The Involvement of the Primate Frontal Cortex-Basal Ganglia System in Arbitrary Visuomotor Association Learning

by

Michelle S. Machon

B.S. Brain and Cognitive Sciences
Massachusetts Institute of Technology, 2004

Submitted to the Department of Brain and Cognitive Sciences in Partial Fulfillment of the Requirements for the Degree of

DOCTOR OF PHILOSOPHY IN NEUROSCIENCE

at the

MASSACHUSETTS INSTITUTE OF TECHNOLOGY

June 2009

Signature of Author: _____

Michelle S. Machon
Department of Brain and Cognitive Sciences
May 26, 2009

Certified by: _____

Earl K. Miller, Ph.D.
Picower Professor of Neuroscience
Thesis Supervisor

Accepted by: _____

Earl K. Miller, Ph.D.
Picower Professor of Neuroscience
Chairman, Committee for Graduate Students

The Involvement of the Primate Frontal Cortex-Basal Ganglia System in Arbitrary Visuomotor Association Learning

by

Michelle S. Machon

B.S. Brain and Cognitive Sciences
Massachusetts Institute of Technology, 2004

Submitted to the Department of Brain and Cognitive Sciences on
June 1, 2009 in Partial Fulfillment of the Requirements for the
Degree of Doctor of Philosophy in Neuroscience

Abstract

It is the goal of this thesis to examine the frontal cortex-basal ganglia system during arbitrary visuomotor association learning, the forming of arbitrary links between visual stimuli and motor responses (e.g. red means stop), a fundamental learning process that underlies much of our complex behavior such as written language. The experiments contained in this thesis investigate the involvement of four components of this system in the acquisition of these associations: dorsolateral prefrontal cortex (dlPFC), caudate nucleus (Cd), frontal eye field (FEF), and the internal segment of the globus pallidus (GPi). Extracellular electrophysiological recordings were performed in awake-behaving primates performing three different learning tasks. In the different behavioral paradigms used in these studies, learning with and without reversals is investigated and compared both directly within the same experiment and indirectly across experiments.

The results of these studies suggest that a complex interplay between brain areas in the frontal cortex-basal ganglia system exists. The study of FEF during reversal learning revealed that FEF contains task-related information from the start of learning, suggesting that it may be passing information onto PFC and Cd to aid the learning process. In addition, GPi is shown to contain more specific information about the learned association during the reversal task, providing evidence for an increase in the complexity of information processing through the basal ganglia.

The in-depth study of dlPFC and Cd suggests that the frontal cortex-basal ganglia system functions only when competition between learning contexts exist. When all competition is eliminated by removing reversal learning from the behavioral task, Cd does not show involvement in the learning process. But when competition exists, both Cd and PFC show learning-related changes in task-relevant information. As determined by coherence analysis of local field potentials, communication between dlPFC and Cd is greater during reversal learning, when competition is heightened. This communication also decreases as learning progresses suggesting a role in the transfer of information between areas in facilitating the learning process. Overall, these studies further the understanding of the role of the frontal cortex-basal ganglia system in arbitrary visuomotor learning and posit that the function of the system is dependent on the existence of competition between learned information.

Thesis Supervisor: Earl K. Miller, Ph.D.
Title: Picower Professor of Neuroscience

Acknowledgements and Dedication

First and foremost, I want to thank my family. Without the support of my mom, Susan, and my brother, Paul, this thesis certainly would never have been written. Through the good times and the bad, they were both there with words of wisdom or encouragement, and they have always supported my life and career decisions. They are the two most important people to me, and I am so grateful for their love and support. This thesis is dedicated to them, as well as to my grandmother, Hilda Lord. Her memory and spirit will always be a part of me and will guide me throughout my life.

My years in the Miller Lab have been an undeniable growing and learning experience. I will always be grateful to my advisor, Earl Miller, for his guidance and support. He has taught me much about the effective communication of information, and no matter where life takes me those are skills that will be crucial for my success. I want to thank him for accepting my career choices and allowing me to explore them as a student in his lab. My two official mentors in the lab also made a big impact on my intellectual, as well as personal, development. My first mentor, Dave Freedman, will always be a sort of scientific father to me. He was so patient in teaching me techniques, and it was as his assistant that I found my passion for science. My second lab mentor was Anitha Pasupathy, and she taught me so much about collecting and then analyzing data. Under her mentorship, I developed confidence as a scientist and the ability to critically analyze scientific work and speak my mind in lab meetings. I honestly don't know what I would have done without all her help throughout the years, and not only am I grateful for her scientific mentorship, but I value our friendship and hope that we find a way to keep in touch.

I also need to thank many other Miller lab members, past and present. The Miller Lab has consistently contained top-notch scientists: smart, caring people who work hard yet always make the time to help a colleague. In the most recent years, Tim Buschman, Jefferson Roy, Jason Cromer, and Markus Siegel have endured my many questions. They have all taught me what being a good scientist and a good lab member means, and I will truly miss my intellectual conversations with them. Past lab members, Mark Histed, Melissa Warden, and Wael Asaad, taught me more than I think they realize. I specifically want to thank Mark for helping with the stimulation studies, and Wael for his help and support at the drop of a hat. Kristin MacCully was a huge support in and out of the lab. I am so grateful for her friendship and her wonderful presence in the lab for so many years. Kristin, lab just isn't the same without you! I will always appreciate the guidance and support of former lab member, Jon Wallis, during my time at UC Berkeley. In a dark time in my life, he was a beacon of light.

There are many people outside the Miller Lab who have significantly influenced my graduate studies. Most obviously, my committee members, Matt Wilson, Chris Moore, and Anitha Pasupathy, deserve a huge thanks. While it goes without saying that the thesis committee is a big part of a graduate student's studies (and fears!), I chose my committee members because I highly respect them as scientists and people. They all care passionately about science and are highly critical, deep thinkers. They have all been role models for me through my years at MIT. I would also be remiss if I did not mention the animal care staff at MIT. Without the success of the animal facility, none of this research would be possible. Specifically, I want to thank Corey Gallo and Bob Marini, two men I admire and respect for their hard work, wisdom, and endless smiles. The administrators in the department often go unnoticed, but they all work extremely hard and make the lives of students much easier and are always willing to lend a hand.

Specifically, I want to thank Denise Heintze, Susan Lanza, and Brandy Baker. These three amazing women are fountains of knowledge and support. They truly care about the students in the department. I will always be grateful to them for the countless conversations, both academic and personal in nature, and for the invaluable help they provided through all my years at MIT.

Finally, I have many people to thank that have contributed to supporting and fostering my pursuit of teaching. Jim DiCarlo taught me much about the importance of careful preparation for teaching and the effect that attention to detail has in the pursuit of any kind of excellence. His mentorship both in and out of the Systems Laboratory class meant a lot to me. Assisting Sebastian Seung in the Introduction to Neuroscience class was also a significant influence on my career decisions. I want to thank Sebastian for always treating me like a colleague and for his support in my pursuit of teaching. My experiences in 9.01 are some of my best memories from my graduate years. Marie Maloof has also been a huge source of support for my passion for teaching. She always found a way to make "it" work, whatever "it" happened to be. I have enjoyed our friendship and wish her the best in her future endeavors. This semester I also had the opportunity to be an assistant teacher for Steve Pinker's class, The Human Mind. His support and thoughtfulness as well as the camaraderie of all the teaching assistants have been a wonderful part of this semester. I have learned so much about psychology and also about great mentorship from Steve, knowledge that has changed the way I think about the teacher-student interaction.

Lastly, I'd like to thank my friends and fellow grad students. Without you guys where would I be? I probably would have dropped out because I didn't think I belonged in a group of such smart, motivated, wonderful people. You have all inspired me to be a critical scientist, a deeper thinker, and a better person. I hope the friendships we've developed will last for many years to come.

Table of Contents

Figures and Tables

Chapter 1:
General Introduction

Chapter 2:
Learning-Related Changes in FEF and GPi During Learning With Reversals

Chapter 3:
Learning-Related Changes in PFC and Cd During Learning Without Reversals

Chapter 4:
Learning-Related Changes in PFC and Cd During Learning With and Without Reversals: Analysis of Single Units and Local Field Potentials

Chapter 6:
General Methodology

Chapter 1: General Introduction

A fundamental question in neuroscience today is how memories are encoded and ultimately stored. From the individual molecules, proteins, receptors, etc. to the large-scale changes in network connections and dynamics involved in learning and memory, knowledge is progressing, but there is still much to discover. The goal of this thesis is to investigate the neural basis of the learning process involved in the formation of a particular type of memory: the arbitrary association between a visual stimulus and a motor response. The experiments performed here take a systems neuroscience approach to investigate this question, as the activity of populations of single neurons and local field potentials are recorded in different brain areas of the frontal cortex-basal ganglia system, and the learning-related dynamics of the different populations are the central focus of study.

The term "arbitrary" is used to characterize a visuomotor association when there is nothing inherent about the visual stimulus that may lead naturally to the associated response (Murray et al., 2000; Wise et al., 1996b). In contrast, a "standard" visuomotor association is what one uses to configure their hand appropriately when reaching to grasp a coffee mug (Wise et al., 1996b). When performing this standard visuomotor association, the visual input is a three-dimensional representation of the mug, and this visual-spatial input is transformed into the correct hand configuration from the inherent spatial configuration of the handle on mug, the mug's size, etc. While this type of standard visuomotor learning is crucial for an organism's successful interaction with the environment, the learning process for this ability develops early in one's life, as humans and other primates can grasp a brand new object correctly in the first interaction with the object. In contrast, arbitrary visuomotor association learning occurs throughout the lifespan and pervades our everyday lives. This ability to learn a particular behavioral response to a certain visual stimulus is what allows us to stop at red traffic lights and go at green. We can learn simple rules for the correct direction to turn a lid of a jar to open it ("righty-tighty, lefty-loosy"), and extremely complex associations between symbolic written language and the sounds we should produce to pronounce them. It is this ability to learn arbitrary visuomotor associations, and the neural basis of the learning process, that is the focus of this thesis.

Neural Basis of Arbitrary Visuomotor Association Learning

The entire process of associating a visual stimulus in the environment with a particular behavioral response involves much of the brain (Figure 1.1). First, sensory information is transmitted from the sensory periphery up to sensory cortex (red), and ultimately, motor structures execute the appropriate motor command (blue). The structures where this sensory-motor integration is thought to occur are the frontal cortical areas and the basal ganglia (purple). The hippocampus has also been implicated in this process (Cahusac et al., 1993; Wirth et al., 2003; Murray et al., 2000); however, arbitrary stimulus-response learning is not dependent on hippocampal function (White & McDonald, 2002; Packard & Knowlton, 2002). Hippocampal responses to arbitrary visuomotor learning have been shown to be weak (Cahusac et al., 1993), and the role of the hippocampus has been hypothesized to be secondary to, in support of, or in competition with, the learning occurring in cortical-basal ganglia pathways (Murray et al., 2000; White & McDonald, 2002; Poldrack et al., 2001). Thus, an investigation of the hippocampus was not included in the current studies.

The areas involved in forming arbitrary sensory-motor associations, the frontal cortex and basal ganglia, are highly interconnected. In particular, multiple, parallel, anatomically and functionally distinct loops connect cortex and basal ganglia (Alexander et al., 1986; Middleton & Strick, 2000b). In these loops, information from cortex is sent to the striatum (caudate nucleus and putamen) of the basal ganglia, processed in basal ganglia circuits, and ultimately passed to thalamus by the globus pallidus or substantia nigra pars reticulata to be sent back up to cortex. It is important to note here that while these circuits are anatomically distinct, there is much overlap in the cortical regions that are involved in each loop. For example, two specific loops described by Alexander et al. (1986) are the oculomotor loop and the dorsolateral prefrontal loop. The oculomotor loop connects the frontal eye field as well as the dorsolateral prefrontal cortex with the basal ganglia, whereas the dorsolateral prefrontal loop connects dorsolateral prefrontal cortex as well as premotor cortex with the basal ganglia.

This convergence of information onto the input structures of the basal ganglia (i.e. striatum) coupled with the output of these loops projecting primarily to frontal cortical areas provides a well-suited substrate for the integration of information for the purpose of controlling behavior. Alexander et al. (1986) described the idea of information "funneling" in this system. Not only

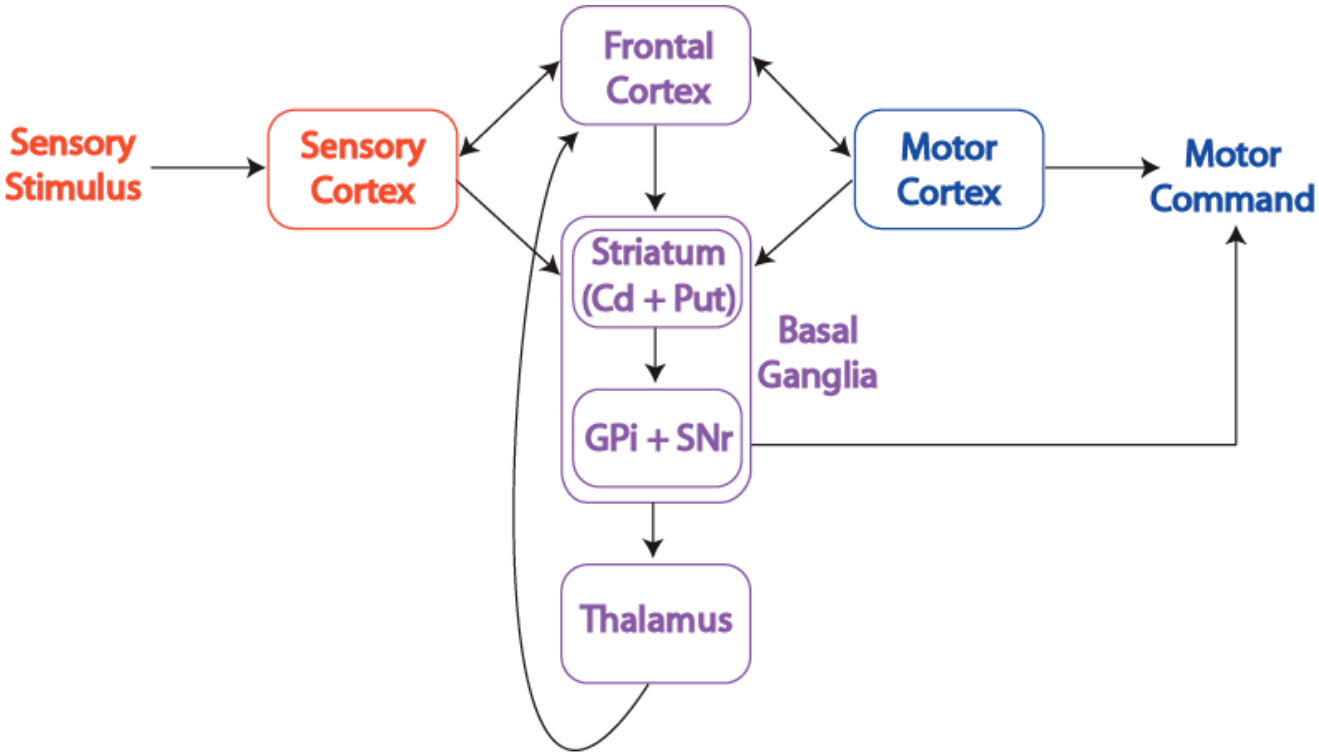

Figure 1.1
Framework for Neural Basis of Arbitrary Visuomotor Association Learning
When an arbitrary visuomotor association is learned, much of the brain is involved in the process. The sensory information about the visual stimulus enters through the periphery and ultimately reaches sensory cortex. This information then gets passed to both frontal cortical areas and the striatum, the input areas of the basal ganglia. The interaction between the frontal cortex and basal ganglia is thought to mediate visuomotor learning as both motor information (blue) and visual information (red) are present in these areas (purple). The frontal cortex and basal ganglia communicate via loops involving the thalamus; thus, the transfer of information is bi-directional between the areas. The output of this system ultimately influences the motor responses of the animal.

does anatomical convergence occur, but the idea here is that functional convergence may occur as well. This theory suggests that functional properties of neurons in the different parts of the system may differ in specificity, with output structures containing more highly specific information than input structures. With these overarching ideas of the frontal cortex-basal ganglia system in mind, I now turn to discuss each area of this system in detail.

Prefrontal Cortex

There is an extensive literature of the role of prefrontal cortex (PFC) in memory processing. Early physiological recordings of single PFC neurons demonstrated their property of sustained increases in firing rate during delay intervals in which specific memories had to be maintained in order to appropriately respond at the end of the trial (Fuster & Alexander, 1971). This maintenance property of PFC neurons provided support for the role of PFC in "working memory": the ability to keep information in mind for a specific goal on a timescale of seconds to minutes (Ungerleider, 1995; Fuster, 1973). An example used often to describe this phenomenon is the now culturally irrelevant task of having to remember a phone number long enough to dial the phone correctly. This type of short-term memory behavior has also been shown to be PFC-dependent; for example, Fuster and Alexander (1970) cooled PFC to inhibit neuronal activity and found a behavioral deficit on a delayed-response task that relies on successfully maintaining information across a delay period.

This maintenance of increased activity in PFC was then further investigated to determine the information contained in this neuronal activity and whether PFC contains any functional organization. Goldman-Rakic was a strong proponent of the idea that each subregion of PFC performs similar functions but for different information domains (Goldman-Rakic, 1996; Romanski, 2004). One study providing evidence for this idea concluded that the "what" and "where" visual pathways are also segregated in PFC (Wilson et al., 1993). This study found that ventrolateral PFC contained cells responsive to visual form and thus obtained more input from the ventral "what" visual pathway, and dorsolateral PFC contained cells responsive to visual spatial location and thus obtained more input from the dorsal "where" visual pathway. However, there is evidence against such a strict delineation of the PFC subregions (Miller and Cohen, 2001; Wallis et al., 2001).

The investigation of the function of the PFC has drifted away from a strictly working memory perspective. Newer ideas stress a higher cognitive role for PFC, such as involvement in response selection (Rowe et al., 2000; Lebedev et al., 2004; Passingham 1993), including the inhibition of unwanted responses (Robbins 1996; Funahashi et al., 1993; McDonald et al., 2007). Miller & Cohen (2001) provided a unified way to think of PFC function: PFC provides context-dependent bias signals to many brain structures that aid in sensory processing as well as response selection for the purpose of successfully producing goal-directed behavior. In order for PFC to have such high level cognitive functions, PFC cells must contain high level task-specific information. Primate physiology studies have provided this evidence, showing, for example, that PFC neurons contain information about behavior-guiding rules (Wallis et al., 2001) and visual category information that is irrespective of actual visual form similarity (Freedman et al., 2001).

With this anatomical and functional understanding of PFC, PFC seems well suited to play an integral role in the learning of arbitrary visuomotor associations. Through a series of lesion experiments performed by Petrides, evidence supported the idea that the frontal cortex is important for the visuomotor learning process (Petrides 1982, 1985a, 1985b, 1997). Specifically, because dorsolateral PFC is more highly connected to motor output regions yet still receives extensive visual information (Miller & Cohen, 2001; Fuster, 2000), this subregion of PFC has been implicated in the learning process. In the context of an arbitrary visuomotor learning task, Asaad et al. (1998) showed that cells in PFC encode both the direction of the animal's behavioral response as well as the identity of the visual objects used in the associations. In addition, learning-related changes in the response-related information encoded in the PFC cell population were presented: information about the animal's response appeared earlier in the trial in the PFC cell population as learning progressed.

Striatum

The striatum is the input structure of the basal ganglia and consists of the caudate nucleus (Cd) and the putamen. Like the PFC, the striatum receives input from much of cortex, both sensory and motor-related areas (Graybiel & Saka, 2004). Information conveyed through these vast input connections is segregated into multiple, parallel modules in the striatum. One such segregation is in the matrisomal versus striosomal compartments of the striatum (Graybiel & Ragsdale, 1978). These compartments are biochemically distinct and are thought to receive

input from distinct cortical areas (Eblen & Graybiel, 1995). Another mode of segregation of inputs occurs through two pathways in the basal ganglia, the direct and indirect pathways (Mink, 1996; Wilson, 2004). The direct pathway connects the striatum with the internal segment of the globus pallidus (GPi) directly. In contrast, striatal cells of the indirect pathway transfer information to the external segment of the globus pallidus (GPe) and the subthalamic nucleus before the information reaches GPi.

This segregation of information in the striatum may serve an important computational purpose and aid in the different functions subserved by the basal ganglia. Proposals for the functions of the basal ganglia are many. Extending and modifying the perspective that the basal ganglia's main function is the initiation of movement, Mink (1996) proposed that the basal ganglia activate desired motor programs while also inhibiting competing programs. Graybiel suggests that the basal ganglia is involved in "chunking" action sequences into discrete units then used to more easily guide behavior (Graybiel & Saka, 2004; Graybiel 1998; Jog et al., 1999, Graybiel et al., 1994). Others have argued for a more cognitive role of the basal ganglia (Middleton & Strick, 2000a, 2000b; Saint-Cyr, 2003). For example, Packard & Knowlton (2002) comprehensively reviewed evidence that the basal ganglia are responsible for stimulus-response learning and memory.

The proposed ideas of the learning and memory functions of the basal ganglia stem from both anatomical and behavioral evidence. Lesions of the dorsal striatum have shown behavioral deficits in the processes involved in learning and reversing associations (Bellebaum et al., 2008; El Massioui et al., 2007). Important anatomical inputs to the basal ganglia for this learning theory are the dopamine projections from the substantia nigra pars compacta (SNpr). These inputs modulate the cortico-striatal synapses and may guide the synaptic changes underlying learning and memory (Schultz, 2002; Kawagoe et al., 1998; Reynolds et al., 2001; Aosaki et al., 1994). Studies using methods to deplete dopamine concentrations have shown effects on stimulus-response learning and reversing (O'Neill & Brown, 2007; Lee et al., 2007). The NMDA receptor coupled with the striatal acetylcholinergic interneuron population may also play a role in the learning process, as Palencia & Ragozzino (2006) observed a learning deficit resulting from the blockage of an increase in acetylcholine efflux in the striatum combined with the infusion of an NMDA antagonist.

Animal electrophysiology studies from Ann Graybiel's group provide evidence for how the information encoding in the striatal population may change through the learning process. For example, Jog *et al*. (1999) showed that during the learning of a T-maze task, where the animal learns to associate a particular sensory instruction cue (e.g. auditory tone) with a direction to turn at the T-maze juncture, striatal neurons changed their firing patterns with learning. Initially the majority of cells fired around the time when the animal executed its behavioral response, but as learning progressed the cell population instead fired more at the start and end of the maze. This dynamic reorganization of striatal firing patterns suggests the striatum may play a crucial role in arbitrary stimulus-response association learning.

Globus Pallidus

The globus pallidus (GP) consists of two parts: the internal (GPi) and external segments (GPe). While GPe is mainly involved in information processing of the indirect pathway, GPi is one of the main outputs of the basal ganglia and is thus involved in processing of both direct and indirect pathway projections. The main output of GPi is the inhibition of thalamus. (For a comprehensive review of basal ganglia anatomy see Wilson (2004))

GP has long been thought to be involved in movement control (DeLong, 1971). Neurons in GP fire at very high rates, and their activity correlates with parameters of movement, such as movement type (e.g. flexion, extension), direction, and amplitude (DeLong, 1971; DeLong et al., 1985; Turner & Anderson, 1997). Both increases and decrease in discharge rate have been observed, although the relative frequencies of these types of firing rate changes vary between studies (DeLong, 1971; Turner & Anderson, 1997).

A functional distinction between the two output pathways of the basal ganglia, one through GP and the other through the substantia nigra pars reticulata (SNpr), is thought to exist. GP has often been studied in the context of skeletal movements (DeLong, 1971; DeLong et al., 1985; Turner & Anderson, 1997) and the study of SNpr has focused on eye movements (Hikosaka & Wurtz, 1983a, 1983b, 1983c, 1983d). However, this strict delineation of function ignores both anatomical and behavioral evidence of the functional overlap of these basal ganglia output structures. All five loops (including both "oculomotor" and "motor") outlined in DeLong et al. (1986) include both GPi and SNpr. In addition, saccade-related deficits in patients with Huntington's Disease have been demonstrated (Lasker & Zee, 1997), and using deep brain

stimulation of GPi has shown to alleviate the severity of these saccadic symptoms (Fawcett et al., 2005).

Due, at least in part, to the precise anatomical link between GPi and PFC (Middleton & Strick, 2002), higher level functions of GPi have also been proposed. Arkadir et al. (2004) provided evidence for the encoding of multiple task-related parameters, including both movement direction and trial outcome (i.e. reward prediction). Turner and Anderson (2005) have shown that GPi activity is context-dependent, where movement-related changes in firing rate differ depending on the context of the movement (e.g. in a memory task versus sensory driven task). Although much evidence exists to implicate GPi in higher level cognitive function, only one report has provided evidence of the role of GPi in arbitrary visuomotor learning (Inase et al., 2001). This study required the monkey to associate three different visual images with specific arm movements. When the animals performed previously-learned associations, firing rate changes were seen during the delay period of the task in which the animal had to wait before executing the movement associated with the previously displayed visual stimulus. During the learning of new associations, these changes in delay activity were enhanced, thus providing evidence that GPi is involved in arbitrary visuomotor learning.

Motor-Related Cortical Areas

Just as the pyramidal tract originating in primary motor cortex (M1) sends output directly to spinal cord circuits to initiate the execution of voluntary skeletal movements, FEF sends output directly to brainstem nuclei that control voluntary eye movements (for review see Krauzlis, 2005). Thus, neural coding in these areas primarily reflects essential movement parameters. For example, activity of cells in M1 has been shown to encode the force required for specific movements (Evarts, 1968) and the direction of the executed movements (Georgopoulos et al., 1982). Paralleling these findings in M1, evidence has shown that FEF encodes saccade direction and amplitude, as well as displaying coding of visual information during visually guided saccades (Bruce & Goldberg, 1985).

The activity of both M1 and FEF are influenced by supporting cortical areas, including the premotor cortex and supplementary eye field (SEF), respectively. Compared to activity in motor cortex, activity in premotor cortex shows neural coding for more complex motor parameters, such as sequences, and often activity significantly precedes movement (Mushiake et al., 1991;

Tehovnik, 2000). Similarly, neuronal coding in SEF has been shown to display higher level cognitive effects, such as attentional modulation, to a greater extend and earlier in time, (Coe et al., 2002) than similar effects seen in FEF (Schall, 2004).

Studies from the laboratory of Steven Wise have demonstrated the involvement of premotor cortex, SEF, and FEF during the acquisition of arbitrary visuomotor associations (Chen & Wise, 1995a; Chen & Wise, 1995b; Mitz et al., 1991). In all areas, both increases and decreases in firing rate were seen as learning arbitrary associations progressed. In a study directly comparing SEF and FEF (Chen & Wise, 1995b), more cells were found to be modulated by learning in SEF than in FEF. However, it is clear that all motor-related cortical areas are involved in some way in the learning process. Now that the main components of the frontal cortex-basal ganglia system have been discussed, focus now turns to the function of the system as a whole, particularly during associative learning.

Theories of the Role of FC-BG System in Arbitrary Visuomotor Association Learning

The work of Richard Passingham has been crucial to the understanding of the neural basis of arbitrary visuomotor learning. Passingham suggests that one main role of the frontal cortex is to select appropriate motor responses, and he offers that the basal ganglia may function to bias the response selection (Rowe et al., 2000; Passingham, 1993). To investigate this hypothesis, numerous human brain imaging studies have been performed, which have been helpful in determining the large-scale brain networks involved in this process. Toni et al. (2001a) described a network for executing arbitrary visuomotor associations that included prefrontal, striatal, and premotor areas. The temporal dynamics of the network were then analyzed during learning, and it was shown that activation of prefrontal cortex decreased as learning progressed, whereas activation of basal ganglia increased with learning (Toni et al., 2001b). In a subsequent study (Toni et al., 2002), the connectivity of the frontal-striatal network was analyzed during learning (using the same data set as Toni et al., 2001b), and the connectivity of the frontal-striatal network was shown to increase as learning progressed. These results suggest that PFC is involved in the initial learning of arbitrary visuomotor associations, but as learning progresses so does the communication between the frontal cortex and basal ganglia. Thus, the basal ganglia may be involved in the memory consolidation process rather than the initial learning phase. The long-term memory stored in the basal ganglia may then be able to bias PFC's response selection when needed.

Building on these ideas, Wise et al. (1996) proposed an alternative view of the function of the frontal cortex-basal ganglia system. Instead of the basal ganglia aiding in PFC's response selection, Wise et al. (1996, p. 342) argued that the PFC may aid the basal ganglia in exerting its influence in potentiating rules already learned:

> "Rather than view the basal ganglia as mediating PF's [PFC's] motor outputs, the present thesis treats PF as mediating much of the basal ganglia's influence on the remainder of the CNS [central nervous system]. [...] For example, the efferent projection from PFv [ventral PFC] to inferior temporal cortex has been proposed to result in a "top-down" suppression of sensory responses to familiar stimuli as well as in enhancement of neuronal responses to anticipated ones. As one component of the larger frontal cortex-basal ganglia system, PFv may be viewed as exerting the basal ganglia's influence over visual information processing."

Wise et al. (1996) suggests that PFC acts when new learning needs to occur or when previously learned information needs to be suppressed and that the basal ganglia acts to "train" PFC when previously learned information is relevant to the current context and thus should guide current behavior. Wise et al. (1996) stresses the role that context plays in this learning process. The basal ganglia are thought to encode the context of specific associations and thus recognize when that context occurs to enforce the context-appropriate association (Houk & Wise, 1995).

A recent study from the Wise laboratory has extended these ideas to include the direct interaction between premotor cortex and the basal ganglia (Brasted & Wise, 2004). Direct comparisons of striatal activity and premotor cortex activity were made during arbitrary visuomotor learning, and it was shown that the firing rates of the population of recorded cells in both areas changed with a similar time course as learning progressed. Thus, it seems that the role of the interaction between motor-related cortical areas and basal ganglia differs from the interaction between prefrontal cortex and basal ganglia during the learning and execution of arbitrary visuomotor associations.

A study from Earl Miller's laboratory (Pasupathy & Miller, 2005) provided evidence in support of this difference in interaction between cortical regions and basal ganglia. Instead of comparing premotor cortex and striatum, Pasupathy & Miller (2005) compared dorsolateral PFC (dlPFC) and striatum (specifically Cd). Results from this study showed that both dlPFC and Cd cell populations exhibit learning-related changes in the motor-related information encoded by the

neurons. In particular, motor-related information encoded in the cell populations appeared earlier in the trial as learning progressed. However, unlike the results of Brasted & Wise (2004), the changes in PFC and Cd showed different time courses. Changes occurred in Cd much earlier in learning than similar changes in dlPFC. Pasupathy & Miller (2005) interpreted this as evidence that basal ganglia may learn these visuomotor associations first and then transfer this information to PFC to guide behavior.

If this in fact is true, that there are slow learning mechanisms occurring in the cortex and fast learning mechanisms occurring in the basal ganglia, what benefits accompany this type of system? In describing this result and placing it in the context of the large body of knowledge regarding the learning of goal-directed behavior, Miller & Buschman (2007) argue that this particular organization benefits the organism. While an organism needs the ability to learn from every behavioral experience, a system that utilizes this information to regulate behavior on a very short timescale is prone to errors. In contrast, a system that needs multiple experiences from which to extract information about the appropriate behavior is slow and costly. However, a slower system is less error-prone and has the potential to guide behavior in complex ways. MIller & Buschman (2007) suggest that the fast learning system of the basal ganglia training the slow, yet complex learning system of the PFC ultimately produces appropriate goal-directed behavior that can be complex or abstract.

While this interpretation of Pasupathy & Miller (2005) is valid and also supported by theoretical modeling results (Daw, et al., 2005), it does disregard one crucial point. The animals in this study were performing a visuomotor learning task with serial reversals: once the stimulus-response association (e.g. image A, saccade right) was learned it was subsequently, and repeatedly, reversed (e.g. image A, saccade left). Thus, this learning task had additional requirements of inhibiting or overwriting the previous association that may be reflected in the neural activity in PFC and Cd. Or perhaps the animal (or brain) considers the two different associations (i.e. original and reversed) as different contexts (ala Houk & Wise, 1995) and the role of the basal ganglia is to recognize the context and relay this information to PFC to guide behavior appropriate to the current context.

T.W. Robbins has done a series of work to disentangle the behavioral processes involved in this sort of learning task (for review see Robbins, 2007). This work originally stemmed from the use

of the Wisconsin Card Sorting Task (Milner, 1963), where subjects are asked to sort a deck of cards according to different rules (e.g. color, shape, or number of the symbols drawn on the cards) that are learned by trial and error. Throughout the sorting process, the appropriate rule to use may change, and, through feedback, the subject must recognize this change and alter behavior accordingly. In this task, patients with damage to PFC can learn to correctly implement the first rule; however, once the rule changes the patients are unable to switch their behavior and instead continue to apply the previously appropriate rule (termed "perseveration") (Milner, 1963). This deficit has been interpreted as an inability of these patients to inhibit previously correct behavior. In the task used by Robbins (Owen et al., 1991; Owen et al., 1992; Owen et al., 1993), subjects are also asked to apply rules to sets of images, based on certain parameters of the images (e.g. curvy lines, rectangular shape). However, Robbins alters the types of changes in behavior that are required to determine if inhibitory mechanisms are disrupted producing perseveration or if new learning processes are disrupted. Findings suggest a dissociation of the contributions of basal ganglia and frontal cortex: patients with basal ganglia dysfunction find new learning processes difficult, whereas patients with frontal lobe dysfunction find inhibiting previously correct behavior difficult (Owen et al., 1993).

Overview of Studies in this Thesis

With all of these different ideas about the function of the frontal cortex-basal ganglia system during arbitrary visuomotor learning, it seems clear that the involvement of this system in the learning process is far from clear. The experiments presented in this thesis attempt to further the understanding of the overall function of the frontal cortex-basal ganglia system during learning, while at the same time also contributing knowledge of the roles of specific parts of the system. Three experiments were performed and are briefly introduced here.

The first aim of this work is to expand on the results of Pasupathy & Miller (2005). Specifically, the roles of the frontal eye field (FEF) and the internal segment of the globus pallidus (GPi) are investigated during the same learning task as Pasupathy & Miller (2005): an arbitrary visuomotor learning task that required monkeys to learn and serially reverse learned associations. Data from this experiment will help determine if other motor-related cortical areas, such as FEF, play a role in the learning process, and if the information acquired during learning is passed through the BG unaffected, or whether each part of the BG contributes unique information relevant to the learning process.

The second experiment investigates the consequences of changing the learning context on the activity in dlPFC and Cd. In this learning task, reversing associations will be eliminated. Thus, the dependence of the learning-related changes seen in Pasupathy & Miller (2005) on the specific learning context will be revealed. Perhaps the fast learning seen in Cd and the slow learning in PFC only occur during the specific task with reversals. This experiment will also shed light on the context-dependency of the BG proposed by Houk & Wise (1995).

In the third experiment, changes during learning with and without reversals will be directly compared in dlPFC and Cd. An analysis of both single cell recordings and local field potentials will be presented in order to determine the function of the frontal cortex-basal ganglia system as a whole during the different types of learning. The study of single cells will provide an understanding of the information encoded in the activity of the two populations of cells, whereas the study of oscillatory activity of dlPFC and Cd may provide insight into how these two areas are processing information locally and communicating globally.

Chapter 2: Learning-Related Changes in FEF and GPi During Learning With Reversals

Introduction

The ability to form associations between arbitrary sensory stimuli and appropriate motor responses pervades our everyday lives. This learning process underlies our ability to drive safely to work in the morning, and also provides the foundation for written language. Previously, the learning of arbitrary visuomotor associations was shown to involve both Cd and PFC (Pasupathy & Miller, 2005). However, the learning-related changes in these brain areas differed in time course: changes in Cd were rapid and occurred early in learning, whereas changes in PFC were gradual and occurred later in learning, more closely paralleling behavioral performance.

PFC and Cd are two key components of the frontal cortex-basal ganglia system that is thought to mediate this type of stimulus-response learning. But there are other components of this system that may be contributing to this learning process. The goal of the experiment described here is to investigate the involvement of two other areas during the learning of arbitrary visuomotor associations: the frontal eye field (FEF) and the internal segment of the globus pallidus (GPi). Data from FEF provides insight into the origins of the learning-related changes seen in Pasupathy & Miller (2005), and data from GPi sheds light on the internal processing occurring in the basal ganglia during the learning process.

Methods

[for more methodolgical details see Chapter 6]

Subjects

One Rhesus macaque monkey (*Macaca mulatta*), Monkey P, was used in these experiments. All animal procedures conformed to the NIH guidelines and were approved by the MIT Committee on Animal Care. It should be noted that data from Monkey P was also used in Pasupathy & Miller (2005), thus providing justification for direct comparisons between the two studies.

Behavioral Task

The animal performed a serial reversal learning task (Figure 2.1). In this task, the animal had to associate a visual stimulus with a particular saccadic response direction (e.g. see picture A,

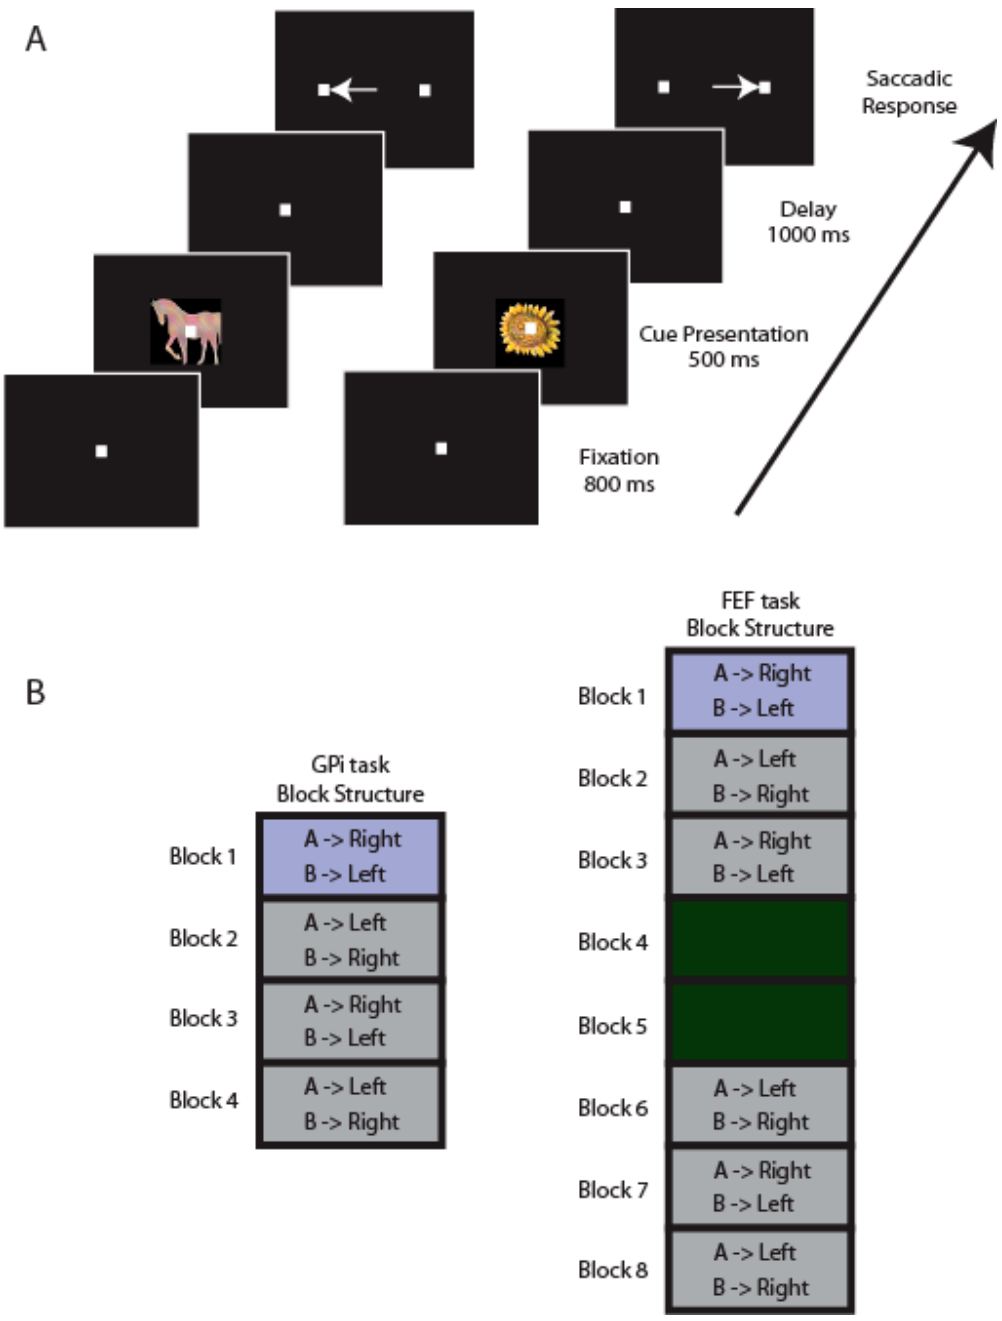

Figure 2.1
Serial Reversal Task Trial Events and Block Structure
The behavioral task used in this experiment requires the animal to associate a particular saccadic response with a particular visual stimulus. ***A***. The animal is shown a visual stimulus for 500ms, and after a one-second delay period, is required to make a saccade to one of two peripherally presented targets. Each visual stimulus is associated with a particular saccade direction. Execution of the correct saccade results in the delivery of a juice reward. ***B***. The block structure for the tasks used in collection of GPi data (left) and FEF data (right). The animal first learns new associations at the start of every session (blue) then repeatedly reverses the learned associations (gray). The FEF task also contained delayed-match-to-sample blocks (green) which were not analyzed here.

saccade right). Once the animal learned the correct associations by trial and error, the associations were reversed (e.g. see picture A, saccade left), and this pattern continued throughout the session. In the FEF recordings, the serial reversal task was interleaved with a delay-match-to-sample task; however, the analyses presented here will focus only on the data collected during reversal learning. The GPi experiment consisted only of serial reversal learning (same as Pasupathy and Miller, 2005).

Electrophysiological Data Collection

Up to 16 Tungsten microelectrodes were acutely implanted into FEF and GPi in a single session (23 sessions for FEF, 30 sessions for GPi). These recordings were performed in separate experiments; thus, these recordings are not simultaneous. Each single unit included in analysis contained at least four full behavioral learning blocks of data. A total of 200 FEF cells and 96 GPi cells were used in analysis.

Microstimulation was used to confirm FEF recording locations. In separate sessions before recordings were performed, electrodes were lowered into hypothesized FEF sites and microstimulation was applied to electrodes as the animal naturally scanned the visual world. Stimulation currents between 50 and 150μA were used for each pulse phase of the biphasic current pulses. Sites in which microstimulation elicited vector saccadic eye movements were confirmed to be located in FEF, and subsequent recordings focused on those sites.

Analytical Techniques

Analysis of the behavioral data focused on accuracy and reaction time across learning. All recording sessions (53 total: FEF, N=23; GPi, N=30) were used in this analysis. All correct and incorrect trials were used in the order in which they occurred in each learning block (trials in which the animal made a fixation error were ignored). Means +/- SE were computed across all blocks from all sessions (319 total blocks: FEF, N=129; GPi, N=190).

In order to identify cells that contained task-related information, a 2-way ANOVA (with object identity and saccade direction as factors) was performed on average firing rate in each of 4 task epochs: 100-600ms after cue onset (“Cue”), 600-1500ms after cue onset (“Delay”), 150ms before to 150ms after saccade initiation (“Saccade”), and 50-300ms after the start of reward delivery (“Reward”). A cell was considered to contain object or saccade direction information if

it *only* had a significant main effect of object or direction, respectively. A cell was considered to contain object-saccade association information if it had a main effect of object *and* direction or a significant interaction term. Significance level was set at $p < 0.05$, and was corrected for multiple comparisons using Bonferroni's correction. To graphically show cells containing task-related information, firing rate histograms of individual cells were used. Firing rate in each trial was computed in a 100ms bin, and the bin was slid in 10ms steps across time in trial. The mean +/- SE across trials was computed and is shown in Figures 2.4 and 2.5.

In order to further characterize the firing rate dynamics of the cells containing task-related information, a comparison between each cell's average firing rate (across all correct trials) in a baseline period of central fixation (from 450-200ms before the onset of the visual cue) and each cell's average firing rate (across all correct trials) in each task epoch (one-sided ttest, $p <0.05$, Bonferroni corrected). It is important to note here that all trials were used in this analysis. Thus, even cells containing task-related information may appear to have "no change" in firing rate as a result of this analysis. Selectivity measures differences between trial types, and this firing rate dynamics analysis is averaging across all trial types, so information specific to trial type is lost.

To investigate the learning-related changes in the activity of the cell populations, two measures were used: normalized firing rate and direction selectivity. [For detailed explanation of the calculation of these measures, see chapter 6.] It is important to stress that all color plots presented here are averages of a *population of cells*. The goal of this type of analysis is to understand the learning-related changes occurring on the population level, and how the observed changes are different across brain areas. Directly comparing three-dimensional plots becomes difficult; thus, two linear measures are used to compare the learning-related changes in direction information carried by these populations of cells: peak selectivity and risetime. Peak, or maximum, selectivity is used as a measure of the *strength* of information. Risetime is defined as the time to half maximum selectivity and is used as a measure of the changes in *time course* of information. The half maximum is calculated as:

$$\text{Minimum } FEV_{dir} + [\ (\text{Maximum } FEV_{dir} - \text{Minimum } FEV_{dir})\ ] / 2$$

One note about the risetime calculation in GPi: the risetimes in GPi during the cue period were calculated as time to 55% selectivity (instead of 50%) due to the very small range of selectivity strength in this population during the early trial period.

Average direction selectivity was compared across different trial types to attempt to link this saccade direction information with behavior. The different trial types used were correct trials performed on the associations that reversed, incorrect trials on these reversal associations, and correct trials performed on the two familiar associations the animal had practiced for months/years. The number of these different trial types is drastically different; thus, a titration procedure was used to equate the number of trials in each category. This procedure used the minimum number of trials performed in each category (most often this was the number of incorrect trials performed). Then, for each incorrect trial used, the closest correct reversal trial and correct familiar trial were used. Since the majority of incorrect trials occur at the start of learning, this procedure attempts to equate the trial types, not just for number, but also for time in learning. This is an important aspect of the analysis since it is shown that direction selectivity may increase or decrease as learning progresses. Once trials were chosen in this manner, average saccade direction information during the cue (100-500ms after cue onset) or saccade period (0-400ms after saccade onset) was averaged across the trials for each cell, and then averaged across cells. To determine if significant information was contained in the cell populations during these time periods of the trial, ttests were performed against zero for each trial type. Comparisons between trial types were performed using paired ttests. All p-values are presented in Tables 2.1 and 2.2.

Results

The behavioral performance of the animal across all recording sessions (N=53) is depicted in Figure 2.2. At the start of learning, mean performance dropped close to zero percent correct (13.8% +/- 1.9%) and slowly increased as learning progressed, and reaction time was greatest early in learning and decreased as learning progressed. Since no cue was given at the start of a new learning block, the animal continued to perform the previously rewarded associations for the first trial in a new block. This first trial phenomenon can also be seen in the animal's reaction times. The first trial of the block (which is usually an incorrect trial) had very low reaction time (170.1ms +/- 0.14), similar to the previously correct trials at the end of the last

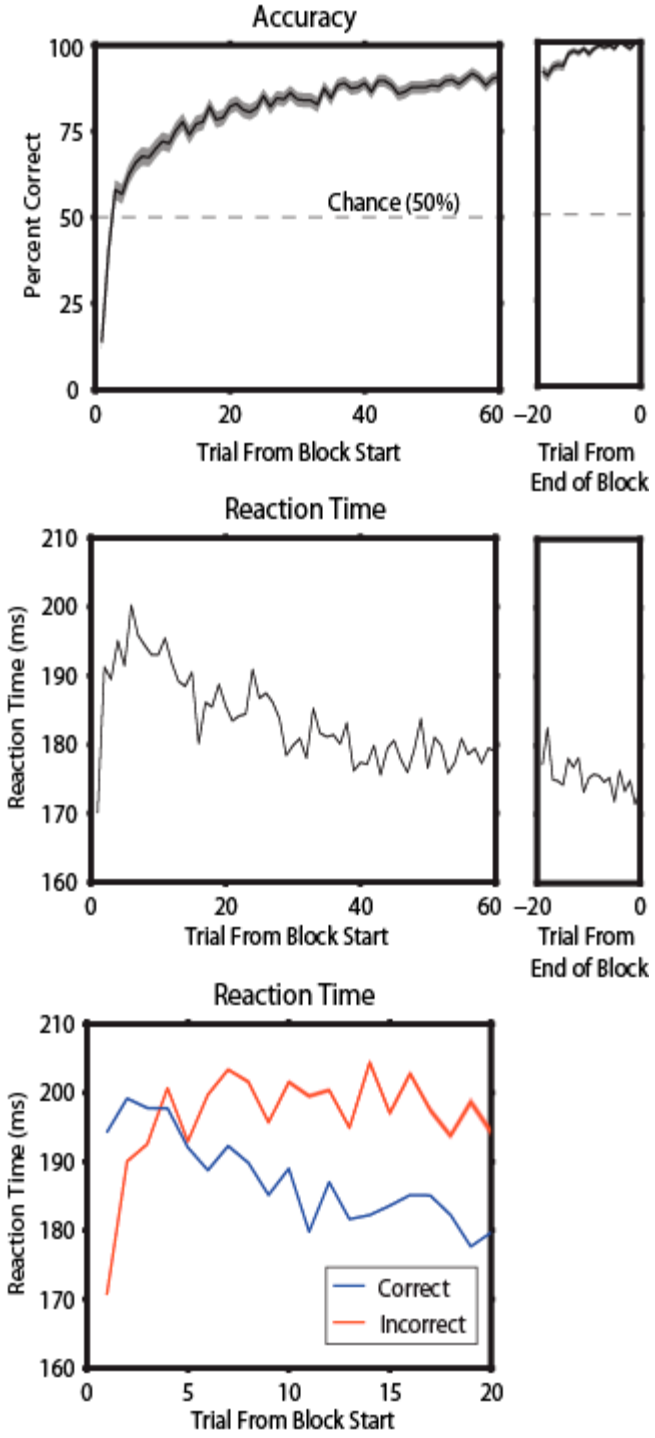

Figure 2.2
Behavioral Performance in the Visuomotor Association Learning Task with Reversals
Percent correct performance (top) and reaction times (middle and bottom) are plotted taking into account all correct and incorrect trials. On the first trial in learning, accuracy is close to 0% and the animal performs very quickly since the previously rewarded associations are still being performed. Subsequently, performance jumps to chance level (50% correct, dotted line) and reaction time sharply increases. As learning progresses, accuracy increases and reaction time decreases.

learning block. The rest of the incorrect trials, at any point during learning, had much higher reaction times, about 30ms higher, than the majority of correct trials.

In total, 200 FEF and 96 GPi cells were analyzed. In order to determine the type of information encoded by each single cell, a 2-way ANOVA with object identity and saccade direction as factors, was performed in each of 4 epochs of the trial: the cue period (100-600ms after cue onset), the delay period (600-1500ms after cue onset), the saccade period (from 150ms before saccade onset to 150ms after saccade onset), and the reward period (from 50ms after reward onset to 300ms after reward onset). A significance value of $p<0.05$ was used and p-values were corrected for multiple comparisons using Bonferroni correction. Cells with only a significant main effect of direction were classified as saccade "direction selective". Cells with only a significant main effect of object identity were classified as "object selective. And cells that had significant main effects of both direction and object or a significant interaction between object and direction were classified as object-saccade "association selective" cells.

Figure 2.3A shows the fraction of cells in each brain area that contains specific task-related information in each of the four time epochs of the trial. The FEF and GPi data are compared with the previous PFC and Cd data from Pasupathy and Miller (2005). Fractions are based on total cell counts: 200 FEF, 96 GPi, 350 PFC, and 250 Cd. Overall, there are very few cells in all of these brain areas that contain information about object identity (yellow bars), and the fraction of cells that do show this information is greatest early in the trial when the object is presented. The majority of the information encoded by the cells is either the direction of the saccade (blue bars) or more specific information about the object-saccade association (purple bars). The distribution of direction information across time in trial appears to be slightly different in the different brain areas. While FEF has a consistently large population of direction selective cells during the delay, saccade, and reward periods, the largest population of direction selective cells in Cd appears during the saccade period. PFC shows large numbers of direction cells during both the saccade and reward periods, and direction cells in GPi are most numerous during the reward period. These timing differences suggest that this saccade direction information may be transferred through the network from FEF to Cd and then to GPi and PFC.

Figure 2.3B summarizes the total fraction of direction (left panel), association (middle panel), and object (right panel) selective cell populations in the four areas. FEF followed by PFC

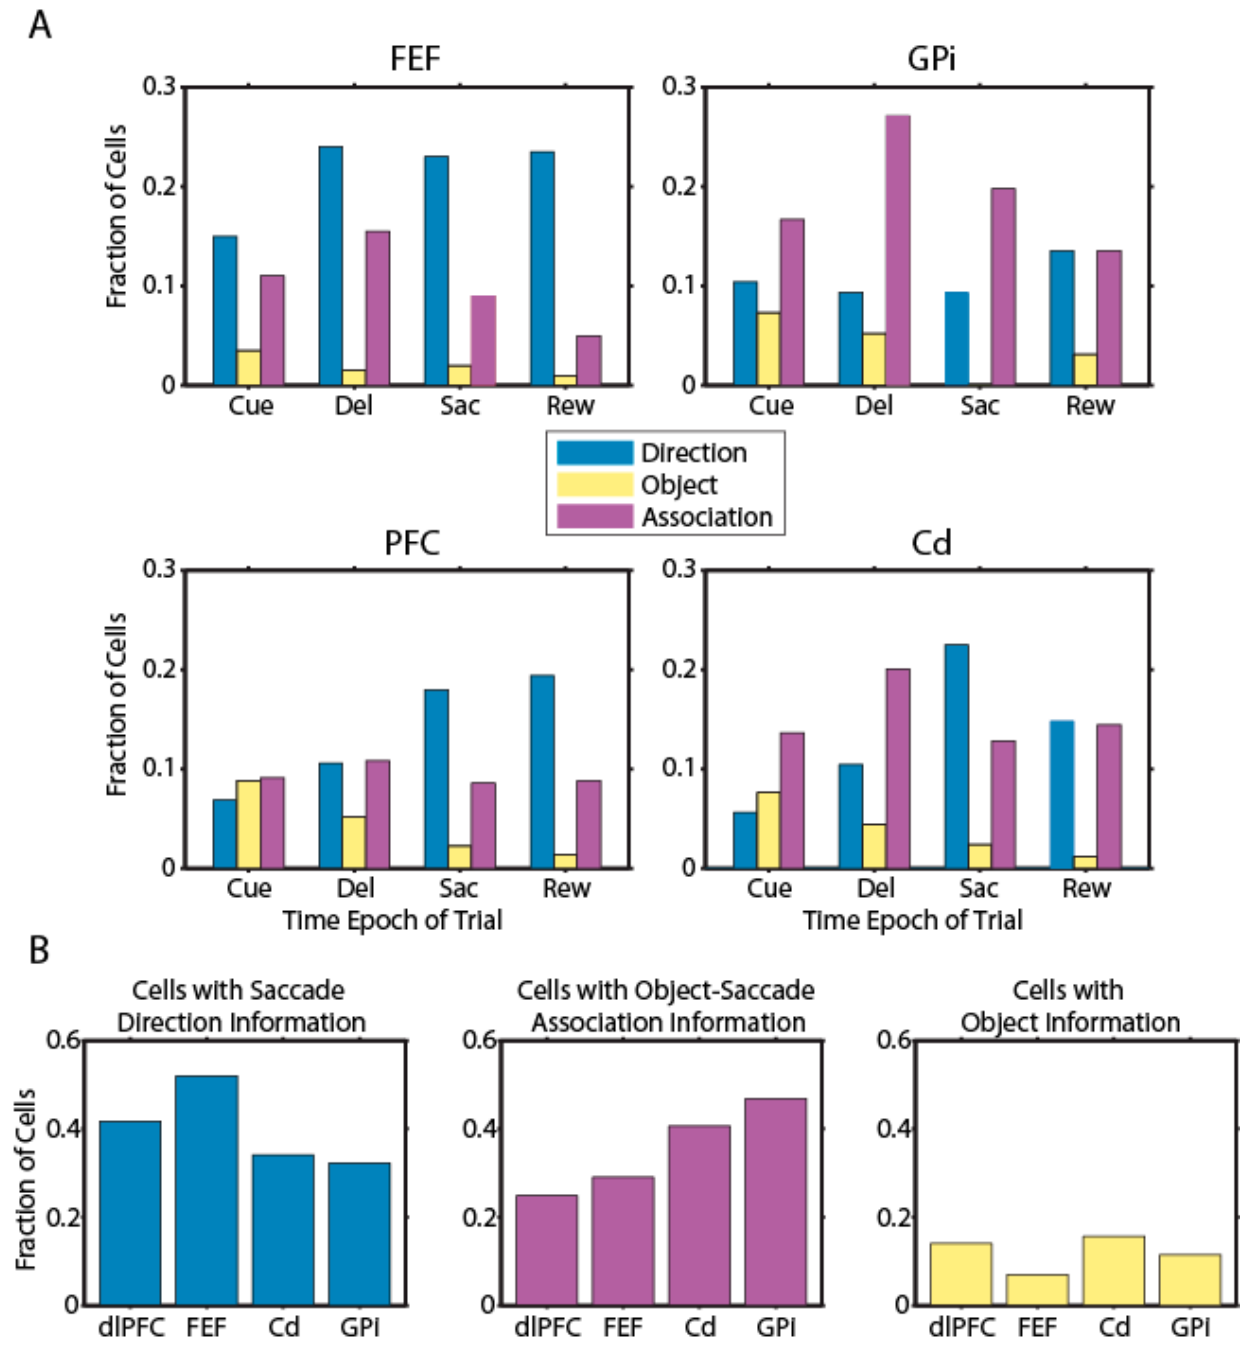

Figure 2.3
Fractions of Cells Containing Task-related Information in FEF, GPi, dlPFC, and Cd
A. The fraction of cells containing saccade direction (blue), object identity (yellow), and object-saccade association (purple) information is plotted by time epoch (Cue, Delay (Del), Saccade (Sac), and Reward (Rew)) in trial for FEF (top left), GPi (top right), dlPFC (bottom left), Cd (bottom right). ***B***. Total fractions of cells in any trial epoch containing direction (left), object-saccade association (middle), and object (right) are compared across brain areas. The cortical areas contain more cells with saccade direction information, the basal ganglia contain more cells with object-saccade association information, and PFC and Cd contain the greatest fraction of cells with object information.

contain the largest fraction of direction selective cells. In contrast, GPi followed by Cd contain the largest fraction of object-saccade association cells. PFC and Cd show the largest object selective population reflecting the larger visual input to these areas.

In order to understand what the firing patterns of cells containing task-related information look like, an example direction selective FEF neuron is shown in Figure 2.4. The rasters in the top panel are provided to show the consistency, and thus the stability of recordings, of the firing patterns of the cell across the entire session. This cell contains data from six learning blocks, as delineated by the horizontal black lines through the rasters. Only data from correct trials is shown in this figure. Information about the direction of the animal's saccadic response is encoded by this cell in the differential firing patterns between trials in which the animal made a leftward saccade (darker colors) and trials in which the animal made a rightward saccade (lighter colors). The left panel of this figure shows the activity of the cell aligned on the onset of object presentation (0ms), whereas the right panel of the figure aligns the end of the trial on saccade onset (0ms). This cell contains saccade direction information during the delay, saccade, and reward epochs of the trial (2-Way ANOVA, $p<0.05$, Bonferroni corrected).

There are many ways in which a cell could be classified as an object-saccade association cell. Thus, to understand precisely what these association cells are encoding, histograms of single cells were analyzed. Two example GPi neurons are provided in Figure 2.5. The top panel depicts a cell classified as association selective during the delay, saccade, and reward trial epochs. This cell seems to treat each association differently, and thus contains significant interactions between object and direction information. The cell depicted in the bottom panel of Figure 2.5 shows a different trend. This cell is classified as an association cell during the cue epoch of the trial, as it exhibited a significant main effect of object and a significant object-direction interaction. It appears to respond at a higher rate during "object A-saccade Right" trials, and it responds the same to the other three associations. Over the population of GPi association cells, there was not a uniform way in which this association information was exhibited.

The cell in the bottom panel of Figure 2.5 is also classified as object selective during the delay period and direction selective during the saccade and reward periods. As a general observation, this type of "multi-tasking" cell is quite prevalent in all the brain areas analyzed.

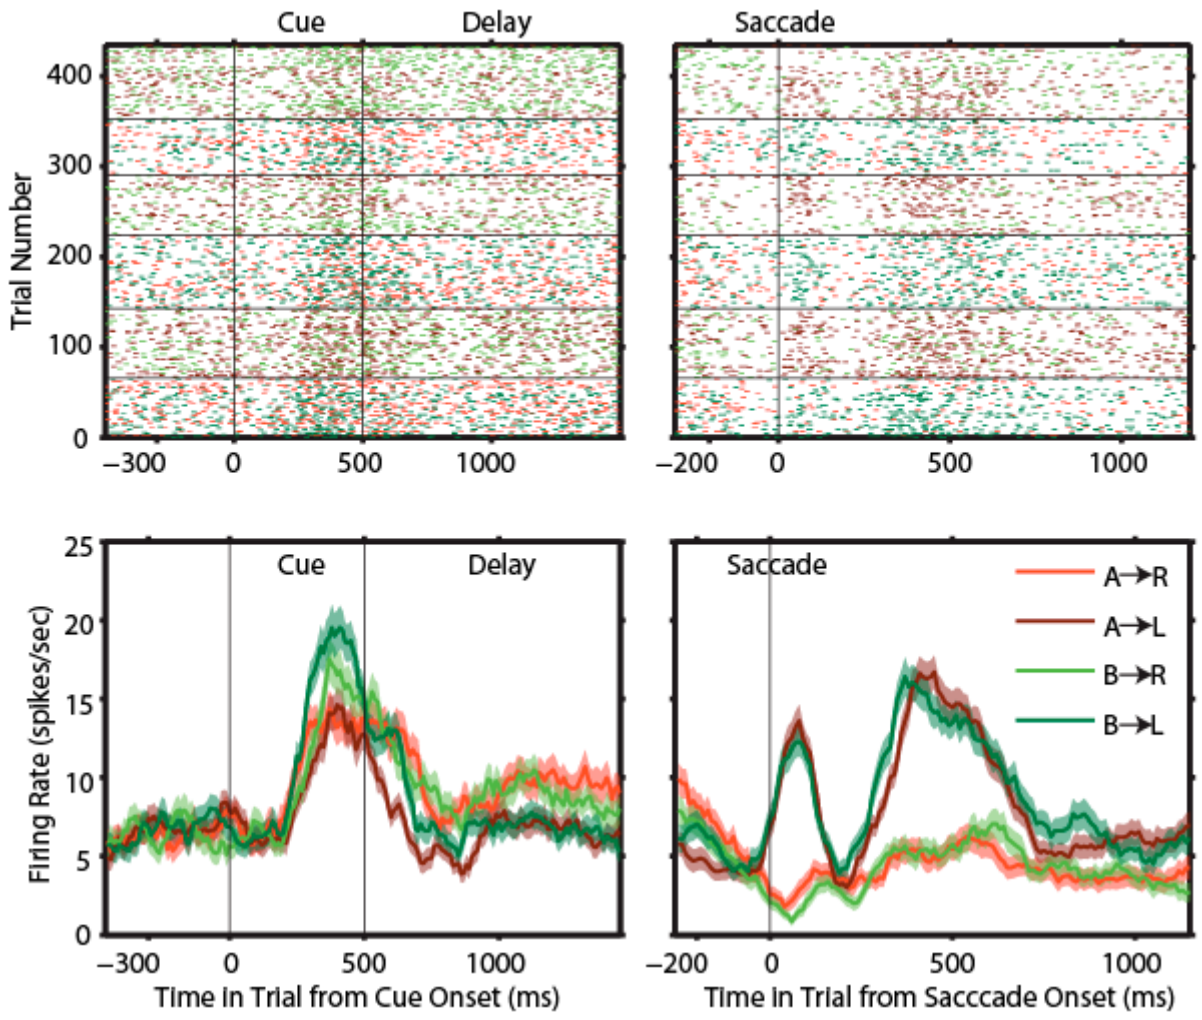

Figure 2.4
FEF Example Cell Containing Saccade Direction Information
This FEF cell displays both a visual response during the presentation of the visual cue and saccade-related activity around the time of saccade and contains information about the direction of the executed saccade during the delay, saccade, and reward periods of the task (2-way ANOVA, $p < 0.05$, Bonferroni corrected). Rasters (top) and histograms (bottom) display activity for all correct trials, and the different colors represent different trial types: trials in which object A was presented (red), trials in which object B was presented (green), trials in which the animal made a rightward saccade (light colors), and trials in which the animal made a leftward saccade (dark colors). Neural activity depicted in the left panel is aligned on the onset of the visual cue, and data in the right panel is aligned on the animal's initiation of the saccadic response. Histograms display mean +/- SE.

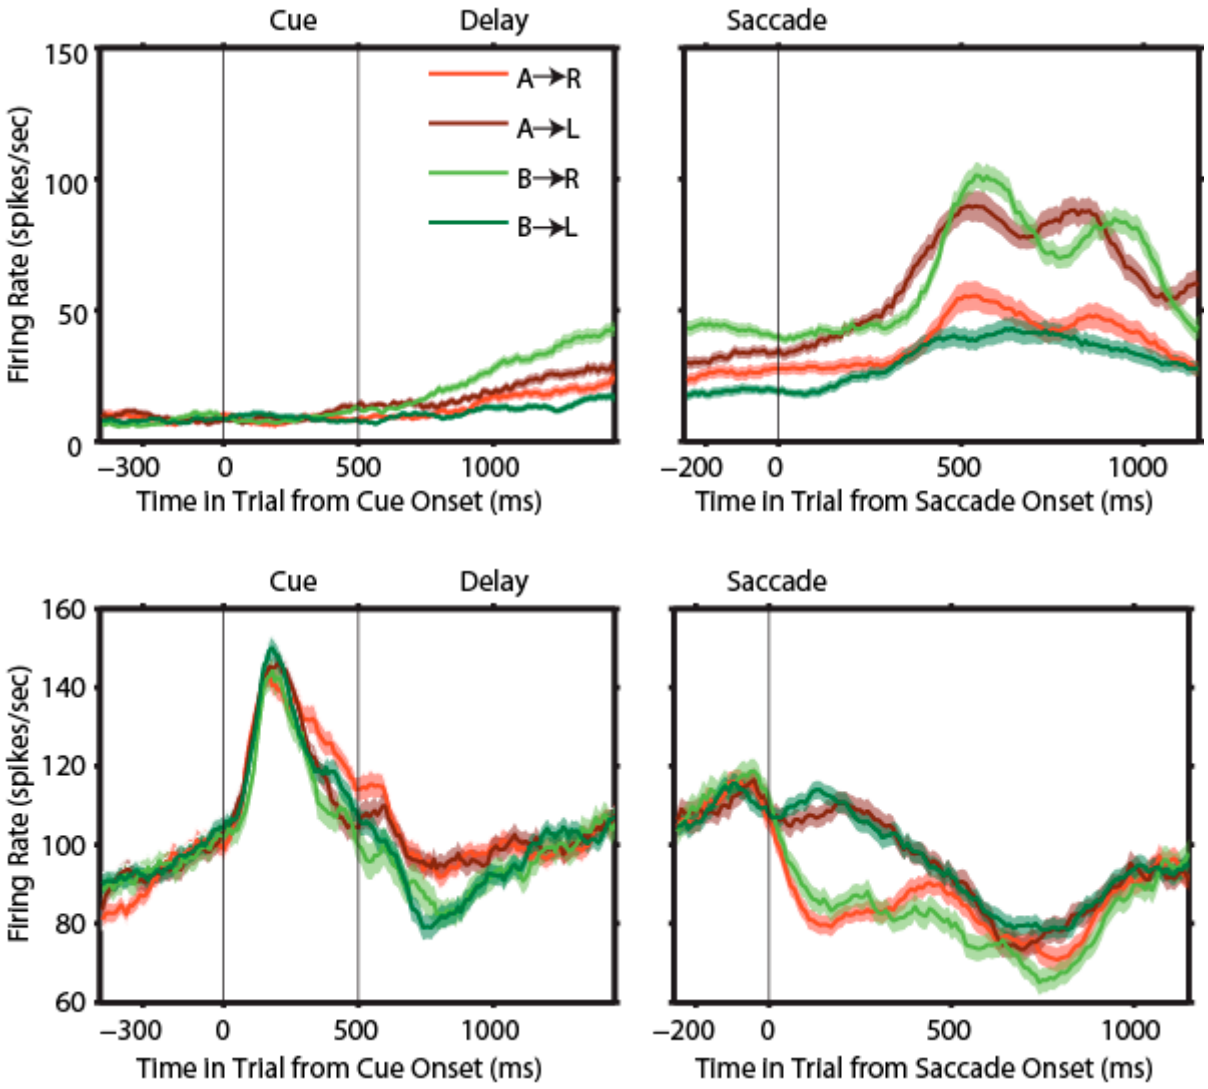

Figure 2.5
Two GPi Example Cells Containing Object-Saccade Association Information
Firing rate histograms of two GPi cells are shown with formatting similar to Figure 2.3. *Top panel*: This cell contains object-saccade association information during the delay, saccade, and reward task periods. *Bottom panel*: This cell contains object-saccade association information during the cue period, object information during the delay period, and saccade direction information during saccade and reward task periods (2-way ANOVA, $p < 0.05$, Bonferroni corrected).

These cells can contain different task-related information at different points in the trial, suggesting that the transfer of information between cells within and between areas is quite complex.

The three cells presented in Figures 2.4 and 2.5 all show increases in firing rate at various times throughout the trial. However, there were cells in the recorded populations that showed inhibitory behavior, or decreases in firing rate from baseline. Figure 2.6 characterizes the population of selective cells in FEF and GPi cells according to these firing rate dynamics. The average firing rate in each trial epoch was compared to the average firing rate during a baseline period from 450ms to 250ms before the onset of the cue image during the fixation period at the start of the trial (ttest, $p<0.05$, Bonferroni corrected). Fractions of cells are based on the total number of selective cells (for either direction, object, or association) for each trial epoch (GPi: cue N=33; delay N=40; saccade N=28; reward N=29. FEF: cue N=59; delay N=82; saccade N=68; reward N=59). Decreases in firing rate of these selective cells were the minority, but in each trial epoch there was a small population that showed this inhibitory behavior.

In order to investigate how the FEF and GPi populations of cells are changing with learning, focus was placed on direction information and those cells classified as direction selective cells in any of the four trial epochs. Direction selectivity was used instead of object-saccade association selectivity for a couple reasons. Firstly, since Pasupathy and Miller (2005) focused on the evolution of saccade direction information with learning and it was my goal to compare the current FEF and GPi data to this previous data, at least first-round analysis needed to be the same. Secondly and more importantly, direction selectivity is consistent from block to block, whereas association selectivity is not necessarily consistent. The association selective cells may only have an altered firing rate for one association, say object A-go Right. Thus, during the blocks where object A is paired with a leftward saccade, the cell may not distinguish between the two relevant associations for that particular block (e.g. cell depicted in Figure 2.5, bottom panel). This specificity of information minimizes the amount of data that can be analyzed for learning effects, and thus makes any statistical analysis much more difficult since statistical power is diminished.

To graphically represent how populations of cells are changing across the trial and across learning, a three-dimensional plot is necessary. Since this type of plot will be presented many

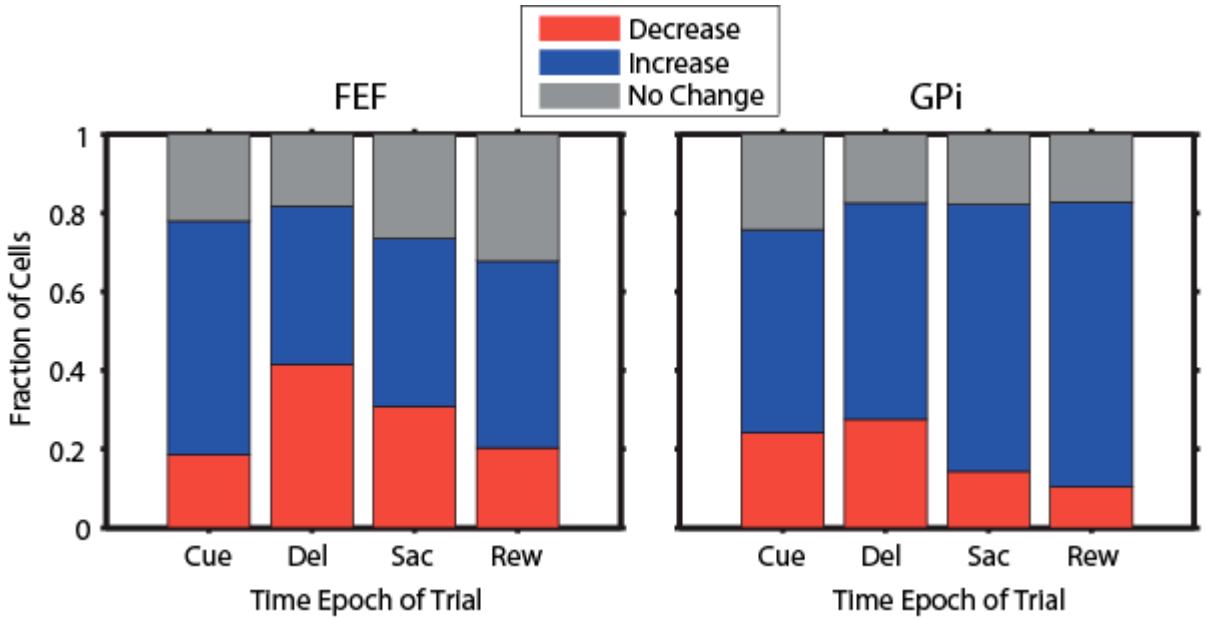

Figure 2.6
Firing Rate Dynamics of FEF and GPi Task-related Cells
Fractions of all FEF (left) and GPi (right) cells containing task-related information that show average increases (blue), decreases (red), or no change (gray) in average firing rate from a baseline period of fixation at the start of the trial. Both areas show more cells with increases versus decreases in firing rate, and the trends across the time epochs of the trial are also consistent between brain areas.

times in this thesis, it is important to understand how to read these figures, so I will take the opportunity here to explain the plots in general. Two dimensions are plotted on the x- and y-axes, and the third dimension is represented by color. The x-axis in these plots is always time in trial, and depending on the particular plot, time in trial may either be aligned on visual cue onset at the start of the trial or saccade onset at the end of the trial. Trial in learning is plotted on the y-axis, and all of these colorplots need to be read from the bottom-up, as the first trial in learning is plotted at the bottom of the figure. Color may be used to represent different dimensions of the data. The two main variables I will use for the color axis are normalized firing rate and direction selectivity (fraction explained variance by the direction factor in an ANOVA, FEV_{dir}). It is important to keep in mind three key points: 1) normalized firing rate is an average *of all correct trials*, 2) direction selectivity is a measure of the differences in firing rate *between trials* when the animal made a correct rightward and leftward saccade, and 3) all colorplots depict aspects of a *population* of cells.

The top panel of Figure 2.7 presents the direction selectivity of the FEF population of direction selective cells (N=104) across learning. The left panel shows data from the entire trial and well into the inter-trial interval. This plot shows two main bands of direction information: one around the time of visual cue presentation and the other around the time of saccade execution. Zooming in on the beginning of the trial, the middle panel shows that saccade direction selectivity is present from the very start of learning. The black dots on this figure show the risetimes for this direction selectivity: the time in each trial where half-maximum selectivity was reached. These risetimes can be used as a measure of the timing of information, and from them it appears that the timing of saccade direction information does not change as learning progresses. The rightward most panel shows the band of direction information around the time of saccade execution, and again it appears that the timing of this information does not change with learning.

In addition to direction selectivity, Figure 2.7 shows averaged normalized firing rate for the same population of FEF cells. Two bands of activity, at the start and end of the trial, are also present in this plot. This data is presented to make the point that the saccade direction information does not necessarily reflect the overall average firing rate, as direction selectivity is a measure of the difference in activity on right versus left trials. Two main differences to point out between the

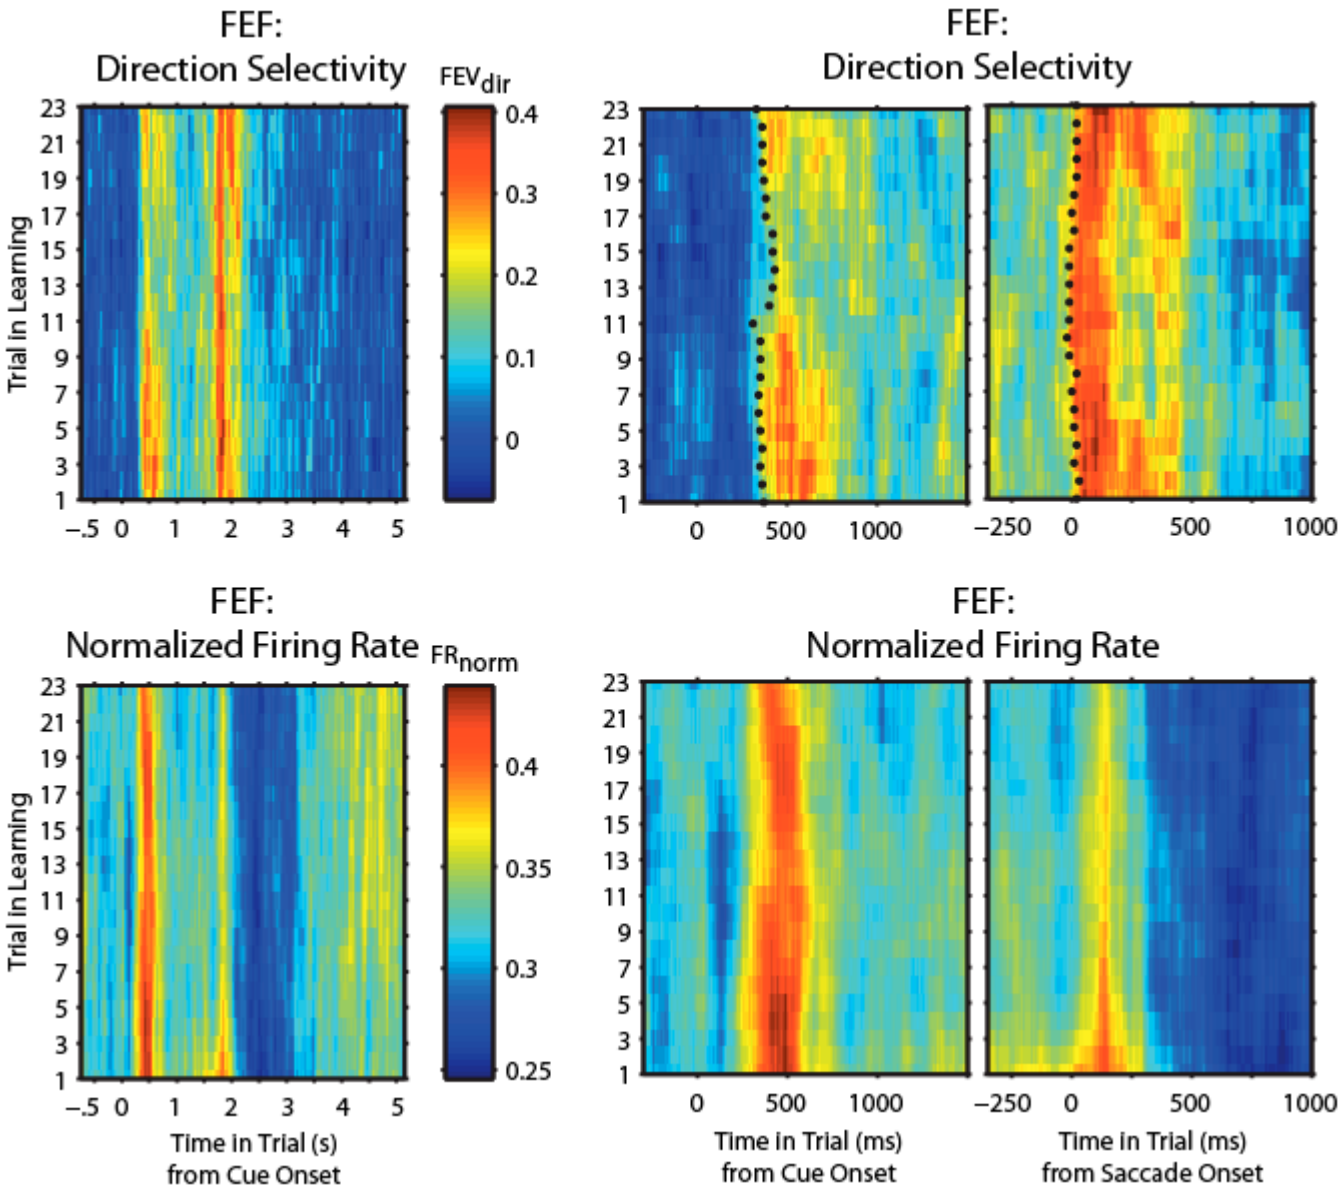

Figure 2.7 Changes in Saccade Direction Selectivity with Learning in FEF

Learning-related changes in information (FEV_{dir}) about saccade direction (top) and in normalized firing rate (bottom) for the entire saccade direction selective population of FEF cells (N = 104) is shown. Data from the entire trial (left) shows two main bands of increased information/activity: one around the time of visual cue presentation (middle) and one around the time of saccade execution (right). Risetimes, the time to half-maximum direction selectivity, are shown as block dots. There are two main learning-related changes apparent: information about saccade direction during the cue period decreases in strength as learning progresses, and firing rate around the time of saccade execution decreases as learning progresses.

firing rate and direction selectivity of this population of cells: 1) while the firing rate does not change with learning around the time of visual cue presentation, the strength of direction selectivity decreases as learning progresses, and 2) while the strength and timing of direction selectivity around the time of saccade execution does not change with learning, the firing rate of this population decreases as learning progresses.

The same analysis was performed on the direction-selective population of GPi cells (N=31), and the results are shown in Figure 2.8. Like FEF, GPi exhibits two bands of direction information during the trial: one around the time of cue presentation and one around the time of saccade execution. However, in GPi, the band of information early in the trial is much weaker than the information around the time of saccade execution. The increases in firing rate of this population of GPi cells (Figure 2.8, bottom) also occur at these time periods in the trial; however, the changes in firing rate do not differ between the different time periods. There is also a third band of firing rate increase during the inter-trial interval, which is not seen in the other brain areas investigated.

Two linear measures of this saccade direction information were used to compare the learning-related changes in information across learning in the different brain areas: risetime, a measure of the timing of information, and peak information, a measure of the strength, or amount, of information. These two variables are plotted across learning in the four brain areas investigated during the early-trial period of the reversal learning paradigm in Figure 2.9. Risetimes in FEF and GPi (Figure 2.9A) are consistent across learning; there is information from the start of learning in both areas and the timing of this information remains constant as learning progresses. This consistency sharply contrasts with the saccade direction information observed in dlPFC and Cd: both PFC and Cd show drastic learning-related changes in the timing of information (Figure 2.9B). Another difference between these brain areas is observed in the learning-related changes in strength of saccade direction information: strength of information in FEF decreases with learning (Figure 2.9C), whereas strength of information in PFC and Cd increases with learning (Figure 2.9D).

But what is the role for this early-trial saccade direction information in FEF and GPi that is present from the start of learning? Since these are primarily motor structures, it is possible that this information is strictly motor-related, simply reflecting the planning of the animal's saccadic

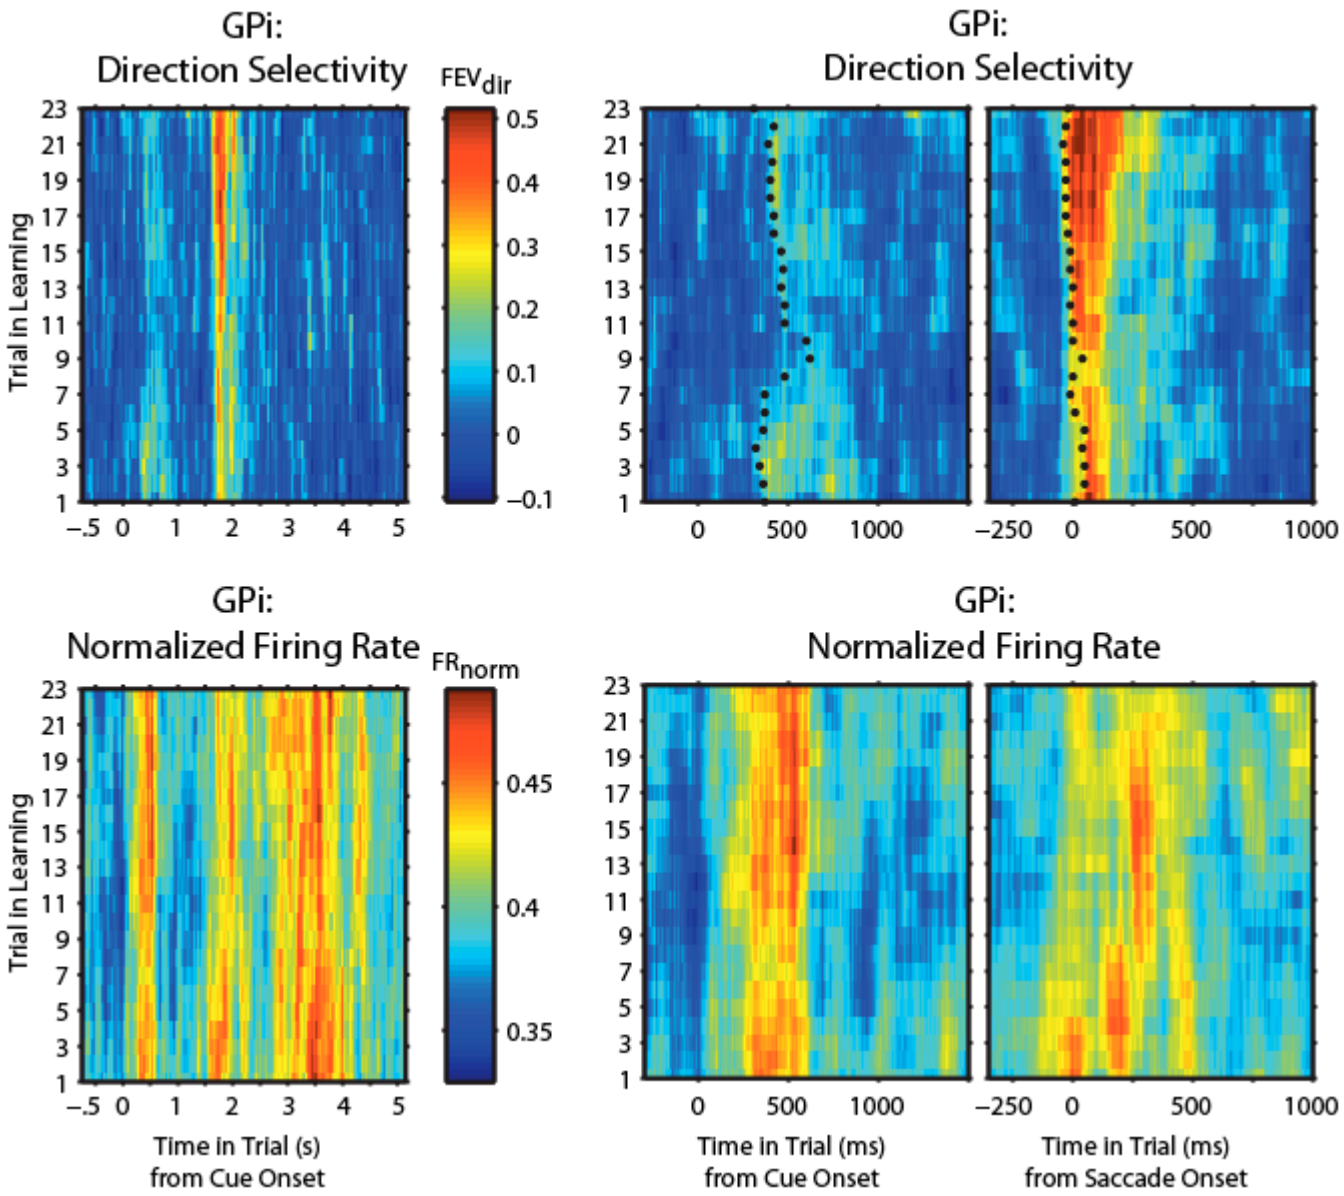

Figure 2.8 Changes in Saccade Direction Selectivity with Learning in GPi

Learning-related changes in information (FEV_{dir}) about saccade direction (top) and in normalized firing rate (bottom) for the entire saccade direction selective population of GPi cells (N = 31) is shown by these colorplots. Data from the entire trial (left) shows two main bands of increased direction information: one around the time of visual cue presentation (middle) and one around the time of saccade execution (right), although the latter band is much stronger than the former. Risetimes, the time to half-maximum direction selectivity, are shown as block dots. Strength of direction selectivity around the time of saccade execution increases as learning progresses, and no learning-related changes in firing rate are seen.

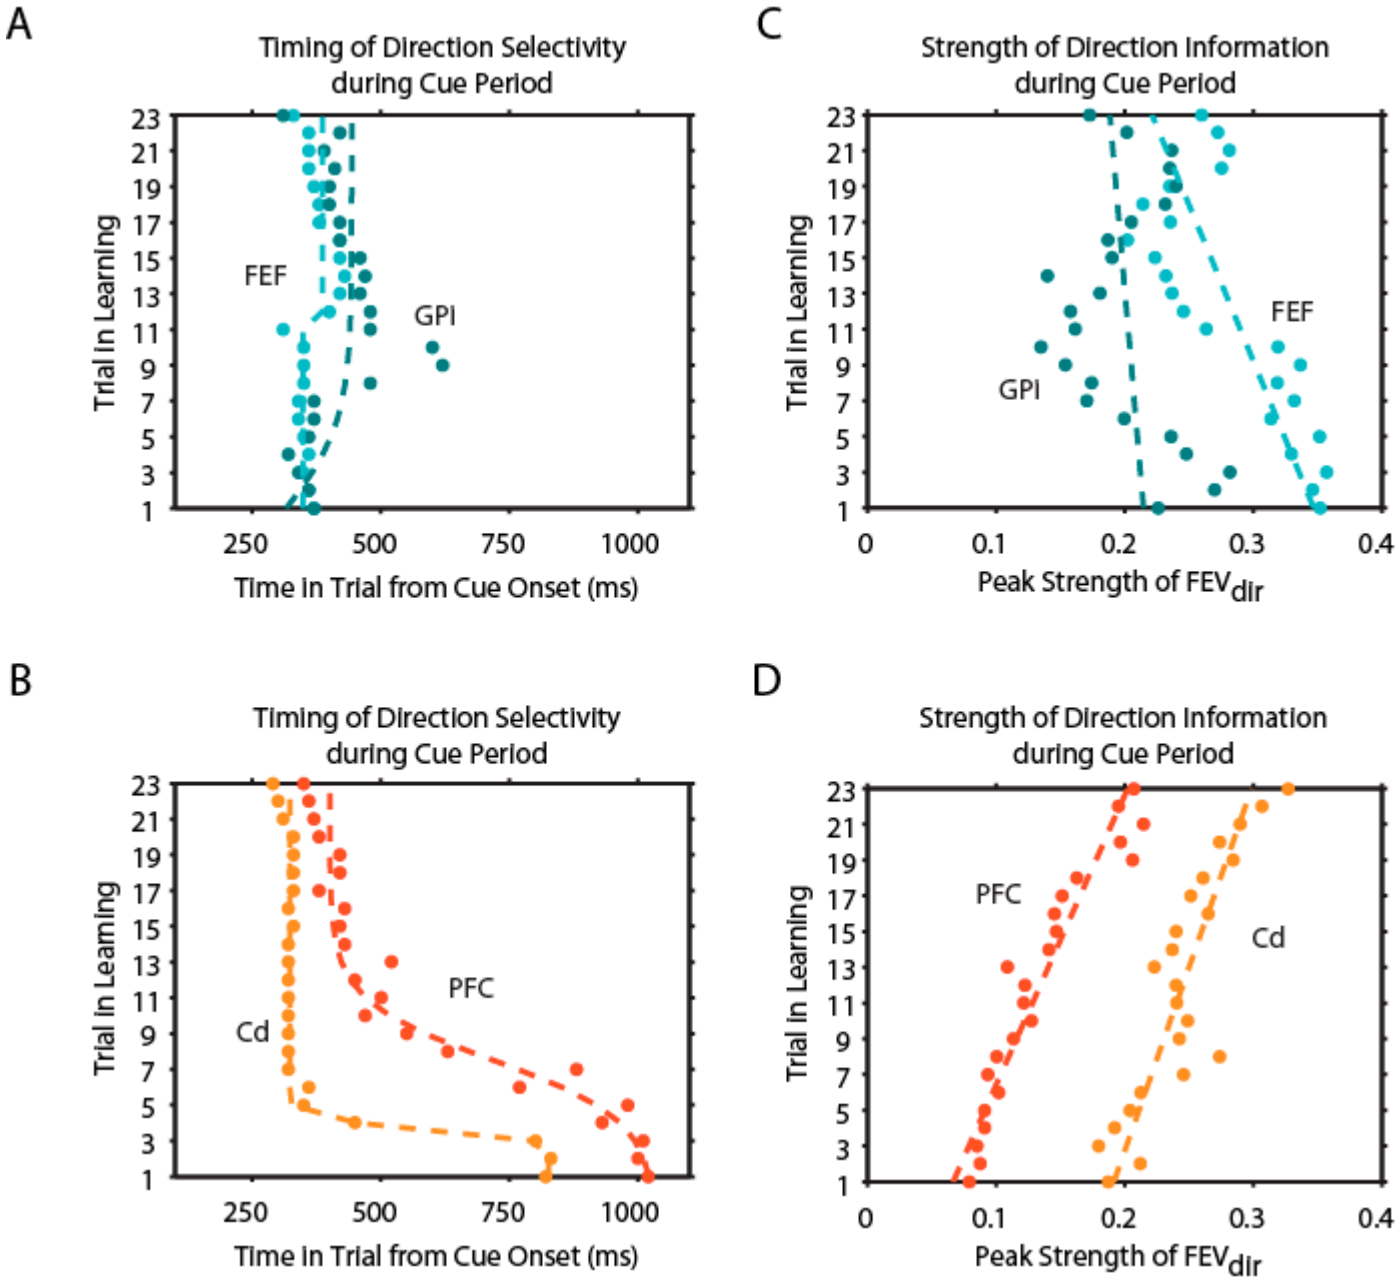

Figure 2.9
Changes in Early-Trial Direction Selectivity with Learning in FEF, GPi, dlPFC and Cd
The timing (left) and strength (right) of information about saccade direction during the cue period at the start of the trial is compared between FEF (top), GPi (top), Cd (bottom), and PFC (bottom). ***A***. Saccade direction information is present from the start of learning in FEF and GPi and the time course of this information doesn't change with learning. ***B***. Saccade direction information in PFC and Cd appears progressively earlier in the trial with learning. This change in timing happens earlier in learning and more abruptly in Cd than in PFC. ***C***. Strength of saccade direction information decreases as learning progresses in FEF. ***D***. Strength of saccade direction information increases as learning progress in PFC and Cd.

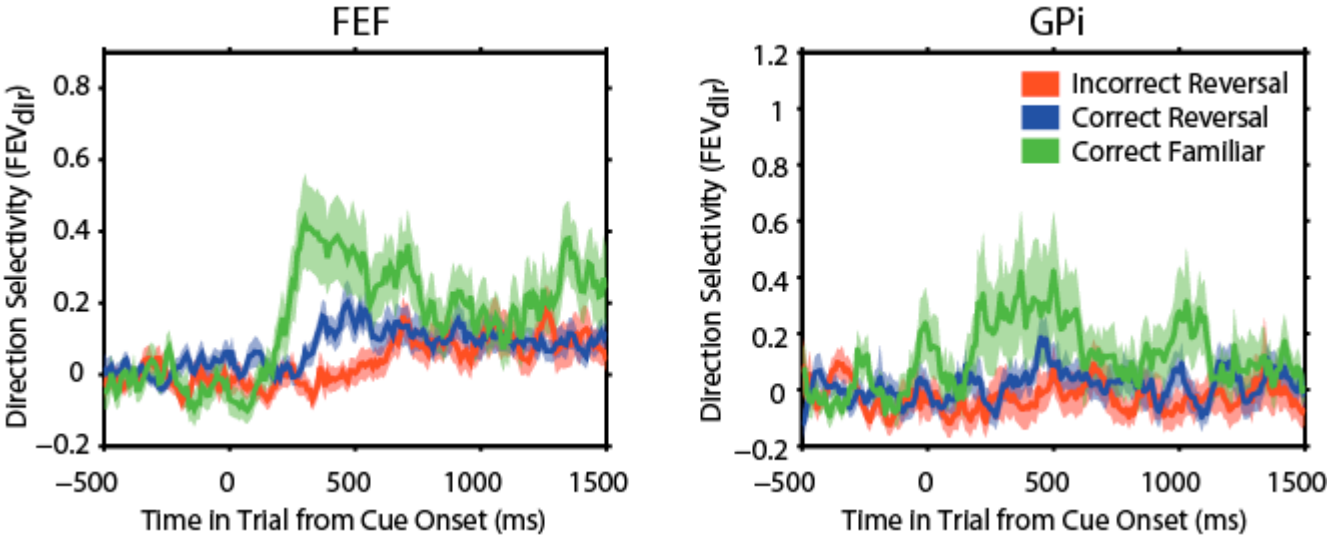

Figure 2.10
The Relationship Between Early-Trial Saccade Direction Selectivity and Behavior
Average saccade direction selectivity in FEF (left) and GPi (right) for incorrect trials on reversing associations (red), correct trials on reversing associations (blue), and correct trials on familiar associations (green) is shown across time in trial.

FEF	Correct Reversal	Incorrect Reversal	Zero
Correct Familiar	P = 0.019	P = 0.0011	P = 0.0011
Correct Reversal		P = 0.0055	P = 0.00025
Incorrect Reversal			P = 0.91

GPi	Correct Reversal	Incorrect Reversal	Zero
Correct Familiar	P = 0.3651	P = 0.1420	P = 0.1316
Correct Reversal		P = 0.3115	P = 0.2046
Incorrect Reversal			P = 0.9974

Table 2.1
P-values for Comparisons of Early-Trial Direction Selectivity Between Trial Types
Paired ttests were done on the average firing rate in FEF (top) and GPi (bottom) during the cue period for the different trial types depicted in Figure 2.9: correct trials of familiar associations, correct trials of reversal associations, and incorrect trials of reversal associations. Average firing rate was calculated from 100-500ms after cue onset. *Top*. FEF exhibits significant information during the cue period (compared with zero, right column) in correct trials of both the familiar and reversal associations, but not for incorrect reversal associations. All pair-wise comparisons between the trial types were significant. *Bottom*. GPi contains no significant information during the cue period (compared with zero, right column) in any trial type, nor are there any significant differences between any trial types.

eye-movement. If this were true, the cell population would express this information on every trial where the animal makes a saccade. To investigate this claim, the saccade direction information on different trial types was compared (Figure 2.10 and Table 2.1). Correct trials on the reversal associations were compared to both incorrect trials of reversal associations as well as correctly performed trials of familiar associations. To perform a fair comparison between the trials, an equal number of each trial type were used: all incorrect trials were used (as this category contains the fewest trials), and correct familiar and reversal trials were chosen to be as close in time to each incorrect trial to minimize effects of the placement of trials within the block. Direction selectivity was computed across the chosen trials and averaged over all direction cells. During the time of cue presentation (100-500ms after cue onset), FEF contains significant saccade direction information in correct familiar and reversal trials, and this information is greatest for familiar trials (Figure 2.10 and Table 2.1). In contrast, GPi contains no significant direction information on any trial type during this early-trial period.

End of trial effects, around the time of saccade execution, show a different trend. FEF shows no changes with learning in either the timing (Figure 2.11A) or strength (Figure 2.11C) of saccade direction information. However, saccade direction information around the time of saccade in GPi appears progressively earlier in the trial (Figure 2.11A) and increases in strength (Figure 2.11C) with learning. These effects observed in GPi parallel those seen in Cd and PFC (Figure 2.11B,D). When comparing the late-trial (0-400ms after saccade initiation) direction selectivity in the different trial types (Figure 2.12 and Table 2.2), FEF does contain significant saccade direction information on all trial types, significantly more information on familiar associations than correct or incorrect reversal associations, and the same amount of information on incorrect and correct reversal associations. GPi only shows significant direction information in correct trials of familiar associations, but, like FEF, this information is greater than the selectivity in either the correct or incorrect trials of reversal associations.

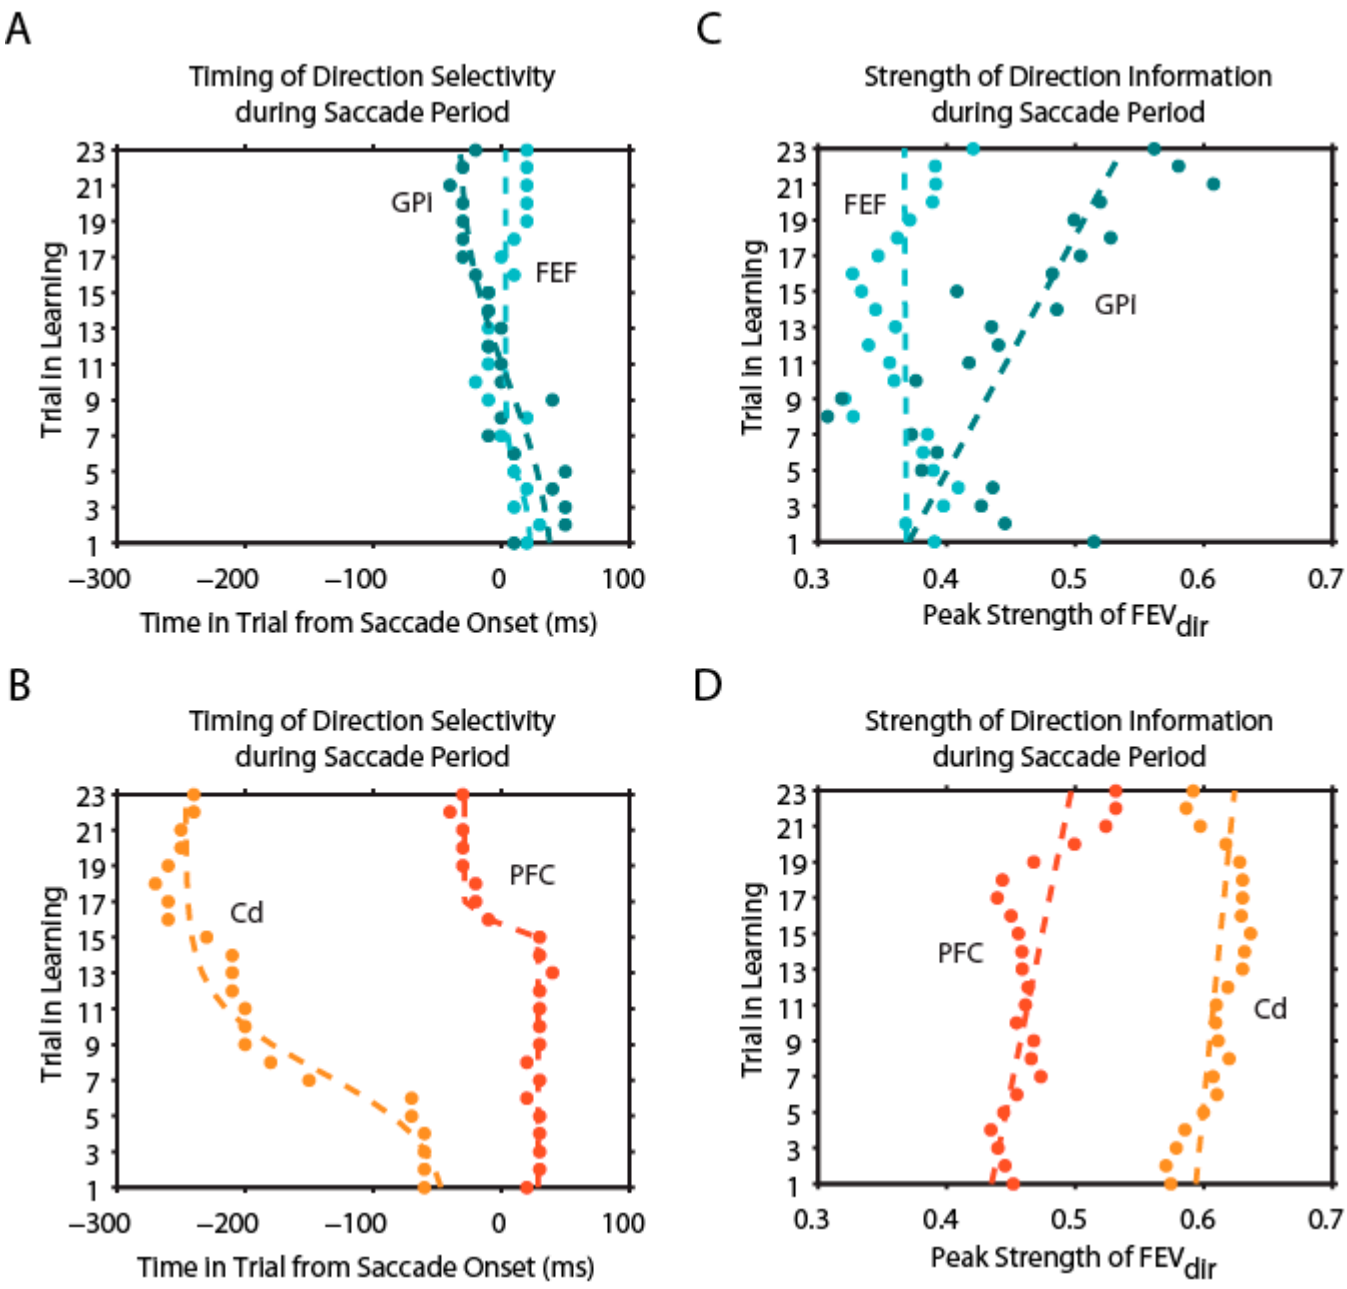

Figure 2.11
Changes in Late-Trial Direction Selectivity with Learning in FEF, GPi, dlPFC and Cd
The timing (left) and strength (right) of information about saccade direction during the saccade period at the start of the trial is compared between FEF (top), GPi (top), Cd (bottom), and PFC (bottom). *A*. Saccade direction information becomes progressively earlier in GPi as learning progresses, but does not change in FEF. *B*. Saccade direction information in PFC and Cd appears progressively earlier in the trial with learning. This change in timing happens earlier in learning in Cd than in PFC. *C*. Strength of saccade direction information increases as learning progresses in GPi. *D*. Strength of saccade direction information increases as learning progress in PFC and Cd.

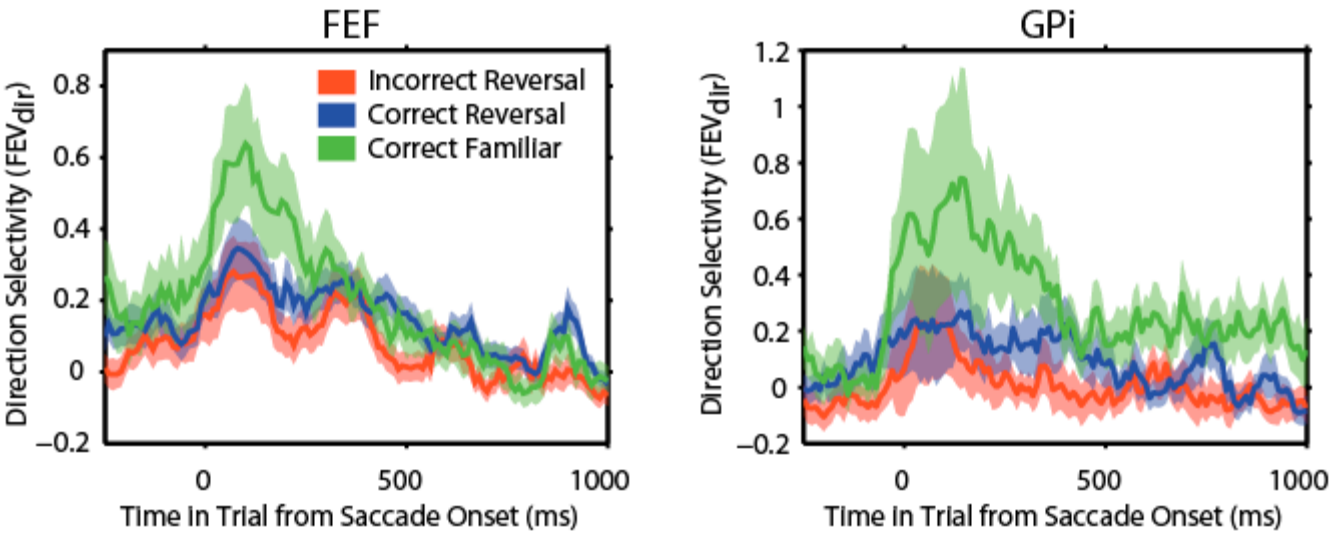

Figure 2.12
The Relationship Between Late-Trial Saccade Direction Selectivity and Behavior
Average saccade direction selectivity in FEF (left) and GPi (right) for incorrect trials on reversing associations (red), correct trials on reversing associations (blue), and correct trials on familiar associations (green) is shown across time in trial.

FEF	Correct Reversal	Incorrect Reversal	Zero
Correct Familiar	P = 0.0016	P = 0.00034	P = 8.7×10^{-7}
Correct Reversal		P = 0.1909	P = 7.6×10^{-6}
Incorrect Reversal			P = 0.00014

GPi	Correct Reversal	Incorrect Reversal	Zero
Correct Familiar	P = 0.0047	P = 0.0059	P = 0.0049
Correct Reversal		P = 0.2847	P = 0.1098
Incorrect Reversal			P = 0.2822

Table 2.2
P-values for Comparisons of Late-Trial Direction Selectivity Between Trial Types
Paired ttests were done on the average firing rate in FEF (top) and GPi (bottom) during the saccade period for the different trial types depicted in Figure 2.11: correct trials of familiar associations, correct trials of reversal associations, and incorrect trials of reversal associations. Average firing rate was calculated from 0-400ms after saccade onset. *Top*. FEF exhibits significant direction information in all trial types during the saccade period, and all comparisons between trials are significant except for the difference between correct and error trials of reversal associations. *Bottom*. GPi contains significant direction information only in correct trials of familiar associations during the saccade period, and significantly more information in familiar trials versus correct and incorrect trials of reversal associations.

Discussion

This study of the activity of FEF and GPi during the acquisition of arbitrary visuomotor associations uncovered two main results. First, the task-related information contained in these structures is different: the FEF population contains more information about the direction of the behavioral response, and GPi contains more information regarding the specific visuomotor association. Second, when investigating the motor response-related information, FEF exhibits this information during the early-trial period from the start of learning and shows a learning-related decrease in the strength, or amount, of information. On the other hand, most of the motor-related information in GPi occurs around the time of the behavioral response and this information increases in strength and appears earlier in the trial as learning progresses.

These main results are graphically represented in Figure 2.13. Two cortico-basal ganglia loops are shown: the oculomotor loop which includes FEF, and the dlPFC loop (Alexander et al., 1986). While these loops do operate in parallel, overlap does exist, as dlPFC is known to contribute to the oculomotor loop as well. The finding from the study presented here that FEF contains response-related information from the start of learning suggests that FEF may be contributing to the accumulation of information in Cd and/or PFC (red arrows) either directly through FEF's projections to these areas, or indirectly through potential crosstalk, or convergence, of information in the BG. Also, the learning-related decrease in the strength of information contained in FEF coupled with the learning-related increases in the strength of information in both PFC and Cd provide further evidence for this information transfer from FEF to these structures.

The grayscale shading of the areas connected in these loops shown in Figure 2.13 represents the processing of information occurring in the frontal cortex-basal ganglia system. Black and white portions represent populations of cells containing different types of information (e.g. visual object identity and saccade direction). While these populations are mainly separate in the cortex, the convergence of information in BG creates increased overlap in this information, ultimately resulting in GPi (outlined in red) containing a majority of cells encoding specific object-saccade associations as was presented here.

[Note: In this figure the populations of cells carrying different information are depicted as arising from the same cortical structures. That does not necessarily need to be so, and actually most likely is not so; however, this representation graphically simplifies the figure to emphasize the main results found in this study.]

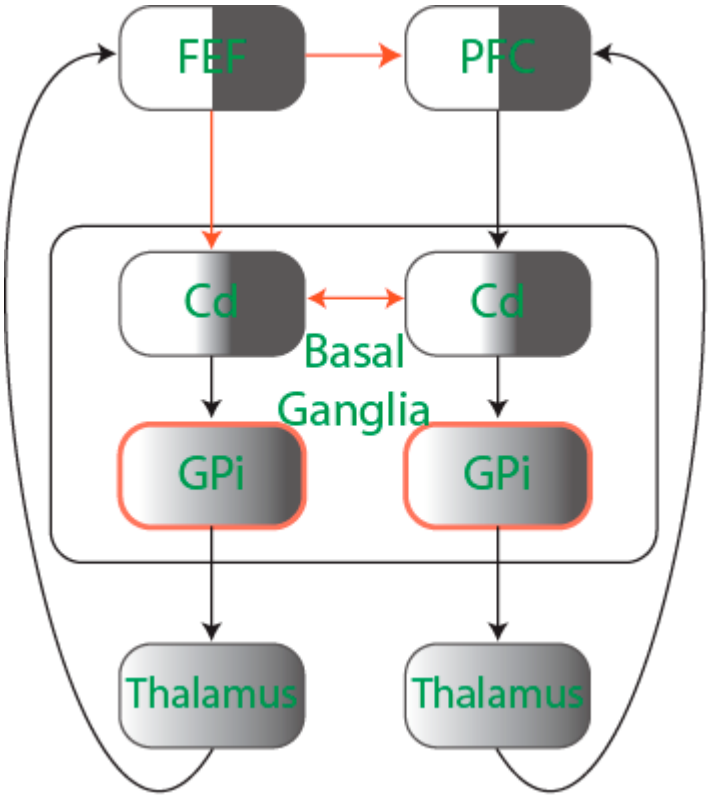

Figure 2.13
Summary and Interpretation of FEF and GPi during Learning With Reversals: A Model
Evidence presented in this chapter suggests that FEF may provide Cd and PFC with saccade direction information directly (red arrows from FEF) or indirectly, through interactions between cells in the basal ganglia (red arrow, middle), to aid the development of this information in PFC which may then be used to guide behavior. GPi was shown to contain many cells with specific object-saccade association information, supporting the information "funneling" idea posed by Alexander et al. (1986) and depicted here with grayscale coloring. Black and white portions in cortical areas are different cell populations containing different types of information (e.g. cell populations encoding saccade direction and object identity). These populations converge slightly in Cd and converge even further in GPi resulting in cells with increasing specificity of information encoding (i.e. association cells).

Chapter 3: Learning-Related Changes in PFC and Cd During Learning Without Reversals

Introduction

The PFC and Cd are two key components of the frontal cortex-basal ganglia system that is thought to mediate arbitrary visuomotor association learning, forming links between arbitrary visual stimuli and motor responses. Previously, learning-related changes in these brain areas were shown to exhibit differential time courses of the expression of task-related information: Cd contained response-specific informatoin early in learning and changes in this information occurred early in the learning process, in contrast to the slow and late changes that occurred in PFC with learning (Pasupathy & Miller, 2005). Crucially, the behavioral performance of the animal more closely paralleled the time course of changes exhibited by PFC, suggesting that PFC is more closely linked with behavioral control in arbitrary visuomotor learning. These differences were interpreted to suggest a transfer of information from BG to PFC; BG learns an association first, trains the PFC and the PFC then guides behavior.

The goal of the present experiment is to investigate the context-dependency of the function of the frontal cortex-basal ganglia system during learning. Perhaps it is only under certain behavioral circumstances that the frontal cortex-basal ganglia system operates in the manner described by Pasupathy & Miller (2005). The previous behavioral task required the animals to learn new associations (e.g. object A, look Right) and then subsequently, and repeatedly, reverse the associations (e.g. object A, look Left). The current study altered the behavioral paradigm by eliminating reversals to determine the effects on the function of the frontal cortex-basal ganglia system.

Methods

[for more methodolgical details see Chapter 6]

Subjects

Two Rhesus macaque monkeys (*Macaca mulatta*), Monkey P and Monkey A, were used in these experiments. All animal procedures conformed to the NIH guidelines and were approved by the MIT Committee on Animal Care. It should be noted that data from the same animals was also used in Pasupathy & Miller (2005), thus providing justification for direct comparisons between the two studies.

Behavioral Task

The animal performed a visuomotor learning task without any reversals of the object-saccade associations (Figure 3.1). In this task, the animal had to associate a visual stimulus with a particular saccadic response direction (e.g. see picture A, saccade right). Once the animal learned the correct associations by trial and error, the objects were discarded and replaced by new objects and the animal had to learn the new object-saccade association (e.g. see picture B, saccade left). This pattern continued throughout the experimental session.

Electrophysiological Data Collection

Up to 16 Tungsten microelectrodes were acutely implanted into dlPFC and Cd in a single session (31 sessions for Monkey A, 5 sessions for Monkey P). PFC data was collected from both animals, and Cd data was collected only from Monkey A. In Monkey A, PFC and Cd much of the data was collected simultaneously, but not all. Each single unit included in analysis contained at least four full behavioral learning blocks of data. A total of 196 PFC cells (99 from Monkey A and 97 from Monkey P) and 52 Cd cells (all from Monkey A) were used in analysis.

Analytical Techniques

Analysis of the behavioral data focused on accuracy and reaction time across learning. All recording sessions (Total: N=36) were used in this analysis (Figure 3.2). All correct and incorrect trials were used in the order in which they occurred in each learning block (trials in which the animal made a fixation error were ignored). Means +/- SE were computed across all blocks from all sessions (187 total blocks: Monkey A, N=156; Monkey P, N=31). For the behavioral analysis shown in Figure 3.4, reaction time and accuracy were computed using a 5-trial window centered on each correct trial in learning. Data shown is mean across all blocks in the task with reversals and without reversals.

In order to identify cells that contained task-related information, a 1-way ANOVA (with saccade direction as the factor) was performed on average firing rate in each of 4 task epochs: 100-600ms after cue onset (“Cue”), 600-1500ms after cue onset (“Delay”), 150ms before to 150ms after saccade initiation (“Saccade”), and 50-300ms after the start of reward delivery (“Reward”). A cell was considered to contain saccade direction information if it had a significant main effect of direction. Significance level was set at $p < 0.05$, and was corrected for multiple comparisons

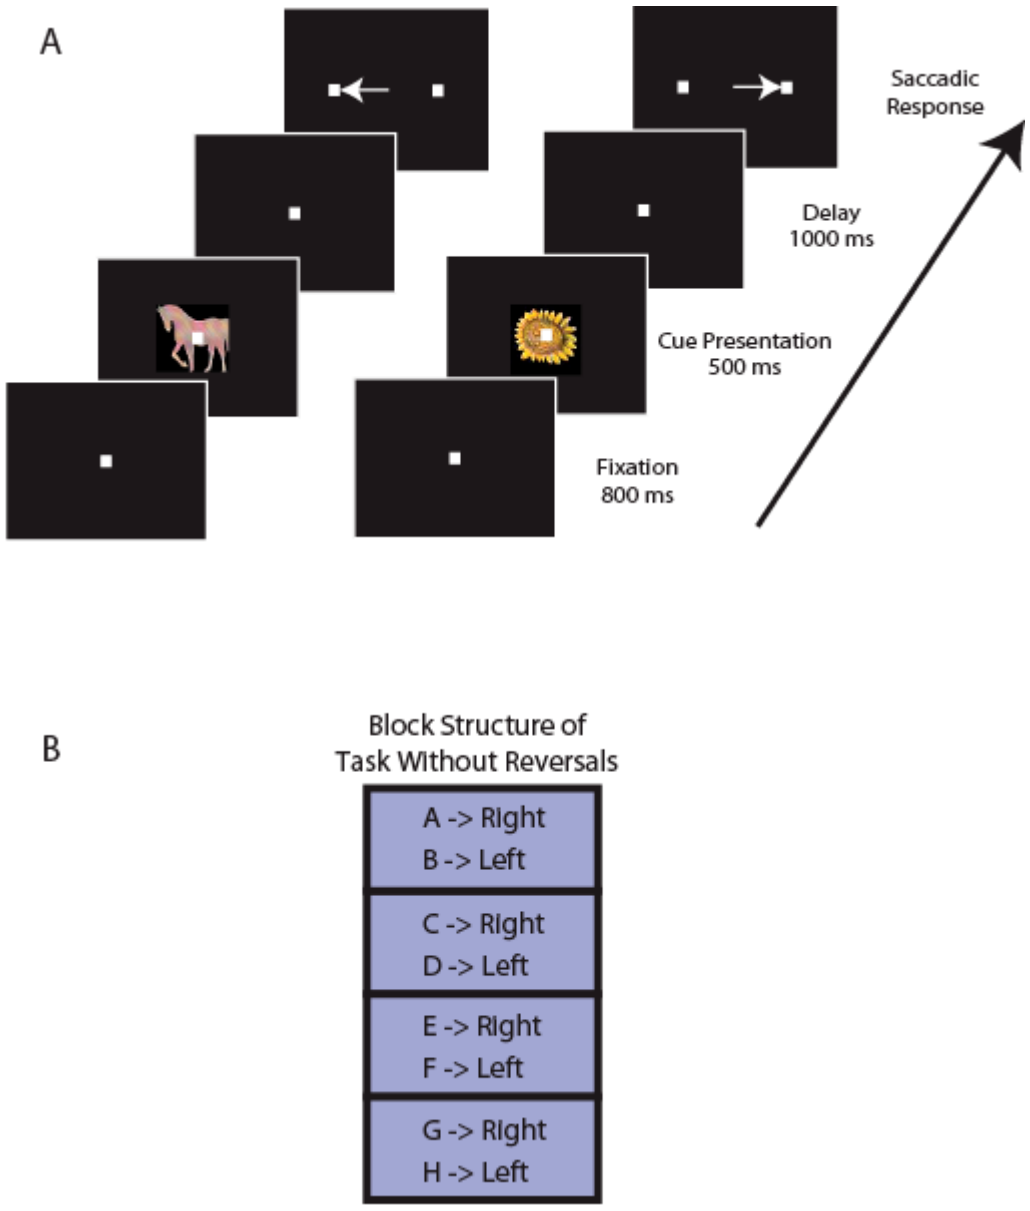

Figure 3.1
Trial Events and Block Structure of the Learning Task Without Reversals
The behavioral task used in this experiment requires the animal to associate a particular saccadic response with a particular visual stimulus. ***A***. The animal is shown a visual stimulus for 500ms, and after a one-second delay period, is required to make a saccade to one of two peripherally presented targets. Each visual stimulus is associated with a particular saccade direction. Execution of the correct saccade results in the delivery of a juice reward. ***B***. The block structure for the task required the animals to learn new associations (blue) in every block. Once the animal learned the associations in a block, the images were discarded.

using Bonferroni's correction. A bar graph showing the fraction of cells in PFC and Cd containing saccade direction information is presented in Figure 3.3

Normalized firing rate, direction selectivity (fraction of explained variance by the direction factor), and risetimes were computed in the same manner as previously described (see Chapters 2 and 6). The color plots are averages across all cells containing saccade direction information (PFC: N=102 (73 from Monkey A, 29 from Monkey P); Cd: N=22).

Results

Unlike the behavioral task with serial reversals (Chapter 2), animals begin this task with no history with the visual images or particular object-saccade associations. Thus, at the start of learning behavioral accuracy is at chance performance (50% correct), and reaction time is at a maximum (Figure 3.2). As learning progresses, accuracy increases to criterion level and reaction time decreases.

To determine which cells of the recorded population (PFC: N=196, Cd: N=52) contained task-related information, a 1-way ANOVA (with direction as the factor; $p<0.05$, Bonferroni corrected) was performed on the firing rate of each cell at each of four time epochs of the trial: 100-600ms after cue onset ("Cue"), 600-1500ms after cue onset ("Delay"), 150ms before to 150ms after saccade initiation ("Saccade"), and 50ms-300ms after the start of reward delivery ("Reward"). Results of this analysis are presented in Figure 3.3. Approximately equal proportions of neurons displayed information related to the direction of the animal's saccadic response in PFC (green bars, total N=102) and Cd (yellow bars, total N=22).

The saccade direction information contained in the population of PFC direction-selective cells (N=102) changes during the early-trial period as learning progresses (Figure 3.4, top). This information appears earlier in the trial as learning progresses, but this is not a gradual change. Very early in learning this change occurs abruptly and continues at asymptote for the rest of learning. In contrast, saccade direction information around the time of saccade execution does not change with learning. In addition, no learning-related changes are observed in the normalized firing rate of this population of PFC cells at any point in the trial (Figure 3.4, bottom).

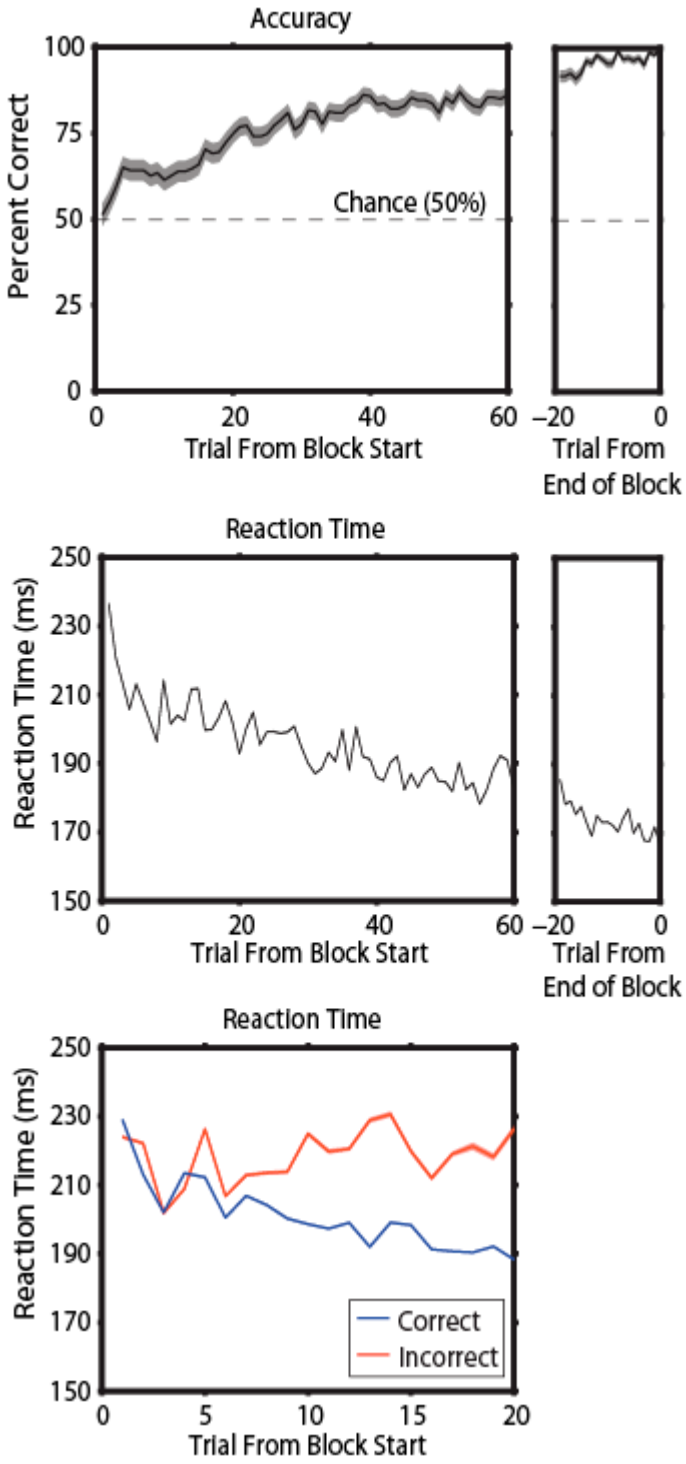

Figure 3.2
Behavioral Performance in the Visuomotor Association Learning Task without Reversals
Percent correct performance (top) and reaction times (middle and bottom) are plotted taking into account all correct and incorrect trials. On the first trial in learning, accuracy is at chance (50%) and the reaction time is very high. As learning progresses, accuracy increases and reaction time decreases.

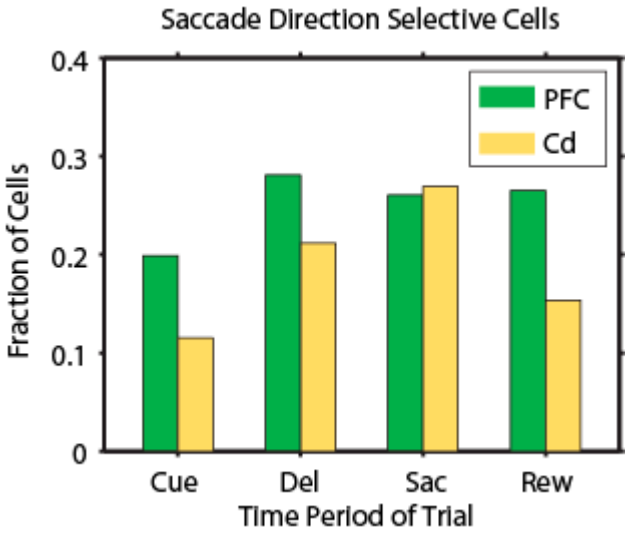

Figure 3.3
Fraction of Cells Containing Task-related Information in PFC and Cd
The fraction of cells containing saccade direction information is plotted by time epoch (Cue, Delay (Del), Saccade (Sac), and Reward (Rew)) in trial for Cd (yellow) and PFC (green). Similar trends are seen across the two areas. Fractions are based on total cell count (PFC: N = 196; Cd: N = 52). Direction cells were defined as having a significant main effect of direction from a one-way ANOVA ($p < 0.05$, Bonferroni corrected).

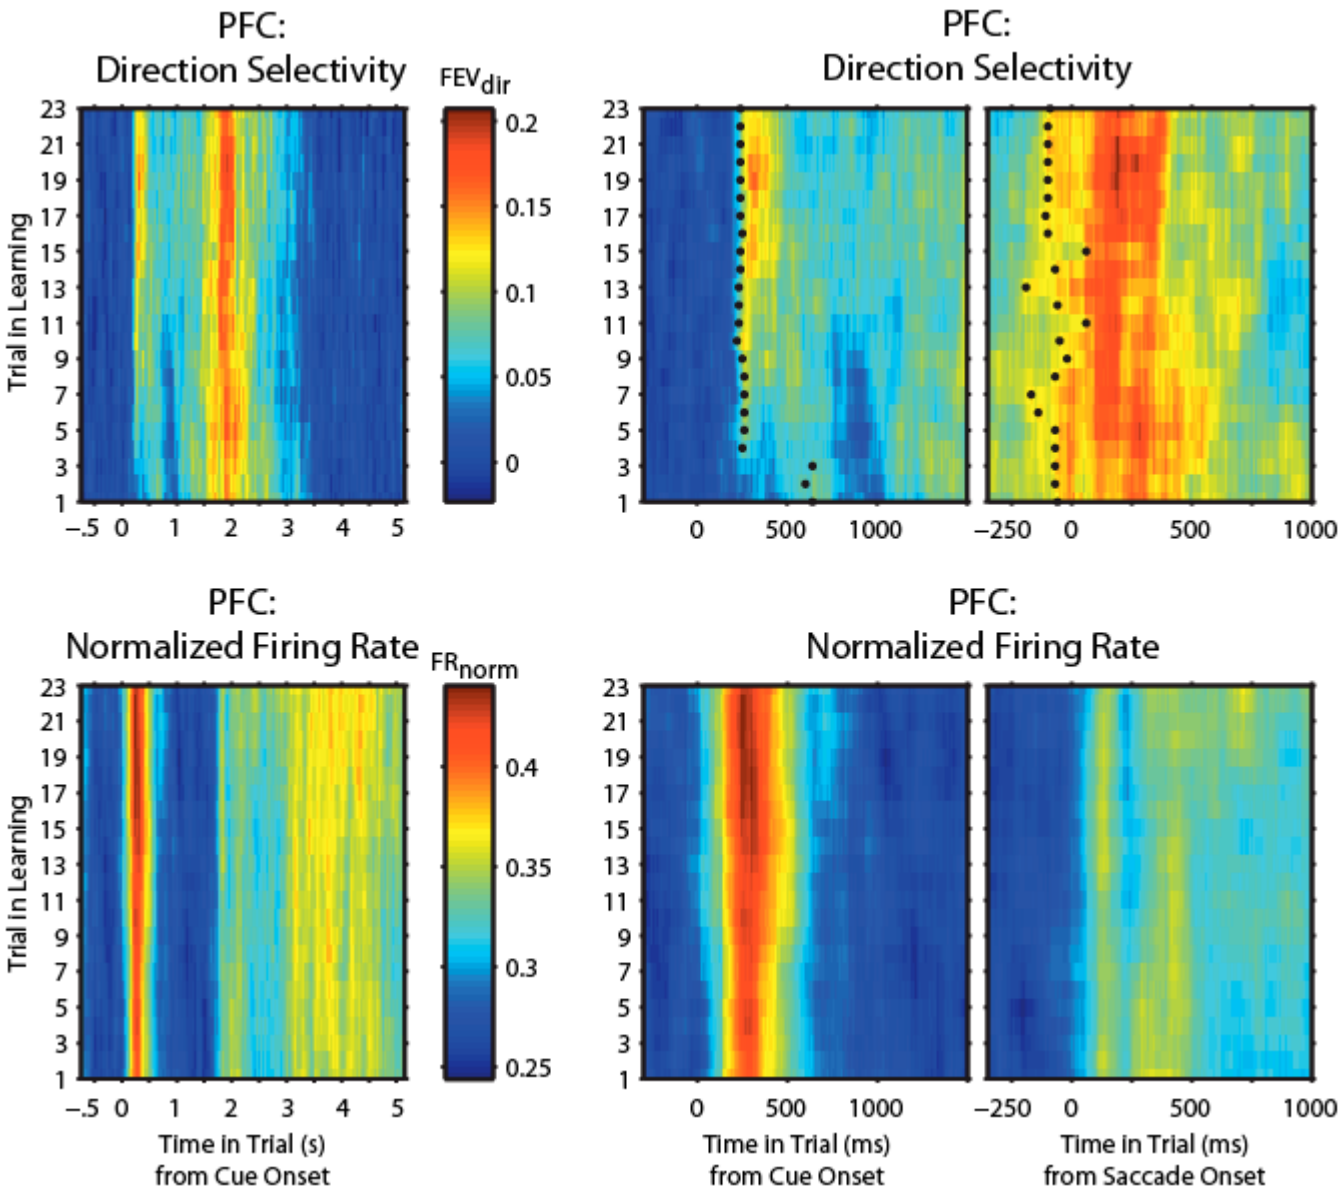

Figure 3.4
Changes in Saccade Direction Selectivity with Learning in PFC

Learning-related changes in information (FEV_{dir}) about saccade direction (top) and in normalized firing rate (bottom) for the entire saccade direction selective population of PFC cells (N = 102) is shown by these colorplots. Data from the entire trial (left) shows two main bands of increased direction information: one around the time of visual cue presentation (middle) and one around the time of saccade execution (right). Risetimes, the time to half-maximum direction selectivity, are shown as block dots. Saccade direction information during the cue period appears progressively earlier in the trial as learning progresses. No learning-related changes in firing rate are seen.

The timing change observed in the early-trial saccade direction information in PFC during the current behavioral task without reversals is compared to the changes observed in PFC and Cd during the task with reversals in Figure 3.5 (top). Whereas PFC changed slowly with learning during the task with reversals, the timing changes occurring in PFC during the task without reversals appear very similar to the changes seen in Cd during the task with reversals. The behavior of the animals is plotted in a similar manner as the risetimes for comparison purposes (Figure 3.5, bottom). The changes in the timing of saccade direction information between the two behavioral tasks qualitatively reflect the differences in behavioral performance, both reaction time and accuracy, between the two tasks.

The saccade direction information in Cd also shows differences between the two tasks. Instead of containing two bands of saccade direction information during the trial (during the cue and saccade periods) as was seen in the task with reversals, Cd shows only one band of information that occurs around the time of saccade execution during the task without reversals (Figure 3.6, top). There is little to no information during the cue period in Cd. However, the firing rate of this population of cells still increases during the cue period (Figure 3.6, bottom). There is just no information contained in this neural activity.

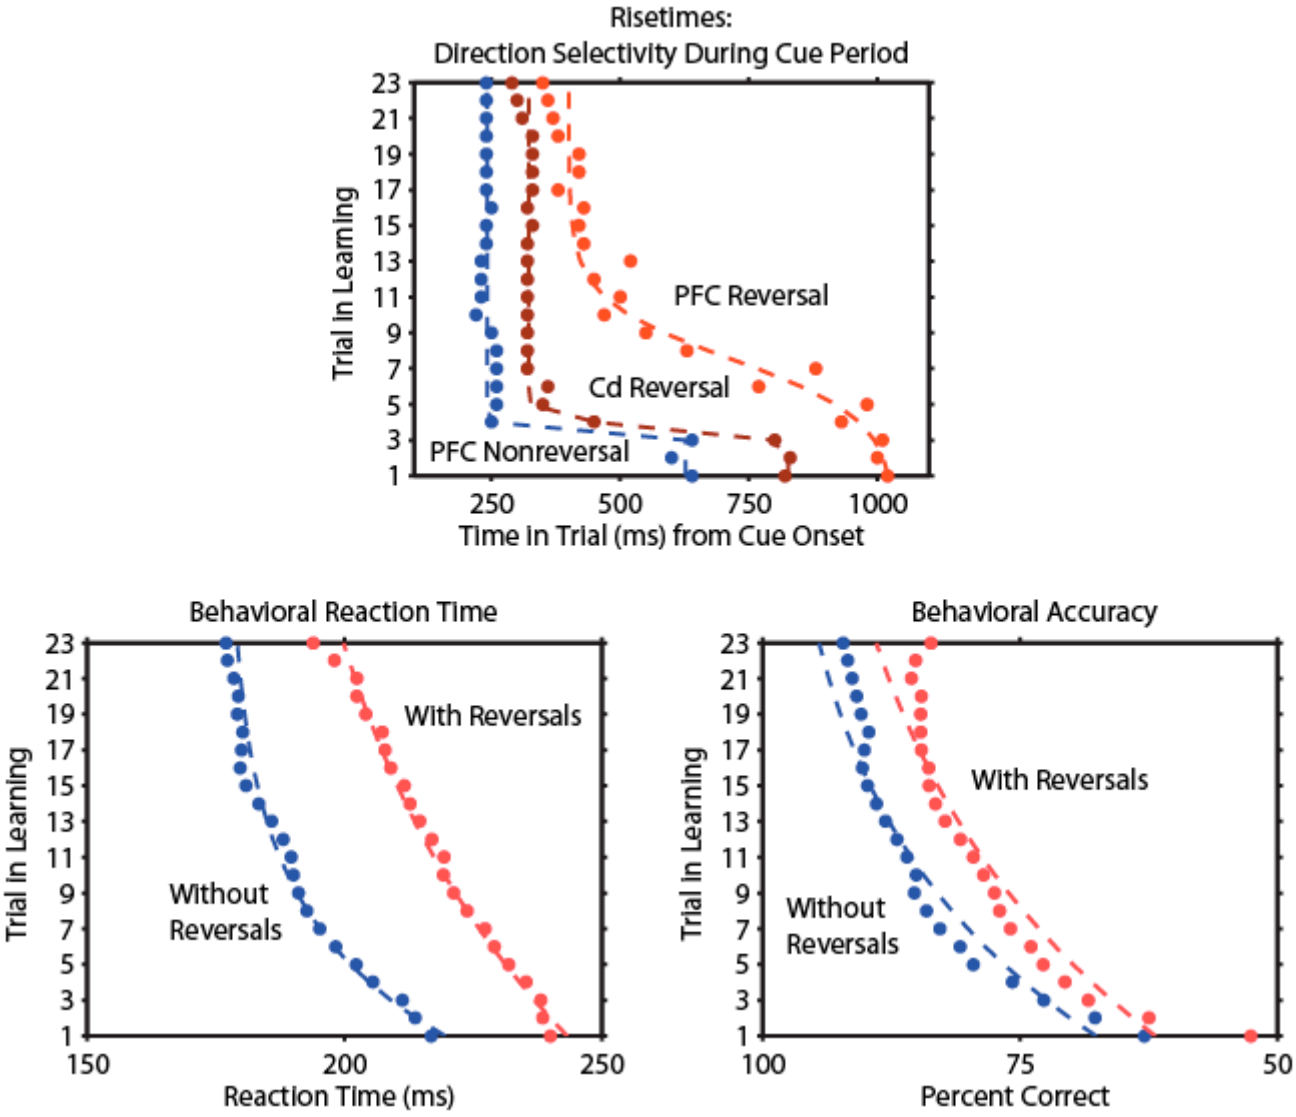

Figure 3.5
Comparing Early-Trial Saccade Direction Selectivity in PFC and Cd in Learning With and Without Reversals
Top. The timing of saccade direction information during the early-trial cue period is plotted for PFC during the learning task with reversals (light red) and the task without reversals (blue) as well as Cd during the learning task with reversals (dark red). In both tasks, saccade direction information in PFC appears progressively earlier in the trial as learning proceeds; however, this change in PFC occurs earlier in learning and more abruptly during the task without reversals, and shows striking similarities to the time course of direction information in Cd during the task with reversals. *Bottom*. Behavioral measures of reaction time (left) and accuracy (right) in the task with reversals (red) and without reversals (blue) are plotted in a similar manner to the risetimes for comparison purposes. The change in the timing of information in PFC between the two tasks qualitatively parallels the differences in behavior between the two tasks.

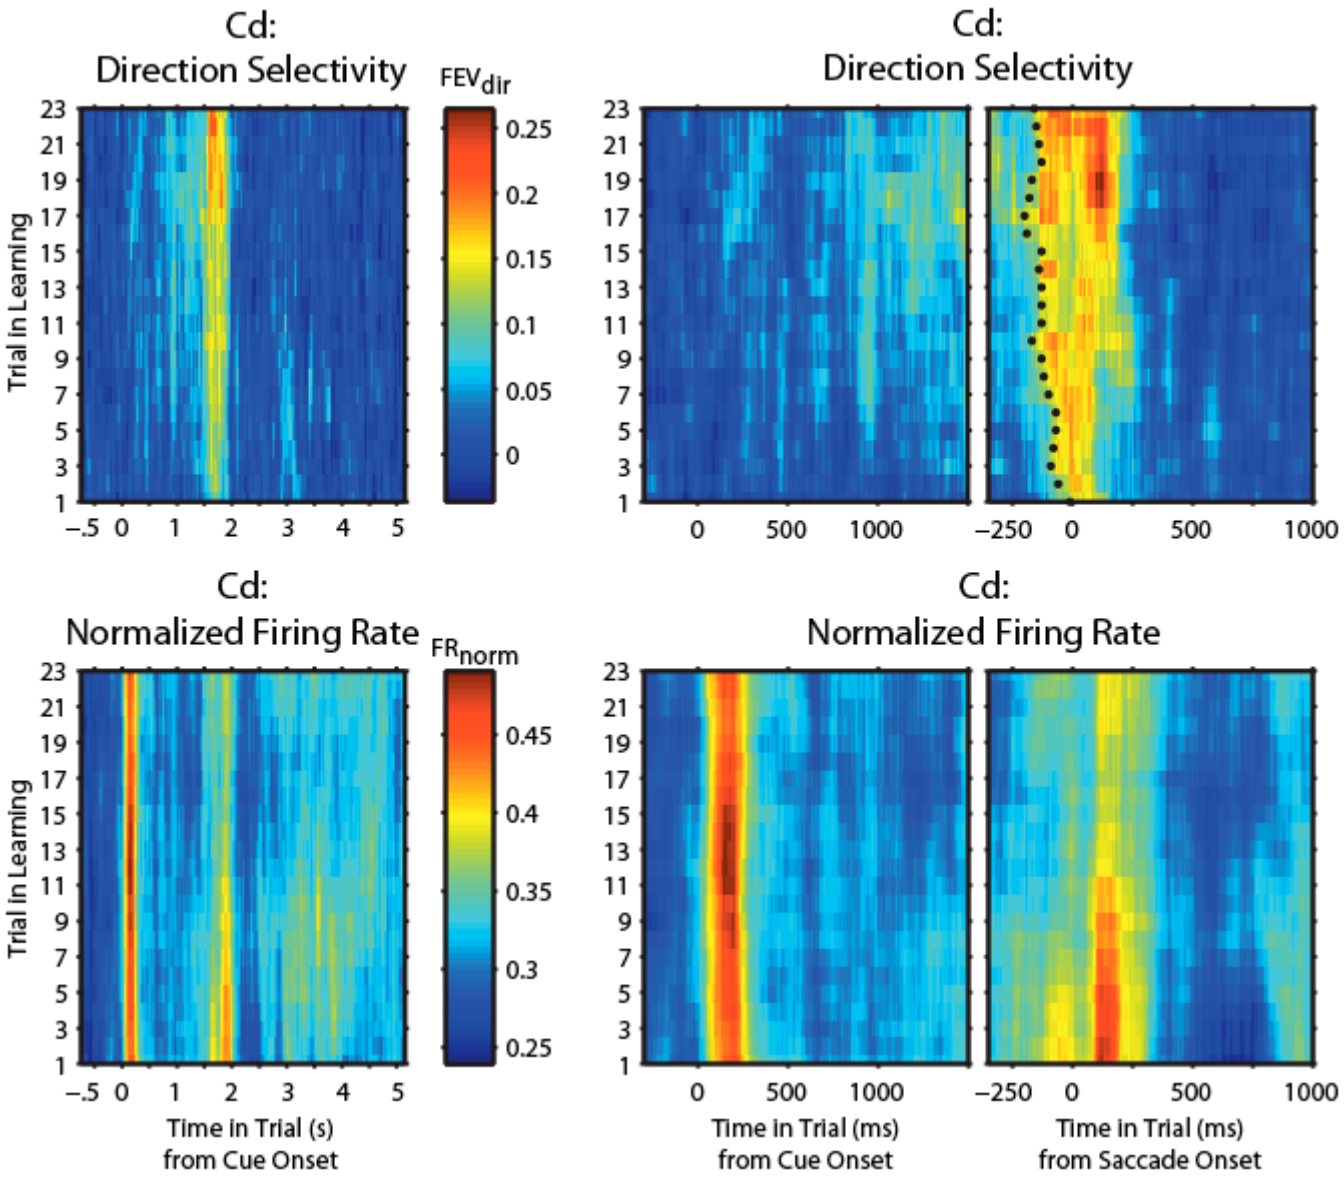

Figure 3.6
Changes in Saccade Direction Selectivity with Learning in Cd

Learning-related changes in information (FEV_{dir}) about saccade direction (top) and in normalized firing rate (bottom) for the entire saccade direction selective population of Cd cells (N = 22) is shown by these colorplots. Data from the entire trial (left) shows only one band of increased direction information that occurs around the time of saccade execution (right). Risetimes, the time to half-maximum direction selectivity, are shown as block dots. Saccade direction information during the cue period is conspicuously absent, even though firing rate increases during the cue period are still apparent.

Discussion

This experiment changed the behavioral context for learning by removing the requirement of reversing the object-saccade associations. The change in the behavioral paradigm produced two striking effects in the neural activity of PFC and Cd: 1) the early-trial saccade direction information in PFC changed abruptly at the start of learning, similar to changes seen in Cd during the task with reversals, and 2) Cd no longer contained information during the early-trial period. These effects suggest that the PFC always parallels behavioral performance, as it may be the driving force behind response selection, and that the basal ganglia's involvement in the visuomotor learning process is context-specific.

There are some important differences between the behavioral learning tasks with and without reversals. For one, the objects are only ever used *once* in the learning task without reversals; after the animal learns the object-saccade association, the object is discarded and replaced by a new object. The animal never sees the learned object again. Perhaps this limited temporal importance of the learned visuomotor association obviates the need to encode the visuomotor association information in the basal ganglia. This interpretation is supported by much of Ann Graybiel's work, as she has argued the importance of the basal ganglia in the formation of habits (Jog et al., 1999; Graybiel & Saka, 2004).

Another distinction between the two behavioral tasks is that the task without reversals has no component of inhibition of previously learned information. It is possible that the role of the basal ganglia in the inhibition of undesired motor output, through the indirect pathway, is more powerful than its role in selective activation of desired motor output, through the direct pathway (Mink, 1996; Graybiel & Saka, 2004). Thus, in a task without the need for this inhibition (i.e. the task without reversals), the role of the basal ganglia may be marginalized.

A third difference between the two tasks is the singularity of context in the task without reversals. In this task, there is only one behavioral context, whereas in the learning task with reversals there are two contexts between which the animal must switch. Perhaps the role of the basal ganglia in visuomotor learning is to provide context-dependent bias signals to cortex for the purpose of effectively guiding behavior (Passingham, 1993). However, if only one context exists, the necessity of encoding in basal ganglia might diminish.

Chapter 4: Learning-Related Changes in PFC and Cd During Learning With and Without Reversals: Analysis of Single Units and Local Field Potentials

Introduction

The previous chapter presented evidence that the role of the basal ganglia is highly dependent on the context when learning arbitrary visuomotor associations. Specifically, it was shown that Cd exhibited no motor-related information during the early-trial period during a task that did not require animals to reverse the object-saccade associations. This result opposes previous work that used a serial reversal task: these previous results suggested that the basal ganglia may learn an association first and then transfer that information to PFC to guide the slow learning process of PFC (Pasupathy & Miller, 2005; Miller & Buschman, 2007).

This comparison of the neural activity of PFC and Cd during these two tasks was performed across experiments. Thus, different cells were recorded in each data set. While the cells were recorded from the same animals, this comparison would be more appropriate when performed in the same experiment, with the same cells being recorded. This is the goal of the study presented here. The behavioral task interleaved learning with and without reversals in order to directly compare the neural activity of PFC and Cd during the two tasks.

Methods

[for more methodolgical details see Chapter 6]

Subjects

Two Rhesus macaque monkeys (*Macaca mulatta*), Monkey P and Monkey A, were used in this experiment. All animal procedures conformed to the NIH guidelines and were approved by the MIT Committee on Animal Care.

Behavioral Task

The animal performed a visuomotor learning task that interleaved learning with and without reversals of the object-saccade associations (Figure 4.1). In this task, the animal had to associate a visual stimulus with a particular saccadic response direction (e.g. see picture A, saccade right). Once the animal learned the correct associations by trial and error, the associations were reversed (e.g. see picture A, saccade left). After the animal learned the reversal, the objects were discarded and replaced by new objects and the animal had to learn the new object-saccade association (e.g. see picture B, saccade left). This pattern continued

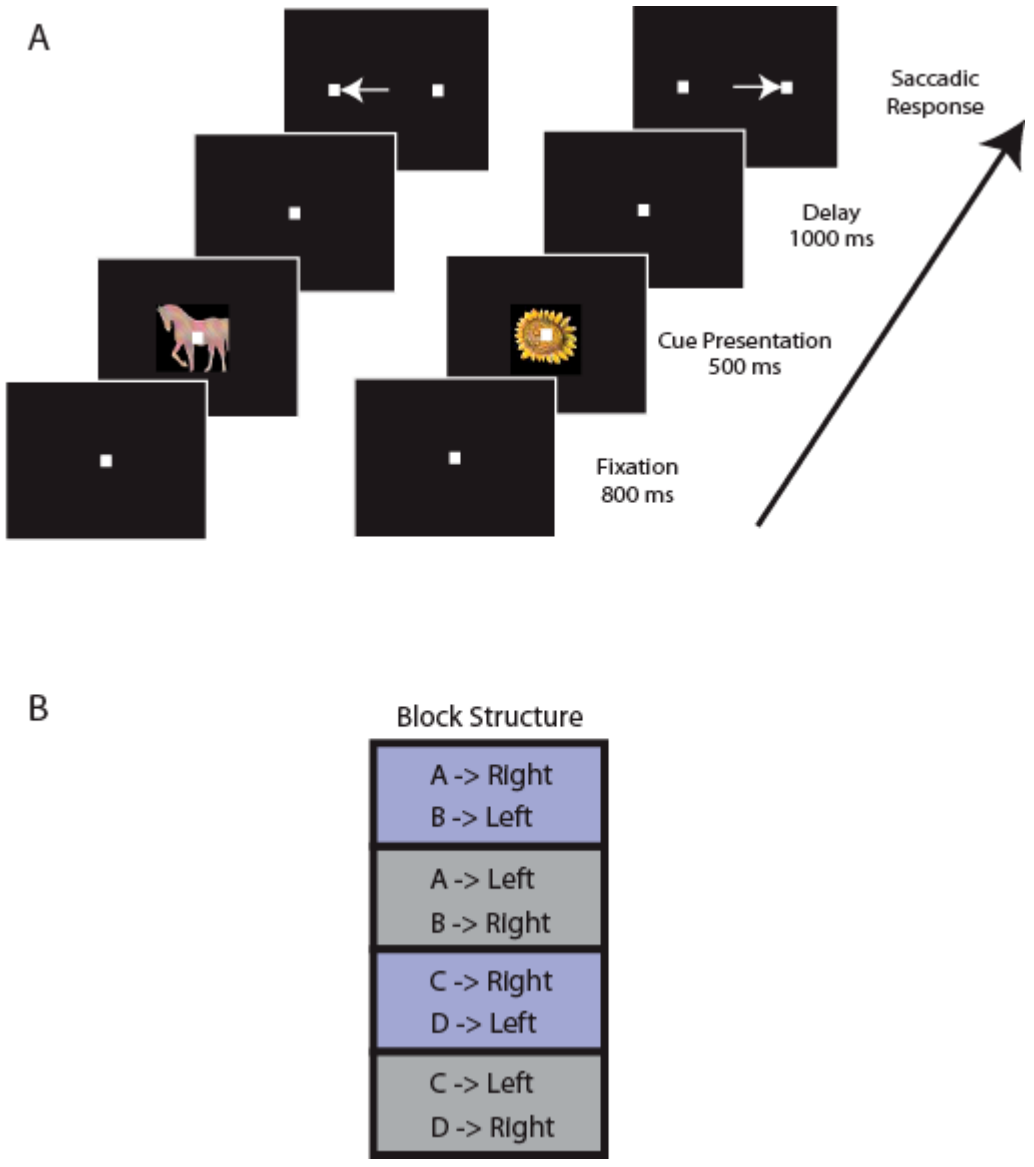

Figure 4.1
Trial Events and Block Structure of the Task Interleaving Learning With and Without Reversals
The behavioral task used in this experiment requires the animal to associate a particular saccadic response with a particular visual stimulus. ***A***. The animal is shown a visual stimulus for 500ms, and after a one-second delay period, is required to make a saccade to one of two peripherally presented targets. Each visual stimulus is associated with a particular saccade direction. Execution of the correct saccade results in the delivery of a juice reward. ***B***. The block structure for the task required the animals to learn new associations (blue) and reverse them only once (grey). Thus, every other block is a reversal of the previous block's associations.

throughout the experimental session, where every other block was a reversal of the previous block's object-saccade associations.

Electrophysiological Data Collection

Up to 16 Tungsten microelectrodes were acutely implanted into dlPFC and Cd in a single session (33 sessions for Monkey A, 31 sessions for Monkey P). In both monkeys, much of the data from the two areas was collected simultaneously, but not all. Each single unit included in analysis contained at least four full behavioral learning blocks of data. A total of 382 PFC cells (185 from Monkey A and 197 from Monkey P) and 130 Cd cells (98 from Monkey A and 32 from Monkey P) were used in analysis.

Local field potentials were also recorded from each electrode and are also analyzed here. Only electrodes from which a cell was recorded and analyzed (thus, the cell was well isolated and contained at least four full blocks of data) were used in analyses. Local field potentials were recorded at 1 kHz, and digitized for offline analysis. All signals were referenced to a common ground. Data is presented from a total of 226 electrodes in PFC (114 in Monkey A, 112 in Monkey P) and 99 electrodes in Cd (71 in Monkey A, 28 in Monkey P). For coherence analysis, 314 PFC-Cd pairs (228 in Monkey A, 86 in Monkey P), 792 PFC-PFC pairs (340 in Monkey A, 452 in Monkey P), and 170 Cd-Cd pairs (140 in Monkey A, 30 in Monkey P) were used in analysis. All electrodes used in the pairings were at least 1mm apart.

Analytical Techniques

Analysis of the behavioral data focused on accuracy and reaction time across learning. All recording sessions (Total: N=64) were used in this analysis (Figure 4.1). All correct and incorrect trials were used in the order in which they occurred in each learning block (trials in which the animal made a fixation error were ignored). Means +/- SE were computed across all blocks from all sessions (420 total blocks with reversal: Monkey A, N=203; Monkey P, N=217; 428 total blocks without reversal: Monkey A, N=208; Monkey P, N=220).

To determine which cells of the recorded populations contained task-related information, a 2-way ANOVA (with object identity and saccade direction as factors) was performed on each cell's average firing rate in each of 4 task epochs: 100-600ms after cue onset (“Cue”), 600-1500ms after cue onset (“Delay”), 150ms before to 150ms after saccade initiation (“Saccade”), and 50-

300ms after the start of reward delivery ("Reward"). A cell was considered to contain object or saccade direction information if it *only* had a significant main effect of object or direction, respectively. A cell was considered to contain object-saccade association information if it had a main effect of object *and* direction or a significant interaction term. Significance level was set at $p < 0.05$, and was corrected for multiple comparisons using Bonferroni's correction. Because this analysis did not seem to accurately capture the information contained in the cell population, a second analysis was performed to extract cells with saccade direction information. A sliding receiver operating characteristic (ROC) analysis was performed across time in trial and it compared all left trials with all right trials. ROC analysis essentially provides a quantification of the separation between two distributions; in this case, the firing rate of left trials versus the firing rate of right trials. The ROC calculation was performed on 50ms bin of data slid by 50ms throughout the trial (thus, the bins are nonoverlapping). The ROC values obtained from each calculation was then converted into a zscore by a randomization procedure: 1) the ROC calculation was performed 1000 times with the left and right trial assignment randomly assigned, 2) then for each time bin the true ROC value was converted to a zscore based on the distribution of ROC values obtained from the randomization procedure. The results of this analysis are shown in Figure 4.3B, where the ROC zscores for each time bin across the trial (x-axis) are plotted for each individual cell (y-axis). A cell was determined to contain saccade direction information if at least two consecutive time bins received a zscore of greater than 3. Cells are aligned in Figure 4.3B on the first significant time bin.

To validate that this ROC analysis was accurately selecting cells with saccade direction information, percent explained variance by the direction factor (using the parsing of variance of the 2-way ANOVA) was plotted for the entire population of cells, as well as for the population of direction selective cells determined by the ROC analysis (Figure 4.4). Since the direction population very accurately portrayed the direction information in the entire population, this direction selective population was used in all further analyses. Changes in saccade direction information with learning were quantified and displayed as in previous experiments, and the same linear measures of the population trends were used: risetime for timing trends and peak value for strength trends. These measures were quantified separately for the learning blocks with reversal versus without reversals and compared in Figures 4.6 and 4.8.

The analysis of local field potentials focused on two main quantitative measures: power and coherence. The power spectrum was calculated using a one-sided fourier transform of the continuous LFP signals. Power is calculated as the square of the magnitude of the transform. For data presented in Figures 4.9 and 4.10, 200ms bins of data were used for the power calculations (thus resulting in 5Hz resolution in the frequency domain). Because of the tradeoff between time and frequency resolution in this analysis, 200ms was chosen to provide adequate resolution in both domains. For the analysis of learning-related changes in the power spectrum, the data in the entire cue (0-500ms after cue onset) or delay periods (0-1000ms after cue offset) were used. Averages were then taken across prominent, high power, frequency bands to quantify learning-related changes.

In order to investigate communication between brain areas, coherence analysis was performed on the LFPs. The coherence statistic provides a measure of the extent that two signals are oscillating together. The idea behind the use of this analysis is that if two signals are oscillating together, the probability is higher that they are in direct or indirect communication. The coherence statistic is normalized by the power in each frequency, thus underlying power changes (as seen in the power analysis) should not affect the results of this coherence analysis. Coherence was calculated as:

$$S_{yw} = (\hat{Y}^*(f) / \rho_Y(f)) (\hat{W}(f) / \rho_W(f))$$
$$Coh = | E \{S_{yw}\} |$$

Where S_{yw} is the normalized cospectrum of two signals Y and W, $\hat{Y}^*$ is the complex conjugate of the transform of Y, $\hat{W}$ is the transform of W, ρ_Y is the magnitude of the transform of Y, ρ_w is the magnitude of the transform of W, and Coh is the coherence statistic that is presented in Figures 4.15 and 4.16.

Results

Analysis of Single Units

The animals' behavioral performance (Figure 4.2) on this task that interleaved learning with and without reversals was consistent with their performance when the two tasks were performed separately. During blocks where the animals were presented with new objects with which to form object-saccade associations (termed "new learning" or "without reversals"), accuracy

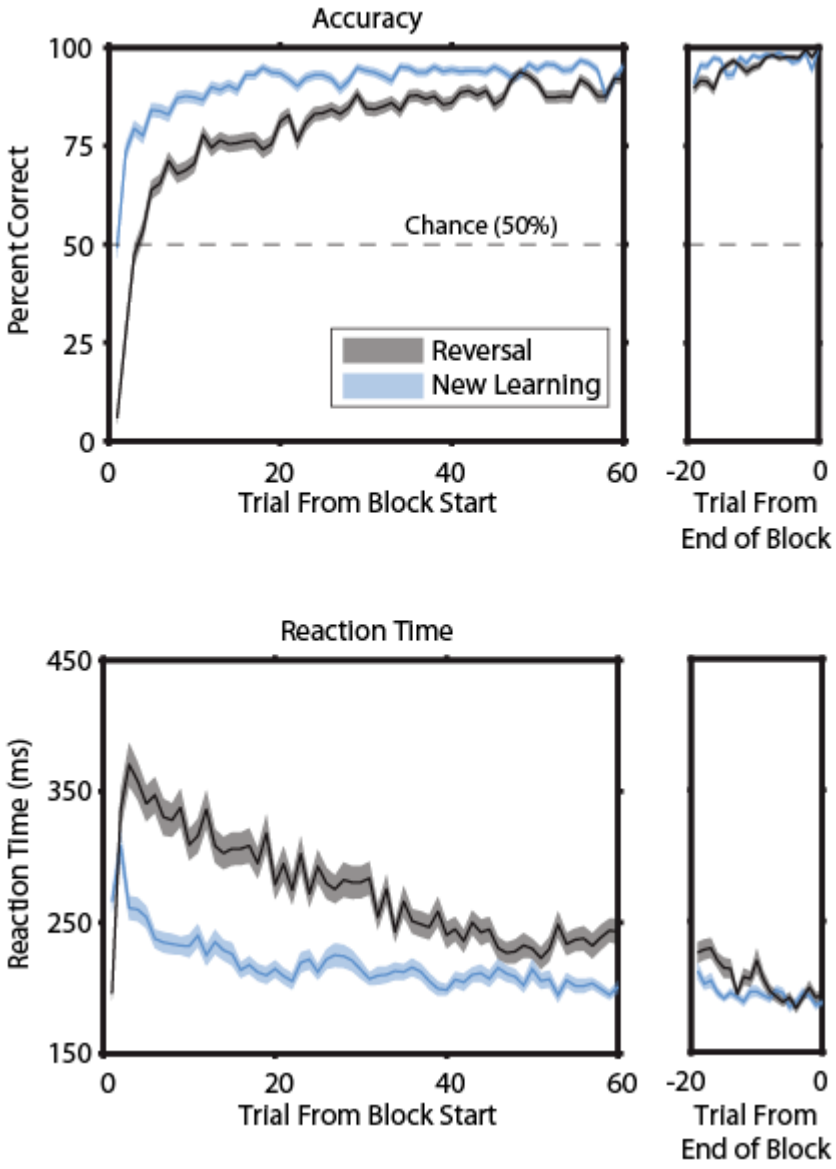

Figure 4.2
Behavioral Performance in the Task Interleaving Learning With and Without Reversals
Percent correct performance (top) and reaction times (bottom) are presented for all correct and incorrect trials during the learning task that interleaves learning with (gray) and without (blue) reversals. As was seen when these tasks were performed separately by the animals, on the first trial in learning, accuracy is close to 0% during reversals and the animals perform very quickly since the previously rewarded associations are still being performed. Subsequently, performance jumps to chance level (50% correct, dotted line) and reaction time sharply increases. During new association learning reaction time at the start of a new block begins high and accuracy starts at chance level (50% correct). In general, as learning progresses, accuracy increases and reaction time decreases.

began at chance level (50% correct) and slowly rose to the learning criterion (90% correct), and reaction time began high and decreased as learning progressed. In contrast, during reversal learning blocks, accuracy began at 0% correct and reaction time was very low, since the animals were still performing the previously rewarded associations. After the first trial of the reversal block, accuracy rose to chance level and reaction time jumped over 100ms. With learning, reaction time decreased and accuracy increased to criterion level.

To determine which cells of the recorded population (PFC: N=382, Cd: N=130) contained task-related information, a 2-way ANOVA (with object and direction as factors; $p<0.05$, Bonferroni corrected) was performed on the firing rate of each cell at each of four time epochs of the trial: 100-600ms after cue onset ("Cue"), 600-1500ms after cue onset ("Delay"), 150ms before to 150ms after saccade initiation ("Saccade"), and 50ms-300ms after the start of reward delivery ("Reward"). Results of this analysis are presented in Figure 4.3A. This analysis did not reveal many cells containing direction information (PFC: 12/382, Cd: 4/130), but a large number of cells containing object-saccade association information. Because an inspection of each cell's histograms revealed a different qualitative result (where it seemed that many cells were carrying information regarding saccade direction), a sliding receiver operating characteristic (ROC) analysis was performed to extract cells with saccade direction information (Figure 4.3B). This analysis revealed many more direction-selective cells: 241 of 382 PFC cells and 49 of 130 Cd cells.

To validate the results of the ROC analysis, the fraction of explained variance explained by the direction factor was calculated across time in trial and also across trials in learning. Figure 4.4 presents this data for the entire recorded population of PFC and Cd cells (top row) and also for the direction-selective populations revealed by the ROC analysis (bottom row). These figures are very similar; that is, the identified population of cells thought to contain saccade direction information does a very good job at representing the information present in the entire recorded population. In PFC, two distinct bands of information are present: one during the cue period and the other around the time of saccade execution. This information is present very early in learning during the early-trial period and does not change drastically as learning progresses. Cd exhibits a slightly different pattern of information. This population of cells does not show two bands of information, but fairly sustained information throughout the trial, up to and including the time of saccade execution.

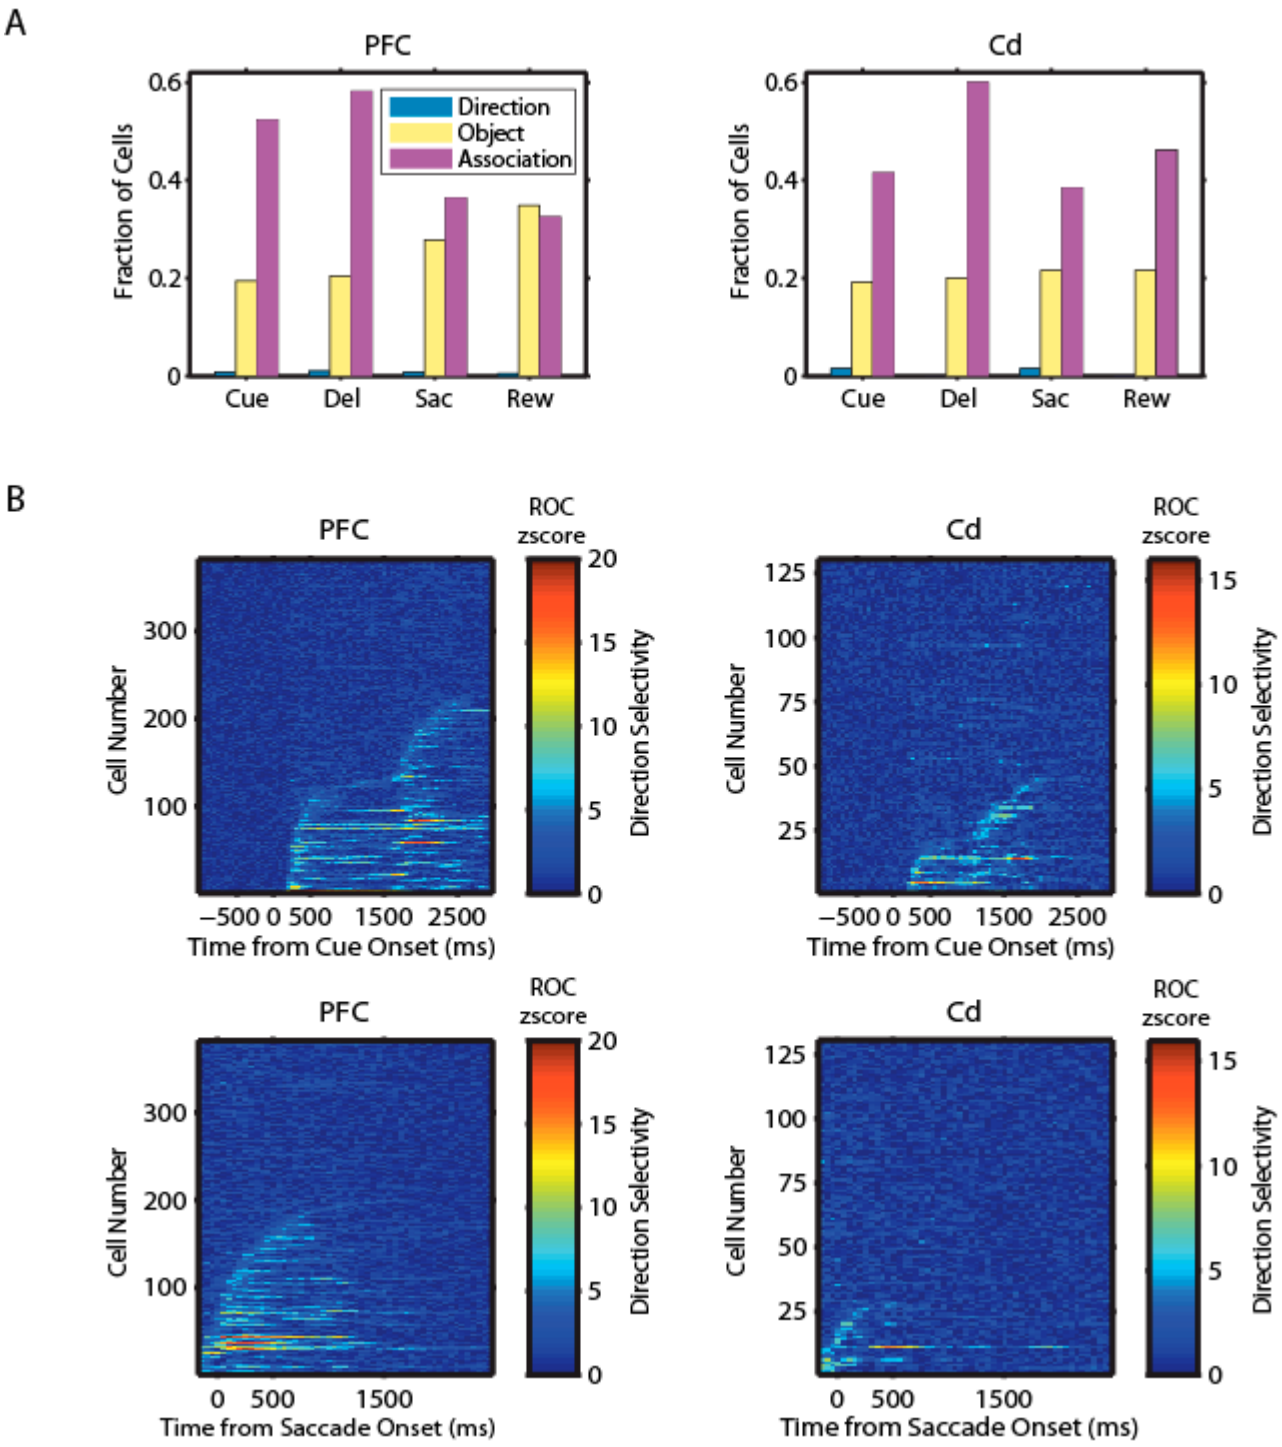

Figure 4.3
Fractions of Cells Containing Task-related Information in dlPFC and Cd
A. The fraction of cells containing saccade direction (blue), object identity (yellow), and object-saccade association (purple) information is plotted by time epoch (Cue, Delay (Del), Saccade (Sac), and Reward (Rew)) in trial for dlPFC (left) and Cd (right). These results were obtained from a two-way ANOVA with object and direction as factors. Very few cells containing saccade direction information were found (PFC: 12/382, Cd: 4/130) *B*. To more precisely extract cells containing saccade-direction information, a sliding ROC analysis was performed and all cells with significant direction information were termed “direction cells” (PFC: 241/382, Cd: 49/130).

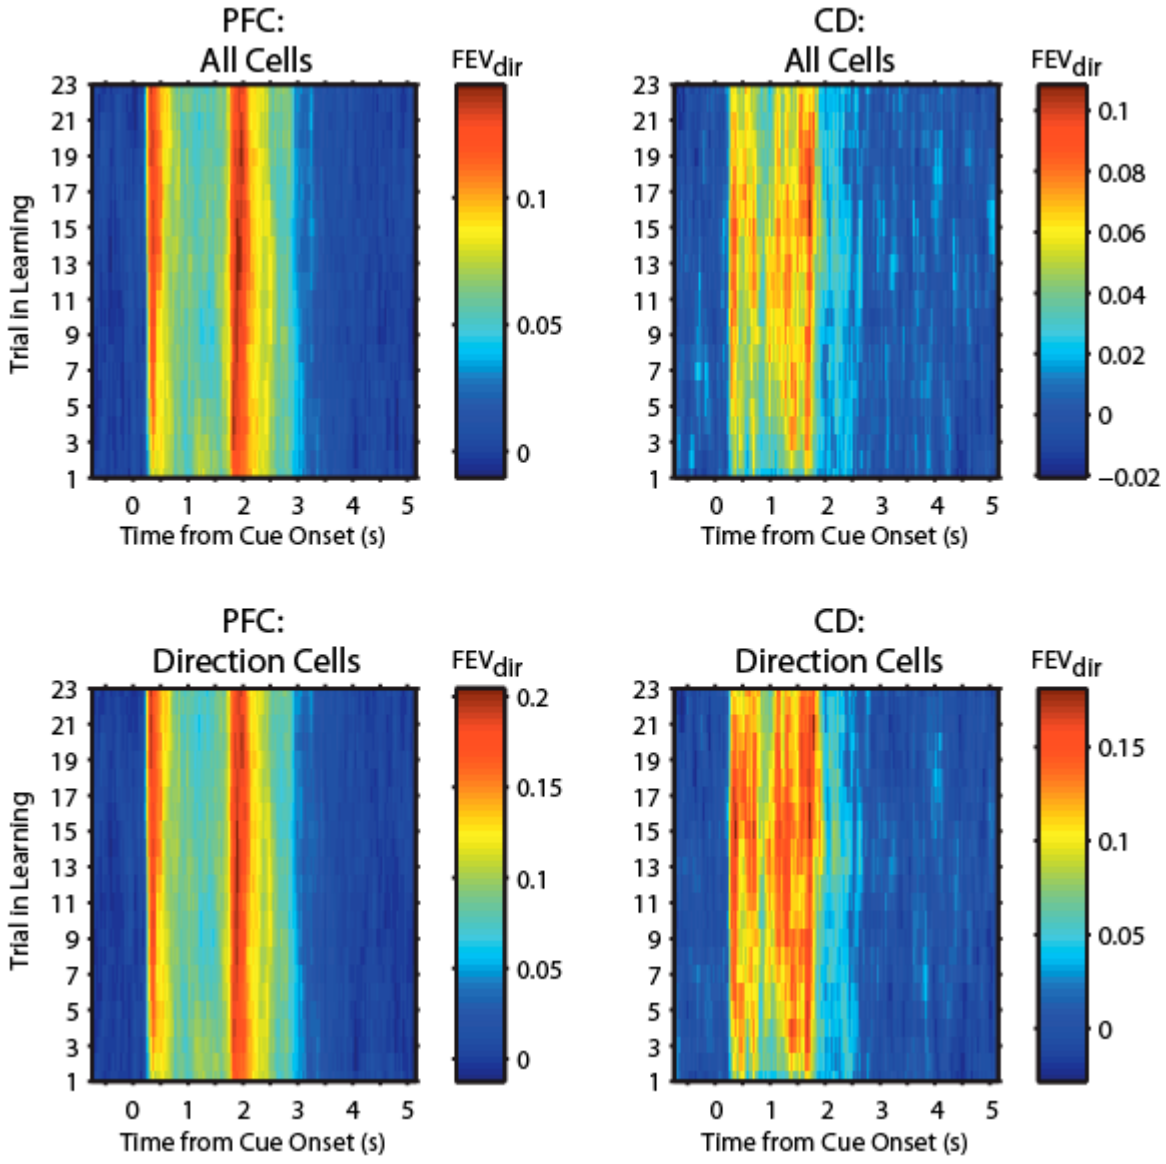

Figure 4.4
Information about Saccade Direction in PFC and Cd
Information about saccade direction is apparent in both PFC and Cd during the performance of the behavioral task. The top row presents the saccade direction information (quantified by fraction of explained variance by the direction factor, FEVdir) in the entire population of recorded cells (PFC: 382, Cd: 130). The bottom row shows the information contained in the population of direction-selective cells revealed by the ROC analysis (PFC: 241, Cd: 49). The population of direction cells qualitatively represents the entire population accurately, where there are two distinct bands of information in PFC and sustained information throughout the trial in Cd.

The crucial point of this experiment, though, is to see if there are differences between learning with and without reversals. Figure 4.5 plots the saccade direction information in the direction-selective population of PFC cells separately for the two types of learning. The information present during the execution of the two tasks is quite similar. However, there are some subtle differences to point out that might be most easily observed in the risetime and strength of selectivity plots of Figure 4.6. During the task with reversals, saccade direction information is stronger and appears earlier in the trial at the start of learning than information in the reversal task. In both tasks, strength of early-trial direction selectivity increases as learning progresses. Information around the time of saccade execution does not differ between the two tasks, although there is an increase in the strength of information in the middle of the learning process without reversals that is not seen in the reversal task.

The saccade direction information in Cd shows different trends in general throughout the trial, although there are similar differences between the two tasks as was present in PFC (Figures 4.7 and 4.8). Instead of two distinct bands of information, Cd shows more sustained information throughout the trial. Like PFC, Cd exhibits greater strength of information during learning without reversals than learning with reversals; however, unlike PFC, this trend is seen in the saccade period in addition to the cue period. Similar timing differences during the early-trial period are also seen in Cd, where saccade direction information is present earlier in the trial at the start of learning in the task without reversals.

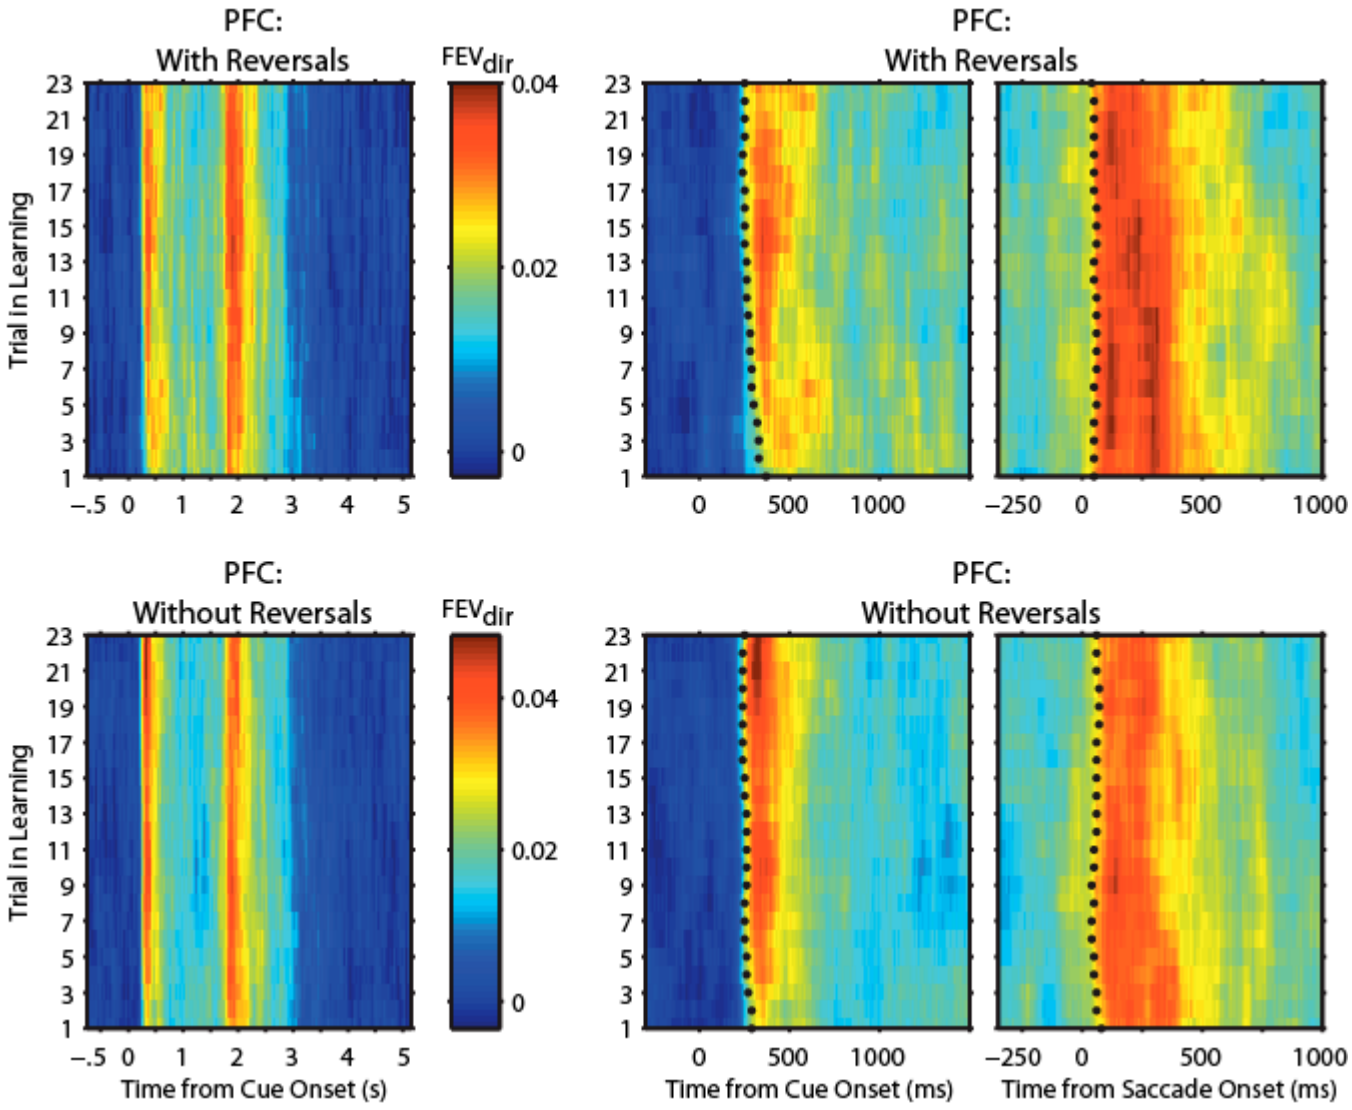

Figure 4.5
Learning-related Changes in PFC Direction Selectivity With and Without Reversals
The saccade direction information for the population of PFC direction-selective cells (N = 241) is show separately for learning with (top) and without (bottom) reversals. In both tasks, two distinct bands of information are present (left). However, when looking closer at the early-trial period (middle), subtle differences between the two tasks in the timing and strength of information are apparent.

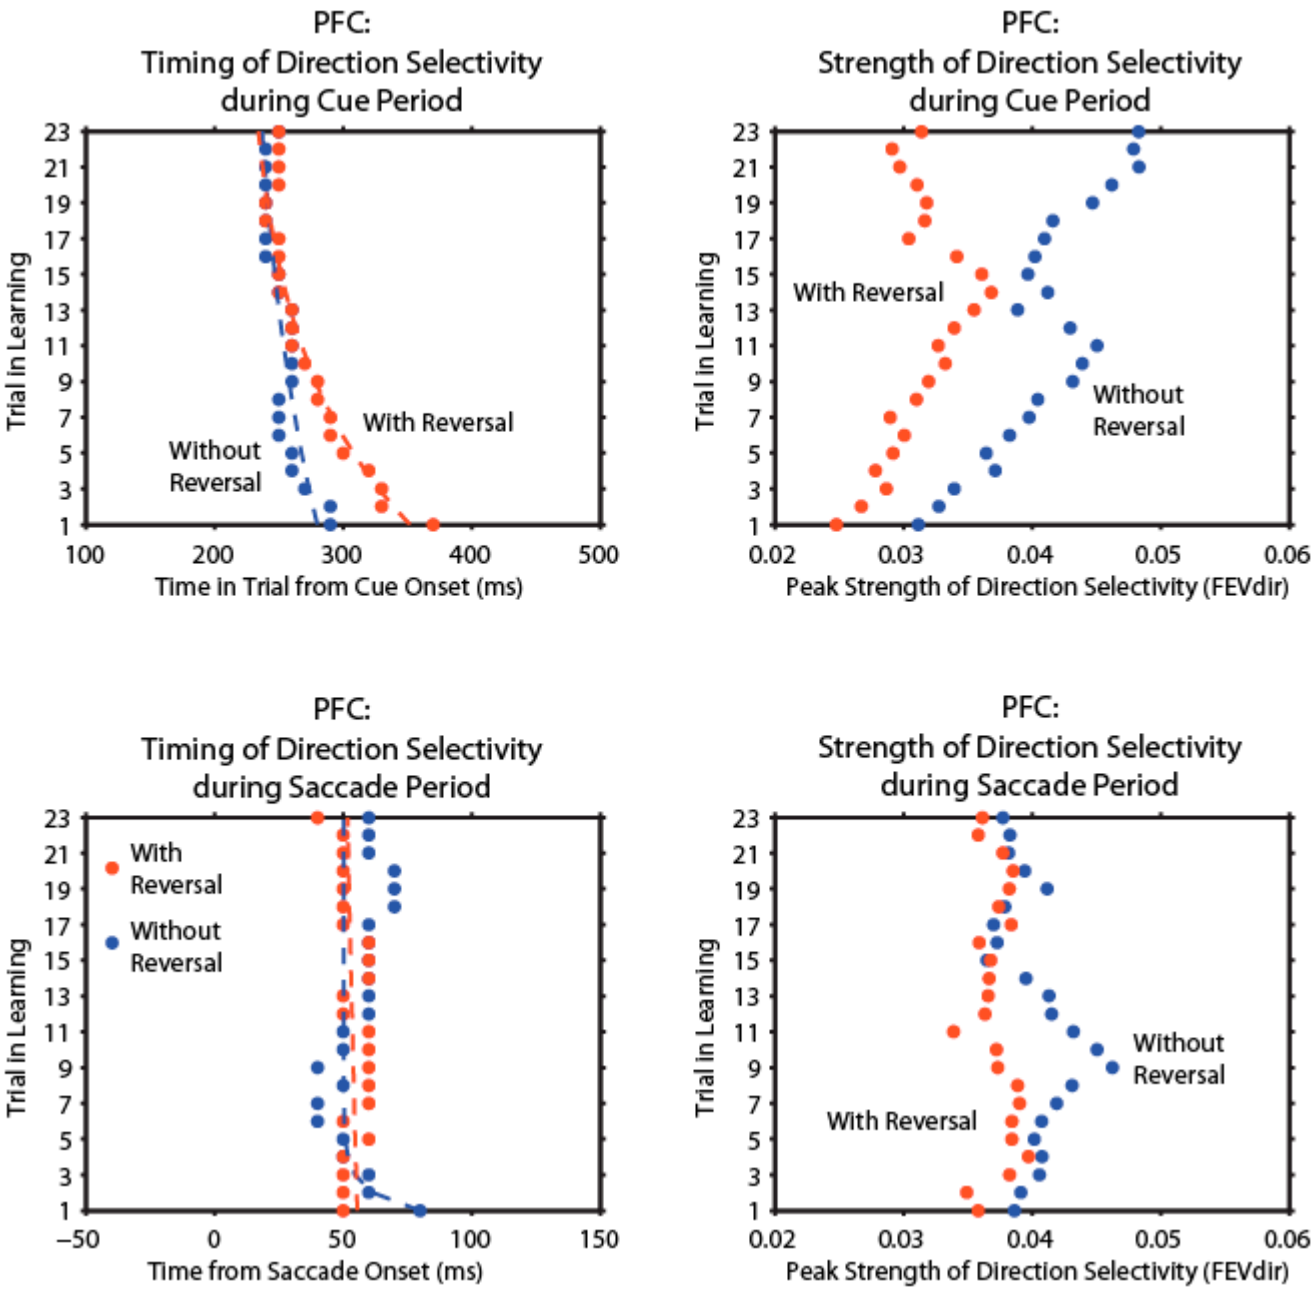

Figure 4.6
Differences in Strength and Timing of PFC Direction Information With and Without Reversals
The timing (left) and strength (right) of saccade direction information in PFC is depicted for both the reversal (red) and nonreversal (blue) tasks. During the cue period (top), information is present earlier in the trial during the task without reversals and information through learning is greater without reversals. Information during the saccade period (bottom) does not change with learning or differ between the two tasks in PFC.

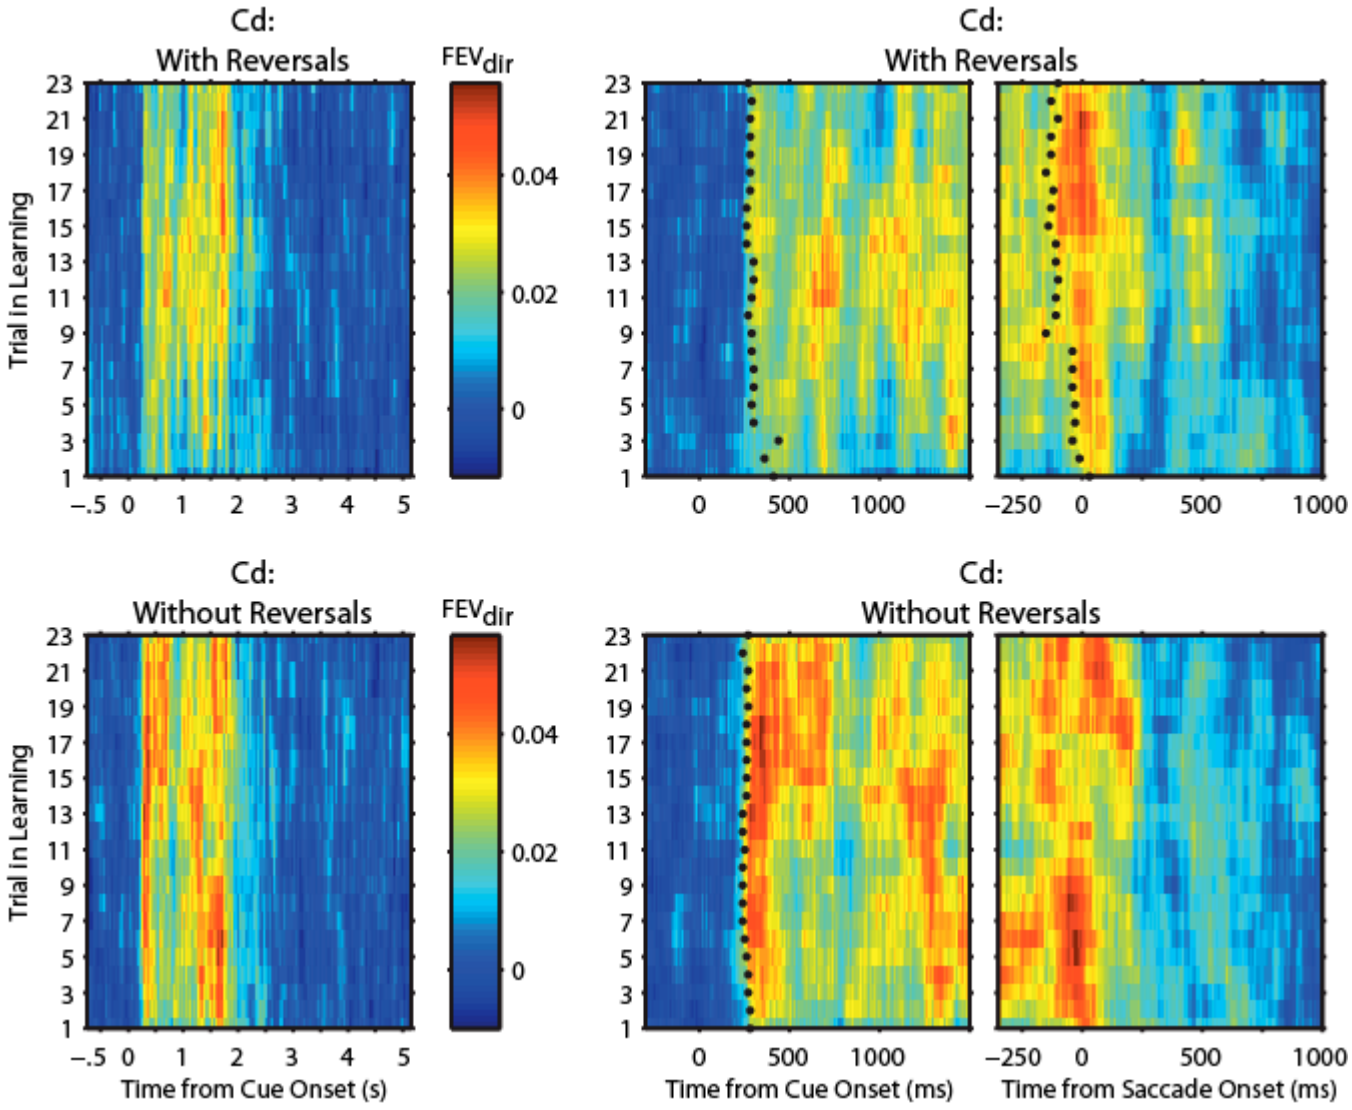

Figure 4.7
Learning-related Changes in Cd Direction Selectivity With and Without Reversals
The saccade direction information for the population of Cd direction-selective cells (N = 49) is show separately for learning with (top) and without (bottom) reversals. This population of cells does not exhibit two distinct bands of information, but more sustained information throughout the trial. The most noticeable difference between the two tasks also occurs during the cue period, where information appears earlier and is stronger in the task without reversals.

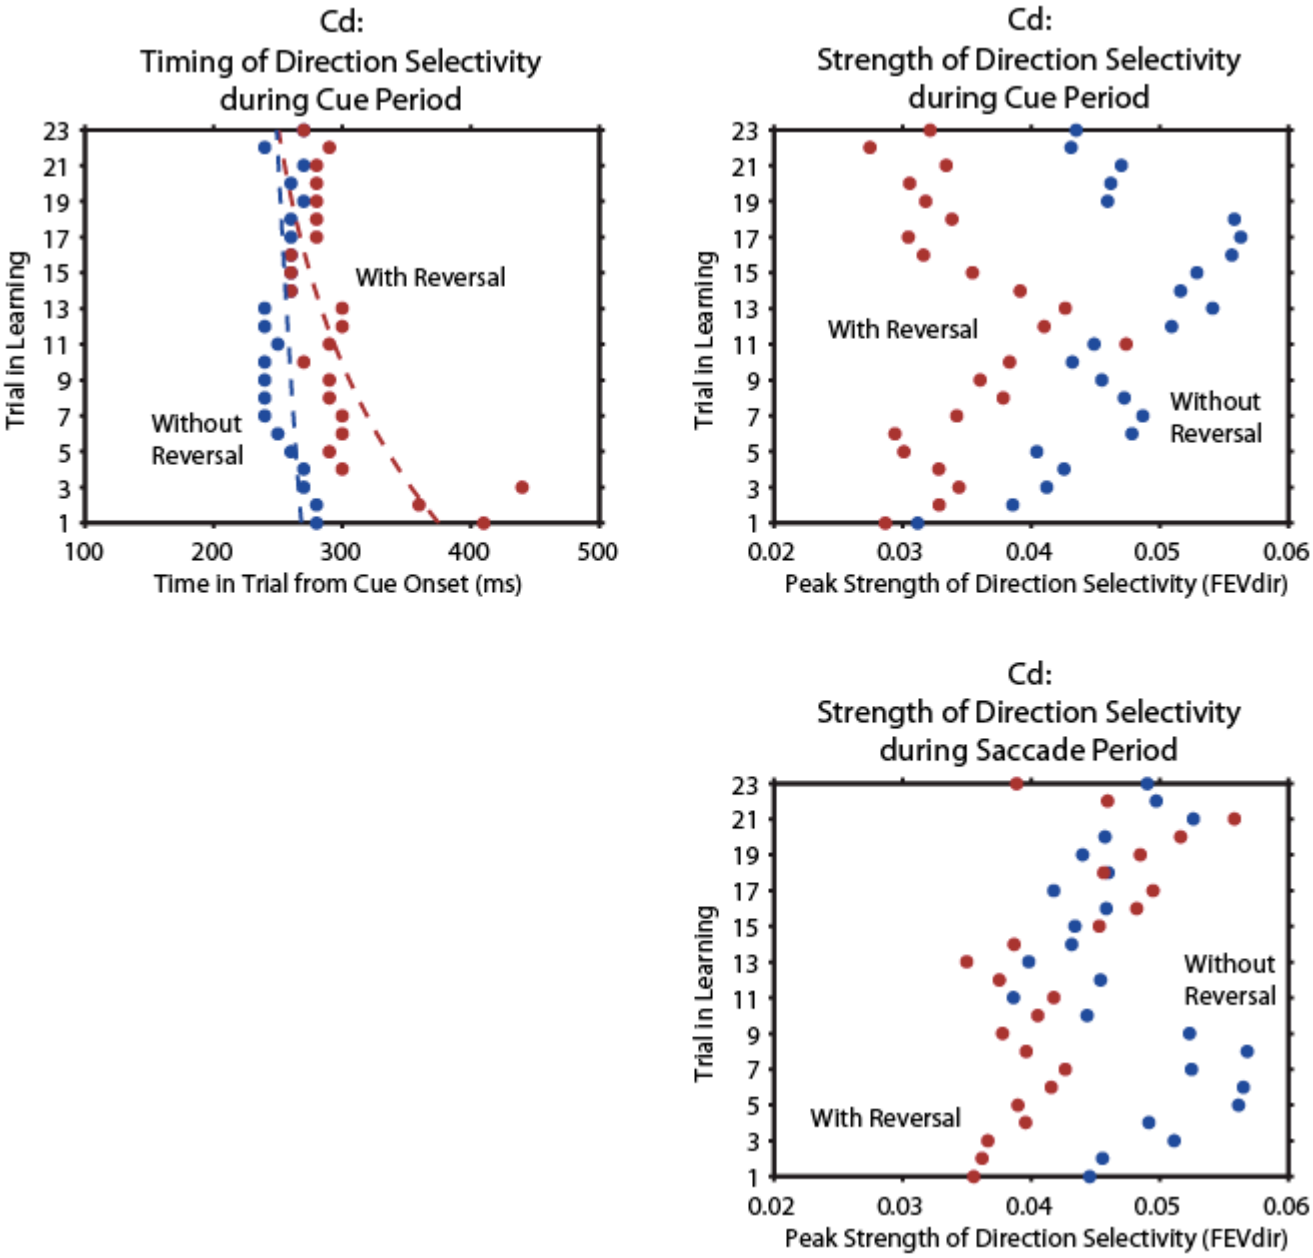

Figure 4.8
Differences in Strength and Timing of Cd Direction Information With and Without Reversals
The timing (left) and strength (right) of saccade direction information in Cd is depicted for both the reversal (red) and nonreversal (blue) tasks. During the cue period (top), information is present earlier in the trial during the task without reversals and information through learning is greater without reversals. Since there was no clear band of information around the time of saccade execution, no timing information could be reasonably extracted from the data. Strength calculations (bottom) do show differences between the two tasks early in learning, where the task without reversals contains higher information than the task with reversals.

Analysis of Local Field Potentials

Figure 4.9 presents the average power spectrum of all PFC (N=226, Monkey A: N=114, Monkey P: N=112) and Cd (N=99, Monkey A: N=71, Monkey P: N=28) local field potential (LFP) recordings plotted across time in trial. In both brain areas the lower frequencies contain the most power, and this power changes as trial events occur. Specifically, a decrease in power is observed during the time of cue presentation and saccade execution in frequencies around 15-25 Hz. There are some prominent differences in the power spectrums of the data collected from the two animals, so, in addition to the spectrograms of the entire data set (top row), spectrograms of the data from each animal are presented separately. Data from Monkey A (middle row) shows an increase of power in high frequencies (~75-115Hz) at the end of the trial in Cd. The power spectrum of LFPs recorded from Monkey P exhibit high power in lower frequencies (~10-15Hz in PFC and ~10Hz in Cd) that is not observed in Monkey A.

The power increase of high frequencies in Cd observed in Monkey A at the end of the trial was hypothesized to be linked to reward (and possibly the influx of dopamine). With this hypothesis in mind, comparing the late-trial power spectrum during correct trials with incorrect trials should show differences if this power is related to reward. Figure 4.10 presents the late-trial power spectrum with the data aligned on the initiation of the animal's saccadic response. The incorrect trials do exhibit a much smaller power increase than the correct trials; however, a power increase of these high frequencies does still occur.

In order to examine the changes in these oscillations across learning, a specific time period of the trial had to be chosen. The first period examined was the entire delay period, from 0-1000ms after the offset of the visual stimulus. These results are presented in Figure 4.11. Again, due to the differences between the animals, the data is presented with both animals combined (top row) and then separately for each animal. In general, the frequencies with the highest power in Monkey A are higher than those in Monkey P. Learning-related changes were not found in the combined data or in Monkey P's data. However, Monkey A shows a decrease in power of frequencies in the beta range (~16-34Hz) as learning progresses. Although this effect is seen in both PFC and Cd, the range of frequencies affected is slightly higher in PFC (~25-34Hz) than in Cd (~16-25Hz).

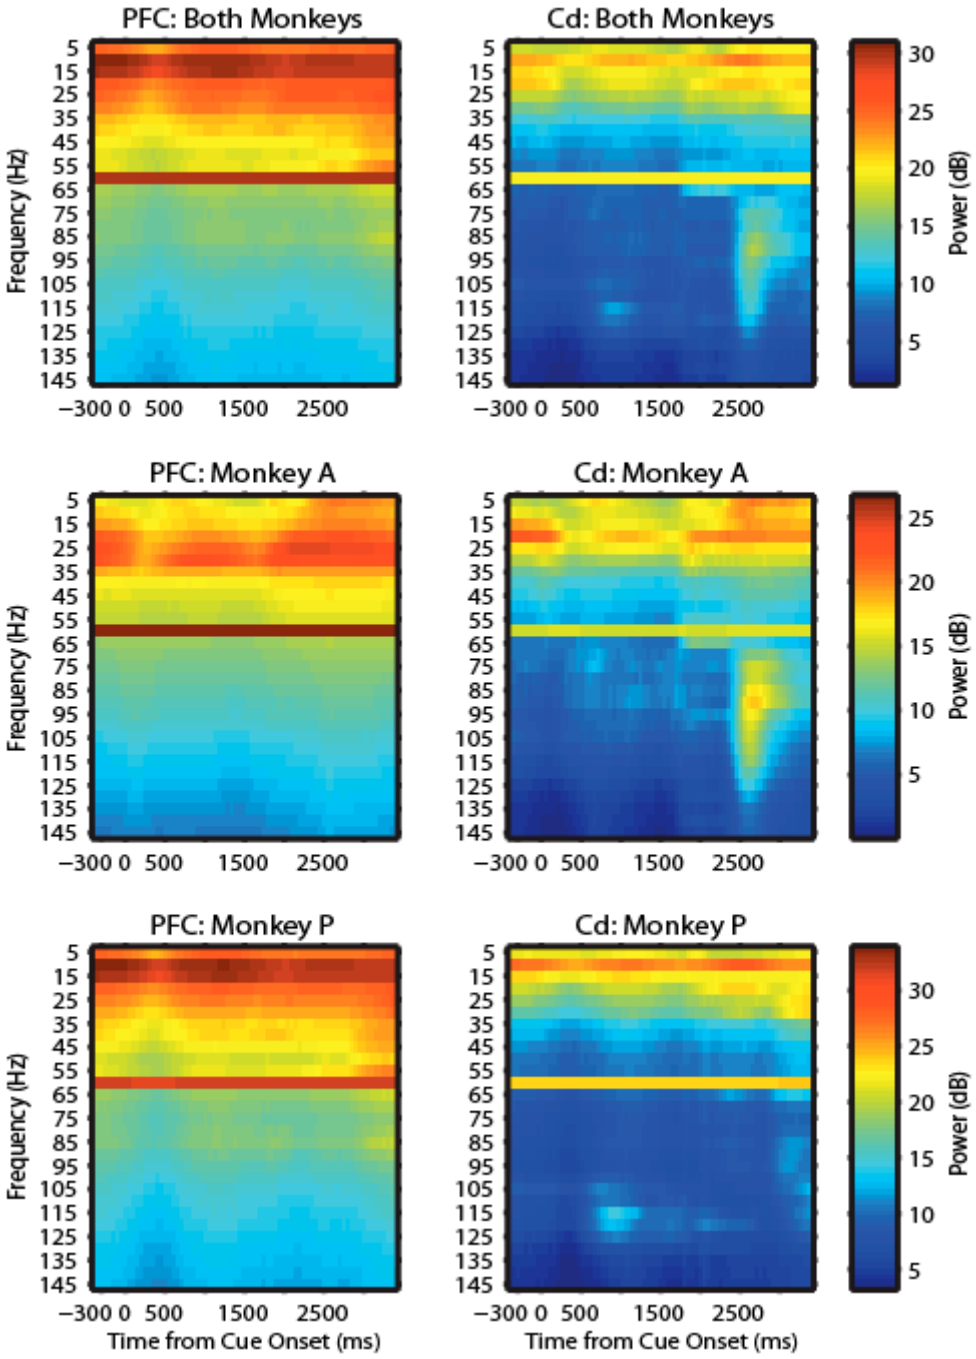

Figure 4.9
Changes in PFC and Cd Power Spectrum through Trial Events
Average power spectrum density for all PFC (left) and Cd (right) electrodes is plotted across time in trial. The highest power is seen in the lower frequencies in both areas, and these show event related desynchronization (i.e. power decreases). Data from each animal is shown combined (top row) and separated (middle and bottom rows) to show differences in the two animals. Monkey A shows an increase in power at relatively high frequencies at the end of the trial in Cd, and in general Monkey P shows higher power in the 5-15Hz range across the trial in both PFC and Cd. The increase in power at 60Hz in all recordings is due to the standard electrical noise.

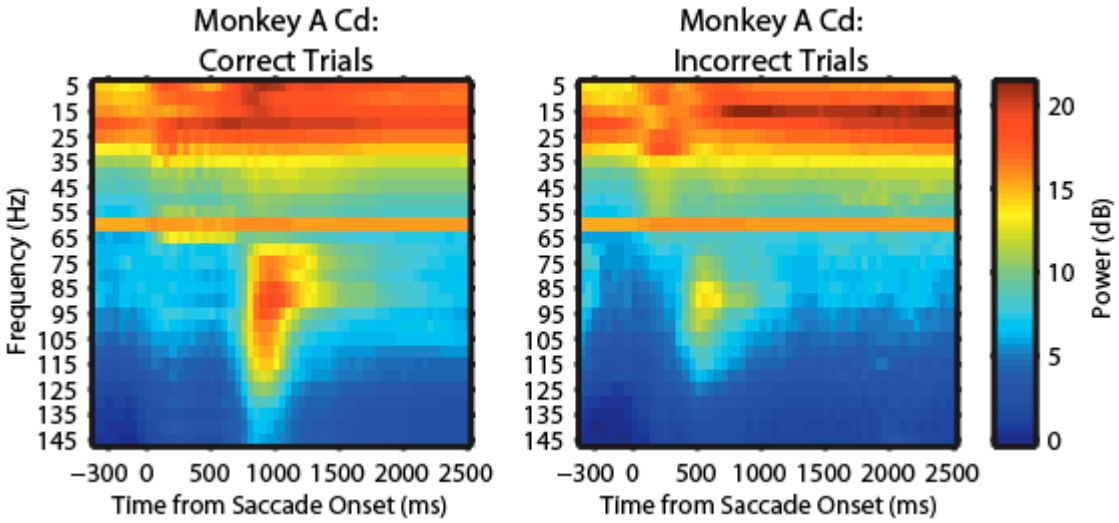

Figure 4.10
The Relationship Between Late-Trial High Frequency Power in Cd and Behavior
The increase in high frequency power present at the end of correct trials (left) is reduced during incorrect trials (right). This result is only apparent in the data from Monkey A, as this power in high frequencies is not observed in Monkey P. Here, data is aligned on saccade onset (0ms).

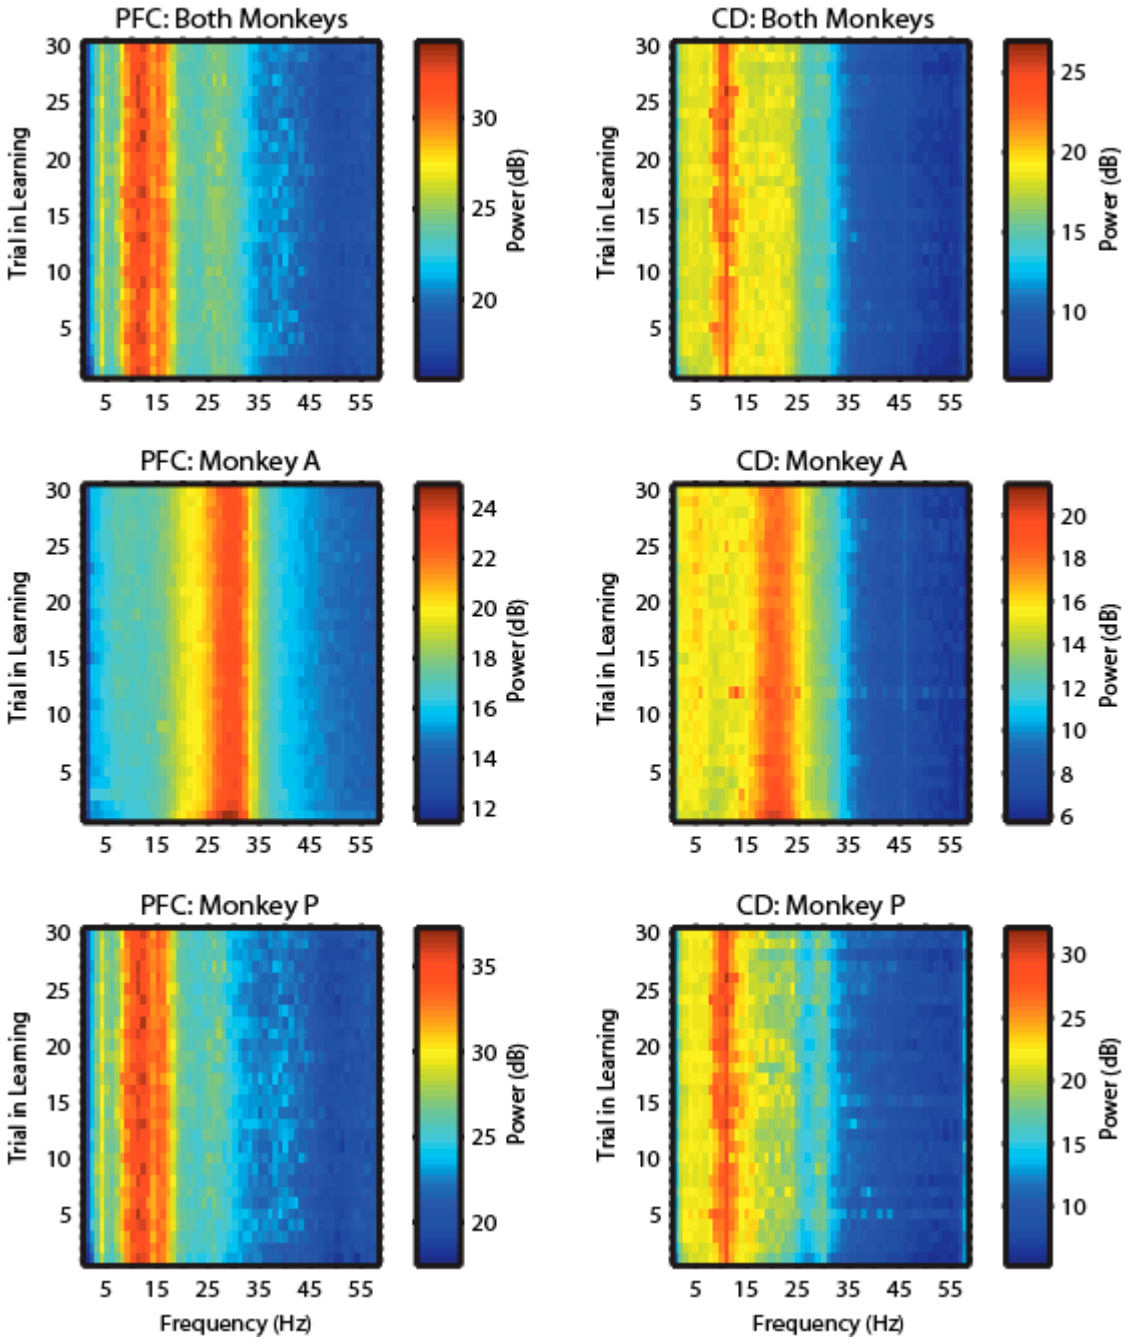

Figure 4.11
Learning-related Changes in Cd and PFC Power Spectrum during the Delay Period
The power spectrum, focused on lower frequencies, for the entire one second delay period (0-1000ms after cue offset) is plotted across trials in learning. The data is presented both combined (top row) and separated by animal (middle and bottom rows). Power spectrum for PFC (left) and Cd (right) are quite similar; however, there are differences seen in the relative frequencies expressed. The only learning-related change in power occurs in the beta frequency range in both PFC and Cd in Monkey A: power in beta decreases as learning progresses.

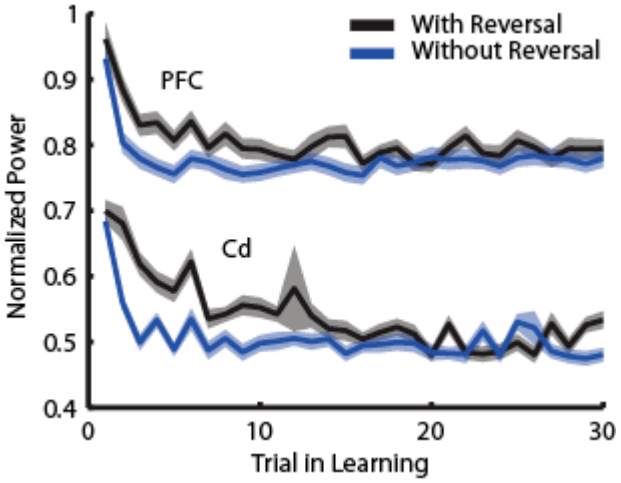

Figure 4.12
Learning-related Decrease in Beta Power during the Delay Period in Monkey A
The decrease in beta (16-34Hz) power is plotted separately for learning with (black) and without (blue) reversals in both PFC (top) and Cd (bottom). The decrease is seen in both areas, and a difference between the two tasks is evident in both areas as well: although power begins the same in both tasks, in the reversal task the decrease takes more trials.

This decrease in beta is broken down by task in Figure 4.12. Because the range of frequencies was slightly different in PFC and Cd, the average power in the 16-34Hz band was used in this analysis to incorporate the bands in both areas. The decrease in beta occurs during learning both with and without reversals (paired ttest of trial 1 versus trial 30, $p<0.05$). However, in the reversal task this power decrease takes more trials. Beta power begins at the same level in both tasks, but power in the reversal task is significantly higher until trial 10 in Cd, with the exception of trial 8 which did not reach significance, and until trial 6 in PFC (paired ttests, $p<0.05$, Bonferroni corrected).

The same analysis was performed on the data during the cue period, from 0-500ms after cue onset (Figure 4.13). Similar to the delay period, a decrease in beta power with learning was found in Monkey A (paired ttest of trial 1 versus trial 30, $p<0.05$), but not in Monkey P. Although this decrease was seen in both PFC and Cd, differences between the two tasks were only observed in Cd (Figure 4.14), where there was higher power in the 16-34Hz band in the reversal task versus the task without reversals (paired ttests, $p<0.05$, Bonferroni corrected. Trials with significant differences: 2, 5, 7-12, 15-17, 22, 24).

Although the power spectrum in the LFPs of PFC and Cd may provide insight into the functioning of the areas on a more macro scale, what is really of importance is understanding the communication between the areas during learning, and how this communication might change during the learning process. To investigate this question, coherence analysis was used. The coherence statistic quantifies the extent that two signals oscillate together: if their phase offset is constant this measure would produce a value of 1, whereas if their phase offset is random coherence would be 0. The assumption here is that there is a higher probability of interaction between two signals if they are oscillating together; thus, higher coherence is interpreted to be greater communication between the signals. Because coherence normalizes for the power in each frequency, the magnitude of power does not effect the coherence calculation. Thus, for this analysis, the data from both animals were compiled together. Although there were slight differences between the animals in which frequencies were most coherent, the overall trends were the same.

There were two main trends observed in the coherence data. First, coherence decreases as learning progresses (paired ttest of trial 1 versus trial 30, $p<0.05$), and second, coherence is

greater during learning with reversals then learning without reversals (paired ttests, $p<0.05$, Bonferroni corrected). Coherence results for the delay period (0-1000ms after cue offset) are presented in Figure 4.15. Coherence was calculated between all signal pairings: between PFC and Cd (PFC-Cd, top row), within PFC (PFC-PFC, middle row), and within Cd (Cd-Cd, bottom row). PFC-Cd coherence in the 16-24Hz band during the delay period decreases with learning for both tasks, and is greater during the task with reversals mostly at the start of learning. PFC-PFC coherence in the 10-22Hz band does not show learning-related decreases during the delay period, but is consistently higher during the task with reversals versus without reversals throughout the learning process. Cd-Cd coherence does not differ between the two tasks, but does show learning-related decreases in coherence in both tasks.

Coherence during the cue period shows similar trends, where coherence is greater for the reversal task than the task without reversals, and decreases in coherence are seen as learning progresses. However, here not only does Cd-Cd coherence not show differences between the two tasks, it also does not decreases with learning. Also in contrast to the delay period, PFC-PFC coherence does show learning-related decreases in both tasks.

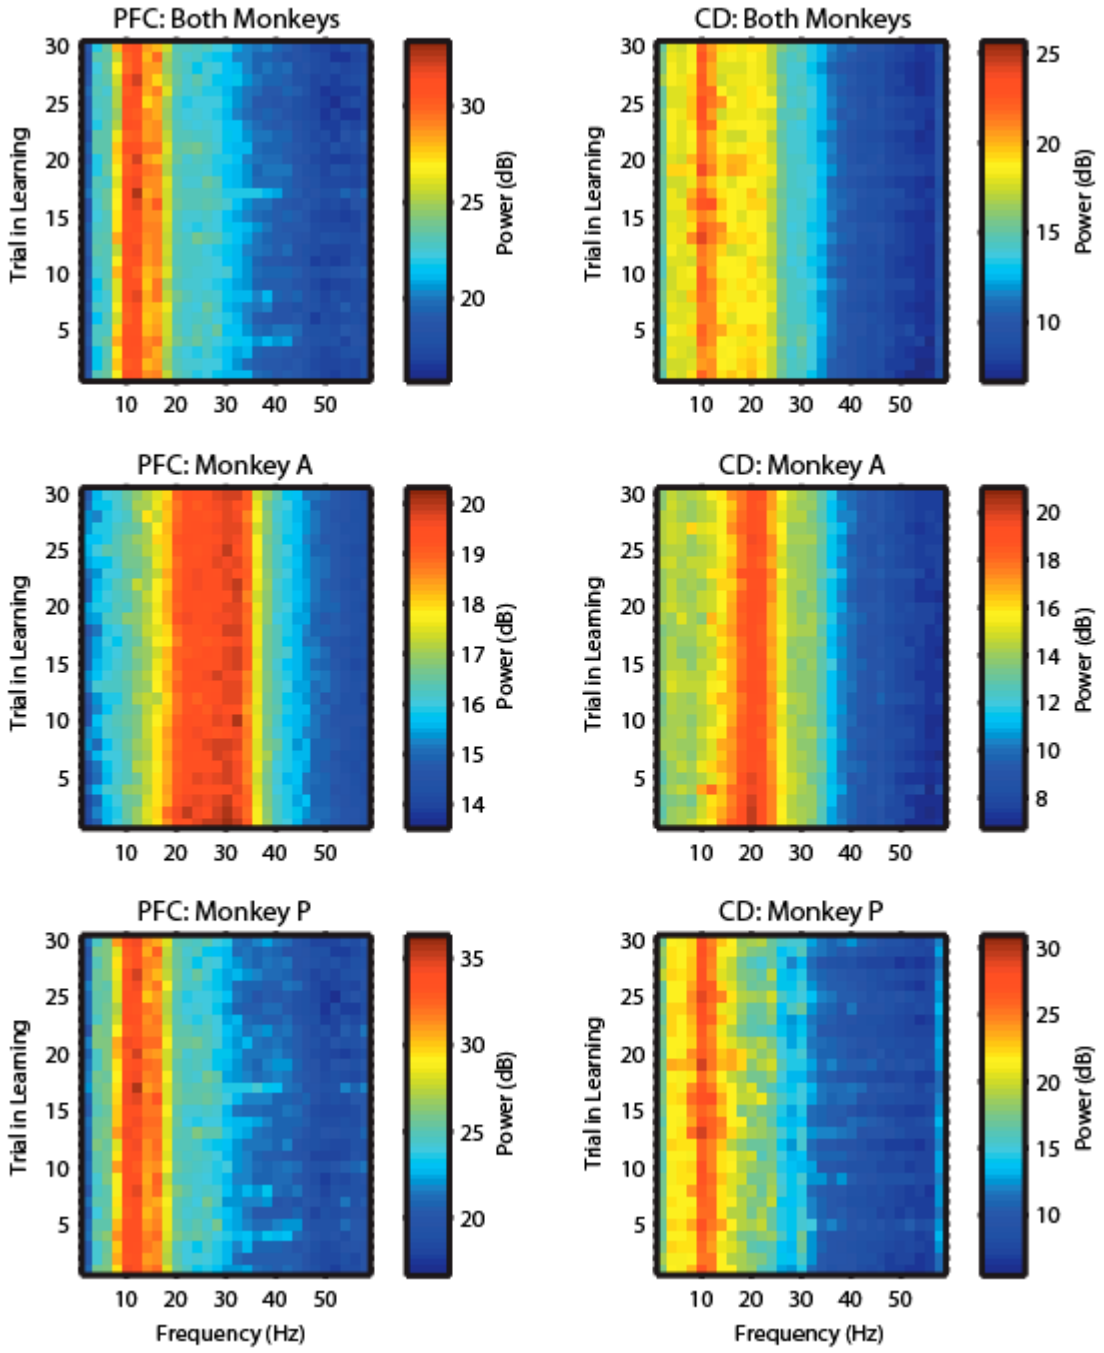

Figure 4.13
Learning-related Changes in Power Spectrum during Cue Period
The power spectrum, focused on lower frequencies, for the cue period (0-500ms after cue onset) is plotted across trials in learning. The data is presented both combined (top row) and separated by animal (middle and bottom rows). Power spectrum for PFC (left) and Cd (right) are quite similar; however, there are differences seen in the relative frequencies expressed. The only learning-related change in power occurs in the beta frequency range in Cd in Monkey A: power in beta decreases as learning progresses.

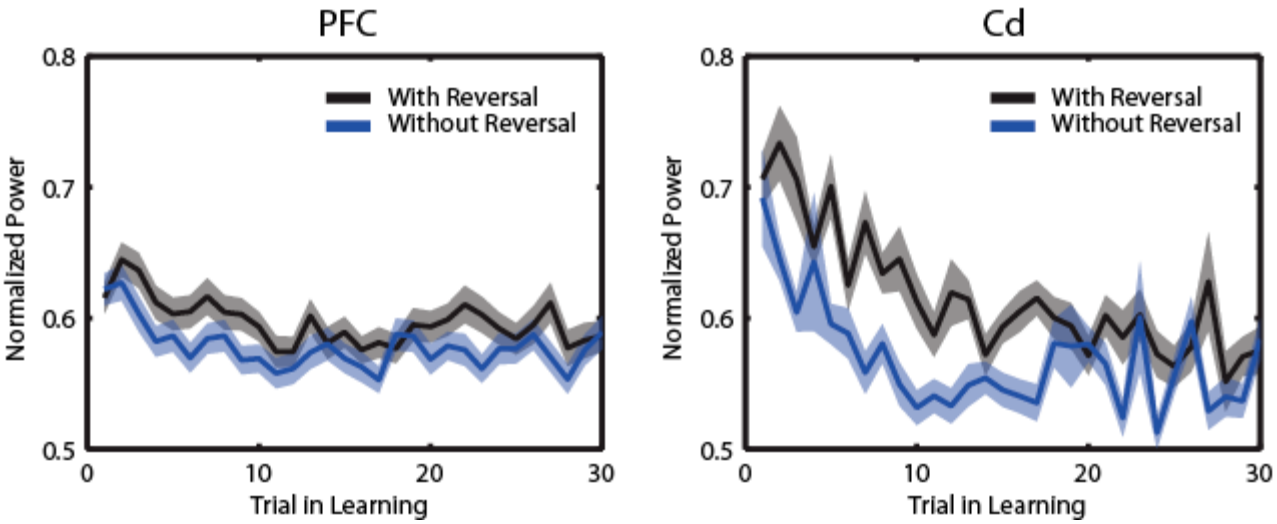

Figure 4.14
Learning-related Changes in Beta Power during the Cue Period in Monkey A
Beta (16-34Hz) power during the cue period is plotted separately for learning with (black) and without (blue) reversals in both PFC (left) and Cd (right). A decrease in beta power is seen in both PFC and Cd, but a difference between the two tasks is only apparent in Cd.

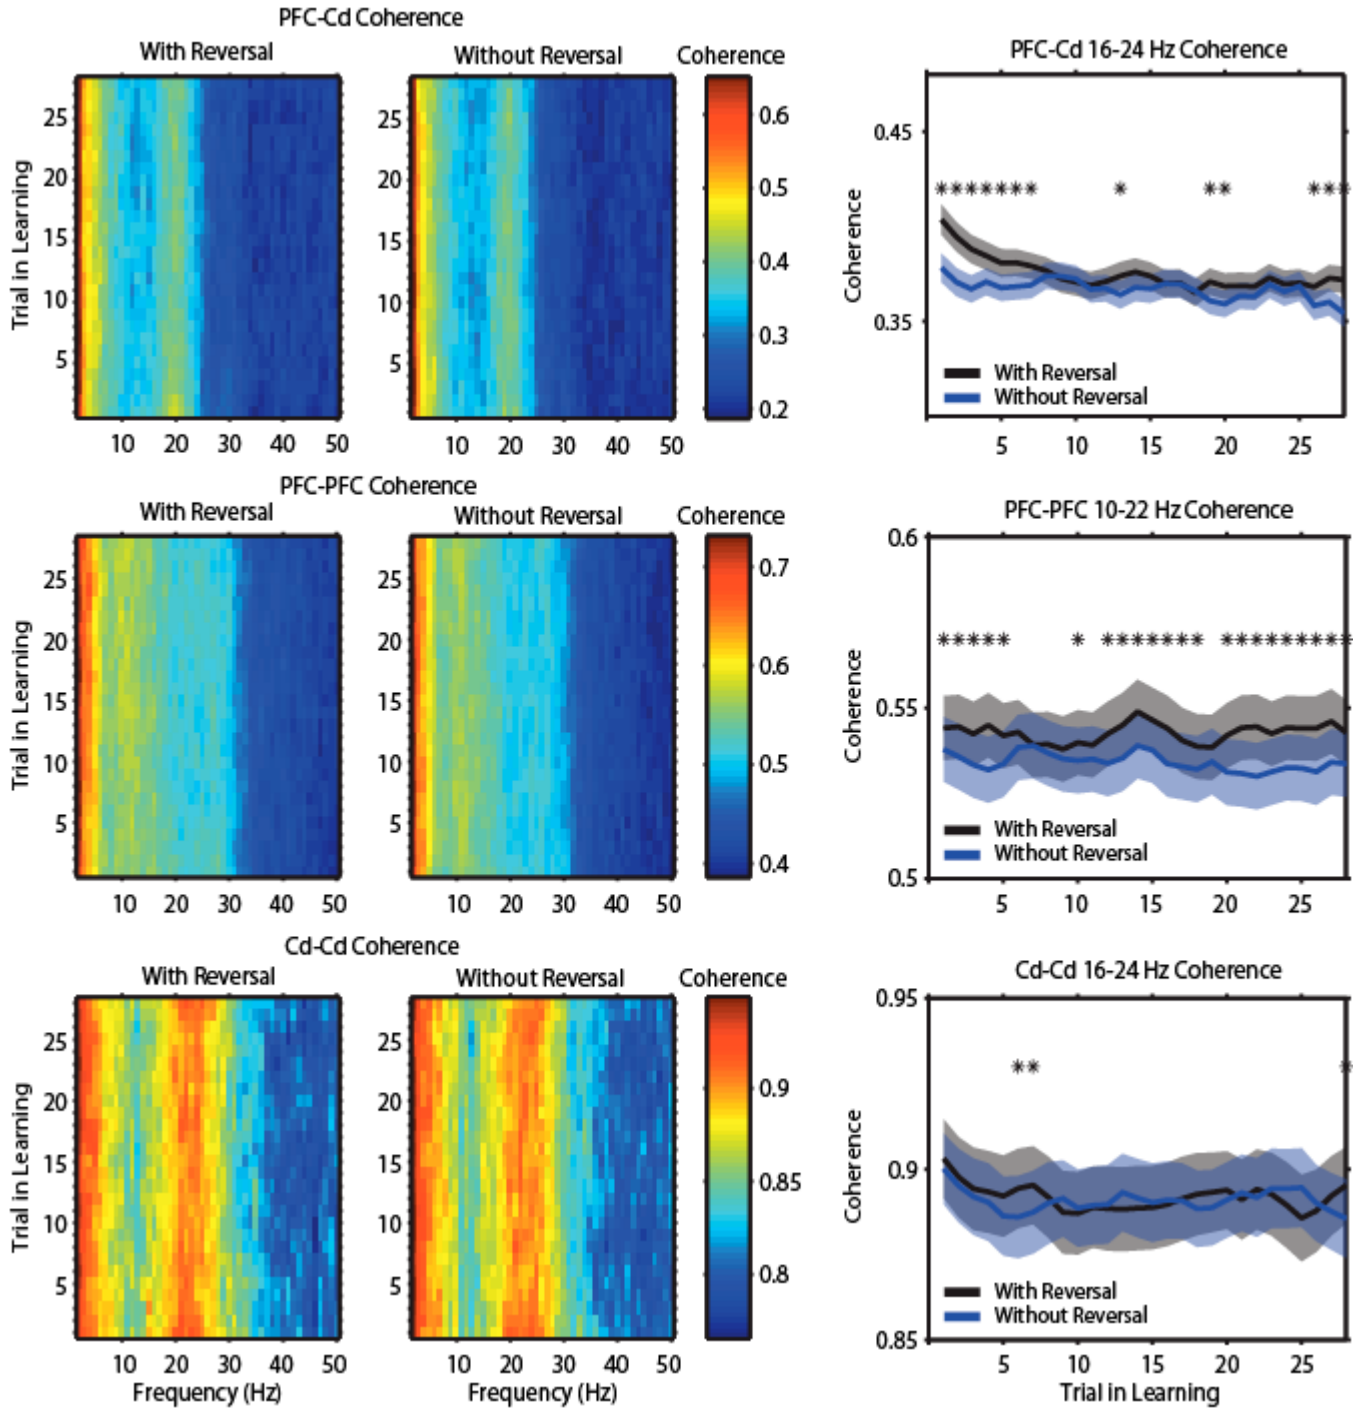

Figure 4.15
Learning-related Changes in Coherence during the Delay Period
Coherence for all frequencies and across learning (colorplots) was calculated on the entire delay period (0-1000ms after cue offset) for all signal pairings: PFC-Cd (top), PFC-PFC (middle), and Cd-Cd(bottom). Two main trends are seen in the coherence data: coherence decreases as learning progresses (except in PFC-PFC interaction), and coherence during the reversal task is higher than coherence during the task without reversals. Asterisks mark significant differences between tasks (paired ttests, $p<0.05$, Bonferroni corrected).

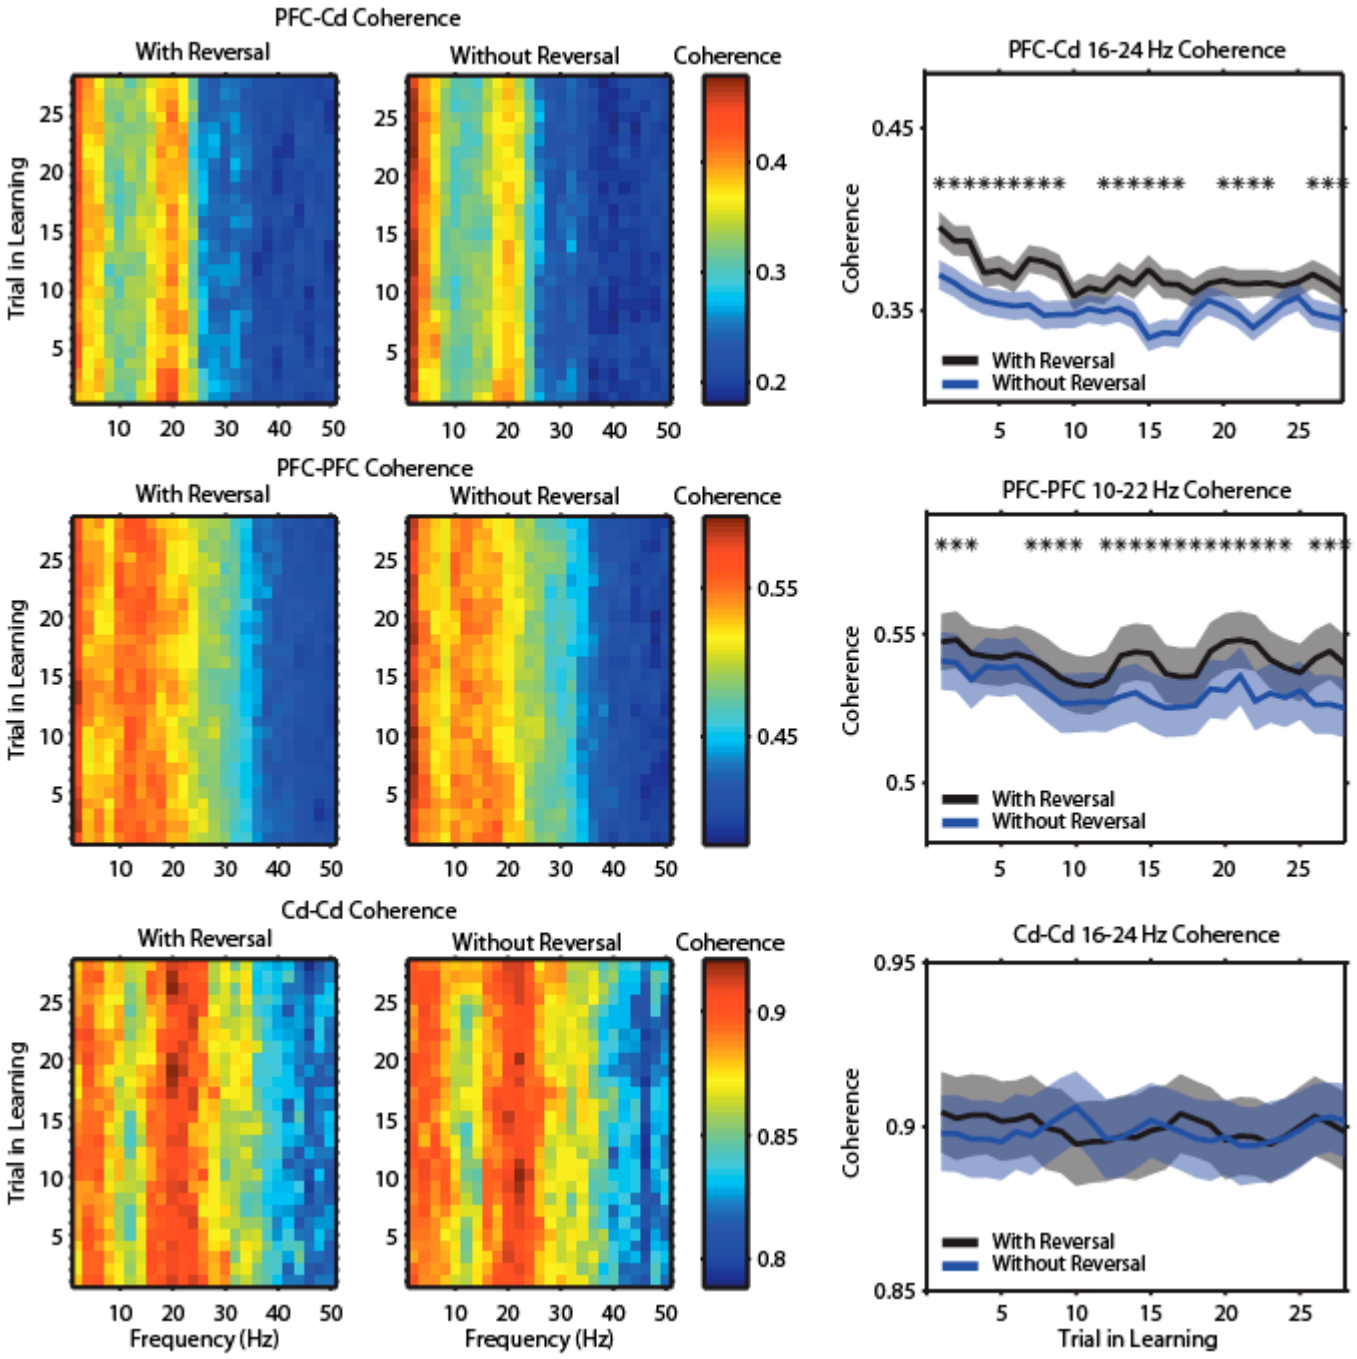

Figure 4.16
Learning-related Changes in Coherence during the Cue Period
Coherence for all frequencies and across learning (colorplots) was calculated on the entire cue period (0-500ms after cue onset) for all signal pairings: PFC-Cd (top), PFC-PFC (middle), and Cd-Cd(bottom). The same two main trends are seen in the coherence data: coherence decreases as learning progresses (except in Cd-Cd interaction), and coherence during the reversal task is higher than coherence during the task without reversals (except in Cd-Cd interaction). Asterisks mark significant differences between tasks (paired ttests, $p<0.05$, Bonferroni corrected).

Discussion

The findings of this study directly comparing learning with and without reversals suggests that the expression of information in and the communication between dlPFC and Cd does underlie learning and differs depending on the specifics of the learning process. PFC and Cd both displayed changes in their expression of task-related information, specifically saccade-direction information, where information appeared earlier in the trial and with greater strength at the start of learning without reversals versus with reversals. This availability of information in these two brain structures suggests that communication between the two areas may not be as crucial during learning without reversals, and this hypothesis was supported by evidence from an analysis of LFPs in PFC and Cd. Both areas showed learning-related decreases in power in the beta band (~16-34Hz), and this decrease took more trials during learning with reversals. This result suggests an increased communication between the areas, which coherence analysis supported, showing that beta coherence (~16-24 Hz) between PFC and Cd is greater during learning with reversals. Coherence was also shown to be higher at the start of learning versus the end of learning in both tasks.

It seems noteworthy that the main frequencies with changes during learning occurred in the beta frequency range. Beta has been implicated in motor preparation (Sanes and Donoghue, 1993) and sensorimotor integration (Murthy and Fetz, 1992), is found in both PFC (Zhang et al., 2008a) and Cd (Courtemanche et al., 2003), and as such seems to be an appropriate candidate for involvement in visuomotor learning. The desynchronization of beta oscillations observed in the current data is also in line with many previous findings that have linked this desynchronization to motor preparation in both cortical and subcortical regions (Zhang et al., 2008b; Kuhn et al., 2004). It is even thought that aberrant beta oscillations may underlie some of the motor deficits observed in diseases such as Parkinson's Disease (Bergman et al., 1988). Recently, it was shown that suppression of beta in the subthalamic nucleus by high frequency stimulation improves motor performance in Parkinson's patients (Kuhn et al., 2008). This provides a more direct link to the function of these oscillations and supports the idea that neuronal oscillations may underlie brain function in very fundamental ways (Singer, 1999). Thus, if we can understand these oscillations and their role in behavior, we may be able to treat diseases more effectively and provide patients with some relief.

Increasingly, the field of neuroscience is moving towards understanding brain function more holistically. In addition to understanding the function and modulation of individual neurons, it is crucial to understand how information is processed within and across brain areas if neuroscience is ever to understand brain function in any depth. Thus, neuronal oscillations have become a prime target of investigation, and their presence throughout the brain, mechanistic explanation, and modulation by behavior are areas of intense research. One common way to measure these oscillations (primarily in animal studies like the present experiment) is the recording of LFPs, which are thought to reflect the synaptic input to the recorded area, in contrast to single or multi-unit recordings which are thought to reflect the neural output of the area (Logothetis et al., 2001; Legatt et al., 1980). LFP recordings have been shown to be modulated by behavior; for example, attention (Fries et al., 2008), movement (Aumann and Fetz, 2004), and memory encoding and retreival (Herrmann et al., 2004; Leiberg et al., 2006). Thus, it will be important for future studies to examine these local oscillations to gain a better understanding of communication between and within brain areas.

Chapter 5: General Discussion

The main goal of this thesis is to provide insight into the function of the frontal cortex-basal ganglia system during arbitrary visuomotor learning. Two tasks were used to study this learning process: learning new associations and learning to reverse previously learned associations. And the differences between the tasks were exploited to investigate underlying neural activity involved in the learning process. Evidence from the studies presented here suggests that the frontal cortex-basal ganglia system underlies the learning process *only* when competition between learning contexts exists.

Two primary structures of the system were investigated in depth: dorsolateral prefrontal cortex (dlPFC) and the caudate nucleus (Cd) of the basal ganglia. Previous work, on which the current thesis is based, from the same laboratory using the same animal subjects, suggested that the basal ganglia may learn an association first, during reversal learning, and transfer this information to dlPFC to aid in the learning process and ultimately guide behavior (Pasupathy and Miller, 2005). When the task was changed to eliminate reversal learning (Chapter 3), the basal ganglia no longer showed functional involvement in the learning process. This suggests that something inherent to the reversal task selectively employs the frontal cortex-basal ganglia system.

There are many differences between learning new associations and learning to reverse previously learned associations. When learning to reverse associations there is a history with the associations that is in competition with learning the reversal. This competition may be thought of as necessitating inhibition or suppression of the previously learned response or a separation of the two learning contexts in some way. In whatever way this competition manifests itself, the fact of additional learning requirements in the reversal process is undeniable, as this difference between the tasks is robustly expressed in the animals' behavior. The reversal task is a much more difficult task for the animals, and this is reflected both in longer reaction times and lower accuracy. Thus, this contextual competition may be the factor guiding the function of the frontal cortex-basal ganglia system in the learning process, and without this competition, as seen in the experiment without reversals in Chapter 3, the basal ganglia need not play a role.

In order to insert this contextual competition back into the behavioral task, while at the same time maintaining the ability to compare the two different learning processes with and without reversals, a new task was designed. This behavioral task alternated between the two learning tasks where every other block was a reversal of the previous block's associations (Chapter 4). Here, Cd again became involved in the learning process, not only during the reversal blocks but also during the learning of new images. This suggests that the basal ganglia might be a key player in encoding the competing contexts.

But encoding the competing contexts is not enough. This information must be communicated in order to control behavior effectively. Results from the analysis of LFPs in dlPFC and Cd presented in Chapter 4 suggests that this communication does occur between the two areas, and that this communication is greater during the learning of reversals, when the contextual competition is the greatest. This communication also decreases as learning progresses supporting the idea that this communication is involved in the transfer of contextual information and also suggesting that communicating the context may place the system in a new state to facilitate the learning process.

This thesis also explored the roles of other members of the frontal cortex-basal ganglia system during the reversal learning process (Chapter 2). FEF was shown to contain task-related information from the start of learning, suggesting that it may be transferring this information to both or either dlPFC and Cd to aid the learning process. GPi was shown to contain more specific information about the learned associations providing evidence for information funneling in the basal ganglia (Alexander et al., 1986). These results suggest that the interplay between many brain areas is responsible for the learning process, and leads one to conclude that the way to truly address questions regarding the function of entire networks of brain areas is in large-scale neural recordings with very fine spatial resolution, as is offered currently by LFP recordings. However, hopefully technology will advance such that invasive methods will be replaced by noninvasive methods with very fine temporal and spatial resolution that allow both an accurate portrayal of the properties of individual neurons as well as the more macroscopic view of the interaction between brain areas. Advances of this kind will revolutionize the field of neuroscience.

Chapter 6: General Methodology

Subjects

For all experiments, two Rhesus monkeys (Macaca mulatta) were used: Monkey A (female, weighing 5 kg) and Monkey P (male, weighing 6-10 kg) [for which of the animals were used for the different experiments see the methods sections of the individual chapters]. All animal procedures conformed to the guidelines of NIH and the MIT Committee on Animal Care. Both animals had titanium head restraints implanted under general anesthesia for use in head immobilization, as well as titanium recording chambers for safe and repeated access to neural tissue. Chambers were positioned using stereotaxic coordinates based on structural MRI scans of each animal. Experiments using only one of the two animals are noted in the methods section for that particular experiment.

Behavioral Tasks

In all behavioral tasks, on each trial a visual image was presented and, after a one-second delay, the animal had to make a saccade to one of two targets that appeared in the periphery (Figure 6.1). If the animal made the correct saccade, a juice reward was given. Each visual stimulus was paired with a single correct saccade direction, and the animals learned this association by trial and error (i.e. no cue signifying the start of a new block of trials was ever presented). Specific criteria were used to determine if learning has occurred: 30 correct trials of each association must have been performed, and 9 of the last 10 trials must have been performed correctly (i.e. 90% correct performance) on each association.

Each learning task differed in its block structure (Figure 6.2). In the serial reversal task (Chapter 2), once learning criteria are met, the animal has to learn associations between the same images and the opposite saccade directions. Thus, in a single session the animal continually reverses associations using the same visual images. However, the task used to collect the FEF data was slightly altered. After every three reversal learning blocks, two delayed-match-to-sample blocks were performed. The analyses presented in this thesis focus only on the reversal blocks of this task. In the task without reversals (Chapter 3), once learning criteria are met the learned images are replaced by new images, and the learned images are never used again. In the task interleaving reversals and new image learning (Chapter 4), each visual image is paired exactly once with each saccade direction and then discarded.

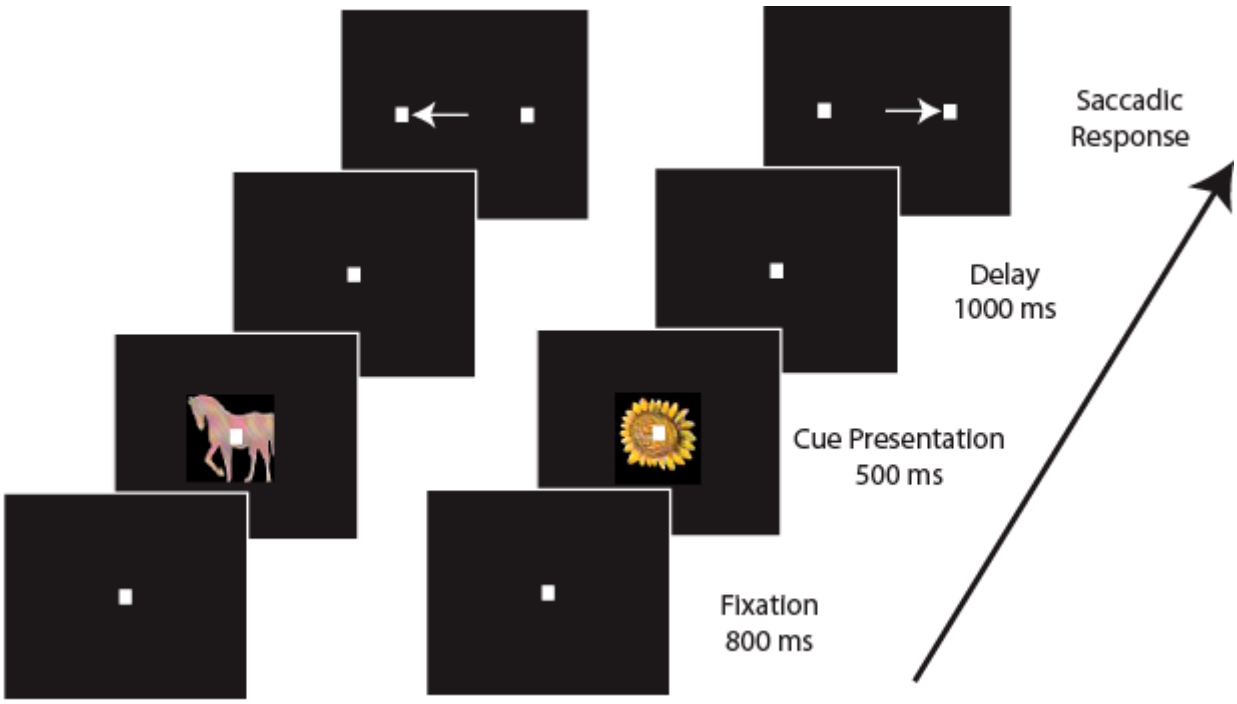

Figure 6.1
Single Trial Events of All Behavioral Tasks
In every trial, the animal was presented with a visual stimulus and, after a one-second delay interval, had to saccade to one of two peripheral targets. The animal was required to fixate a centrally located, white fixation point throughout the trial. The onset of the two peripheral targets occurred simultaneously with the disappearance of the central fixation point, which signaled to the animal to make a response. This figure depicts two different trials: in this example, when the sunflower (right) was presented a saccade to the right was required, whereas when the horse (left) was presented a saccade to the left was required. If the animal executed a saccade in the correct direction, a juice reward was given at the end of the trial. If the animal executed a saccade in the incorrect direction, or did not maintain fixation at any point in the trial, reward was withheld and an increased inter-trial interval was given (as a "timeout").

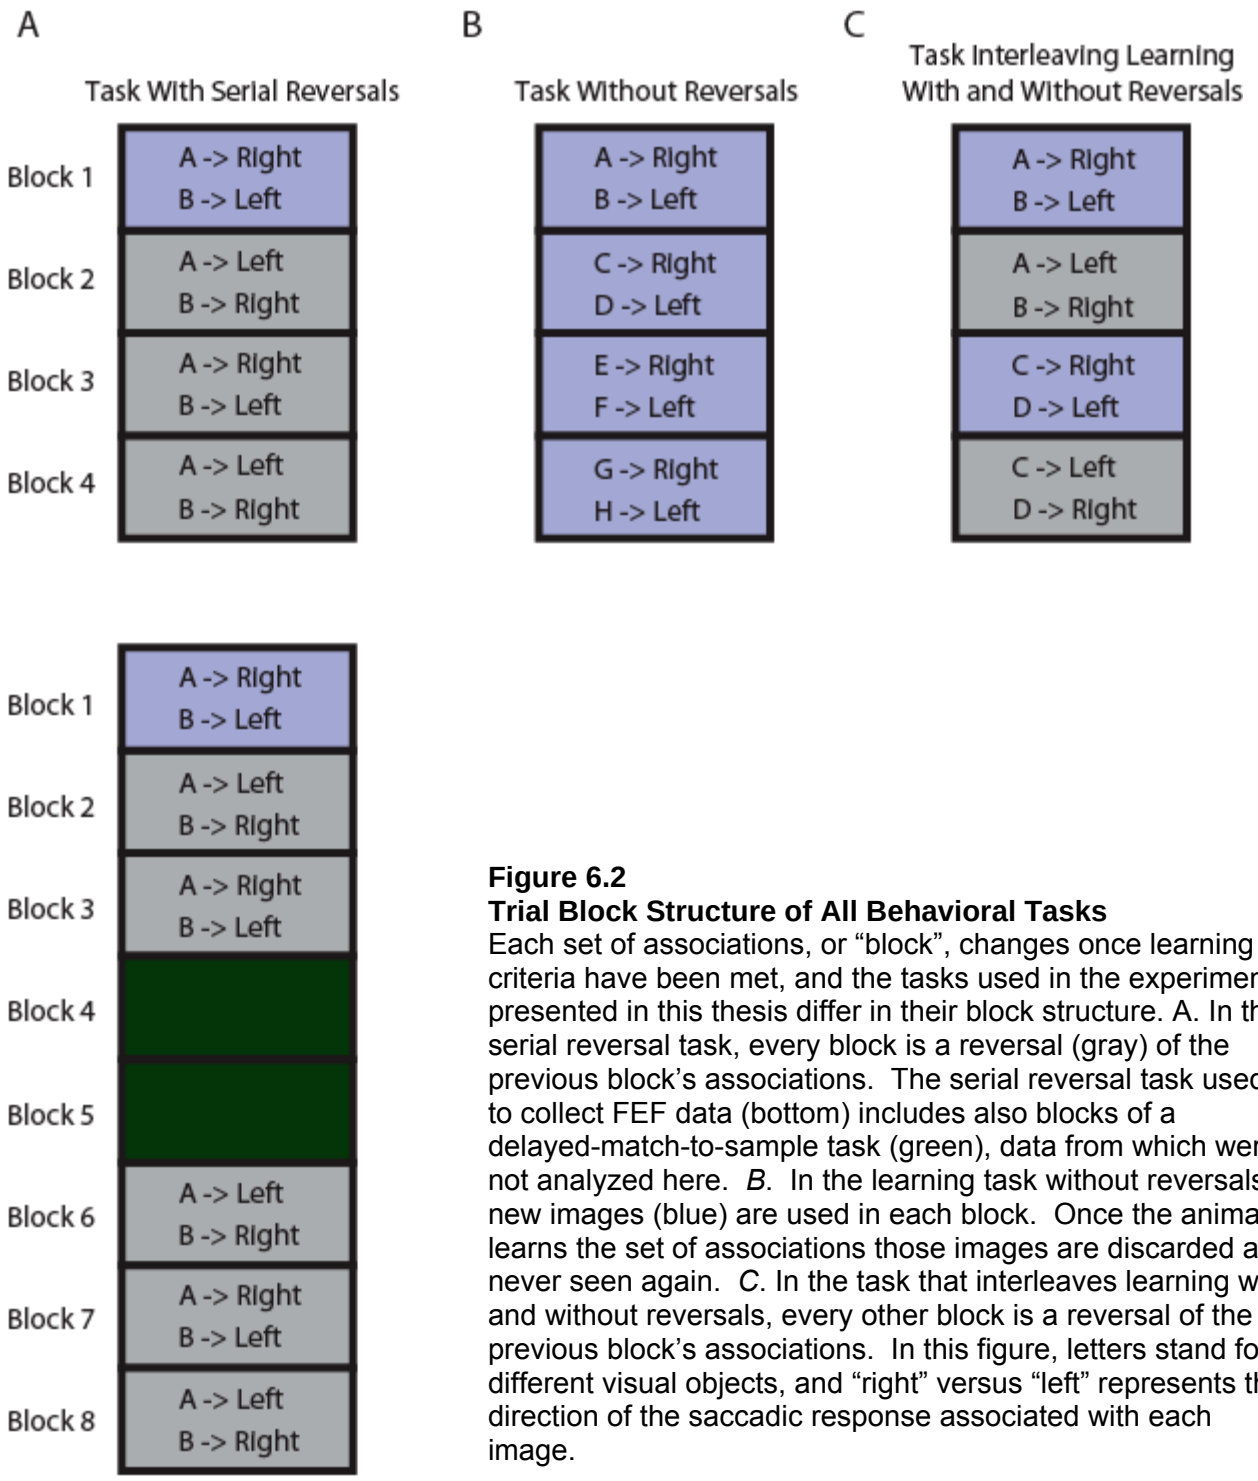

Figure 6.2
Trial Block Structure of All Behavioral Tasks
Each set of associations, or "block", changes once learning criteria have been met, and the tasks used in the experiments presented in this thesis differ in their block structure. A. In the serial reversal task, every block is a reversal (gray) of the previous block's associations. The serial reversal task used to collect FEF data (bottom) includes also blocks of a delayed-match-to-sample task (green), data from which were not analyzed here. *B*. In the learning task without reversals, new images (blue) are used in each block. Once the animal learns the set of associations those images are discarded and never seen again. *C*. In the task that interleaves learning with and without reversals, every other block is a reversal of the previous block's associations. In this figure, letters stand for different visual objects, and "right" versus "left" represents the direction of the saccadic response associated with each image.

All visual images used in the experiments were complex, colored shapes or scenes. They were compiled through image databases (such as Corel) or internet image searches. All images were scaled to be squares with sides of length four degrees of visual angle. Reward given at the end of a trial was titrated according to each animal's motivation through the training process such that the animal would perform enough trials and work at a high enough rate to complete at least four blocks in a session, yet still receive enough liquid each day to maintain health and body weight. Once these reward levels were determined they were kept constant throughout all recording sessions.

Behavioral Monitoring

All behavioral tasks were coded in C and the Cortex (www.cortex.salk.edu) behavioral monitoring system was used to collect behavioral data and control the task display. The animals' eye position was tracked using an infrared system developed by ISCAN (Woburn, MA: www.iscaninc.com) in which a camera monitored the reflections of an infrared beam aimed at the animals' eye to calculate eye position. The animal needed to maintain fixation within a square of 3-4 degrees on a side.

Electrophysiological Techniques

All electrophysiological recordings were performed with epoxy fully coated tungsten microelectrodes from FHC (Bowdoin, ME) with 250 μm diameter. Since dura was left in-tact, electrodes had to be thick and stiff enough to penetrate the dura without bending or losing all impedance. Electrodes were lowered using custom built microdrive arrays that required the manual rotation of screws to change the depth of the electrode. For a small fraction of the deeper recordings performed (GPi and Cd), a hydraulic microdrive was used (either from FHC or Plexon, Inc). All voltage signals were collected using the data aquisition system from Plexon, Inc. Both spike waveforms and local field potentials were recorded from every electrode. Spike waveforms were captured using an online threshold, but were sorted more accurately offline using Offline Sorter (Plexon, Inc.). Standard spike sorting methods were used including but not limited to: principal components analysis, measures of energy, and specifics of waveform shape such as peak and trough levels and waveform duration.

Recording locations were determined using structural MRIs of each animal. Well placement was maximized for access to relevant brain areas and also for stabile placement on the skull.

The recording wells were cylindrical with a 19mm diameter. All PFC recordings were concentrated in the dorsolateral portion of prefrontal cortex which lies dorsal (medial) to the principal sulcus. Cd recordings were concentrated in the head and body and primarily in the dorsal portions of the structure since recording electrodes were lowered from the vertical.

For Monkey A, both PFC and Cd recordings were performed in the right hemisphere. The PFC well was placed 29.6 mm anterior to the interaural line and 16mm lateral to the midline, and the Cd well was placed 16.3mm anterior to the interaural line and 0.8mm lateral with no angle from the vertical. For Monkey P, FEF right hemisphere recordings were performed from a well centered 18.2mm anterior and 17mm lateral. The right hemisphere GPi well was centered 12.5 mm anterior and 24.3 mm lateral at a 42 degree angle from the horizontal. The PFC well used to collect data for Chapter 3 was centered 26.7mm anterior and 19.1 mm lateral on the right hemisphere. The PFC well used to collect data for Chapter 4 was centered 28.2mm anterior and 24.6 mm lateral on the left hemisphere. Monkey P's Cd well was centered 14 mm anterior and 1.7 mm lateral and recordings targeted the right hemisphere.

Analytical Techniques

Cell and Trial Inclusion

Not all data that was collected could be used in analysis due to many factors. First, all cells must have been properly isolated using standard cluster cutting techniques. Second, all included cells had to contain at least four full learning blocks of data. This criterion ensures enough trials are present to obtain an accurate account of the behavior of the cell. To minimize bias in analysis, trials in learning blocks not completed were not analyzed. For example, if a cell was recorded throughout the entire behavioral session, trials in the last learning block were discarded if the block was not completed.

Identifiying Cells with Task-Related Information

In order to identify cells that contained task-related information, an ANOVA was performed on the average firing rate of each cell in all included trials broken down into four trial time epochs: 100-600ms after cue onset (“Cue”), 600-1500ms after cue onset (“Delay”), 150ms before to 150ms after saccade initiation (“Saccade”), and 50-300ms after the start of reward delivery (“Reward”). Thus, the average firing rate of each trial was defined by four single numbers, one number (firing rate) for each of the four time epochs. Then within each time epoch, an ANOVA

was performed to extract the differences in firing rate between the different trial types. The specifics of this ANOVA varied depending on task. For example, in the serial reversal task used in Chapter 2, a 2-way ANOVA was used with object and direction as factors, and in the learning task without reversals (Chapter 3) a 1-way ANOVA was performed with saccade direction as the only factor.

A cell was considered to contain task-related information if it *only* had a significant main effect of the task factor. For example, a cell was considered to contain information about saccade direction if it had a significant main effect of direction (and *not* a significant main effect of object or a significant interaction if a 2-way ANOVA was performed). For the results from the 2-way ANOVAs, a cell was considered to contain object-saccade association information if it had a main effect of object *and* direction *or* a significant interaction term. Significance level was set at alpha = 0.05, and was corrected for multiple comparisons using Bonferroni's correction.

In the task interleaving learning with and without reversals used in the experiment presented in Chapter 4, a sliding receiver operating characteristic (ROC) analysis on the firing rate of left versus right trials was performed to determine direction-selective cells. The area under the ROC curve (which here I'm just terming ROC for short) is a measure of the distance between two distributions. The simplest way to explain and understand this analysis is to consider ROC to be the probability that an ideal observer would be able to categorize each data point in the two distributions correctly. So if the two distributions are nonoverlapping ROC has a value of 1 (the ideal observer would be correct every time), and if two distributions completely overlap ROC has a value of 0.5 (chance probability level for the ideal observer since there are two choices). The ROC value is calculated based on the true positive categorizations versus the false positive categorizations. It is the area under the curve plotting these two variables.

<u>Learning-related Changes in Normalized Firing Rate and Saccade Direction Information</u>

The focus of analyses in this thesis, in general, is how the recorded cell populations are changing as the animal learns the visuomotor associations. Two measures of activity were used here: normalized firing rate, and direction selectivity. Both measures were computed in a 100ms bin slid across the trial in 10ms steps, and each trial bin in learning contained 8 trials of each association in each learning block. The trial bins were aligned on the first trial in the bin, so the first trial bin used trials 1-8 of each association in each block of that session, and the

second trial bin used trials 2-9 of each association in each block of that session. Thus, since the minimum number of trials per association in a block is 30, 23 trial bins were used to ensure an equal number of trials in all trial bins. This explains why all colorplots have a maximum y-value of 23.

Firing rate (FR) was computed for each trial, according to the details above, and once the two-dimensional FR matrix was computed (with dimensions of time in trial and trial in learning), normalization was performed for each individual cell. Normalized firing rate was computed as follows:

$$[\text{(FR)} - \text{(Minimum FR)}] / [\text{(Maximum FR)} - \text{(Minimum FR)}]$$

This normalization method scales the firing rate on a 0-1 scale, where 0 represents the minimum firing rate and 1 represents the maximum firing rate of that particular cell. After each cell's FR was normalized, averages were computed across the direction-selective cell populations revealed by the ANOVA described above.

Direction selectivity was quantified as the fraction of explainable variance by the direction factor in an ANOVA. The variance (Var) in the firing rate of each cell was partitioned according to the partitions of an ANOVA. For example, in a 2-way ANOVA with object and direction as factors, the partitioning would be done as follows:

$$\text{Total Var} = \text{Var}_{dir} + \text{Var}_{obj} + \text{Var}_{int} + \text{Var}_{err}$$

where Var_{dir} is the variance explained by the direction factor, Var_{obj} is the variance explained by the object factor, Var_{int} is the variance explained by the interaction of the two factors, and Var_{err} is the unexplained variance (or variance due to error).

The raw fraction of explained variance for direction was calculated as:

$$\text{Raw FEV}_{dir} = \text{Var}_{dir} / \text{Total Var}$$

This quantity was then normalized in two ways. First, the mean (computed across both time and trial bins) fraction of explained variance by direction in the baseline fixation period (from 450-200ms before cue onset) was subtracted out.

$$\text{FEV}_{dir} = (\text{Raw FEV}_{dir}) - (\text{Mean Raw FEV}_{dir} \text{ during baseline period})$$

Then, it was divided by the the maximum fraction of explainable variance (computed across all time and trial bins):

$$FEV_{dir} = (FEV_{dir}) / (\text{Maximum of } (1 - Var_{err}))$$

This normalization procedure was used for a few reasons. First, the differences in firing rate between cells will impact the amount of explainable variance; this is particularly seen in cells with low firing rates, since the dearth of spikes places the explainable variance at extremes. Second, this procedure scales each cell so that comparisons of the changes in direction information can be quantitatively compiled (through averaging) without biasing the population in favor of certain cells with very high explained variance.

Risetime and Peak Selectivity Calculations

Because directly comparing the three dimensional direction selectivity data presented in the colorplots is extremely difficult, two linear measures were used. A measure of the timing of information was used, termed "risetime", and a measure of the strength, or amount, of information was calculated. Risetime was defined as the time in each trial bin in learning where half-maximum selectivity was reached. The strength of information was defined as the peak (i.e. maximum) selectivity in each trial bin. It is important to note that the purpose of these measures is to pull out the trends observed in the colorplots of direction selectivity. Thus, these linear measures were calculated in time windows encompassing the band of information being studied. These linear measures were fit with either sigmoid, exponential, or linear functions depending on the particular trend of the data. The method of least squares was used to determine the optimal parameters for the significant fits of each function.

Abbreviations

PFC:	prefrontal cortex
dlPFC:	dorsolateral prefrontal cortex
Cd:	caudate nucleus
GPi:	internal segment of the globus pallidus
FEF:	frontal eye field
BG:	basal ganglia
FC:	frontal cortex
FEV:	fraction of explained variance
FEV_{dir}:	fraction of explained variance that is explained by the direction factor
FR:	firing rate
FR_{norm}:	normalized firing rate
LFP:	local field potential

References

Alexander GE, DeLong MR, Strick PL. (1986). Parallel orgnaization of functionally segregated circuits linking basal ganglia and cortex. *Ann. Rev. Neurosci*., 9, 357-381.

Aosaki T, Graybiel AM, Kimura M. (1994). Effect of nigrostriatal dopamine system on acquired neural responses in the striatum of behaving monkeys. Science, 265, 412-415.

Arkadir D, Morris G, Vaadia E, Bergman H. (2004). Independent coding of movement direction and reward prediction by single pallidal neurons. *J. Neurosci*., 24 (45), 10047-10056.

Asaad WF, Rainer G, Miller EK. (1998). Neural activity in the primate prefrontal cortex during associative learning. *Neuron*, 21, 1399-1407.

Aumann TD, Fetz EE. (2004). Oscillatory activity in forelimb muscles of behaving monkeys evoked by microstimulation in the cerebellar nuclei. *Neurosci. Letters*, 361, 106-110.

Bellebaum C, Koch B, Schwarz M, Daum I. (2008). Focal basal ganglia lesions are associated with impairments in reward-based reversal learning. *Brain,* 131, 829-841.

Bergman H, Feingold A, Nini A, Raz A, Slovin H, Abeles M, Vaadia E. (1988). Physiological aspects of information processing in the basal ganglia of normal and parkinsonian primates. *Trends In Neurosci*., 21, 32-38.

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

Cahusac PMB, Rolls ET, Miyashita Y, Niki H. (1993) Modification of the responses of hippocampal neurons in the monkey during the learning of a conditional spatial response task. *Hippocampus*, 3 (1), 29-42.

Chen LL, Wise SP. (1995a). Neuronal activity in the supplementary eye field during acquisition of conditional oculomotor associations. *J. Neurophys*., 73 (3), 1101-1121.

Chen LL, Wise SP. (1995b). Supplementary eye field contrasted with the frontal eye field during acquisition of conditional oculomotor associations. J. Neurophys., 73 (3), 1122-1134.

Coe B, Tomihara K, Matsuzawa M, Hikosaka O. (2002). Visual and anticipatory bias in three cortical eye fields of the monkey during an adaptive decision-making task. J. Neurosci., 22 (12), 5081-5090.

Courtemanche R, Fujii N, Graybiel AM. (2003). Synchronous, focally modulated beta oscillations characterize local field potential activity in the striatum of awake behaving monkeys. *J. Neurosci*., 23 (37), 11741-11752.

Daw ND, Niv Y, Dayan P. (2005). Uncertainty-based competition between prefrontal and dorsolateral striatal systems for behavioral control. *Nature Neuroscience*, 8, 1704-1711.

DeLong MR. (1971). Activity of pallidal neurons during movement. *J. Neurophys*, 34 (3), 414-427.

DeLong MR, Crutcher MD, Georgopoulos AP. (1985). Primate globus pallidus and subthalamic nucleus: functional organization. *J. Neurophys*., 53 (2), 530-543.

Eblen F, Graybiel AM. (1995). Highly restricted origin of prefrontal cortical inputs to striosomes in the macaque monkey. *J. Neurosci*., 15 (9), 5999-6013.

El Massioui N, Cheruel F, Faure A, Conde F. (2007). Learning and memory dissociation in rats with lesions to the subthalamic nucleus or to the dorsal striatum. *Neurosci*., 147, 906-918.

Evarts EV. (1968). Relation of pyramidal tract activity to force exerted during voluntary movement. *J. Neurophys*., 31 (1), 14-27.

Fawcett AP, Moro E, Lang AE, Lozano AM, Hutchinson WD. (2005). Pallidal deep brain stimulation influences both reflexive and voluntary saccades in huntington's disease. *Movement Disorders*, 20 (3), 371-386.

Freedman DJ, Riesenhuber M, Poggio T, Miller EK. (2001). Categorical representation of visual stimuli in the primate prefrontal cortex. *Science*, 291, 312-316.

Fries P, Womelsdorf T, Oostenveld R, Desimone R. (2008). The effects of visual stimulation and selective visual attention on rhythmic neuronal synchronization in macaque area V4. *J. Neurosci.*, 28 (18), 4823-4835.

Funahashi S, Bruce CJ, Goldman-Rakic PS. (1989). Mnemonic coding of visual space in the monkey's dorsolateral prefrontal cortex. *J. Neurophysiol*., 61, 331-349.

Funahashi S, Chafee MV, Goldman-Rakic PS. (1993). Prefrontal neuronal activity in rhesus monkeys performing a delayed anti-saccade task. *Nature*, 365, 753-756.

Fuster JM, Alexander GE. (1970). Delayed response deficity by cryogenic depression of frontal cortex. *Brain Research*, 20, 85-90.

Fuster JM, Alexander GE. (1971). Neuron activity related to short-term memory. *Science*, 173, 652-654.

Fuster JM. (1973). Unit activity in prefrontal cortex during delayed-response performance: neuronal correlates of transient memory. *J. Neurophysiol*., 36, 61-78.

Fuster JM. (2000). Executive frontal functions. Exp. Brain Res., 133, 66-70.

Georgopoulos AP, Kalaska JF, Caminiti R, Massey JT. (1982). On the relations between the direction of two-dimensional arm movements and cell discharge in primate motor cortex. *J. Neurosci*., 2 (11), 1527-1537.

Goldman-Rakic PS. (1996). The prefrontal landscape: implications of functional architecture for understanding human mentation and the central executive. *Phil. Trans. R. Soc. Lond. B*, 351, 1445-1453.

Graybiel AM, Ragsdale CW (Jr.). (1978). Histochemically distinct compartments in the striatum of human, monkey, and cat demonstrated by acetylthicholinesterase staining. *Proc. Natl. Acad. Sci*., 75 (11), 5723-5726.

Graybiel AM, Aosaki T, Flaherty AW, Kimura M. (1994). The basal ganglia and adaptive motor control. *Science*, 265, 1826-1831.

Graybiel AM. (1998). The basal ganglia and chunking of action repertoires. *Neurobio. of Learning and Memory*, 70, 119-36.

Graybiel AM, Saka E. (2004). The basal ganglia and the control of action. In MS Gazzaniga (Ed.), *The New Cognitive Neurosciences*, 3rd Ed. Cambridge, MA: MIT Press, 495-510.

Harlow JM. (1848). Passage of an iron rod through the head. *Boston Medical and Surgical Journal*, 39, 389–393.

Herrmann CS, Munk MHJ, Engel AK. (2004). Cognitive functions of gamma-band activity: memory match and utilization. *Trends in Cog. Sci*., 8, 347-355.

Hikosaka O, Wurtz RH. (1983a). Visual and oculomotor functions of monkey substantia nigra pars reticulata. I. Relation of visual and auditory responses to saccades. J. Neurophys., 49 (5), 1230-1253.

Hikosaka O, Wurtz RH. (1983b). Visual and oculomotor functions of monkey substantia nigra pars reticulata. II. Visual responses related to fixation of gaze. J. Neurophys., 49 (5), 1254-1267.

Hikosaka O, Wurtz RH. (1983c). Visual and oculomotor functions of monkey substantia nigra pars reticulata. III. Memory-contingent visual and saccade responses. J. Neurophys., 49 (5), 1285-1301.

Hikosaka O, Wurtz RH. (1983d). Visual and oculomotor functions of monkey substantia nigra pars reticulata. IV. Relation of substantia nigra to superior colliculus. J. Neurophys., 49 (5), 1230-1253.

Houk JC, Wise SP. (1995). Distributed modular architectures linking basal ganglia, cerebellum, and cerebral cortex: their role in planning and controlling action. *Cerebral Cortex*, 2, 95-110.

Inase M, Li B-M, Takashima I, Iijima T. (2001). Pallidal activity is involved in visuomotor association learning in monkeys. *Euro. J. Neurosci*., 14, 897-901.

Jog MS, Kubota Y, Connolly CI, Hillegaart V, Graybiel AM. (1999). Building neural representations of habits. *Science*, 286, 1745-1749.

Kawagoe R, Takikawa Y, Hikosaka O. (1998). Expectation of reward modulates cognitive signals in the basal ganglia. *Nature Neurosci*., 1 (5), 411-416.

Krauzlis RJ. (2005). The control of voluntary eye movements: new perspectives. The Neuroscientist, 11 (2), 124-137.

Kuhn AA, Williams D, Kupsch A, Limousin P, Hariz M, Schneider GH, Yarrow K, Brown P. (2004). Event-related beta desynchronization in human subthalamic nucleus correlates with motor performance. *Brain*, 127, 735-746.

Kuhn AA, Kempf F, Brucke C, Gaynor Doyle L, Martinez-Torres I, Pogosyan A, Trottenberg T, Kupsch A, Schneider GH, Hariz MI, Vandenberghe W, Nuttin B, Brown P. (2008). High frequency stimulation of the subthalamic nucleus suppresses oscillatory beta activity in patients with parkinson's disease in parallel with improvement in motor performance. *J. Neurosci*., 28 (24), 6165-6173.

Lasker AG, Zee DS. (1997). Ocular motor abnormalities in huntington's disease. *Vision Res*., 37 (24), 3639-3645.

Lebedev MA, Messinger A, Kralik JD, Wise SP. (2004). Representation of attended versus remembered locations in prefrontal cortex. *PLOS Biology*, 2 (11), 1919-1935.

Lee B, Groman S, London ED, Jentsch JD. (2007). Dopamine D_2/D_3 receptors play a specific role in the reversal of a learned visual discrimination in monkeys. *Neuropsychopharm*., 32, 2125-2134.

Legatt AD, Arezzo J, Vaughan HG. (1980). Averaged multiple unit activity as an estimate of phasic changes in local neuronal activity: effects of volume-conducted potentials. *J. Neurosci. Methods*, 2, 203-217.

Leiberg S, Lutzenberger W, Kaiser J. (2006). Effects of memory load on cortical oscillatory activity during auditory pattern working memory. *Brain Research*, 1120, 131-140.

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

McDonald RJ, Foong N, Ray C, Rizos Z, Hong NS. (2007). The role of medial prefrontal cortex in context-specific inhibition during reversal learning of a visual discrimination. *Exp. Brain Res*., 177, 509-519.

Middleton FA, Strick PL. (2000a). Basal Ganglia Output and Cognition: Evidence from Anatomical, Behavioral, and Clinical Studies. *Brain and Cognition,* 42, 183-200.

Middleton FA, Strick PL. (2000b). Basal ganglia and cerebellar loops: motor and cognitive circuits. *Brain Res. Rev*., 31, 236-250.

Middleton FA, Strick PL. (2002). Basal-ganglia 'projections' to the prefrontal cortex of the primate. *Cerebral Cortex*, 12, 926-935.

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

Miller EK, Buschman TJ. (2007). Bootstrapping your brain: how interactions between the frontal cortex and basal ganglia may produce organized actions and lofty thoughts. In Neurobiology of Learning and Memory, 2nd Ed. Kaster RP and Martinez JL (eds), Elsevier, Inc.

Milner B. (1963). Effects of different brain lesions on card sorting. *Archives of Neurology*, 9, 90-100.

Mink JW. (1996). The basal ganglia: focused selection and inhibition of competing motor programs. *Prog. Neurobio.*, 50, 381-425.

Murray EA, Bussey TJ, Wise SP. (2000). Role of prefrontal cortex in a network for arbitrary visuomotor mapping. *Exp. Brain Res.*, 133, 114-129.

Murthy VN, Fetz EE. (1992). Coherent 25- to 35-Hz oscillations in the sensorimotor cortex of awake behaving monkeys. *Proc. Nat. Acad. of Sci.*, 89, 5670-5674.

Mushiake H, Inase M, Tanji J. (1991). Neuronal activity in the primate premotor, supplementary, and precentral motor cortex during visually guided and internally determined sequential movements. *J. Neurophys.*, 66 (3), 705-718.

O'Neill M, Brown VJ. (2007). The effect of striatal dopamine depletion and the adenosine A_{2A} antagonist KW-6002 on reversal learning in rats. *Neurobio. of Learn. and Memory*, 88, 75-81.

Owen AM, Roberts AC, Polkey CE, Sahakian BJ, Robbins TW. (1991). Extra-dimensional versus intra-dimensional set shifting performance following frontal lobe excisions, temporal lobe excisions or amygdala-hippocampectomy in man. *Neuropsychologia,* 29, 993-1006.

Owen AM, James M, Leigh PN, Summers BA, Marsden CD, Quinn NP, Lange KW, Robbins TW. (1992). Fronto-striatal cognitive deficits at different stages of parkinson's disease. *Brain*, 115, 1727-1751.

Owen AM, Roberts AC, Hodges JR, Summers BA, Polkey CE, Robbins TW. (1993). Contrasting mechanisms of impaired attentional set-shifting in patients with frontal lobe damage or parkinson's disease. *Brain*, 116, 1159-1175.

Packard MG, Knowlton BJ. (2002). Learning and Memory Functions of the Basal Ganglia. *Ann. Rev. Neurosci.*, 25, 563-593.

Palencia CA, Ragozzino ME. (2006). The effect of n-methyl-d-aspartate receptor blockage on acetylcholine efflux in the dorsomedial striatum during response reversal learning. *Neurosci.*, 143, 671-678.

Passingham RE. (1993). The frontal lobes and voluntary action. Oxford: Oxford University Press.

Pasupathy A, Miller EK. (2005). Different time courses of learning-related activity in the prefrontal cortex and striatum. *Nature,* 433, 873-876.

Petrides M. (1982). Motor conditional associative-learning after selective prefrontal lesions in the monkey. *Beh. Brain Res*., 5, 407-413.

Petrides M. (1985a). Deficits in non-spatial conditional associative learning after periarcuate lesions in the monkey. *Beh. Brain Res*., 16, 95-101.

Petrides M. (1985b). Deficits on conditional associative-learning tasks after frontal- and temporal-lobe lesions in man. *Neuropsychologia*, 23 (5), 601-614.

Petrides M. (1997). Visuo-motor conditional associative learning after frontal and temporal lesions in the human brain. *Neuropsychologia*, 35 (7), 989-997.

Poldrack RA, Clark J, Pare-Blagoev EJ, Shohamy D, Creso Moyano J, Myers C, Gluck MA. (2001). Interactive memory systems in the human brain. *Nature*, 414, 546-550.

Reynolds JNJ, Hyland BI, Wickens JR. (2001). A cellular mechanism of reward-related learning. *Nature*, 413, 67-70.

Robbins TW. (1996). Dissociating executive functions of the prefrontal cortex. *Phil. Trans. R. Soc. Lond. B.*, 351, 1463-1471.

Robbins TW. (2007). Shifting and stopping: fronto-striatal substrates, neurochemical modulation and clinical implications. *Phil. Trans. R. Soc. B*., 362, 917-932.

Romanski LM. (2004). Domain specificity in primate prefrontal cortex. *Cog., Affect., Beh. Neurosci*., 4 (4), 421-429.

Rowe JB, Toni I, Josephs O, Frackowiak RSJ, Passingham RE. (2000). The prefrontal cortex: response selection or maintenance within working memory? *Science*, 288, 1656-1660.

Saint-Cyr JA. (2003). Frontal-striatal circuit functions: context, sequence, and consequence. *J. Internat. Neuropsych. Soc*., 9, 103-28.

Sanes JN, Donoghue JP. (1993). Oscillations in local field potentials of the primate motor cortex during voluntary movement. *Proc. Nat. Acad. of Sci.*, 90, 4471-4474.

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

Schultz W. (2002). Getting formal with dopamine and reward . *Neuron*, 36, 241-263.

Singer W. (1999). Neuronal Synchrony: A Versatile Code for the Definition of Relations? *Neuron*, 24, 49-65.

Tehovnik EJ, Sommer MA, Chou I-H, Slocum WM, Schiller PH. (2000). Eye fields in the frontal lobes of primates. *Brain Res. Reviews*, 32, 413-448.

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

Toni I, Rushworth MFS, Passingham RE. (2001a). Neural correlates of visuomotor associations: spatial rules compared with arbitrary rules. Exp. Brain Res., 141, 359-369.

Toni I, Ramnani N, Josephs O, Ashburner J, Passingham RE. (2001b). Learning arbitrary visuomotor associations: temporal dynamics of brain activity. NeuroImage, 14, 1048-1057.

Turner RS, Anderson ME. (1997). Pallidal discharge related to the kinematics of reaching movements in two dimensions. *J. Neurophys*., 77 (3), 1051-1074.

Turner RS, Anderson ME. (2005). Context-dependent modulation of movement-related discharge in the primate globus pallidus. J. Neurosci., 25 (11), 2965-2976.

Ungerleider LG. (1995). Functional brain imaging studies of cortical mechanisms for memory. *Science*, 270, 769-775.

Wallis JD, Anderson KC, Miller EK. (2001). Single Neurons in Prefrontal Cortex Encode Abstract Rules. *Nature*, 411, 953-956.

White NM, McDonald RJ. (2002). Multiple parallel memory systems in the brain of the rat. *Neurobiology of Learning and Memory*, 77, 125-184.

Wilson CJ. (2004). Basal ganglia. In GM Shephard (Ed.), The Synaptic Organization of the Brain, 5th Ed., Oxford: Oxford University Press, 361-413.

Wilson FAW, O Scalaidhe SP, Goldman-Rakic PS. (1993). Dissociation of object and spatial processing domains in primate prefrontal cortex. *Science*, 260, 1955-1958.

Wirth S, Yanike M, Frank LM, Smith AC, Brown EN, Suzuki WA. (2003). Single neurons in the monkey hippocampus and learning of new associations. *Science*, 300, 1578-1581.

Wise SP, Murray EA, Gerfen CR. (1996a). The frontal cortex-basal ganglia system in primates. *Crit. Rev. Neurobio*., 10, 317-356.

Wise SP, di Pellegrino G, Boussaoud D. (1996b). The premotor cortex and nonstandard sensorimotor mapping. *Can. J. Physiol. Pharmacol*., 74, 469-482.

Zhang Y, Wang X, Bressler SL, Chen Y, Ding M. (2008a). Prestimulus cortical activity is correlated with speed of visuomotor processing. *J. Cog. Neurosci*., 20 (10), 1915-1925.

Zhang Y, Chen Y, Bressler SL, Ding M. (2008b). Response preparation and inhibition: the role of the cortical sensorimotor beta rhythm. *Neurosci*., 156, 238-246.