with a simple command:

```
>> su = db.r(1).c(10).s(2);  
  
>> su.putData('meanFiringRate', meanRate);
```

This will store the data in the correct location on the hard drive and make it available for access with index notation in the future. Various other programming conveniences are provided including referencing recordings by name rather than number, functions to "open" and "close" the objects (loading their data to RAM for rapid access and removing it), and functions for working easily with certain types of behavioral data, among other things. Further documentation is available upon request.

4.3.4.4 *Waveform duration*

Average waveform durations across all spike events recorded for each neuron were computed in order to distinguish between “broad” and “narrow” spiking waveforms (e.g. Cohen et al., 2008; Mitchell et al., 2007). The waveform duration for each neuron was defined as the median of the time from trough to peak for each calculated for each waveform. This procedure resulted in a bimodal distribution of waveform durations, as expected, split at 250 μ sec (Figure 4-3). This procedure also conveniently identified the subpopulation of neurons I call “peak-first” neurons, in reference to the fact that their waveform peaks precede the troughs. The biophysical interpretation of these peak-first waveforms is unclear, but I conjecture that they may reflect recordings from myelinated axons rather than from cell bodies. For some analyses they have been excluded.




Figure 4-3. Waveform duration distribution. Red line indicates the dividing line for considering a neuron to be broad or narrow spiking. Neurons of these two classes were found in the proportions indicated, approximately as expected, though with somewhat greater representation of narrow spiking neurons than other authors have reported for area V4 (Mitchell et al., 2007).

4.3.4.5 *Firing rate*

Firing rate was computed in three different ways: for the average rate across the whole trial, for the baseline, pre-stimulus period, and for the stimulus period for non-preferred stimulus trials. The average rate across the whole trial was determined as the

mean rate on all attempted trials (correct or incorrect) between the time of stimulus onset and 500ms after the stimulus reappearance at the end of the trial (this window included the time of saccades for almost all trials with saccades). The baseline rate was computed for the window from 100ms prior to stimulus onset until 30ms after stimulus onset. The rate on nonpreferred stimulus trials was computed for the period from 25ms to 275ms after stimulus onset. The nonpreferred stimulus was defined as the orientation that elicited the minimum response from the neuron during that window, regardless of whether the neuron was significantly tuned for orientation. For this analysis, only isolated single neurons were considered.

4.3.4.6 Stimulus onset latency

Stimulus onset latency was computed in two ways: manually, and as the time to peak response. In the manual method, the time point at which the stimulus-onset aligned PSTH appeared to first diverge from baseline was identified for each neuron. In the time to peak method, the time of the peak of the stimulus-aligned PSTH was determined automatically. In both cases, PSTHs were smoothed with a 20ms box filter.

4.3.4.7 Orientation tuning

Significance of orientation tuning was determined with a Kruskal-Wallis test on the stimulus driven firing rates grouped by the orientation of the stimuli on those trials. Units with K-W $p < 0.0001$ were considered significantly tuned.

4.3.4.8 Modulation by cue direction

Firing rates and modulation indices were computed as described in Section 3.3.5.

4.4 Results

We recorded the activity of 696 single and multi-units within area V4 during a cued change-detection task with antisaccade response. We performed these recordings with linear array electrodes such that the recorded units were distributed in depth across the cortical layers. We considered how basic response properties and modulation by cognitive state differed between neurons recorded at different depths.

4.4.1 RF alignment

We employed an MRI targeting approach so that electrodes could be inserted perpendicular to the cortical layers. With this approach, RFs of all simultaneously recorded units were overlapping, so that a single visual stimulus placed on the screen would drive all of the units similarly (Figure 4-4).




Figure 4-4. Example RF alignment using the MRI targeting procedure. Each colored contour indicates the RF border from one channel of the recording and the corresponding colored dots indicate the RF centers. The color of the RF border designates the depth in cortex at which that contour was recorded, from most superficial (orange) to deepest (magenta). The mapped region of space was the lower-left quadrant, i.e., the monkey fixated at the upper-right corner of the depicted area.

4.4.2 Depth registration

In all included recordings, a prominent current sink was identified near the middle of the electrode, approximately 40-50ms after flash onset. This was followed by another sink just below the first, peaking approximately 100ms after flash onset. These two sinks appeared in every included recording, and we therefore aligned the recordings on these functional markers of cortical laminae (Figure 4-5). In many recordings, further sinks were observed near the top of the probe at ~150ms and near the bottom of the probe at ~50ms. Because the widths of all four of these sinks, when present, was highly consistent from recording to recording, we assigned each channel a depth relative to this central feature. Depths were measured in millimeters, and positive depths indicate channels superficial relative to the CSD feature.

In some sessions, further CSD recordings at deeper locations revealed that no further current sources or sinks of comparable magnitude could be identified below these CSD features, assuring us that our electrode covered the depth of cortex (data not shown).




Figure 4-5. Example current source density (CSD) with alignment feature. CSD responses were computed from the LFP response to full-field flash stimuli. Across recording sites and monkeys similar features are visible. These features allow alignment of the channels in each recording to a consistent functional depth.

4.4.3 Firing rate

We measured the mean firing rate of isolated single neurons recorded at different depths within area V4 and sorted the units according to their depth (Figure 4-6). Different measures of firing rate were computed: mean baseline (pre-stimulus) rate; mean rate across the whole duration of the trial; and mean rate during trials on which a nonpreferred stimulus was presented. Firing rates were distributed roughly exponentially, and the majority of units had average rates $<5\text{Hz}$. The least-responding units were preferentially found in the deep channels, such that average firing rates were higher for superficial units (superficial median = 4.09Hz ; deep median = 1.97Hz ; ranksum test $p=0.003$). This same difference across depth was true also for firing rates

computed only on trials with non-preferred stimuli (2.88 vs 1.17Hz, $p=0.01$), to rule out an influence of differences across depth in the tuning for the stimulus set. While it was not significant for baseline, pre-stimulus firing rates, the trend was the same (superficial vs. deep 2.28 vs 1.62Hz, $p=0.10$).




Figure 4-6. Firing rates as a function of depth in cortex. Each blue circle corresponds to the depth and latency of one recorded V4 unit. Red circles and black bars represent mean and S.E.M. for units recorded at each depth. For display purposes, depths with ≤ 5 recorded units are not displayed (depths 1.5, 1.2, 1.05, -1.5, -1.65mm), and one other individual unit is omitted (depth=0, rate=99.42Hz).

4.4.4 Stimulus onset latency

We measured the time from stimulus onset to the time when firing rate began to increase or decrease for each unit and considered how this quantity varied across depth (Figure 4-7). The latency of responses was on average later in the deeper layers

(median 56.7ms in deep; 54.2ms in superficial; ranksum test, $p=0.02$). The latest-responding units tended to be found in the deepest layers (latest 7 all deep). However, the distributions of response latencies at all depths were highly overlapping, with some units at all depths with ~ 40 ms latencies, the earliest seen in our sample. The results were similar with other definitions of stimulus onset latency, including time to peak response (superficial vs. deep, 70.5 vs. 74.7ms, $p=0.004$).




Figure 4-7. Stimulus onset latency as a function of depth. Each blue circle corresponds to the depth and latency of one recorded V4 unit. Red circles and black bars represent mean and S.E.M. for units recorded at each depth. Depths with ≤ 5 recorded units are not displayed.

4.4.5 Orientation tuning

We determined, for each unit, whether it was significantly tuned for the orientation of the RF stimulus (Figure 4-8). We compared the proportion of significantly tuned

units in the superficial portion of cortex to that in the deep. Superficial units were significantly more likely to be tuned than deep units ($\sim 2/3$ to $\sim 1/3$, chi-square $p=0.001$).




Figure 4-8. Orientation tuning across depth. A, Example orientation functions for two significantly tuned and one untuned unit. Blue circles and black lines correspond to mean and S.E.M. of firing rate elicited during the stimulus onset period by each stimulus orientation. Red lines correspond to the von Mises fit to the data. B, Histogram of number of tuned (red) and untuned (black) units by depth. C, Proportion of tuned and untuned units by depth.

4.4.6 Firing rate modulation during attention and saccade preparation

We computed firing rate modulation during the cue-RF and cue-opposite conditions compared to the cue-orthogonal condition, and asked how this quantity varied across depth. Neither cue-RF nor cue-opposite modulation was significantly

affected by depth (Kruskal-Wallis $p > 0.05$), and neither cue-RF nor cue-opposite modulation was different between superficial and deep units. The two forms of modulation did not have different patterns of modulation across depth (cue-RF MI minus cue-opposite MI, Kruskal-Wallis $p > 0.05$).




Figure 4-9. Firing rate modulation by cue direction across depth. Mean and S.E.M. shown for cue-RF (cyan) and cue-opposite (red) MIs for units recorded at each depth.

4.5 Discussion

We recorded from neurons across layers in area V4 using acutely-inserted linear array electrodes. We recorded from neurons within a single cortical column, confirmed by measured RF alignment, and registered the depth of channels recorded on different days with a functional register of depth. These recordings allowed us to assess the variation of basic response properties across depth as well as the modulation due to attention- and saccade-related feedback.

4.5.1 Overlapping RFs and recording from a single “column”

This served both a practical and theoretical purpose. Practically, with overlapping receptive field we were able to show task-relevant stimuli in the RF of all neurons simultaneously, thus maximizing the yield of channels driven by the task conditions. Theoretically, we considered that given substantial data about the mostly vertically oriented connectivity in many cortical areas and the sharp fall-off in connection probabilities with lateral distance (Boucsein et al., 2011), we were more likely to find neurons functionally connected by recording from vertically separated sites. Definitions of the cortical column and estimations of its width are contentious, but some data exists on this topic for area V4 of the macaque. In one study, investigators demonstrated with focal biocytin injections into V4 that neurons project laterally in a patchy manner, such that axon terminals can be found elsewhere in V4 clustered into column-like cylinders of ~250-450um in diameter (Yoshioka et al., 1992). Therefore we sought to record only from neurons separated by this distance or less, as defined by distance between RFs and the cortical magnification factor (see Section 4.3.4.1).

4.5.2 Correspondence of CSD features with anatomical layers

Two alignments of these functionally defined layers with anatomical cortical layers seem possible, and at present it has been impossible to perform the histological examination necessary to disambiguate them. The late, uppermost sink could correspond to layer 2/3 (together), the early sink just above depth=0mm layer 4, the later sink just below depth=0mm layer 5, and the deepest early sink layer 6 (Figure 4-10, middle panel). Alternately, the four visible sinks could correspond to layer 2, 3,

4, and 5 in order from superficial to deep (with no layer 6 in the range of the recording; Figure 4-10, right panel). On the one hand, the first assignment seems reasonable as the thickness of the layers known histologically matches the thickness of these CSD features reasonably well, and our expectation from primary sensory areas is that layers 4 and 6 will have the earliest responses (Hansen and Dragoi, 2011; Schroeder et al., 1998; Swadlow et al., 2002). However, the cortex may well be compressed around the electrode as it is inserted thus skewing the measured layer thicknesses. Layer 2 and 3 are well-differentiated cytoarchitecturally in V4 unlike in V1, suggesting they may not appear as a single sink. Finally, the earliest driving visual inputs into V4 are probably not from the ventral stream (Chen et al., 2007), which project into layer 4 (Ungerleider et al., 2008), and may instead arrive from the pulvinar nucleus of the thalamus (Guillery and Sherman, 2002; Shipp, 2003), which synapses into deep layer 3 (Jones, 2007).




Figure 4-10. Two possible alignments between the anatomical layers in area V4 (left) and the CSD as measured in response to the full-field flash stimulus (right). Anatomical image taken from www.brainmaps.org.

4.5.3 Differences in response properties across depth

We observed differences in mean firing rate, response latency, and orientation tuning across depth in area V4. None of these properties have been quantified as a function of depth for single neurons in this area before. The latency and tuning results together suggest a particular functional structure within area V4. Since deep layers are

activated later than superficial, it may be that descending projections from the superficial layer neurons primarily drives the deep. Indeed this projection is a large component of the canonical neocortical microcircuit (Douglas and Martin, 2004). The differences in proportion of neurons significantly tuned for the stimuli in this experiment may thus reflect that a computation takes place between the superficial tuned neurons and the deep untuned neurons as a result of this projection, perhaps a computation that combines responses of the superficial neurons in such a way as to confer on the deep neurons the more complex stimulus selective properties characteristic of many V4 neurons, such as curvature or boundary contour tuning (see Section 1.3.1). Future experiments could specifically test the hypothesis that deep neurons preferentially express selectivity for more complex stimulus properties with the recording techniques described here. Furthermore, the hypothesis that superficial layer neurons drive deep neurons could be more rigorously tested with cross-correlation analysis or more advanced system-analysis methods (Blomquist et al., 2009).

The greater firing rates observed in superficial neurons are perhaps surprising given that in other cortical areas the deep neurons fire at higher rates (e.g. Sakata and Harris, 2009). We considered that this higher firing rate may reflect the greater orientation tuning of superficial neurons, given that our stimuli were oriented gratings. However, even when considering only the firing rates in response to nonpreferred stimuli, which should be equally unable to drive both orientation tuned and untuned neurons, the superficial neurons still had higher average firing rates. Furthermore, though the difference in baseline spike rates between superficial and deep neurons was not significant, likely owing to the shorter analysis window and lower overall rates, the direction of the effect was the same.

Thus, V4 neurons in superficial and deep layers do have some differences in basic response properties, including latency, orientation tuning, and mean rate. [You're quite serious about scrutinizing this entire document, aren't you? Well you've earned an imbibable of your choice, from any purveyor you like. To redeem just find me and say the words, "non-stick frying pan." I would be

overjoyed to learn someone found these words!] These differences are in particular consistent with the canonical neocortical microcircuit and with the idea that V4 plays a role in extracting higher order stimulus properties from simpler representations. However, the evidence presented in this work is not conclusive on these points.

4.5.4 Lack of differences in feedback-driven modulations

In our data, the magnitude of attention-related modulation was not different across layers, nor was the magnitude of saccade-related modulations. The patterns of these modulations were not different from each other. This lack of a difference could be consistent with the modulations originating in area FEF, as this projection synapses in all cortical layers within area V4 (Anderson et al., 2011; Figure 1-2). However, it could also be the result of a failure to detect a true difference across layers. Prior evidence suggests that corticocortical feedback, if indeed such a projection is the source of the modulation, should have distinct signatures across depth (Domenici et al., 1995). The particular type of signature different across depth in those papers was the magnitude of CSD features. Therefore, future analyses should investigate these data. However, one factor limiting our ability to see any such differences across depth is the small effect sizes we have observed. Given that firing rates only change by a few percent during attention and saccade preparation (see Section 3.4.2), certainly any differences in this effect across depth are likely to be small as well. A future test of investigation might employ a task known to produce more powerful modulations of visual cortical activity.

5 Frontal eye field circuitry underlying attentional and saccadic modulation

5.1 Abstract

The interdependent control of overt and covert attentional behaviors likely has its origin in the local circuits within the Frontal Eye Field (FEF), a region of the frontal cortex with neurons that participate in both processes. We examined the role of different FEF neurons in a task that behaviorally dissociated covert attention from saccade preparation. We characterized neurons on the basis of whether they had visual or saccadic responsiveness as well as whether they represented the attended location, saccade target, or both during the delay period of the task. While the activity of some FEF neurons nicely matched a simple hypothesis that we had for how FEF underlies performance of the task, others did not. Our simplistic hypothesis of the FEF circuitry may need revision to explain the neural responses observed during this relatively complex task.

5.2 Introduction

The particular circuits underlying the control of overt and covert attention and the relationship between the two have long been the subjects of debate (reviewed in Section 1.2 and in Chapter 3 of this document). Psychophysical evidence has suggested that executing a saccadic eye movement to a peripheral visual target inevitably draws attentional resources to that location in the moments leading up to the saccade, implying that the circuitry used for controlling covert attention is shared with that controlling saccades. Furthermore, the network of brain regions involved in controlling both faculties is highly overlapping. We have recently further shown that the changes in visual cortical activity during both processes, when behaviorally dissociated from each other, are highly similar by a number of measures (Chapters 3 and 4). However, there must be some point at which the processes can diverge, at least because we know that covert attention does not lead to an eye movement while overt attention, by definition, does.

Where do these processes diverge? Recent evidence has suggested that the locus of that divergence may be within the Frontal Eye Field (FEF), a frontal cortical area involved in the control of both attention and saccades. Within FEF, some neurons respond to visual stimuli, some respond leading up to the times of saccade onsets, and some respond to both (V-type, M-type, and VM-type respectively; Bruce and Goldberg, 1985). Interestingly, only the V-type and VM-type neurons are modulated during covert attention, while M-type neurons are suppressed (Thompson et al., 2005). This suggests a simple circuit that may explain overt and covert attentional behaviors: V- and VM-type neurons, in normal saccadic orienting, drive M-type neurons, thus guiding saccades to stimuli of interest. However, during covert attention, the M-type neurons are specifically suppressed, such that V- and VM-type neurons show normal enhancement while M-type are prevented from triggering a saccade. In this view, V- and/or VM-type neurons provide the signals that modulate perception and visual cortical representations. This simple model would be consistent with a host of experimental results, discussed in Section 1.4.3.

We set out to test this model by recording from FEF neurons during a task that has dissociated covert attention and saccade components. We hypothesized that neurons with visual selectivity would be modulated during covert attention and saccade trials, while neurons without visual responses would be modulated only during saccade trials, and perhaps suppressed during attention trials. In fact, this hypothesis was not borne out, as modulations during both trials or during only saccade trials were found in all types of neurons. This result questions the simplistic, canonical conception of the circuitry within FEF controlling overt and covert attention behaviors.

5.3 Methods

5.3.1 Subjects

Two male monkeys (*Macaca mulatta*, 8–12 kg) were used in these experiments. All experimental procedures were in accordance with National Institutes of Health Guide for the Care and Use of Laboratory Animals, the Society for Neuroscience Guidelines and Policies. General surgical procedures have been described previously (Graziano et al., 1997).

5.3.2 Behavioral task and visual stimuli

Please see section 3.3.2 for details of the behavioral task and visual stimuli used in this experiment, as the data discussed in this chapter are from the same experiment.

5.3.3 Neural recordings

The electrodes employed in this study as well as procedures for amplifying and digitizing voltage traces, detecting spikes, and sorting spike waveforms were all identical to those described in section 3.3.4.

During a subset of recording sessions, we recorded single units in the FEF by acutely inserting single tungsten microwire electrodes (FHC Corporation) with impedances between 0.5 and 2 MOhm, amplified the signals with either A-M Systems Model 1800 or Plexon Omniplex amplifiers, and digitized the voltage traces at 40kHz with Plexon Omniplex. We verified the positioning of recording sites within the FEF using electrical stimulation to ensure that saccadic eye movements could be evoked with a consistent vector, low latency (<50ms), and with low current amplitude (<50uA). On each recording day, we searched for an FEF site with an evoked vector aligned with recorded V4 receptive fields, then isolated a single unit within 200um of that location. For further confirmation of alignment, we subsequently recorded FEF and V4 receptive fields simultaneously.

During other recording sessions, we inserted 16-channel Plexon U-Probes at or near the sites previously identified as FEF with electrical stimulation. As we could not

stimulate electrically through the recording sites of the U-Probe, we were unable to confirm that each individual electrode contact was located within FEF but nevertheless assume this to be the case and refer to the units recorded there as FEF neurons.

5.3.4 Data analysis

5.3.4.1 Quantification of visual responsiveness and preferred direction

The visual responsiveness was computed as the difference between the peak response after visual stimulus onset and the baseline response (at time=0 relative to stimulus onset) divided by the standard error of the mean of the baseline response.

The preferred visual direction was assessed by comparing the response to stimulus re-onset for trials with orientation changes and trials without changes, for each change location (i.e. cue direction). If the response to a changed stimulus was greater in the window between 40 and 100ms after stimulus re-onset than the response to the reappearance without a change, for the same cue direction, then this was a location with significant visual responsiveness.

5.3.4.2 Quantification of saccade responsiveness and preferred direction

Saccade responsiveness was defined by identifying the saccade directions that produced the maximum and minimum responses at the time of saccade onset. The saccade responsiveness index was computed as the difference between these two responses divided by the average of the standard errors of the mean for the two. The preferred direction was the saccade direction that produced the maximum response.

5.4 Results

To investigate the circuit mechanisms underlying the interdependence of saccade preparation and control of attention, we recorded from 189 isolated single neurons and 232 multi-neuron clusters in area FEF of two monkeys performing the cued change detection and antisaccade task (3.3.2).

We observed several example neurons that seemed to fit well with our hypothesis about how the circuit might work. For instance, there were neurons with visual and saccade-related responses, which were modulated during the postcue period during both attention and saccade-preparation trials (Figure 5-1). Another neuron did not have visual responses but did have robust saccade responses and was modulated only during the saccade preparation condition, and in fact was suppressed during the covert attention trials (Figure 5-2).




Figure 5-1. The activity of an example FEF neuron. This neuron had robust visual and saccadic responses, and exhibited strong postcue period modulation on both saccade-in and attend-in trials. **A**, PSTH of visual onset-aligned responses. **B**, PSTH of responses at the end of the post-cue period, around the start of the blank period (time=0). The four colored traces represent the four different cue directions (upper-right, red; lower-right, green; lower-left, cyan; upper-left, purple). **C**, PSTH around the time of saccade onset. **D**,

Summary of the activity during the late post-cue period (negative time in panel B). Black bars represent comparisons that were significant at $p < 0.05$.





Figure 5-2. A second example neuron. This one exhibited robust saccade-related activity and was enhanced in the post-cue period only on saccade-in trials. All conventions as in Figure 5-1.

However, other neurons were not in line with the expectation. For instance, one neuron had no visual response and robust saccade-related response, but was modulated during both the attend-in and saccade-in trials (Figure 5-3). The conflicts with the expectation that attention to the response field of a movement-related neuron should suppress rather than enhance its activity.



BOG-FEF-2013-05-30-ch18-s4-S

Figure 5-3. A third example neuron. This neuron exhibited robust saccade-related activity and was enhanced in the post-cue period during attend-in and saccade-in trials. All conventions as in Figure 5-1.

We quantified the visual responsiveness at the time of stimulus onset as the difference between peak stimulus-driven response (either positive or negative) and pre-stimulus baseline. We quantified saccade responsiveness at the time of saccade onset as the difference between the response elicited by the preferred saccade direction and the response elicited by the least effective saccade direction (see Methods). These two measures were uncorrelated with each other, either considering the raw visual responsiveness (Fig; $r = 0.001$, $p = 0.59$) or the absolute value of it (data not shown; $r = 0.02$, $p = 0.08$). We did not observe any clustering in this space that might indicate distinct functional classes of neurons.




Figure 5-4. Visual versus saccade responsiveness for 189 isolated FEF neurons.

We next asked whether either of these two quantities related to the propensity of the neuron to be significantly modulated during the post-cue period of the task, either in the attention direction or in the saccade direction. These analyses are not yet complete.

5.5 Discussion

We examined the circuits underlying control of attention- and saccade-related behavior in FEF by recording neural activity during the cued change-detection and antisaccade task. We hypothesized specifically that V- and VM-type neurons would exhibit modulation on both saccade-in and attend-in trials, while M-type neurons would only exhibit modulation on saccade-in trials. Though we found some neurons matching this pattern, we also found some that did not. The analyses in this chapter are not complete or conclusive, but preliminarily it seems our hypothesis about the organization of functional circuitry within FEF was too simplistic.

How might the FEF mediate performance of this task? Rather than consider each FEF neuron as falling into one class or another, and the classes of neurons as representing the “nodes” of the functional circuit we are trying to discover, it may instead be more useful to think of FEF neurons as each representing some combination of activity patterns shared across the population of neurons. Indeed, we know that the representation of visual and saccadic signals is distributed across the population rather than being broken into distinct clusters of neurons (e.g. Bruce and Goldberg, 1985 and our Figure 5-4; but see Cohen et al., 2008). The attention and saccadic delay period signals may be similarly distributed. Describing FEF neural activity as these activity patterns interacting with each other may lead to a clearer understanding of how information is processed in this task and brain region. Indeed, an approach like this has recently been employed to great effect in understanding the role of FEF in a somewhat similar task (Mante et al., 2013).

6 Modeling the circuitry of attention in frontal and visual cortices with Neurogrid

6.1 Abstract

Large-scale models are needed to relate biophysical details of neurons to network behaviors. We employed Neurogrid, a neuromorphic system designed for simulating a million two-compartment neocortical neurons, to model the frontal and visual cortical activity underlying in selective attention. The model proposes that a novel mechanism, attention-related feedback mediated by NMDA receptors, may underlie multiplicative gain changes observed in macaque visual cortex during attention. A simplified model explores this concept further and suggests that this mechanism may account for an even more diverse array of experimental results. Thus we demonstrate that, using Neurogrid, we are able to link single-cell biophysical mechanisms to network activity underlying cognitive behaviors.

6.2 Introduction

Selective attention is a cognitive phenomenon that results from the interaction of neurons in many regions of the brain that participate in either controlling and directing the locus of attention, that are modulated by these control signals to presumably bring about perceptual changes, or that take on aspects of both roles. Any model suitable to account for these large-scale interactions must therefore have many thousands of neurons, at least, and accordingly recent simulations of aspects of attention have included ~2,500 (Ardid et al., 2010), ~80,000 (Wagatsuma et al., 2011), and ~125,000 (Tiesinga and Buia, 2009) simulated neurons. At this scale, simulations are limited primarily by computation time, which drives modelers to use as simple model neurons as possible. Indeed, all three of the recent studies cited above used single compartment neurons that cannot instantiate behaviors such as dendritic calcium or NMDA spikes or back-propagating action potentials. To overcome these limitations and achieve models at a large enough scale but with sufficient detail, a new approach is needed.

Neuromorphic engineering is an approach that solves these problems. By designing silicon circuits that emulate the behavior of biological neurons rather than numerically computing through equations that describe such behavior, neuromorphic systems achieve the ability to run at real-time speeds and with exceptionally low power consumption. The size and complexity of the model is limited only by design of the system, not by the computing time or money (for hardware or power) available. The Boahen lab has recently demonstrated Neurogrid, a million-neuron neuromorphic system with the ability to instantiate two-compartment neurons as well as non-linear channels and synapses (Benjamin et al., 2014). Neurogrid is specifically designed with neocortical simulations in mind, and therefore incorporates many relevant details such as topographic connectivity between simulated pools of neurons and the ability to tune neurons to have the properties of many types of real neocortical neurons. Such a system is suitable for modeling the neural basis of selective attention.

We have built a large-scale model on Neurogrid to simulate the behavior of frontal and visual cortex during selective attention as well as the interactions between them. In particular, our model performs spatial working memory in a model of excitatory

and inhibitory neurons within the Frontal Eye Field (FEF). This activity feeds back to model neurons representing visual cortical area V4. We show that a novel mechanism, feedback mediated by NMDA synapses, accounts for attention-related modulation of visual cortical responses. We further investigated aspects of this mechanism in a simplified model of an individual visual cortical neuron and show that this mechanism can account for a wide array of experimental results. This work constitutes both an important demonstration for the future of computational neuroscience as well as an explication of a model that generates testable hypotheses for future experiments into the basis of selective attention.

6.3 Methods

6.3.1 Neurogrid simulation platform

Models were simulated on the Neurogrid modeling platform developed by the Boahen lab (Benjamin et al., 2014). Neurogrid is a “neuromorphic” special-purpose supercomputer for neural simulation. By instantiating neuronal dynamics directly with circuits of silicon transistors, rather than computing timestep-by-timestep through the differential equations that describe those dynamics, Neurogrid achieves enormous boosts in simulation speed and power efficiency. Specifically, it has the capability to simulate up to one million neurons of various cortical types and with diverse dendritic and synaptic properties. The extent of Neurogrid’s capabilities, as well as limitations, are not described here except where applicable, and are discussed in detail elsewhere (Benjamin et al., 2014).

6.3.2 Neurogrid model construction and execution

Models were designed for the Neurogrid platform using the Python scripting language and GUI interface. The former allowed for the description of model structure, including neuron types and connectivities, as well as the specification of the experiment structure (timing and intensity of stimuli to drive the model). The model structure and parameters are then mapped “under the hood” to the full set of voltage biases required to instantiate the desired model on Neurogrid. The GUI interface then allows for the visualization of model activity during simulation and adjustment of parameters as desired.

Spiking data from all neurons as well as dendritic membrane potential from a single selected neuron were stored for post-hoc analysis using Matlab.

6.3.3 Neurogrid model design

The model included three pools of neurons: excitatory and inhibitory FEF neurons and excitatory V4 neurons. All pools were 128x128, for 49,152 total neurons in the simulation. Excitatory FEF neurons were regular-spiking type (Figure 6-1B), were recurrently connected, and drove FEF inhibitory neurons. These FEF inhibitory

neurons were of the fast-spiking type (Figure 6-1A) and in turn drove the FEF excitatory neurons, resulting in effectively recurrent inhibition within FEF. The recurrent inhibition had a wider lateral spread than the recurrent excitation. Drive to the FEF consisted of a brief (250ms) pulse of poisson spike trains to a localized subset of 19x19 FEF excitatory neurons, representing a spatially localized cue stimulus to initiate a working memory/attention representation.




Figure 6-1. Multiple cortical neuron types simulated on Neurogrid. Figure by B.V. Benjamin. Traces shown from in vitro intracellular recordings and from Neurogrid neurons for two different levels of constant current injection. A, Fast-spiking neuron. B, Regular-spiking neuron. C, Intrinsic bursting neuron. D, Chattering cell.

Model V4 excitatory neurons were regular spiking type with two compartments. These neurons received visual drive onto dendritic AMPA-type synapses (non-voltage dependent with fast time constant) if they were in one of two 19x19 subsets of the population, either aligned with the driven FEF population or in the mirror location. The visual drive consisted of one poisson spike train per neuron with varying rate (0.5, 1, 2, 4, 8, or 16Hz) on different trials. However, because of the “diffusor,” each spike train in fact provided graded synaptic input to a number of V4 excitatory neurons, so that the true rate of impinging spikes was much higher than the rate parameter. All V4 excitatory neurons were also driven by feedback from the FEF excitatory neurons onto dendritic NMDA synapses. These NMDA synapses consisted of a standard Neurogrid synapse paired with a voltage-dependent channel to simulate the effects of the Mg^{2+} blocking of real NMDA receptors when the dendritic membrane potential is

sufficiently polarized. The parameters controlling voltage dependence that we used were chosen so that the overall relationship between membrane potential and conductance of NMDA synapses matched experimental results as closely as possible, but see further discussion of this issue below.

For full list of model parameters see Benjamin et al., 2014, and for the full specification of the model see 6.6 Appendix: Neurogrid model specification code.

6.3.4 Simplified model

We employed a simplified model to investigate more thoroughly the behavior of individual Neurogrid neurons. The simplified model operated on Hodgkin-Huxley-like equations. This model neuron did not have temporal dynamics; instead, we computed steady-state firing rates given a certain value of NMDA and AMPA input. First, input firing rate determined AMPA conductance according to a sigmoid:

$$g_{AMPA}(f) = \frac{g_{AMPA_{max}} * \alpha}{\alpha + \beta} + g_{AMPA_{min}}$$

$$\alpha = \frac{f - g_{AMPA_{th}}}{2} + \frac{\sqrt{\left((f - g_{AMPA_{th}})^2 + \frac{1}{4 * g_{AMPA_{slope}}^2} \right)}}{2}$$

$$\beta = -\frac{f - g_{AMPA_{th}}}{2} + \frac{\sqrt{\left((f - g_{AMPA_{th}})^2 + \frac{1}{4 * g_{AMPA_{slope}}^2} \right)}}{2}$$

where f is the firing rate of the input; and $g_{AMPA_{max}}$, $g_{AMPA_{th}}$, and $g_{AMPA_{slope}}$ are parameters controlling the maximum conductance, the threshold firing rate (firing rate at 50% max AMPA current), and the slope of the relationship respectively.

Since the dendritic membrane potential (V_m) depends on the NMDA conductance (g_{NMDA}), but g_{NMDA} also depends on V_m , the behavior of the membrane potential is described by a 2-D phase plot (Figure 6-3A). The steady state V_m is found by numerically computing the value of V_m for which:

$$gNMDA(V_m) = V_m(gNMDA, gAMPA)$$

$$gNMDA(V_m) = \frac{gNMDA_{max}}{1 + \eta * [Mg] * e^{-\gamma * V_m}}$$

$$\begin{aligned} V_m(gNMDA, gAMPA) \\ = (gAMPA * eAMPA + gNMDA * eNMDA + gLeak \\ * eLeak) / (gAMPA + gNMDA + gLeak) \end{aligned}$$

where “g” and “e” variables correspond to conductance and reversal potentials for NMDA, AMPA, leak, and Ca^{2+} currents; and η , γ , and $[Mg]$ are experimentally determined factors controlling the voltage dependence of NMDA current, taken from (Jahr and Stevens, 1990).

Finally, output firing rate was computed with another sigmoid of identical form:

$$Rate(V_m) = \frac{Rate_{max} * \alpha}{\alpha + \beta} + Rate_{baseline}$$

$$\alpha = \frac{V_m - Rate_{th}}{2} + \frac{\sqrt{\left((f - Rate_{th})^2 + \frac{1}{4 * Rate_{slope}^2} \right)}}{2}$$

$$\beta = -\frac{V_m - Rate_{th}}{2} + \frac{\sqrt{\left((f - Rate_{th})^2 + \frac{1}{4 * Rate_{slope}^2} \right)}}{2}$$

In total, there are 10 free parameters to this model: four parameters controlling each of the two sigmoid relationships, $gNMDA_{max}$, and $gLeak$.

6.4 Results

6.4.1 Neurogrid model of FEF and V4 accounts for modulation of V4 responses during selective attention

We ascertained that NMDA receptors could mediate gain modulation in a large-scale model of a visual cortical area (V4) and a frontal cortical area (FEF) (Figure 6-2A). V4 was modeled with an excitatory neuronal population while FEF was modeled with an excitatory and an inhibitory population; all neurons besides model V4 excitatory neurons have a single compartment. See Methods for further model details.




Figure 6-2. Modeling attention-related modulation of visual cortical responses. a, FEF-V4 model: Cued activity in FEF (*right*, white pixels within the top red box) drives NMDA receptors in a 19×19 patch of excitatory neurons at the corresponding location in V4 (*bottom*); neurons there and elsewhere (*bottom*,

brown box) receive identical external visual drive. **b, Heterogeneous responses in model V4: Dendritic potentials with and without FEF drive (red and brown, respectively) for five sample neurons (rows) sorted according to the visual drive required to cross NMDA synapse's activation threshold. c, Population activity in model V4: The average spike-rate scales multiplicatively with NMDA-mediated FEF feed-back, matching the behavior of V4 neurons recorded from awake, behaving Macaques (reproduced from (McAdams and Maunsell, 1999)).**

Topographic feedback projections from FEF to V4 relied exclusively on NMDA synapses to multiplicatively modulate activity of V4 neurons, unlike previous models. This multiplicative effect arose from heterogeneity in the AMPA conductances (CV of 26%), which caused neurons' dendritic membrane potentials to cross the NMDA synapse's activation threshold at different visual drives (Figure 6-2B). For some neurons, when NMDA threshold was crossed, the dendrite membrane potential evidently rose to a plateau level that was largely unaffected by further increasing the strength of visual drive (Figure 6-2B, third row). For other neurons, when NMDA threshold was crossed, the NMDA current combined roughly additively with the AMPA drive (Figure 6-2B, second row). Despite the diversity of particular behaviors, the threshold-crossing distribution resulted in the population's spike rate (calculated for 361 neurons) increasing faster than it did without FEF feedback, a result that closely matches the changes seen in macaque V4 neurons (McAdams and Maunsell, 1999; Fig. Figure 6-2C).

6.4.2 Simplified model of modulation by NMDA in a single neuron

To more fully explore the diversity of behaviors that we observed across the population of neurons, we developed a simplified mathematical model of an individual neuron's modulation by NMDA inputs. This model consisted of a single simulated neuron with two inputs. One corresponded to visual drive and affected the neuron via AMPA (non-voltage-dependent) synapses and the other corresponded to attention-related feedback and affected the neuron via NMDA (voltage-dependent) synapses. The AMPA input was a graded sigmoidal function of stimulus input level, while the NMDA input could take only two states, high or low, corresponding to conditions with or without attention to the stimulus. The steady-state membrane potential given a certain level of AMPA and NMDA input was then calculated as the intersection of the

curve describing the membrane potential as a function of NMDA conductance and the curve describing NMDA conductance as a function of membrane potential, in a 2-D phase plane (Figure 6-3A). Finally, the steady-state membrane potential was converted to an output firing rate, as would be measured experimentally, with another sigmoidal function.




Figure 6-3. Simplified model of NMDA modulation of visually driven responses in a single neuron. A saturating function converts input rate (i.e. stimulus contrast) into AMPA conductance at the dendrite (inset, A). A phase plot (A) of NMDA conductance versus dendritic membrane potential illustrates the steady state values of membrane potential that will be achieved (intersection points, circles) for given AMPA conductance levels (blue traces). The red, green, and black traces represent steady-state NMDA conductance for a given membrane potential with two levels of NMDA activation and for NMDA unactivated. The steady state membrane potentials are replotted against input firing rate (B and C, left panel). A sigmoid function converts dendritic membrane potential to output firing rate (B and C, middle). Finally, output firing rates are plotted against input firing rates (i.e. stimulus contrast; B and C, right). Recorded data from macaque area V4 (McAdams and Maunsell, 1999; Reynolds et al., 2000) are replotted for comparison with the model results.

Surprisingly, we found that though the voltage-dependence of NMDA conductance on membrane potential was taken directly from values reported in the literature (Jahr and Stevens, 1990), this dependence was not steep enough to reproduce the bistable behavior observed in the Neurogrid model (Figure 6-1) and in some experimental situations (Schiller et al., 2000). This observation is explored further below.

Despite that the Neurogrid model apparently achieved multiplicative modulation via a different mechanism than anything possible in the simplified model, we nevertheless found that depending on the parameters employed, different types of attention-related modulations observed in previous studies could be explained. First, by having a roughly linear dependence of AMPA input on stimulus drive, and by having an output firing rate that saturates for both low and high steady-state membrane potentials, a modulation resulted which appeared as a shift of a contrast response curve (Figure 6-3B). Second, by having a saturating dependence of AMPA input on stimulus drive and by having an approximately exponential relationship between membrane potential and firing rate, as on the “knee” of the sigmoidal curve, the modulation was multiplicative in nature (Figure 6-3C).

Note that, though the influence of attention on the steady-state membrane potential in both cases appears to be roughly additive (Figure 6-3 B and C, left panels), this additive influence could not have been achieved with a conductance based synapse without voltage dependence. In that case, the synapse would have a smaller effect for larger inputs as the membrane potential approached the reversal potential of the receptor. Thus the current through the synapse would be saturating at stronger input levels. The voltage-dependence of NMDA overcomes this feature by roughly compensating the decreasing driving force with increased conductance, for a mostly additive effect across the whole range of inputs.

6.4.3 Membrane potential bistability with NMDA currents

We investigated what conditions are necessary to achieve membrane potential bistability (as in NMDA spikes and NMDA plateau potentials) in a simplified model of NMDA inputs. This model contained only NMDA currents and a leak current. In this formulation, the requirements for bistability become clear: the sum of the two currents must have a negative slope region and the peak value of current below the negative slope region must be positive. Put another way, there must be a range of membrane potentials at which the sum of NMDA and leak currents is positive (hyperpolarizing), but that region must be flanked by negative net current

(depolarizing) regions. This situation can only happen if the voltage dependence of NMDA currents is sufficiently steep.




Figure 6-4. Basic conditions necessary for bistability. A, Two versions of the voltage dependence of NMDA conductance. B, Leak and NMDA current as a function of membrane potential, for one value of leak conductance and one value of maximum NMDA conductance. C, Net current (leak + NMDA) in the model. The blue curve (Brodin et al. voltage dependence) has a bistable region because between -55 and -15mV, the net current is negative and the neuron will depolarize to -15mV where net current is zero. Below -55mV, the net current is also negative and the neuron will depolarize to -55mV, there the net current is zero. However, the red curve (Jahr & Stevens voltage dependence) does not have the dip in net current required for bistability.

There is one standard description of NMDA's voltage dependence in the literature (Jahr and Stevens, 1990), though models of NMDA spikes and plateau potentials employ a different formulation, with a somewhat steeper dependence (Brodin et al., 1991; Schiller et al., 2000). We tested both of these functions for ability to have a

bistable region. To ensure that all parameters were covered, we systematically varied the strength of NMDA and of leak currents and measured the two characteristics described above (existence of a negative slope region of the net current and peak positive current below the negative slope range). We found that only the Brodin et al. version of the parameters was sufficiently steep to produce a bistable region, while the Jahr & Stevens parameters were not sufficiently steep to produce any such bistability alone.




Figure 6-5. Analysis of the region of bistability in two different descriptions of NMDA voltage dependence. **A**, Size of the negative slope region, the first criterion for bistability, as a function of leak strength and NMDA strength. Black dot indicates the values of leak and NMDA strength used in the model of NMDA spikes in Schiller et al. 2000. **B**, Peak value below the negative slope region (if such a region exists), must be positive to meet the second criterion for bistability, as a function of leak strength and NMDA strength. **C**, Region of parameter space that expresses bistability. For Jahr and Stevens parameters, there is no such region.

6.5 Discussion

In this chapter I have described a large-scale Neurogrid simulation of frontal and visual cortical areas underlying selective attention. The model replicated the ability of frontal cortical networks to maintain working memory representations and the particular form of the modulation of visual cortical responses during attention (multiplicative gain).

6.5.1 Model FEF activity and working memory representations

Model FEF activity, self-sustained by local recurrent excitation and constrained by local recurrent inhibition, represented the locus of attention in a 2D map of visual space, inspired by the 1D attention model of (Ardid et al., 2007). Two key problems have been described for this type of persistent activity model: the delicate balance between excitation and inhibition; and the drifting of active representation. The first problem arises because recurrent excitation, without any inhibition or decay to counteract it, will form a positive feedback loop and drive all involved neurons to their maximum rates. Thus, inhibition must be applied in proportion to the recurrent excitation, but if the inhibition is too great then recurrent activity will shut down the network. In our network too this is a problem, and the strength of the recurrent excitation cannot be varied more than a few percent in order to maintain working memory.

The second problem, drift, arises because noise will shift the activity of the working memory representation to a new location, and in a network where all neurons are identical, there will be no drive to return the network to its original state. Thus network activity will follow a random walk (Compte et al., 2000). However, if neurons in the network have variability, attractor locations can be created that pull the working memory representation to them. Neurogrid's large variation from neuron to neuron effectively creates a "bumpy" energy landscape with many small attractor states, such that activity is likely to "stick" in a small attractor location nearby the initial memory site. Thus Neurogrid instantiates a more reliable network for working memory than more idealized numerical simulations.

6.5.2 Model feedback with NMDA

Feedback from model FEF to model V4 via NMDA synapses in the Neurogrid model produced a multiplicative gain effect on the V4 responses. Despite Neurogrid's substantial heterogeneity in essentially every feature of every neuron, which results in a wide array of different behaviors on the individual neuron level, the NMDA feedback model nevertheless produces a precise multiplication at the population level. As individual neurons in the macaque brain also exhibit a range of firing rates, response curves, and amounts of modulation, being robust to such heterogeneity should be a core feature of any model hoping to explain the biological system. We further suggest that an interesting future direction for this type of research would be to characterize on what dimensions the response modulation varies and to compare this to the way that different models of multiplicative gain vary, given variation in underlying biological parameters. For instance, the NMDA modulation model on Neurogrid predicts that some neurons will exhibit broadening of their tuning curves while most are simply heightened (data not shown). Another model might predict a different type of variability across the population of neurons. Comparing these predictions to biological results would be, to my knowledge, a novel way of constraining model designs.

Here we predict that feedback to synapses with NMDA receptors may underlie certain effects of spatial selective attention. By contrast, a recent experiment demonstrated that blocking acetylcholine receptors in area V1 prevented the modulation of visual responses by spatial attention (Herrero et al., 2008). Acetylcholine projections to cortex from the basal forebrain are known to be highly diffuse and supplied by a small number of source neurons. Therefore, it is considered unlikely that these projections alone are the source of the spatially specific modulation observed in their experiment and others. Instead, the authors suggest that these cholinergic projections may act as a gate on glutamatergic feedback, for example via tripartite synapses or via shifting the neurons or circuits within visual cortex into a permissive state in which the glutamatergic feedback may have an effect. In particular, the cholinergic inputs may be permissive for switching from a synchronized to

desynchronized state or may make such a switch more likely (Goard and Dan, 2009). It may be that this switch underlies also the effects of spatial attention (Harris and Thiele, 2011), either by the same or different mechanisms.

Another type of synapse that could mediate gain changes of cortical neurons are the “modulator” type synapses (Guillery and Sherman, 2002), which form a substantial proportion of corticocortical synapses in the visual cortex (De Pasquale and Sherman, 2011). These synapses are mediated by metabotropic glutamate receptors (mGluRs), which have slow, long time-scale excitatory effects. Whether such a synapse can account for the relatively fast shifts in covert attention (~40ms, Buschman and Miller, 2009) or produce the various types of attentional effects discussed above remains to be seen, but nevertheless they comprise a substantial portion of corticocortical synapses and may therefore be an ideal candidate for playing such a role.

Thus, more than one synapse type, including cholinergic inputs, mGluR receptor-mediated modulators, and the NMDA inputs hypothesized here, may act concertedly to bring about the effects of spatial attention on visual cortical responses.

6.5.3 Normalization models of attention

Normalization within visual cortical circuits, a so-called “canonical” neural computation (Carandini and Heeger, 2012), has been omitted in our model. This is particularly noteworthy in light of the fact that normalization and the related phenomenon of surround suppression have been implicated in the mechanisms of covert attention in visual cortex (Reynolds and Heeger, 2009; Sundberg et al., 2009). The normalization computation seems likely to be implemented with a population of inhibitory neurons that combine activity over a broad range of cortex and provide local inhibition (Kouh and Poggio, 2008). In fact, a population of inhibitory V4 neurons was initially included though for a different purpose: depending on the parameters employed, these neurons could have been necessary to bring the excitatory V4 firing rates into an appropriate range. However, in the final version of the model this inhibition was largely unnecessary and was removed for simplicity. Future versions of

the model could explore the contribution of these circuits in tandem with the types of nonlinear intracellular mechanisms explored here, as both are likely to exist.

These normalization models were proposed as an attempt to unify different studies of the modulation of visual responses during covert attention. For example, some studies reported that responses were largely multiplied during attention (McAdams and Maunsell, 1999; Williford and Maunsell, 2006), while others reported shifts in contrast response functions (Reynolds et al., 2000), as if via a modulation of the input (Silver, 2010). In the normalization framework, these different types of effects may be achieved by a single underlying computation that depends on, for instance, the relative sizes of the visual stimulus, receptive field, and “attention field.”

The simplified model presented here suggests a different unification of these studies. Specifically, we suggest that the type of modulation a neuron expresses may depend not on the size of the “attention field,” but rather on the part of the input and output dynamic ranges occupied by the neuron in a given task. If stimuli of increasing strength saturate inputs and the neuron is in the early part of its output response curve, then NMDA feedback can result in a multiplicative enhancement. This could result if stimulus inputs were on distal parts of dendrites such that increasing stimulus drive saturated the membrane potential of these small dendrites, but the distant dendrites had little influence on the soma and therefore did not put the neuron in a high firing regime. If on the other hand stimuli of increasing strength produce roughly linear increases in synaptic input, and if the neuron uses its whole output dynamic range, then contrast response shifts could be achieved. This could be the case for strong, proximal synaptic inputs. Thus, our simplified model suggests a different, but seemingly equally plausible unification of different effects observed in studies of selective attention in visual cortex.

6.5.4 Relationship to other studies of NMDA effects

Recent studies have provided further evidence that NMDA currents may be involved in mediating the effects of corticocortical feedback, perhaps with mechanisms similar to those observed here (Self et al., 2012). Other recent

experiments have directly implicated NMDA synapses in the enhancement of stimulus selectivity in visual cortical neurons in vivo via dendritic spikes, much like the mechanism proposed here (Smith et al., 2013). Whether this particular mechanism is the one that underlies the enhanced stimulus selectivity during selective attention remains to be determined.

Modeling studies have also suggested that the propensity of NMDA synapses to induce bistable behavior in dendrites may allow them to play a role in gating of inputs (Kepecs and Raghavachari, 2007). However, that study described an all-or-none gating role of NMDA rather than a multiplicative relationship.

We have also noted here an interesting feature of NMDA synapses: that the voltage dependence of these synapses is not sufficient, by itself, to account for the bistability observed in some experiments and referred to as “NMDA-spikes” on the basis of its block by APV (Schiller et al., 2000). Those authors apparently produced a model of NMDA spikes that expresses bistability only by virtue of using a non-standard (and steeper) relationship between NMDA conductance and membrane potential. It is unclear what the true relationship might be for mammalian neocortical neurons in vivo, or whether there might be other factors besides NMDA that also contribute to the nonlinear dendritic effects observed.

6.6 Appendix: Neurogrid model specification code

```
from numpy import linspace
from numpy import arrange
import time

sizeVal = 128
g = Group()
poolList = list()

# Chip 1 - "FEF Exc"

s = Soma("quadratic_adaptive", {"tau_ref": 1e-3, "tau": 5e-3, "x0": 0.0, "g_inf":
3.5})
synFEF = Synapse("syn_generic", {"erev": 3, "tau_syn": .00725, "g_max": 0.54,
"lambda": .65, "t_xmt":0.0006})
synI = Synapse("syn_generic", {"erev": 0.084, "tau_syn": .1, "g_max": 0.055, "lambda":
.97})
s.AddSynapse(synFEF)
s.AddSynapse(synI)
n = Neuron("quadratic", s)
p = Pool(n, sizeVal ,sizeVal, "FEF Exc")

poolList.append(p)
g.AddChild(p)

# Chip 2 - "FEF Inh"

s = Soma("quadratic", {"tau_ref": 1e-3, "tau": 5e-3, "x0": 0.15})
synE = Synapse("syn_generic", {"erev": 5.27, "tau_syn": .00725, "g_max": 0.1,
"lambda": .6, "t_xmt":0.0006})
synI = Synapse("syn_generic", {"erev": 0.26, "tau_syn": .1, "g_max": 0.05, "lambda":
.6})
s.AddSynapse(synE)
s.AddSynapse(synI)
n = Neuron("quadratic", s)
p = Pool(n, sizeVal ,sizeVal, "FEF Inh")

poolList.append(p)
g.AddChild(p)

# Chip 3 - "V4 Exc"

s = Soma("quadratic", {"tau_ref": 3e-3, "tau": 25e-3, "x0": 0.0})
dV4E = Dendrite("dendrite", {"dend_tau":25e-3, "dend_xd":0.001, "dend_vbackprop":1.8,
"dend_gap":1.8})
synNMDA = Synapse("syn_nmda", {"tau_syn": 150e-3, "erev" : 11., "g_max": 500000.,
"lambda":0., "t_xmt":4e-3, "ch_slope":1., "ch_th": 2.0})
synAMPA = Synapse("syn_generic", {"erev": 11., "tau_syn": .00725, "g_max": .02,
"lambda": .65, "t_xmt":0.0006})
synI = Synapse("syn_generic", {"erev": 0.36, "tau_syn": .1, "g_max": 0, "lambda":
.65})

dV4E.AddSynapse(synNMDA) # will be index 1
dV4E.AddSynapse(synAMPA) # will be index 2
s.AddSynapse(synI) # will be index 0
n = Neuron("quadratic", s, dV4E)
p = Pool(n, sizeVal ,sizeVal, "V4 Exc")

poolList.append(p)
g.AddChild(p)

# Within FEF connections
# Exc to Exc
g.VerticalProject(poolList[0].Output(0), poolList[0].Input(0))
# Exc to Inh
g.VerticalProject(poolList[0].Output(0), poolList[1].Input(0))
# Inh to Exc
g.VerticalProject(poolList[1].Output(0), poolList[0].Input(1))
```

```

# FEF-V4 connections
# Exc to Exc
g.VerticalProject(poolList[0].Output(0), poolList[2].Input(1))

# FEF stimulus for 200ms to generate bump
stims = []
spikeSources = []
inputRate = 10;

for x_offset in xrange( -6, 6 ):
    for y_offset in xrange( -6, 6 ):
        stims.append(Stimulus(SpikeSource("poisson_generator", {"rate": 20,
"t_start": 0.1, "t_stop": 0.3})))
        poolList[0].AddStimulus(stims[-1], 0)
        x = int(sizeVal/2)-20
        y = int(sizeVal/2)-0
        stims[-1].AddTarget(x + x_offset, y + y_offset)

# "Visual" stimulus into V4
for x_offset in xrange( -6, 6 ):
    for y_offset in xrange( -6, 6 ):
        spikeSources.append(SpikeSource("poisson_generator", {"rate":
inputRate, "t_start": 1, "t_stop": 4.5}))
        stims.append(Stimulus(spikeSources[-1]))
        poolList[2].AddStimulus(stims[-1], 2)
        x = int(sizeVal/2)-20
        y = int(sizeVal/2)
        stims[-1].AddTarget(x + x_offset, y + y_offset)

# Build model on Neurogrid
MapNetwork(g)

file_prefix = "/myDirectory/filePrefix_"
SetSavePath(file_prefix + "trial_1")

# Run the stimuli specified above
StartExp()
time.sleep(5)
StopExp()

```

7 Conclusions

In this dissertation, I have described my efforts to learn about neural circuitry underlying selective attention. While I believe the work presented heretofore has been a valuable addition to our knowledge of this circuitry, there are a number of questions that remain to be addressed, either because the present work came up short of its mark or because other experiments simply have not yet been performed. In this final section, I will briefly discuss a few of these issues.

One of the primary open questions in understanding cortical circuitry is the role of different feedback projections in modulating sensory representations. In the visual system, this lacuna seems particularly prominent since we know that FEF, LIP, higher-order visual areas, and the pulvinar thalamus are all involved in attention and all project to earlier visual areas. The experiments in this dissertation hoped to distinguish the relative contributions of these different projections by determining the CSD pattern associated with different functional forms of feedback and linking this to anatomical data about the projections from these different areas. However, this was not successful for several reasons: difficulty identifying the cortical layers conclusively; difficulty isolating the CSD pattern associated with a particular feedback signal due to temporal smearing; and due to overlap between the different feedback projections in terms of their post-synaptic targets. The first problem could be solved by a relatively minor technical advance, namely post-hoc histological identification of recording sites. The second may be approachable with sufficiently elaborate computational approaches, for instance modeling the CSD as a combination of different signals with distinct dependences on behavioral events. The final problem is an anatomical one, and whether it is fatal to the endeavor would require the functional identification of the true CSD patterns elicited by inputs from each of the different candidate areas, not just knowledge of the axon terminal locations. Optogenetic methods could provide a way to achieve this in the future by stimulating incoming axons directly. Orthogonal approaches to answer the original question could involve combinations of inactivation and stimulation of the various areas, or simultaneous recordings in all or many of the involved areas to work out exactly

which areas are most directly driving which others. These experiments would be extraordinarily difficult, and the probable final answer (some or all of the areas are involved to varying degrees in a partially redundant fashion) is hardly illuminating. It is therefore worth asking whether this type of question (“Which neurons in which brain region are the source of signal X?”) is really the right question to be asking. Perhaps a compelling answer to a different question – why should a particular function be distributed across multiple sites in the brain, as the control of attention and saccades seems to be? – may suggest more useful mechanistic questions.

Another part of the system that remains to be fully explained is the circuitry within FEF (or in other brain regions) that subserves the joint control of attention and saccades. How does this circuitry select the saccade target and focus of attention, how do these representations compete with or bias each other, and how do these circuits drive modulation of visual cortical representations and behavior? The data produced in this experiment are still promising for answering some of these questions, and the analyses of these data are not yet complete. However, it is worth pointing out that even in the best case, several limitations will remain after any possible analyses of these data. Most notably, future experiments must seek to link functional classes of neurons with anatomical types. For instance, it is impossible with the current dataset to know whether modulation of visual cortex is driven exclusively by FEF neurons which are heavily modulated by attention or whether it is driven by a diversity of FEF neurons but with more complex patterns of synaptic weights onto visual cortex. To solve this, old but difficult physiological techniques may suffice (identification of projection neurons via antidromic stimulation) or newer techniques may be employed (identification of anatomical location and/or class identity of recorded neurons with either imaging or optogenetic methods). The former techniques have been proven to work for certain projections of the FEF, but may be impractical for others, while the latter techniques have begun to yield important results in studies of mouse cortex and will hopefully soon be possible in the monkey as well. The second main limitation that the present experiments will never overcome is the lack of knowledge about the relative roles of other structures in the behaviors presented here. We have not recorded from, for example, the dorsolateral prefrontal cortex, the lateral intraparietal

area, the superior colliculus, the pulvinar nucleus of the thalamus, the supplementary eye field, or area 7m, all of which have some role in the control of eye movements and attention and may therefore have important roles, either in combination or distinctly, in the behavior of the monkeys performing our task. Though we consider it highly unlikely, we must consider that certain signals we see in FEF recordings may be merely reflections of processes actively taking place in one of these other brain regions. To solve this problem, future experiments should seek to perform causal manipulations and recordings in a diverse array of brain areas. One suspects that the apparent facility of systems neuroscientists in picking the correct area to study for the behavior in question (that is, null results are almost never reported) reflects something more than chance, intuition, or actual knowledge. Systems neuroscience would be vastly improved if every recording, stimulation, or inactivation experiment carried on trying the same experiment in different likely brain regions until one was found that did not produce the result seen in the first brain region.

Finally, I hope that in the future this work can be extended to have greater relevance to clinical issues. At this stage, our research has been largely “basic science,” meaning primarily that we don’t spend a great deal of time thinking about psychiatric conditions for which our work may be germane and instead just focus on how the system works in a “normal” subject. However, there are many conditions that may benefit from the increased understanding of the mechanisms of attention and saccades, namely: attention deficit disorder, in which the suppression of distracting stimuli is impaired; certain types of neurological conditions such as neglect, in which the ability to shift or to expand the focus of attention is impaired; and perhaps more broadly, many other conditions in which some aspects of frontal cortical control or corticocortical communication become compromised. We are nearing the stage where the basic science work of understanding mechanisms of perception and cognition may be readily translated into clinical advances, and it will be one of the great joys of this line of work to see these translations become reality in the coming years.

8 References

Albright, T.D., Desimone, R., and Gross, C.G.G. (1984). Columnar organization of directionally selective cells in visual area MT of the macaque. *J. Neurophysiol.* *51*, 16.

Andersen, R.A., Asanuma, C., Essick, G., and Siegel, R.M. (1990). Corticocortical connections of anatomically and physiologically defined subdivisions within the inferior parietal lobule. *J. Comp. Neurol.* *296*, 65–113.

Anderson, J.C., Kennedy, H., and Martin, K.A.C. (2011). Pathways of Attention: Synaptic Relationships of Frontal Eye Field to V4, Lateral Intraparietal Cortex, and Area 46 in Macaque Monkey. *J. Neurosci.* *31*, 10872–10881.

Angelucci, A., and Bullier, J. (2003). Reaching beyond the classical receptive field of V1 neurons: horizontal or feedback axons? *J. Physiol. Paris* *97*, 141–154.

Ardid, S., Wang, X.-J., and Compte, A. (2007). An integrated microcircuit model of attentional processing in the neocortex. *J. Neurosci.* *27*, 8486.

Ardid, S., Wang, X.-J., Gomez-Cabrero, D., and Compte, A. (2010). Reconciling coherent oscillation with modulation of irregular spiking activity in selective attention: gamma-range synchronization between sensory and executive cortical areas. *J. Neurosci.* *30*, 2856–2870.

Armstrong, K.M., and Moore, T. (2007). Rapid enhancement of visual cortical response discriminability by microstimulation of the frontal eye field. *Proc. Natl. Acad. Sci.* *104*, 9499–9504.

Armstrong, K.M., Fitzgerald, J.K., and Moore, T. (2006). Changes in visual receptive fields with microstimulation of frontal cortex. *Neuron* *50*, 791–798.

Armstrong, K.M., Chang, M.H., and Moore, T. (2009). Selection and maintenance of spatial information by frontal eye field neurons. *J. Neurosci.* *29*, 15621–15629.

Awh, E., Armstrong, K.M., and Moore, T. (2006). Visual and oculomotor selection: links, causes and implications for spatial attention. *Trends Cogn. Sci.* *10*, 124–130.

Barash, S., Bracewell, R.M., Fogassi, L., Gnadt, J.W., and Andersen, R. a (1991). Saccade-related activity in the lateral intraparietal area. I. Temporal properties; comparison with area 7a. *J. Neurophysiol.* *66*, 1095.

Barone, P., Batardiere, A., Knoblauch, K., and Kennedy, H. (2000). Laminar distribution of neurons in extrastriate areas projecting to visual areas V1 and V4 correlates with the hierarchical rank and indicates the operation of a distance rule. *J. Neurosci.* *20*, 3263.

Bartos, M., Vida, I., Frotscher, M., Meyer, A., Monyer, H., Geiger, J.R.P., and Jonas, P. (2002). Fast synaptic inhibition promotes synchronized gamma oscillations in hippocampal interneuron networks. *Proc. Natl. Acad. Sci.* *99*, 13222–13227.

Beierlein, M., Gibson, J.R., and Connors, B.W. (2003). Two dynamically distinct inhibitory networks in layer 4 of the neocortex. *J. Neurophysiol.* *90*, 2987–3000.

Benevento, L.A., and Rezak, M. (1976). The cortical projections of the inferior pulvinar and adjacent lateral pulvinar in the rhesus monkey (*Macaca mulatta*): an autoradiographic study. *Brain Res.* *108*, 1.

Benjamin, B., Steinmetz, N.A., Gao, P., McQuinn, E., Choudhary, S., Chandrasekaran, A., Bussat, J.-M., Alvarez-Icaza, R., Arthur, J. V., Merolla, P.A., et al. (2014). Energy-efficient simulation of large-scale neural models. *Nature* *submitted*.

Benjamini, Y., and Hochberg, Y. (1995). Controlling the false discovery rate: a practical and powerful approach to multiple testing. *R. Stat. Soc. Ser. B* *57*, 289.

Bichot, N.P., Rossi, A.F., and Desimone, R. (2005). Parallel and Serial Neural Mechanisms for Visual Search in Macaque Area V4. *Science* (80-.). *308*, 529–534.

Binzegger, T., Douglas, R.J., and Martin, K.A.C. (2004). A Quantitative Map of the Circuit of Cat Primary Visual Cortex. *J. Neurosci.* *24*, 8441.

Bisley, J.W., and Goldberg, M.E. (2003). Neuronal activity in the lateral intraparietal area and spatial attention. *Science* (80-.). *299*, 81–86.

Blomquist, P., Devor, A., Indahl, U.G., Ulbert, I., Einevoll, G.T., and Dale, A.M. (2009). Estimation of Thalamocortical and Intracortical Network Models from Joint Thalamic Single-Electrode and Cortical Laminar-Electrode Recordings in the Rat Barrel System. *PLoS Comput. Biol.* *5*, 55–80.

Borg-Graham, L.J., Mohler, C.W., and Frégnac, Y. (1998). Visual input evokes transient and strong shunting inhibition in visual cortical neurons. *Nature* *393*, 369.

Börgers, C., Epstein, S., and Kopell, N.J. (2005). Background gamma rhythmicity and attention in cortical local circuits: a computational study. *Proc. Natl. Acad. Sci.* *102*, 7002–7007.

Bosman, C.A., Womelsdorf, T., Desimone, R., and Fries, P. (2009). A Microsaccadic Rhythm Modulates Gamma-Band Synchronization and Behavior. *J. Neurosci.* *29*, 9471–9480.

Boucsein, C., Nawrot, M.P., Schnepel, P., and Aertsen, A. (2011). Beyond the cortical column: abundance and physiology of horizontal connections imply a strong role for inputs from the surround. *Front. Neurosci.* *5*, 32.

Brodin, L., Traven, H.G., Lansner, A., Wallen, P., Ekeberg, O., Grillner, S., Traven, H.G., Lansner, A., Wallen, P., and Ekeberg, O. (1991). Computer simulations of N-methyl-D-

- aspartate receptor-induced membrane properties in a neuron model. *J. Neurophysiol.* *66*, 473–484.
- Bruce, C.J., and Goldberg, M.E. (1985). Primate frontal eye fields. I. Single neurons discharging before saccades. *J. Neurophysiol.* *53*, 603–635.
- Bruce, C.J., Goldberg, M.E., Bushnell, M.C., and Stanton, G.B. (1985). Primate frontal eye fields. II. Physiological and anatomical correlates of electrically evoked eye movements. *J. Neurophysiol.* *54*, 714–734.
- Buffalo, E. a, Bertini, G., Ungerleider, L.G., and Desimone, R. (2005). Impaired filtering of distracter stimuli by TE neurons following V4 and TEO lesions in macaques. *Cereb. Cortex* *15*, 141–151.
- Buffalo, E. a, Fries, P., Landman, R., Buschman, T.J., and Desimone, R. (2011). Laminar differences in gamma and alpha coherence in the ventral stream. *Proc. Natl. Acad. Sci.* *108*, 1–6.
- Buia, C.I., and Tiesinga, P.H.E. (2008). Role of interneuron diversity in the cortical microcircuit for attention. *J. Neurophysiol.* *99*, 2158.
- Burrows, B.E., and Moore, T. (2009). Influence and limitations of popout in the selection of salient visual stimuli by area V4 neurons. *J. Neurosci.* *29*, 15169–15177.
- Buschman, T.J., and Miller, E.K. (2007). Top-down versus bottom-up control of attention in the prefrontal and posterior parietal cortices. *Science* (80-.). *315*, 1860–1862.
- Buschman, T.J., and Miller, E.K. (2009). Serial, Covert Shifts of Attention during Visual Search Are Reflected by the Frontal Eye Fields and Correlated with Population Oscillations. *Neuron* *63*, 386–396.
- Bushnell, M.C., Goldberg, M.E., and Robinson, D.L. (1981). Behavioral enhancement of visual responses in monkey cerebral cortex. I. Modulation in posterior parietal cortex related to selective visual attention. *J. Neurophysiol.* *46*, 755–772.
- Cadieu, C., Kouh, M., Pasupathy, A., Connor, C.E., Riesenhuber, M., and Poggio, T. (2007). A model of V4 shape selectivity and invariance. *J. Neurophysiol.* *98*, 1733–1750.
- Carandini, M. (2012). From circuits to behavior: a bridge too far? *Nat. Neurosci.* *15*, 507–509.
- Carandini, M., and Heeger, D.J. (2012). Normalization as a canonical neural computation. *Nat. Rev. Neurosci.* *13*, 51–62.
- Carandini, M., Heeger, D.J., and Movshon, J. a (1997). Linearity and normalization in simple cells of the macaque primary visual cortex. *J. Neurosci.* *17*, 8621–8644.
- Carandini, M., Heeger, D.J., and Senn, W. (2002). A synaptic explanation of suppression in visual cortex. *J. Neurosci.* *22*, 10053.

Cardin, J.A., Carlén, M., Meletis, K., Knoblich, U., Zhang, F., Deisseroth, K., Tsai, L.-H.H., and Moore, C.I. (2009). Driving fast-spiking cells induces gamma rhythm and controls sensory responses. *Nature* *459*, 663–667.

Carrasco, M. (2011). Visual attention: the past 25 years. *Vision Res.* *51*, 1484–1525.

Cauli, B., Audinat, E., Lambolez, B., Angulo, M.C.C., Ropert, N., Tsuzuki, K., Hestrin, S., and Rossier, J. (1997). Molecular and physiological diversity of cortical nonpyramidal cells. *J. Neurosci.* *17*, 3894.

Cavanaugh, J.R., and Wurtz, R.H. (2004). Subcortical modulation of attention counters change blindness. *J. Neurosci.* *24*, 11236.

Chance, F.S., Abbott, L.F., and Reyes, A.D. (2002). Gain modulation from background synaptic input. *Neuron* *35*, 773–782.

Chang, M., Xian, S., Rubin, J., and Moore, T. (2013). Latency of chromatic information in area V4. *J. Physiol. Paris*.

Chelazzi, L., Miller, E.K., Duncan, J., and Desimone, R. (1993). A neural basis for visual search in inferior temporal cortex. *Nature* *363*, 345–347.

Chen, C.-M.M., Lakatos, P., Shah, A.S., Mehta, A.D., Givre, S.J., Javitt, D.C., and Schroeder, C.E. (2007). Functional anatomy and interaction of fast and slow visual pathways in macaque monkeys. *Cereb. Cortex* *17*, 1561.

Churchland, M.M., Yu, B.M., Ryu, S.I., Santhanam, G., and Shenoy, K. V (2006). Neural Variability in Premotor Cortex Provides a Signature of Motor Preparation. *J. Neurosci.* *26*, 3697–3712.

Churchland, M.M., Yu, B.M., Sahani, M., and Shenoy, K. V (2007). Techniques for extracting single-trial activity patterns from large-scale neural recordings. *Curr. Opin. Neurobiol.* *17*, 609–618.

Churchland, M.M., Yu, B.M., Cunningham, J.P., Sugrue, L.P., Cohen, M.R., Corrado, G.S., Newsome, W.T., Clark, A.M., Hosseini, P., Scott, B.B., et al. (2010). Stimulus onset quenches neural variability: a widespread cortical phenomenon. *Nat. Neurosci.* *13*, 369–378.

Clark, K.L., Armstrong, K.M., and Moore, T. (2011). Probing neural circuitry and function with electrical microstimulation. *Proc. R. Soc. B Biol. Sci.* *278*, 1121–1130.

Cohen, M.R., and Maunsell, J.H.R. (2009). Attention improves performance primarily by reducing interneuronal correlations. *Nat. Neurosci.* *12*, 1594–1600.

Cohen, M.R., and Maunsell, J.H.R. (2010). A neuronal population measure of attention predicts behavioral performance on individual trials. *J. Neurosci.* *30*, 15241–15253.

- Cohen, J.Y., Pouget, P., Woodman, G.F., Subraveti, C.R., Schall, J.D., and Rossi, A.F. (2007). Difficulty of visual search modulates neuronal interactions and response variability in the frontal eye field. *J. Neurophysiol.* *98*, 2580–2587.
- Cohen, J.Y., Pouget, P., Heitz, R.P., Woodman, G.F., and Schall, J.D. (2008). Biophysical support for functionally distinct cell types in the frontal eye field. *J. Neurophysiol.* *101*, 912–916.
- Compte, A., Brunel, N., Goldman-Rakic, P.S., and Wang, X.-J.J. (2000). Synaptic mechanisms and network dynamics underlying spatial working memory in a cortical network model. *Cereb. Cortex* *10*, 910–923.
- Connor, C.E., Preddie, D.C., Gallant, J.L., and Van Essen, D.C. (1997). Spatial attention effects in macaque area V4. *J. Neurosci.* *17*, 3201–3214.
- Cook, E.P., and Maunsell, J.H.R. (2002). Dynamics of neuronal responses in macaque MT and VIP during motion detection. *Nat. Neurosci.* *5*, 985–994.
- Corbetta, M., Akbudak, E., Conturo, T.E., Snyder, A.Z., Ollinger, J.M., Drury, H.A., Linenweber, M.R., Petersen, S.E., Raichle, M.E., and Van Essen, D.C. (1998). A common network of functional areas for attention and eye movements. *Neuron* *21*, 761–773.
- Craighero, L., and Rizzolatti, G. (2005). The premotor theory of attention. *Neurobiol. Atten.* 181–186.
- Crick, F.C. (1984). Function of the thalamic reticular complex: The searchlight hypothesis. *Proc. Natl. Acad. Sci.* *81*, 4586–4590.
- Cutrell, E.B., and Marrocco, R.T. (2002). Electrical microstimulation of primate posterior parietal cortex initiates orienting and alerting components of covert attention. *Exp. Brain Res.* *144*, 103–113.
- David, S. V, Hayden, B.Y., and Gallant, J.L. (2006). Spectral receptive field properties explain shape selectivity in area V4. *J. Neurophysiol.* *96*, 3492–3505.
- David, S. V, Hayden, B.Y., Mazer, J. a, and Gallant, J.L. (2008). Attention to Stimulus Features Shifts Spectral Tuning of V4 Neurons during Natural Vision. *Neuron* *59*, 509–521.
- Desimone, R., and Duncan, J. (1995). Neural mechanisms of selective visual attention. *Annu. Rev. Neurosci.* *18*, 193–222.
- Desimone, R., and Schein, S.J. (1987). Visual properties of neurons in area V4 of the macaque: sensitivity to stimulus form. *J. Neurophysiol.* *57*, 835–868.
- Deubel, H., and Schneider, W.X.X. (1996). Saccade target selection and object recognition: Evidence for a common attentional mechanism. *Vision Res.* *36*, 1827–1837.

Dias, E.C., and Segraves, M.A. (1999). Muscimol-Induced Inactivation of Monkey Frontal Eye Field: Effects on Visually and Memory-Guided Saccades. *J. Neurophysiol.* *81*, 2191–2214.

Disney, A.A., Domakonda, K.V. V, and Aoki, C. (2006). Differential expression of muscarinic acetylcholine receptors across excitatory and inhibitory cells in visual cortical areas V1 and V2 of the macaque monkey. *J. Comp. Neurol.* *499*, 49–63.

Disney, A.A., Aoki, C., and Hawken, M.J. (2007). Gain modulation by nicotine in macaque v1. *Neuron* *56*, 701–713.

Domenici, L., Harding, G.W.W., and Burkhalter, A. (1995). Patterns of synaptic activity in forward and feedback pathways within rat visual cortex. *J. Neurophysiol.* *74*, 2649.

Douglas, R.J., and Martin, K.A.C. (2004). Neuronal Circuits of the Neocortex. *Annu. Rev. Neurosci.* *27*, 419–451.

Douglas, R.J., and Martin, K.A.C. (2007). Recurrent neuronal circuits in the neocortex. *Curr. Biol.* *17*, 496–500.

Douglas, R.J., Martin, K.A.C., and Whitteridge, D. (1989). A canonical microcircuit for neocortex. *Neural Comput.* *1*, 480–488.

Edelman, J. a, and Keller, E.L. (1996). Activity of visuomotor burst neurons in the superior colliculus accompanying express saccades. *J. Neurophysiol.* *76*, 908–926.

Ekstrom, L.B., Roelfsema, P.R., Arsenault, J.T., Kolster, H., and Vanduffel, W. (2009). Modulation of the Contrast Response Function by Electrical Microstimulation of the Macaque Frontal Eye Field. *J. Neurosci.* *29*, 10683.

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

Felleman, D.J., and Van Essen, D.C. (1991). Distributed hierarchical processing in the primate cerebral cortex. *Cereb. Cortex* *1*, 1–47.

Fischer, B., and Boch, R. (1981). Enhanced activation of neurons in prelunate cortex before visually guided saccades of trained rhesus monkeys. *Exp. Brain Res.* *44*, 129–137.

Fries, P., Reynolds, J.H., Rorie, a E., and Desimone, R. (2001). Modulation of Oscillatory Neuronal Synchronization by Selective Visual Attention. *Science* (80-.). *291(5508)*, 1560–1563.

Gallant, J.L., Connor, C.E., Rakshit, S., Lewis, J.W., and Van Essen, D.C. (1996). Neural responses to polar, hyperbolic, and Cartesian gratings in area V4 of the macaque monkey. *J. Neurophysiol.* *76*, 2718–2739.

- Gattass, R., Sousa, A.P., and Gross, C.G. (1988). Visuotopic organization and extent of V3 and V4 of the macaque. *J. Neurosci.* *8*, 1831–1845.
- Ghose, G.M., and Maunsell, J.H.R. (2002). Attentional modulation in visual cortex depends on task timing. *Nature* *419*, 616–620.
- Girard, P., Lomber, S.G., and Bullier, J. (2002). Shape discrimination deficits during reversible deactivation of area V4 in the macaque monkey. *Cereb. Cortex* *12*, 1146–1156.
- Givre, S.J., Schroeder, C.E., and Arezzo, J.C. (1994). Contribution of extrastriate area V 4 to the surface-recorded flash VEP in the awake macaque. *Vision Res.* *34*, 415–428.
- Goard, M., and Dan, Y. (2009). Basal forebrain activation enhances cortical coding of natural scenes. *Nat. Neurosci.* *12*, 1444–1449.
- Goldberg, M.E., and Wurtz, R. (1972). Activity of superior colliculus in behaving monkey. II. Effect of attention on neuronal responses. *J Neurophysiol* *35*, 560–574.
- Gottlieb, J.P., Kusunoki, M., and Goldberg, M.E. (1998). The representation of visual salience in monkey parietal cortex. *Nature* *391*, 481–484.
- Graziano, M.S., Hu, X.T., and Gross, C.G. (1997). Visuospatial properties of ventral premotor cortex. *J. Neurophysiol.* *77*, 2268–2292.
- Gregoriou, G.G., Gotts, S.J., Zhou, H., and Desimone, R. (2009). High-Frequency, Long-Range Coupling Between Prefrontal and Visual Cortex During Attention. *Science* (80-.). *324*, 1207–1210.
- Gregoriou, G.G., Gotts, S.J., and Desimone, R. (2012). Cell-type-specific synchronization of neural activity in FEF with V4 during attention. *Neuron* *73*, 581–594.
- Guillery, R.W., and Sherman, S.M. (2002). Thalamic Relay Functions and Their Role in Corticocortical Communication: Generalizations from the Visual System. *Neuron* *33*, 163–175.
- Hamker, F.H., and Zirnsak, M. (2006). V4 receptive field dynamics as predicted by a systems-level model of visual attention using feedback from the frontal eye field. *Neural Networks* *19*, 1371–1382.
- Han, X., Xian, S.X., and Moore, T. (2009). Dynamic sensitivity of area V4 neurons during saccade preparation. *Proc. Natl. Acad. Sci.* *106*, 13046–13051.
- Hanes, D.P., and Schall, J.D. (1996). Neural Control of Voluntary Movement Initiation. *Science* (80-.). *274*, 427–430.
- Hansen, B.J., and Dragoi, V. (2011). Adaptation-induced synchronization in laminar cortical circuits. *Proc. Natl. Acad. Sci.* *108*, 10720.

- Harris, K.D., and Thiele, A. (2011). Cortical state and attention. *Nat. Rev. Neurosci.* *12*.
- Hegd , J., and Van Essen, D.C. (2007). A comparative study of shape representation in macaque visual areas V2 and V4. *Cereb. Cortex* *17*, 1100.
- Heinzle, J., Hepp, K., and Martin, K.A.C. (2007). A Microcircuit Model of the Frontal Eye Fields. *J. Neurosci.* *27*, 9341.
- Helminski, J.O., and Segraves, M.A. (2003). Macaque frontal eye field input to saccade-related neurons in the superior colliculus. *J. Neurophysiol.* *90*, 1046.
- Herrero, J.L., Roberts, M.J., Delicato, L.S., Gieselmann, M. a, Dayan, P., and Thiele, A. (2008). Acetylcholine contributes through muscarinic receptors to attentional modulation in V1. *Nature* *454*, 1110–1114.
- Heywood, C.A., and Cowey, A. (1987). On the role of cortical area V4 in the discrimination of hue and pattern in macaque monkeys. *J. Neurosci.* *7*, 2601.
- Hikosaka, O., and Wurtz, R.H. (1986). Saccadic eye movements following injection of lidocaine into the superior colliculus. *Exp. Brain Res.* *61*, 531–539.
- Hilgetag, C., O’Neill, M., and Young, M. (1996). Indeterminate organization of the visual system. *Science* (80-.). *271*, 776–777.
- Hill, D.N., Mehta, S.B., and Kleinfeld, D. (2011). Quality metrics to accompany spike sorting of extracellular signals. *J. Neurosci.* *31*, 8699–8705.
- Hinkle, D. a, and Connor, C.E. (2002). Three-dimensional orientation tuning in macaque area V4. *Nat. Neurosci.* *5*, 665–670.
- Hinkle, D. a, and Connor, C.E. (2005). Quantitative characterization of disparity tuning in ventral pathway area V4. *J. Neurophysiol.* *94*, 2726–2737.
- Hirsch, J. a, and Martinez, L.M. (2006). Laminar processing in the visual cortical column. *Curr. Opin. Neurobiol.* *16*, 377–384.
- Hoffman, J.E., and Subramaniam, B. (1995). The Role of Visual Attention in Saccadic Eye Movements. *Percept. Psychophys.* *57*, 787.
- Holt, G.R., and Koch, C. (1997). Shunting inhibition does not have a divisive effect on firing rates. *Neural Comput.* *9*, 1001–1013.
- Hubel, D.H., and Wiesel, T. (1962). Receptive fields, binocular interaction and functional architecture in the cat’s visual cortex. *J. Physiol.* *160*, 106–154.
- Hubel, D.H., and Wiesel, T.N. (1968). Receptive fields and functional architecture of monkey striate cortex. *J. Physiol.* *195*, 215–243.

- Huerta, M.F., Krubitzer, L. a, and Kaas, J.H. (1986). Frontal eye field as defined by intracortical microstimulation in squirrel monkeys, owl monkeys, and macaque monkeys: I. Subcortical connections. *J. Comp. Neurol.* *253*, 415–439.
- Hunt, A.R., and Kingstone, A. (2003). Covert and overt voluntary attention: linked or independent? *Cogn. Brain Res.* *18*, 102–105.
- Hupé, J.M., James, A.C., Payne, B.R., Lomber, S.G., Girard, P., and Bullier, J. (1998). Cortical feedback improves discrimination between figure and background by V1, V2 and V3 neurons. *Nature* *394*, 784–787.
- Ignashchenkova, A., Dicke, P.W., Haarmeier, T., and Thier, P. (2004). Neuron-specific contribution of the superior colliculus to overt and covert shifts of attention. *Nat. Neurosci.* *7*, 56–64.
- Ipata, A.E., Gee, A.L., Goldberg, M.E., and Bisley, J.W. (2006). Activity in the lateral intraparietal area predicts the goal and latency of saccades in a free-viewing visual search task. *J. Neurosci.* *26*, 3656–3661.
- Jahr, C., and Stevens, C. (1990). Voltage Dependence of NMDA-Activated Macroscopic Conductances Predicted by Single-Channel Kinetics. *J. Neurosci.* *10*, 3178–3182.
- Janssen, P., and Shadlen, M.N. (2005). A representation of the hazard rate of elapsed time in macaque area LIP. *Nat. Neurosci.* *8*, 234–241.
- Jones, E.G. (2007). *The Thalamus* (Cambridge University Press).
- Kalwani, R.M., Bloy, L., Elliott, M. a, and Gold, J.I. (2009). A method for localizing microelectrode trajectories in the macaque brain using MRI. *J. Neurosci. Methods* *176*, 104–111.
- Keller, E.L., and Edelman, J. a (1994). Use of interrupted saccade paradigm to study spatial and temporal dynamics of saccadic burst cells in superior colliculus in monkey. *J. Neurophysiol.* *72*, 2754–2770.
- Kepecs, A., and Raghavachari, S. (2007). Gating information by two-state membrane potential fluctuations. *J. Neurophysiol.* *97*, 3015.
- Koch, C., and Ullman, S. (1985). Shifts in selective visual attention: towards the underlying neural circuitry. *Hum Neurobiol* *4*, 219–227.
- Kodaka, Y., Mikami, A., and Kubota, K. (1997). Neuronal activity in the frontal eye field of the monkey is modulated while attention is focused on to a stimulus in the peripheral visual field, irrespective of eye movement. *Neurosci. Res.* *28*, 291–298.
- Kotake, Y., Morimoto, H., Okazaki, Y., Fujita, I., and Tamura, H. (2009). Organization of color-selective neurons in macaque visual area V4. *J. Neurophysiol.* *102*, 15–27.

- Kouh, M., and Poggio, T. (2008). A canonical neural circuit for cortical nonlinear operations. *Neural Comput.* *20*, 1427–1451.
- Kowler, E. (2011). Eye movements: the past 25 years. *Vision Res.* *51*, 1457–1483.
- Kunzle, H., and Akert, K. (1977). Efferent connections of cortical area 8 (frontal eye field) in *Macaca fascicularis*. A reinvestigation using the autoradiographic technique. *J. Comp. Neurol.* *173*, 164.
- Kustov, A.A., and Robinson, D.L. (1996). Shared neural control of attentional shifts and eye movements. *Nature* *384*, 74–77.
- Larkum, M.E., Senn, W., and Lüscher, H.-R. (2004). Top-down dendritic input increases the gain of layer 5 pyramidal neurons. *Cereb. Cortex* *14*, 1059–1070.
- Leichnetz, G.R., Spencer, R.F., Hardy, S.G., and Astruc, J. (1981). The prefrontal corticotectal projection in the monkey; an anterograde and retrograde horseradish peroxidase study. *Neuroscience* *6*, 1023–1041.
- Leichnetz, G.R., Smith, D.J., and Spencer, R.F. (1984). Cortical projections to the paramedian tegmental and basilar pons in the monkey. *J. Comp. Neurol.* *228*, 388–408.
- Leopold, D.A., and Logothetis, N.K. (1998). Microsaccades differentially modulate neural activity in the striate and extrastriate visual cortex. *Exp. Brain Res.* *123*, 341–345.
- Li, C.S., Mazzone, P., and Andersen, R.A. (1999). Effect of reversible inactivation of macaque lateral intraparietal area on visual and memory saccades. *J. Neurophysiol.* *81*, 1827–1838.
- Liu, Y., Yttri, E.A., and Snyder, L.H. (2010). Intention and attention: different functional roles for LIPd and LIPv. *Nat. Neurosci.*
- Lovejoy, L.P., and Krauzlis, R.J. (2010). Inactivation of primate superior colliculus impairs covert selection of signals for perceptual judgments. *Nat. Neurosci.* *13*, 261–266.
- Ludwig, K. a, Miriani, R.M., Langhals, N.B., Joseph, M.D., Anderson, D.J., and Kipke, D.R. (2009). Using a Common Average Reference to Improve Cortical Neuron Recordings From Microelectrode Arrays. *J. Neurophysiol.* *101*, 1679–1689.
- Lynch, J.C., and McLaren, J.W. (1989). Deficits of visual attention and saccadic eye movements after lesions of parietooccipital cortex in monkeys. *J. Neurophysiol.* *61*, 74–90.
- Lynch, J.C., and Tian, J.-R.R. (2006). Cortico-cortical networks and cortico-subcortical loops for the higher control of eye movements. *Prog. Brain Res.* *151*, 461–501.
- Lynch, J.C., Hoover, J.E., and Strick, P.L. (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.

- Mante, V., Sussillo, D., Shenoy, K. V., and Newsome, W.T. (2013). Context-dependent computation by recurrent dynamics in prefrontal cortex. *Nature* *503*, 78–84.
- Marr, D. (1982). *Vision: A computational investigation into the human representation and processing of visual information* (New York, NY: Henry Holt & Co).
- Maunsell, J.H.R. (1987). Physiological evidence for two visual subsystems. In *Matters of Intelligence*, pp. 59–87.
- Mazer, J. a, and Gallant, J.L. (2003). Goal-Related Activity in V4 during Free Viewing Visual Search Evidence for a Ventral Stream Visual Saliency Map. *Neuron* *40*, 1241–1250.
- McAdams, C.J., and Maunsell, J.H.R. (1999). Effects of Attention on Orientation-Tuning Functions of Single Neurons in Macaque Cortical Area V4. *J. Neurosci.* *19*, 431–441.
- McAlonan, K., Cavanaugh, J.R., and Wurtz, R.H. (2008). Guarding the gateway to cortex: attention in visual thalamus. *Nature* *456*, 391.
- Mehta, A.D., Ulbert, I., and Schroeder, C.E. (2000). Intermodal selective attention in monkeys. II: physiological mechanisms of modulation. *Cereb. Cortex* *10*, 359–370.
- Mendelson, M.J., Haith, M.M., and Goldman-Rakic, P.S. (1982). Face scanning and responsiveness to social cues in infant rhesus monkeys. *Dev. Psychol.* *18*, 222–228.
- Merigan, W.H., and Pham, H. (1998). V4 lesions in macaques affect both single-and multiple-viewpoint shape discriminations. *Vis. Neurosci.* *15*, 359–367.
- Messinger, A., Lebedev, M. a, Kralik, J.D., and Wise, S.P. (2009). Multitasking of attention and memory functions in the primate prefrontal cortex. *J. Neurosci.* *29*, 5640–5653.
- Mitchell, S.J., and Silver, R.A. (2003). Shunting inhibition modulates neuronal gain during synaptic excitation. *Neuron* *38*, 433–445.
- Mitchell, J.F., Sundberg, K.A.A., and Reynolds, J.H. (2007). Differential attention-dependent response modulation across cell classes in macaque visual area V4. *Neuron* *55*, 131–141.
- Mitchell, J.F., Sundberg, K. a, and Reynolds, J.H. (2009). Spatial Attention Decorrelates Intrinsic Activity Fluctuations in Macaque Area V4. *Neuron* *61*, 879–888.
- Mitzdorf, U. (1985). Current source-density method and application in cat cerebral cortex: investigation of evoked potentials and EEG phenomena. *Physiol. Rev.* *65*, 37–100.
- Monosov, I.E., and Thompson, K.G. (2009). Frontal eye field activity enhances object identification during covert visual search. *J. Neurophysiol.* *102*, 3656–3672.

Monosov, I.E., Trageser, J.C., and Thompson, K.G. (2008). Measurements of simultaneously recorded spiking activity and local field potentials suggest that spatial selection emerges in the frontal eye field. *Neuron* *57*, 614–625.

Monteon, J.A., Constantin, A.G., Wang, H., Martinez-Trujillo, J., and Crawford, J.D. (2010). Electrical stimulation of the frontal eye fields in the head-free macaque evokes kinematically normal 3D gaze shifts. *J. Neurophysiol.* *104*, 3462–3475.

Moore, T. (1999). Shape Representations and Visual Guidance of Saccadic Eye Movements. *Science* (80-.). *285*, 1914–1917.

Moore, T. (2006). The neurobiology of visual attention: finding sources. *Curr. Opin. Neurobiol.* *16*, 159–165.

Moore, T., and Armstrong, K.M. (2003). Selective gating of visual signals by microstimulation of frontal cortex. *Nature* *421*, 370–373.

Moore, T., and Chang, M.H. (2009). Presaccadic discrimination of receptive field stimuli by area V4 neurons. *Vision Res.* *49*, 1227–1232.

Moore, T., and Fallah, M. (2001). Control of eye movements and spatial attention. *Proc. Natl. Acad. Sci.* *98*, 21549498.

Moore, T., and Fallah, M. (2004). Microstimulation of the Frontal Eye Field and Its Effects on Covert Spatial Attention. *J. Neurophysiol.* *91*, 152–162.

Moore, T., Tolias, A.S., and Schiller, P.H. (1998). Visual representations during saccadic eye movements. *Proc. Natl. Acad. Sci.* *95*, 8981–8984.

Moore, T., Armstrong, K.M., and Fallah, M. (2003). Visuomotor Origins of Covert Spatial Attention. *Neuron* *40*, 671–683.

Moran, J., and Desimone, R. (1985). Selective attention gates visual processing in the extrastriate cortex. *Science* (80-.). *229*, 782.

Motter, B.C. (1993). Focal attention produces spatially selective processing in visual cortical areas V1, V2, and V4 in the presence of competing stimuli. *J. Neurophysiol.* *70*, 909.

Mountcastle, V.B. (1997). The columnar organization of the neocortex. *Brain* *120*, 701.

Muller, J.R., Philiastides, M.G., and Newsome, W.T. (2005). Microstimulation of the superior colliculus focuses attention without moving the eyes. *Proc. Natl. Acad. Sci.* *102*, 524.

Murphy, B.K., and Miller, K.D. (2003). Multiplicative gain changes are induced by excitation or inhibition alone. *J. Neurosci.* *23*, 10040–10051.

- Nakamura, K., and Colby, C.L. (2002). Updating of the visual representation in monkey striate and extrastriate cortex during saccades. *Proc. Natl. Acad. Sci.* *99*, 4026–4031.
- Nandy, A.S., Sharpee, T.O., Reynolds, J.H., and Mitchell, J.F. (2013). The fine structure of shape tuning in area V4. *Neuron* *78*, 1102–1115.
- Noudoost, B., and Moore, T. (2011a). Control of visual cortical signals by prefrontal dopamine. *Nature* *474*, 372–375.
- Noudoost, B., and Moore, T. (2011b). The role of neuromodulators in selective attention. *Trends Cogn. Sci.* *15*, 585–591.
- Noudoost, B., Chang, M.H., Steinmetz, N.A., and Moore, T. (2010). Top-down control of visual attention. *Curr. Opin. Neurobiol.* *20*, 183–190.
- O'Regan, J.K., Rensink, R. a, and Clark, J.J. (1999). Change-blindness as a result of “mudsplashes”. *Nature* *398*, 34.
- Okun, M., Naim, A., and Lampl, I. (2010). The subthreshold relation between cortical local field potential and neuronal firing unveiled by intracellular recordings in awake rats. *J. Neurosci.* *30*, 4440–4448.
- Olsen, S.R., Bortone, D.S., Adesnik, H., and Scanziani, M. (2012). Gain control by layer six in cortical circuits of vision. *Nature* *483*, 47–52.
- De Pasquale, R., and Sherman, S.M. (2011). Synaptic properties of corticocortical connections between the primary and secondary visual cortical areas in the mouse. *J. Neurosci.* *31*, 16494–16506.
- Pasupathy, A., and Connor, C.E. (2001). Shape representation in area V4: position-specific tuning for boundary conformation. *J. Neurophysiol.* *86*, 2505.
- Petersen, S.E., Robinson, D.L., and Morris, J.D. (1987). Contributions of the pulvinar to visual spatial attention. *Neuropsychologia* *25*, 97–105.
- Pettersen, K.H., Devor, A., Ulbert, I., Dale, A.M., and Einevoll, G.T. (2006). Current-source density estimation based on inversion of electrostatic forward solution: effects of finite extent of neuronal activity and conductivity discontinuities. *J. Neurosci. Methods* *154*, 116–133.
- Posner, M.I. (1980). Orienting of attention. *Q. J. Exp. Psychol.* *32*, 3–25.
- Posner, M.I., Snyder, C.R., and Davidson, B.J. (1980). Attention and the detection of signals. *J. Exp. Psychol.* *109*, 160–174.
- Pouget, P., Stepniewska, I., Crowder, E. a, Leslie, M.W., Emeric, E.E., Nelson, M.J., and Schall, J.D. (2009). Visual and motor connectivity and the distribution of calcium-binding proteins in macaque frontal eye field: implications for saccade target selection. *Front. Neuroanat.* *3*, 2.

Poulet, J.F.A., and Petersen, C.C.H. (2008). Internal brain state regulates membrane potential synchrony in barrel cortex of behaving mice. *Nature* 454, 881–885.

Raizada, R.D.S., and Grossberg, S. (2003). Towards a theory of the laminar architecture of cerebral cortex: Computational clues from the visual system. *Cereb. Cortex* 13, 100–113.

Reynolds, J.H., and Chelazzi, L. (2004). Attentional modulation of visual processing. *Annu. Rev. Neurosci.* 27, 611–647.

Reynolds, J.H., and Heeger, D.J. (2009). The Normalization Model of Attention. *Neuron* 61, 168–185.

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

Riehle, a, and Requin, J. (1993). The predictive value for performance speed of preparatory changes in neuronal activity of the monkey motor and premotor cortex. *Behav Brain Res* 53, 35–49.

Riesenhuber, M., and Poggio, T. (1999). Hierarchical models of object recognition in cortex. *Nat. Neurosci.* 2, 1019–1025.

Robinson, D. (1972). Eye movements evoked by collicular stimulation in the alert monkey. *Vision Res.* 12, 1795–1808.

Robinson, D.A., and Fuchs, A.F. (1969). Eye movements evoked by stimulation of frontal eye fields. *J. Neurophysiol.* 32, 637–648.

De Ruyter van Steveninck, R.R., Lewen, G.D., Strong, S.P., Koberle, R., and Bialek, W. (1997). Reproducibility and Variability in Neural Spike Trains. *Science* (80-.). 275, 1805–1808.

Sakata, S., and Harris, K.D. (2009). Laminar structure of spontaneous and sensory-evoked population activity in auditory cortex. *Neuron* 64, 404–418.

Sato, T.R., and Schall, J.D. (2003). Effects of stimulus-response compatibility on neural selection in frontal eye field. *Neuron* 38, 637–648.

Schafer, R.J., and Moore, T. (2007). Attention Governs Action in the Primate Frontal Eye Field. *Neuron* 56, 541–551.

Schafer, R.J., and Moore, T. (2011). Selective attention from voluntary control of neurons in prefrontal cortex. *Science* (80-.). 1568, 1568–1571.

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

- Schein, S.J., and Desimone, R. (1990). Spectral properties of V4 neurons in the macaque. *J. Neurosci.* *10*, 3369–3389.
- Schiller, P.H., and Lee, K. (1991). The role of the primate extrastriate area V4 in vision. *Science* (80-.). *251*, 1251.
- Schiller, P.H., and Stryker, M. (1972). Single-unit recording and stimulation in superior colliculus of the alert rhesus monkey. *J. Neurophysiol.* 915–924.
- Schiller, J., Major, G., Koester, H.J., and Schiller, Y. (2000). NMDA spikes in basal dendrites of cortical pyramidal neurons. *Nature* *404*, 285–289.
- Schiller, P.H., True, S.D., and Conway, J.L. (1980). Deficits in eye movements following frontal eye-field and superior colliculus ablations. *J. Neurophysiol.* *44*, 1175–1189.
- Schiller, P.H., Sandell, J.H., and Maunsell, J.H.R. (1987). The effect of frontal eye field and superior colliculus lesions on saccadic latencies in the rhesus monkey. *J. Neurophysiol.* *57*, 1033–1049.
- Schroeder, C.E., Mehta, A.D., and Givre, S.J. (1998). A spatiotemporal profile of visual system activation revealed by current source density analysis in the awake macaque. *Cereb. Cortex* *8*, 575–592.
- Segraves, M. (1992). Activity of monkey frontal eye field neurons projecting to oculomotor regions of the pons. *J. Neurophysiol.* *68*, 1967–1985.
- Self, M.W., Kooijmans, R.N., Supèr, H., Lamme, V. a, and Roelfsema, P.R. (2012). Different glutamate receptors convey feedforward and recurrent processing in macaque V1. *Proc. Natl. Acad. Sci.* *109*, 11031–11036.
- Sheinberg, D.L., and Logothetis, N.K. (2001). Noticing Familiar Objects in Real World Scenes: The Role of Temporal Cortical Neurons in Natural Vision. *J. Neurosci.* *21*, 1340–1350.
- Shibutani, H., Sakata, H., and Hyvärinen, J. (1984). Saccade and blinking evoked by microstimulation of the posterior parietal association cortex of the monkey. *Exp. Brain Res.* *55*, 1–8.
- Shipp, S. (2003). The functional logic of cortico-pulvinar connections. *Philos. Trans. R. Soc. B Biol. Sci.* *358*, 1605–1624.
- Silver, R.A. (2010). Neuronal arithmetic. *Nat. Rev. Neurosci.* *11*, 474–489.
- Simons, D.J., and Rensink, R. a (2005). Change blindness: past, present, and future. *Trends Cogn. Sci.* *9*, 16–20.
- Smith, S.L., Smith, I.T., Branco, T., and Häusser, M. (2013). Dendritic spikes enhance stimulus selectivity in cortical neurons in vivo. *Nature*.

Snyder, L.H., Dickinson, A.R., and Calton, J.L. (2006). Preparatory delay activity in the monkey parietal reach region predicts reach reaction times. *J. Neurosci.* 26, 10091–10099.

Sommer, M. a, and Tehovnik, E.J. (1997). Reversible inactivation of macaque frontal eye field. *Exp. Brain Res.* 116, 229–249.

Sommer, M.A., and Wurtz, R.H. (2001). Frontal eye field sends delay activity related to movement, memory, and vision to the superior colliculus. *J. Neurophysiol.* 85, 1673–1685.

Sporns, O., and Zwi, J. (2004). The small world of the cerebral cortex. *Neuroinformatics* 2, 145–162.

Spratling, M.W., and Johnson, M.H. (2004). A feedback model of visual attention. *J. Cogn. Neurosci.* 16, 219–237.

Squire, R.F., Steinmetz, N.A., and Moore, T. (2012). Frontal Eye Field. *Scholarpedia* 7, 5341.

Squire, R.F., Noudoost, B., Schafer, R.J., and Moore, T. (2013). Prefrontal contributions to visual selective attention. *Annu. Rev. Neurosci.* 36, 451–466.

Sridharan, D., Steinmetz, N.A., Moore, T., and Knudsen, E.I. (2013). A mathematical model for distinguishing bias from sensitivity effects in multialternative detection tasks. *Arxiv* 1310.4219.

Stanton, G., Deng, S.Y., Goldberg, M.E., and McMullen, N. (1989). Cytoarchitectural characteristic of the frontal eye fields in macaque monkeys. *J. Comp. Neurol.* 282, 415–427.

Stanton, G.B., Bruce, C.J., and Goldberg, M.E. (1995). Topography of projections to posterior cortical areas from the macaque frontal eye fields. *J. Comp. Neurol.* 353, 291–305.

Steinmetz, N.A., and Moore, T. (2010). Changes in the response rate and response variability of area V4 neurons during the preparation of saccadic eye movements. *J. Neurophysiol.* 103, 1171–1178.

Steinmetz, N.A., and Moore, T. (2012). Lumping and splitting the neural circuitry of visual attention. *Neuron* 73, 410–412.

Sundberg, K. a, Mitchell, J.F., and Reynolds, J.H. (2009). Spatial attention modulates center-surround interactions in macaque visual area v4. *Neuron* 61, 952–963.

Supèr, H., and Lamme, V. a F. (2007). Strength of figure-ground activity in monkey primary visual cortex predicts saccadic reaction time in a delayed detection task. *Cereb. Cortex* 17, 1468–1475.

- Swadlow, H.A. (2003). Fast-spike interneurons and feedforward inhibition in awake sensory neocortex. *Cereb. Cortex* *13*, 25–32.
- Swadlow, H.A., Gusev, A.G., and Bezdudnaya, T. (2002). Activation of a cortical column by a thalamocortical impulse. *J. Neurosci.* *22*, 7766.
- Tanaka, M., Lindsley, E., Lausmann, S., and Creutzfeldt, O.D. (1990). Afferent connections of the prelunate visual association cortex (areas V4 and DP). *Anat. Embryol. (Berl)*. *181*, 19–30.
- Tanigawa, H., Lu, H.D., and Roe, A.W. (2010). Functional organization for color and orientation in macaque V4. *Nat. Neurosci.* *13*, 1542–1548.
- Tehovnik, E.J., Sommer, M. a, Chou, I.H., Slocum, W.M., and Schiller, P.H. (2000). Eye fields in the frontal lobes of primates. *Brain Res. Rev.* *32*, 413–448.
- Thompson, K.G., Biscoe, K.L., and Sato, T.R. (2005). Neuronal basis of covert spatial attention in the frontal eye field. *J. Neurosci.* *25*, 9479–9487.
- Thomson, A.M., and Bannister, A.P. (2003). Interlaminar connections in the neocortex. *Cereb. Cortex* *13*, 5–14.
- Thomson, A.M., and Lamy, C. (2007). Functional maps of neocortical local circuitry. *Front. Neurosci.* *1*, 19.
- Tiesinga, P.H.E., and Buia, C.I. (2009). Spatial attention in area V4 is mediated by circuits in primary visual cortex. *Neural Networks* *22*, 1039–1054.
- Tiesinga, P.H.E., Fellous, J.M.M., Salinas, E., José, J.V. V, and Sejnowski, T.J.J. (2004). Synchronization as a mechanism for attentional gain modulation. *Neurocomputing* *58*, 641–646.
- Tolias, A.S., Moore, T., Smirnakis, S.M., Tehovnik, E.J., Siapas, A.G., and Schiller, P.H. (2001). Eye Movements Modulate Visual Receptive Fields of V4 Neurons. *Neuron* *29*, 757–767.
- Trojanowski, J.Q., and Jacobson, S. (1976). Areal and laminar distribution of some pulvinar cortical efferents in rhesus monkey. *J. Comp. Neurol.* *169*.
- Tu, T., and Keating, E. (2000). Electrical stimulation of the frontal eye field in a monkey produces combined eye and head movements. *J. Neurophysiol.* *84*, 1103–1106.
- Ungerleider, L.G., Galkin, T.W., Desimone, R., and Gattass, R. (2008). Cortical connections of area V4 in the macaque. *Cereb. Cortex* *18*, 477–499.
- Vishwanath, D., and Kowler, E. (2003). Localization of shapes: eye movements and perception compared. *Vision Res.* *43*, 1637–1653.

- Wagatsuma, N., Potjans, T.C., Diesmann, M., and Fukai, T. (2011). Layer-dependent attentional processing by top-down signals in a visual cortical microcircuit model. *Front. Comput. Neurosci.* 5, 1–15.
- Walsh, V., Carden, D., Butler, S.R., and Kulikowski, J.J. (1993). The effects of V4 lesions on the visual abilities of macaques: hue discrimination and colour constancy. *Behav. Brain Res.* 53, 51–62.
- Wardak, C., Olivier, E., and Duhamel, J.-R. (2004). A deficit in covert attention after parietal cortex inactivation in the monkey. *Neuron* 42, 501–508.
- Wardak, C., Ibos, G., Duhamel, J.R., and Olivier, E. (2006). Contribution of the monkey frontal eye field to covert visual attention. *J. Neurosci.* 26, 4228–4235.
- De Weerd, P., Peralta, M.R., Desimone, R., and Ungerleider, L.G. (1999). Loss of attentional stimulus selection after extrastriate cortical lesions in macaques. *Nat. Neurosci.* 2, 753–758.
- De Weerd, P., Iii, M.R.P., Desimone, R., and Ungerleider, L.G. (2000). Loss of attentional stimulus selection after extrastriate cortical lesions in macaques. *Nat. Neurosci.* 3, 409.
- De Weerd, P., Desimone, R., and Ungerleider, L.G. (2003). Generalized deficits in visual selective attention after V4 and TEO lesions in macaques. *Eur. J. Neurosci.* 18, 1671–1691.
- Welch, K., and Stuteville, P. (1958). Experimental production of unilateral neglect in monkeys. *Brain* 81, 341–347.
- Williford, T., and Maunsell, J.H.R. (2006). Effects of spatial attention on contrast response functions in macaque area V4. *J. Neurophysiol.* 96, 40.
- Womelsdorf, T., Fries, P., Mitra, P.P., and Desimone, R. (2006). Gamma-band synchronization in visual cortex predicts speed of change detection. *Nature* 439, 733–736.
- Yantis, S., and Jonides, J. (1984). Abrupt visual onsets and selective attention: Evidence from visual search. *J. Exp. Psychol. Hum. Percept. Perform.* 10, 601–621.
- Yarbus, A.L. (1967). *Eye Movements and Vision* (New York: Plenum Press).
- Yoshioka, T., Levitt, J.B., and Lund, J.S. (1992). Intrinsic lattice connections of macaque monkey visual cortical area V4. *J. Neurosci.* 12, 2785.
- Zagha, E., Casale, A.E., Sachdev, R.N.S., McGinley, M.J., and McCormick, D.A. (2013). Motor Cortex Feedback Influences Sensory Processing by Modulating Network State. *Neuron* 1–12.
- Zeki, S. (1983). The distribution of wavelength and orientation selective cells in different areas of monkey visual cortex. *Neuroscience* 9, 741.

Zénon, A., and Krauzlis, R.J. (2012). Attention deficits without cortical neuronal deficits. *Nature*.

Zhang, Y., Wang, X., Bressler, S.L., Chen, Y., and Ding, M. (2008). Prestimulus cortical activity is correlated with speed of visuomotor processing. *J. Cogn. Neurosci.* *20*, 1915–1925.

Zhou, H., Friedman, H.S., and von der Heydt, R. (2000). Coding of border ownership in monkey visual cortex. *J. Neurosci.* *20*, 6594.