Exploration and exploitation of action selection in the motor cortex and basal ganglia

by

Gregory J. Gage

A dissertation submitted in partial fulfillment
of the requirements for the degree of
Doctor of Philosophy
(Biomedical Engineering)
in The University of Michigan
2010

Doctoral Committee:

Assistant Professor Joshua D. Berke, Co-Chair
Associate Professor Daryl R. Kipke, Co-Chair
Professor Emeritus David J. Anderson
Associate Professor Edward L. Ionides

Dedicated to :
My Father

ACKNOWLEDGEMENTS

I would like to thank my committee members, specifically Daryl Kipke for his guidance, generosity, and mentorship on how to be a leader both in academics and business. Josh Berke for his commitment to detail, and teaching me the value of being skeptical, especially of your own results.

I wish to thank my family for their support during the time I spent working on this dissertation, both professionally and personally. To my Mom and Dad, for helping me through the toughest of times. To Sarah Banks for giving me the courage to start this new challenge. To my sisters Karen, Kelley and Cathy (and to my brother-in-law Bill) for always being there for me, and for keeping me in check when needed. And to my niece and nephews for all the great memories, and reminding me of why I love Michigan.

Much of the things one learns in a graduate training program, is via the interchange between labmates and colleagues. For this I am eternally grateful to members of the Kipke Lab and Berke Lab. Tim Marzullo for readiness to provide honest feedback. Kip Ludwig, Hirak Parikh, and Azadeh Yazmen for our many intellectual disagreements and constructive conversations over the years. And also I wish to thank Matt Johnson, John Seymour, Matt Gibson, Taneev Escamilla, Rachel Maraini, Nick Langhals, Paras Patel and T.K. Kozai for the great memories of the NEL. Alex Wiltschko, Colin Stoetzner, Dan Leventhal, Jeff Pettibone, and Vaughn Hetrek for their friendship and all of the support over the years in East Hall.

Thanks also to Tonya Brown, Vera Williams, Karen Coulter, Mary Mohrbach, Susan Douglas and the rest of the University of Michigan's administrative staff for all their help.

Finally, I wish to acknowledge Dr. Matt O'Donnell, the former chair of the University of Michigan BME department and current dean of the College of Engineering at the University of Washington. I initially was rejected by the admissions board when I applied to the UM BME program; however, Matt looked at my application personally. He noticed I was a non-traditional grad student coming out of industry with a strong engineering skills, and decided to overturn the decision and take a chance on me. For this, I am eternally grateful.

TABLE OF CONTENTS

LIST OF FIGURES

Figure

LIST OF TABLES

Table

ABSTRACT

The basal ganglia (BG) have been proposed as a possible neural substrate for action selection in the vertebrate brain. This hypothesis has been developed primarily through pathological observations. Human neurological disorders of the basal ganglia can result in movements that are slowed or eliminated (bradykinesia/akinesia in Parkinson's Disease) or conversely, uncontrolled or unwanted (e.g. in Huntington's Disease and Tourette Syndrome). However, the precise mechanisms by which BG circuits influence behavior remain to be understood. In this thesis, I have focused on determining the role of BG circuits in selection of well-trained actions, and how these findings can be applied for use in neuroprosthetic devices.

In the first study, I investigated one proposed mechanism to help resolve competition between actions in the BG: feedforward inhibition of striatal medium spiny cells (MSNs) by fast-spiking interneurons (FSIs). I recorded single unit activity from presumed MSNs and FSIs together with motor cortex and globus pallidus (GP), in rats performing a simple choice task. My findings support the idea that FSIs contribute to action selection processes within striatal microcircuits, but suggest that the feedback pathway from GP to FSIs may be particularly important for the suppression of highly trained yet unwanted actions.

In my second study, I examined the role of large neuronal ensembles of the BG and motor cortex during two variations on a simple action selection task. Analysis of local field potential (LFP) oscillations revealed that ~20Hz rhythms (β_{20}) were prominent during the hold period, but only if subjects were instructed on which direction to move during the hold period. This finding is consistent with the hypothesis that β_{20}

is involved with withholding specific selected actions, and agrees with pathological observations of increased β_{20} in Parkinson's Disease.

In the third study, I examined how action selection circuitry can be exploited to aid in the development of a neuroprosthetic system. This system is one solution to a long standing problem in neural engineering: central nervous system neurons do not regenerate after traumatic injury, and can lead to paralysis. By bypassing injured neurons, we can allow for direct motor control from non-injured neurons. I developed an algorithm that observes the pattern of activity in cortical ensembles and allows both the subjects and control system to co-adapt their behavior to allow naïve rats to use a neuroprosthetic device. The results of this study show that subjects can learn to select discrete actions in real-time using the neural activity of the cortex.

By developing a deeper understanding of the mechanisms behind selecting motor actions, we will provide further insight into such neurological diseases as Parkinson's Disease or Tourette Syndrome. In this thesis, I investigate action selection at the single-unit and multi-unit levels, while studying neural ensembles both within and across brain structures. Further knowledge in this field will also yield more sophisticated, yet more natural control of neuroprosthetic devices which will rely on native BG and cortical roles in action selection.

CHAPTER I

Introduction

> The chief function of the body is to carry the brain around.
>
> *Thomas Edison*: Inventor, Scientist, Businessman, fellow Michigander

The brain is one of least understood entities in science. Packed into this 1.4kg organ of fatty tissue lies the hopes, dreams, and desires that make us human. The brain provides us the ability to coordinate movements, generate speech, internalize thoughts, and solve complex problems. We are just beginning to understand the brain, and the plethora of functions that it serves. While Mr. Edison may be correct, one could argue that the chief function of the brain is to move the body. The brain survives (as do we) by generating movements that avoid predators, gather food, and eventually reproduce. Intentional communication using motor actuators in the tongue, lips and jaw allow us to repair failed communication, teach or persuade others, and build shared goals that help us to live in community. Motor output is the brain's only communication to the outside world, and it is most readily available for scientific understanding.

1.1 Overview

The chief goal of this dissertation is to explore the underlying mechanisms behind the output of brain's decision making process: action selection. I will investigate

the decision-making ability using micro-scale (cell-to-cell), aggregate-scale (neurual ensembles), and macro-scale (brain region-to-brain region) levels. By understanding the neural mechanisms of action-selection, next-generation neuroprosthetics can be developed that exploit the brains innate ability to quickly switch between actions. Towards this end, I employee engineering principles to develop a control paradigm that allows for arbitrary output states to be selected quickly using electrophysiological signals.

1.2 The Action Selection Problem

The basal ganglia were first implicated in motor actions in 1912 when S. A. Kinnier Wilson described a disease with pathological changes in the liver and basal ganglia that was characterized by tremors, muscular rigidity, and weakness (Wilson, 1912). Wilson made the observation that these motor deficits occurred in the absence of damage to the corticospinal ("pyramidal") motor system, and therefore were the result of an "extrapyramidal" motor system dysfunction. He postulated that the basal ganglia comprised a large portion of this extrapyramidal motor system and that it was independent of the pyramidal system. Wilson developed monkey models of basal ganglia impairment (Wilson, 1914), and would eventually hypothesize that the extrapyramidal motor system was evolutionarily "old" and used primarily for automatic, non-modifiable movements; while the "new" corticospinal system was responsible for voluntary and modifiable movements (Wilson, 1928)[1].

Further evidence for the basal ganglia's involvement in movements came with the advent of modern anatomical techniques. Researchers found that a vast majority of the basal ganglia output went via thalamus to motor cortical areas (for review,

[1] Coincidentally, this paper was the first neurology paper to use the term "déjà vu" when describing an experience of familiarity (Eadie, 1998).

see: Nauta and Mehler, 1966). And while other output pathways have been found that are directed to the more-cognitive frontal lobes (Middleton and Strick, 1994) and brain stem (Parent and Bellefeuille, 1982), the motor control centers appear to dominate in volume, and will be the focus of this dissertation.

When discussing motor control, we need to address the "action selection problem" that arises whenever two or more competing motor programs seek simultaneous access to a restricted resource (Redgrave et al., 1999). Effective behavior requires that the conflicts between "selected" actions are resolved uniquely from the many possible actions, and that these disputes are resolved rapidly. An eloquent example of this conflict resolution is described by Jonathan Mink (1996) in his review of action selection:

> Consider as an example the act of reaching to pick an apple from a tree. If asked to describe the movement, most observers would describe the reaching movement. Yet, during the reach, multiple other motor mechanisms act together to maintain the upright posture of the rest of the body. Prior to the reach, these mechanisms were also active in the reaching arm to maintain its posture. When the arm reaches toward the apple the postural mechanisms must be turned off selectively in the arm while they remain active in the rest of the body. When the reach is completed, the reaching motor pattern generators must be turned off and the postural mechanisms must be turned back on. If the competing posture-holding and reaching mechanisms were inappropriately active at the same time, the result might be instability of posture or slowing of movement, or both.

A number of different architectures have been proposed to deal with the selection problem in motor output of a specific action. One attractive way to implement

action selection is to have all motor plans inhibit other motor plans with competing resources (Fukai and Tanaka, 1997). This recurrent reciprocal inhibition can support a "winner-take-all" functionality making only one action viable at a time. Over time, the strengths of incoming excitatory links, and of the inhibitory links between competitors, can be tuned to support a complex patterns such that resource allocation among the competitors can be optimized. However, this flexibility comes at a high cost. For n competing actions, a fully connected cortical network with reciprocal inhibition requires $n(n-1)$ connections (assuming one neuron for every action) and adding a new competitor costs a further $2n$ connections (McFarland, 1965). Another candidate architecture is a centralized selection mechanism that has overall behavioral control of the output motor system. Snaith and Holland (1991) demonstrated that McFarland's n competing actions requires only two connections for each competitor, resulting in a total of $2n$ connections $[<< n(n-1)]$ when an architecture with centralized selection is used. Since evolution should normally prefer architectures which achieve comparable functionality with fewer connections (at lower levels of activity), the centralized selection mechanism has been posited to mediate competing actions. The basal ganglia, due to the recurrent connections to the motor cortices, are hypothesized to provide the brain with a specialized, central selection mechanism to resolve conflict between competing systems at different functional levels (Redgrave et al., 1999).

1.3 Possible mechanisms for Action Selection in the BG

Clues as to how the BG may function to serve a conflict resolution role come from both anatomical and physiological studies. As the main input structure of the basal ganglia, the striatum receives massive excitatory glutamatergic inputs from all

regions of the cortex (Gerfen and Wilson, 1996; Voorn et al., 2004). This cortical information is integrated and "compressed" by the striatum which in turn generates activities that eventually disinhibit premotor nuclei in the thalamus and brainstem. The striatum is almost entirely composed of inhibitory GABAergic neurons. The vast majority of these, at least 95%, are medium-sized spiny projection neurons which also serve as the only source of output from the nucleus (Wilson, 2004). The remaining cell types are comprised of large aspiny cholinergic interneurons, and at least 3 distinct types of GABAergic interneurons.

To avoid conflicts, the striatum should suppress (inhibit) all but strong, focally converging cortical inputs. There are two potential sources of the GABAergic inhibition in striatum: feedforward inhibition from the GABAergic interneurons and feedback inhibition from the axon collaterals of the spiny neurons themselves. MSNs were found to create a large number of GABAergic synapses onto nearby MSNs (Bolam et al., 1983). This lead to the suggestions that MSNs form a mutual inhibitory network in which each output neuron makes symmetric reciprocal inhibitory synapses onto its neighbors. With such a configuration, action selection could be achieved via "winner-take-all" interactions. However, the axon collaterals do not appear to have a noticeable effect on nearby cells /emin vivo/em. Jaeger et al. (1994) used sharp electrode intracellular recordings, but were unable to detect collateral inhibitory postsynaptic potentials (IPSPs) resulting from antidromic activation of neostriatal spiny neurons in vivo or in vitro. They also were unable to detect synaptic interactions between simultaneously recorded pairs of nearby spiny neurons in striatal slices, even when their axonal and dendritic fields were seen to overlap.

The lack of lateral inhibition of spiny cells has increased interest in the contributions of other elements of intrastriatal circuitry. Striatal parvalbumin-containing

fast-spiking interneurons (FSIs) receive cortical inputs, and form strong perisomatic GABAergic synapses onto hundreds of surrounding MSNs (Bennett and Bolam, 1994). This FSI-mediated feed-forward inhibition has been argued to make an important contribution to action selection and execution, via the broadly tuned, distributed suppression of MSNs representing unwanted actions (Kita et al., 1990; Parthasarathy and Graybiel, 1997). In Chapter 2, I will investigate the role that FSIs play during the selection of (and consequently, the suppression of other) actions.

1.4 Oscillations in the Basal Ganglia

The action potentials of single cells are not the only signal captured by microelectrodes placed in the brain. Local Field Potentials (LFP), which correspond to the lower frequency range of the extracellular signal, can carry information about a large number of neurons. Historically, it is the LFP that preceded all other electophysiological signals in neuroscience. On July 6, 1924 Hans Berger was able for the first time to record, with a small Edelmann galvanometer, the electrical brain potentials directly from the cortical surface in a 17-year-old young man undergoing an operation for a suspected brain tumor. The results of this early work would be published 5 years later (Berger, 1929), after Berger painstakingly re-examined his data and obtain further recordings to verify his results[2]. The first human electroencephalographic (EEG) pattern described was a smooth, regular electrical oscillation at $8-12$ Hz, which Berger called an alpha wave. This early work was followed by a barrage of intensive clinical and basic research. From scalp recordings, investigators identified various further oscillatory patterns that were particularly obvious during a resting state and while asleep, and discovered that these oscillations would change as the subjects became conscious and awake.

[2]Including 73 electroencephalograms from his own son, aged 15-17 years, and 56 from himself.

It was clear from even these early studies that neural oscillations are dependent on the brain state and brain region. In more recent years, biophysical studies have revealed that even single neurons are endowed with complex dynamics, including their intrinsic abilities to resonate and oscillate at multiple frequencies (Llinás, 1988; Hutcheon and Yarom, 2000). This finding suggested that the precise timing of cell activity within neuronal networks could represent information. Around the same time, the neuronal assembly structures of the oscillatory patterns found during sleep were related to the experiences of the previous awake period (Wilson and McNaughton, 1994). These results led to the idea that perception, memory, and even consciousness could be the result from synchronized networks (Engel et al., 2001; Varela et al., 2001).

If synchronized activity is required for information processing, then diseases that eliminate or exaggerate oscillation mechanisms could have devastating effects. Such is the case with Parkinson's Disease (PD). Recordings in humans as a result of functional neurosurgery have revealed a tendency for basal ganglia neurons to oscillate and synchronize their activity, giving rise to a rhythmic population activity, manifested as oscillatory local field potentials in the beta band (β, 13-30 Hz) (Kühn et al., 2005). These exaggerated oscillations in PD are reduced with therapeutic interventions that restore motor abilities (Brown et al., 2001; Levy et al., 2002; Kühn et al., 2006). Accordingly synchronization in the beta band has been hypothesized to be essentially antikinetic in nature and pathophysiologically relevant to bradykinesia.

The lack of motor output in PD can be put in terms of an inability to perform action selection. In Chapter 3, I will investigate the role that β activity plays during the suppression of a cued actions. I will compare this to a similar antikinetic "idle", or waiting, state to determine if β-band activity is associated specifically with general

absence of action, or (as I hypothesize) the suppression of a specific selected action plan.

1.5 Exploiting Action Selection in Motor Neuroprosthetics

Research on Neuroprosthetics (NPs), also known as Brain-Machine Interfaces, has seen a revival over the last decade, and these efforts can roughly be divided into four categories (Sanchez et al., 2009). First is the most common NP, the sensory NP, which substitutes for natural sensory inputs such as visual (Chelvanayagam et al., 2008; Zrenner, 2002) or auditory (Rouger et al., 2007). More than 120,000 people have been implanted worldwide with cochlear implants. Motor NPs (Serruya et al., 2002; Taylor et al., 2002) substitute for parts of the body to convey intent of motion to prosthetic limbs; while cognitive NPs that repair communication between brain areas such as the hippocampus (Berger et al., 2001) that mediates short term to long term memories. Finally, there are clinical NPs that stimulate specific brain areas to repair normal function, such as deep brain stimulation for Parkinsons disease (Lozano and Mahant, 2004) or to avoid or abort epileptic seizures (Ludvig et al., 2005).

Motor neuroprosthetics have have improved steadily since the pioneering work of Fetz (1969) that showed it was possible to operantly condition single neurons, and by Chapin et al. (1999) that demonstrated a real-time, one-dimensional, goal-driven NPs that extracted intended movements using electrodes in the brain. Within a few years, this would be extended into two-and three dimensional control (Serruya et al., 2002; Taylor et al., 2002; Carmena et al., 2003; Shenoy et al., 2003).

I wished to expand this body of work by developing a motor NP that would exploit the cortical ability to perform action selection. In chapter 4, I demonstrate a novel cortical control paradigm that allows the motor neurons to develop their own

internal strategy to achieve discrete output states. These output states are decoded in real-time to allow feedback to modify the brain's behavior on a trial-by-trial basis.

1.6 Dissertation Organization

This dissertation covers three components of action selection: the role that striatal fast spiking interneurons play during the selection of actions (chapter 2), β_{20} oscillations during the withholding of selected actions (chapter 3), and a neuroprostetic control model to select discrete actions (chapter 4). At the time of this dissertation completion, chapter 2 has been submitted and is in the review process, chapter 3 is in preparation for manuscript submission, and chapter 4 has been published in the Journal of Biomedical Engineering (Gage et al., 2005).

CHAPTER II

Selective Activation of Striatal Fast Spiking Interneurons during Choice Execution

Abstract

Basal ganglia circuits are thought to participate in decision-making by facilitating one option while suppressing alternatives. One proposed mechanism to help resolve competition between actions is feedforward inhibition of striatal medium spiny cells (MSNs) by fast-spiking interneurons (FSIs). To test this, we recorded presumed MSNs and FSIs together with motor cortex and globus pallidus (GP), in rats performing a simple choice task. MSN activity was broadly distributed across the task sequence, especially near rewards. By contrast, FSIs were preferentially active as subjects initiated their chosen action, in conjunction with a sharp decrease in GP activity. Both MSNs and FSIs were movement-selective, but neighboring MSNs and FSIs showed opposite selectivity. Our findings support the idea that FSIs contribute to action selection processes within striatal microcircuits, but suggest that the feedback pathway from GP to FSIs may be particularly important for the suppression of highly trained yet unwanted actions.

2.1 Introduction

Effective action requires unequivocal decision-making—the prompt and appropriate selection of one alternative while others are suppressed. The neurobiological mechanisms of choice behavior have been intensely investigated, and many brain areas appear to contribute [for review, see (Gold and Shadlen, 2007)]. However, the basal ganglia (BG) are thought to play an especially critical role, for a wide range of decision types (Redgrave et al., 1999; Samejima and Doya, 2007). Human neurological disorders of the basal ganglia can result in movements that are slowed or eliminated (bradykinesia/akinesia in Parkinson's Disease) or conversely, uncontrolled or unwanted (e.g. in Huntington's Disease and Tourette Syndrome). Despite much experimental and theoretical progress (Lo and Wang, 2006; Leblois et al., 2006), the precise mechanisms by which BG circuits influence behavior remain less than clear.

A number of attempts to relate BG structure to function have focused on microcircuitry within the largest BG nucleus, the striatum. The great majority of striatal neurons are medium-spiny projection neurons (MSNs) that integrate many convergent cortical and thalamic inputs and also provide the striatal output to other BG nuclei (Wilson, 2004). MSNs form a mutually inhibitory network, that may achieve action selection via "winner-take-all" interactions [for reviews see (Wilson, 2000; Wickens et al., 2007)]. However, MSN-MSN synapses are typically weak and far from the cell body (Jaeger et al., 1994; Koós et al., 2004), which has increased interest in the contributions of other elements of intrastriatal circuitry. Striatal parvalbumin-containing fast-spiking interneurons (FSIs) receive cortical inputs, and form strong perisomatic GABAergic synapses onto hundreds of surrounding MSNs (Bennett and Bolam, 1994). This FSI-mediated feed-forward inhibition has been

argued to make an important contribution to action selection and execution, via the broadly tuned, distributed suppression of MSNs representing unwanted actions (Kita et al., 1990; Parthasarathy and Graybiel, 1997). Consistent with such a role, a reduced number of striatal FSIs has been found in a rodent model of paroxysmal dystonia [co-contractions of opposing muscle groups; (Gernert et al., 2000)] and in postmortem tissue from human Tourette Syndrome patients [who have difficulty suppressing unwanted, highly learned actions (Kalanithi et al., 2005)].

Despite these ideas, few studies have investigated the activity of likely striatal FSIs in awake behaving animals (Berke et al., 2004; Schmitzer-Torbert and Redish, 2008; Berke, 2008). Recently it was observed that FSIs had highly individualized patterns of responding, and did not appear to act as a coordinated population in rats performing a radial maze task (Berke, 2008). However, this maze task was not ideal for investigating the fine temporal evolution of decision-related neural activity. We therefore examined patterns of FSI activity in rats performing a simple choice task, in which the timing of key events was closely monitored. To gain greater insight into these patterns we compared them to other, simultaneously recorded elements of cortical-BG circuits: striatal MSNs and neurons in primary motor cortex (M1) and globus pallidus (GP).

2.2 Results

To study striatal FSI contributions to choice behavior, we designed a simple conditional discrimination task (Fig. 2.1) that we expected to require intact function of the lateral (sensorimotor) striatum. Lateral striatum is involved in the acquisition and expression of cue-guided responses (McDonald and White, 1993; Adams et al., 2001; Berke et al., 2009), particularly movements to contraversive space (Cook

Table 2.1: *Distinct properties of neuronal populations.* Data are presented as mean (SD). Firing rate is the session-wide mean.

	n	Firing Rate / Hz	Peak Width /μs	Peak to Valley /μs
MI	72	2.7 (4.4)	250.3 (68.5)	695.7 (246.5)
MSN	257	1.1 (4.5)	270.7 (48.7)	873.2 (120.0)
FSI	38	18.3 (18.3)	128.1 (31.8)	259.9 (86.7)
O	7	22.1 (15.4)	231.8 (29.8)	423.8 (50.4)
GP	25	17.8 (13.4)	153.3 (47.8)	294.0 (96.6)

and Kesner, 1988; Carli et al., 1989; Brown and Robbins, 1989; Brasted et al., 1997; Packard and McGaugh, 1996), and this subregion also has the highest density of FSIs [*e.g.*(Kita et al., 1990; Berke et al., 2004)]. Hungry rats were placed in an operant box with five nosepoke holes, and each trial began with the illumination of one of the three more-central holes. The rat placed and held its nose in that hole while a brief instruction tone played, then performed a rapid nosepoke to one of the immediately adjacent holes, on either the left or right side depending on the instruction tone.

To verify that the striatum is important for the left-right choice in this task, in a group of well-trained rats (n=6), we performed unilateral striatal infusions of the $GABA_A$ agonist muscimol, or artificial cerebrospinal fluid (ACSF) as a control (Fig. 2.2). Infusion of ACSF did not affect task performance ($p=0.926$; all comparisons ANOVA with Tukey post hoc test). Infusion of muscimol caused a selective, reversible reduction in cued contraversive responding (responses towards the side opposite to the infusion, $p<0.001$) without interfering with ipsiversive performance ($p=0.990$). Hence the lateral striatum is preferentially involved in contraversive responding in this task, and $GABA_A$ receptor stimulation in this subregion can powerfully affect choice behavior.

To examine the activity of FSIs during the performance of this choice task, ad-

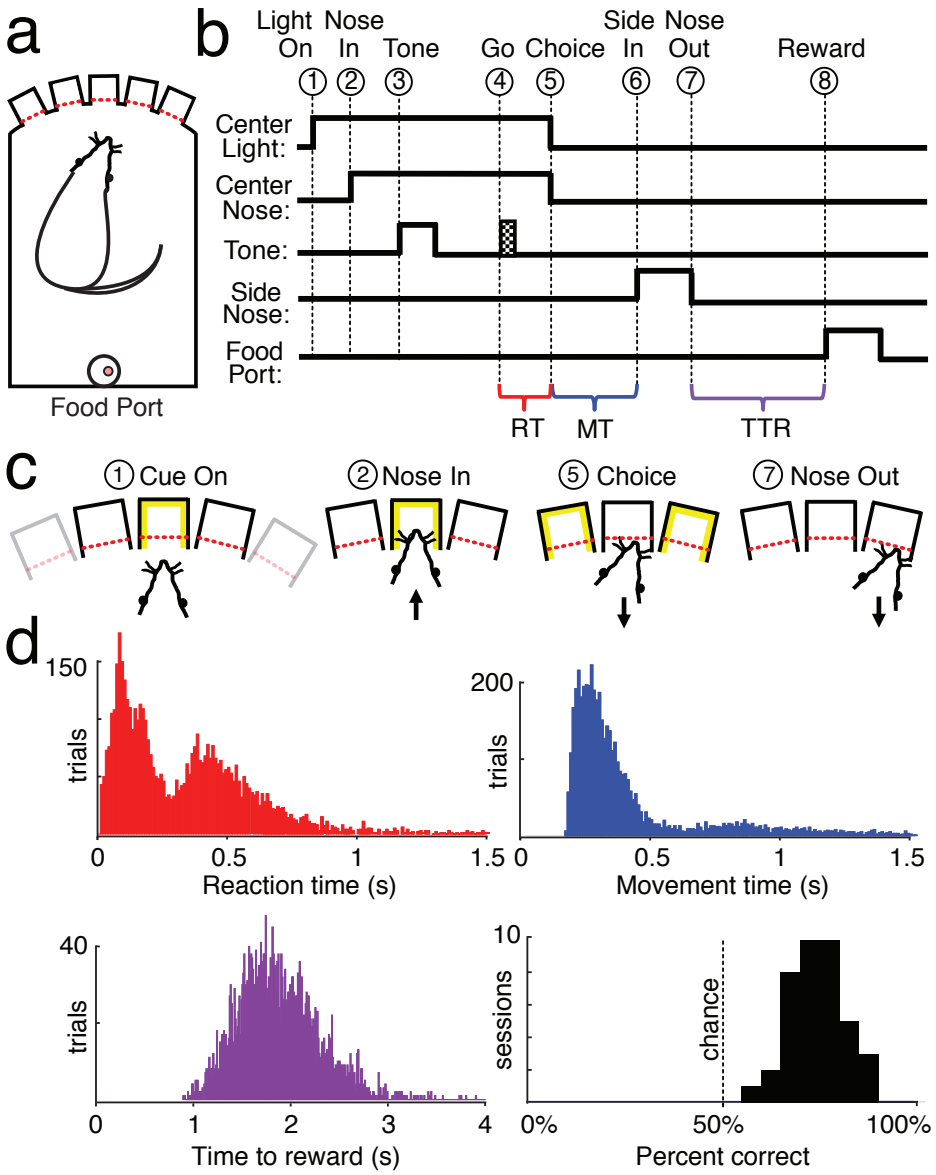

Figure 2.1: *Behavioral task and performance.* a) Depiction of operant chamber, with five nosepoke holes opposite food delivery port. b,c) Task event sequence, for correct performance. Each trial began with illumination of one of the three most central nosepoke holes ("Light On", event 1). The rat had to place his nose in the illuminated hole ("Nose In", event 2) and stay there (total hold duration = 900-1250 ms). During the hold window, a 250 ms instruction cue was played ("Tone", event 3), followed after a variable delay (600-950 ms) by a Go cue (125 ms white noise burst; "Go", event 4). The rat then pulled his nose out of the center hole ("Choice", event 5) and poked an immediately adjacent hole ("Side In" event 6). If the direction of movement matched the instructional tone (learned arbitrary mapping: 1kHz, go left, 4kHz go right) then a sugar pellet was immediately delivered with an audible food hopper click, and could be collected by moving out of the side hole ("Nose Out", event 7) and to the food port on the rear wall ("Reward", event 8). Brackets indicate time epochs used to measure reaction time (RT), movement time (MT), and time to reward (TTR). d) Distribution of RT, MT, and TTR times (10 ms bins) and session performance (5% bins) from all animals. The mean time between events 6 and 7 was 337 ms (SD 311 ms). All four subjects had bimodal RT distributions, consistent with rats sometimes anticipating, and sometimes reacting to, the Go cue (Figure 2.3).

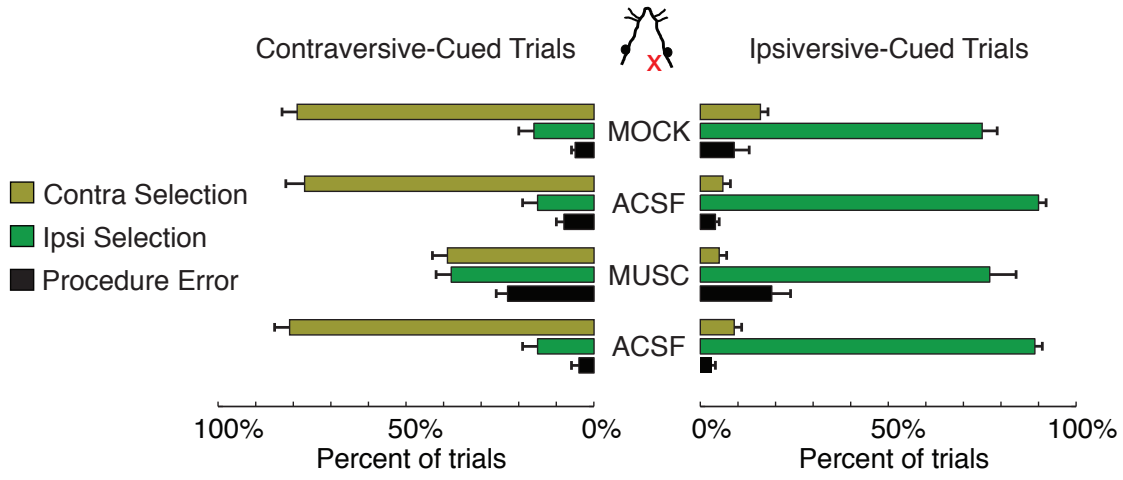

Figure 2.2: *Unilateral striatal injection of a GABA-A agonist causes a selective, reversible deficit in contraversive responding.* Results shown are from each of four succesive daily behavioral sessions in which rats (n=6) were tethered to an injection cannula but not infused ("MOCK" infusion, session 1), infused with artificial cerebral spinal fluid ("ACSF", Sessions 2,4) or infused with 0.05 μg muscimol ("MUSC", session 3). "Contraversive-Cued Trials", "Ipsiversive-Cued Trials" indicate whether the rat received an instruction cue to move to the nosepoke hole contralateral or ipsilateral to the injection side, respectively. Error bars, S.E.M.

ditional well-trained rats (n=4) were implanted with tetrodes into multiple target regions simultaneously (Fig. 2.4). Most tetrodes were aimed towards lateral portions of striatum (Fig. 2.6), though for comparison we also recorded neurons in other striatal subregions (including nucleus accumbens, NAc), GP, and "neck" regions of M1 (Sanes et al., 1990). To help distinguish between sensory and motor aspects of neural coding, the task variant used in the electrophysiological studies had a brief, variable delay between the instruction cue and a "go" cue for movement onset (as shown in Fig. 2.1). From the intended targets a total of 437 distinct, well-isolated cells (striatum, 339; M1, 73; GP, 25) were obtained from 39 sessions (mean number of trials/session:125.8; mean % correct: 74.4, range: 64.3 - 87.1, Fig. 2.1e).

Examination of striatal neuron waveforms revealed distinct clusters of cell properties that closely resembled those seen in our previous studies of striatum [in dif-

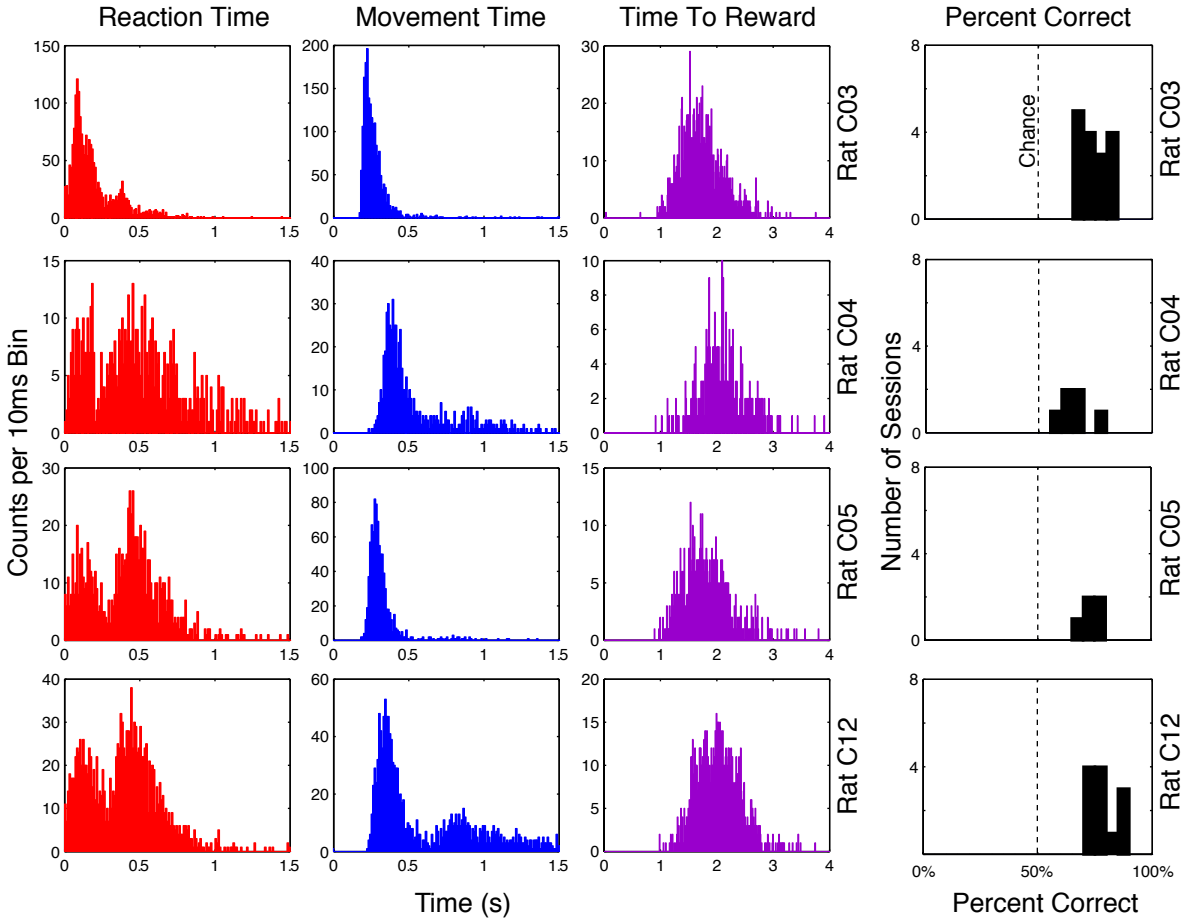

Figure 2.3: *Reaction time, movement time, time to reward, and percent correct for each subject of the electrophysiology experiments.* Behavioral data shown is only for those sessions for which we included neurons in the electrophysiological analyses. Note the prominent bimodal distribution in the reaction times. Regression analysis revealed that short RTs (< 300ms) inversely varied with the hold time ("Go" - "Nose In"; p=0.00004, regression t-test), while long RTs did not (p=0.39, regression t-test) This suggest that that the rats were sometimes anticipating the Go cue (short RT peak) and sometimes reacting to it (long RT peak). The percent correct session histograms are divided into 5% bins.

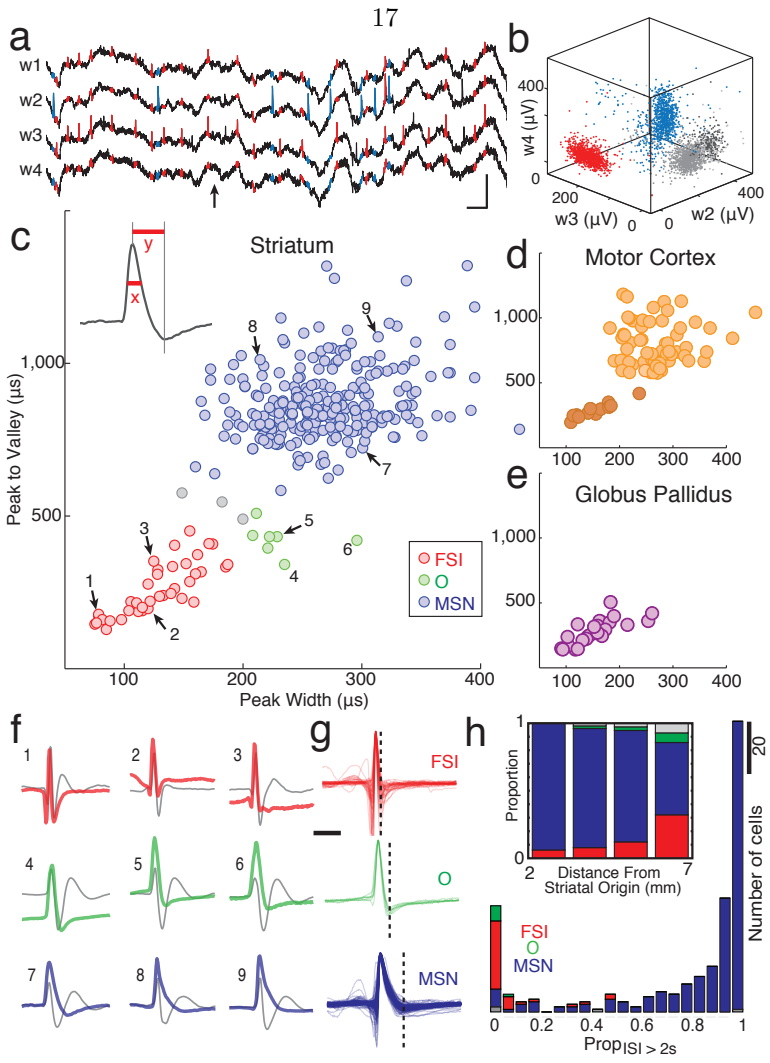

Figure 2.4: *Classification of neurons.* a) Example of wide-band recording (1-9,000 Hz) from a tetrode (four wires, w1-4) in striatum (arrow = arrival at food port; blue color highlights spikes from a presumed MSN, red highlights spikes from a presumed FSI). Scale bars: 0.5mV,10ms. b) Single unit identification based on peak filtered spike voltage; three of the four tetrode wires shown. The red, blue clusters correspond to the spikes in a). c-e) Scatter plots of mean spike waveform durations (x, peak half-maximum; y, peak-to-valley time) for each single-unit. c) Striatal cells. Presumed MSNs are in blue, FSIs in red, O cells in green. d) M1 cells. Darker color indicates possible interneuron population. e) GP cells; all had brief spike durations. f) Mean wide-band spike waveforms for nine representative striatal cells. Numbers (1-9) correspond to cells indicated in c. For comparison to prior studies, digitally filtered versions of the waveforms are also shown (gray, 600-6,000 Hz Butterworth filter). g) All wide-band waveforms for the FSI, O, and MSN striatal cell classes, superimposed to show inter-cell variability. Vertical dashed bar = mean time of detected valleys. h) Phasic vs. tonic activity of striatal cell types. Histograms show, for each cell, the proportion of time spent in long interspike intervals (ISI>2s). Inset: Presumed FSIs were more common in sensorimotor (lateral/dorsal/caudal) striatum. Bars show proportion of each cell type by distance from an origin point near the medial-ventral-rostral tip of the striatum [AP 3.13, ML 0, DV 8 mm below bregma; compare to (Berke et al., 2004)].

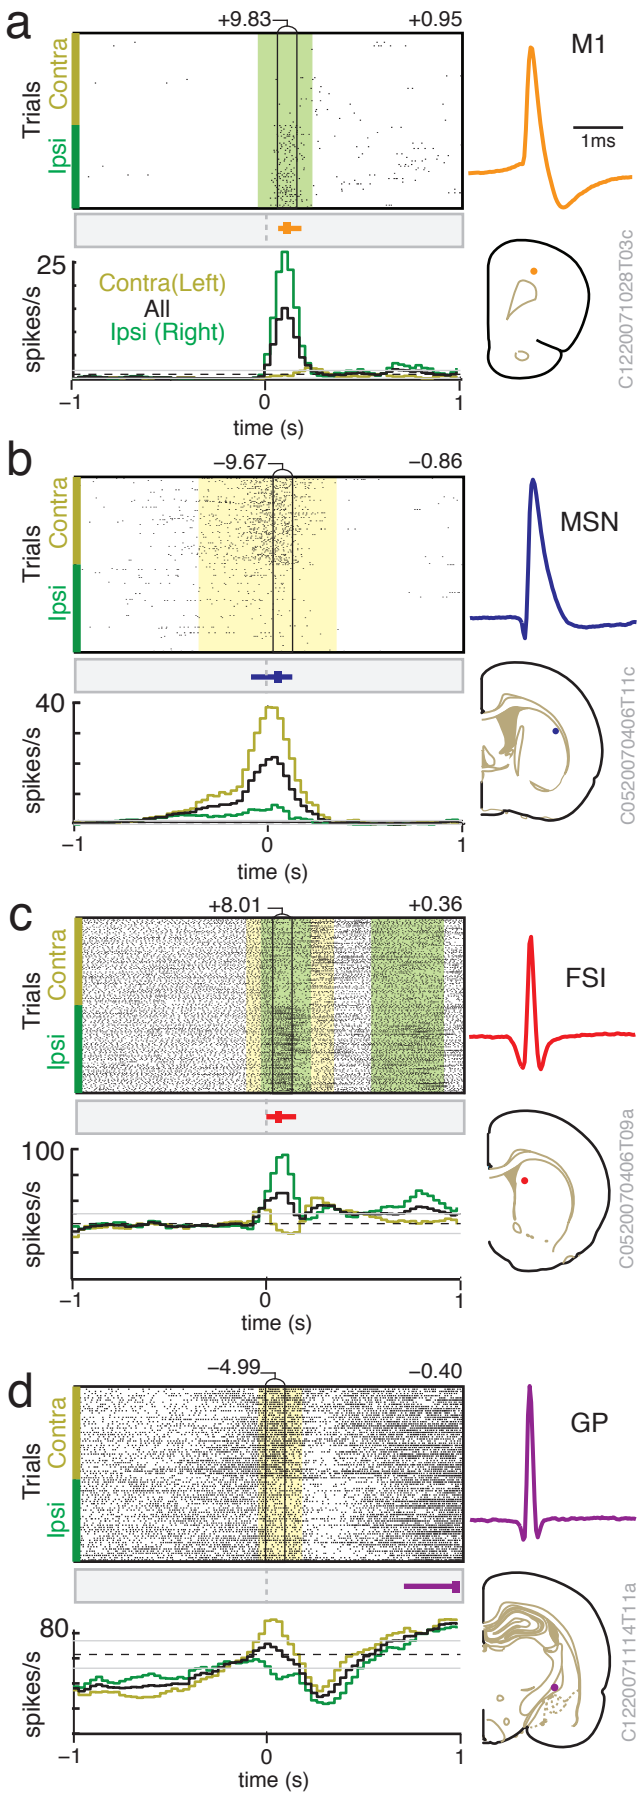

Figure 2.5: *Analysis of peak firing rate and movement selectivity.* a-d) Examples of movement-related cells. In each case mean wide-band spike waveform is at top right and recording location at bottom right. Center panels show spike rasters for all trials, aligned on movement onset (event 5) and separated by movement direction (contraversive on top, ipsiversive on bottom). Epochs with a significant ($p < 0.01$) contra/ipsi firing rate difference are indicated by color shading (contraversive = gold, ipsiversive = green). Vertical black lines indicate the 100 ms period with the most significant contra/ipsi difference (peak directional selectivity is shown above bin). Bottom panels show corresponding peri-event time histograms (PETHs) for contraversive (gold), ipsiversive (green), and all trials (black). Bin size = 30 ms, with 3-point moving average smoothing. Above the histograms, the point of maximum firing rate across all PETHs is shown with a vertical tick, and the period of greater than half-maximal response (i.e. > (mean rate + peak rate)/2) by a horizontal colored bar. The selectivity index value is shown above the rasters, on the right side.

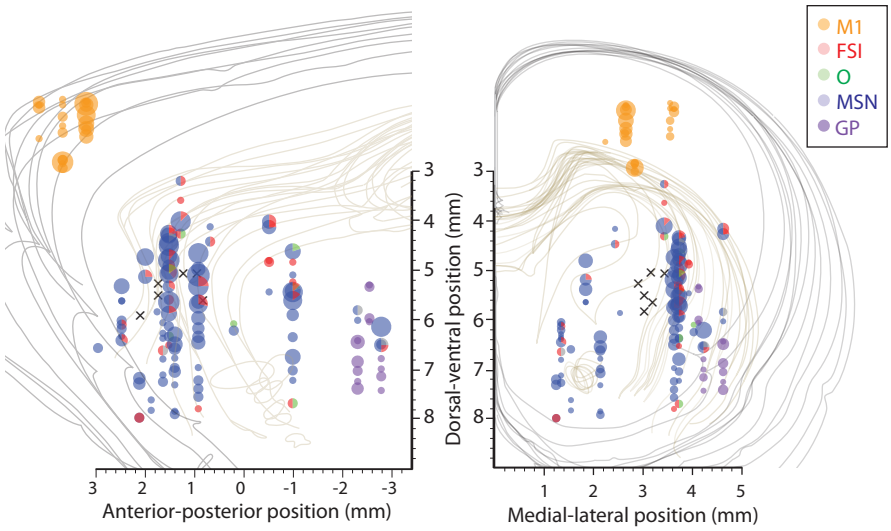

Figure 2.6: *Locations of electrophysiological recordings and drug infusions.* Locations are shown as projections onto sagittal (left) and coronal (right) stacked atlas sections (Paxinos and Watson, 2005). Area of circles is proportion to number of simultaneously recorded cells, segment colors indicate the class of cells from each location. The location of each muscimol injection in the separate behavioral experiment is indicated with an "x". Note that the GP cells were recorded from caudal parts of that structure.

ferent rats; (Berke et al., 2004; Berke, 2008)]. The largest class of striatal cells (n=257) had relatively long duration waveforms (Fig. 2.4g; Table 1) and typically also had phasic firing patterns (Fig. 2.4h); these were presumed to be MSNs, which comprise the great majority of striatal neurons. The second most numerous group (n=38) had the very brief waveforms (Kawaguchi et al., 1995; Mallet et al., 2005) and graded intrastriatal distribution [Fig. 2.4h inset; (Kita et al., 1990)] characteristic of parvalbumin-positive fast-spiking interneurons [FSIs; for discussion see (Berke et al., 2004)]. Compared to the MSN population, these cells had higher firing rates (Table 2.1) and were usually tonically active (Fig. 2.4h). A final class of striatal cells also had high baseline firing rates, but a characteristic waveform shape with intermediate peak and valley widths; as before (Berke, 2008) we labeled these as "O" cells for other, currently unknown phenotype. Because they were few in number (n=7),

O cells were excluded from most analyses. Distinct clusters were also observed for waveforms from motor cortex cells (Fig. 2.4d). The briefer waveforms are very likely to be interneurons (Barthó et al., 2004; Cardin et al., 2009), though for our analysis M1 cells were treated as one group except where noted below. GP cells all had high firing rates and relatively narrow waveforms (Fig. 2.4e), consistent with prior observations [*e.g.*(Turner and Anderson, 1997)].

2.2.1 FSIs are disproportionately active around choice execution

We next wished to determine if FSIs are preferentially active at any particular moment during the performance of the choice task, in comparison to other neuronal populations. To generate a temporal response profile for each neuron we calculated perievent time histograms (PETHs; see Fig. 2.5 for examples) around each task event, and normalized this event-related firing by the peak response across all PETHs. Figure 2.7a shows this profile for all task-responsive cells of each class, sorted by moment of peak response. Since there are variable delays between task events, the highest firing rate obtained across all PETHs allows us to determine which task event produces the strongest response for each neuron. This epoch of peak response is shown for each cell in Fig. 2.7b.

Each studied brain area contained many cells with task-related changes in firing rate, especially near arrivals at the baited food port. In this study we did not attempt to distinguish between motoric and hedonic aspects of reward retrieval and consumption, but simply refer to all cells with maximal firing around reward receipt as "reward-related". However, in marked contrast to the MSN population, FSIs were disproportionately active around the earlier time at which the rats initiated their left/right choice ("choice execution"; event 5). For units active during task performance, 35.1% (13/37) of FSIs showed maximal firing when aligned to this

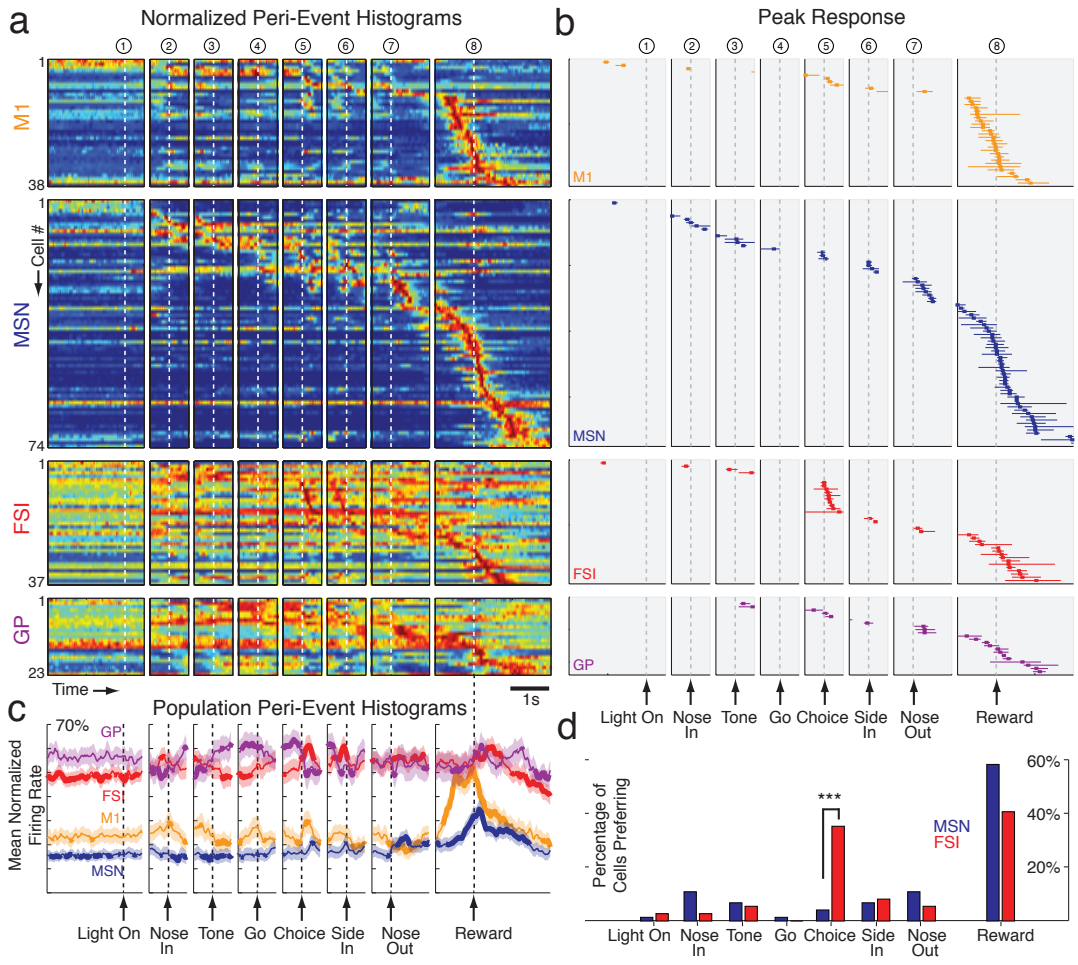

Figure 2.7: *FSIs preferentially increase firing rate during choice execution.* a) Peri-event time histograms (PETHs) for each cell, aligned on each of the eight events. For each cell, firing rate is shown by a color scale ranging from blue (zero) to red (peak rate), and within each class of cells neurons are rank ordered by time of peak firing.To be included, a cell had to have a firing rate of at least 5Hz in at least one PETH (using 30 ms bins). For display purposes, only a portion of each 3 s PETH window is shown (see Methods). For additional cell classes, see Fig. 2.8. b) Events associated with peak firing rate. Order of cells is the same as (a). As in Fig. 3, vertical tick marks indicate the time bin with peak firing rate, and horizontal lines indicate the epoch for which firing rate was elevated more than half-way between mean rate and peak rate (shown only for the PETH containing peak rate). c) Mean normalized firing rate for each cell population in (a). Normalization before averaging helps to emphasize the population response, by minimizing the contribution of particular cells with especially strong responses. Bin size = 30 ms, smoothed with 3-point moving average. Shaded region = S.E.M., and bold lines indicate that population mean is outside 5% and 95% confidence intervals (see Methods). d) Proportions of MSN and FSI cells with peak firing for each event. The choice execution event was associated with a significantly different proportion of FSI and MSN peak firing (*** p=0.0002, two-sample proportion test, adjusted for multiple comparisons). All other events did not show a significant difference (p>0.05, adjusted for multiple comparisons). The inclusion criterion of at least 5Hz peak firing did not substantially change our results: with all neurons included, a significantly higher proportion of FSIs than MSNs still showed peak firing in association with the"choice" event (Z=3.48, p=0.0009, two-sample proportion test, corrected for multiple comparisons).

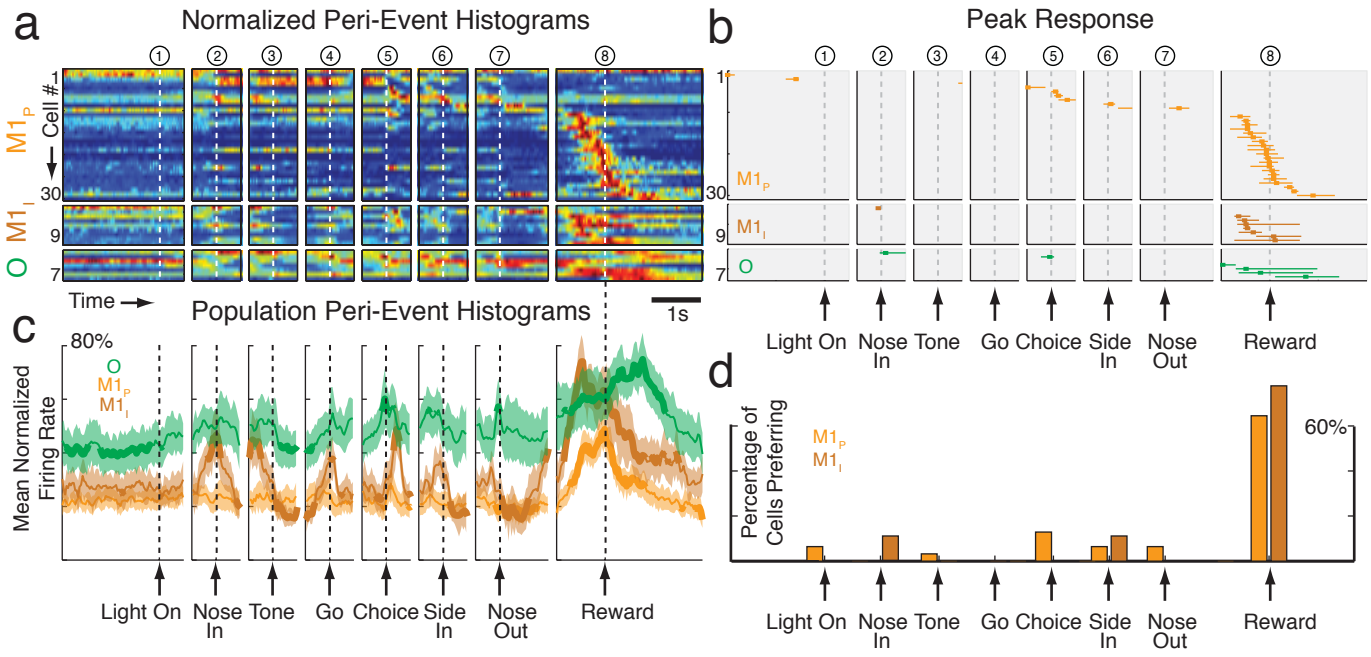

Figure 2.8: *Analysis of O, $M1_P$, and $M1_I$ neuronal classes.* (a-d) Additional cell classes using analysis described in Fig. 2.7. O cells may represent an additional class of striatal interneurons [see (Berke, 2008)], $M1_P$ and $M1_I$ are motor cortex presumed projection and interneuron populations (see Fig 2.4d). Presumed cortical interneurons did not show the same pattern of event preference as striatal FSIs, and there were no significant differences in the events that cortical $M1_I$ and $M1_P$ populations preferred ($p>0.05$ for all comparisons; note however that the number of cells analyzed is not large). As in previous work (Berke, 2008), O cells preferred the rewarding event, although slightly less uniformly in this task.

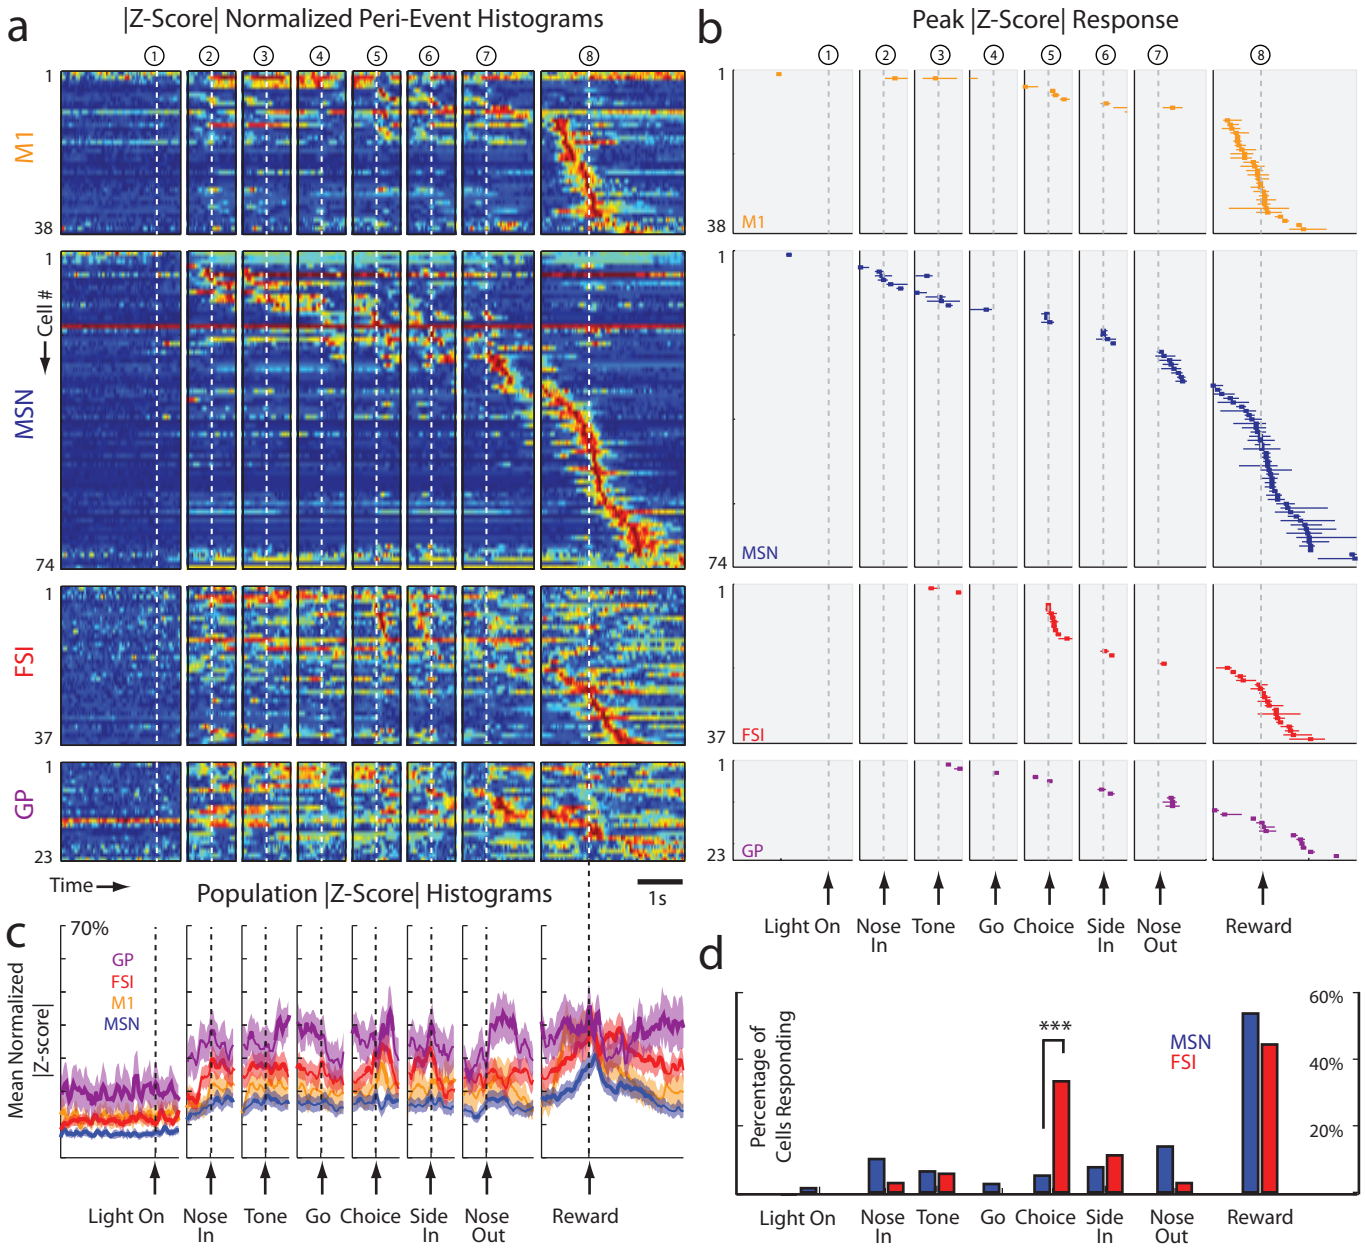

Figure 2.9: *|Z-score| normalized response profile.* Format is identical to Figure 2.7a, except that instead of firing rate the absolute value of the firing rate Z-scores are used instead (see Methods). This serves to highlight both potential increases and decreases in firing rate. PETHs are displayed as a heat map from blue (no change from baseline) to red (maximum change from baseline, either positive or negative). Using this variant of the analysis, the difference between the proportion of FSIs and MSNs responding to the choice execution event remained significant ($p=0.00065$, two-sample population proportions test, adjusted for multiple comparisons), and all other pairs remained not significant ($p>0.05$, adjusted for multiple comparisons).

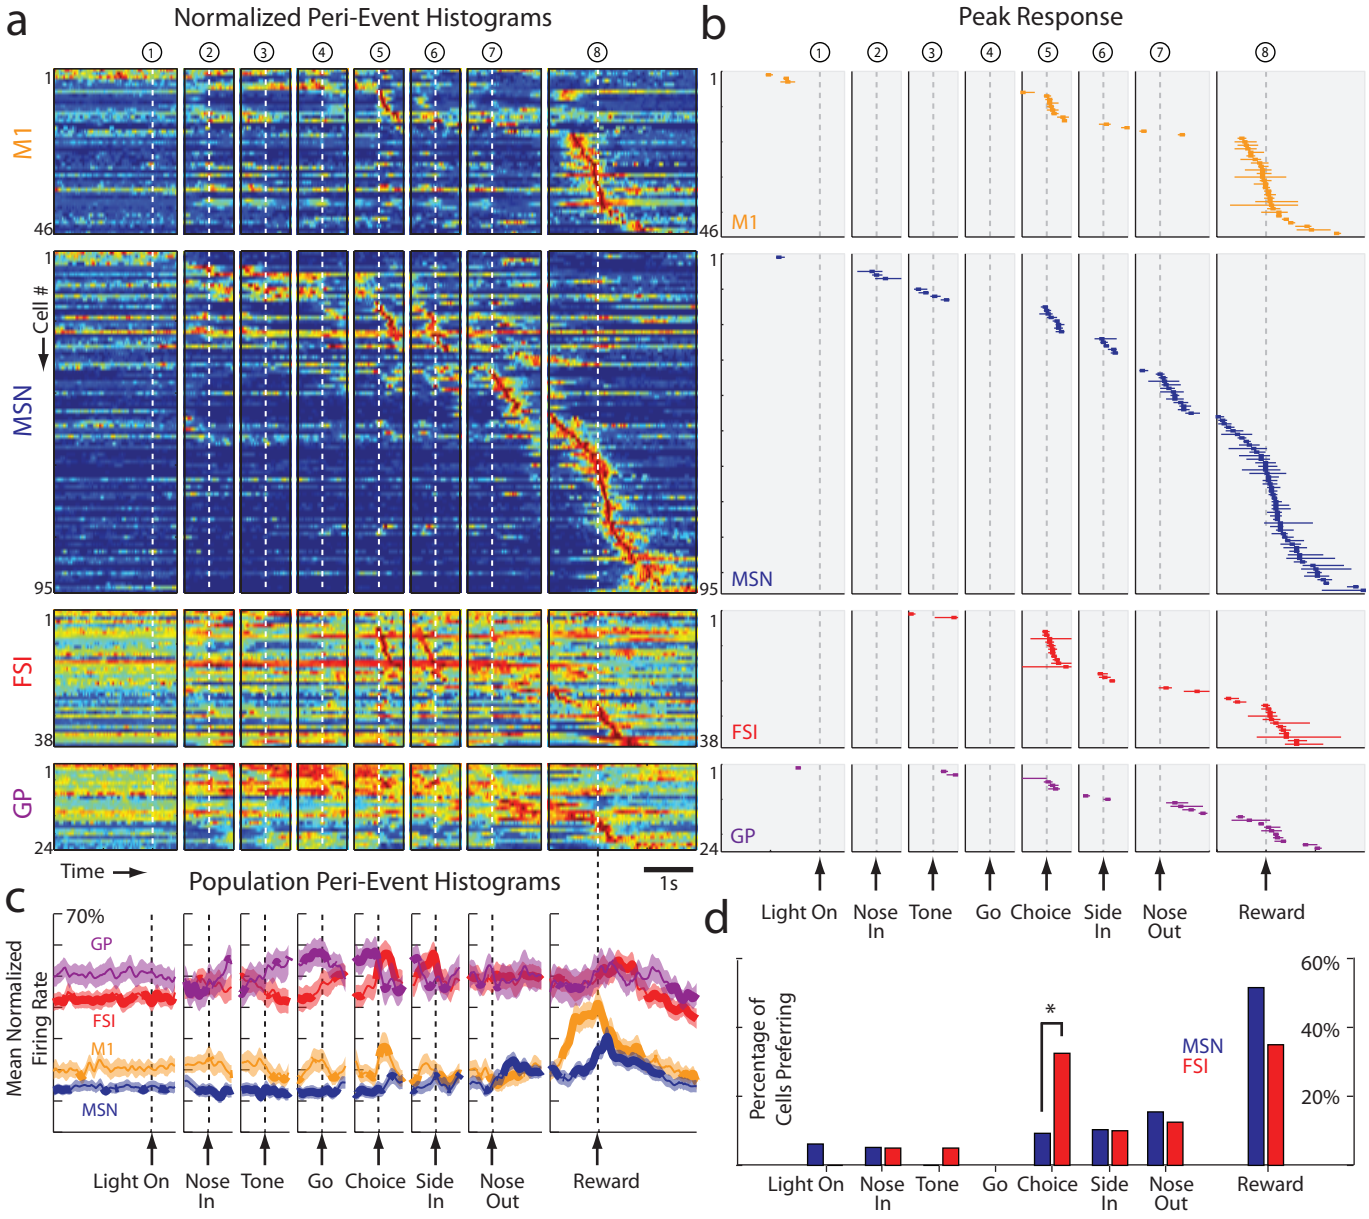

Figure 2.10: *Selective PETH response profiles.* Format is identical to Figure 2.7a, except that the peak firing rates were calculated using subsets of trials, during which (1) High tone played, (2) Low tone played, (3) the subject moved contraversive, or (4) the subject moved ipsiversive. This serves two purposes. First, it helps to includes cells for which the 5Hz peak inclusion criterion was reached, but only on a subset of trials. Secondly, it helps to find the preferred event for some neurons, in which firing rate increases on one subset of trials were matched by decreases on another subset. Under these conditions, the choice execution event remained associated with a significantly different proportion of FSI and MSN peak firing (* p=0.012, two-sample proportion, adjusted for multiple comparisons), and all other pairs remained not significant (p>0.05, adjusted for multiple comparisons).

event - a significantly higher proportion than the 4.1% (3/74) of MSNs (Z=4.39, p=0.0002, two-sample proportion test corrected for multiple comparisons). Since we were interested in both increases and decreases in firing rate, we repeated these analyses using an alternative form of PETH normalization, based on the absolute value of firing rate Z-scores; very similar results were obtained (Fig. 2.9). Although there are several other movement-aligned events in this task (nose in, side in, nose out), only choice execution showed a significantly higher proportion of FSIs active over MSNs (Fig. 2.7d). We note that this is the moment at which the sequence of actions performed within a trial bifurcates along two highly learned paths. This produces a need to suppress an alternative action, that is prepotent and brought close to execution "threshold", yet unintended on this trial. Thus enhanced FSI activity at this time is at least consistent with a role in the inhibition of alternative actions.

We considered several other reasons why FSIs might be more likely to show this "choice-related" firing than MSNs. Firstly, FSIs and MSNs have very different average firing rates - could this be affecting our analyses? We think not, since both GP cells and presumed cortical interneurons had similar high firing rates to FSIs, yet neither group showed a comparable preference for the choice event (Figs. 2.7 and 2.8). Secondly, FSIs tend to be found more often in lateral striatum. We therefore examined whether the preferential activity of FSIs with choice execution was simply a reflection of the information processing occurring in that brain subregion. While the small number of choice-related MSNs were all found in dorsal-lateral striatum, choice-related FSIs were much more broadly distributed (Fig. 2.11). This indicates that the different balance of choice-related and reward-related firing for FSIs and MSNs is not caused by the increased FSI density in lateral striatum, and suggests

that a choice-related increase in striatal FSI activity may act as a relatively global signal. Finally, averaging across the whole session may diminish some strong MSN responses that occur only on certain trial types. To assess this possibility we repeated our analysis, this time assigning cells to events on the basis of the strongest PETH response during either low tone trials, high tone trials, leftward trials or rightward trials. Although this did change the assignation of some specific neurons to events, FSIs still disproportionately preferred the choice execution event (Fig. 2.10).

Assigning neurons to a single event actually underestimates the proportion of FSIs that increase firing around choice execution, since many such FSIs showed even greater activity at another point in the trial . When we examined the overall response of each neuronal population a clear "pulse" of enhanced FSI activity was observed around choice execution, while peak MSN population activity was instead found around reward retrieval (Fig. 2.7c). This population level analysis further revealed a striking pattern of opposite changes in the FSI and GP populations. GP cells tended to have elevated firing rates as the rats maintained their head-fixed position while waiting to make their choice; as they finally initiated an action, GP population activity fell sharply as FSI activity increased. This result is especially interesting as there is a specific direct GABAergic projection from GP to striatal FSIs (Bevan et al., 1998), suggesting that disinhibition may contribute to the FSI pulse.

2.2.2 FSIs are selective for movement direction

A coordinated pulse of striatal FSI activity is, by itself, consistent with theories that view these cells as providing broadly-tuned, blanket suppression of MSNs via feed-forward inhibition [*e.g.* (Wickens and Arbuthnott, 1993; Parthasarathy and Graybiel, 1997)]. However, individual FSIs clearly have diverse patterns of firing rate change, both in the present data and in our prior results (Berke, 2008). To

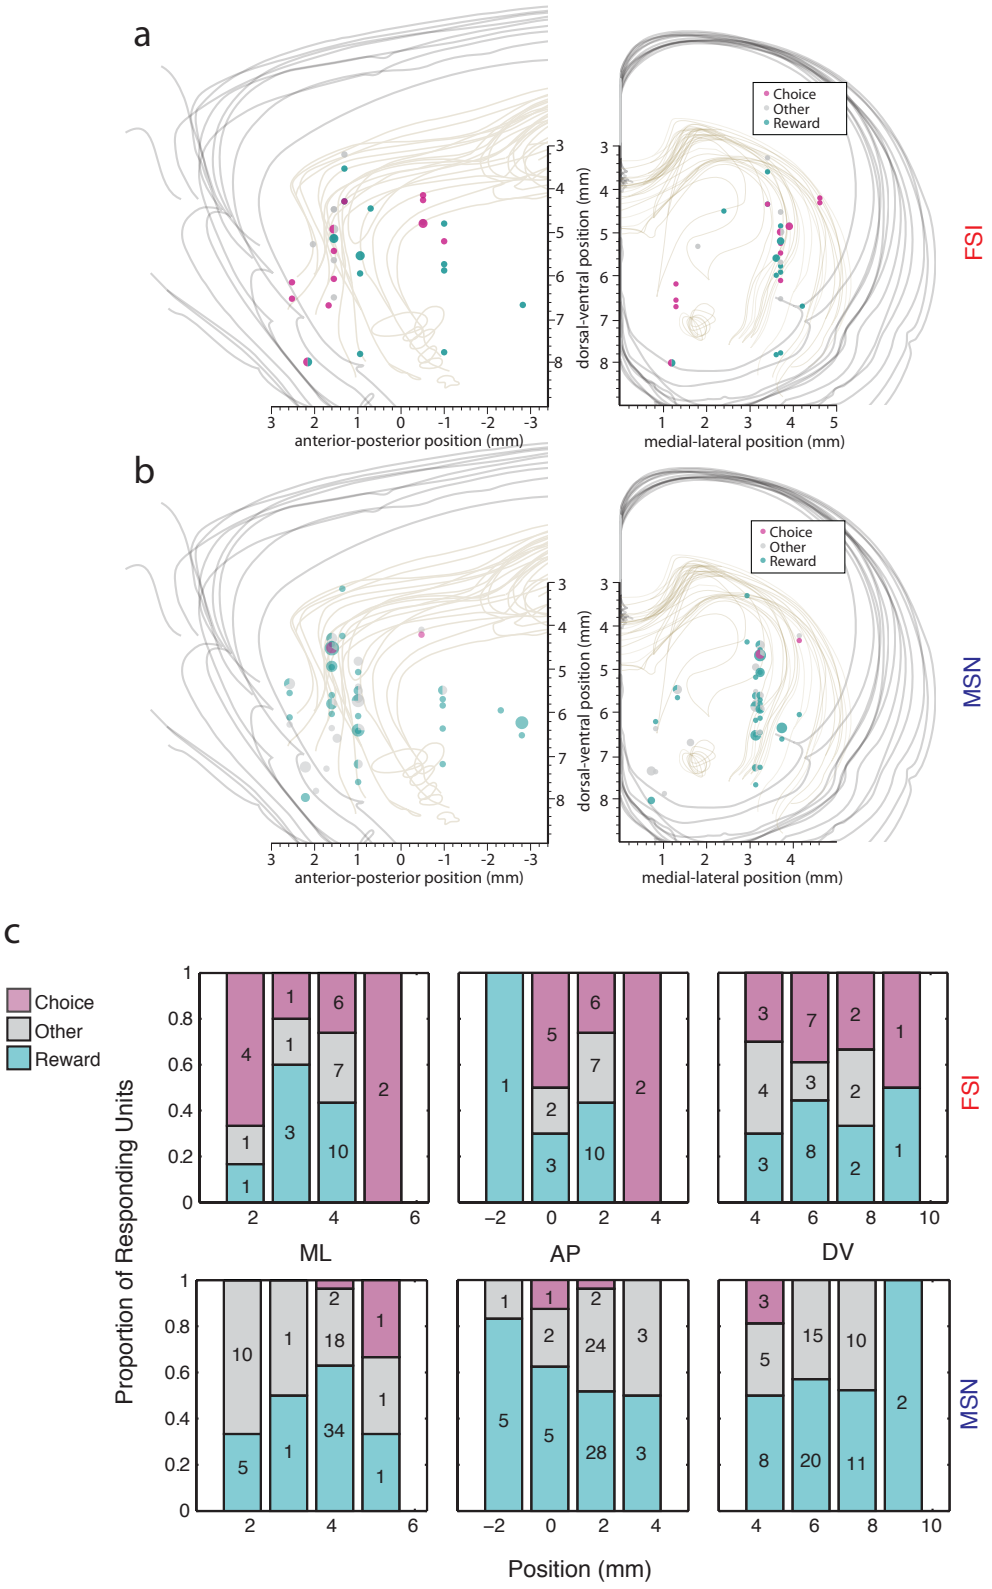

Figure 2.11: *Location of cells with peak activity during the "choice" or "reward" epochs.* (a,b) Sagittal (left) and coronal (right) view of cell locations of FSI cells (a) and MSN cells (b). c) Location histogram of reward- and choice-related cells for FSI (top) and MSN (bottom) populations. Histological locations are given relative to bregma in the ML (left), AP (center), and DV (right) directions.. The choice-related FSIs and MSNs did not show a significant relationship to position ($p>0.05$, regression t-test for both choice and reward preference regressed against the recording location: ML, DV, AP, and distance from origin).

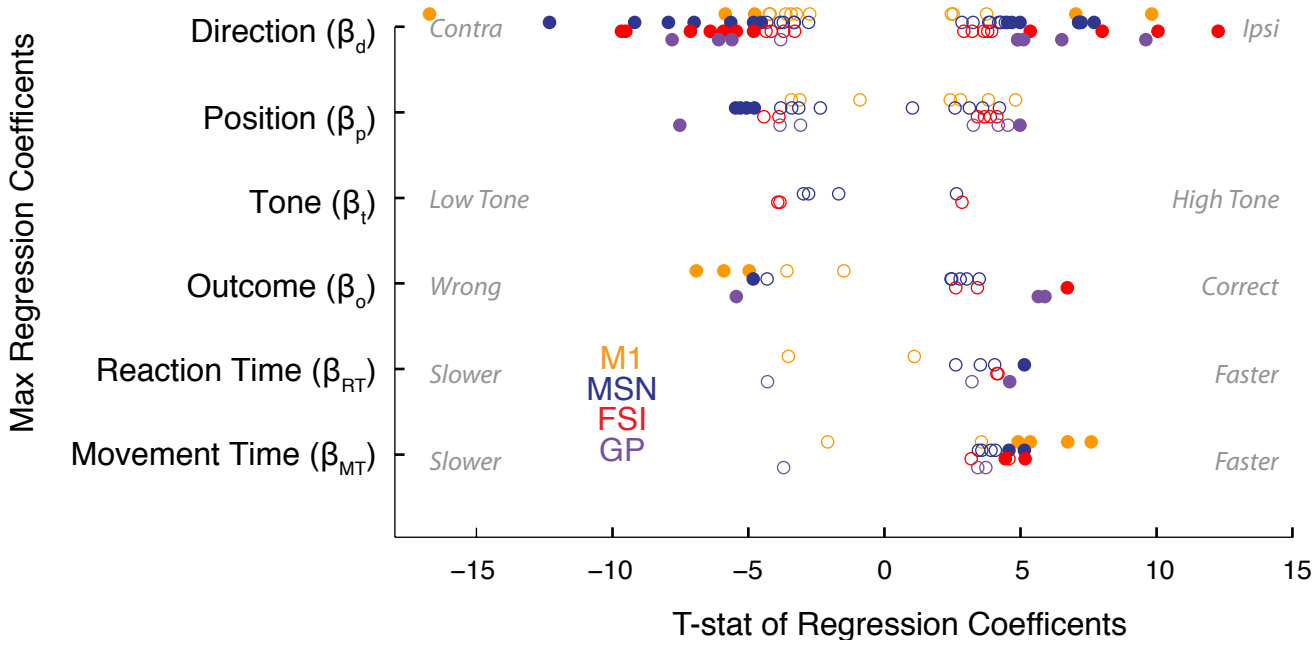

Figure 2.12: *Regression Analysis of Choice Event.* Maximum T-statistics are displayed for four cell classes (M1, MSN, FSI, GP) for six regression coefficients: Direction (β_d), Position (β_p), Tone (β_t), Outcome (β_o), Reaction Time (β_{RT}), Movement Time (β_{MT}). Only the maximum t-stat for each cell is plotted. Solid markers indicate that the coefficient was significant ($p<0.01$, regression t-test, corrected for multiple comparisons).

explore information processing by individual FSIs, we next examined selectivity for one chosen action over the other. The great majority of FSIs had high movement selectivity, and both contraversive and ipsiversive-preferring neurons were observed in similar numbers (Fig. 2.13). To tease apart other factors that may contribute to the firing rate of FSIs and other subpopulations during action selection, we performed multiple regression analysis using a range of variables including the instruction tone, movement direction, the spatial position at which the choice was executed, reaction time, movement time, and trial outcome (Fig. 2.12). For each of the MSN, FSI, M1 and GP populations, movement direction was the most common dominant factor, with very few cells more concerned with other task aspects such as the specific tone or the rat's spatial position (e.g. Fig. 2.13d. In particular, among FSIs that had significant factors in the multiple regression 11/14 (78.6 %) were most concerned with the specific movement direction. These results indicate that, rather than FSIs acting constantly as a single global signal, the transient coordination of FSI activity is superimposed on a background of idiosyncratic individual firing rate time courses (Berke, 2008) that are highly influenced by movement direction.

2.2.3 Nearby FSIs and MSNs show opposite selectivity, but weak interactions

The muscimol injection experiment indicates that the lateral striatum is particularly important for contraversive movements in our task. Yet, both contraversive- and ipsiversive-preferring neurons were found intermingled in lateral striatum (and other striatal subregions) in similar numbers, and there was no gross relationship between recording location within striatum and direction preference ($p>0.05$ for both MSNs and FSIs, regression t-tests for directional selectivity vs. each [AP,ML,DV] dimension of recording location). On a finer scale, MSNs and FSIs are each components of local microcircuits (Gustafson et al., 2006) that may be functional modules (Wilson,

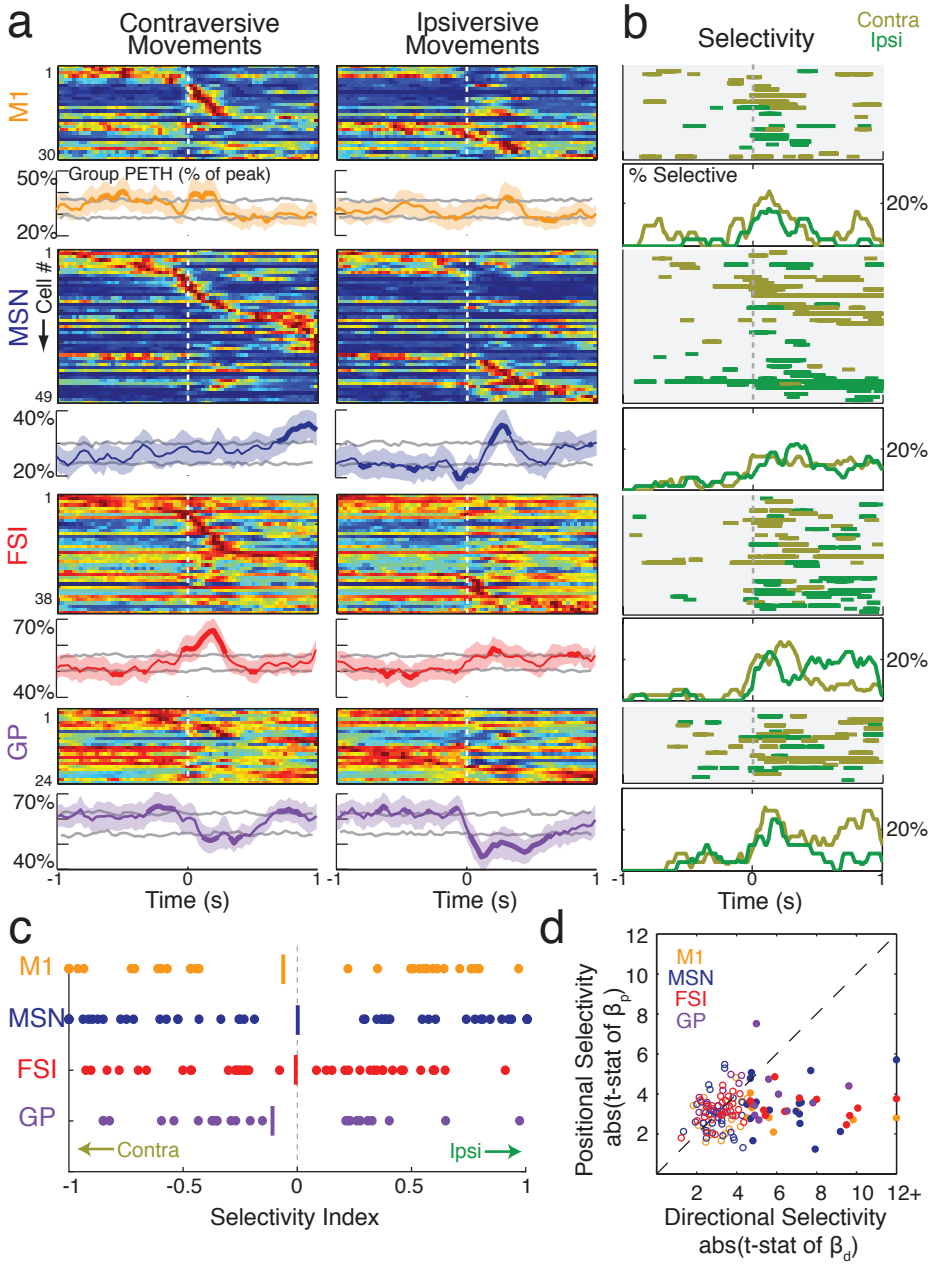

Figure 2.13: *Directional selectivity around choice execution.* a) Normalized PETHs for contraversive (leftward) and ipsiversive (rightward) movements for four cell classes. To be included, a cell had to have a peak firing rate of at least 5Hz within 1s of movement onset (toward either direction). For each cell class, the top plot shows normalized PETHs for individual neurons (rank ordered by the time of peak firing) and the bottom plots show population PETH. Grey lines indicate the 5th and 95th centile of confidence intervals; portions of the mean line that extend outside of this interval for at least two consecutive bins are indicated by increased thickness. Shaded area = S.E.M. b) Time epochs of significant directional selectivity. For each cell class, the top plots indicate epochs for which each neuron fired at a significantly higher rate on trials with contraversive (gold) or ipsiversive (green) movements ($p<0.01$, based on t-stat of regression β_d; see Methods); bottom plots indicate the instantaneous percentage of cells showing significantly higher firing rates for each movement direction (time bin = 5ms with 3-point smoothing). c) Selectivity index values for different cell classes. Vertical line indicates mean. d) Scatter plot of directional selectivity vs spatial (i.e. which of the three central holes) selectivity. Filled circles indicate that either the peak directional or spatial selectivity was significant ($p<0.01$).

2000). Since FSIs provide GABAergic input to nearby MSNs they might be expected to have opposite direction preferences (Diester and Nieder, 2008). To test this, we examined all pairs of cells recorded simultaneously from the same tetrode, for which both neurons had significant direction selectivity. Consistent with an inhibitory role of FSIs within local microcircuits, all such FSI-MSN pairs had opposite direction preferences (8/8 pairs opposite; p=0.0039, 50% binomial distribution). This was not the case for MSN-MSN pairs, which tended to have the same direction preference (4/15 pairs opposite).

Finally, we looked for additional evidence of the inhibitory influence of FSIs on MSNs using crosscorrelograms. A suppression in target cell firing for 2-3 ms after a reference cell spike has been used to identify presumed monosynaptic $GABA_A$-mediated inhibition (Barthó et al., 2004). Once again we focused on FSI-MSN pairs recorded from the same tetrode, since for these pairs the MSN was likely to be in range of the FSI axonal field (Berke, 2008). Although in many cases we had limited power to detect inhibition due to the low baseline firing rates of MSNs, of 91 such pairs none showed evidence of FSI inhibition of MSNs in session-wide crosscorrelograms (data not shown). This absence of obvious spike suppression suggests that FSI:MSN interactions in awake, behaving animals may operate differently to the strong shunting inhibition of MSNs seen following direct somatic current injection into single FSIs *in vitro* (Koós and Tepper, 1999).

2.3 Discussion

Cortical-basal ganglia circuits use multiple internal control signals, such as the neuromodulators dopamine and acetylcholine, to achieve their overall functions of action selection and reinforcement-based learning. Here we have found another po-

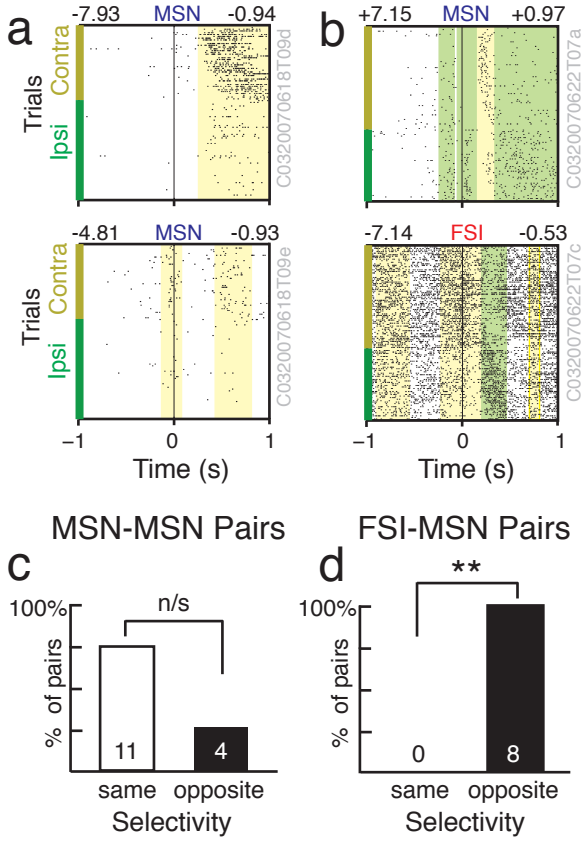

Figure 2.14: *Nearby FSI-MSN cell pairs have opposite direction selectivity* a,b) Examples of two simultaneously recorded MSN cells (a) and a MSN and FSI cell pair (b). Raster format is the same as in Fig. 3. Note the inverse relationship between peak directional selectivity (above rasters, right) in the MSN/FSI pair. c,d) Selectivity directions for all 15 simultaneously recorded MSN-MSN pairs (c) and all 8 MSN-FSI pairs (d) for which both cells showed direction selectivity within 1s of choice execution. Double asterisks indicates significance, p=0.0039. n/s, not significant: p=0.0592.

tential internal control signal: a pulse of enhanced striatal FSI activity that occurs at moments of choice execution. This FSI pulse did not occur in conjunction with every performed action, but rather was specific to times at which an alternative, highly-learned action required suppression. Additional studies will be required to further delineate the circumstances that produce enhanced FSI activity. However, our results suggest a role for striatal FSIs in the suppression of unwanted actions, that seems quite compatible with recent findings of an FSI deficit in Tourette Syndrome (Kalanithi et al., 2005).

The specific firing rate time courses of individual FSIs during operant task performance were reproducible from trial to trial, but highly variable between different (even neighboring) FSIs. This is consistent with our prior work, in which we found highly idiosyncratic FSI activity in a radial maze task (Berke, 2008). Both tasks demonstrate that the patterns of FSI firing are far more complex than had been expected, given their interconnection by gap junctions and inhibition of many nearby MSNs (Koós and Tepper, 1999). However, in the radial maze we found no clear evidence for any moment of coordinated FSI firing rate change, despite related task demands such as cue-guided decision-making. The most important difference may be that the operant task used here was specifically designed to obtain greater temporal definition of behavioral events, including the moment of choice execution. Unlike the maze task, the operant task also includes an enforced hold period that helps to define just when the chosen action is initiated, and this may have introduced additional demands for inhibitory control involving FSIs.

What accounts for these two aspects of striatal FSI firing - idiosyncratic individual activity time courses, but transiently coordinated firing rate increases at choice execution? We propose that this reflects two different types of input to FSIs. On the

one hand, FSIs are receiving complex combinations of sensory and motor information from multiple cortical regions (Ramanathan et al., 2002). They seem to be active participants in the fine details of information processing within striatal microcircuits, influencing the spiking of MSNs in a subtle manner rather than by constant strong inhibition. On the other hand, FSIs receive a continuous barrage of GABAergic inputs from high-firing-rate GP neurons (Bevan et al., 1998). The pallidostriatal pathway is more divergent than the striatopallidal pathway (Spooren et al., 1996), allowing GP neurons to coordinate FSI activity over more widely distributed regions of striatum. GP cells themselves receive inputs from subthalamic nucleus, which can provide a broad brake over behavior (Aron and Poldrack, 2006), and increases in GP activity have been previously noted under hold conditions, in which a specific movement is programmed but not yet executed [*e.g.*(Turner and Anderson, 2005)]. The sharp reduction in GP firing at the end of the hold period is thus consistent with the broad disinhibition of striatal FSIs. GP control of coordinated FSI firing may also account for the transient nature of the FSI pulse, since FSI-MSN-GP-FSI connections form a negative feedback circuit. Biophysically-accurate simulations of such basal ganglia dynamics would be helpful for evaluating this proposal.

Within cortical microcircuits, the activity of excitatory projection cells and inhibitory interneurons is tightly linked by recurrent connections [*e.g.* (Shu et al., 2003)]. This local feedback sculpts oscillatory activity (Cardin et al., 2009) and allows maintenance of active representations without runaway excitation. Striatal microcircuits have a distinct architecture, in which the projection cells (MSNs) do not directly send recurrent connections to interneurons. Since the striatum receives excitation exclusively from extrinsic sources and the MSN projection cells are GABAergic, tightly coupled FSI activity is not required to prevent runaway excitation. This

may help to explain the quite different population time courses of FSIs and MSNs in the present experiments. Our findings are generally compatible with models in which basal ganglia circuits serve as non-linear gating or threshold elements in decision-making, rather than (for example) representing the accumulated evidence supporting one decision over another [*e.g.* (Lo and Wang, 2006)].

Although our working hypothesis is that a broadly distributed FSI pulse helps to suppress prepotent but currently inappropriate actions, some alternative possibilities should be explored in future work. These include a network reset, that facilitates the transition between ensembles representing distinct components of an action sequence components (Wickens and Arbuthnott, 1993; Carrillo-Reid et al., 2008), and a role in guiding striatal plasticity, as broad signals about overall population response can assist reinforcement-based learning (Urbanczik and Senn, 2009). Nonetheless, our results suggest a circuit arrangement in which specific complex patterns of information feed-forward through largely parallel, segregated striatal-pallidal channels, while less information-specific, divergent control signals flow in the opposite direction.

2.4 Methods

Experimental Procedures

2.4.1 Behavioral task and drug infusions

All animals were housed on a 12 hr:12 hr light/dark cycle, with experiments performed during the light phase. For daily training sessions, adult male Long-Evans rats (∼350 g) were placed in a recording chamber (MED-NPW-5L; Med Associates Inc., St Albans, VT, USA; modified to accommodate large headstages) with stainless steel grid floors, five nosepokes, a pellet dispenser, a speaker, and a video camera (Fig. 2.1a). Infrared photobeams detected the presence of the rat's nose at each nosepoke hole and the food port. Rats were initially trained to nosepoke an illuminated hole

to receive a 45mg sucrose pellet delivered to a receptacle in the rear of the chamber. Rats were then trained to hold in the nosepoke and wait for a brief burst of white noise to get a reward. The time delay to the white noise was gradually increased until the rats waited for 900-1250 ms for >85% of the trials. In the next phase of training (see Fig. 2.1b,c), rats waited for the white noise cue as before; however, now either a high (4 kHz) or a low (1 kHz) 250 ms tone was played during the hold period. The time between nose in and tone onset (pre-tone delay) varied between 250-350 ms. The white noise burst instructed the animals that they were free to choose one of the adjacent nosepokes. For the 1 kHz tone, trials were rewarded for leftward movements, while 4 kHz tones rewarded rightward movements. The total hold time required to correctly complete the trial was pseudo-randomly selected to be between 900-1250 ms (uniform distribution). If the rats failed to hold until the white noise burst, trials were aborted and a 10-15 s timeout began (houselight on). To discourage the development of a side preference, rats were cued to move in a given direction only if at least one of the three previous responses was to the opposite side. Intertrial intervals were selected pseudo-randomly from the range 15-30 s. Roughly 10% of the session consisted of "free-choice" catch trials, in which both tones were played simultaneously and left and right choices were each rewarded at p=0.5. Catch trials were not analyzed here. After each training session rats were fed 14g of standard chow, which kept them at approximately 90% of free feeding weight.

For drug infusions, six rats were trained to perform the behavioral task above, with the exception that the 1 kHz or 4 kHz tone indicated the end of the hold period (i.e. no separate "Go" cue). Once performance had asymptotically stabilized, a guide cannula was implanted unilaterally into the striatum (target coordinates AP +0.5, ML +3.5, DV 4.5 mm, including the additional 1mm ventral protrusion

of the infusion cannula) on either the left (n=3) or right side (n=3). After two weeks recovery rats resumed a series of behavioral testing sessions which included (on different days, in order) a mock injection (in which the cannula was connected to the infusion apparatus but without infusion); an injection of ACSF (0.5 μl over 5min; ion concentrations in mM: Na 150; K 3.0; Ca 1.4; Mg 0.8; P 1.0; Cl 155); a muscimol injection (0.5 μg/0.5 μl over 5 min); and another ACSF injection.

Electrophysiological data were obtained from four rats, each implanted with 21 individually drivable tetrodes [four 12.5 μm Ni-Cr wires twisted together; (Wilson and McNaughton, 1993)]. Tetrodes were directed toward the dorsal lateral striatum, the nucleus accumbens, the globus pallidus, and toward the primary motor cortex (M1, target region: +3.0mm AP, 3.0mm ML). All tetrodes were in the right hemisphere. Three skull screws were placed in contact with frontal, parietal, and motor cortical regions to record ECoG signals. Additional skull screws served as ground (posterior lateral skull ridge) or as a reference for LFP/ECoG signals (on the midline, approximately 1 mm posterior to lambda). Data acquisition was performed using a 96 channel system built around custom amplifiers and LABVIEW^{TM} software (National Instruments, Inc.). This system also acquired synchronized digital video images (640x480 pixels, 15 frames/s). Neural signals were recorded in wide-band (1 to 9,000 Hz) to reduce distortions of waveform shape (Wiltschko et al., 2008) and digitized continuously at 31,250 Hz. The digital signals that controlled the behavioral chamber and detected nose pokes were also sampled at the same frequency (32 μs resolution). Following implantation and one week of recovery, recordings were made for several weeks to several months during performance of the delayed choice task. At the end of the experiment, each tetrode site was marked with a small lesion by passing 25 μA of current for 10s. Following perfusion and Nissl staining, final

tetrode locations were mapped onto coordinates in a reference brain atlas (Paxinos and Watson, 2005) using Sqirlz Morph software (Xiberpix, Inc.), and the location of prior recording days were estimated from screw turns. Cells that were not unequivocally in the motor cortex, striatum, or globus pallidus were not included in analyses. To avoid introducing biases into the activity of neuronal subpopulations, we wished not to repeat analysis of the same cells. Thus, neurons were only included from one session for each tetrode, unless the tetrode had been moved by a minimum of 100μm between sessions.

2.4.2 Spike Sorting and Classification

Spike-sorting was performed manually using Offline Sorter (Plexon Inc, Dallas TX), following digital high-pass (512 Hz) filtering of the continuous data. Differences in the waveform size and shape across the four tetrode wires were used for separating single units. The reliability of spike cluster separation was quantitatively determined by the refractory period in the auto-correlograms (Harris et al., 2000). Across all cells in our database, the mean proportion of inter-spike-intervals <1 ms was 0.00073, suggesting well-separated neurons. Once spike times were obtained for each single unit, the mean wide-band waveforms were obtained simply as a spike-triggered average of the wide-band continuous signals. Striatal cells were further classified as either a putative medium spiny cell (MSN), fast-spiking interneuron (FSI), or an "other" presumed interneuron (O) based on three distinct clusters found in a scatter plot of two measurements of the wide-band spike waveform: (1) the peak width at one-half maximum (FSI: 50-200 μs; MSN: 150-450 μs; O: 200-300 μs), and (2) the time from peak to valley (FSI: 100-455 μs; MSN: 560-1500 μs; O: 300-550 μs). All cells are shown with negative voltage upward. Cells that were inverted (n=20) or did not show a clear valley (n=14) were not classified.

2.4.3 Data Analysis

All analyses were performed using MATLAB (Mathworks, Inc.; Natick, MA) or SPSS (SPSS, Inc.; Chicago, IL). To measure the proportion of time spent in long inter-spike intervals (ISIs), $\text{Prop}_{\text{ISI>Xs}}$ (Schmitzer-Torbert and Redish, 2008), for each spike train we found all ISIs which exceeded a criterion (here $X = 2$ s), summed those ISIs, and divided by the total session time. We characterized each neuron's responsiveness to task events by constructing peri-event time histograms (PETHs) around each of the eight events shown in Fig. 2.1. For each PETH, we analyzed a 3 s window with a bin size of 30 ms, followed by smoothing by a 3-point moving average. To restrict the analysis to cells active during the task, we adopted an inclusion criterion that the peak of at least one PETH must be greater than 5 Hz (given the bin size, this roughly corresponds to a minimum of 1 spike every seven trials, or more spikes on fewer trials). The peak firing rate for each neuron was estimated from the maximum bin across all eight PETHs, and the peak response period was determined as the time between half-peak maximums from the baseline firing rate (horizontal line, Fig. 2.5). For Z-score-based analyses, PETHs were normalized by subtracting the mean firing rate from each time bin, dividing by the standard deviation of firing from the session, and taking the absolute value. For contraversive and ipsiversive PETHs (Fig 2.13), a 2 s window was used centered on the choice execution event, and the results were rank-ordered by peak time of contralateral, then ipsilateral trials. The selectivity index (SI) was derived from these contraversive and ipsiversive PETHs (Fig. 2.13c). The SI of the nth bin was calculated by:

$$\text{SI}_\text{n} = \frac{\text{Ipsi}_\text{n} - \text{Contra}_\text{n}}{\text{Ipsi}_\text{n} + \text{Contra}_\text{n}},$$

and we report the maximum SI across all 30 ms bins.

2.4.4 Statistical Analysis

For the muscimol manipulation data, we used a mixed-model ANOVA with a subject factor (RAT), and repeated measures factors of CUE (ipsi vs. contra), ACTION (ipsi, contra, error), and SESSION (MOCK, ACSF1, MUSC, ACSF2). The OUTCOME (proportion of trials selecting the action) was our dependent variable. The analysis indicated significant main effects of ACTION (F=352.1, df=2, P<0.001) but not of SESSION or CUE (F=0.0, P=1 for both). The results showed a significant SESSION × ACTION interaction (F=17.7, df=6, P<0.001), a significant CUE × ACTION interaction (F=531.1, df=2, P<0.001), and a significant three-way interaction of SESSION × CUE × ACTION (F=10.107, df=6, P<0.001).

When comparing the proportion of cell classes, we used a two-sample test of proportions (Crewson, 2006) in which the standard error ($S_{p_1-p_2}$) is

$$S_{p_1-p_2} = \sqrt{\hat{p}\hat{q}}\sqrt{\frac{n_1+n_2}{n_1 n_2}}, \hat{p} = \frac{c_1+c_2}{n_1+n_2}, \hat{q} = 1-\hat{p},$$

where c_1 and c_2 are the number of occurrences in the two groups, and n_1 and n_2 are the total number in each group. We then computed the test statistic of the proportion difference,

$$Z = \frac{(c_1/n_1 - c_2/n_2)}{S_{p_1-p_2}}.$$

To correct for multiple comparisons, we simply multiplied the resulting p-value by 8 (the number of task events examined) and considered it significant if it remained below 0.05.

For the multiple regression analysis, we analyzed the residual component, $\epsilon(i)$, using the animal's direction of movement d(i), the location of the starting position of each trial p(i), the tone that played t(i), the trial outcome (correct/wrong) o(i),

the reaction time RT(i), the movement time MT(i), and the trial number n(i). For each bin from time t to $t + \Delta t$, the magnitude of firing rate, F(i), for the i^{th} trial were fitted by the following multiple regression model:

$$\mathrm{F}(i) = \beta_0 + \beta_d \cdot \mathrm{d}(i) + \Sigma_{n=1}^{N-1} \beta_{p_n} \cdot \mathrm{p_n}(i) + \beta_t \cdot \mathrm{t}(i) + \beta_o \cdot \mathrm{o}(i) + \beta_{RT} \cdot \mathrm{RT}(i) + \beta_{MT} \cdot \mathrm{MT}(i) + \beta_n \cdot \mathrm{n}(i) + \epsilon(i).$$

The regression slopes $\beta_0, \beta_d, \beta_{p_1}, \beta_{p_2}, \beta_t, \beta_o, \beta_{RT}, \beta_{MT}, \beta_n$ (for $N = 3$ positions) and their t-values were estimated by the REGSTATS functionof the MATLAB Statistical Toolbox. Analysis was performed $T = 380$ times using a sliding time window of $\Delta t = 100$ ms that stepped in 5 ms intervals from $t = -1$ s before execution of choice movement to $t = 1$ s after. For each neuron, the peak movement selectivity was defined as the maximum t-statistic for β_d across all T timesteps (Fig. 2.5, black rasters). Similarly, the peak position selectivity was defined as the maximum t-statistic of β_{p_1} and β_{p_2} across T timesteps. For Figs. 2.5, 2.13, 2.14 the time bins in which the t-statistic of β_d were significant are shaded gold for negative (contraversive) and green for positive (ipsiversive) selectivity. Confidence intervals for population PETHs were constructed using a resampling method, with 100 shuffled datasets. In order to obtain estimates that were not dominated by a few outlier cells, shuffling was performed within each cell's normalized PETH by randomizing the order of 30 ms bins. This preserved the peak event response of that cell while randomizing the time at which this peak contributed to the population PETH. A given bin within the population PETH was considered to be significant if the value from the real data was either higher or lower than at least 95 of the shuffled data sets.

For identification of putative monosynaptic connections, short-latency and short-duration sharp peaks in the cross-correlograms were used as described in (Csicsvari

et al., 1998). To reduce the effect of random interactions, a shuffled cross-correlation histogram was calculated by shifting the spike train of the second cell with a fixed (100 ms) time interval. This operation retains the internal dynamics of spiking for both trains while eliminating the causal relationship between them. The shuffled histograms were then subtracted from the originals. Peaks (1 ms bin width) within 3 ms of the center bin were defined as significant excitatory interactions when at least one of the bins exceeded 99.9^{th} percentile of the mean. Similarly, short-latency troughs were considered to be due to inhibition when at least two neighboring bins were $<0.1^{\text{th}}$ percentile of the mean. For cell pairs recorded on the same electrode, the artificial trough from 0 to 1 ms bin (created by our clustering method) was excluded from the analysis.

Acknowledgements

We thank members of the Berke laboratory for comments and assistance with experiments. GG was supported by a National Institutes of Health grant (R21-HD049842) awarded to D.R. Kipke. This project was funded by the University of Michigan, and by grants from the National Institutes of Health (R01-DA14318) and the Tourette Syndrome Association to JB.

CHAPTER III

β-Band (15-25Hz) Oscillation Synchronize in Basal Ganglia and Cortical Local Field Potential Activity of Rats During the Selection of Motor Movements

Abstract

Beta oscillations (15-25Hz, β_{20}) are exaggerated in the output of the Basal Ganglia of Parkinson's Disease (PD) patients. Mounting experimental evidence also point to β_{20} oscillations forming a signal to "stop" movements in behaving animals. However, it remains to be determined what the physiological role of β_{20} activity is in the suppression of movements in normal subjects. We designed a task to test the hypothesis that β_{20} is related to the suppression of a selected (but not yet executed) action sequence. Nine rats were trained to hold in a nose poke for $\sim$1s and move quickly to one of two side ports for a food reward. In 4 rats, we instructed the upcoming direction within the hold period. The remaining 5 were given the instruction at the end of the hold period. The onset of β_{20} synchronization was tightly coupled to the onset of the instructional tone, independent of when it was delivered; and desynchronized after the action had completed. Rats that were instructed early had a increase in β_{20} power during the hold period that was coherent across the striatum, globus pallidus and motor cortices. Well-isolated single units were found to be entrained to the β_{20} rhythm in all recorded regions during active participation in the task.

Taken together, these findings support the idea that beta oscillations accompany the selection of a motor plan, and remain until the plan has been executed.

3.1 Introduction

Synchronized oscillations are known to be correlated with many behaviors and brain states. Most investigations have focused on their role in the sensory domain, but oscillations in local field potentials (LFPs) are also apparent at multiple levels of the human and non-human motor system. In particular, rhythmic activity of Beta oscillations ($15-25$Hz, β_{20}) can be recorded from the motor cortex (Murthy and Fetz, 1992; Sanes and Donoghue, 1993; Baker et al., 1997), cerebellar system (Marsden et al., 2000), and basal ganglia (Brown et al., 2001; Levy et al., 2002; Silberstein et al., 2003; Courtemanche et al., 2003; Berke et al., 2004). The consequence of β_{20} is readily picked up in the peripheral nervous system through the synchronization of motor units within and between muscles (Farmer et al., 1993; Baker et al., 1997; Kilner et al., 1999). However, the function of such synchronization within the motor system remains less understood.

One consistent finding is that β-band LFP activity is related to a decrease in movement activity. During movements, β_{20} is desynchronized and replaced by higher frequency oscillations (Sanes and Donoghue, 1993; Donoghue et al., 1998). Conversely, patients with Parkinson's Disease (PD); characterized by muscle rigidity, the slowing (bradykinesia) or removal of movements (akinesia); have excessive β_{20} synchronization in the cortex, subthalamic nucleus and globus pallidus (Brown, 2006; Hammond et al., 2007). These exaggerated oscillations in PD are reduced with therapeutic interventions that restore motor abilities (Brown et al., 2001; Levy et al., 2002; Kühn et al., 2006). Applying stimulation at β_{20} frequencies directly to the cortex leads to

a slowing of movement in healthy (Pogosyan et al., 2009) and PD subjects (Chen et al., 2007). Taken together, these data point to β_{20} being related to a decrease in motor activity.

This inverse relationship between β_{20} activity and the processing of movements is, however, open to two very different interpretations. First, beta-band activity may just be a passive characteristic of the cortical network when not engaged in active processing. Pfurtscheller et al. (1996) demonstrated a post movement β_{20} synchronization following movements, which they interpreted as a sign of "idling" in the motor cortex. Alternatively, β_{20} activity may impair movement by actively promoting processes related to postural stability, including tonic holding contractions (Androulidakis et al., 2007). Supporting this view, Gilbertson et al. (2005) found that physiological periods of beta synchrony are associated with a cortical state in which a postural set is reinforced and the speed of new movements impaired.

Here, we test the hypothesis that β_{20} activity becomes synchronized during the active suppression of a selected motor action. By varying the delivery of an instructional cue in a simple reaction time task, we aim to separate the difference between holding the suppression of a selected motor plan, and holding while awaiting the selection of a motor plan.

3.2 Methods

3.2.1 Behavioral task

All animals were housed on a 12 hr:12 hr light/dark cycle, with experiments performed during the light phase. For daily training sessions, adult male Long-Evans rats (∼350 g) were placed in a recording chamber (MED-NPW-5L; Med Associates Inc., St Albans, VT, USA; modified to accommodate large headstages) with stainless steel grid floors, five nosepokes, a pellet dispenser, a speaker, and a video camera.

Infrared photobeams detected the presence of the rat's nose at each nosepoke hole and the food port. Rats were initially trained to nosepoke an illuminated hole to receive a 45mg sucrose pellet delivered to a receptacle in the rear of the chamber. Rats were then trained to hold in the nosepoke and wait for a brief burst of white noise to get a reward. The time delay to the white noise was gradually increased until the rats waited for 900-1250 ms for >85% of the trials. In the next phase of training, rats were divided into two groups. In the first set of rats, the white noise "go" cue was replaced with a high (4 kHz) or a low (1 kHz) 250 ms tone. The tone instructed the animals that they were free to choose one of the adjacent nosepokes. For the 1 kHz tone, trials were rewarded for leftward movements, while 4 kHz tones rewarded rightward movements. The total hold time required to correctly complete the trial was pseudo-randomly selected to be between 900-1250 ms (uniform distribution). If the rats failed to hold until the white noise burst, trials were aborted and a 10-15 s timeout began (houselight on). To discourage the development of a side preference, rats were cued to move in a given direction only if at least one of the three previous responses was to the opposite side. Inter-trial intervals were selected pseudo-randomly from the range 15-30 s. A second set of rats was trained on a similar task, however the high or low tone was played 250-350 ms (uniform distribution) into the 900-1250 ms hold period, and a brief white noise burst indicated the end of the hold period. Roughly 10% of the session consisted of "free-choice" catch trials, in which both tones were played simultaneously and left and right choices were each rewarded at p=0.5. Catch trials were not analyzed here. After each training session rats were fed 14g of standard chow, which kept them at approximately 90% of free feeding weight.

3.2.2 Recordings and Histology

Electrophysiological data were obtained from nine rats, each implanted with 21 individually drivable tetrodes [four 12.5 μm Ni-Cr wires twisted together; (Wilson and McNaughton, 1993)]. Tetrodes were directed toward the dorsal lateral striatum, the nucleus accumbens, the globus pallidus, and toward the primary motor cortex (M1, target region: +3.0mm AP, 3.0mm ML). All tetrodes were in the right hemisphere. Three skull screws were placed in contact with frontal, parietal, and motor cortical regions to record ECoG signals. Additional skull screws served as ground (posterior lateral skull ridge) or as a reference for LFP/ECoG signals (on the midline, approximately 1 mm posterior to lambda). Data acquisition was performed using a 96 channel system built around custom amplifiers and LABVIEW^{TM} software (National Instruments, Inc.). This system also acquired synchronized digital video images (640x480 pixels, 15 frames/s). Neural signals were recorded in wide-band (1 to 9,000 Hz) to reduce distortions of waveform shape (Wiltschko et al., 2008) and digitized continuously at 31,250 Hz. The digital signals that controlled the behavioral chamber and detected nose pokes were also sampled at the same frequency (32 μs resolution). Following implantation and one week of recovery, recordings were made for several weeks to several months during performance of the delayed choice task. At the end of the experiment, each tetrode site was marked with a small lesion by passing 25 μA of current for 10s. Following perfusion and Nissl staining, final tetrode locations were mapped onto coordinates in a reference brain atlas (Paxinos and Watson, 2005) using Sqirlz Morph software (Xiberpix, Inc.), and the location of prior recording days was estimated from screw turns. Electrophysiological data that were not unequivocally in the motor cortex, striatum, or globus pallidus were excluded from analyses. To avoid introducing biases into the activity of neuronal

subpopulations, we wished not to repeat analysis of the same cells. Thus, electrophysiological data were included for analysis only if the tetrode had been moved by a minimum of 100μm from a previously analyzed session.

3.2.3 LFP Analysis

The analysis for this study was performed using custom software developed in the Matlab (MathWorks, Natick, MA; Version 7.9) environment. Local field potentials were extracted from the original wide-band recording files by down-sampling the 31,250 Hz data by a factor of 61, resulting in a new sampling rate of ∼512.3 Hz. Initial visual scanning allowed for the removal of noise events. Sites within the striatum were classified into nucleus accumbens (NAcc) if locations were within the core [as defined by Paxinos and Watson (2005)], and caudoputamen (CPu) if striatal sites were further dorsolateral.

To estimate the length of beta oscillation events in the LFP data, we sought out epochs of high power in the 15-25 Hz range using a method similar to the "temporal spectral evolution" (Salmelin and Hari, 1994) as follows. The LFP signal was zero-phase bandpass filtered from 15-25 Hz using digital 102-sample length equiripple FIR filter with odd symmetry. These data were then squared (Siapas and Wilson, 1998) and smoothed using a 20-point moving average filter. Peaks in the time line of the squared signal were detected when the data became greater than a defined amplitude threshold for a minimal episode length of 150 ms (three cycles). The threshold level was chosen heuristically and corresponded to the 60th percentile of the range of values of the squared signal.

To determine the overall rhythmicity of the field potentials, we constructed full-session power spectra and mean triggered scalograms around task events of interest. Power spectra were calculated using the 512-point discrete Fourier transform (DFT).

For time varying spectral analysis, we computed scalograms using the Gabor transform (Antoine et al., 2000). While the DFT is a useful tool for understanding the frequencies present in a digital signal, it can only indicate whether a frequency is present in the signal, but not when it began or for how long it lasted. To examine transient brain rhythms, there are many different time-frequency analysis techniques may be used to visualize a signal's time-varying frequency content. We chose the Gabor analysis, which is optimized for simultaneous localization of spatial and frequency information. In this transform, each frequency band is created by convolving the original signal with a frequency-dependent Gabor kernel, which is a complex sinusoid windowed by a Gaussian. The frequency of a kernel's sine wave determines its pass band, or center frequency. By constructing a bank of many Gabor kernels with center frequencies evenly spanning the spectral range of interest (e.g. 120 kernels between 1-120 Hz), we may reveal the sub-second timing of transient rhythms in LFP data.

For an LFP segment of length n, the Gabor kernel ($\mathbf{G}_{f,t}$) was defined for frequency f (in Hz) as the $1 \times n$ array:

$$\mathbf{G}_{f,t} = \frac{1}{\sqrt{2\pi\sigma^2}} e^{\frac{x^2}{2\sigma^2}} e^{2i\pi f \mathbf{x}}, \tag{3.1}$$

where f ranges from 1 to 120, σ controls the width of the Gaussian ($\sigma = 0.2$ in our analysis), and $\mathbf{x}$ is a $1 \times n$ time array ranging from -1 to 1. We then computed our scaleogram as

$$\mathbf{S}_{f,t} = \mathbf{G}_{f,t} * \mathbf{LFP}_t, \tag{3.2}$$

where $\mathbf{S}_{f,t}$ is a $120 \times n$ matrix, $\mathbf{G}_{f,t}$ is defined in 3.1, and $\mathbf{LFP}_t$ is the LFP time-domain data of length n.

Average event-triggered scalograms were made for each channel and for each behavioral event within the trial. The scalograms were computed one-second around 50 randomly selected events of interest in the trail (*e.g.* 50 x "Nose In" events) and the results were averaged together. To calculate the baseline activity of a channel we performed the same procedure, only we averaged over 100 randomly selected epochs within the recording session. When displaying the normalized scalograms in figure 3.4, we divided the mean power of the event-triggered scalograms between $15-25$ Hz, by the mean power of the baseline-scalograms between $15-25$ Hz. Confidence intervals for population normalized β-band activity were constructed using a resampling method, with 100 shuffled datasets. In order to obtain estimates that were not dominated by a few outlier channels, shuffling was performed within each channel's normalized β_{20} activity by randomizing the time in which each bin appeared. This preserved the peak amount of beta activity for that channel while randomizing the time at which this peak contributed to the population average. A given time bin within the median population beta was considered to be significant if the median value from the real data was either higher or lower than at least 95 of the shuffled data sets.

Coherence analysis, which measures whether oscillations in two sites show a consistent phase relationship (Nunez and Srinivasan, 2006), was estimated using Welch's method from from 1 to 120 Hz with *mscohere* in Matlab. Comodulation analysis, which measures the correlation in power of frequencies across trials, was calculated as follows. For each trial, a power spectra was calculated using Welch's method for each wire during the hold time (events 2-5). For every pair of LFP channels, we calculated the trial-by-trial correlation coefficient (*corrcoef* in Matlab) of the power in each frequency from the first channel, to that of the power in each frequency of

second channel. The results of this analysis produced a 120 × 120 matrix ranging from -1 to 1, and were plotted using a pseudo-color heat map.

3.2.4 Unit Analysis

Spike-sorting was performed manually using Offline Sorter (Plexon Inc, Dallas TX), following digital high-pass (512 Hz) filtering of the continuous data. Differences in the waveform size and shape across the four tetrode wires were used for separating single units. The reliability of spike cluster separation was quantitatively determined by the refractory period in the auto-correlograms (Harris et al., 2000). Across all cells in our database, the mean proportion of inter-spike-intervals <1 ms was 0.00021, suggesting well-separated neurons. Once spike times were obtained for each single unit, the mean wide-band waveforms were obtained simply as a spike-triggered average of the wide-band continuous signals. Striatal cells were further classified as either a putative medium spiny cell (MSN), fast-spiking interneuron (FSI), or an "other" presumed interneuron (O) based on three distinct clusters found in a scatter plot of two measurements of the wide-band spike waveform: (1) the peak width at one-half maximum (FSI: 50-200 μs; MSN: 150-450 μs; O: 200-300 μs), and (2) the time from peak to valley (FSI: 100-455 μs; MSN: 560-1500 μs; O: 300-550 μs). All cells are shown with negative voltage upward. Cells that were inverted or did not show a clear valley were not classified (n=62) .

To determine the phase relationship between single cells and beta activity, we used a Hilbert transform to get the instantaneous phase of the detected 15-25 Hz intervals described above. Each action potential occurring between two successive peaks was assigned a corresponding phase, and the cycles were superimposed to compute a mean phase. The significance for each unit and for each population were calculated using standard circular statistical methods (Durand and Greenwood, 1958; Zar, 1974).

3.3 Results

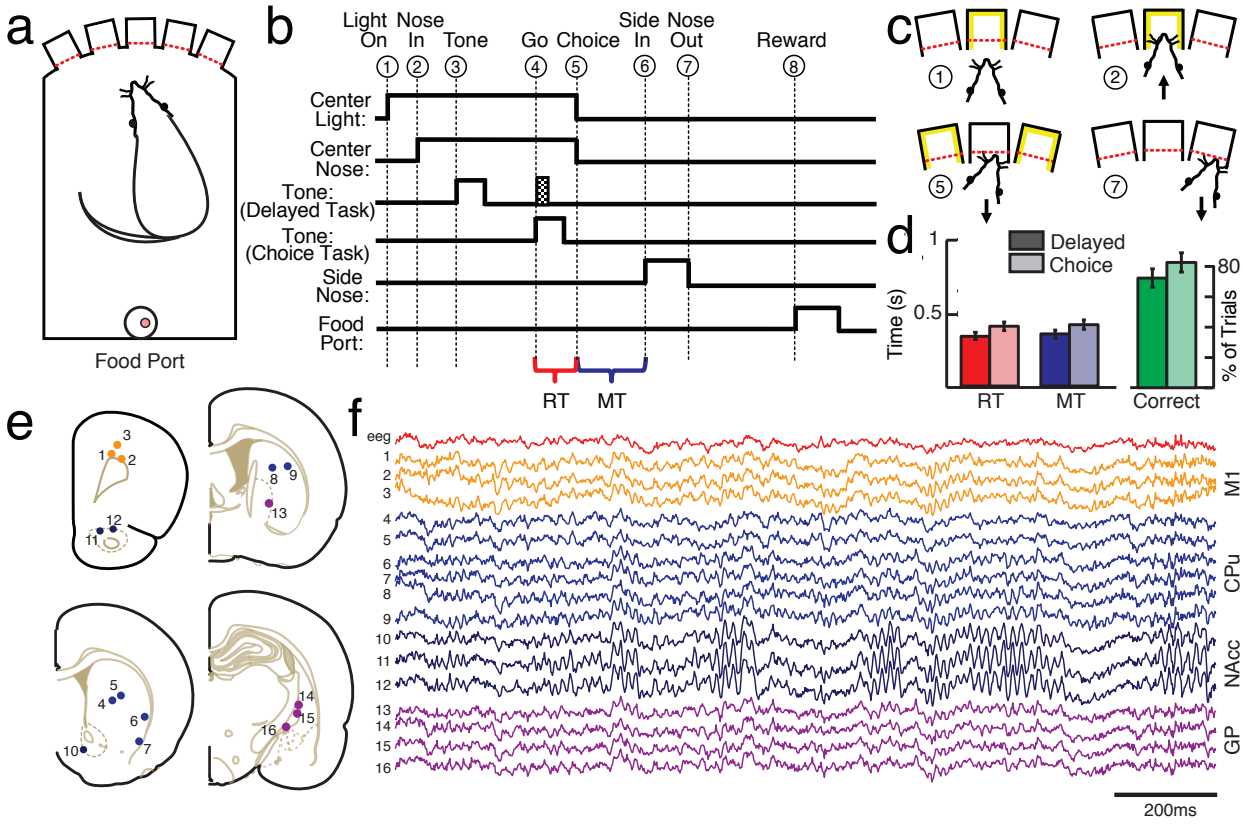

Figure 3.1: Behavioral Task and Recordings. a) Depiction of operant training box showing the position of the 5 nosepokes and food receptacle. b) Design of early (delayed task) and late (choice task) tone delivery versions of the reaction time tasks. Seven events are marked with numbers that were used in this analysis, a subset of events are shown in (c). d) Behavioral results of the two tasks. Time measurements as in (b): RT=Reaction Time, MT=Movement Time. Error bars, SEM. e-f) Example of tetrode recordings from one session. e) The location of one wire from each tetrode f) Downsampled LFP data, numbers indicate position in (e).

We recorded single unit and LFPs from up to 21 tetrodes implanted in the striatum, globus pallidus, and motor cortex to compare activity across two choice reaction tasks (Fig 3.1a-f). In the first task (choice task), rats (n=5) were trained to hold still in a operant nosepoke until a tone (1kHz or 4Khz) instructed them to move to an adjacent nosepoke. Rats were rewarded with a sugar pellet in the rear of the cage (Fig. 3.1a) if the selected nosepoke was in the direction that the tone instructed

(left=1kHz, right=4kHz). A second group of rats (n=4) performed a similar task (delayed task), however the instructional tone played during the hold period, and the rats were signaled to begin movements with a white noise burst. The total hold time was kept consistent across both groups of subjects (Fig. 3.1b 900−1250ms, uniform distribution). Across the two tasks, the success rate was consistent (Fig. 3.1d) and above chance [81.7% (SD 6.3%), choice; 71.1% (SD 6.1%), delayed] but the reaction times and movement times were faster ($p < 0.05$, t-test) for the delayed group [RT: 407ms (333ms) choice vs. 339ms (288ms) delayed; MT: 418ms (155ms) choice vs. 355ms (172ms) delayed].

Local field potentials were recorded from 425 locations (237 choice task, 188 delayed task) across 68 recording sessions (32 choice, 36 delayed), see table 3.1. Across these recording sessions, prominent episodes of β_{20} $(15 - 25$ Hz) and γ_{50} $(45 - 55$ Hz) oscillatory activity could be seen in some striatal and GP locations (Fig. 3.2a). Low frequency (< 10 Hz) oscillations were also present in the signal but were not analyzed in this study. The location of the recording electrode was a determining factor on the ratio of β_{20} to γ_{50} $(\frac{P(\beta_{20})-P(\gamma_{50})}{P(\beta_{20})+P(\gamma_{50})})$. Figure 3.2b shows the distribution of ratio of $\beta_{20} : \gamma_{50}$ power as a function of location. While the dorsoventral axis showed weak correlations, the mediolateral and anterioposterior axes had significant trends ($p< 0.01$, regression t-test). When measured on a rotated axes (Voorn et al., 2004), defined by the distance to the striatal origin point (Berke et al., 2004), the ratio favored γ_{50} ventromedially. If we analyzed the power spectrum only during the hold period (events 2-5), the β_{20} peaks became more defined (Fig. 3.2c). For this study, we concentrated on β_{20} both during the hold period and task performance.

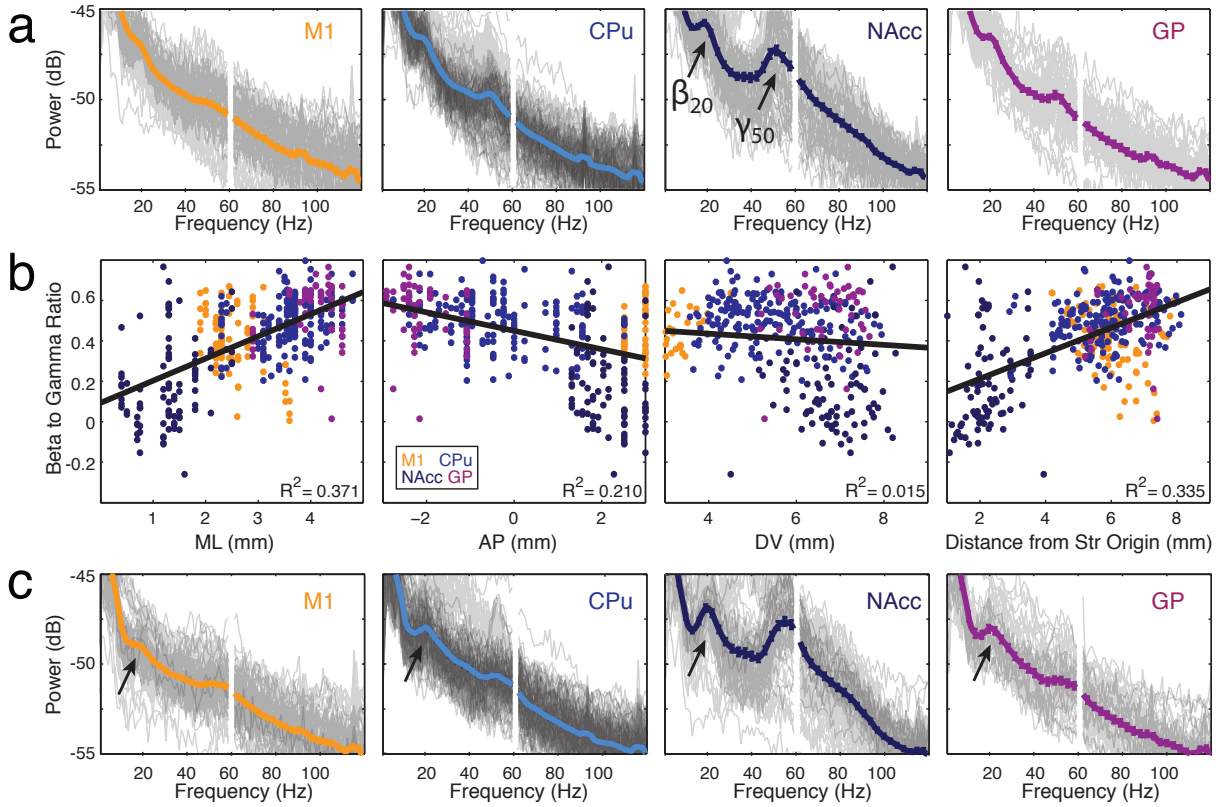

Figure 3.2: Characteristics of Local Field Potentials in four brain regions. a) Individual LFP channel (thin lines) and mean of all channels (colored) for all channels recorded in the two versions of the task. Error bars = SD b) Spatial distribution of β_{20} to γ_{50} ($\frac{P(\beta_{20})-P(\gamma_{50})}{P(\beta_{20})+P(\gamma_{50})}$) across 4 axes: (left) mediolateral (left center) anterioposterior (right center) dorsoventral (right) distance from the origin from the striatum (AP 3.13 mm anterior, ML 0, DV 8 mm below bregma). c) Same as (a), however power spectra are from hold period only (events 2-5). Arrows indicate peaks in β_{20} in all regions.

3.3.1 15-25 Hz oscillation epochs during task performance have similar properties across regions.

We first investigated the properties of β_{20} during the task performance (events 2-7). We identified periods of high β_{20} activity which revealed that 15-25 Hz oscillations occurred in brief epochs lasting ~250 ms, and were not a continuous part of the LFP structure during the task. Figure 3.3a shows a typical β_{20} epoch from a striatal LFP recording site, along with a representative sample of 15 overlaid epochs. The trial-to-trial consistency can be seen in these epochs, as well as in 3.3b which depicts all

of the hold periods for the session. Identified β_{20} epochs are highlighted in red. The mean epoch length for this session was 255.5 ms (SD 108.5 ms), which was typical of CPu locations (see Table 3.1). Note that the β_{20} events, while brief, occurred with high probability (0.92) during a single trial and were likely to occur multiple times within a given trial (averaging 1.62 epoch/trial in this example). Across the regions recorded, epochs had similar properties (Table 3.1).

	n	Cycles/Epoch	Epochs/Trial	% of Trials
M1 (Choice)	59	4.54 (0.28)	1.34 (0.31)	78.3 (10.5)
M1 (Delayed)	34	4.59 (0.46)	1.61 (0.36)	80.7 (10.0)
M1	93	4.56 (0.35)	1.44 (0.35)	79.2 (10.3)
CPu (Choice)	113	4.81 (0.35)	1.42 (0.21)	82.3 (6.0)
CPu (Delayed)	77	4.96 (0.41)	1.52 (0.23)	85.7 (6.4)
CPu	190	4.87 (0.38)	1.46 (0.22)	83.7 (6.4)
NAcc (Choice)	35	5.34 (0.24)	2.16 (0.18)	96.2 (2.0)
NAcc (Delayed)	56	5.33 (0.55)	2.01 (0.32)	92.4 (4.8)
NAcc	91	5.34 (0.45)	2.06 (0.28)	93.8 (4.4)
GP (Choice)	30	4.97 (0.31)	1.58 (0.19)	86.3 (5.0)
GP (Delayed)	21	5.13 (0.59)	1.62 (0.19)	86.6 (4.6)
GP	51	5.04 (0.45)	1.59 (0.19)	86.4 (4.8)
Total Choice	237	4.84 (0.40)	1.44 (0.25)	82.8 (6.4)
Total Delayed	188	5.02 (0.54)	1.64 (0.36)	86.3 (7.0)
Total	425	4.92 (0.48)	1.53 (0.32)	84.3 (7.0)

Table 3.1: *Properties of β_{20} epochs during trial performance.* Data are presented as mean (SD). % of Trials indicates the percentage of trials that contained at least one epoch during task performance (event 2 − event 7).

3.3.2 LFP oscillations during delayed task have higher β_{20} power

To access the temporal evolution of the β_{20} epochs during behavior, we constructed event triggered scalograms using the Gabor transform to reveal the frequency dynamics within the LFP. Figure 3.3c-d shows the average results of seven events from the two behavior sessions. These channels were recorded from different animals, but in similar striatal locations (1.4 mm separation across subjects) during the (3.3c)

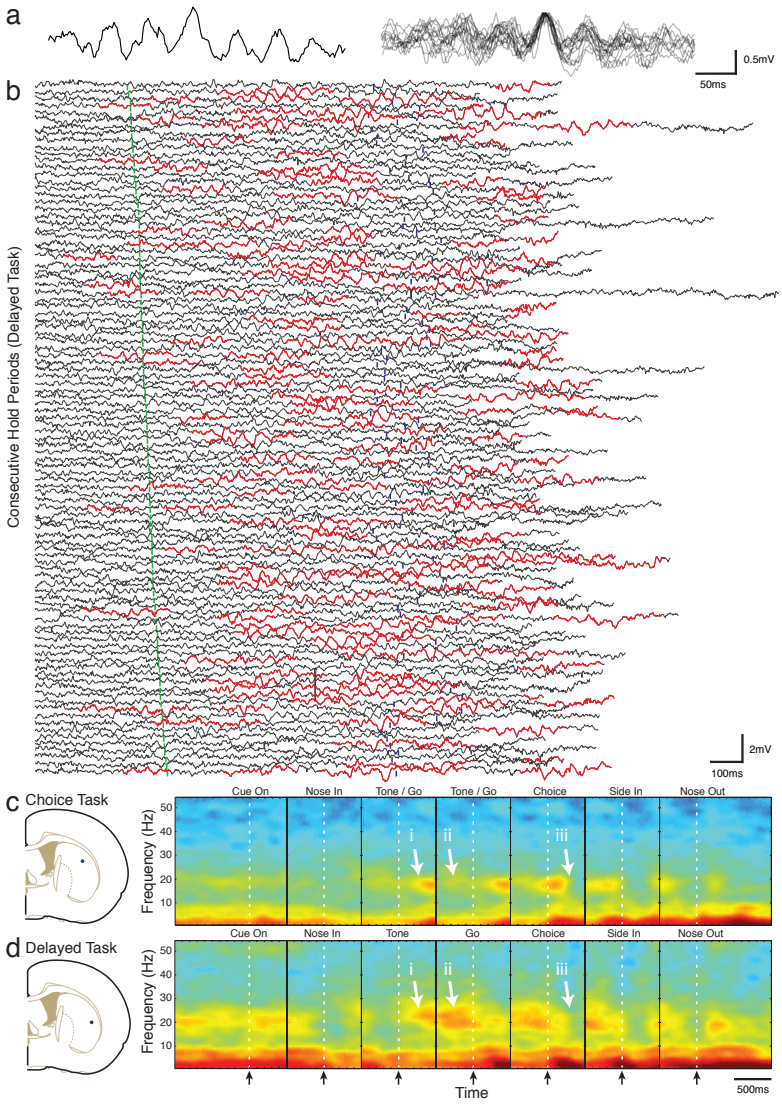

Figure 3.3: Beta Oscillations during hold period. a) Left: An isolated ∼7 cycle β_{20} epoch from a striatal recording site [see (c) for location] during task performance. Right: Overlay of 15 randomly selected epochs aligned to maximum peak. b) LFP traces from one recording site across all trial hold periods, event (2) to event (5), during the delayed task. Identified β_{20} epochs are highlighted in red. Trials are ordered by tone onset (green ticks). Blue ticks indicate "go" cue. In the data shown, the mean (SD) cycles/epoch was 5.11 (2.17) Range: 3.01−15.38. Ninety-two % of trials had at least one epoch, averaging 1.62 epoch/trial. c-d) spectral activity during the hold period from striatal channels during the (c) choice and (d) delayed task. Left: estimated recording location. Data from (a-b) is from same channel and session shown in (d). Right: Event triggered scalograms from epochs 1−7. The "Tone/Go" event for the choice task is repeated twice for comparison. Note the prominent β_{20} activity after the instructional cue (i), the increase in β_{20} activity (ii) during the hold period, and similar β_{20} desynchronization timescale after the hold period. Channel locations of (c) and (d) are [3.7,-1.0,5.5] and [4.3,-1.3,6.8] in [ML,AP,DV] (mm), respectively.

choice and (3.3d) delayed versions of the task. Since choice trials did not have an early instructional tone (event 3), the go cue (event 4, the instructional tone for the choice task) is repeated for event 3 here for comparison. The initiation of the trial (and subsequent hold period) is shown in the second epoch, followed directly by the tone delivery event. In both tasks, an increase synchronization in the 15−25 Hz range occurs ∼200 ms after the instructional tone onset (i). While the β_{20} onset occurs simultaneously, it should be noted that these events occur at distinctly different times between the two tasks (temporal separation between 650 ms to 1,000 ms across tasks). This was clearly visible in the following set of go-event triggered analyses, where there was sustained β-band activity leading up to the "go" signal for the delayed task when compared to the choice (ii). Finally, the desynchronization of β_{20} occurs simultaneously in these examples (iii), occurring ∼300 ms after the onset of choice execution (event 5).

We next wished to determine if the sustained β_{20} after the tone onset, continuing to the choice movement was prominent feature across our database of LFPs. To visualize this, we generated a temporal response profile for each channel by calculating the average change in 15−25 Hz power in each of the 7 event-triggered scalograms. To compare across channels, we then normalized the β_{20} average of each event by dividing through with the session background β_{20} activity. Figure 3.4a shows this profile for all channels of each region for the choice version of the task. Each region is sorted by moment of peak β_{20} activity. The population activity can visually be compared to the delayed version of the task (Fig. 3.4b). The hold period is displayed as a gray band beneath the response profiles of both tasks. There are many similarities between the tasks. β-band activity desynchronized similarly after the nose in and side in events. A burst of β_{20} activity can be seen around the "nose

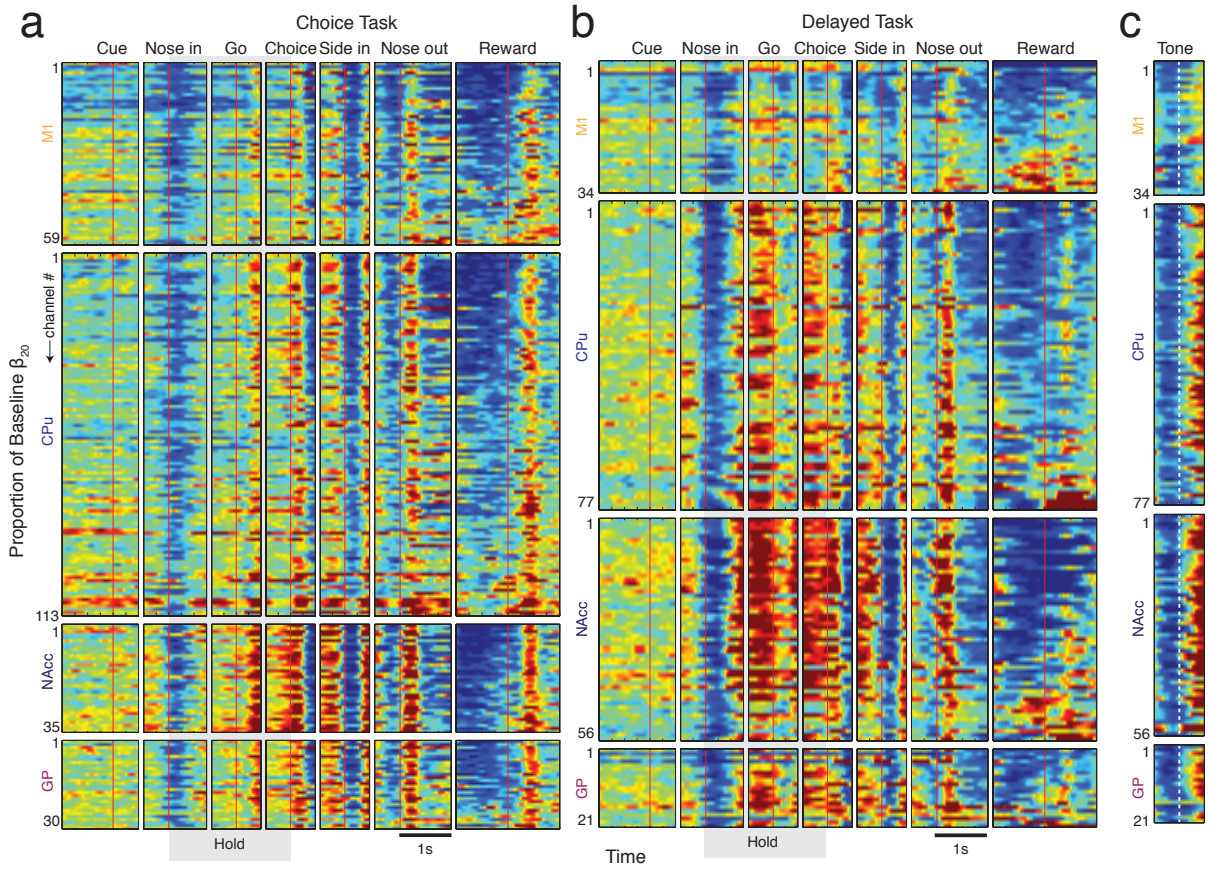

Figure 3.4: Population comparison of β_{20} power during performance of the choice and delayed tasks. a-b) Mean activity in the 15−25 Hz range of the scalograms triggered on the six shared events of the (a) choice and (b) delayed version of the task. Channels are normalized by dividing each event by the session mean of β_{20} power. For each channel, β_{20} power is shown by a color scale ranging from blue (a decrease in power by 40%) to red (a increase of β_{20} power by 60%). Within each class, channels are rank ordered by time of peak β_{20} activity. Note the increase in activity during the hold period of the delayed task. c) Tone event (event 3) from the delayed task. Same time scale as in (a-b).

out" event in both cases. The striking difference is during the hold period, where there is an increased probability of β_{20} activity during the delayed version of the task. The onset of this synchronization is visible in the "tone" event of the delayed task (Fig. 3.4c), and is similar to the "go" event of the choice task (Fig. 3.4a, third column).

A summary of the these analyses are shown in figure 3.5a. Here, the population data is plotted as the median of the normalized β_{20} activity. The choice task is shown

as a solid line, while the delayed task is a broken line. Bold lines indicate epochs where the β_{20} signal was outside of the 95% confidence intervals, generated across the population (see methods). Note that β_{20} significantly increased power during the hold period in all recorded brain regions (3.5a,i). We also wish to test if the timing of synchronization could be related to the onset of the movement instruction. Figure 3.5b compares the trial events of the two tasks in which the information about the upcoming movement (the tone) is presented. Note the timing that each region increases β_{20} significantly is the same across the two versions of the task (3.5b,iii). Similarly, the β_{20} decrease from both tasks is locked onto the end of the movement action (3.5a,ii).

3.3.3 LFP activity is highly synchronous across striatum, globus pallidus and motor cortex

We next looked to see if the increase in 15-25 Hz across the regions was the result of one process, or of multiple and independent activity. We used coherence and comodulation analysis to ask whether the phase and amplitude, respectively, were related across brain structures. Figure 3.6a shows one trial from a recording session of the delayed task. LFP traces from 6 recording channels and their respective locations are shown. Synchrony of a brief β_{20} epoch can be seen across the motor cortex and medial striatum. The presence of β_{20} can also be seen in the power spectra of the channels during the hold period (Fig. 3.6b, top). The coherence of this epoch across channels can be seen directly in the example (vertical lines), and across all trials by the increase near 20Hz in the coherency plots (Fig. 3.6b, bottom).

These results were confirmed across the entire population of recordings. Figure 3.7 shows the results of the comodulation and coherence analysis for pairs of simultaneously recorded LFPs during the delayed task. During the hold period, there is

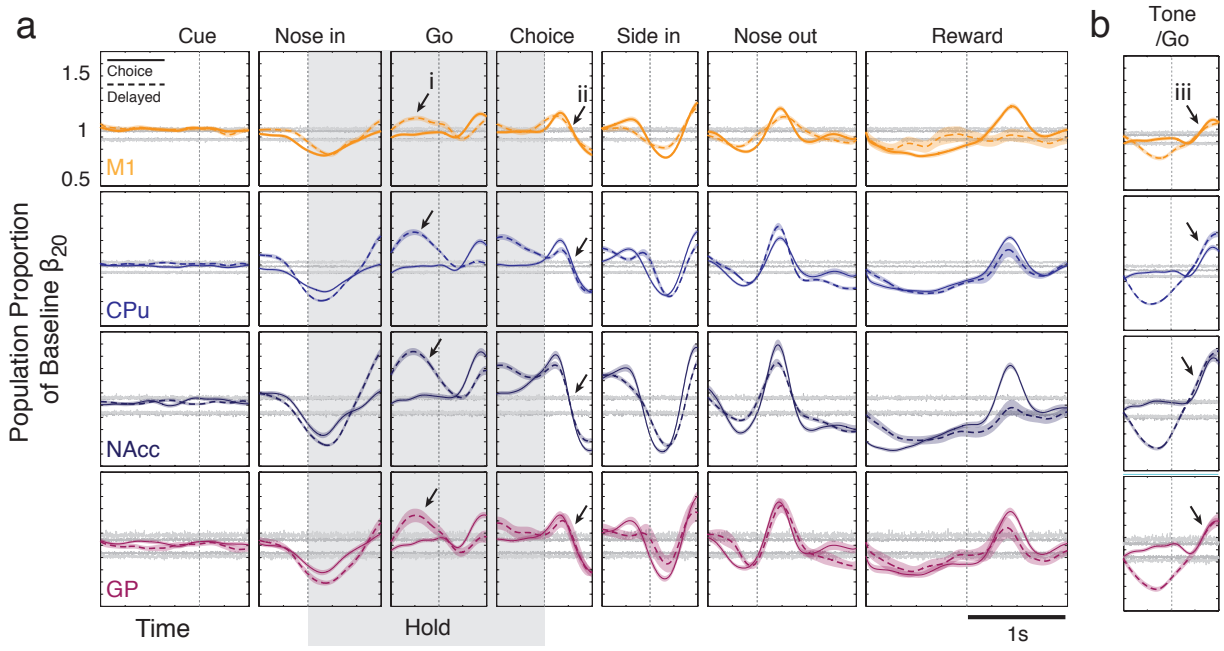

Figure 3.5: Population summary of β_{20} between the choice and delay tasks. a) Median 15-25 Hz oscillation activity of LFPs across 4 regions. Data are normalized by the baseline activity of each channel from the entire session. Solid lines = choice task. Broken line = delay task. Shaded region indicates ±1 s.e.. Gray vertical lines indicate the 95% confidence intervals of the population. Regions of data that fall outside the confidence intervals are bold. Note the significant increase in β_{20} activity across all regions before the go cue is only for the delayed version of the task (i). Beta becomes desynchronized 250 ms after choice execution onset (ii) in both tasks. b) Comparison of β_{20} onset. The choice task is aligned to "go" (4) event, while the delayed task is aligned to the tone (3) event. In both versions of the task, these events denote the onset of directional information. Note the close time evolution of the onset (iii) across the two tasks.

a prominent peak at ∼20 Hz both within and across regions in the comodulation analysis, indicating that the power spectra in the β-band is correlated between the two regions across trials. There were also ∼20 Hz peaks in the average coherence analysis indicating that the β_{20} oscillations were phase-locked across the regions. The choice task had similar results (data not shown).

3.3.4 The spiking activity of some neurons are entrained to β_{20} oscillations

To examine the relationship between the spike activity of single units and the LFP 15-25 Hz oscillations recorded simultaneously, we analyzed the relative timing

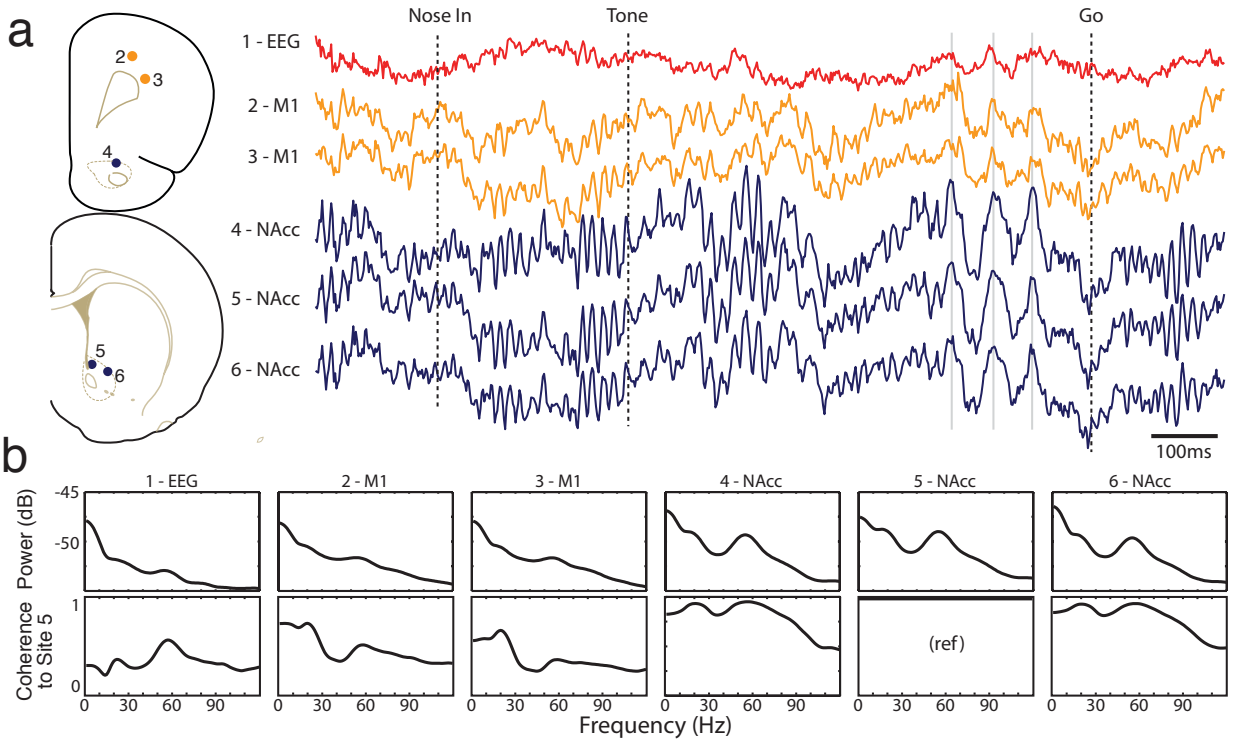

Figure 3.6: β_{20} is coherent across the striatum and motor cortex. a) Simultaneous field potentials recorded from 6 locations during the delayed task. Reconstructed locations are shown (left). Note epoch of β_{20} (15 Hz - 25Hz) is coherent across both structures and in EEG. Gray vertical lines indicate peaks of β_{20} from recording site 5. b) Power spectrum of recording sites in (a) during the hold period. c) Coherence between striatum and motor cortex. Striatum (site 5) is used as a reference. Note the increase at ~20 Hz in EEG and Motor Cortex.

of spikes and LFPs for the well-isolated neurons in all regions recorded (n=205, choice task; n=381, delayed task). The instantaneous phase angle of each β_{20} was calculated using the Hilbert transform, with the peaks of the oscillation denoted as 0°/360°. We plotted phase histograms (30° bin size) of the probability of spike firing as a function the phase within each oscillation cycle. We found 102 pairs of unit and field potentials that were significantly entrained to β_{20} activity ($p<0.05$, Rayleigh's test for circular uniformity). Of these pairs, 43 were from unique units. Four units (examples of cortical, striatal, GP neurons) are shown in Fig. 3.8 together with their respective phase histograms. Entrained units were synchronous with 2.37 (SD 1.66)

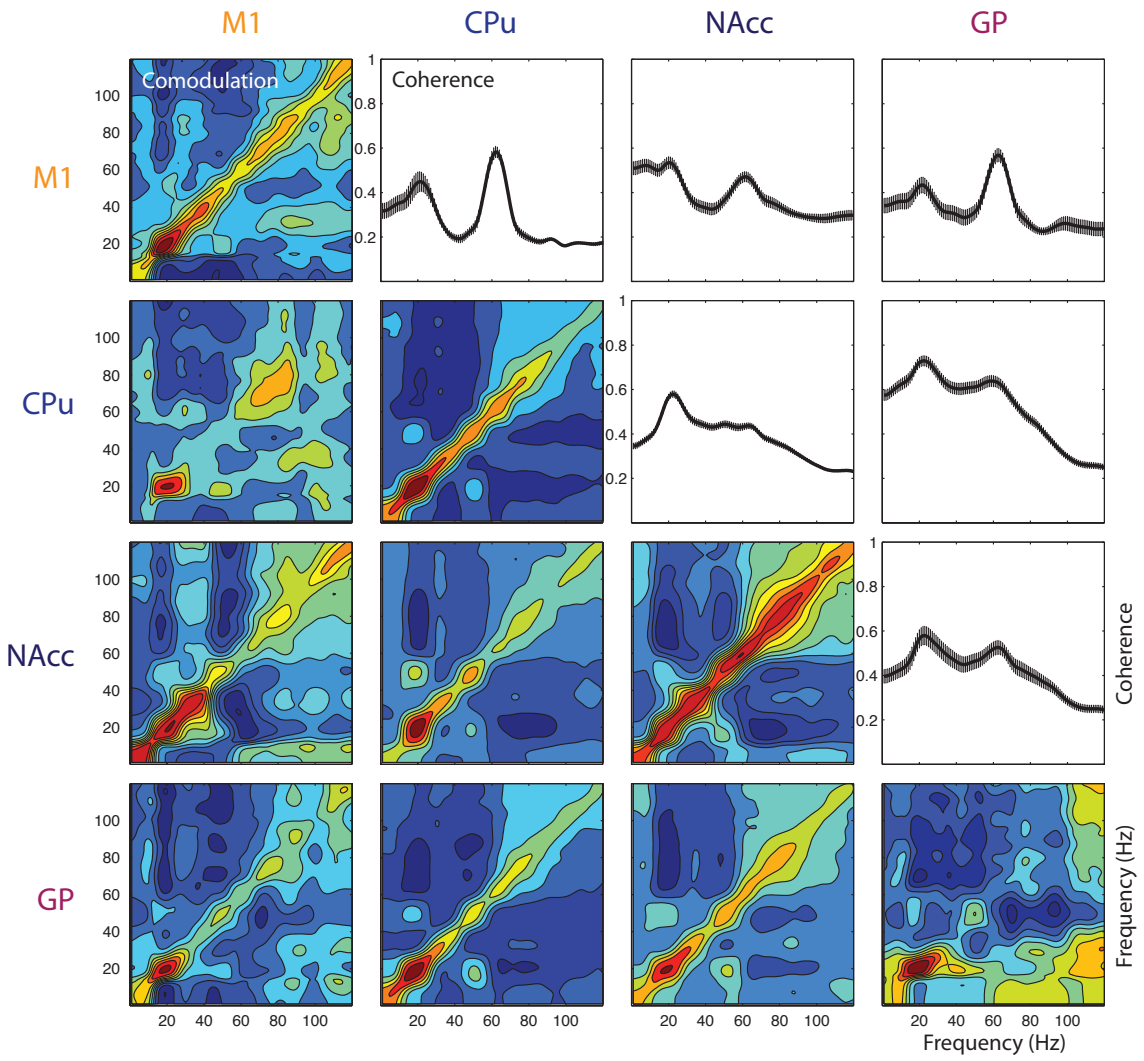

Figure 3.7: β_{20} is coherent across the Motor Cortex, Striatum, and GP during the hold period. Comodulation and coherence of all simultaneously recorded LFPs during the hold period of the delayed task. The diagonal shows the comodulation of pairs of LFPs in the same structure as labeled across the top and side. Beneath the diagonal are comodulation of pairs of LFPs across structures. Note the peak in comodulation at ∼20 Hz across all recorded structures. The average coherency from LFP channels recorded simultaneously are shown above the diagonal. Note the prominent increase in coherency at ∼20 Hz. Error bars indicate ± 1 s.e.m.

LFP channels on average. Of the 40/43 cells (93.02 %) that were entrained to a remote LFP site, the mean distance was 3.89 mm (SD 1.78 mm) apart, and ranged of from 0.74 mm up to 7.80 mm (GP LFP to M1 Unit). Cells entrained to multiple LFP sites (n=24) maintained a similar phase relationship across the LFP channels,

with a mean difference of $-2.22°$ (SD 9.88°).

The striatal neuron waveforms had distinct features that allow us to further classify them as either presumed FSI or MSN neurons. The waveforms and firing properties from these tasks closely resembled those seen in our previous studies of striatum (Berke et al., 2004; Berke, 2008). The overall phase preference of striatal unit firing was dependent of the cell-phenotype of unit recorded. Figure 3.9 shows the phase summary of the 43 unique cells that showed entrainment. For this analysis, we used the median value of phase angle when a cell was entrained to multiple LFPs. Motor cortical units (a), and GP units (d) did not show prominent phase preferences. However, the population of recorded FSI-units (b) maximized their firing just before the peak in beta activity at 0°. The population of MSN cells (c) recorded from the NAcc and CPu fired early in the phase cycle, just after the peak in the field potential oscillation.

3.4 Discussion

To better understand the significance of physiological β oscillations seen during decreases in movement activity, we adopted a combined 'macrocircuit/microcircuit' approach, investigating both relationships to activity in other, interconnected brain regions and the participation of specific neuron types. Consistent with the idea that β_{20} actively promotes neural processes related to maintaining tonic contractions, we found that synchronization occurred during the hold period only if the upcoming movement direction was known. However, this simple hypothesis does not explain why β_{20} is up-regulated over a similar timescale after the instructional tone in the choice task. Another hypothesis is that the β-band activity acts to "lock in" the current motor plan and avoid switching to others until such time as the current

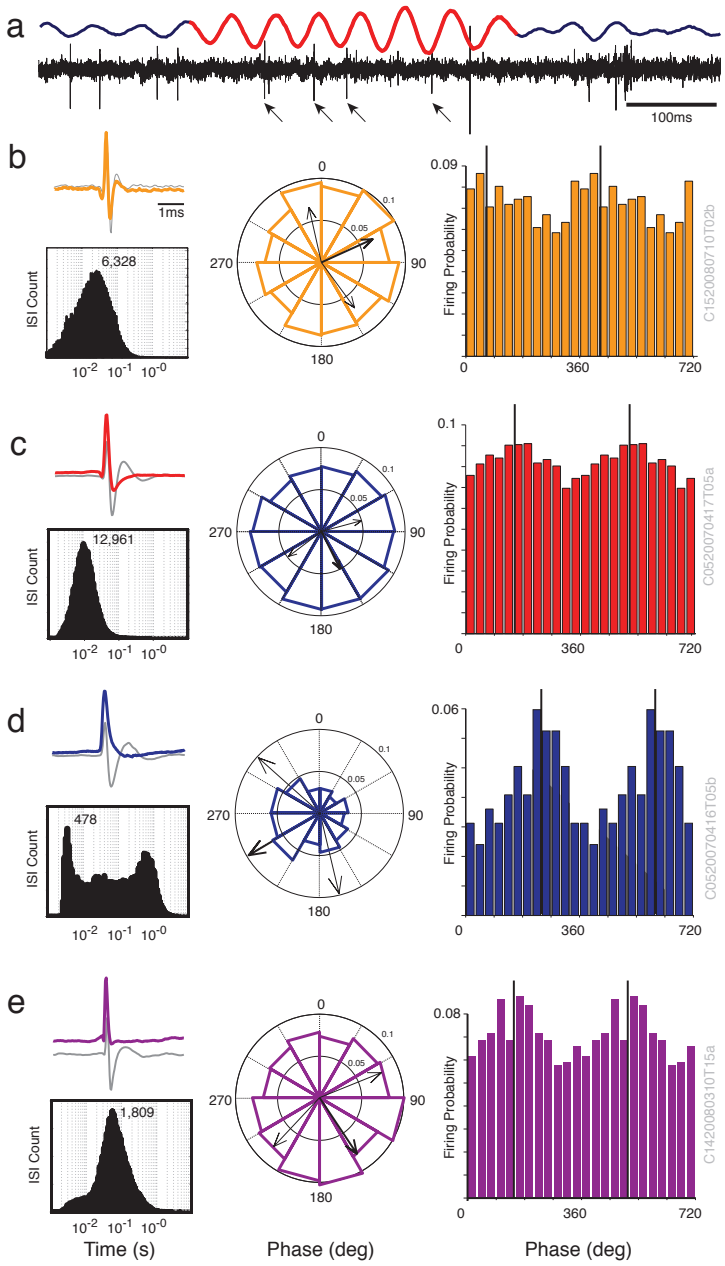

Figure 3.8: Unit phase locking during β_{20} epochs during hold period. a) Example of unit entrainment. (top) 15−25 Hz filtered LFP signal with β_{20} epoch in red. (bottom) 600−6,000 Hz filtered signal with entrained FSI (arrows). b-e) Top left: 5 ms trace of ∼unfiltered (colored) and Butterworth filtered (gray, 600−6,000 Hz) spike waveshapes. Bottom Left: ISI histogram, peak bin count is noted. Center-Right: Rose plot (center) and phase histogram (right) of spike firing probability as a function of phase of β_{20}. 30° bins. Arrows indicate mean phase (thick) and ± angular deviations (thin). (b) Motor cortical neuron. Rayleigh's z=3.198, p=0.041 (c) FSI from NAcc. Rayleigh's z=9.967, p=0.00005 (d) MSN from NAcc. Rayleigh's z=14.867, p<0.00001 (e)) GP unit. Rayleigh's z=3.953, p<0.019. Examples show LFPs from same tetrode as units. Peak= 0°, Valley= 180°.

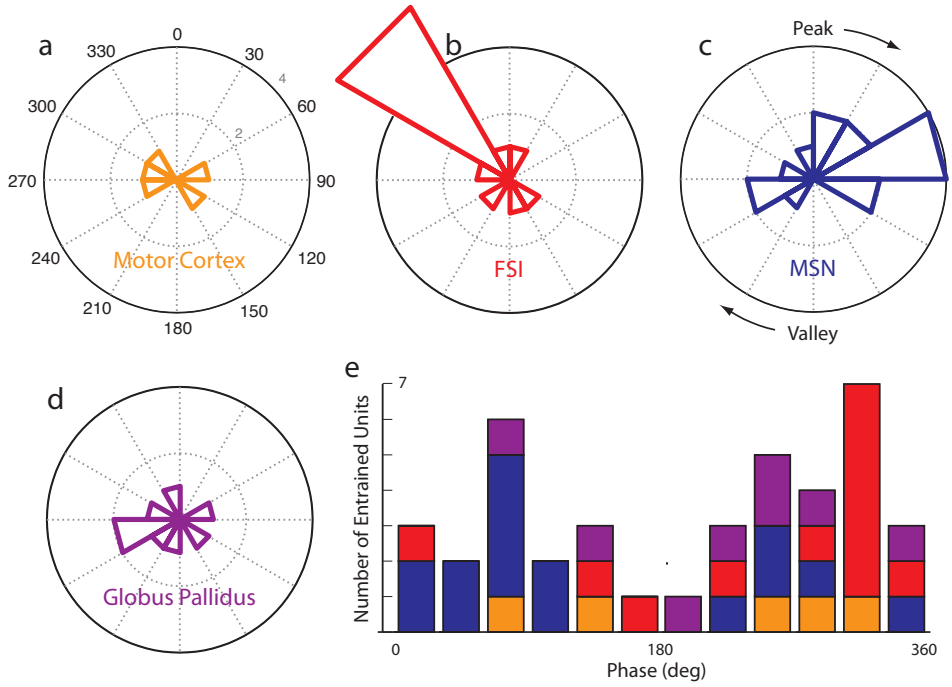

Figure 3.9: Summary of 43 β_{20} entrained units. a-d) Population phase histograms for (a) motor cortical, (b-c) striatal FSIs and MSNs respectively, and (d) GP units. Phase histogram with bin size of 30°. FSI cells tended to fire before the peak, while MSNs unit firing followed peaks in β_{20} oscillations. e) Summary of all cell classes are shown on one histogram. Same data as in (a-d). FSI and MSNs include cells from both the NAcc and CPu regions.

sequence is executed (Brown and Marsden, 1998). The desynchronization seen ∼250 ms after movement execution could act as a neural "clutch" allowing the system to change motor plans.

Consistent with this idea, studies of the motor cortex have found oscillations most often occur before movement onset and are much less evident during motor actions (Murthy and Fetz, 1996a,b; Sanes and Donoghue, 1993; Rubino et al., 2006). Prominent epochs of β_{20} in the precentral gyrus (Donoghue et al., 1998) and the caudate nucleus/putamen (Courtemanche et al., 2003) occur between the instructional pre-cue and the go cue during tasks where a preparatory periods are provided. The β_{20} oscillations have been shown to travel in propagating waves across the primary motor and dorsal premotor cortices and these waves contain information on the upcoming

(planned) movements (Rubino et al., 2006).

One question that needs addressing is where the β oscillations come from in our task. LFPs, in general, are somewhat ambiguous reflections of the underlying neuronal activity, and often better reflect synchronized subthreshold postsynaptic potentials (*i.e.*, correlated input activity) than suprathreshold output activity (Mitzdorf, 1985). In our data, β_{20} was the weakest in the motor cortex (Fig. 3.2a,c). The NAcc and CPu had higher power, but without the aide of a laminar structure that can amplify postsynaptic potentials and generate large oscillations, these appears to be a poor candidates. After dopamine loss, inappropriate reverberatory interactions between GP and subthalamic nucleus (STN) have been shown to actively promote the emergence of excessively synchronized β oscillations at the network level (Mallet et al., 2008). It is possible, therefore, that GP-STN reciprocal connections may be the source of the transient β_{20} epochs seen in our task. While figure 3.2 indicates GP had the highest power in the β_{20} range; further studies will need to determine if GP-STN contribute to physiological beta activity, or alternatively if they arise from deeper laminar structures such as the piriform cortex (Litaudon et al., 2008; Chapuis et al., 2009). Coherent β and γ activity have been in cortical and subcortical areas prior to olfactory-guided movement initiation (Hermer-Vazquez et al., 2007).

Wether or not these rhythms are volume conducted from adjoining structures, the fact that a number of cells were significantly entrained to β_{20} demonstrates they are participating in the information processing in our recorded regions. Interestingly, fast-spiking interneurons were shown to increase their spike probability earlier in the β_{20} phase than the MSN projection neurons. Suggesting that β epochs may serve a role in coordinating ensembles of striatal cells, perhaps maintaing the existing motor plan.

CHAPTER IV

Naïve Coadaptive Cortical Control

Abstract

The ability to control a prosthetic device directly from the neocortex has been demonstrated in rats, monkeys, and humans. Here we investigate whether neural control can be accomplished in situations where (1) subjects have not received prior motor training to control the device (naïve user) and (2) the neural encoding of movement parameters in the cortex is unknown to the prosthetic device (naïve controller). By adopting a decoding strategy that identifies and focuses on units whose firing rate properties are best suited for control, we show that naïve subjects mutually adapt to learn control of a neural prosthetic system. Six untrained Long-Evans rats, implanted with silicon micro-electrodes in the motor cortex, learned cortical control of an auditory device without prior motor characterization of the recorded neural ensemble. Single and multi-unit activity was decoded using a Kalman filter to represent an audio "cursor" (90ms tone pips ranging from 250Hz−16kHz) which subjects controlled to match a given target frequency. After each trial, a novel adaptive algorithm trained the decoding filter based on correlations of the firing patterns with expected cursor movement. Each behavioral session consisted of 100 trials and began with randomized decoding weights. Within 7±1.4 (mean±SD) sessions, all subjects

were able to significantly score above chance ($P<0.05$, randomization method) in a fixed target paradigm. Training lasted 24 sessions in which both the behavioral performance and signal to noise ratio of the peri-event histograms increased significantly ($P<0.01$, ANOVA). Two rats continued training on a more complex task using a bilateral, two target control paradigm. Both subjects were able to significantly discriminate the target tones ($P<0.05$, Z-test), while one subject demonstrated control above chance ($P<0.05$, Z-test) after 12 sessions and continued improvement with many sessions achieving over 90% correct targets. Dynamic analysis of binary trial responses indicated that early learning for this subject occurred on session 6. This study demonstrates that subjects can learn to generate neural control signals that are well suited for use with external devices without prior experience or training.

4.1 Introduction

There are over 250,000 cases of spinal cord injuries in the United States of America, with a majority of these injuries resulting in quadriplegia: the loss of movement and sensation in both the arms and legs (Lucas et al., 2004). Electroencephalographic (Lacourse et al., 1999) and functional magnetic resonance imaging (Shoham et al., 2001) studies have shown that spinal cord injured patients who imagine movements in their paralyzed limbs can still produce activation of the motor cortex, even after an extended period of time post trauma. Early studies demonstrating that single units of the motor cortex could be operantly conditioned (Olds, 1965; Fetz, 1969; Fetz and Finocchio, 1971) led Edward Schmidt to propose in 1980 that unit recordings from the motor cortex may constitute a viable control signal for external devices (Schmidt, 1980). Recent technology advances have enabled several groups to begin investigating the possibility of cortically controlled neural prostheses (Kennedy

et al., 2000; Serruya et al., 2002; Taylor et al., 2002; Carmena et al., 2003; Shenoy et al., 2003; Musallam et al., 2004; Olson et al., 2005). Many of the current cortical control paradigms consist of analyzing the relationship between cortical activity and measured motor parameters (Chapin et al., 1999; Wessberg et al., 2000; Sanchez et al., 2004). This known relationship is then used to transform the neuronal population signals into real-time prosthetic device movements. Unfortunately, such motor information cannot be obtained in patients with traumatic spinal lesions or neurological disorders which prevent movement.

Several groups have investigated the possibility of neural prostheses that could adapt to cell tuning properties (Taylor et al., 2002; Eden et al., 2004; Musallam et al., 2004). One study (Musallam et al., 2004) built a database of cortical responses to motor reaches that were subsequently used to decode brain-controlled tasks. Data collected during brain-controlled reaches to targets were used to continually update the database which eventually contained only brain-controlled trials. They reported that this could be done without leading to a loss in performance. Taylor *et al.* (Taylor et al., 2002) demonstrated that an *a priori* neural database was not required for an adaptive algorithm to allow brain-controlled movements. By starting with random tuning properties and allowing these estimates to be iteratively refined, subjects could make long sequences of three-dimensional movements using a brain-controlled cursor. These algorithms mutually adapted to the learning-induced changes in cell tuning properties, thus creating a *coadaptive* neural prosthetic system.

One common thread in these neural prosthetic research models is that the animals are first trained on a motor task. The experiment is then repeated with a brain-controlled task that represents the trained motor control. It is not clear if this *a priori* motor training is required to learn control of a neural prosthetic device. There

may also be devices that do not have inherent correlates with physical motor control (*e.g.* powered wheelchairs, trolleys, communication boards, and various adapted vocational tools). Fine control of such neural prosthetic devices would not come from motor training, but rather through adapting arbitrary control signals from the neocortex.

Our objective in this study was to determine if untrained subjects were able to learn control of an unfamiliar neural prosthetic device. We hypothesized that by coadapting to the subject's neuronal responses, we could allow previously naïve subjects to gain control of a foreign neural prosthesis. Here we show that subjects can learn one-dimensional neural control of a novel auditory device without prior motor training.

4.2 Methods

4.2.1 Surgical Procedure

Six Long-Evans rats weighing 275-300g (Charles River Laboratories) were used during this study. Animals were kept on a reversed light schedule and housed within the animal facility of the University of Michigan. Subjects were handled before surgery, however training did not occur until the cortical control experiments began. Implantation methods have been discussed previously in more detail (Kipke et al., 2003; Vetter et al., 2004) and will only briefly be described here. Prior to surgery, anesthesia was induced through an intraperitoneal injection of a mixture of 50 mg/ml ketamine, 5 mg/ml xylazine, and 1 mg/ml acepromazine at an injection volume of 0.125 ml/100 g body weight. Anesthesia was maintained with hourly intraperitoneal injections of 0.1 ml ketamine (50 mg/ml). Subjects were placed in stereotaxic ear bars (MyNeuroLab.com, St. Louis, MO) and an incision was made down the midline of the head. Tissue was removed to reveal the reference fissures on the skull and 3

bone screw holes were drilled using a surgical bit. A craniotomy was created over the forelimb area of the primary motor cortex (MI) of the left hemisphere. The dura was resected to allow insertion of the penetrating electrode. One 16-channel chronic silicon-substrate microelectrode array (Kipke et al., 2003; Vetter et al., 2004) was implanted by hand using fine forceps into the brain (see 4.1 for stereotaxic locations). Each electrode (Center of Neural Communication Technology, Ann Arbor, MI; catalog 4x4 4mm200 chronic) had four separate shanks (200 μm inter-shank spacing) with four recording sites spaced evenly along each shank (200 μm inter-site spacing). The crshortciteaniotomy was filled with a hydrogel polymer (ALGEL^{TM}, Neural Intervention Technologies, Ann Arbor, MI) to anchor the electrode, and a silicone polymer (Kwik-Sil^{TM}, World Precision Instruments) was applied to protect the electrode ribbon cable. Finally, a protective headcap was created using dental acrylic (Co-Oral-Ite, Dental Mfg. Co., Santa Monica, CA). The animals were allowed 48-72 hours to recover from surgery. All surgical and animal care procedures were in accordance with the National Institute of Health guidelines and were approved by the University of Michigan Institutional Animal Care and Use Committee.

4.2.2 Data Acquisition

During each experimental session neural electrophysiological data from the 16 electrode channels were sampled at 40 kHz. These signals were simultaneously amplified and bandpass filtered (450−5000 Hz) on a Multichannel Neuronal Acquisition Processor (MNAP; Plexon Inc, Dallas, TX). At the beginning of each recording session, units on each electrode channel were separated and identified using online thresholding, template matching, and principal component analysis. Spike times were transmitted with nominal delays over a local TCP/IP connection to a second computer running custom software (MATLAB, Mathworks Inc., Natick, MA) for

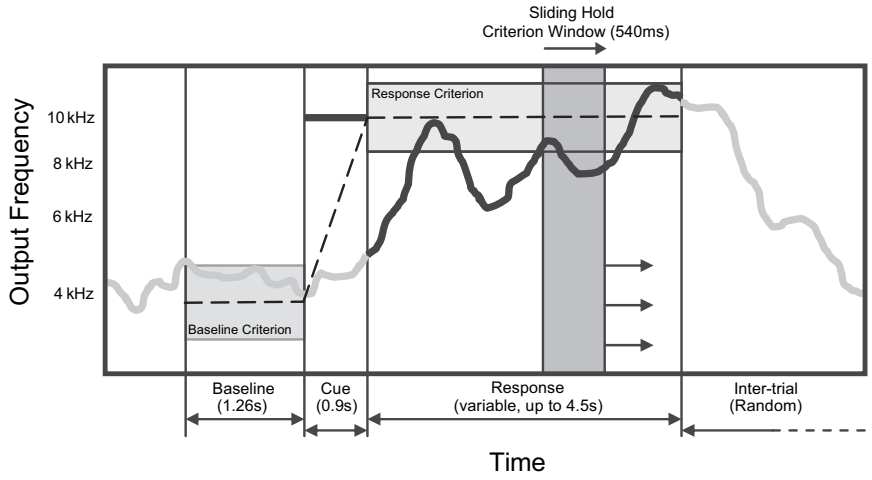

Figure 4.1: Behavioral Paradigm. Dark line indicates auditory tone frequency played back to the subjects (initial cue followed by 90 ms feedback pips). Horizontal shaded regions represent the criterion windows for the baseline and response frequencies. Responses were determined correct and were rewarded if the feedback frequency was maintained within the response criterion of the target frequency, for the duration of a 540 ms sliding window, incrementing in 90 ms steps. The unobserved ideal response or "intended" response used during the training of the adaptive filter is indicated by a thin dotted line. The result of the illustrated trial would be determined as a "correct target" trial.

neural decoding and the environmental hardware control (Tucker-Davis Technologies, Gainesville, FL). Spike times and waveshapes were stored to disk for offline analysis. Event timings for target tone onset and food delivery were captured and stored by pulse signals that synchronized the hardware events with spike timing.

4.2.3 Behavioral Paradigm

The cortical control system used in these experiments is a one-dimensional auditory analog of the center-out reaching task (Georgopoulos et al., 1986; Schwartz et al., 1988). In center-out reaching experiments, the hand (or cursor, in the case of brain-controlled tasks) is held at the center of a circle until a target cue is placed in one of a fixed number of points on the perimeter of the circle. The subject's task is to move the hand (cursor) into the target position and hold it for a fixed amount of time. Target acquisition must be completed within an allowable response

period. In our auditory version, an audio cursor is represented by 90ms sound pips representing the one-dimensional location within the logarithmically-spaced 250Hz to 16kHz frequency spectrum. Baseline firing rates were mapped to the center of the frequency space, and trials began with the presentation of a target tone at a given frequency. Subjects had a fixed amount of time to match the target frequency using the auditory cursor. The movement of the auditory cursor was dependent on the real-time decoding of the cortical firing rates as described below. As with the center-out paradigm, trials are marked either as correct (held at correct target frequency for the hold period), wrong (held at an incorrect target frequency for the hold period), or late (no answer within the response period).

Subjects were kept at 85% of their free feeding weight and were tested using either a fixed target task (10kHz tone, N=6) or a target discrimination task (1.5kHz or 10kHz tone, N=2). The fixed target task was run for 2−3 sessions a day for 8 days. Each session consisted of 100 trials. 4.1 illustrates trial timing and sequence for this task. Trials began when subjects held the auditory cursor within the criterion window of the baseline level (4kHz) for 540ms. The criterion window was set as ±17% of the logarithmic workspace (1.5kHz-10kHz). Pilot studies determined that this criterion window allowed naïve users to acquire the target in 15-20% of trials by chance. An auditory cue (10kHz) was then presented for 900ms to indicate the target frequency. Auditory feedback of the predicted cursor position was presented to the subjects in 90ms pips. Subjects had a 4.5s response window to maintain the cursor within the criterion window of the 10kHz target for 540ms. Correct responses were reinforced with a food pellet (45mg; BioServe #F0021, Laurel, MD). A random inter-trial interval (5−15s) separated each trial.

The behavioral trial for the target discrimination task was similar to the fixed

target, with the exception of an additional target tone. The two target tones were equally spaced (log scale) from the baseline tone. Baseline, target1, and target2 frequencies were set to 4kHz, 10kHz, and 1.5kHz respectively. Subjects of the discrimination task were initially trained on a fixed target before being presented with two targets. Experimental runs were constructed as either training or testing sessions. Training sessions repeated missed targets up to 4 times, while testing sessions pseudo-randomly presented an equal number of the 1.5kHz or 10kHz targets. All sessions began with randomized weights and consisted of 200−300 trials. Subjects ran 1 training session and 1−2 testing sessions per day. Only testing sessions were used for analysis in this study. At the end of each day of training, supplemental dry food was provided (if necessary) to maintain body weight near 85% *ad lib*.

Our behavioral paradigm provided a contingency between the stimulus (target tone), the desired response (reaching the correct target), and the presentation of a reinforcer (food pellet) which allowed subjects to associatively learn the cortical control task. Initially, the probability of a correct response was low, but through this operant conditioning paradigm, the number of correct trials occurred with higher frequency.

4.2.4 Ensemble Decoding

To enable real-time neural control of the auditory cursor, we used a Kalman filter to infer the cursor frequency from the neural recording data (Wu et al., 2004). A detailed description of the Kalman filter can be found in Maybeck (Maybeck, 1979). Briefly, the Kalman filter is a mathematical procedure that provides an efficient computational means to estimate the state of a process based on noisy Gaussian observations. Here, the observations are the binned spike times, which can be assumed to be Gaussian distributed provided that unit firing rates are sufficiently high. In our

paradigm, the cursor frequency was modeled as a system state variable x_{tk}, where t is the trial index ($t = 1, 2, ..., T$; where T is the total number of trials in the session), and k is an index of time ($k = 1, 2, ..., K$; where K is the number of 90ms time steps in trial t). For each trial, the cursor x_{tk} was assumed to propagate in time according to an unobserved difference equation

$$x_{tk} = Ax_{t(k-1)} + w_{t(k-1)}, \tag{4.1}$$

where A relates the prior cursor position to the current position ($A = 1$ in our experiments) and w_{tk} is a white noise term that was assumed to have a normal probability distribution, $w_{tk} \sim N(0, W_t)$. Note that x_{tk} is a scalar in our one-dimensional paradigm. However, similar equations to those that follow can be written when x_{tk} is a vector of multiple dimensions.

Unit firing times were collected in 90ms bins and modeled as an observed noisy response to the unobserved state process 4.1. We define our measurement difference equation that describes the relationship between the cursor frequency (x_{tk}) and recorded spike bins ($\mathbf{z}_{tk}$) as:

$$\mathbf{z}_{tk} = \mathbf{H}_t x_{tk} + \mathbf{q}_{tk} \tag{4.2}$$

where $\mathbf{z}_{tk}$ is a $C \times 1$ vector of spike bins from C cells, $\mathbf{H}_t$ is a $C \times 1$ vector that linearly relates the frequency state to the neural firing. Again, we assume the noise in the observations has zero mean and is normally distributed, i.e. $\mathbf{q}_{tk} \sim N(0, \mathbf{Q_t})$.

As 4.2 illustrates, the Kalman filter uses the above state equations 4.1 and 4.2 as a model to infer, or predict, the cursor position, $\hat{\mathbf{x}}_{tk}$, given only the observed spike bins, $\mathbf{z}_{tk}$, of a trial. In order to use this model, we must first define 3 parameters: $\mathbf{H}_t$, W_t, and $\mathbf{Q}_t$. It is with these parameters that we adapt the Kalman filter to the neural response of our subjects.

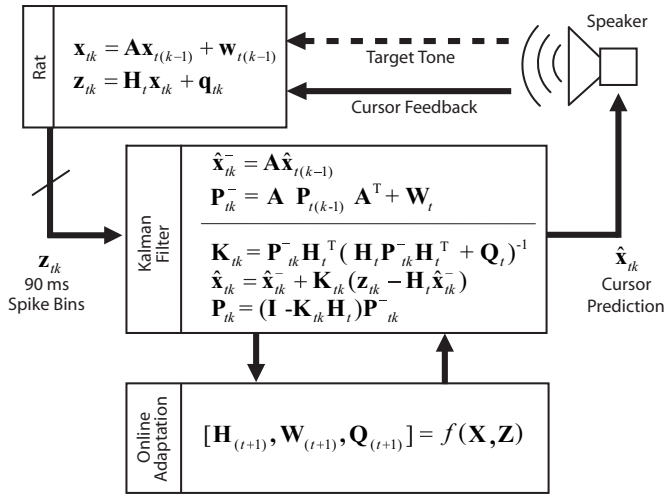

Figure 4.2: Closed loop cortical control schematic. Spike bins ($\mathbf{z}_{tk}$) from the motor cortex were decoded using a Kalman filter to predict the cursor frequency ($\hat{\mathbf{x}}_{tk}$). The predicted frequency was fed back to subjects via a speaker every 90ms of the response window.

4.2.5 Filter Adaptation

The challenge presented to a "naïve" prosthetic controller is to predict the cursor position, while simultaneously estimating the weights that should be assigned to an evolving neural ensemble experiencing learning-induced changes. For assistance we turn to stochastic control theory, a well-established engineering discipline for tracking non-stationary control signals (Åström, 1970; Maybeck, 1979). Stochastic control theory offers many tools for dealing with parameter estimation of dynamical state-space systems. Algorithms such as recursive least squares (RLS) or recursive Newton-Raphson provide a way for filter coefficients and cursor predictions to be simultaneously calculated on a bin-by-bin basis (Davis and Vinter, 1984). Additionally, block estimation allows for standard system identification techniques to be used on non-stationary signals (Haykin, 1996). In block estimation, the available data are divided into individual blocks which are small enough to assume pseudo-stationarity. The filter coefficients are then computed and updated on a block-by-block basis.

We selected the block estimation approach as it allowed the filter coefficients to be trained on select regions within the block that maximized our selection criteria (see eq. 4.3). This selective training technique provided an opportunity to identify and focus on units whose firing rate properties were best suited for control.

The blocks used in this study consisted of data obtained from the past 10 trials, and were updated on a trial-by-trial basis. Data from new trials were added to a block by selecting the appropriate time lag in the response window where the unit firing rates had the largest correlation to the expected cursor movement. The time lag at which this occurred was determined by calculating the correlation coefficient of a sliding window of the recorded unit responses with a window of the "intended" target frequency movement for the given trial. The lag chosen for training (l) was calculated across C cells via the formula

$$l = \text{argmax}_j \left(\text{argmax}_{c \in (1,2,\ldots,C)} corr([\mathbf{x}_{t(1:n)}\ \mathbf{x}_{t(j:j+n)}], [\mathbf{z}_{t(1:n),c}\ \mathbf{z}_{t(j:j+n),c}]) \right), \tag{4.3}$$

where $z_{tk,c}$ is the spike bin count of cell c at time k, and j is the index of a sliding window of length n across R response bins, $j \in [1, 2, ..., R-n]$. In our experiments, $R = 50$ and $n = 6$. The colon operator indicates concatenation, for example a vector containing the values of x for the first 5 time steps of trial t is written as $\mathbf{x}_{t(1:5)} = [x_{t1}\ x_{t2}\ x_{t3}\ x_{t4}\ x_{t5}]$. The function *corr* is defined as:

$$corr(\mathbf{x}, \mathbf{z}) = \frac{\sum_i (\mathbf{x}_i - \overline{\mathbf{x}})(\mathbf{z}_i - \overline{\mathbf{z}})}{\sqrt{\sum_i (\mathbf{x}_i - \overline{\mathbf{x}})^2 \sum_i (\mathbf{z}_i - \overline{\mathbf{z}})^2}}, \quad \overline{\mathbf{x}} = \frac{1}{n} \sum_{i=0}^{n} \mathbf{x}_i. \tag{4.4}$$

Note that the vector $\mathbf{x}$ of equation 4.3 does not refer to the predicted cursor position of trial t, but rather to the ideal unobserved cursor movement (see dashed line of 4.1). The vectors used in the correlation function are a concatenation of the baseline $(1:n)$ and response $(j:j+n)$ windows of the trial, thus allowing for one-dimensional movements. By correlating vectors of the ideal movements with vectors of the ob-

served firing rates $\mathbf{z}_t$, we can adapt the decoding filter to focus on unit responses that are potentially suited for control.

The lag calculated from 4.3 over the past $M+1$ trials was used to train the parameters of the decoding filter for the immediately upcoming trial ($M = 9$ in our experiments). The transformation matrix $\mathbf{H}$ was estimated online by the regression equation

$$\hat{\mathbf{H}}_{(\mathbf{t+1})} = (\mathbf{Z}\mathbf{X}^T)(\mathbf{X}\mathbf{X}^T)^{-1}, \tag{4.5}$$

where $\mathbf{X}$ is the *intended cursor block* that concatenates $\mathbf{X_t}$ of the past $M+1$ trials

$$\mathbf{X} = [\mathbf{X}_t\ \mathbf{X}_{(t-1)}\ \mathbf{X}_{(t-2)}\ ...\ \mathbf{X}_{(t-M)}], \tag{4.6}$$

and $\mathbf{X_t}$ is defined as the ideal cursor movement from baseline to response

$$\mathbf{X_t} = [\mathbf{x}_{t(1:n)}\ \mathbf{x}_{t(l:l+n)}]. \tag{4.7}$$

Similarly, the matrix $\mathbf{Z}$ of 4.5 is the *observation block* which is defined as the concatenation of the ensemble spike bins from the past $M+1$ trials at their appropriate lag, l, that best correlated to the ideal movement

$$\mathbf{Z} = [\mathbf{Z}_t\ \mathbf{Z}_{(t-1)}\ \mathbf{Z}_{(t-2)}\ ...\ \mathbf{Z}_{(t-M)}] \tag{4.8}$$

$$\mathbf{Z}_t = [\mathbf{z}_{t(1:n)}\ \mathbf{z}_{t(l:l+n)}]. \tag{4.9}$$

By using a trial-by-trial sliding block to obtain the least squared estimate of relationship between spike rates and cursor frequency ($\hat{\mathbf{H}}$) and by operantly rewarding correct trials, we hypothesized that subjects would learn the behavioral task and develop a strategy to control the auditory cursor.

The noise matrices $\hat{\mathbf{Q}}_{(t+1)}$ and $\hat{W}_{(t+1)}$ were then estimated using the equations

$$\hat{W}_{(t+1)} = (\mathbf{X}_{1:(N-1)} - A\mathbf{X}_{1:(N-1)})(\mathbf{X}_{2:N} - A\mathbf{X}_{1:(N-1)})^T/(N-1), \tag{4.10}$$

and

$$\hat{\mathbf{Q}}_{(t+1)} = (\mathbf{Z} - \hat{\mathbf{H}}_{(t+1)}\mathbf{X})(\mathbf{Z} - \hat{\mathbf{H}}_{(t+1)}\mathbf{X})^T/N, \tag{4.11}$$

where N is the block length, $N = (M+1)2n$.

The updated filter parameters were used to decode the cursor frequency for the immediately following trial. The end of one trial was the beginning of the next, so $x_{t0} = x_{(t-1)K}$. On the initial trial of each session $\mathbf{H_1}$, $\mathbf{Q_1}$, and $\mathbf{W_1}$ were randomized.

4.2.6 Behavioral Performance

The behavioral performance for each subject was investigated by comparing the percentage of correct targets with the percentage that would be expected by chance. For the target discrimination task, late trials were discarded and each session was treated as a two-state forced choice paradigm in which chance was 50%. However, the fixed target paradigm consisted of trials that could not end in a "wrong target" state. Therefore, a stimulus randomization method was employed to determine the amount of correct targets that could have been selected by chance during the response period.

The stimulus randomization method is described as follows. If the predicted frequency ($\hat{x}$) was not related to the stimulus (target) tones, then the tone times can be randomized without affecting the chance that the rat meets the reward criterion. For each training session, 300 sequences of random tone times were drawn. Each random sequence had the same number of tone times and the same tone time distribution as the experimental session. The number of times that the subject would have been rewarded for each of the 300 sequences was used to determine the distribution of chance. For each behavioral session, the *percent correct above chance* was calculated.

For the fixed target task, we defined the *early learning session* as the first session

on which there was reasonable certainty (P<0.05) that the subject performed better then chance. For the target discrimination task, we also performed a dynamic analysis of learning using a state-space framework that analyzes binary observations (Smith et al., 2005). Correct and wrong trials from all testing sessions were arranged to form a time series of binary trial responses and were used to compute the learning curve and its confidence intervals using the state-space smoothing algorithm described in Smith *et al.* (Smith et al., 2004). We defined the early learning trial as the first trial on which there is reasonable certainty (P<0.05) that a subject performs better than chance for the following 500 trials.

An additional measure of chance was estimated through the use of catch trials. Catch trials consisted of stimuli in which the intensity of the tone and feedback were set to 0 dB to determine the number of trials which reached the correct target by chance.

4.2.7 Unit Analysis

To monitor the evolution of the tuning properties over sessions, we measured the signal strength of the peri-event histograms (PEH) relative to background noise for each unit for every session. The PEH, or signal histogram, was centered on target tone onset and had the width of a full trial window (6.66s, 90ms bins). Noise histograms were generated by selecting random times from the recording session, and centering the trial window at this random location. We calculated the root mean square (RMS) of both the signal and noise histograms as

$$\mathrm{RMS_k} = \sqrt{\frac{1}{\mathrm{N}}\sum_{\mathrm{n=1}}^{\mathrm{N}}(\mathbf{s}_{\mathrm{k,n}} - \mu)^2}, \tag{4.12}$$

where k is a label, k $\in$ {signal, noise}, N is the total number of bins for the histogram, $\mathbf{s}$ is the firing rate value of the $\mathrm{n^{th}}$ bin, and μ is the mean firing rate. Using the RMS

calculations of both signal and noise histograms, we then calculated the signal to noise ratio (SNR) as

$$\text{SNR(dB)} = 20\log_{10}\left[\frac{\text{RMS}_{\text{signal}}}{\text{RMS}_{\text{noise}}}\right]. \tag{4.13}$$

To determine if a recorded unit showed a significant inhibitory or excitatory response, confidence intervals were calculated for each target onset-centered PEH. The computations for these intervals were based on the null hypothesis that spike trains are the realization of independent Poisson-point processes as described in the literature (Abeles, 1982). Unit responses that crossed the upper 99% confidence interval during the response window were labeled as *excitatory* while those that crossed the lower 99% confidence interval during the response window were labeled *inhibitory*. Responses crossing both confidence intervals were excluded from this analysis.

To test if discrimination had occurred in our subjects, we looked for distinct changes in the unit response pattern for each target. Mean firing rates during the response window of all trials were loaded into MATLAB (Mathworks Inc., Natick, MA) and were separated and labeled according to the target for each trial. Trial by trial analysis was then performed using a Support Vector Machine (SVM) classification toolbox (Ma et al., 2002). The SVM filter was used to predict the target frequency given the test data from each trial using training data (mean firing rates and target labels) from all other trials. This analysis was repeated for each trial in the session (*"leave-one-out"* method) and the total percent correct was noted. Sessions where the classification percent correct was significantly greater than the calculated binomial distribution ($P<0.05$, Z-test) were labeled as demonstrating *target discrimination*.

4.2.8 Histology

Upon completion of training, rats were transcardially perfused with 4% formaldehyde. The brains were removed, sectioned into 40 μm coronal slices, and stained with a conventional cresyl violet Nissl stain. The sections were then analyzed using a Leica MZFLIII light microscope (Leica Microsystems, Inc., Germany) to determine probe placement.

4.3 Results

We summarize two sets of experiments. The first involves data from 6 naïve subjects in which the task was to move an auditory cursor to a fixed target. Next, we considered the case in which multiple targets were presented and the task was to discriminate and move to the correct target. Two subjects from the initial study were used during the later stage.

4.3.1 Fixed Target Task

After recovery from surgery, populations of 5−23 (mean, 11.5) single and multi-unit clusters were discriminated from the recording electrodes and used for the fixed target control task. Subjects were able to control the feedback cursor significantly above chance (early learning session) within 7±1.4 (mean±SD) behavioral sessions ($P<0.05$, stimulus randomization method, N=6). 4.3 shows the group performance across 24 successive sessions. The 95% confidence intervals indicate 2 standard deviations above and below the mean chance distribution (randomization method). Horizontal lines indicate the number of days that each subject participated in the study. Tick marks indicate the early learning sessions for each of the subjects. Two subjects were not included after session 9 due to illness (KCD−06) and the loss of all recording units (KCD−03). The average percent of trials on which the target was

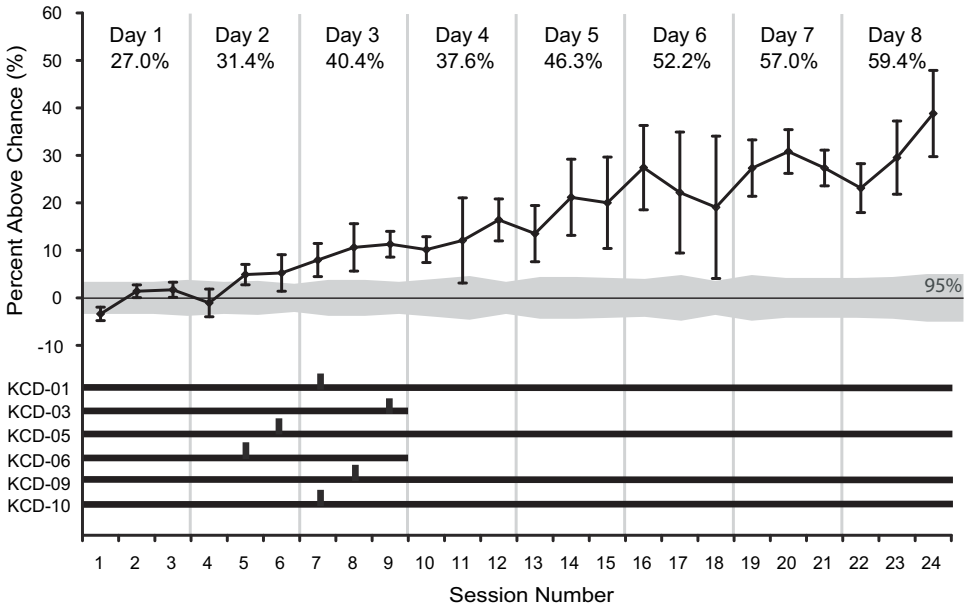

Figure 4.3: Performance in the fixed target task for 6 rats across 8 days. The top line shows the mean of the group's performance above chance. Error bars represent standard error. Gray band indicates the 95% confidence interval. Bottom lines indicate the length of time each subject was performing the task. Tick marks indicate early learning sessions. Days of training and mean percent correct for all subjects for each day are indicated across the top.

Subject	Coordinates(mm)	Sessions	Reg. Coeff.	P
KCD−01	AP:+3.8 ML:2.0	20	1.6	< 0.001
KCD−03	AP:+2.5 ML:3.0	9	1.3	< 0.05
KCD−05	AP:+2.6 ML:2.5	24	2.4	< 0.02
KCD−06	AP:+3.0 ML:2.5	9	0.24	> 0.58
KCD−09	AP:+1.2 ML:2.9	23	1.1	< 0.02
KCD−10	AP:+1.5 ML:1.6	24	0.96	< 0.01

Table 4.1: Electrode placement and performance for all subjects in a single target task. Total number of sessions and the regression coefficients of subject's performance as a function of session number are given. P values <0.05 indicates that this coefficient was significant.

successfully acquired is shown for each day. This percentage increased from the first session (21.7%±6.5%, mean±SD) to the last (69.6%±17.6%, mean±SD; P<0.01, one-way ANOVA).

Regression coefficients of the subjects' performance over chance as a function of session number are given in 4.1. These coefficients were all positive and in 5 out of 6 subjects (83.3%) significant (P<0.05, one-way ANOVA), indicating that performance increased with training. The coefficient for the entire group of subjects was 1.53 percentage points per session and was significant (P<0.01, one-way ANOVA).

4.4 shows a typical output of the fixed target experiments. Two trials (20-21) from subject KCD−01 are shown from session 24. A raster of the spike train outputs (a) from the 14 sorted single and multi-unit clusters are aligned with the behavioral output (b). After each trial, the coadaptive algorithm updated the weights (c) of the decoding matrix ($\hat{\mathbf{H}}$) to minimize the output frequency error based on the neural ensemble responses from the past 10 trials. The processing time for this updating algorithm to run is indicated in the figure (typically between 2-8s). PEHs averaged across all 100 trials are shown in (d-g). Each PEH is centered at the time that the 10kHz tone was presented (indicated by arrows). Excitatory and inhibitory responses can be observed in the response window in both the raster plot (a) and the averaged

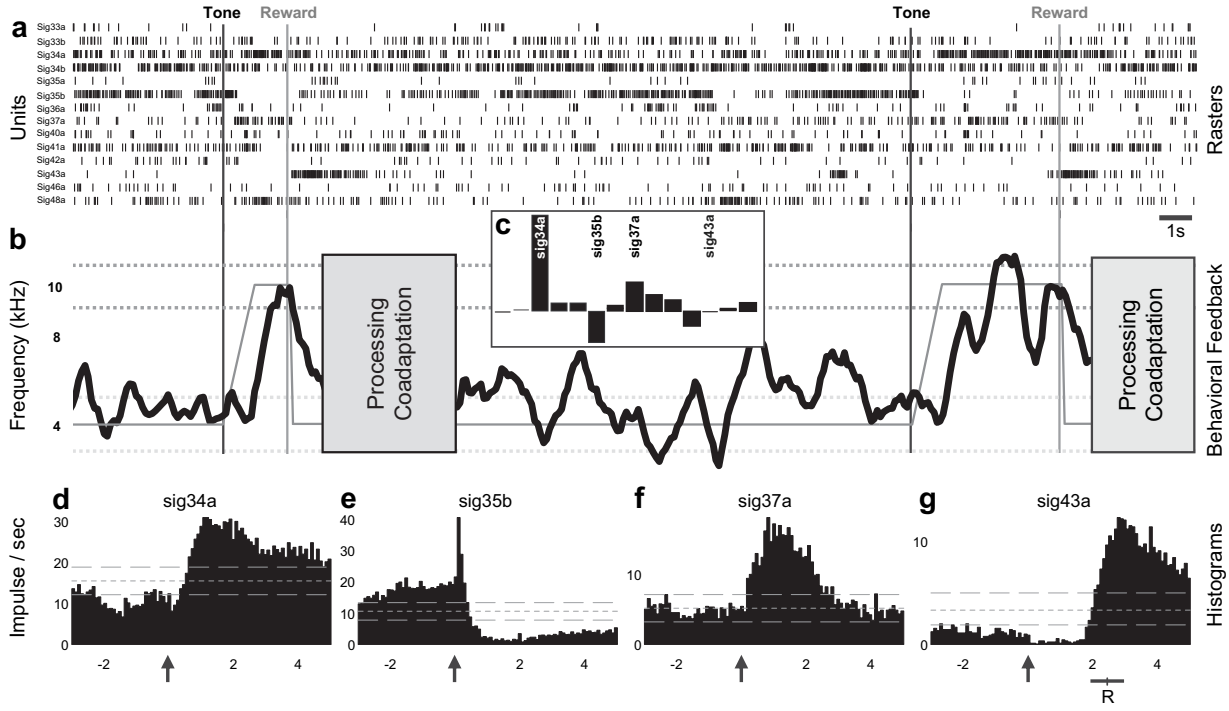

Figure 4.4: Example of trial output aligned with neural ensemble response. a) Raster plot of spike trains from 14 units over two trials. b) Behavioral output over two trials. The estimated cursor frequency ($\hat{x}$) is shown as a thick line, while the expected frequency (x) is shown as a thin line. The criterion windows for the target and baseline frequencies are indicated as horizontal dashed lines. The beginning of each target tone is marked with a black vertical line, while gray vertical lines indicate when the subject produced a correct response and was positively reinforced. c) Relatively weighted translation decoding $\hat{\mathbf{H}}$ matrix. d-g) Peri-Event Histograms over 100 trials, centered on target onset (indicated by arrow). Bottom bar (R) indicates mean±SE of the reward distribution for the PEH plots. Time scale for a and b is indicated by the 1s bar. Time scale for d-g is shown in seconds.

PEH (d-g), which allowed for control of the auditory cursor along a one-dimensional frequency axis. Upon tuning, units that responded in an excitatory fashion to the tone during the response window received high positive decoding weights, while inhibitive responses during the response window received larger negative weights. Units that did not respond, or only responded outside the trial window (g), received low weights as they were not useful for device control. During this session, the estimated auditory cursor ($\hat{\mathbf{x}}$) was strongly correlated ($\rho = 0.707$, P<0.01) to the ideal target and baseline frequencies ($\mathbf{x}$).

To visualize the cursor across an entire session, we examined the distribution of the cursor predictions during all of the baseline and response windows. 4.5 shows a histogram of $\hat{\mathbf{x}}$ predictions in 100 logarithmically-spaced bins for both the baseline window (4kHz, gray) and target response window (10kHz, black) from all trials (N = 100) of three sessions of KCD−01. During the first session (a), the distributions for both the baseline and response are centered on the baseline frequency. The response distribution begins to spread into the target window during the early learning session (b) as the subject learns to control the cursor towards the target. By the late learning session (c), the distribution becomes bimodal indicating that the subject was able to hold the cursor at the baseline target (4kHz) during the baseline window, then quickly acquire the target (10kHz) and hold it for the required criterion (540ms) during the response period. The change in scaling of the response distribution is an indication of the increased correct responses (thus shorter response windows) during learning. The mean response window length was 4.02s, 2.82s, and 2.05s for the naïve session, the early learning session, and the late learning session respectively.

To ensure that the subject's behavior was based on the target and feedback tones, we conducted an experiment using randomized intensities of these tones from a set

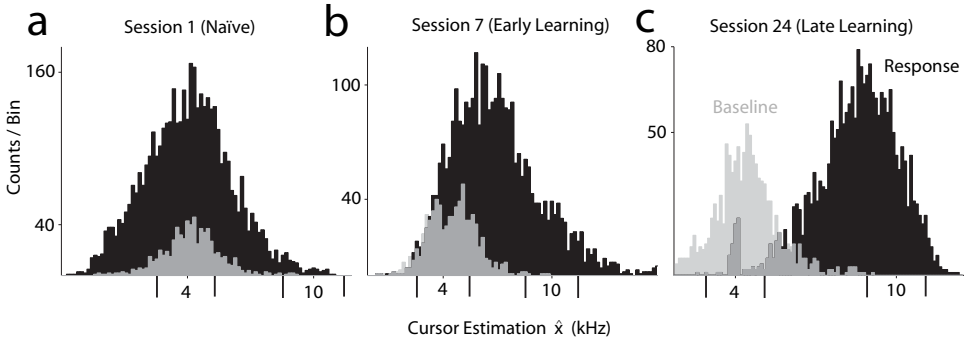

Figure 4.5: Histogram of auditory cursor frequency estimation ($\hat{\mathbf{x}}$) during two trial windows: baseline (gray) and response (black), for three sessions: naïve, early learning, and late learning of KCD−01. Baseline period was fixed at 1.2s, while the response window was variable (up to 4.5s) depending on response time. Criterion windows for the baseline and response frequencies are indicated by vertical bars on the $\hat{\mathbf{x}}$ axis.

of 5 levels (0, 17.5, 35, 52.5, and 70dB). 4.6 shows the mean behavioral results from 341 random intensity trials across 2 sessions (KCD−01). The results demonstrate that the subject's behavior was a function of the stimulus intensity, indicating that the target stimuli (and not random fluctuations in cursor predictions) were driving the observed behavior. The sigmoidal shape of 4.6 is consistent with psychophysical measurements of tone detection tasks (Thomas and Setzer, 1972). Furthermore, the 0dB catch trials provided an alternative online determination of chance that we could measure against the randomization method. During these sessions, 0dB trials received almost no correct responses (5.1%±1.1%, mean±SD). Chance for this task using the randomization method was calculated to be 19.9%±2.2% (mean±SD), indicating that the randomization method provides a conservative method for determining chance.

The adaptive algorithm selected units that showed potential for control based on the criterion that they were able to modulate from their baseline firing rate during the response window. This selection criterion resulted in changes to the signal to noise ratio of the PEHs across behavioral sessions. The mean SNR increased sig-

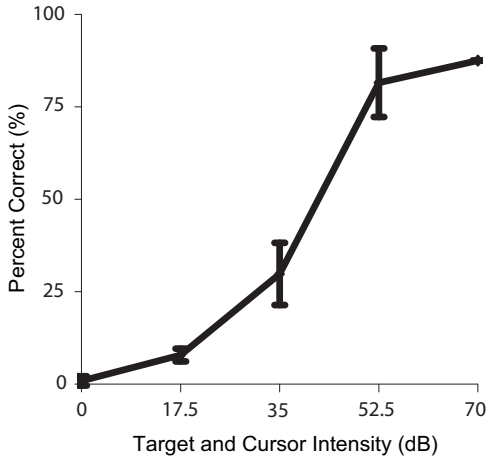

Figure 4.6: Intensity randomization test. Behavioral results of a well trained subject (KCD−01) from 341 trials across 2 sessions where the intensity of both target and feedback tones was randomized between 0−70dB. Error bars indicate standard error.

nificantly for the unit PEHs of all rats (P<0.01, ANOVA), indicating that at least one unit in the ensemble had developed the ability to control the system via changes in the firing rates. Moreover, the median SNR also increased significantly (P<0.01, ANOVA) across all subjects indicating that many units in the ensemble were active in control, a desirable property for establishing a robust control signal. 4.7 shows the distribution of SNR calculations for the PEHs of all units of KCD−05 across 24 training sessions. Median and mean are indicated, and increased significantly with session number (P<0.001, ANOVA). Representative PEHs and the SNR values are shown from sessions 2 and 23. Across all subjects, units from sessions in which the subject did not show significant performance above chance had a mean SNR of 1.7±2.1dB (mean±SD, N=377 units) with median 1.2dB, while units from sessions where control was significantly above chance measured 3.7±4.0dB (mean±SD, N=509 units) with a median of 4.4dB.

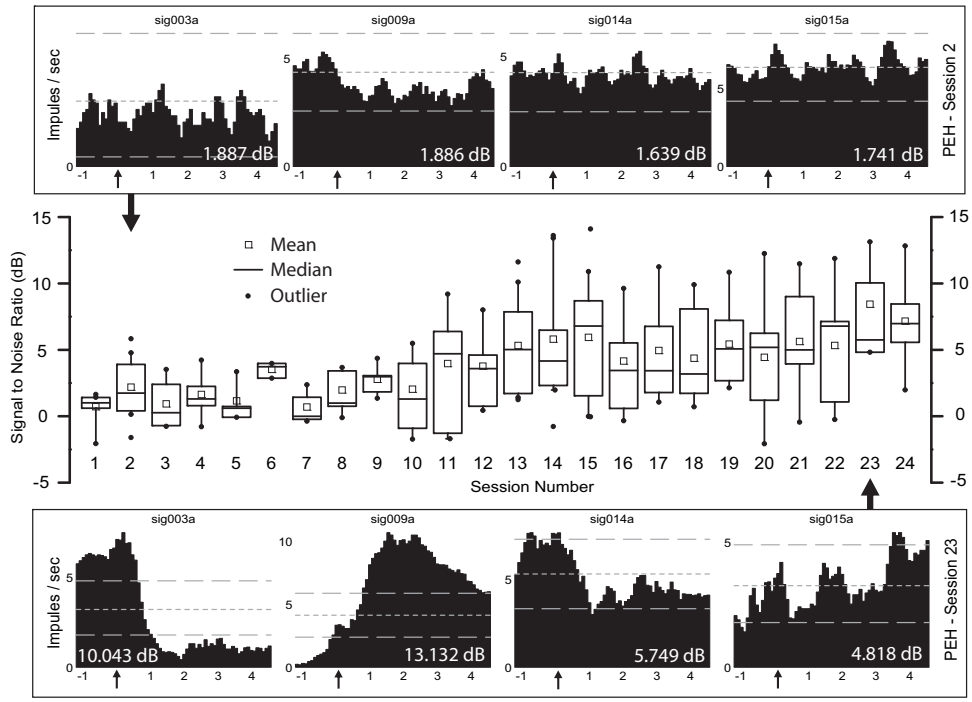

Figure 4.7: Signal to noise ratio (SNR) results for KCD-05 and examples of 8 units from 2 sessions. Boxes indicate the upper and lower quartile of the distribution, whiskers indicate the 10th and 90th percentile. The regression coefficients of both the mean and median SNR across sessions were positive and significant (mean, 0.26 dB/session, $P<0.001$, ANOVA; median, 0.27 dB/session, $P<0.001$, ANOVA). Representative PEHs and their 99% confidence intervals are shown from sessions 2 and 23. The SNR calculation for each PEH is shown in white.

4.3.2 Target Discrimination Task

Two subjects (KCD−09 and KCD−10) were trained on a two target discrimination control task for 30 and 8 testing sessions, respectively. 4.8 shows the behavioral results for both subjects. The gray band indicates the 95% confidence intervals of chance as calculated from a binomial distribution based on the number of trials for each session. KCD−09 showed a positive learning trend (0.8 percentage points per session, $P<0.01$, ANOVA) which allowed for several sessions where >90% of aquired targets were correct. KCD−10 remained at chance for most sessions, however offline SVM classification (Ma et al., 2002) of the recorded unit activity was able to correctly classify trials above chance ($P<0.05$, Z-test) for both subjects indicating that KCD−10 had learned to discriminate the targets, but had not yet learned control of the cursor to select both targets. The behavioral performance of KCD−09 rises and stays above the 95% confidence interval of chance for several sessions after session 12. However, dynamic analysis using the method of Smith *et al.* (Smith et al., 2004) in which correct and wrong trials (N=3708) from all sessions were used as binary observations of the unobservable learning state process indicated that early learning occurred much earlier, on trial 715 (session 6). Dynamic analysis of KCD−10 trials indicated that early learning had not occurred.

Three sessions from KCD−09 fell to chance after the early learning session. During session 19, the units previously used for control could not be sorted from the noise. Offline unit analysis showed that the subject was able to adopt an alternative strategy involving units from other channels in the array on session 20, leading to a recovery in performance.

In order for subjects to control the cursor to reach targets in opposite directions, ideally the unit responses to each target should also be tuned to deviate from the

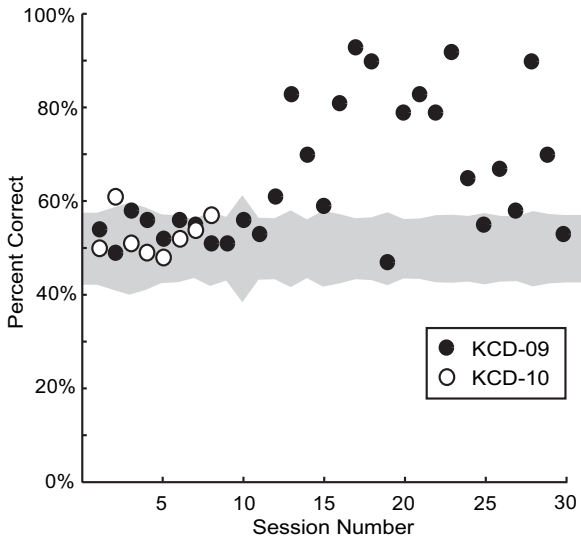

Figure 4.8: Performance of two subjects in discrimination control task. Gray band indicates the 95% confidence interval of chance as calculated from a binomial distribution based on the total number of trials. Changes in band width are due to slight variations in the number of trials per session. In sessions where both subjects are present, the gray band indicates the larger of the two chance intervals.

baseline firing rate in opposite directions. Across the 39 behavioral sessions of both subjects, 438 units were manually sorted and used for control (mean±SD, 11.1±2.4 units per session). Of these units, 235 (53.7%) showed significant responses immediately following the 10kHz target tone. These responses were further classified as either excitatory (135 units, 57.5%) or inhibitory (100 units, 42.6%). For the 1.5kHz tone, 228 units (52.1%) showed significant modulation during the response window, where 176 (77.2%) of these were excitatory and 52 (22.8%) were inhibitory. Of the units that responded significantly to either of the targets (371 units, 87.9%), 92 (24.8%) did not show a discrimination between the tones. These units were either excitatory (71 units, 77.2%) for both of the presented targets, or inhibitory for both (21 units, 22.8%). However, 232 units (62.5%) selectively showed a significant response for one but not the other target tone. Moreover, 47 units (12.7% of all behaviorally responding units) selectively had significant excitatory responses for one

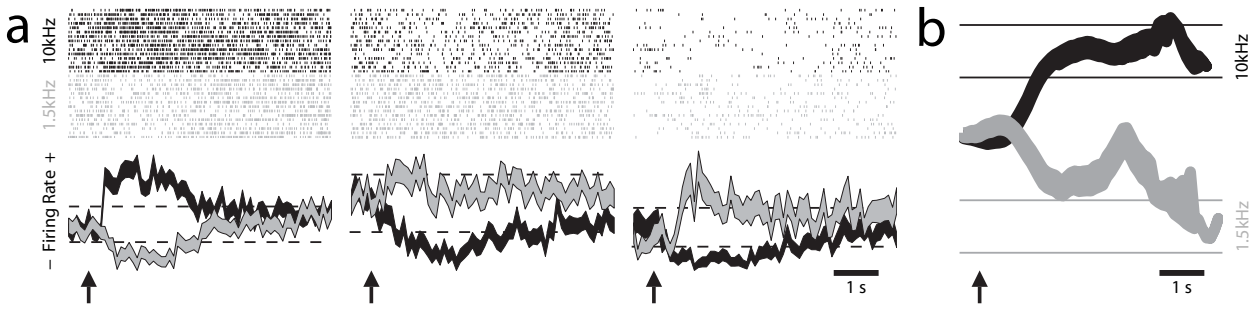

Figure 4.9: a) Raster and PEH plots of 3 units (selected from 10 units) from KCD−09 on session 17 of the two target task. All plots are centered on target onset (indicated by arrows). The first 15 trials for both 1.5kHz target (gray) and 10kHz target (black) are shown in the raster plots, the PEH includes all 200 trials of the session. Thickness of PEH indicates standard error. Dashed lines indicate the upper and lower 99% confidence intervals. b) Trajectory plot for 10kHz target (black) and 1.5kHz target (gray). Late trials were ignored. Horizontal lines indicate the criterion window for the respective target. Thickness indicates standard error.

target, while having significantly inhibitory responses for the other.

These bilateral target discriminating units, combined with units that responded for only one of the targets, allowed one subject (KCD−09) to control the cursor to both targets significantly above chance ($P<0.05$, randomization method) after 13 testing sessions. 4.9 shows the response of three units (selected from 10) in the target discrimination task from session 17. Raster plots of the first 15 pseudo-random trials for each target have been sorted into 10kHz (black) and 1.5kHz (gray). The PEHs for all trials of both tones are plotted directly beneath. The arrow indicates when the target was presented. One unit shows an excitatory response for the 10kHz tone and inhibitory response for the 1.5kHz tone, while the other units show the opposite response. The trajectory plot of all trials is shown in (b). The subject in this session was able to reach the correct target in 93% of the trials.

4.4 Discussion

This study demonstrates that naïve subjects can learn closed-loop, real-time control of one-dimensional cursor movements using single-unit and multi-unit activity of the neocortex. Several investigators have demonstrated similar degrees of adaptation and robustness (Musallam et al., 2004; Wolpaw and McFarland, 2004; Taylor et al., 2002). Previous studies have demonstrated that one-dimensional control signals could be inferred from simultaneously recorded neurons in rats that were first trained on a lever task (Chapin et al., 1999; Olson et al., 2005). Our results suggest that this initial training period may not be required for neural prosthetic control.

We did observe stereotyped motor behaviors in some animals performing the cortical control task. We are unable to determine whether these stereotyped behaviors were related to the recorded units' firing rate modulation, or whether the motor behaviors were merely superstitious behaviors. We can state however, that all feedback and rewards were based solely on the unit activity of the motor cortex. Other neural prosthetic studies which have looked at EMG signals from the muscle groups of the cortical region used for control, report that EMG modulations were eventually absent in brain-controlled control tasks (Taylor et al., 2002; Carmena et al., 2003). These and other studies (Fetz and Finocchio, 1971; Chapin et al., 1999) indicate that cortical control is possible without direct motor movements.

Our model contains assumption violations that do not affect our conclusions but should be addressed. First, the Kalman filter implicitly assumes a linear Gaussian relationship between cursor movement and unit firing rate. While this assertion is not exact, a linear assumption has been shown to be a reasonable approximation by several investigators for real time control (Taylor et al., 2002; Serruya et al., 2002).

The Gaussian assumption only holds for binned data when spike rates are sufficiently high. An alternative approach to avoid this violation would be to use a filter based not on binned data, but on the point process observations of spike times (Eden et al., 2004). Additionally, the variance of the noise used in state equation 4.1 was not an accurate representation of the feedback cursor, as this was determined by regressing intended movements, which were specified as ideal (noiseless) movements. Nonetheless, the resulting equations allowed our subjects to quickly learn control of the modeled feedback cursor. Further work needs to be done to establish whether decoding filters that remove these assumptions will assist subjects to learn faster or achieve higher performance.

One limitation of our adaptive algorithm is that it is based on a supervised learning paradigm in which both the intended target and neural activity were used to estimate each neuron's receptive field. This paradigm works well for goal-directed tasks but has disadvantages for longer, free-ranging tasks where receptive fields may change or new neurons become available after the training stage. Our work could be extended to allow the adaptive algorithm to simultaneously predict the cursor estimates while tracking the evolution of the receptive field parameters ($\hat{\mathbf{H}}$) on a bin-by-bin basis. This technique would allow newly added neurons to contribute to the decoding, even when they were not present during the encoding (training) stage. Eden *et al.* (Eden et al., 2004) provide an example of receptive field tracking in their development of an adaptive stochastic state point process filter for neural decoding.

Subjects were able to control the one-dimensional cursor in the fixed target task using a few (11.5±3.5, mean±SD) MI units per session. This is comparable to similar experiments in monkeys where 7-30 MI neurons were used for two-dimensional control (Serruya et al., 2002) and 18 units for three-dimensional control (Taylor et al., 2002).

While the number of units used per session did not increase or decrease significantly, there were changes in the number of units sorted and used for control from session to session (coefficient of variance = 30.1%). The regression coefficient of the subjects' performance (percent over chance) as a function of the number of units used for control was calculated to be an increase of 2.4 percent per unit, and was significant ($P<0.03$, ANOVA). While this analysis indicates that performance increased with the number of units, it does not mean subjects could not perform the task with fewer units. There were sessions where subjects reached >90% of targets (>50% above chance) using as few as 5 units.

Our control algorithm requires no prior knowledge of the recorded neuron's tuning properties for adequate system performance. For the initial trial of each session, random weights were assigned to each recorded unit. Subsequently, the system adapted based on stereotyped unit responses across the microelectrode array to enable control. This coadaptive process allowed the rats in our paradigm to learn neural control within 3 days of training. These findings suggest that the described adaptive decoding filter may be a means of training paralyzed human patients where motor tuning properties of neurons are unobtainable. Several groups are investigating brain computer interfaces in which human subjects must learn to derive control signals using electroencephalographic recordings (Wolpaw and McFarland, 2004; Fabiani et al., 2004), electrocorticographic recordings (Leuthardt et al., 2004a), cortical local field potentials (Kennedy et al., 2004a), and single-unit activity (Serruya et al., 2004; Kennedy et al., 2000). As each of these systems must derive tuning properties from an initial naïve state, a coadaptive decoding filter (Taylor et al., 2002) may decrease the time to learn brain control or may maximize control performance by allowing the brain to explore all possible neural responses and to adopt the strategy that is

the easiest for control.

4.5 Acknowledgements

This work was supported by DARPA contract number N66001-02-C-8059. Probes were provided by the Center for Neural Communication Technology sponsored by NIH NIBIB grant P41-EB002030-11. Aspects of this work were previously presented at the 26th Annual International Conference of the IEEE Engineering in Medicine and Biology Society, San Francisco, USA, Sept. 2004 and at the Society for Neuroscience Annual Meeting, San Diego, CA, Oct. 2004. The authors would like to thank Rio Vetter for surgical assistance, Kristina Shalizi for data analysis, Tim Marzullo for helpful suggestions on the behavioral paradigm, Eugene Daneshvar and Luis Salas for their assistance with animal experiments, and Matt Johnson for histology.

CHAPTER V

Conclusion

In this dissertation, I developed a further understanding of several key processes involved in the generation of intended actions. As this work focused on the prominent role the basal ganglia play in action selection, it is of no surprise that deeper brain structures, including the striatum, have been implicated in generating motor commands in non-invasive brain computer interface (BCI) experiments Hinterberger et al. (2005). I wish to conclude my dissertation, by restating the major findings of this thesis, and relating how these can help advance the development of the next generation neuromotor prosthesis (NMP).

While the identification of cell types (as performed in the analysis of Chapter 2) has yet to be performed for neuromotor prosthesis, there have been attempts to use neural oscillations (Chapter 3) as a control signal. A number of studies (Wolpaw et al., 1991; del R Millán et al., 2004; Cincotti et al., 2003) have shown that EEG signals in humans can be used in a large range of applications, including the control computer of cursors to rudimentary control of a wheelchair. The EEG signal on one electrode represents the summed activity of thousands of neurons recordings and can only resolve low frequencies of neural activity (for a review of recording characteristics, see Buzsaki, 2004). These recordings suffer from a low signal-to-noise ratio

and have very low spatial resolution (Schwartz et al., 2006). Electrocorticographic (ECoG) recordings are made from direct contact with the surface of the brain, and are therefore more robust and are a noise-free alternative to EEG recordings. ECoGs have been used quite successfully for neuromotor prosthesis applications (Leuthardt et al., 2004b; Schalk et al., 2007; Pistohl et al., 2008; Gunduz et al., 2009). However, both EEG and ECoG may not be practical solutions for NMPs, since normal motor activities such as eye movements, talking and chewing may cause high noise artifacts in both signals.

Despite the many promising, high-profile results in the realm of intra-cortical invasive NMP (Santhanam et al., 2006; Velliste et al., 2008), single unit recordings can be difficult to maintain for long periods of time (Hochberg et al., 2006). Micromotion of implanted micro-electrodes affects recording stability and insertion injury, chronic inflammation and glial encapsulation increases electrode impedance and limits the ability to record spikes over time (Schwartz et al., 2006). LFPs can be recorded from intra-cortical electrodes even under low impedance conditions, and could be used to supplement single unit decoders. Kennedy et al. (2004b) used the LFP to control a computer cursor in one dimension and a virtual finger, with each LFP channel controlling one degree of freedom. Mehring et al. (2003) used a support vector machine to predict movement trajectories using LFP and showed that the discriminatory power of the signal for a 2-D task was similar to that of single-unit data. This work has lead to LFPs being used for several NMP devices (Rickert et al., 2005; Scherberger et al., 2005).

5.1 FSI and MSN cell identification for neuromotor prosthesis control

In the first study, I identified two major classes of cells using the shape of the waveform of averaged spikes. A neuromotor prosthesis with the ability to distinguish individual cell types, either based on waveshapes, or by other means (optogenetics, for example) would be able to provide additional information to the neuromotorprosthetic controller. One finding that could be exploited is that when compared to other neuronal populations, FSIs were preferentially active at moment of the choice event. I also showed that at this event, FSI neurons were highly tuned to the direction of the selected movement. Identified FSI cells recorded simultaneously could be pooled together such that a coordinated pulse of activity could be detected indicating that an action plan has been selected, and that the information is ready to be inferred from their respective weights.

5.2 Detection of β_{20} oscillations to indicate neuromotor "pausing"

Local field potential control remains largely unexplored in the neuroprosthetic community, though groups continue to show that the LFP is co-modulated along with spikes during a given behavioral task (Mehring et al., 2003; Heldman et al., 2006). Rickert et al. (2005) reported that movement direction (for a 2D center out task) is encoded in the LFP in both the time and frequency domains. They also found that the amplitude of the peaks of movement evoked potentials (mEPs) and cue evoked potentials (cEPs) varied with direction. Evoked potentials are a characteristic waveform that appears in an LFP signal when it is averaged over a time-locked event. More recently, Rubino et al. (2006) demonstrated that amplitude and phase of 15−45 Hz (β) predicted the directions of upcoming arm movements. Rickert et al. (2005) reported two positive (P1, P2) and two negative (N1, N2) peaks

in the mEP and cEP, and found that across all 419 LFP recordings over all sessions in the contralateral and ipsilateral hemispheres, 16% of the mEPs were tuned during P1, 38% during N1, 48% during P2 and 50% during N2. They also found that three frequency bands ($\leq$ 4, 6−13 and 63−200 Hz) were directionally modulated. The spectra of the individual LFP signals were very similar on average in all of the frequency bands. However, they did not report the percentage of LFP signals tuned in the frequency domain, or the distribution of preferred directions in either the frequency or time domains.

Heldman et al. (2006) also reported directional modulation in the 18−26, 30−80 and 60−200 Hz bands (10.5%, 12.5% and 14.9% of recorded M1 LFPs were significantly tuned), in a 3D reach task. Scherberger et al. (2005) showed that LFPs in the posterior parietal cortex modulated with the direction of reaches and saccades prior to execution.

In the second study of this dissertation (Chapter 3), I observed that beta oscillations in the range of 15−25 Hz (β_{20}) were related to the suppression of a selected motor plan. This can immediately be put to use to solve one engineering issue facing the development of practical neuroprosthetics: unintended movements. The presence of unwanted movements in neuromotor prosthesic devices can cause usability or even safety concerns. Currently, thresholding of the decoding algorithm is an often-used method of eliminating uninitiated movements. However, this method is prone to many sources of error, including the elimination of desired movements. By recording the field potentials in addition to the spiking activity used for NMP control, real-time information concerning the stopping of signals can be read out along side the directional information. The finding that β_{20} oscillations was selectively active during the suppression of an intended action can be applied directly to help address this

problem. By digitally filtering the LFP in real-time, an instantaneous estimate of the β_{20} could inversely be applied to the output speed of the NMP. Where as an absence of β_{20} would allow for normal control at designed speeds, and the up-regulation of β_{20} would cause the NMP to stop entirely. One could imagine this forming a graded potential allowing for speed of motions to be controlled more easily.

One basic question remains to be explored: can we use LFP to control a NMP on a single trial basis? The findings that demonstrated increased β_{20} oscillations during motor plan suppression were based on the *average of all trials taken together*. As shown in the raw traces of figure 3.3b, the β_{20} epoch is present on most trials, but not all. It is not a continuous signal during the hold period that can easily be detected above the background noise, rather it occurs in multiple brief $\sim$7 cycle spindles. We may still be able to harness the observed modulation of β_{20} to predict suppression of moments in a BMI. Further research may investigate the optimal parameter set based on the delay post-detection of a beta event to allow an action to occur. Each β_{20} epoch could reset an internal counter to release the next movement. Each following epoch would delay the movement further until the suppression of beta activity eventually releases the timer, to allow the pre-programmed movement to occur.

Bibliography

Abeles, M. (1982). Quantification, smoothing, and confidence limits for single-units' histograms. *J Neurosci Methods*, 5(4):317–25.

Adams, S., Kesner, R. P., and Ragozzino, M. E. (2001). Role of the medial and lateral caudate-putamen in mediating an auditory conditional response association. *Neurobiology of Learning and Memory*, 76(1):106–16.

Androulidakis, A. G., Doyle, L. M. F., Yarrow, K., Litvak, V., Gilbertson, T. P., and Brown, P. (2007). Anticipatory changes in beta synchrony in the human corticospinal system and associated improvements in task performance. *Eur J Neurosci*, 25(12):3758–65.

Antoine, J. P., Coron, A., and Dereppe, J. M. (2000). Water peak suppression: time-frequency vs time-scale approach. *J Magn Reson*, 144(2):189–94.

Aron, A. R. and Poldrack, R. A. (2006). Cortical and subcortical contributions to stop signal response inhibition: role of the subthalamic nucleus. *J Neurosci*, 26(9):2424–33.

Åström, K. J. (1970). *Introduction to Stochastic Control Theory*, volume 70 of *Mathematics in Science and Engineering*. Academic Press, Inc., New York.

Baker, S. N., Olivier, E., and Lemon, R. N. (1997). Coherent oscillations in monkey

motor cortex and hand muscle emg show task-dependent modulation. *J Physiol (Lond)*, 501 (Pt 1):225–41.

Barthó, P., Hirase, H., Monconduit, L., Zugaro, M., Harris, K. D., and Buzsaki, G. (2004). Characterization of neocortical principal cells and interneurons by network interactions and extracellular features. *J Neurophysiol*, 92(1):600–8.

Bennett, B. D. and Bolam, J. P. (1994). Synaptic input and output of parvalbumin-immunoreactive neurons in the neostriatum of the rat. *Neuroscience*, 62(3):707–19.

Berger, H. (1929). Uber das elektrenkephalogramm des menschen. *Arch Psychiatr Nervenkr*, 87:527–570.

Berger, T. W., Baudry, M., Brinton, R., and Liaw, J. (2001). Brain-implantable biomimetic electronics as the next era in neural prosthetics. *Proceedings of the IEEE*, 89(7).

Berke, J. D. (2008). Uncoordinated firing rate changes of striatal fast-spiking interneurons during behavioral task performance. *J Neurosci*, 28(40):10075–80.

Berke, J. D., Breck, J. T., and Eichenbaum, H. B. (2009). Striatal versus hippocampal representations during win-stay maze performance. *J Neurophysiol*, 101(3):1575–87.

Berke, J. D., Okatan, M., Skurski, J., and Eichenbaum, H. B. (2004). Oscillatory entrainment of striatal neurons in freely moving rats. *Neuron*, 43(6):883–96.

Bevan, M. D., Booth, P. A., Eaton, S. A., and Bolam, J. P. (1998). Selective innervation of neostriatal interneurons by a subclass of neuron in the globus pallidus of the rat. *J Neurosci*, 18(22):9438–52.

Bolam, J. P., Somogyi, P., Takagi, H., Fodor, I., and Smith, A. D. (1983). Localization of substance p-like immunoreactivity in neurons and nerve terminals in the neostriatum of the rat: a correlated light and electron microscopic study. *J Neurocytol*, 12(2):325–44.

Brasted, P. J., Humby, T., Dunnett, S. B., and Robbins, T. W. (1997). Unilateral lesions of the dorsal striatum in rats disrupt responding in egocentric space. *J Neurosci*, 17(22):8919–26.

Brown, P. (2006). Bad oscillations in parkinson's disease. *J Neural Transm Suppl*, (70):27–30.

Brown, P. and Marsden, C. D. (1998). What do the basal ganglia do? *Lancet*, 351(9118):1801–4.

Brown, P., Oliviero, A., Mazzone, P., Insola, A., Tonali, P., and Lazzaro, V. D. (2001). Dopamine dependency of oscillations between subthalamic nucleus and pallidum in parkinson's disease. *J Neurosci*, 21(3):1033–8.

Brown, V. J. and Robbins, T. W. (1989). Elementary processes of response selection mediated by distinct regions of the striatum. *J Neurosci*, 9(11):3760–5.

Buzsaki, G. (2004). Large-scale recording of neuronal ensembles. *Nat Neurosci*, 7(5):446–51.

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

Carli, M., Jones, G. H., and Robbins, T. W. (1989). Effects of unilateral dorsal

and ventral striatal dopamine depletion on visual neglect in the rat: a neural and behavioural analysis. *J Neurosci*, 29(2):309–27.

Carmena, J. M., Lebedev, M. A., Crist, R. E., O'Doherty, J. E., Santucci, D. M., Dimitrov, D. F., Patil, P. G., Henriquez, C. S., and Nicolelis, M. A. L. (2003). Learning to control a brain-machine interface for reaching and grasping by primates. *PLoS Biology*, 1(2):e2.

Carrillo-Reid, L., Tecuapetla, F., Tapia, D., Hernández-Cruz, A., Galarraga, E., Drucker-Colin, R., and Bargas, J. (2008). Encoding network states by striatal cell assemblies. *J Neurophysiol*, 99(3):1435–50.

Chapin, J. K., Moxon, K. A., Markowitz, R. S., and Nicolelis, M. A. (1999). Real-time control of a robot arm using simultaneously recorded neurons in the motor cortex. *Nat Neurosci*, 2(7):664–70.

Chapuis, J., Garcia, S., Messaoudi, B., Thevenet, M., Ferreira, G., Gervais, R., and Ravel, N. (2009). The way an odor is experienced during aversive conditioning determines the extent of the network recruited during retrieval: a multisite electrophysiological study in rats. *J Neurosci*, 29(33):10287–98.

Chelvanayagam, D. K., Vickery, R. M., Kirkcaldie, M. T. K., Coroneo, M. T., and Morley, J. W. (2008). Multichannel surface recordings on the visual cortex: implications for a neuroprosthesis. *J Neural Eng*, 5(2):125–32.

Chen, C. C., Litvak, V., Gilbertson, T., Kühn, A., Lu, C. S., Lee, S. T., Tsai, C. H., Tisch, S., Limousin, P., Hariz, M., and Brown, P. (2007). Excessive synchronization of basal ganglia neurons at 20 hz slows movement in parkinson's disease. *Exp Neurol*, 205(1):214–21.

Cincotti, F., Mattia, D., Babiloni, C., Carducci, F., Salinari, S., Bianchi, L., Marciani, M. G., and Babiloni, F. (2003). The use of eeg modifications due to motor imagery for brain-computer interfaces. *IEEE TNSRE*, 11(2):131–3.

Cook, D. and Kesner, R. P. (1988). Caudate nucleus and memory for egocentric localization. *Behav Neural Biol*, 49(3):332–43.

Courtemanche, R., Fujii, N., and Graybiel, A. M. (2003). Synchronous, focally modulated beta-band oscillations characterize local field potential activity in the striatum of awake behaving monkeys. *J Neurosci*, 23(37):11741–52.

Crewson, P. (2006). Applied statistics handbook. *AcaStat Software*.

Csicsvari, J., Hirase, H., Czurko, A., and Buzsaki, G. (1998). Reliability and state dependence of pyramidal cell-interneuron synapses in the hippocampus: an ensemble approach in the behaving rat. *Neuron*, 21(1):179–89.

Davis, M. H. A. and Vinter, R. B. (1984). *Stochastic Modelling and Control*. Monographs on Statistics and Applied Probability. Chapman and Hall, New York.

del R Millán, J., Renkens, F., Mouriño, J., and Gerstner, W. (2004). Noninvasive brain-actuated control of a mobile robot by human eeg. *IEEE transactions on bio-medical engineering*, 51(6):1026–33.

Diester, I. and Nieder, A. (2008). Complementary contributions of prefrontal neuron classes in abstract numerical categorization. *J Neurosci*, 28(31):7737–47.

Donoghue, J. P., Sanes, J. N., Hatsopoulos, N. G., and Gaál, G. (1998). Neural discharge and local field potential oscillations in primate motor cortex during voluntary movements. *J Neurophysiol*, 79(1):159–73.

Durand, D. and Greenwood, J. (1958). Modifications of the rayleigh test for uniformity in analysis of two-dimensional .... *The Journal of Geology*.

Eadie, M. J. (1998). Dreamy mental states in late 20th century neurology. *Journal of clinical neuroscience : official journal of the Neurosurgical Society of Australasia*, 5(2):157–60.

Eden, U. T., Frank, L. M., Barbieri, R., Solo, V., and Brown, E. N. (2004). Dynamic analysis of neural encoding by point process adaptive filtering. *Neural Comput*, 16(5):971–98.

Engel, A. K., Fries, P., and Singer, W. (2001). Dynamic predictions: oscillations and synchrony in top-down processing. *Nat Rev Neurosci*, 2(10):704–16.

Fabiani, G. E., McFarland, D. J., Wolpaw, J. R., and Pfurtscheller, G. (2004). Conversion of eeg activity into cursor movement by a brain-computer interface (bci). *IEEE Trans Neural Syst Rehabil Eng*, 12(3):331–8.

Farmer, S. F., Bremner, F. D., Halliday, D. M., Rosenberg, J. R., and Stephens, J. A. (1993). The frequency content of common synaptic inputs to motoneurones studied during voluntary isometric contraction in man. *J Physiol (Lond)*, 470:127–55.

Fetz, E. E. (1969). Operant conditioning of cortical unit activity. *Science*, 163:955–8.

Fetz, E. E. and Finocchio, D. V. (1971). Operant conditioning of specific patterns of neural and muscular activity. *Science*, 174:431–5.

Fukai, T. and Tanaka, S. (1997). A simple neural network exhibiting selective activation of neuronal ensembles: from winner-take-all to winners-share-all. *Neural computation*, 9(1):77–97.

Gage, G. J., Ludwig, K. A., Otto, K. J., Ionides, E. L., and Kipke, D. R. (2005). Naive coadaptive cortical control. *J Neural Eng*, 2(2):52–63.

Georgopoulos, A. P., Schwartz, A. B., and Kettner, R. E. (1986). Neuronal population coding of movement direction. *Science*, 233(4771):1416–9.

Gerfen, C. R. and Wilson, C. J. (1996). The basal ganglia v. *Handbook of Chemical Neuroanatomy, Vol 12: Integrated Systems of the CNS, Part III*, L. W. Swanson, A. Bjorklund, & T. Hokfelt (Eds.) Amsterdam: Elsevier:371–468.

Gernert, M., Hamann, M., Bennay, M., Löscher, W., and Richter, A. (2000). Deficit of striatal parvalbumin-reactive gabaergic interneurons and decreased basal ganglia output in a genetic rodent model of idiopathic paroxysmal dystonia. *J Neurosci*, 20(18):7052–8.

Gilbertson, T., Lalo, E., Doyle, L., Lazzaro, V. D., Cioni, B., and Brown, P. (2005). Existing motor state is favored at the expense of new movement during 13-35 hz oscillatory synchrony in the human corticospinal system. *J Neurosci*, 25(34):7771–9.

Gold, J. I. and Shadlen, M. N. (2007). The neural basis of decision making. *Annu Rev Neurosci*, 30:535–74.

Gunduz, A., Sanchez, J. C., Carney, P. R., and Principe, J. C. (2009). Mapping broadband electrocorticographic recordings to two-dimensional hand trajectories in humans motor control features. *Neural networks : the official journal of the International Neural Network Society*, 22(9):1257–70.

Gustafson, N., Gireesh-Dharmaraj, E., Czubayko, U., Blackwell, K. T., and Plenz, D. (2006). A comparative voltage and current-clamp analysis of feedback and feed-

forward synaptic transmission in the striatal microcircuit in vitro. *J Neurophysiol*, 95(2):737–52.

Hammond, C., Bergman, H., and Brown, P. (2007). Pathological synchronization in parkinson's disease: networks, models and treatments. *Trends Neurosci*, 30(7):357–64.

Harris, K. D., Henze, D., Csicsvari, J., Hirase, H., and Buzsaki, G. (2000). Accuracy of tetrode spike separation as determined by simultaneous intracellular and extracellular measurements. *J Neurophysiol*, 84(1):401–14.

Haykin, S. (1996). *Adaptive filter theory (3rd ed.)*. Prentice-Hall, Inc., Upper Saddle River, NJ, USA.

Heldman, D. A., Wang, W., Chan, S. S., and Moran, D. W. (2006). Local field potential spectral tuning in motor cortex during reaching. *IEEE TNSRE*, 14(2):180–3.

Hermer-Vazquez, R., Hermer-Vazquez, L., Srinivasan, S., and Chapin, J. K. (2007). Beta- and gamma-frequency coupling between olfactory and motor brain regions prior to skilled, olfactory-driven reaching. *Experimental brain research Experimentelle Hirnforschung Expérimentation cérébrale*, 180(2):217–35.

Hinterberger, T., Veit, R., Wilhelm, B., Weiskopf, N., Vatine, J.-J., and Birbaumer, N. (2005). Neuronal mechanisms underlying control of a brain-computer interface. *Eur J Neurosci*, 21(11):3169–81.

Hochberg, L. R., Serruya, M. D., Friehs, G. M., Mukand, J. A., Saleh, M., Caplan, A. H., Branner, A., Chen, D., Penn, R. D., and Donoghue, J. P. (2006). Neuronal ensemble control of prosthetic devices by a human with tetraplegia. *Nature*, 442(7099):164–71.

Hutcheon, B. and Yarom, Y. (2000). Resonance, oscillation and the intrinsic frequency preferences of neurons. *Trends Neurosci*, 23(5):216–22.

Jaeger, D., Kita, H., and Wilson, C. J. (1994). Surround inhibition among projection neurons is weak or nonexistent in the rat neostriatum. *J Neurophysiol*, 72(5):2555–8.

Kalanithi, P. S. A., Zheng, W., Kataoka, Y., DiFiglia, M., Grantz, H., Saper, C. B., Schwartz, M. L., Leckman, J. F., and Vaccarino, F. M. (2005). Altered parvalbumin-positive neuron distribution in basal ganglia of individuals with tourette syndrome. *Proc Natl Acad Sci USA*, 102(37):13307–12.

Kawaguchi, Y., Wilson, C. J., Augood, S. J., and Emson, P. C. (1995). Striatal interneurones: chemical, physiological and morphological characterization. *Trends Neurosci*, 18(12):527–35.

Kennedy, P. R., Andreasen, D., Ehirim, P., King, B., Kirby, T., Mao, H., and Moore, M. (2004a). Using human extra-cortical local field potentials to control a switch. *Journal of Neural Engineering*, 1(2):72–7.

Kennedy, P. R., Bakay, R. A., Moore, M. M., Adams, K., and Goldwaithe, J. (2000). Direct control of a computer from the human central nervous system. *IEEE Trans Rehabil Eng*, 8(2):198–202.

Kennedy, P. R., Kirby, M. T., Moore, M. M., King, B., and Mallory, A. (2004b). Computer control using human intracortical local field potentials. *IEEE TNSRE*, 12(3):339–44.

Kilner, J. M., Baker, S. N., Salenius, S., Jousmäki, V., Hari, R., and Lemon, R. N.

(1999). Task-dependent modulation of 15-30 hz coherence between rectified emgs from human hand and forearm muscles. *J Physiol (Lond)*, 516 (Pt 2):559–70.

Kipke, D. R., Vetter, R. J., Williams, J. C., and Hetke, J. F. (2003). Silicon-substrate intracortical microelectrode arrays for long-term recording of neuronal spike activity in cerebral cortex. *IEEE Trans Neural Syst Rehabil Eng*, 11(2):151–5.

Kita, H., Kosaka, T., and Heizmann, C. W. (1990). Parvalbumin-immunoreactive neurons in the rat neostriatum: a light and electron microscopic study. *Brain Res*, 536(1-2):1–15.

Koós, T. and Tepper, J. M. (1999). Inhibitory control of neostriatal projection neurons by gabaergic interneurons. *Nat Neurosci*, 2(5):467–72.

Koós, T., Tepper, J. M., and Wilson, C. J. (2004). Comparison of ipscs evoked by spiny and fast-spiking neurons in the neostriatum. *J Neurosci*, 24(36):7916–22.

Kühn, A. A., Kupsch, A., Schneider, G.-H., and Brown, P. (2006). Reduction in subthalamic 8-35 hz oscillatory activity correlates with clinical improvement in parkinson's disease. *Eur J Neurosci*, 23(7):1956–60.

Kühn, A. A., Trottenberg, T., Kivi, A., Kupsch, A., Schneider, G.-H., and Brown, P. (2005). The relationship between local field potential and neuronal discharge in the subthalamic nucleus of patients with parkinson's disease. *Exp Neurol*, 194(1):212–20.

Lacourse, M. G., Cohen, M. J., Lawrence, K. E., and Romero, D. H. (1999). Cortical potentials during imagined movements in individuals with chronic spinal cord injuries. *Behav Brain Res*, 104(1-2):73–88.

Leblois, A., Boraud, T., Meissner, W., Bergman, H., and Hansel, D. (2006). Competition between feedback loops underlies normal and pathological dynamics in the basal ganglia. *J Neurosci*, 26(13):3567–83.

Leuthardt, E. C., Schalk, G., Wolpaw, J. R., Ojemann, J. G., and Moran, D. W. (2004a). A brain-computer interface using electrocorticographic signals in humans. *Journal of Neural Engineering*, 1(2):63–71.

Leuthardt, E. C., Schalk, G., Wolpaw, J. R., Ojemann, J. G., and Moran, D. W. (2004b). A brain-computer interface using electrocorticographic signals in humans. *J Neural Eng*, 1(2):63–71.

Levy, R., Ashby, P., Hutchison, W. D., Lang, A. E., Lozano, A. M., and Dostrovsky, J. O. (2002). Dependence of subthalamic nucleus oscillations on movement and dopamine in parkinson's disease. *Brain*, 125(Pt 6):1196–209.

Litaudon, P., Garcia, S., and Buonviso, N. (2008). Strong coupling between pyramidal cell activity and network oscillations in the olfactory cortex. *Neuroscience*, 156(3):781–7.

Llinás, R. R. (1988). The intrinsic electrophysiological properties of mammalian neurons: insights into central nervous system function. *Science*, 242(4886):1654–64.

Lo, C.-C. and Wang, X.-J. (2006). Cortico-basal ganglia circuit mechanism for a decision threshold in reaction time tasks. *Nat Neurosci*, 9(7):956–63.

Lozano, A. and Mahant, N. (2004). Deep brain stimulation surgery for parkinson's disease: mechanisms and .... *Parkinsonism and Related Disorders*.

Lucas, J. W., Schiller, J. S., and Benson, V. (2004). Summary health statistics for u.s. adults: National health interview survey, 2001. *Vital Health and Statistics*, 10(218):60–2.

Ludvig, N., Kovacs, L., Medveczky, G., Kuzniecky, R. I., and Devinsky, O. (2005). Toward the development of a subdural hybrid neuroprosthesis for the treatment of intractable focal epilepsy. *Epilepsia*, 46(270).

Ma, J., Zhao, Y., and Ahalt, S. (2002). OSU SVM Classifier MATLAB Toolbox (v3.00) [`http://www.ece.osu.edu/~maj/osu_svm`]. The Ohio State University, Ohio Supercomputer Center (OSC), Columbus, Ohio 43210.

Mallet, N., Moine, C. L., Charpier, S., and Gonon, F. (2005). Feedforward inhibition of projection neurons by fast-spiking gaba interneurons in the rat striatum in vivo. *J Neurosci*, 25(15):3857–69.

Mallet, N., Pogosyan, A., Márton, L. F., Bolam, J. P., Brown, P., and Magill, P. J. (2008). Parkinsonian beta oscillations in the external globus pallidus and their relationship with subthalamic nucleus activity. *J Neurosci*, 28(52):14245–58.

Marsden, J. F., Ashby, P., Limousin-Dowsey, P., Rothwell, J. C., and Brown, P. (2000). Coherence between cerebellar thalamus, cortex and muscle in man: cerebellar thalamus interactions. *Brain*, 123:459 –1470.

Maybeck, P. S. (1979). *Stochastic models, estimation, and control*, volume 141 of *Mathematics in Science and Engineering*. Academic Press, Inc., New York.

McDonald, R. J. and White, N. M. (1993). A triple dissociation of memory systems: hippocampus, amygdala, and dorsal striatum. *Behav Neurosci*, 107(1):3–22.

McFarland, D. J. (1965). Flow graph representation of motivational systems. *The British journal of mathematical and statistical psychology*, 18:25–43.

Mehring, C., Rickert, J., Vaadia, E., de Oliveira, S. C., Aertsen, A., and Rotter, S. (2003). Inference of hand movements from local field potentials in monkey motor cortex. *Nat Neurosci*, 6(12):1253–4.

Middleton, F. A. and Strick, P. L. (1994). Anatomical evidence for cerebellar and basal ganglia involvement in higher cognitive function. *Science*, 266(5184):458–61.

Mink, J. W. (1996). The basal ganglia: focused selection and inhibition of competing motor programs. *Prog Neurobiol*, 50(4):381–425.

Mitzdorf, U. (1985). Current source-density method and application in cat cerebral cortex: investigation of evoked potentials and eeg phenomena. *Physiol Rev*, 65(1):37–100.

Murthy, V. N. and Fetz, E. E. (1992). Coherent 25- to 35-hz oscillations in the sensorimotor cortex of awake behaving monkeys. *Proc Natl Acad Sci USA*, 89(12):5670–4.

Murthy, V. N. and Fetz, E. E. (1996a). Oscillatory activity in sensorimotor cortex of awake monkeys: synchronization of local field potentials and relation to behavior. *J Neurophysiol*, 76(6):3949–67.

Murthy, V. N. and Fetz, E. E. (1996b). Synchronization of neurons during local field potential oscillations in sensorimotor cortex of awake monkeys. *J Neurophysiol*, 76(6):3968–82.

Musallam, S., Corneil, B. D., Greger, B., Scherberger, H., and Andersen, R. A. (2004). Cognitive control signals for neural prosthetics. *Science*, 305(5681):258–62.

Nauta, W. J. and Mehler, W. R. (1966). Projections of the lentiform nucleus in the monkey. *Brain Res*, 1(1):3–42.

Nunez, P. L. and Srinivasan, R. (2006). A theoretical basis for standing and traveling brain waves measured with human eeg with implications for an integrated consciousness. *Clin Neurophysiol*, 117(11):2424–35.

Olds, J. (1965). Operant conditioning of single unit responses. In *Proc XXIII Int Congress Physiol Sci*, Tokyo.

Olson, B., Si, J., Hu, J., and He, J. (2005). Closed-loop cortical control of direction using support vector machines. *IEEE Trans Neural Syst Rehabil Eng*, 13(1):72–80.

Packard, M. G. and McGaugh, J. L. (1996). Inactivation of hippocampus or caudate nucleus with lidocaine differentially affects expression of place and response learning. *Neurobiology of Learning and Memory*, 65(1):65–72.

Parent, A. and Bellefeuille, L. D. (1982). Organization of efferent projections from the internal segment of globus pallidus in .... *Brain Res.*

Parthasarathy, H. B. and Graybiel, A. M. (1997). Cortically driven immediate-early gene expression reflects modular influence of sensorimotor cortex on identified striatal neurons in the squirrel monkey. *J Neurosci*, 17(7):2477–91.

Paxinos, G. and Watson, C. (2005). The rat brain in stereotaxic coordinates, 5th ed. *Elsevier Academic Press.*

Pfurtscheller, G., Stancák, A., and Neuper, C. (1996). Post-movement beta synchronization. a correlate of an idling motor area? *Electroencephalogr Clin Neurophysiol*, 98(4):281–93.

Pistohl, T., Ball, T., Schulze-Bonhage, A., Aertsen, A., and Mehring, C. (2008). Prediction of arm movement trajectories from ecog-recordings in humans. *J Neurosci Methods*, 167(1):105–14.

Pogosyan, A., Gaynor, L. M. F. D., Eusebio, A., and Brown, P. (2009). Boosting cortical activity at beta-band frequencies slows movement in humans. *Curr Biol*.

Ramanathan, S., Hanley, J. J., Deniau, J.-M., and Bolam, J. P. (2002). Synaptic convergence of motor and somatosensory cortical afferents onto gabaergic interneurons in the rat striatum. *J Neurosci*, 22(18):8158–69.

Redgrave, P., Prescott, T. J., and Gurney, K. N. (1999). The basal ganglia: a vertebrate solution to the selection problem? *Neuroscience*, 89(4):1009–23.

Rickert, J., de Oliveira, S. C., Vaadia, E., Aertsen, A., Rotter, S., and Mehring, C. (2005). Encoding of movement direction in different frequency ranges of motor cortical local field potentials. *J Neurosci*, 25(39):8815–24.

Rouger, J., Lagleyre, S., Fraysse, B., Deneve, S., Deguine, O., and Barone, P. (2007). Evidence that cochlear-implanted deaf patients are better multisensory integrators. *Proc Natl Acad Sci USA*, 104(17):7295–300.

Rubino, D., Robbins, K. A., and Hatsopoulos, N. G. (2006). Propagating waves mediate information transfer in the motor cortex. *Nat Neurosci*, 9(12):1549–57.

Salmelin, R. and Hari, R. (1994). Spatiotemporal characteristics of sensorimotor neuromagnetic rhythms related to thumb movement. *Neuroscience*, 60(2):537–50.

Samejima, K. and Doya, K. (2007). Multiple representations of belief states and action values in corticobasal ganglia loops. *Ann N Y Acad Sci*, 1104:213–28.

Sanchez, J. C., Carmena, J. M., Lebedev, M. A., Nicolelis, M. A., Harris, J. G., and Principe, J. C. (2004). Ascertaining the importance of neurons to develop better brain-machine interfaces. *IEEE Trans Biomed Eng*, 51(6):943–53.

Sanchez, J. C., Mahmoudi, B., DiGiovanna, J., and Principe, J. C. (2009). Exploiting co-adaptation for the design of symbiotic neuroprosthetic assistants. *Neural networks : the official journal of the International Neural Network Society*, 22(3):305–15.

Sanes, J. N. and Donoghue, J. P. (1993). Oscillations in local field potentials of the primate motor cortex during voluntary movement. *Proc Natl Acad Sci USA*, 90(10):4470–4.

Sanes, J. N., Suner, S., and Donoghue, J. P. (1990). Dynamic organization of primary motor cortex output to target muscles in adult rats. i. long-term patterns of reorganization following motor or mixed peripheral nerve lesions. *Experimental brain research Experimentelle Hirnforschung Expérimentation cérébrale*, 79(3):479–91.

Santhanam, G., Ryu, S. I., Yu, B. M., Afshar, A., and Shenoy, K. V. (2006). A high-performance brain-computer interface. *Nature*, 442(7099):195–8.

Schalk, G., Kubánek, J., Miller, K. J., Anderson, N. R., Leuthardt, E. C., Ojemann, J. G., Limbrick, D., Moran, D. W., Gerhardt, L. A., and Wolpaw, J. R. (2007). Decoding two-dimensional movement trajectories using electrocorticographic signals in humans. *J Neural Eng*, 4(3):264–75.

Scherberger, H., Jarvis, M. R., and Andersen, R. A. (2005). Cortical local field potential encodes movement intentions in the posterior parietal cortex. *Neuron*, 46(2):347–54.

Schmidt, E. M. (1980). Single neuron recording from motor cortex as a possible source of signals for control of external devices. *Ann Biom Eng*, 8:339–49.

Schmitzer-Torbert, N. and Redish, A. D. (2008). Task-dependent encoding of space and events by striatal neurons is dependent on neural subtype. *J Neurosci*, 153(2):349–60.

Schwartz, A. B., Cui, X. T., Weber, D. J., and Moran, D. W. (2006). Brain-controlled interfaces: movement restoration with neural prosthetics. *Neuron*, 52(1):205–20.

Schwartz, A. B., Kettner, R. E., and Georgopoulos, A. P. (1988). Primate motor cortex and free arm movements to visual targets in three-dimensional space. i. relations between single cell discharge and direction of movement. *J Neurosci*, 8(8):2913–27.

Serruya, M. D., Caplan, A. H., Saleh, M., Morris, D. S., and Donoghue, J. P. (2004). The braingate pilot trial: Building and testing a novel direct neural output for patients with severe motor imparments. Program No. 190.22, 2004 Abstract Viewer/Itinerary Planner, San Diego, CA, Society for Neuroscience, 2004.

Serruya, M. D., Hatsopoulos, N. G., Paninski, L., Fellows, M. R., and Donoghue, J. P. (2002). Instant neural control of a movement signal. *Nature*, 416(6877):141–2.

Shenoy, K. V., Meeker, D., Cao, S., Kureshi, S. A., Pesaran, B., Buneo, C. A., Batista, A. P., Mitra, P. P., Burdick, J. W., and Andersen, R. A. (2003). Neural prosthetic control signals from plan activity. *Neuroreport*, 14(4):591–6.

Shoham, S., Halgren, E., Maynard, E. M., and Normann, R. A. (2001). Motor-cortical activity in tetraplegics. *Nature*, 413(6858):793.

Shu, Y., Hasenstaub, A., and McCormick, D. A. (2003). Turning on and off recurrent balanced cortical activity. *Nature*, 423(6937):288–93.

Siapas, A. G. and Wilson, M. A. (1998). Coordinated interactions between hippocampal ripples and cortical spindles during slow-wave sleep. *Neuron*, 21(5):1123–8.

Silberstein, P., Kühn, A. A., Kupsch, A., Trottenberg, T., Krauss, J. K., Wöhrle, J. C., Mazzone, P., Insola, A., Lazzaro, V. D., Oliviero, A., Aziz, T., and Brown, P. (2003). Patterning of globus pallidus local field potentials differs between parkinson's disease and dystonia. *Brain*, 126(Pt 12):2597–608.

Smith, A. C., Frank, L. M., Wirth, S., Yanike, M., Hu, D., Kubota, Y., Graybiel, A. M., Suzuki, W. A., and Brown, E. N. (2004). Dynamic analysis of learning in behavioral experiments. *J Neurosci*, 24(2):447–61.

Smith, A. C., Stefani, M. R., Moghaddam, B., and Brown, E. N. (2005). Analysis and Design of Behavioral Experiments to Characterize Population Learning. *J Neurophysiol*, 93(3):1776–1792.

Snaith, M. and Holland, O. (1991). An investigation of two mediation strategies suitable for behavioural control in .... *Proceedings of the first international conference on ....*

Spooren, W. P., Lynd-Balta, E., Mitchell, S., and Haber, S. N. (1996). Ventral pallidostriatal pathway in the monkey: evidence for modulation of basal ganglia circuits. *J Comp Neurol*, 370(3):295–312.

Taylor, D. M., Helms Tillery, S. I., and Schwartz, A. B. (2002). Direct cortical control of 3d neuroprosthetic devices. *Science*, 296(5574):1829–32.

Thomas, D. R. and Setzer, J. (1972). Stimulus generalization gradients for auditory intensity in rats and guinea pigs. *Psychonomic Science*, 28:22–4.

Turner, R. S. and Anderson, M. E. (1997). Pallidal discharge related to the kinematics of reaching movements in two dimensions. *J Neurophysiol*, 77(3):1051–74.

Turner, R. S. and Anderson, M. E. (2005). Context-dependent modulation of movement-related discharge in the primate globus pallidus. *J Neurosci*, 25(11):2965–76.

Urbanczik, R. and Senn, W. (2009). Reinforcement learning in populations of spiking neurons. *Nat Neurosci*, 12(3):250–2.

Varela, F., Lachaux, J. P., Rodriguez, E., and Martinerie, J. (2001). The brainweb: phase synchronization and large-scale integration. *Nat Rev Neurosci*, 2(4):229–39.

Velliste, M., Perel, S., Spalding, M. C., Whitford, A. S., and Schwartz, A. B. (2008). Cortical control of a prosthetic arm for self-feeding. *Nature*, 453(7198):1098–101.

Vetter, R. J., Williams, J. C., Hetke, J. F., Nunamaker, E. A., and Kipke, D. R. (2004). Spike recording performance of implanted chronic silicon-substrate microelectrode arrays in cerebral cortex. *IEEE Trans Neural Syst Rehabil Eng*, 52(1):896–904.

Voorn, P., Vanderschuren, L. J. M. J., Groenewegen, H. J., Robbins, T. W., and Pennartz, C. M. A. (2004). Putting a spin on the dorsal-ventral divide of the striatum. *Trends Neurosci*, 27(8):468–74.

Wessberg, J., Stambaugh, C. R., Kralik, J. D., Beck, P. D., Laubach, M., Chapin, J. K., Kim, J., Biggs, S. J., Srinivasan, M. A., and Nicolelis, M. A. (2000). Real-

time prediction of hand trajectory by ensembles of cortical neurons in primates. *Nature*, 408(6810):361–5.

Wickens, J. R. and Arbuthnott, G. W. (1993). The corticostriatal system on computer simulation: an intermediate mechanism for sequencing of actions. *Prog Brain Res*, 99:325–39.

Wickens, J. R., Arbuthnott, G. W., and Shindou, T. (2007). Simulation of gaba function in the basal ganglia: computational models of gabaergic mechanisms in basal ganglia function. *Prog Brain Res*, 160:313–29.

Wilson, C. J. (2000). Striatal circuitry: Categorically selective or selectively categorical? *Brain Dynamics and the Striatal Complex*, (edited by R.Miller and J.R. Wickens. CRC).

Wilson, C. J. (2004). Basal ganglia. *Synaptic Organization of the Brain, 5th Edition, G.M. Shepherd editor*.

Wilson, M. A. and McNaughton, B. L. (1993). Dynamics of the hippocampal ensemble code for space. *Science*, 261(5124):1055–8.

Wilson, M. A. and McNaughton, B. L. (1994). Reactivation of hippocampal ensemble memories during sleep. *Science*, 265(5172):676–9.

Wilson, S. A. K. (1912). Progressive ienticular degeneration: A familial nervous system disease associated with cirrhosis of the liver. *Brain*, 34:295–507.

Wilson, S. A. K. (1914). An experimental research into the anatomy and physiology of the corpus striatum. *Brain*, 36:427–492.

Wilson, S. A. K. (1928). The psychical components of temporal (uncinate) epilepsy. *Modern Problems in Neurology*, (edited by SAK Wilson (Arnold, London)):51–75.

Wiltschko, A. B., Gage, G. J., and Berke, J. D. (2008). Wavelet filtering before spike detection preserves waveform shape and enhances single-unit discrimination. *J Neurosci Methods*, 173(1):34–40.

Wolpaw, J. R. and McFarland, D. J. (2004). Control of a two-dimensional movement signal by a noninvasive brain-computer interface in humans. *PNAS*, 101(51):17849–54.

Wolpaw, J. R., McFarland, D. J., Neat, G. W., and Forneris, C. A. (1991). An eeg-based brain-computer interface for cursor control. *Electroencephalogr Clin Neurophysiol*, 78(3):252–9.

Wu, W., Black, M. J., Mumford, D., Gao, Y., Bienenstock, E., and Donoghue, J. P. (2004). Modeling and decoding motor cortical activity using a switching kalman filter. *IEEE Trans Biomed Eng*, 51(6):933–942.

Zar, J. (1974). Biostatistical analysis. Prentice Hall, Upper Saddlw River, New Jersey.

Zrenner, E. (2002). Will retinal implants restore vision? *Science*, 295(5557):1022–5.