UNIVERSITY OF CALIFORNIA, SAN DIEGO

On the Control of Selective Attention in the Primate Superior Colliculus

A dissertation submitted in partial satisfaction of the
requirements for the degree Doctor of Philosophy
in
Neurosciences

by

Lee Phipps Lovejoy

Committee in charge:

Professor Richard J. Krauzlis, Chair
Professor Thomas D. Albright
Professor Steven A. Hillyard
Professor Harvey J. Karten
Professor David Kleinfeld
Professor John H. Reynolds

2010

The dissertation of Lee Phipps Lovejoy is approved, and it is acceptable in quality and form for publication on microfilm and electronically:

_____

_____

_____

_____

_____

_____
Chair

University of California, San Diego

2010

DEDICATION

To all committed to the endeavor.

EPIGRAPH

Millions of items of the outward order are present to my senses which never properly enter into my experience. Why? Because they have no *interest* for me. *My experience is what I agree to attend to.* Only those items which I *notice* shape my mind—without selective interest, experience is an utter chaos.

William James, Principles of Psychology (1890)

TABLE OF CONTENTS

LIST OF ABBREVIATIONS

CM-Pf	Centromedian-Parafascicular Complex
CN	Caudate Nucleus
DLPFC	Dorsolateral Prefrontal Cortex
DMF	Dynamic Memory-Field
FEF	Frontal Eye Fields
ips	impulses per second
LGN	Lateral Geniculate Nucleus
LIP	Lateral Intraparietal Area
MD	Mediodorsalis, Medial Dorsal Nucleus
MF	Memory-Field
MT	Mediotemporal area of Visual Cortex
SC	Superior Colliculus
SGI	Substantia Griseum Intermedium
SGP	Substantia Griseum Profundum
SGS	Substantia Griseum Superficiale
SN, SNpc, SNpr	Substantia Nigra, pars compacta, pars reticulata
SO	Stratum Opticum
TRN	Thalamic Reticular Nucleus

LIST OF FIGURES

LIST OF TABLES

ACKNOWLEDGMENTS

I would like to thank all of my committee members, in particular Rich Krauzlis, for his support, guidance, tolerance, and patience throughout my graduate career, and Harvey Karten, for encouraging me to ground my thinking in precise consideration of the anatomy. In addition, I thank Paul Insel and the UC San Diego Medical Scientist Training Program for bringing me to UCSD, and the Bert and Ethel Aginsky Scholar Award and the Institute for Neural Computation for their support during my tenure as a graduate student. This work would not have been possible without the extensive help of our laboratory technicians, including Natalie Dill, Eileen Boehle, and Krista Kornylo, the staff of the Salk Institute, including the Animal Resources Department and Facility Services for construction of equipment, and the CNL computer support personnel. Also, much of the analysis would not have been possible without access to the computing resources of the Computational Neurobiology Laboratory of Terry Sejnowski. Finally I thank my mother, Betty Lovejoy, without whom I would not have been possible.

Chapter 2, in full, appears as it was published in *Nature Neuroscience*, 2009, Lovejoy, L. P., Kruazlis, R. J. The dissertation author was the primary investigator and author of this paper.

Chapter 3, in full, appears as it may be submitted for publication. The dissertation author was the primary investigator and author of this paper.

VITA

1995–1999	Dean's Honor Scholarship, Tulane University
1996	Tau Beta Pi
1997–1999	Barry M. Goldwater Scholar
1999	B. S. E. in Biomedical Engineering, Tulane University
1999–2000	Tau Beta Pi Graduate Fellowship
2000	M. S. E. in Biomedical Engineering, Tulane University
2000	Medical Scientist Training Program, University of California, San Diego
2002–2003	Merck Fellowship
2006–2009	Aginsky Scholar
2007–2008	Institute for Neural Computation Predoctoral Fellow
2008	Elsevier/Vision Sciences Society Travel Award
2010	Ph. D. in Neurosciences, University of California, San Diego

PUBLICATIONS

Lovejoy, L. P., Krauzlis, R. J., 2009: Inactivation of primate superior colliculus impairs covert selection of signals for perceptual judgements. *Nature Neuroscience*.

Lovejoy, L. P., Fowler, G. A., Krauzlis, R. J., 2009: Spatial allocation of attention during smooth pursuit eye movements. *Vision Research*, **49**(10), 1275-1285.

Krauzlis, R. J., Lovejoy, L. P., 2008: Walk this way. *Neuron*, **60**(1), 7-8.

Nummela, S. U., Lovejoy, L. P., Krauzlis, R. J., 2008: Saccade selection when reward probability is dynamically manipulated using Markov chains. *Experimental Brain Research*, **187**(2), 321-330.

Lovejoy, L. P., Shepard, P. D., Canavier, C. C., 2001: Apamin-induced irregular firing *in vitro* and irregular single-spike firing observed *in vivo* in dopamine neurons is chaotic. *Neuroscience*, **104**(3), 829-840.

PRESENTATIONS

Lovejoy, L. P. and Krauzlis, R. J., 2009: Primate superior colliculus regulates selection of signals for perceptual judgments, *Society for Neuroscience Abstracts*.

Lovejoy, L. P. and Krauzlis, R. J., 2008: Primate superior colliculus gates sensory information for perceptual decisions, *Society for Neuroscience Abstracts*.

Lovejoy, L. P. and Krauzlis, R. J., 2008: Inactivation of superior colliculus causes visual extinction, *Vision Sciences Society Abstracts*.

Lovejoy, L. P. and Krauzlis, R. J., 2007: Inactivation of superior colliculus biases the allocation of voluntary spatial attention, *Society for Neuroscience Abstracts*.

Krauzlis, R. J. and Lovejoy, L. P., 2007: Neural correlates of voluntary attention in the superior colliculus, *Society for Neuroscience Abstracts*.

Lovejoy, L. P. and Krauzlis, R. J., 2006: Superior colliculus activity during manipulations of voluntary and stimulus-driven attention, *Society for Neuroscience Abstracts*.

Lovejoy, L. P., Chukoski, L., and Krauzlis, R. J., 2005: Attention capture alters motion discrimination, *Vision Sciences Society Abstracts*.

Lovejoy, L. P., Fowler, G. A., and Krauzlis, R. J., 2004: Attention is allocated to the intended target during smooth pursuit, *Society for Neuroscience Abstracts*.

Lovejoy, L. P., Fowler, G. A., and Krauzlis, R. J., 2004: Allocation of spatial attention during fixation and smooth pursuit, *Vision Sciences Society Abstracts*.

Lovejoy, L. P., Liston, D., Chukoskie, L., and Krauzlis, R. J., 2003: Optimal information transmission in superior colliculus neurons, *Society for Neuroscience Abstracts*.

Lovejoy, L. P., Shepard, P. D., and Canavier, C. C., 2000: Irregular firing observed in vivo in dopamine neurons is chaotic and dynamically similar to apamin-induced irregular firing, *Society for Neuroscience Abstracts*.

Lovejoy, L. P., Shepard, P. D., and Canavier, C. C., 1999: Irregular firing induced by apamin in dopamine neurons is chaotic, *Society for Neuroscience Abstracts*.

ABSTRACT OF THE DISSERTATION

On the Control of Selective Attention in the Primate Superior Colliculus

by

Lee Phipps Lovejoy
Doctor of Philosophy in Neurosciences
University of California, San Diego, 2010
Professor Richard J. Krauzlis, Chair

Selective attention is a fundamental aspect of our experience of both the internal and external environment. I deal here with its specific incarnation as spatial attention, in which behavior is guided by one of several competing visual signals. The superior colliculus (SC), a brainstem structure best known in primates for its role in the control of voluntary eye movements, has long been suspected of a role in the control of attention. Nevertheless, proof of such a connection has remained elusive since we normally selectively process visual information by directing gaze to objects of interest. I here present evidence that causally links neural activity in the SC to the control of selective attention independent of its role in the control of eye movements.

We conducted reversible lesion studies in primates trained on a selective attention task. The task required the monkeys to covertly attend to some stimuli without moving their eyes and to ignoring misleading stimuli. Information at the attended location instructed the monkeys where to move their eyes or which buttons to push in order to receive rewards. Reversible inactivation of the SC caused the monkeys to ignore information in the cued location and instead to base decisions on the misleading signals. This change in behavior could not be attributed to a primary sensory deficit nor to the motor consequences of lesion of the SC, thus proving that normal SC function is required for selective attention.

In addition, we performed neurophysiological recordings in the SC during the task. I describe the responses of functional classes of neurons suspected to regulate

attention and how these groups may contribute to the control of selective attention. In particular, I describe a group of cells with oscillatory activity associated with working memory and with selective attention. The characteristics of these neurons during the task suggest that they could be critical for selecting which stimuli will be perceived or will guide behavior.

Chapter 1

Introduction

We performed a set of experiments to determine whether or not the circuitry controlling the selection of stimuli for visual processing has components located in the superior colliculus (SC). Prior to describing these experiments in subsequent chapters, I will begin with a review of our understanding of the basic phenomena under consideration and some relevant background on the superior colliculus and its connection to attention. I continue with a review of a functional deficit known as unilateral neglect in which the selection of signals that guide behavior and perception is severely impaired, followed by a more precise description of the known functional connections between neglect and the superior colliculus. This leads to a proposal about how the superior colliculus might be integrated into selection circuitry that would allow it to influence both the control of orienting movements and the selection of stimuli for visual processing. I end with the main hypotheses tested in the experiments.

1.1 Selective attention

Simply put, the purpose of any nervous system is to control movement. All interactions with the environment involve movement of one form or another, and under any given circumstance some movements are much more likely to lead to the survival of the organism than others. Sensory systems allow organisms to

determine the current state of the environment and, accordingly, prepare and execute an appropriate movement in response. Even the environment of the smallest, simplest organism is extraordinarily complex, and the organism must select which aspects of the environment will guide its behavior. This need is the fundamental motivation of the selection in selective attention—given a limited set of possible behaviors and a rich and complex environment, which features of the environment should be selected, or attended to, in order to choose that behavior which will maximize the odds of the organism's survival? If we consider that the environment is a dynamic place and that the animal has at least some ability to retain in memory a subset of relevant features of its immediately preceding experience, the question recursively becomes even more complicated: given a rich and complex, changing environment, which features should be retained in memory in order to make decisions about which stimuli would be reasonably expected to be relevant and, as a result, should be attended to in order to make decisions about which behaviors should be performed, and so on. It quickly becomes apparent that there are deep functional links between the notions of memory and attention, expectation and prediction, and that all of these cannot occur without the other. Selective attention is as much about a relationship between an organism and its internal environment as it is about its relationship with its external environment. At some level, there is little distinction between the two.

In Principles of Psychology (1890), William James fully embraced the ephemeralness of the distinction between external and internal environments and spoke of attention as an action applied to both. In an appeal to common sense and shared experience, he said:

> Everyone knows what attention is. It is the taking possession by the mind, in clear and vivid form, of one out of what seem several simultaneously possible objects or trains of thought. Focalization, concentration, of consciousness are of its essence. It implies withdrawal from some things in order to deal effectively with others, and is a condition which has a real opposite in the confused, dazed, scatterbrained state.

In James' view, the agent capable of selection is the mind, and it acts upon objects, or trains of thought. Thus his conception of *object* was not something in the material

world with form and substance; instead, it has a more procedural sense—the objects of attention are those things that can be subjected to the selection process of the mind, so mental images, percepts, ideas, potential action, and the like could all be objects. James drew a distinction between *sensory* objects and *ideal* or *representational* objects; sensory objects are elicited by stimuli external to the mind, and ideal or representational objects exist only in the mind. Yet a key connection was that sensory objects existed only fleetingly, and if they were to persist in the mind, they would be replaced by a corresponding representational object.

James went on to describe attention in terms of dichotomies of the motivations underlying the relationship between attention and any given object. James took a rather Darwinian stance in his views on attention—he believed that "the practical and theoretical life of whole species, as well as of individual beings, results from the selection which the habitual direction of their attention involves." Although such a theme is relatively implicit throughout his explication of attention, the dichotomies are nonetheless rooted in the relevance of an object to both personal and species survival.

The first of these dichotomies was that attention is either *immediate* or *derived*, a distinction which drew upon the heritage of one's interest in the objects. Attention is *immediate* when the object is inherently interesting, not in the sense that the stimulus has some intrinsic value but instead because the nature of the basic machinery of the mind is to regard as of unique importance a set of properties of those objects which are of special significance to our survival independent of any other consideration. Attention is *derived* when attention to an object is due to its association with some other object of immediate interest. This distinction appears to invoke the notion that immediate attention is reflexive and the objects to which it responds are characteristic of the species and the environment in which it has evolved. It is juxtaposed against derived attention, which is specific to the individual and depends critically on learned associations.

The second dichotomy was that attention is *passive*, *reflexive*, or *non-voluntary*,

in contrast to being *active* or *voluntary*. The key notion here is that passive, immediate attention to a sensory object is directly related to the intensity or suddenness of the sensory input, and will only persist if it is continuously engaged by a changing object; in contrast, passive, derived attention to a sensory object occurs when the stimulus, no matter how slight, is strongly associated with something of intense interest. Attention is *voluntary* when it takes effort, most often experienced when we attempt to discriminate some faint aspect of a sensory stimulus. In contrast, *passive* attention does not take effort. In James' construction, voluntary attention occurs when an effort must be made to repeatedly bring an object back to mind, which then momentarily engages passive attention.

The similarities between much of the conception of attention advanced by James and more current descriptions is apparent. In particular, James's notion of immediate attention for sensory objects would appear to be equivalent to *salience*. As a term, *salience* refers to the notion of importance and thereby to the properties of an object in relation to a sensory system. Furthermore, the distinction between immediate and derived attention is very similar to that advanced by Posner (1980), namely *endogenous* versus *exogenous* attention. In that conception, selective attention due to factors having an internal cause or origin, such as learned associations or memory, is endogenous, whereas if it develops due to external factors, it is exogenous.

Unlike the clear equivalence between endogenous versus exogenous and derived versus immediate attention, distinctions based on voluntary or passive attention have continued to bedevil experimental psychologists. From a modern perspective, the idea that any attention could be "effortless" might seem puzzling, as would the idea that any form of attention would not be active. We must keep in mind that James still thought of the mind as something potentially separate from the brain and, in a sense, in conflict with the reflexive operation of the physical corpus. Of course we are quite comfortable with the notion of conflict between different impulses in the control of attention, and so we can relate James's distinction as one between transient, or temporally limited, attention in response to external events, like James's

passive, and sustained attention independent of external events, like James's active (e.g. Nakayama and Mackeben, 1989). The question of reflexiveness has a modern form in the characterization of automaticity (Yantis and Jonides, 1990), which describes the degree to which a process is automatic in terms of the degree to which it can be suppressed. Similarly, the observation that external events can become more noticeable due to a learned association (passive derived attention, in other words, contingent upon a learned association) is known as contingent attention capture (Folk et al., 1994). Both of these formulations attempt to describe the relationship between systematic contributions to selective attention. The experimental approach to both of these invokes the notion that to actually determine if attention were automatically controlled or captured, one would have to observe the phenomenon occur in an *involuntary* way, as in, contrary to volition or desirability. In other words, in order to know that attention were fully automatic or that attention were allocated purely based on stimulus features rather than in response to an interest in the stimulus, we would have to construct a circumstance in which the subject should be entirely disinterested in the stimulus. As sensible as such an approach might seem, it proves to be quite impractical in non-human primates since subjects inevitably learn to quash the effects if they are disadvantageous to obtaining a reward. We might consider that this approach conflates the notion of voluntary with the notion of advantageous: is there a sense in which we do not elect to survive? If an automatic process, which initially allowed an organism to respond advantageously to elements of its environment, later became disadvantageous, then that organism would be at a significant disadvantage if it could not learn to suppress the process.

An astute reader may point out that the distinctions between transient and sustained attention or between voluntary and automatic attention seem to incorporate phenomenological aspects such as timing into something little more than a riff on endogenous versus exogenous attention. In contrast, James regarded attention has having substantially more varieties. We might consider that empirical psychologists have sought to ground theories of attention in terms of empirically observable

elemental mechanisms, whereas James was describing the procedural and purposeful mechanics of attention to both observable phenomena and unobservable phenomena such as thoughts themselves. Therefore experimental psychologists refer to attention to a sensory stimulus, or describe a stimulus as being salient, whereas James might refer to attention to the sensory *object* or the representational *object*, and the sensory *object* as having immediate interest. One might ponder which of these two levels of description would be more appropriate to neuroscience, in which we attempt to directly observe the neural representations of stimuli and computations performed on those representations.

Through a mixture of subjective introspection and consideration of emerging experimental evidence, James postulated that the observable effects of selective attention would be enhancements in our ability to *perceive*, *conceive*, *distinguish*, and *remember*, and that attention would *shorten reaction time*. By *perceive*, James meant that our sense impressions would be stronger, and that our ability to detect and discriminate features of sensory objects would be enhanced. By *conceive*, James meant that our ability to form mental images of sensory stimuli would be enhanced in a similar manner. Interestingly enough, James pondered the perennial paradox that the intensity of our impression of a stimulus might increase, but our impression of the intensity of the stimulus does not. He ascribed this to an enhancement in our ability to *distinguish*, in the sense of intellectual discrimination, or an increase in the perception of relations or subdivisions within our conceptions of things. Also, as is intuitively clear, attention is critically related to our ability to *remember*: "an object once attended to will remain in the memory, whilst one inattentively allowed to pass will leave no traces behind." Finally, James defers to the analysis of Wilhelm Wundt, the father of empirical psychology, in his assessment of the shortening of reaction time (Wundt, 1874), with italics added for emphasis:

> When we wait with strained attention for a stimulus, it will often happen that instead of registering the stimulus, we react upon some entirely different impression, and this not through confounding the one with the other. On the contrary, we are perfectly well aware at the moment of making the movement that we respond to the wrong stimulus... *We can-*

> *not well explain these results otherwise than by assuming that the strain of the attention towards the impression we expect coexists with a preparatory innervation of the motor centre for the reaction, which innervation the slightest shock then suffices to turn into an actual discharge.* This shock may be given by any chance impression, even by one to which we never intended to respond. When the preparatory innervation has once reached this pitch of intensity, the time that intervenes between the stimulus and the contraction of the muscles which react, may become vanishingly small.

It is interesting to note that the postulated link between attention and the "innervation of the motor centre for the reaction" justified the use of reaction time as the primary measured quantity in many of the experimental studies which are currently regarded as foundational to our modern understanding of selective attention. In order for this link to exist, attention would have to simultaneously lead to an increase in perception of the stimulus, distinction between the potential forms which the stimulus could take, and both conception and preparation of potential responses. Nevertheless, using the decrease in reaction time is currently less favored because many investigators believe that it is, in some sense, merely a motor effect, as if attention could only apply to sensory stimuli and not to our own potential behaviors. Such an approach would seem to necessitate that attention to a sensory stimulus, including the formation of the perceptual categories into which that stimulus could fall, was naturally separate from attention to the potential responses, and that in experiments the two are unnaturally confounded. I am unaware of any evidence to suggest that such a separation is necessarily the case, in fact, quite the opposite. For example, the formation of perceptual categories appears to be intrinsically linked to the generation of responses in monkeys (Gold and Shadlen, 2000). Working memory appears to be similarly related to oculomotor control, leading to the proposal that working memory is related to oculomotor "rehearsal," meaning the preparation of eye movements without their execution (Smyth and Scholey, 1994; Awh et al., 1998). An analogous outcome is observed in the control of selective attention and eye movements (Deubel and Schneider, 1996; Kowler et al., 1995; Hoffman and Subramaniam, 1995). In fact, despite the seeming balkanization of the study of attention into the disparate com-

ponents that James identified as all facets of the same process, little evidence exists to indicate that they are normally separable (reviewed in Olivers, 2008).

Given this constellation of changes, James deduced that the internal mechanisms of selective attention would depend on a pair of functions: first, the accommodation or adjustment of the sensory organs, and second, the anticipatory preparation from within of the ideational centers concerned with the object to which one attends. When interpreting these ideas we should keep in mind that historically, the term *organ* can refer not simply to a self-contained part of an organism with a vital function such as, in this case, the eye, but also to a region of the brain held to be the seat of a particular faculty. Although in 1890, notions about how the machinery of mind was incarnated in the biological processes of the brain were at best murky and incipient, the recognition was present that the mind was most likely manifest in the operation of the brain. Therefore we might safely presume that James was presaging the idea that attention was associated both with changes in the function of neural circuitry related to the detection and processing of stimuli and with operation of neural circuitry associated with controlling which stimuli will be incorporated into ongoing decisions about behavior.

In recent decades, the study of attention has seemed to focus nearly exclusively on changes in performance on the discrimination and detection of sensory stimuli, probably because these are empirically observable and can be systematically explained with frameworks such as Signal Detection Theory (Green and Swets, 1966). Changes in detection and discrimination are associated with the control of sensory processing, and some of the most influential theories attribute selective attention to control at this level. For example, in Broadbent's Filter Model (Broadbent, 1958), a filter controls the access of sensory stimuli into working memory; the selected stimuli is passed into working memory and subjected to further processing for meaning, while the ignored stimuli are routed into a buffer and held until working memory is free. Such a model is a paradigmatic example of the tendency to restrict *attention* to the control of sensory processing and to segregate it conceptually from other related

phenomena such as working memory. This distinction between attention and working memory persists despite the fact that the experiments used to investigate both are extremely similar and seem to differ only in the theoretical lens through which the results are interpreted (reviewed in Olivers, 2008).

A large amount of experimental work in non-human and human primates has identified changes in the properties of neurons during selective attention. The increases in discrimination and detection performance (e.g. Yeshurun and Carrasco, 1998) are associated both with changes in the response properties of neurons in many areas of visual cortex (Reynolds and Chelazzi, 2004; Cook and Maunsell, 2002; Williford and Maunsell, 2006) and with changes in the relationship between the firing properties of neurons (Mitchell et al., 2009; Cohen and Maunsell, 2009). These changes in response properties would be expected to increase the efficiency and rate with which information about visual stimuli are encoded by those neurons. Although the work has been performed under the rubric of working memory, the associated effects of attention have also been studied in frontal and parietal cortex, as will be discussed below and in Chapter 3.

Before moving on to the primate superior colliculus and the control of selective attention, I must introduce one final distinction in the control of selective attention—*overt* versus *covert* attention. In general, we selectively attend to visual stimuli by directly orienting to them, but as we are all well aware, we can orient in one direction while attending to stimuli elsewhere in the visual field. This occurs when it is socially inappropriate to directly stare at a person, for example, or when we wish to conceal the object of our attention from others who might attempt to infer it from the direction of gaze (hence the term covert). These are daily experiences for us all. Although covert attention has long been considered in theoretical treatments of attention (in fact, James describes nineteenth century efforts to study the phenomenon), a basic inability to record the position of the eyes provided a substantial technical obstacle to its experimental investigation. Technological advances in our ability to measure the direction of gaze led to the formalization of the study

of covert attention (Posner, 1980). Covert attention can be subjected to quantitative measurement in an experimental setting by asking subjects to attend to stimuli without moving their eyes and then by asking them to perform fine discrimination or detection judgements about those stimuli. If the performance on the task is improved by instructions about spatial locations, then we conclude that the subject must have been covertly attending to the stimulus in response to the cue. This approach is one of the most commonly used paradigm for studying attention. An alternative to this approach would be to provide the subject with multiple stimuli that they are trained to associate with certain behavioral responses other than directly orienting to the stimuli. If the subjects respond appropriately to the presented stimuli, then they must have been selectively attending to the cued stimuli rather than the distracters. This is the basic construction that we have used to monitor selective attention.

1.2 Superior colliculus and selective attention

The primate superior colliculus is a brainstem structure most often associated with the control of orienting movements such as saccades. It has a superficial layer receiving direct retinal input and a deeper compartment that has circuitry related to the generation of saccades and the selection of targets for orienting movements. Historically, the structure has been understood primarily in terms of the visual guidance of eye movements (Wurtz and Albano, 1980). The superior colliculus contains a map of visual space in the sense that neurons in a particular portion of the map respond to the presence of a target at the corresponding position in the visual field (the receptive field) or to a saccade made to that position (the response, or movement field), and stimulation of activity at the appropriate depth can evoke a saccade to that location (Robinson, 1972).

The superficial layers of the SC, the stratum griseum superficiale, (SGS), and the stratum opticum (SO) have neurons which appear to be most related to visually-guided orienting behavior. In addition to a large retinal input into the SGS

(Schiller and Malpeli, 1977), the superior colliculus receives substantial input from the cerebral cortex. This input is in register with the retinotopic map corresponding to the retinal input. Striate cortex and prestriate recipient areas project to the superficial layers, and the motion sensitive area MT projects into both the SGS and SO (Lock et al., 2003; Ungerleider et al., 1984). This pattern makes sense in light of the fact that just ventral to the SO is a group of cells that appear to mediate orienting movements triggered by visual stimuli (Mohler and Wurtz, 1976). Visual cortex substantially enhances the processing capacity which would be possible with the SGS alone, and so projections into the SGS and SO might allow the SC to produce orienting movements as if the processing were accomplished in the SGS. This pattern of connectivity would seem to make ethological sense because it would allow the animal to orient to biologically important stimuli, in particular small moving object which may correspond to prey.

Neurons in the superficial layers may also be involved in the regulation of visual processing at the target of an upcoming saccade. Many neurons in the superficial layers appear to project through pulvinar to visual cortical areas (Chomsung et al., 2008; Kaas and Lyon, 2007; Raczkowski and Diamond, 1978; Robson and Hall, 1977; Trojanowski and Jacobson, 1975; Casagrande and Daimond, 1974). Since the discharges of superficial layer neurons in response to a visual onset are enhanced if the monkey is about to make a saccade to that stimulus (Wurtz and Mohler, 1976; Goldberg and Wurtz, 1972b), and because small lesions of the pulvinar can lead to deficits on an attention task (Petersen et al., 1987), it has been believed that the SC to pulvinar to visual cortex loop is a primordial attention control system that enhances visual processing at the targets of upcoming eye movements. Such a system might substantially increase the efficiency of orienting and associated behaviors such as hunting or feeding.

The intermediate and deeper layers of the SC, the stratum griseum intermediate (SGI) and the stratum griseum profundum (SGP) have neurons which appear to be associated with the production of saccades, including neurons which are activated

long before the saccade is actually produced. As a result, these neurons were suspected of being involved in selecting targets for saccades (Wurtz and Albano, 1980). Cortical input to these cells is supplied largely by frontal and parietal cortical areas such as dorsolateral prefrontal cortex (DLPFC), frontal eye fields (FEF), and lateral intrapariatal area (LIP), and visual areas such as V4 (Lock et al., 2003; Sommer and Wurtz, 2000; Paré and Wurtz, 1997; Johnston and Everling, 2006). Since these regions accomplish complex transformations of information and incorporate factors such as memory and other goals into representations of space, these cortical inputs would allow the SC to produce eye movements *based* on visual stimuli rather than simply *towards* visual stimuli. In essence, the cortical inputs provide a synthetic target to which the SC can generate an orienting response. Furthermore, this function would seem to require selective attention for the stimuli which inform the generation of the synthetic target, regardless of whether those stimuli appear in the visual field or are remembered.

If such selection of stimuli and the transformation of their information into a synthetic target is to take place, then the normal orienting response generated by the SGS and visual cortex module would have to be suspended in favor of the action of cortical regions. Indeed, when visual cortex is ablated in monkeys, SGS neurons retain some responsiveness to visual stimuli but neurons in the SGI become almost entirely unresponsive (Schiller et al., 1974), which led to the erroneous belief that no functional connection existed between the visually-guided orienting module of the SGS and the intermediate and deep layers. Instead, these connections do in fact exist (reviewed in Isa and Hall, 2009) but would seem to require cortical facilitation to contribute to orienting behavior. A more accurate model of the relationship may be that the SC is under substantial inhibition to prevent automatic orienting behavior, and that automatic orienting can only occur when released from inhibition by the intervention of cortical circuits. This view is supported by the observation that cats with ablation of visual cortical areas can still orient to spots of light, but not dim gratings, as long as the SC is released from inhibition (Loop and Sherman, 1977).

Although it would seem clear that selective attention is important for transforming information in the visual field (or from memory or another sensory modality) into a orienting response to another location, whether or not the SC would itself be involved in this selection is largely an open question. That it would be is suggested by the anatomy: in addition to the descending projections to the saccadic brain-stem circuitry (Moschovakis et al., 1988a,b), neurons in the intermediate and deep layers project to a variety of thalamic relay nuclei such as mediodorsalis (MD), which is thought to carry corollary discharge about eye movements to the FEF (Sommer and Wurtz, 2004b, 2002). As I will elaborate upon later, the MD also has a variety of functional properties related to memory (Tanibuchi and Goldman-Rakic, 2003), which could influence those aspects of selective attention (i.e. the anticipatory preparation of the ideational centers, in the words of James). In addition, projections from the SGI reach other thalamic relays such as the thalamic reticular nucleus (Harting et al., 1980), which has been speculated to be the primary seat of action of attentional focus (Crick, 1984) due to its position as a primary modulator of transmission of visual information through the lateral geniculate (LGN) en route to visual cortex. Lesions of the TRN cause attentional deficits in rats (Weese et al., 1999), and cells in monkey are known to be modulated by selective attention (McAlonan et al., 2006). Therefore neurons in the SGI could influence the rate at which visual information is transmitted or the specificity of visual channels for certain stimuli (i.e. the accommodation or adjustment of the sensory organs).

The ascending connectivity of the SGI and SGP would seem to be consistent with the notion that the SGI could promote the orienting of attention to the target of an upcoming eye movement. This proposition is consistent with psychophysical data that suggests that oculomotor control and orienting of attention share a common set of resources (Deubel and Schneider, 1996; Kowler et al., 1995; Hoffman and Subramaniam, 1995), which led to the resurgence of a notion that preparation of an orienting movement was synonymous with the orienting of attention, and that it was only covert if the movement was withheld. In its strongest form, this notion is known

as the Premotor Theory of Attention (Rizzolatti et al., 1987). In addition, human patients with progressive supranuclear palsy (PSP), a variant of Parkinsonism that involves degeneration of the superior colliculus, pretectum, periaqueductal gray and mesencephalic raphe, in addition to the substantia nigra, have deficits in orienting attention to regions of space that are associated with deficits in initiating eye movements to the same regions (Posner et al., 1982). Moreover, the patients appear to be unable to disengage attention from the unaffected portions of visual space even in the absence of eye movements. This observation led Posner, Cohen, and Rafal to conclude that the degeneration of the SC caused hyperactive orienting of attention to the unaffected visual space, implying that the SC was normally involved in the orienting of attention in response to exogenous factors. Furthermore, response time is delayed in a covert attention task following reversible lesion of the SGI in monkeys (Robinson and Kertzman, 1995). Even when saccades are not produced, neurons in the SGI show elevations in firing rate related to covert attention (Kustov and Robinson, 1996; Ignashchenkova et al., 2004). Finally, microstimulation of the superior colliculus causes improvements in motion discrimination and detection tasks as if attention were driven to the corresponding portion of the visual field (Cavanaugh and Wurtz, 2004; Müller et al., 2005). These converging lines of evidence all support the notion that the SGI can induce orienting of attention to the target of an upcoming orienting movement, just as appears to be the case for the SGS.

All of these results would seem to indicate that the primate superior colliculus can orient attention to a position in space, particularly in response to exogenous factors such as transiently appearing stimuli, even if the saccade is not made, and that cerebral cortex may be required for more complicated control of selective attention. Following work such as Posner's studies of human patients with neurological disorders, (Posner and Petersen, 1990), the control of attention in response to both exogenous and endogenous factors is often assumed to be a cortical function. In humans, imaging studies show that covert attention and eye movements activate a common set of areas in the parietal and frontal cerebral cortex (Corbetta and Shulman, 2002). Likewise,

in non-human primates, cortical areas important for eye movements have also been implicated in the control of covert attention. For example, reversible inactivation of FEF causes temporary deficits in performance on covert attention tasks (Wardak et al., 2006), and electrical stimulation of the FEF improves performance on detection tasks and promotes the enhancement of visual processing in area V4 (Moore et al., 2003; Moore and Armstrong, 2003; Moore and Fallah, 2001). The LIP is implicated in the control of both saccadic eye movements and spatial attention (Snyder et al., 1997; Bisley and Goldberg, 2003). Overall, the predominance of evidence would seem to indicate that both the visual circuitry in the SGS and the oculomotor circuitry in the SGI could orient attention to the target of an upcoming saccade even if the saccade is not made. But whether or not the activity of the SC is necessary for these functions or merely supplements that of the cerebral cortex has remained, until now, a matter of speculation.

A summary of the relationship between the SC and cortex regarding the control of orienting movements and attention would seem to distill down to the following basic architecture. The superficial layers of the SC, including SGS and SO, mediate visually-guided orienting movements in vertebrates, in particular those without well-developed cerebral cortex. This function is largely suspended in primates in favor of cortical visual processing, which is fed back into the superficial SC in register with the retinal projections. The intermediate and deep layers of the SC, including the SGI and SGP, are involved in the generation of orienting movements including, but not limited to, those that are visually-guided. In addition, the SGI is involved in the selection of targets for these movements. The SGI receives cortical input from regions which transform visual and other information into synthetic targets which are operated on by the SC as if they were visual targets. An aggressive statement of this idea would be that the primate SC does not even truly respond to visual targets in the traditional sense, and that instead visual targets are replaced by synthetic targets via cortical input into the SC. Ascending projections from the SGS to the pulvinar may alter visual processing in striate visual cortex (V1) and striate recipient areas

such as MT. This may provide a primordial control of attention based on exogenous factors, and this system may remain active in primates even if this module no longer is primarily responsible for visually-guided orienting. Ascending projections from the SGI to MD and other thalamic relay nuclei may influence the selective processing of stimuli which are transformed by frontal and parietal cortex into the targets for orienting movements, or this function could be entirely accomplished within cerebral cortex. In addition, projections from SGI to TRN could be involved in the regulation of sensory processing in visual cortical areas by modulating the transmission of visual information in LGN. Overall, this model supports the notion that the SC controls selective attention at the level of orienting responses, that it promotes orienting of attention to the potential targets of eye movements, and that it might possibly also influence the processing of stimuli which inform the targets of orienting movements without themselves becoming the targets of those movements.

Although such a model may provide a relatively parsimonious relationship between known anatomy, physiology, and behavior, and has unknown components which would be tractable to further experimental elaboration, it ignores some other major functional connections. In the following sections I will describe a disorder known as unilateral neglect, which will provide a better context in which these ideas can be explored. As I describe below, neglect appears to be a generalized disorder in spatial cognition and occurs in a variety of reference frames. One hypothesis would be that activity in the SC is required for the generation of those reference frames. Another possibility, based on the premise that the superior colliculus does not appear to be able to perform its role in the selection of targets for orienting movements without the participation of the basal ganglia, in particular the substantia nigra, is that selection deficits for both orienting behavior and attention arise out of a dysfunction in a basal ganglia circuit. Overall, investigation of these factors will reveal inadequacies in the preceding description of how the SC and cerebral cortex interact in the control of selective attention and will demand the development of a more sophisticated explanation.

1.3 Unilateral neglect and visual extinction

1.3.1 The neglect syndrome

In Chapter 2, I describe the outcome of a series of experiments which appeared to induce visual extinction, a feature of the neglect syndrome, via reversible lesion of the superior colliculus. Therefore an explanation of what the neglect syndrome is and how the appearance of this phenomena would influence notions of the function of the SC, deserves at least a brief treatment.

Although one might readily imagine that the behavior of a patient with a neurological disorder reflects operation of the circuitry which is so outside the range of normal so as to be uninformative about typical function, inferring function from deficit has become a central tool in the theory underlying neurology. The principle is well expressed in the words of the nineteenth-century French physician Claude Bernard, that symptoms of disease do not reflect a morbid state, but merely a normal physiological process that has become exaggerated. Such a theory is the foundation of allopathic medicine as practiced in the western world. It justifies the notion that observing which factors are required to maintain homeostasis allows one to infer that the damaged portion of the brain provided the restorative function, or that manipulation of the circuitry such as by compensatory lesions would restore normal function by achieving balance between opposing functional forces. As a result, we seek to learn about the function of regions by studying the effects of destructive lesions of those regions, either naturally occurring in patients or experimentally induced in animal models. Therefore we study unilateral neglect to understand the normal control of selective attention.

The neglect syndrome includes an extremely diverse range of spatially specific deficits in attention and movement that cannot be attributed to a primary motor or sensory deficit (e.g. Driver and Mattingley, 1998; Heilman and Valenstein, 2003). Neglect is generally unilateral, or specific to one sensory hemisphere, and is thus distinct from generalized disorders of arousal or lethargy. Neglect patients fail to explore

regions of visual space and ignore stimuli presented in those regions. In some cases this neglect appears to be primarily a defect in the ability to either orient into a region or to initiate movements into a region, particularly when competing targets for movements are present. This variety of neglect is sometimes referred to as *intentional neglect* and overlaps with movement disorders known as akinesia (Heilman, 2004), particularly when no other sensory impairments are seen. Intentional neglect arises from damage to frontal cortex and subcortical structures such as the basal ganglia, and on tasks such as visual search can appear to have sensory components which are difficult to distinguish from the defects in orienting responses (Husain and Kennard, 1997; Karnath et al., 2002).

Neglect patients often experience a wide range of sensory and cognitive deficits, particularly following right-parietal cortex damage. In an infamous litany of symptoms dealing with spatial relationships, patients sometimes eat food from only the right side of the plate, miss words on the left of the page or letters to the left within words when reading, ignore people and objects on the left, only copy the right hand side of symmetric images, and only recall the right hand geography of places and terrains visited prior to injury. Some neglect patients are only capable of right-sided visual imagery and even have right-sided visual hallucinations (Rafal, 1994). Patients sometimes ignore the left half of their bodies, failing to dress or shave themselves appropriately. The deficit is not simply retinotopic; it can appear in any conceivable reference frame, including personal body space and object centered space (patients sometimes ignore the right side of flipped pictures of faces, for example). The striking variability in these effects should really come as no surprise when one considers that parietal cortex seems to be involved in the construction of arbitrary reference frames for spatial information and contains generalized maps of intentions for movement (e.g. Andersen and Buneo, 2002). Hence the precise overlap between any particular patient's lesion with his individual variant of parietal cortex functional localization would lead to enormous variability in symptoms between patients. Overall, it seems as if the symptoms of neglect could be appropriately described as defects

in attention, orienting, internal representation of space, spatial working memory, spatial cognition, or any other cognitive process that impinges in any way on a notion of space or spatial action selection. Yet what is perhaps most striking about these patients is that, in general, they are initially completely unaware that they have any problem whatsoever, meaning that there has been a fundamental perturbation in their very ability to conceive of spatial information within the particular reference frame best matching their symptoms.

Regardless of the peculiarities of the symptomatology, one of the key conceptual distinctions to make when diagnosing neglect is that the deficits do not arise as a result of a primary motor or sensory deficit. However this should not be taken to mean that there are no primary motor or sensory deficits in neglect patients. Provided that we could adequately define what a primary motor or sensory deficit is, unilateral neglect could be so severe as to preclude our ability to make such an assessment. Alternatively, neglect could be comorbid with a primary sensory deficit. For example, cortical damage could be so severe as to eliminate both primary visual cortex and parietal cortex, but this does not somehow rule out the presence of unilateral neglect. Similarly, most neglect patients have at least some defects in movement, particularly oculomotor control, which may or may not be due to the sensory problems or sympatric localization of sensory and motor control circuits in the damaged region of the brain. Therefore, in an animal model of neglect, it would seem that great care should be taken to precisely characterize any sensory deficit which might exist and exclude that as the sole contributing factor in the appearance of neglect.[1]

Since spatial cognition appears to be impaired in unilateral neglect, it stands to reason that a potential means by which an SC lesion could lead to neglect would be

[1]In addition to ruling out sensory deficits, neurologists must also rule out contributing factors such as dementia when diagnosing unilateral neglect. Dementia is often associated with dissociative disorders from limbs or regions of the body or space, and could be confused with neglect. It is intriguing to note that in some earlier treatments of animal models of neglect, investigators found it necessary to attempt to determine that the experimental subjects were not demented. For instance, Raczkowski, Casagrande, and Diamond (1976) report of their tree shrews: "Yet they were not blind; they did not suffer from some general dementia; nor did they show any apparent defect in eye movements, or pupillary response or any other motor defect." Nevertheless, it begs the question of how exactly one might recognize a demented tree shrew.

if the normal action of parietal cortex depended in some way on the normal function of the SC. One model of neglect holds that the symptoms arise from defects in the construction of reference frames by the loss of spatial transformations from sensory into other coordinate systems (Pouget and Sejnowski, 1997, 2001). Perhaps those reference frames depend on the presence of a retinotopic map in the SC. The retinotopic map provides a space with an origin and spatial relation (left-ness or right-ness with respect to straight ahead). Perhaps more abstract spaces are constructed from this activity. The idea that SC activity contributes to the construction of spatial relationships in other systems is not new; it was suggested at least by Cooper et al. (1998) in the context of rat navigation. Supporting evidence in primates is sparse. For example, some SC neurons with response fields on the right hand side of visual space will be active when the monkey is preparing to make a saccade to the rightmost of two response dots on the left hand side of visual space (Horwitz et al., 2004), which would not be expected if the SC neurons responded solely to the direction of the saccade. Instead, it would be consistent with a model in which the decision was made in a perceptual space which mapped the rostral pole of the SC to its origin and the left and right hand sides of the SC onto the left and right hand sides of the space with respect to that arbitrary origin. In addition, the response properties of SC neurons to saccades made to a particular location in space varied based on whether the location corresponded to the center, corner, or edges of an object (Edelman and Goldberg, 2003). Although this could be attributed to the variation in the configuration of the visual stimulus corresponding to the position of the object in the neuron's response field, or to variation in the exact metrics of the saccade based on these features, it is also possible that the variations reflected a contribution of the position of the saccade with respect to the object's reference frame. This could occur if the SC were part of a sequence of gain fields (Salinas and Sejnowski, 2001) required to accomplish coordinate transformations within parietal cortex. Taken together, these observations would at least be consistent with the idea that SC activity contributes to the generation of reference frames in parietal cortex, and that loss of this activity

could cause enough of a disruption in that reference frame so as to induce unilateral neglect.

1.3.2 Visual extinction

Of particular interest for the present investigation is visual extinction, a disorder within the neglect syndrome. Visual extinction is characterized by the scenario in which a patient is capable of detecting and responding to a stimulus in the affected portion of the visual field as long as it is the only stimulus in the field of view. When a competing stimulus is present in the unaffected portion of the visual field, the patient only responds to this stimulus. In practice, visual extinction occurs in the presence of multiple stimuli to which attention could be allocated for the purpose of providing a response. Furthermore, temporal order appears to be very important in inducing visual extinction—when asked about the temporal order of simultaneously presented stimuli, patients typically report the stimulus in the unaffected side as appearing prior to the stimulus in the affected side. If the stimulus in the affected side appears 200 ms prior to the stimulus in the unaffected side, the visual extinction is overcome and the patient reports the two as simultaneous (Rorden et al., 1997; Di Russo et al., 2008; Marzi et al., 2001; Angelelli et al., 1996; Chou and Schiller, 1999). Lastly, visual extinction seems to be stronger if the competing stimuli require the same class of response, and unattended stimuli cause delays in reaction time when subjects respond to those in the unaffected field (Rafal, 1994; Vuilleumier and Rafal, 2000; Rafal et al., 2006). This indicates that visual processing is still occurring in the affected field and that visual extinction occurs in the selection of stimuli for both awareness and action.

Like neglect, visual extinction appears to be primarily a disorder of spatial cognition. Substantial evidence exists to suggest that visual processing of signals in the affected region still continues (Husain and Kennard, 1997). For instance, in a line cancellation task, the patients will often cross out lines as if they are seeing phantom lines transposed from the affected hemifield into the unaffected hemifield

(Toraldo et al., 2005). Alternatively, when reproducing a drawing, some patients copy both sides of a symmetric image when presented individually, but when the whole is presented, the patients will only represent half of the global structure of the original but will add in the local features from the neglected side (Halligan et al., 1992; Lepore, 2003). In remarkably strange examples, the objects in the neglected hemifield appear in the perception of the unaffected side as hallucinations (Nys et al., 2008).

Just as with neglect, visual extinction does not require that a primary sensory or motor deficit be completely absent, only that if it is present, that it not be sufficient to explain the visual extinction. Visual extinction has been convincingly induced in non-human primate studies, see, for instance that of De Weerd, Desimone, and Ungerleider (2003), where sensory deficits were ruled out as contributing causes. In these cases, care was taken to measure psychometric curves and to use stimuli easily detected by the animal in the absence of distracters. Then, when the animal is unable to respond to the cued stimuli in the presence of distracters, a claim of visual extinction appears valid. In our inactivation studies, we took a similar approach. Once we determined that the monkeys were neglecting signals in the affected region of visual space, we measured psychometric curves to determine a signal strength at which the stimuli were equivalently visible in the affected and unaffected regions. When this strength was exceeded in the affected region, the animals still neglected the signals in that region. We therefore concluded that visual extinction had occurred.

Numerous models of visual extinction have been proposed, with more or less predictive power, and at least two are well suited for interpreting the outcome of the experiments described in Chapter 2. The first is the biased competition model (Rafal, 1994; Rafal et al., 2006), essentially equivalent to the attention model of the same name (Desimone and Duncan, 1995; Reynolds et al., 1999; Reynolds and Heeger, 2009). According to this model, a biasing signal weights stimuli in a spatially specific manner prior to a divisive normalization stage. When applied to neglect, it states that there is a lateralized bias toward orienting for both attention and oculomotor control in each hemisphere, and when the right hemisphere of cerebral cortex is damaged, the

intact side drives attention and orienting toward its preferred visual field. The exact symptoms experienced by the patient depend on what functional module is impacted by the damage. Subcortical, particularly thalamic, neglect is explained as a causative agent in inducing imbalance between the hemispheres. This model is quite consistent with the early signs of neglect after right parietal stroke, such as deviation of the eyes and head away from the affected visual field and complete inattention to that region. Extinction and selection deficits emerge because the biasing drives of both sides pass through a nonlinearity such as divisive normalization as in the biased competition account of attention. Furthermore, it or its theoretical cousins such as the Theory of Visual Attention of Bundesen (1990), have been used to quantify deficits in visual attention in neglect patients (e.g. Duncan et al., 1999). What is rarely mentioned is that damage to the left cerebral cortex in the region which would be expected to induce neglect according to this model does so only exceedingly rarely in humans. This observations begs the question of how biased competition could occur when the competitor of the right hemispheric drive is conspicuously and demonstrably not present.

Nevertheless, support for a biased competition account comes from physiological studies in monkeys in which the neural responses of visual neurons to multiple stimuli in their response fields is biased by which of the two has been cued (Desimone and Duncan, 1995; Reynolds et al., 1999; Reynolds and Heeger, 2009). Although this model is specifically in reference to how a neuron may select which of multiple stimuli within its response field will most drive its responses, a number of lesion and physiology studies have appealed to its explanatory power regardless of whether or not it is mechanistically applicable. The basic notion is that the model describes the resolution of conflict between conflicting sources of information as competition in a mutually inhibitory network. Some examples of lesion studies for which the outcome would seem to be described by a biasing of competition include those in the rat thalamic reticular nucleus (Weese et al., 1999), monkey DLPFC (Rossi et al., 2007) , and monkey area V4 (De Weerd et al., 2003). In addition to those mentioned above, the

model has been used to describe the effect of selective attention on neural responses in area MT (Womelsdorf et al., 2008) and area LIP (Bisley and Goldberg, 2003). The model has attractive properties because it is tunable in the sense of a control on the weighting and the degree of nonlinearity on those weights. Although it may be quite adequate in predicting the neural response to competing stimuli or of the response distribution of a subject during a psychophysical experiment, it would require substantially more machinery to handle competition based on temporal ordering, which has emerged as a potentially major component in visual extinction.

An alternative model for extinction that embraces temporal ordering is the doctrine of prior entry. This doctrine holds that attended stimuli are perceived before unattended stimuli (Shore et al., 2001). Rather than controlling selection via changing the weights applied to stimuli, selection is controlled via the sequence in which stimuli reach the selection process. This model seems to be very reminiscent of Broadbent's Filter Model, discussed above. In visual extinction patients, it is often entirely possible to eliminate the deficit by simply changing the temporal order in which stimuli are presented (Rorden et al., 1997) with a smooth transition from complete extinction to a complete lack of extinction simply based on the asynchrony in stimulus presentation. This model is supported by the observation of delays in the visually evoked potentials on the side of the lesion (Di Russo et al., 2008; Angelelli et al., 1996). In his studies of monkeys with FEF lesions, Schiller found that changing the asynchrony of target presentation could completely reverse biases in saccades made to visual targets despite the neglect induced by the FEF lesion (Schiller and Chou, 2000a,b, 1998). This led him to propose that neglect may be the result of temporal delays rather than underactivation of brain regions (Chou and Schiller, 1999). In addition, in intentional neglect, deficits often emerge only when simultaneous movements are planned; rather than occurring in parallel, the movements occur in sequence (Heilman, 2004).

There is a very obvious problem with stating that changing temporal order of presentation mitigated visual extinction, however. One might argue that visual

extinction only exists when more than one stimulus is presented simultaneously. If they are presented sequentially, then extinction cannot occur. Of course the counter to this is the doctrine of prior entry itself—stimuli that are attended are *perceived* prior to those that are unattended, even if they occur simultaneously, and we are truly working in the domain of the patient's or experimental subject's *perception* of the appearance of the stimuli, rather than the domain of the experimental stimulus. Visual extinction occurs in the perceptual domain, and so the fact that the stimuli are not simultaneously presented by the experimenter does not undermine a claim of extinction.

The idea that selection could be controlled by temporal order has received little investigation and only ambiguous support from electrophysiological recordings in monkeys. For example, Lee et al. (2007) found that the latency of a response to a stimulus of a given contrast was marginally higher when the stimulus appeared in the cued location. Nevertheless, it seems reasonable that temporal order could be an important component of control, particularly when the number of distracters is high and competition between them is exaggerated, and if access into working memory were limited. Overall, the biased competition model is tractable to computational analysis, intuitive, and has good predictive power until one wishes to account for phenomena such as temporal ordering. Perhaps the doctrine of prior entry could be instantiated using a diffusion or race model (Reddi and Carpenter, 2000; Ratcliff, 2006), and there is no reason the two would necessarily stand in opposition to each other.

Finally, the phenomena of intentional neglect (Heilman, 2004) are most obvious when simultaneous responses picked from several competing options is required. For example, simultaneous pointing actions are forced to occur in sequence, implying a substantial decrease in the throughput of a selection mechanism. Also, selection from between several similar actions can be delayed, as if dithering in the selection process is increased. It would seem as if the dysfunction of intentional neglect is related to the spatially specific reduction in computational capacity of a selection

mechanism, and this dysfunction can be overcome or elicited by changing the order in which competing movements reach the selection mechanism. That the same should be true by analogy for the neglect syndrome may not have been more thoroughly recognized simply because we habitually have not posed the question in the right manner. The support for neglect as a defect in a selection mechanism are more thoroughly investigated in the next section.

1.4 Target selection and the Sprague effect

The most obvious and empirically observable method by which the superior colliculus is involved in the selection of signals that guide behavior is by directly orienting the animal to relevant stimuli. In addition to its major contribution to the motor circuits that control how the eyes and head are oriented, (Wurtz and Albano, 1980; Sparks, 1999), the superior colliculus is important in selecting which stimulus will be the target of an upcoming eye movement. When the region of the SC representing a target position is reversibly inactivated, orienting responses are often misdirected to distracters appearing in unaffected parts of the visual field. This phenomena has been observed in monkeys (Mcpeek and Keller, 2004) and rodents (Felsen and Mainen, 2008). Conversely, electrical stimulation of the SC at currents too weak to directly evoke eye movements can nonetheless bias target selection toward the stimulus in the activated location, regardless of whether the target is acquired with a pursuit or saccadic eye movement (Carello and Krauzlis, 2004). These effects on target selection seem to be very similar to intentional neglect (Heilman, 2004) since neither inactivation nor stimulation of the superior colliculus in the aforementioned studies abolished the basic ability to produce orienting movements to the affected regions. Instead, the movements associated with single stimuli were delayed, as has been well established for reversible lesions of the superior colliculus (Hikosaka and Wurtz, 1985a; Quaia et al., 1998).

The appearance of target selection deficits would appear to be related to

the classic Sprague effect, in which cortical blindness and neglect caused by lesions of occipital and parietal cortex on one side of the brain were relieved by lesion of the SC on the other side (Sprague, 1966). This led to the hypothesis that the SC ipsilateral to the cortical lesion, which lacked normal corticotectal input, was suppressed by inhibition from the opposite SC via tectotectal connections. Disinhibition of the SC by ablation of the SC contralateral to the lesion or tectal commissurotomy permits the use of SC visual information and restores orienting behavior to visual stimuli in the previously neglected field. An important component of this theory is that the the SC is a source of visual information for guiding orienting behavior via the retinal projection to the SGS and then on to cortex via pulvinar (Daimond and Hall, 1969; Trojanowski and Jacobson, 1975; Robson and Hall, 1977; Harting et al., 1980; Kaas and Lyon, 2007).

This theory about the origins of the Sprague effect would seem to be inconsistent with more modern conceptions of neglect, and subsequent lesion experiments raised the possibility that it was not simply the visual information of the SC underlying the Sprague effect. For example, although ablation of the superficial superior colliculus in tree shrew caused deficits in pattern discrimination, lesion of both the superficial and deeper layers caused profound visual neglect (Casagrande and Daimond, 1974). Furthermore, transection of the predorsal bundle, carrying descending projections from the deep layers of the superior colliculus, caused profound inattention while sparing pattern vision (Raczkowski et al., 1976). In cat (Loop and Sherman, 1977), bilateral ablation of occipito-temporal cortex impaired visual orienting behavior in a number of visual discrimination tasks, such as luminance, grating, and pattern detection, but transection of the SC commissure restored orienting behavior in the brightness task and not the others. Interpreted at the time as an example of the Sprague effect, this study showed that the SC on its own is sufficient to guide orienting behavior and that visual cortex is required for more refined visual processing. Finally, in monkeys (Latto, 1978), bilateral removal of the SC increased the manual response time on visual search tasks without explicitly changing saccade scan paths,

but only when distracters were present in the visual field. All of these studies pointed towards a potential role for the deeper layers of the SC in the Sprague effect due to the manifestation of neglect-like findings.

Wallace et al. (1989) revisited the issue of relief of neglect by collicular lesion and showed that the effect did not depend at all on tectotectal interactions. Cutting the nontectotectal fibers in the commissure of the superior colliculus was sufficient to relieve parietal neglect. Ibotenic acid lesions in the superior colliculus, sparing crossing fibers, had no effect on neglect, nor did cutting the tectotectal fibers in the rostral portion of the commissure. Therefore the source of inhibition was not tectal. In a subsequent study, Wallace et al. (1990) showed that the source of the inhibitory input was the substantia nigra pars reticulata (SNpr), which provides the majority of the inhibitory input into the superior colliculus through both ipsilateral projections and contralateral projections passing through the caudal portion of the collicular commissure. Destruction of either the contralateral SC along with the caudal commissure or the contralateral SNpr relieved neglect. Small lesions in the ipsilateral SNpr also relieved neglect, although large lesions of the entire ipsilateral SN did not. Cutting the caudal commissure and lesion of either SNpr would be expected to release both sides of the SC from inhibition, which begs the question of whether a balance of activity is involved in neglect in anything more than a metaphorical sense.

Since the SNpr projections to the SC terminate exclusively in the SGI, the importance of the nigrotectal path indicated that the SGS was either uninvolved or not exclusively involved in the Sprague effect and that it was instead a function of inhibition on the SGI. Both the contralateral and ipsilateral nigrotectal projections terminate almost exclusively in the intermediate layers of the superior colliculus (SGI) with some near stratum opticum (SO) (Huerta et al., 1991; Karabelas and Moschovakis, 1985; Beckstead, 1983; Graybiel, 1978). When the results of Wallace et al. (1989) and Wallace et al. (1990) are interpreted in terms of the nigrotectal paths, it becomes more evident that the Sprague effect could be unrelated to the visual functions of the SGS and instead depend on the disinhibition of orienting cir-

cuitry in the SGI ipsilateral to the cortical lesion. It would seemingly remain possible that even though the SNpr does not directly inhibit SGS, that the loss of SGI would prevent the enhancement of SGS visual responses as previously observed (Goldberg and Wurtz, 1972b; Wurtz and Mohler, 1976). What might happen to the response properties of SGS neurons following an isolated SGI lesion would be a matter of speculation. Although we might retain an indirect effect on the SGS-pulvinar-visual cortex pathway, of already uncertain significance in primates, as a possibility, we will continue with examination of the SNpr-SGI circuit as the principal in the Sprague effect.

Although it may be established that the relevant players in the Sprague effect are the SGI and the SNpr, the results would still seem to lead to an apparent contradiction: if the SNpr is the major source of inhibition ipsilaterally, how could both an SC lesion contralateral to the cortical lesion *and* an SNpr lesion contralateral to the cortical lesion independently relieve neglect? Furthermore, while the small ipsilateral SNpr lesions would relieve inhibition on the ipsilateral SC, why would much larger lesions of the entire SN not relieve neglect? Do these results contradict or confirm the hypothesis that disinhibition of the ipsilateral SC is necessary to relieve neglect? A potential explanation for this quandry was provided by Ciaramitaro et al. (1997), who directly tested the disinhibition hypothesis by injection of bicuculline, a GABA(A) antagonist, into the SGI. Bicuculline injection into the ipsilateral SGI relieved neglect in most cases, confirming that disinhibition of the SC ipsilateral to the cortical lesion was necessary and sufficient to relieve neglect. Moreover, since loss of the contralateral SC was not sufficient to relieve neglect, it would seem that a strict balance of activity between the colliculi would be, on its own, not a component of the Sprague effect.

This issue is far from clear, however. Bilateral inactivation of parietal cortex or of the superior colliculus appears to relieve neglect symptoms, supporting a balance of activity view (Lomber and Payne, 1996; Payne et al., 2003). Payne et al. (2003)

showed that neglect from posterior suprasylvian sulcus (pMS)[2] lesion in cat could be relieved either by cooling the contralateral SC or the contralateral pMS, consistent with a balance of activity theory. This result confirms that the lesion of visual cortex as still performed in previous studies (Wallace et al., 1989, 1990) is actually unnecessary to elicit the Sprague effect and that it is, truly, an issue of neglect rather than of cortical blindness. Furthermore, in a 2-deoxyglucose (2DG) uptake study (Rushmore et al., 2006), lesion of pMS was associated with a higher than average uptake of 2DG in the contralateral SGI and not in the SGS nor in the contralateral pMS; 2DG uptake was diminished in the ipsilateral SGI, and cancellation of neglect was associated with reduction of contralateral 2DG uptake in the SGI to normal or sightly below normal levels. Therefore lesion of pMS and the consequent neglect led to an increase in activity as measured by metabolic rate in the contralateral SGI, thus implying disinhibition, and a decrease in metabolic rate in the ipsilateral SGI, implying inhibition. Furthermore, there was no change in the contralateral pMS. Likewise, when neglect was relieved by cooling the contralateral pMS, the activity levels in both sides of the superior colliculus were restored to normal or slightly lower, consistent with a loss of excitatory corticotectal input on both sides. These results would indicate that if the Sprague effect is a matter of balance of activity between competing orienting drives, then that balance is not in the cortex. Instead, it is in the superior colliculus. Of course, the appearance of "balance" in the activity could be merely a side-effect of symmetrical normal function rather than a causative agent in and of itself.

In addition, it remains unresolved why large lesions of the ipsilateral SN were unable to relieve neglect whereas small lesions of the SNpr were. One possible explanation could be that the large lesions of the SN also included portions of the substantia nigra pars compacta (SNpc), a major recipient of descending input from the SC which appears to play a role in alerting the dopaminergic reward system to the presence of biologically salient visual stimuli (May et al., 2009; Comoli et al., 2003). This outcome would be consistent with the previous observation that loss

[2]pMS is the feline homolog of primate parietal cortex.

of descending output from the SC caused profound inattention (Raczkowski et al., 1976).

Another factor which may not have been recognized in previous descriptions of the Sprague effect is that the intermediate and deep layers of the SC are connected to the caudate nucleus (CN, striatum) via a disynaptic connection through the centromedian-parafascicular (CM-Pf) complex, at least in rats (Kobayashi, 2003). The SGI supplies projections to the CM-Pf in rats, cats, and monkeys (Krout et al., 2001; Royce et al., 1991; Harting et al., 1980). In monkeys, CM-Pf lesion causes a loss of the validity effect in a cued attention task (Matsumoto et al., 2001). Many of the neurons appear to respond to alerting stimuli and unexpected onsets, but quickly habituate. Therefore the SGI would appear to relay a signal to caudate via CM-Pf that supplies information about behaviorally significant sensory events which can activate striatal neurons. Furthermore, neurons in the striatum have responses correlated with reward and bias for some directions over others (Lauwereyns et al., 2002; Takikawa et al., 2002). The spatially selective reward dependent activity in this study seemed to create a motivation bias for contralateral visual space, an outcome which, if unopposed, would appear very much like a selection deficit or bias towards the unaffected side. Therefore an additional component to the Sprague effect may very well be an SC to CM-Pf to CN projection. Inhibition of the SC by the SNpr would reduce the activity on the SC to CM-Pf limb of the circuit, thus in turn decreasing stimulation of caudate by CM-Pf. If caudate is necessary for creating motivation and biases toward the contralateral visual space, then increasing its activation level on one side and decreasing it on the other could produce an effect very much like neglect.

I summarize the current understanding of the Sprague effect along with additional possible contributing factors in Figure 1.1. The contributions are numbered and correspond to the following events and changes to the circuit. All terms (ipsilateral and contralateral) are with respect to the cortical lesion. Parietal cortex is lesioned (1). This leads to loss of excitatory input onto the ipsilateral caudate nucleus, which is consequently inhibited with respect to its normal activation level (2). The

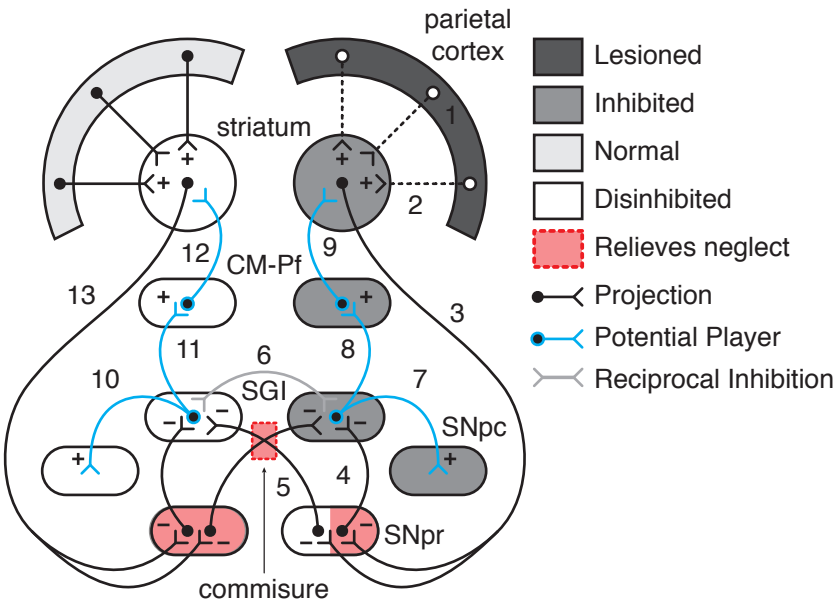

Figure 1.1: Circuit diagram of Sprague effect. Darkest gray crescent represents the lesion area of parietal cortex. Intermediate gray regions represent inhibited portions of the circuit. Light gray represents unaltered portions of the circuit. White regions are disinhibited. The pink areas indicate portions of the circuit which can be lesioned to relieve neglect. Black lines represent projections between nodes of the circuit. Gray line represents tectotectal projections which include both inhibitory and excitatory connections. Blue line represents additional circuitry which may contribute to both neglect and relief following manipulation of the circuit. Numbering as in text.

decrease in activation of the striatum causes a disinhibition of the ipsilateral SNpr (3). The SNpr increases inhibition on both the ipsilateral (4) and contralateral (5) SGI. Increasing inhibition on the ipsilateral SGI may lead to a decrease in inhibitory drive on the contralateral SGI, and while a precise prediction of the outcome is unclear since the tectotectal connections are both inhibitory and excitatory, it appears that the contralateral SGI is disinhibited (6). I suggest that additional contributing factors could include a loss of excitatory drive from the ipsilateral SGI into the SNpc, leading to a failure to detect relevant stimuli in the affected visual hemifield (7) and, presumably, a failure to orient attention to those stimuli. The inhibition of the SGI leads to a decrease in excitatory drive on the ipsilateral CM-Pf (8) and in turn on the striatum (9). Symmetrically, an increase in activity in the contralateral SGI increases

excitatory drive on the contralateral SNpc (10) and on the CM-Pf (11), leading to additional drive on the contralateral striatum (12).

The minimal portions of this circuit which can be eliminated to relieve neglect are also shown in Figure 1.1. These include the caudal portion of the collicular commissure, which carry crossing nigrotectal projections; the contralateral SNpr, which inhibits both ipsilateral and contralateral SGI, and portions of the ipsilateral SNpr. Why only small lesions of the ipsilateral SNpr relieve neglect whereas large lesions do not remains a mystery of the Sprague effect. In addition, not shown on this figure, direct disinhibition of the ipsilateral SGI can also relieve neglect. When the SGI ipsilateral to the lesion is disinhibited, the cascade of effects is to restore excitation to the SNpc along path (7), to the CM-Pf along path (8), and to striatum along path (9). This restores the animal's ability to detect relevant stimuli on both sides and to orient attention symmetrically, along with the equalization of drive on the striatum, despite the continued cortical lesion.

Finally, in a return to the initial issue of target selection in the superior colliculus, electrophysiological evidence has linked target selection in the SC with activity in the SNpr. In a series of papers (Hikosaka and Wurtz, 1985a,b, 1983a,b,c), Hikosaka and Wurtz showed that neurons in SNpr were linked in a retinotopic fashion with neurons in the ipsilateral SC. These neurons inhibit their counterparts in the SC, and just as bicuculline relieves inhibition and promotes irrepressible saccades in the SC, inhibition of SNpr neurons by injection of a GABA(A) agonist, muscimol, disinhibits the corresponding SC neurons. Later, Basso and Wurtz (2002) showed that SNpr neurons have firing rates inversely related to their SC neuron counterparts during a task in which the monkey must select one of several visual stimuli for a saccade. It would appear that a pause in their activity immediately prior to the saccade releases the SC from inhibition and allows a saccade to occur. The functional subunit of target selection would seem to include not simply the superior colliculus, but also the topographically linked SNpr and the rest of the basal ganglia circuit required to induce the disinhibition of the SC at the exact moment and location

required to select one target over many.

1.5 A proposal: superior colliculus as retinotopic interface into basal ganglia selection circuitry

All things considered, the evidence from lesion studies and from physiological studies of SC and SNpr neurons during eye movements and target selection would seem to indicate a fundamental link between the target selection deficits seen after SC inactivation and the parietal neglect relieved by disinhibition of the SC ipsilateral to the cortical lesion. It would seem that the neglect phenomena arise not so much as the result of the direct loss of function from the parietal cortex lesion, but instead as a result of the suppression of the ipsilateral circuit including the SC and the basal ganglia. Nevertheless, it remains unknown what exactly the impact on behavior the cortical lesion makes in the context of a compensating lesion that relieves the observable effects of neglect.

Based on the apparent importance of the SC-basal ganglia circuit in supporting proper selection between targets, I propose the following: first, that the superior colliculus serves as a retinotopic interface into the basal ganglia selection circuitry (Redgrave et al., 1999); second, that the selection circuitry is capable of flexibly mediating competition within both SGI even if one SNpr is eliminated; and third, that the selection circuitry does not operate merely on targets for orienting responses, but instead on all spatial information appearing in the SC, regardless of whether that represents the target of an orienting response or a working memory or a signal which will guide behavior via another output modality. In this view, neglect involves a disruption in the basal ganglia selection circuitry that causes stimuli in the affected hemifield to be "locked-out" of the selection process by inhibition of the retinotopic interface (the SC). Restoration of the retinotopic interface permits not only orienting behavior but also the access of any spatial information in the retinotopic SC map into the selection circuitry. The precise features of the neglect would be related to

how frontal and parietal cortical structures manage spatial information and which of these functions were impaired.

The idea that the SC is a retinotopic interface into the basal ganglia circuitry would be consistent with the observation that the SC projects into the centro-medial parafasciular complex (CM-Pf) (Harting et al., 1980; Kobayashi, 2003; Royce et al., 1991) and that these neurons in turn transmit information about behaviorally significant sensory events to the caudate (Matsumoto et al., 2001). In essence, the superior colliculus lies squarely inside a basal ganglia feedback loop. Furthermore, extensive inactivation of the SC has been shown to have a variety of impacts on behavior other than target selection, such as irrepressible approach and avoidance behaviors (Ciaramitaro et al., 1997). While these effects might at first seem quite strange, they make more sense if the SC is more closely integrated into basal ganglia selection circuitry than previously appreciated—lesions of the SC cause disfunction in the circuitry that leads to a jam in the selection for either inclusion or exclusion of appetitive and aversive behavior. Of course these effects have also been linked to SGI projections through paraventricular thalamic nuclei which project to the amygdala (Krout et al., 2001), but one might argue that the amygdala is in fact a shunned sibling of the basal ganglia that is not normally considered as such because of our own philosophical biases against considering emotion to be merely another feature of decision and selection processes.

While it is well known that the basal ganglia play an important function in the control of cognitive processes (e.g. Middleton and Strick, 2000), the link between the SC and working memory and other cognitive functions may at first seem unusual. It is well known that many SC neurons have working memory related activity during a memory-guided saccade task, but this activity is often conceived of as motor preparatory or related to target selection (none of which are exclusive possibilities). There may be a more direct link to both spatial and nonspatial working memory, however. SC neurons are essentially endowed with a variety of properties inherited from their cortical inputs (one might regard this as the purpose of cortical structures,

in fact), and the SGI receives a large bilateral input from the DLPFC (Johnston and Everling, 2006), which is associated both with working memory and with the context-specific control of selective attention. For example, the DLPFC projects to the striatum in a segregated pattern based on spatial and nonspatial working memory (Levy et al., 1997). In addition, the DLPFC connects with the thalamic mediodorsal nucleus (Rouiller et al., 1999; Giguere and Goldman-Rakic, 1988), a major output of the SGI which has been already explored in terms of corollary discharge to FEF (Sommer and Wurtz, 2004b,a). What may not be appreciated from this is that MD has a variety of functional properties related to memory (Tanibuchi and Goldman-Rakic, 2003). The precise details of the connections in the circuitry between SC, DLPFC, MD, and caudate, in addition to LIP and FEF, is unknown, but the fingerprints of such a connection are apparent.

When one considers that the list of usual suspects in unilateral neglect are all connected to the superior colliculus and the basal ganglia selection circuits through thalamic relays, that the SC receives input from all of these regions in spatial register, that the relief of neglect symptoms as in the Sprague effect do not appear to be related to overall balances of activity between cortical regions but instead to the alleviation of hyperactive suppression of selection in the SC-SNpr circuit, that despite the huge variety of reference frames and factors associated with the various manifestations of the neglect syndrome, at their root they all involve defects in selection processes, it should begin to become apparent that some fundamental connection exists between the basal ganglia selection circuitry and selection of signals for regulating perception. I have proposed the somewhat extreme possibility that the core features of unilateral neglect emerge from a "lock out" of certain regions of space from the basal ganglia selection circuitry, that other regions are, as a result, "locked in," and that these effects are mediated by the SC. Therefore the SC acts as a retinotopic interface into selection circuitry such that a variety of types of spatial information, ranging from memory to movement, can be placed in competition and selected. If this is the case, then the consequences would be that a fundamental resource limitation in attention

and behavior is the bandwidth of the SC and its ability to simultaneously represent alternatives, and that cognition and action compete for access into selection schemes with either partially or totally overlapping circuitry.

1.6 Main hypotheses tested

In this introduction I have elaborated three models of the relationship between the superior colliculus and selective attention. The *first model* is based primarily on a legacy of work focusing on the role of the SC in the control of visually-guided orienting movements. It holds that the transition from visually-guided orienting behavior to more sophisticated orienting behavior relies essentially on the co-option of the circuitry in the SC by a cortical circuit that uses information in the periphery to synthesize targets for orienting movements. The central question in this model is whether or not the selection of which stimuli in the visual field contribute to the generation of the synthetic target relies on circuitry resident in the superior colliculus or if instead it is a purely cortical phenomenon. The *second model* holds that the SC provides an important input into parietal cortex for establishing reference frames, and that removal of that input causes disruption in the parietal cortex representation of space, thus leading to both intentional and sensory neglect. The third model is grounded in consideration of the Sprague effect and evidence that the selection mechanisms of the superior colliculus depend largely on its placement within a basal ganglia circuit. This *third model* has the specific hypothesis that the selection between neural representations including those of both orienting movements and attention is mediated through the same selection mechanism localized in the basal ganglia. The latter two models both would predict that SC inactivation should not be restricted to the selection of which stimuli guide orienting behavior, and instead predict that generalized sensory and intentional neglect will result from SC inactivation.

The *first hypothesis* tested in the following experiments is that the SC is required for the selection of which signals in the visual field are used to guide eye

movements. This will resolve whether the SC simply contributes to the shift of attention to a location in space in anticipation of an orienting movement or if instead it mediates the spatial selection of which visual signals are used to generate targets for orienting movements. In Chapter 2, we find that after inactivation of the SC, monkeys are unable to use information in the corresponding portion of the visual field to generate targets for saccades; instead, they rely entirely on erroneous information presented elsewhere in the visual field. Furthermore, this pattern of selection persisted even when primary sensory deficits could be eliminated as explanatory factors. This result demonstrates that the SC is necessary for selective attention beyond the selection of targets for orienting movements, and is consistent with all three proposed models.

The *second hypothesis* tested is that the SC is required for the selection of which signals in the visual field are used to guide movement in general and, as would seemingly be necessarily implied, perception. This will resolve whether or not the function of the SC in selective attention is specific to the generation of orienting behavior. In Chapter 2, we find that after inactivation of the SC, monkeys are unable to report the properties of visual stimuli with button press, and instead base decisions on erroneous information. The results for manual response were qualitatively identical to those for saccadic response. Therefore the SC has a central role in the selection of which signals inform action in general in addition to perception. This result would seem to rule out the first model and support the latter two.

Since the SC is involved in the selection of both targets for saccades and which signals inform perception, the question of how the different options are represented in the SC and thereby subjected to selection mechanisms becomes paramount. Therefore the *third hypothesis* tested is that the SC has neural correlates of more abstract representational forms. In Chapter 3, we find that oscillatory activity in some cells is associated with the spatial components of working memory. This result implies that the SC contains circuitry which can represent more abstract forms such as working memory, and likely could subject those representations to selection

mechanisms.

Finally, previous conceptions of the role of the SC in selective attention identified maintained firing rate in neurons such as buildup cells as the source of a drive that directed attention to the corresponding location in visual space. Therefore the *fourth hypothesis* tested is that neurons in the SC are predictive of which signals the monkey was instructed to select and that variability of the neurons determines the variability in performance on the task. In Chapter 6, we find that many classes of neurons show substantial changes in mean firing rate associated with the cue position, but that these neurons are not informative of the cue position because of excessive variability in the maintained firing rate. Instead, only a single group of neurons were informative about the cue position, and these were active in a brief period of time when the monkey had to commit to memory which signal to select. This outcome would be more in accord with a model in which the SC is involved at discrete points in time when selections must occur, but is not involved otherwise.

Overall, these results show that the SC is required for selecting which visual signals inform perception. We specifically reject models in which the SC is involved only in the selection of which signals are used to generate targets for orienting movements. In addition, the results are not consistent with the notion that the SC maintains a drive that orients attention to the location of upcoming eye movements. Instead, the results are more consistent with models in which the SC is involved with spatial selection between more abstract representational forms at distinct points in time, which seems most consistent with the third, basal ganglia model.

Chapter 2

Inactivation of Primate Superior Colliculus Impairs Covert Selection of Signals for Perceptual Judgements

2.1 Abstract

Primates base perceptual judgments on some sensory inputs while ignoring others. The covert selection of sensory information for perception is often thought to be accomplished mostly by the cerebral cortex, whereas the overt orienting toward relevant stimuli involves a variety of additional structures such as the superior colliculus, a subcortical region involved in the control of eye movements. Contrary to this view, we show that the superior colliculus is necessary for determining which stimuli will inform perceptual judgments, even in the absence of orienting movements. Reversible inactivation of the superior colliculus in monkeys performing a motion discrimination task caused profound inattention for stimuli in the affected visual field, but only when distracters containing counter-informative signals appeared in the unaffected field. When distracting stimuli contained no information, discrimination performance was

largely unaffected. These results indicate that the superior colliculus is a bottleneck in the covert selection of signals for perceptual judgments.

2.2 Introduction

The primate superior colliculus (SC) has long been implicated in the mechanisms of visual attention. For overt attention, the role of the SC is well established–we often look directly at attended objects and the SC is a major component of the motor circuits that control how we orient our eyes and head (Wurtz and Albano, 1980; Sparks, 1999). In addition to its well-known role in guiding the motor output, the SC is also important for the preceding step of selecting which stimulus will be the target of an eye movement. When the region of the SC representing the target is reversibly inactivated, saccades are often misdirected to distracters appearing in unaffected parts of the visual field (Mcpeek and Keller, 2004). Conversely, electrical stimulation of the SC at currents too weak to directly evoke eye movements can nonetheless bias target selection toward the stimulus in the activated location, regardless of whether the target is acquired with a pursuit or saccadic eye movement (Carello and Krauzlis, 2004). We consider these effects on target selection to be a form of intentional neglect (Heilman, 2004) that may be related to the classic Sprague effect, in which deficits in orienting caused by lesion of parietal cortex on one side of the brain can be relieved by lesion of the SC on the other side (Sprague, 1966). Thus, the SC contributes to overt attention both by controlling the motor output and by participating in the selection process that determines where we look next.

However, we also attend covertly–without directing our gaze toward the attended object–and it is now clear that at least some components of the oculomotor system play a role in covert attention. In humans, imaging studies show that covert attention and eye movements activate a common set of areas in the parietal and frontal cerebral cortex (Corbetta and Shulman, 2002). The symptoms of spatial neglect in human patients are thought to arise from an imbalance of activity within this fronto-

parietal regulatory network that drives attention to the unaffected side (Driver and Mattingley, 1998; Mort et al., 2003; Corbetta et al., 2005). Likewise, in non-human primates, cortical areas important for eye movements have also been implicated in the control of covert attention. Reversible inactivation of the frontal eye fields (FEF) causes temporary deficits in performance on covert attention tasks (Wardak et al., 2006), and electrical stimulation of the FEF improves performance on detection tasks and promotes the enhancement of visual processing in area V4 (Moore et al., 2003; Moore and Armstrong, 2003; Moore and Fallah, 2001). The lateral intraparietal area (LIP) is implicated in the control of both saccadic eye movements and spatial attention (Snyder et al., 1997; Bisley and Goldberg, 2003), and is one of the few cortical areas known to contain neurons that represent the spatial decision variables (Gold and Shadlen, 2007) important for guiding both overt and covert orienting.

Nevertheless, it is not yet known whether the control of covert attention is restricted to the cerebral cortex, or extends to subcortical structures such as the SC. The results to date are ambiguous. Some neurons in the SC increase their activity level when a monkey attends into their response fields, even when the attended stimulus is not the target of a saccade (Kustov and Robinson, 1996; Ignashchenkova et al., 2004); this activity could be related to the control of covert attention, or it could be related to oculomotor planning by the SC that occurs during covert attention. Microstimulation of the SC drives attention to a location in space almost as if the monkey had been cued to attend to that location (Cavanaugh and Wurtz, 2004; Müller et al., 2005); these effects show that the SC is part of the circuit for covert attention–perhaps via its connections to frontal and parietal cortex–but they do not distinguish whether the SC is crucial for the control of covert attention or simply updated about its current state.

We addressed this ambiguity by performing reversible inactivation in the SC of monkeys trained to perform a selective attention task. The task required subjects to ignore distracting stimuli while covertly attending to a cued stimulus that instructed them where to orient, thus distinguishing between control of gaze and

control of attention. Inactivation of the SC caused extinction-like deficits–subjects ignored cued signals in the inactivated region when they competed with distracting foils placed elsewhere, but discrimination ability was largely intact when the cued signal appeared alone. These effects reflected a generalized impairment in covert attention, because they were also observed using a manual-response version of the task without eye movements. Together, these results demonstrate a causal role for the SC in the control of covert attention.

2.3 Results

2.3.1 Monkeys normally attended to cued stimuli and ignored foils

Monkeys performed a motion discrimination task that required them to judge the direction of motion in one of four peripheral stochastic motion stimuli while ignoring distracting signals (Figure 2.1). The odd colored ring cued the monkey to attend to one of the four stimuli. After a delay, brief pulses of coherent motion appeared simultaneously at both the previously cued location and the diametrically opposite foil location. The direction of motion in the cued location was drawn at random from any of the four diagonal directions, and the direction in the foil was drawn from any of the remaining three. Monkeys were required to maintain fixation throughout the presentation of the motion stimuli, and in separate versions of the task, reported the direction of the motion pulse either by making a saccade or by pushing a button corresponding to the direction of motion. In the button-press version they were required to maintain fixation for the entire duration of the trial, including the response interval. After extensive training, both monkeys based their responses on the cued signal on a substantial majority of the trials ($\sim$75%), indicating that they were able to selectively attend to the cued stimulus (Figure 2.2). In order to achieve equivalent performance between the animals, we set the motion coherence at 0.1875 for subject F and at 0.25 for subject M. A preponderance of errors was consistent

with the foil signal, indicating that mistakes in the task were usually due to selecting the wrong stimulus, rather than simply guessing. Figure 2.2 shows behavioral data collected in the months prior to the inactivation experiments. In the saccade task, this includes 50,651 trials over 78 sessions with subject F and 33,170 trials over 82 sessions for subject M. In the button-press task, this includes 8,756 trials in 28 sessions for subject F and 20,265 trials in 44 sessions for subject M.

We assessed the degree to which the SC regulates selective attention by inactivating portions of the intermediate and deep layers of the SC corresponding to parts of the visual field in which motion stimuli were presented. Immediately prior to each inactivation, we collected control behavior. We then injected 0.5 μL of muscimol (0.5μg μL^{-1}), a GABA agonist, into the intermediate and deep layers of the SC to temporarily inactivate neurons in those regions (Figure 2.9 and Methods), although it is possible that neurons in the superficial layer might also have been affected. Muscimol spread laterally through these layers, and we assessed the spatial extent of the resulting inactivation effects by observing the decrease in peak velocity of visually guided saccades (Quaia et al., 1998) (Figure 2.3A). In each session, the effects were restricted to a portion of the visual field that overlapped with either the cued or foil stimulus (Figure 2.3B). Monkeys were cued to attend either into the affected quadrant or the diametrically opposite quadrant in alternating blocks of 40 trials. Since the motion direction was independent of the cued location, the locus of attention was independent of the orienting response. Any impact of SC inactivation on attention was therefore distinct from effects on saccades and could be distinguished by examining the subset of trials in which neither the cued signal nor the foil signal pointed into the affected quadrant of the visual field. In these trials, no response should be made into the affected quadrant and thus few responses should be affected by the inactivation. Consequently we included only this subset of trials in further analysis. If the SC is necessary for selective attention, inactivation should decrease the ability of the monkeys to base judgments on cued signals in the affected visual field; the degree to which the foil intrudes upon the judgments will depend on the degree

to which selective attention is biased into the unaffected visual field. SC inactivation impaired selective attention

After inactivation of the SC, both monkeys ignored the cued signal when it was presented in the affected quadrant. In a representative pair of cued and foil motion directions from the inactivation in Figure 2.3A, subject F was initially able to base the majority of his judgments on the cued signal and ignore the foil signal (Figure 2.4A). After inactivation, he was unable to base judgments on the cued signal when it appeared in the affected quadrant of visual space (Figure 2.4B). This outcome cannot be explained in terms of a deficit in target selection or saccade execution (Mcpeek and Keller, 2004; Carello and Krauzlis, 2004), because the cued signal required a rightward response into an unaffected quadrant. Instead, SC inactivation caused the monkey to base its decisions on the foil signal even though this signal instructed a leftward movement toward the affected side of the visual field. Inactivation also did not cause the monkeys to become unmotivated or to begin randomly guessingthe majority of decisions were still based on the presented stimuli, but now incorrectly followed the foil. Although SC inactivation did cause deficits in saccade production, such as a reduction in the peak velocity of saccades made into the affected quadrant (Figure 2.3), the changes in behavior indicate that the monkey could not selectively attend into the affected quadrant.

SC inactivation caused monkeys to ignore the signal in the affected quadrant of visual space in all experimental sessions. We summarize these effects for each session by plotting the overall proportion of choices driven by each signal after inactivation against the proportions prior to inactivation (Figure 2.5). A decrease in the proportion of choices driven by a particular signal causes the corresponding points to lie below the unity line, whereas an increase causes them to lie above the unity line. When the cue appeared in the affected quadrant (Figure 2.5A and C for each subject), the proportion of choices with the cued signal decreased from $\sim$0.7 pre-injection to only $\sim$0.2 post-injection (red symbols). Conversely, the proportion of choices with the foil signal increased from $\sim$0.2 pre-injection to $\sim$0.7 post-injection (yellow sym-

bols). Notably, the proportion of choices with neither signal did not increase (gray symbols). This indicates that the monkeys did not engage in more random guessing after the inactivation, but instead still based their choices on signals present in the display.

A complementary effect often occurred when the foil signal was presented in the affected quadrant. The proportion of choices with the cued signal increased slightly post-injection in many sessions, and the proportion of choices with the foil decreased, (Figure 2.5B and D for each subject). The amplitude of these improvements was limited by the high level of the monkeys pre-injection performance, and was more evident with subject M. Again, the proportion of choices with neither signal did not change. Thus, SC inactivation caused the monkeys to preferentially base decisions on the signal in the unaffected quadrant, regardless of the cue, as if selective attention were biased toward the signal in the unaffected region.

2.3.2 Deficits also occurred for manual-response

To test whether SC inactivation caused a general impairment of selective attention rather than an effect restricted to the guidance of eye movements, we trained both subjects to perform the task by pushing buttons rather than by making saccades. Just as in the saccade version of the task, SC inactivation again caused the subjects to ignore the cued signal when it was presented in the affected quadrant, and instead to base their responses on the foil (Figure 2.6A and C). When the foil signal was presented in the affected quadrant, performance was sometimes improved (Figure 2.6B and D), although this effect was observed less often than in the saccade version of the task. Additionally, we performed a saline control experiment in subject M, which led to no changes in performance regardless of where the cued and foil signals were placed (open symbols in Figure 2.6C and D). Overall, these results rule out explanations based on the role of the SC in the control of saccades, because the changes in performance occurred while the subjects maintained fixation for the entire duration of each trial. Instead, they indicate that the changes in task performance

after SC inactivation reflect a generalized impairment in covert attention.

2.3.3 SC inactivation slightly impaired motion discrimination

Although inactivation of the SC appeared to bias selective attention into the unaffected field, a local impairment of motion processing could have contributed to the change in behavior. To address this possibility, we tested motion discrimination on single stimuli concurrently with the inactivation sessions. Subjects reported the direction of motion in a single peripheral stimulus with a saccade (Figure 2.7A). We pooled trials from all inactivation sessions and appropriately associated directions to maintain the spatial relationship of each direction to the inactivated region (contra, opp, and ipsi as in Figure 2.7A). This included for subject F 4,638 trials in 23 control sessions and 1,283 trials in 9 inactivation sessions; for subject M, 3,873 trials in 20 control sessions and 1,293 trials in 8 inactivation sessions. Subject Fs motion discrimination performance for each direction of motion as a function of coherence during both control and inactivation conditions appears in Figure 2.7B-E. Curves were fitted to the data using multinomial logistic regression (McCullagh and Nelder, 1998). The corresponding data and curve fits for subject M appear in Figure 2.7F-I.

SC inactivation impaired performance on the task, as evident by the decrease in proportion of correct responses. We quantified this impairment in terms of an increase in bias and a decrease in sensitivity to motion coherence. In probit analysis for two-alternative discrimination tasks, bias describes the relative preference for one option over another independent of signal strength, and sensitivity describes how the relative gain on signal strength contributes to the preference for one option over another. Similarly, in four-alternative discrimination, three relative bias terms and three relative sensitivity terms describe preference for each option over a single reference, which in this case was defined as the direction corresponding to the inactivated region. The insets in Figure 2.7 show relative bias and sensitivity to coherent motion for each direction, pre- and post-injection.

We found that subjects were biased to respond away from the affected quadrant of visual space, as expected from previous studies showing that SC inactivation affects saccade selection3. Subjects were normally biased slightly away from the region of the motion stimulus (open bars in Figure 2.7 insets), and this tendency was exaggerated by inactivation of the SC (gray bars). In addition, subjects showed a decrease in sensitivity to coherent motion after SC inactivation (red bars compared to black bars). The changes in sensitivity amount to a three-fold decrease in the gain on the signal for subject F and a two-fold decrease in the gain for subject M. Hence, SC inactivation caused an impairment in local motion discrimination, in addition to biasing saccade responses.

2.3.4 Sensory impairment did not cause deficits in attention

Although SC inactivation decreased sensitivity to local motion, discrimination performance could be recovered to control levels by simply increasing signal strength (Figure 2.7). For example, discrimination performance using maximal signal strength during SC inactivation ($\sim$0.95 correct with 0.5 coherence, see Figure 2.7) was equivalent or better for unaffected directions than that using near-threshold signal strengths without inactivation ($\sim$0.7 correct with $\sim$0.2 coherence, see Figure 2.2). The impact of SC inactivation on local motion discrimination with single stimuli could therefore be equated with a reduction in effective signal strength. This equivalence provided a method for testing if local changes in motion sensitivity contributed to the effects of SC inactivation on selective attention. Specifically, if the behavioral effects in the attention task were due to a local impairment in motion discrimination, then they should be entirely reversed by appropriately boosting the strength of the motion signal. In a set of additional inactivation experiments, we therefore increased the coherence of stimuli in the affected quadrant to the maximum presentable while leaving the foil strength the same (0.5 cued / 0.1875 foil for subject F, 0.5 / 0.25 for subject M). Despite the increase in cued signal strength, we observed a qualitatively identical pattern of results as in the initial experimentthe monkeys ignored the cued

signal and instead based responses on the foil (Figure 2.8). Therefore increasing the strength of the cued stimulus could not reverse the behavioral changes observed after SC inactivation. It follows that a primary sensory impairment is insufficient to explain the observed effects, and instead that the deficits observed after SC inactivation were caused by a disruption of selective attention.

2.3.5 Only misleading sensory information impaired performance

Finally, we tested whether the impairments in selective attention required the presence of a foil signal. Spatial neglect sometimes appears to involve an inability to disengage attention, implying that the deficit may involve not simply an impairment in directing attention into the affected field, but instead an inappropriately hyperactive drive to attend to stimuli in the so-called unaffected field (Corbetta et al., 2005; Halligan and Marshall, 1994; Joanette and Brouchon, 1984). We therefore considered whether the impairment in covert attention caused by SC inactivation could be due to an inability to disengage attention from stimuli in the unaffected field. In this scenario, the foil drove choices after SC inactivation because attention was unavoidably engaged at that stimulus, thus permitting whatever signal was present there to drive the choice. If this were the case, then the mere presence of incoherent motion patches in the unaffected field should cause the monkeys to ignore the cued stimulus in the affected quadrant. Contrary to this prediction, however, when no foil signal competed with the cued signal, we observed no significant decrease in performance (Figure 2.8). This indicates that the presence of stimuli in the unaffected field was insufficient to impair performance. Thus the deficits in covert attention caused by SC inactivation were due to an inability to filter distracting or misleading sensory content, not simply the presence of distracting stimuli.

2.4 Discussion

The primate SC has long been implicated in the control of attention and eye movements. Some of the first recordings in the primate SC showed that visual responses were enhanced with the shift of attention presumed to precede a saccadic eye movement (Goldberg and Wurtz, 1972b), and older models of attention have outlined a role for the SC in the orienting of attention (Posner and Petersen, 1990). More recently, neural activity in the superior colliculus has been correlated with both voluntary (Kustov and Robinson, 1996) and stimulus-driven allocation of attention (Fecteau et al., 2004), although these observations could be due to an obligatory preparation of saccades that occurs concurrently with covert orienting. In addition, electrical stimulation of the superior colliculus caused mild enhancements in performance on a change blindness task (Cavanaugh and Wurtz, 2004) and motion discrimination task (Müller et al., 2005), which shows that the SC is at least part of the circuit for covert attention. By showing that reversible inactivation of the SC causes a profound impairment in correctly selecting which visual stimulus will inform perceptual judgments, even in the absence of eye movements, our results demonstrate that the primate SC is not only part of this circuit, but that its activity is crucial for the normal control of selective attention. Hence, the SC may form a bottleneck in the control of both overt orienting and covert attention.

The primary deficit we observed after SC inactivation was an inability to filter out distracting or misleading sensory information. These effects on selective attention have a somewhat different character than those typically observed in areas of cerebral cortex, which are normally associated with increases in the detectability or discriminability of stimuli at the attended location (Reynolds and Chelazzi, 2004; Cook and Maunsell, 2002; Yeshurun and Carrasco, 1998; Luck et al., 1997; Williford and Maunsell, 2006). Instead, our results are reminiscent of visual extinction, a phenomenon observed in less florid cases of spatial neglect in which stimuli in the affected side of the visual field are perceived as long as there are no competing stimuli in the unaffected side (Rafal, 1994). In the present experiment, misleading informa-

tion, not simply distracters, was necessary to induce the deficits in performance. This observation suggests that a competition occurs, not simply between visual stimuli, but between the potential sources of information to guide the subjects perceptual judgment. This type of mechanism is compatible with the biased competition model of spatial attention, which holds that a biasing signal weights stimuli in a spatially specific manner prior to a divisive normalization stage (Desimone and Duncan, 1995; Reynolds et al., 1999; Reynolds and Heeger, 2009). The extinction-like deficits we observed suggest that SC inactivation disrupts the normal weighting of signals in the affected quadrant, but leaves the process of divisive normalization intact. The SC is not necessarily the source of the biasing signal, but given its anatomical connections (Wurtz and Albano, 1980), it could serve as a site of convergence and integration of many potential biasing signals from various cortical and subcortical sources, and then broadcast the results to appropriate targets in thalamus, cortex and elsewhere (Posner and Petersen, 1990; De Weerd et al., 2003; Rossi et al., 2007).

We also found that SC inactivation caused a mild impairment in the ability to discriminate motion signals. This more subtle effect could be due to changes in sensory processing. The discrimination of motion signals depends on neurons in the middle temporal area, which show modulation with spatial attention (Treue and Maunsell, 1996; Cook and Maunsell, 2002) and receive indirect inputs from neurons in the more superficial layers of the SC (Berman and Wurtz, 2008). However, lesion of the SC has been shown not to lead to changes in the response properties of neurons in area MT (Rodman et al., 1990), suggesting that other mechanisms might be involved. In particular, changes in discrimination performance could be due to changes in how sensory signals are pooled when making decisions, with no changes in sensory processing. For instance, in monkeys learning a motion discrimination task, increases in discrimination performance were not associated with changes in the response properties of neurons in MT but instead with the association of motion signals with response directions in area LIP (Law and Gold, 2008). In our experiments, disruption of the normal spatial weighting of signals by SC inactivation might have

allowed perceptual judgments to be influenced by irrelevant activity originating from outside the motion patch, possibly illustrating the hazards of internal noise that is not properly suppressed by selective attention (Dosher and Lu, 2000).

This study demonstrates that the primate SC is not simply updated about covert selective attention, but is necessary for its normal operation. The pattern of extinction-like deficits provides an outline of the possible mechanisms by which the SC makes its contribution, and shows that the fronto-parietal network in cerebral cortex is insufficient on its own to allocate selective attention. Given its evolutionary history and its interconnectedness with other brain regions, identifying how the SC exerts its control over perceptual judgments is likely to be central to understanding what selective attention is and how it works.

2.5 Methods

2.5.1 Behavioral tasks

The selective attention task required monkeys to maintain fixation while attending to a stochastic motion stimulus (Figure 2.1). Each trial began with the appearance of a fixation dot. Four colored circles appeared for 480 ms. The odd colored ring indicated the cued location. Stochastic motion stimuli appeared within the circles, and after an additional 480 ms, the cue circles disappeared. Stimuli remained on the screen for 480 ms plus a geometrically distributed delay of mean 480 ms. The hazard function remained flat for nearly the entire duration of the trial so that the monkeys were not provided with information that could allow them to predict the onset of the motion pulses. Coherent motion pulses (160 ms) occurred in both the cued location and the diametrically opposite location. Motion was transitioned from incoherent to coherent by assigning newly appearing dots into the pool of coherently moving dots. The monkeys task was to report the direction of motion in the cued location (by saccade or button press, in separate experiments). Monkeys received a liquid reward only for correct responses in completed trials. If the monkey broke

fixation mid-trial, such as making a saccade toward one of the motion stimuli, the trial was aborted and repeated later in the session. Motion coherence of the pulses was titrated for each subject to maintain 65-70% correct performance (0.1875 for subject F and 0.25 for subject M).

During both control behavior and inactivation, monkeys fixated a small, stationary spot presented at eye-level directly in front of them at the center of a CRT display with a refresh rate of 75 Hz. The fixation spot consisted of a single pixel of background luminance surrounded by a 1-pixel-thick white square. With our display geometry and distance, this 3x3 pixel stimulus corresponded to ∼9x9 min arc of visual angle. The background luminance of the monitor was 14 cd m^{-2}. Luminance of the fixation dot was 50 cd m^{-2}. The cue rings appeared at an eccentricity of 8.2 degrees of visual angle (fixation point to ring center), had a diameter of 8.8 degrees of visual angle, and had a thickness of 0.25 degrees of visual angle. The ring luminance was 25 cd m^{-2} green or red. The stochastic motion stimuli appeared centered within the cue rings and had a radius of 4.25 degrees of visual angle. The stimuli consisted of limited lifespan dots (Cook and Maunsell, 2002). The dot lifespan was 2 refreshes. At each refresh, each dot either appeared at a random location within the patch or was displaced by 4 pixels (∼0.2 degrees of visual angle). The coherence was the proportion of total dots moving in the same direction; the remainder moved in uniformly distributed random directions. Each dot had a peak luminance of 50 cd m^{-2}. Responses were conveyed with saccades made to choice dots appearing at an eccentricity of 8.2 degrees of visual angle and with a radius of 0.2 degrees of visual angle and peak luminance of 50 cd m^{-2}.

On the same day as each inactivation session, we collected pre-injection control behavior on both the single stimulus motion discrimination task (192 trials) and the selective attention task (352-528 trials). During this phase of the session, we advanced the tip of the injection cannula into the quadrigeminal cistern. After collecting pre-injection control data, we advanced the injection cannula into the intermediate and deep layers of the SC and verified our position as described below

under Muscimol injections. After completing the injection, we mapped the extent of the inactivation effects with visually-guided saccades (80 trials). We then assessed motion discrimination performance with the single stimulus discrimination task (192 trials) before proceeding with the selective attention task (352-528 trials). For the single stimulus discrimination task, we randomly interleaved stimuli with a range of motion coherences (0.625, 0.125, 0.1875, 0.25, 0.375, 0.50); for the selective attention task, we used a single coherence just above threshold for each of the monkeys (0.1875 for subject F, 0.25 for subject M).

In the button-press version of the task, subjects pushed buttons mounted at waist level on the left side of the chair within easy reach of the left hand. Four buttons were arranged in a square and each button corresponded to a direction of motion. Each subject used only his left hand to push buttons. As in the initial version of the attention task, monkeys were required to maintain fixation during the trial; in the button-press version, they were required to maintain fixation during the response period and push the button corresponding to the judgment of motion direction. They could only push buttons during this interval and were rewarded for the first button pushed.

2.5.2 Muscimol injections

We injected muscimol (0.5 μL, 5 μg μL^{-1}), a GABA agonist, into the intermediate and deep layers of the SC using an injection cannula with an electrode threaded down its barrel Chen et al. (2001). Three methods allowed us to localize the cannula tip within the SC prior to injection (Figure 2.9). First, we advanced the cannula to a depth (1.5-3 mm below the SC surface) corresponding to the intermediate and deep layers based on a history of microelectrode recordings and histological studies (Wurtz and Albano, 1980). Second, we recorded activity during saccades consistent with known responses in the SC, thus confirming the depth in the SC. The location of the units movement fields also indicated our placement within the SCs retinotopic map. Third, we evoked saccades with microstimulation. The current required to evoke sac-

cades (typically 10 μA in the intermediate and deep layers) provided an additional indication of depth, and the direction and amplitude of the evoked saccades indicated the position within the map. Nonetheless, we cannot be certain that our effects are due solely to inactivation of neurons in the intermediate and deep layers of the SC, because some drug may have diffused vertically through the layers or tracked up the shaft of the injection cannula to affect neurons in the overlying superficial layers.

We performed a total of 27 SC inactivation experiments in the two monkeys. The saccade response version included 7 for subject F and 4 for subject M; this version of the task also included the single-patch control experiment. For the button-press version of the task, we inactivated at 4 sites for subject F and 4 sites for subject M; we also conducted one saline control experiment with the button-press task in subject M. Finally, for the experiments varying the coherence in both the foil and cue patches, we conduced 4 sessions in each of the two subjects, and pooled the data across those sessions separately for each subject.

2.5.3 Animal preparation

We performed reversible inactivation of the intermediate and deep layers of the superior colliculus in two adult rhesus monkeys (subjects F and M) that were 10-15 years of age and weighed 12-15 kg. The monkeys were prepared using standard surgical techniques described in detail previously (Hafed et al., 2008). All experimental protocols were approved by the Institutional Animal Care and Use Committee and complied with United States Public Health Service policy on the humane care and use of laboratory animals. The laboratory setup for behavioral control and monitoring was identical to that described previously (Hafed et al., 2008).

Chapter 2, in full, is a reprint of the material as it appears in *Nature Neuroscience*. Lovejoy, LP; Kruazlis, RJ, 2010. The dissertation author was the primary investigator and author of this paper.

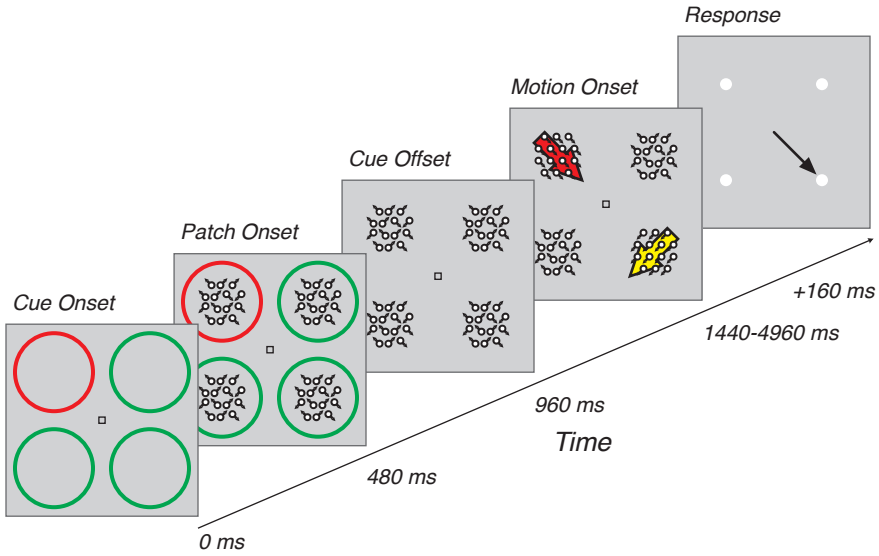

Figure 2.1: Selective attention task design. After a brief fixation period, colored cue rings were presented. Stochastic motion patches appeared next, and then the cues disappeared. Following a delay, brief coherent motion pulses occurred in both the cued location (red arrow) and the diametrically opposite location (yellow arrow). When responding by saccade, monkeys reported the direction of the cued motion signal by making an eye movement to a response dot in the same direction; when responding by button push, monkeys pressed a button corresponding to the motion direction.

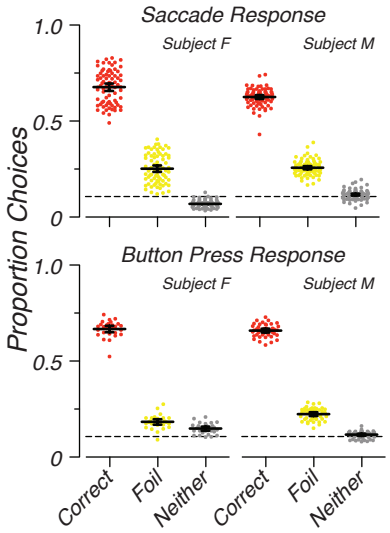

Figure 2.2: Performance on selective attention task. Summary of behavioral performance for both subjects in the saccade response version (top) and the button press response version of the task (bottom). Red dots represent proportion of correct choices (based on cued signal) in each session. Errors could be either driven by the foil signal (yellow dots) or by neither signal (gray dots). Scatter indicates variability across control sessions collected over several months prior to the inactivation experiments. Black lines and error bars indicate population average and 95% multinomial confidence intervals. Dashed line indicates the proportion of responses consistent with the foil that would be expected by chance.

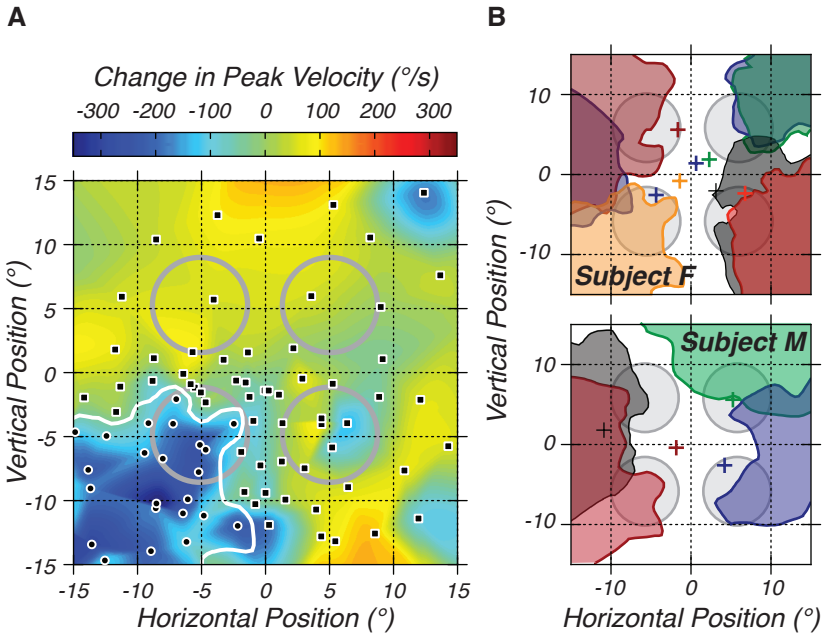

Figure 2.3: Map of inactivation effects. A) Single session data from subject F. Black dots indicate saccade end-points and interpolated color map indicates the changes in peak-velocity after muscimol injection. Cooler colors in the lower left quadrant indicate the decrease in peak-velocity caused by SC inactivation, and the white contour delineates the affected region. Gray circles indicate the positions of the four stochastic motion stimuli, which were at fixed locations throughout the set of experiments. B) Summary of SC injections for the saccade response task. Crosses indicate the average end-points of saccades evoked by microstimulation at the injection sites. Shaded regions indicate the extent of the visual field affected by each injection experiment.

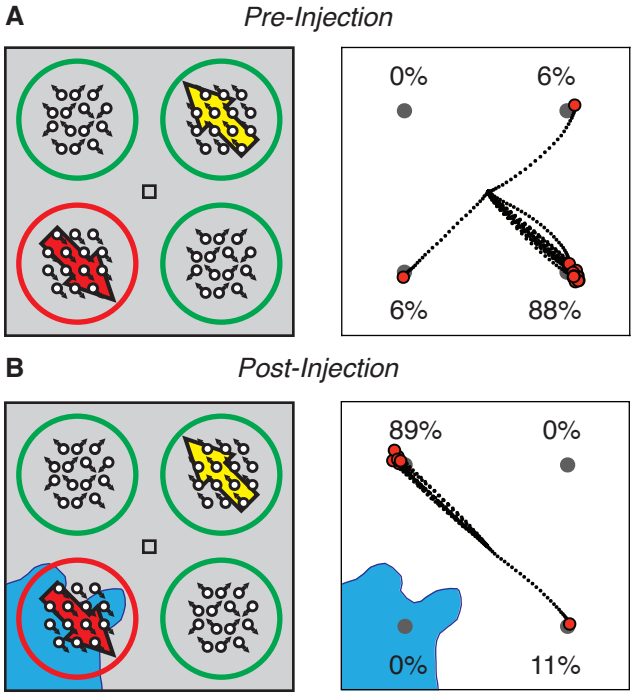

Figure 2.4: Sample data from one inactivation session. A) Behavior of subject F in a representative stimulus condition prior to inactivation. Schematic of stimulus indicates cued signal position with red ring and directions of motion in the cued signal (red arrow) and foil signal (yellow arrow). Red dots indicate end-points of saccade trajectories. A majority of responses were correctly directed by the cued signal. B) Behavior in the same condition after inactivation of the SC. Schematic of stimulus now indicates the affected portion of visual space as a blue shaded region. Only a minority of responses were correctly guided by the cued signal; instead, the majority of decisions were based on the foil.

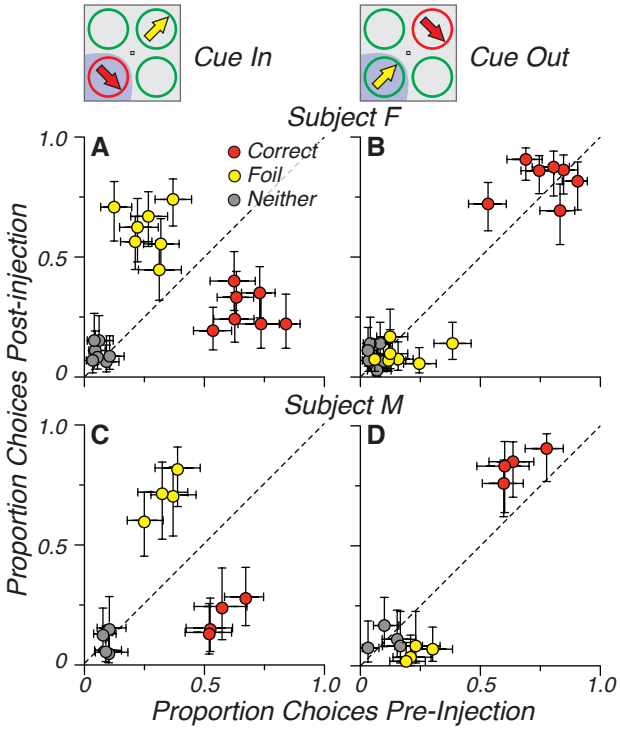

Figure 2.5: Summary results from inactivation sessions in saccade-response version of the task. The proportion of choices post-injection is plotted against the proportion pre-injection. Red circles indicate correct choices matching the cued signal. Yellow and gray symbols represent errors driven by either the foil signal or neither signal. Error bars indicate 95% multinomial confidence intervals for each of the sessions, which included 176-264 trials per session. Data for subject F are shown in panels A and B; data for subject M appear in panels C and D. When the cued signal was in the affected region (A, C), subjects ignored this signal and instead based their choices on the foil. Conversely, when the foil signal appeared in the affected region (B, D), subjects tended to base their choices on the cued signal and ignore the foil.

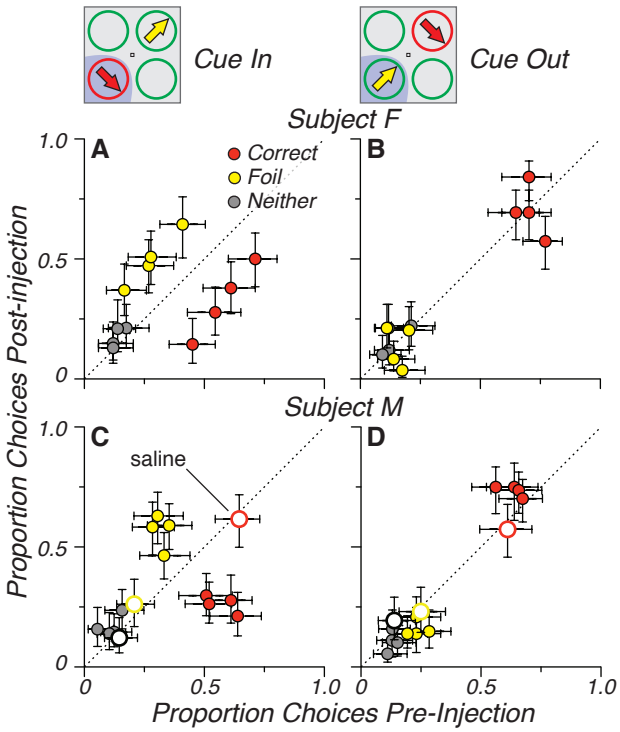

Figure 2.6: Summary results from inactivation sessions in button-press version of the task. Same conventions as in Figure 4. Data for subject F appear in panels A and B; data for subject M in panels C and D. When the cued signal was in the affected region (A, C), subjects ignored this signal and instead based their choices on the foil, similar to the results from the saccade-response version. The injection of saline during a control experiment, depicted with large open symbols in panels C and D, produced no significant changes in performance.

Figure 2.7 *(following page)*: The effects of SC inactivation on local motion discrimination. A) Single stimulus motion discrimination task. After a brief fixation period, a stochastic motion patch appeared in one quadrant of the visual field. After a random delay (480 ms plus a geometrically distributed interval with mean 1000 ms), a brief (160 ms) coherent motion pulse occurred. The direction of motion was drawn at random from any of the four diagonal directions. Monkeys reported the direction of motion by making a saccade to a response dot in the same direction. Data were pooled across sessions based on the direction of motion with respect to the affected quadrant: ipsilateral, other quadrant on the same side; opposite, diagonally opposite quadrant on the other side; contralateral, directly opposite quadrant on the other side. B-E) Performance on motion discrimination task for subject F for each of the four motion directions. Black circles show correct task performance in control sessions. Red circles show correct performance post-injection. Error bars represent 95% multinomial confidence intervals. Solid lines are fits by multinomial logistic regression and dashed lines indicate 95% confidence intervals on the fits. Insets show the bias (control in white, post-injection in gray) and sensitivity (control in black, post-injection in red) relative to choices into the quadrant containing the stimulus. Error bars represent standard error of the fitted parameters. After inactivation, bias significantly increased away from the injection site and sensitivity significantly decreased for all directions of motion. F-I) Performance on motion discrimination task for subject M.

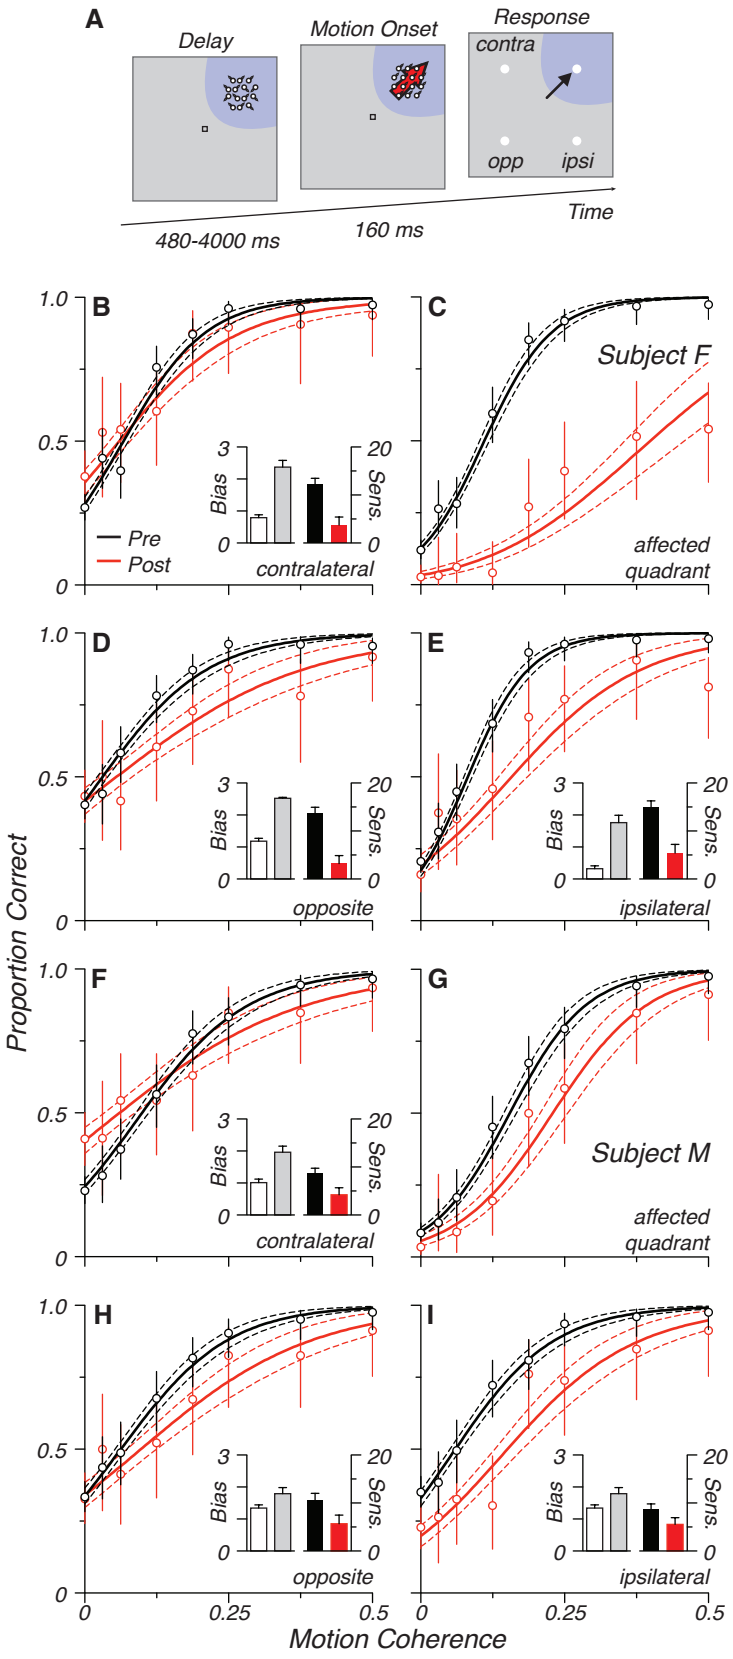
A
Delay
Motion Onset
Response
contra
opp
ipsi
480-4000 ms
160 ms
Time
B
C
D
E
F
G
H
I
Subject F
Subject M
Pre
Post
Bias
Sens.
contralateral
affected quadrant
opposite
ipsilateral
Proportion Correct
Motion Coherence
1.0
0.5
0
0.25
3
20

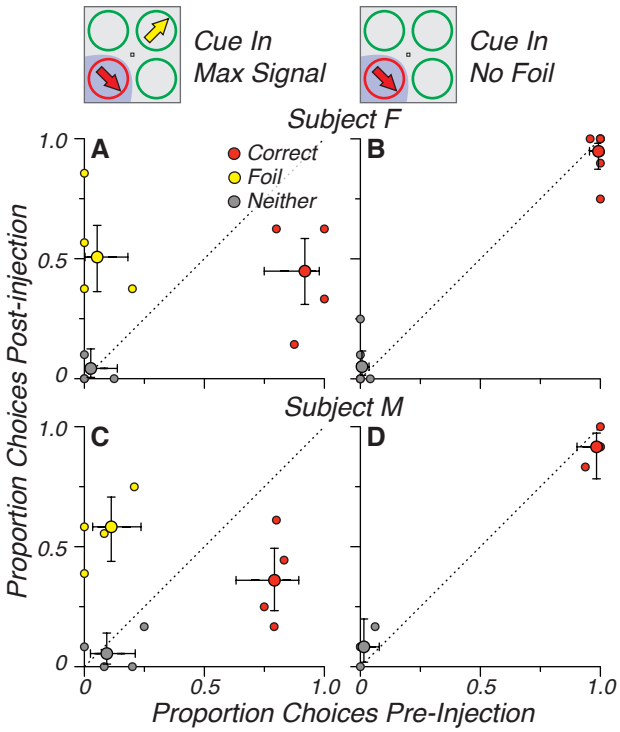

Figure 2.8: Impairments in selective attention after SC inactivation required the presence of a foil signal. A, C) Post-injection, monkeys tended to base judgments on the near-threshold foil signal in the unaffected region even when the cued signal in the affected region was set to maximal coherence. Individual symbols without error bars indicate performance on individual sessions; symbols with error bars indicate pooled performance. B, D) When no foil signal appeared in the unaffected region, however, monkeys successfully ignored the three distracter stimuli with 0% motion coherence and based their choices on the cued signal.

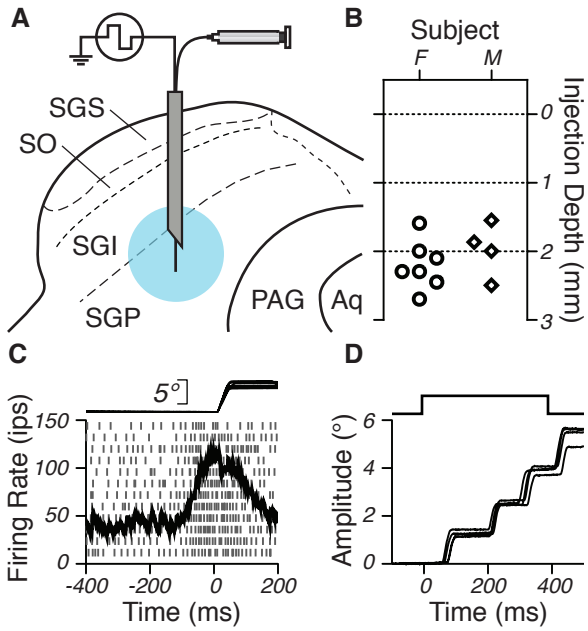

Figure 2.9: Identifying inactivation sites in the Superior Colliculus. A) Schematic of injection experiments. We injected muscimol into the intermediate and deep layers of the SC through an injection cannula with an electrode threaded down its barrel (See Methods Summary). We targeted the intermediate and deep layers of the superior colliculus (SGS, stratum griseum superficiale; SO, stratum opticum; SGI, stratum griseum intermedium; SGP, stratum griseum profundum; PAG, peri-aquaductal gray; Aq, cerebral aqueduct); scale as in panel B. B) Depth of injections. Circles show the 7 inactivation sites for subject F; diamonds show the 4 sites for subject M. C) Example saccade related activity from a unit recorded prior to inactivation. Traces represent radial eye position aligned on saccade onset; spike density function is mean firing rate, and raster represents spike times. Saccade related units produce bursts of activity for saccades made into their response fields and not for those made elsewhere in the visual field. Since the units are arranged in a retinotopic map, observation of activity at a particular saccade target localizes the injection. D) Staircase of saccades evoked my microstimulation. Line above saccade traces indicates duration of microstimulation. Since units code for particular saccade vectors, stimulation evokes saccades to the corresponding position in visual space.

Chapter 3

Oscillatory Activity in Primate Superior Colliculus Related to Spatial Working Memory

3.1 Abstract

Working memory is thought to be stored by neural circuits with reverberant oscillatory activity. Oscillatory neural activity and synchrony in the gamma-band (30-80 Hz) is a ubiquitous property of neural circuits in cerebral cortex underlying a variety of functions including not only working memory, but also attention and movement planning. We now report gamma-band oscillatory activity outside of cortex in the superior colliculus, a brainstem structure associated with the control of orienting movements. We find that neurons in the intermediate and deep layers of the superior colliculus have oscillatory spiking activity in the gamma band specifically when a spatial location must be remembered. These results show that the superior colliculus contains dynamic memory-fields like those seen in cortex and suggests that the SC is part of a distributed circuit for spatial working memory.

3.2 Introduction

Memory-contingent behavior requires an animal to transform a sensory stimulus into a persistent neural representation prior to making a response (Fuster, 2000). Understanding how these representations are stored and how they contingently shape neural processing is of fundamental importance to understanding neural computation. Reverberant oscillatory activity has been proposed as a mechanism for storing working memory in neural circuits (Amit, 1995) (Bodner et al., 2005; Goldman-Rakic, 1995; Fuster, 2000). Such circuits are reentrant in the sense that a pattern of activity is sufficient to re-excite the circuit, thus perpetuating the pattern in time and allowing the circuit to hold its current state (Bodner et al., 2005). The regular repetition of a pattern of spikes should lead to an observable oscillation in the activity of the neuron. Indeed, oscillatory activity has been observed in a variety of contexts involving working memory and context-dependent computation, in particular attention and action selection (Fries, 2009; Howard et al., 2003; Osipova et al., 2006; Pesaran et al., 2008; Womelsdorf and Fries, 2006).

The potential importance of reverberant oscillatory activity in the maintenance of spatial memory is underscored by the appearance of dynamic memory-fields in parietal cortex. In monkeys performing memory-guided saccade and reaching tasks, neurons in the lateral intraparietal area (LIP) (Pesaran et al., 2002) and the parietal reach region(Buneo et al., 2003) have oscillatory activity associated with the location of a remembered goal. The region of space for which this oscillatory activity appears is the dynamic memory-field of the cell (Pesaran et al., 2002).

In this study we investigated the potential presence of dynamic memory-fields in the superior colliculus (SC) of monkeys performing a memory-guided saccade task (Hikosaka and Wurtz, 1983b). The superior colliculus is a brainstem region associated with the generation of saccades and stimulus-driven orienting (Fecteau et al., 2004). The SC therefore may not seem the most likely place to find circuits that store spatial memory or allow contingent processing. Indeed, the primacy of the frontal eye fields (FEF) in memory-guided saccades is well established. Inactivation

of the FEF increases the latency and inaccuracy of memory-guided saccades into the effected region of visual space based on the duration of time over which the memory must be maintained (Dias and Segraves, 1999), as if the location is rapidly forgotten. When multiple distracters are present, memory-guided saccades are eliminated (Keller et al., 2008). Microstimulation of the FEF alters the memory for spatial location (White and Snyder, 2007). The FEF provides information about spatial memory to the superior colliculus via a direct projection (Sommer and Wurtz, 2001), as does the dorsolateral prefrontal cortex (Johnston and Everling, 2008). These projections could explain the presence of persistent elevations in activity during memory-guided saccades in the superior colliculus (Munoz and Wurtz, 1995). The functional importance of this persistent activity is uncertain, although inactivation of the SC does increase the latency and increase the variance in end-point of memory-guided saccades (Hikosaka and Wurtz, 1985a).

Nevertheless, the FEF does not appear to be the only locus of spatial memory storage for saccades, because dynamic memory-fields have been found in parietal cortex. An alternative hypothesis is that spatial memory for saccades is distributed throughout the oculomotor system, and that the FEF is necessary for maintenance of spatial memory across the circuit. In this case the role of the FEF may be similar to that seen during attention, in which neurons in the FEF appear to induce oscillatory activity in visual cortical areas such as V4 (Gregoriou et al., 2009). We therefore speculate that dynamic memory-fields might also be present in the superior colliculus, as well as in parietal cortex. If reverberant oscillatory activity is indeed the mechanism of storing this information at the circuit level, then oscillatory activity associated with spatial memory would be found in the SC. Therefore a first step to identifying the presence of a distributed working memory system that includes the SC would be observing memory-contingent oscillations in spiking activity.

We now report the presence of oscillatory spiking activity in the superior colliculus during a memory-guided saccade task. We examined persistent memory-related activity in SC neurons using spectral analysis (Mitra and Bokil, 2007). Al-

though most neurons studied had temporally unstructured activity, a subset had distinct oscillatory peaks in the power spectra. Most of these peaks were in the gamma band (30 to 80 Hz). Moreover, this oscillatory activity was related to the presentation of a target within the cells' movement field. It appeared only during working memory and not during target presentation or saccade execution. These results support the conclusion that the superior colliculus contains dynamic memory-fields and indicate that memory for spatial location is encoded by neural circuits with reverberant oscillatory activity that include the SC.

3.3 Results

We evaluated oscillatory activity in a total of 271 neurons recorded in the intermediate and deep layers of the superior colliculus in 3 monkeys that performed a memory-guided saccade task (Figure 3.1A). Oscillatory activity was detected by estimating power spectra for spiking activity during the working memory portion of the saccade task and evaluating the spectra for the presence of appropriately shaped peaks consistent with oscillatory modes in the firing rate (see Methods). Out of the total, 148 neurons (55%) had both saccade-related responses on the memory-guided saccade task as well as maintained activity during the memory period (Methods). Of these, 63 (43%) were buildup units and 13 (9%) had oscillatory modes; 11 (17%) of the buildup units had oscillatory modes.

3.3.1 Oscillatory activity appeared during a memory-guided saccade task

Some units in the SC had oscillatory activity during the memory-guided saccade task. This oscillatory activity mostly appeared in the gamma band and was exclusively associated with working memory. An example of a buildup unit with oscillatory activity appears in Figure 3.1. When the target appeared in its movement field, the neuron began firing after target onset (Figure 3.1E) and maintained its

discharge rate until fixation offset (Figure 3.1F), at which time its firing rate rose and peaked immediately prior to the saccade. In contrast, the neuron was mostly quiescent when the target appeared outside of its movement field. The spike spectrogram (Figure 3.1G and 1H) showed an increase in power within the gamma band starting near the time of target offset and continuing until the saccade. The spectrum during the memory period (Figure 3.1I) had a peak in the gamma band (45 Hz) that was significantly greater than the mean firing rate (dashed line). When the target appeared outside of the movement field, few spikes occurred during the memory period and the power was consequently negligible. The peak in the spectrum when the monkey remembered a target in the movement field confirms the presence of spatially specific oscillatory activity during the memory period. Spectra for spiking activity with refractory periods tend to have dips at low frequencies and then rise to a limiting value corresponding to the mean firing rate as expected for a homogenous Poisson processes (Jarvis and Mitra, 2001), but the peak in this example indicates that this neuron consistently had an oscillation in its instantaneous firing rate at 45 Hz.

Most neurons in the SC with sustained memory related activity did not have oscillatory peaks. Data from three neurons, illustrating the range of SC neuron types with memory-related activity, appear in Figure 3.2. These neurons include a buildup unit with oscillatory activity (I, Panel A), a buildup unit without an oscillatory peak (II, Panel B), and a unit with prelude activity preceding a saccade burst (III, Panel C). Units I and II show the typical rise in activity well in advance of the saccade as seen in buildup units. The spectrograms for each unit (Panels D-F) show that unit I had elevated power in the gamma band during the memory period preceding the saccade, whereas units II and III did not. In these panels, the power is shown as a ratio of power in the window centered at a given time and frequency over the mean firing rate within that window. A unit with purely Poisson firing would have a ratio of 1 at all frequencies and would have a spectrogram that was uniformly green. Concentration of power at a certain frequency appears as hot colors rising to red, and suppression, due to factors such as the refractory period, appears as cool colors sinking to blue.

Buildup unit I had oscillatory activity during the memory period, which appears as the reddish ribbon near 45 Hz. Units II and III did not have oscillatory activity, and their spectrograms appear more uniform during the memory period, although they also had large saccade bursts, which produced the local increases in low-frequency power at the time of the saccade.

Spectra during the memory period confirm the presence of oscillatory activity in the first example unit and its absence in the remainder (Figure 3.2G, H, and I). Spectra are shown as the ratio of the power at a given frequency to the high frequency limit (see Methods). Unit I had a statistically significant peak (the ratio exceeds the 95% confidence threshold, Fcrit) that is appropriately narrow, demonstrating the presence of an oscillatory mode at a frequency of 45 Hz in its spiking activity. In contrast, Units II and III did not have oscillatory peaks. Unit II shows a low frequency increase in power indicative of bursting activity; since closely spaced spikes are separated by wide intervals, low frequency power emerges when the spectra is calculated for long windows (Bair et al., 1994). Unit III lacks the low frequency peak and has power rising to the mean firing rate as expected for a neuron with unstructured, Poisson firing. The candidate peaks for both units II and III are marked by f_p although neither was significant.

3.3.2 Units with oscillatory peaks form a distinct subset of buildup units

Nearly all units with significant oscillatory activity also met the criteria for buildup activity, as shown in a scatter plot of power ratio against the firing rate 100 ms prior to the saccade (Figure 3.3A). Since different numbers of trials were collected for different neurons, we normalized the power ratio by the threshold so that the ratio could be referenced off the same standard for all. Points lying above 1.0 (solid symbols) represent units with statistically significant oscillatory peaks, whereas points lying below this line (open symbols) did not have significant peaks. In addition, points lying to the right of the horizontal dashed line (30 ips) meet the standard for buildup

units (Munoz and Wurtz, 1995). Example units from Figure 3.3 are indicated with enlarged symbols and labeled (I, II and III).

The significant peaks present in the spectra of these neurons represent the presence of lines, or discrete oscillatory modes, rather than broadband elevations in firing rate. All of these points represent neurons with narrow peaks consistent with oscillatory modes (Figure 3.4A). The power spectra for all of these units appears in Figure 3.4B.

Neurons with oscillatory peaks appeared to form a distinct subset of buildup units. The histogram of power ratio over critical threshold was not unimodal (Figure 3.2B). The distribution appears to have two modes, the first near 0.9 for units without an oscillatory peak, and near 1.1 for neurons with an oscillatory peak. This distribution failed a Hartigan's dip test for unimodality (dip = 0.04, p =0.6) (Hartigan and Hartigan, 1985) and was better fit with a mixture model of two distributions than a single distribution according to the Akaike Information Criteria (-86 vs. -117) (Akaike, 1974). Hence, units with an oscillatory peak appear to be a distinct group rather than the tail of a unimodal distribution. Most neurons with significant peaks had oscillatory activity in the gamma band (30 to 80 Hz) and were buildup units (Figure 3.2C). Only four had significant peaks in the beta band (15 to 30 Hz), and two of these were buildup units.

In addition to the presence of oscillatory activity during the memory period, these neurons tended to have a characteristic set of response properties. These neurons tended to have muted saccade bursts (mean of 117 ips as opposed to 164 ips for all buildup units), see Figure 3.5. The movement fields as mapped with delayed saccades (see Methods) tended to be large and we were generally unable to find the distal edge of the field. These characteristics were not unique to buildup units with oscillatory activity, however; the distribution of mean saccade bursts fell within the range typical of buildup units, and the shape of the movement field was within the normal variation for buildup units. Overall these neurons appear to be a group identified by the presence of an oscillatory peak in the power spectrum during the

memory period.

3.3.3 Oscillatory activity is specifically associated with working memory

On average, oscillatory activity during the saccade task was specifically associated with working memory. In neurons with oscillatory activity in the gamma band (n=11), the oscillatory mode emerged near target offset, after which the monkey was required to remember the target location (Figure 3.6A). It continued until fixation offset, after which the monkey made a saccade to the remembered location (Figure 3.6B). In particular, the oscillatory activity vanished well before saccade execution (Figure 3.6C). A similar pattern held for four neurons with oscillatory activity in the beta band (Figure 3.6D-F). Since spectrograms were normalized by firing rate prior to averaging, concentration of power at a specific frequency is indicative of an oscillatory peak and not simply an increase in firing rate.

The average spectrogram reveals that oscillatory activity is temporally specific to the memory periodit begins at target offset, when the monkey must commit to memory where to saccade, and ends at fixation offset, when that saccade is made. Moreover, it disappears at fixation offset, rather than continuing through the saccade, indicating that the presence of oscillatory activity is specifically related to those aspects of saccade generation requiring the temporary maintenance of a spatial memory. Finally, since the neurons do not respond to targets presented outside of their movement fields, the oscillatory activity was spatially specific.

3.4 Discussion

In this report we have described two principle findings. First, we examined the activity of superior colliculus neurons during a memory-guided saccade task and found that a subset of neurons had significant oscillatory activity in the gamma band (30-80 Hz). Second, we determined that this activity was restricted to the period of

time in which the monkey was required to temporarily remember the target location. Thus the oscillatory activity appears to be specifically linked to working memory. In addition, the oscillatory activity was spatially specific since the neurons responded only to targets presented within the movement field. These results demonstrate that oscillatory activity in the gamma band is a feature of neural computation outside of cerebral cortex and that dynamic memory-fields exist in subcortical structures such as the superior colliculus.

Although oscillatory activity appeared in a task in which the monkey remembered a spatial location in the absence of visual stimuli, it seems likely that the activity would occur during other behaviors in which spatial working memory is required. The memory-guided saccade task involves turning off the stimulus in order to exclude sensory responses, but spatial memory is still important, and perhaps more frequently used, in situations when stimuli remain visible but some form of memory is needed to properly prioritize or identify which stimuli should be processed. Neurons with temporally structured activity are co-localized with neurons with temporally unstructured activity in the SC; such a dichotomy may represent a segregation of function. Neurons with oscillatory activity might represent the presence of an object of interest at a specific spatial location in a more or less discrete or categorical fashion. In contrast, neurons with temporally unstructured activity might represent other aspects of stimuli at the corresponding location in a more graded fashion (Brody et al., 2003). In this scheme, cells with dynamic memory-fields would serve as a spatial index, whereas other neurons in the circuit might process features such as salience, which is a more widely accepted interpretation of maintained activity of SC neurons (Basso and Wurtz, 1997; Fecteau et al., 2004). Doing otherwise might conflate the presence of an object with the degree of its behavioral relevance.

If spatial memory is encoded in dynamic memory-fields distributed across multiple oculomotor regions, then coordination of those representations is necessary. Gamma-band synchronization between different cortical areas has been observed during both working memory and attention (Gregoriou et al., 2009; Pesaran et al., 2008),

and this could reflect a coordination of memory-fields. Such synchronization has been proposed to underlie the consistency or stability of visual perception and the binding of features (Singer and Gray, 1995). Furthermore, disruption of these coordinated representations at any stage might be expected to cause the coherent representation to lose stability. Indeed, inactivation of oculomotor structures such as the SC, LIP, or FEF causes deficits in the ability of the animal to properly select targets for saccades from amongst distracters (Keller et al., 2008; Mcpeek and Keller, 2004; Wardak et al., 2002). In these cases the targets in the effected regions of visual space seem to disappear from the decision process. Overall, reverberant oscillatory activity in oculomotor structures might not be so much a system for storing spatial memory, but instead a common system for categorical representation of objects of interest in perception and action selection.

3.5 Methods

3.5.1 Behavioral task

The monkeys viewed stimuli presented on a CRT display with a refresh rate of 75 Hz at a viewing distance of 410 mm. The displayed screen size was 389 mm by 293 mm at a resolution of 1024 by 768 pixels. Each target consisted of a single pixel of background luminance surrounded by a 1-pixel-thick white border. Given the display geometry and the viewing distance, this 3x3 pixel stimulus corresponded to 9 by 9 minutes of arc of visual angle. The background luminance of the monitor was 14 cd m^{-2} and the luminance of the targets was 50 cd m^{-2}.

The memory-guided saccade task proceeded as follows (see Figure 3.1A). Monkeys first fixated a centrally appearing dot (F on) for 500 to 1000 ms. A second dot appeared at the location to which a saccade should be made (T on) and remained on the screen for 500 to 900 ms (the overlap period). The target then disappeared (T off), and the monkey maintained fixation while remembering the location of the dot until the fixation dot disappeared (F off) after 500 to 800 ms (the memory period).

The monkey then made a saccade to the remembered location upon fixation offset. Monkeys were rewarded for making saccades landing within a 2° window of the target position. Monkeys maintained fixation throughout the entire trial (Figure 3.1C) and only made saccades following fixation offset (Figure 3.1D). We placed the target either at the maximally responsive location within the unit's movement field as mapped with the delayed saccade task or at the symmetric location in the opposite quadrant (Figure 3.1B). We randomly interleaved these target positions and performed between 20 and 40 trials per location per neuron.

3.5.2 Animal preparation

We recorded from three adult male rhesus monkeys that were 10-15 years of age and weighted 15-16 kg. The monkeys were prepared using standard surgical techniques described in detail previously (Krauzlis et al., 2000). In particular, each animal had a recording chamber for SC single neuron recording affixed to the skull with dental acrylic and additional titanium screws. The chamber was angled 38° to the posterior of vertical and directed at the midline 15 mm above and 1 mm posteror to the interaural line.

All experimental protocols were approved by the Institutional Animal Care and use Committee and complied with United States Public Health Service policy on the humane care and use of laboratory animals. The laboratory setup for behavioral control and monitoring was identical to that described previously (Hafed and Krauzlis, 2008).

3.5.3 Electrophysiology and neuron classification

We used tungsten microelectrodes (Frederick Haer, FHC) with impedances: 1-2.5 MΩ to record extracellular action potentials of individual neurons in the intermediate and deep layers of the SC (1.5 2.9 mm below surface). Electrodes were advanced through stainless steel guide tubes (23 gauge) with a microdrive mounted on top of the recording chamber. The guide tubes were held fixed in the chamber

with a delrin grid system (Crist et al., 1988). Extracellular neuron activity was passed through a standard head stage, amplified, and converted into trigger pulses with a window discriminator applying both time and amplitude criteria (Plexon Systems, Inc.). The time of each action potential was stored with 1-ms resolution.

We mapped the movement field of each neuron using a delayed saccade task. Movement field mapping typically required 50-150 trials in which the target was selected on each trial by the experimenter in order to maximally elucidate the extent of the field and the location to which the neuron maximally responded. We calculated the movement response as the mean discharge rate within a window starting 8 ms prior to saccade onset and ending 8 ms prior to saccade offset. We placed the target for memory-guided saccades at the location to which the neuron was maximally responsive during the movement field mapping.

Our criteria for inclusion in this study were that in the memory-guided saccade task, neurons had both saccade-related responses and maintained discharge during the memory period. We defined a saccade-related response as discharge greater than 60 impulses per second (ips) in the same interval used for delayed saccades. Response during the memory period was at least 5 ips. Out of the 271 neurons recorded, 148 met both of these standards. We further distinguished between neurons with prelude activity prior to the saccade burst and neurons with buildup activity, which is a functional characterization that describes an increase in firing rate hundreds of milliseconds prior to the saccade. A common threshold for classifying a unit as buildup is if the discharge rate meets or exceeds 30 ips at time 100 ms prior to the saccade onset (Munoz and Wurtz, 1995). Out of 148 neurons, 63 met the standard for buildup. Finally, we examined all 148 neurons for oscillatory activity.

3.5.4 Data Analysis

We used spectral analysis to detect the presence of oscillatory activity and to characterize the temporal structure in the data. We used multitaper methods as implemented in the Chronux(Mitra and Bokil, 2007) package for MATLAB to

estimate the spectrum of the spike counts, or the spike spectrum. When estimating spectra, we used a frequency resolution of 7 Hz on a 500 ms window aligned on fixation offset, leading to six Slepian data tapers. When estimating the spectrograms, we used a frequency resolution of 14 Hz on a 250 ms moving window at steps of 10 ms with the time index aligned to the center of the window. Again, this time-frequency resolution led to six tapers. We constructed confidence intervals on the spectral estimates using estimates of variance from jackknife over both tapers and trials, although we used F-distribution statistics to detect the presence of oscillatory modes (Jarvis and Mitra, 2001; Thomson, 1982). When averaging across neurons, we normalized power by the mean firing rate (specifically, the high frequency limiting power as described below). This transformed intensity values from representing power, which rises monotonically at all frequencies as firing rate increases, to representing the concentration of power in an oscillatory mode.

Detecting oscillatory modes included two stages. The first stage was identification of potential peaks in the power spectrum, and the second was evaluation of those peaks to determine if they were oscillatory modes. We identified potential peaks as follows. First, we identified the lowest point in the spectrum at low frequency. For an otherwise Poisson processes with a refractory period, this point should be at zero. However if the spiking activity is characterized by burst firing, in which spikes are clustered in widely spaced groups, power will be elevated at low frequency (representing the inter-burst interval) followed by a dip in the power spectrum before it rises back to the mean firing rate (see, for instance, Figure 3.2H) (Bair et al., 1994). Second, we found the first peak above this dip that exceeded the mean firing rate. We then evaluated these potential peaks to determine if they corresponded to oscillatory modes.

We used two features to determine if potential peaks were oscillatory modes: the significance of the peak and its width. The basic premise is that an oscillatory mode should appear as a significant and narrow elevation in power above noise. To evaluate the significance of the peak above noise, we examined the ratio of power at

a given frequency to that at the high frequency limit. The high frequency limit is the total variance in the spiking activity and, at least for homogenous Poisson processes, is equivalent to the mean firing rate (Jarvis and Mitra, 2001). The high frequency limit is calculated as the mean square of the projection of the spike trains on the data tapers. If I is the power spectrum and h_k is the kth data taper, then

$$I(f \to \infty) = \frac{1}{N_T K} \sum_{k=0}^{K-1} \sum_{n=1}^{N_T} \sum_{j=1}^{N_n(T)} h_k(t_j^n)^2 \tag{3.1}$$

where t_j^n is the jth spike in the nth trial and $N_n(T)$ is the total number of spikes in the nth trial. Since power is a measure of variance at a given frequency, its variability is described by a central χ^2 distribution with, in this case, $2KC$ degrees of freedom where K is the number of data tapers and C is the number of channels. Consequently the ratio of power at a given frequency to that at the high frequency limit follows an F distribution with $2KC$ degrees of freedom in both the numerator and the denominator (Thomson, 1982). We identified a potential oscillatory peak as significant if the power ratio exceeded the 95% confidence interval for the $F_{2KC,2KC}$ distribution.

Not all significant elevations in power correspond to oscillatory modes, however. A broadband elevation in power might produce a statistically significant peak, but a single oscillatory mode should appear as a line in the spectrum. Nevertheless, the multitaper spectral estimator will introduce bias in frequencies adjacent to the line, creating a peak with a width no more than twice the bandwidth of the estimator (Figure 3.4A) (Jarvis and Mitra, 2001; Thomson, 1982). The expected value of the spectral estimator with K tapers for the true spectrum S at a frequency f is

$$E\{I(f)\} = \int_{-\infty}^{\infty} S(f') \frac{1}{K} \sum_{k=0}^{K-1} |H_k(f - f')|^2 \, df' - \frac{1}{KT} \sum_{k=0}^{K-1} |H_k(0)|^2 S(0) \tag{3.2}$$

Since our model of the true spectrum has a line at f_p and is zero everywhere else, the shape of the estimated spectrum around the line should be described by

$$E\{I(f)\} \approx \frac{1}{K} \sum_{k=0}^{K-1} |H_k(f - f_p)|^2 \tag{3.3}$$

Therefore we excluded peaks with width greater than 14 Hz. A scatter plot of frequency and width of peaks along with the spectra of units with detected modes

appears in Figure 3.5A and B.

Chapter 3, in full, has been submitted for publication of the material as it may appear in *Neuron*, 2010, Lovejoy, LP; Kruazlis, RJ. The dissertation author was the primary investigator and author of this paper.

Figure 3.1 *(following page)*: Oscillatory activity during a memory-guided saccade task. A) Memory-guided saccade task. After the monkey fixated a centrally appearing dot (F on), a target appeared (T on) for 500 to 900 ms (overlap period). The monkeys' task was to remember the location of that target after it disappeared (T off) during the 500-800 ms memory period until the fixation dot disappeared (F off). The offset of the fixation spot signaled the monkey to make a saccade (S) to the remembered location. B) Schematic of movement field and target placement for memory-guided saccades. Targets were placed at the position in the movement field with maximum response and at the symmetric location in the opposite quadrant. C) Radial eye position aligned to target onset. The monkey maintained fixation throughout the overlap and memory periods. D) Amplitude of eye position aligned to saccade onset (S). Distribution of the delay by which fixation-offset precedes saccade onset is shown by the whisker plot. Dotted line represents average time of fixation offset with respect to saccade onset. Saccades into movement field are in black; saccades out of movement field have negative sign and are in blue. E) Mean firing rate and spike raster for example neuron, aligned to target onset. Raster plot for saccade made into the movement field is above that saccade made out of the movement field. F) Mean firing rate and spike raster aligned to saccade onset. G) Spike spectrogram for example neuron aligned at target onset. Time-frequency bandwidth was 14 Hz by 250 ms. Heat map represents the power at a given frequency and time. H) Spectrogram aligned on saccade onset. I) Spike spectrum during memory period. Dot indicates position of statistically significant oscillatory peak in spectrum. Error bars are 95% confidence intervals by jackknife. Dashed line is high frequency limit on power (mean firing rate). Frequency bandwidth is 7 Hz and spectrum is calculated for the 500 ms of each trial immediately preceding fixation offset.

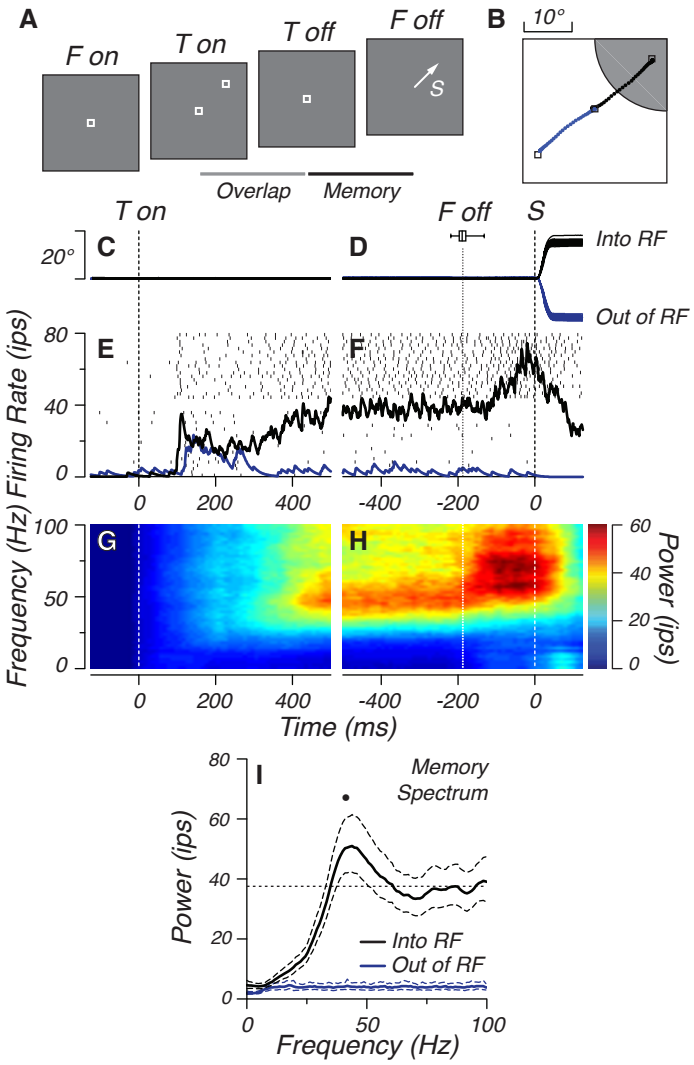
A
F on
T on
T off
F off
S
Overlap
Memory
B
10°
T on
F off
S
C
D
20°
Into RF
Out of RF
E
F
80
40
0
0
200
400
-400
-200
0
Frequency (Hz) Firing Rate (ips)
G
H
100
50
0
60
40
20
0
Power (ips)
Time (ms)
I
Memory
Spectrum
Power (ips)
Into RF
Out of RF
Frequency (Hz)
50
100

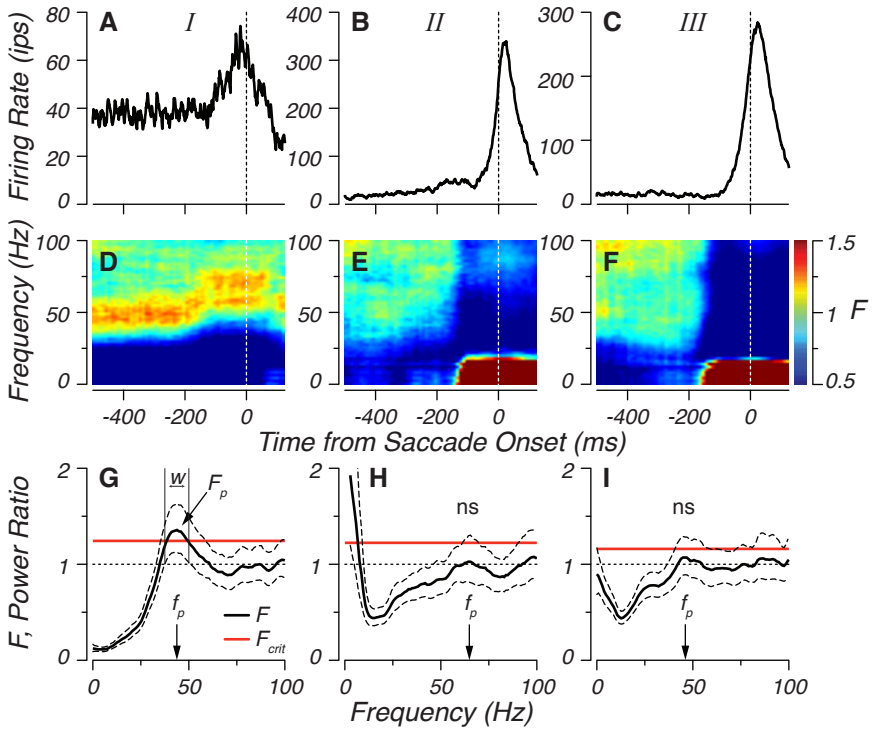

Figure 3.2: Example units and spectra. A-C) Firing rate aligned to saccade onset for buildup unit with significant oscillatory activity (I), buildup unit (II) and saccade burst unit with prelude activity (III), both without significant oscillatory activity. Roman numerals correspond to labeled points in Figure 3.2. D-F) Power-ratio spectrograms aligned to saccade onset. Heat map represents concentration of power in oscillatory modes. G-I) Power-ratio spectra from memory period for example units. Spectra are plotted as the ratio of power to the high frequency limit (mean firing rate). Red line indicates critical ratio for significant peaks. Panel G shows spectrum for the buildup unit with an oscillatory peak in the gamma band (f_p), with a significant power ratio (F_p), and a width (w). Panels H and I show spectra for units without significant power ratio at candidate peaks (f_p) marked as ns (not significant).

Figure 3.3 *(following page)*: Summary of oscillatory peaks in SC neurons. A) Scatter plot of power ratio of potential oscillatory peak versus mean firing rate 100 ms prior to saccade onset. Vertical dashed line marks 30 ips criteria for buildup units. Open symbols represent units without a significant peak. Solid symbols represent units with a significant peak. Roman numerals indicate example units shown in Figure 3.2. B) Histogram of power ratio for buildup units. Solid bars represent population of units with significant oscillatory peak. This distribution failed tests for unimodality (Hartigan's dip test). C) Scatter plot of frequency of oscillatory peak versus mean firing rate. Vertical dashed line represents buildup criteria; horizontal dashed lines represent boundaries of frequency bands (gamma, γ, and beta, β). For most buildup units with significant oscillatory peaks, the reverberant frequency was within the gamma band (30 to 80 Hz).

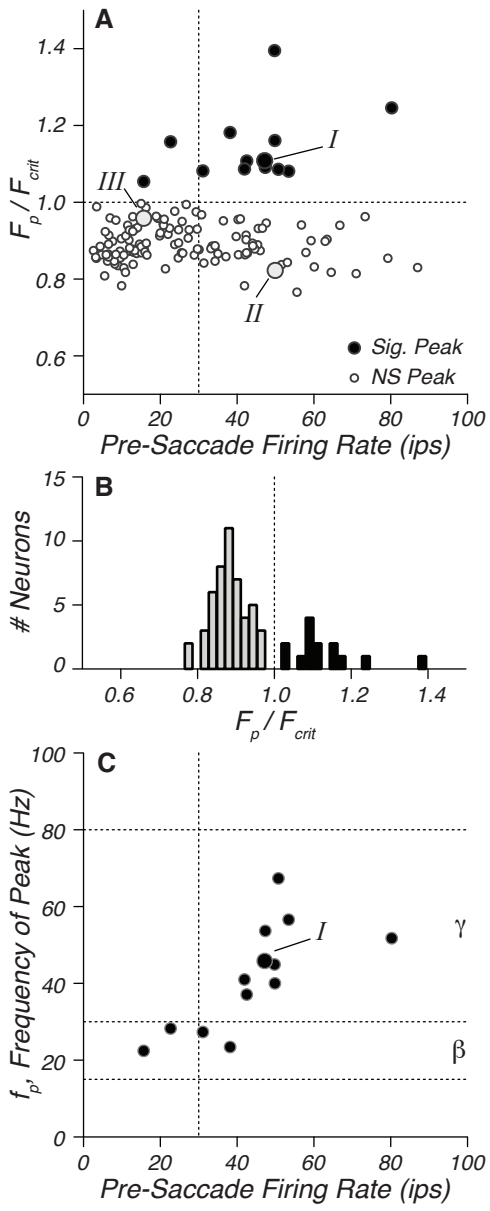
A
1.4
1.2
1.0
0.8
0.6
F_p / F_{crit}
I
III
II
Sig. Peak
NS Peak
0
20
40
60
80
100
Pre-Saccade Firing Rate (ips)
B
15
10
5
0
Neurons
0.6
0.8
1.0
1.2
1.4
F_p / F_{crit}
C
100
80
60
40
20
0
f_p, Frequency of Peak (Hz)
I
γ
β
0
20
40
60
80
100
Pre-Saccade Firing Rate (ips)

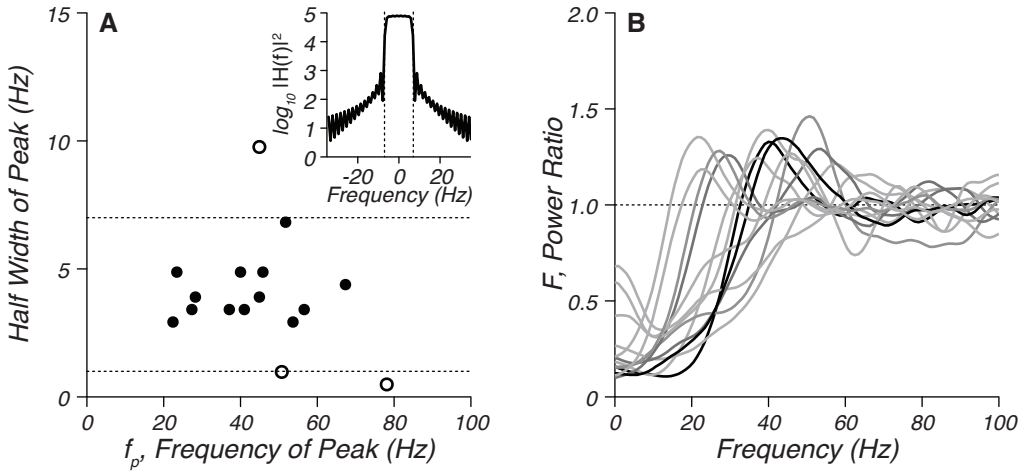

Figure 3.4: Frequency and widths of spectral peaks. A) Scatter plot of half width of oscillatory peak against frequency of peak. Dashed lines represent 1 Hz and 7 Hz bounds on width based on the bandwidth of the spectral estimator. Solid symbols show units that satisfy the conjunction of requirements on significant power and half width boundaries. Inset shows average magnitude of the Fourier transform of the data tapers. The dashed lines represent the bandwidth of the spectral estimator, 7 Hz. B) Spectra for all neurons with significant oscillatory peaks and appropriate half widths.

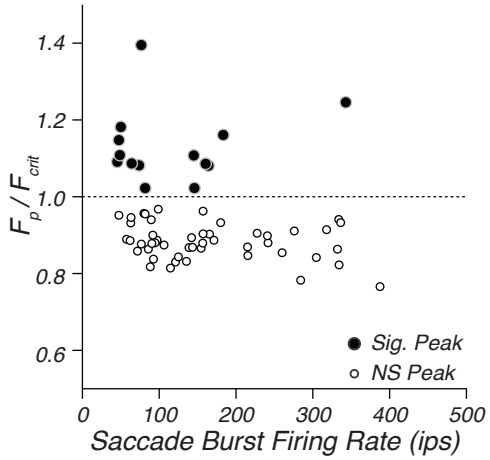

Figure 3.5: Scatter plot of Power ratio against firing rate during saccade burst for buildup units. Units with a significant oscillatory peak (black symbols) tend to have lower magnitude bursts (mean 117 ips) than buildup units without significant peaks (mean 164 ips, open symbols), although they lie within the distribution of burst magnitudes typical of buildup units.

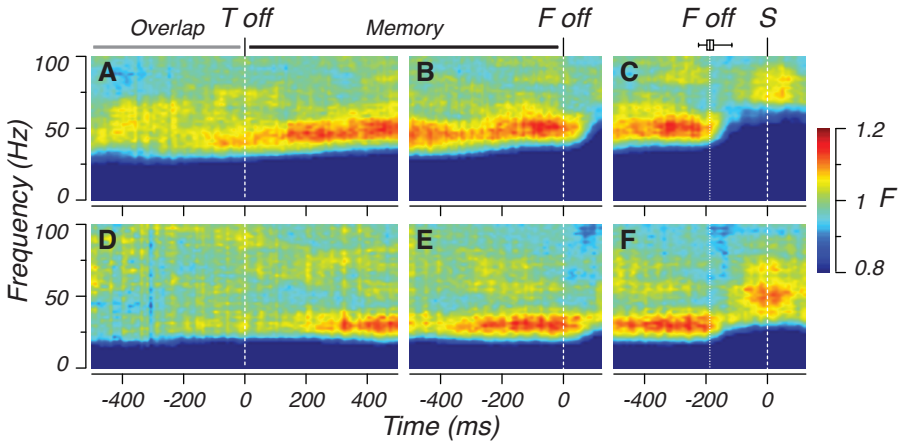

Figure 3.6: Mean power-ratio spectrogram for units with oscillatory peaks. A) Target onset aligned spectrogram for units (n=11) with peak in the gamma-band. Time-frequency bandwidth was 14 Hz by 250 ms. Overlap period spans time starting with target onset (not shown) and ends with target offset at time zero. Memory period begins with target offset and ends with fixation offset in Panel B. Heat map represents concentration of power into an oscillatory mode. The dark red region of plot indicates the presence of an oscillatory mode beginning after target onset and continuing through the memory period. B) Fixation offset aligned spectrogram for units with peak in the gamma-band. Power is concentrated in an oscillatory peak throughout the memory period and ends at fixation offset. C) Saccade onset aligned spectrogram for units with peak in the gamma-band. Whisker plot indicates distribution of time of fixation offset with respect to saccade onset. Concentration of oscillatory power ends near fixation offset, prior to saccade onset, and does not continue through saccade execution. D) Target onset aligned spectrogram for units (n=4) with peak in the beta-band. E) Fixation offset aligned spectrogram for units with peak in the beta-band. F) Saccade onset aligned spectrogram for units with peak in the beta-band.

Chapter 4

Normative Behavior on Selective Attention Task

4.1 Introduction

The goal of this chapter is to examine the monkeys' normative performance on the selective attention task in terms of accuracy, response distributions, and response time, and, by so doing, to gain qualitative insight into the factors influencing performance and the strategies which the animals may employ to complete the task. In particular, I will attempt to find answers for the following questions: first, were there spatial asymmetries in performance that were systematically related to the position of the cue; second, were there spatial asymmetries in response and reaction time distributions that were not systematically related to the stimuli or the configuration of the cues; third, is it possible to determine if the monkeys employed a strategy in which decisions were based on pooling of sensory evidence across the visual field, or could they exclusively select cued signals and ignore foils; and fourth, was there a dependence of performance on the delay between cue offset and stimulus onset. Prior to delving into these issues I will describe the task in detail and clarify what we may take a *strategy* to mean in this context. Finally I will examine the accuracy, response, and reaction time distributions and attempt to inductively arrive at a characteriza-

tion of the monkeys' possible strategies, which when combined with the single-unit data to be presented in Chapter 6, will provide insight into why inactivation of the superior colliculus led to visual extinction as described in Chapter 2.

4.1.1 Task design

Monkeys performed a motion discrimination task that required them to judge the direction of motion in one of four peripheral stochastic motion stimuli while ignoring distracting signals (Figure 4.2). After the monkey maintained fixation for at least 500 ms, a set of cue rings appeared. The odd colored ring indicated which of the four stimuli to which the monkey should attend. For Subject F, the cued position was red and the other three were green. For Subject M, the cued position was green and the other three were red. After 480 ms, stochastic motion patches appeared within each of the four rings. After an additional 480 ms, the cue rings disappeared. Following a delay, brief (160 ms) pulses of coherent motion appeared simultaneously at both the previously cued location and the diametrically opposite foil location. The direction of motion in the cued location was drawn at random from any of the four diagonal directions, and the direction in the foil was drawn from any of the remaining three. The delay was 480 ms with an additional random portion which was geometrically distributed with a mean of 480 ms. Following the motion pulses, choice dots appeared where the patches had previously been centered.

A geometric distribution is a discrete-time version of the Poisson distribution for which the probability of the event (in this case motion onset) occurring at the next time step does not depend on the time elapsed to the current time step. As a result, the hazard function is flat. The selection of such a timing distribution was motivated by the desire to remove any possibility that the monkeys might incorporate the structure of the timing distribution in their vigilance or attention to the task. A quantitative description of a geometric distribution is as follows. If the number of frames of incoherent motion shown prior to the motion pulse is $k, k \in \{0, 1, \cdots, K\}$ and the probability of the motion pulse appearing on the next frame is p, and the

number of frames of incoherent motion shown, Y, then a probability mass function describing the likelihood that k frames of incoherent motion will appear prior to the motion pulse is

$$P(Y = k) = (1 - p)^k \, p \tag{4.1}$$

The hazard function is the conditional probability of the motion pulse occurring on the subsequent frame given that it has not yet occurred, and this requires both the probability mass function and the survival function $S(Y)$, which is the probability that the number of frames of incoherent motion will be greater than Y:

$$S(Y) = \sum_{k>Y} P(k) \tag{4.2}$$

The hazard function, $h(Y)$, is then

$$h(Y) = \frac{P(Y)}{S(Y)} \tag{4.3}$$

For the geometric distribution on an unbounded support, meaning that k could be any non-negative integer, $h(Y) = p$. In the experiment, we set p to be 0.0278, leading to a mean delay of 36 frames or 480 ms. We set the maximum number of frames at 264 to yield a maximum random delay interval of 3520 ms. Therefore the hazard function starts to diverge from the theoretical prediction between 1500 ms and 2000 ms. The empirical and predicted probability mass function for the delay and the hazard function of the delay are shown in Figure 4.1. Without a strong prediction for how the animal will incorporate the deviation from flatness in his behavior, we are left to rely on the fact that delays for which the hazard function is deviating from the ideal are extremely rare and that the animals are not exposed to enough of these trials to warrant concern.

Monkeys received a liquid reward only for correct responses in completed trials. Motion coherence of the pulses was titrated for each subject to maintain 65-70% correct performance on the task (0.1875 for Subject F and 0.25 for Subject M). The response was recorded as the first of the choice dots fixated within a 4° square surrounding the dot after the monkey's eye position was no longer detected as being

within a 1.5° window around fixation. Monkeys were free to make the response saccade anytime after the end of the motion pulse. However the monkeys appeared to initiate saccades prior to the end of the motion pulse so that their eyes were leaving the fixation window at extremely short delays, and so I measured the reaction time from the onset of the motion pulse. All responses had to be completed within 600 ms of the motion onset.

In addition to the cue rings, information about the cued location was available in the block structure of the task. The first four trials of each block of forty-four trials had the highest possible strength of motion (0.5) in the cued location and no motion signal in the foil signal. The subsequent forty trials included four trials in which motion appeared in neither the cued nor the foil signal location and thirty-six trials representing the combinations of signal and foil direction. If the monkeys broke fixation during the trial or failed to reach fixation, the trial was aborted and the combination of signal and foil for that trial was repeated later. In addition, the delay that would have preceded the motion pulse on that trial was repeated later. In this way the monkeys were unable to avoid performing trials with long durations. Furthermore, fixation breaks were penalized with wait times of five seconds to discourage the monkeys from breaking fixation during long trials in the hope of initiating a shorter trial.

Stimuli were presented to the monkeys on a CRT display mounted at eye level at a viewing distance of 410 mm. The viewable portion of the screen was 389 mm by 293 mm at a resolution of 1024 by 768 pixels. The CRT was a BARCO Reference Calibrator V driven at a refresh rate of 75 Hz. At this refresh rate, the delay between frames was $13.\bar{3}$ ms, so in order to insure exact frame counting most intervals are multiples of 40 ms (hence the unusual delay of 480 ms). Monkeys fixated a small, stationary spot presented at eye-level directly in front of them at the center of the display. The fixation spot consisted of a single pixel of background luminance surrounded by a 1-pixel-thick white square. With our display geometry and distance, this 3x3 pixel stimulus corresponded to ∼9x9 min arc of visual angle. The background

luminance of the monitor was 14 cd m^{-2}. Luminance of the fixation dot was 50 cd m^{-2}. The cue rings appeared at an eccentricity of 8.2° of visual angle (fixation point to ring center), had a diameter of 8.8° of visual angle, and had a thickness of 0.25° of visual angle. The ring luminance was 25 cd m^{-2} green or red. The stochastic motion stimuli appeared centered within the cue rings and had a radius of 4.25° of visual angle. The stimuli consisted of limited lifespan dots (Cook and Maunsell, 2002). The dot lifespan was 2 refreshes. At each refresh, each dot either appeared at a random location within the patch or was displaced by 4 pixels ($\sim 0.2°$). The coherence was the proportion of total dots moving in the same direction; the remainder moved in uniformly distributed random directions. Each dot had a peak luminance of 50 cd m^{-2}. Responses were conveyed with saccades made to choice dots appearing at an eccentricity of 8.2° of visual angle and with a radius of 0.2° and peak luminance of 50 cd m^{-2}.

4.1.2 Data analysis

The record of performance on each trial is the subject's response and the latency of the response. The latency of the response is measured as the time delay between the onset of the motion pulse and the initiation of the response saccade. For the purposes of this treatment, trials were pooled across all sessions.

Although responses are strictly either correct or not, two classes of errors are possible. Therefore there are three possible outcomes for accuracy, and when response distributions are examined, four outcomes are possible. Thus multinomial proportions are the appropriate statistic for the response and accuracy distributions. Just as with binomial distributions, the sufficient statistic for the multinomial distribution is not simply the proportion vector but both the proportion vector and the number of trials performed in each session. Pooling data across sessions to obtain the population average involves adding the number of responses in each condition and the number of trials in each session rather than averaging the proportions (Johnson et al., 2005). Furthermore, joint confidence intervals, which account for the covariation of elements

within the multinomial proportion vector, are assessed based on the multinomial distribution. Joint confidence intervals for multinomial distributions are not currently part of the standard MATLAB Statistics Toolbox, and so I computed joint confidence intervals on multinomial proportions following the prescriptions provided by Bailey (1980) and Kwong and Iglewicz (1996).

The reciprocal of latency tends to follow a normal distribution (Reddi and Carpenter, 2000), and so latency data is presented on a reciprocal latency scale. Distributions are visualized with whisker plots, for which the central line represents the median value, the box represents the upper and lower quartiles, the error bars represent the upper and lower limits of the central 95% of the distribution, and individual dots represent outliers. When necessary, standard t-tests can be performed on reciprocal latency to establish the significance of differences between the population means.

4.1.3 Monkey behavioral strategies

In common usage, a *strategy* is a plan of action designed to achieve an overall aim, and use of such a word to describe a monkey's behavior may strike some readers as an unnecessary anthropomorphism. The correct approach to solving a selective attention task would seem to be obvious: maintain fixation, remember which stimulus had been cued, ignore all the rest, and report the qualia of the discriminandum at the appropriate time. However each one of these steps is substantially easier said than done. First of all, maintaining fixation for long periods of time is difficult. Second, even if they understand the meaning of the cue, monkeys may have poor spatial memory or may have difficulty remembering spatial locations longer than a moment. Third, monkeys may have difficulty ignoring stimuli. Fourth, the discrimination tasks are intentionally made quite challenging. Fifth, the monkeys may become impatient or frustrated when made to wait for long periods of time. Hence each stage of the task may be assigned a different subjective cost by each animal, and these costs may, for whatever reason, depend on the spatial aspects of the stimulus. For example, human

observers sometimes appear to be better at motion discrimination in the inferior visual field (Rezec and Dobkins, 2004) and better at contrast discrimination in the inferior visual field (Silva et al., 2007). If monkeys decide on any given trial whether or not to participate in the task based on anticipated difficulty, then the degree to which they attend to the cues could depend on spatial location; such a difference would constitute a difference in strategy and could have led to odd results such as differences in the effectiveness of SC inactivation based on which hemifield had been inactivated. Since the impact of SC inactivation on response accuracy appeared to be invariant to which quadrant had been inactivated (see Chapter 2) and did not differ between the animals, then we may reasonably conjecture that the SC contains circuits which are part of a common pathway in any strategy employed by the animal to solve the task.

Ambiguity in the appropriate timing of events could lead to different strategies. In our task, for instance, the monkeys are allowed to move their eyes the moment the choice dots appear. However the detection of the position of the eyes slightly lags the actual eye position, so it is possible that a saccade could begin prior to the appearance of the choice dots. In fact, it is conceivably possible that a monkey could learn to initiate his saccades upon motion detection, whereas another could mistakenly believe that he was required to wait until he had detected the appearance of the choice dots. Not only would this sort of difference place the two animals on very different positions on a speed-accuracy trade-off, it could potentially fundamentally change the nature of the task. An animal that initiates saccades at motion detection would solve the task by detecting the motion in the cued location first, in which case errors might be associated with early detection of the foil. Such an animal could possibly never notice that the cued and foil signal appeared simultaneously and may not gain exposure to the foil signal sufficient that he would learn that it is counter-informative. In contrast, a monkey that intentionally delays responses until after the choice dots observes signals for longer and has a consequently expanded opportunity to incorporate sensory evidence from the entire visual field. Moreover, he would have

greater exposure to both signals and would have the opportunity to incorporate that additional information in his strategy, and thus improve performance.

A less obvious but more insidious source of variation in strategy between different monkeys is the provision of alternative sources of information in the same task. From some perspectives, this could constitute a fairly substantial design flaw. In our task, the appropriate patch to select is conveyed in two ways—by both cueing and by blocking. As a result, the monkeys have two sources of information about where to attend. Since the two are perfectly correlated, knowing one does not contribute additional information about the cue location. Unless the monkeys are forgetful or lazy about noticing the cue, there is no possible benefit to employing both. It would seem prima facie to be easier for the monkeys to use the cue on each trial, but blocking trials with instructors is a commonly used method of cueing since monkeys can readily learn its meaning (e.g. Lee et al., 2007). We have no guarantee that the monkey uses one source of information over another, or that they use a mixed strategy, or that their strategy is constant in time, or that different monkeys have settled on the same strategy. Furthermore, it is possible that a monkey who has learned to use the block design to keep track of of the cue would be completely unaware of the meaning of the cues. What's worse, since using a color singleton as a cue or a block design as a cue would seem to involve very different strategies (visual discrimination versus short term memory between trials), and very different neural circuits, it is possible that variation between the strategies of different animals would lead to differences in the outcome of experiments in which a component of one of those neural circuits is selectively eliminated.

As we will see, the monkeys did show substantial differences in their reaction time distributions but not in their accuracy. This is likely because the only feature of their behavior which we attempted to shape was their accuracy. I will attempt to highlight points where we can deduce that aspects of their strategies were substantially divergent. In addition, I will attempt to determine critical points in task performance where their strategies were convergent. Since the results of SC inactivation were nearly

identical between the two animals, the points of convergence in their performance and strategy will be the most significant in identifying possible mechanisms for solving the task that involve circuitry with components in the SC.

4.2 Results

4.2.1 Accuracy and reaction time distributions

After completion of training on the task, which I took to be stabilization of accuracy for at least a week of training sessions, Subject F performed 39,835 trials in 76 sessions between February 28, 2007 and December 19, 2007. Subject M performed 26,914 trials in 80 sessions between January 1, 2008 and December 31, 2008. These numbers exclude both instructor trials and trials without signal. The target performance level was 65% to 75% correct. Accuracy (proportion of correct trials) and error types across all sessions for both subjects is shown in Figure 4.3A-B with the scatter of individual dots representing variability across sessions and the bars (with joint 95% multinomial confidence intervals) representing population mean performance. Sessions are pooled across all cue conditions. The dashed line indicates the proportion of trials which would be expected, on average, to be consistent with the foil signal if the errors were distributed uniformly. Since the proportion of trials consistent with the foil signal is substantially greater than this level, I infer that while some choices in the same direction as the foil could be unrelated coincidences, most errors are the result of an inappropriate intrusion of the foil signal on the decision process.

Pooled reaction time distributions (Figure 4.3C-D) indicate that mean reaction time depended on whether or not the monkeys based decisions on the presented signals. For Subject F, there was no statistical difference in the mean latency between correct choices and those consistent with the foil under a t-test corrected for multiple comparisons, while both were significantly faster than choices consistent with neither signal. For Subject M, the fastest responses were correct responses, followed

by choices with the foil and lastly by choices consistent with neither signal. These differences were statistically significant by t-tests corrected for multiple comparisons, and were far more pronounced in Subject M than in Subject F. Subject M also had markedly faster response times; recall that the choice dots appeared at 160 ms after the onset of the motion pulse, and so responses less than 200 ms after motion onset occur within 40 ms of the choice dot appearance. Therefore the fastest of Subject M's responses were likely initiated prior to the onset of the choice dots. In contrast, nearly all of Subject F's responses were over 150 ms after the appearance of the choice dots, which may indicate that the monkey initiated his response only upon the appearance of the choice dots.

When the accuracy distributions are divided by cue location as in Figure 4.4, asymmetries emerge. Although the pattern in accuracy was equivalent for all quadrants, that is, correct choices predominating and choices with the foil signal a substantially higher proportion of errors than expected, both Subject M and Subject F had asymmetries in accuracy across the visual field. Subject F was much better at selecting cued signals and ignoring foils when they appeared on the QI-QIII axis than if the cues and foils appeared on the QII-QIV axis (Figure 4.4A). In contrast, Subject M had higher accuracy when the cue appeared in the right hemifield (QI and QII) than when the cue appeared in the left hemifield (QIII and QIV). Although he was much more likely to base decisions on the cued signal regardless of the quadrant, a distinct rightward preference is apparent (Figure 4.4B). In contrast, Subject F's preference does not follow a strict spatial organization.

Reaction time distributions divided by cue location showed greater segregation based on the selected signal than when pooled across quadrants (Figure 4.5), but the patterns were somewhat inconsistent. Subject F (Figure 4.5A) had faster responses for correct responses when the cue appeared on the right (QI and QII), equal mean speed for the lower left (QIII), and faster responses with the foil for the upper right (QIV). In contrast, Subect M (Figure 4.5B) had a preference for the right hemifield even more pronounced than in the accuracy distributions. Correct choices

(based on the signal in that quadrant) occurred at a shorter latency when the cue appeared on the right (QI and QII) but occurred at a longer latency when the cue appeared on the left (QIII and QIV). Taken with the accuracy results, the response time distributions would seem to indicate that Subject M had a predilection, or bias, to select signals in the right hemifield. If he failed to suppress this tendency when the cue appeared in the left hemifield, he would quickly base decisions on the foil signal. If he could overcome this tendency, which he most often did, he would make longer latency responses based on the cued signal. When viewed in light of Subject M's distributions, Subject F could also have a predilection to select signals in the right hemifield and simply was better able to suppress this tendency, perhaps by adopting a different position on a speed-accuracy trade-off (Reddi and Carpenter, 2000; Reed, 1973).

Taken alone, accuracy distributions would seem to paint an incomplete picture of the animal's ability to select signals. When combined with the response time distributions, a more complete image emerges—both monkeys appeared to have a bias to select signals appearing in the right hemifield over signals appearing in the left. However they generally succeeded in overcoming this tendency, but only if the choices were delayed, as if the monkeys intentionally withheld responses in order to allow the signal in the left hemifield to overcome that in the right. The apparent lack of a spatial asymmetry in Subject F's accuracy but its presence in his reaction time distributions would seem to support the notion that he was better able to overcome the spatial asymmetry than was Subject M.

4.2.2 Spatial influence on response and reaction time distributions

In addition to spatial asymmetries in their tendency to select one signal over another, both monkeys had asymmetries in both their response distributions and reaction time distributions based on the cue position. In order to visualize the general dependence of response distributions and reaction time distributions on cue

location, I pooled data across cue locations by maintaining the geometric relationship between cue locations and response directions. The direction toward the cued location is referred to as *Cued*, the direction in the ipsilateral hemifield is *Ipsi*, the direction diametrically opposite the cued location is *Opp*, and the direction in the same vertical hemifield on the contralateral side is *Contra*. This relationship is shown schematically in Figure 4.6.

Both subjects tended to make fewer responses toward the cued location than to other locations. This could reflect the fact that monkeys were initially trained to perform motion discrimination with single patches and had to learn not to look directly at the patch. Hence from the outset they were trained to avoid making saccades to the peripheral and attended stimuli, and this tendency could have carried over into the response distribution. Subject F (Figure 4.6A) tended to distribute choices to the other locations fairly evenly, whereas Subject M (Figure 4.6B) tended to distribute choices to the ipsilateral target, almost as if the cue endowed him with a hemispheric preference that was countered by his tendency to avoid the cue only at that location.

Nevertheless, the shortest latency trials were in the same direction as the cued signal. Despite the fact that the monkeys tended to avoid orienting in the direction of the cued signals, the saccades made in those directions tended to be the fastest responses. As with the response distributions, Subject F (Figure 4.6C) did not have particularly strong preferences for any direction over any other (other than the cued direction), and these distributions were not significantly different from each other on t-tests corrected for multiple comparisons. In contrast, Subject M (Figure 4.6D) had a distinct preference for the ipsilateral target after the cued signal, and the mean of the distributions of response time for responses to the opposite hemifield were not different.

The trend toward shorter latency responses toward the cue location and overall bias away from the cued location did not cleanly apply to all cue locations. Subject F (Figure 4.7) tended to have a preponderance of choices directed towards

the inferior visual field (QII and QIII), and his response time distributions for those saccades had significantly shorter means than for saccades into the superior visual field. This tendency overwhelmed any dependence on cue location, although when the cue was in the inferior right visual field, responses to the cued quadrant had a marginally, but significant, shorter mean latency. Conversely, when the cue was in the superior visual field, responses toward the cued quadrant were the faster of the longer mean latency distributions.

In contrast, Subject M (Figure 4.8) appeared to have a balance between the trend to be biased away from the cued location and a spatial preference for choices. When the cue appeared in the right hemifield, Subject M was nearly entirely unbiased; however when the cue was in the left hemifield, Subject M tended to produce responses away from the upper right quadrant (one of his better quadrants for accuracy, Figure 4.4). His fastest responses were generally towards the cued location, except for the lower right quadrant, in which case there were no differences between the means of the distributions.

4.2.3 Interaction between spatial location and signal direction

Thus far we have seen that the response time distributions tend to have shorter mean latencies for correctly completed trials and for responses made toward the cued locations. Nevertheless, the differences in reaction time distributions related to cue location appear to be strongly influenced, if not outright overridden, by substantial spatial asymmetries. Furthermore, the pattern of accuracy and response time distributions would seem to indicate that the monkeys have a selection bias for signals in the right hemisphere. When those quadrants are cued, the decisions proceed rapidly and the correctly completed trials have shorter latencies than the trials in which the monkeys base decisions on the foils. When those quadrants are not cued, the monkeys must inhibit responses to allow sufficient time to overcome the selection bias. In that case the correct responses occur at a longer latency than the choices

with the foil signal.

An additional important question to ask is whether or not there is an interaction in reaction time between which signal is selected and the location of the cue. If the choices made toward the cued signal and based on the cued signal have the same distribution as choices made toward the foil signal and based on the foil signal, then this would indicate that the error arose as a failure in selection; if the monkey in some sense accidently intended to select the foil signal, then the synergistic interaction been the spatial location of the cue and the direction of the choice on reaction time would cause these trials to have the same reaction time distribution as trials in which the monkey correctly selected the cued signal and made his response in that direction. Conversely, responses made toward the cued signal and based on the foil should have longer latencies as if they were correct trials made toward the foil. If instead the reaction time distributions were still faster for correctly completed trials, then this would imply that the impact of reaction time of cueing carried through to the response even if the inappropriate signal had been selected, as if the monkey still intended to select the cued signal.

Pooled across quadrants, Subject F had a pattern of responses which appeared consistent with the model that choices with the foil signal occurred due to failures in the intended selection, whereas Subject M did not (Figure 4.9). There was no difference in the mean of the distributions for Subject F's responses correctly made toward the cue location and responses based on the foil signal made toward the foil signals' location (Figure 4.9A). Both occurred at shorter latencies than correct trials made toward the foil or choices made toward the cued location and based on the foil. In contrast, for Subject M (Figure 4.9B), the fastest responses were correct responses made toward the cued location, followed by responses made towards the foil based on the foil. Choices based on the cued signal towards the foil were no different on average than choices made toward the cued signal based on the foil. Such a mixed pattern would suggest that in Subject M, the impact of cueing on reaction time was somewhat separate from the impact of cueing on which signal was selected, whereas

in Subject F the impacts were not separate (or at least equivalent as far as can be distinguished with reaction time).

These patterns are not preserved, however, when the reaction time distributions are segregated by cue locations. In Subject F (Figure 4.10), the fastest responses are cued and correct responses directed into the inferior visual field (QII and QIII). Likewise when the cue was in the superior visual field, the fastest responses are those made into the inferior visual field, particularly if those responses were based on the foil signal. In Subject M (Figure 4.11), the fastest average responses were correctly completed responses toward the cued location in the left hemifield. In the right hemifield, there is no statistically significant difference in latency between correctly completed choices toward the cue and choices toward the foil and based on the foil signal.

Overall, the appearance in Subject F that choices made toward the foil and based on the foil signal had the same reaction time distribution as correct choices made towards the cue would seem to be almost entirely artifactual. Instead, reaction time seems to be primarily a function of whether or not the saccade is made toward a cued location and whether or not the cue is sympatric with the monkeys' preferred locations for both saccades and selection.

4.2.4 Conflict between cued and foil signal directions

A fundamentally important question to understanding why inactivation of the superior colliculus led to visual extinction for stimuli in the corresponding visual field is whether or not the monkeys normally pool motion signals across the visual field. On the one hand, monkeys could select one signal or the other and completely ignore whichever signals were not selected; in this case, the choice would be driven by the signal that had been selected, and no pooling would be performed. The variance in which signal the monkey selected to drive his choice would be explained primarily as variance in which signal had been selected, as if the monkey flipped a weighted coin on each trial to determine which signal he would select. In this case, differences in reaction time would indicate that a selection process occurred as a competition

between the stimuli and that whichever process finished first determined which signal would guide the decision. After the winner is determined, the monkey consults which direction of motion was present at that location and this direction determines the response. On the other hand, monkeys could pool sensory input across the visual field. In this case the competition is between the response directions rather than the stimuli, where each response direction receives contributions from the entire visual field (or, at the very least, both the cued and the foil signal). In this case, variance in behavior would be explained in terms of, possibly, variance in the weight applied, and by noise in the sensory stimulus leading to the foil signal direction occasionally dominating the cued signal direction.

Although a range of pooling schemes might be possible, any linear or non-linear pooling operation in which signals are combined would stand in contrast to a rule in which the signals were not combined. Whether or not pooling occurs leads to very different predictions about how choice and reaction time distributions might depend on the directions of the cued and foil signals. If no pooling occurs, then the choice and reaction time distributions should be the same regardless of the combination of directions in the cued and foil signals. If instead pooling occurs, then both the choice and reaction time distributions should depend on the combination of directions in these signals. It is conceivably possible that the animals could use a strategy in which each response direction and signal location are considered entirely separately, in which case signals are not pooled but instead each signal's contribution to a certain decision outcome is considered as a separate process from the other signal's contribution to the same outcome. In that case there would essentially be two ways to achieve each decision outcome, and this would be most apparent if the cued and foil signals coincided. In our design this case does not occur, and so as a result it would seem that we could not detect such a strategy if it were in place. However, my goal here is not to identify any particular strategy but instead to determine whether or not signals are pooled across the visual field. Therefore I will examine response and reaction time distributions based on the relationship between cued and foil signals in

order to rule out the absence of pooling. Moreover, the pattern of errors that emerges based on the relationship between these signals could inform the geometry of the perceptual space for motion discrimination (e.g. Ashby and Perrin, 1988; Bundesen, 1990).

Presuming that pooling is present, how should it change response and reaction time distributions? In order to answer this, we need to make some assumptions about how the signals might interact. First, sensory evidence in support of one direction is evidence against the others, but the strength of this counter-evidence may depend on the angular relationship between the options. In particular, evidence in one direction should be most antagonistic to evidence provided in the direct opposite direction, certainly more so than evidence from a signal pointing in the orthogonal direction. Second, the choice rendered is that corresponding to the maximum evidence combined with a signal-independent offset which adds as if it were evidence; this offset could correspond to such factors as reward expectations or biases. Third, the reaction time is determined in part by the length of time that the decision takes to complete, which is when either the evidence itself or a function on the evidence reaches some criteria. These assumptions are the basic propositions underlying most decision models currently in use (e.g. Gold and Shadlen, 2007; Palmer et al., 2005; Churchland et al., 2008; Britten et al., 1992; Ratcliff, 2006; Ratcliff and McKoon, 2008; Niwa and Ditterich, 2008; Usher and McClelland, 2001; McMillen and Holmes, 2006; Churchland et al., 2008; Jazayeri and Movshon, 2007).

Based on these notions of how pooling might occur, we can predict the interaction between cued and foil signals for the cases in which they are in direct opposition to each other and when they are orthogonal to each other. When the cued and foil signals are in direct opposition, both the cued and foil direction should be suppressed by the other, leading to the relative exaggeration of evidence for either of the orthogonal directions (which, presumably, also antagonize each other). This could lead to a lower than average proportion correct, and the errors would be expected to be more associated with neither of the signals than with the foil direction, which is

directly suppressed by the cued signal direction. Furthermore, since there is relatively less evidence for either the cued or foil signal directions, the decision criteria may take longer to be met, leading to overall longer reaction times for both of these choice outcomes. Choices with neither signal may be relatively faster since their outcome reaches the decision criteria more often and more quickly. When the cued and foil signals are orthogonal to each other, then the signals do not directly suppress one another. Evidence from the cued signal is not directly antagonized by evidence from the foil; instead, the foil signal contributes to the orthogonal directions. As a result, the errors should be more associated with the foil signal than with neither signal. Furthermore, since neither of the directions directly suppress one another, the decision criteria may be more rapidly met than when the signals are in direct opposition. Hence the choices with the cued and foil signal occur at shorter latencies, and choices with neither option occur at longer latencies.

Examination of the response distributions and reaction time distributions from Subject F (Figure 4.12) and Subject M (Figure 4.13) reveals the predicted pattern of response and reaction time distributions exactly as expected if the signals are pooled across the visual field. Pooled across quadrants and for each quadrant, the proportion of choices with the foil signal are greater when the signals are orthogonal than when they are antiparallel. In addition, the proportion of choices with neither signal is lower when the signals are orthogonal than when they are antiparallel. Since the trials are pooled across signal directions and each signal direction is paired with one antiparallel direction, all response directions are included equivalently in the distributions. Similarly, each signal direction is associated with a pair of orthogonal directions, which represent the same directions as the antiparallel case. Therefore we have averaged across the spatial asymmetries influencing response and reaction time distributions described above, which apply equally to both the antiparallel and orthogonal signal combinations.

The differences between the orthogonal and antiparallel signal combinations are summarized in Figure 4.14, which more clearly indicates the relationship in popu-

lation average response distributions and response time distributions, sorted by quadrant, between the conditions. Both subjects had the greater proportion of choices with the foil signal for the orthogonal case than the antiparallel case, and both had the lesser proportion of choices with neither signal for the orthogonal case than the antiparallel case. Subject F had little difference in proportion correct, but Subject M had a significantly greater proportion correct in all but one quadrant when the signals were orthogonal. Both subjects had slightly shorter responses when the signals were orthogonal than when they were antiparallel for correct and choices with the foil. In Subject M, the choices with neither signal occurred at a systematically longer latency.

Overall, the systematic variation of response and reaction time distributions based on whether or not the cued and foil signals were orthogonal or antiparallel strongly supports the notion that evidence is accumulated as a weighted sum of the signals. Therefore inactivation of the SC would remove one component of this sum, allowing the other to fully penetrate.

4.2.5 Impact of delay on accuracy and reaction time

The monkeys' performance on the task may have depended on the delay between the cue offset and the presentation of the coherent motion stimulus. We used a delay between cue offset and stimulus presentation that had a fixed portion of 480 ms and a random portion of mean 480 ms. The random portion followed a geometric distribution, so that the hazard function after 480 ms was flat and the animal could not predict when the signals were about to appear. This choice was motivated by concerns over anticipation and transient shifts attention induced by the visual cues.

Predictable task timing influences neural activity related to both attention and oculomotor planning. When stimulus events are predictable, attentional modulation of the activity of neurons in macaque visual cortex area V4 occurs in close association with the timing of the stimulus appearance (Ghose and Maunsell, 2002); similarly, neurons in macaque parietal cortex increase their firing rate in anticipation

of stimuli when the timing is predictable (Janssen and Shadlen, 2005). Since both attention-related and oculomotor-related activity in cortical areas that project to the SC are so dependent on the timing of task events, it stands to reason that these same factors could influence responses in the SC. Since these features of the response are not the focus of the current study, we elected to attempt to eliminate them as confounding factors by making the timing entirely unpredictable (i.e. flat hazard function). It is possible that monkeys could have an erroneous model of the task timing, but if monkeys are so acutely aware of the timing distribution that their neural activity is closely matched to task timing, then it would seem that such an erroneous strategy would be disadvantageous and that the monkey should soon dispense with it. In any event, when their is no predictability in the stimulus timing, any correlation of activity with an erroneous timing model would be completely invisible to us unless we had some independent measure of the monkey's temporal expectation.

We were also concerned that the visual cues would transiently attract the monkey's attention to the region with a color singleton, but that the animals would not continue to attend to that location (e.g. Nakayama and Mackeben, 1989), and so delaying stimulus presentation by a long and variable period of time would necessitate the animals to continue to attend to the cued stimuli even in the absence of visual cues. If the animal's attention was drawn simply based on the visual transient, then the animals should selectively attend to the cued stimulus only momentarily before reverting to either a non-selective strategy or a natural bias. Since the distribution of delays is geometric, it would place most of the trials in a short interval of time after cue offset, and this would reduce our ability to determine that the monkey was not transiently attending. Therefore we added a fixed delay longer than the typical transient shift in attention. Here we evaluate the success of this design in the production of a steady-state behavior in which the animals attend to the cued stimuli in the absence of visual cues.

Both monkeys maintained performance long after the offset of the visual cues (Figure 4.15A-B). Although there was a decline in performance over the course

of seconds, both monkeys were able to successfully maintain their accuracy at a high level for the initial portion of the delay distribution, which encompassed a substantial majority of trials. Furthermore, when performance did decline, it was in terms of a trade from the proportion of correctly completed trials to that associated with the foil signal, rather than a substantial increase in the proportion associated with neither signal. Decreases in vigilance or motivation might have been associated with an overall tendency to transfer choices from signal driven to guesses; however, this did not occur. Therefore the basic approach of the animals did not seem to change—the responses are associated with the presented signals, but the ability of the animals to base their decision on the appropriate signal seems to have declined for the longest trials. Such a change would be consistent with a decrease in the animal's ability to asymmetrically apply weighting to the signals in the visual field.

A similar pattern is observed for the reaction time distributions (Figure 4.15C-D and E-F). The scatter plots are shown with linear regression fits and 95% prediction bounds. The slopes of the fits, while small, are significantly greater than zero, indicating that the longer the delay, the longer the reaction time even for correctly completed trials. Parameters of the fit with 95% confidence intervals are shown in Table 4.1. If the weights applied to the signals had remained constant, then the reaction time distributions should have the same mean independently of the delay. Instead, the delays increased, which would be caused by a decrease in the weight on the cued signal relative to that on the foil signal.

Table 4.1: Slope and intercept for linear regression of reaction time on delay time.

	Subject F	Subject M
Correct Trials		
slope (ms/ms) $\times 10^{-3}$	3.934 (2.925,4.942)	10.38 (9.036,11.72)
intercept (ms)	345.3 (344.2,346.4)	279.2 (277.8,280.6)
Foiled Trials		
slope (ms/ms) $\times 10^{-3}$	3.62 (1.925,5.314)	5.705 (3.633,7.777)
intercept (ms)	347.3 (345.4,349.2)	297.7 (295.3,300)

4.3 Conclusions

The objective of this section was to inductively derive qualitative assessments of the factors influencing performance on the selective attention task and the strategies that monkeys might have used to solve the task. After examination of accuracy, response, and reaction time distributions, I arrive at the following six conclusions:

1. Monkeys were able to base decisions on cued signals while ignoring foil signals, although most errors appeared to be driven by the foil signal.

2. Both monkeys had a bias for basing decisions on signals in the right hemifield, and choices based on signals in the right hemifield occurred at a shorter mean latency than those in the left hemifield regardless of their direction.

3. Subject F's responses had longer average latency than Subject M's by an amount consistent with the speculation that Subject F waited for the appearance of the choice dots to initiate responses, whereas Subject M may have initiated responses upon detection of the motion pulses.

4. The response distributions for both monkeys were biased away from the cued position. Responses toward the cue tended to have the shortest mean latency, although this pattern was obscured by a competing tendency for saccades made into the inferior visual field to have shorter mean latency than those made into the superior visual field.

5. The response and reaction time distributions systematically varied based on whether cued and foil signals were orthogonal or antiparallel. The variation was consistent with a strategy in which sensory evidence is pooled across the visual field. This observation almost surely means that monkeys pooled sensory evidence across the visual field in what appears to be a weighted sum, rather than strictly selecting one stimulus to the exclusion of another.

6. Performance on the task was not substantially dependent on the duration of the delay, indicating that the monkeys did not merely transiently attend to the cued stimuli in response to the visual cues. Nevertheless, accuracy did fall and response time did increase as if the animals had at least slight difficulty maintaining asymmetric weighting across the visual field for very long delays.

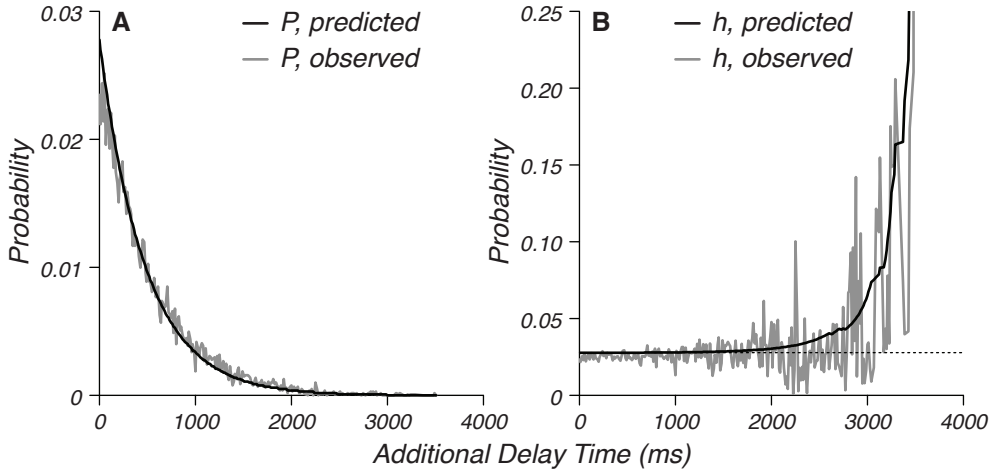

Figure 4.1: Probability mass function and hazard function for randomized delay. A) Probability mass function predicted (black) and observed (gray) of delay time. B) Hazard function predicted (black) and observed (gray) of delay time. Theoretical prediction is shown in dashed line.

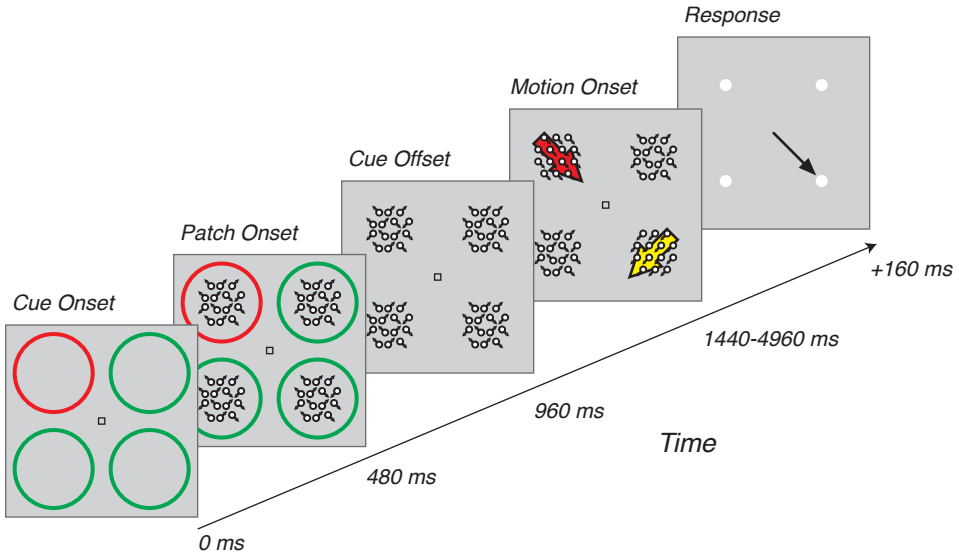

Figure 4.2: Selective attention task. After a brief fixation period, colored cue rings were presented (*Cue Onset*). Stochastic motion patches appeared next (*Patch Onset*), and then the cues disappeared (*Cue Offset*). Following a delay, brief coherent motion pulses occurred (*Motion Onset*) in both the cued location (red arrow) and the diametrically opposite location (yellow arrow). Monkeys reported the direction of the cued motion signal by making an eye movement to a response dot in the same direction as the motion pulse.

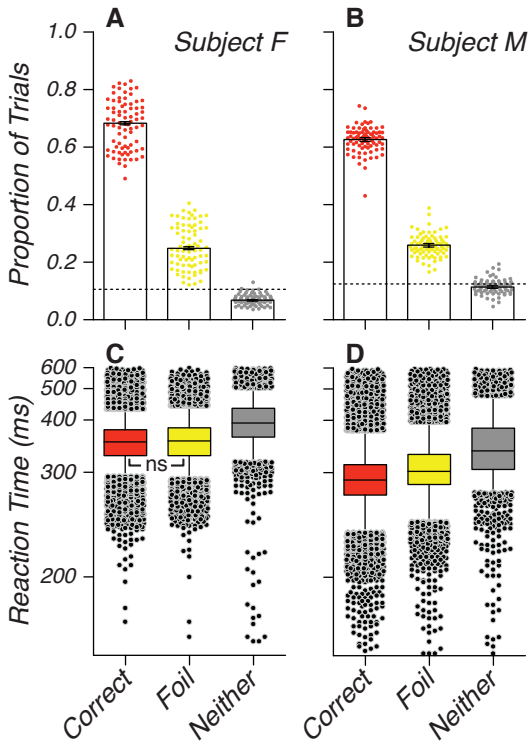

Figure 4.3: Accuracy and reaction time in behavioral sessions conducted in the months prior to and concurrent with the recording and inactivation experiments. Bars indicate population average and error bars indicate 95% multinomial confidence intervals. Dashed line represents the proportion of choices which would be expected to coincide with the foil signal if errors were uniformly distributed. C-D) Summary of reaction time for Subjects F and M. Whisker plots represent the median reaction time with upper and lower quartile. Black dots symbols are outliers exceeding the 95% range of the distributions. Reaction time distributions are plotted on a reciprocal latency scale. Population mean latencies for correct and foiled trials for Subject F were not significantly different, indicated by the bracket and designation ns.

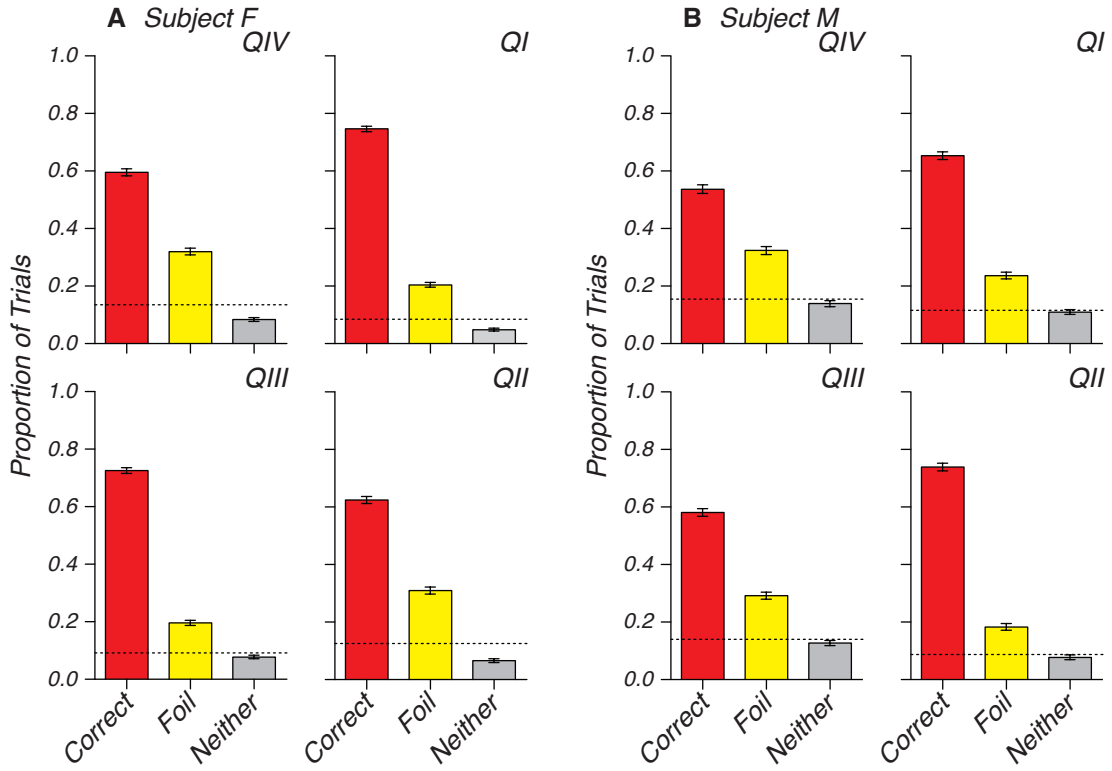

Figure 4.4: Accuracy distributions sorted by cue position for Subject F (A) Subject M (B). Cue locations are designated clockwise from upper right quadrant as QI, QII, QIII, and QIV. Color conventions as in Figure 4.3

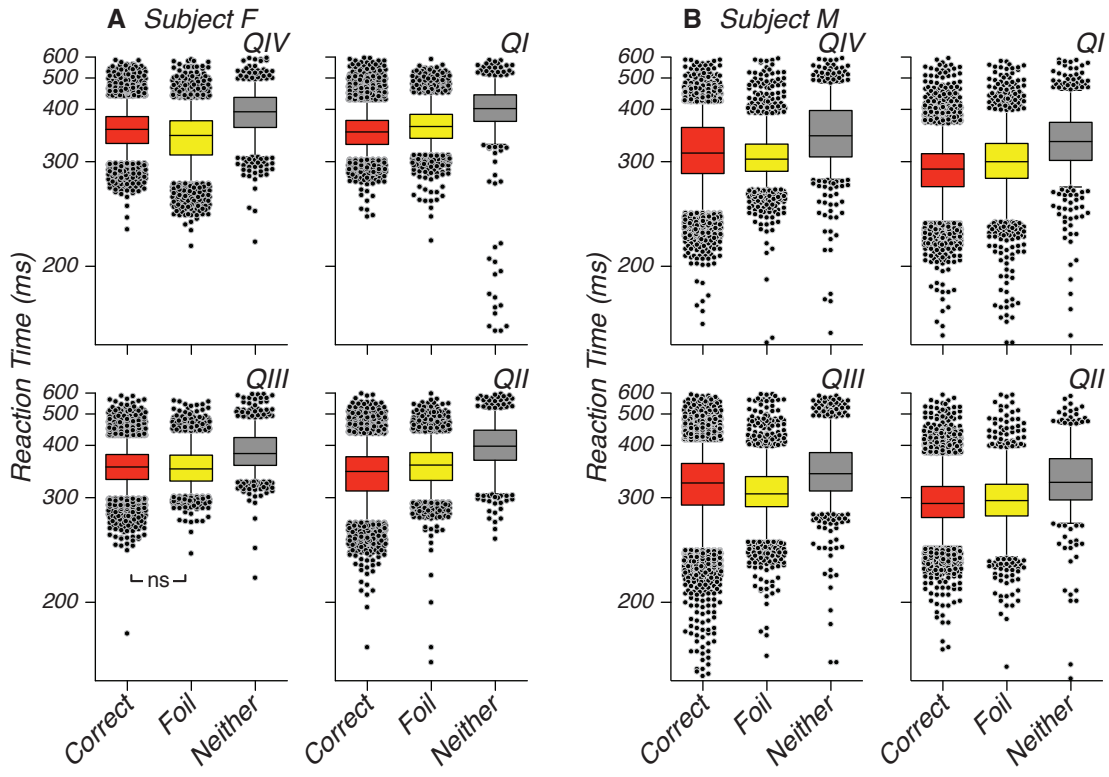

Figure 4.5: Reaction time distributions sorted by cue position for Subject F (A) and Subject M (B). Cue locations are designated clockwise from upper right quadrant as QI, QII, QIII, and QIV. Whisker plots represent the median reaction time with upper and lower quartile. Black dots symbols are outliers exceeding the 95% range of the distributions. Reaction time distributions are plotted on a reciprocal latency scale. Color conventions as in Figure 4.3

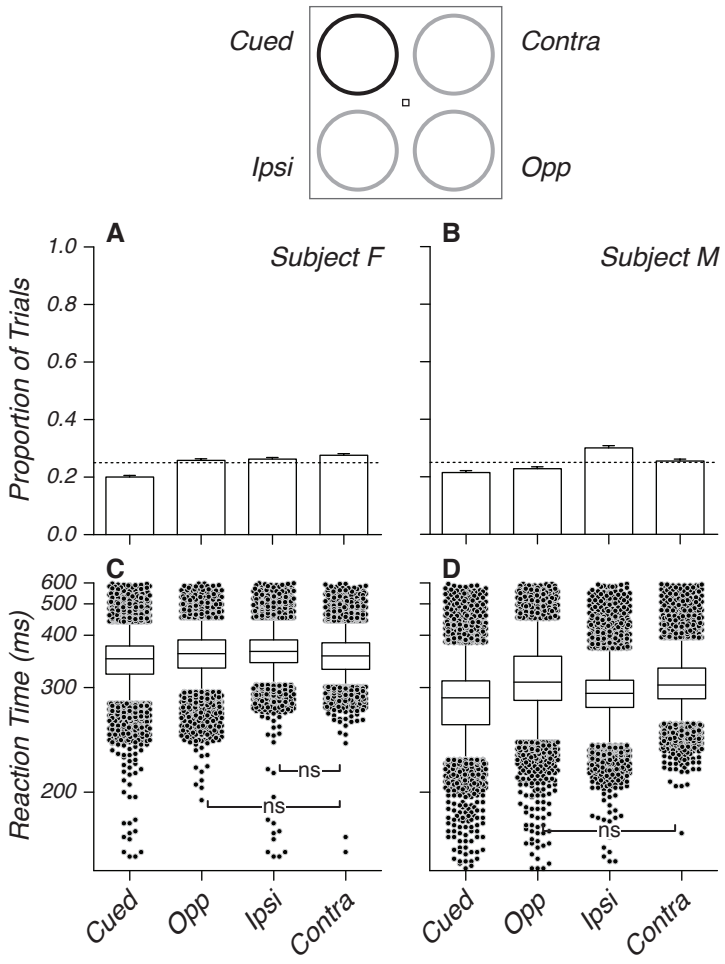

Figure 4.6: Response and reaction time distributions by direction. A-B) Distribution of responses independent of signal direction and cue location for Subjects F and M. C-D) Whisker plots of response time distributions sorted by direction for Subjects F and M. Whisker plots represent the median reaction time with upper and lower quartile. Black dots symbols are outliers exceeding the 95% range of the distributions. Reaction time distributions are plotted on a reciprocal latency scale.

Figure 4.7 *(following page)*: Response (A) and reaction time distributions (B) by cue location for Subject F. Cue locations are designated clockwise from upper right quadrant as QI, QII, QIII, and QIV. Bars for cue locations are shaded gray. Whisker plots represent the median reaction time with upper and lower quartile. Black dots symbols are outliers exceeding the 95% range of the distributions. Reaction time distributions are plotted on a reciprocal latency scale.

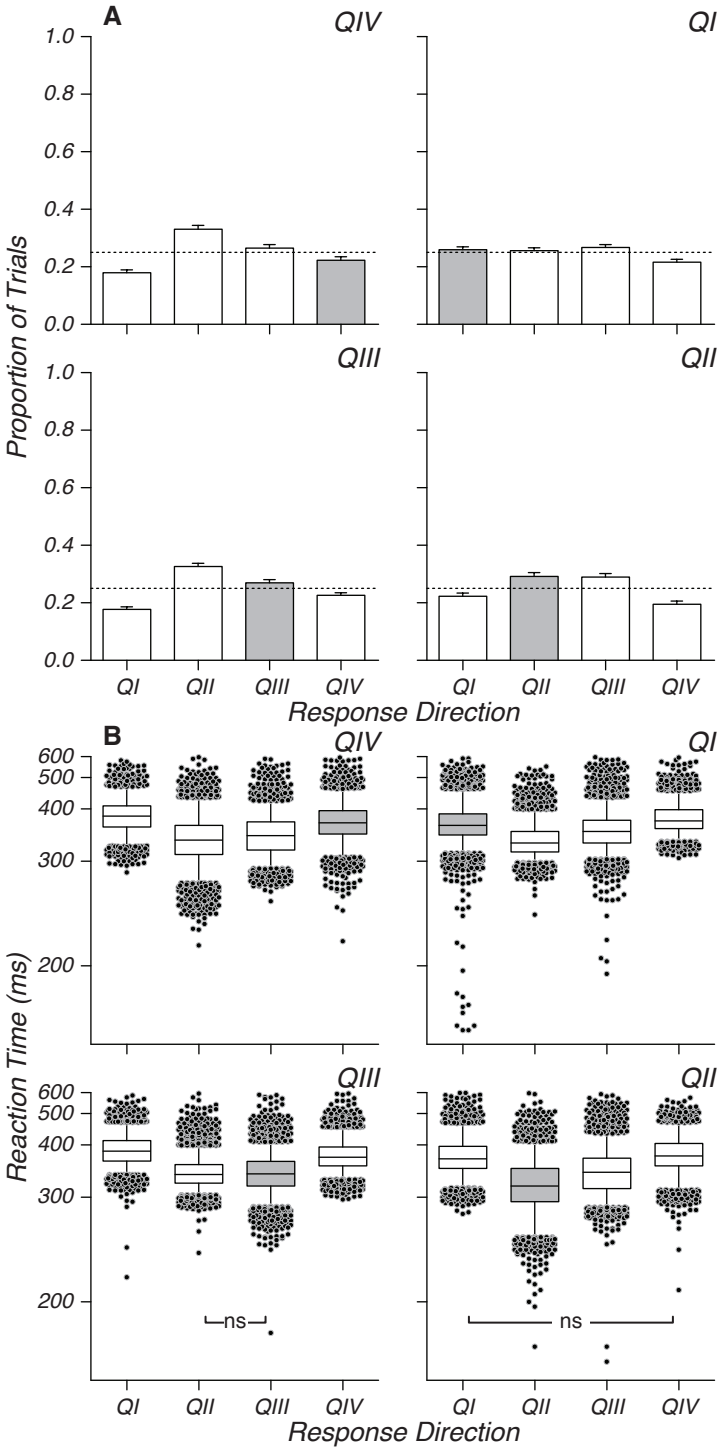
A
QIV
QI
QIII
QII
Proportion of Trials
1.0
0.8
0.6
0.4
0.2
0.0
QI
QII
QIII
QIV
Response Direction
B
Reaction Time (ms)
600
500
400
300
200
ns
ns

Figure 4.8 *(following page)*: Response (A) and reaction time distributions (B) by cue location for Subject M. Cue locations are designated clockwise from upper right quadrant as QI, QII, QIII, and QIV. Bars for cue locations are shaded gray. Whisker plots represent the median reaction time with upper and lower quartile. Black dots symbols are outliers exceeding the 95% range of the distributions. Reaction time distributions are plotted on a reciprocal latency scale.

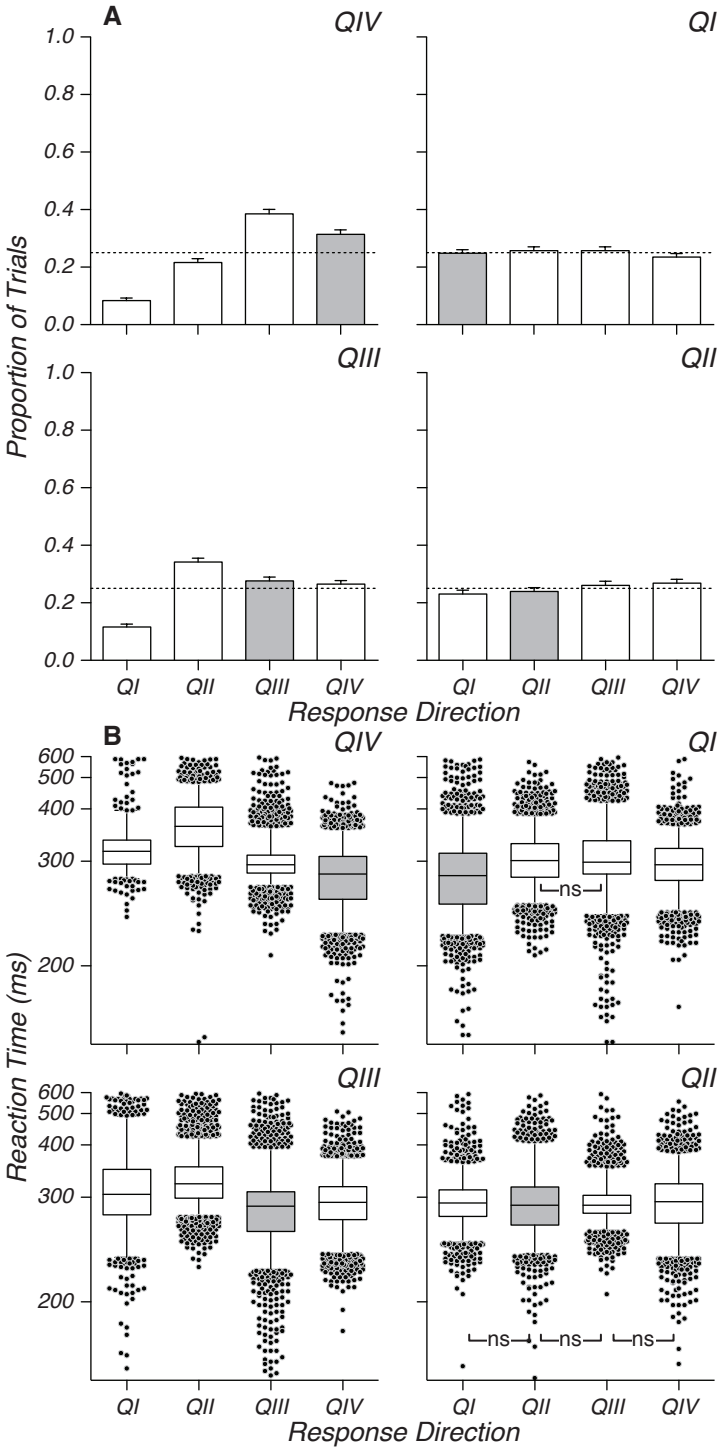
A
QIV
QI
QIII
QII
1.0
0.8
0.6
0.4
0.2
0.0
Proportion of Trials
QI
QII
QIII
QIV
Response Direction
B
600
500
400
300
200
Reaction Time (ms)
ns
ns
ns
ns

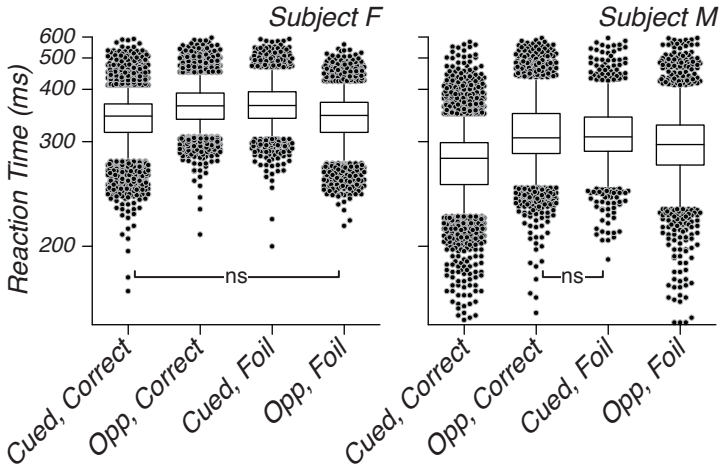

Figure 4.9: Reaction time distributions by signal direction relative to cue position for Subjects F (left) and M (right). Combinations shown are for correct responses made towards the location of the cued signal (Cued, Correct), correct choices made directly away from the location of the cued signal (Opp, Correct), choices based on the foil signal and made toward the location of the cued signal (Cued, Foil), and choices based on the foil signal and made toward the foil signal (Opp, Foil). Whisker plots represent the median reaction time with upper and lower quartile. Black dots symbols are outliers exceeding the 95% range of the distributions. Reaction time distributions are plotted on a reciprocal latency scale.

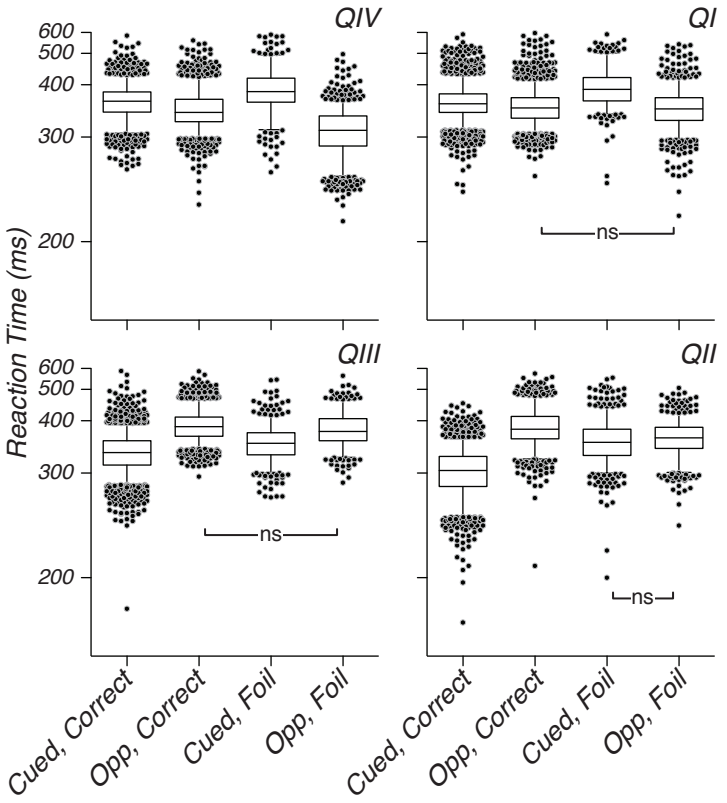

Figure 4.10: Reaction time distributions by signal direction relative to cue position for Subject F. Cue locations are designated clockwise from upper right quadrant as QI, QII, QIII, and QIV. Combinations shown are for correct responses made towards the location of the cued signal (Cued, Correct), correct choices made directly away from the location of the cued signal (Opp, Correct), choices based on the foil signal and made toward the location of the cued signal (Cued, Foil), and choices based on the foil signal and made toward the foil signal (Opp, Foil). Whisker plots represent the median reaction time with upper and lower quartile. Black dots symbols are outliers exceeding the 95% range of the distributions. Reaction time distributions are plotted on a reciprocal latency scale.

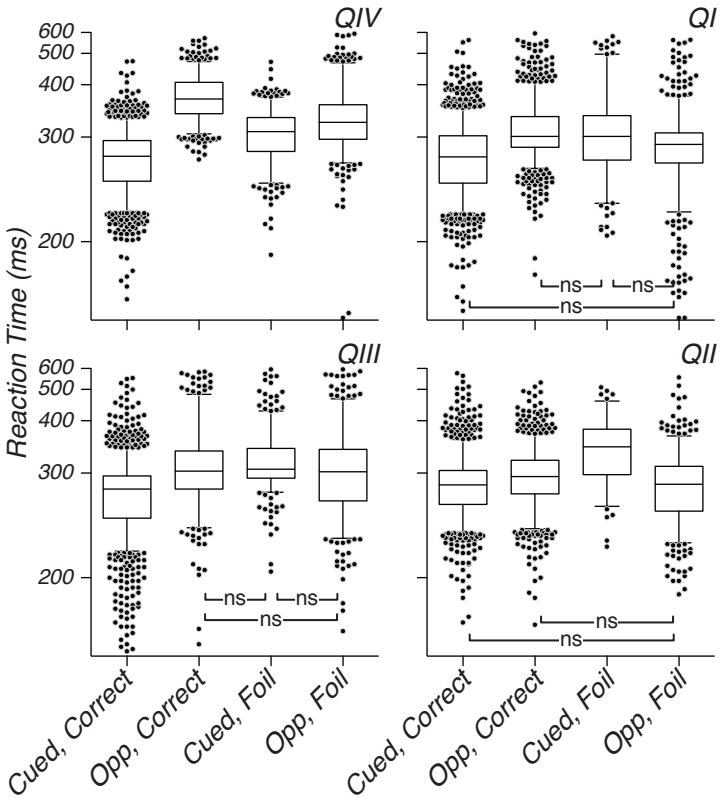

Figure 4.11: Reaction time distributions by signal direction relative to cue position for Subject M. Cue locations are designated clockwise from upper right quadrant as QI, QII, QIII, and QIV. Combinations shown are for correct responses made towards the location of the cued signal (Cued, Correct), correct choices made directly away from the location of the cued signal (Opp, Correct), choices based on the foil signal and made toward the location of the cued signal (Cued, Foil), and choices based on the foil signal and made toward the foil signal (Opp, Foil). Whisker plots represent the median reaction time with upper and lower quartile. Black dots symbols are outliers exceeding the 95% range of the distributions. Reaction time distributions are plotted on a reciprocal latency scale.

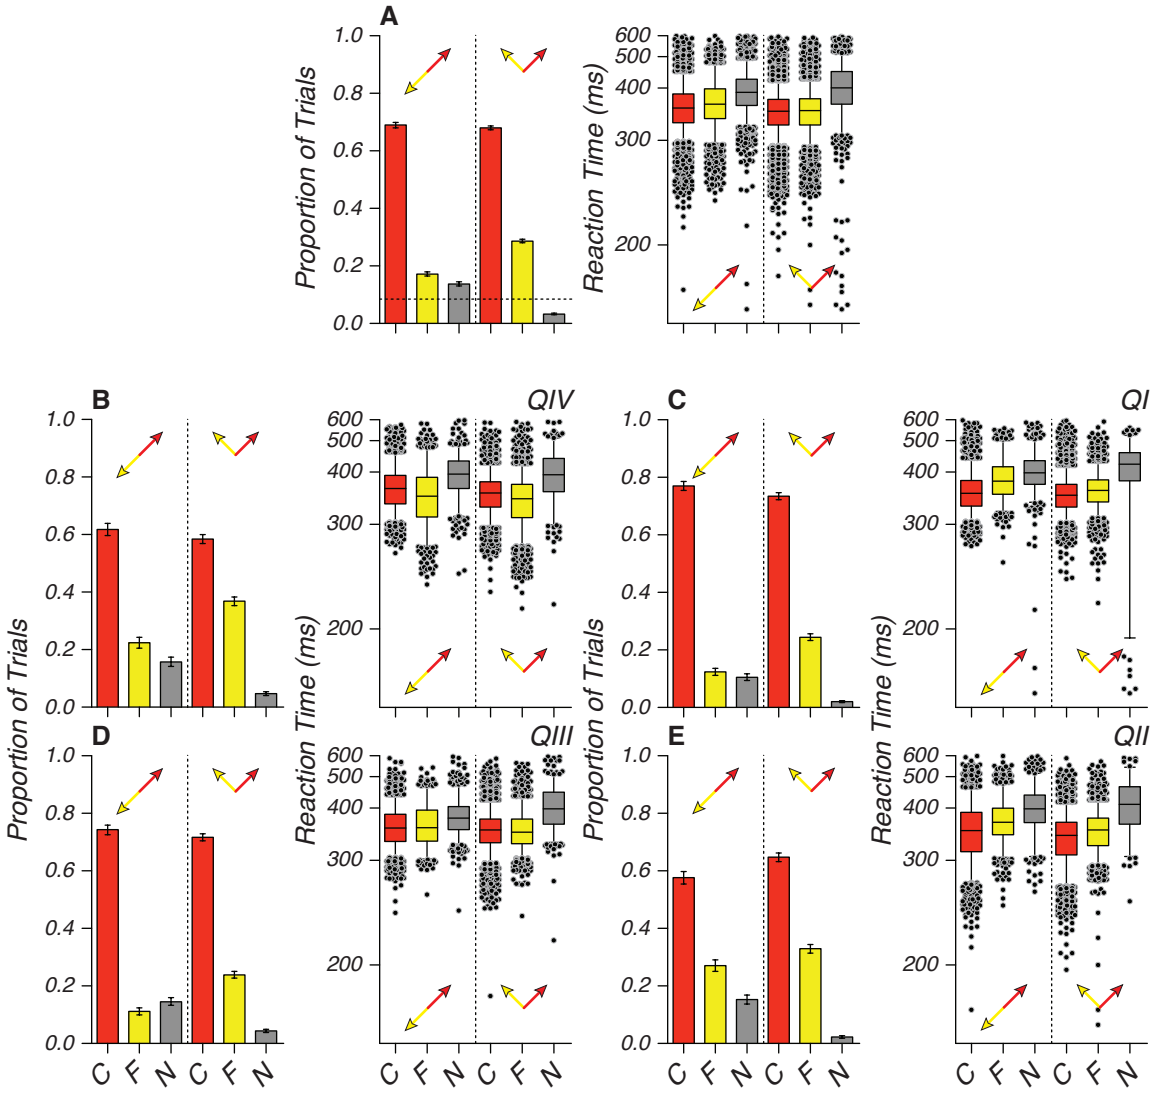

Figure 4.12: Accuracy and reaction time distributions reflecting conflict between cued and foil signals for Subject F. Data pooled across all quadrants (A) and separated by quadrants (B-D). Columns correspond to correct trials (C), trials consistent with the foil signal (F), and trials consistent with neither signal (N). Whisker plots represent the median reaction time with upper and lower quartile. Black dots symbols are outliers exceeding the 95% range of the distributions. Reaction time distributions are plotted on a reciprocal latency scale.

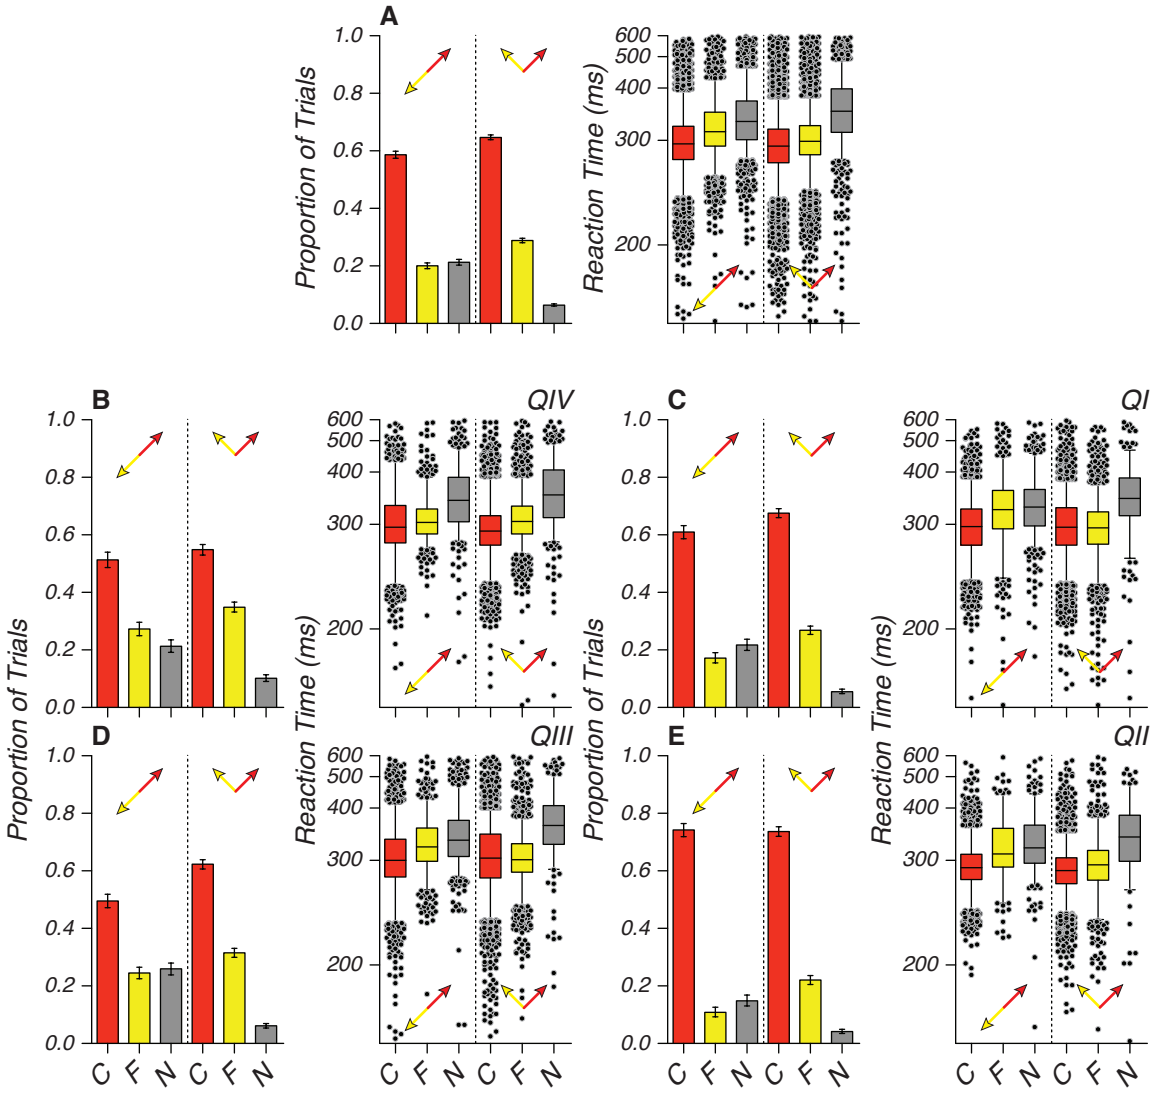

Figure 4.13: Accuracy and reaction time distributions reflecting conflict between cued and foil signals for Subject M. Data pooled across all quadrants (A) and separated by quadrants (B-D). Columns correspond to correct trials (C), trials consistent with the foil signal (F), and trials consistent with neither signal (N). Whisker plots represent the median reaction time with upper and lower quartile. Black dots symbols are outliers exceeding the 95% range of the distributions. Reaction time distributions are plotted on a reciprocal latency scale.

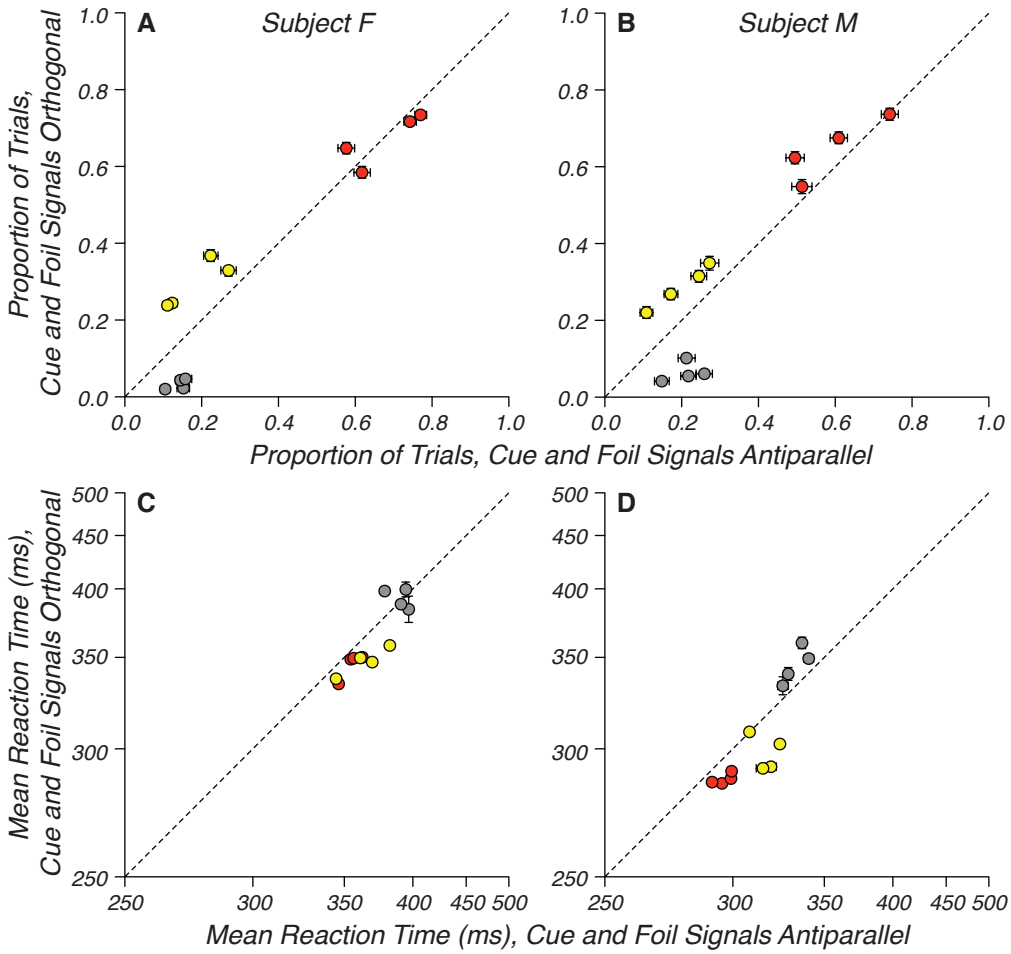

Figure 4.14: Population average of accuracy and reaction time distributions by quadrant. Red symbols are correct choices, yellow choices with the foil, and gray choices with neither signal. Error bars are 95% confidence intervals on population averages. Population mean of accuracy distributions for Subject F (A) and Subject M (B) are plotted with the proportion of trials for the case in which the cue and foil signals are orthogonal against the corresponding case in which the cue and foil signals are antiparallel. Analogous plots for reaction time in (C) and (D).

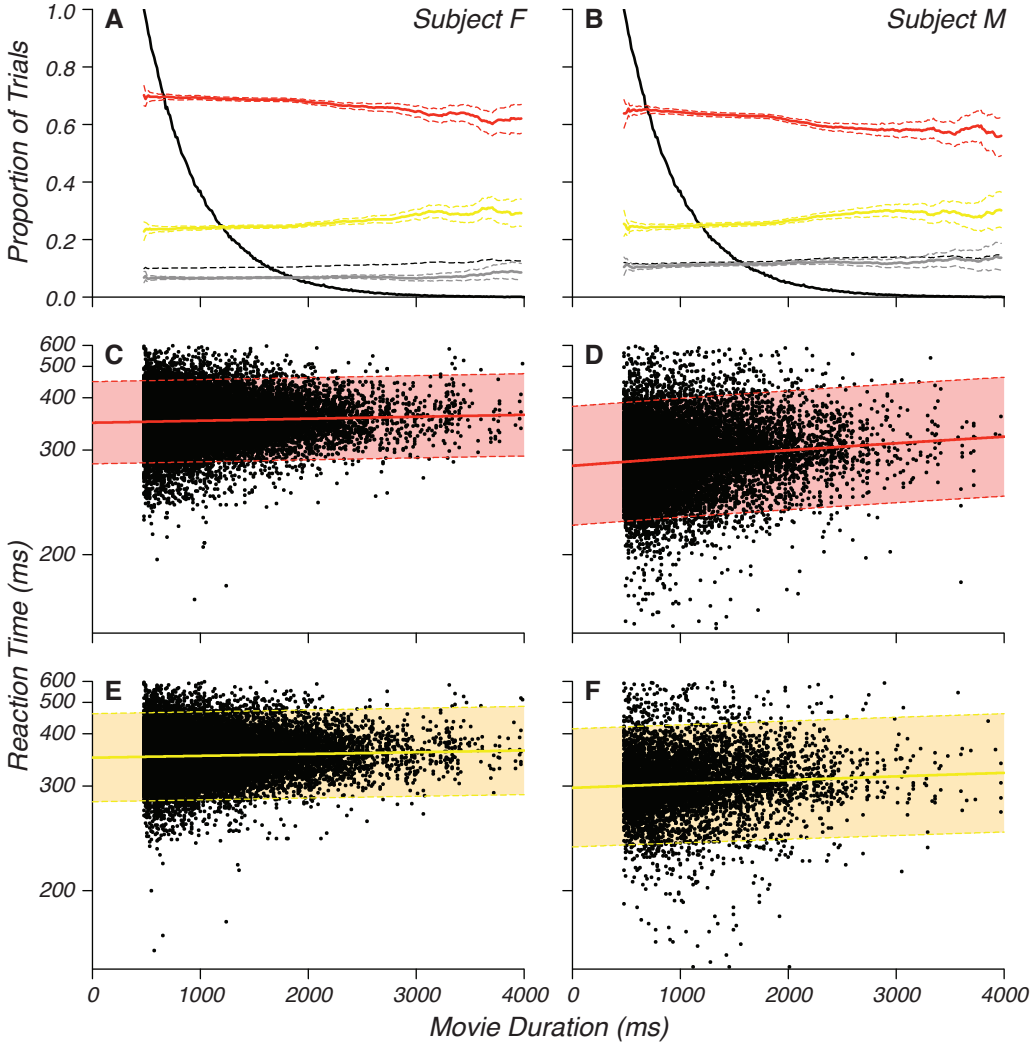

Figure 4.15: Accuracy and reaction time as a function of movie duration. Accuracy for Subjects F and M (A-B) as a function of movie duration. Red lines represent the proportion of correct trials, yellow lines represent the proportion of choices consistent with the foil signal, and gray lines represent the proportion of trials consistent with neither signal. the dashed line represents the proportion of trials expected to be consistent with the foil if the errors were uniformly distributed. Black line is empirical distribution of trial durations. Reaction time as a function of movie duration for Subjects F and M for correct responses (C-D) and choices with the foil signal (E-F). Black symbols represent individual trials. Solid lines are the result of linear regression, and error lines and filled region represents the 95% non-simultaneous prediction bounds for the regression.

Chapter 5

Recording and Classification of Single-Units in the Superior Colliculus

5.1 Introduction

With some notable exceptions focusing on arm-movement related activity (e.g. Werner et al., 1997b,a; Nagy et al., 2006) or multi-sensory input (Wallace et al., 1996; Groh and Sparks, 1996a,b), most studies of the primate superior colliculus (SC) have focused on the circuitry related to the control of saccadic eye movements in response to visual stimuli. This work has sought to understand visual-motor guidance under the general theory that the circuitry in the SC provides an architecture for the visuo-motor transformation (reviewed by Wurtz and Albano (1980)). The basic notion is that there is a horizontal retinotopic organization, but the vertical organization is a matter of debate (see Isa and Hall (2009) for a discussion and connections to the cellular histology and circuitry of the superior colliculus). One influential proposal for the vertical organization, known as the visual-movement convergence model (Mohler and Wurtz, 1976), holds that direct retinal input and visual input from striate cortex enters the colliculus by way of the superficial layers and descends to meet motor,

context, and mixed sensory input ascending through the deep layers at a point of convergence. An alternative is that individual layers have separate functional efferents. In any case, since the discourse on the primate superior colliculus focuses on the visuo-motor transformation, the descriptive language we use to identify neurons is an assessment of their responses to both visual stimuli and the saccades made to those stimuli.

When recording in the primate superior colliculus, we first enter through the superficial layers from the quadrigeminal cistern (see the schematic representation in Chapter 2 Figure 2.9). The superficial layers are often taken to be the stratum zonale, stratum griseum superficiale (SGS), and stratum opticum (SO) (Wurtz and Albano, 1980). Neurons in the superficial layers receive direct retinal inputs and respond, to the best of our knowledge, exclusively to visual input and not to saccades. Neurons in these layers tend to have responses to visual transients and moving stimuli within circumscribed receptive fields (e.g. Goldberg and Wurtz, 1972a). These receptive fields are retinotopically organized; each SC represents mostly contralateral visual space with the foveal region represented at the rostral pole of the SC and the periphery caudally. The horizontal meridian runs down the midline with temporal targets represented laterally and minimal ipsilateral representation at the medial margin.

After passing through the superficial layers we enter the deep layers, notable for neurons with saccade related activity. The deep layers are divided into the stratum griseum intermediate (SGI), the startum album intermediate, the stratum griseum profundum (SGP), and the stratum album profundum, which for our purposes are lumped under the rubric intermediate and deep layers. Some neurons in these layers have *receptive* fields for visual transients (Wurtz and Goldberg, 1972; Mohler and Wurtz, 1976; Wurtz and Albano, 1980) in addition to a variety of other sensory inputs (Wallace et al., 1996), but are most marked by their responses to saccades. These neurons begin firing prior to the onset of the saccade and discharge in response to saccades made to specific locations in space (called interchangeably either the *movement* or *response* field to distinguish it from the visual receptive field). The

movement fields are retinotopically organized and in rough register with the map in the superficial layers. The receptive fields of the neurons are generally not coterminous with the movement field; instead, they correspond to a similar region of space but generally cover less of an area.

In order to identify the response properties of neurons in the SC, investigators have used a combination of the responses of neurons to spontaneous saccades, made either in an illuminated or darkened room, and to saccades made during simple saccade tasks. The most simple of these is the visually-guided saccade task, in which a monkey first fixates a bright spot of light and then makes a saccade to a new point of light, which appears at the same time that the spot at fixation disappears. Several variations of the basic visually-guided saccade task have been used to investigate the properties of saccade-related discharges. These include the delayed visually-guided saccade task, in which the fixation dot persists after the appearance of the target and the monkey withholds saccades until the fixation point disappears, and the double-step saccade task, in which two targets briefly appear in short order and the monkey makes saccades to those two locations (e.g. Mays and Sparks, 1980). These tasks separate the onset of the target from the later generation of the saccade and so isolate visual responses from saccade-related discharge. The double-step saccade requires the monkey to remember the second saccade location and so bears some similarity to a more contemporary saccade task, the memory-guided saccade. In memory-guided saccades, the monkey must remember the location of a briefly presented target dot and make the saccade only after the fixation dot is extinguished. Introduced by Hikosaka and Wurtz (1983b), this task is commonly used in neuron classification despite the fact that most of the categories were established prior to its introduction, and to the best of my knowledge, no effort has been made to systematically determine whether or not it would produce the same classification scheme as the now uncommonly used double-step paradigm.

Following Mays and Sparks (1980) and based on these tasks, neurons in the SC are generally classified as being either *visual* or *saccade related* (SR). Visual

neurons have little or no response to the saccade in a delayed visually-guided saccade task and are usually found in the superficial layers. Saccade related neurons begin discharge in anticipation of the saccade and produce bursts aligned to saccade onset, but produce no or minimal responses following the onset of a visual target. Other neurons are called *visuo-motor* (VM) if they produce both visual and saccade related activity. Based on the double-step saccade task, Mays and Sparks (1980) also identified a group of cells called *quasi-visual* (QV) which maintained a sustained discharge corresponding to the second target location almost as if they held a memory of that second location, having been identified by a visual target, without producing a discrete saccade burst.

Cells with saccade-related discharge are further subdivided. While the saccade-related discharge begins around 50 ms prior to the saccade in many cells, in others it begins hundreds of milliseconds prior, earning these the prepended designation *long-lead*. When these neurons produce discrete bursts at saccade onset, they are known as long-lead burst neurons (LLBN). The initiation of the saccade discharge depends on when the monkey can determine that a saccade will be made into the neuron's response field. This initiation occurs even when there are multiple potential saccade goals, but only when the eventual saccade is made into the neuron's response field. As a result, this group of neurons has long been suspected of having some role in the more general issue of target selection (Wurtz and Albano, 1980).

A specific subset of long-lead burst neurons associated with the selection of targets for saccades have been identified. These cells, known as *buildup* neurons(Munoz and Wurtz, 1995), have an increasing firing rate leading up to the initiation of the saccade. The buildup of activity is predictive of the eventual selection for the saccade target. For example, the firing rates of buildup neurons prior to saccades are proportional to the probability that the monkey will make a saccade to that location (Basso and Wurtz, 1997, 1998; Dorris and Munoz, 1998), in particular when the likelihood is determined by context other than the number of targets on the screen, thus ruling out lateral inhibition from the increasing number of visual targets.

Buildup neurons have activity which is predictive of which saccade will be made long before the initiation, and the firing rate appears to be related to the accumulation of evidence as in a diffusion process (Ratcliff et al., 2003, 2007); and the initiation of that saccade appears to be related to when that representation of evidence reaches a threshold (Paré and Hanes, 2003). Moreover, the neural responses of buildup neurons associated with target selection, as in a motion discrimination task, appear to be separable from the responses associated with actual planning of the saccade and specification of its metrics, which appear to be the domain of saccade-related burst neurons (Horwitz and Newsome, 1999, 2001a,b). In addition, the eventual selection can be decoded from single unit data recorded from buildup units (Krauzlis and Dill, 2002; Mcpeek and Keller, 2002). Finally, based in part on the intuition that buildup units would be involved in target selection, they are also one of the primary candidates for selection of stimuli for visual attention, and have been shown to have changes in activity correlated with both the preparation of saccades and with covert attention (Kustov and Robinson, 1996).

Although the association with target selection would seem to make the buildup neuron group an important functional classification, it does lead to some ambiguity when applied to neurons which might also be called visuo-motor. The buildup is a property of the discharge preceding the eventual saccade, and some buildup neurons have visual responses while others do not (among other widely varying properties for these neurons) (Munoz and Wurtz, 1995). Therefore some buildup neurons may be considered to be visuo-motor neurons, whereas others may not. The presence or absence of visual responses does not appear to be a necessarily important factor when considering the pre-saccadic target selection related activity. Nonetheless, the designation of visuo-motor is still in use, and many of the properties specifically attributed to buildup neurons have been studied and documented in visuo-motor neurons without concern for the buildup designation, such as target selection related activity (Mcpeek and Keller, 2002, who referred to them as visuo-motor units with prelude activity) and covert attention (Ignashchenkova et al., 2004; Fecteau et al.,

2004).

An additional group of neurons which is ambiguously related to visuo-motor neurons are the *visually-triggered movement* (VTM) cells. These are neurons with saccade-related discharges that require a visual target in order to respond. These cells are observed directly underneath stratum opticum (Mohler and Wurtz, 1976; Mays and Sparks, 1980) and in a darkened room do not produce discharges for spontaneous saccades or the second of a double-step saccade; instead the saccade-related discharge only occurred if the saccade was made to a visual target (Mohler and Wurtz, 1976; Mays and Sparks, 1980). Whether or not these cells would produce a response for a memory guided saccade is a matter of speculation, but the fact that they did not respond for the second, remembered, saccade of the double-step paradigm would indicate that they should not. Nevertheless, it is common to use only the memory-guided saccade in neuron classification, and so it is unclear whether or not VTM cells would have been classed as visual or visuo-motor neurons in previous work (e.g. Ignashchenkova et al., 2004).

Finally, when additional features of the neural activity are examined, such as the spectral characteristics of the spiking activity during wait periods, additional neuron classes may be seen. In particular, I have found neurons in the superior colliculus which appear to have oscillatory spiking activity related to short-term spatial working memory (Chapter 3), a property which is spatially restricted to the movement field and temporally restricted to the portion of the trial in which the monkeys must remember the position of a target on the screen in the absence of a visual stimuli. Following Pesaran et al. (2002), I refer to this property as a *dynamic memory-field* and refer to these neurons as memory-field (MF).

5.2 Neuron classification scheme

In order to classify neurons, I used the delayed visually-guided and the memory-guided saccade tasks. I mapped visual receptive fields and movement fields

using the delayed visually-guided saccade task and then characterized the discharge prior to the saccade with the memory-guided saccade tasks. In the delayed visually-guided saccade task (Figure 5.1A), the monkey initially fixates a central dot. After the target appears, the fixation dot persists for a brief delay (500-900 ms). The monkey makes a saccade to the target at fixation offset. I randomly selected targets from a pre-defined grid and also chose target positions by hand in order to elucidate the extent of the response field. Mapping normally involved 50-150 trials. In the memory-guided saccade task (Figure 5.1B), the monkey initially fixated a central dot. The target appears for a brief period (500-900 ms) and then disappears. The monkey must remember this location for an additional delay (500-1000 ms), after which the fixation dot disappears and the monkey makes a saccade to the remembered location. I chose the location for the target based on the position of maximum saccade discharge in the delayed visually-guided saccade task, and then randomly interleaved trials in which the target appeared at this location and at the diametrically opposite location in the visual field. The monkeys performed 40-80 trials on this task.

Aside from delayed visually-guided and memory-guided saccades, I used spontaneous saccade activity to aid in detection of single-unit activity but neither recorded this activity nor systematically investigated it. In addition, I did not use a double-step paradigm or any other systematic behavioral tasks (such as those for anti-saccades or to detect fixation related activity). As a result, it would not be possible for me to determine if a neuron was quasi-visual. In addition, the fact that I did not record spontaneous activity means that I cannot with certainty characterize a neuron as a visually-triggered movement cell as defined by Mohler and Wurtz (1976), but I argue that when the authors characterized these cells, they conceived of the designation as referring to the need for a visual stimulus guiding the movement rather than an internally generated goal (such as a remembered location) since spontaneous saccades occur for entirely unknown reasons. Hence, the memory-guided saccade task will take the place of spontaneous saccades in the dark for the purposes of classification. Finally, I only isolated and recorded from neurons which had a spatially

segregated response (in other words, either receptive or movement fields, or both) on the delayed visually-guided saccade task. This practice, while it simplified later classification, probably significantly limited my ability to find any neurons other than the sort traditionally studied in the primate superior colliculus.

I classified neurons based on the following measures of response. I quantified visual responses as the increase in mean firing rate over baseline in an interval 50 to 150 ms after target onset. I quantified the saccade discharge as a the mean firing rate during a period starting 8 ms prior the saccade onset and ending 8 ms after the saccade offset. I quantified the prelude, or lead, portion of the saccade discharge as the mean firing rate during the 500 ms interval immediately prior to fixation offset in the memory-guided saccade task. As in (Munoz and Wurtz, 1995), I detected buildup activity as mean activity 100 ms prior to the saccade onset greater than 30 ips. As described in Chapter 3, I used the spiking spectrum in the 500 ms interval prior to fixation offset in the memory-guided saccade task to detect oscillatory activity.

Rather than using the ambiguous designation visuo-motor, I divided neurons which may have been classified as such in a more descriptive manner based primarily on the presence or absence of visual responses and the magnitude of prelude (long-lead) activity prior to the saccade discharge. Neurons with no visual responses or prelude activity were *saccade burst* (SB) neurons. These would be the same as the short-lead saccade-related discharge neurons of Mays and Sparks (1980). Neurons with no visual response but prelude activity that did not meet the criteria for buildup were *saccade burst neurons with prelude* (SBP), which would be examples of long-lead saccade-related discharge neurons of Mays and Sparks (1980). Neurons with a visual response but not prelude activity were *visually responsive saccade burst* (VSB). Neurons with both a visual response and prelude activity that did not meet the buildup criteria were *visually responsive saccade burst neurons with prelude* (VSBP). Both of these (VSB and VSBP) would reasonably be called visuo-motor units as in Mays and Sparks (1980), the latter visuo-motor units with prelude activity as in Mcpeek and Keller (2002). Furthermore, only BU units would be classified as such

according to Munoz and Wurtz (1995), whereas SB, VSB, VSBP, and SBP would be burst units. In the single unit study, I found it most practical due to similarity in response and the small number of examples found to fold SBP into the VSBP designation. Therefore, in addition to the traditional classifications of VTM, BU and SB, I employed VSB, VSBP, and MF.

Based on these measures and classes, I used a hierarchical (tree) sorting scheme so that once a neuron was classified, it was no longer eligible for consideration in another class. Therefore the sequence of classification is intrinsically important for distinguishing between the groups, and a different outcome on the same set of neurons would be possible were the scheme applied in the wrong order. First, I classified neurons as VTM if they had movement fields with saccade-related discharges greater than 10 ips in the delayed visually-guided saccade task but did not have a memory-guided saccade response[1] (in other words, a response less than 10 ips). Second, I classified neurons as MF if they had significant oscillatory spiking activity as defined in Chapter 3. Third, I classified neurons as BU if they had greater than 30 ips saccade related discharge 100 ms prior to the saccade onset in the memory-guided saccade task. Fourth, I classified neurons as VSBP if their mean visual response was greater than 10 ips and their mean prelude activity was greater than 4 ips, both in the memory-guided saccade task. Fifth, I classified neurons as SBP if their mean prelude activity was greater than 4 ips and their mean visual response was less than 10 ips, both in the memory-guided saccade task. I classified neurons as VSB if their mean prelude activity was less than 4 ips and their mean visual response was greater than 10 ips, both in the memory-guided saccade task. Finally, any remaining neurons with mean saccade responses greater than 60 ips on the memory-guided saccade task were classified as SB. Any neurons unclassified at the end of this procedure were not included in the single-unit study.

Examples of each of these classes are shown in Figure 5.2. Mean firing rate

[1]All such neurons in this study had a receptive field, although Mohler and Wurtz (1976) and Mays and Sparks (1980) reported neurons without receptive fields that still only had saccade-related discharges in the presence of visual targets.

and spike rasters are shown for target-onset aligned responses and saccade aligned responses for the memory-guided saccade task on the left, and receptive and response fields mapped using delayed visually-guided saccades appear on the right. Rasters for trials in which the target appeared in the response field appear above those for trials in which the target appeared in the diametrically opposite location. Mean firing rate for targets appearing in the response field and saccades made into the response field are blue, and mean firing rate for targets appearing out of the response field and saccades made out of the response field are in green. In maps of receptive and response fields, the mean firing rate of the visual response and of the saccade-related discharge appear as heat maps. The white dots in the receptive fields correspond to the target positions. The large yellow dots are the positions of the memory-guided saccade targets. In maps of response fields, the white dots correspond to the saccade end-points. The gray disks represent the positions of the stochastic motion stimuli in the single-unit study.

Examples neurons in Figure 5.2 appear in the rough order that these neurons would be encountered in a recording session with a penetration from most superficial to most deep. A VTM cell appears in A and B. Note that its receptive field is substantially larger and of higher magnitude than its response field for delayed visually-guided saccades. An SB cell appears in C and D. Its response field was extremely compact and is obscured by the saccade end-points; its burst had a peak firing rate close to 500 ips. A VSB cell appears in E and F. Its receptive field was substantially smaller and of lower magnitude than its movement field, and is partially obscured by the symbols representing target position. A VSBP cell appears in G and H. A BU appears in I and J. For whatever reason, the cell had minimal response to the target onset during the memory-guided saccade but small responses to other targets during the delayed visually-guided saccade. An MF cell appears in K and L.

5.3 Classification results

I isolated and characterized 119 single-units in the superior colliculi of two monkeys (Subjects F and M), 71 from Subject F and 48 from Subject M. All but one of these neurons were classified according to the scheme described above. These units were all isolated after mapping of the grid locations had been performed to focus our study only on those locations overlying portions of the map overlapping the stimuli in the single-unit study. Nevertheless, not all neurons could be included due primarily to loss of isolation prior to collection of sufficient trials on the selective attention task, which required usually 1.5-2 hours to complete. In some cases behavior was so poor, particularly if a neuron was not isolated until very late in the recording session, so that data was excluded. In a few rare instances the neurons had no response for any phase of the task because the receptive fields were so far from the stimuli. The counts of each type of neuron and the number of each type included in the selective attention study appear in Table 5.1.

Table 5.1: Classifications of neurons isolated for the selective-attention task.

Class	Subject F	Subject M	Total	Included
VTM	18	5	23	7
SB	5	1	6	2
SBP	8	1	9	2
VSB	7	15	22	13
VSBP	18	14	32	15
BU	11	3	14	7
MF	1	4	5	4
	71	48	119	50

The relative proportions of these neurons should in no way be taken as an indication of the prevalence of these neurons in the SC. The relative proportion is primarily a function of sampling bias. The most superficial neurons are those most likely to be sampled unless I made a concerted effort to find and isolate buildup neurons. In many cases a neuron appeared qualitatively to be buildup during the memory-guided saccade task but did not afterwards meet the criterion. Hence there

is a potential over-representation of VSB and VSBP neurons in the study. Finally, we did not establish the designation of MF until after all data had been collected, and since these neurons subjectively appeared to be deeper in the SC and we did not have reason to make a concerted effort to locate them, they are likely underrepresented.

A scatter plot of the depths of isolated neurons appears alongside structural images of Subject F's brain in Figure 5.3. A saggital section with the approach path at 38° of the recording electrode shown in red appears in Figure 5.3A, and a coronal section appears in Figure 5.3B. The area enclosed in the black box is the region presumably targeted in our recording studies. A scatter plot of the depths of the neurons isolated for these studies appears in Figure 5.3C. It is equivalent in its aspect ratio to the region shown in A and B, but only its vertical axis is intended to convey depth within the SC. Neurons are sorted by class (VSBP and SBP have been pooled as in the recording study), and their depths are represented on the y-axis as relative to surface of the superior colliculus. Since progressive penetrations cause some damage to the tissue that impairs our ability to accurately estimate surface depth, I used the initial detected depths at each grid location as the reference for the corresponding neurons. Mean and standard error are represented by the black lines and error bars. Approximate depths of the layers are shown in dashed lines; this is in rough correspondence to the observation that VTM cells are located just underneath stratum opticum (SO).

5.4 Methods

5.4.1 Stimulus presentation

The monkeys viewed stimuli presented on a CRT display mounted at eye level at a viewing distance of 410 mm. The viewable portion of the screen was 389 mm by 293 mm at a resolution of 1024 by 768 pixels. The CRT was a BARCO Reference Calibrator V driven at a refresh rate of 75 Hz. Each target consisted of a single pixel of background luminance surrounded by a 1-pixel-thick white border. Given

the display geometry and the viewing distance, this 3x3 pixel stimulus corresponded to 9 by 9 minutes of arc of visual angle. The background luminance of the monitor was 14 cd m^{-2} and the luminance of the targets was 50 cd m^{-2}.

5.4.2 Animal preparation

We recorded from three adult male rhesus monkeys that were 10-15 years of age and weighted 15-16 kg. The monkeys were prepared using standard surgical techniques described in detail previously (Krauzlis et al., 2000; Krauzlis, 2003). In particular, each animal had a head-post which allowed us to fix the head in the standard stereotaxic position during experiments. Each animal had a recording chamber for SC single neuron recording affixed to the skull with dental acrylic and additional titanium screws. The chamber was angled 38° to the posterior of vertical and directed at the midline 15 mm above and 1 mm posterior to the interaural line. In addition, each animal had implanted a scleral search coil around each eye (Judge et al., 1980) which were monitored with the electromagnetic induction technique (Fuchs and Robinson, 1966). All experimental protocols were approved by the Institutional Animal Care and Use Committee and complied with United States Public Health Service policy on the humane care and use of laboratory animals. The laboratory setup for behavioral control and monitoring was identical to that described previously (Hafed and Krauzlis, 2008).

5.4.3 Electrophysiology and neuron classification

We used tungsten microelectrodes (Frederick Haer, FHC) with impedances 1-2.5 MΩ to record extracellular action potentials of individual neurons in the intermediate and deep layers of the SC (1.5–2.9 mm below surface). Electrodes were advanced through stainless steel guide tubes (23 gauge) with a microdrive mounted on top of the recording chamber. The guide tubes were held fixed in the chamber with a delrin grid system (Crist et al., 1988). Extracellular neuron activity was passed through a standard head stage, amplified, and converted into trigger pulses

with a window discriminator applying both time and amplitude criteria (Plexon Systems, Inc.). The time of each action potential was stored with 1-ms resolution. We confirmed the position of recorded sites in the SC and classified neurons based on responses to a standard delayed and the memory-guided saccade task. Instantaneous firing rate was estimated by convolving each spike train with an acausal FIR filter with an impulse response function shaped like an excitatory post-synaptic potential (EPSP) (Figure 5.4) prior to averaging across trials.

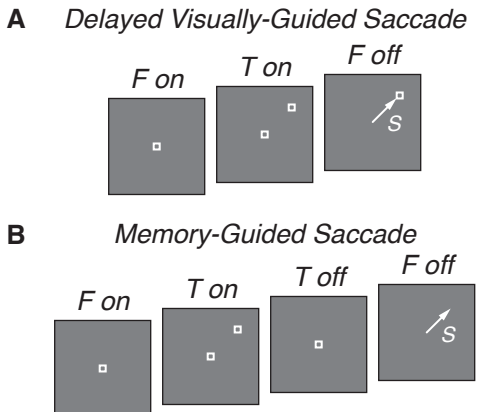

Figure 5.1: Delayed visually-guided and memory-guided saccade task. A) In the delayed visually-guided saccade task, the monkey fixates (F *on*) a dot. A Target appears in the periphery (T *on*). Once the fixation dot disappears (F *off*), the monkey makes a saccade to the target (S). B) In the memory-guided saccade task, the monkey fixates a dot appearing at the center of screen (F *on*). A target appears in the periphery (T *on*) and then disappears (T *off*). The monkey must remember the location of the target until fixation offset (F *off*) at which time he makes a saccade to that location (S).

Figure 5.2 *(following page)*: Example neural responses during memory-guided and delayed visually-guided saccade tasks. A) Rasters and mean firing rate of visually-triggered movement cell (VTM) during memory-guided saccade task aligned to target onset (T *on*) and saccade onset (S) Rasters of trials in which the saccade was made into the response field are plotted above those in which the saccade was out of the response field. B) Map of response of VTM cell to target onset and saccade onset on delayed visually-guided saccade task. Color represents average response in windows surrounding the target onset and offset (see Methods). White dots represent target position in map of visually evoked response and saccade end-point in map of saccade response. Large yellow dots represent end-points for memory-guided saccades. C-D) Responses for saccade burst cell (SB). Note that the saccade end-points in the saccade response map obscure the response field. E-F) Responses for visually-responsive saccade burst cell (VSB). G-H) Responses for visually-responsive saccade burst cell with prelude activity (VSBP). I-J) Responses for buildup cell (BU). K-L) Responses for a buildup cell with dynamic memory-field (MF).

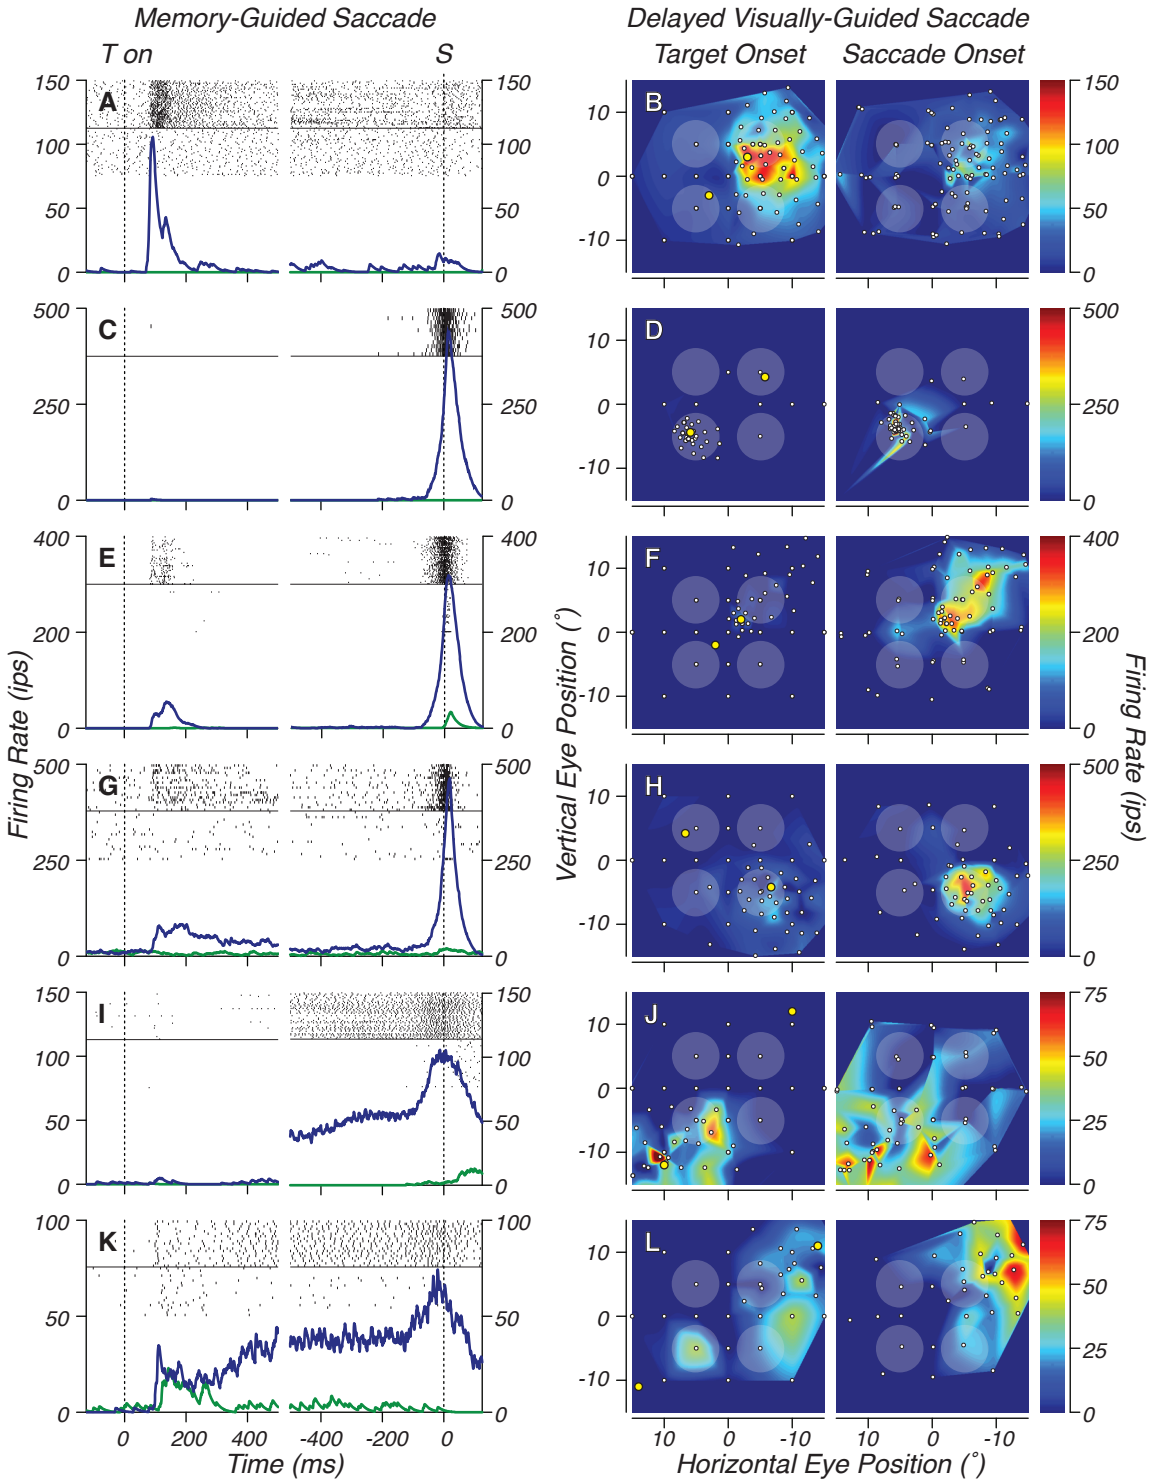
Memory-Guided Saccade
T on
S
Delayed Visually-Guided Saccade
Target Onset
Saccade Onset
A
B
C
D
E
F
G
H
I
J
K
L
Firing Rate (ips)
Vertical Eye Position (°)
Horizontal Eye Position (°)
Time (ms)

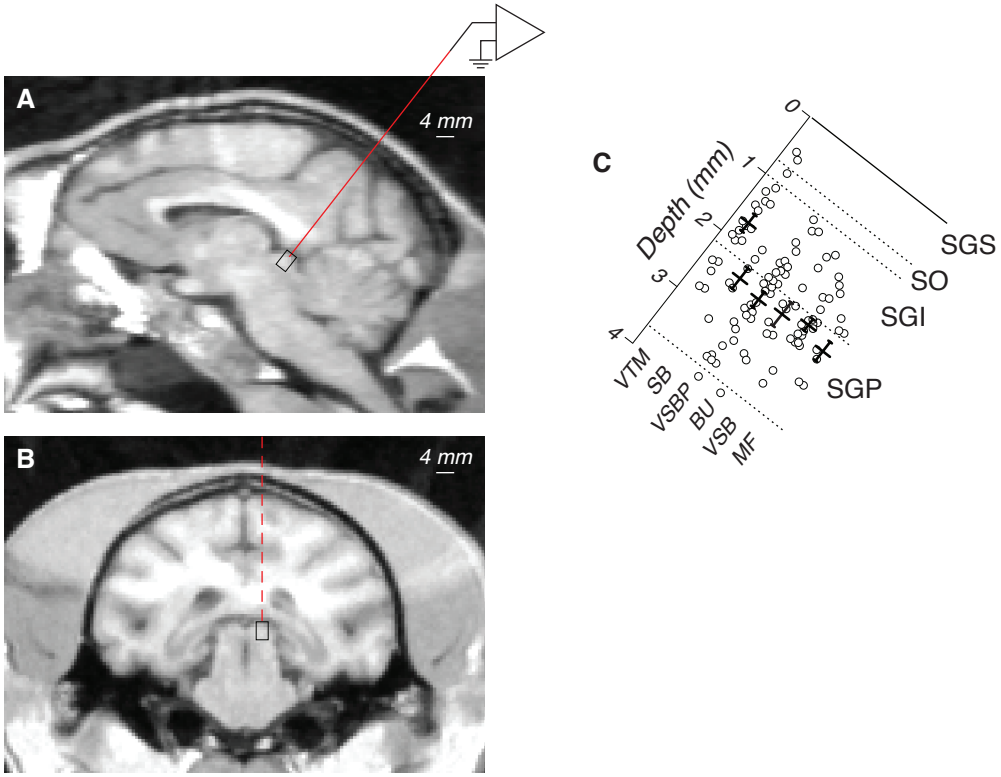

Figure 5.3: Location of superior colliculus on MRI and depths of recorded cells. A) Sagittal section of subject F. Plane of section is indicated by red dashed line in B. Black square encloses target recording area, and red line represents path of electrode from recording chamber to recording site. B) Coronal section through superior colliculus. C) Scatter plots of depths of cells relative to surface of the superior colliculus pitched back to standard approach angle of 38°.

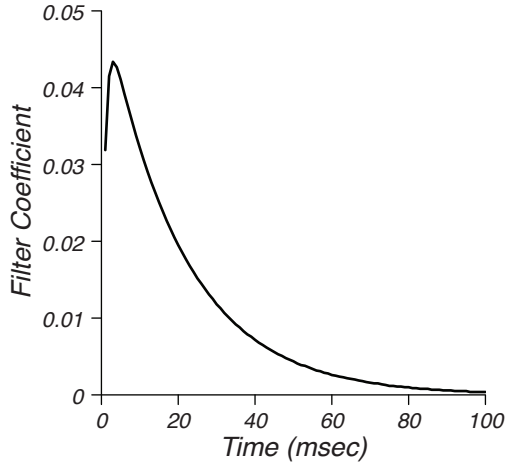

Figure 5.4: Filter for firing rate estimation.

Chapter 6

Single-Unit Responses in the Superior Colliculus to the Selective Attention Task

6.1 Introduction

Based on a history of recording studies on the visual guidance of eye movements (e.g. Goldberg and Wurtz, 1972b,b; Wurtz and Albano, 1980; Mohler and Wurtz, 1976; Wurtz and Mohler, 1976; Mountcastle, 1978) combined with a history of human patients with deficits in attention related to optic tectum lesions Posner and Petersen (1990); Posner et al. (1982); Posner (1980), the SC has long been suspected of a role in spatial attention. The basic characterization of this role was that the SC mediated the shift of attention preceding an eye movement to a location in space, and this view was bolstered by psychophysical observations that interference emerged when subjects attempted to make saccades and attend to different locations (e.g. Deubel and Schneider, 1996), indicating that some aspect of saccade programming and selective attention was shared. More recent recording studies have found mounting evidence of a role in the superior colliculus in covert attention (Kustov and Robinson, 1996; Ignashchenkova et al., 2004; Fecteau et al., 2004). Causal evidence

was found in microstimulation studies that showed that manipulating activity in the SC could mimic the selection of targets for enhanced visual processing (Cavanaugh and Wurtz, 2004; Müller et al., 2005; Cavanaugh et al., 2006). Finally, in light of our results that reversible lesion of the superior colliculus causes spatially specific visual extinction (described in Chapter 2), we are revisiting the issue of neural correlates of selective attention in the primate superior colliculus to investigate whether the activity of any groups of neurons can be related to the selection of stimuli that guide behavior.

I used the selective attention task described in Chapter 4 with which we showed that inactivation of the SC caused visual extinction. I studied five groups of neurons, including visually-triggered movement (VTM) cells, visually-responsive saccade bursting neurons with prelude activity (VSBP), visually-responsive saccade bursting neurons (VSB), buildup units (BU), and the memory-field (MF) cells described in Chapter 3. I investigated the responses to task events and the modulation of those responses by the presence of the cued position either in or out of the response fields of the neurons, including the cue onset, patch onset, the delay period prior to motion onset, motion onset, and the saccade response. I investigated how the timing and magnitude of responses to the motion onsets were related to cueing. I investigated whether any groups of neurons could be used to encode the location of the cued stimulus. Having discovered that only the memory-field cells were highly informative about the cue condition, I investigated what properties of their spiking activity appeared to be related to encoding the cue position and predicting task performance.

I found that different groups of neurons in the SC appear to have different functional roles in selective attention, including identification of the time of onset and spatial location of relevant stimuli in the visual field, selection of appropriate targets for saccadic eye movements, and encoding memory about spatial locations relevant for guiding behavior. The connections of the findings to these functions are discussed.

6.2 Data analysis methods

6.2.1 Detection of onset of discharge evoked by motion pulses

Some neurons may produce a sudden volley of spikes above a given tonic rate in response to the onset of the coherent motion pulses. Detection of this event is complicated by the fact that the firing rates of the neurons are not stationary in time, and so detection of the event must account for the immediate prior spiking activity. The basic approach was to generate a model of the firing rate at a given moment in time and to detect a deviation away from this expected response.

The model of firing rate was a mean subtracted and normalized null distribution of firing rates aligned on the onset of the motion pulse. Since the delay between cue offset and motion patch onset is geometrically distributed, there was a substantial overrepresentation of short trials. These came from a period of time in which most of the neurons were only beginning to show differences in firing rate based on the cue condition. Longer trials were sparsely represented and, on any given session, there were few or no trials with similar durations. Therefore I found it necessary to pool snippets of the trials across time. Each snippet was an interval at least 160 ms long starting at motion onset. Snippets from short trials had different firing rate statistics at any given moment, and the time varying statistics were different from that of snippets from longer trials in terms of both the mean and the variance. In order to generate the signal model, I start by taking the mean firing rate as a function of time but included only that portion of the trial that preceded the motion pulse. These mean firing rates are shown in Figure 6.3. At longer and longer trial durations, the number of trials in the average declined. At any given moment, the variance of the firing rate also increased because the firing rate increased. Therefore in addition to calculating the mean firing rate, I also calculated the standard deviation in firing rate as a function of time. In order to correct for the changing mean and variance, I subtracted the mean from each trial and divided each trial by the time varying standard deviation. This converted each trial to a sequence with mean zero and unit

variance.

To generate the null distribution, for each of these mean subtracted, normalized trials, I found all trials that were at least as long and contained a snippet at the same time of the appropriate length that did not include a motion pulse. I aligned all of these comparator snippets at the same starting time. This means that snippets of a given trial duration are appropriately represented in the null distribution. Since there are more short trials than long trials, there are more short trials in the null distribution than long trials. To generate the test distribution, I took each snippet at motion onset and aligned them all at the same starting time. Finally, to detect the evoked response to motion onset, I performed a two sample Kolmogorov-Smirnov test at each moment in time in the snippet. The onset time is the first moment that the distributions are significantly different at a 95% confidence level for three consecutive milliseconds.

6.2.2 Calculating per-trial periodicity

In order to quantify the structure within a single trial, I used a measure based on the model that each trial is generated from a point process with a mean described by a constant plus a sinusoid (Jarvis and Mitra, 2001). The spectra of these trials should be lines, and so I use the complex regression of the multi-taper spectra of each trial on this noise model as a statistic to measure that trial's periodicity.

The point process had the following time varying firing rate:

$$\lambda(t) = \lambda_0 + \lambda_1 \cos(2\pi f_1 t + \phi) \tag{6.1}$$

where λ_0 is the constant portion of the firing rate and λ_1 is the magnitude of the oscillatory component. The frequency of the oscillation is f_1, and I generally refer to this as the *peak frequency*.

The Fourier transform of the mean-subtracted point process and its expected value for the kth data taper are

$$J_k(f) \;=\; \int_{-\infty}^{\infty} h_k(t) e^{-2\pi i f t} \lambda_1 \cos(2\pi f_1 t + \phi)\, dt \tag{6.2}$$

$$E\{J_k(f)\} = c_1 H_k(f - f_1) + \bar{c_1} H_k(f + f_1) \quad (6.3)$$

$$E\{J_k(f)\} \simeq c_1 H_k(f - f_1) \quad (6.4)$$

where $c_1 = \lambda_1 e^{i\phi}/2$ and f is positive frequency.

Since each of the data tapers are orthogonal, the J_k provide an uncorrelated set of estimates of $c_1 H_k(0)$, and so the estimate of c_1 is found by complex regression:

$$\hat{c}_1 = \frac{\sum_k J_k(f_1) H_k(0)}{\sum_k |H_k(0)|^2} \quad (6.5)$$

The residual spectrum is

$$\hat{S}(f) = \frac{1}{K} \sum_k |J_k(f) - \hat{c}_1 H_k(f - f_1)|^2 \quad (6.6)$$

The ratio of the spectrum to the residual at f_1 is

$$F = \frac{|\hat{c}_1|^2 \sum_k |H_k(0)|^2 (K-1)}{\sum_k |J_k(f_1) - \hat{c}_1 H_k(0)|^2} \quad (6.7)$$

where F is distributed according to the F distribution with 2 and $2(K-1)$ degrees of freedom. F is the periodicity.

6.3 Results

6.3.1 Neural responses during selective attention task

Onset of the cues evoked a response in both VTM and VSBP cells. When the color-singleton was in the response field of the neuron, the peak of the evoked response was higher than when the color-singleton was out of the response field. This difference in peak response is shown in Figure 6.1A for VTM cells and in Figure 6.1C for VSBP cells. A suppression following the cue-evoked response was seen in both VTM and VSBP cells, and a second wave of spikes was seen in some neurons as in Figure 6.1C. The tonic response of both VTM and VSBP cells during the persistence of the cue was often larger when the cued position was in their response field. Both the suppression and the second onset were inconsistent features, whereas the large

difference in peak response consistently occurred, as seen in the population average for VTM cells (Figure 6.1B). The difference in peak response was not consistent on average for VSBP cells (Figure 6.1D).

Onset of the incoherent motion patches evoked a second response in both VTM (Figure 6.2A) and VSBP (Figure 6.2C) cells. In addition, VSB cells (Figure 6.2E) often responded to the onset of the patches with a short burst of spikes. The peak response to the patches for VTM cells when the cued position was in the response field was larger than when the cued position was out of the response field. This difference in peak response was consistent across the population as seen in the population average (Figure 6.2B). In contrast, VSBP cells did not have a consistently higher peak response to the onset of the patches when the cued position was in the response field. This was seen in the example shown in Figure 6.2C and was consistent across the population (Figure 6.2D). Sustained responses to the patches were often higher when the cued position was in the response field for VSBP cells (Figure 6.2D). Similarly, despite the response evoked by the cue onsets in VSB cells, the response was symmetric across cue positions as seen in the example shown in Figure 6.2E and the population average (Figure 6.2F).

Following offset of the cues, large sustained differences in firing rate emerged for nearly all groups of cells. Motion pulses did not occur until at least 480 ms after the cue offset, and the differences in firing rate emerged prior to this time for VTM (Figure 6.3A), VSBP (Figure 6.3B), VSB (Figure 6.3C), and buildup (Figure 6.3D) cells. The difference began emerging hundreds of milliseconds prior to the cue offset for MF cells (Figure 6.3E), and no difference emerged for saccade burst neurons, which had minimal responses to any phase of the task (Figure 6.3F). This difference in firing rate was associated with increasing firing rates both when the cued position was in the cell's response field and when the cued position was not in the response field. Firing rates began increasing during the initial delay between cue offset and the first time that the motion pulses could occur and continued to increase as trials continued. In MF cells (Figure 6.3E), however, the difference in firing rate peaked

at around the time that stimuli could first occur and then slowly declined over the course of the trials. At any given moment, the calculation of mean firing rate for each neuron includes only those trials for which the motion pulse has yet to occur.

On average, the firing following cue-offset rose to a sustained level around the mean of the wait time, represented by the dashed lines in Figure 6.3 and Figure 6.4, and then persisted at a relatively sustained level. This increase in firing rate emerged for both the condition in which the cued position was in the response field and when it was out of the response field. Just as with the example shown in Figure 6.3E, the population average response of MF cells (Figure 6.4E) peaked at the first moment that motion pulses could occur and then declined from that point forward. The response for VTM cells (Figure 6.4A) may have in fact began to decrease in magnitude and collapse back to a sustained level, although this would have occurred on a small proportion of total trials. In contrast, VSBP (Figure 6.4B), VSB (Figure 6.4C), and buildup (Figure 6.4D) cells reached a relatively sustained level of firing. The population average for (all two) saccade burst neurons recorded reflected the near complete absence of response for these neurons during the task.

Firing rate aligned on motion pulse onset reflects both the difference based on cue position and the fairly linear increase over the course of the delay period (Figure 6.5). In Figure 6.5, the second dashed line represents the offset of the coherent motion pulses and the onset of the choice targets, and the firing rate seemed to increase prior to the onset of the choice dots in examples from the more visual cell groups, including VTM (Figure 6.5A), VSBP (Figure 6.5B), and VSB (Figure 6.5C). In contrast, a substantial deviation in firing rate did not occur until after the onset of the choice dots in the examples from cell groups that were not visually responsive, including BU (Figure 6.5D), MF (Figure 6.5E), and SB (Figure 6.5F). Although this response may appear to be related to the motion onset, it could possibly also reflect the beginning of activity related to saccade target selection or saccade generation, and so additional investigation will be required to disentangle the potential response to motion onset and the response associated with the saccade.

On average, the firing rate aligned to motion pulse onset reflects the difference in response based on cue position and the increase over the course of the delay period (Figure 6.6). On average, VTM (Figure 6.6A) cells had an increase in firing rate up until the motion pulse and a large transient increase in firing rate that had a higher peak for trials in which the cued position was in the response field than out of the response field. Similar results were seen for VSBP (Figure 6.6B) and VSB cells (Figure 6.6C). Unlike the example in Figure 6.5D), buildup units on average had a surge in firing rate during the motion pulse that was followed by a second peak likely related to the saccade (Figure 6.6D). MF cells did not on average have a substantial change in response until after the appearance of the choice dots (Figure 6.6E), as did SB cells (Figure 6.6F).

Following the motion pulses, the monkeys made saccades to the choice dots to convey their response. For the purpose of visualization and analysis, firing rate is shown for trials in which the monkeys made saccades into the response fields and for trials in which the monkeys made saccades out of the response field; in this case, the response is averaged across all three directions out of the response field. Although a choice dot appeared in its response field, the example VTM cell (Figure 6.7A) did not appear to produce a saccade-selective response. The large burst in activity around 200 ms prior to the saccade onset is the motion response seen in Figure 6.5A and Figure 6.6A. Following the motion onset response, firing rate was relatively constant across cue conditions and later response directions until after the initiation of the saccade. In contrast, examples from all other types of cells had distinct saccade-related discharge that was much greater for saccades made into the response field than for those made out of the response field. The firing rate in these examples increased prior to the saccade for all response directions and for both cue positions prior to separation of the trajectories based on the response directions. The initial increase in activity could reflect a response to the motion onset prior to the saccade, in particular for the VSB example (Figure 6.7C). In addition, the response to the saccade in both the buildup (Figure 6.7D) and MF (Figure 6.7E) cell examples appeared to have a larger response

for saccades made into the response field when the cued position was in the response field than when it was not.

Population average responses to saccade onset showed the same pattern seen for the examples (Figure 6.8). VTM cells did not show a pronounced saccade response, and the putative response to motion onset was larger when the cued position was in the response field than not (Figure 6.8A). VSBP, VSB, and BU cells (Figure 6.8B-D) showed a transition from reflecting the cue position to the saccade with no difference, on average, in the saccade response based on the cue position. MF cells on average maintained a difference in firing rate based on cue position until after the saccade onset (Figure 6.8E). SB cells, as expected, produced substantial responses for saccades made into the response field and were otherwise not responsive to the task (Figure 6.8F).

Prior to examining the modulation of firing rate based on attention during the different measurement intervals, I will focus on establishing whether or not the putative response to the motion onset observed in VTM, VSBP, VSB, and buildup cells as seen in Figure 6.5 and Figure 6.6 was in fact a response to the motion onset and not merely the early appearance of saccade-related activity.

6.3.2 Separating motion onset-evoked discharge from saccade-related discharge

The response to motion onset seen in some groups of neurons is either a response to the motion onset or part of the initial response preceding the saccade. Although neurons in the superficial layers of the superior colliculus are known to respond to moving stimuli in their response fields (Wurtz and Albano, 1980), VTM cells have not been investigated for this property. In monkeys trained to discriminate the direction of coherent motion stimuli, buildup units are known to have motion selective responses even when the monkeys are not making saccades into those neurons' response fields (Horwitz and Newsome, 1999, 2001a,b). In monkeys not trained to discriminate coherent motion stimuli, however, buildup units are well-driven by

stochastic motion stimuli but exhibit essentially no direction selectivity (Krauzlis, 2004). Therefore, if the response has any direction selectivity, one may reasonably suspect that it is related to saccade planning; in any case, any effort to distinguish a motion response which is direction selective from the immediately succeeding saccade-related activity would be ambiguous at best. The primary effort will be to distinguish a motion response which is not direction selective from a response which is predictive of the saccade direction.

In addition to distinguishing the motion response from saccade-related activity, the effort to determine that cells respond to motion onset is complicated by the fact that there is no set of trials in which a direct comparison is possible; in other words, trials in which the motion pulses occur at a given time are not explicitly paired with otherwise identical trials with incoherent motion pulses at the same time. Furthermore, the firing rate of the neurons increase in time, so what we are essentially tasked with doing is to detect a transient increase in firing rate which is already riding on top of a non-stationary response. As described in more detail in Methods, my strategy was to compare trials in which the motion pulse occurred at a given time with all those trials for which the motion pulse had yet to occur at that same time.

The standards for determining that the response is to motion onset depend on the relationship between saccade response latency and the neural response to motion onset. If I examine trials of different response latencies, such as the first and second half of the distribution for each neuron, then the time of detectable motion onset response will be independent of mean response latency within those groups. In contrast, the time of detectable selectivity for the saccade direction will be later in longer-latency trials since it is directly related to the saccade onset time.

In VTM neurons, the time of the evoked neural discharge to motion onset is independent of the latency of the monkeys response saccades. Moreover, there is no time at which selectivity for the saccade direction emerges prior to the saccade (Figure 6.9A). In this graph, each point represents the time at which the motion response

was detectable for each cell, and there were no examples in which selectivity for the saccade emerged. In an example VTM cell, the firing rate for conditions in which the response is either into or out of the response field and the cued signal is either in or out of the response field diverge from the expected firing rate prior to motion onset and do not show saccade selectivity near the time of the saccade (indicated by the whisker plots for cue in and cue out (blue and green, respectively). The apparent timing of the motion response does not change for trials in which the response time is shorter than average (Figure 6.9D) or longer than average (Figure 6.9F). A similar pattern emerged for the population average, which had no discernible difference in neural discharge based on saccadic response latency (Figure 6.9C,E, and G).

As in VTM neurons, the time of the evoked neural discharge to the motion onset is independent of the latency of the monkeys' response saccades for VSBP neurons. In addition, the time of the emergence of saccade selective activity was dependent on the latency of the response saccades (Figure 6.10A). In the example neuron shown, detailed examination of the firing rate reveals that the discharge for saccades made both into and out of the response field appears to have the evoked response to motion onset followed by responses associated with the saccade (Figure 6.10B). Segregating trials based on the response latency leaves the response to motion onset essentially unchanged but the response to the saccade shifted with the average latency (Figure 6.10D and F). In the population average, the response to the motion onset appears to be concurrent with an emerging saccade selective response. Regardless of the later direction of the saccade, there is a transient increase in firing rate based on the motion onset followed by subsequent separation of firing rates reflecting the eventual saccade. This occurs for both the condition in which the cued position is in the neurons' response fields and when it is not.

Similar to the VSBP cells, the time of evoked neural discharge to the motion onset is independent of the latency of the monkeys' response saccades for VSB neurons (Figure 6.11A). In addition, the time at which saccade selectivity emerges depends on the latency of the response saccade. As with VSBP neurons, the responses to the

motion onset and to the saccade in the example neuron appear to be superimposed (Figure 6.11B), and in longer latency saccades, the responses appear to be almost entirely separate (Figure 6.11C). Presumably if the response could be delayed even farther from motion onset, there would be no ambiguity that these responses were entirely independent. On average, the motion onset response occurs simultaneously with the increase in firing rate associated with the saccade, and only near the peak response to the motion onset does the firing rate diverge.

Buildup cells appear to have a more ambiguous dissociation between the motion onset response and the saccade-related discharge (Figure 6.12). As for the other cell groups, a response to motion is detectable before saccade selective activity emerges (Figure 6.12A). However the average firing rate in both the example cell and the population average (Figure 6.12B-C) shows that the activity related to saccades made into the neurons' response fields emerges at the same time as the motion onset response. This response has such great variability that the difference between saccades made into the response field and out of the response field does not emerge until later. When the saccadic response latency is long, the response to the motion onset becomes more pronounced and the separation in saccade-related activity is delayed (Figure 6.12F-G).

Overall, the earliest response to motion was that of the buildup cells at 112 ms, followed by VTM cells at 115 ms, VSBP cells at 128 ms, and VSB cells at 124 ms. In addition, the onset time was earlier when the cued position was in the response field than when the cued position was out of the response field. The motion-onset evoked discharge times are tallied in Table 6.1. When the cue was in the response field of VTM neurons, the onset of the discharge was nearly 10 ms faster than when the cued position was out of the response field. Cue position did not change the onset time for VSBP, VSB, and BU cells.

Table 6.1: Time of motion onset evoked discharge.

Class	Onset time (ms)	
VTM	Cue In	115 ± 5
	Cue Out	122 ± 7.5
VSBP	Cue In	128 ± 5.2
	Cue Out	127 ± 5.3
VSB	Cue In	124 ± 9.2
	Cue Out	124 ± 16
BU	Cue In	112 ± 9.2
	Cue Out	112 ± 4.5

6.3.3 Summary of modulation of firing rate by cue condition and signal selection

I now examine changes in firing rate during the intervals shown in the time course of neural response to task events. For cue and patch onset, I will examine changes in mean firing rate associated with the cue position. In addition, I will look at normalized differences in firing rate between trials in which the cued position was in the response field and when it was out of the response field, and for each of these conditions, when the monkeys based decisions on the cued signal or on the foil signal. In these latter to cases, I will adopt the convention that positive differences are associated with choices based on the signal in the response field of the neuron. One important question to consider is whether or not the data reflects the following proposition: if the differences in firing rate observed based on cue position are involved in controlling the signal to which the monkey attends, then the actual performance of the monkey should be related to the variability in the neural response. As in, higher firing rates are associated with the cued position in the response field of the neuron, and the monkeys most often base decisions on the cued signal; therefore, the monkey should be more likely to base decisions on the signal in the neuron's response field if the neuron's response is larger.

As seen in the example neuron and population averages shown in Figure 6.1, VTM cells had a higher peak response when the cued position was in their

response field than when it was not in their response field. This trend was true for all observed VTM neurons and across the population (Figure 6.13A and C). In contrast, VSBP cells did not have as large a difference in their peak firing rate based on the cue position, but the normalized difference in firing rate based on cue position does reflect this difference for cells with lower overall firing rates. Despite the higher firing rate for trials in which the cued position was in the response field, however, correct performance was associated with little or no difference or a suppression of the peak firing rate (Figure 6.13E) for VTM cells and no difference for VSBP cells (Figure 6.13F). Choices with the foil signal were associated with little change in the difference in firing rate for VTM cells and for an increase in firing rate for VSBP cells (Figure 6.13G-H) when the cued signal was not in the cells' response fields.

Similar to the sample neuron and population averages shown in Figure 6.2, the response to patch onset in VTM cells was elevated but the response to patch onset in VSBP and VSB cells was either unchanged or depressed by the presence of the cued position in the cells' response fields. Correct responses were associated with almost no difference in activity for either VTM, VSBP, or VSB cells when the cued position was in the response field (Figure 6.14G-I). When the cued position was not in the neurons' response field, the response of VTM cells was slightly suppressed if the monkey based his decision on the cued signal, but the response of the VSBP and VSB cells was slightly higher; in other words, higher firing rates when the foil signal was present in the response field were associated with later choices based on that signal.

Firing rates for most neurons were higher during the initial memory phase following cue offset when the cued position was in their response fields than when it was not (Figure 6.15A-C, Figure 6.16A-B). This trend was most pronounced for MF cells, which had their peak firing rate in this interval (Figure 6.16B,D). It was least pronounced for VTM cells, which had little difference in activity in this interval (Figure 6.15A,D). When the cued position was in the response field, little if any difference in firing rate existed between trials in which the monkey based his decisions

on the cued signal or on the foil (Figure 6.15G-I, Figure 6.16E,F). When the foil signal was in the neurons' response fields, slight elevations in firing rate were associated with choices based on that signal for VSBP, VSB, BU, and MF cells (Figure 6.15K-L, Figure 6.16G-H).

Larger differences in firing rate based on cued position emerged in VTM, VSBP, and VSB cell groups prior to motion onset (Figure 6.17A–F), but the difference in BU and MF cells (Figure 6.18A-D) had at that time begun to decline. As in previous intervals, when the cued position was in the neuron's response field, correct choices were not associated with a particular change in the firing rate of the neuron for any cell group (Figure 6.17G-I and Figure 6.18E-F). On the other hand, when the foil signal was in the response field, most neurons had a marginally elevated firing rate if the monkey based his decision on those signals (Figure 6.17J-L and Figure 6.18G-H).

Finally, nearly all neurons had an elevated response based on the presence of the cued position in their response fields following motion onset (Figure 6.19A-C, Figure 6.20A-B). The differences in firing rate were large for nearly all neurons (Figure 6.19D-F and Figure 6.20C-D). Although VTM cells appeared to have a slight suppression of their activity when the cued position was in the neuron's response field and the monkey based his decision on that signal, nearly all neurons had elevations in activity for this case. Furthermore, when the foil signal was in the neurons' response fields, the firing rate was higher if the monkey based his decision on that signal.

In summary, neurons consistently had elevated firing rates for both transient and sustained responses when the cued position was in the neurons' response field. Visually responsive cell groups VTM and VSBP cells had elevated responses to cue onset. VTM cells had an elevated response to patch onset whereas VSBP and VSB cells, in general, did not. In the initial memory period following cue offset but preceding the onset of any stimuli, MF cells had the largest sustained difference in firing rate of any group. Later, prior to the motion onset, all neurons had elevations in firing rate based on the presence of the cue, and this trend continued through to

responses associated with the onset of the motion pulse.

Whether or not the neurons displayed changes in mean firing rate based on whether or not the monkey eventually based his decision on the signal in that location depended on the temporal proximity of that interval to the appearance of the stimuli and whether or not the cued or foil signal was in the response field of the neuron. Initially, the response of VTM cells to the cue onset was suppressed slightly if the monkey was correct, regardless of the cue position. The response of VSBP neurons was slightly elevated when the foil signal appeared in their response field and if the monkey based his decision on that signal. Similarly, the response of VTM cells to patch onset was slightly suppressed for correct performance regardless of the cue position, and VSBP cells and VSB cells had elevations in their firing rate when the foil signal was in their response field and the monkey based his decision on that signal. This appeared to be the case for the period of time following cue offset and became more pronounced immediately prior to the motion onset. Finally, only in the response to motion onset itself were elevations in firing rate seen associated for choices based on the signal in the neurons' response fields for VSBP, VSB, BU, and MF cells, although VTM cells remained marginally suppressed in their visual responses if the monkey correctly based his decision on the cued signal in their response field.

6.3.4 Dependence of motion onset response on signal strength

In individual cells and in the population average, the discharge evoked by motion onset in VTM, VSBP, VSB, and BU cells was present for both the cued signal and the foil signal, although its magnitude was larger for the cued signal. Innumerable possibilities exist for what significance this discharge may have for the circuit and the animal's ability to select the appropriate signal to base his decisions, some of which I will consider in the Discussion. What we can here determine about the response is how it depends on signal strength. In addition to trials in which the cued and foil signals were matched in strength, instructor trials in which the signal

in the cued position had 0.5 coherence, the strongest presentable, and no signal in the foil location, appeared at the start of each block. Although the number of these trials is quite small in each neuron, I can determine whether the population average for the discharge evoked by the instructor trials is any different from that evoked by the trials in which the cued and foil signals are matched.

The response to the instructor signal was much larger than the response to the cued signal in both VTM and VSBP cells, but essentially the same in VSB and BU cells. In VTM neurons, the response to the instructor signal when it is in the response field of the neuron is more pronounced than the response to any other signal. It occurs at roughly the same time but has a faster rate of rise and a higher magnitude (Figure 6.21A). When the instructor signal is out of the response field, no additional response is evoked. In VSBP neurons, the response to the instructor signal when it is in the response field of the neuron is similarly more pronounced than either the response to the cued signal or the response to the foil signal (Figure 6.21B). The response appears to begin at the same time as that for the cued and foil signal, but its rate of rise is larger and its peak is greater than for either of these. In contrast, the response of VSB cells, which had a fairly distinct response to the motion pulse, is no different between to the instructor than to the cued signal (Figure 6.21C). Likewise, buildup neurons have no difference between the response to the instructor trial and to the cued signal.

6.3.5 Interaction between selective attention and saccade-related discharge

A hallmark of neurons in the intermediate and deep layers of the SC is that their firing rate, in particular that of buildup neurons, is predictive of saccade choices well before the saccade is executed (e.g. Krauzlis, 2003; Krauzlis and Dill, 2002). In FEF neurons with prelude activity prior to a saccade, many cells appear to reach a fixed threshold after which the saccade direction can be distinguished from the firing rate but prior to the onset of the saccade (Hanes and Schall, 1996). This

characteristic would seem to be requisite in a group of neurons in which the firing rate reflects the evolution of a decision process–it may start at a different point based on initial conditions, but in the end it follows the same trajectory. So, if the neurons have a substantial change in their baseline firing rate prior to the saccade, one may reasonably ask if their relationship to saccade onset remains the same. Therefore I examined when the saccade direction (either into or out of the response field) could be determined based on the firing rate of the neurons leading up to the saccade.

A common analysis strategy to obtain such a discrimination time is to compute receiver operating characteristic (ROC) curves on a millisecond timescale and to use the area under that curve as a measure of the separability of firing rate distributions between cases in which the saccade is made into or out of the response field. A key assumption underlying this analysis is that there is a mechanistic connection between the neuron and the saccade onset: each neuron is paired with an antineuron across the SC that has a mirror symmetric response field, otherwise identical in every other way. When the difference in firing rate between these two reaches a fixed threshold, the saccade is triggered (see Britten et al. (1992), which perpetrated such a theoretically inelegant and biologically unlikely model upon the field in their study of MT encoding of motion direction, and example of studies in which it nevertheless has proven useful: Basso and Wurtz (1998); Mcpeek and Keller (2002); Van Wezel and Britten (2002); Krauzlis and Dill (2002); Kim and Basso (2008)). If the relevant control signal is the difference in activity between the neuron and its anti-neuron pair, then what impact would attention have on the mechanism of saccade initiation?

The failure of the neuron/antineuron pair as a model of saccade initiation becomes clear when we visualize how the pair might be responding prior to a saccade into the response field of a buildup neuron in this task. If the monkey has been attending to the response field of a buildup neuron and is preparing to make a saccade into the response field of the neuron, then it and its anti-neuron pair in the unattended, diametrically opposite location would have the firing rates shown in Figure 6.22A. The mean firing rate of the neuron is everywhere greater than the mean

firing rate of its antineuron pair, and there is no ambiguity about the target of the saccade, nor any relationship of the difference between these firing rates and the time of saccade initiation. An outside observer with perfect knowledge of the statistics of the firing rate would instead employ distributions with means shown in Figure 6.22B, in which we can see that the neuron's expected response prior to a saccade into its response field when the cue is in its response field does not begin to diverge until around 150 ms prior to saccade onset. Nevertheless, this hypothetical distribution could be called the nonexistent neuron (similar to the construction of antineuron) since it corresponds to a hypothetical response which nowhere occurs. Instead, it is a construct that we may use to determine when we could discriminate the target of the saccade, regardless of the actual circuitry underlying saccade initiation.

Instead of the neuron/antineuron pair model, I will examine the divergence between the distributions of firing rates when the saccade is made either into or directly out of the response field for a given cue condition. The questions I wish to pose then about the saccade-related discharge are when can the target of the saccade first be determined from the firing rate preceding the saccade, and what is the firing rate of the neuron at this moment for both the case in which the cued position is in the response field of the neuron and when it is not. I determined the discrimination time as the third consecutive millisecond that the distribution of firing rate for saccades made into the response field was significantly different from the distribution of firing rates for saccades made out of the response field by a two-sample Kolmogorov-Smirnov test. The results of this analysis are detailed in table Table 6.2 and Figure 6.23 and Figure 6.24.

The target of the saccade could be discriminated from the firing rate distribution earlier and at a higher firing rate when the cued position was in the response field of the neuron than when it was out for VSBP and VSB neurons. These differences in discrimination time were at best marginal for both classes. In contrast, the target of the saccade could be discriminated from the firing rate distribution at essentially the same time and at the same firing rate for buildup units regardless of

Table 6.2: Discrimination time and firing rate preceding saccades.

Class		Discrimination Time (ms)	Firing Rate (ips)
VSBP	Cue In	-70.125 ± 9.59	51.23 ± 9.28
	Cue Out	-65.50 ± 11.07	43.6 ± 9.39
VSB	Cue In	-69.46 ± 9.15	32.5 ± 6.5
	Cue Out	-57.6 ± 8.48	25.69 ± 6.62
BU	Cue In	-74.86 ± 17.03	61.3 ± 13.45
	Cue Out	-71.14 ± 19.08	63.18 ± 15.24
MF	Cue In	-88.75 ± 9.68	26.60 ± 5
	Cue Out	-82.75 ± 19.92	12.32 ± 3.47

cue condition. Lastly, the discrimination time for MF cells was essentially the same irrespective of cue condition, but the firing rate was much higher when the cued signal was in the response field. These differences are summarized in Table 6.2. Overall, these results would seem to be consistent with the notion that the buildup units are part of the decision mechanism for the selection of the saccade goal and the initiation of the saccade, whereas the saccade related activity is in a sense imposed upon the MF cells since it is superimposed upon the differences in firing rate based on cue position.

6.3.6 Encoding cue position in spike count

Although there were substantial differences in mean firing rate in SC neurons related to the cue position at different phases of the task, this alone is insufficient to indicate whether or not these neurons could be components of a circuit which regulates selective attention. The issue is not simply that the differences in firing rate are only correlated with both the cue condition and, as a result, the monkey's behavior. Instead, other statistics of the response, such as the variability of neural activity, must be considered. For example, if the spiking activity of the neuron is Poisson, then increases in firing rate are associated with increases in variability. This increase in variability would undercut any ability to encode information in the spiking activity which might have been implied by the increase in firing rate. One method

of assessing this variability is to examine the Fano Factor, the ratio of variance in spike count to mean spike count in an interval (Rieke, 1997). When the Fano Factor is large, in particular greater than 1, the increases in firing rate can actually degrade the degree to which information could be encoded in spiking activity because the variance grows faster than the change in firing rate. A second method to assess how efficiently the neurons could encode information is to calculate the mutual information between the spike count and the cue condition. Mutual information is a measure of the reduction in uncertainty in one variable after observation of another (Cover and Thomas, 2006). As such, it would seem to be the most direct means of probing the issue under consideration, that is, how efficiently the spiking activity of different SC neurons could be employed in encoding the position of the cue if all the code were based on was the number of spikes in an interval of time. The units of this measure are typically bits, so in this case there is 1 bit of uncertainty about cue position because it can be either directly in or out of the response field, requiring only a single binary digit to specify in which of those two positions it is in.

Despite the fairly large differences in peak firing rate in the discharge evoked by cue onset, the number of spikes in that response was, in general, only marginally informative about cue position for VTM neurons and statistically uninformative for VSBP neurons (Figure 6.25A). This is because the actual changes in spike count were not large for VTM neurons and because the Fano Factor for many of them is greater than 1 (Figure 6.25B,D). Similarly, the changes in mean spike count for VSBP neurons were not large and the Fano Factor for many of these was much greater than 1 (Figure 6.25C,E).

Responses to patch onset were even less informative about the cue position than were responses to the cue onset (Figure 6.26A). Again, despite the difference in firing rate observed in VTM neurons (Figure 6.26B), the Fano Factor was quite high (Figure 6.26E). VSBP and VSB neurons were statistically uninformative about the cue position due to a lack of change in firing rate (Figure 6.26C,D) and Fano Factors which were greater than 1 (Figure 6.26F,G).

During the initial memory period after cue offset, neither VTM nor VSB neurons were informative about the cue condition (Figure 6.27A). In contrast, VSBP and buildup neurons were, on average, informative, and MF cells were extremely informative about the cue condition. This is because, again, despite the differences in firing rate during these intervals observed in both groups of cells (Figure 6.27B and D), the responses were quite variable between trials and the Fano Factors were quite high (Figure 6.27E and G). The substantial differences in mean spike count in VSBP neurons (Figure 6.27C) were undercut by the relatively large Fano Factors (Figure 6.27F) for many of the neurons. MF cells, however, had both large differences in firing rate based on the cue condition and had Fano Factors less than 1 when the cued position was in the response field of the neuron (Figure 6.28D) and low (although still greater than 1) Fano Factors when the cued position was out of the response field of the neuron.

The memory-field cells would appear to have the optimal properties for encoding information in spike countlarge differences in firing rate coupled with decreases in Fano Factor based on cue condition. The time course of the efficiency of this encoding is shown in Figure 6.29. Time is with respect to cue offset, and Figure 6.29A shows the mutual information for each of the example neurons shown in Figure 6.3. The inset shows the empirical distribution of trial durations and the performance of the monkeys on the selective attention task as a function of time. The MF cell increases in how informative it is about cue condition until it reaches its peak value at cue offset, then retains that value until the first moment that the signals could appear. Then, as the firing rate declines, the information declines along with it. Other example neurons do not increase in how informative they are about the cue condition to anywhere near a similar degree. The inset serves to remind us that the all of these neurons are either not informative or declining in how informative they are about the cue condition over the same period of time that the monkeys maintain their memory of the cued location. The population average of the time course of information emphasizes that these trends are not properties of individual examples but properties of

the population of each group.

Prior to motion onset, all groups of neurons other than MF cells would seem to be more informative about the cue position than in previous intervals (Figure 6.30A). Nevertheless, the variability in spike count is so large that the chance of seeing a random grouping of trials with a substantial amount of information about the cue position by chance is quite large. Thus the amount of information in the spiking activity expected by chance is larger than the average amount of information for the classes except VSBP and MF cells, which are both only marginally informative about the cue condition. This is because the large changes in firing rate as seen for all groups are accompanied by extremely high Fano Factors, which means that the cells become even more variable as their firing rate increases, a feature which drastically impairs the efficiency with which they could encode information about the cue position in advance of the motion onset. In addition, the reason that MF cells become less informative despite their relatively substantial difference in firing rate (Figure 6.31B) is because the Fano Factor (Figure 6.31D) is much higher than during the period immediately following cue offset (Figure 6.28D). One reason that the Fano Factor is so high is that the firing rate changes in time and the trials are aligned on motion onset. Therefore the variance is reflecting not simply variability in spiking activity but also variability in the trial duration. Therefore additional observation beyond what is available, such as the firing rate of other neurons or the trial duration, would be necessary to decode the cue position from the neuron's response.

Lastly, the evoked response to the motion onset (or the maintained firing during the same period) also is only marginally informative of the cued position (Figure 6.32 and Figure 6.33). Again, the average information in VTM cells is at the chance level, whereas VSBP, VSB, BU, and MF cells are just above chance information. Despite the substantial differences in firing rate in VSBP and VSB cells (Figure 6.32C,D), the Fano Factors are quite high (Figure 6.32F,G). In contrast, the cells with lower Fano Factors, namely BU and MF cells (Figure 6.33) have smaller differences in firing rate in this interval despite their relatively smaller Fano Factors.

The reader may reasonably question how relevant it is to calculate mutual information for single cells and then draw conclusions about the coding efficiency of the population. If the firing rates of many neurons are pooled, he may propose, then would that not increase the total information about cue condition? The degree to which several observations can be pooled to increase coding efficiency depends critically on several unknowns, in particular, the degree of correlation between the observations (the more correlated the less informative are additional observations) and the similarity of the coding properties of the different neurons from which the observations are made. In other words, if several neurons have precisely overlapping response fields and are completely uncorrelated from each other, then the variability of each independently can be overcome by decoding based on the population. However if the neurons are highly correlated, then once one value is known, all are known, and no further gains can be made. On the other hand, if the response fields are very different, then pooling would increase encoding efficiency at the expense of significant loss in spatial precision. Examining this data, we see that MF cells are extremely informative about the cue condition for a brief period of time, and no other groups neurons are anywhere near as informative. Nevertheless, these neurons are not informative later in the trial, in particular in the period of time in which the stimuli appear. While the responses from these neurons could be pooled, the informativeness of this pooling scheme would likewise be far less later in the trial than directly after cue offset.

A second question to pose is why I do not attempt to quantify how informative the cells are about the choice that the monkey will eventually make. The answer to this is that we are already fairly certain that the monkey will base his decision on the cued signal, and so we would be attempting to reduce this low level of uncertainty by observation of the spiking activity. However we are, in general, unable to reduce uncertainty about the cue position given the spike count. Since these two variables are so tightly linked, and because we have so few trials per neuron in which the monkey based his decision on the foil signals, the spike count is not informative about the actual signal on which the monkey based his decision. I verified this by

calculating the conditional information between the spike count and performance on the task given that the cue position is known, and none of the cells were informative. I have omitted this analysis in the hope of reducing the length and tediousness of an already lengthy explication.

6.3.7 Spectral properties of memory-field cells

Given that the only group of neurons which was particularly informative about the cue condition were the memory-field cells, I ask what property of these neurons made them so informative about the cue condition. Alternatively, I ask, why was the Fano Factor so low. The Fano Factor was low because the variability of spike count between trials was low. Therefore the question is, why was the spike count in memory-field cells during the initial memory period so consistent. One possible answer is that the spiking activity is highly regular. Since these neurons are identified by an oscillatory peak in the spiking power spectrum, I examined the power spectra and spectrograms from the period of time following cue offset for evidence of oscillatory spiking activity and for any differences between trials in which the monkey successfully based his decision on the cued signal and when he based it on the foil signal.

memory-field cells had broad peaks in the power spectra between 40 and 60 Hz when the monkey correctly based his decision on the cued signal and the cued signal was in the response field of the neuron (Figure 6.34). In each cell there were too few trials to determine if there was or was not a peak in the power spectra when the monkey based his decision on the foil signal. Averaged across the four examples, the power spectra when the cue was in the response field and the monkey based his decision on that signal had a peak around 60 Hz; when the monkey based his decision on the foil signal, there appeared to be a peak near 40 Hz.

Part of the reason why the peaks were broad in the spectra for correctly completed trials (and not present for trials in which the monkey based his decision on the foil) is that the frequency of the peak was not constant over the interval. Average

spectrograms appear in Figure 6.35 for trials in which the cue was in the response field and the monkey correctly based his decision on that signal and when he based it on the foil signal. Starting around cue offset, the frequency at which the peak appears seems to sweep upward from close to 25 Hz to over 50 hz prior to vanishing.

What appears to be the case for the cells on average is roughly true for each of the observed examples. The properties of these cells are tabulated in Figure 6.36. The mean firing rate is shown in the top row (A-D) for all cells. Fano Factor when the cue was in and when it was out of the response field are shown in the second row (E-H). In general, the Fano Factor was greater than 1 but dropped momentarily following cue offset. The mutual information between spike count and cue position is shown in the third row (I-L). It rose to a peak around cue offset and then decreased over the course of the trial. The spectrograms for correctly completed trials are shown in the fourth row (M-P). The frequency at which the peak of the specta occured in each of the time windows is shown in the fifth row (Q-T), and these had the property that the peak frequency swept up during the initial memory period. Finally the power ratio, or elevation of the peak above the high-frequency limiting power, appeared in the sixth row (U-X). The strength of the peak also increased over the course of the memory period before becoming inconsistent.

In addition to examining mean power spectra across conditions, I also attempted to decode whether or not the monkey based his decision on the cued signal by examining the periodicity of individual trials. As described in Methods, I performed complex regression of the multi-taper Fourier transform of each trial on a signal model consisting of a point process with a mean firing rate plus a periodic component. If the spiking activity within the interval is very regular, then it will have a large magnitude regression coefficient at a frequency best matching the periodic component of the spiking activity (Figure 6.37A). If it is not, then the component will be small at all frequencies (Figure 6.37B). The periodic component of the Fourier representation of the spike train can be subtracted out, yielding a residual spectrum that represents what the spectra of that trial would have been had there not been periodic activity

(Figure 6.37C). The ratio of the spectrum at the peak frequency to the power at the same frequency in the residual spectrum provides an F-ratio statistic that can be used as a statistical measure of the periodicity of the spiking activity. When there is no large complex regression coefficient, the residual spectrum is very similar to the original spectrum (Figure 6.37D). Finally, when the significant periodic components are removed from the spectra, as expected, the average power spectra lacks a significant peak and appears to be that of a Poisson process with a refractory period (Figure 6.37E).

Application of this procedure yielded a distributions of per-trial periodicity that was significantly different between conditions in which the monkey correctly selected the cued signal and when he did not. A scatter plot of per-trial periodicity versus the peak frequency for each cell appears in Figure 6.38. The peak frequencies were quite variable, ranging from 20 Hz to 100 Hz. There was no significant difference in the spike count in the interval (250 to 500 ms post cue offset) or the peak frequency by a two sample Kolmogorov-Smirnov test, but the per-trial periodicity was significantly greater for trials in which the monkey based his decision on the cued signal than on the foil signal. For all cells, the proportion of trials in which the periodicity exceeded a 99% confidence level was greater when the animal based his decision on the cued signal than on the foil signal.

6.4 Discussion

In this investigation I sought to characterize the responses of functional groups of neurons within the primate SC during the selective attention task which we had used show that inactivation of the SC caused visual extinction. These neurons included the groups as described in Chapter 5: the visually-triggered movement cells (VTM), visually responsive saccade bursting cells with prelude activity (VSBP), visually responsive saccade bursting neurons (VSB), buildup cells (BU), memory-field cells (MF), and saccade bursting cells (SB). I investigated the discharge of the neurons

evoked by events in the selective attention task and how that activity was altered by the placement of the cued signal either in or directly out of the neurons' response fields; I investigated how these responses were related to animal behavior; and I investigated how well these signals could be used in encoding information about the cue position.

Overall, the changes in mean firing rate associated with the cue condition seen in both the evoked responses and the maintained firing rate prior to stimulus onset are broadly consistent with the results of previous studies, in particular recent work on covert attention in primates (Kustov and Robinson, 1996; Ignashchenkova et al., 2004; Fecteau et al., 2004). Taken with other work, it extends the long-standing observation that the discharge rate of neurons in the SC is enhanced by the shift of attention to a location preceding a saccade (Goldberg and Wurtz, 1972b; Wurtz and Mohler, 1976; Mohler and Wurtz, 1976). The fact that neither the evoked responses nor the sustained firing rate during the delay period, particularly that of buildup units, is informative of the cue condition begs the question of what, exactly the purpose of either these transient events or the sustained difference in firing rate might be. The main findings of the work, organized by cell group, were as follows.

6.4.1 Visually-triggered movement cells

VTM neurons had large discharges evoked by the onset of the cues and the patches, and the peak firing rate of this discharge was greater when the cued signal was in the response field of the neuron than when it was not. During the delay period, the firing rate of the neurons rose, and the magnitude of this maintained discharge was greater when the cued signal was in the response field. The neurons produced a large response evoked by the onset of the coherent motion pulses. This response was briskest and of highest magnitude for the high-strength instructor trials and greater for the cued signal than the foil signal. Furthermore, the response was earlier for the cued signal than the foil signal. Following this response, these neurons did not produce a saccade selective discharge. For each of these epochs, successful completion of the

task was associated with slightly lower firing rates, and errors were associated with slightly higher firing rates for the foil signal. Nevertheless, these responses were not particularly informative about the cued position because the responses were highly variable. This suppression for correct responses and exaggeration for choices with the foil would seem to indicate that the neurons were actively suppressed, as if the psychomotor set of the animal was not to allow them to drive behavior during the trial and failures were associated with the transient response breaking through.

6.4.2 Visually responsive saccade bursting cells with prelude activity

VSBP neurons also produced large discharges evoked by the onset of the cues and the patches, and the peak firing rate of the discharge for cue onset was greater when the cued signal was in the response field of the neuron than when it was not. The magnitude of evoked responses to the patches were not particularly dependent on the cue condition. During the delay period, the firing rate of the neurons rose, and the magnitude of this maintained discharge was greater when the cued signal was in the response field. Like the VTM neurons, these cells produced a large discharge evoked by the onset of the coherent motion pulses; while not as large or as rapid in rise as that of the VTM neurons, it was largest for the instructor trials, then the cued, and then the foil signals. These neurons produced vigorous bursts in response to saccades made into their response fields. In these neurons, the magnitude of the responses was slightly higher prior to the motion pulses and in response to the motion pulse if the monkey based his decision on the signal in the neurons' response fields. Despite the large differences in maintained firing rate prior to the onset of the motion stimuli, the variability between trials was so high that these neurons were only marginally informative about the cue condition.

6.4.3 Visually responsive saccade bursting cells without prelude activity

VSB neurons did not produce responses evoked by the onset of the cues, but did produce a response evoked by the onset of the patches. This response was not modulated by the cue condition. The neurons produced large maintained discharges during the delay period which were higher when the cued position was in their response fields than when it was not, and the neurons produced a large discharge evoked by the onset of the motion pulse. This discharge was higher when the cued signal was in the response field than when it was not, but unlike VTM and VSBP cells, it did not appear to depend on the strength of the signal and only on the presence or absence of a signal in the neurons' response field. VSB neurons produced vigorous bursts in response to saccades made into their response fields. As for VSBP neurons, their firing rate was marginally higher if the monkey based his decision on the signal in their response field during the period of time immediately preceding the motion pulses and during the response evoked by the motion pulse. Nevertheless, the variability in these responses was so high that the neurons were not particularly informative about the cue condition.

6.4.4 Buildup cells

Buildup neurons did not produce evoked responses for either the cue onset or the patch onset. Following cue offset, their firing rate rose and was higher when the cued position was in their response field than when it was not. They produced the earliest response to the onset of the motion pulse, and like the VSB neurons, this response did not depend on the strength of the motion pulse, only on its presence or absence. As for the other neurons, the firing rate was so variable between trials that the buildup neurons were not particularly informative about the cue condition. Buildup neurons produced large discharges for saccades made into their response fields. The firing rate appeared to pass through a fixed point prior to saccade onset,

which might be expected if the firing rate of these neurons reflects an evolving decision process which must pass through a particular trajectory en route to the decision about the target of the saccade and the initiation of the movement.

6.4.5 memory-field cells

Finally, MF cells began firing shortly after the onset of the patches, rose to their peak firing rate at the end of the memory period following cue offset, and then slowly tapered off over the course of the trial. Their peak firing rate occurred at the first moment that the motion pulses could appear. These neurons did not produce evoked discharges in response to the onset of the motion pulses. During the initial memory period, the variability of the firing rate across trials dropped substantially when the cued position was in the neurons' response fields, and the difference in firing rate between cue conditions was large. These properties made these neurons extremely informative about the cue condition prior to the appearance of the coherent motion pulses. The same period of time in which the neurons were extremely informative about the cue condition was associated with the emergence of a broad peak in the power spectrum of their spiking activity, and in the spectrogram this emerged as a peak that started around 25 to 30 Hz and swept up to around 60 Hz. Moreover, individual trials tended to be more periodic, based on a rough measure of how consistently spaced spikes were in an interval, when the monkey based his decision on the signal in the neurons' response fields.

6.4.6 Relationship between activity prior to saccade onset and cue condition

In addition, I examined how the change in firing rate based on cue condition changed the relationship of the saccade-related discharge to the saccade. For VSBP and VSB neurons, the direction of the saccade became evident at slightly higher firing rates and slightly earlier when the cued position was in the neurons' response

fields than when it was not. However these differences were extremely small and could be artifactual; the ability to discriminate the distribution of firing rates prior to the saccade could be increased at higher firing rate. Nevertheless, it appears as if the modulation in the neurons' firing rates based on cued position was effectively added on to the neurons' discharge preceding the saccade. In contrast, in buildup cells the time at which the direction of the saccade became discriminable and the firing rate at which this point occurred was invariant to the cue condition, as if the neurons had to pass through the same fixed firing rate point prior to selection of the saccade target. This property is similar to that seen in some neurons in the FEF (Hanes and Schall, 1996), and would seem to be consistent with buildup neurons reflecting the evolving selection process in their firing rate. The time at which the direction was discriminable in the responses of MF cells was also the same across cue conditions but occurred at substantially higher firing rates when the cued position was in the neurons' response fields; indeed, the elevation in firing rate persisted during and through the saccade-related discharge for these neurons. This difference between MF cells and buildup units would indicate that they are indeed separate functional classes and bear entirely different functional connections to the initiation of saccades and the selection of their targets.

6.4.7 Visual transients responses and the doctrine of prior entry

Boehnke and Munoz (2008) have suggested that the visual transient response of neurons in the SC is a priority signal which is transmitted from the SC into recipient cortical structures via connections such as the inferior pulvinar (Kaas and Lyon, 2007), and to the dopaminergic reward system via direct connections into the substantia nigra (Krout et al., 2001; Comoli et al., 2003). This priority signal is defined as a mixing of bottom-up salience (i.e. Stimulus related) information with the top-down relevance an event has for the observer based on expectations and prior experience. Such a definition invokes a familiar idea but defines it in an ambiguous

way in that it is unclear how relevance and prior experience are not stimulus related information. Really, salience is a property of the visual system when applied to a stimulus, and not of stimuli themselves. Furthermore, the word "priority" can only properly be understood when it is clarified what the priority is for. We might modify the characterization of the transient as a "priority signal" to be a signal which has only moderate shaping of the outcome of initial visual processing by contingent features such as expectations and prior experiencing. Therefore a priority signal should transmit the appearance of a *potentially* behaviorally important stimulus, allowing other areas which are concerned with other computations to apply other contextual factors when interpreting that transient.

Control of these transients for the purpose of influencing selection of signals to guide behavior would come not in changing their magnitude so much as changing their order; if the primary limitation in processing is the number of items that can be simultaneously processed, then control of access into the selection process would be of utmost importance in controlling which stimuli guide behavior (e.g. Olivers, 2008). Broadly speaking, this would appear to be an accurate description of the visual responses evoked by the cue onset and the motion pulse onset in the VTM neurons, which have already been proposed as the ideal candidate for the regulation of visually-guided orienting of both saccades and attention (Wallace et al., 1990). Furthermore, this observation is broadly consistent with the doctrine of prior entry, which holds that attended stimuli reach visual awareness prior to unattended stimuli (Shore et al., 2001), and the observation that changing the delay between stimulus onsets can overcome visual extinction effects in human patients (Rorden et al., 1997) or in monkeys following FEF lesions (Schiller and Chou, 2000a,b, 1998). Temporal ordering has been proposed as a fundamental mechanism of visual neglect (Chou and Schiller, 1999), even more important than underactivation of detection processes, and loss of VTM neuron responses as in the lesion could be the root of visual extinction subsequent to SC inactivation.

If the VTM cells are transmitting a priority signal in their transient re-

sponses, then by definition the response would not be expected to be shaped by the cue condition. Instead, the order matters. Quantitatively speaking, the magnitude of the transient response should be highly informative about the spatial location and the time of onset of the visual event, but not be informative about the cue condition. Only the relative timing of events should be informative about the importance of the stimuli. This assessment seems to be well-matched to our observations of the visual transient response of VTM, VSBP, and VSB neurons; rather than being informative about the cue location, they are informative about the onset of the stimulus by dint of their temporally precise discharge, and they are informative about the spatial location of the stimulus because of their small and spatially circumscribed response field (see Chapter 5). The perennial question of how can a transient response be both a priority signal and a saccade-related discharge could be answered by the simple proposition that it is notthe VTM cells provide the priority signal only, and the discharge seen in VSBP and VSB neurons, which lags that of the VTM cells by 10 ms, is the consequence of that same signal on oculomotor planning in the SC. This conjecture is also supported by the fact that VTM and VSBP cells had substantially larger responses to the high-strength instructor trials than to the trials in which the cued and foil signals were matched. Furthermore, If the primary functional role of the VTM cells is as a transmitter of a transient priority signal, then the variability of the maintained discharge during the delay period could be explained. These neurons normally produce transient responses, not phasic responses, but they are continuously driven by the stochastic motion stimuli. Hence, their response is chattery and inconsistent in time, leading to large variability across trials on short time scales. The same logic could apply to the responses of VSBP and VSB cells, and even the buildup units.

This begs the question of why VTM neurons produce saccade-related discharges on a visually-guided saccade task but not on a memory-guided saccade task nor in this task; perhaps it is because the discharge is not saccade-related per se, but instead related to the sudden change in the priority of the visual target when it becomes the saccade goal, whereas in the selective attention task, the motion pulses

themselves are the high priority events which determine the target of the saccade.

Another unexplained aspect is that the response of the VTM cells evoked by motion onset occurs at a latency of 115 ms, whereas the latency of response to motion stimuli in area MT is as low as 88 ms(Raiguel et al., 1999; Schmolesky et al., 1998). While this difference in timing would seem to be a necessity since the evoked response of the VTM cells likely depends on cortical input, it also makes It substantially unclear what functional role the discharge could play if it is occurring almost after the stimulus appears on the screen and its presence is already old news in cortex. A conservative explanation is that the signal facilitates the association of an event with a reward by alerting the dopaminergic reward system that a decision has been made. A mostly speculative possibility is that the transient response is necessary to signal that an important event *has already* occurred rather than is *currently* occurring. Models of decision as diffusion to bound have met with great success in predicting both performance and reaction time on motion discrimination tasks Palmer et al. (2005); Churchland et al. (2008), but a potential fatal flaw in these models that limits their applicability to any realistic situation is that the point at which the diffusion process *starts* is entirely experimentally set. If the animals had no idea when the stimuli were appearing, then how could the accumulation of evidence even begin? Or, if the diffusion process was constantly ongoing, only resetting at saccade onset, then how would irrelevant evidence ever be forgotten? This problem is similar to one already solved in the engineering and economics literature, in which it is known as the change or fault detection problemLai (2000); Baum and Veeravalli (1994a,b) and the Bermudan look-back option pricing problem, respectively (Guo and Liu, 2005). In these problems, a event occurs at an unpredictable time in a system being continuously monitored, and a response is required upon detection of the event. The event could occur at any moment, may be relatively unlikely, and the cost of missing the event or reacting prior to the event is extremely costly. Given this scenario, the optimal solution involves maintaining a memory buffer to hold a portion of the signal. Once the evidence in the buffer reaches a threshold, specifically, once the evidence

for one of the potential changes is sufficiently divergent from the rest, then the event is detected. The error introduced by a buffer length is vanishingly small at finite buffer lengths. Perhaps, then, the evoked response of the VTM cells signals that the event has been detected, and the result of the motion processing must be poled from working memory

6.4.8 Sustained activity in buildup units is not a biasing signal for attention

The temporally and spatially specific transient responses of the VTM, VSBP, and VSB cells may be at least in part explained, but the non-informative, maintained activity of the buildup units is not. The buildup units have historically been one of the prime candidates for a role in selective attention because of their role in the selection of targets for saccadic eye movements (Kustov and Robinson, 1996). The maintained activity of buildup neurons has been related to the probability of a saccade made to a location in space (Basso and Wurtz, 1997, 1998), but as we saw in examination of the behavioral data in Chapter 4, the monkeys are not biased to produce responses into the cued location. Hence this cannot be an explanation for that activity. I suggest that the variability in activity may, in fact, be functionally critical for preventing a complete conflation of oculomotor planning and attention. Just as is seen for some neurons in FEF that appear to be involved in the decision about movement direction (Hanes and Schall, 1996), the buildup units appear to pass through a consistent fixed point in firing rate prior to saccade initiation. So what stops them from initiating a saccade if they reach that same firing rate during the delay period? In addition to local excitatory connections, inhibitory interneurons link cells in the SC (Isa and Hall, 2009). Perhaps the network of buildup units is maintained in an anti-correlated, mutually inhibiting state during attention so that the population as a whole does not reach the trigger point for a saccade. At any given moment when one buildup unit begins to fire, it suppresses its neighbors and the local excitatory positive feedback loop cannot induce recruitment of its neighbors to initiate a saccade burst, even

though average firing rate across the population at any given moment is higher with attention than without. This explanation would be consistent with the observation that synchronous stimulation of the SC at more than one site is necessary to produce a saccade which is the vector average of the saccades evoked by stimulation of each site alone, whereas as asynchronous stimulation leads to saccades which are the sum of the vectors alone (Brecht et al., 2004). This result suggests that synchronous firing between neurons in the SC is necessary for the normal readout of SC population activity.

Another aspect of the buildup units which was unique was that they responded to the onset of the motion stimuli at approximately the same time as the VTM cells. It is the case that cells in the intermediate and deep layers receive input from a variety of cortical structures other than prestriate visual cortex and area MT, which project into SGS and into the vicinity of the VTM cells near stratum opticum (Collins et al., 2005). Some of this input from area FEF and LIP has been identified as visual in nature (Sommer and Wurtz, 2001; Paré and Wurtz, 2001, 1997). It stands to reason that the input they receive is related to the initiation of saccade planning and is very different from the priority signal of VTM cellsit has been stripped of the salience aspect and reflects only the unshaped relevance of the signals for signaling the location of potential eye movements. This is somewhat similar to another suggestion that buildup neurons are primarily reflective of top-down biasing rather than bottum-up control (Bell and Munoz, 2008).

6.4.9 memory-field cells encode memories for cued locations

Lastly, the memory-field cells were the most informative of all groups of neurons about the cue conditions. They were most informative early in the trial, when the monkey had to commit to memory where to pay attention, and prior to the appearance of any stimuli. In addition, their response to the saccade seemed to interrupt a maintained difference in firing rate as if it were superimposed on top of an ongoing activity level, and the difference in firing rate based on cue condition

persisted past the saccade. I described these neurons in Chapter 3 based on their oscillatory activity in relation to spatial working memory; perhaps here they are performing essentially the same function. These cells are involved in establishing the working memory for the cued position and are most important early in the trial, after which point the memory for the cued location has been relocated elsewhere or is distributed in a less visible way throughout the circuit. In particular, I would speculate that the synchronous activity in the gamma band between frontal cortex and visual area V4 recently observed in awake behaving monkeys performing an attention task Gregoriou et al. (2009) is actually related to spatial working memory for the cue instead of the changes in tuning properties of cells in V4 as a consequence of attention. This synchronous activity could be necessary to maintain the coordination of representations across multiple areas. Perhaps the SC plays a part in initiating these long range connections and thus is involved in organizing working memory for regions of space.

6.5 Conclusions

Overall, the recording results suggest that neurons in the SC participate in the selection of stimuli for guiding behavior in multiple ways. First, the visual transient responses as in the VTM cells alert both cortical and subcortical areas that a biologically important event has occurred. The magnitude of the response was related to the physical signal strength, and the timing of the response may have reflected the relative importance of the signals on the screen as if it could be involved in controlling the sequence of processing of stimuli. Second, despite changes in firing rate related to the cued position, buildup units appeared to be primarily involved in the selection of targets for saccades rather than encoding a cue position. Third, memory-field cells were the most informative about the cue condition and may be necessary for encoding working memory about which location has been cued.

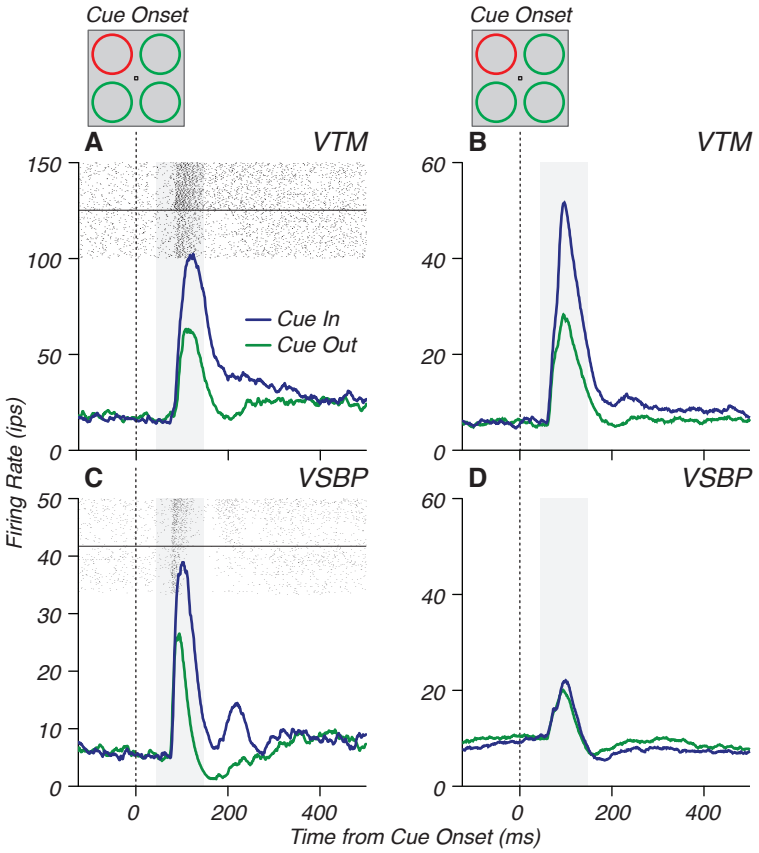

Figure 6.1: Example and average responses to cue onset. A) Raster and mean firing rate of response evoked by cue onset for VTM cell (f171207). Rasters for trials in which the cue appeared in the response field are above those in which the cue appeared in the diametrically opposite location. Mean firing rate for cue in trials is blue and green for cue out. Gray bar represents interval over which mean responses are measured, 50-150 ms after cue onset. B) Average response to cue onset for VTM cells. C) Raster and mean firing rate for VSBP cell (m011008). D) Average response to cue onset for VSBP cells.

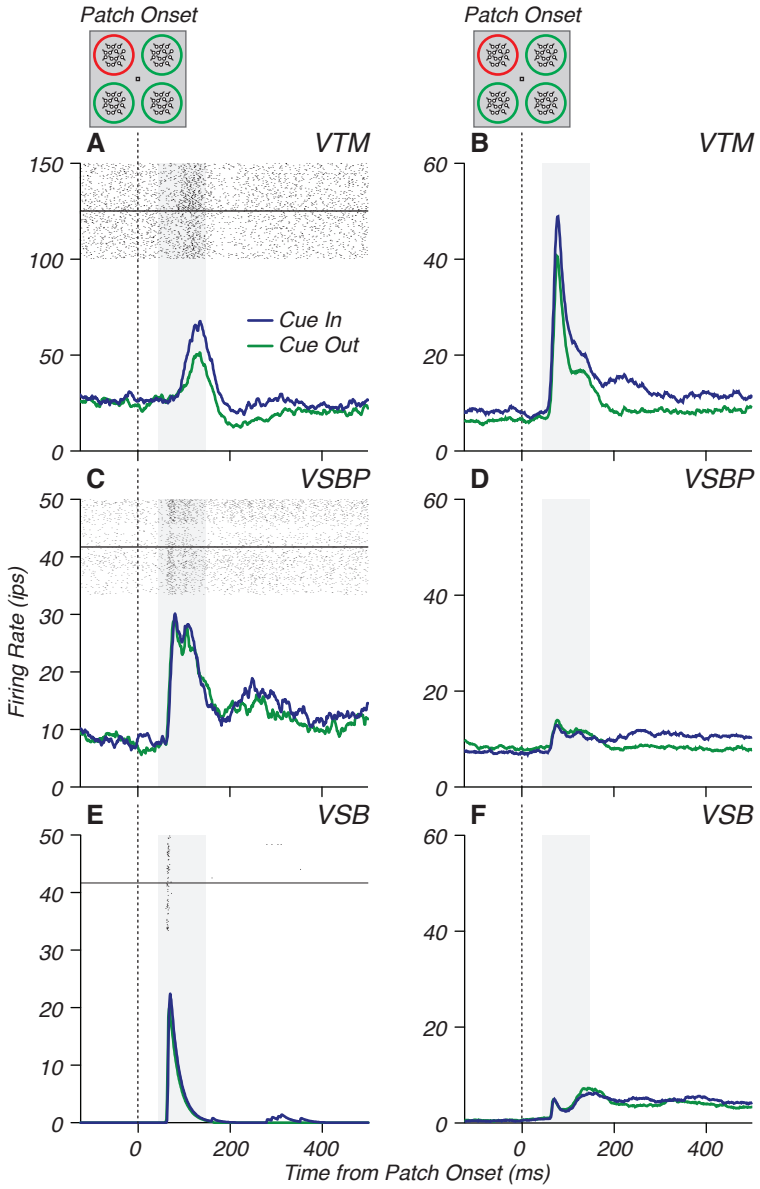

Figure 6.2: Example and average responses to patch onset. A) Raster and mean firing rate of response evoked by patch onset for VTM cell (f171207). Rasters for trials in which the cue appeared in the response field are above those in which the cue appeared in the diametrically opposite location. Mean firing rate for cue in trials is blue and green for cue out. Gray bar represents interval over which mean responses are measured, 50-150 ms after patch onset. B) Average response to patch onset for VTM cells. C) Raster and mean firing rate for VSBP cell (m011008). D) Average response to cue onset for VSBP cells. E) Raster and mean firing rate for VSB cell (f180607b). F) Average response to patch onset for VSB cells.

Figure 6.3 *(following page)*: Example responses to cue offset A) Raster and mean firing rate of response following cue offset for VTM cell (f171207). Rasters for trials in which the cue appeared in the response field are above those in which the cue appeared in the diametrically opposite location. Rasters are sorted based on movie duration and include only that portion of the trial preceding the motion pulse. Mean firing rate for cue in trials is blue and green for cue out. At any given moment, the mean is taken over only those trials still preceding the motion pulse. Gray bar represents interval over which mean responses are measured, 250-500 ms after cue offset. Dashed line at 960 ms is mean motion pulse onset time with respect to cue offset. B) Raster and mean firing rate following cue offset for VSBP cell (m011008). C) Raster and mean firing rate following cue offset for VSB cell (f180607b). D) Raster and mean firing rate following cue offset for BU cell (f150807). E) Raster and mean firing rate following cue offset for MF cell (m301208). F) Raster and mean firing rate following cue offset for SB cell (m190908).

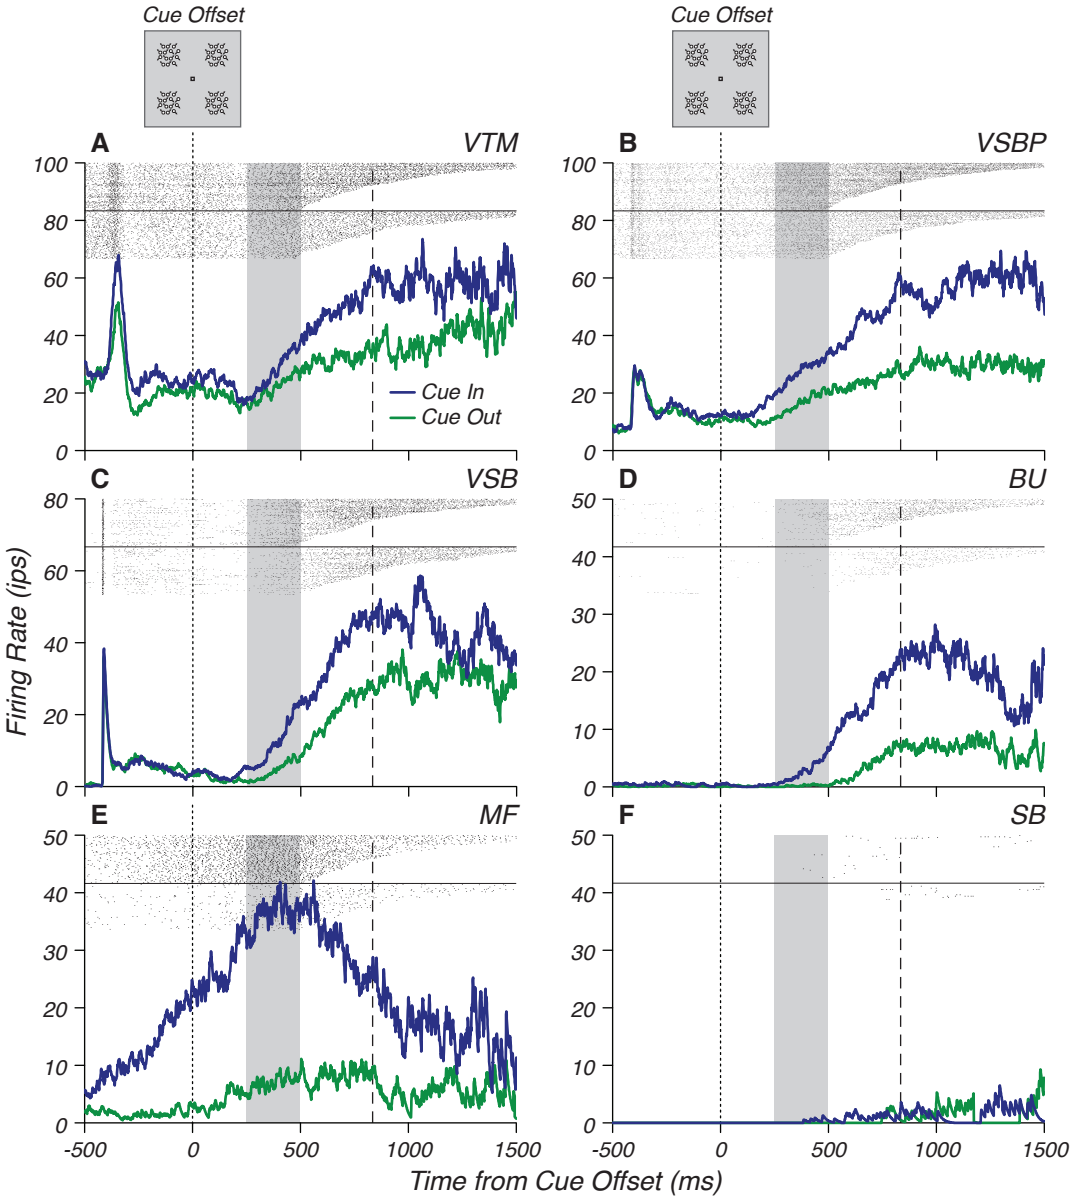
Cue Offset
Cue Offset
A
VTM
B
VSBP
C
VSB
D
BU
E
MF
F
SB
Cue In
Cue Out
Firing Rate (ips)
Time from Cue Offset (ms)
-500
0
500
1000
1500

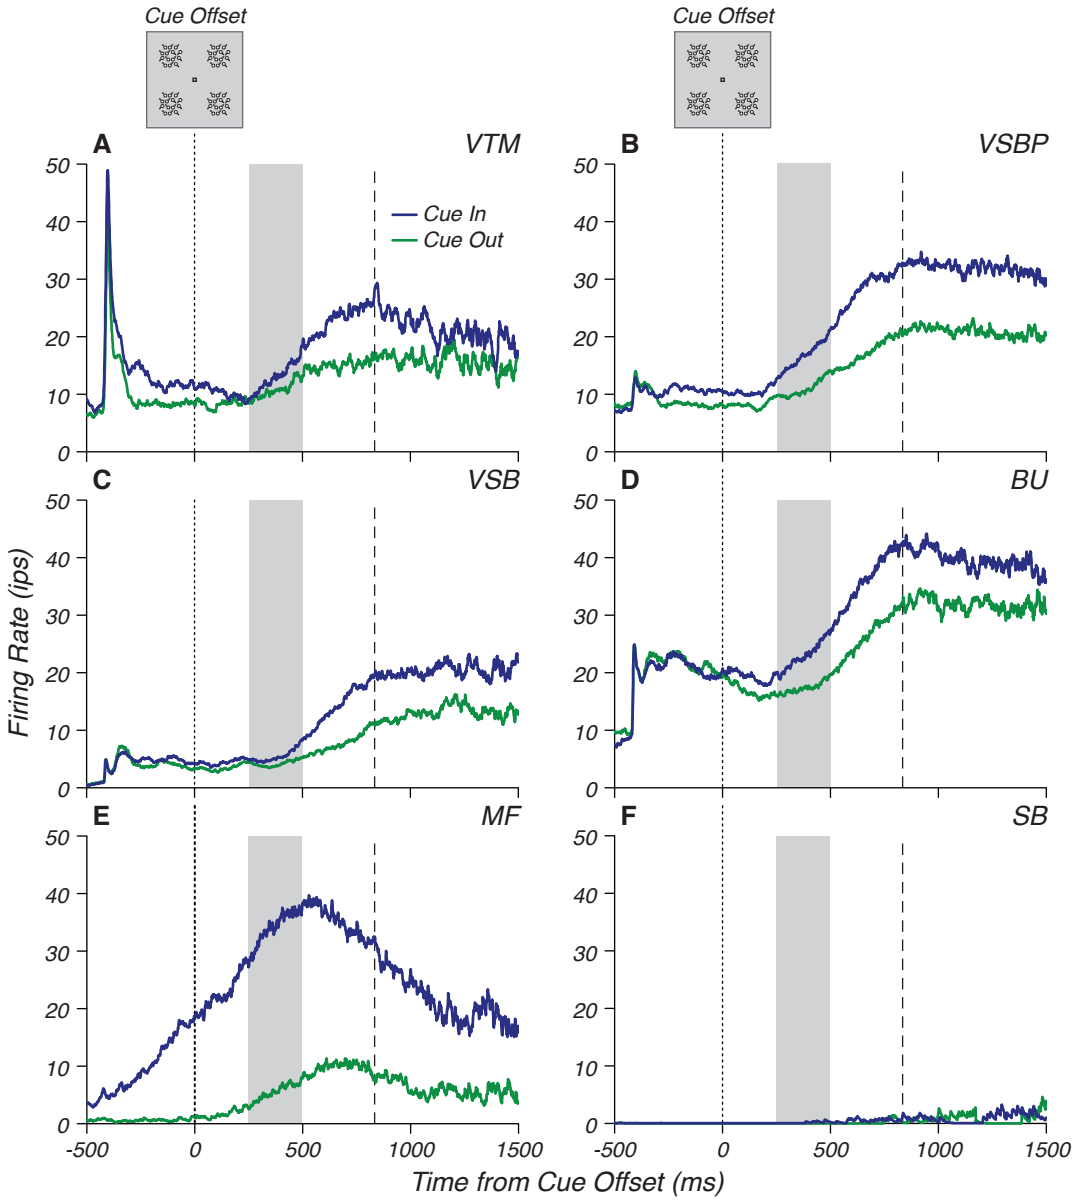

Figure 6.4: Average responses to cue offset. A) Average response following cue offset for VTM cells. Mean firing rate for cue in trials is blue and green for cue out. At any given moment, the mean firing rate includes only those trials in each cell still preceding the motion pulse. Gray bar represents interval over which mean responses are measured, 250-500 ms after cue offset. Dashed line at 960 ms is mean motion pulse onset time with respect to cue offset. B) Average response following cue offset for VSBP cells. C) Average response following cue offset for VSB cells. D) Average response following cue offset for BU cells. E) Average response following cue offset for MF cells. F) Average response following cue offset for SB cells.

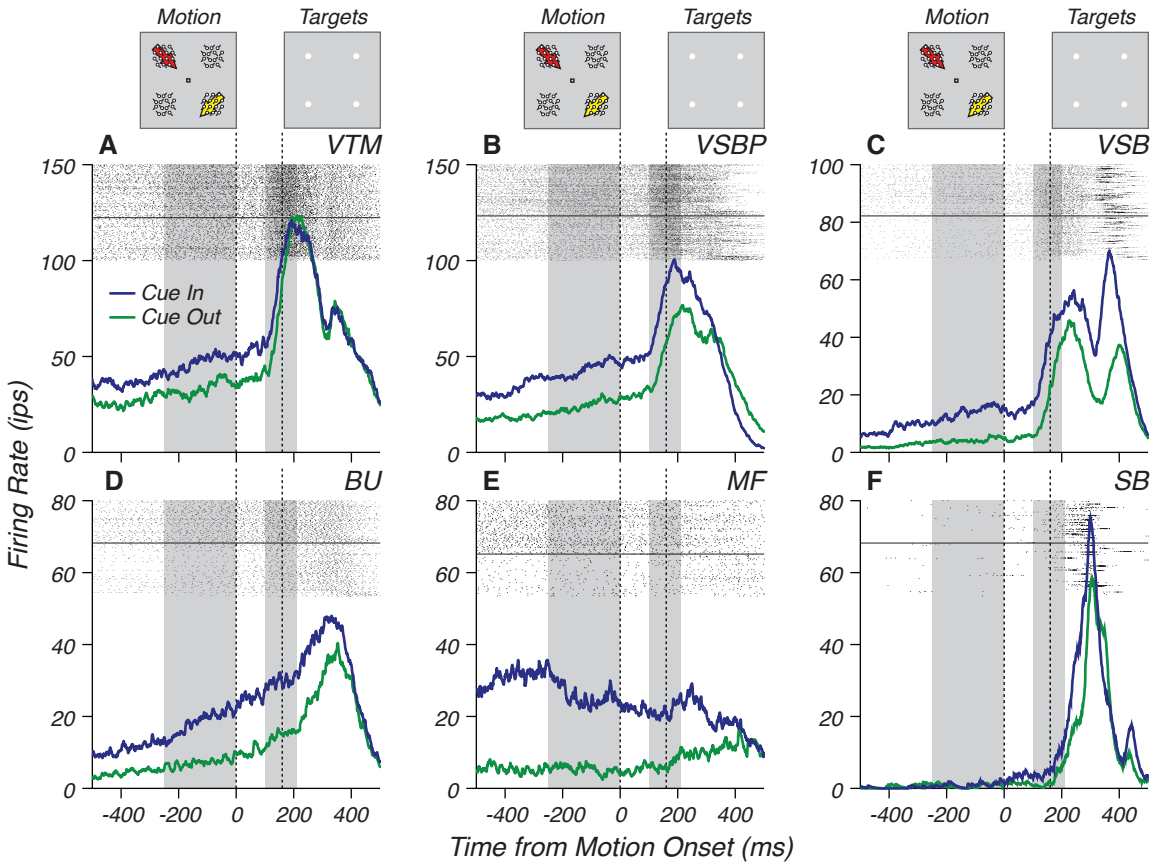

Figure 6.5: Example responses preceding and evoked by motion onset A) Raster and mean firing rate of response preceding and evoked by motion onset for VTM cell (f171207). Rasters for trials in which the cue appeared in the response field are above those in which the cue appeared in the diametrically opposite location. Mean firing rate for cue in trials is blue and green for cue out. Gray bars represents intervals over which mean responses are measured, 250 ms prior to motion onset and 100-200 ms after motion onset. B) Raster and mean firing rate of response for VSBP cell (m011008). C) Raster and mean firing rate of response for VSB cell (f180607b). D) Raster and mean firing rate of response for BU cell (f150807). E) Raster and mean firing rate of response for MF cell (m301208). F) Raster and mean firing rate of response for SB cell (m190908).

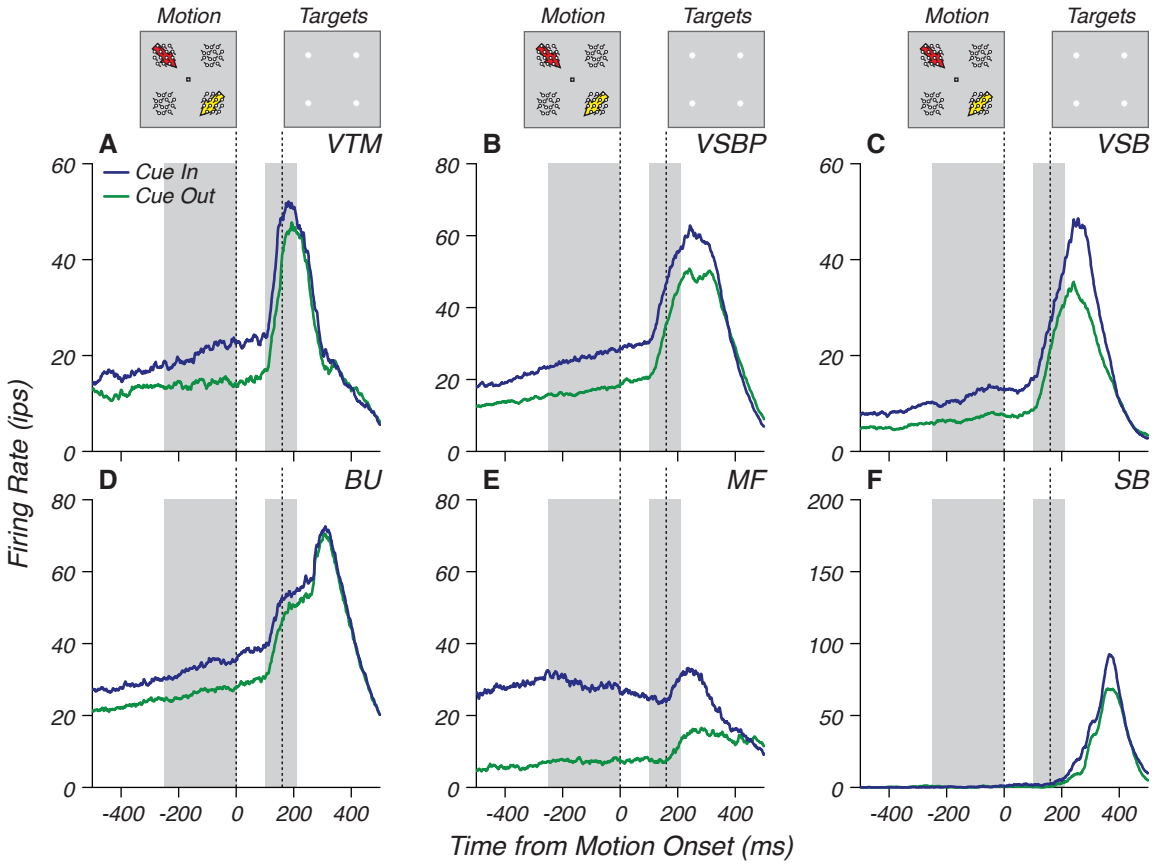

Figure 6.6: Average responses preceding and evoked by motion onset. A) Average response preceding and evoked by motion onset for VTM cells. Mean firing rate for cue in trials is blue and green for cue out. Gray bars represents interval over which mean responses are measured, 250 ms prior to motion onset and 100-200 ms after motion onset. B) Average response for VSBP cells. C) Average response for VSB cells. D) Average response for BU cells. E) Average response for MF cells. F) Average response for SB cells.

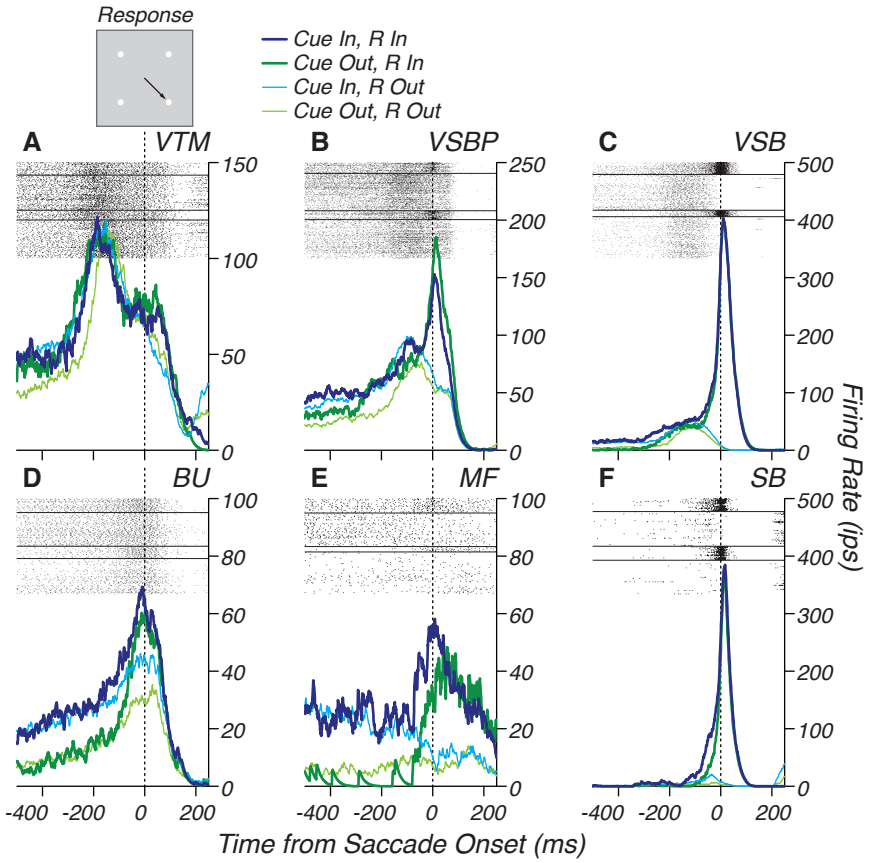

Figure 6.7: Example responses preceding and evoked at saccade onset A) Raster and mean firing rate of response preceding and evoked at saccade onset for VTM cell (f171207). Rasters for trials in which the cue appeared in the response field are above those in which the cue appeared in the diametrically opposite location, and within these groups, rasters for trials in which the saccade was into the response field are above those in which the saccade is out of the response field. Mean firing rate for cue in trials is blue and green for cue out. B) Raster and mean firing rate of response for VSBP cell (m011008). C) Raster and mean firing rate of response for VSB cell (f180607b). D) Raster and mean firing rate of response for BU cell (f150807). E) Raster and mean firing rate of response for MF cell (m301208). F) Raster and mean firing rate of response for SB cell (m190908).

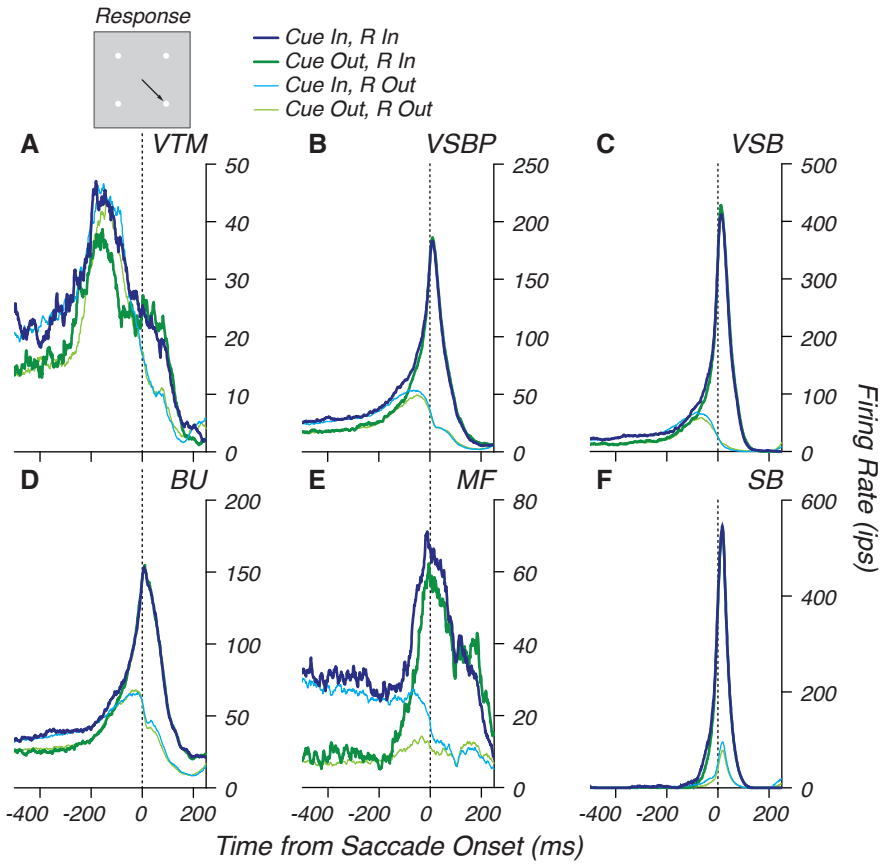

Figure 6.8: Average responses preceding and evoked at saccade onset. A) Average response preceding and evoked at saccade onset for VTM cells. Mean firing rate for cue in trials is blue and green for cue out. B) Average response for VSBP cells. C) Average response for VSB cells. D) Average response for BU cells. E) Average response for MF cells. F) Average response for SB cells.

Figure 6.9 *(following page)*: VTM cells have a response evoked by the motion pulse onset but not a response to the saccade. A) Response onset times for all VTM cells sorted based on the position of the cue and whether the response latency was shorter or longer than the median. Individual symbols correspond to individual neurons. Black lines and error bars are means and standard error. Gray bar represents interval over which mean response is measured, 100-200 ms after onset. B) Example of response evoked by motion pulse (f171207). Black dashed line represents the expected firing rate of the neuron if the cue is in the response field and had a motion pulse not occurred. Gray dashed line represents the expected firing rate of the neuron if the cue is out of the response field and had a motion pulse not occurred. Whisker plots represent the saccade time with respect to the motion pulse onset. Vertical dashed line marks the appearance of the choice dots. C) Average response evoked by motion pulse. D-E) Example (f171208) and average response evoked by motion pulse for short-latency trials. F-G) Example (f171208) and average response evoked by motion pulse for short latency trials.

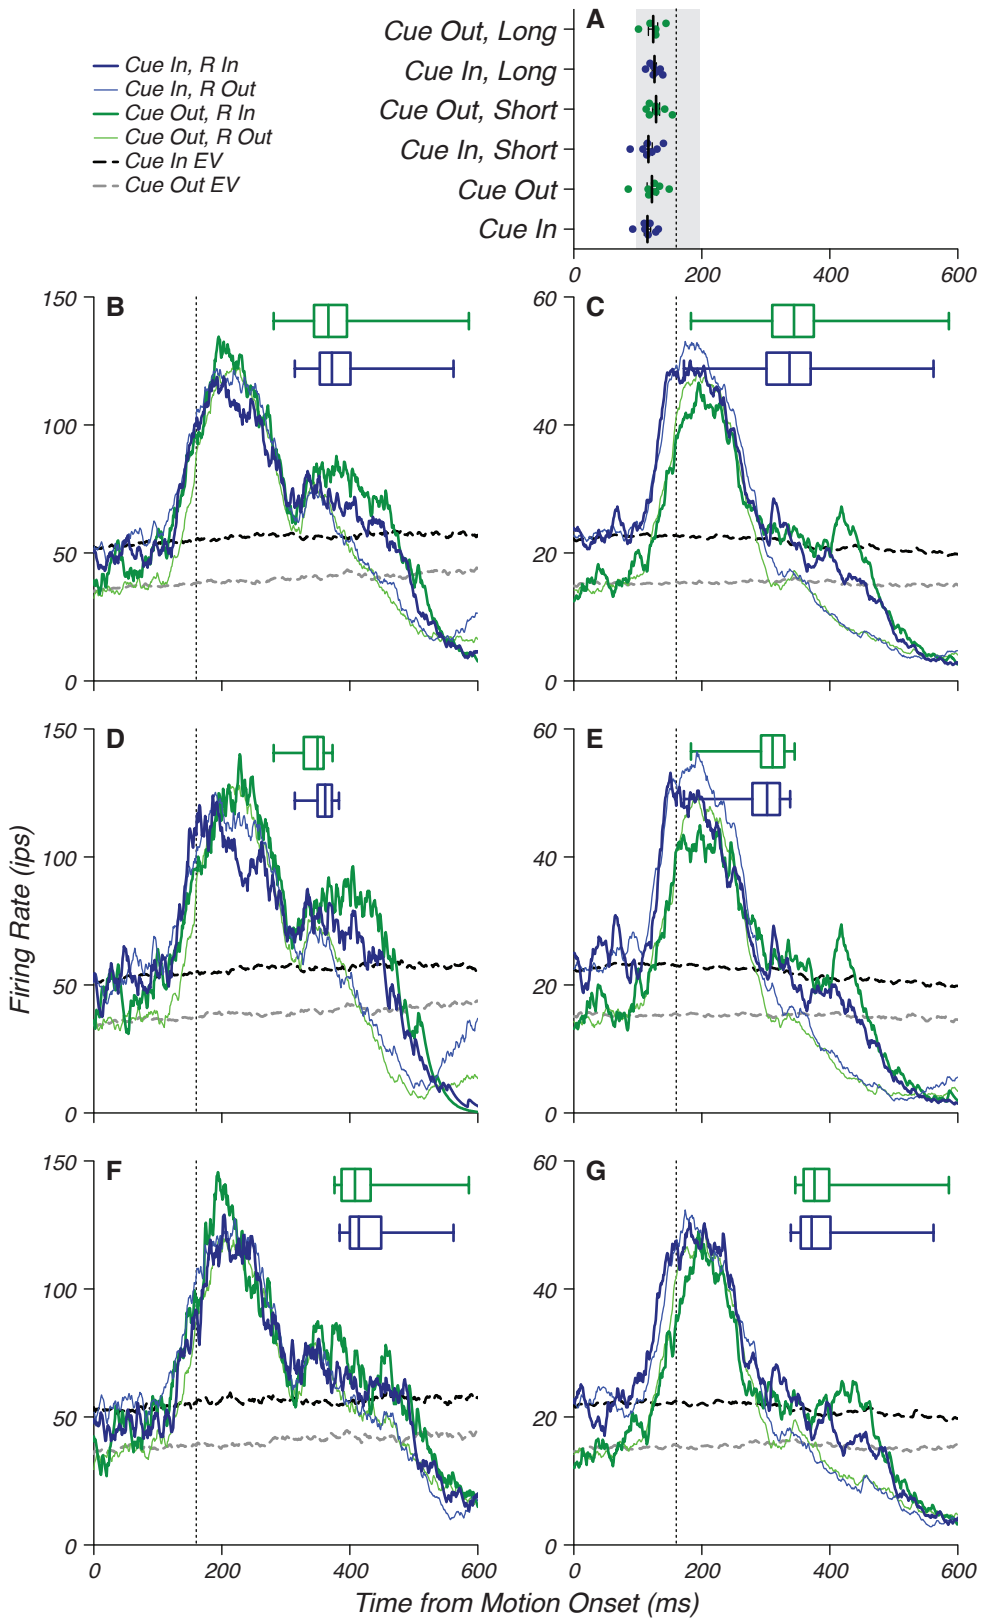
Cue In, R In
Cue In, R Out
Cue Out, R In
Cue Out, R Out
Cue In EV
Cue Out EV
A
Cue Out, Long
Cue In, Long
Cue Out, Short
Cue In, Short
Cue Out
Cue In
0
200
400
600
B
C
D
E
F
G
150
100
50
0
60
40
20
Firing Rate (ips)
Time from Motion Onset (ms)

Figure 6.10 *(following page)*: VSBP cells have a response evoked by the motion pulse onset which precedes the response to the saccade onset. A) Response onset times for all VSBP cells sorted based on the position of the cue and whether the response latency was shorter or longer than the median. Individual symbols correspond to individual neurons. Solid symbols are onset times for a direction-insensitive motion response and open symbols are the onset times for a saccade selective response. Black lines and error bars are means and standard error. Gray bar represents interval over which mean response is measured, 100-200 ms after onset. B) Example of response evoked by motion pulse (m011008). Black dashed line represents the expected firing rate of the neuron if the cue is in the response field and had a motion pulse not occurred. Gray dashed line represents the expected firing rate of the neuron if the cue is out of the response field and had a motion pulse not occurred. Whisker plots represent the saccade time with respect to the motion pulse onset. Vertical dashed line marks the appearance of the choice dots. C) Average response evoked by motion pulse. D-E) Example (m011008) and average response evoked by motion pulse for short-latency trials. F-G) Example (m011008) and average response evoked by motion pulse for short latency trials.

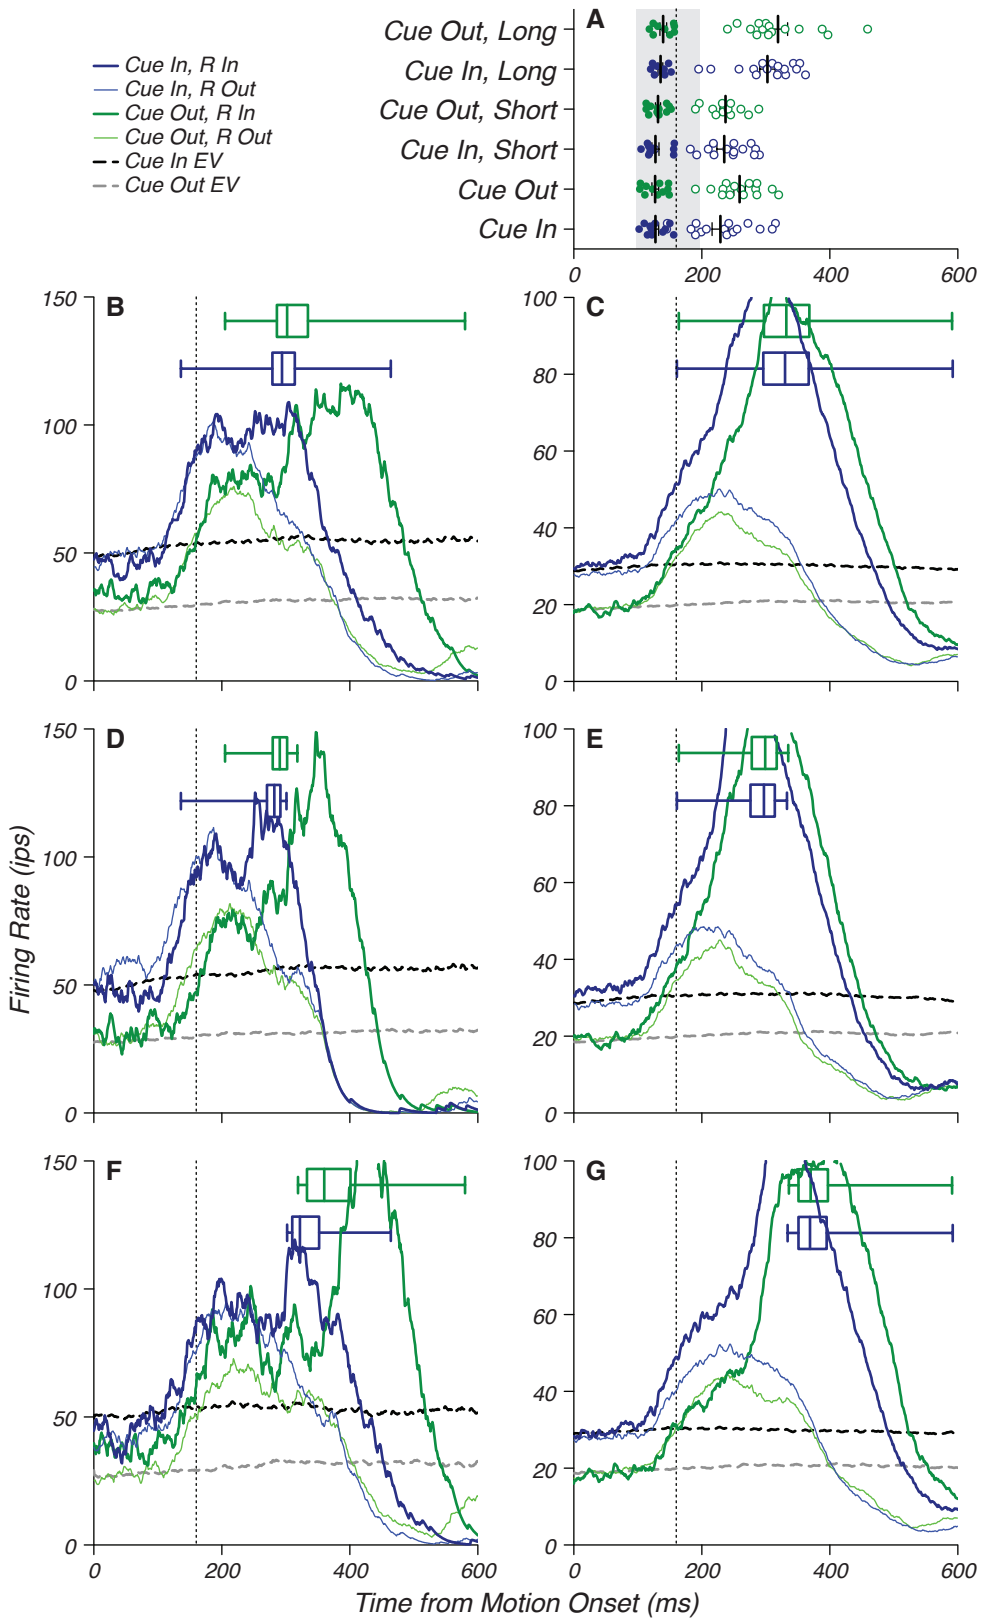
Cue In, R In
Cue In, R Out
Cue Out, R In
Cue Out, R Out
Cue In EV
Cue Out EV
A
Cue Out, Long
Cue In, Long
Cue Out, Short
Cue In, Short
Cue Out
Cue In
0
200
400
600
B
C
D
E
F
G
150
100
50
0
80
60
40
20
Firing Rate (ips)
Time from Motion Onset (ms)

Figure 6.11 *(following page)*: VSB cells have a response evoked by the motion pulse onset which precedes the response to the saccade onset. A) Response onset times for all VSB cells sorted based on the position of the cue and whether the response latency was shorter or longer than the median. Individual symbols correspond to individual neurons. Solid symbols are onset times for a direction-insensitive motion response and open symbols are the onset times for a saccade selective response. Black lines and error bars are means and standard error. Gray bar represents interval over which mean response is measured, 100-200 ms after onset. B) Example of response evoked by motion pulse (f180607b). Black dashed line represents the expected firing rate of the neuron if the cue is in the response field and had a motion pulse not occurred. Gray dashed line represents the expected firing rate of the neuron if the cue is out of the response field and had a motion pulse not occurred. Whisker plots represent the saccade time with respect to the motion pulse onset. Vertical dashed line marks the appearance of the choice dots. C) Average response evoked by motion pulse. D-E) Example (f180607b) and average response evoked by motion pulse for short-latency trials. F-G) Example (f180607b) and average response evoked by motion pulse for short latency trials.

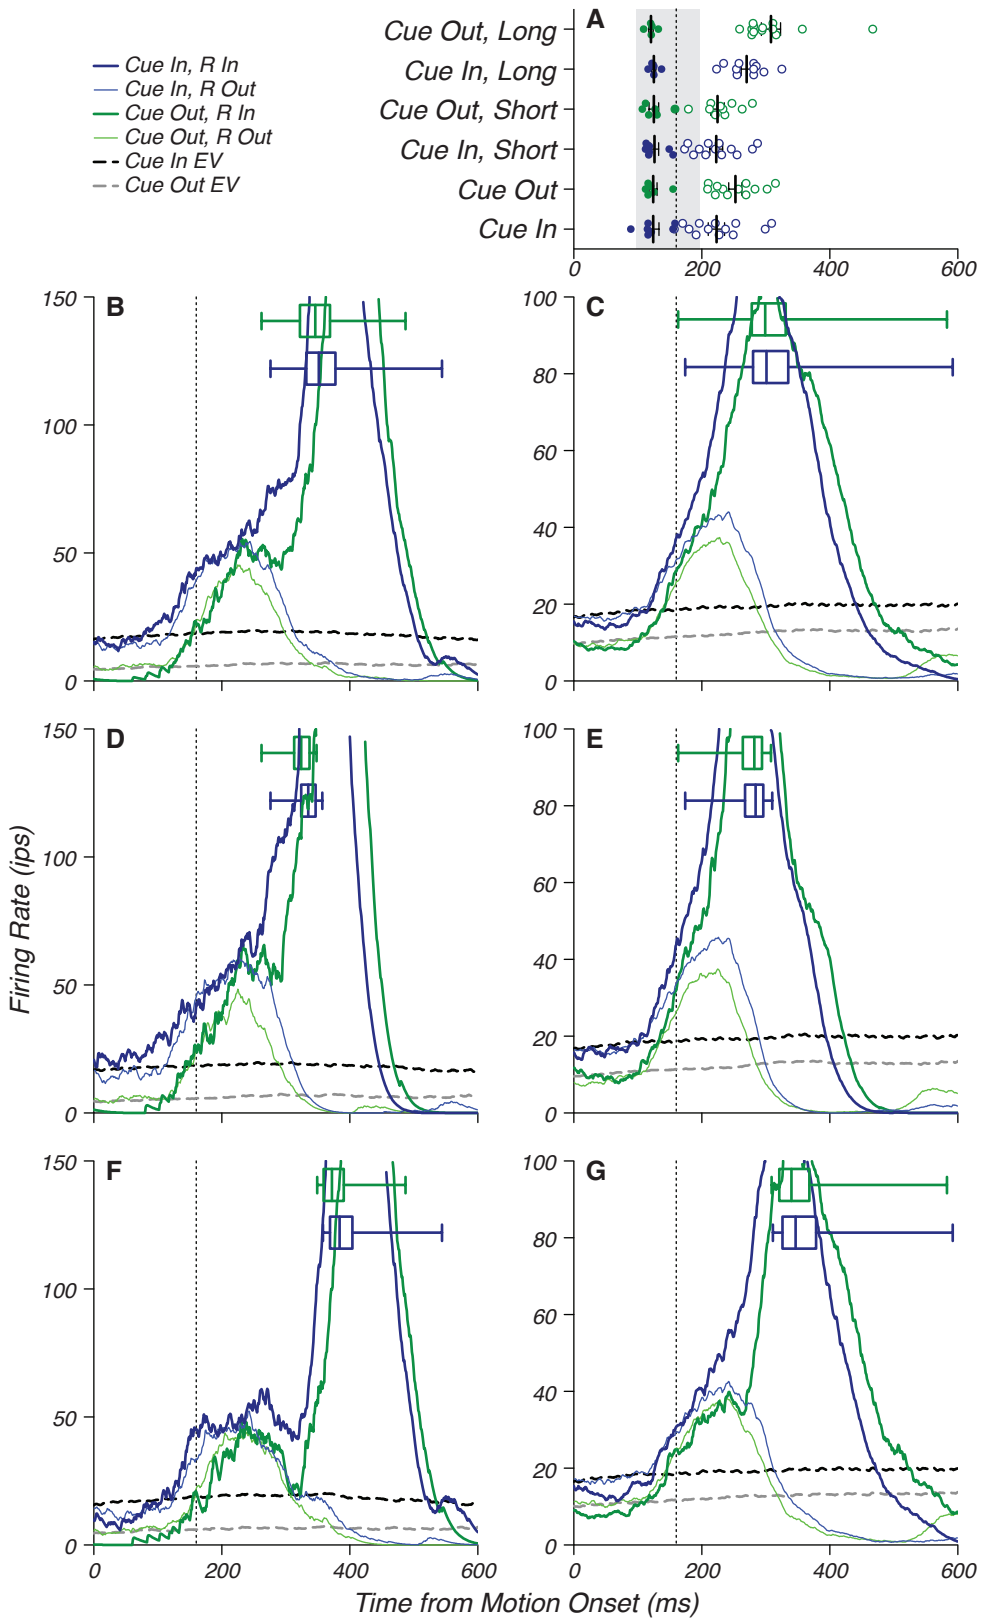
Cue In, R In
Cue In, R Out
Cue Out, R In
Cue Out, R Out
Cue In EV
Cue Out EV
A
Cue Out, Long
Cue In, Long
Cue Out, Short
Cue In, Short
Cue Out
Cue In
0
200
400
600
B
C
D
E
F
G
150
100
50
0
80
60
40
20
Firing Rate (ips)
Time from Motion Onset (ms)

Figure 6.12 *(following page)*: buildup cells have a response evoked by the motion pulse onset which precedes the response to the saccade onset. A) Response onset times for all buildup cells sorted based on the position of the cue and whether the response latency was shorter or longer than the median. Individual symbols correspond to individual neurons. Solid symbols are onset times for a direction-insensitive motion response and open symbols are the onset times for a saccade selective response. Black lines and error bars are means and standard error. Gray bar represents interval over which mean response is measured, 100-200 ms after onset. B) Example of response evoked by motion pulse (f150807). Black dashed line represents the expected firing rate of the neuron if the cue is in the response field and had a motion pulse not occurred. Gray dashed line represents the expected firing rate of the neuron if the cue is out of the response field and had a motion pulse not occurred. Whisker plots represent the saccade time with respect to the motion pulse onset. Vertical dashed line marks the appearance of the choice dots. C) Average response evoked by motion pulse. D-E) Example (f150807) and average response evoked by motion pulse for short-latency trials. F-G) Example (f150807) and average response evoked by motion pulse for short latency trials.

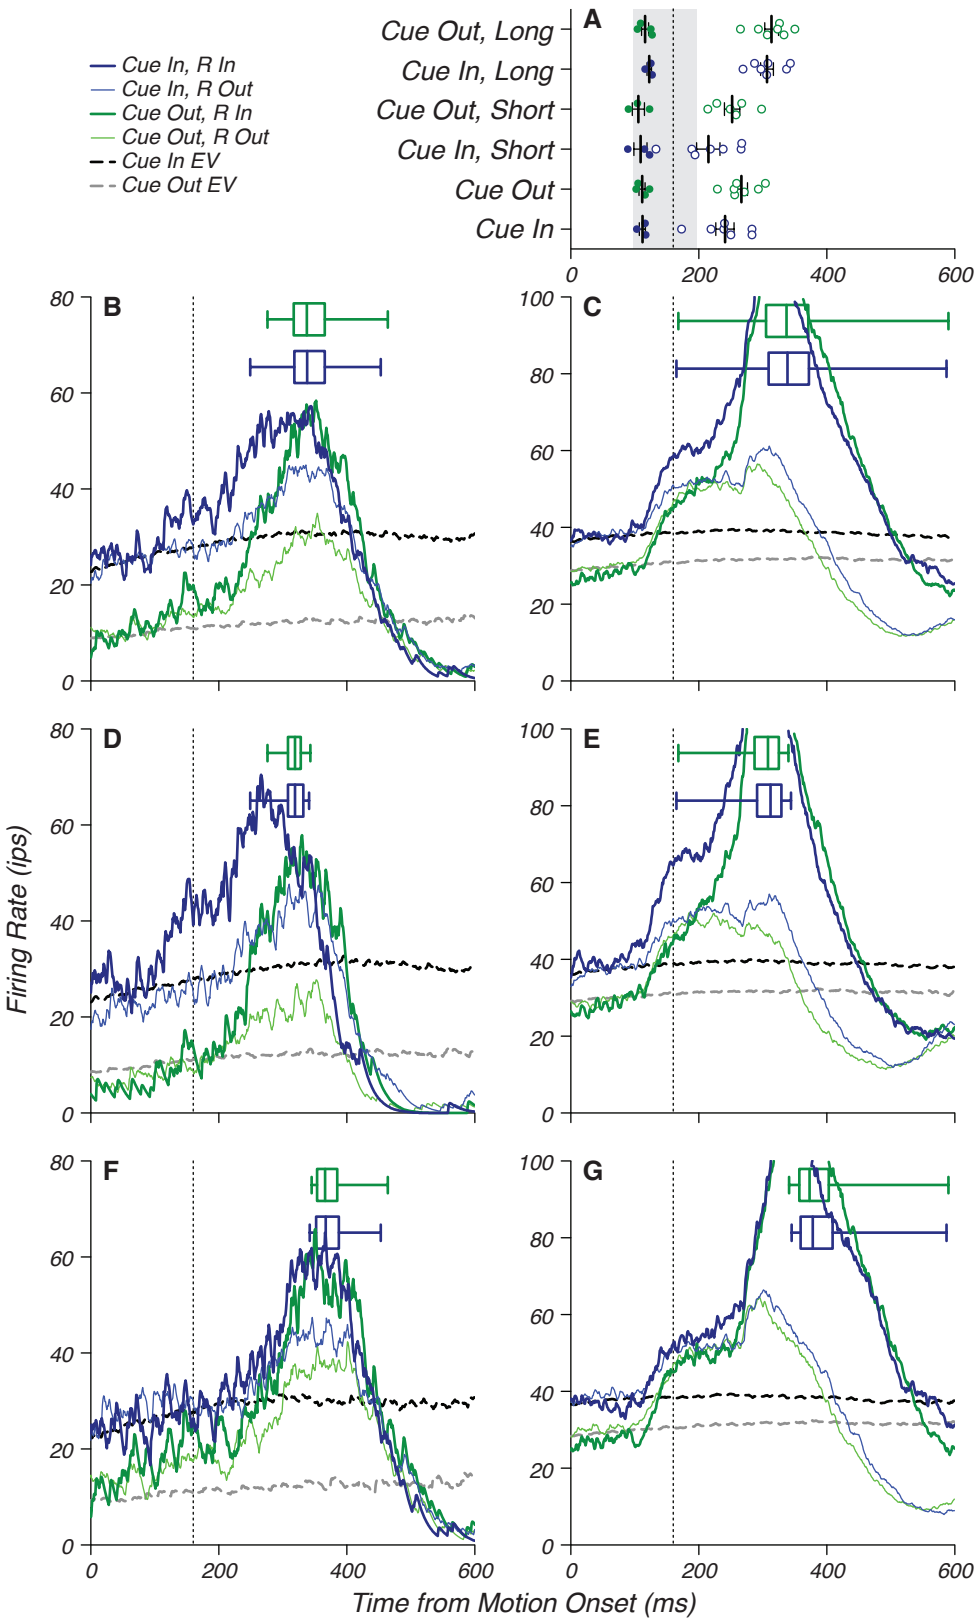
Cue In, R In
Cue In, R Out
Cue Out, R In
Cue Out, R Out
Cue In EV
Cue Out EV
A
Cue Out, Long
Cue In, Long
Cue Out, Short
Cue In, Short
Cue Out
Cue In
0
200
400
600
B
C
D
E
F
G
Firing Rate (ips)
Time from Motion Onset (ms)

Figure 6.13 *(following page)*: Summary of responses to cue onset. A-B) Peak mean firing rate for cue out of response field versus peak mean firing rate for cue in response field for VTM and VSBP cells. Black dots are peak average response for individual cells, and the white symbol is the population average peak response. Peak is the maximum average time-varying firing rate in an interval 50-150 ms after onset as shown in Figure 6.1. Error bars are standard error of the mean. C-D) Change in firing rate based on cue condition for VTM and VSBP cells. Responses are normalized by peak average mean firing rate across cue conditions. E-F) Change in firing rate based on which signal drove choice when the cue was is in the response field for VTM and VSBP cells. Responses are normalized by mean response when cue was in the response field. G-H) Change in firing rate based on signal driving choice when cue is out of the response field for VTM and VSBP cells. Responses are normalized by mean response when cue is out of the response field.

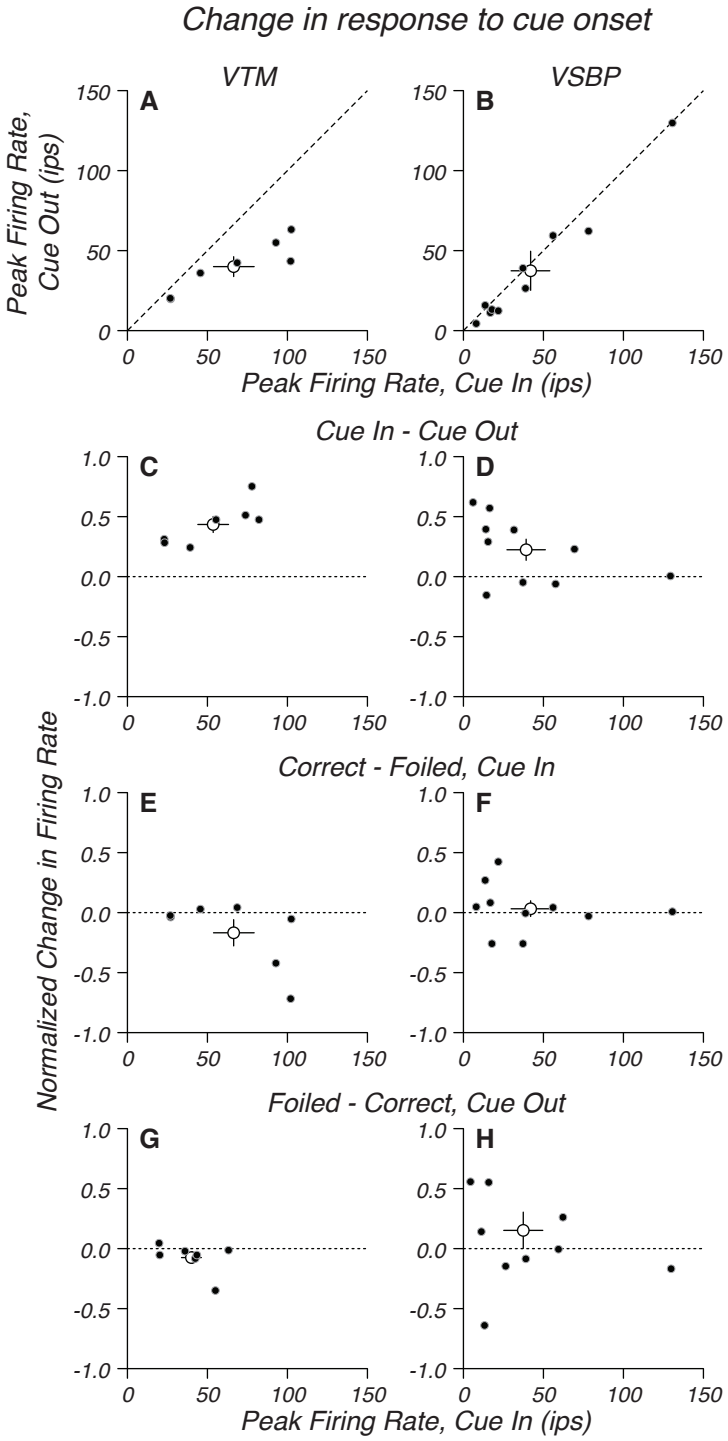
Change in response to cue onset
VTM
VSBP
A
B
Peak Firing Rate, Cue Out (ips)
Peak Firing Rate, Cue In (ips)
Cue In - Cue Out
C
D
Correct - Foiled, Cue In
E
F
Normalized Change in Firing Rate
Foiled - Correct, Cue Out
G
H
Peak Firing Rate, Cue In (ips)

Figure 6.14 *(following page)*: Summary of responses to patch onset. A-C) Peak mean firing rate for cue out of response field versus peak mean firing rate for cue in response field for VTM, VSBP, and VSB cells. Black dots are peak average response for individual cells and white symbol is the population average peak response. Peak is the maximum average time-varying firing rate in an interval 50-150 ms after patch onset as shown in Figure 6.2. Error bars are standard error of the mean. D-F) Change in firing rate based on cue condition for VTM, VSBP, and VSB cells. Change is normalized by peak mean firing rate across cue conditions. G-I) Change in firing rate based on signal driving choice when the cue is in the response field. Change is normalized by peak mean response when cue is in the response field. J-L) Change in firing rate based on signal driving choice when cue is out of the response field. Change is normalized by peak mean response when cue is out of the response field.

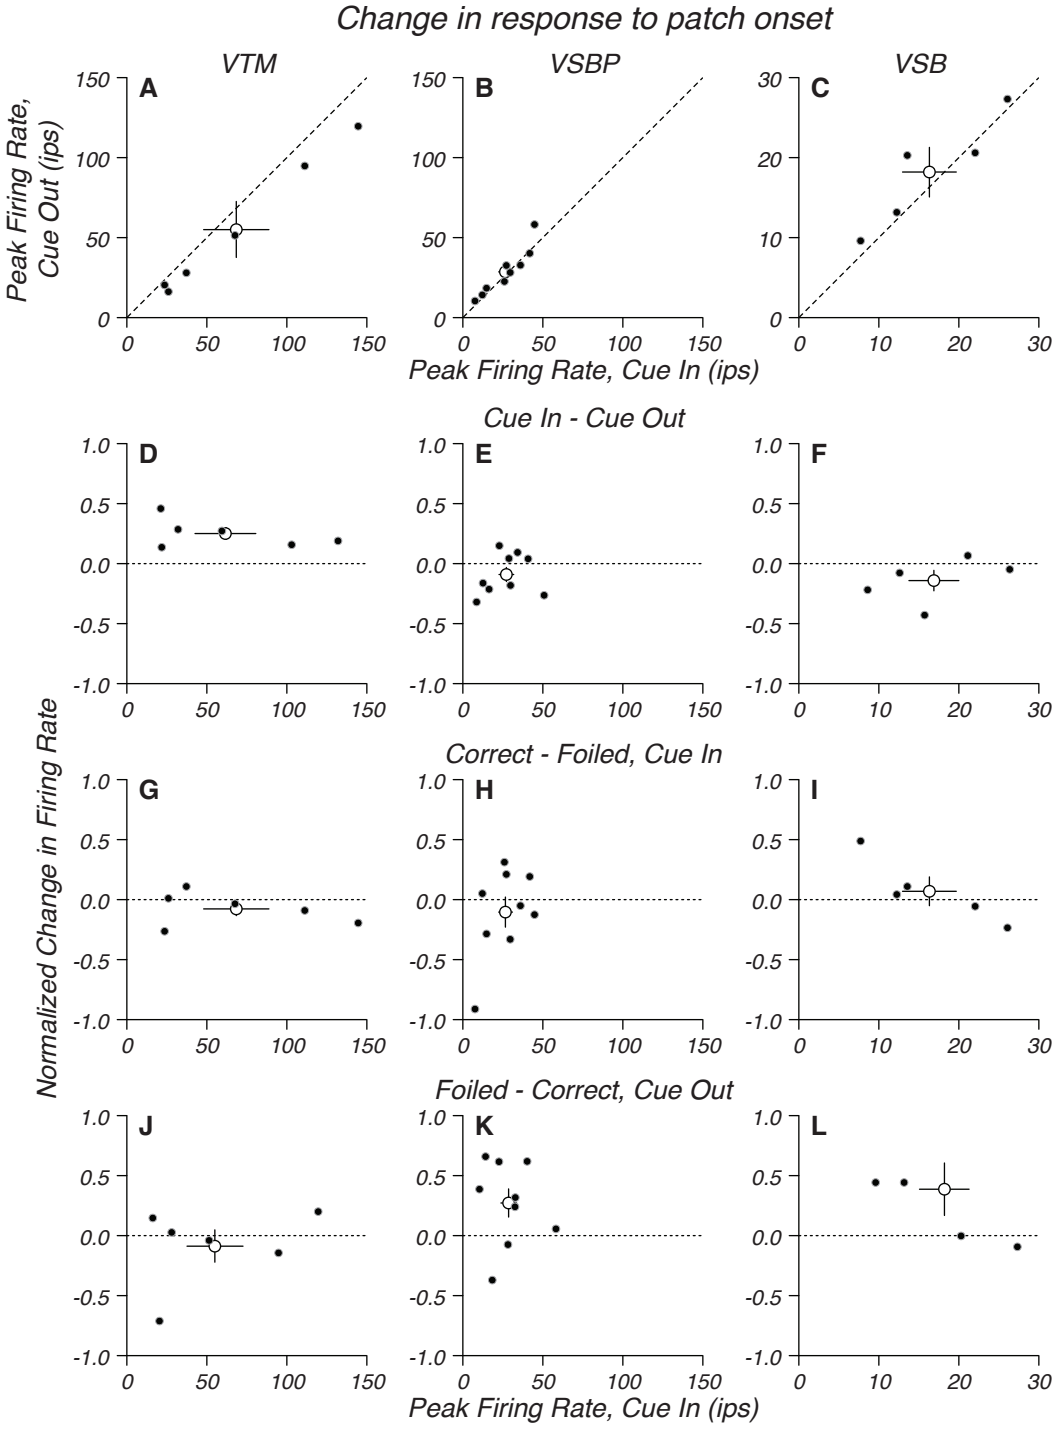
Change in response to patch onset
VTM
VSBP
VSB
A
B
C
Peak Firing Rate, Cue Out (ips)
Peak Firing Rate, Cue In (ips)
Cue In - Cue Out
D
E
F
Correct - Foiled, Cue In
G
H
I
Foiled - Correct, Cue Out
J
K
L
Normalized Change in Firing Rate
Peak Firing Rate, Cue In (ips)

Figure 6.15 *(following page)*: Summary of activity following cue offset, part 1. A-C) Mean firing rate for cue out of response field versus mean firing rate for cue in response field for VTM, VSBP, and VSB cells. Black dots are mean activity for individual cells and white symbol is population average. Error bars are standard error of the mean. D-F) Change in firing rate based on cue condition for VTM, VSBP, and VSB cells. Change is normalized by mean firing rate in interval represented in Figure 6.3. G-I) Change in firing rate based on signal driving choice when the cue is in the response field. Change is normalized by mean response when cue is in the response field. J-L) Change in firing rate based on signal driving choice when cue is out of the response field. Change is normalized by mean response when cue is out of the response field.

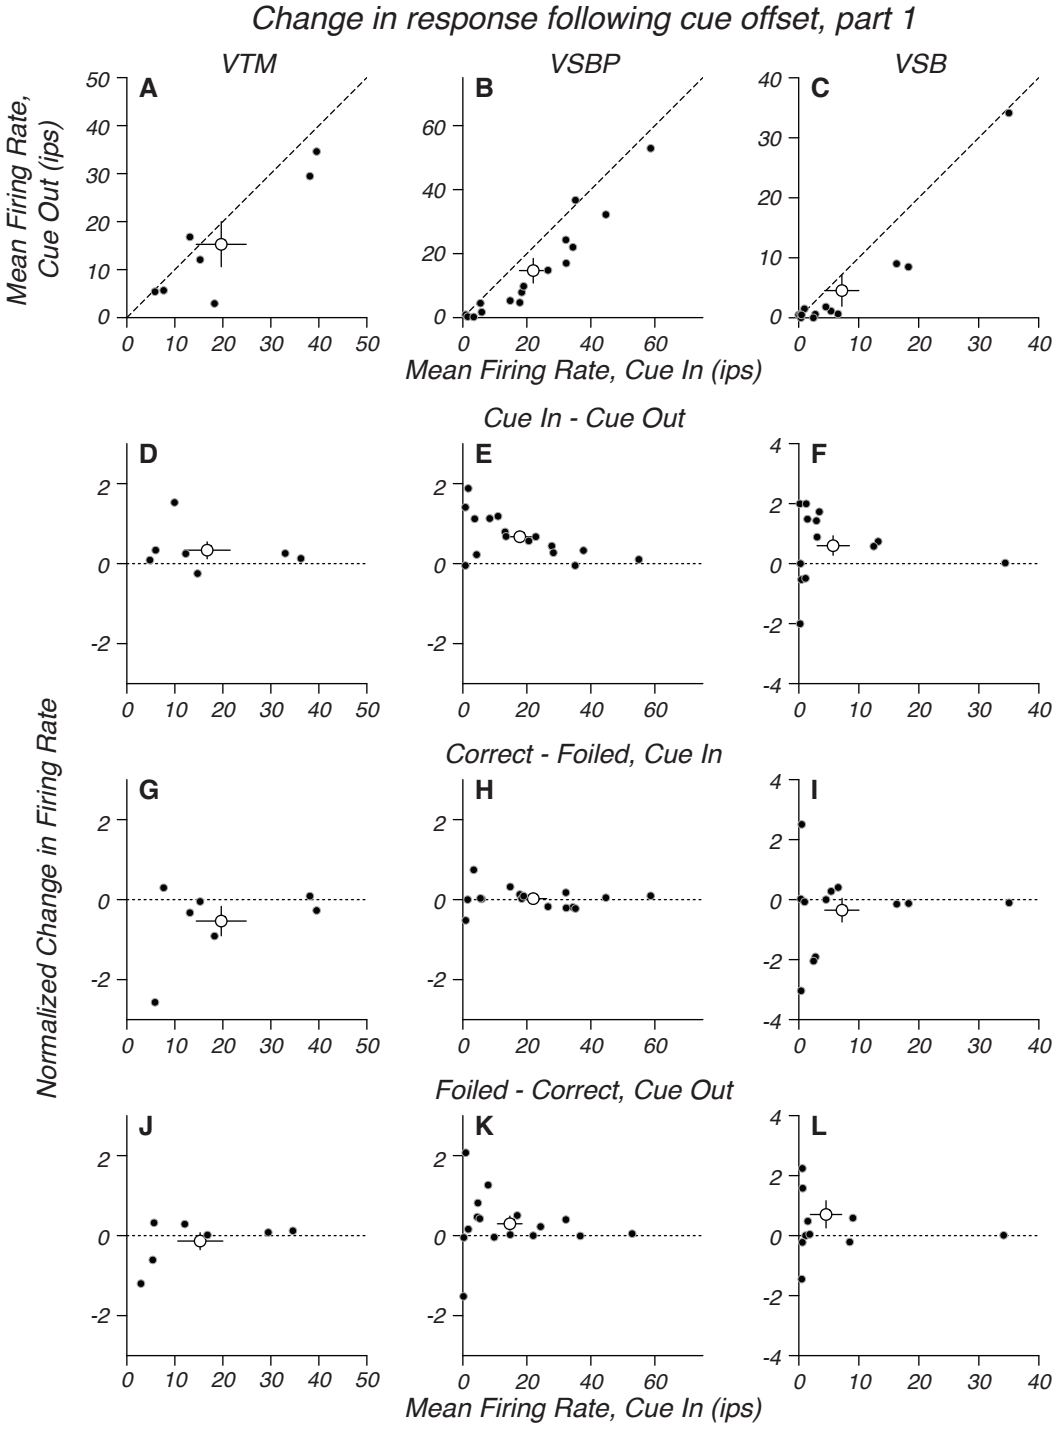
Change in response following cue offset, part 1
VTM
VSBP
VSB
A
B
C
Mean Firing Rate, Cue Out (ips)
Mean Firing Rate, Cue In (ips)
Cue In - Cue Out
D
E
F
Correct - Foiled, Cue In
G
H
I
Foiled - Correct, Cue Out
J
K
L
Normalized Change in Firing Rate
Mean Firing Rate, Cue In (ips)

Figure 6.16 *(following page)*: Summary of activity following cue offset, part 2. A-B) Mean firing rate for cue out of response field versus mean firing rate for cue in response field for BU and MF cells. Black dots are mean activity for individual cells and white symbol is population average. Error bars are standard error of the mean. C-D) Change in firing rate based on cue condition for BU and MF cells. Change is normalized by mean firing rate in interval represented in Figure 6.3. E-F) Change in firing rate based on signal driving choice when the cue is in the response field. Change is normalized by mean response when cue is in the response field. G-H) Change in firing rate based on signal driving choice when cue is out of the response field. Change is normalized by mean response when cue is out of the response field.

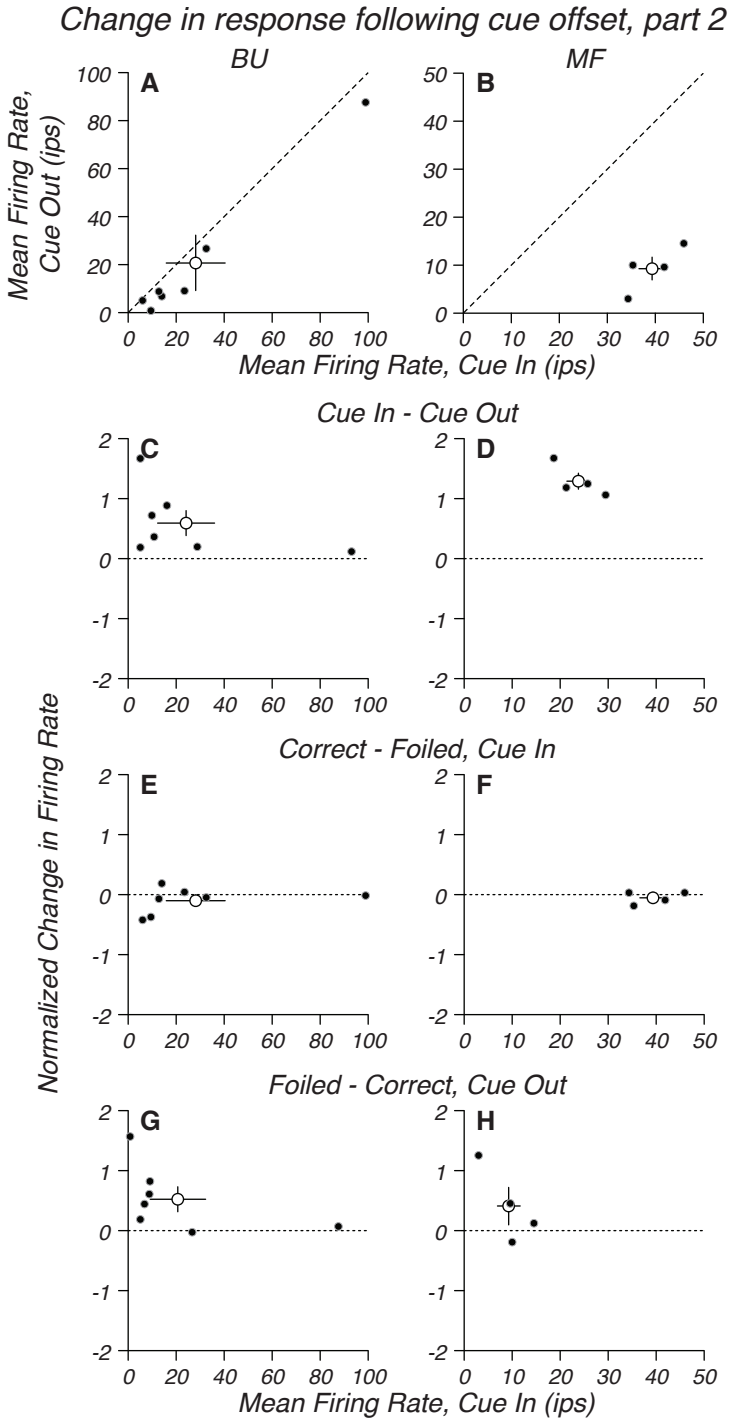
Change in response following cue offset, part 2
BU
MF
A
B
Mean Firing Rate, Cue Out (ips)
Mean Firing Rate, Cue In (ips)
Cue In - Cue Out
C
D
Correct - Foiled, Cue In
E
F
Normalized Change in Firing Rate
Foiled - Correct, Cue Out
G
H
Mean Firing Rate, Cue In (ips)

Figure 6.17 *(following page)*: Summary of activity prior to motion onset, part 1. A-C) Mean firing rate for cue out of response field versus mean firing rate for cue in response field for VTM, VSBP, and VSB cells. Black dots are mean activity for individual cells and white symbol is population average. Error bars are standard error of the mean. D-F) Change in firing rate based on cue condition for VTM, VSBP, and VSB cells. Change is normalized by mean firing rate in interval represented in Figure 6.5. G-I) Change in firing rate based on signal driving choice when the cue is in the response field. Change is normalized by mean response when cue is in the response field. J-L) Change in firing rate based on signal driving choice when cue is out of the response field. Change is normalized by mean response when cue is out of the response field.

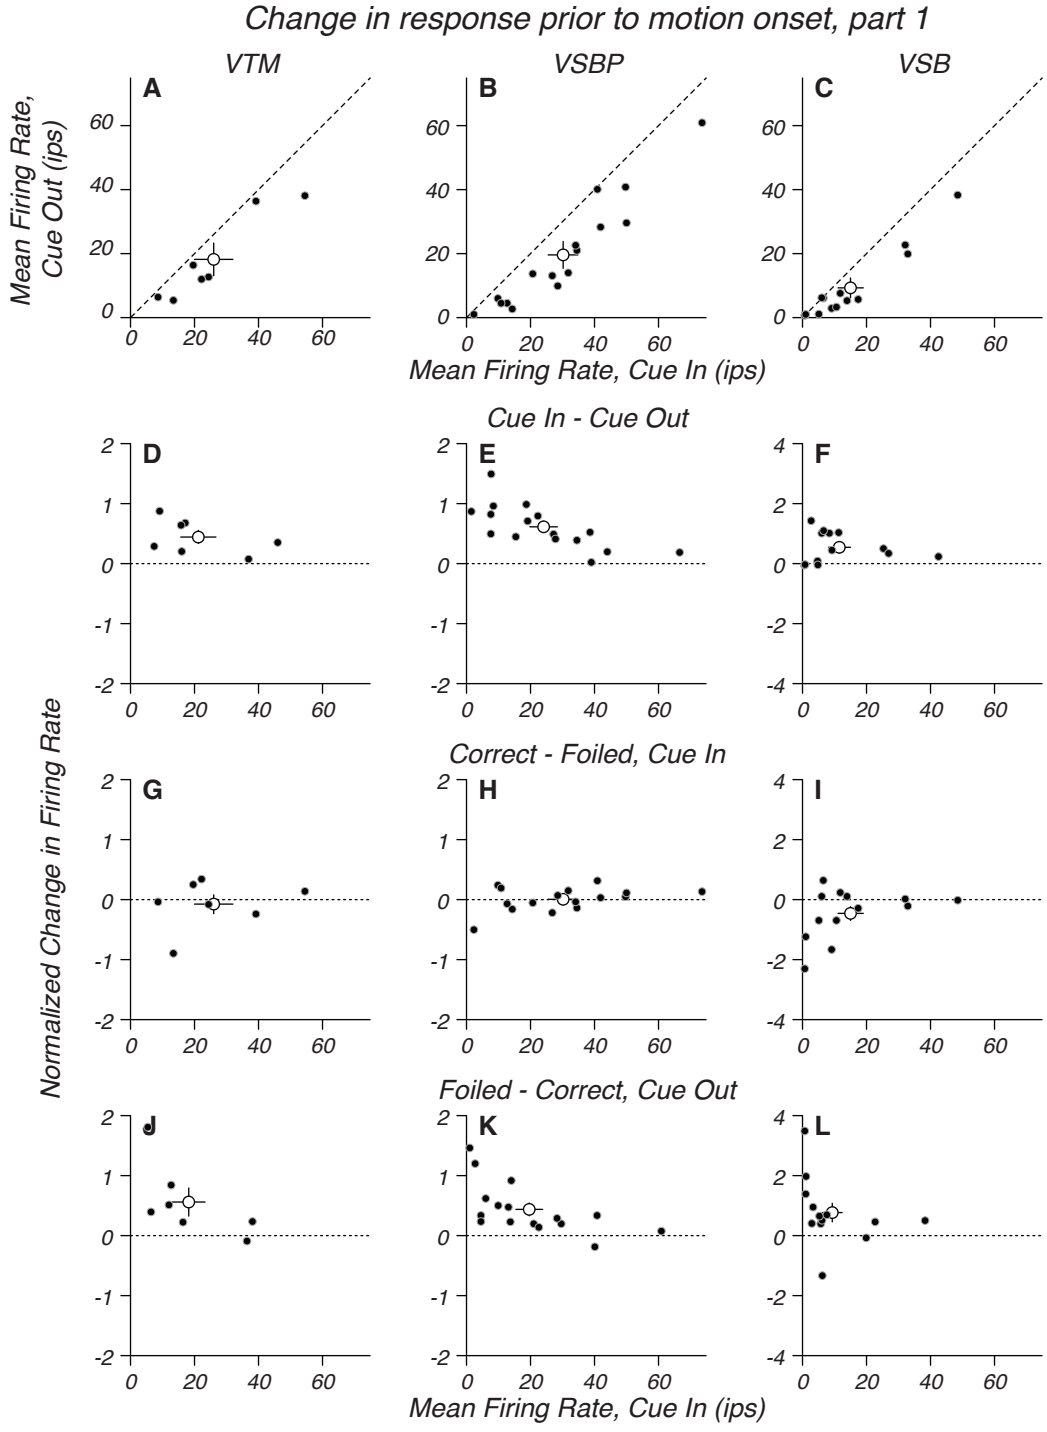
Change in response prior to motion onset, part 1
VTM
VSBP
VSB
A
B
C
Mean Firing Rate, Cue Out (ips)
Mean Firing Rate, Cue In (ips)
Cue In - Cue Out
D
E
F
Correct - Foiled, Cue In
G
H
I
Foiled - Correct, Cue Out
J
K
L
Normalized Change in Firing Rate
Mean Firing Rate, Cue In (ips)

Figure 6.18 *(following page)*: Summary of activity prior to motion onset, part 2. A-B) Mean firing rate for cue out of response field versus mean firing rate for cue in response field for BU and MF cells. Black dots are mean activity for individual cells and white symbol is population average. Error bars are standard error of the mean. C-D) Change in firing rate based on cue condition for BU and MF cells. Change is normalized by mean firing rate in interval represented in Figure 6.5. E-F) Change in firing rate based on signal driving choice when the cue is in the response field. Change is normalized by mean response when cue is in the response field. G-H) Change in firing rate based on signal driving choice when cue is out of the response field. Change is normalized by mean response when cue is out of the response field.

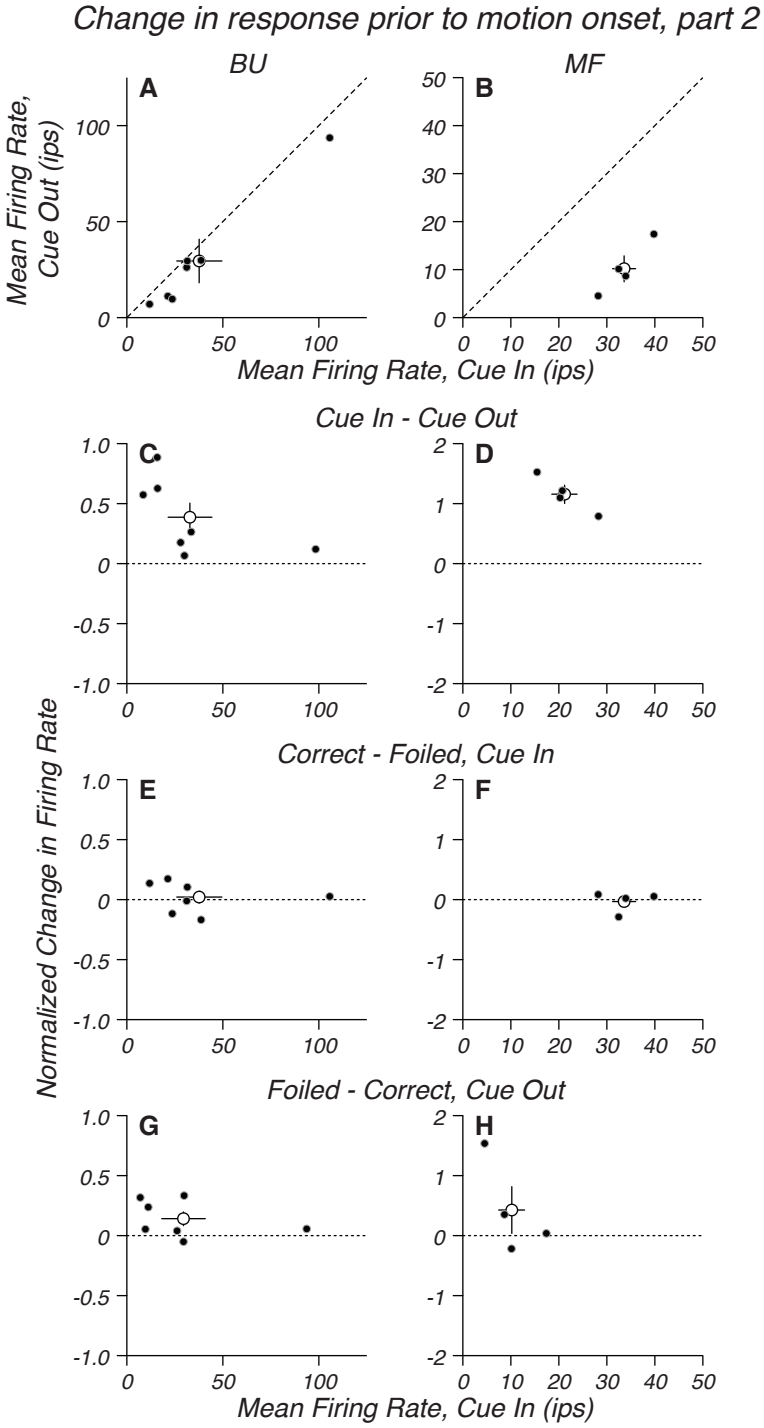
Change in response prior to motion onset, part 2
BU
MF
A
B
Mean Firing Rate, Cue Out (ips)
Mean Firing Rate, Cue In (ips)
Cue In - Cue Out
C
D
Correct - Foiled, Cue In
E
F
Normalized Change in Firing Rate
Foiled - Correct, Cue Out
G
H
Mean Firing Rate, Cue In (ips)

Figure 6.19 *(following page)*: Summary of response to motion pulse, part 1. A-C) Mean firing rate for cue out of response field versus mean firing rate for cue in response field for VTM, VSBP, and VSB cells. Black dots are mean activity for individual cells and white symbol is population average. Error bars are standard error of the mean. D-F) Change in firing rate based on cue condition for VTM, VSBP, and VSB cells. Change is normalized by mean firing rate in interval represented in Figure 6.5. G-I) Change in firing rate based on signal driving choice when the cue is in the response field. Change is normalized by mean response when cue is in the response field. J-L) Change in firing rate based on signal driving choice when cue is out of the response field. Change is normalized by mean response when cue is out of the response field.

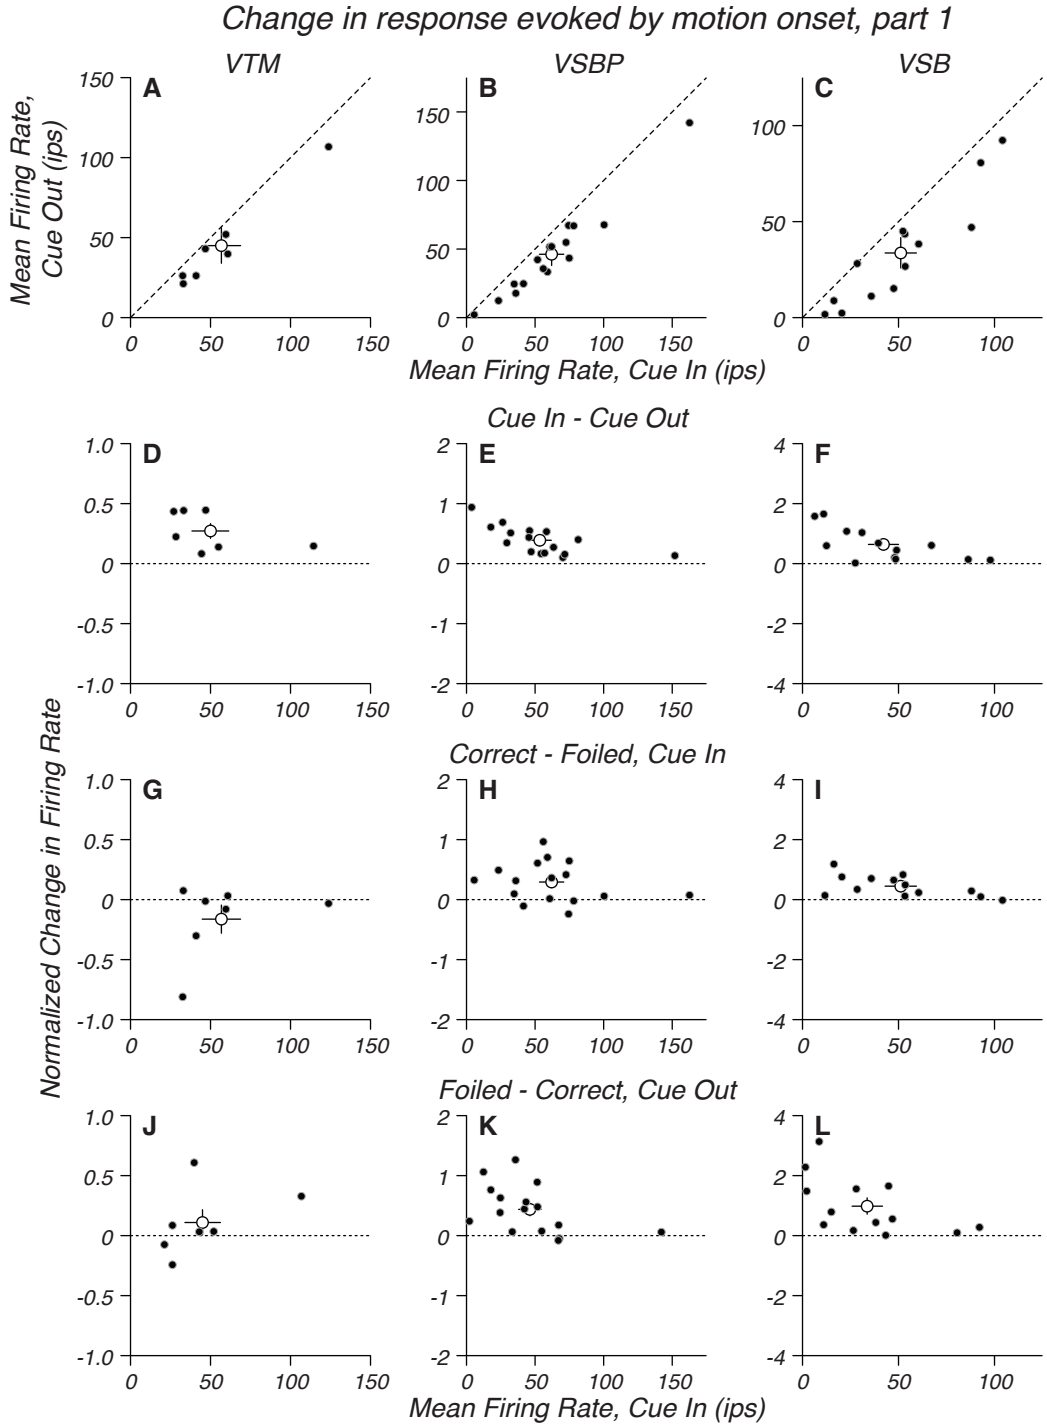
Change in response evoked by motion onset, part 1
VTM
VSBP
VSB
A
B
C
Mean Firing Rate, Cue Out (ips)
Mean Firing Rate, Cue In (ips)
Cue In - Cue Out
D
E
F
Correct - Foiled, Cue In
G
H
I
Foiled - Correct, Cue Out
J
K
L
Normalized Change in Firing Rate
Mean Firing Rate, Cue In (ips)

Figure 6.20 *(following page)*: Summary of response to motion pulse, part 2. A-B) Mean firing rate for cue out of response field versus mean firing rate for cue in response field for BU and MF cells. Black dots are individual cells and white symbol is average response. Error bars are standard error of the mean. C-D) Change in firing rate based on cue condition for BU and MF cells. Change is normalized by mean firing rate in interval represented in Figure 6.5. E-F) Change in firing rate based on signal driving choice when the cue is in the response field. Change is normalized by mean response when cue is in the response field. G-H) Change in firing rate based on signal driving choice when cue is out of the response field. Change is normalized by mean response when cue is out of the response field.

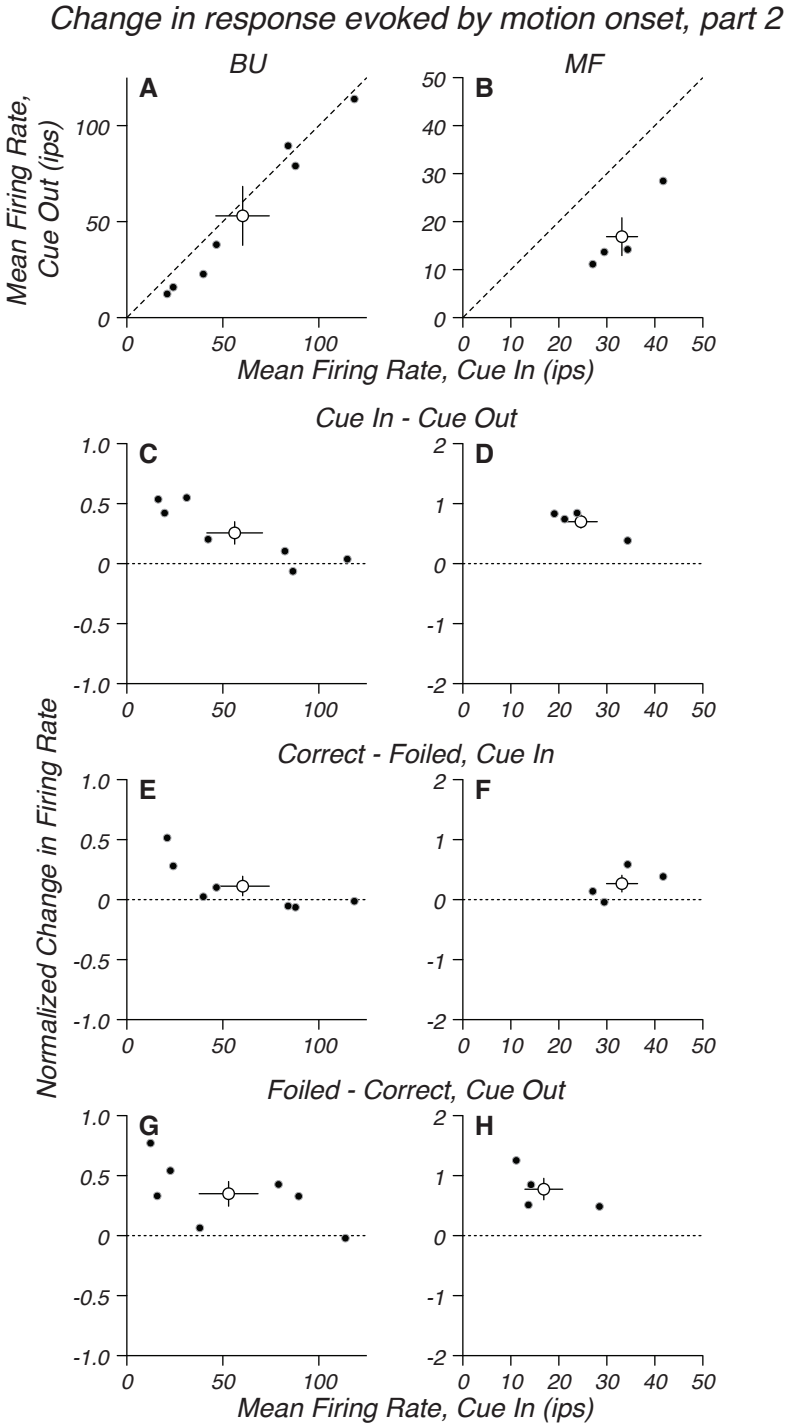
Change in response evoked by motion onset, part 2
BU
MF
A
B
Mean Firing Rate, Cue Out (ips)
Mean Firing Rate, Cue In (ips)
Cue In - Cue Out
C
D
Correct - Foiled, Cue In
E
F
Normalized Change in Firing Rate
Foiled - Correct, Cue Out
G
H
Mean Firing Rate, Cue In (ips)

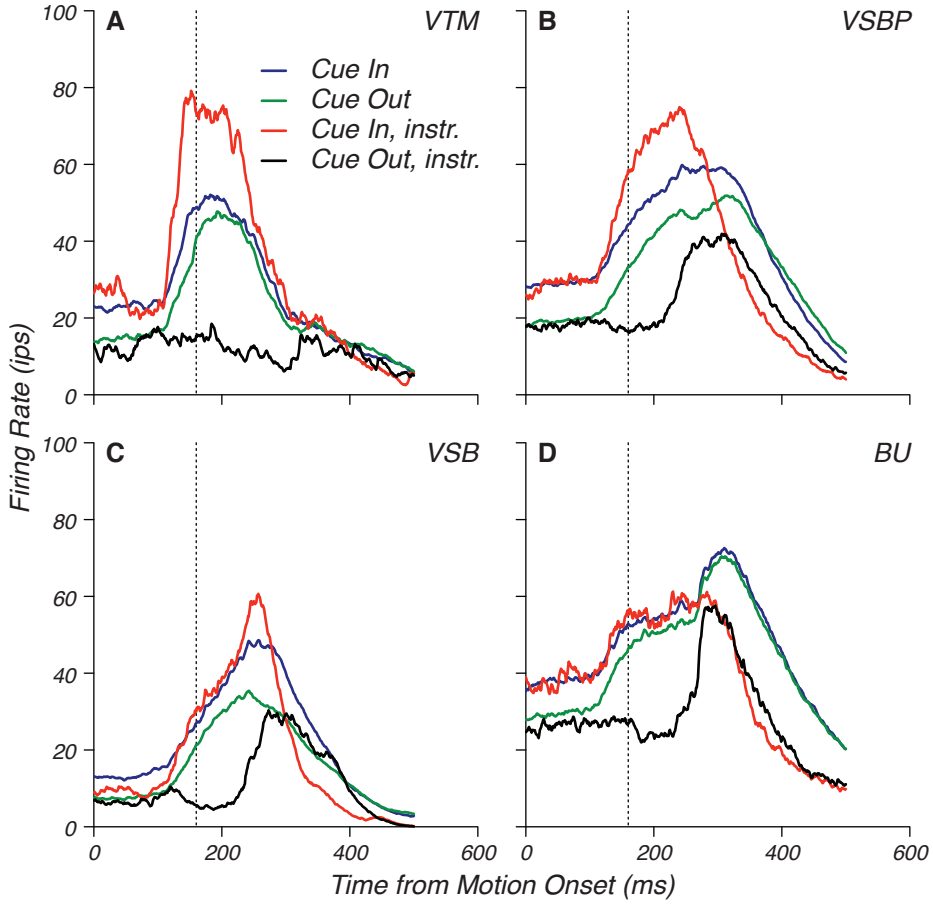

Figure 6.21: Average response evoked by motion pulse in instructor trials. Blue and green lines represent population average response evoked by motion pulses when the cue was in or out of the response field and the strength of the pulses was equal in both locations. In instructor trials, the motion pulse in the cued location had a coherence of 0.5 and the foil signal remained incoherent. The red line represents response to the instructor pulse when the cue was in the response field, and the black line represents the response to incoherent motion when the cue was out of the response field. Average responses are shown for VTM cells (A), VSBP cells (B), VSB cells (C), and buildup cells (D).

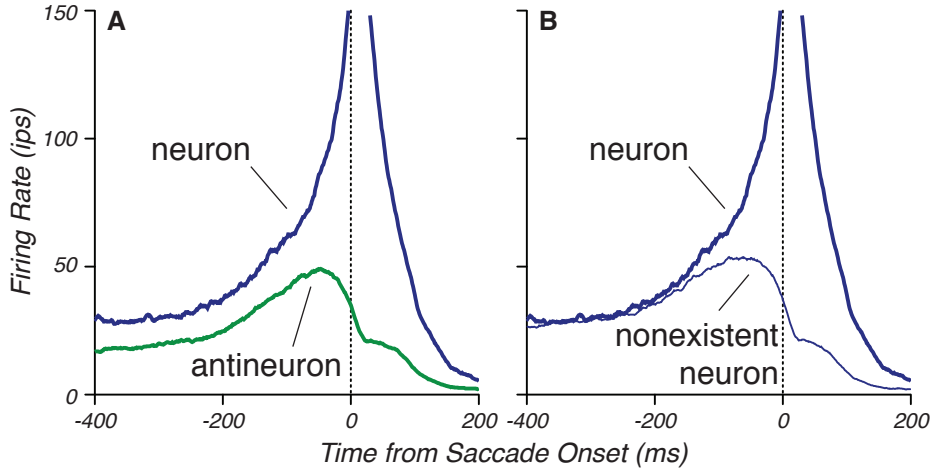

Figure 6.22: The neuron/antineuron pair (A) versus the neuron/nonexistent neuron pair (B).

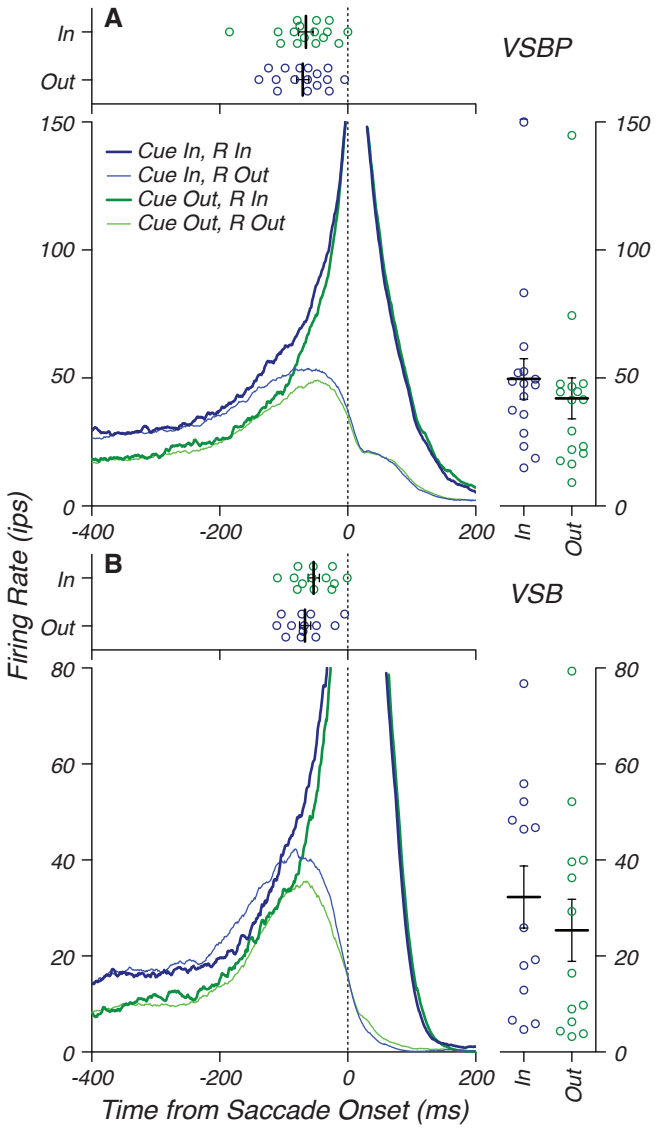

Figure 6.23: Impact of cue on saccade discrimination time and firing rate at discrimination time for VSBP (A) and VSB (B) cells. Population average of response aligned to saccade onset. Thick blue and green lines represent firing rate when saccades were made into the response field and the cue was either in or out of the response field. Thin blue and green lines represent firing rate when saccades were made out of the response field and the cue was either in or out of the response field. Horizontal scatter plots represent discrimination time for saccades for each cell, defined as in Methods. Black lines and error bars are population means and standard error. Vertical scatter plots represent average firing rate at saccade discrimination time for each cell. Black lines and error bars are population means and standard error.

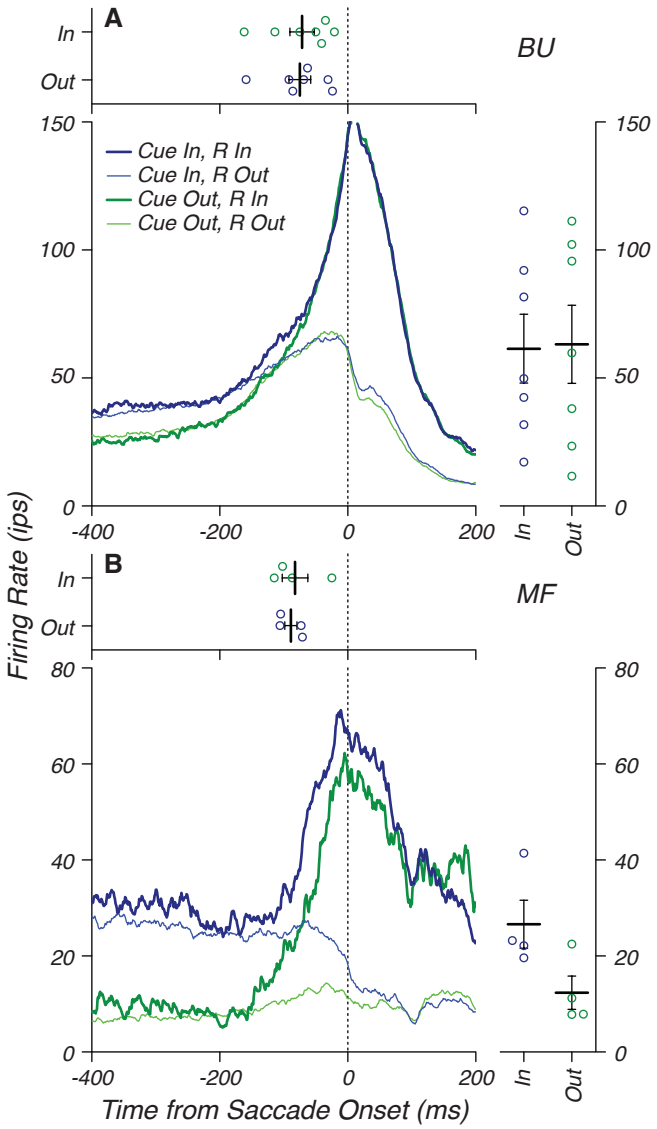

Figure 6.24: Impact of cue on saccade discrimination time and firing rate at discrimination time for BU (A) and MF (B) cells. Population average of response aligned to saccade onset. Thick blue and green lines represent firing rate when saccades were made into the response field and the cue was either in or out of the response field. Thin blue and green lines represent firing rate when saccades were made out of the response field and the cue was either in or out of the response field. Horizontal scatter plots represent discrimination time for saccades for each cell, defined as in Methods. Black lines and error bars are population means and standard error. Vertical scatter plots represent average firing rate at saccade discrimination time for each cell. Black lines and error bars are population means and standard error.

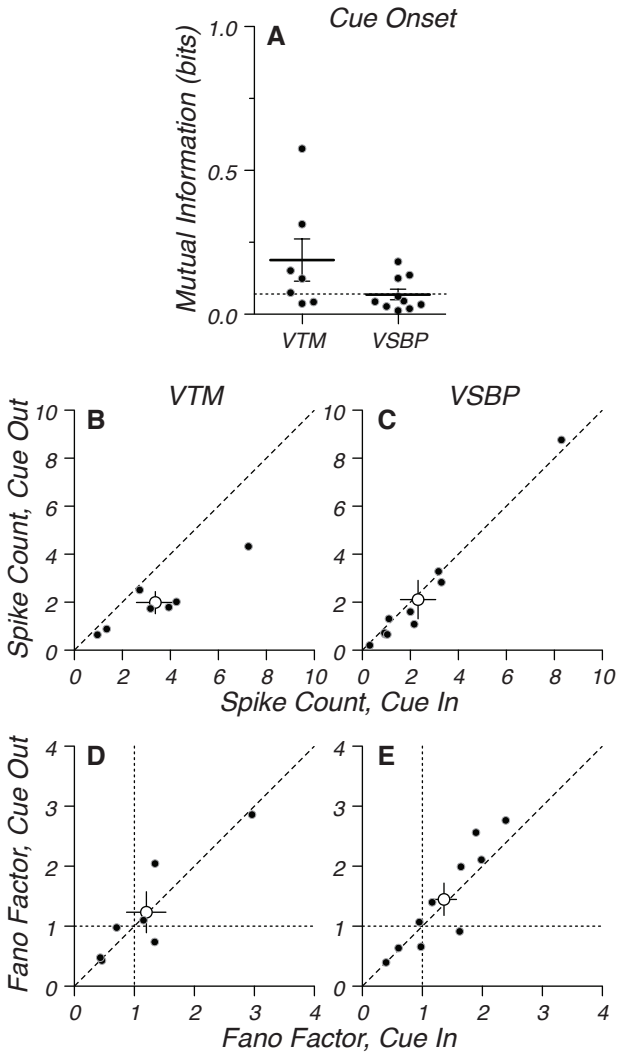

Figure 6.25: Encoding cue position in spike count in interval 50-150 ms after cue onset. A) Mutual information between spike count and cue position for VTM and VSBP cells. Individual symbols represent individual cells; bars and error bars are population mean and standard error. Horizontal dashed line is 95% upper confidence limit on chance information by permutation test. B-C) Spike count for trials in which the cue was out of the response field versus spike count when the cue was in the response field for VTM and VSBP cells. Open symbol is population average with standard error bars. D-E) Fano Factor when cue was out of response field versus Fano Factor when cue was in response field for VTM and VSBP cells.

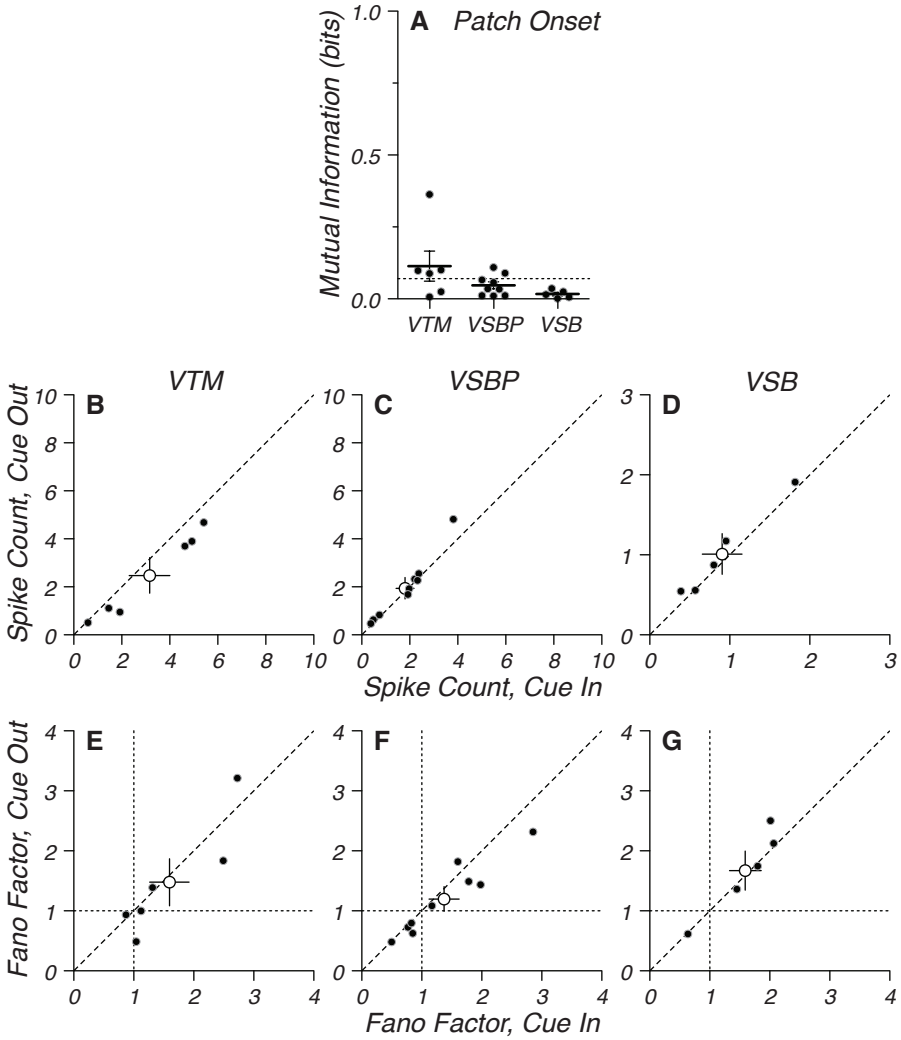

Figure 6.26: Encoding cue position in spike count in interval 50-150 ms after patch onset. A) Mutual information between spike count and cue position for VTM, VSBP, and VSB cells. Individual symbols represent individual cells; bars and error bars are population mean and standard error. Horizontal dashed line is 95% upper confidence limit on chance information by permutation test. B-D) Spike count for trials in which the cue was out of the response field versus spike count when the cue was in the response field for VTM, VSBP, and VSB cells. Open symbol is population average with standard error bars. D-E) Fano Factor when cue was out of response field versus Fano Factor when cue was in response field for VTM, VSBP, and VSB cells.

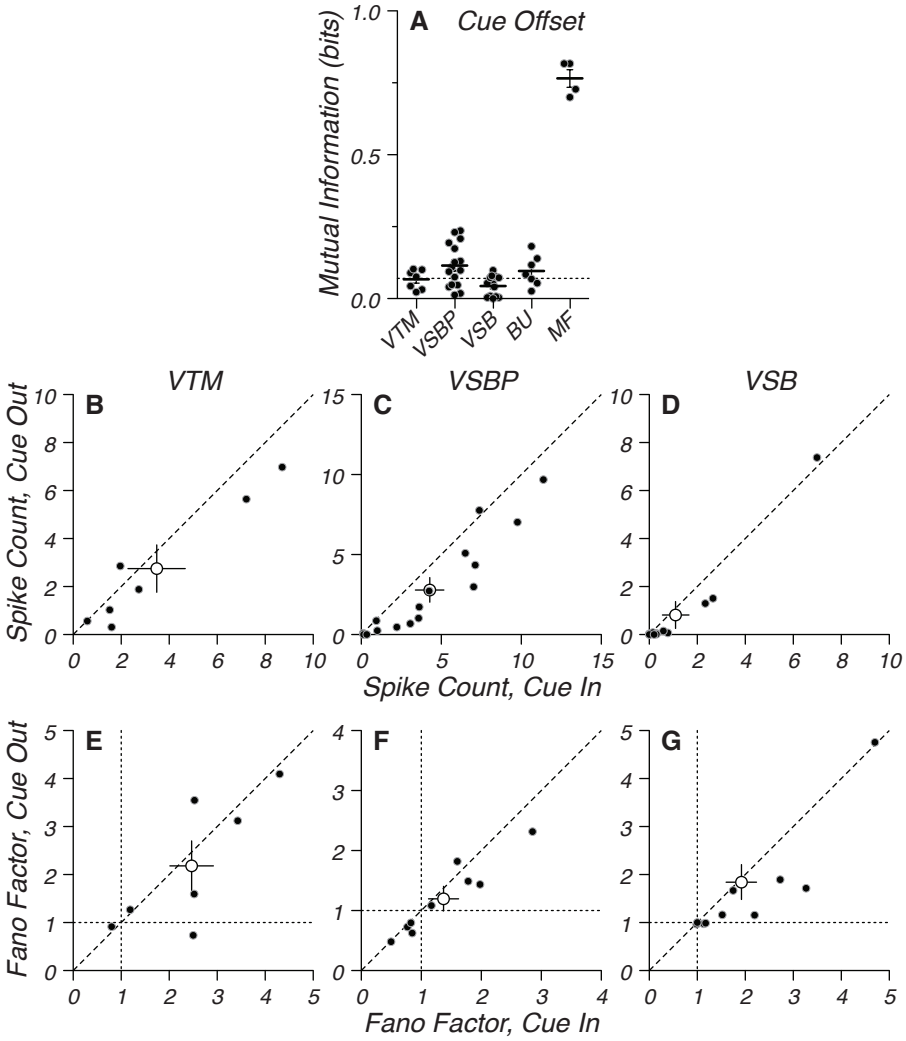

Figure 6.27: Encoding cue position in spike count in interval 250-500 ms after cue offset, part 1. A) Mutual information between spike count and cue position for VTM, VSBP, VSB, BU and MF cells. Individual symbols represent individual cells; bars and error bars are population mean and standard error. Horizontal dashed line is 95% upper confidence limit on chance information by permutation test. B-D) Spike count for trials in which the cue was out of the response field versus spike count when the cue was in the response field for VTM, VSBP, and VSB cells. Open symbol is population average with standard error bars. D-E) Fano Factor when cue was out of response field versus Fano Factor when cue was in response field for VTM, VSBP, and VSB cells.

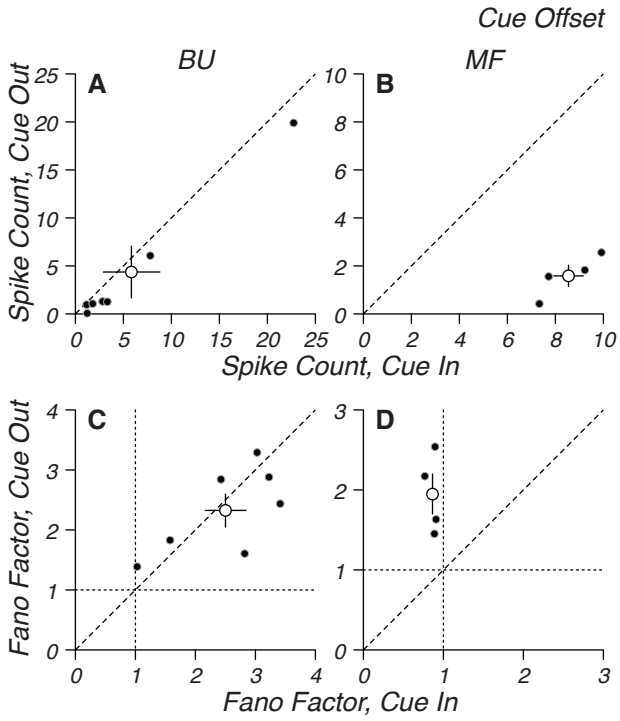

Figure 6.28: Encoding cue position in spike count in interval 250-500 ms after cue offset, part 2. A-B) Spike count for trials in which the cue was out of the response field versus spike count when the cue was in the response field for BU and MF cells. Open symbol is population average with standard error bars. C-D) Fano Factor when cue was out of response field versus Fano Factor when cue was in response field for BU and MF cells.

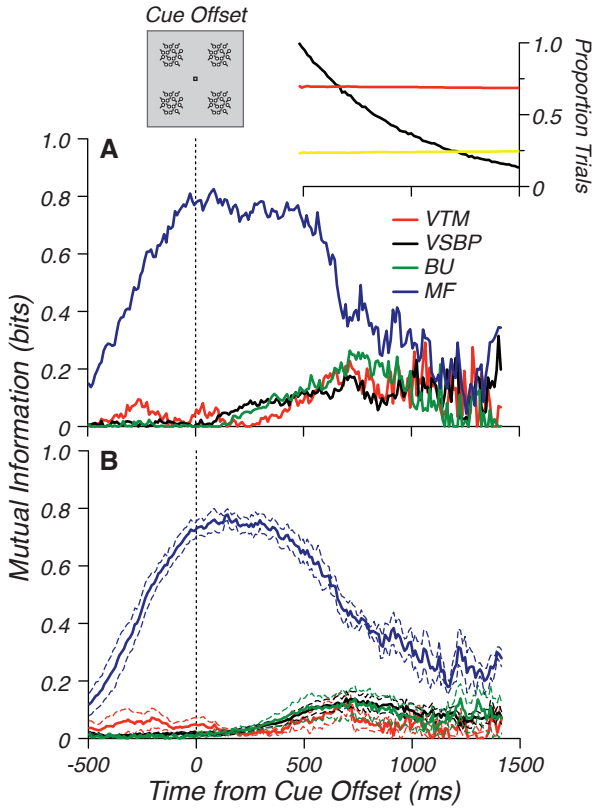

Figure 6.29: Time course of mutual information aligned on cue offset. A) Mutual information for example VTM (f171207), VSBP (m011008), BU (f150807), and MF (m301208) cells. Inset shows performance as a function of delay (proportion of choices with cued signal in red, with foil in yellow) and empirical distribution of delays (black line). B) Population average mutual information for cell groups as in (A).

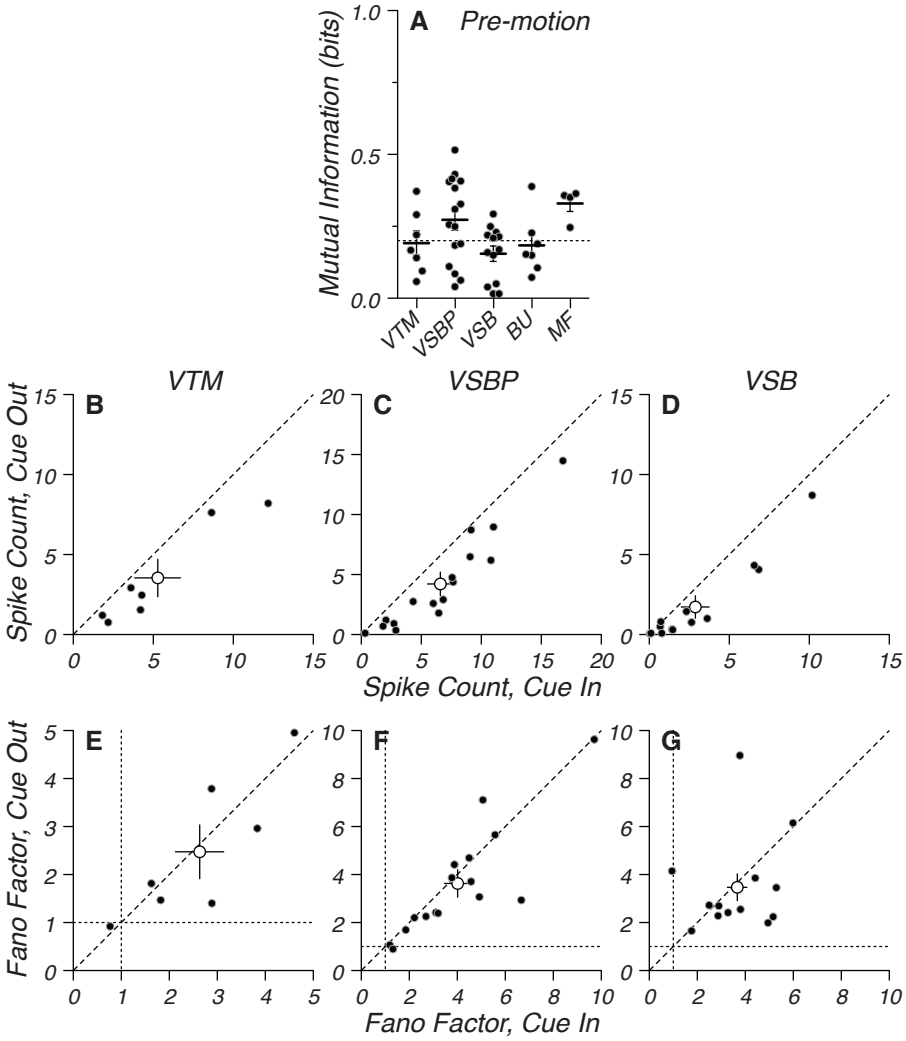

Figure 6.30: Encoding cue position in spike count in 250 ms interval preceding motion onset, part 1. A) Mutual information between spike count and cue position for VTM, VSBP, VSB, BU and MF cells. Individual symbols represent individual cells; bars and error bars are population mean and standard error. Horizontal dashed line is 95% upper confidence limit on chance information by permutation test. B-D) Spike count for trials in which the cue was out of the response field versus spike count when the cue was in the response field for VTM, VSBP, and VSB cells. Open symbol is population average with standard error bars. D-E) Fano Factor when cue was out of response field versus Fano Factor when cue was in response field for VTM, VSBP, and VSB cells.

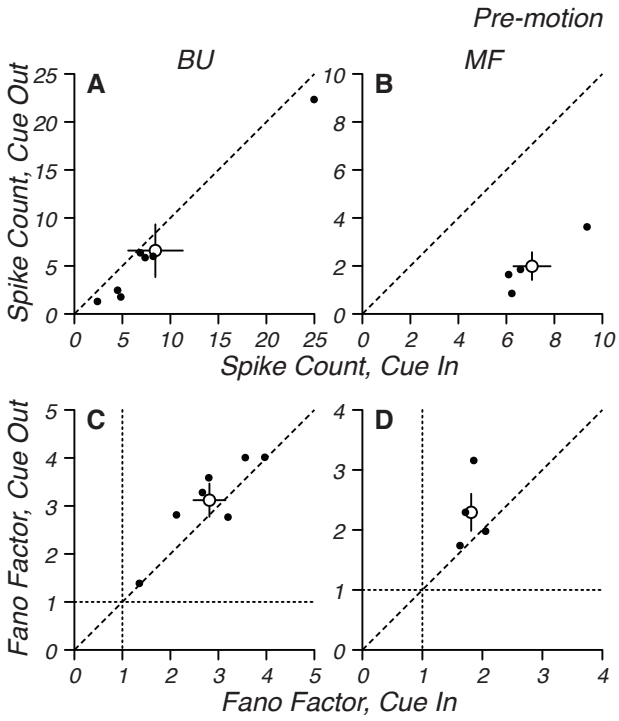

Figure 6.31: Encoding cue position in spike count in 250 ms interval preceding motion onset, part 2. A-B) Spike count for trials in which the cue was out of the response field versus spike count when the cue was in the response field for BU and MF cells. Open symbol is population average with standard error bars. C-D) Fano Factor when cue was out of response field versus Fano Factor when cue was in response field for BU and MF cells.

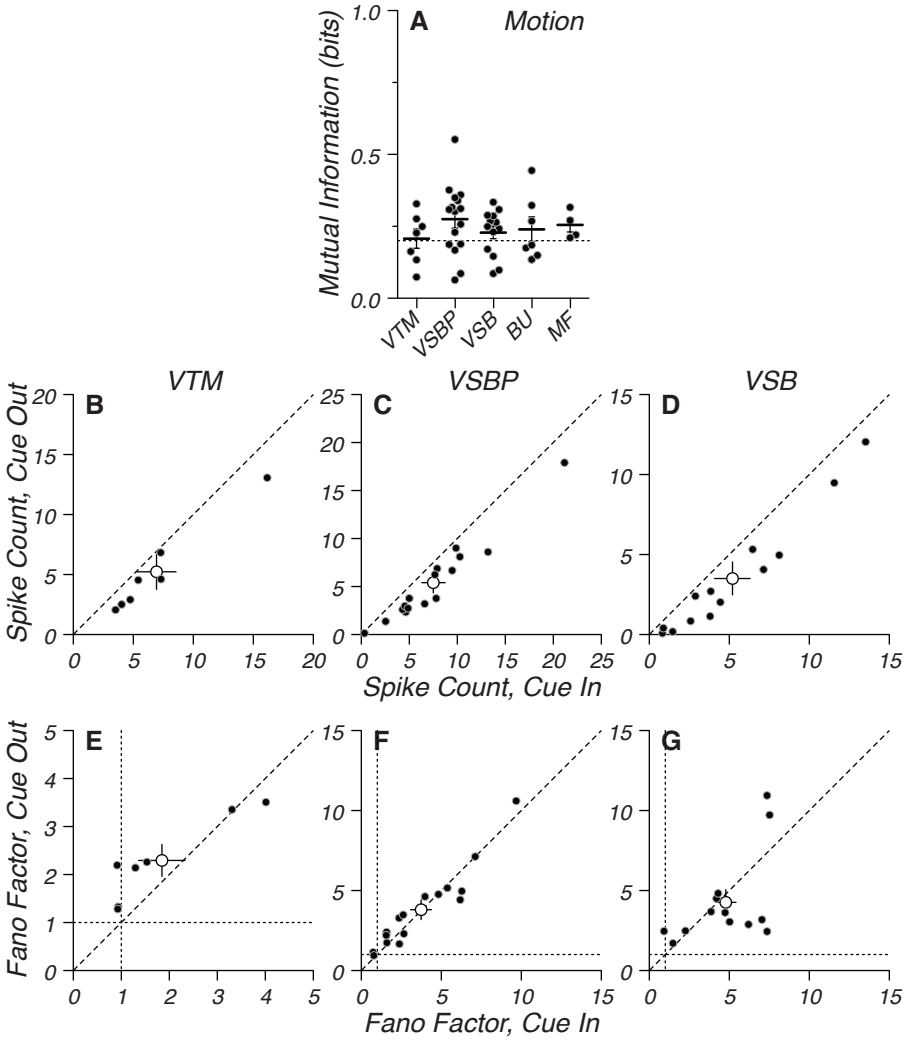

Figure 6.32: Encoding cue position in spike count in interval 100-200 ms following motion onset, part 1. A) Mutual information between spike count and cue position for VTM, VSBP, VSB, BU and MF cells. Individual symbols represent individual cells; bars and error bars are population mean and standard error. Horizontal dashed line is 95% upper confidence limit on chance information by permutation test. B-D) Spike count for trials in which the cue was out of the response field versus spike count when the cue was in the response field for VTM, VSBP, and VSB cells. Open symbol is population average with standard error bars. D-E) Fano Factor when cue was out of response field versus Fano Factor when cue was in response field for VTM, VSBP, and VSB cells.

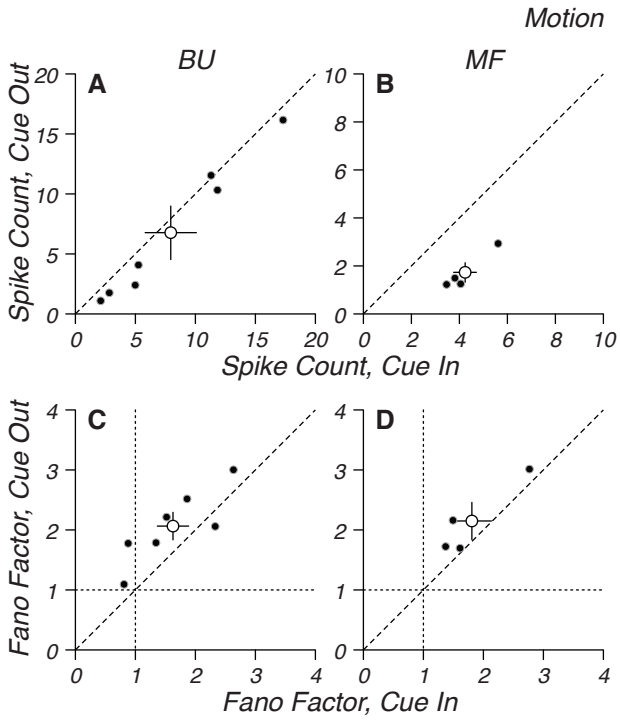

Figure 6.33: Encoding cue position in spike count in interval 100-200 ms following motion onset, part 2. A-B) Spike count for trials in which the cue was out of the response field versus spike count when the cue was in the response field for BU and MF cells. Open symbol is population average with standard error bars. C-D) Fano Factor when cue was out of response field versus Fano Factor when cue was in response field for BU and MF cells.

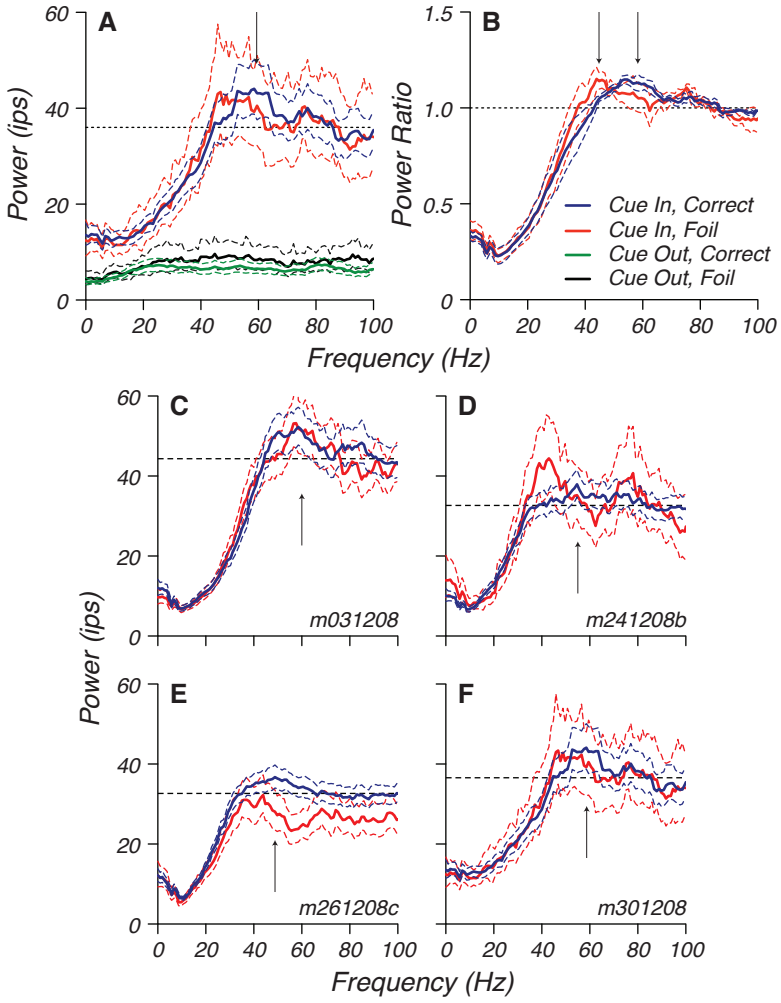

Figure 6.34: Spiking spectra for memory-field cells during initial memory period. A) Example memory-field cell spectra for cue in response field. Solid lines are averages across trials and dashed envelopes are 95% jackknife confidence intervals. Trials are grouped according to whether the subject chose with the cued signal (solid blue line or solid green line) or with the foil (solid red line or solid black line) when the cue was either in or out of the response field. A significant peak in the power spectrum is indicated by the arrow, which denotes the power and frequency of a peak significantly greater than the high-frequency limiting power (dashed line). B) Population average spectra for memory-field cells. Power for each cell is normalized by the high-frequency limit prior to averaging. Dashed envelopes represent the standard error of the mean. C-F) Spiking spectra for all four MF cells from which data was collected during the task.

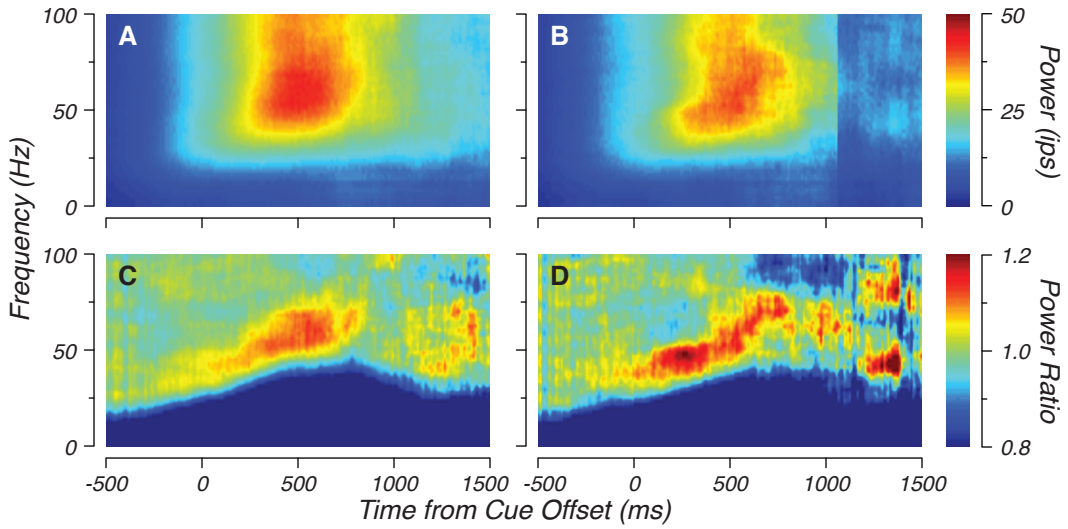

Figure 6.35: Population average spectrograms for memory-field cells aligned on cue offset. Time frequency bandwidth is 250 ms by 14 Hz. A) Spectrogram for cue in response field and choices guided by the cued signal. B) Spectrogram for cue in response field and choices guided by the foil signal. C) Population average power ratio for cue in response field and choices guided by the cued signal. Power is normalized by the mean firing rate within the 250 ms moving window of the spectrogram. D) Normalized population average power for cue in response field and choices guided by the foil signal.

Figure 6.36 *(following page)*: Coding properties of memory-field cells recorded during task. A-D) Mean firing rate aligned to cue offset; blue lines represent cue in response field, green represents cue out. E-H) Fano Factor in a 250 ms moving window. Blue lines represent cue in response field, green out. I-L) Mutual Information between cue condition and spike count in 250 ms moving window. M-P) Normalized spectrograms; color scale as in Figure 6.34. Time frequency bandwidth is 250 ms by 14 Hz. Q-T) Peak frequency within each time window in spectrogram. U-X) Power ratio at peak frequency within each time window in spectrogram.

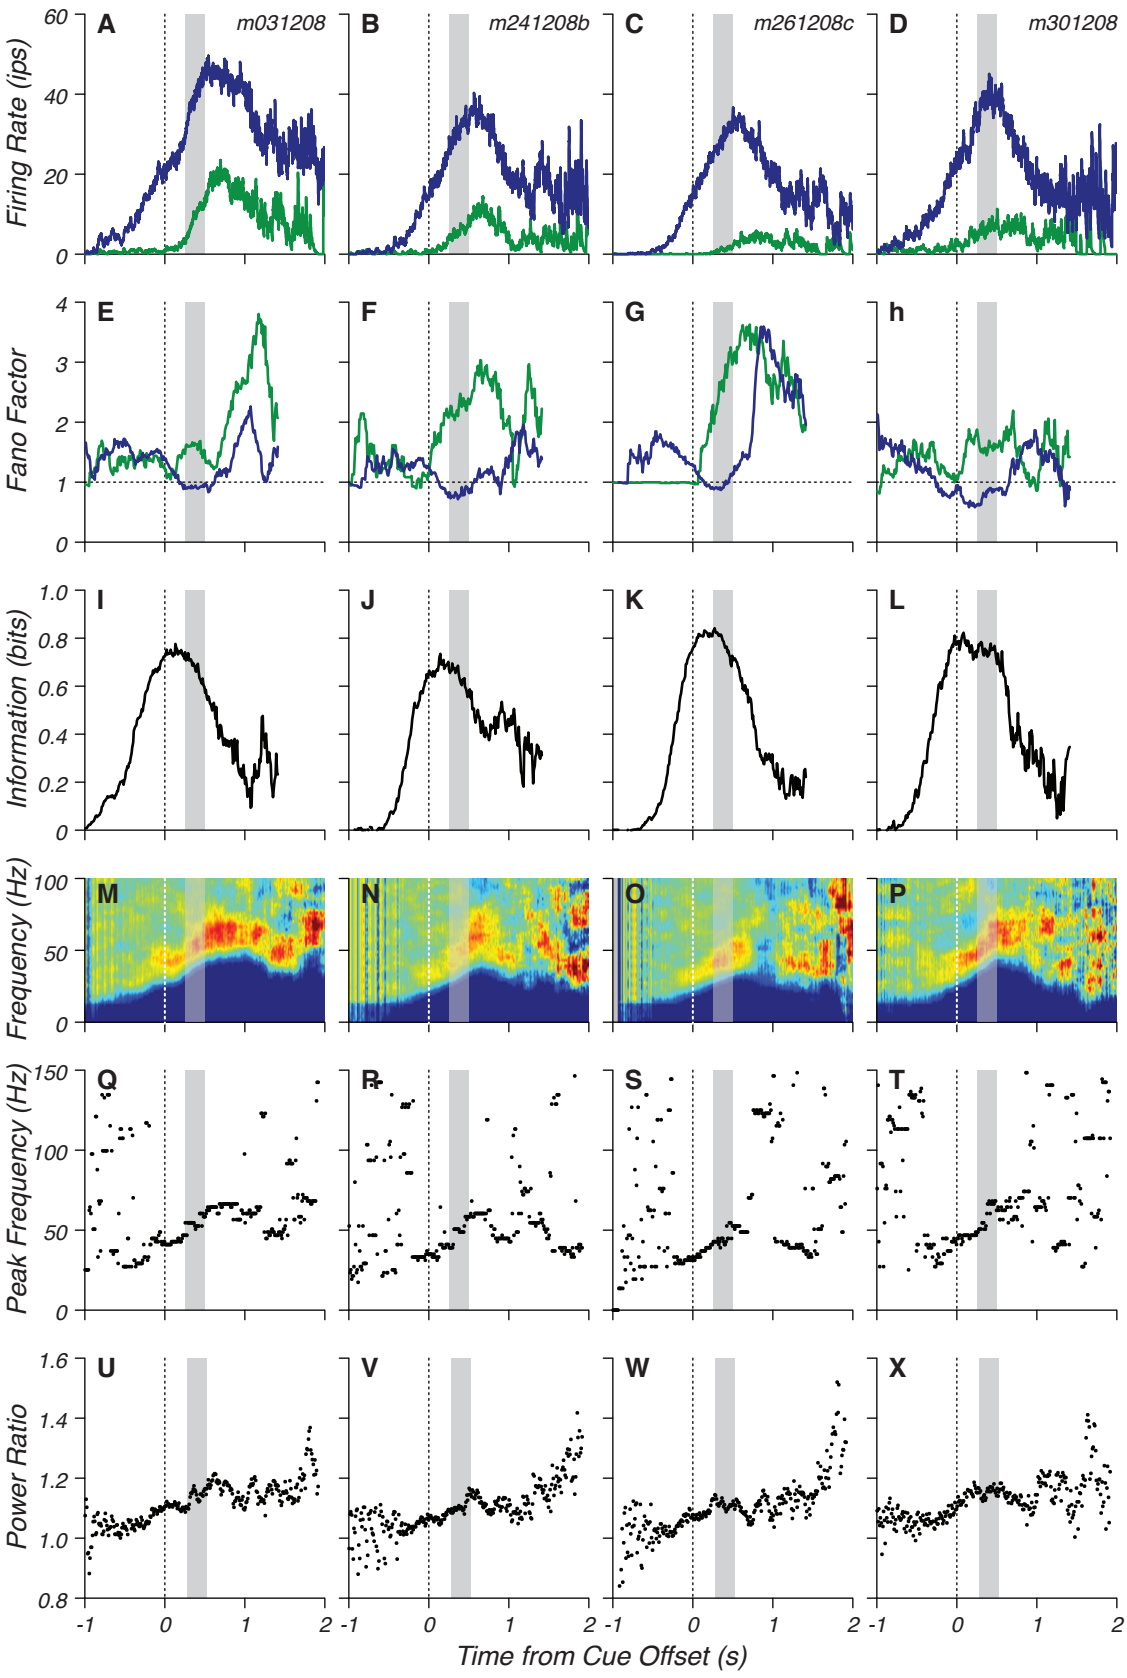
A
m031208
B
m241208b
C
m261208c
D
m301208
Firing Rate (ips)
60
40
20
0
E
F
G
h
Fano Factor
4
3
2
1
0
I
J
K
L
Information (bits)
1.0
0.8
0.6
0.4
0.2
0
M
N
O
P
Frequency (Hz)
100
50
0
Q
R
S
T
Peak Frequency (Hz)
150
100
50
0
U
V
W
X
Power Ratio
1.6
1.4
1.2
1.0
0.8
-1
0
1
2
Time from Cue Offset (s)

Figure 6.37 *(following page)*: Quantifying periodicity in each trial during memory period for each trial. A-B) Magnitude of complex regression coefficient as a function of frequency for a trial in which there was a significant peak near 60 Hz and for a trial in which there was no significant peak. C-D) Spiking spectrum for trial with significant peak and for trial without significant peak. Spectra are in black; residual spectra, with the oscillatory component removed, are in gray. E) Mean power spectra across trials (black) and across residual spectra (gray). Dashed lines are 95% jackknife confidence intervals.

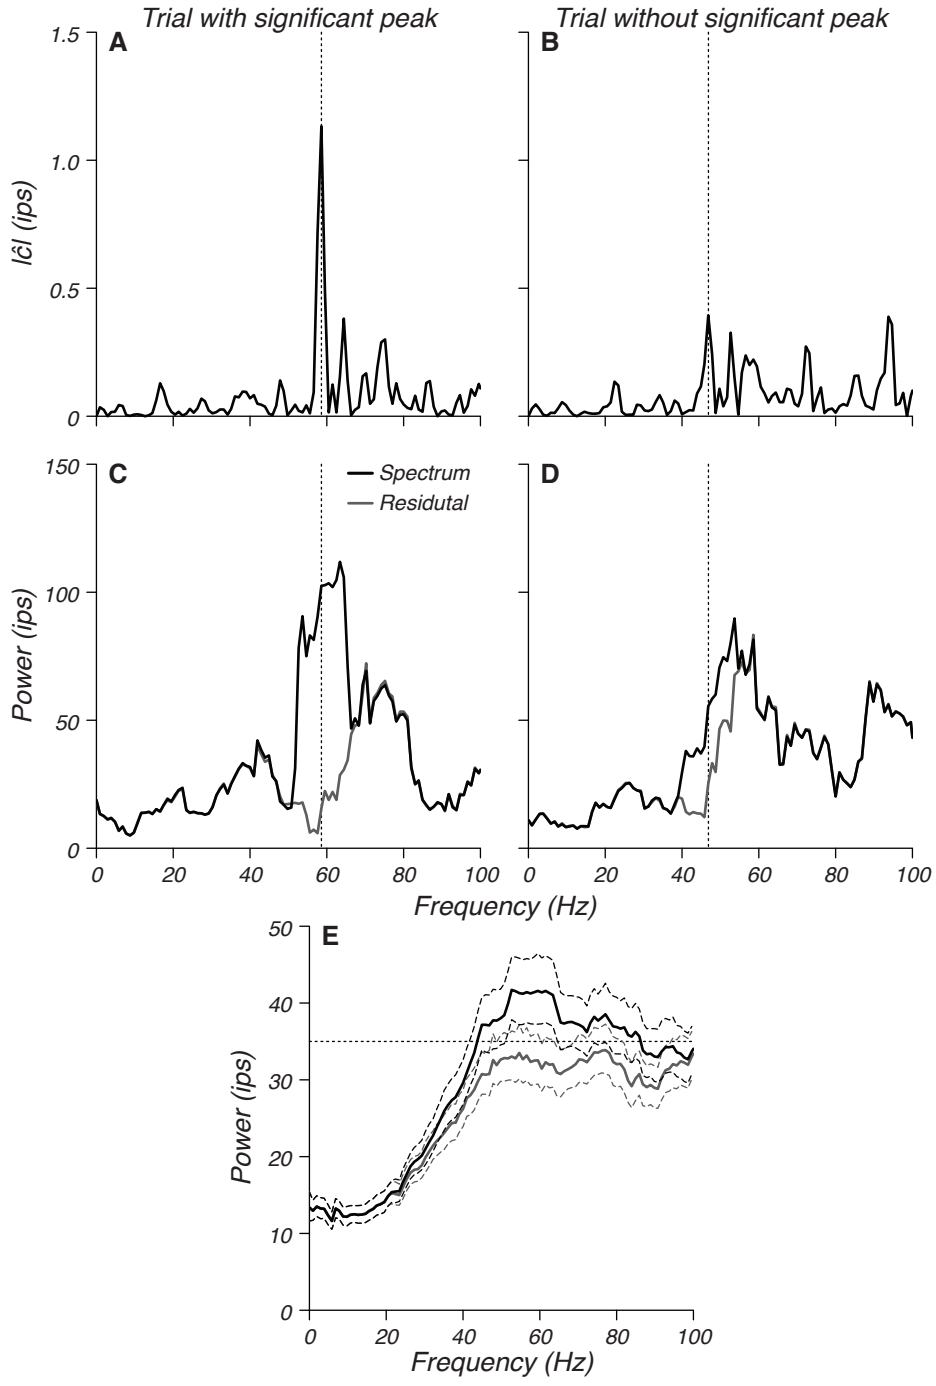
Trial with significant peak
Trial without significant peak
A
B
C
D
E
Spectrum
Residutal
|ĉl| (ips)
Power (ips)
Frequency (Hz)
1.5
1.0
0.5
0
150
100
50
0
50
40
30
20
10
0
0
20
40
60
80
100

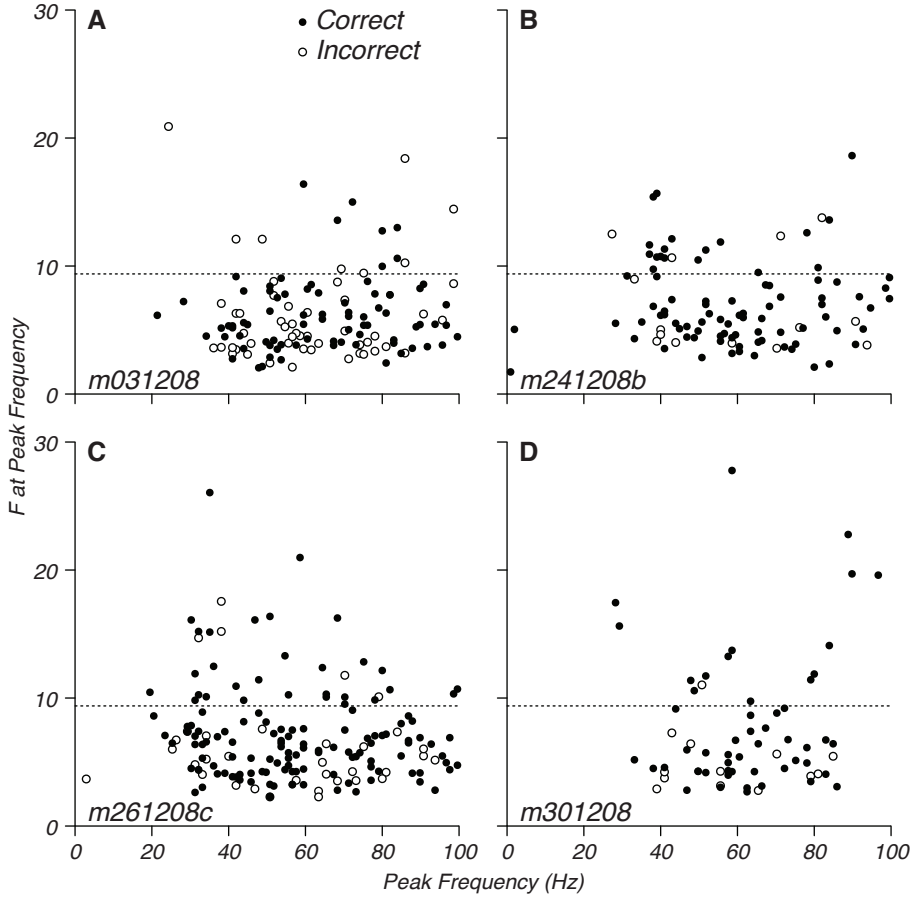

Figure 6.38: A-D) Per-trial periodicity as a function of peak frequency for trials in which the monkey based decisions on the cued signal and based choices on the foil signal. Solid symbols are correct choices and open symbols are choices with the foil. Horizontal dashed-line is 99% confidence threshold for F statistic on per-trial periodicity.

Chapter 7

Discussion

7.1 Restatement of key findings

The primary objective of this work was to determine whether or not the primate superior colliculus is necessary for selecting which signals guide perceptual judgements when multiple possible stimuli are in the visual field. As described in Chapter 2, we found that inactivating the superior colliculus caused the monkeys to base decisions on whichever signals were present in the unaffected visual field. In addition, this deficit in selection required the presence of signal in the unaffected field and did not occur if there was only a stochastic motion patch in that location. This result would seem to militate against an explanation based on irrepressible orienting of attention driven by the presence of a visual stimulus, as might have been expected from a lesion of the superior colliculus. Instead, the lesions of the SC caused a selection deficit. In addition, we found that motion discrimination was marginally impaired by the SC lesion, even when the effects on saccade selection were accounted for, which was itself a surprising result in light of some previous work that had shown that MT response properties in anesthetized monkeys were largely unchanged by an isolated lesion of the SC.

In order to explain these findings, we conducted additional detailed analysis of behavior on the task and performed a single-unit recording study to determine

what functional groups of neurons might underly the control of selective attention in the primate superior colliculus.

Detailed examination of the monkeys' performance on the selective attention task revealed that the monkeys appeared to pool sensory evidence across the visual field when deciding which direction to report. I reached this conclusion because the degree of conflict between the cued stimulus and the foil stimulus changed the response and reaction time distributions exactly as would be expected if the monkeys based decisions on a weighted sum of the evidence presented in the two locations. At the very least, considerable distortion of the information at the selected stimulus is contributed by the foil signal. It would seem possible that the algorithm by which the monkeys appear to solve this task is some form of a weighted average or a biased competition. Although additional modeling of the choice distributions and reaction times would probably yield quantitative support for different accounts of the monkeys' selection scheme, the large idiosyncratic asymmetries found in both monkeys would make this effort challenging.

In the course of the single-unit study, we found a group of neurons in the SC which appear to encode working memory in the form of maintained oscillatory activity. During a memory-guided saccade task, these neurons only fired when the target was present in their movement fields, when the monkey was remembering the target location (but only if it had been presented in their response fields) and during the saccade into their movement field. When the monkey was remembering the location, but not when the target was present on the screen, the spiking activity was oscillatory. This would seem to indicate that the oscillatory activity is associated with working memory, and is similar to response patterns seen in parietal cortex. By analogy to the phenomena in parietal cortex, we named this group *memory-field* cells. During the selective attention task, the memory-field cells were the most informative about the cue condition in a brief period of time immediately following the offset of the cues. This was associated with a substantial difference in firing rate and a large decrease in inter-trial variability, which would indicate that the spiking activity was

temporally structured. In fact, during this interval there was a time-limited increase in oscillatory spiking activity.

I also recorded from buildup units, a group frequently studied for covert attention tasks because of their link with target selection for eye movements. I found that although their mean firing rate increased with attention, the activity was so variable that the cells were not informative about cue condition. I also found that the cells passed through a fixed firing rate point at the same time prior to the saccade regardless of the cue condition, which would seem to be consistent with the view that these cells are primarily involved in target selection for saccades. The mean difference in firing rate is probably associated with the general tendency to orient to attended locations and the large variability is an adaptive mechanism to prevent the population of buildup units from reaching a state in which the saccade is automatically triggered.

Finally, I recorded from a group of units called visually-triggered movement cells, which produced large transient discharges in response to the cue onsets and the onsets of the motion pulses, but did not respond to the saccadic eye movements. Although the mean firing rates of the neurons were higher when the cue was in the neuron's response field, the firing rate was so variable that only during the visual transient was the activity consistent enough to encode information about the cue location. In addition, the time of the onset of the visual transient was more rapid by about 10 ms when the cue was in the neuron's response field, which could be consistent with a model that the neurons are controlling the timing of stimulus processing as a method of selecting which signal will drive choice.

Taken together, these results suggest at least three possible mechanisms underlying extinction and control of attention. First is that the transient discharge of the visually-triggered movement cells acts as an additive offset in a gain on the stimuli prior to their summation. If this offset is lost, then one stimulus outweighs the other, but only if there is signal present. This explanation could account for the extinction effects. Second, the temporal order of the discharges could control the sequence in which the stimuli are processed. Correct responses were often faster

than choices with the foil, which may reflect that failures in the control of processing sequence were related to errors. Therefore the sequence of the discharges, which were also influenced by the cues, could control which signal is selected to guide behavior. If one of the discharges is eliminated, then that on the unaffected side is always first, and that signal is always processed first. Finally, the third possibility is that the memory-field cells are critically important in encoding the working memory for which location in space has been cued. Without those cells, the monkey effectively has an amnesia for a location in space, and can rarely even remember that the cued signal is in that location.

7.2 Relation to models of SC involvement in selective attention

In the Introduction, I proposed three different models of how the superior colliculus might contribute to the control of selective attention. The *first model* was based primarily on a legacy of work focusing on the role of the SC in the control of visually-guided orienting movements. It held that the transition from visually-guided orienting behavior to more sophisticated orienting behavior relies essentially on the co-option of the circuitry in the SC by a cortical circuit that uses information in the periphery to synthesize targets for orienting movements. The selection of which stimuli in the visual field contribute to the generation of the synthetic target relies on circuitry resident in the superior colliculus. The *second model* held that the SC provides an important input into parietal cortex for establishing reference-frames, and that removal of that input causes disruption in the parietal cortex representation of space, leading to both intentional and sensory neglect. The *third model* held that the selection mechanisms of the superior colliculus depend largely on its placement within a basal ganglia circuit and that the selection between neural representations including those of both orienting movements and attention is mediated through the same selection mechanism localized in the basal ganglia. The latter two models both

would predict that SC inactivation should not be restricted to the selection of which stimuli guide orienting behavior, and instead predict that generalized sensory and intentional neglect will result from SC inactivation. I detailed four main hypotheses addressed in the work, and that the evidence collected regarding them would either support or contradict the three models under consideration.

The *first main hypothesis* addressed was that the SC would be required for the selection of which signals in the visual field are used to guide eye movements. The results described in Chapter 2 unequivocally demonstrate this to be the case since inactivation of a restricted region of the SC map led to neglect for the corresponding region of visual space. Therefore any of the three models could potentially be an accurate description of the role of the SC in the control of selective attention.

The *second main hypothesis* was that the SC would be required for the selection of which signals in the visual field are used to guide behavior in general and, potentially, perception. The results on the button press version of the task described in Chapter 2 demonstrate that this is the case since the neglect caused by SC inactivation applied to both manual and saccade responses. Therefore the contribution of neurons within the SC to the control of selective attention is not circumscribed to the selection of which signals guide eye movements, as held by the first model. Therefore the role of these neurons in the support of selective attention must be more general.

The *third main hypothesis* was that the SC would contain correlates of the representation of spatial information more abstractly, such as working memory, if it is a generalized retinotopic interface into the selection circuitry. The presence of oscillatory activity associated with spatial working memory is broadly in support of this conjecture. While it does not rule out the second, reference-frame based, model, it would appear to be more consistent with the third, selection based, model. Furthermore, since these neurons were active during the attention task at a discrete point in time, it seems as if the oscillatory activity represents the memory or intention of which stimuli will be selected and that it is, at the appropriate moment, proffered

up to the selection circuitry of the basal ganglia.

The *fourth main hypothesis* was that neurons in the SC would be informative about the position of the cues if they are involved in maintaining the selection of a particular location to guide behavior. Historically, the primary group of neurons suspected in this role are the buildup units, although the broad class of visuomotor units has also been studied in this context. I found that neither of these groups of neurons were informative about the spatial location of the cues. Therefore this hypothesis is substantially undermined.

Overall, two main groups of neurons appear to be involved in the control of selection mechanisms. The first, the memory-field cells, appear to be involved in the selection of which stimuli will later be employed to guide behavior. This is because they became active for a brief period of time when the monkey would have to have committed to memory where to attend, because they were the only neurons strongly informative of the cue position, and because their oscillatory activity was predictive of success or failure in the monkey's later selection. This would seem to suggest that they play a causative role in selection of stimuli and influence processing by inducing a state-change in the system corresponding to memory for the location of the cue at a given moment in time. The memory of where to attend does not appear to be encoded in the SC, so must be encoded elsewhere, such as in cortical structures. This hypothesis would lead to the conjecture that the synchronization between brain regions seen during working memory and attention is, at least in part, induced by the SC oscillatory activity. The second group, visually-triggered movement cells, produced bursts of spikes in response to the appearance of relevant stimuli. The burst was delayed when the cued signal was not in the neuron's response field. In addition, the bursts occurred *after* the stimuli had appeared. Therefore, the neurons may control the *order* in which stimuli are presented to the selection mechanism *after* they have been processed and associated with a response direction.

The group of neurons most often hypothesized to underly selection of stimuli for visual processing, the buildup cells, did not appear to be related to this function in

this task. Instead, the highly variable increase in activity associated with attention to their response fields would seem to bring the population of neurons closer to the state in which a decision about the saccade is made, thus decreasing reaction times. This could reflect the "innervation of the motor centre for the reaction" as described by James and Wundt. Of course, it also implies that a substantially more complicated mechanism exists for making decisions about saccade initiation and direction than previously appreciated; the decision is not simply based on an average of activity across the network of buildup units as compared to some other region or population, but instead must depend critically on the covariance or synchrony of activity across that network. Whether or not this is true, or what the precise mechanism is, must be addressed in future experiments.

I found no evidence in the single-unit data to support the model in which SC activity contributes to the generation of reference-frames in parietal cortex, but I also did not find any evidence in contradiction of this model. Since the absence of evidence is not the evidence of absence, the second model remains an extant possibility and is neither ruled in nor out despite the fact that much of the single-unit data is very consistent with the selection-based model. Since the single-unit data appears to support the idea that neurons in the SC are involved in discrete selection steps, I regard this as broadly supportive of the notion that the SC inactivations produced spatially specific defects in selection mechanisms. While none of the evidence specifically deals with the contribution of basal ganglia, the context of previous investigations, in particular those regarding the Sprague effect, strongly supports the notion that the SC lesions are altering the operation of a functional unit involving both the basal ganglia and the SC. Therefore the data collected to date is most consistent with the model in which the superior colliculus acts as a retinotopic interface into the basal ganglia selection circuitry for a variety of different representational forms, including working memory and movement planning.

Appendix A

Use of Muscimol for Reversible Lesion Experiments

Muscimol (3-hydroxy-5-aminomethyl-isoxazole) is a hallucinogenic or psychotomimetic alkaloid derived from the mushroom *Amanita muscaria*, and is a structural analog of γ-aminobutyric acid (GABA), a ubiquitous inhibitory neurotransmitter in the mammalian central nervous system (Snodgrass, 1978; Andrews and Johnston, 1979). The 3-isoxazolol moiety acts as a masked carboxyl group which is recognized efficiently by GABA(A) receptors but not by other macromolecules that interact with GABA in the brain (Figure A.1). It is not a substrate for GABA transaminase and is transported inefficiently by the GABA uptake system. The dissociation constant (Kd) is 16 nM for GABA and 1.8 nM for muscimol (Andrews and Johnston, 1979; Defeudis, 1980). Because of these properties, tritiated muscimol ([^{3}H]-muscimol) has gain widespread use in autoradiographic studies of the GABA receptor system. Furthermore, its long-lasting effects as a physiological agonist of GABA have made it ideal in physiological studies of GABA receptors or more generally as a method of temporarily inactivating neural activity (Krogsgaard-Larsen and Falch, 1981; Martin and Ghez, 1999).

When used to induce reversible lesions, muscimol diffuses away from the injection site and causes a suppression of neural activity roughly equivalent to its

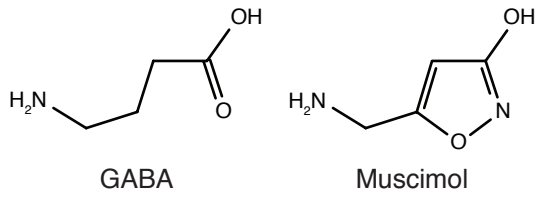

Figure A.1: Chemical structure of γ-aminobutyric acid (A) and muscimol (B).

detectable spread through the tissue. Martin (1991) measured spread of muscimol injection (1 μL of a 1 μg μL^{-1} solution injected over 4 minutes) by autoradiography and the effect on neural activity by radio-labeled glucose ([^{14}C]-glucose) uptake. Uptake of [^{14}C]-glucose was decreased maximally in a 1 mm sphere around the injection site along with a penumbra of up to 3 mm, attributed to reduced activity of neurons receiving projections from the core region. Maximum spread of muscimol to 1 mm was reached at 10 minutes and dropped to less than 20% of initial peak radioactivity by 120 minutes. In contrast, [^{14}C]-lidocaine was undetectable by 60 minutes. The duration of behavioral effects depend on the dose and can be observed even 24 hours after injection (Hikosaka and Wurtz, 1985a). Similar results were found by Arikan et al. (2002), who injected [^{3}H]-muscimol into the rat cerebellum and measured spontaneous neural activity with a grid of electrodes surrounding the injection site along with the spread of muscimol by autoradiography. Suppression of spontaneous activity was roughly equivalent to the spread of activity of muscimol, reached its maximum effect within 1.5 hours and lasted for at least 5 hours. The duration of muscimols effect and the fact that it does not blockade fibers of passage, as does lidocaine, has made muscimol the agent of choices when performing reversible lesion experiments.

Behavioral effects of muscimol last for hours after injection, and the exact length of time appears to vary based on a number of factors including the volume and concentration of the injection and the amount of ultrastructural damage to the tissue caused by repeated penetrations. The offset of behavioral effects could be due to a number of factors. Muscimol is removed from the synaptic cleft by a small basic amino-acid transporter more prevalent in neuron terminals than in glia (Johnston

et al., 1978). In addition, at the concentration of muscimol injected (a 5 μg μL^{-1} solution is 43 mM), agonist-induced endocytosis of the GABA(A) receptor could lead to the removal of muscimol-bound GABA(A) receptors from the synaptic cleft (Barnes, 2000; Kittler and Moss, 2003; Luscher and Keller, 2004). Furthermore, agonist-incuded phosphorylation of the receptor could lead to desensitization of the receptor (Luscher and Keller, 2004; Kittler and Moss, 2003). It is conceivable that such desensitization could contribute to a decrease in the apparent behavioral effects normally attributed to clearance of muscimol.

In studies of fluorescent muscimol, the diffusion properties are substantially different (in particular, fluorescent muscimol does not appear to cross myelinated fiber tracts), and the label can be detected in a tear-drop shaped pattern stretching up the shaft of the injection cannula (Allen et al., 2008). This corroborates some previous observations that tracers co-injected with muscimol can be detected along the track of the injection cannula (unpublished observations of Laurent Goffart, Centre National de la Recherche Scientifique, Aix-Marseille Universitès, France). Consequently, despite the reports of spread of muscimol into a 1 mm sphere surrounding the injection site, it is entirely possible that muscimol could spread up the shaft of the injection cannula, potentially via wicking of the solution. In the hope of minimizing such an effect, we conducted injections extremely slowly (0.5 μL in ten minutes, even slower than those described in the autoradiographic studies discussed above), and decreased the total volume of injection while increasing the concentration. Although we have every expectation that following such a strategy would make our injections the most like those verified via autoradiography, we unfortunately have no direct measure of the spread of muscimol in the superior colliculus.

Appendix B

Schematics of Injection Cannula with Internal Electrode

Schematics of the injection cannula and its connectors appear in Figure B.1 and Figure B.2.

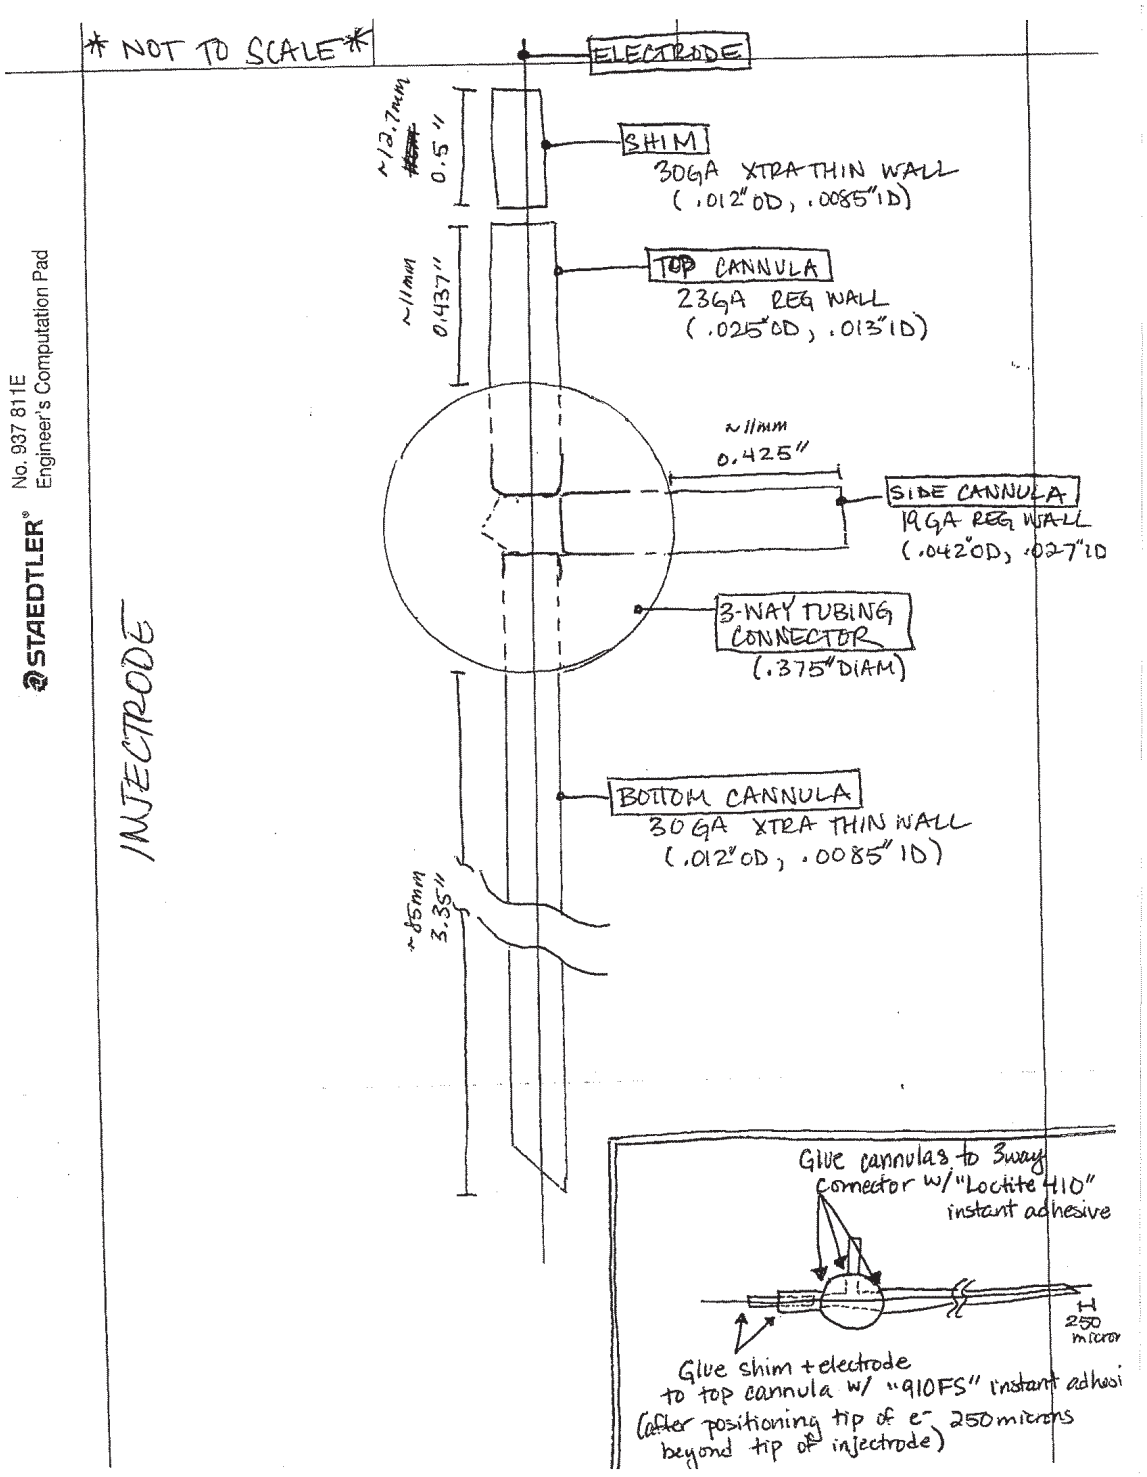

Figure B.1: Schematics of injectrode

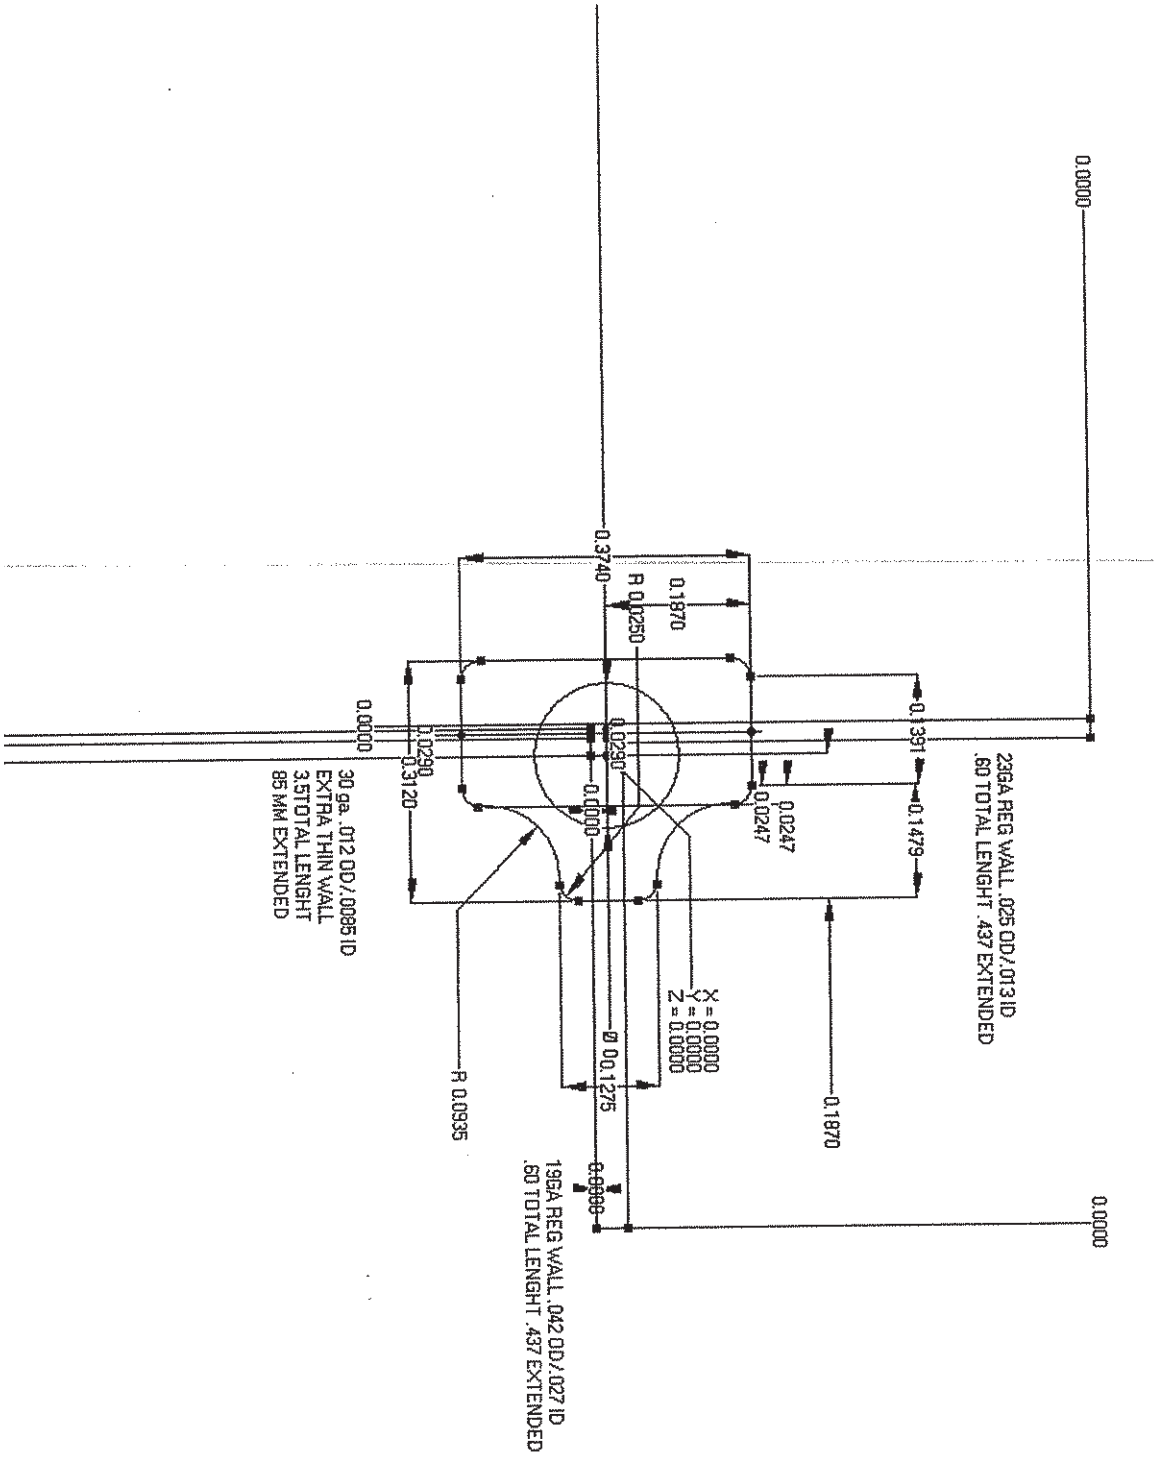

Figure B.2: Schematics of three-way tubing connector

References

Akaike, H., 1974: A new look at the statistical model identification. *IEEE Transactions on Automatic Control*, **19**(6), 716–723.

Allen, T. A., Narayanan, N. S., Kholodar-Smith, D. B., Zhao, Y., Laubach, M., and Brown, T. H., 2008: Imaging the spread of reversible brain inactivations using fluorescent muscimol. *Journal of Neuroscience Methods*, **171**(1), 30–38.

Amit, D., 1995: The Hebbian paradigm reintegrated: Local reverberations as internal representations. *Behavioral and Brain Sciences*, **18**(4), 617–625.

Andersen, R. A., and Buneo, C. A., 2002: Intentional maps in posterior parietal cortex. *Annual Review of Neuroscience*, **25**(1), 189–220.

Andrews, P. R., and Johnston, G. A. R., 1979: GABA agonists and antagonists. *Biochemical Pharmacolology*, **28**(18), 2697–2702.

Angelelli, P., DeLuca, M., and Spinelli, D., 1996: Early visual processing in neglect patients: A study with steady-state VEPs. *Neuropsychologia*, **34**(12), 1151–1157.

Arikan, R., Blake, N. M. J., Erinjeri, J. P., Woolsey, T. A., Giraud, L., and Highstein, S. M., 2002: A method to measure the effective spread of focally injected muscimol into the central nervous system with electrophysiology and light microscopy. *Journal of Neuroscience Methods*, **118**(1), 51–57.

Ashby, F. G., and Perrin, N. A., 1988: Toward a unified theory of similarity and recognition. *Psychological Review*, **95**(1), 124–150.

Awh, E., Jonides, J., and Reuter-Lorenz, P. A., 1998: Rehearsal in spatial working memory. *Journal of Experimental Psychology: Human Perception and Performance*, **24**(3), 780–790.

Bailey, B. J. R., 1980: Large sample simultaneous confidence-intervals for the multinomial probabilities based on transformations of the cell frequencies. *Technometrics*, **22**(4), 583–589.

Bair, W., Koch, C., Newsome, W., and Britten, K., 1994: Power spectrum analysis of bursting cells in area MT in the behaving monkey. *Journal of Neuroscience*, **14**(5), 2870–2892.

Barnes, E. M., 2000: Intracellular trafficking of GABA(A) receptors. *Life Sciences*, **66**(12), 1063–1070.

Basso, M. A., and Wurtz, R. H., 1997: Modulation of neuronal activity by target uncertainty. *Nature*, **389**(6646), 66–69.

Basso, M. A., and Wurtz, R. H., 1998: Modulation of neuronal activity in superior colliculus by changes in target probability. *Journal of Neuroscience*, **18**(18), 7519–7534.

Basso, M. A., and Wurtz, R. H., 2002: Neuronal activity in substantia nigra pars reticulata during target selection. *Journal of Neuroscience*, **22**(5), 1883–1894.

Baum, C., and Veeravalli, V., 1994a: Sequential multihypothesis testing and nonuniform costs with applications in hybrid serial search. *Proceedings, 1994 IEEE International Symposium on Information Theory*, 256.

Baum, C. W., and Veeravalli, V. V., 1994b: A sequential procedure for multihypothesis testing. *IEEE Transactions on Information Theory*, **40**(6), 1994–2007.

Beckstead, R. M., 1983: Long collateral branches of substantia nigra pars reticulata axons to thalamus, superior colliculus and reticular-formation in monkey and cat - multiple retrograde neuronal labeling with fluorescent dyes. *Neuroscience*, **10**(3), 767–779.

Bell, A. H., and Munoz, D. P., 2008: Activity in the superior colliculus reflects dynamic interactions between voluntary and involuntary influences on orienting behaviour. *European Journal of Neuroscience*, **28**(8), 1654–1660.

Berman, R. A., and Wurtz, R. H., 2008: Exploring the pulvinar path to visual cortex. *Progress in Brain Research*, **171**, 467–473.

Bisley, J. W., and Goldberg, M. E., 2003: Neuronal activity in the lateral intraparietal area and spatial attention. *Science*, **299**(5603), 81–86.

Bodner, M., Shafi, M., Zhou, Y. D., and Fuster, J. M., 2005: Patterned firing of parietal cells in a haptic working memory task. *European Journal of Neuroscience*, **21**(9), 2538–2546.

Boehnke, S. E., and Munoz, D. P., 2008: On the importance of the transient visual response in the superior colliculus. *Current Opinion in Neurobiology*, **18**(6), 544–551.

Brecht, M., Singer, W., and Engel, A. K., 2004: Amplitude and direction of saccadic eye movements depend on the synchronicity of collicular population activity. *Journal of Neurophysiology*, **92**(1), 424–432.

Britten, K. H., Shadlen, M. N., Newsome, W. T., and Movshon, J. A., 1992: The analysis of visual motion: a comparison of neuronal and psychophysical performance. *Journal of Neuroscience*, **12**(12), 4745–4765.

Broadbent, D. E., 1958: *Perception and Communication*. Pergamon Press, New York.

Brody, C. D., Romo, R., and Kepecs, A., 2003: Basic mechanisms for graded persistent activity: discrete attractors, continuous attractors, and dynamic representations. *Current Opinion in Neurobiology*, **13**(2), 204–211.

Bundesen, C., 1990: A theory of visual attention. *Psychological Review*, **97**(4), 523–547.

Buneo, C. A., Jarvis, M. R., Batista, A. P., and Andersen, R. A., 2003: Properties of spike train spectra in two parietal reach areas. *Experimental Brain Research*, **153**(2), 134–139.

Carello, C. D., and Krauzlis, R. J., 2004: Manipulating intent: evidence for a causal role of the superior colliculus in target selection. *Neuron*, **43**(4), 575–583.

Casagrande, V. A., and Daimond, I. T., 1974: Ablation study of superior colliculus in tree shrew *(Tupaia-glis). Journal of Comparative Neurology*, **156**(2), 207–237.

Cavanaugh, J., Alvarez, B. D., and Wurtz, R. H., 2006: Enhanced performance with brain stimulation: attentional shift or visual cue? *Journal of Neuroscience*, **26**(44), 11347–11358.

Cavanaugh, J., and Wurtz, R. H., 2004: Subcortical modulation of attention counters change blindness. *Journal of Neuroscience*, **24**(50), 11236–11243.

Chen, L. T. L., Goffart, L., and Sparks, D. L., 2001: A simple method for constructing microinjectrodes for reversible inactivation in behaving monkeys. *Journal of Neuroscience Methods*, **107**(1-2), 81–85.

Chomsung, R. D., Petry, H. M., and Bickford, M. E., 2008: Ultrastructural examination of diffuse and specific tectopulvinar projections in the tree shrew. *Journal of Comparative Neurology*, **510**(1), 24–46.

Chou, I. H., and Schiller, P. H., 1999: Temporal delays, not underactivation of detection processes may be responsible for neglect. *Behavioral and Brain Sciences*, **22**(4), 675–676.

Churchland, A. K., Kiani, R., and Shadlen, M. N., 2008: Decision-making with multiple alternatives. *Nature Neuroscience*, **11**(6), 693–702.

Ciaramitaro, V. M., Todd, W. E., and Rosenquist, A. C., 1997: Disinhibition of the superior colliculus restores orienting to visual stimuli in the hemianopic field of the cat. *Journal of Comparative Neurology*, **387**(4), 568–87.

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

Collins, C. E., Lyon, D. C., and Kaas, J. H., 2005: Distribution across cortical areas of neurons projecting to the superior colliculus in new world monkeys. *The Anatomical Record Part A*, **285A**(1), 619–627.

Comoli, E., Coizet, V., Boyes, J., Bolam, J. P., Canteras, N. S., Quirk, R. H., Overton, P. G., and Redgrave, P., 2003: A direct projection from superior colliculus to substantia nigra for detecting salient visual events. *Nature Neuroscience*, **6**(9), 974–980.

Cook, E. P., and Maunsell, J. H. R., 2002: Attentional modulation of behavioral performance and neuronal responses in middle temporal and ventral intraparietal areas of macaque monkey. *Journal of Neuroscience*, **22**(5), 1994–2004.

Cooper, B., Miya, D., and Mizumori, S., 1998: Superior colliculus and active navigation: Role of visual and non-visual cues in controlling cellular representations of space. *Hippocampus*, **8**(4), 340–372.

Corbetta, M., Kincade, M. J., Lewis, C., Snyder, A. Z., and Sapir, A., 2005: Neural basis and recovery of spatial attention deficits in spatial neglect. *Nature Neuroscience*, **8**(11), 1603–1610.

Corbetta, M., and Shulman, G. L., 2002: Control of goal-directed and stimulus-driven attention in the brain. *Nature Reviews. Neuroscience*, **3**(3), 201–215.

Cover, T. M., and Thomas, J. A., 2006: *Elements of Information Theory*. Wiley-Interscience, Hoboken, N.J., 2nd edition.

Crick, F., 1984: Function of the thalamic reticular complex: the searchlight hypothesis. *Proceedings of the National Academy of Sciences USA*, **81**(14), 4586–4590.

Crist, C. F., Yamasaki, D. S. G., Komatsu, H., and Wurtz, R. H., 1988: A grid system and a microsyringe for single cell recording. *Journal of Neuroscience Methods*, **26**(2), 117–122.

Daimond, I. T., and Hall, W. C., 1969: Evolution of neocortex. *Science*, **164**(3877), 251–262.

De Weerd, P., Desimone, R., and Ungerleider, L. G., 2003: Generalized deficits in visual selective attention after V4 and TEO lesions in macaques. *European Journal of Neuroscience*, **18**(6), 1671–1691.

Defeudis, F. V., 1980: Binding-studies with muscimol – relation to synaptic gamma-aminobutyrate receptors. *Neuroscience*, **5**(4), 675–688.

Desimone, R., and Duncan, J., 1995: Neural mechanisms of selective visual attention. *Annual Review of Neuroscience*, **18**, 193–222.

Deubel, H., and Schneider, W. X., 1996: Saccade target selection and object recognition: evidence for a common attentional mechanism. *Vision Research*, **36**(12), 1827–37.

Di Russo, F., Aprile, T., Spitoni, G., and Spinelli, D., 2008: Impaired visual processing of contralesional stimuli in neglect patients: a visual-evoked potential study. *Brain*, **131**, 842–854.

Dias, E. C., and Segraves, M. A., 1999: Muscimol-induced inactivation of monkey frontal eye field: effects on visually and memory-guided saccades. *Journal of Neurophysiology*, **81**(5), 2191–2214.

Dorris, M. C., and Munoz, D. P., 1998: Saccadic probability influences motor preparation signals and time to saccadic initiation. *Journal of Neuroscience*, **18**(17), 7015–7026.

Dosher, B. A., and Lu, Z. L., 2000: Mechanisms of perceptual attention in precuing of location. *Vision Research*, **40**(10-12), 1269–92.

Driver, J., and Mattingley, J. B., 1998: Parietal neglect and visual awareness. *Nature Neuroscience*, **1**(1), 17–22.

Duncan, J., Bundesen, C., Olson, A., Humphreys, G., Chavda, S., and Shibuya, H., 1999: Systematic analysis of deficits in visual attention. *Journal of Experimental Psychology: General*, **128**(4), 450–478.

Edelman, J. A., and Goldberg, M. E., 2003: Saccade-related activity in the primate superior colliculus depends on the presence of local landmarks at the saccade endpoint. *Journal of Neurophysiology*, **90**(3), 1728–1736.

Fecteau, J. H., Bell, A. H., and Munoz, D. P., 2004: Neural correlates of the automatic and goal-driven biases in orienting spatial attention. *Journal of Neurophysiology*, **92**(3), 1728–37.

Felsen, G., and Mainen, Z. F., 2008: Neural substrates of sensory-guided locomotor decisions in the rat superior colliculus. *Neuron*, **60**(1), 137–148.

Folk, C. L., Remington, R. W., and Wright, J. H., 1994: The structure of attentional control: contingent attentional capture by apparent motion, abrupt onset, and color. *Journal of Experimental Psychology: Human Perception and Performance*, **20**(2), 317–29.

Fries, P., 2009: Neuronal gamma-band synchronization as a fundamental process in cortical computation. *Annual Review of Neuroscience*, **32**, 209–224.

Fuchs, A. F., and Robinson, D. A., 1966: A method for measuring horizontal and vertical eye movement chronically in monkey. *Journal of Applied Physiology*, **21**(3), 1068–1070.

Fuster, J., 2000: Cortical dynamics of memory. *International Journal of Psychophysiology*, **35**(2-3), 155–164.

Ghose, G. M., and Maunsell, J. H. R., 2002: Attentional modulation in visual cortex depends on task timing. *Nature*, **419**(6907), 616–620.

Giguere, M., and Goldman-Rakic, P. S., 1988: Mediodorsal nucleus - areal, laminar, and tangential distribution of afferents and efferents in the frontal-lobe of rhesus-monkeys. *Journal of Comparative Neurology*, **277**(2), 195–213.

Gold, J., and Shadlen, M., 2000: Representation of a perceptual decision in developing oculomotor commands. *Nature*, **404**(6776), 390–394.

Gold, J. I., and Shadlen, M. N., 2007: The neural basis of decision making. *Annual Review of Neuroscience*, **30**, 535–574.

Goldberg, M. E., and Wurtz, R. H., 1972a: Activity of superior colliculus in behaving monkey. I. Visual receptive fields of single neurons. *Journal of Neurophysiology*, **35**(4), 542–559.

Goldberg, M. E., and Wurtz, R. H., 1972b: Activity of superior colliculus in behaving monkey. II. Effect of attention on neuronal responses. *Journal of Neurophysiology*, **35**(4), 560–574.

Goldman-Rakic, P. S., 1995: Cellular basis of working-memory. *Neuron*, **14**(3), 477–485.

Graybiel, A. M., 1978: Organization of nigrotectal connection - experimental tracer study in cat. *Brain Research*, **143**(2), 339–348.

Green, D. M., and Swets, J. A., 1966: *Signal detection theory and psychophysics*. R. E. Krieger Pub. Co, Huntington, N.Y.

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*, **324**(5931), 1207–1210.

Groh, J. M., and Sparks, D. L., 1996a: Saccades to somatosensory targets .II. Motor convergence in primate superior colliculus. *Journal of Neurophysiology*, **75**(1), 428–438.

Groh, J. M., and Sparks, D. L., 1996b: Saccades to somatosensory targets .III. Eye-position-dependent somatosensory activity in primate superior colliculus. *Journal of Neurophysiology*, **75**(1), 439–453.

Guo, X., and Liu, J., 2005: Stopping at the maximum of geometric brownian motion when signals are received. *Journal of Applied Probability*, **42**(3), 826–838.

Hafed, Z. M., Goffart, L., and Krauzlis, R. J., 2008: Superior colliculus inactivation causes stable offsets in eye position during tracking. *Journal of Neuroscience*, **28**(32), 8124–8137.

Hafed, Z. M., and Krauzlis, R. J., 2008: Goal representations dominate superior colliculus activity during extrafoveal tracking. *Journal of Neuroscience*, **28**(38), 9426–9439.

Halligan, P. W., and Marshall, J. C., 1994: Toward a principled explanation of unilateral neglect. *Cognitive Neuropsychology*, **11**(2), 167–206.

Halligan, P. W., Marshall, J. C., and Wade, D. T., 1992: Left on the right - allochiria in a case of left visuospatial neglect. *Journal of Neurology, Neurosurgery, and Psychiatry*, **55**(8), 717–719.

Hanes, D. P., and Schall, J. D., 1996: Neural control of voluntary movement initiation. *Science*, **274**(5286), 427–430.

Hartigan, J. A., and Hartigan, P. M., 1985: The dip test of unimodality. *Annals of Statistics*, **13**(1), 70–84.

Harting, J. K., Huerta, M. F., Frankfurter, A. J., Strominger, N. L., and Royce, G. J., 1980: Ascending pathways from the monkey superior colliculus - an autoradiographic analysis. *Journal of Comparative Neurology*, **192**(4), 853–882.

Heilman, K. M., 2004: Intentional neglect. *Frontiers in Bioscience*, **9**, 694–705.

Heilman, K. M., and Valenstein, E., 2003: *Clinical Neuropsychology*. Oxford University Press, Oxford, 4th edition.

Hikosaka, O., and Wurtz, R. H., 1983a: Visual and oculomotor functions of monkey substantia nigra pars reticulata. II. Visual responses related to fixation of gaze. *Journal of Neurophysiology*, **49**(5), 1254–1267.

Hikosaka, O., and Wurtz, R. H., 1983b: Visual and oculomotor functions of monkey substantia nigra pars reticulata. III. Memory-contingent visual and saccade responses. *Journal of Neurophysiology*, **49**(5), 1268–1284.

Hikosaka, O., and Wurtz, R. H., 1983c: Visual and oculomotor functions of monkey substantia nigra pars reticulata. IV. Relation of substantia nigra to superior colliculus. *Journal of Neurophysiology*, **49**(5), 1285–1301.

Hikosaka, O., and Wurtz, R. H., 1985a: Modification of saccadic eye movements by GABA-related substances. I. Effect of muscimol and bicuculline in monkey superior colliculus. *Journal of Neurophysiology*, **53**(1), 266–91.

Hikosaka, O., and Wurtz, R. H., 1985b: Modification of saccadic eye-movements by GABA-related substances. II. effects of muscimol in monkey substantia nigra pars reticulata. *Journal of Neurophysiology*, **53**(1), 292–308.

Hoffman, J. E., and Subramaniam, B., 1995: The role of visual attention in saccadic eye movements. *Perception & Psychophysics*, **57**(6), 787–795.

Horwitz, G. D., Batista, A. P., and Newsome, W. T., 2004: Representation of an abstract perceptual decision in macaque superior colliculus. *Journal of Neurophysiology*, **91**(5), 2281–96.

Horwitz, G. D., and Newsome, W. T., 1999: Separate signals for target selection and movement specification in the superior colliculus. *Science*, **284**(5417), 1158–1161.

Horwitz, G. D., and Newsome, W. T., 2001a: Target selection for saccadic eye movements: direction-selective visual responses in the superior colliculus. *Journal of Neurophysiology*, **86**(5), 2527–2542.

Horwitz, G. D., and Newsome, W. T., 2001b: Target selection for saccadic eye movements: prelude activity in the superior colliculus during a direction-discrimination task. *Journal of Neurophysiology*, **86**(5), 2543–2558.

Howard, M. W., Rizzuto, D. S., Caplan, J. B., Madsen, J. R., Lisman, J., Aschenbrenner-Scheibe, R., Schulze-Bonhage, A., and Kahana, M. J., 2003: Gamma oscillations correlate with working memory load in humans. *Cerebral Cortex*, **13**(12), 1369–1374.

Huerta, M. F., Vanlieshout, D. P., and Harting, J. K., 1991: Nigrotectal projections in the primate galago-crassicaudatus. *Experimental Brain Research*, **87**(2), 389–401.

Husain, M., and Kennard, C., 1997: Distractor-dependent frontal neglect. *Neuropsychologia*, **35**(6), 829–841.

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. *Nature Neuroscience*, **7**(1), 56–64.

Isa, T., and Hall, W. C., 2009: Exploring the superior colliculus in vitro. *Journal of Neurophysiology*, **102**(5), 2581–2593.

James, W., 1890: *The Principles of Psychology*. Henry Holt and Company, New York.

Janssen, P., and Shadlen, M. N., 2005: A representation of the hazard rate of elapsed time in macaque area lip. *Nature Neuroscience*, **8**(2), 234–241.

Jarvis, M. R., and Mitra, P. P., 2001: Sampling properties of the spectrum and coherency of sequences of action potentials. *Neural Computation*, **13**(4), 717–749.

Jazayeri, M., and Movshon, J. A., 2007: Integration of sensory evidence in motion discrimination. *Journal of Vision*, **7**(12), 1–7.

Joanette, Y., and Brouchon, M., 1984: Visual allesthesia in manual pointing - some evidence for a sensorimotor cerebral organization. *Brain and Cognition*, **3**(2), 152–165.

Johnson, N. L., Kemp, A. W., and Kotz, S., 2005: *Univariate Discrete Distributions*. Wiley, Hoboken, N.J., 3rd edition.

Johnston, G. A. R., Kennedy, S. M. E., and Lodge, D., 1978: Muscimol uptake, release and binding in rat-brain slices. *Journal of Neurochemistry*, **31**(6), 1519–1523.

Johnston, K., and Everling, S., 2006: Monkey dorsolateral prefrontal cortex sends task-selective signals directly to the superior colliculus. *Journal of Neuroscience*, **26**(48), 12471–12478.

Johnston, K., and Everling, S., 2008: Task-relevant output signals are sent from monkey dorsolateral prefrontal cortex to the superior colliculus during a visuospatial working memory task. *Journal of Cognitive Neuroscience*, **21**(5), 1023–1038.

Judge, S. J., Richmond, B. J., and Chu, F. C., 1980: Implantation of magnetic search coils for measurement of eye position - an improved method. *Vision Research*, **20**(6), 535–538.

Kaas, J. H., and Lyon, D. C., 2007: Pulvinar contributions to the dorsal and ventral streams of visual processing in primates. *Brain Research Reviews*, **55**(2), 285–296.

Karabelas, A. B., and Moschovakis, A. K., 1985: Nigral inhibitory termination on efferent neurons of the superior colliculus - an intracellular horseradish-peroxidase study in the cat. *Journal of Comparative Neurology*, **239**(3), 309–329.

Karnath, H. O., Himmelbach, M., and Rorden, C., 2002: The subcortical anatomy of human spatial neglect: putamen, caudate nucleus and pulvinar. *Brain*, **125**(2), 350–360.

Keller, E. L., Lee, K. M., Park, S. W., and Hill, J. A., 2008: Effect of inactivation of the cortical frontal eye field on saccades generated in a choice response paradigm. *Journal of Neurophysiology*, **100**(5), 2726–2737.

Kim, B., and Basso, M. A., 2008: Saccade target selection in the superior colliculus: A signal detection theory approach. *Journal of Neuroscience*, **28**(12), 2991–3007.

Kittler, J. T., and Moss, S. J., 2003: Modulation of GABA(A) receptor activity by phosphorylation and receptor trafficking: implications for the efficacy of synaptic inhibition. *Current Opinion in Neurobiology*, **13**(3), 341–347.

Kobayashi, S., 2003: Synaptic organization of the rat parafascicular nucleus, with special reference to its afferents from the superior colliculus and the pedunculopontine tegmental nucleus. *Brain Research*, **980**(1), 80–91.

Kowler, E., Anderson, E., Dosher, B., and Blaser, E., 1995: The role of attention in the programming of saccades. *Vision Research*, **35**(13), 1897–1916.

Krauzlis, R. J., 2003: Neuronal activity in the rostral superior colliculus related to the initiation of pursuit and saccadic eye movements. *Journal of Neuroscience*, **23**(10), 4333–4344.

Krauzlis, R. J., 2004: Activity of rostral superior colliculus neurons during passive and active viewing of motion. *Journal of Neurophysiology*, **92**(2), 949–958.

Krauzlis, R. J., Basso, M. A., and Wurtz, R. H., 2000: Discharge properties of neurons in the rostral superior colliculus of the monkey during smooth-pursuit eye movements. *Journal of Neurophysiology*, **84**(2), 876–91.

Krauzlis, R. J., and Dill, N., 2002: Neural correlates of target choice for pursuit and saccades in the primate superior colliculus. *Neuron*, **35**(2), 355–363.

Krogsgaard-Larsen, P., and Falch, E., 1981: GABA agonists. *Molecular and Cellular Biochemistry*, **38**(1), 129–146.

Krout, K. E., Loewy, A. D., Westby, G. W., and Redgrave, P., 2001: Superior colliculus projections to midline and intralaminar thalamic nuclei of the rat. *Journal of Comparative Neurology*, **431**(2), 198–216.

Kustov, A. A., and Robinson, D. L., 1996: Shared neural control of attentional shifts and eye movements. *Nature*, **384**(6604), 74–77.

Kwong, K. S., and Iglewicz, B., 1996: On singular multivariate normal distribution and its applications. *Computational Statistics & Data Analysis*, **22**(3), 271–285.

Lai, T. L., 2000: Sequential multiple hypothesis testing and efficient fault detection-isolation in stochastic systems. *IEEE Transactions on Information Theory*, **46**(2), 595–608.

Latto, R., 1978: The effects of bilateral frontal eye-field, posterior parietal or superior collicular lesions on visual search in the rhesus monkey. *Brain Research*, **146**(1), 35–50.

Lauwereyns, J., Watanabe, K., Coe, B., and Hikosaka, O., 2002: A neural correlate of response bias in monkey caudate nucleus. *Nature*, **418**(6896), 413–417.

Law, C.-T., and Gold, J. I., 2008: Neural correlates of perceptual learning in a sensory-motor, but not a sensory, cortical area. *Nature Neuroscience*, **11**(4), 505–513.

Lee, J., Williford, T., and Maunsell, J. H. R., 2007: Spatial attention and the latency of neuronal responses in macaque area V4. *Journal of Neuroscience*, **27**(36), 9632–9637.

Lepore, M., 2003: On the different mechanisms of spatial transpositions: a case of representational allochiria in clock drawing. *Neuropsychologia*, **41**(10), 1290–1295.

Levy, R., Friedman, H., Davachi, L., and Goldman-Rakic, P., 1997: Differential activation of the caudate nucleus in primates performing spatial and nonspatial working memory tasks. *Journal of Neuroscience*, **17**(10), 3870–3882.

Lock, T. M., Baizer, J. S., and Bender, D. B., 2003: Distribution of corticotectal cells in macaque. *Experimental Brain Research*, **151**(4), 455–470.

Lomber, S. G., and Payne, B. R., 1996: Removal of two halves restores the whole: Reversal of visual hemineglect during bilateral cortical or collicular inactivation in the cat. *Visual Neuroscience*, **13**(6), 1143–1156.

Loop, M. S., and Sherman, S. M., 1977: Visual discriminations of cats with cortical and tectal lesions. *Journal of Comparative Neurology*, **174**(1), 79–88.

Lovejoy, L. P., and Krauzlis, R. J., 2009: Inactivation of primate superior colliculus impairs covert selection of signals for perceptual judgments. *Nature Neuroscience*.

Luck, S. J., Chelazzi, L., Hillyard, S. A., and Desimone, R., 1997: Neural mechanisms of spatial selective attention in areas V1, V2, and V4 of macaque visual cortex. *Journal of Neurophysiology*, **77**(1), 24–42.

Luscher, B., and Keller, C. A., 2004: Regulation of GABA(A) receptor trafficking, channel activity, and functional plasticity of inhibitory synapses. *Pharmacology & Therapeutics*, **102**(3), 195–221.

Martin, J. H., 1991: Autoradiographic estimation of the extent of reversible inactivation produced by microinjection of lidocaine and muscimol in the rat. *Neuroscience Letters*, **127**(2), 160–164.

Martin, J. H., and Ghez, C., 1999: Pharmacological inactivation in the analysis of the central control of movement. *Journal of Neuroscience Methods*, **86**(2), 145–159.

Marzi, C. A., Girelli, M., Natale, E., and Miniussi, C., 2001: What exactly is extinguished in unilateral visual extinction? neurophysiological evidence. *Neuropsychologia*, **39**(12), 1354–1366.

Matsumoto, N., Minamimoto, T., Graybiel, A. M., and Kimura, M., 2001: Neurons in the thalamic CM-Pf complex supply striatal neurons with information about behaviorally significant sensory events. *Journal of Neurophysiology*, **85**(2), 960–76.

May, P. J., McHaffie, J. G., Stanford, T. R., Jiang, H., Costello, M. G., Coizet, V., Hayes, L. M., Haber, S. N., and Redgrave, P., 2009: Tectonigral projections in the primate: a pathway for pre-attentive sensory input to midbrain dopaminergic neurons. *European Journal of Neuroscience*, **29**(3), 575–587.

Mays, L. E., and Sparks, D. L., 1980: Dissociation of visual and saccade-related responses in superior colliculus neurons. *Journal of Neurophysiology*, **43**(1), 207–232.

McAlonan, K., Cavanaugh, J., and Wurtz, R. H., 2006: Attentional modulation of thalamic reticular neurons. *Journal of Neuroscience*, **26**(16), 4444–4450.

McCullagh, P., and Nelder, J. A., 1998: *Generalized Linear Models*. Chapman and Hall/CRC, Boca Raton, 2nd edition.

McMillen, T., and Holmes, P., 2006: The dynamics of choice among multiple alternatives. *Journal of Mathematical Psychology*, **50**(1), 30–57.

Mcpeek, R. M., and Keller, E. L., 2002: Saccade target selection in the superior colliculus during a visual search task. *Journal of Neurophysiology*, **88**(4), 2019–2034.

Mcpeek, R. M., and Keller, E. L., 2004: Deficits in saccade target selection after inactivation of superior colliculus. *Nature Neuroscience*, **7**(7), 757–763.

Middleton, F. A., and Strick, P. L., 2000: Basal ganglia output and cognition: Evidence from anatomical, behavioral, and clinical studies. *Brain and Cognition*, **42**(2), 183–200.

Mitchell, J. F., Sundberg, K. A., and Reynolds, J. H., 2009: Spatial attention decorrelates intrinsic activity fluctuations in macaque area v4. *Neuron*, **63**(6), 879–888.

Mitra, P., and Bokil, H., 2007: *Observed Brain Dynamics*. Oxford University Press.

Mohler, C. W., and Wurtz, R. H., 1976: Organization of monkey superior colliculus: intermediate layer cells discharging before eye movements. *Journal of Neurophysiology*, **39**(4), 722–744.

Moore, T., and Armstrong, K. M., 2003: Selective gating of visual signals by microstimulation of frontal cortex. *Nature*, **421**(6921), 370–373.

Moore, T., Armstrong, K. M., and Fallah, M., 2003: Visuomotor origins of covert spatial attention. *Neuron*, **40**(4), 671–683.

Moore, T., and Fallah, M., 2001: Control of eye movements and spatial attention. *Proceedings of the National Academy of Sciences USA*, **98**(3), 1273–1276.

Mort, D. J., Malhotra, P., Mannan, S. K., Rorden, C., Pambakian, A., Kennard, C., and Husain, M., 2003: The anatomy of visual neglect. *Brain*, **126**(9), 1986–1997.

Moschovakis, A. K., Karabelas, A. B., and Highstein, S. M., 1988a: Structure-function relationships in the primate superior colliculus. I. Morphological classification of efferent neurons. *Journal of Neurophysiology*, **60**(1), 232–262.

Moschovakis, A. K., Karabelas, A. B., and Highstein, S. M., 1988b: Structure-function relationships in the primate superior colliculus. II. Morphological identity of presaccadic neurons. *Journal of Neurophysiology*, **60**(1), 263–302.

Mountcastle, V. B., 1978: Brain mechanisms for directed attention. *Journal of the Royal Society of Medicine*, **71**(1), 14–28.

Müller, J. R., Philiastides, M. G., and Newsome, W. T., 2005: Microstimulation of the superior colliculus focuses attention without moving the eyes. *Proceedings of the National Academy of Sciences USA*, **102**(3), 524–529.

Munoz, D. P., and Wurtz, R. H., 1995: Saccade-related activity in monkey superior colliculus. I. Characteristics of burst and buildup cells. *Journal of Neurophysiology*, **73**(6), 2313–2333.

Nagy, A., Kruse, W., Rottmann, S., Dannenberg, S., and Hoffmann, K., 2006: Somatosensory-motor neuronal activity in the superior colliculus of the primate. *Neuron*, **52**(3), 525–534.

Nakayama, K., and Mackeben, M., 1989: Sustained and transient components of focal visual attention. *Vision Research*, **29**(11), 1631–1647.

Niwa, M., and Ditterich, J., 2008: Perceptual decisions between multiple directions of visual motion. *Journal of Neuroscience*, **28**(17), 4435–4445.

Nys, G., Nijboer, T., and de Haan, E., 2008: Incomplete ipsilesional hallucinations in a patient with neglect. *Cortex*, **44**, 350–352.

Olivers, C. N. L., 2008: Interactions between visual working memory and visual attention. *Frontiers in Bioscience*, **13**, 1182–1191.

Osipova, D., Takashima, A., Oostenveld, R., Fernandez, G., Maris, E., and Jensen, O., 2006: Theta and gamma oscillations predict encoding and retrieval of declarative memory. *Journal of Neuroscience*, **26**(28), 7523–7531.

Palmer, J., Huk, A. C., and Shadlen, M. N., 2005: The effect of stimulus strength on the speed and accuracy of a perceptual decision. *Journal of Vision*, **5**(5), 376–404.

Paré, M., and Hanes, D. P., 2003: Controlled movement processing: Superior colliculus activity associated with countermanded saccades. *Journal of Neuroscience*, **23**(16), 6480–6489.

Paré, M., and Wurtz, R. H., 1997: Monkey posterior parietal cortex neurons antidromically activated from superior colliculus. *Journal of Neurophysiology*, **78**(6), 3493–3497.

Paré, M., and Wurtz, R. H., 2001: Progression in neuronal processing for saccadic eye movements from parietal cortex area lip to superior colliculus. *Journal of Neurophysiology*, **85**(6), 2545–62.

Payne, B. R., Lomber, S. G., Rushmore, R. J., and Pascual-Leone, A., 2003: Cancellation of visuoparietal lesion-induced spatial neglect. *Experimental Brain Research*, **150**(3), 395–398.

Pesaran, B., Nelson, M. J., and Andersen, R. A., 2008: Free choice activates a decision circuit between frontal and parietal cortex. *Nature*, **453**(7193), 406–409.

Pesaran, B., Pezaris, J. S., Sahani, M., Mitra, P. P., and Andersen, R. A., 2002: Temporal structure in neuronal activity during working memory in macaque parietal cortex. *Nature Neuroscience*, **5**(8), 805–811.

Petersen, S. E., Robinson, D. L., and Morris, J. D., 1987: Contributions of the pulvinar to visual spatial attention. *Neuropsychologia*, **25**(1A), 97–105.

Posner, M. I., 1980: Orienting of attention. *Quarterly Journal of Experimental Psychology*, **32**, 3–25.

Posner, M. I., Cohen, Y., and Rafal, R. D., 1982: Neural systems control of spatial orienting. *Philosophical Transactions of the Royal Society of London. Series B, Biological Sciences*, **298**(1089), 187–198.

Posner, M. I., and Petersen, S. E., 1990: The attention system of the human brain. *Annual Review of Neuroscience*, **13**, 25–42.

Pouget, A., and Sejnowski, T. J., 1997: Spatial transformations in the parietal cortex using basis functions. *Journal of Cognitive Neuroscience*, **9**(2), 222–237.

Pouget, A., and Sejnowski, T. J., 2001: Simulating a lesion in a basis function model of spatial representations: Comparison with hemineglect. *Psychological Review*, **108**(3), 653–673.

Quaia, C., Aizawa, H., Optican, L. M., and Wurtz, R. H., 1998: Reversible inactivation of monkey superior colliculus. II. Maps of saccadic deficits. *Journal of Neurophysiology*, **79**(4), 2097–2110.

Raczkowski, D., Casagrande, V. A., and Diamond, I. T., 1976: Visual neglect in the tree shrew after interruption of the descending projections of the deep superior colliculus. *Experimental Neurology*, **50**(1), 14–29.

Raczkowski, D., and Diamond, I. T., 1978: Cells of origin of several efferent pathways from superior colliculus in galago-senegalensis. *Brain Research*, **146**(2), 351–357.

Rafal, R., 1994: Neglect. *Current Opinion in Neurobiology*, **4**(2), 231–236.

Rafal, R., Ward, R., and Danziger, S., 2006: Selection for action and selection for awareness: Evidence from hemispatial neglect. *Brain Research*, **1080**(1), 2–8.

Raiguel, S., Xiao, D., Marcar, V., and Orban, G., 1999: Response latency of macaque area MT/V5 neurons and its relationship to stimulus parameters. *Journal of Neurophysiology*, **82**(4), 1944–1956.

Ratcliff, R., 2006: Modeling response signal and response time data. *Cognitive Psychology*, **53**(3), 195–237.

Ratcliff, R., Cherian, A., and Segraves, M., 2003: A comparison of macaque behavior and superior colliculus neuronal activity to predictions from models of two-choice decisions. *Journal of Neurophysiology*, **90**(3), 1392–1407.

Ratcliff, R., Hasegawa, Y. T., Hasegawa, R. P., Smith, P. L., and Segraves, M. A., 2007: Dual diffusion model for single-cell recording data from the superior colliculus in a brightness-discrimination task. *Journal of Neurophysiology*, **97**(2), 1756–1774.

Ratcliff, R., and McKoon, G., 2008: The diffusion decision model: theory and data for two-choice decision tasks. *Neural Computation*, **20**(4), 873–922.

Reddi, B., and Carpenter, R., 2000: The influence of urgency on decision time. *Nature Neuroscience*, **3**, 827–830.

Redgrave, P., Prescott, T. J., and Gurney, K., 1999: The basal ganglia: a vertebrate solution to the selection problem? *Neuroscience*, **89**(4), 1009–1023.

Reed, A. V., 1973: Speed-accuracy trade-off in recognition memory. *Science*, **181**(4099), 574.

Reynolds, J. H., and Chelazzi, L., 2004: Attentional modulation of visual processing. *Annual Review of Neuroscience*, **27**(1), 611–647.

Reynolds, J. H., Chelazzi, L., and Desimone, R., 1999: Competitive mechanisms subserve attention in macaque areas V2 and V4. *Journal of Neuroscience*, **19**(5), 1736–1753.

Reynolds, J. H., and Heeger, D. J., 2009: The normalization model of attention. *Neuron*, **61**(2), 168–85.

Rezec, A., and Dobkins, K., 2004: Attentional weighting: a possible account of visual field asymmetries in visual search? *Spatial Vision*, **17**(4-5), 269–293.

Rieke, F., 1997: *Spikes: Exploring the Neural Code*. MIT Press, Cambridge, Mass.

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

Robinson, D. A., 1972: Eye-movements evoked by collicular stimulation in alert monkey. *Vision Research*, **12**(11), 1795–1808.

Robinson, D. L., and Kertzman, C., 1995: Covert orienting of attention in macaques. iii. contributions of the superior colliculus. *Journal of Neurophysiology*, **74**(2), 713—721.

Robson, J. A., and Hall, W. C., 1977: The organization of the pulvinar in the grey squirrel *(Sciurus carolinensis)*. I. Cytoarchitecture and connections. *Journal of Comparative Neurology*, **173**(2), 355–88.

Rodman, H. R., Gross, C. G., and Albright, T. D., 1990: Afferent basis of visual response properties in area mt of the macaque. II. Effects of superior colliculus removal. *Journal of Neuroscience*, **10**(4), 1154–1164.

Rorden, C., Mattingley, J. B., Karnath, H. O., and Driver, J., 1997: Visual extinction and prior entry: Impaired perception of temporal order with intact motion perception after unilateral parietal damage. *Neuropsychologia*, **35**(4), 421–433.

Rossi, A. F., Bichot, N. P., Desimone, R., and Ungerleider, L. G., 2007: Top down attentional deficits in macaques with lesions of lateral prefrontal cortex. *Journal of Neuroscience*, **27**(42), 11306–11314.

Rouiller, E. M., Tanne, J., Moret, V., and Boussaoud, D., 1999: Origin of thalamic inputs to the primary, premotor, and supplementary motor cortical areas and to area 46 in macaque monkeys: A multiple retrograde tracing study. *Journal of Comparative Neurology*, **409**(1), 131–152.

Royce, G. J., Bromley, S., and Gracco, C., 1991: Subcortical projections to the centromedian and parafascicular thalamic nuclei in the cat. *Journal of Comparative Neurology*, **306**(1), 129–155.

Rushmore, R. J., Valero-Cabre, A., Lomber, S. G., Hilgetag, C. C., and Payne, B. R., 2006: Functional circuitry underlying visual neglect. *Brain*, **129**(7), 1803–1821.

Salinas, E., and Sejnowski, T. J., 2001: Gain modulation in the central nervous system: where behavior, neurophysiology, and computation meet. *Neuroscientist*, **7**(5), 430–440.

Schiller, P. H., and Chou, I. H., 1998: The effects of frontal eye field and dorsomedial frontal cortex lesions on visually guided eye movements. *Nature Neuroscience*, **1**(3), 248–253.

Schiller, P. H., and Chou, I. H., 2000a: The effects of anterior arcuate and dorsomedial frontal cortex lesions on visually guided eye movements in the rhesus monkey: 1. Single and sequential targets. *Vision Research*, **40**(10-12), 1609–1626.

Schiller, P. H., and Chou, I. H., 2000b: The effects of anterior arcuate and dorsomedial frontal cortex lesions on visually guided eye movements in the rhesus monkey: 2. Paired and multiple targets. *Vision Research*, **40**(10-12), 1627–1638.

Schiller, P. H., and Malpeli, J., 1977: Properties and tectal projections of monkey retinal ganglion-cells. *Journal of Neurophysiology*, **40**(2), 428–445.

Schiller, P. H., Stryker, M., Cynader, M., and Berman, N., 1974: Response characteristics of single cells in monkey superior colliculus following ablation or cooling of visual-cortex. *Journal of Neurophysiology*, **37**(1), 181–194.

Schmolesky, M., Wang, Y., Hanes, D., Thompson, K., Leutgeb, S., Schall, J., and Leventhal, A., 1998: Signal timing across the macaque visual system. *Journal of Neurophysiology*, **79**(6), 3272–3278.

Shore, D. I., Spence, C., and Klein, R. M., 2001: Visual prior entry. *Psychological Science*, **12**(3), 205–212.

Silva, M. F., Maia-Lopes, S., Mateus, C., Guerreiro, M., Sampaio, J., Faria, P., and Castelo-Branco, M., 2007: Retinal and cortical patterns of spatial anisotropy in contrast sensitivity tasks. *Vision Research*, **48**, 127–135.

Singer, W., and Gray, C., 1995: Visual feature integration and the temporal correlation hypothesis. *Annual Review of Neuroscience*, **18**(1), 555–586.

Smyth, M. M., and Scholey, K. A., 1994: Interference in immediate spatial memory. *Memory & Cognition*, **22**(1), 1–13.

Snodgrass, S. R., 1978: Use of ^{3}H-muscimol for GABA receptor studies. *Nature*, **273**(5661), 392–394.

Snyder, L. H., Batista, A. P., and Andersen, R. A., 1997: Coding of intention in the posterior parietal cortex. *Nature*, **386**(6621), 167–170.

Sommer, M., and Wurtz, R., 2000: Composition and topographic organization of signals sent from the frontal eye field to the superior colliculus. *Journal of Neurophysiology*, **83**(4), 1979–2001.

Sommer, M. A., and Wurtz, R. H., 2001: Frontal eye field sends delay activity related to movement, memory, and vision to the superior colliculus. *Journal of Neurophysiology*, **85**(4), 1673–1685.

Sommer, M. A., and Wurtz, R. H., 2002: A pathway in primate brain for internal monitoring of movements. *Science*, **296**(5572), 1480–1482.

Sommer, M. A., and Wurtz, R. H., 2004a: What the brain stem tells the frontal cortex. I. Oculomotor signals sent from superior colliculus to frontal eye field via mediodorsal thalamus. *Journal of Neurophysiology*, **91**(3), 1381–1402.

Sommer, M. A., and Wurtz, R. H., 2004b: What the brain stem tells the frontal cortex. II. Role of the SC-MD-FEF pathway in corollary discharge. *Journal of Neurophysiology*, **91**(3), 1403–1423.

Sparks, D. L., 1999: Conceptual issues related to the role of the superior colliculus in the control of gaze. *Current Opinion in Neurobiology*, **9**(6), 698–707.

Sprague, J. M., 1966: Interaction of cortex and superior colliculus in mediation of visually guided behavior in cat. *Science*, **153**(3743), 1544–1547.

Takikawa, Y., Kawagoe, R., and Hikosaka, O., 2002: Reward-dependent spatial selectivity of anticipatory activity in monkey caudate neurons. *Journal of Neurophysiology*, **87**(1), 508–515.

Tanibuchi, I., and Goldman-Rakic, P., 2003: Dissociation of spatial-, object-, and sound-coding neurons in the mediodorsal nucleus of the primate thalamus. *Journal of Neurophysiology*, **89**(2), 1067–1077.

Thomson, D. J., 1982: Spectrum estimation and harmonic-analysis. *Proceedings of the IEEE*, **70**(9), 1055–1096.

Toraldo, A., Gandola, M., Loffredo, S., Rancati, A., Zanardi, G., and Bottini, G., 2005: Canceling out both the real and the spectral lines. *Brain and Cognition*, **57**(3), 226–235.

Treue, S., and Maunsell, J. H. R., 1996: Attentional modulation of visual motion processing in cortical areas mt and mst. *Nature*, **382**(6591), 539–541.

Trojanowski, J. O., and Jacobson, S., 1975: Peroxidase labeled subcortical afferents to pulvinar in rhesus monkey. *Brain Research*, **97**(1), 144–150.

Ungerleider, L. G., Desimone, R., Galkin, T. W., and Mishkin, M., 1984: Subcortical projections of area MT in the macaque. *Journal of Comparative Neurology*, **223**(3), 368–386.

Usher, M., and McClelland, J. L., 2001: The time course of perceptual choice: the leaky, competing accumulator model. *Psychological Review*, **108**(3), 550–592.

Van Wezel, R. J. A., and Britten, K. H., 2002: Motion adaptation in area MT. *Journal of Neurophysiology*, **88**(6), 3469–3476.

Vuilleumier, P. O., and Rafal, R. D., 2000: A systematic study of visual extinction: between-and within-field deficits of attention in hemispatial neglect. *Brain*, **123**(6), 1263.

Wallace, M., Wilkinson, L., and Stein, B., 1996: Representation and integration of multiple sensory inputs in primate superior colliculus. *Journal of Neurophysiology*, **76**(2), 1246–1266.

Wallace, S. F., Rosenquist, A. C., and Sprague, J. M., 1989: Recovery from cortical blindness mediated by destruction of nontectotectal fibers in the commissure of the superior colliculus in the cat. *Journal of Comparative Neurology*, **284**(3), 429–450.

Wallace, S. F., Rosenquist, A. C., and Sprague, J. M., 1990: Ibotenic acid lesions of the lateral substantia-nigra restore visual orientation behavior in the hemianopic cat. *Journal of Comparative Neurology*, **296**(2), 222–252.

Wardak, C., Ibos, G., Duhamel, J. R., and Olivier, E., 2006: Contribution of the monkey frontal eye field to covert visual attention. *Journal of Neuroscience*, **26**(16), 4228–4235.

Wardak, C., Olivier, E., and Duhamel, J. R., 2002: Saccadic target selection deficits after lateral intraparietal area inactivation in monkeys. *Journal of Neuroscience*, **22**(22), 9877–9884.

Weese, G. D., Phillips, J. M., and Brown, V. J., 1999: Attentional orienting is impaired by unilateral lesions of the thalamic reticular nucleus in the rat. *Journal of Neuroscience*, **19**(22), 10135–10139.

Werner, W., Dannenberg, S., and Hoffmann, K., 1997a: Arm-movement-related neurons in the primate superior colliculus and underlying reticular formation: comparison of neuronal activity with emgs of muscles of the shoulder, arm and trunk during reaching. *Experimental Brain Research*, **115**(2), 191–205.

Werner, W., Hoffmann, K. P., and Dannenberg, S., 1997b: Anatomical distribution of arm-movement-related neurons in the primate superior colliculus and underlying reticular formation in comparison with visual and saccadic cells. *Experimental Brain Research*, **115**(2), 206–216.

White, R. L., and Snyder, L. H., 2007: Subthreshold microstimulation in frontal eye fields updates spatial memories. *Experimental Brain Research*, **181**(3), 477–492.

Williford, T., and Maunsell, J. H. R., 2006: Effects of spatial attention on contrast response functions in macaque area V4. *Journal of Neurophysiology*, **96**(1), 40–54.

Womelsdorf, T., Anton-Erxleben, K., and Treue, S., 2008: Receptive field shift and shrinkage in macaque middle temporal area through attentional gain modulation. *Journal of Neuroscience*, **28**(36), 8934–8944.

Womelsdorf, T., and Fries, P., 2006: Neuronal coherence during selective attentional processing and sensory–motor integration. *Journal of Physiology-Paris*, **100**(4), 182–193.

Wundt, W., 1874: *Grundzüge der Physiologischen Psychologie*. Wilhelm Engelmann, Leipzig.

Wurtz, R. H., and Albano, J. E., 1980: Visual-motor function of the primate superior colliculus. *Annual Review of Neuroscience*, **3**, 189–226.

Wurtz, R. H., and Goldberg, M. E., 1972: Activity of superior colliculus in behaving monkey. III. Cells discharging before eye movements. *Journal of Neurophysiology*, **35**(4), 575–586.

Wurtz, R. H., and Mohler, C. W., 1976: Organization of monkey superior colliculus: enhanced visual response of superficial layer cells. *Journal of Neurophysiology*, **39**(4), 745–765.

Yantis, S., and Jonides, J., 1990: Abrupt visual onsets and selective attention: voluntary versus automatic allocation. *Journal of Experimental Psychology: Human Perception and Performance*, **16**(1), 121–134.

Yeshurun, Y., and Carrasco, M., 1998: Attention improves or impairs visual performance by enhancing spatial resolution. *Nature*, **396**(6706), 72–75.