Characterizing Corticostriatal Circuit Function During Performance of Habitual Action Sequences

by

Nuné Martiros

B. S., Massachusetts Institute of Technology 2008

SUBMITTED TO THE DEPARTMENT OF BRAIN AND COGNITIVE SCIENCES
IN PARTIAL FULFILLMENT OF THE REQUIREMENTS FOR THE DEGREE OF:

DOCTOR OF PHILOSOPHY
AT THE
MASSACHUSETTS INSTITUTE OF TECHNOLOGY
MAY 2016
[June 2016]

Signature of Author ..... Signature redacted

Department of Brain and Cognitive Sciences
April 9, 2016

Certified by ..... Signature redacted

Ann M. Graybiel
Institute Professor and Professor in Brain and Cognitive Sciences
Thesis Supervisor

Accepted by ..... Signature redacted .....

Matthew A. Wilson
Sherman Fairchild Professor of Neuroscience and Picower Scholar
Director of Graduate Education for Brain and Cognitive Science

Characterizing Corticostriatal Circuit Function During Performance of Habitual Action Sequences

by

Nuné Martiros

Submitted to the Department Of Brain and Cognitive Sciences on May, 2016 in Partial Fulfillment of the Requirements for the Degree of Doctor of Philosophy in Neuroscience.

ABSTRACT

The striatum is the largest nucleus in the basal ganglia and the recipient of dense dopamine input. Multiple cortico-basal ganglia-thalamic loops are thought to function together during the learning and performance of reinforced behaviors, with the dorsolateral circuit being particularly critical for the learning of habitual chains of action sequences. However, how this circuit works to generate such behavior is poorly understood. To explore the nature of striatal neural representations during learned action sequences, I designed a task targeted at disambiguating movement-related responses from habit representations in striatum. In combination with this task, I employed electrophysiology and optogenetics techniques to characterize task-related neuronal activity in the corticostriatal circuit. I found that, unlike in motor cortex, neurons in striatum did not respond simply to particular individual actions, but responded preferentially at the initiation and termination of learned action sequences. These experiments provide a test for the existence of a generalized striatal signal marking the start and end of units of habitual behaviors which may be produced with the contribution of striatal interneurons, providing a mechanism by which striatum can control the encoding and performance of chunked action sequences. In a separate set of experiments, I explored the effect of dopamine depletion on local field potential oscillations in the same region of striatum. My goal was to investigate the interaction between abnormal oscillations caused by dopamine depletion in Parkinson's disease and the functional task-related oscillations that normally occur in healthy striatum. Against our expectations, I found that local unilateral dopamine depletion in dorsolateral striatum did not result in changes in pre-task baseline strength of oscillations, but rather in the overexpression of the normal task-related oscillations. These studies add support to theories of striatal function and dysfunction that emphasize selective network modulation by learned behaviors.

Thesis supervisor: Ann Graybiel
Title: Institute Professor and Professor in Brain and Cognitive Sciences

TABLE OF CONTENTS

ACKNOWLEDGEMENTS

There are certain transitional moments in one's lifetime that urge one to stop and consider current circumstances from a large distance. To imagine the infinite universe around me and the speck of stardust from which I arose, and my brain arose, which in turn convinced me to study itself is one such exercise that helps me realize the scope of the incredible circumstances under which I was able to complete this chapter in life, and is both bewildering and humbling. When I started this scientific and life journey toward a PhD, I had passion in thinking about the brain and behavior but little idea of the reality of the scientific pursuit. As it turns out, being a neuroscientist is much more about supportive people, knowledgeable people, technology, and experience than I would have ever guessed. I am thankful to have this opportunity to express my gratitude, which I often find difficult to express, toward the many people who have in numerous ways helped me learn these lessons, inspired me, given me confidence in myself, and given me the courage to think differently from others. Of the many people I would like to thank, the first is my husband Martin Lemaire, with whom we have bravely faced adulthood. He has inspired me with his persistence, discipline (in which I am sorely lacking), and with his intense focus on priorities at hand. In fact, he is sitting across from the table from me at this moment, intensely focused on studying for his most difficult MIT Masters in Finance final exam tomorrow. Moreover, Martin has been an essential force in my life in his unfaltering belief in my abilities and his trust in me. Second, I would like to thank my son Remi who has taught me to be selfless, and who has brought a previously unimaginable joy into my life and my heart each day since he was born and even while he was in the womb. I am infinitely grateful to my parents Art and Ana who have always acted selflessly to take care of me, who taught me to think critically, and instilled a sense of independence in me. I would like to thank my PhD advisor Ann Graybiel for her endless support of me and for inspiring me. As one of the first women in neuroscience, she has overcome enormous obstacles, yet she is one who does not care to mention them. She is always in the present, always looking forward, never still or indifferent. Most importantly, she is one of the many people that have helped me understand that the first determinant of success is persistence and hard work. I always wondered what is was about me that made me lucky enough to catch her attention, but I am certain that I have enjoyed the pleasure of her mentorship due to this luck. I would like to thank my classmates for their comradery while we navigated becoming scientists and being in our twenties. I was lucky enough to fall into a group of wonderful people who constituted my classmates, the kinds of friends with whom I am quite sure I could easily pick up where we left off easily ten years from now. One huge force in providing perspective in our lives throughout the many years of graduate school have been our undergraduate students for whom we have served as graduate resident tutors. We are some of the few lucky people to have been friends with many of our future leaders at the ages of 18, and I am hopeful from knowing the compassion that they possess. Additionally, I am grateful for the patience of Jessica Pourian, Alexandra Burgess, Devin Ahern, and Jennifer Bustamante who spent their hours and years meticulously assembling recording drives under microscopes and training rats. I am grateful also for the guidance of my thesis committee members Matthew Wilson, Jerry Schneider, and Naoshige Uchida who were kind enough to donate their time to me, who saw value in my work, and who helped me generate ideas for new ways of looking at the data. I am incredibly obliged to very many other current and past members of our laboratory for their support and help, including Henry Hall, Jannifer Lee, Michael Riad, Min Jung Kim, Eric Burguiere, Christiane Schreiweis, Jill Crittenden, Kyle Smith, Leif Gibb, Hisham Atallah, Alexander Friedman, Mark Howe, Ledia Hernandez, Carolyn Lacey, Yasuo Kubota, Bernard Bloem, Daniel Gibson, and others.

CHAPTER 1: GENERAL INTRODUCTION

The primary area of study of this thesis work and of our laboratory are the mechanisms of function of basal ganglia circuits in the learning and performance of learned behaviors including chains of actions performed in a sequence, and decision making. While a clear set of mechanisms for striatal based learning including the involvement of striatal microcircuitry and striatal inputs and outputs is still far from reach, it is critical to make use of all levels of description to work toward such a theory of striatal function. Here, I will summarize some findings from anatomical work, lesioning work, and neuronal activity monitoring which have led to the views and hypotheses relating to striatal function and especially of the function of the dorsolateral/habit/sensorimotor striatum that have motivated the work in this thesis.

Striatum: inferring function from structure

The striatum is the largest nucleus in the basal ganglia and one of the most evolutionarily preserved brain structures specializing in action selection by reinforcement learning by being the recipient of the densest dopamine input in the brain and serving as an interface between many sensory and motor control brain regions (Medina and Reiner, 1995, Striedter, 2005, Stephenson-Jones et al., 2011, Schneider, 2014). While the initial sensory and motor regions providing input to striatum were subcortical, as neocortical sensory and motor regions developed they became among those providing input to striatum with almost every cortical region sending inputs to striatum (Alexander et al., 1986, Schneider, 2014). The neocortex also become one of the primary targets of basal ganglia output through the thalamus. While the basic microcircuitry of the striatum appears to be fairly homogeneous throughout the large structure, and preserved through evolution (Medina and Reiner, 1995), the input and output pathways to and from different striatal areas are very diverse. As such, the striatum appears to provide a multi-purpose microcircuitry for the integration of many inputs and the result has been several parallel circuits that have been roughly characterized as sensorimotor, associative, and limbic (Alexander et al., 1986, Voorn et al., 2004).

Cortical input to the striatum follows a complex organization. Generally, cortical input to the striatum is topographically organized such that the ventral striatum receives input from limbic corticies, the dorsomedial striatum receives input from more medial associative corticies, and the dorsolateral striatum receives input from more lateral sensorimotor corticies (McGeorge and Faull, 1989, Sesack et al., 1989, Ebrahimi et al., 1992, Brown et al., 1998, Hoffer and Alloway, 2001, Voorn et al., 2004, Wu et al., 2009). On the anterior-posterior axis, there is also a rough topographic organization with anterior corticies such as motor cortex sending inputs to the anterior striatum and posterior corticies such as visual cortex sending inputs to the posterior striatum. On a smaller scale, a different level of functional organization is prevalent; some cortical areas or layers project preferentially to striosome or matrix compartments of the striatum (Eblen and Graybiel, 1995, Trytek et al., 1996) and there is overlap in the projection zones of functionally related cortical areas, for example overlap in projections from somatosensory and motor cortices representing the same body part (Parthasarathy et al., 1992, Brown et al., 1998). Axons to these overlapping sensory-motor projection areas terminate on discrete groups of matrix and striosome clusters of neurons called matrisomes which in turn send axons that converge in globus pallidus (Flaherty and Graybiel, 1994) presenting the possibility that they are modular processing units of the striatum (Amemori et al., 2011). Furthermore, the striatum receives two types of cortical projections, those that are collaterals of corticospinal axons arising from layer 5B pyramidal cells and projecting only ipsilaterally, and a separate set of short-range axons to striatum arising from both layer 2/3 and layer 5B neurons and projecting bilaterally (Akintunde and Buxton, 1992, Anderson et al., 2010). Individual striatal neurons receive inputs from tens of thousands of cortical neurons with each individual cortical neuron contributing few synapses (Shepherd, 2003). How this anatomical connectivity translates to functional connectivity is unknown with only a handful of studies finding rough agreement between the known topography of corticostriatal projections and the location of local field potential responses or c-Fos expression in striatum in response to cortical electrical stimulation (Parthasarathy and Graybiel, 1997, Sgambato et al., 1997, Glynn and Ahmad, 2002). However, there is some agreement between the types of known functions of these corresponding cortical and striatal areas (Bailey and Mair, 2007, Boulougouris et al., 2007, Jonkman et al., 2009), which further suggests that

parallel corticostriatal circuits are functionally important for different aspects of learning and behavior.

The organization of outputs from the striatum at a basic level is similarly homogeneous throughout the large structure with spiny projection neurons expressing D1 type dopamine receptors and substance P (direct pathway) projecting primarily to globus pallidus internal segment (or in rodents entopeduncular nucleus) and substantia nigra pars reticulata, which in turn send another set of inhibitory connections to thalamus disinhibiting it. These same neurons often have collaterals to the external segment of globus pallidus and subthalamic nucleus, the primary target structures of the indirect pathway. The spiny projection neurons expressing mostly D2 type dopamine receptors and enkephalin (indirect pathway) do not normally project to the basal ganglia output structures, but to the intermediate structures of the globus pallidus external segment and subthalamic nucleus (Parent et al., 2000).

Major theories of striatal function have been developed from the anatomical structure of these direct and indirect pathways, with the direct pathway involving two inhibitory links to the thalamus thus providing disinhibition to the thalamus and the indirect pathway having three inhibitory links thus providing inhibition to the thalamus. As such, it is thought that while direct and indirect pathway spiny projection neurons may receive similar inputs providing information about sensory stimuli, context, and motivational levels; the direct pathway can use such information to potentiate appropriate behaviors by disinhibiting thalamus while the indirect pathway can use it to inhibit behaviors (Albin et al., 1989). However, in recent years many aspects of this model have been challenged and the model has been revised as will be discussed in a later section. One additional clue about basal ganglia function from its structure comes from the reduction in size and apparent funneling down of information from the input stage (striatum) to the output stages (globus pallidus internal segment/entopeduncular nucleus and substantia nigra pars reticulata) (Oorschot, 1996). Together, the vast diversity of inputs to the striatum, the funneling down of this information through the basal ganglia structures, and the potential for opposite influences of the direct and indirect pathways on behavior, as well as the highly dense dopamine input to the striatum indicate that the striatum is likely to play an important role in

learning and action selection based on sensory/contextual information and a history of reinforcement (Graybiel, 2008, Redgrave et al., 2011, Da Cunha et al., 2012).

Striatum: Inferring function from lesion deficits

Many attempts have been made at determining the role of the striatum in behavior by creating lesions, knocking-out genes, or blocking neurotransmitters. While there have been trends of results arising from these endeavors, their interpretation has been difficult owing to several different factors. A primary factor making it difficult to pinpoint behavioral deficits is the apparent expansive redundancy of learning systems in the mammalian brain. A learning function that may be primarily fulfilled by the striatum may be able to be accomplished to a satisfactory level by other brain circuits in the case of striatal malfunction. In this case, the importance of the basal ganglia in the learning process will be underestimated and could be identified only by using sensitive tests for subtle features of the behavior rather than overall performance, as will be discussed later. Thus, using a permanent lesion strategy works well if a strong effect of the lesion is observed, but may not be reliable if such an effect is not observed. Alternatively, in recent years, the ability to acutely manipulate subsets of neurons using optogenetic techniques has led to the possibility of probing the role of those neuronal subsets in real-time with little possibility left for compensation. However, this strategy poses an equally large risk for misinterpretation given a positive result due to the major acute changes introduced to the brain circuits by the sudden manipulations which are likely to disrupt homeostasis and affect many associated brain regions, resulting in a behavioral change even in the absence of a meaningful functional role of the directly disturbed neurons (Otchy et al., 2015). In this case, a positive result in the form of a behavioral disturbance may not be reliable; however, a negative result may be more informative.

The vast majority of the research on the role of the striatum in behavior has been conducted using long-term lesioning strategies, and sensitive tests have been required to identify features of behavior that are altered by these manipulations. Some basic behaviors that appear to be preserved in case of striatal lesions are basic motor behaviors and stimulus-response association learning. However, subtle features of such behaviors may be altered based on the area of damage. As an example, rats with lesions in dorsal striatum or in hippocampus can both learn a

two-choice maze task; however, they seem to use different strategies to learn it. By using a cross-maze and changing the starting location of the animal on the cross-maze it was shown that rats with an intact hippocampus but not striatum used a place guided strategy to learn the maze, and rats with an intact striatum but not hippocampus used a response (left or right) guided strategy to learn it (McDonald and White, 1994, Packard and McGaugh, 1996, Compton, 2004). Such distinctions between learning strategies employed by different striatal regions and by hippocampal or prefrontal cortical circuits have been under deep investigation (Johnson et al., 2007, van der Meer et al., 2010, van der Meer and Redish, 2011). Actor-critic models are another framework which has been proposed to be implemented by striatal sub-regions for the learning of instrumental behaviors. In such models, the actor part of the circuitry carries out the selected action and the critic part of the circuitry provides feedback about the desirability of the selected action to train the actor. In some cases, the ventral striatum has been suggested to play the role of the critic in instrumental conditioning while the dorsal striatum plays the role of the actor (Atallah et al., 2007, van der Meer and Redish, 2011).

A major behavioral feature that has been thought to be at least partially under striatal control is the axis of goal-directed versus habitual behavior. In both goal-directed and habitual modes, the outward behavior of the animal in the form of choice and running time is likely to look similar. However, the motivations guiding those behaviors may be very different. Common wisdom argues that younger students are more adaptable in their behavior and that people who have practiced the same behaviors for many years are less capable of adaptation. Supporting this observation are human and animal behavior studies in which it was found that in early stages of learning behavior tends to be goal-directed and in late stages of training after extensive repetition it transitions to being habitual (Dickinson and Adams, 1983, Dickinson, 1985, Balleine and O'Doherty, 2010). Goal-directed or action-outcome behavior is one that is guided by the desire for a particular outcome and is capable of adjustment based on changing conditions in the animal's environment, but is usually slower and less efficient (Keramati et al., 2011). Habitual or stimulus-response behavior is defined as behavior that is inflexible and well-practiced such that it is performed automatically, in a stereotyped manner, and is less dependent on the outcome of the action. Each of these types of behavior is thought to be competitive in different conditions;

goal-directed in uncertain new conditions, and habitual in well-learned behaviors in certain conditions (Daw et al., 2005). However, when the optimal balance between these behaviors is disturbed such as in drug addiction (Everitt and Robbins, 2005, Belin et al., 2009), stressful conditions (Schwabe and Wolf, 2011), or obsessive compulsive disorder (Gillan et al., 2011) the result is maladaptive and detrimental to quality of life.

A behavior can be determined to be habitual in an experimental setting if an animal continues to perform the behavior after the reward is taken away or no longer desirable (Holland and Straub, 1979, Dickinson, 1985). Researchers have used this approach to learn which brain areas are necessary for goal-directed and habitual behavior. Numerous lesion studies have found dissociated roles of the dorsomedial striatum (DMS) and the dorsolateral striatum (DLS) in goal-directed behavior and habitual behavior. When the DMS was subjected to excitotoxic lesions, the process of the transition from goal-directed to habitual behavior was accelerated; but with lesions of the DLS behavior continued to remain goal-directed even after extensive training (Ragozzino et al., 2002, Yin et al., 2004, Yin et al., 2005b, Bailey and Mair, 2006, Yin et al., 2006, Balleine and O'Doherty, 2010). When NMDA transmission was selectively blocked in one of these areas, similar effects were observed implicating the glutamatergic input and synaptic plasticity to the striatum in these functional roles of DMS and DLS (Yin et al., 2005a, Dang et al., 2006). Moreover, potentiation of synaptic strengths in the DMS seems to occur only in the early stages of training on the rotarod and returns to native levels in late stages of training, while potentiation of synaptic strengths in the DLS occurs in the late stages of training (Yin et al., 2009).

A behavioral process that is likely highly related to habit formation is the learning and chunking of action sequences, a phenomenon where actions performed as part of a functional fixed sequence become perceptually grouped together (Perruchet and Amorim, 1992, Koch and Hoffmann, 2000, Verwey, 2001, Verwey et al., 2010). One behavioral measure reflecting action sequence chunking is the increase of the reaction time for the first element of the action sequence to allow for pre-planning the full sequence and a decrease in the reaction time for the remaining elements of the action sequence, as well as a slowing of performance when the order of the required actions is changed in a serial reaction time task (Koch and Hoffmann, 2000, Kennerley et al., 2004, Bailey and Mair, 2006, Acuna et al., 2014). Experiments using lesioning

strategies and these behavioral measures of action sequence chunking have found an important causal role for the striatum in this behavioral process in mouse nose poke sequences and rodent grooming sequences; as well as neural correlates of chunking in human putamen in sequential key press experiments (Van den Bercken and Cools, 1982, Berridge and Whishaw, 1992, Berridge et al., 2005, Bailey and Mair, 2006, Wymbs et al., 2012, Graybiel and Grafton, 2015). Although habits can range from simple to complex series of behaviors, one feature that is shared by chunked action sequences and habits is their rigidity, automaticity, and slow development through training – suggesting a fundamental link between these behavioral processes and the involvement of dorsolateral striatum.

In summary, the behavioral effects of lesions in different striatal subregions can be subtle and complicated likely due to multiple redundant neural systems learning experimental tasks in parallel, and using different strategies to do so. However, there is much evidence pointing to differential roles of the ventral striatum possibly as the critic teaching the dorsal striatum, the dorsomedial striatum in early stage goal-directed behavior, and the dorsolateral striatum in later stage rigid stimulus-response habitual behavior and action sequence chunking. How the circuits in these striatal subregions participate in biasing or controlling behaviors in such ways and how the striatal inputs and outputs reflect these functions is still very poorly understood.

Movement correlates in dorsolateral striatum

In this part of the introduction, I will focus primarily on results from neuronal recordings in the sensorimotor or dorsolateral striatum – the focus of this thesis work. The activity of neurons in the dorsolateral striatum has been monitored in primates and rodents under a large variety of experimental conditions including slice physiology, under anesthesia, and in awake animals performing complicated behavioral tasks. One of the primary outcomes of this body of literature has been the understanding that neuronal activity in the striatum is often extremely heterogeneous, that the sources and function of this heterogeneity are difficult to explain, and that there are not often cleanly separable sub-types of neuronal responses but a continuous distribution of them. Descriptions of neural correlates in the dorsolateral striatum may be

roughly divided into those reflecting basic movement or sensory related parameters, and those reflecting higher-order parameters related to decision making, habits, or rewards.

One prominent line of electrophysiology research in the dorsolateral striatum has focused on identifying movement-related neuronal responses. Three of the most common strategies used to explore such movement correlates in dorsolateral striatum have been (1) moving or stimulating specific body parts while recording from neurons in DLS, (2) recording head direction and overall locomotion speed in freely moving animals, (3) stimulating the direct or indirect pathway in DLS and observing the effects of this stimulation on locomotion. A tremendous amount of work on this topic has been conducted by Mark West and colleagues, their careful video observations in freely moving rats identifying moments of movement and direction of movement of particular body parts and identifying neurons in the DLS that appear to respond preferentially during the movement of those body parts. Carelli, West, and others found neurons in dorsolateral but not dorsomedial striatum that were responsive to the movement, active manipulation, or touch of ipsilateral/contralateral forelimbs, hindlimbs, vibrissae, shoulders, trunk, head, neck, snout, and chin (DeLong, 1973, West et al., 1990, Carelli and West, 1991, Mittler et al., 1994). The organization of those body part related neurons in DLS was not similar to that of the homunculus in the motor cortex, but clusters of neurons that appeared related to particular body parts were intermingled throughout all three dimensions of the DLS and arranged in longitudinal strips as had been previously shown in some anatomical work (McGeorge and Faull, 1989). This group later demonstrated similar properties in the DLS of mice (Coffey et al., 2016). It is important to note, however, that during these observations there were no analyses conducted on the context within which these single movements of body parts and the large number of confounding factors that likely co-varied with the movement of the individual body parts were not considered.

The same authors later conducted an experiment in which they trained rats to press a lever after a cue (Carelli et al., 1997). They found DLS neurons that appeared to fire in relation to the movement of the contralateral forelimb used to press the lever in the early session of training, but stopped doing so in the later sessions. Most of the forelimb-responsive neurons they recorded lost responsiveness to the movement after 4-9 days of training on the lever with 70

trials per day. They demonstrated that this decrease in lever-press related firing could not be accounted for by changes in the movements the rats did, a reduction in the force they used to press the lever, sampling differences across sessions, or tissue damage. One of the authors' conclusions is that "gradual disappearance of striatal firing suggest that movement-related activity may cease during certain movements that have become automatic or habitual, but not before that activity may have contributed to the formulation in other areas (e.g., premotor areas) of computations needed to carry out the automatic movement." (Carelli et al., 1997) In studies with strong parallels to that by Carelli and West, Tang et al. found decreased activity of 89% of DLS head-movement related neurons as rats' behavior became stereotyped in a task requiring vertical head movements (Tang et al., 2007); and found decreased activity in licking-related DLS neurons during the acquisition and overtraining of a licking task (Tang et al., 2009).

In addition, careful study of the timing of these movement-related responses in the primate DLS revealed that the firing onset of most forelimb related neurons lags behind the onset of the electromyogram (EMG) activity of the forelimb muscles, indicating that it is unlikely for the striatal neurons to be controlling the muscle movement (Crutcher and DeLong, 1984, Liles, 1985, Kimura, 1990). Direct blockade of neural activity with the use of muscimol targeted to putamen, the globus pallidus internal segment, and external segment resulted in decreased efficiency and increased variability in arm movements but did not inhibit movement (Kato and Kimura, 1992).

Neural correlates of head movement, head direction, and overall velocity of the animal in DLS have also been a focus of study (Kim et al., 2014, Rueda-Orozco and Robbe, 2015), with overall body acceleration and velocity in an open field recently becoming a popular simple measure of activity levels in rodents (Venkatraman et al., 2010, Barter et al., 2015, Rueda-Orozco and Robbe, 2015), and especially in animal models of basal ganglia disorders. However, it can be argued that based on the primary role of the striatum in learning and habit formation, the complexity of movement representations in the striatum (DeLong, 1973, Carelli and West, 1991) along with the modulation of those movement representations with learning and context (Kimura, 1990, Carelli et al., 1997, Tang et al., 2007, Tang et al., 2009), we need to consider the relationship of striatal neuronal activity to many factors in addition to movement.

Correlates of habits and rewards in dorsolateral striatum

How the neural activity in the DMS and DLS contribute to habit learning is still not understood, but some studies with recordings from these areas during learning provide initial clues. Several chronic tetrode recording experiments in our laboratory have shown that the activity of the neurons in the striatum changes as a rat learns to perform a T-maze task. We found that with learning, neurons in the dorsolateral striatum (DLS) transition from firing throughout the task to firing mostly at the beginning and end of the task (Jog et al., 1999, Barnes et al., 2005, Thorn et al., 2010). Because DLS is required for habitual behavior, this beginning and end activity was proposed to be a neural basis for action sequence chunking. The DLS activity could provide an initiation signal that triggers downstream areas holding the memory of the learned behavior, such as motor or associative motor cortices where plasticity is known to occur after extensive motor training (VandenBerg et al., 2002, Kleim et al., 2004) or subcortical motor areas, to execute the entire sequence of actions required to perform the task. This interpretation was supported by the added finding that when the task was altered to allow for complete pre-planning of the entire motor sequence (the cue about which direction to turn on the T-maze was provided before the animal started running rather than after), the degree of depression of firing mid-task increased accentuating the beginning and end pattern (Barnes et al., 2011). In contrast, in the same task, DMS neurons tended to be active in the middle of the trial at the decision point of whether to turn left or right early in learning and the activation disappeared late in learning when DLS activation was strongest (Thorn et al., 2010). The DMS neural activity was also consistent with the known role of DMS early in learning when behavior is goal-directed and decisions are made with the purpose of obtaining the related outcome. Recently, start-stop activity similar to that seen in our laboratory in the DLS was found in the dorsal striatum and substantia nigra in a fixed ratio schedule lever press task (Jin and Costa, 2010). The contrast between these task and learning dependent neural activities in DMS and DLS seems to be consistent with the roles that lesion studies suggested for these areas. However, it is unclear whether the fleeting decision-related activity in DMS and the beginning and end activity in DLS that persists into late training are a general phenomenon occurring in a range of tasks and what the true relationship is of the neural firing to the task performance.

Why is this work important?

With the expansion of neocortex in humans, these newly emerging cortical areas have continued to provide input to the striatum utilizing this evolutionarily ancient learning structure for new learning capacities while maintaining much of the organization and functional divisions of earlier mammals. Research on the mammalian striatum has established its critical function in the learning and performance of reinforcement-based instrumental and habitual behaviors and disturbances in these behaviors in a range of neurological disorders including Tourette's syndrome, Parkinson's disease, Huntington's disease, drug addiction, dystonias, and obsessive compulsive disorder (Belin et al., 2009, DeLong and Wichmann, 2010, Koob and Volkow, 2010, Gillan et al., 2011). Outside of clinically diagnosable deviations from normal basal ganglia function, these circuits likely play a big role in the individual variability present among children and adults in learning and various behaviors. A primary common factor in these disorders and non-optimal behavior patterns is the imbalance between flexible goal-directed behavior and automated habitual behavior which are thought to be controlled by these complimentary associative cortical-dorsomedial striatum circuits and sensorimotor cortical- dorsolateral striatum circuits. Understanding the principles of operation of these circuits is critical for developing strategies of therapies and for understanding non-optimal habit behaviors. In addition to the therapeutic implications of understanding the neural basis of these processes, understanding how habitual action sequences are learned and executed by the brain is an important step toward the long-term goals of fundamental neuroscience. Towards these goals, we can start by addressing how cortical input to the striatum is integrated and transformed within the striatum, how this can occur differently in associative and sensorimotor striatal areas, and the role of dopamine input in this process.

CHAPTER 2: GENERALIZED STRIATAL SIGNAL FOR CHUNKED BEHAVIORAL REPERTOIRES

SUMMARY

Habits consisting of a series of actions control much of our waking behavior. The dorsolateral striatum is known to be important for habitual behaviors and action sequence chunking, but the circuits underlying this function are still mysterious. To explore the nature of striatal neural representations during learned action sequences, I designed a task specifically targeted at disambiguating movement-related responses from habit representations in dorsolateral striatum. I found that, unlike in motor cortex, neurons in the dorsolateral striatum did not respond simply to particular individual actions, but encoded the start and end of the learned action sequence within which the individual actions occurred. This activity pattern generalized across a wide variety of movement sequences learned by different rats. Remarkably, when rats did unreinforced action sequences containing similar sub-movements, the same neurons failed to exhibit the task-boundary activity. Motor cortex did not recapitulate the task-boundary activity seen in striatum and did not appear to be a primary driver of the striatal start and end signal, further suggesting that the task-boundary activation was not controlled by movement itself. In contrast, I found that inhibitory interneurons are likely involved in shaping the task-boundary activity by increasing firing rates mid-task when striatal projection neuron activity was suppressed. These experiments provide a definitive test for the existence of a striatal signal that selectively marks behavioral units consisting of learned action sequences which could enable the encoding and expression of such chunked action sequences by the basal ganglia.

INTRODUCTION

The capacity to string together behavioral repertoires is critical in human and animal behaviors. Some such repertoires are hard-coded into the nervous system and are termed fixed action patterns, while others are learned throughout a lifetime and serve as the building blocks of our daily activities. When a series of actions are repeatedly performed together, these actions become "chunked" into a single behavioral unit (Lashley, 1951, Rosenbaum et al., 2007). Such action sequences are often carried out in a stereotyped manner upon being triggered, and become inflexible, automatic, and habitual after extensive repetition (Graybiel, 1998, 2008). Importantly, such behavioral repertoires vary widely in terms of the nature and number of the movements involved. Libraries of learned action sequences underlie much of our daily behaviors, allowing allocation of attention to new or priority tasks at hand while minimizing the effort toward accomplishing well-rehearsed optimized tasks. As a result, such habitual action sequences can increase efficiency, but they can also dominate behavior and can be extremely difficult to alter even when required (Daw et al., 2005, Graybiel, 2008, Dolan and Dayan, 2013, Dayan and Berridge, 2014, Graybiel and Grafton, 2015). Overexpression of such stereotyped behaviors is maladaptive and is thought to be a defining feature in many psychiatric disorders including obsessive compulsive disorder, Tourette's syndrome, dyskinesias, drug addiction, and can also result in perseverative maladaptive habits in otherwise healthy people (Miltenberger et al., 1998, Leckman and Riddle, 2000, Berridge et al., 2005, Koob and Volkow, 2010). Despite the importance of this behavioral process, how the brain represents such chunked units of behavior has not been well characterized; however, most prominent models of basal ganglia suggest a prominent role of the striatum in the selection of such action programs (Stephenson-Jones et al., 2011, Friend and Kravitz, 2014, Graybiel and Grafton, 2015). The dorsolateral striatum, in particular, found to be important for the transition of behavior from goal-directed to habitual, stereotypies, and action sequence chunking, may be involved in representing such behavioral units (Aldridge and Berridge, 1998, Jog et al., 1999, Aldridge et al., 2004, Yin et al., 2004, 2005a, Yin et al., 2005b, Yin et al., 2006, Graybiel, 2008, Graybiel and Grafton, 2015). Recordings from the dorsolateral striatum during the running of a T-maze in rats (Jog et al., 1999, Barnes et al.,

2005) and during an FR8 lever press task in mice (Jin and Costa, 2010) have shown that some neurons in dorsolateral striatum fire in a manner that accentuates the beginning and end of the trial.

In parallel lines of study, dorsolateral striatal spiny projection neurons (SPNs) have long been thought to have patterns of activity that correlate with movement and motor behaviors (DeLong, 1973, Carelli and West, 1991, Kim et al., 2014, Rueda-Orozco and Robbe, 2015), but this activity has been known to change with learning (Carelli et al., 1997, Tang et al., 2007, Tang et al., 2009) and which aspects of behavior are represented by the striatal neurons has not been pinned down due to the many different factors that can co-vary with movement such as the intention of action and history of reinforcement.

To tease apart such factors, I designed a task consisting of several ordered steps requiring similar movements. In this task, each rat learned to perform one specific sequence of three lever presses on a set of two levers. Lever-press related movements were thus present in the beginning, middle, and end of the learned sequence, as well as in the unrewarded incorrect sequences that they performed (Fig. 1A). As such, I could use the lever presses performed within different time points in reinforced and unreinforced sequences to determine the relationship of the striatal spiking activity to the behavior. Within single training sessions the behavior of trained rats oscillated between "in the zone" periods during which they performed a high proportion of correct sequences, and random or exploratory periods during which they performed below chance level with many incorrect sequences, as well as self-initiated rest periods (Fig. 1B).

It also remains to be explored how representations of motor behaviors in dorsolateral striatum are different from motor cortex and other motor-related regions, and what kind of transformations in neural representations occur with each link in the cortico-basal-ganglia-thalamic circuits. I compared the activity of dorsolateral striatal neurons and motor cortex neurons during these periods and in relation to lever pressing movements. I found that during in-zone periods, the activity of striatal neurons, but not motor cortex neurons, peaked at the initiation and termination of the learned sequence of presses and that this pattern of activity was not present at times in which they were performing incorrect sequences of presses. Thus,

striatum signaled the boundaries of chunked behavioral repertories in a manner that appeared independent of the exact motor components of the behavior.

RESULTS

Rats learn individualized stereotyped movement patterns to execute correct sequence

Given the lack of cues for which lever to press and the eight possible three-step lever press sequences they could perform for each trial, the rats' performance began below chance level (12.5%) and gradually improved over the course of 35+ days of training ($p < 0.001$, Mann-Whitney test). Within single sessions, well trained rats exhibited their potential for high levels of performance of the correct sequence during periods in which they performed many consecutive correct trials. During such high-performance periods, rats' performance was usually above 80% correct (Fig. 1D) and they were performing a correct trial every 11 seconds including several seconds required for reward consumption (Fig. 1E).

However, during other times within the same sessions rats would often enter periods of poor performance or exploratory behavior, or would enter self-initiated rest periods (Fig. 1B). In the first 10 days of training, the rats spent 50% of the total active pressing time in the session performing below chance level, usually with a large number of trials in which they repetitively pressed the same lever, 48% of the total active time performing between 12.5%-50% correct trials, and 2% of the time performing majority correct trials. They also spent 30% of the total time in the operant chamber resting with no lever pressing (Fig. 1C). In the 31-40 days of training rats spent 24% of their active time performing at chance levels or below, 21% of the active time performing between 12.5%-50% correct trials, and 55% of the time performing majority correct trials. They continued to spend a large proportion of the time in the operant chamber, 43% of the total time, resting which was increased from earlier sessions likely due to large volumes of reward they received during high-performance periods (Fig. 1C).

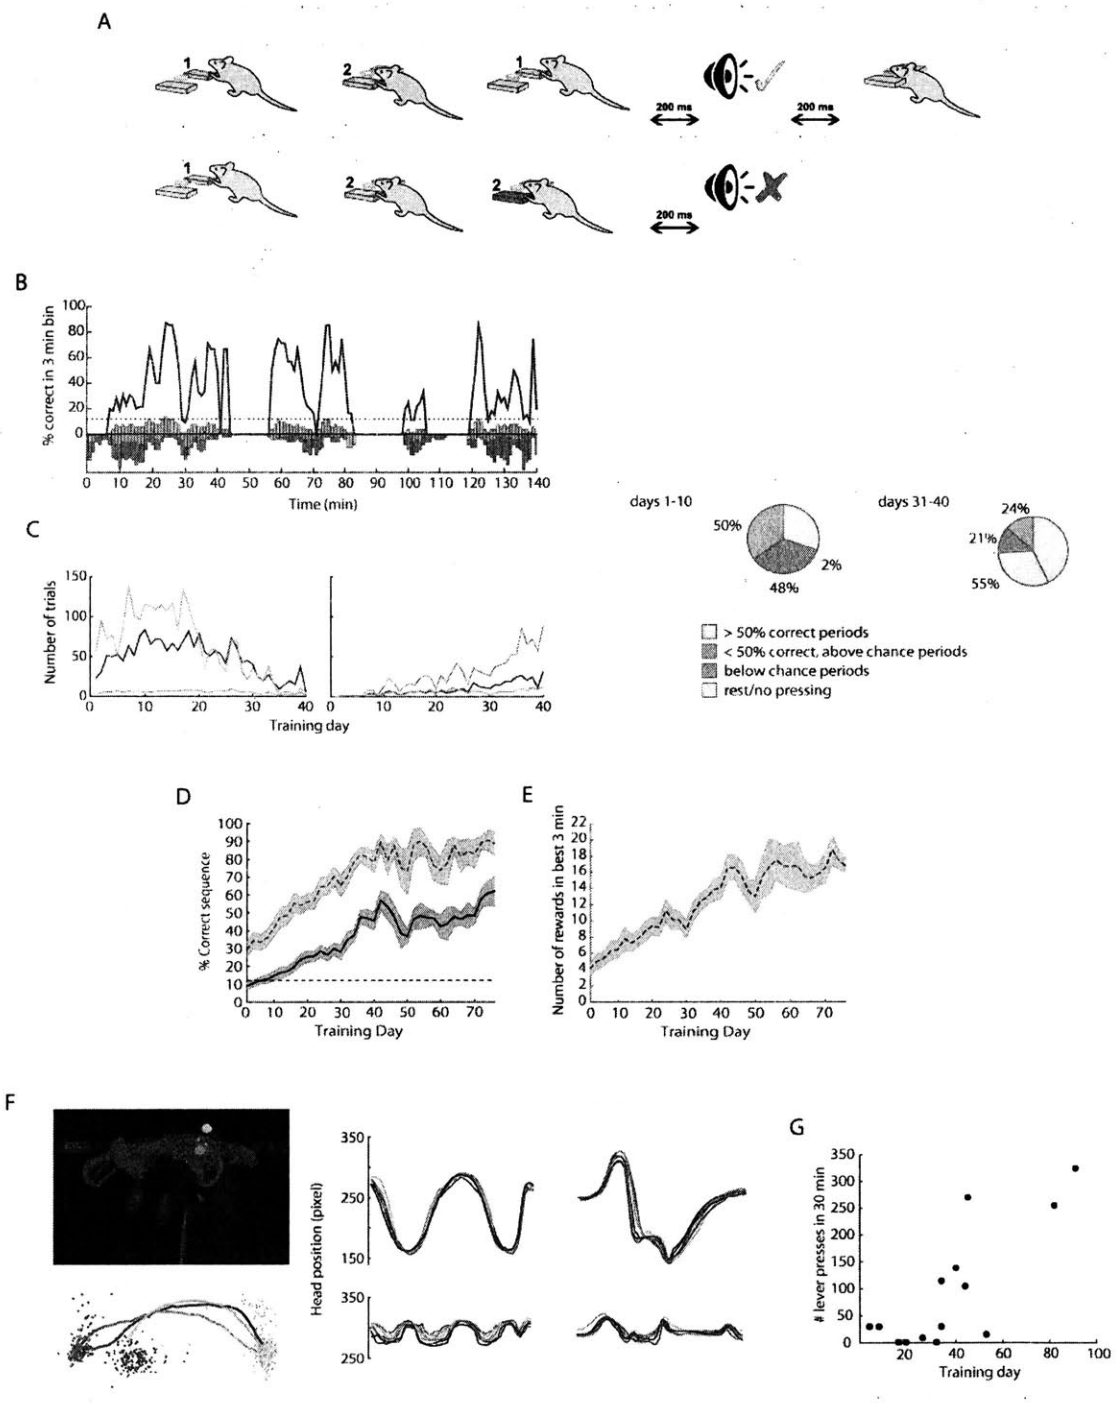
A
1
2
1
200 ms
200 ms
1
2
2
200 ms
B
% correct in 3 min bin
Time (min)
C
Number of trials
Training day
days 1-10
50%
2%
48%
days 31-40
24%
21%
55%
> 50% correct periods
< 50% correct, above chance periods
below chance periods
rest/no pressing
D
% Correct sequence
Training Day
E
Number of rewards in best 3 min
Training Day
F
Head position (pixel)
G
lever presses in 30 min
Training day

Figure 1

Figure 1. Rats learn to perform specific three-step lever press sequences.

(A) Each rat was assigned a specific 3-step sequence to be performed on two levers located on the right (lever 1) and on the left (lever 2) of the reward well. After each three presses, if the sequence was correct a click auditory stimulus was played with a 200ms delay and a chocolate milk reward was delivered 200ms after the click. If any other sequence of levers was pressed a white noise stimulus was played with a 200ms delay.

(B) A single training session during which the rat oscillated between periods of high correct performance, periods of low correct performance, and rest periods. Blue line indicates percent correct within 3 min bin. Green histogram indicates number of correct trials within each 3 min bin. Stacked yellow and red histograms indicate number of repeat press (1-1-1 or 2-2-2) or non-repeat press incorrect trials within each 3 min bin.

(C) On left, number of total repeat press incorrect trials (yellow), non-repeat incorrect trials (red), and correct trials (green) within periods of below chance level (12.5%) performance across training. On right, number of total such trials during periods of majority correct performance across training. In pie charts, The distribution of <12.5% correct performance periods, 12.5%-50% correct performance periods, and > 50% correct performance periods in the first 10 days of training and in the days 31-40 of training.

(D) Solid line, the total session percent correct across training days (n = 13 rats). Only 4 unimplanted rats were trained past day 40. Dashed line, percent correct in the best 3 min performance period in session across training days in same rats. Error bars indicate ±SEM.

(E) Number of correct trials performed during best 3 min performance period in each session across training days (n = 13 rats). Error bars indicate ±SEM.

(F) On top left, head tracking data from consecutive correct trials in one training session. Bottom left, the average session trajectory and of a rat performing the 1-2-1 sequence with head locations for the lever press events and reward delivery. On right, head side-to-side (top, x) session average trajectories and forward-and-back trajectories (bottom, y) from 7 consecutive training days in a rat performing the 1-2-1 sequence (left column) and a rat performing the 1-2-2 sequence (right column).

(G) Number of lever presses done in a 30 min devaluation probe session as a function of number of days trained. In devaluation probe sessions rats had access to unlimited chocolate milk for 2 hours prior to training. No milk was delivered during probe sessions.

Each of the rats developed specific movement patterns to complete the correct learned sequence, as was evidenced by head tracking data. While performing the learned sequence, the rats' head position moved in stereotyped manner that was similar trial-to-trial, and session-to-session. These movement patterns varied greatly between rats based on the sequence they

learned and the specific movements they developed to successfully execute the sequence (Fig. 1F).

I tested a separate group of four behavior-only rats for devaluation resistance by providing them with unlimited chocolate milk for two hours prior to placing them in the operant chamber. I found that after extensive training, these rats became devaluation resistant to the reward devaluation (Mann-Whitney test, $p < 0.05$) (Fig. 1G). The slow trial-and-error learning process, the development of stereotyped movement patterns, and the development of devaluation resistance further suggests that the rats' were developing a habit throughout the course of the lever press sequence training. However, there were periods in which even well trained rats changed their behavior and appeared to be anxiously and repetitively pressing the same lever or exhibiting exploratory behavior by attempting many different sequences of presses. These out-of-zone periods provided an opportunity to compare neuronal activity in dorsolateral striatum and motor cortex during the performance of the learned action sequence and during the performance of other unrewarded sequences of lever presses.

Motor cortex, but not dorsolateral striatum task-related neuronal activity can be accounted for by individual motor actions

The question of how neuronal representations are transformed with each successive node in the cortico-basal ganglia-thalamic circuit is a fundamental one for understanding the mechanisms of function of these circuits. In this study, I recorded from two such nodes, a forelimb area of the motor cortex (MC) and its target striatal projection zone in the dorsolateral striatum (DLS). For each recorded single unit, I assessed the pattern of firing rate modulation in the correct sequence and any incorrect sequences the rat performed within the given training session. Specifically, I addressed whether the neuron responded during lever 1 or lever 2 presses, or during transitions from one of the levers to the other lever, and whether the response to those events was similar whether they occurred at any point during the rewarded sequence or during other sequences of presses (See methods). If the unit did fire similarly in relation to a given lever press regardless of

when it occurred within the rewarded sequence and unrewarded sequences then I could argue that the spiking response of the unit could be accounted for by the occurrence of the given motor action alone. Many units in motor cortex did fit this criterion – one such example motor related unit is shown in Fig. 2A. However, a majority of the units in dorsolateral striatum, did not fulfill this criterion: they were more strongly modulated during the same motor events during some sequences than others. Two examples of such striatal neurons are shown in Fig. 2B and C.

I assessed the relative distribution for such simple motor units in motor cortex and DLS by identifying sessions in which I had simultaneously recorded at least 10 putative motor cortex pyramidal neurons and 10 DLS spiny projection neurons (SPNs). I found that in across all the motor timepoints I analyzed (the start of the shift from lever 1 to lever 2, the end of the shift from lever 2 to lever 1, the completion of the resulting lever 1 press, and corresponding events for the transition from lever 1 to lever 2) motor cortex neurons were significantly more likely to fit the simple motor unit criterion than were the DLS units ($p < 0.001$, Mann-Whitney test) (Fig. 1D). This suggests different levels of task representations in motor cortex and striatum where motor cortex neurons are more likely to directly represent movement parameters, whereas DLS neurons, although in the sensorimotor region of striatum, are more likely to modulate their responses based on the context of the behavior.

Dorsolateral striatum SPN firing reflects the boundaries of the learned action sequence across a variety of different learned lever press sequences

Given the minority of simple motor neurons in the DLS, I considered alternative accounts for the task-related firing rate modulations in DLS. One observation that became clear after examination of a large number of single unit response profiles was that there were neurons that responded in a similar manner across rats each of whom had learned a different lever press sequence. Most apparent were neurons that were highly and phasically activated around the time of the first lever press in the learned sequence, around the time of the last lever press in the learned sequence, or both. These neurons were present in rats which learned different sequences of lever presses (Fig. 3A).

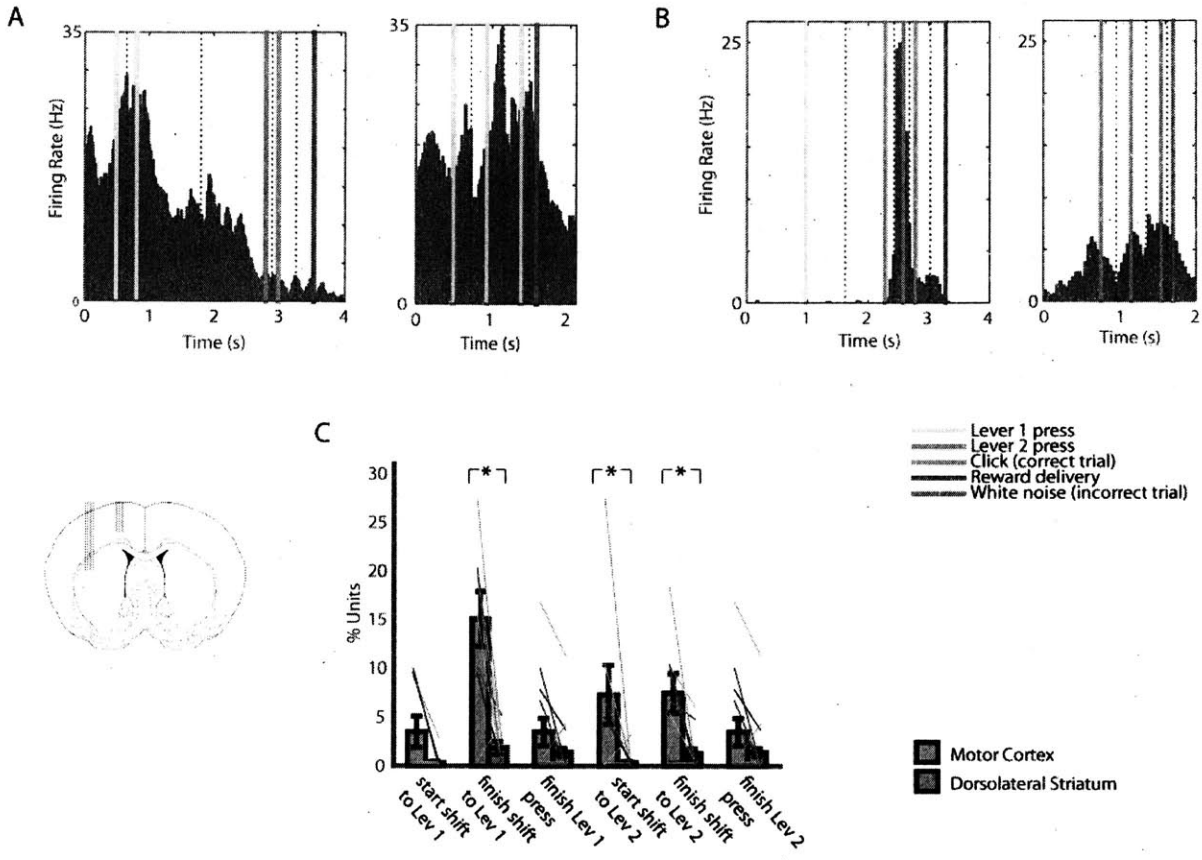

Figure 2. Motor-type neurons are prevalent in forelimb motor cortex but not dorsolateral striatum.

(A) An example of a motor type neuron recorded from motor cortex. Peri-event histograms for each of the events in the trial were pasted together with window sizes according to the median time between each successive event. Event markers are indicated with colored lines and the pasted window edges are indicated with dashed lines. This unit fired during lever 1 presses both in the correct sequence and during repetitive lever 1 pressing.

(B) An example of a non-motor type neuron in dorsolateral striatum. This neuron responded selectively during the initiation of the correct sequence, but not during the same lever press during repeated presses of the lever.

(C) The proportions of motor-type neurons in n = 14 sessions with 10 or more neurons recorded simultaneously in both motor cortex and DLS. Error bars indicate ±SEM. Asterisks indicate significant difference between motor cortex and DLS cell proportions.

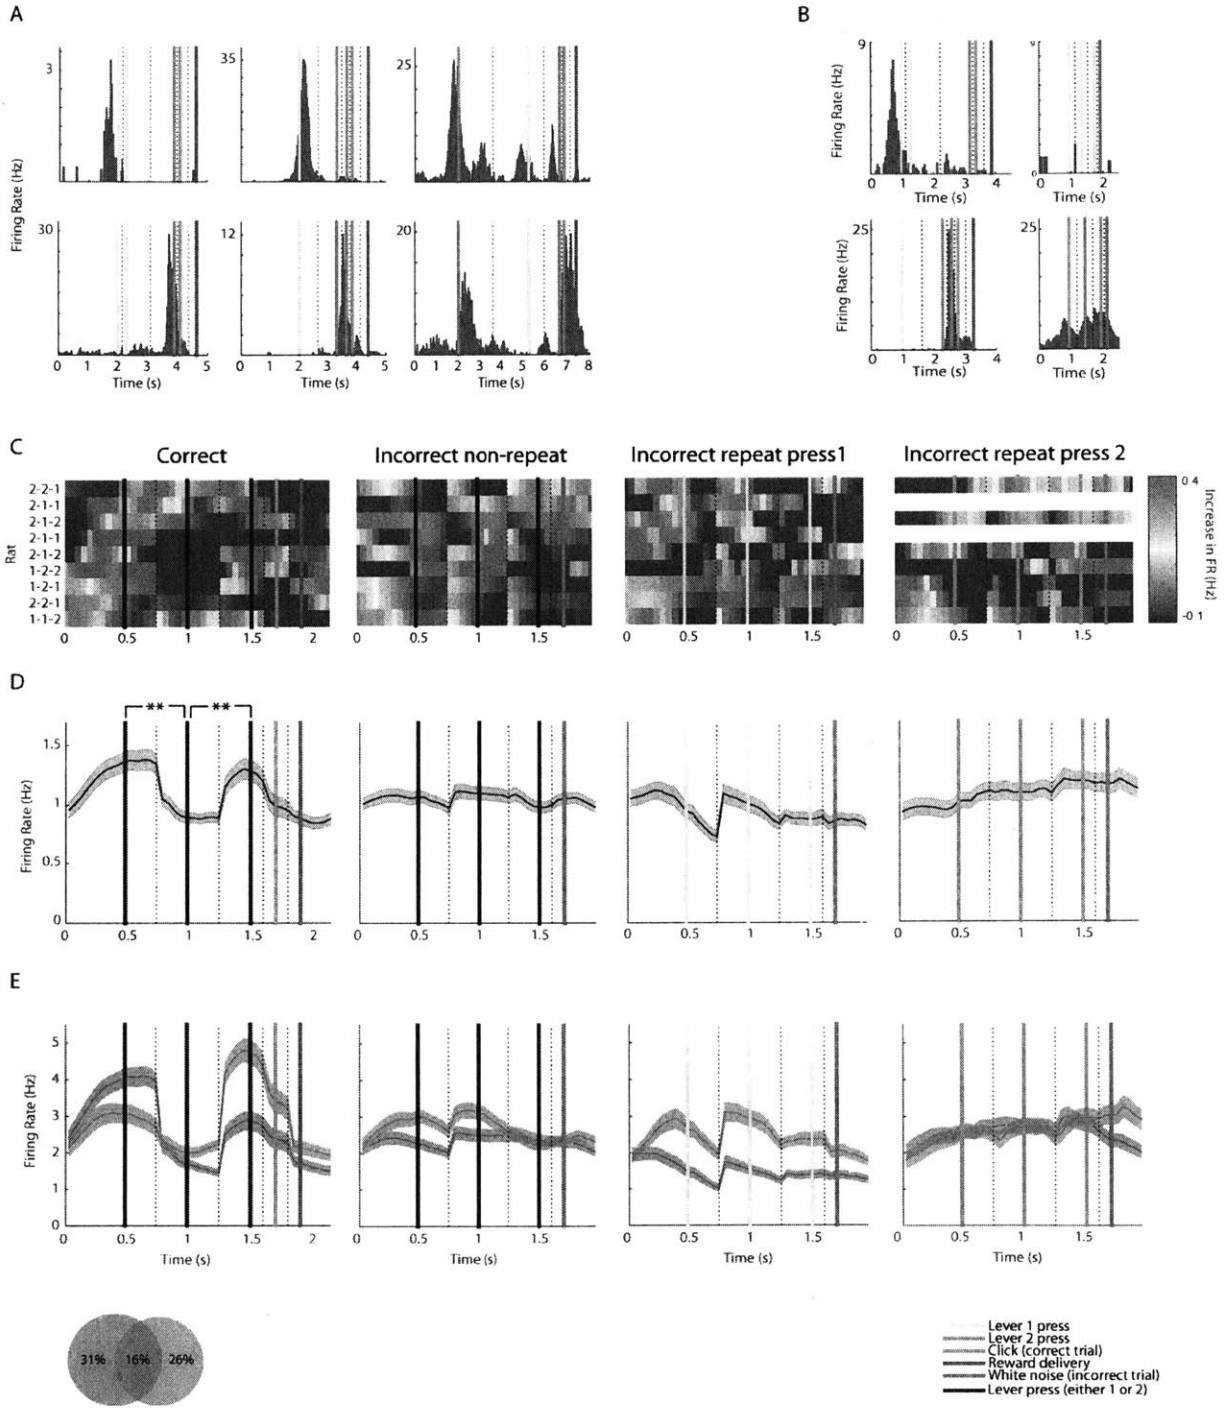
A
B
C
D
E
Firing Rate (Hz)
Time (s)
Correct
Incorrect non-repeat
Incorrect repeat press1
Incorrect repeat press 2
Rat
Increase in FR (Hz)
**
**
31%
16%
26%
Lever 1 press
Lever 2 press
Click (correct trial)
Reward delivery
White noise (incorrect trial)
Lever press (either 1 or 2)

Figure 3.

Figure 3. DLS SPN population spiking is concentrated around the initiation of the learned lever press sequence but not during incorrect lever press sequences, regardless of the lever press sequence learned.

(A) Examples of beginning-type (top tow) and end-type (bottow row) neurons in rats who learned different lever press sequences (1-1-2 first column, 1-2-2 second column, 2-1-2 third column).

(B) Top row, an example of a beginning-type neuron which fired during the first lever 1 press in the correct sequence, but not during the middle lever 1 press. This neuron did not respond during incorrect trials with repetitive lever 1 pressing. Bottom row, an example of an end-type neurons with fired during the last lever 2 press in the correct sequence but not during the middle lever 2 press during the correct sequence. This neuron responded weakly during repetitive lever 2 pressing.

(C) Average peri-event DLS SPN spiking during correct (first column), non-repeat incorrect (second column), lever 1 repeat (third column), and lever 2 repeat (fourth column) in 9 implanted rats who learned different lever press sequences indicated on the left.

(D) Average peri-event DLS SPN spiking across all rats during correct sequence performance (first column, n = 2501 putative SPNs), and non-repeat incorrect trials (second column, same n = 2501 putative SPNs), incorrect lever 1 repeat press trials (third column, n = 1143 SPNs), and incorrect lever 2 repeat press trials (fourth column, n = 1338 SPNs).

(E) Sub-groups of beginning-type (blue) and end-type (red) responsive SPNs in the same trial types. On right, the proportion of total beginning and end-type neuron sub-types and their overlap.

All error bars indicate ±SEM

Commonly these beginning or end neurons which fired in relation to the first or last lever press in the learned sequence, did not fire when the same lever was pressed in the middle of the trial. Similarly, when the same lever was pressed repetitively in an incorrect trial, these neurons were often not modulated or weakly modulated (Fig. 3B).

As a result of these strongly modulated beginning and end neurons the population firing pattern across nine rats, each of whom learned a single sequence, was concentrated around the time of the first lever press and around the time of the last lever press (Fig. 3C, first column). The population of all the recorded DLS SPNs increased firing rates from a baseline inter-trial interval rate of 0.4Hz to 1.3-1.4Hz during the first and last presses in the learned lever press sequence while the average firing rate during the middle lever press was 0.8Hz ($p < 0.01$ for first v. second

press, and second v. third press, Mann-Whitney tests) (Fig. 3C, first column). Although these rats learned different lever press sequences and developed individualized movement patterns in order to correctly complete the rewarded sequence, the population of SPNs accentuated the task-boundaries of the learned sequence across animals.

This task-boundary activity developed quickly, as soon as there were a minimum of 30 correct trials within training sessions and continued throughout training (Fig. S1), similarly to the early development of the beginning and end pattern found in the T-maze task (Barnes et al., 2005, Thorn et al., 2010). In addition to findings that this task-boundary activity occurred in rats each of whom learned different stereotyped movement sequences and in the T-maze task which required locomotion behavior distinct from the movements used during the lever press task; I found that overall speed of the animal was unlikely to account for the task-boundary activation due to the decrease of speed that occurred prior to each time the animal pressed a lever (Fig. S2).

The timing of the peak population spiking varied across rats. This variability could have been produced by a multitude of factors, including the subsets of sampled neurons as the timing of the bursts of firing varied across single units within rats. How the learned sequence was encoded or the perceptual start point of the chunked action sequence could also vary across rats and be reflected in the variability in the spiking. However, in each of the rats, the DLS neuronal population was more active around the time of the first and last lever presses than around the time of the middle lever press, regardless of whether the middle lever press was a lever 1 or lever 2 press, or whether it matched the identity of the first or last lever in the sequence.

By contrast, these same units had a remarkably different pattern of firing when the animals were performing incorrect sequences within the same recording session. The beginning and end activity was absent across nine rats in cases when they were performing an alternative non-repeat sequence (Fig. 3C and D, second column) and in cases when they were repeatedly pressing lever 1 or lever 2 (Fig. 3C and D, third and fourth column). When the rats were pressing the levers in a sequence that was different from the rewarded lever press sequence, the population firing

was weakly modulated prior to or during each of the individual presses in the sequence with no significant differences between the firing rates during the first, second, or third press ($p > 0.6$ for first v. second press, and second v. third press, Mann-Whitney test). The firing rates during the first and third lever press in incorrect trials were significantly lower than in correct trials in the same units, and firing rates during the middle press were significantly higher than in correct trials ($p < 0.000001$ for all three, Mann-Whitney test).

To look specifically at the units which were contributing to the task-boundary activity, I identified neurons that increased their firing rate more than 2 standard deviations above their firing rate during baselines defined as the periods after reward consumption and until 2 seconds prior to the next lever press. According to this criterion, 15% of recorded putative SPNs were responsive around the time of the first lever press, but not the last press, 9% of the SPNs were responsive around the time of the last lever press but not the first press, and 15% of all the units were responsive at both times. Thus, a total of 30% of all SPNs were significantly activated at the time of the first press and 24% of them were activated at the time of the last press. The firing rate modulations of these sub-groups of units were much greater than the modulations of the population average with firing rate increases to 4-5Hz at the times of the first and last lever presses (Fig. 3E, first column). Those units that were most responsive during the time of the first lever press were often also responsive but more weakly during the time of the last lever press, and those that were most responsive during the last lever press were often also responsive but more weakly during the time of the first lever press. The same neurons fired in a different manner when the rats performed incorrect sequences of lever presses or pressed lever 1 or lever 2 repetitively (Fig. 3E, second, third, and fourth columns respectively). A separate group of DLS SPNs that were generally unresponsive during the pressing period, but were activated during the reward consumption period was also present across rats (Fig. S3). This group of SPNs made up 13% of the total population, and were similar in characteristics and proportions to a group of reward responsive neurons recorded in the DLS in the T-maze task (Smith and Graybiel, 2016).

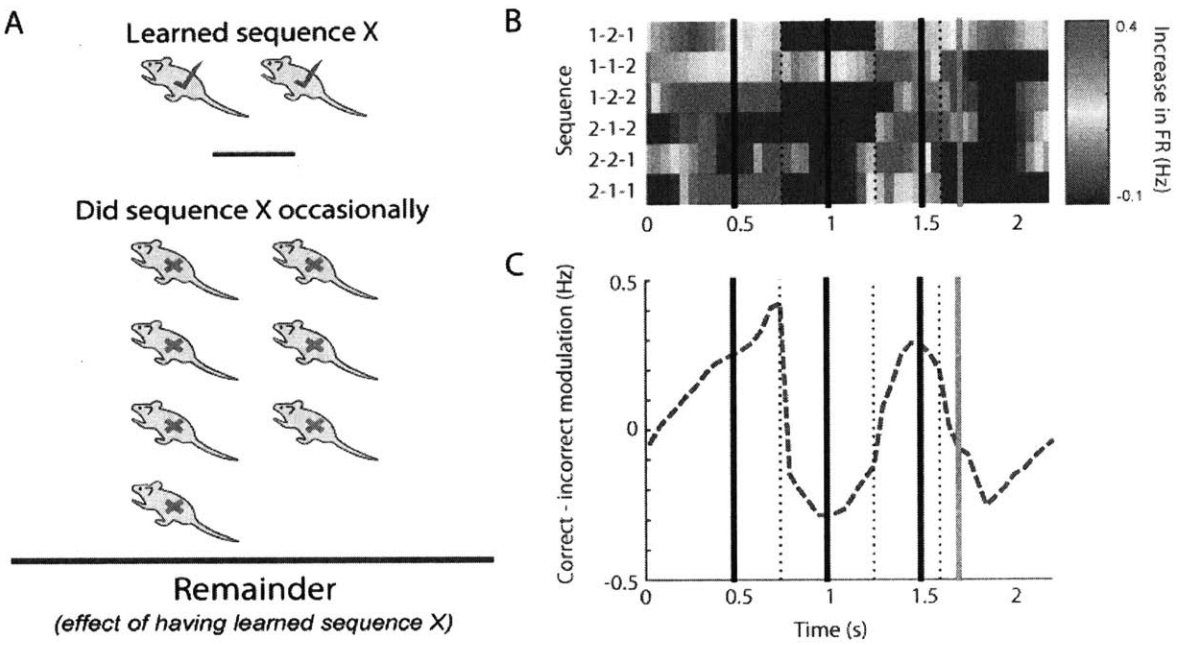

Figure 4. Start and end of specific sequence is preferentially emphasized in rats who were trained on the sequence.

(A) The procedure used to separate effect of having learned a sequence from neural activity related to executing sequence. Neural activity of rats who occasionally incorrectly performed sequence X (bottom) but were trained on a different sequence was subtracted from the neural activity of rats who were trained on sequence X (top). The remainder of this subtraction represents the effect of having learned sequence X.

(B) The effect of having learned each of the six possible sequences using the operation from (A). In each of the 6 possible sequences rats who were trained on the sequence had higher DLS SPN activity at the time of the first and last lever press, and lower DLS SPN activity at the time of the middle lever press.

(C) Average difference in firing rate across the first, second, and last lever press in the sequence in rats who were trained on the sequence v. those that were not.

Experience of reinforcement on behavioral sequence results in accentuation of DLS SPN firing at task boundaries.

I could take advantage of this task design in which different rats were trained on different lever press sequences to compare neural activity in rats trained on a particular sequence to that in other rats who were not trained on that sequence but happened to perform it occasionally without receiving reward (Fig. 4A). By using this strategy for each of the six possible (non-repeat) sequences that rats were trained on, I could identify the effect of having been trained on a particular sequence on the activity in DLS SPNs while attempting to cancel out population activity related to the particular sequence itself. I found that, for each of the six possible sequences, the effect of having been trained on the sequence was an increase in the firing of DLS SPNs at the time of the first and last lever press, and a depression in their firing rate during the middle lever press (Fig. 4B and C). This analysis provides a powerful way to assess the effect of training and reinforcement on a particular set of movements on the activity in DLS, by enabling me to compare the neural activity that occurs in rats who were not trained on the same set of movements but happened to perform them spontaneously. The result reinforces the notion that the task boundary activity in DLS is a signal that is specifically related to marking the start and end of learned behavioral programs.

Task-boundary activity occurs in partial correct sequences but not in the most commonly performed incorrect sequence

To further investigate the nature of this task-boundary representation, I asked whether it (1) occurs when animals do a specific incorrect sequence that they commonly perform, and (2) in circumstances in which they perform a partial correct sequence. First, I considered whether this task-boundary activation might occur in action sequences that are frequently performed but are lacking in history of reinforcement. In many cases, I encountered sessions in which rats favored performing particular incorrect sequences but which were not reinforced. In these cases, I analyzed the population firing pattern in DLS and found that despite the repeated occurrence of

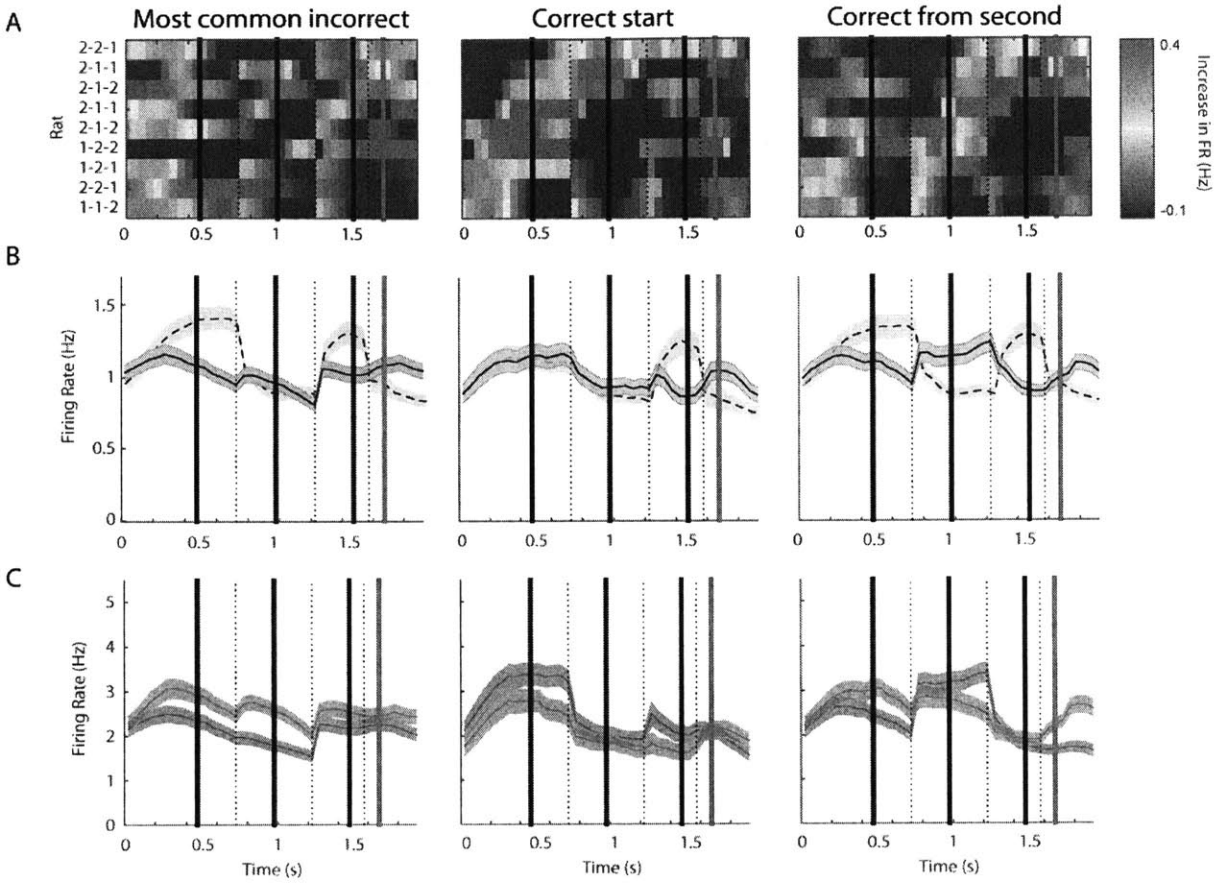

Figure 5. Most commonly performed unreinforced sequence does not produce task-boundary activity in DLS, but partially complete correct sequences result in activation of beginning-type neurons.

(D) Average peri-event DLS SPN spiking during the most common incorrect sequence (first column), trials in which the first two presses were correct but the last was incorrect (second column), and trials in which the second and third press were the first two presses of the correct sequence (third column) in 9 implanted rats who learned different lever press sequences indicated on the left.

(E) Average peri-event DLS SPN spiking across all rats during the same sets of trials. In dotted lines, the spiking pattern is shown for the identical neurons during correct sequence performance (n = 2501 units, n = 1920 units, and n = 2501 units respectively).

(F) Sub-groups of beginning-type (blue) and end-type (red) responsive SPNs in the same trial types (n = 822 beginning-type cells and 693 end type units, n = 551 and 494 units, and n = 769 and 654 units, respectively)

All error bars indicate ±SEM

the given incorrect lever press sequence, the pattern of neuronal responses was similar to other incorrect trials and task-boundary representations were not prevalent (Fig. 5A and B, first column). This result suggests that the dorsolateral striatum may be selectively involved in encoding and chunking of action sequences with previous history of reinforcement. However, it could also be the case that these commonly performed incorrect sequences did not undergo long term practice as the rewarded sequence, and thus were not subject to becoming habitual to the same degree as the rewarded sequence.

In the cases in which rats started the correct sequence correctly by performing the first two lever presses correctly but pressed the incorrect lever for the last press in the trial, I found that the population firing of the SPNs at the trial initiation was not significantly different from that in correct trials. However, the typical peak at the last lever press in the sequence was missing (Fig. 5A and B, second column). If instead, the rat self-corrected by beginning the correct sequence in the middle of a trial by pressing the first two steps of the correct sequence starting from the second press in the trial, the population activity appeared to shift with a peak occurring at the second lever press and a dip in activity occurring at the time of the last lever press, which would normally correspond to the middle press in the correct sequence (Fig. 5A and B, third column). Thus, if the rat fails to successfully finish the correct sequence, the completion signal is omitted and if the rat begins the correct sequence mid-trial after an initial mistake DLS beginning-responsive neurons become active. These results suggest an active role of the DLS in responding or controlling the successful initiation and completion of the correct action sequence, rather than passive encoding of trial structure; however, further study is required to distinguish between the possible interpretations of these results.

DLS activity marks beginning and end of learned sequence similarly in isolated correct trials and correct trials within high performance periods

Given the within session fluctuations in the rats' accuracy in performance, I asked how DLS SPN firing might differ during isolated correct trials in random or exploratory periods in which the rats' accuracy was below chance level (12.5%) compared to periods in which they were performing a majority of correct trials. To do so, I selected correct trials performed in only such

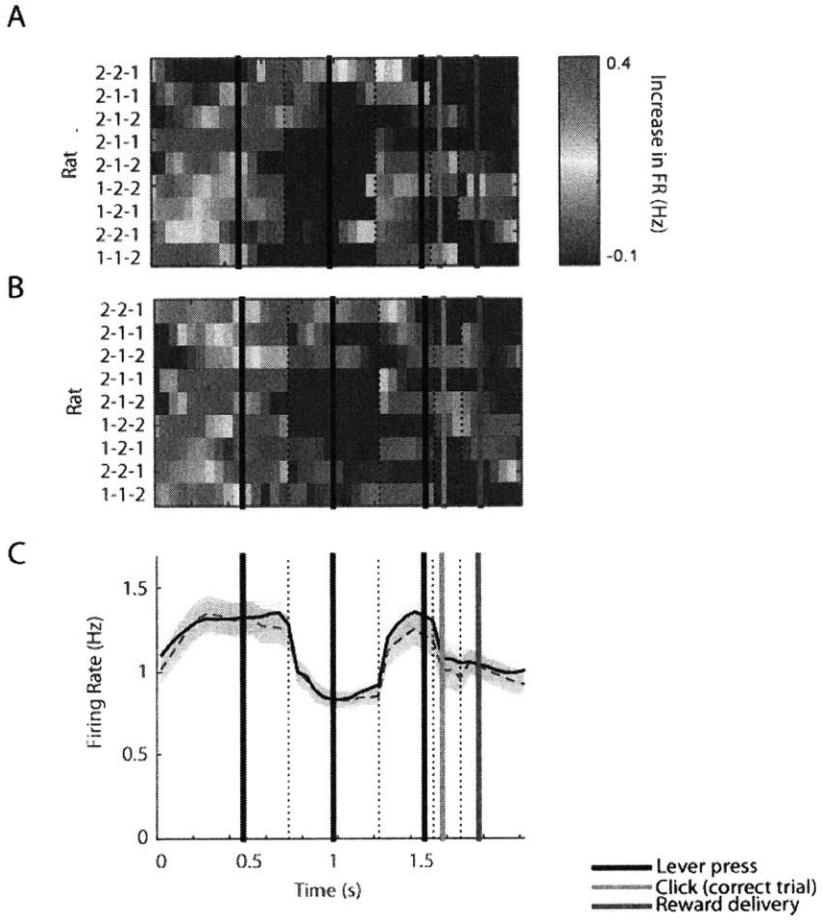

Figure 6. DLS SPNs represent correct trial boundaries in isolated correct trials during low-performance sessions similarly to high-performance periods.

(A) Average peri-event DLS SPN spiking during correct trials selected only from periods of majority correct performance in 9 implanted rats who learned different lever press sequences indicated on the left.

(B) SPN population activity in the same rats and sessions but in periods of below chance level (12.5%) performance.

(C) Overlay of average activity of the identical neurons during correct trials in periods of high-performance (solid line) and isolated correct trials in periods of below chance performance in the same sessions. (n = 1883 units).

All error bars indicate ±SEM

low or high-performance periods in sessions in which both of those conditions existed to directly compare the neuronal responses. I found a robust replication of the task-boundary SPN spiking even in isolated trials performed in low-performance periods. The population firing pattern of the SPNs was indistinguishable during these two different sets of trials (Fig. 6A, B, and C). However, preliminary analysis shows this task-boundary activity in isolated correct trials may be less clear in sessions in which no high-performance periods existed (Fig. S4), although this difference was not statistically significant. Thus, it appeared that at least when the animals had learned the sequence well enough that they could at their optimal performance achieve high accuracy, the striatum encoded the beginning and end of the learned sequence even when it was performed during exploratory or random performance periods. However, before the rats developed the capacity to perform at high levels of accuracy, and likely often performed the correct sequence by chance, and this learned action sequence beginning and end representation may not have been fully present.

Inhibiting cortical cell bodies has weak effects on DLS SPN and FSI firing

Addressing the relationship between motor cortex neuronal activity and striatal neuronal activity is important for understanding the nature of the cortico-basal ganglia-thalamic circuit and for understanding the differences these neuronal recordings revealed in the nature of the task representations in these areas. To address the basic influence of motor cortex firing on the spiking of DLS SPNs I conducted optogenetic silencing of motor cortex cell bodies and motor cortex axon terminals in striatum while recording from DLS single units. I injected CamKII-NpHR3-YFP virus into the forelimb motor cortex recording site and allowed 4 weeks for halorhodopsin expression to reach high levels in cell bodies and in the terminals within striatum. After this period, recording drives were implanted and optogenetic manipulations were conducted throughout the 5-6 week training period. Six of the nine rats with striatal recordings also had halorhodopsin manipulations and tetrode implantation in motor cortex (Fig. 7A). I confirmed that the post-experimental lesions marking the tips of the striatal recording tetrodes were within the halorhodopsin expressing YFP labeled termination region within striatum (Fig. 7B).

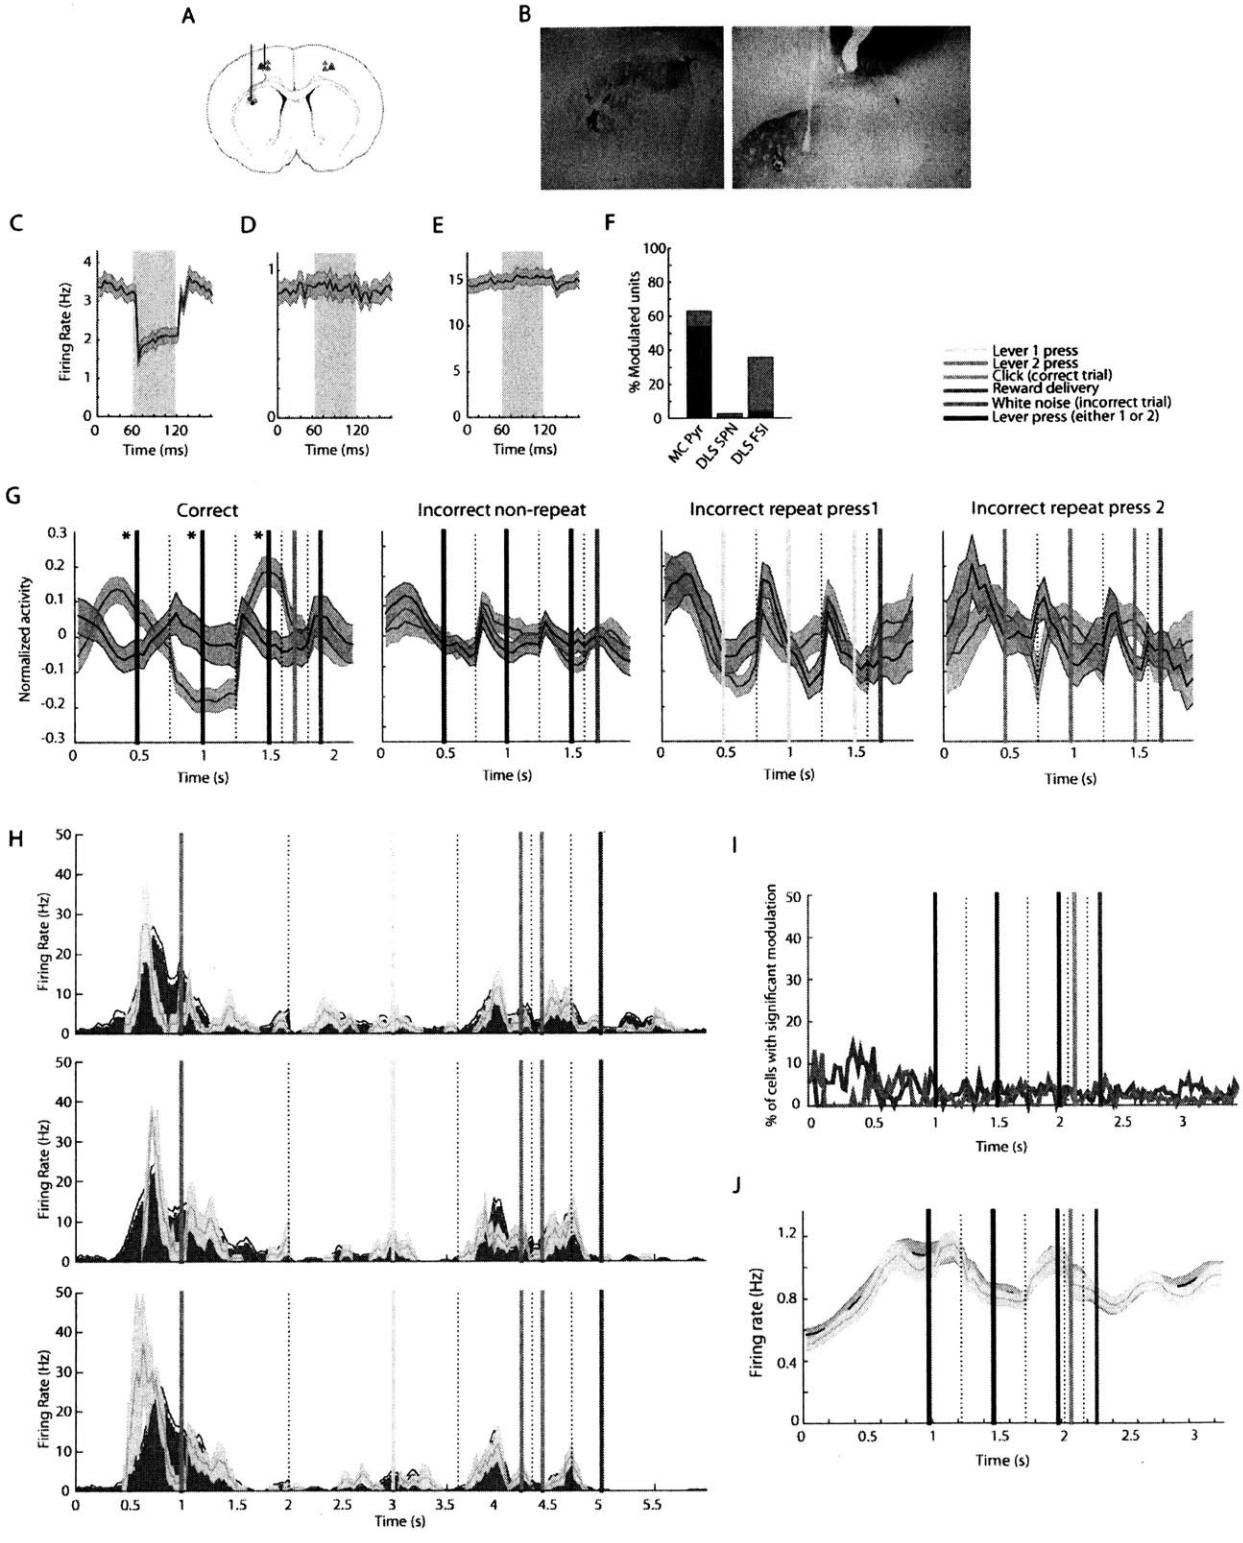
A
B
C
D
E
F
Firing Rate (Hz)
Time (ms)
% Modulated units
MC Pyr
DLS SPN
DLS FSI
Lever 1 press
Lever 2 press
Click (correct trial)
Reward delivery
White noise (incorrect trial)
Lever press (either 1 or 2)
G
Correct
Incorrect non-repeat
Incorrect repeat press1
Incorrect repeat press 2
Normalized activity
Time (s)
H
I
% of cells with significant modulation
J
Firing rate (Hz)

Figure 7.

Figure 7. Motor cortex is unlikely to be source of task-boundary activation in dorsolateral striatum.

(A) Halorhodopsin was injected into motor cortex bilaterally. After 4 weeks for virus expression, a recording drive with bundles of tetrodes surrounding an optical fiber targeted to left DLS and left motor cortex was implanted.
(B) Halorhodopsin expressing cortical terminals were seen in dorsolateral striatum (brown). Lesion marks (red) indicating the tips of striatal tetrodes were within this termination zone.
(C) Average firing rate of putative cortical pyramidal neurons when pulses of yellow light were delivered to cortical cell bodies in the freely moving rat (n = 383 units). Light on period indicated with yellow shading.
(D) Average firing rate of putative DLS SPNs recorded at the same time (n = 366 units).
(E) Average firing rate of putative DLS FSIs recorded at the same time (n = 106 units).
(F) Proportions of significantly inhibited (blue) and activated (red) units for each cell group.
(G) Session average activity of simultaneously recorded neurons in motor cortex and striatum in correct trials (n = 51 sessions), non-repeat incorrect trials (n = 57 sessions), lever 1 repeat press trials (n = 34 sessions), and lever 2 repeat press trials (n = 20 sessions).
(H) An example of a DLS SPN that was modulated in a task-time selective manner as a result of cortical terminal inhibition by an optical fiber placed in striatum. Black histograms indicate firing rate of unit during laser-off periods and yellow lines indicate firing rate during laser-on periods. Rows illustrate replicated effect of the laser across three days of recording.
(I) Proportions of DLS SPNs inhibited (blue) and activated (red) by laser targeting motor cortex terminals in DLS throughout task performance.
(J) The population SPN firing during correct trials during laser-off periods (gray) and laser-on periods (yellow) by an optical fiber placed in striatum (n = 1998 units).

All error bars indicate ±SEM

During post-training periods, I silenced cortical cell bodies in the freely moving rats using pulses of light targeted to an optical fiber placed in the motor cortex and surrounded by the tetrodes. I found that the average firing rate of all putative cortical pyramidal neurons was reduced from 3.3Hz to 1.5Hz at the onset of the optogenetic inhibition (Fig. 7C) and that 54% of the 383 putative cortical pyramidal neurons were significantly inhibited while 8% were significantly activated (Fig. 7F). However, the average firing rate of 366 striatal SPNs recorded at the same time from tetrodes located within the halorhodopsin expressing termination region in DLS was not changed by the cortical cell body silencing (Fig. 7D). Of the striatal SPNs, only 2%, below chance level for

multiple comparisons, were significantly modulated by the manipulation (Fig. 7F). Even when I considered only the 106 striatal SPNs firing at 1Hz or above, I observed no significant effect of the cortical cell body inhibition (Fig. S5A). Similarly, the mean firing rate of 109 putative narrow-waveform FSIs in DLS was unchanged by the manipulation. However, 32% of these FSIs were subtly but significantly activated by the manipulation (Fig. 7F, Fig. S5B). I repeated this manipulation using an optical fiber placed in the striatum surrounded by DLS tetrodes within the halorhodopsin expressing terminal region in DLS. As with the cell body inhibition, I found that the effect on DLS SPNs and FSIs was weak, but a similar proportion of FSIs was subtly activated (Fig. S5C-F). Based on the strong effects of the optogenetic silencing on the cortical cell bodies, but insignificant effects of this manipulation on the SPNs located within the halorhodopsin expressing terminals in DLS, I suggest that in my rats each of whom underwent a substantial training to press single levers prior to recording drive implantation, motor cortex does not appear to be a strong driver of SPN spiking in DLS.

Motor cortex neuronal population lacks task-boundary activity seen in DLS

I compared the simultaneous recordings of motor cortex and striatum by choosing sessions in which at least 5 putative DLS SPNs (Fig. 8A) and 5 putative motor cortex pyramidal neurons (Fig. S10) were recorded simultaneously and comparing the normalized session average trajectories in motor cortex and DLS. I found that the population firing in motor cortex typically increased approximately 250ms prior to each lever press and was weaker at the lever press completion timepoint (Fig. 7G). This was true in correct trials as well as incorrect non-repeat sequences (Fig. 7G, first and second columns respectively), but was even more evident during repetitive pressing of lever 1 or lever 2 (Fig. 7G, third and fourth columns). Thus, I did not see evidence that motor cortex neurons responded in the same task-boundary accentuating manner during correct sequence execution as striatal neurons. This finding is potentially consistent with the different types of neuronal task representations in motor cortex and striatum presented earlier (Fig. 2D).

Motor cortex terminal silencing does not affect striatal task-boundary response

In addition to using pulses of light to inhibit cortical cell bodies and terminals, I used an optical fiber that was placed in striatum within the halorhodopsin-expressing terminal region and surrounded by recording tetrodes to inhibit the corticostriatal terminals mid-task (see methods). In most DLS SPNs in which sufficient numbers of laser-on and laser-off trials and sufficiently high firing rates existed to statistically compare task-related firing during laser-on and laser-off times, I found that the task-related firing in laser-off times was robustly replicated during laser-on times (Fig. S6). I found that in some rare instances the inhibition of the terminals had a subtle but significant effect in individual unit task-related firing that was replicated across several recording sessions (Fig. 7H, Fig. S7). When I analyzed the subpopulation of SPNs that were significantly modulated by the optogenetic inhibition of the terminals, I found that there were groups of units recorded on particular tetrodes that appeared to be modulated at the same time in the task and in a similar manner across days (Fig. S8). However, the prevalence of such effects was low and there was a relatively uniform distribution of time periods in which SPNs were significantly inhibited or activated during laser-on times (Fig. 7I). Interestingly, the task time in which SPNs were most likely to be inhibited by the optogenetic silencing of the cortical terminals was prior to initiation of the trial and not during the trial. In total, the population SPN firing rates during laser-off times and laser-on times were very similar (Fig. 7J). The effects of the laser on DLS FSIs were also relatively uniformly distributed and the population FSI firing in-task was not significantly affected by the terminal inhibition (Fig. S9).

Together, the different kinds of task representations in motor cortex and DLS, the lack of the task-boundary activation in motor cortex, and the weak effects of cortical cell body and terminal inhibition on SPN spiking suggest that while motor cortex is likely providing important relevant information about movement to striatal neurons it is not likely to be the primary driver of the task-boundary signal in DLS SPNs.

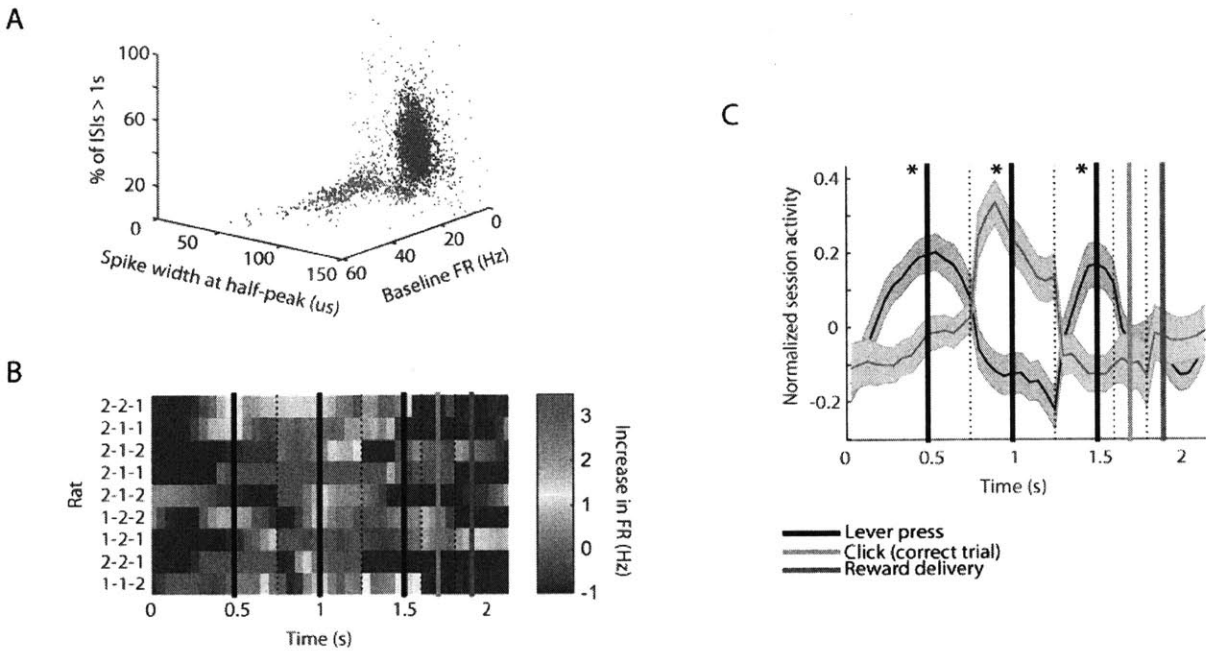

Figure 8. Narrow spike waveform fast spiking interneurons in DLS fire in an opposing pattern to SPNs during correct sequence performance.

(A) Single units recorded in DLS appear to fall into two distinct clusters along features of spike width, firing rate, and % ISIs greater than 1s. Units classified as putative narrow-waveform FSIs are displayed in red and putative SPNs as used in previous analyses are displayed in blue.
(B) Normalized FSI activity for each of the nine animals in the correct sequence.
(C) Normalized session average activity of simultaneously recorded SPNs (gray) and FSIs (red) in n= 69 sessions. Error bars indicate ±SEM

Striatal narrow-waveform fast spiking interneurons may be shaping the task-boundary activity in the spiny projection neurons

Among the single units recorded in the striatum, I found a clear bimodal distribution in spike waveform widths with a smaller cluster of faster firing narrow waveform units and a larger cluster of lower firing wide waveform units (Fig. S11A-B). This division in spike widths was not driven by unit distance from the tetrode as reflected in spike amplitude (Fig. S11C) and was likely reflective of a distinction in physiological properties of the DLS neurons. I used spike width and firing rate

criteria (see methods) to classify putative fast spiking interneurons (FSIs) and putative spiny projection neurons (SPNs) (Fig. 8A). I found that typically, the narrow-waveform FSI population firing pattern during the learned sequence performance across the nine implanted rats was concentrated in the mid-task period (Fig. 8B) and this pattern sharply contrasted with the task-boundary activation in the SPNs of the same animals. To compare directly the task related firing patterns of the FSIs and SPNs I selected sessions in which at least 2 FSIs and 5 SPNs were recorded simultaneously, and I compared the normalized session averaged activity of these FSIs and SPNs. I found that the FSIs had a clearly antagonistic firing patterns to those of the SPNs in which the peak of the FSI firing during the middle lever press in the correct sequence corresponded to a dip in SPN firing (Fig. 8C). In contrast, units firing at baseline rates greater than 3.5Hz but with wider spike widths similar to that of SPNs were most activated during the first and last lever press similar to the SPN population (Fig. S11D). Thus, I found that spike width and not high-firing rate was the property in the single units that distinguished DLS units firing at the task-boundaries and those firing mid-task.

I conducted a closer analysis of the relationship between DLS SPN and narrow-waveform FSI firing by identifying 265 pairs of SPNs and FSIs that were recorded on the same tetrode on the same recording session and assessing the correlation between the trial average firing rates of the pairs of units during correct sequence performance. I found that there were pairs of units with positive and negative correlations and 54% of the pairs were significantly correlated ($p < 0.05$); however, the mean r value of the correlations was negative (Fig. S12A). When I plotted the task-related firing pattern of the SPNs that were negatively correlated with an FSI on the same tetrode and SPNs that were positively correlated with an FSI on the same tetrode, I found that the negatively correlated SPNs were more responsive during trial time and were strongly activated during the first and last lever press in the sequence. The SPNs which were positively correlated with an FSI recorded on the same tetrode responded more weakly during task time and did not fire preferentially around the time of the first or last lever press (Fig. S12B). Based on these results, and the known inhibitory influence of FSIs on SPN firing, I suggest that the narrow waveform FSIs are likely to contribute to increased SPN activation at the task-boundaries and decreased SPN activity mid-task.

Two groups of striatal SPNs, both active at task-boundaries, may exert opposing influence on motor cortex.

Although I did not directly manipulate striatal neuronal activity, I could address the possible influence of striatal firing on motor cortex neurons by analyzing the relationship between pairs of simultaneously recorded DLS SPNs – motor cortex units. To do this, I identified spike bursts in each of the recorded SPNs (see methods) during correct sequence performance. I then aligned the spike activity of any simultaneously recorded putative pyramidal cells in the motor cortex to those SPN spike bursts. I found that there were two clear groups of significantly modulated motor cortex units among the pairs that I recorded. The first group of units fired at low rates prior to the SPN spike burst but decreased their firing rates after the burst. In the second group of striatal SPN – motor cortex cell pairs, high motor cortex unit firing was reduced after the SPN burst (Fig. 9A). Principal components analysis followed by k-means clustering identified two clusters of SPN-motor cortex cell pairs (Fig. 9B). The two clusters were almost identical in size (49% of inhibitory interaction pairs and 51% of excitatory interaction pairs). The latency of the peak of the firing in the activated motor cortex units and the peak of the firing rate reduction in the inhibited motor cortex units was approximately 200ms. However, I found that despite their opposing relationship or possible influence on motor cortex activity, striatal SPNs from the two groups both had the same task-boundary activity as the general population of SPNs. The division in these opposing influences of SPNs on motor cortex activity, and the long latency of this effect is potentially consistent recent findings indicating opposing influence of stimulating D1 expressing MSNs and D2 expressing MSNs on motor cortex activity and with recent findings that these cell types tend to be concurrently active at action initiation and motor program selection (Cui et al., 2013).

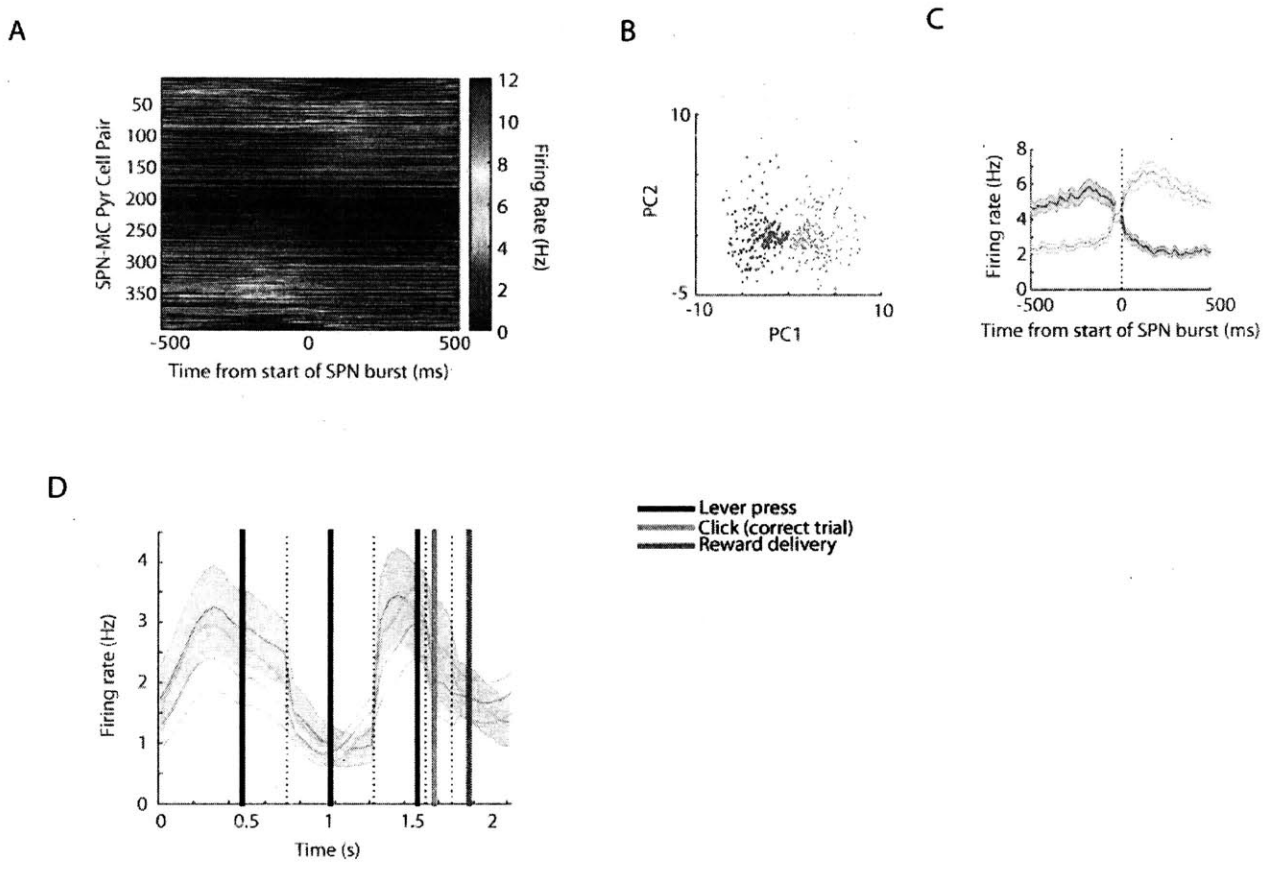

Figure 9. Equal sized groups of DLS SPNs with opposing relationships with motor cortex firing fire concurrently at task-boundaries.

(A) Responses of all significantly modulated motor cortex pyramidal neurons by bursts of spiking in DLS SPNs (left).
(B) Results of PCA analysis and k-means clustering on these responses.
(C) Average responses in relation to the SPN spike burst of the two clustered cell pair groups.
(D) Task-boundary activity in both sets of striatal SPNs which produced inhibitory and excitatory responses in motor cortex.

All error bars indicate ±SEM

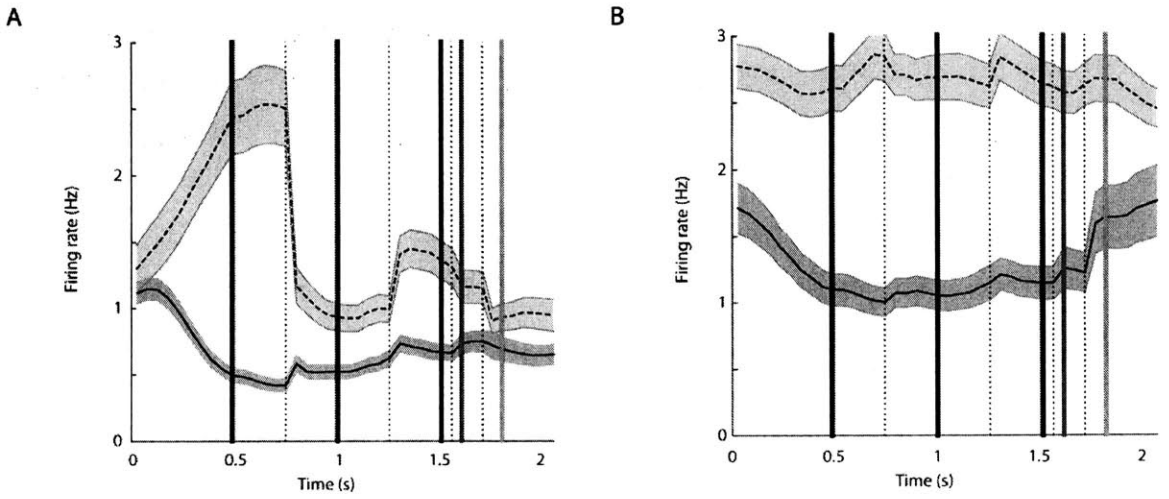

Figure 10. Dorsomedial striatum and prelimbic cortex are weakly activated during learned lever sequence performance as compared to dorsolateral striatum and motor cortex.

(A) Average peri-event DMS SPN firing rates in three rats during correct trials. In dotted lines, the DLS firing rate is shown for the same rats (n = 431 units, n = 584 units respectively).

(B) Average peri-event prelimbic cortex putative pyramidal cell firing rates in one rat during correct trials (n = 187 units). In dotted lines, the DLS firing rate is shown for six rats (n = 817 units).

All error bars indicate ±SEM

Performance of learned lever press sequence elicits strong activation in lateral but not medial corticostriatal circuit

Based on previous literature detailing functional distinctions between the medial and lateral dorsal striatal circuits, I explored the degree of engagement this fixed action pattern task elicited in the medial and lateral corticostriatal circuits. In three rats, I recorded from dorsomedial striatum (DMS) in addition to dorsolateral striatum (DLS) and compared the activity in simultaneously recorded neurons in each site. I found that the baseline inter-trial interval firing rates in DMS and DLS SPNs were 0.47Hz and 0.4Hz respectively and were not significantly different ($p = 0.53$, Mann-Whitney test). However, the mean firing rates of SPNs during the learned sequence performance in DMS remained at baseline levels while the firing rates of the simultaneously recorded DLS neurons peaked at the initiation and termination of the learned

sequence (Fig. 10A) and were significantly higher in-task than the DMS SPN firing rates ($p < 0.00001$, Mann-Whitney test). In one rat, I also recorded 187 putative pyramidal neurons in a region of prelimbic cortex which projects to DMS (Fig. S13), making up part of the medial associative corticostriatal loop (McGeorge and Faull, 1989). The mean baseline firing rates of the putative pyramidal neurons in motor cortex and prelimbic cortex were 2Hz and 1.9Hz respectively and were not significantly different ($p = 0.68$, Mann-Whitney test). However, similarly to the distinction between DMS and DLS activity levels, prelimbic cortex population firing was suppressed below baseline levels while motor cortex neurons were active throughout the trial (Fig. 10B) and had significantly higher in-trial firing rates than prelimbic cortex neurons ($p < 0.00001$, Mann-Whitney test). This discrepancy in the degree of engagement of the medial and lateral corticostriatal circuit likely reflects the nature of this fixed lever press sequence task which was designed to be repetitive, performed in a similar manner for many trials, and to not require flexibility or decision making. The preferential engagement of the lateral corticostriatal circuit in this task lends support to the idea of differential functions in the roles of the associative corticostriatal circuit in goal-directed behaviors and decision making, and the sensorimotor corticostriatal circuit in fixed behavioral patterns.

Discussion

In this study, I designed a behavioral paradigm to systematically test the nature of the neuronal responses related to learned action sequences in the corticostriatal circuit. In order to identify neuronal responses related to individual actions, I required the animals to learn a sequence consisting of three consecutive lever presses that took place in the beginning, middle, and end of the required sequence as well as any incorrect sequences they performed. In order to identify any neuronal signals for action sequences that generalize across different types of learned movement sequences I required different rats to learn different sequences of lever presses. I found that the motor cortex neurons I recorded were more likely than striatal neurons to fire in response to singular motor events such as pressing one of the levers or shifting from one lever

to another. Striatal neurons were much less likely to represent individual motor actions in a manner that was insensitive to the sequence of lever presses within which they were imbedded. Although striatal neurons had a great heterogeneity of responses, one of the most evident kinds of task-related response, which was so prevalent that it was dominant in the average population spiking, was a high level of spiking during the initiation and termination of the learned and reinforced lever press sequence. The identical neurons recorded on the same day during times when the rats pressed the levers in alternative incorrect sequences did not respond in the same manner.

My study was motivated by the detection of beginning and end signals in the dorsolateral striatum in a T-maze task (Jog et al., 1999, Barnes et al., 2005, Thorn et al., 2010) and in mice in a fixed-ratio lever press task (Jin and Costa, 2010) leading to the emergence of the hypothesis that these neuronal responses could be a representation of chunked action sequences. A primary motivation in designing this behavioral paradigm was to test the hypothesis that the habit-related dorsolateral striatum represents the task-boundaries of chunked action sequences regardless of the movement content of those action sequences and is not directly reflective of simple movement parameters or task events – neither of these questions having been directly tested in previous studies. Strikingly, I did find that the DLS SPN population spiking peaked at the action sequence boundaries across rats that learned different sequences of lever presses and developed very different movement patterns in order to execute the required series of lever presses. The proportions of the beginning and end responsive units (30% and 24% of 2501 units, respectively) and a distinct set of reward-period responsive units (13% of 2501 units) were similar to those found previously in the T-maze task (Barnes et al., 2005, Smith and Graybiel, 2016), a task involving locomotor movements qualitatively different from those required during the lever pressing task. The finding that the task-boundary activity does not occur during repetitive pressing of the same lever differs from the finding of such beginning and end activity in mice trained on an FR8 lever press task (Jin and Costa, 2010). However, there are two important differences in these sets of studies, (1) that the mice in the FR8 task were rewarded every 8 lever presses and thus were trained to perform this behavior whereas rats in this study were not rewarded for repetitive pressing of the same lever, and (2) while beginning and end neurons were

identified in mice in the FR8 task, it is not clear whether the activity of the full population of SPNs was modulated in the task-boundary manner. However, given the apparent ubiquitous nature of the task-boundary activity, it is likely that it can develop even in a repetitive lever pressing task which is practiced and rewarded. Thus, based on my experiments and the findings of task-boundary signals in other tasks, I suggest that there is a high-level representation of behavioral units consisting of chunked action sequences in the dorsolateral striatum that is in a large part dissociable from the movement content of those action sequences. This finding is important for furthering the theory that dorsolateral striatum may fulfill its known function of promoting habitual behavior by merging the elements making up useful action sequences into a single rigid unit (Graybiel, 1998, Graybiel and Grafton, 2015).

I further propose that the presence of this robust task-boundary activation together with the low numbers of DLS SPNs whose spiking pattern could be easily accounted for by single motor events supports a shift away from considering striatal neuronal activity within the realm of locomotion or movement alone. Instead, the function of DLS in behavior should be considered within context specific representations of motor programs including proposed center-surround action selection models of direct and indirect pathways (Mink and Thach, 1993, Cui et al., 2013, Friend and Kravitz, 2014), and the proposed modular organization of striatum that would enable selection of appropriate action programs based on context and sensory input (Amemori et al., 2011). One initial piece of evidence supporting the possible concurrent influence of the direct and indirect pathways during action selection is the reported simultaneous activation of direct and indirect pathway SPNs during action initiation (Cui et al., 2013) and my observation that DLS SPNs which have opposing relationships with motor cortex neurons are similarly activated at task-boundaries.

To assess possible sources of the task-boundary activity in dorsolateral striatum, I addressed the potential roles of motor cortex and of striatal interneurons. The population of motor cortex putative pyramidal neurons that I recorded was not preferentially active at the beginning or end of the correct sequences but throughout performance of the entire sequence, and as discussed previously motor cortex neurons were more likely to represent individual actions similarly across

correct and incorrect sequences. I conducted several optogenetic tests to further probe the role of motor cortex in driving striatal neuronal firing. In the freely moving rat, optogenetic silencing of the motor cortex cell bodies produced a reliable inhibition in motor cortex neurons, but failed to inhibit firing in striatal neurons recorded from tetrodes within the halorhodopsin expressing termination zone in DLS. It did however have a very subtle excitatory effect on 32% of putative FSIs – an effect which may have been a result of circuit level effects or possible high sensitivity from low-levels of light reaching striatum. I found similarly small effects when using light pulses targeted directly at the cortical terminals within striatum. I also conducted optogenetic inhibition of the corticostriatal terminals in-task, a strategy which has been previously demonstrated to effectively inhibit excitatory inputs and produce behavioral changes in other pathways including a different corticostriatal pathway (Tye et al., 2011, Friedman et al., 2015), and found that the firing of a small proportion of the striatal SPNs were reliably modulated in a task-dependent manner. However, the timing of these effects was widely distributed across task time, and the population task-boundary activity was robustly present during laser-on times and was indistinguishable from laser-off times.

Together, the different kinds of neuronal task-representations in striatum and motor cortex, the lack of task-boundary activation in the motor cortex in the units I sampled, and the weak effects of cortical cell body and also corticostriatal terminal silencing on striatal SPNs and FSIs suggest that while inputs from motor cortex, together with other cortical inputs, are likely providing important converging information about ongoing movement to dorsolateral striatum, they are not likely to be the source of the task-boundary activity.

This finding is aligned with several recent results on the functional role of motor cortex and its influence on striatum. It was recently demonstrated that motor cortex is required for the learning but not execution of learned motor sequences in a task not requiring highly dexterous movement of the digits (Kawai et al., 2015). Moreover, secondary but not primary motor cortex input to striatum was found to be required for serial order performance (Rothwell et al., 2015). Thirdly, the activity of motor cortex terminals in dorsolateral striatum has been found to be greatly reduced after training on the rotarod (unpublished data, David Lovinger Laboratory NIH). In my recordings, the task-boundary activity was present very early in the recording, indicating that DLS

may have already undergone plasticity resulting in the reduced influence of motor cortex. It is possible that in these experiments, the extensive shaping procedure with rats learning to press the levers prior to the recording drive implantation was sufficient to reach the training induced decreases in the motor cortex inputs or that postsynaptic effects of motor cortex terminals upon DLS SPNs as measured by me diminish at a faster rate than the presynaptic calcium transients measured on the rotarod in mice, especially considering that most of the presented optogenetic manipulations did not begin until the second week of sequence training.

Based on this set of findings, it is possible that subcortical structures are the primary controllers or initiators of learned action sequences, and they may do so by exerting influence on cortical or other subcortical motor structures. In these recordings, it appeared that there were two opposite types of DLS SPN and motor cortex neuron interaction pairs. In one case, when the SPN released a burst of spikes, the motor cortex neuron which had been previously active decreased in activity; and in another case a low-firing motor cortex neuron increased in activity. While the current set of experiments does not provide a direct test for the polysynaptic influence of striatum on motor cortex firing, the two kinds of equally distributed interaction pairs in this dataset are similar in distribution and latency (200ms) to the opposing influences of direct and indirect pathway SPNs on motor cortex directly demonstrated in mice (Oldenburg and Sabatini, 2015) and is predicted by the connectivity of the two SPN types. However, even in the case of striatal control of motor cortex, the authors found that the influence of striatum on the motor cortex neurons was diminished when the animal performed a motor action. This result further suggests a decoupling between motor cortex and striatum and the possibility of striatal control of subcortical motor structures for tasks not requiring highly dexterous movement of digits which has been suggested as a primary advantage of motor cortex over subcortical motor structures (Bortoff and Strick, 1993, Lemon, 2008).

Given the electrophysiological and optogenetic evidence in these experiments, and from lesioning and imaging experiments from other studies indicating a weak influence of motor cortex on striatum; I must consider other circuit elements that could be shaping the strong task-boundary firing in DLS, including striatal microcircuitry. I found that a key physiological feature in striatal single units was spike width which was bimodally distributed with a larger cluster of low-

firing units which I classified as SPNs and a smaller cluster of narrow-waveform high firing units which I classified as putative FSIs. This group of narrow-waveform FSIs had a distinctly contrasting task-related firing pattern than the SPNs and fired mid-task at the time when SPN firing was depressed. Thus, due to the known inhibitory influence of FSIs onto SPNs, I suggest that FSIs may be playing a major role in shaping the task-boundary activation in DLS. How such a firing pattern arises in the FSIs and potential other sources creating the task-boundary activity in the SPNs are major questions for further research.

With this set of experiments, I attempted to establish whether the task-boundary activity in striatum reflects movement parameters or whether it reflects a high-level signal indicating the initiation and termination of chunked action sequences. I found that the task-boundary activity was present in dorsolateral striatum of rats who learned different sequences of lever presses involving different movements. I also demonstrated that unlike motor cortex neurons, striatal neurons activity was less likely to be accounted for by individual actions. I suggest that the finding of such a ubiquitous signal for task-boundaries that was selective for the learned and reinforced action sequence and was not present in other random series of lever presses demonstrates the existence of a fundamental signal in striatum that reflects the initiation of chunked action sequences. This furthers the notion that dorsolateral striatum can drive behavior toward the habitual stereotyped mode by encoding series of actions as a single behavioral unit, parsing behaviors into chunks of functional units.

The existence of such a general purpose signal in dorsolateral striatum for the task-boundaries of chunked action sequences raises numerous questions about the function of this activity and the possible role it may have in driving behaviors toward a rigid habitual mode. One simple interpretation of the reduced DLS activity in the middle of the task performance could be that striatum becomes disengaged during this period. However, it would be difficult to rule out the role of DLS in these behaviors due to two factors, (1) the inordinately high-levels of engagement and high firing rates of large proportions of DLS neurons, especially as compared to dorsomedial striatum, concentrated at the start and end of the learned sequences indicates that DLS is not becoming less responsive but rather that the neuronal activity becomes highly stereotyped and concentrated at specific task points, and (2) the known role of DLS in habit formation informs our

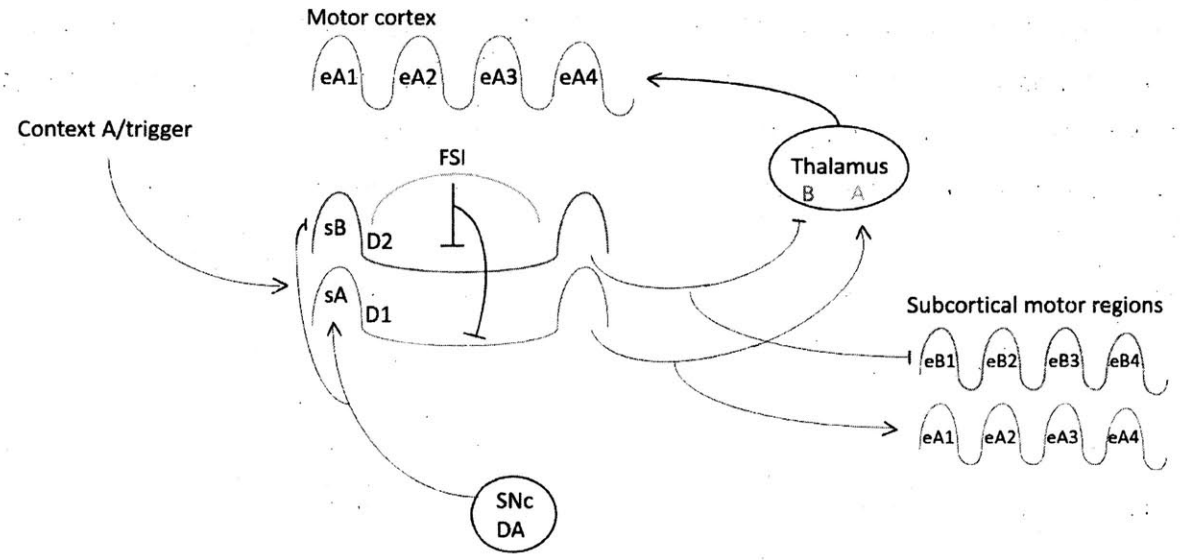

Figure 11. Basic testable model for the learning and execution of motor programs. Sensory stimuli information about context and triggers for a particular behavior (in this case, behavioral program A) is delivered from cortical and thalamic areas not pictured to the direct and indirect pathway SPNs in dorsolateral striatum. Within dorsolateral striatum this input activates subsets of D1 SPNs that encode behavioral sequence A (sA) and also activates D2 SPNs that encode alternative behavioral sequences (in this case, sB). These D1 SPNs activate subsets of neurons that carry out each individual action (eA1-eA4) in behavioral program A either in subcortical motor regions or in the motor cortex through the thalamus. Indirect pathway SPNs inhibit neurons that are involved producing alternative behaviors (in this case, behavioral program B). During the execution of the behavioral program FSIs in DLS inhibit further activation of DLS SPNs to allow full program to become executed before another behavioral program can be triggered. In the case of a positive outcome, dopamine input to DLS from SNc reinforces synaptic links between current sensory inputs and active D1 SPNs further reinforcing the behavior.

interpetation of the task-boundary activity as having a potentially critical role in action sequence chunking. One alternative possibility is that the beginning and end activation in DLS serves to initiate learned action programs in cortical or subcortical motor regions holding the memory of the learned sequence to carry out the behavior. This possibility would suggest a role of DLS in the direct on-line control of learned action programs as has been proposed in action-selection models of the striatum (Da Cunha et al., 2012, Friend and Kravitz, 2014). In another alternative, striatum may not be involved in direct on-line control of the behavior but provide a teaching

signal that allows the chunking of action repertoires. These alternatives may be tested using targeted inhibition of dorsolateral striatum during the acquisition of behavioral repertoires or during the performance of well-learned behaviors.

This leads me to propose a basic testable model basal ganglia involvement in the learning and execution of learned behavioral programs (Fig. 11) in which stimulus information from cortical and thalamic sensory areas activate appropriate sets of direct pathway DLS SPNs to initiate the appropriate behavioral program. These SPNs can start a chain of activations in subcortical motor regions or in cortical motor regions through the thalamus which hold the memory of how to execute each step of the behavioral program and have neural activity reflective of those individual steps. Simultaneously, sets of indirect pathway SPNs can inhibit competing behavioral programs. One important testable prediction from this model is the possibility that DLS interneurons serve to inhibit activation of the DLS SPNs after the behavioral program is initiated to prevent SPNs from being further activated until the current behavioral program is finished executing. The pairings of sensory cues and behavioral programs that produce positive outcomes can result in dopamine release in DLS arising from the SNc which would reinforce the synaptic links between those sensory inputs and the active D1 SPNs further promoting the behavior.

Future directions

A clear new direction for research involving neuronal monitoring in the striatum has been reliable identification of relevant cell types and their manipulation, which will ultimately be a critical addition to the recording of neuronal activity for linking anatomically and histologically defined components of the basal ganglia circuit to function. This was recently made possible by transgenic mouse and rat models with expression of opsins and fluorescent proteins in neuronal subtypes. There are at least four such divisions of neuronal types among striatal projection neurons:

D1 expressing striosome SPNs	D1 expressing matrix SPNs
D2 expressing striosome SPNs	D2 expressing matrix SPNs

Various transgenic mouse lines are now available which express opsins selectively in D1 or D2 SPNs. However, due to lack of full penetrance of the genes used in these mouse lines, the absence of the opsin or fluorescent marker does not indicate a true negative. To identify D1 or D2 expressing SPNs in the same animal, would require distinct markers for these two neuron types and there are no mouse lines available yet that would make this possible. Similarly, there are lines of transgenic mice that would allow the identification of striosome or matrix SPNs with varying levels of success. However, work is still in progress that would allow the identification of both neuronal types in the same animal. Identifying all four of the main classes of SPNs would require the development of a mouse line expressing four distinct markers. Such a transgenic mouse would be the most powerful tool for the continuation of this research. While these kinds of transgenic mice are under development, it is possible to use the existing mouse lines to separately identify D1 or D2 SPNs or striosome or matrix SPNs while recording from single units in the striatum as mice perform habitual action sequences. This would allow for further testing of the theory that D1 and D2 SPNs are both active during the initiation of the action sequence and identify differences between the responses of the direct and indirect pathway SPNs using a between-animal experimental design. Using mice which express opsins in striosome or matrix neurons would be especially informative about the role of a minority of the SPNs in the DLS that are part of striosomes. One possibility that could be tested using this design is that the minority of SPNs which rarely spike during task execution but respond during the feedback period may be part of striosomes. A similar strategy could be used to confirm the identity of the FSIs which were putatively identified in this dataset using electrophysiological properties.

The same transgenic mouse lines could be used to establish the causal role of these neuronal types in the behavior. Inhibition of direct SPNs at the time of the intended initiation of a learned action sequence could help confirm the action selection hypothesis of the basal ganglia. Inhibition of the indirect pathway SPNs could be used to test the circle-surround theory of action selection in which the indirect pathway inhibits the performance of competing motor programs. Alternatively, if these manipulations do not clearly affect the real time behavioral choices of the animal, we could hypothesize that the role of the dorsolateral striatum is not critical for on-line control but for the learning of these action sequences, and test this by chronically inhibiting direct

or indirect pathway activity through the course of learning which may disrupt the automatization of learned action sequences.

A different set of questions arising from these findings is how the many different habitual action sequences that are usually within an animal's repertoire are represented together in DLS. I propose that a logical continuation for this line of inquiry would be to design a behavioral paradigm in which animals would learn to express one of two or more possible learned action sequences depending on context mimicking the manner in which cues and context normally trigger appropriate behavioral patterns. There may be multiple ways in which DLS neurons may respond during such different learned movement sequences, and one possibility, if we subscribe to the action-selection model of the striatum, is that start-responsive neurons for the alternate behaviors may become differentiated over the course of learning, but end responsive neurons and reward or lever press responsive neurons may continue to be shared among the sequences (Fig. S14). However, given the large proportion of start neurons active in DLS for a single sequence, it may be unlikely for each habitual behavior to have a dedicated set of start neurons. An alternative could be that there is a substantial overlap between the start neurons of different habits, but that this overlap depends on the degree of similarity of the behaviors and the stimuli that normally trigger their expression.

A second natural continuation of this line of study is to ask how the DLS task-boundary activity evolves if two well-learned behaviors are strung together to create a new larger chunked action sequence. This question can be tested directly in experimental animals by training them on two action sequences separately as suggested above, and after achieving high performance on both sequences require the animal to do them in quick succession. One possible outcome of this merger is that the "end" responses of the first sequence and the "start" responses of the second sequence could be reduced leaving primarily the "start" activity from the first sequence and the "end" activity from the second sequence (Fig. S15A). Another possibility is that the original task-boundary activity relating to the two separate sequences will be preserved but another set of task-boundary responses will develop around the merged sequences. Both of these outcomes could occur as a result of changes in the activity level of task-related neurons or changes in the numbers of task-related neurons (Fig. S15B).

Finally, a critical step toward the understanding the mechanisms of circuit function in basal ganglia and their relationship to habitual behaviors is addressing the transformation of neural representations from the striatum to the basal ganglia output nuclei and thalamus. Due to the complicated polysynaptic nature of this circuitry, this is a highly challenging endeavor. A promising possibility is the combined use of optogenetic silencing and excitation of specific sets of connections along with neuronal activity monitoring. Optogenetic excitation of terminals while recording or imaging activity in the upstream and downstream cell bodies in behaving animals can establish the functional connectivity of specific neuronal subsets in anatomically connected regions. Optogenetic silencing of the same terminals while recording or imaging task-related activation in upstream or downstream cell bodies can provide evidence of causality in the influence of the upstream neuronal activity on downstream neuronal activity and of the influence of this set of connections on behavior. Using this strategy, particularly in combination with cell type specific labeling of direct/indirect pathway neurons and striosome/matrix neurons, it may be possible to dissect the cortico-basal ganglia- thalamic circuitry step by step using a standardized behavioral paradigm.

SUPPLEMENTARY FIGURES

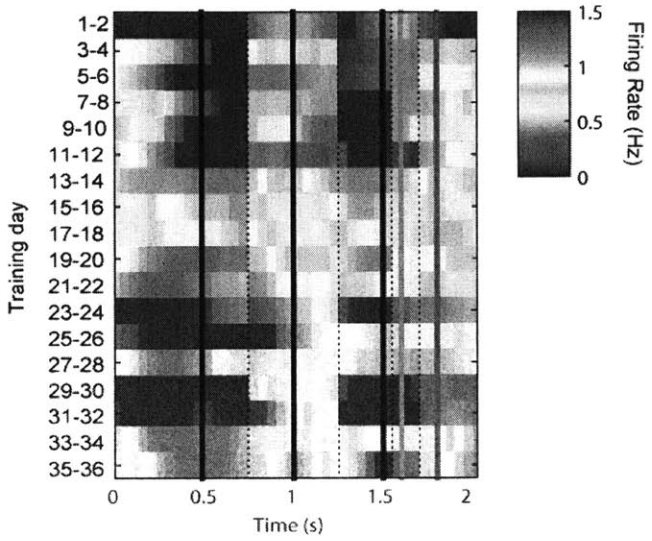

Supplementary Figure 1. Task-boundary activity develops early in training

Firing rate of DLS SPNs during correct sequence performance across pairs of training days. Black lines indicate lever presses.

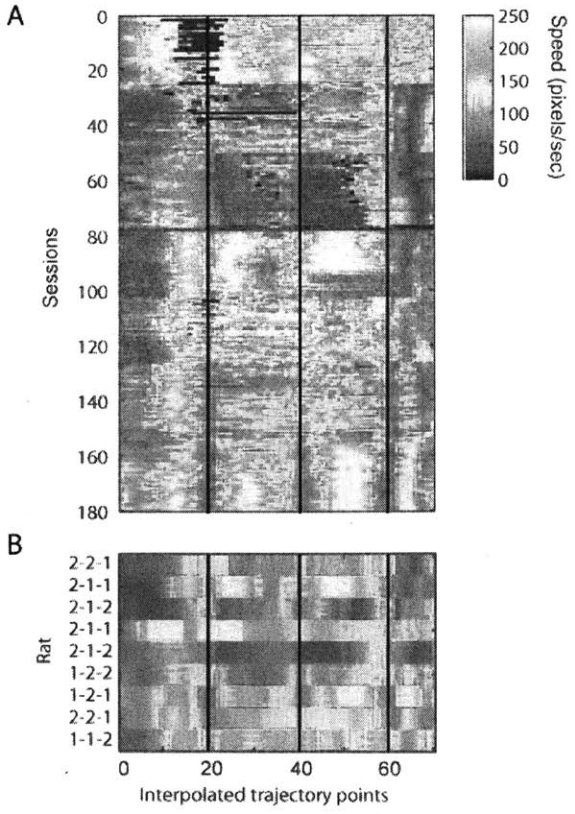

Supplementary Figure 2. Lever presses are associated with decrease in speed

(A) Average speed of head movement in 180 sessions from 8 rats. Each trial's trajectory was interpolated to 70 points before being averaged and plotted. Lever press events are indicated by black lines.

(B) Average speed of head movement in all nine rats displayed in the same order as in neural results figures.

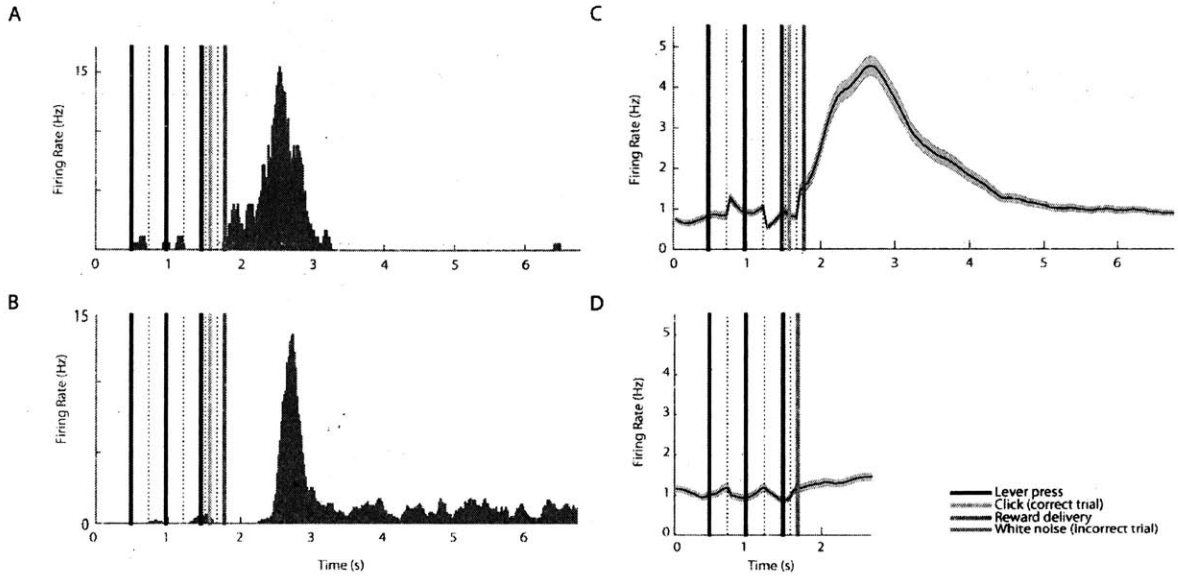

Supplementary figure 3.

(A) An example DLS SPN which fires preferentially after the reward delivery
(B) A second example reward-responsive unit
(C) Mean firing rate of all reward period selective DLS SPNs
(D) Mean firing rate of the same SPNs during and after incorrect trials

All error bars indicate ±SEM

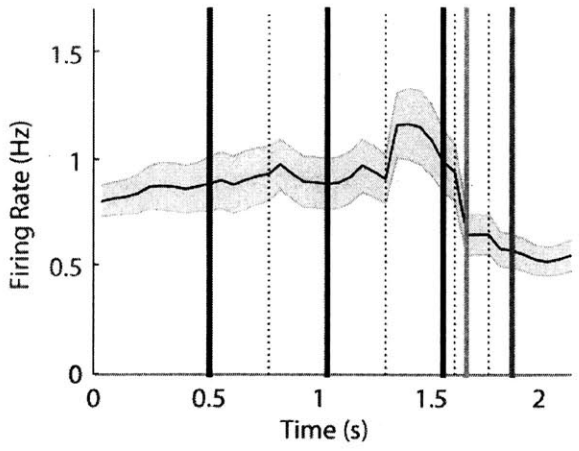

Supplementary Figure 4. DLS SPN activity in correct trials, in sessions with no high-performance periods.

Mean firing rate of n = 404 DLS SPNs in sessions with at least 5 correct trials performed, but no high performance periods occurred.

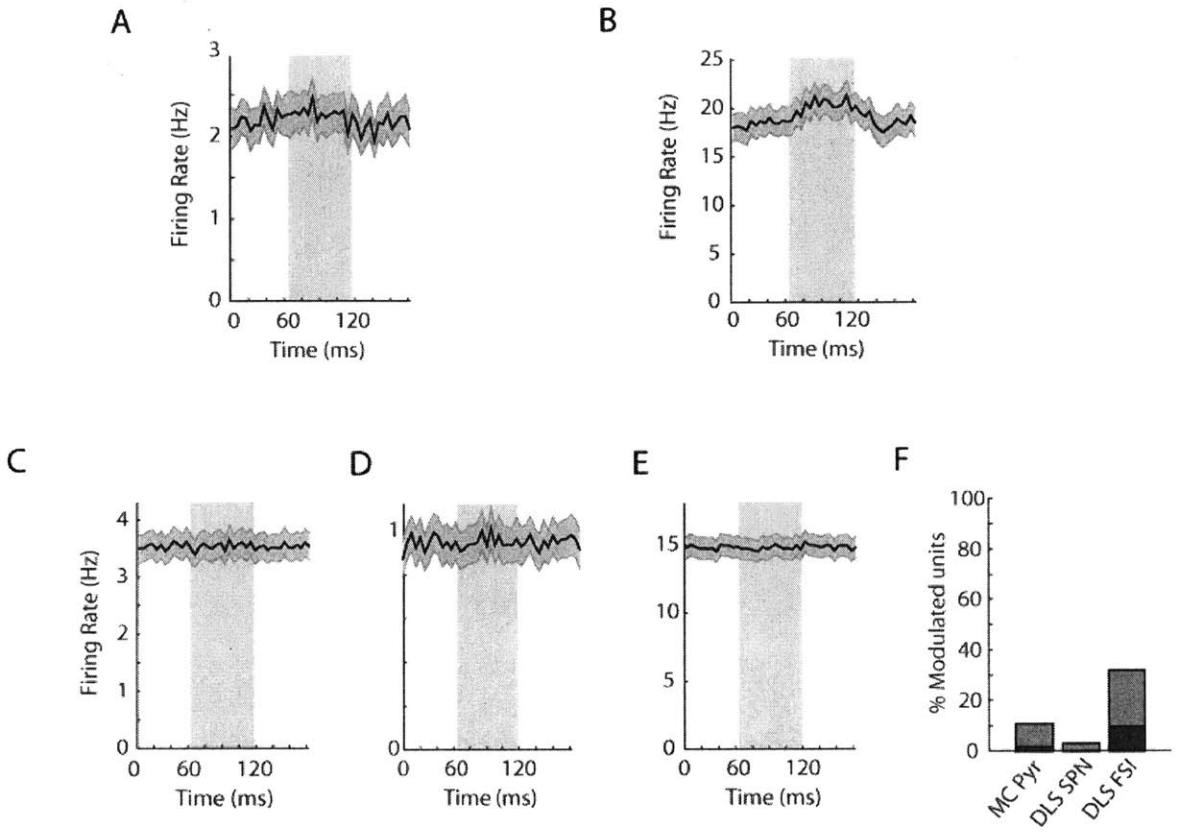

Supplementary Figure 5.

(A) Average firing rate of all putative DLS SPNs with firing rates > 1Hz (n = 96 units) when pulses of yellow light were delivered to cortical cell bodies. Light on period indicated with yellow shading.

(B) Average firing rate of the significantly activated putative DLS FSIs (n = 33 units out of 106 units) when pulses of yellow light were delivered to cortical cell bodies.

(C) Average firing rate of putative cortical pyramidal neurons when pulses of yellow light were delivered to the halorhodopsin expressing terminals *in striatum* in the freely moving rat (n = 383 units). Light on period indicated with yellow shading.

(D) Average firing rate of putative DLS SPNs recorded at the same time (n = 311 units).

(E) Average firing rate of putative DLS FSIs recorded at the same time (n = 108 units).

(F) Proportions of significantly inhibited (blue) and activated (red) units for each cell group for yellow light pulses delivered in striatum to halorhodopsin expressing terminals.

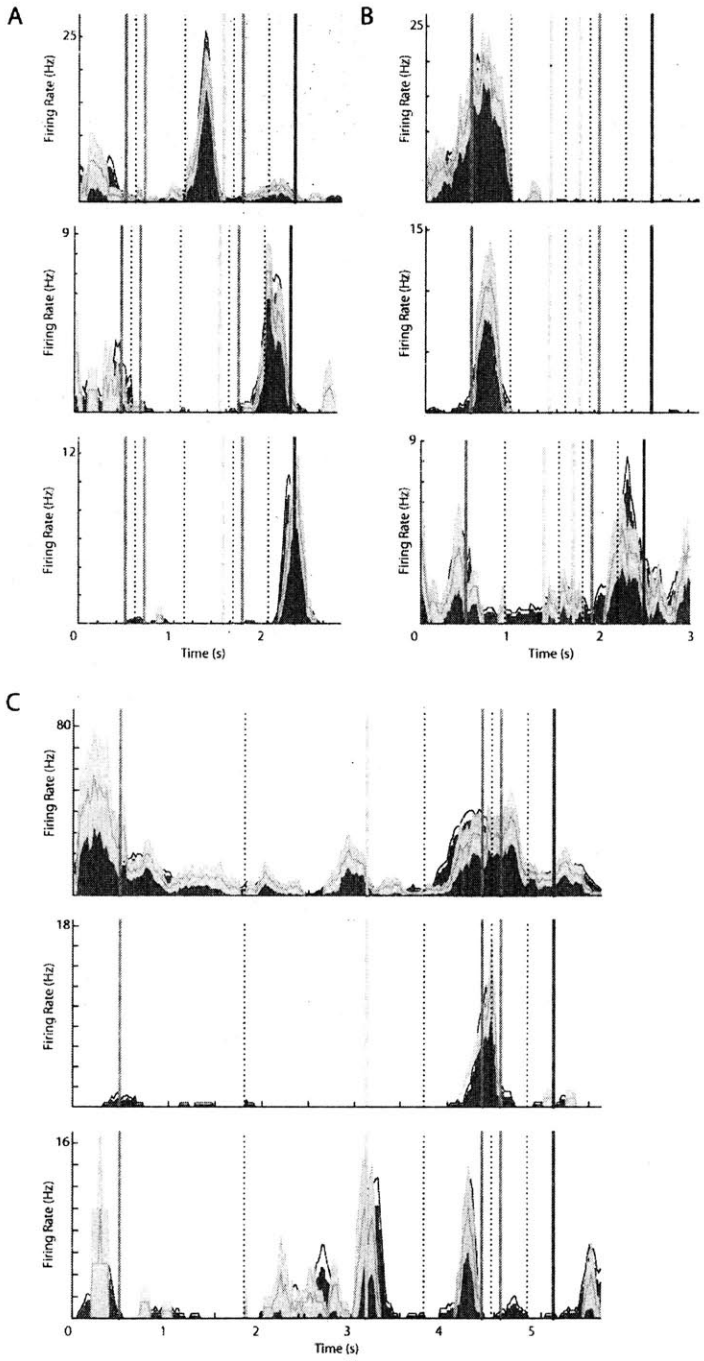

Supplementary figure 6.

(A) Examples of single units recorded from one animal which fired similarly during laser-off times (gray) and laser-on times (yellow).

(B) And (C) Other single units from two more animals which fired similarly during laser-off and laser-on times.

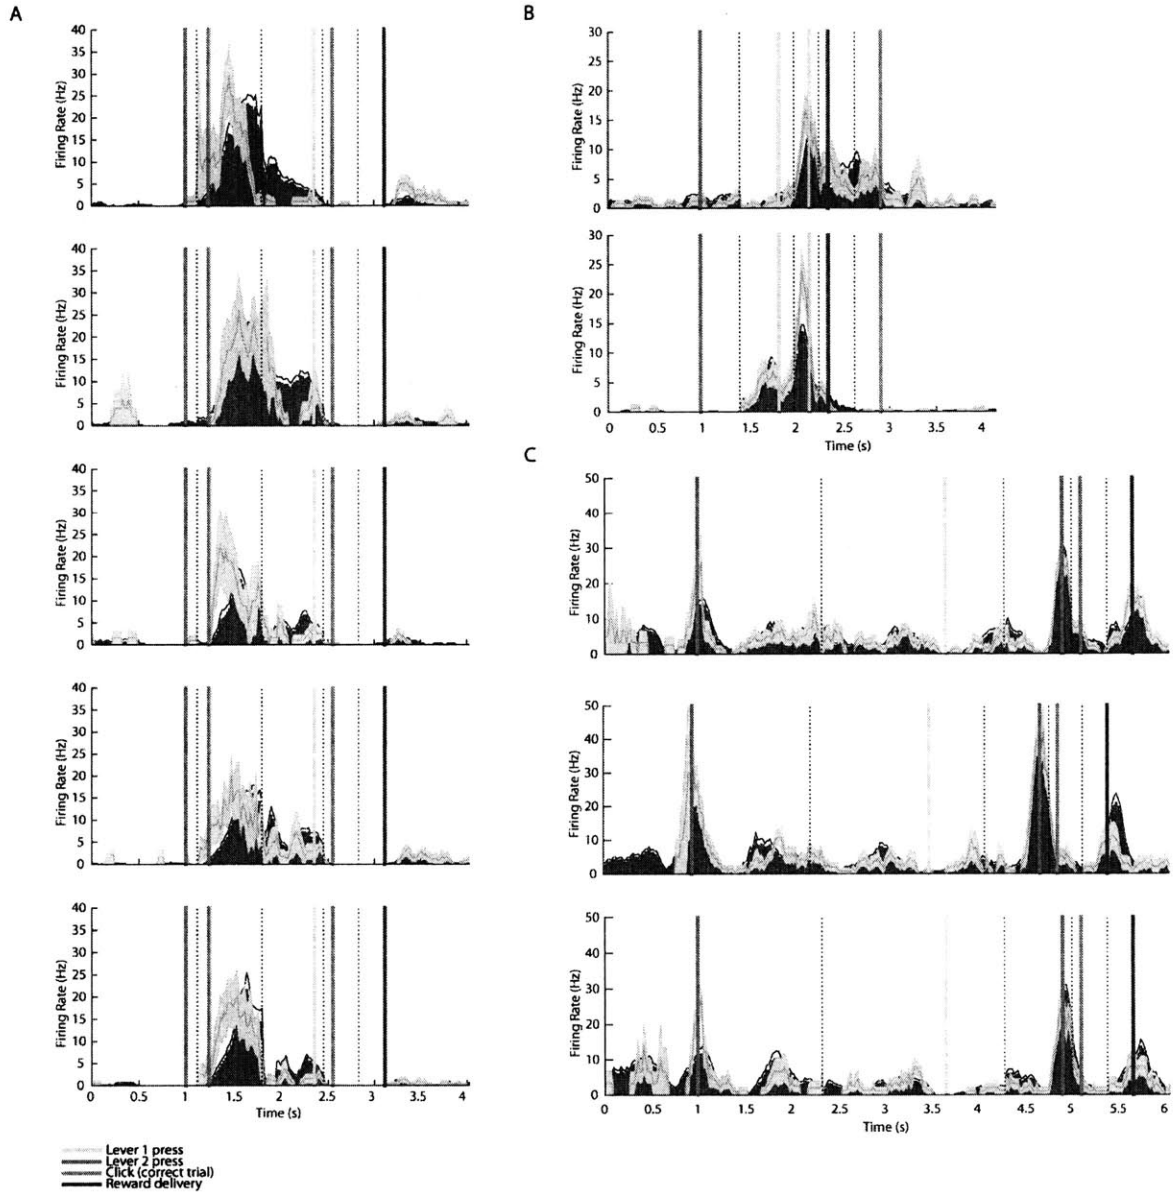

Supplementary figure 7.

(A) An example of a DLS SPN recorded across several days (rows) which was modulated by the laser light directed at motor cortex terminals in DLS. Laser-off firing rate is shown in black and laser-on firing rate is shown in yellow.

(B) And (C) Other examples of laser modulated DLS SPNs recoded across days.

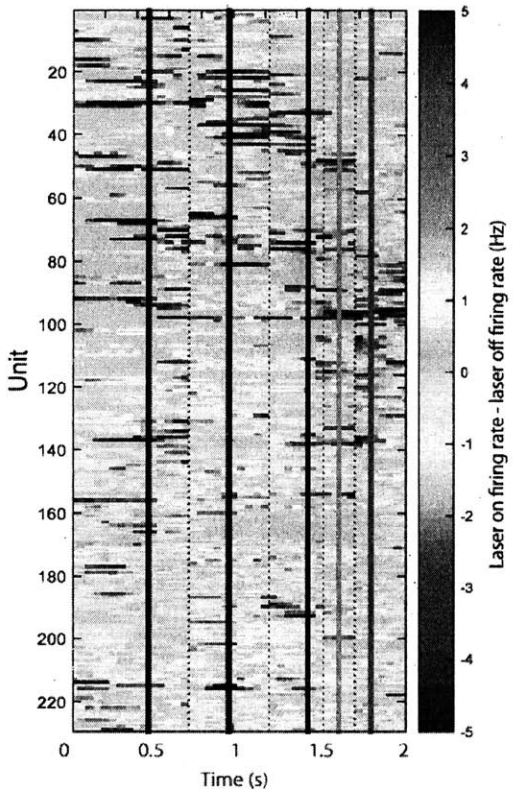

Supplementary figure 8.

All DLS MSNs which were significantly modulated by laser targeted to halorhodopsin expressing motor cortex terminals in the DLS during task time. Warm colors indicate activation during laser-on times and cool colors indicate inhibition during laser on times. Units are grouped by tetrodes, revealing some tetrodes with consistent timing and direction of modulation by the laser.

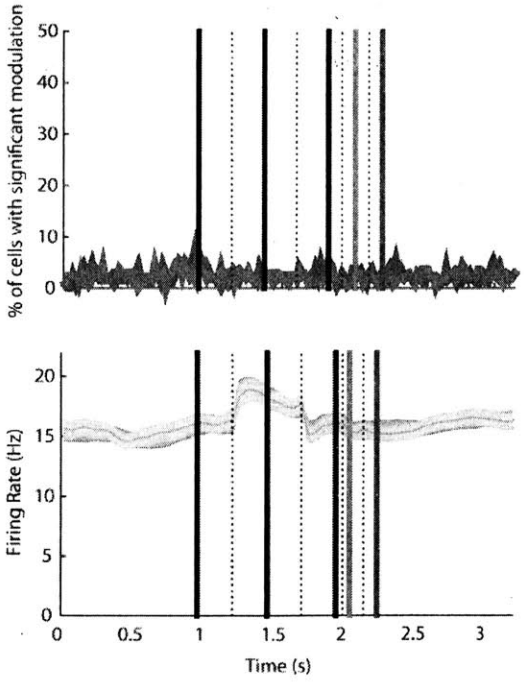

Supplementary figure 9.

(Top) Proportions of DLS FSIs inhibited (blue) and activated (red) by laser targeting motor cortex terminals in DLS throughout task performance.

(Bottom) The population FSI firing in during correct trials during laser-off periods (gray) and laser-on periods (yellow) by an optical fiber placed in striatum (n = 308 units).

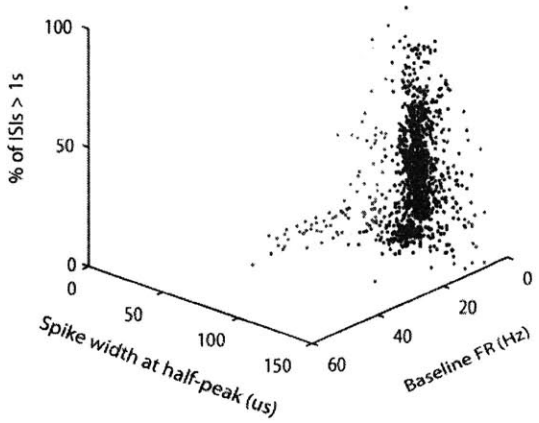

Supplementary Figure 10.

Properties of units recorded from motor cortex. Units in the blue cluster were classified as putative cortical projection neurons and were used in further analyses.

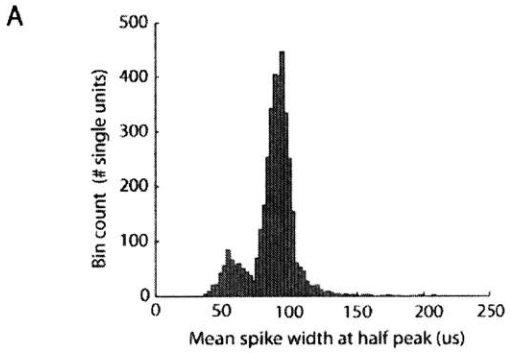

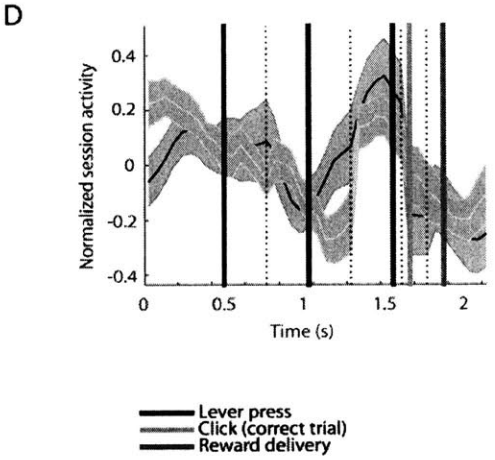

Supplementary figure 11.

(A) The distribution of mean spike width for all units recorded in DLS

(B) Clusters of DLS units classified as MSNs (blue, wide spike waveform and low firing rates), FSIs (red, narrow spike waveform and high firing rates), wide-waveform fast spiking units (green), and unclassified (gray).

(C) The relationship of spike amplitude and spike width in the same groups of single units. The bimodal distribution of spike width in the dataset does not appear to be driven by spike amplitude differences.

(D) Comparison of normalized session activity during correct sequence performance in DLS SPNs (gray) and wide-waveform fast spiking units (green) in n = 32 sessions.

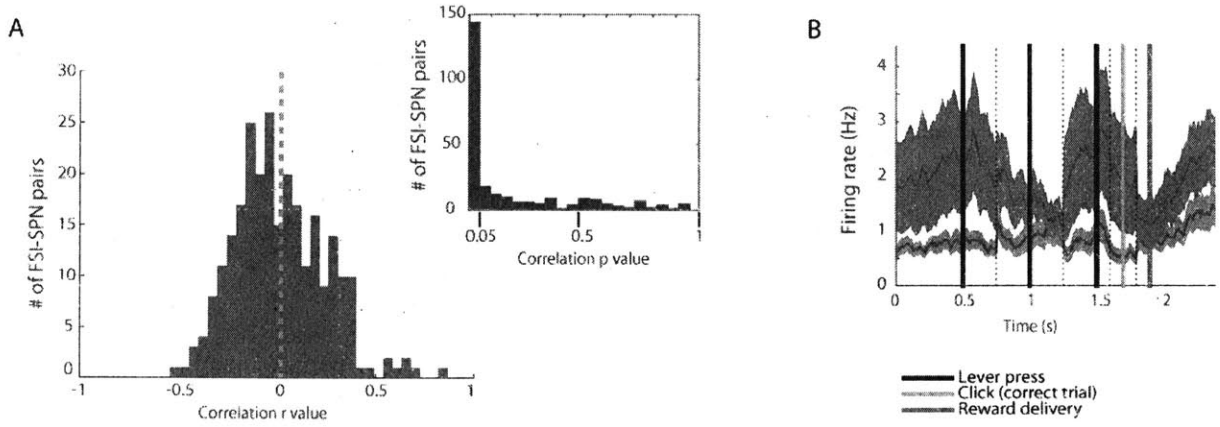

Supplementary figure 12.

(A) The distribution of r values and p values (inset) of correlation between firing rates of 265 pairs of narrow-waveform FSIs and SPNs recorded on the same tetrode.

(B) Task-related firing of SPNs that were negatively correlated with an FSI on the same tetrode (blue, n = 72 units) and SPNs that were positively correlated with an FSI on the same tetrode (red, n = 72 units).

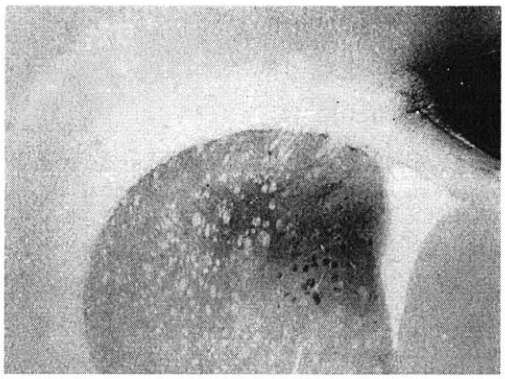

Supplementary figure 13. A test injection of virus expressing YFP into prelimbic cortex results in YFP expressing terminals in dorsomedial striatum.

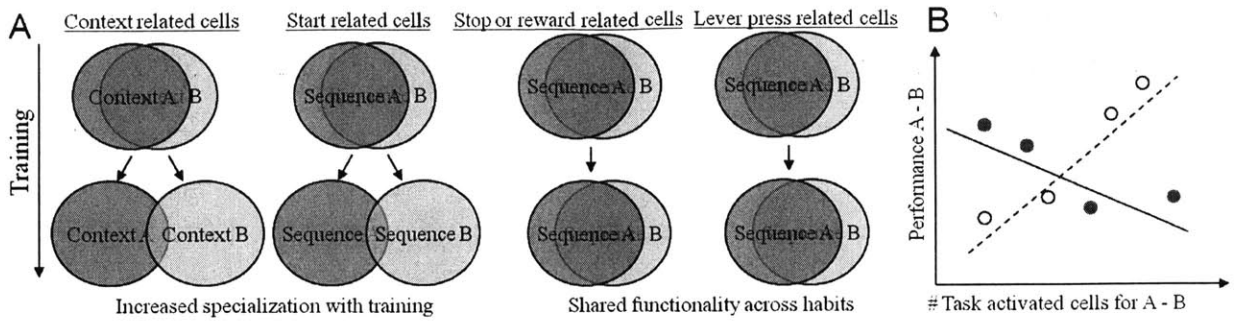

Supplementary figure 14. Possible hypotheses about the neuronal representation of two different lever press sequences before and after learning.

(A) With training some neurons could become specialized for one of the two habits while others will continue to be shared across them. Context and start related neurons may become specialized with training. Stop or reward related cells and cells related to lever press or other motor components of behavior may continue to be shared across sequences.

(B) Performance may be positively (empty circles) or negatively (filled circles) correlated with the proportion of neurons responsive to the task.

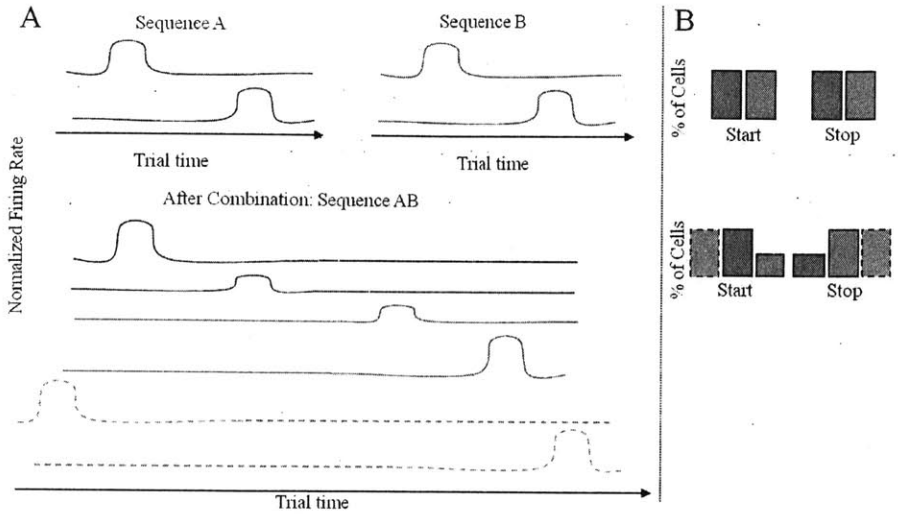

Supplementary figure 15.

Possible hypotheses for the effect of combining two well-learned lever press sequences on the task-boundary activity in DLS.

(A) Firing rate changes: Before combination, each sequence will activate some start and stop cells in DLS. After the combination, it is possible that the activation of the start stop cells that now fall in the middle of the sequence is reduced. There could also be the appearance of a new set of start stop cells (dashed orange line).

(B) Changes in number of neurons: The number of start stop cells for each sequence may be similar for equally well learned sequences or could depend on other factors. With the combination, the number of the cells that respond to events that are now in the middle of the sequence can decrease. New neurons can also appear (in orange) that represent start/stop of the new sequence. A combination of firing rate changes and number of neuron changes could occur.

MATERIALS AND METHODS

Animals

Twelve adult male Long Evans rats (eight wildtype and four Chat-Cre rats) were used for recording experiments and four additional wildtype Long Evans rats were used only for behavior experiments. Three of these twelve were not included in the data set - one due to the loss of the implanted drive, one due to poor striatal recordings, and one due to very few correct trials performed. They were housed under a 12hr reverse dark/light cycle and trained during the dark cycle. After the start of behavioral training, rats were kept on mild food restriction with 15g of food/day and allowed to reduce up to 85% of their free feeding weight.

Recording drive

Custom-build recording drives with 28 independently moveable microdrives was assembled to hold 24 tetrodes house-made with 4 twisted tungsten wires and two 200um optical fibers with a zirconia ferrule from Doric Lenses. Tubes running from the microdrives to drive tip were arranged such that each of them contained an optical fiber in the center of the bundle and was surrounded by 10 tetrode tubes all positioned ~150um away from the edge of the optical fiber.

Surgical procedures

Rats were initially anesthetized with 3% isoflurane gas and 0.2mL of ketamine IP, and thereafter kept under 1-2% isoflurane gas anesthesia. Rats were secured in a stereotaxic frame and small drill holes (1.5mm wide) were made over left striatum (AP 0.5, ML 3.5, tetrodes lowered to DV -4) and left motor cortex (AP 0, ML 1.5, tetrodes lowered to DV -1). Some of the rats additionally had drill holes and tetrodes in left prelimbic cortex (AP 3, ML 0.4, tetrodes lowered to DV -3) and left dorsomedial striatum (AP 1.7, ML 1.7, tetrodes lowered to DV -4). Several small burr holes were also made on other areas of the skull for anchor screws before the recording drive was attached to the skull. All tetrodes were lowered to within ~200um of the desired depth on the day of the implantation. Optical fibers were lowered 1mm on the day of the implantation and

additional turns were done on the days following surgery to reach the target depth. Tetrodes and optical fibers were left in place after reaching the target depth except for rare small adjustments of < 125um to maintain a high yield or single unit recordings.

Four weeks prior to the recording drive implantation, the halorhodopsin virus injection procedure was done in six rats. 0.5uL of AAV-CaMKII-eNpHR3.0-EYFP injections were made into motor cortex bilaterally (AP 0, ML 1.5, DV -1) at a rate of 0.05uL/min. The virus was obtained from the University of North Carolina vector core facility.

Behavioral-training

Rats were habituated to being handled, to drinking chocolate milk, and to the operant chamber prior to the start of lever press training. They were then taught to associate the auditory click tone with chocolate milk delivery and learned to press the levers in two training days in which they were rewarded for pressing either of the two levers. Following this, rats were trained for two days in which they were rewarded after any three lever presses. Rats were then assigned a correct sequence and were rewarded 0.2mL of chocolate milk after performing the correct sequence. They received a small 0.05mL chocolate milk reward in a random 20% of the trials which was gradually decreased to a random 5% as their performance improved.

Devaluation procedure

Four rats were used in the devaluation test. Each of the rats received 3-4 probe sessions across training in which unlimited chocolate milk was made accessible to them for 2 hours prior to the training session. The devaluation probe sessions were 30 minutes in duration and during this time the normal training setup was available to the rat; however, while there was chocolate milk in the reward tube, the chocolate milk syringe was not placed in the pump and therefore no chocolate milk was delivered to the rat in case of correct presses.

Optogenetic Inhibition of cortical terminals in-task

In six rats, I used the 200um optical fiber placed in center of the array of 12 tetrodes in the dorsolateral striatum to inhibit halorhodopsin expressing cortical terminals in striatum. The laser patch cord, from the laser located in an adjacent room was routed to the recording room and through the commutator on the ceiling of the behavior chamber. The laser shutter was controlled by TTL pulses initiated by custom MATLAB behavior software. TTL pulses were copied to two channels in the Neuralynx Data Acquisition system and saved timestamps for laser On/Off times. The laser shutter was opened in one of four periods (1) after the end of the previous trial until the first press of the next trial, (2) after the first press in the trial until the second press, (3) after the second press in the trial until the third press, (4) after the third press in the trial for 3 seconds. These laser trial types were randomly intermixed within the training session, and in 20% of the trials no laser light was used. In any case, when the laser light was illuminated for 3s, it was automatically turned off to prevent heating. For the analysis of the effect of the laser manipulation on neural activity in the striatum, all laser-off periods are combined and shown together, and all laser-on periods are combined and shown together. Post-hoc, any periods where the laser had been on for longer than 2s were not included in the analysis.

Optogenetic inhibition of cortical cell bodies and cortical terminals out-of-task

After each behavioral training and recording session, I tested the responsiveness of the recorded units to cortical cell body and terminal inhibition. I used 3mW of yellow light targeted to the 200um optical fiber located in dorsolateral striatum and then on the optical fiber located in motor cortex. In each test session, 500 trials of 70ms pulses were used.

Data acquisition

Lever press events were recorded using Med Associates hardware and Med Associates MATLAB toolboxes along with custom MATLAB behavioral control software which received timestamps from the Neuralynx Data Acquisition system. Two video cameras connected to the Neuralynx system were mounted above the operant chamber. One was at the top of the chamber and

gathered head position data with the use of red and green LEDs mounted onto the preamplifiers of the recording system. The second was attached to the wall closer to the levers and the reward well and recorded video of the behavior. Unit activity (gain: 200-10,000, filter: 600-6,000 Hz) was recorded with the Neuralynx Data Acquisition System. Spikes exceeding a preset voltage threshold on any of the four tetrode channels triggered the waveform to be sampled at 32 kHz and stored on all four channels.

Spike sorting and quality assessment

Spike data were manually sorted into single units using Plexon Offline Sorted. Additionally, 1/5 of the sessions were spike sorted using an automated clustering procedure (Friedman *et al.*, 2015). All sorted clusters were graded according to a custom algorithm. Each cluster received a grade on waveform quality based on similarity in the waveforms and presence of a valley, and a cluster quality grade based on L-ratio, the distance from other sorted units and from the noise cluster, percent of the short interspike intervals, % of the cluster below threshold, and the continuity of the spikes throughout the recording session. Based on these measures, each cluster was assigned a grade of 1-5 and only clusters with grades >= 3 were used in the analysis (71% of the manually sorted units yielding 2892 DLS units and 777 MC units, and 32% of the automatically sorted units yielding 620 DLS units and 133 MC units). I found for these highly rated units which were classified as one of the cell types I later identified, the electrophysiological properties and task responses of the manually sorted units and of the units sorted using the automated clustering procedure were very similar.

Classifying putative cell types

I used a custom toolbox to visualize the distribution of various electrophysiological properties of the single units in each brain region. I found that waveform width was one primary factor along which distinct clusters appeared in the set of recorded units. In the striatum, I assigned the larger cluster with waveform width at half-peak of greater than 70 us and less than 125us, of baseline firing rates of < 3.5Hz as putative Spiny Projection Neurons. I assigned the smaller cluster with waveform width at half-peak of < 70ms and less than 10% of interspike intervals longer than 1s

to be putative Fast Spiking Interneurons (Fig. 3A). As a precaution against the potential of distant spikes to appear narrower, low-amplitude narrow-spike units were not included in the analysis. I isolated high-firing rate, wide-waveform units as a possible third cell type. Based on comparison with data previously gathered in our laboratory, I found that very few of the units that I recorded satisfied the criterion of putative tonically active neurons (putative cholinergic interneurons) (Atallah *et al.,* 2014) and these neurons were not included in the analysis.

In the motor cortex, I also found a division of units across spike width and assigned units with spike width at half peak of greater than 70ms and less than 125ms as putative pyramidal neurons. The small number of narrow waveform putative interneurons were not included in the analysis.

Analysis of neural activity

Peri-event spike histograms were created with custom MATLAB code using 50ms sliding window and 250ms bin width for each event (lever presses, click or noise feedback, and reward delivery) in each trial type. In individual examples, the peri-event histograms were pasted together using window sizes corresponding to the median time between each set of events. In the population activity plots, peri-event histograms were pasted together using +-0.5s around each event due to the variability across rats and sessions in the time between successive lever presses.

For the analysis of the neural activity pattern across all animals, the average neural activity in each rat was normalized by subtracting the mean firing rate of the population activity in the correct trials. For the comparison of simultaneously recorded SPN and FSI neural activity, and SPN and motor cortex neural activity, I used session averages to eliminate possible confounds of cell number and behavioral differences. Average session activity from each session with at least 5 recorded SPNs and 2 recorded FSIs was normalized and the means of the session activities of the two groups were compared.

Comparison of task representations in motor cortex and striatum

To assess whether the task-related firing of single units could be accounted for with the occurrence of single motor events, I compared the task responses of each neuron during lever 1

press and lever 2 press events, and the approach of these levers. To assess whether the units responded similarly in every instance of these events I identified all the different contexts in which the lever press occurred – such as the first, second, or third press in a correct sequence, or the first, second, or third press in each of the incorrect sequences. If the unit fired 2SD above baseline during each of these scenarios and the firing rates were similar in each case, I included this unit in the list of units whose task-related firing could be well accounted for by single event occurrences. To compare the incidence of such units in motor cortex and striatum, I identified sessions in which I recorded at least 10 units simultaneously in both motor cortex and striatum and compared the proportion of such motor units for in both areas.

Assessing cortical neuronal responses to SPN bursts

The procedure for finding spike bursts DLS SPNS was to, (1) look for instantaneous firing rates (1/ISI) of over 1 STD over the mean, (2) find times when this happens at least 3X in a row, and (3) within that train of spikes, identify if the whole train or part of the train has a firing rates 3 STD > than the mean firing rate of the cell. Such bursts were identified during correct trial execution in SPNs and peri-SPN burst histograms were created for each putative motor cortex pyramidal cell that was simultaneously recorded. Of all such DLS SPN – MC Pyr pairs, I identified pairs in which there was a significant effect of the SPN burst by comparing the numbers of spikes in the motor cortex neuron in the 200ms period prior to and 200ms period after each SPN burst using a Wilcoxon Rank Sum test. For all such significant pairs of SPN – MC Pyr neurons, I conducted a principal components analysis on the peri-event firing histograms of the MC units and used k-means clustering to separate two groups of response types.

Histology

At the end of training, rats were deeply anesthetized (Nembutal, 50-100 mg/kg), and lesions were made to mark the final recording sites (25 μA, 10 s). Rats were then perfused with 4% paraformaldehyde (PFA) in 0.1 M sodium-potassium phosphate buffered saline (PBS), and 30 μm thick transverse frozen sections were stained for CD11 to identify lesion sites and by immunohistochemistry for GFP protein to visualize viral expression in cell bodies and terminals.

CHAPTER 3: EFFECTS OF DOPAMINE DEPLETION ON LFP OSCILLATIONS IN STRIATUM ARE TASK- AND LEARNING-DEPENDENT AND ARE SELECTIVELY REVERSED BY L-DOPA

Nuné Martiros[1], Ledia F. Hernandez[1], Dan Hu, Yasuo Kubota, Mark W. Howe, and Ann M. Graybiel[2]

SUMMARY

A major physiologic sign in Parkinson's disease is the occurrence of abnormal oscillations in cortico-basal ganglia circuits, which can be normalized by L-DOPA therapy. Under normal circumstances, oscillatory activity in these circuits is modulated as behaviors are learned and performed, but how dopamine depletion affects such modulation is not yet known. We here induced unilateral dopamine depletion in the sensorimotor striatum of rats and then recorded local field potential (LFP) activity in the dopamine-depleted region and its contralateral correspondent as we trained the rats on a conditional T-maze task. Unexpectedly, the dopamine depletion had little effect on oscillations recorded in the pre-task baseline period. Instead, the depletion amplified oscillations across delta (~3 Hz), theta (~8 Hz), beta (~13 Hz) and low gamma (~48 Hz) ranges selectively during task performance times when each frequency-band was most strongly modulated, and only after extensive training had occurred. High gamma activity (65-100 Hz), in contrast, was weakened independent of task-time or learning stage. The depletion also increased spike-field coupling of fast-spiking interneurons to low-gamma oscillations. L-DOPA therapy normalized all of these effects except those at low-gamma. Our findings suggest that the task-related and learning-related dynamics of LFP oscillations are the primary targets of dopamine depletion, resulting in over-expression of behaviorally relevant oscillations. L-DOPA normalizes these dynamics except at low-gamma, linked by spike-field coupling to fast-spiking interneurons, now known to undergo structural changes after dopamine depletion and to lack normalization of spike activity following L-DOPA therapy.

[1] *These authors contributed equally to this work.*

INTRODUCTION

Loss of the dopamine-containing innervation of the basal ganglia is a primary pathology in Parkinson's disease, resulting, in addition to its behavioral effects, in abnormal local field potential (LFP) oscillations within cortico-basal ganglia circuits (Brown, 2003, Eusebio and Brown, 2007, Jenkinson and Brown, 2011, Jenkinson et al., 2012). Clinical evidence suggests that successful therapies for Parkinson's disease reduce these abnormal LFP oscillations (Ray et al., 2008, Jenkinson and Brown, 2011, Eusebio et al., 2012, Jenkinson et al., 2012), establishing them as a central feature of Parkinson's disease. In particular, abnormally strong beta-range oscillations (12-30 Hz) and weakened high frequency gamma oscillations (> 70 Hz) have been found in basal ganglia structures. The 'anti-movement' beta-band oscillations are reduced by both L-DOPA therapy and by deep brain stimulation (DBS) (Ray et al., 2008, Jenkinson and Brown, 2011, Eusebio et al., 2012, Jenkinson et al., 2012). How these observations relate to the proposed network functions of oscillatory neural activity is not yet clear. LFP oscillations have been linked not only to motor control but also to sensory perception, attention, learning and memory formation, and inter-regional communication (Buzsaki and Draguhn, 2004, Engel and Fries, 2010, Wang, 2010, Buzsaki and Wang, 2012, Siegel et al., 2012). In Parkinson's disease models, abnormal patterns of synchrony have been found in rest and locomotion (Fuentes et al., 2009, Avila et al., 2010, Brazhnik et al., 2012), but the effect of dopamine loss on LFP oscillations during complex tasks requiring learning and decision-making has not been explored.

Here we report that dopamine depletion in the sensorimotor striatum has striking effects both on oscillatory power in multiple frequency ranges and on spike-field synchrony, but that the abnormal patterns of synchronization are behaviorally regulated and are not omnipresent features of the dopamine-depleted state.

RESULTS

Sixteen rats were given unilateral 6-hydoroxydopamine (6-OHDA) injections to produce local dopamine depletion in the dorsolateral striatum. Fast-scan cyclic voltammetric measurements were made in 4 other rats; these demonstrated that the localized 6-ODHA lesions reduced evoked dopamine release by about 75% relative to levels in the contralateral striatum or ipsilateral striatal regions outside of the injected zone (Ledia Hernandez, 2012) After a 5-week lesion stabilization period, the rats were trained on a T-maze task while we recorded spike and LFP activity in the dopamine-depleted dorsolateral striatum and the contralateral intact dorsolateral striatum (Fig. 1*A*). The spike activity patterns are reported in a companion study (Ledia Hernandez, 2012). The T-maze task began with a click signal indicating the trial initiation, followed by the opening of the gate allowing rats to run down the maze toward one of the two end-arms. Before the rats reached the choice point, one of two auditory cues instructed which maze goal was baited with chocolate reward (Fig. 1*B*). The localized 6-OHDA lesions did not produce apparent difficulties with learning of the task: the rats reached the 72.5% correct learning criterion in an average of 8.3 sessions, a rate comparable to that of normal rats in a similar maze task (Barnes et al., 2005). The percentage of correct trials rose from chance levels to above 85% late in training ($P < 0.0001$, ANOVA, Fig. 1*C*). Running times (from start to goal-reaching), which early in training were slightly longer than those observed for normal rats, decreased significantly as the rats acquired the T-maze task (Fig. 1*D*, $P < 0.0001$) and became comparable to those of normal rats (Barnes et al., 2005).

Dopamine Depletion Amplifies Oscillations Selectively only during Task Times in Which They Are Actively Modulated.

Confirming earlier observations (DeCoteau et al., 2007), task-related modulation of oscillatory LFP activity was apparent in the sensorimotor striatum of the intact hemisphere. Oscillations in the theta range (6-9 Hz) strengthened as the animals started to run down the maze, subsided at the end of the run near goal and were often replaced by a slightly more rapid low beta oscillation (11-15 Hz) at goal-reaching (Fig. 1 *E* and *F*). Delta (2-5 Hz) and gamma oscillations (40-53 Hz) also

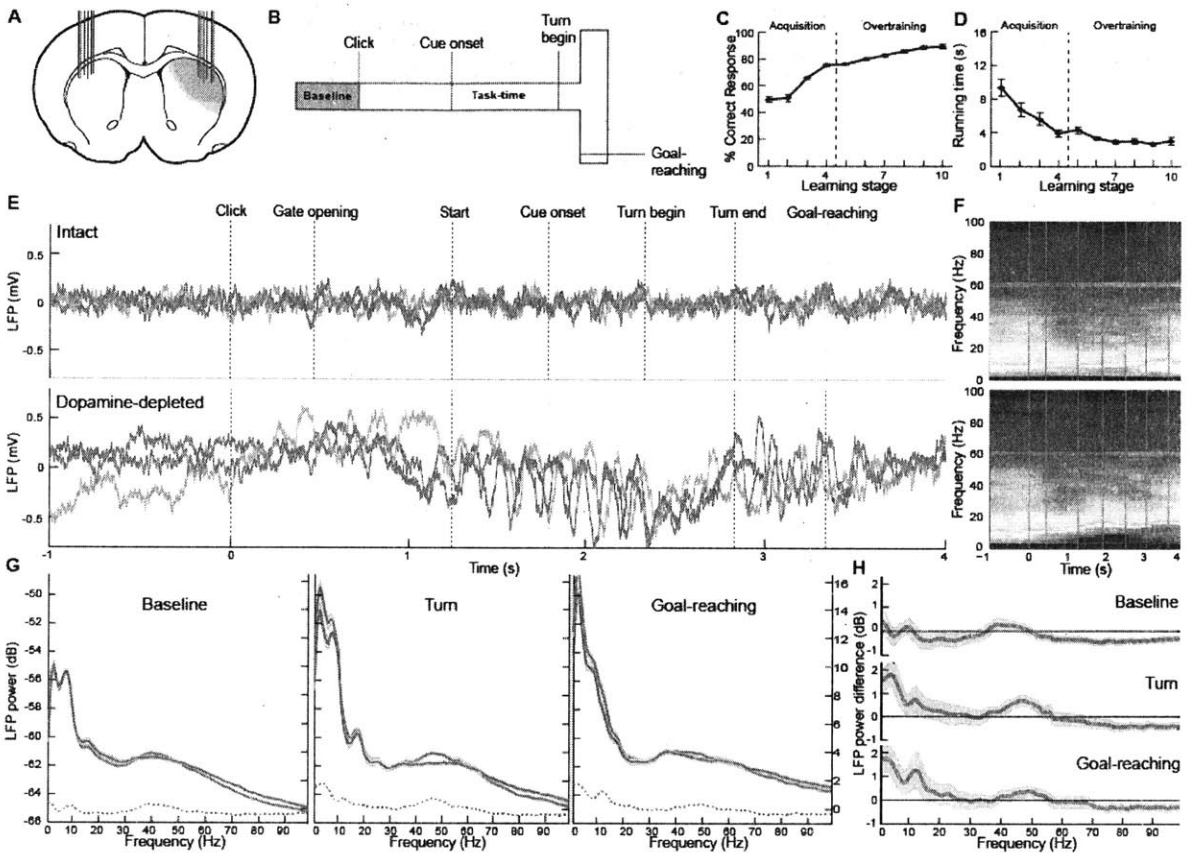

Fig. 1. Dopamine loss amplifies LFP oscillations during performance of a T-maze task. (*A*) The extent of the local, unilateral dopamine-depletion and location of bilateral recording sites in the dorsolateral striatum (DLS). (*B*) A cue-instructed turning task in a T-maze. Baseline period began 2 s before a click to indicate trial start. (*C* and *D*) Mean % correct responses (*C*) and running time (*D*) across training (n = 16 rats). (*E*) Raw LFP trace from three trials in a single training session recorded simultaneously by tetrodes in the intact (*Top*) and dopamine-depleted (*Bottom*) DLS. (*F*) Trial average spectrograms for LFP activity shown in *E*. (*G*) Aggregate session average spectra (single taper), calculated from overtraining sessions of 16 rats. Each plot shows power of oscillations during 1 s periods during pre-trial baseline (n = 114 sessions), around turn beginning (n = 112 sessions), and around goal-reaching (n = 107 sessions) recorded in the intact (blue) and dopamine-depleted (red) DLS. Shading indicates SEM. Dashed lines show power difference between the intact and depleted sides (scale at right). (*H*) Session average differences in LFP power between the intact and dopamine-depleted DLS. Positive values indicate depleted side > intact side. Shading indicates ±2 SEM (95% confidence interval for difference from zero). Only sessions with simultaneous recordings in both sides were used.

occurred during task execution (Fig. 1 *F* and *G*), as well as a harmonic of the 8 Hz theta oscillation at 16 Hz (Fig. 1*G*, middle, and Fig. S1). Similar task-related oscillatory dynamics were present in the dopamine-depleted sensorimotor striatum, but these were markedly enhanced only at specific task points and in select task-relevant ranges of frequencies (Fig. 1 *G* and *H*). We focused on the prominent delta, theta, beta, and low gamma oscillations to determine the role of dopamine depletion on modulations of these rhythms. For all but the high-gamma oscillations, the oscillatory patterns were strongly enhanced as the rats ran in the maze (delta, theta and low gamma rhythms, $P < 0.01$ for all, ANOVA), or at goal-reaching (low beta rhythm, $P < 0.03$; Figs. 1*H* and 2*D*). Remarkably, no significant differences in power were observed during the pre-trial baseline period in any of these frequency bands ($P > 0.3$ for all; Figs. 1*G*, 1*H*, and 2*D*), a period prior to the initiation of each trial in which the rats were at rest while waiting for the trial to begin. LFP power in the high gamma range (65-100Hz) was reduced by dopamine depletion across the entire task and pre-task baseline periods ($P < 0.05$ at all task periods; Figs. 1*G*, 1*H*, and 2).

Effects of Dopamine Depletion on LFP Oscillations Emerge after Learning on the Associative T-Maze Task.

During training, power rose for the delta, theta and beta bands on both the intact side and dopamine-depleted sides ($P < 0.05$ for all, ANOVA; Fig. S2), indicating that the strength of the oscillations increases with task experience. However, this process was significantly augmented in the absence of dopamine. During acquisition training, power in these bands was similar on the two sides ($P > 0.1$, ANOVA; Fig. 2 *A* and *C*). During the overtraining period, however, the power in each of these frequency bands was enhanced on the dopamine-depleted side, relative to the intact side ($P < 0.05$; Fig. 2 *B* and *D*). Exceptionally, low gamma rhythms were elevated from early in training and remained so throughout training ($P < 0.05$ for each phase). High gamma oscillations were reduced in all training stages and all trial periods ($P < 0.02$ for all).

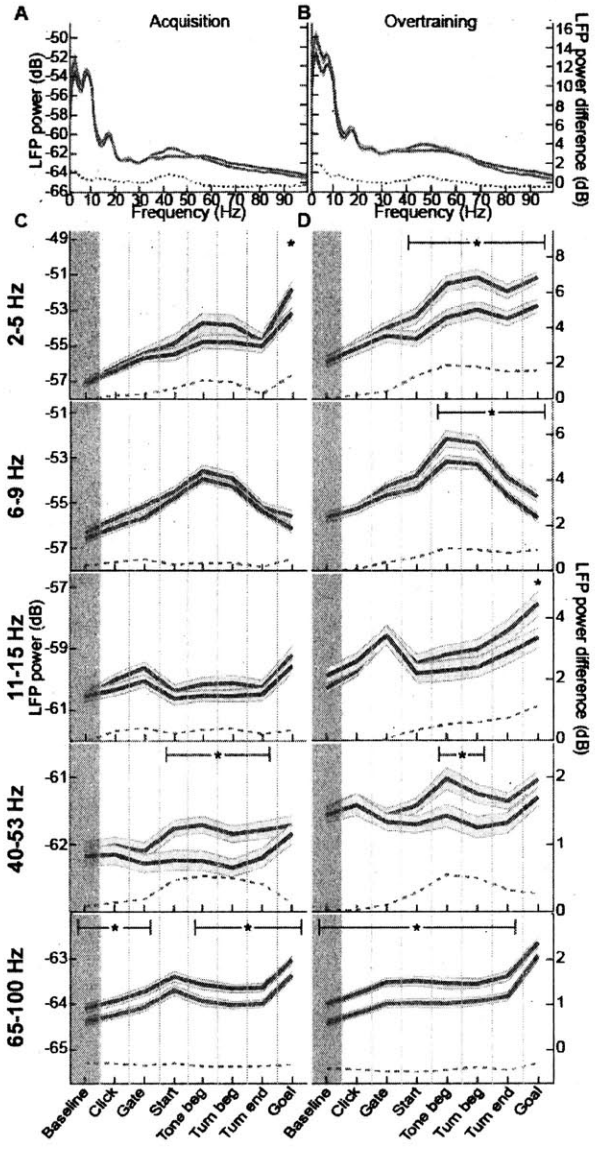

Fig. 2. Effects of dopamine depletion on LFP oscillations are dependent on learning phase and task time. (*A* and *B*) Session average spectra (single taper) for LFPs recorded in the intact (blue) and dopamine-depleted (red) DLS during ±500 ms around turn beginning in sessions before (*A*, n = 65 sessions from 10 rats) and after (*B*, n = 90 sessions from the same 10 rats) acquisition criterion was reached. Shading indicates ±SEM. Dashed lines show power difference between depleted and intact sides (scale at right). (*C* and *D*) Session average power (single taper) of oscillations in the delta (2-5 Hz), theta (6-9 Hz), low beta (11-15 Hz), low gamma (40-53 Hz) and high gamma (65-100 Hz) bands recorded in the intact (blue) and dopamine-depleted (red) DLS during acquisition (*C*) and overtraining (*D*) sessions. Each plot shows power during 1 s windows centered on successive task events. Brown shading indicates the baseline period. Red and blue shading indicates SEM. Dashed lines show differences between the two sides (scale at right). Only sessions with simultaneous recordings in both sides were used. $*P < 0.05$.

L-DOPA Normalizes Power in All Oscillations with the Exception of the Low Gamma Oscillation.

We compared the effects of systemic L-DOPA treatment on the LFP oscillations after extended overtraining. After L-DOPA administration, power in theta, beta, and high gamma frequencies was no longer different from that in the simultaneously recorded intact hemisphere in any of the task periods ($P > 0.1$, ANOVA), and delta band power was nearly normalized (Fig. 3). However, low gamma oscillations remained significantly elevated after L-DOPA administration ($P < 0.005$). This normalization of power in the low frequencies occurred due to a decrease in oscillation strength in the depleted dorsolateral striatum by L-DOPA ($P < 0.05$; Fig. S3) and not due to changes to LFP power in the intact dorsolateral striatum ($P > 0.1$). The power differences at high frequencies could have reflected a decrease in high frequency power in the intact sensorimotor striatum, or a combination of effects in the intact and depleted dorsolateral striatum (Fig. S3). L-DOPA treatment produced no significant differences in low gamma power in the dopamine-depleted dorsolateral striatum ($P > 0.2$), further indicating that this oscillation was not corrected by the treatment (Fig. S3).

Spike-LFP Relationships Are Selectively Affected by Dopamine Depletion.

Single unit spiking was recorded along with the LFP oscillations, and recorded units were putatively separated into neuronal subtypes according to previously used criteria (Barnes et al., 2005, Howe et al., 2011). The average firing rate of medium spiny projection neurons (MSNs) and fast-spiking interneurons (FSIs) was increased by the dopamine depletion as previously seen (Chen et al., 2001, Burkhardt et al., 2009, Ledia Hernandez, 2012). To address the effect of dopamine depletion on the spike-LFP relationships in the dorsolateral striatum, we analyzed the phase coupling of spikes to LFP oscillations, measured during the maze runs in overtraining sessions – the time during which we found strongest effects of dopamine depletion on LFP power. We exclude the high gamma oscillations in this analysis due to the likelihood that spike artifact (Fig. S6) is reflected in this high frequency range; however, we did not observe any effect of dopamine depletion in the spike-LFP coupling in the high gamma range.

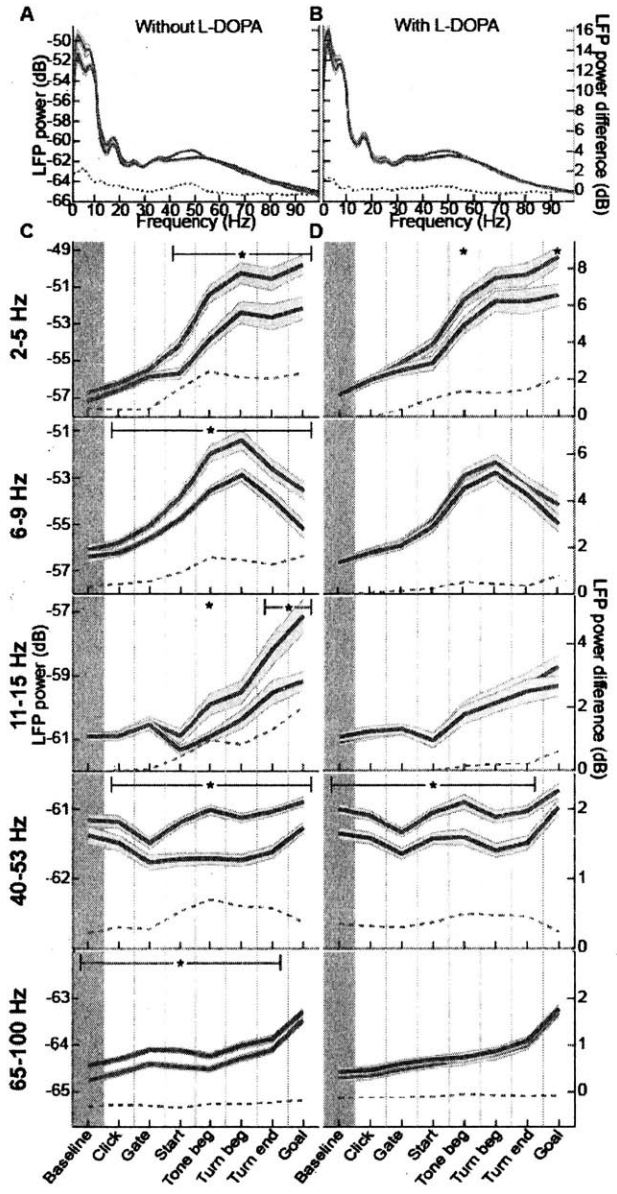

Fig. 3. L-DOPA treatment normalizes LFP oscillations in the dopamine-depleted DLS, with the exception of the low gamma oscillation. (*A* and *B*) Session average spectra calculated and shown as in Fig. 2 *A* and *B*, for overtraining sessions before (*A, n* = 61 sessions from 10 rats) and during (*B, n* = 70 sessions from same 10 rats) daily L-DOPA treatment. (*C* and *D*) Session average power of oscillations in different frequency bands, as in Fig. 2 *C* and *D*, for sessions before (*C*) and during (*D*) L-DOPA treatment. *$P < 0.05$.

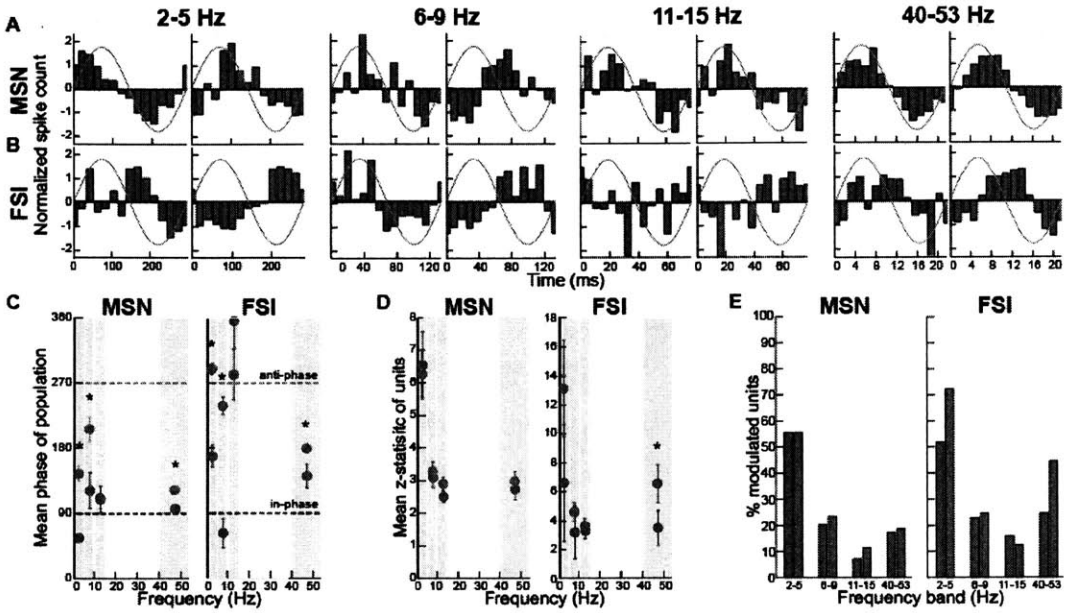

Fig. 4. Dopamine depletion shifts phase preferences of MSN and FSI spiking and strengthens FSI spike coupling to low gamma oscillations. (*A* and *B*) Normalized population spike phase histograms for MSNs (*A*) and FSIs (*B*) in the intact (blue, all spikes from 662 MSNs and 66 FSIs) and dopamine-depleted (red, all spikes from 712 MSNs and 82 FSIs) DLS. LFP activity was recorded during a ±1.5 s window around cue onset in overtraining sessions and filtered for frequencies for each band. (*C*) Mean phase of population spiking of MSNs and FSIs in the intact (blue) and dopamine-depleted (red) DLS. Error bars indicate 75% confidence interval. $*P < 0.05$ (ANOVA for circular distributions) between intact and depleted sides. (*D*) (*Left*) the average Z-statistic (strength of phase locking) from Rayleigh's test of individual MSNs in the intact (blue) and dopamine-depleted (red) DLS. Only units with at least 100 spikes in analyzed window were used and were normalized to 100 spikes each ($n = 227$ in intact DLS, $n = 317$ units in depleted DLS). (*Right*) FSI average Z-statistics for FSI units similarly normalized to 1000 spikes each ($n = 57$ units and 82 units, respectively). Error bars indicate SEM. $*P < 0.05$ (Mann-Whitney test between the two sides) (*E*) Proportions of MSNs and FSIs shown in *D* that are significantly phase-modulated ($P < 0.01$, Rayleigh's test) by oscillations in the given frequency bands.

The population spiking of MSNs and FSIs in the control and depleted striatum displayed distinct phase preferences with respect to the on-going LFP oscillations in the relevant frequency ranges ($P < 0.02$, Rayleigh's test, for all except FSI-low beta; Fig. 4 *A* and *B*). Phase locking of MSNs and FSIs was strongest to delta oscillations (Fig. 4*D*). After normalizing for firing rate, spike-LFP coupling in individual units was subtly, but not significantly, stronger in the dopamine-depleted

hemisphere ($P > 0.05$, Man-Whitney test; Fig. 4*D*). Notably, however, FSIs displayed significantly stronger entrainment to low gamma oscillations in the dopamine-depleted hemisphere relative to the intact ($P < 0.02$, Man-Whitney test; Fig. 4*D*). The percentage of significantly phase-modulated individual MSNs and FSIs ($P < 0.05$, Rayleigh's test) was further consistent with these findings, as the lesion produced a substantial increase in the proportion of FSIs entrained to the low gamma rhythm from 33% to 57% (Fig. 4*E*).

During the baseline period, we found weaker, but still significant, spike-LFP coupling (Fig. S4), and small significant increases in MSN spike-LFP coupling due to dopamine depletion (Fig. S4*D*). Largely different sets of MSNs were active during the baseline period and the task-time (Ledia Hernandez, 2012), such that the MSN population included in the baseline analysis was different than that included for the in-task analysis. This difference could account for differences in baseline and in-task coupling effects. After L-DOPA treatment, similar trends in the MSNs and FSIs remained (Fig. S5). The low number of FSIs in this recording stage made it is difficult to assess the effect L-DOPA on spike-LFP coupling for FSIs, but we continued to see a trend of an increased proportion of FSIs modulated by the low gamma oscillations in the dopamine-depleted hemisphere (Fig. S5*E*). Despite the lack of strong global effects of the local dopamine depletion on the strength of phase-locking, dopamine depletion did result in a broad-band forward shift in the preferred phase of firing of MSNs and FSIs in-task ($P < 0.01$ for delta, theta, and low gamma in MSNs and FSIs, ANOVA for circular distributions; Fig. 4*C*) and at baseline (Fig. S4*C*), which persisted with L-DOPA treatment (Fig. S5*C*).

We constructed spike-triggered waveform averages (STWAs) with LFP traces filtered in the same frequency ranges and normalized for amplitude differences (see Supplementary Methods). The STWAs displayed prominent oscillations at frequencies in the range of the filtered LFPs, verifying the presence of significant spike-LFP relationships (Fig. 5). The delta band STWAs had the highest amplitude waveforms, indicating strong spike-LFP coupling in that range as seen in the spike phase histograms. In most frequency bands, the preferred phase of the MSNs and FSIs tended to shift forward on the dopamine-depleted side (i.e., the spike occurred at a later phase of the LFP; Fig. 5), also consistent with the spike phase histograms. Finally, the amplitude of the FSI STWA in the low gamma band was significantly larger in the dopamine-depleted dorsolateral striatum

than the intact dorsolateral striatum, indicating stronger phase-locking of the FSIs to low gamma oscillations (Fig. 5*B*).

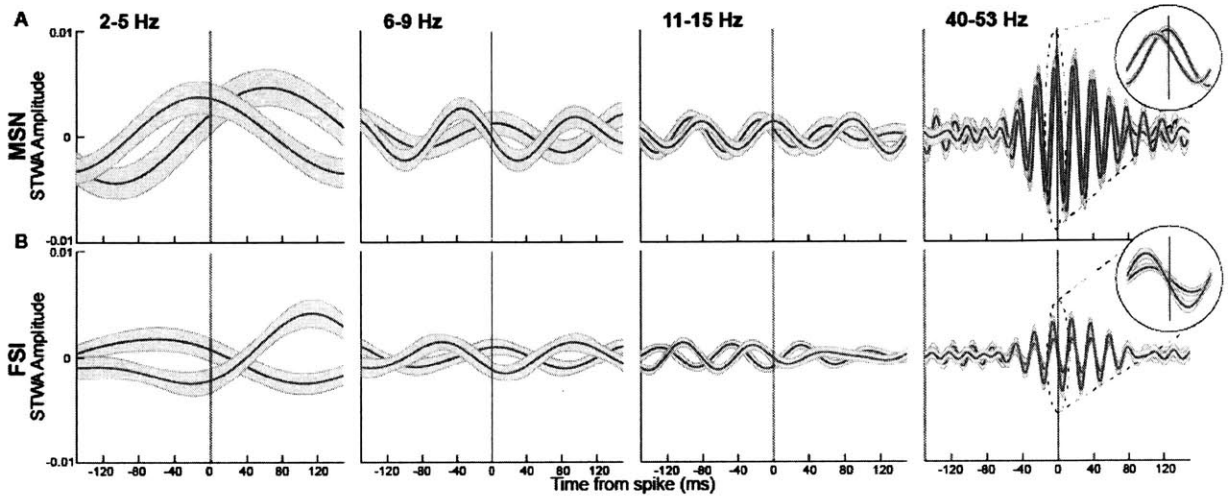

Fig. 5. Spike-LFP coupling is altered by dopamine depletion. Average LFP oscillations aligned MSNs (*A*) and FSIs (*B*) spikes, calculated for activity in the intact (blue, n = 227 MSNs and 57 FSIs) and dopamine-depleted (red, n = 317 MSNs and 81 FSIs) DLS during a ±1.5 s window around cue onset in overtraining sessions. Spike-triggered waveform averages were calculated after randomly selecting 100 spikes for each MSN and 1000 spikes for each FSI to normalize for firing rate differences. Shading indicates SEM.

DISCUSSION

Our findings suggest that dopamine depletion in the sensorimotor striatum does not abolish normal oscillatory patterns nor create oscillations at new frequencies, but instead, amplifies intrinsically occurring oscillations in a selective manner. Enhancement of low frequency (< 55 Hz) oscillations occurred only in particular task-modulated frequency bands, only during periods in which those oscillations were actively task-modulated, and only after the behavior had become highly trained. By contrast, high frequency (> 65 Hz) oscillations were statically and uniformly diminished with dopamine depletion. L-DOPA therapy normalized the oscillations in most frequency ranges but did not reduce the elevated low gamma oscillations, suggesting that certain

aspects of compromised circuit function may not be correctable by L-DOPA treatment. Further, the effects of dopamine depletion on spike-LFP coupling were also highly frequency-band and cell-type specific. With these findings, we begin building a link between the abnormal LFP oscillations found in Parkinson's disease and the field of research characterizing the behavioral relevance of oscillatory activity and dopamine to behavior in normal organisms.

Dopamine Depletion Amplifies Low Frequency LFP Oscillations Only During the Performance of a Well Learned Task.

Our findings support those of other electrophysiological studies, in human Parkinson's patients and in rodent models, in which low frequency oscillatory activity in basal ganglia circuits were increased by dopamine depletion (Brown, 2003, Eusebio and Brown, 2007, Fuentes et al., 2009, Dejean et al., 2012). However, we found that this effect is strongest when performing a well-trained task and that the increases are most evident during the task-times in which specific oscillations are active. These findings accord with studies in which the effects of dopamine depletion on LFP oscillations in the substantia nigra pars reticulata (SNpr) and motor cortex were different during inattentive rest and treadmill walking, demonstrating state dependency of the effect of dopamine depletion (Avila et al., 2010, Brazhnik et al., 2012).

We extend previous observations by addressing the effect of dopamine depletion on LFP oscillations occurring during a complex learned task involving action initiation, decision-making, and reward. We demonstrate that these effects are highly dependent on learning and on task-time. We suggest that these remarkably selective effects reflect previously described synaptic plasticity of dopamine-dependent responses during learning, and the desynchronizing influence that dopamine appears to have on the basal ganglia networks (Calabresi et al., 2007, Costa, 2007, Burkhardt et al., 2009, Cruz et al., 2011, Gittis et al., 2011). The lack of this desynchronizing influence, by this view, could be most evident late in learning and during task performance times at which oscillations in the sensorimotor striatum are strongest. This task selective effect of dopamine loss may be related to the proposed role of dopamine in behavioral effort and

incentive salience (Berridge, 2007, Salamone et al., 2012) whereby the most pronounced effects of dopamine loss could be seen in situations where learned salient stimuli are present, motivation is high, and effort is required. The lack of significant effect of the depletion during the baseline period could thus reflect the level of engagement of these circuits, as well as the levels of dopamine depletion achieved by our local intrastriatal intervention, or to the fact that these were local, not global depletions.

L-DOPA Fails to Normalize Elevated Low Gamma Oscillations Which May Be Specifically Linked to FSI Firing in the Dopamine-Depleted Dorsolateral Striatum.

L-DOPA treatment almost completely restored normal LFP oscillatory power in all of the frequency bands analyzed, with the clear exception of the low gamma (40-53 Hz) oscillation. Notably, FSI, but not MSN, spike coupling to low gamma oscillations, as measured by the strength of spike-LFP coupling, percent of modulated units, and amplitude of STWAs, was enhanced by dopamine depletion. Moreover, the low gamma oscillation was the only prominent rhythm that was elevated in early learning stages, indicating that network alterations affecting the oscillations likely occurred during the 5 weeks from the 6-ODHA injection to the initiation of recording. Dopamine depletion has been found to increase dendritic arborization of FSIs, a structural change which may not be readily reversed by acute L-DOPA treatment (Gittis et al., 2011). We suggest that remodeling in FSI circuits may produce chronic, treatment-resistant alterations in striatal network function which may result in elevated synchrony in the low gamma range.

Dopamine Depletion Alters the Phase Relationships between Spiking and LFPs.

With the notable exception of the low gamma range, the degree of spike-LFP coupling for MSNs and FSIs was not strongly affected by the dopamine depletion after we controlled for the effects of increased firing rate in our measures. Surprisingly, however, we found a broadband shift in the preferred phase of spiking of the MSNs and FSIs on the dopamine-depleted side. We favor the

possibility that this is a small forward shift (< 90º) in the preferred phase (the spikes occur slightly later with respect to the LFP) rather than the alternative of a large backwards shift (> 270º). The causal relationship between spiking and LFP oscillations is not yet clear (Buzsaki et al., 2012), so we cannot ascertain whether this change may represent a shift in spike timing with respect to equivalent LFP signals or a shift in the LFPs themselves. L-DOPA treatment did not reverse this phase shift, despite its reversal of the LFP power differences in the same frequency bands. These findings suggest, along with previous work (Parr-Brownlie et al., 2007, Burkhardt et al., 2009), that the effects of dopamine depletion and L-DOPA on spike-LFP coupling are not a direct reflection of their effects on LFP power.

Dopamine Depletion in Dorsolateral Striatum Affects a Broad Range of Behaviorally Relevant LFP Oscillations in a Dynamic, Task, and Learning Dependent Manner.

We explored the effect of dopamine depletion during the acquisition of a conditional T-maze task with the goal of assessing the full range of changes in striatal LFP oscillations brought about by the loss of dopamine. The effects we found were widespread, and all task-related LFP oscillations were affected in a dynamic manner. We conclude that the loss of dopamine has widespread effects on LFP oscillations beyond the prominent increase in beta range oscillations and decrease in high gamma power as extensively studied in the subthalamic nucleus, globus pallidus, and substantia nigra in animal models of Parkinson's disease (Brown, 2003, Avila et al., 2010, Jenkinson and Brown, 2011). The unifying pattern we observe is an exaggeration of the response profile of task-related LFP oscillations that exist in the normal system. These changes could reflect temporally restricted increases in network synchrony or increased responsiveness to task modulated inputs. The striatum, as a primary source of input to downstream regions in basal ganglia including subthalamic nucleus, globus pallidus, and substantia nigra, is well placed to induce and propagate such oscillations, and has been proposed as a potential source of the abnormal oscillations in Parkinson's disease (Kumar et al., 2011, McCarthy et al., 2011). Finally, we find that L-DOPA therapy normalizes these oscillations with the exception of the gamma oscillation, which could reflect long-term structural changes that occur in the fast spiking

interneurons (Gittis et al., 2011). The failure of L-DOPA to normalize the low gamma oscillations makes this oscillation an important potential target for Parkinson's disease treatments and early interventions to prevent potentially irreversible structural changes.

Acknowledgments

This work was supported by NIH/NINDS P50 NS-38372, National Parkinson Foundation, Parkinson Disease Foundation Fellowship, NIMH Graduate Student Fellowship, and Stanley H. and Sheila G. Sydney Fund. We thank Christine Keller-McGandy and Henry F. Hall for their help.

SUPPLEMENTARY FIGURES

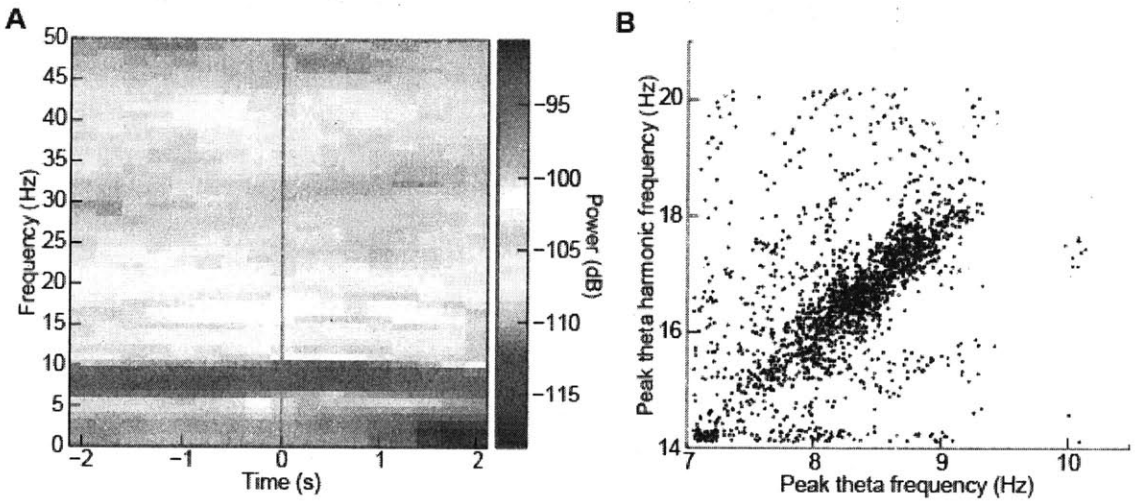

Fig. S1. The theta oscillation produces a harmonic at 15-19 Hz. (*A*) A session average spectrogram of a single recorded channel aligned around the turn beginning event. A strong theta oscillation is present and accompanied by a weaker harmonic. (*B*) Peak frequency in the theta range (7-10 Hz) plotted against peak frequency at double that range (14-20 Hz) taken in non-overlapping 1 s windows for activity recorded on 11 electrode channels during a single session. The peak frequencies are correlated (Pearson's coefficient $r = 0.52$, $P = 0$) and fall along the line theta frequency = 2 X harmonic frequency.

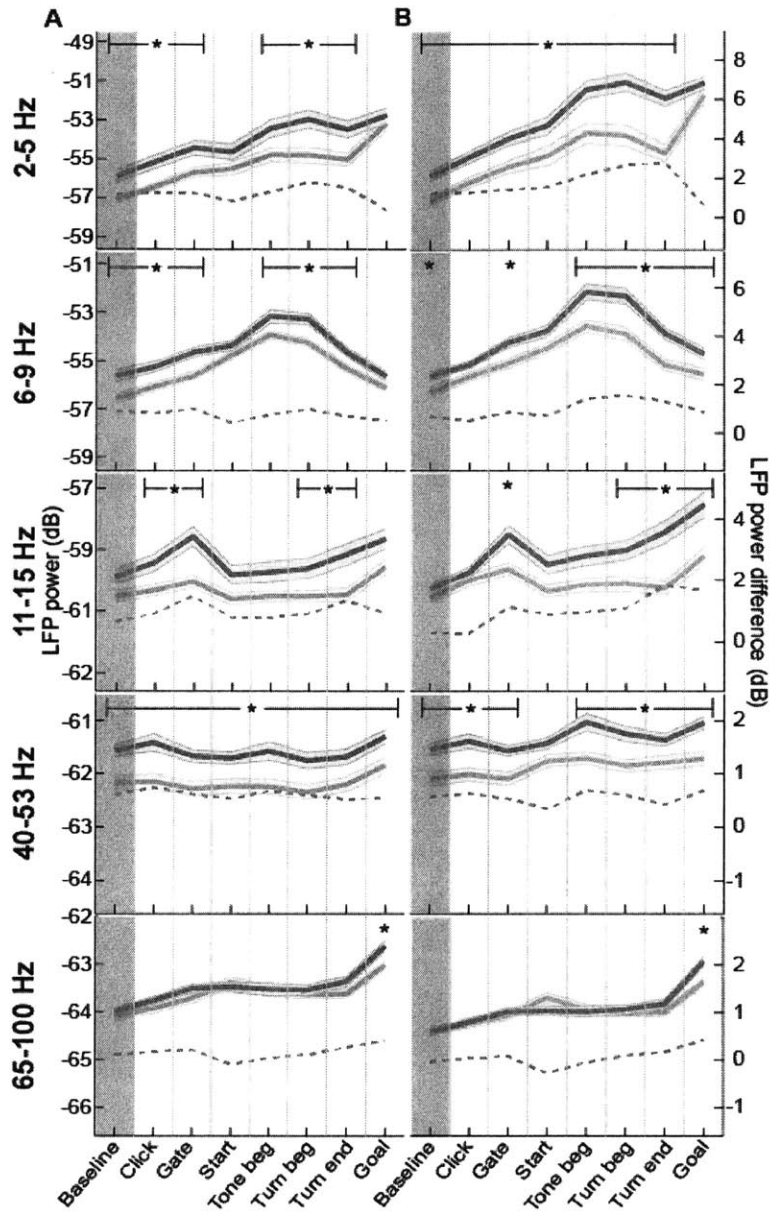

Fig. S2. LFP oscillations strengthen with learning. (*A*) Session average power of oscillations in delta, theta, beta and gamma frequency bands in the intact DLS across trial time. Each plot shows average power during ±500 ms windows centered at successive task events in sessions before acquisition criterion was reached (light blue, n = 65 sessions from 10 animals) and sessions after acquisition criterion was reached (medium blue, n = 90 sessions from same 10 animals). (*B*) Oscillation power in the same sessions recorded in the dopamine-depleted DLS before (light red) and after (medium red) acquisition criterion. Error bars indicate SEM. *$P < 0.05$ (ANOVA).

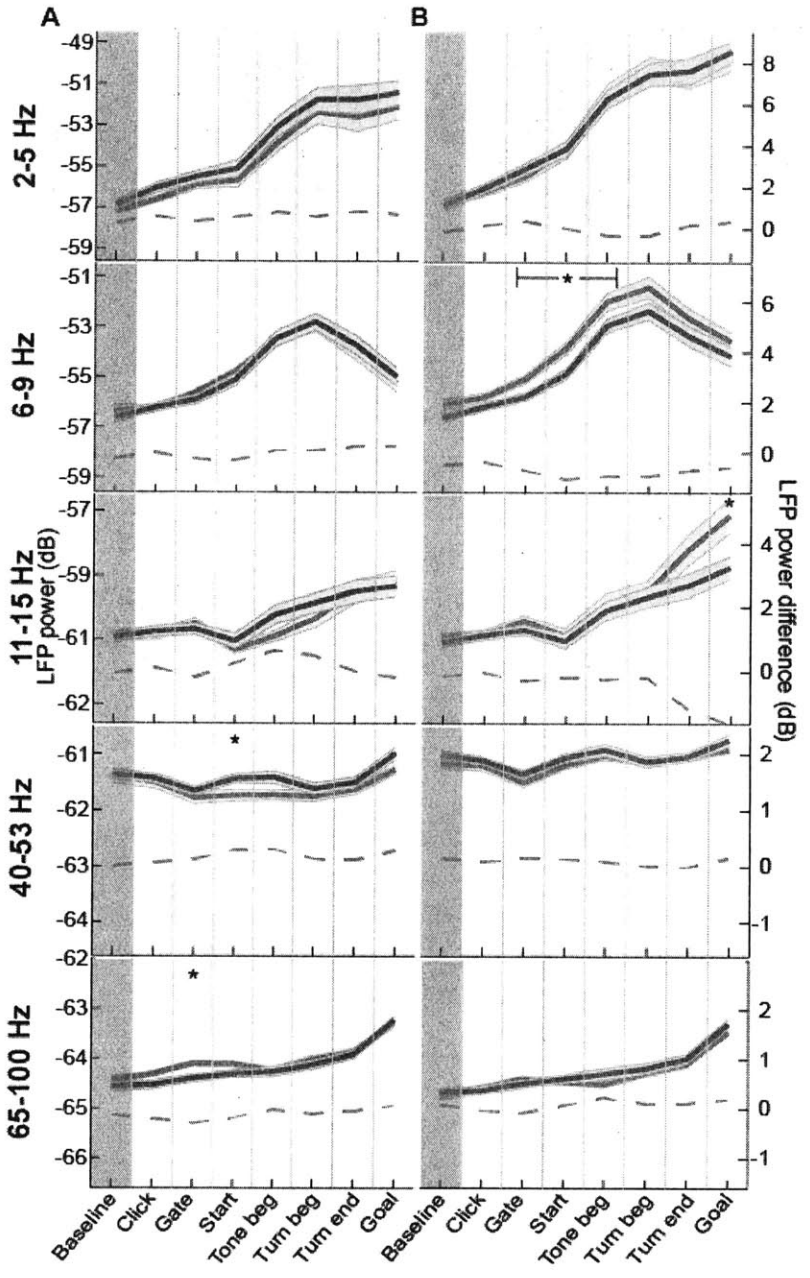

Fig. S3. L-DOPA reduces theta and beta oscillations in the dopamine-depleted DLS, but not in the intact DLS. (*A*) Session average power of relevant frequency bands in the intact DLS. Data from sessions before L-DOPA treatment (medium blue, n = 61 sessions from 10 animals) and from sessions with daily L-DOPA treatment (dark blue, n = 70 sessions from same 10 animals) are shown as in Fig. S2. (*B*) Oscillation power in the same sessions in the dopamine depleted DLS before (medium red) and during (dark red) L-DOPA treatment. Error bars indicate SEM. *$P < 0.05$ (ANOVA).

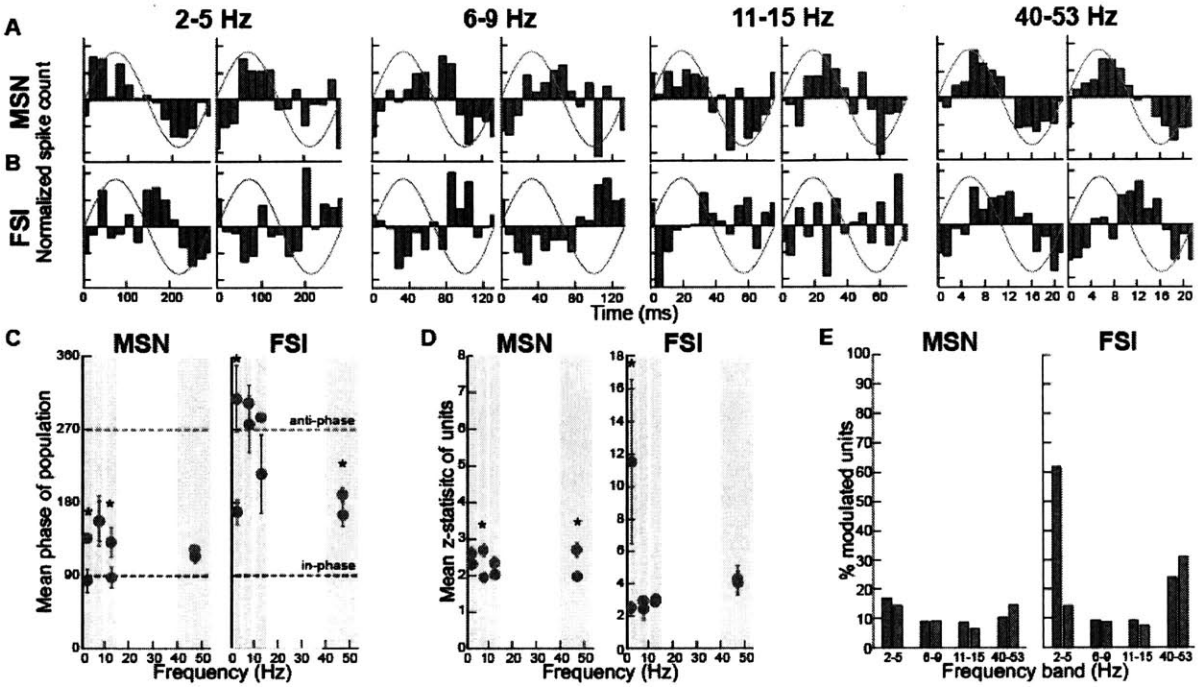

Fig. S4. Dopamine depletion affects spike-LFP coupling in the baseline period. (*A* and *B*) Normalized population spike phase histograms to LFP waves filtered in relevant frequencies in a 1.8 sec. window in the baseline period in overtraining stages. MSN population spike histograms are shown in *A* for MSNs in the intact DLS (blue, all spikes from 662 units) and those in the dopamine-depleted DLS (red, all spikes from 712 units). The same is shown in *B* for FSIs for the intact DLS (blue, from 66 units) and depleted DLS (red, from 82 units). (*C*) Mean phase of population spiking of MSNs and FSIs. Error bars indicate 75% confidence interval. $*P < 0.05$ (ANOVA for circular distributions). (*D*) (*Left*) the average Z-statistic (strength of phase locking) from Rayleigh's test of individual MSNs in the intact (blue) and dopamine-depleted (red) DLS. Only units with at least 60 spikes in analyzed window were used and were normalized to 60 spikes each ($n = 253$ units in intact DLS, $n = 322$ units in depleted DLS). (*Right*) FSI average Z-statistics for FSI units similarly normalized to 600 spikes each ($n = 54$ units and 80 units, respectively). Error bars indicate SEM. $*P < 0.05$ (Mann-Whitney test). (*E*) Proportions of significantly phase modulated ($P < 0.01$, Rayleigh's test) MSNs and FSIs of units in *D* in each of the frequency bands.

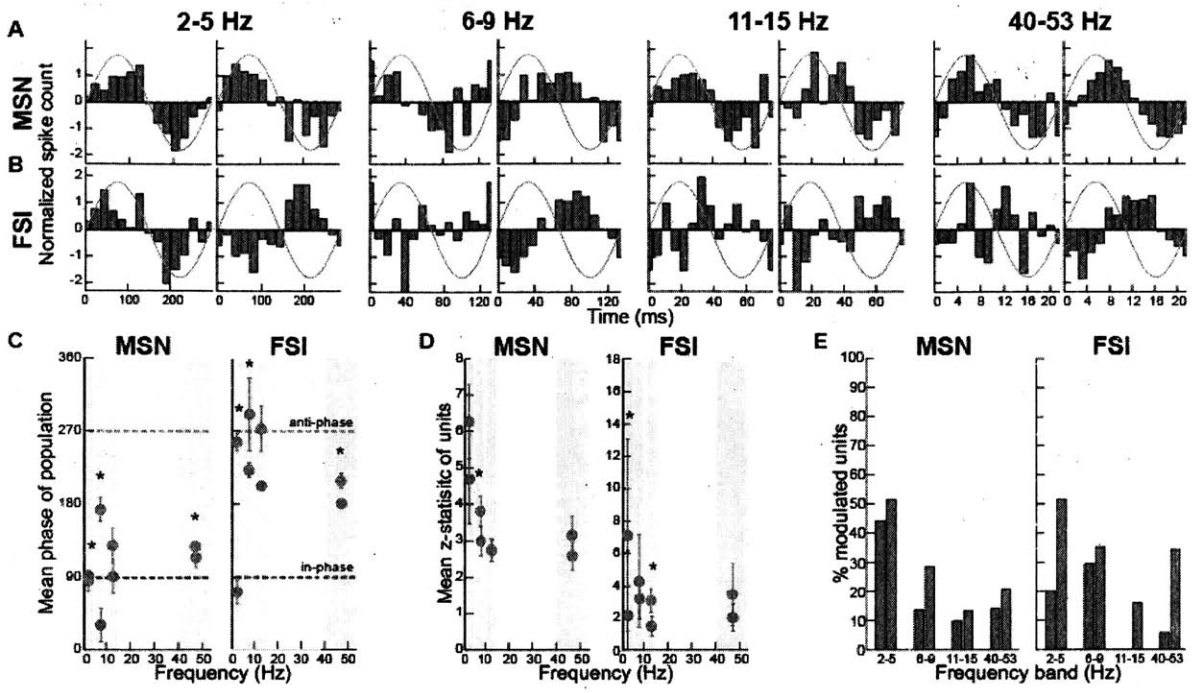

Fig. S5. L-DOPA treatment does not reverse effects of dopamine depletion on spike-LFP coupling. (*A* and *B*) Normalized population spike phase histograms to LFP waves filtered in relevant frequencies in a ±1.5sec. window around cue onset in sessions with L-DOPA treatment. MSN population spike histograms are shown in *A* for MSNs in the intact DLS (blue, all spikes from 248 units) and those in the dopamine-depleted DLS (red, all spikes from 402 units). The same is shown in *B* for FSIs for the intact DLS (blue, from 21 units) and depleted DLS (red, from 39 units). (*C*) Mean phase of population spiking of MSNs and FSIs. Error bars indicate 75% confidence interval. **P* < 0.05 (ANOVA for circular distributions). (*D*) (*Left*) the average Z-statistic (strength of phase locking) from Rayleigh's test of individual MSNs in the intact (blue) and dopamine-depleted (red) DLS. Only units with at least 100 spikes in analyzed window were used and were normalized to 100 spikes each (n = 113 units in intact DLS, n = 158 units in depleted DLS). (*Right*) FSI average Z-statistics for FSI units similarly normalized to 1000 spikes each (n = 17 units and 39 units, respectively). Error bars indicate SEM. **P* < 0.05 (Mann-Whitney test). (*E*) Proportions of significantly phase modulated (*P* < 0.01, Rayleigh's test) MSNs and FSIs of units in *D* in each of the frequency bands.

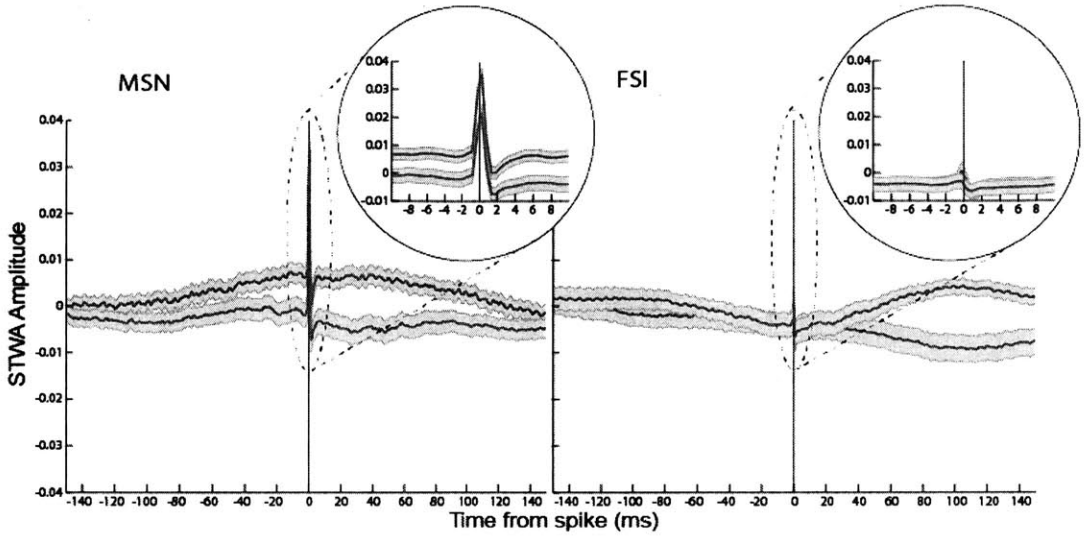

Fig. S6. Spike triggered waveform averages of MSN (*Left*) and FSI (*Right*) spikes recorded in the dopamine-depleted DLS (red) and the intact DLS (blue). The delta oscillation is most prominent in the STWA in both MSNs and FSIs and is in phase with MSN spiking. A quick deflection also occurs in the LFP with MSN spikes and to a lesser extent FSI spikes with a latency of 1 ms.

MATERIALS AND METHODS

Procedures were approved by the MIT Committee on Animal Care and accorded with the National Research Council's Guide for the Care and Use of Laboratory Animals. Male Sprague-Dawley rats (n = 16) received injections of 6-OHDA in the dorsolateral striatum unilaterally. Four weeks later, a recording drive was implanted, targeting the dorsolateral striatum bilaterally, each side with 6 tetrodes. Rats were then trained on a T-maze task (~40 trials per session), in which the rats were instructed to turn left or right in response to a tone to receive a chocolate reward at the end of the indicated end-arm. Rats were trained to a criterion of 72.5% correct trials and then overtrained for 10 consecutive days. They were then tested in 10 additional sessions with daily systemic L-DOPA administration. Spike and LFP activity was recorded throughout training. LFP spectral analysis was conducted using Chronux algorithms (http://chronux.org), the MATLAB CircStat Toolbox, and in-house MATLAB software. Detailed methods are available in Supplementary Methods.

Animals, Dopamine Depletion and Tetrode Implantation.

Experimental procedures were approved by the committee on Animal Care of MIT and follow the National Research Council's Guide for the Care and Use of Laboratory Animals. We housed sixteen male adult Sprague-Dawley rats (310-420g) in a reverse light dark cycle (lights on 9pm-9am). We conducted training and recording during the active dark cycle of the rats. After surgical implantation of the recording drive, the rats were transferred from paired cage housing to single cage housing and were places on food restriction to reach 90% of their free feeding weight before beginning of training.

To induce depletion of dopamine terminals in dorsolateral striatum unilaterally, rats were anesthetized with a mixture of ketamine (100 mg/kg) and xylazine (10 mg/kg) and 6-hydroxydopamine (10μg/3μl per site) was injected at a rate of 1 μl per minute in the right dorsolateral striatum at two locations to cover the extent of the dorsolateral striatum (AP = 1 mm; ML = −3 mm; DV = −4.5 mm and AP = 0.2 mm; ML = −3.6 mm; DV = −5 mm).

Four weeks after the injection surgery, rats were anesthetized again in the same manner and were implanted with a recording drive containing 6 tetrodes targeting the dopamine depleted striatum and 6 tetrodes targeting the contralateral intact striatum (approx. AP = 0.5 mm, ML = ±3.6 mm). The tetrodes were constructed from 4 twisted 12-μm nickel-chromium wires which were heat fused and reinforced by coating with a layer of cyanoacrylate glue resulting in a total diameter of approx. 50 μm. The tetrodes were placed into independently movable microdrives. During the implantation surgery, the recording drive was secured to the skull with dental cement and small screws attached to the skull surface. A wire attached to one of these screws served as the animal ground for the recording drive. The rats were allowed to recover during a week period after the implantation surgery and during this time tetrodes were lowered in small increments to the target depth (3.5-4 mm). After this period, tetrodes were moved only in small increments of less than 100 μm as needed to maintain high quality recordings.

Behavioral Training and Data Collection.

Rats were trained in a custom T-maze (Fig. 1*B*) constructed of 2 raised polycarbonate pieces painted in anti-static black paint joined in a T-shape. The long arm of the maze was 7.6 x 121.9 cm, the short arms were 7.6 x 73.7 cm, and the maze height was 22.9 cm. 40.6 cm high wood walls painted in black anti-static paint surrounded the maze at a distance of 11-20 cm. A gate next to the starting block of the maze could be manually raised or lowered to allow the animal to run down the maze. Photobeam units were placed in the walls of the maze to track the position of the animal to provide input to the Med-PC behavioral control system (Med Associates). Movements and head direction of the rats were also monitored with a video tracker (Neuralynx) by tracking the location of red and green LEDs attached to the implanted drive and recorded by an overhead CCD camera in video captured at a 30 Hz rate.

Each rat was acclimated to the T-maze and the chocolate sprinkle reward in five 30 min sessions in the week prior to recording drive implantation by allowing them to freely explore the maze and eat the chocolate sprinkles. In the last two or three of these sessions the chocolate sprinkles were limited to the goal sites. In the last one or two of these sessions, rats experienced up to 10 trials in which they waited in the start location while both goal sites held chocolate, and one of

two cue tones (1 kHz or 8 kHz) were played in the middle of the run to allow the rats to acclimate to the sounds without building any cue-tone associations.

One week after recording drive implantation, training on the T-maze task began (Fig 1*B*). In each trial, rats waited on the starting block until a warning click was played and the gate was lowered. The rats then ran down the long arm of the maze and after they broke a photobeam placed halfway down the long arm an auditory (1 kHz or 8 kHz pure tone) began to play and stayed on until one of the end goal site photobeams was broken. If the rats turned the correct direction and reached the goal site, they received chocolate sprinkles, and if not they did not receive any reward. They were then guided back to the start position for the next trial. Each session consisted of up to 40 trials and training continued at least 10 sessions after rats initially reached the 72.5% learning criterion ($P < 0.01$, chi square test).

After this initial acquisition training and the following overtraining sessions, 10 rats received 10 further sessions with an i.p. systemic injection of L-DOPA (6 mg/kg) and Benserazide (15 mg/kg) administered 30 min before the start of the session. This low dose of L-DOPA was used to prevent dyskinesia side effects.

Neuronal Recordings.

In each recording session, spike activity and local field potential activity was recorded using a Cheetah Data Acquisition System (Neuralynx). Recording began 2 seconds before the beginning of each trial and stopped 1 second after the goal photobeam was broken before reward consumption to prevent electrical noise from chewing. Spike activity was recorded with a gain of 1000-10000, filtered at 600-6000 Hz, and sampled at 30 or 32 kHz. Spike waveforms were recorded on all four tetrode channels in 1ms windows around the time of a pre-set threshold crossing on any one of the four tetrode channels. LFP activity was recorded with a gain of 1000, filtered at 1–475 Hz, and sampled at 1 kHz continuously from the beginning of recording to the stopping of recording on only one of the channels of each tetrode. LFP channels were referenced to the animal ground attached to a screw on the animal's skull or to panel ground in the room.

Spike data was referenced to one of these references or to a tetrode channel lacking spike activity.

Fast-Scan Cyclic Voltammetry.

Recording drives as those used for electrophysiology were loaded with 8 carbon fiber probes (Clark et al., 2010) for voltametric measurements and 2 tungsten bi-polar electrodes for stimulation of the medial forebrain bundle (MFB) (75 um in diameter, FHC) with independently movable microdrives. Four rats with unilateral 6-ODHA lesions in dorsolateral striatum were implanted with the recording drive as described above along with an Ag/AgCl reference electrode (500 µm diameter, AM Systems) implanted in posterior cortex (AP = −0.3 mm, ML = 2.1 mm, DV = 0.5 mm). The probes were then lowered to the same dorsolateral striatal coordinates (AP = +0.5 mm, ML = ±3.6 mm, DV = 3.5-4.0 mm) and the stimulation electrodes to the MFB (AP = −4.9 mm, ML = ±1.3 mm, DV = 6.5-7.5 mm). A triangular voltage waveform (−0.4 V to 1.3 V relative to the Ag/AgCl reference) was applied to the carbon fiber probes every 100 ms until stable current measurements were reached at the dopamine oxidation potential (~0.6 V). Stimulation parameters were controlled and data collected with the use of two PCI data acquisition cards (National Instruments) and software written in LabView. Currents generated by dopamine oxidation at the surface of the carbon fiber electrode were measured in the intact and dopamine depleted striatum after bi-phasic stimulation (60 Hz, 2 ms pulse width, 200 µA) was applied to the electrodes in MFB. The currents measured at the dopamine oxidation potential were converted to dopamine concentration using electrode calibrations obtained *in vitro* post-stimulation. After measurements were made in the intact and depleted hemispheres, L-DOPA (6 mg/kg + Benserazide 15 mg/kg) was administered i.p. and the stimulations were repeated.

Behavioral Data Analysis.

Performance was measured by response accuracy (number of correct trials / total trials), run time (time from locomotion start to goal reaching), and reaction time (time from gate opening to locomotion start). How these measures change across training was assessed by using ANOVAs. For

the purposes of aggregating behavioral data from rats learning the task at different rates, training sessions were grouped by learning stages where stages 1-4 were defined as the acquisition stages and stages 5-10 were defined as the overtraining stages. Stages 1 and 2 were days 1 and 2 of training. Stage 3 was the first day the rats performed > 60% correct trials and stage 4 was the first day they reached > 70% correct. Each of the stages 5-10 was composed of pairs of consecutive sessions with > 72.5% correct (learning criterion) performance. For the purposes of aggregating LFP recordings across rats, all sessions prior to criterion reaching were grouped as acquisition sessions and all sessions after the beginning of consistent > 72.5% correct performance (two or more days) were grouped as overtraining sessions. All following sessions with L-DOPA treatment were grouped into L-DOPA treatment sessions.

LFP Data Analysis.

LFP activity was analyzed by computing spectral power distributions using Chronux algorithms (http://chronux.org) and in-house MATLAB software. Spectral power was computed using a single taper in a 1 second window in the pre-task baseline period and in consecutive ±500 ms peri-event (1 s) windows centered on each recorded task event. LFP power was compared in the depleted and intact DLS by taking session averages of LFP power in each session containing recordings from both the intact and depleted DLS. In frequency spectrum plots, log power (dB) was normalized for the "1/f factor" by subtracting 10*log(1/frequency).

Session averages of LFP power were determined in the frequency bands 2-5 Hz, 6-9 Hz, 11-15 Hz, 40-53 Hz, and 65-100 Hz. These ranges were chosen by their prominence in the LFP spectra. Power in each band was determined by taking a weighted average of the power in the frequencies (1 Hz resolution) within the frequency band with the most weight given to the center of the band and the least weight to the outer frequencies in the band (for example, for the 11-15 Hz band, the weights for 11, 12, 13, 14, 15 Hz were [1, 2, 3, 2, 1]/9). In the high frequency range (65-100 Hz), no single frequency was prominent, so all frequencies were weighted equally. The effect of dopamine depletion on the LFP power was then assessed in each of these frequency bands by performing a one-way ANOVA on the session averages of simultaneously recorded LFPs in the intact and dopamine depleted DLS. The effect of learning stage and L-DOPA treatment

within the intact and depleted hemispheres was assessed by computing a one-way ANOVA on the session averages from these groups of sessions. In this case, due to the differing numbers of sessions across groups, we reduced the number of sessions in the larger group by taking a random set equal to the number of sessions in the smaller group and computing the ANOVA. This was done iteratively 10 times and the average *P* value was taken from the 10 computations.

Spike-LFP Coupling Analysis.

Spike waveforms were clustered into different units manually using Offline Sorter (Plexon). Units were rated based on cluster separation and the absence of spikes falling within the refractory period. Units with acceptable rating were separated into putative MSNs, FSIs, and TANs based on firing rate, spike waveform shape, interspike interval distributions, and peri-event raster plots (Barnes et al., 2005). The spike times belonging to each of the units of MSNs and FSIs were used in the spike-LFP coupling analyses.

To assess spike-LFP relationships, LFP waves were band-pass filtered in the same frequency bands as in the analyses described above and spike phase histograms were constructed with respect to LFP waves in each frequency band within a 3 s in-task window centered around cue onset. It was assessed whether the population spike phase distributions were non-uniform by performing Rayleigh's test on the population spiking. Whether the preferred phases of the population spiking in the depleted DLS and intact DLS were different was then tested by using ANOVA for circular distributions. The strength of the phase locking to different frequencies of oscillations was then calculated by using the average Z-statistic of individual units. For this analysis, only MSN units with at least 100 total spikes within the analyzed window in the session were used to ensure high quality estimates of spike phase coupling. To control for firing rate effects on the measure of the Z-statistic, 100 randomly chosen spikes from each unit were chosen to compute the Z-statistic. This was done 100 times for each unit and the average Z-statistic of the 100 computations was used for each unit. For FSIs, only units with 1000 or more spikes in the analyzed windows were used and the number of spikes was similarly normalized to 1000 for each computation. To compare the strength of phase locking in the intact and dopamine depleted DLS, the Mann-Whitney test was performed on the Z-statistics of units from the two hemispheres due to the

non-normal distribution of the Z-statistics. These same analyses were performed for spiking in the baseline period in the overtraining sessions. In this case, due to the shorter duration of the baseline period, a 1.8 s window prior to trial start was used and the number of MSN spikes was normalized to 60 and the number of FSI spikes to 600.

For constructing spike triggered waveform averages (STWAs), LFP waves in the same 3 s in-task window as above were filtered and the maximum amplitude of each filtered wave was normalized to 1 to control for differences in LFP amplitude across channels and hemispheres. The STWA for each unit was then computed by taking the average of the normalized filtered waves around spike events. As in the prior analysis, only MSNs with 100 or more spikes within the analysis window were used and a random set of 100 spikes from each unit was used to control for sampling effects between high firing and low firing units. Similarly, 1000 spikes from each FSI were used. For computing the raw STWA, the same analysis was performed without band-pass filtering the LFP waves.

Histology.

After the recordings were completed, lesion marks were made to mark the location of the tetrode tips by anesthetizing the rats with sodium pentobarbital solution (ca. 40-50 mg/kg) and passing current through each tetrode (25 μA, 10 s). After two days, rats were perfused and the brain fixed and removed by applying deep anesthesia with a lethal dose of sodium pentobarbital (ca. 100-145 mg/kg) and transcardially perfusing the rats with 4% paraformaldehyde in 0.1M KNaPO4 buffer. Brains were cut into 30 μm coronal sections. Every other section was stained with Cresyl Violet for visualizing tetrode marks and the alternate sections were stained for tyrosine hydroxylase (TH) to visualize the extent of the lesion. Units recorded on tetrodes with tips that were found to be outside of the lesion in the depleted dorsolateral striatum or not in the dorsolateral striatum in either hemisphere were disregarded.

References

Acuna DE, Wymbs NF, Reynolds CA, Picard N, Turner RS, Strick PL, Grafton ST, Kording KP (2014) Multifaceted aspects of chunking enable robust algorithms. Journal of neurophysiology 112:1849-1856.

Akintunde A, Buxton DF (1992) Origins and collateralization of corticospinal, corticopontine, corticorubral and corticostriatal tracts: a multiple retrograde fluorescent tracing study. Brain Res 586:208-218.

Albin RL, Young AB, Penney JB (1989) The functional anatomy of basal ganglia disorders. Trends in neurosciences 12:366-375.

Aldridge JW, Berridge KC (1998) Coding of serial order by neostriatal neurons: a "natural action" approach to movement sequence. J Neurosci 18:2777-2787.

Aldridge JW, Berridge KC, Rosen AR (2004) Basal ganglia neural mechanisms of natural movement sequences. Canadian journal of physiology and pharmacology 82:732-739.

Alexander GE, DeLong MR, Strick PL (1986) Parallel organization of functionally segregated circuits linking basal ganglia and cortex. Annual review of neuroscience 9:357-381.

Amemori K, Gibb LG, Graybiel AM (2011) Shifting responsibly: the importance of striatal modularity to reinforcement learning in uncertain environments. Front Hum Neurosci 5:47.

Anderson CT, Sheets PL, Kiritani T, Shepherd GM (2010) Sublayer-specific microcircuits of corticospinal and corticostriatal neurons in motor cortex. Nature neuroscience 13:739-744.

Atallah HE, Lopez-Paniagua D, Rudy JW, O'Reilly RC (2007) Separate neural substrates for skill learning and performance in the ventral and dorsal striatum. Nature neuroscience 10:126-131.

Avila I, Parr-Brownlie LC, Brazhnik E, Castaneda E, Bergstrom DA, Walters JR (2010) Beta frequency synchronization in basal ganglia output during rest and walk in a hemiparkinsonian rat. Exp Neurol 221:307-319.

Bailey KR, Mair RG (2006) The role of striatum in initiation and execution of learned action sequences in rats. J Neurosci 26:1016-1025.

Bailey KR, Mair RG (2007) Effects of frontal cortex lesions on action sequence learning in the rat. Eur J Neurosci 25:2905-2915.

Balleine BW, O'Doherty JP (2010) Human and rodent homologies in action control: corticostriatal determinants of goal-directed and habitual action. Neuropsychopharmacology : official publication of the American College of Neuropsychopharmacology 35:48-69.

Barnes TD, Kubota Y, Hu D, Jin DZ, Graybiel AM (2005) Activity of striatal neurons reflects dynamic encoding and recoding of procedural memories. Nature 437:1158-1161.

Barnes TD, Mao JB, Hu D, Kubota Y, Dreyer AA, Stamoulis C, Brown EN, Graybiel AM (2011) Advance cueing produces enhanced action-boundary patterns of spike activity in the sensorimotor striatum. Journal of neurophysiology 105:1861-1878.

Barter JW, Li S, Lu D, Bartholomew RA, Rossi MA, Shoemaker CT, Salas-Meza D, Gaidis E, Yin HH (2015) Beyond reward prediction errors: the role of dopamine in movement kinematics. Frontiers in integrative neuroscience 9:39.

Belin D, Jonkman S, Dickinson A, Robbins TW, Everitt BJ (2009) Parallel and interactive learning processes within the basal ganglia: relevance for the understanding of addiction. Behavioural brain research 199:89-102.

Berridge KC (2007) The debate over dopamine's role in reward: the case for incentive salience. Psychopharmacology (Berl) 191:391-431.

Berridge KC, Aldridge JW, Houchard KR, Zhuang X (2005) Sequential super-stereotypy of an instinctive fixed action pattern in hyper-dopaminergic mutant mice: a model of obsessive compulsive disorder and Tourette's. BMC Biol 3:4.

Berridge KC, Whishaw IQ (1992) Cortex, striatum and cerebellum: control of serial order in a grooming sequence. Exp Brain Res 90:275-290.

Bortoff GA, Strick PL (1993) Corticospinal terminations in two new-world primates: further evidence that corticomotoneuronal connections provide part of the neural substrate for manual dexterity. J Neurosci 13:5105-5118.

Boulougouris V, Dalley JW, Robbins TW (2007) Effects of orbitofrontal, infralimbic and prelimbic cortical lesions on serial spatial reversal learning in the rat. Behavioural brain research 179:219-228.

Brazhnik E, Cruz AV, Avila I, Wahba MI, Novikov N, Ilieva NM, McCoy AJ, Gerber C, Walters JR (2012) State-dependent spike and local field synchronization between motor cortex and substantia nigra in hemiparkinsonian rats. J Neurosci 32:7869-7880.

Brown LL, Smith DM, Goldbloom LM (1998) Organizing principles of cortical integration in the rat neostriatum: corticostriate map of the body surface is an ordered lattice of curved laminae and radial points. J Comp Neurol 392:468-488.

Brown P (2003) Oscillatory nature of human basal ganglia activity: relationship to the pathophysiology of Parkinson's disease. Mov Disord 18:357-363.

Burkhardt JM, Jin X, Costa RM (2009) Dissociable effects of dopamine on neuronal firing rate and synchrony in the dorsal striatum. Front Integr Neurosci 3:28.

Buzsaki G, Anastassiou CA, Koch C (2012) The origin of extracellular fields and currents--EEG, ECoG, LFP and spikes. Nat Rev Neurosci 13:407-420.

Buzsaki G, Draguhn A (2004) Neuronal oscillations in cortical networks. Science 304:1926-1929.

Buzsaki G, Wang XJ (2012) Mechanisms of gamma oscillations. Annu Rev Neurosci 35:203-225.

Calabresi P, Picconi B, Tozzi A, Di Filippo M (2007) Dopamine-mediated regulation of corticostriatal synaptic plasticity. Trends Neurosci 30:211-219.

Carelli RM, West MO (1991) Representation of the body by single neurons in the dorsolateral striatum of the awake, unrestrained rat. J Comp Neurol 309:231-249.

Carelli RM, Wolske M, West MO (1997) Loss of lever press-related firing of rat striatal forelimb neurons after repeated sessions in a lever pressing task. J Neurosci 17:1804-1814.

Chen MT, Morales M, Woodward DJ, Hoffer BJ, Janak PH (2001) In vivo extracellular recording of striatal neurons in the awake rat following unilateral 6-hydroxydopamine lesions. Exp Neurol 171:72-83.

Clark JJ, Sandberg SG, Wanat MJ, Gan JO, Horne EA, Hart AS, Akers CA, Parker JG, Willuhn I, Martinez V, Evans SB, Stella N, Phillips PE (2010) Chronic microsensors for longitudinal, subsecond dopamine detection in behaving animals. Nat Methods 7:126-129.

Coffey KR, Nader M, West MO (2016) Single body parts are processed by individual neurons in the mouse dorsolateral striatum. Brain Res 1636:200-207.

Compton DM (2004) Behavior strategy learning in rat: effects of lesions of the dorsal striatum or dorsal hippocampus. Behav Processes 67:335-342.

Costa RM (2007) Plastic corticostriatal circuits for action learning: what's dopamine got to do with it? Ann N Y Acad Sci 1104:172-191.

Crutcher MD, DeLong MR (1984) Single cell studies of the primate putamen. II. Relations to direction of movement and pattern of muscular activity. Exp Brain Res 53:244-258.

Cruz AV, Mallet N, Magill PJ, Brown P, Averbeck BB (2011) Effects of dopamine depletion on information flow between the subthalamic nucleus and external globus pallidus. J Neurophysiol 106:2012-2023.

Cui G, Jun SB, Jin X, Pham MD, Vogel SS, Lovinger DM, Costa RM (2013) Concurrent activation of striatal direct and indirect pathways during action initiation. Nature 494:238-242.

Da Cunha C, Gomez AA, Blaha CD (2012) The role of the basal ganglia in motivated behavior. Rev Neurosci 23:747-767.

Dang MT, Yokoi F, Yin HH, Lovinger DM, Wang Y, Li Y (2006) Disrupted motor learning and long-term synaptic plasticity in mice lacking NMDAR1 in the striatum. Proc Natl Acad Sci U S A 103:15254-15259.

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

Dayan P, Berridge KC (2014) Model-based and model-free Pavlovian reward learning: revaluation, revision, and revelation. Cogn Affect Behav Neurosci 14:473-492.

DeCoteau WE, Thorn C, Gibson DJ, Courtemanche R, Mitra P, Kubota Y, Graybiel AM (2007) Oscillations of local field potentials in the rat dorsal striatum during spontaneous and instructed behaviors. J Neurophysiol 97:3800-3805.

Dejean C, Nadjar A, Le Moine C, Bioulac B, Gross CE, Boraud T (2012) Evolution of the dynamic properties of the cortex-basal ganglia network after dopaminergic depletion in rats. Neurobiol Dis 46:402-413.

DeLong M, Wichmann T (2010) Changing views of basal ganglia circuits and circuit disorders. Clinical EEG and neuroscience : official journal of the EEG and Clinical Neuroscience Society 41:61-67.

DeLong MR (1973) Putamen: activity of single units during slow and rapid arm movements. Science 179:1240-1242.

Dickinson A (1985) Actions and habits: the development of behavioral autonomy. Philos Trans R Soc Lond B Biol Sci 67–78.

Dickinson A, D. J. Nicholas, Christopher D., Adams D (1983) The effect of the instrumental training contingency on susceptibility to reinforcer devaluation. The Quarterly Journal of Experimental Psychology

Dolan RJ, Dayan P (2013) Goals and habits in the brain. Neuron 80:312-325.

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

Ebrahimi A, Pochet R, Roger M (1992) Topographical organization of the projections from physiologically identified areas of the motor cortex to the striatum in the rat. Neurosci Res 14:39-60.

Engel AK, Fries P (2010) Beta-band oscillations--signalling the status quo? Curr Opin Neurobiol 20:156-165.

Eusebio A, Brown P (2007) Oscillatory activity in the basal ganglia. Parkinsonism Relat Disord 13 Suppl 3:S434-436.

Eusebio A, Cagnan H, Brown P (2012) Does suppression of oscillatory synchronisation mediate some of the therapeutic effects of DBS in patients with Parkinson's disease? Front Integr Neurosci 6:47.

Everitt BJ, Robbins TW (2005) Neural systems of reinforcement for drug addiction: from actions to habits to compulsion. Nature neuroscience 8:1481-1489.

Flaherty AW, Graybiel AM (1994) Input-output organization of the sensorimotor striatum in the squirrel monkey. J Neurosci 14:599-610.

Friedman A, Homma D, Gibb LG, Amemori K, Rubin SJ, Hood AS, Riad MH, Graybiel AM (2015) A Corticostriatal Path Targeting Striosomes Controls Decision-Making under Conflict. Cell 161:1320-1333.

Friend DM, Kravitz AV (2014) Working together: basal ganglia pathways in action selection. Trends in neurosciences 37:301-303.

Fuentes R, Petersson P, Siesser WB, Caron MG, Nicolelis MA (2009) Spinal cord stimulation restores locomotion in animal models of Parkinson's disease. Science 323:1578-1582.

Gillan CM, Papmeyer M, Morein-Zamir S, Sahakian BJ, Fineberg NA, Robbins TW, de Wit S (2011) Disruption in the balance between goal-directed behavior and habit learning in obsessive-compulsive disorder. The American journal of psychiatry 168:718-726.

Gittis AH, Hang GB, LaDow ES, Shoenfeld LR, Atallah BV, Finkbeiner S, Kreitzer AC (2011) Rapid target-specific remodeling of fast-spiking inhibitory circuits after loss of dopamine. Neuron 71:858-868.

Glynn G, Ahmad SO (2002) Three-dimensional electrophysiological topography of the rat corticostriatal system. Journal of comparative physiology A, Neuroethology, sensory, neural, and behavioral physiology 188:695-703.

Graybiel AM (1998) The basal ganglia and chunking of action repertoires. Neurobiol Learn Mem 70:119-136.

Graybiel AM (2008) Habits, rituals, and the evaluative brain. Annual review of neuroscience 31:359-387.

Graybiel AM, Grafton ST (2015) The striatum: where skills and habits meet. Cold Spring Harb Perspect Biol 7:a021691.

Hoffer ZS, Alloway KD (2001) Organization of corticostriatal projections from the vibrissal representations in the primary motor and somatosensory cortical areas of rodents. J Comp Neurol 439:87-103.

Holland PC, Straub JJ (1979) Differential effects of two ways of devaluing the unconditioned stimulus after Pavlovian appetitive conditioning. J Exp Psychol Anim Behav Process 5:65-78.

Howe MW, Atallah HE, McCool A, Gibson DJ, Graybiel AM (2011) Habit learning is associated with major shifts in frequencies of oscillatory activity and synchronized spike firing in striatum. Proc Natl Acad Sci U S A 108:16801-16806.

Jenkinson N, Brown P (2011) New insights into the relationship between dopamine, beta oscillations and motor function. Trends Neurosci 34:611-618.

Jenkinson N, Kuhn AA, Brown P (2012) Gamma oscillations in the human basal ganglia. Exp Neurol.

Jin X, Costa RM (2010) Start/stop signals emerge in nigrostriatal circuits during sequence learning. Nature 466:457-462.

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

Johnson A, van der Meer MA, Redish AD (2007) Integrating hippocampus and striatum in decision-making. Current opinion in neurobiology 17:692-697.

Jonkman S, Mar AC, Dickinson A, Robbins TW, Everitt BJ (2009) The rat prelimbic cortex mediates inhibitory response control but not the consolidation of instrumental learning. Behavioral neuroscience 123:875-885.

Kato M, Kimura M (1992) Effects of reversible blockade of basal ganglia on a voluntary arm movement. Journal of neurophysiology 68:1516-1534.

Kawai R, Markman T, Poddar R, Ko R, Fantana AL, Dhawale AK, Kampff AR, Olveczky BP (2015) Motor cortex is required for learning but not for executing a motor skill. Neuron 86:800-812.

Kennerley SW, Sakai K, Rushworth MF (2004) Organization of action sequences and the role of the pre-SMA. Journal of neurophysiology 91:978-993.

Keramati M, Dezfouli A, Piray P (2011) Speed/accuracy trade-off between the habitual and the goal-directed processes. PLoS computational biology 7:e1002055.

Kim N, Barter JW, Sukharnikova T, Yin HH (2014) Striatal firing rate reflects head movement velocity. Eur J Neurosci 40:3481-3490.

Kimura M (1990) Behaviorally contingent property of movement-related activity of the primate putamen. Journal of neurophysiology 63:1277-1296.

Kleim JA, Hogg TM, VandenBerg PM, Cooper NR, Bruneau R, Remple M (2004) Cortical synaptogenesis and motor map reorganization occur during late, but not early, phase of motor skill learning. J Neurosci 24:628-633.

Koch I, Hoffmann J (2000) Patterns, chunks, and hierarchies in serial reaction-time tasks. Psychol Res 63:22-35.

Koob GF, Volkow ND (2010) Neurocircuitry of addiction. Neuropsychopharmacology : official publication of the American College of Neuropsychopharmacology 35:217-238.

Kumar A, Cardanobile S, Rotter S, Aertsen A (2011) The role of inhibition in generating and controlling Parkinson's disease oscillations in the Basal Ganglia. Front Syst Neurosci 5:86.

Lashley KS (1951) The problem of serial order in behavior. Cerebral mechanisms in behavior pp. 112–131.

Leckman JF, Riddle MA (2000) Tourette's syndrome: when habit-forming systems form habits of their own? Neuron 28:349-354.

Ledia Hernandez YK, Dan Hu, Mark Howe, Nune Lemaire, and Ann Graybiel (2012) Dopamine Depletion and L-DOPA Therapy Affect Learning-Related Firing Dynamics of Striatal Neurons. In: submitted.

Lemon RN (2008) Descending pathways in motor control. Annual review of neuroscience 31:195-218.

Liles SL (1985) Activity of neurons in putamen during active and passive movements of wrist. Journal of neurophysiology 53:217-236.

McCarthy MM, Moore-Kochlacs C, Gu X, Boyden ES, Han X, Kopell N (2011) Striatal origin of the pathologic beta oscillations in Parkinson's disease. Proc Natl Acad Sci U S A 108:11620-11625.

McDonald RJ, White NM (1994) Parallel information processing in the water maze: evidence for independent memory systems involving dorsal striatum and hippocampus. Behav Neural Biol 61:260-270.

McGeorge AJ, Faull RL (1989) The organization of the projection from the cerebral cortex to the striatum in the rat. Neuroscience 29:503-537.

Medina L, Reiner A (1995) Neurotransmitter organization and connectivity of the basal ganglia in vertebrates: implications for the evolution of basal ganglia. Brain Behav Evol 46:235-258.

Miltenberger RG, Fuqua RW, Woods DW (1998) Applying behavior analysis to clinical problems: review and analysis of habit reversal. J Appl Behav Anal 31:447-469.

Mink JW, Thach WT (1993) Basal ganglia intrinsic circuits and their role in behavior. Current opinion in neurobiology 3:950-957.

Mittler T, Cho J, Peoples LL, West MO (1994) Representation of the body in the lateral striatum of the freely moving rat: single neurons related to licking. Exp Brain Res 98:163-167.

Oldenburg IA, Sabatini BL (2015) Antagonistic but Not Symmetric Regulation of Primary Motor Cortex by Basal Ganglia Direct and Indirect Pathways. Neuron 86:1174-1181.

Oorschot DE (1996) Total number of neurons in the neostriatal, pallidal, subthalamic, and substantia nigral nuclei of the rat basal ganglia: a stereological study using the cavalieri and optical disector methods. J Comp Neurol 366:580-599.

Otchy TM, Wolff SB, Rhee JY, Pehlevan C, Kawai R, Kempf A, Gobes SM, Olveczky BP (2015) Acute off-target effects of neural circuit manipulations. Nature 528:358-363.

Packard MG, McGaugh JL (1996) Inactivation of hippocampus or caudate nucleus with lidocaine differentially affects expression of place and response learning. Neurobiol Learn Mem 65:65-72.

Parent A, Sato F, Wu Y, Gauthier J, Levesque M, Parent M (2000) Organization of the basal ganglia: the importance of axonal collateralization. Trends in neurosciences 23:S20-27.

Parr-Brownlie LC, Poloskey SL, Flanagan KK, Eisenhofer G, Bergstrom DA, Walters JR (2007) Dopamine lesion-induced changes in subthalamic nucleus activity are not associated with alterations in firing rate or pattern in layer V neurons of the anterior cingulate cortex in anesthetized rats. Eur J Neurosci 26:1925-1939.

Parthasarathy HB, Graybiel AM (1997) Cortically driven immediate-early gene expression reflects modular influence of sensorimotor cortex on identified striatal neurons in the squirrel monkey. J Neurosci 17:2477-2491.

Parthasarathy HB, Schall JD, Graybiel AM (1992) Distributed but convergent ordering of corticostriatal projections: analysis of the frontal eye field and the supplementary eye field in the macaque monkey. J Neurosci 12:4468-4488.

Perruchet P, Amorim MA (1992) Conscious knowledge and changes in performance in sequence learning: evidence against dissociation. J Exp Psychol Learn Mem Cogn 18:785-800.

Ragozzino ME, Ragozzino KE, Mizumori SJ, Kesner RP (2002) Role of the dorsomedial striatum in behavioral flexibility for response and visual cue discrimination learning. Behavioral neuroscience 116:105-115.

Ray NJ, Jenkinson N, Wang S, Holland P, Brittain JS, Joint C, Stein JF, Aziz T (2008) Local field potential beta activity in the subthalamic nucleus of patients with Parkinson's disease is associated with improvements in bradykinesia after dopamine and deep brain stimulation. Exp Neurol 213:108-113.

Redgrave P, Vautrelle N, Reynolds JN (2011) Functional properties of the basal ganglia's re-entrant loop architecture: selection and reinforcement. Neuroscience 198:138-151.

Rosenbaum DA, Cohen RG, Jax SA, Weiss DJ, van der Wel R (2007) The problem of serial order in behavior: Lashley's legacy. Hum Mov Sci 26:525-554.

Rothwell PE, Hayton SJ, Sun GL, Fuccillo MV, Lim BK, Malenka RC (2015) Input- and Output-Specific Regulation of Serial Order Performance by Corticostriatal Circuits. Neuron 88:345-356.

Rueda-Orozco PE, Robbe D (2015) The striatum multiplexes contextual and kinematic information to constrain motor habits execution. Nature neuroscience 18:453-460.

Salamone JD, Correa M, Nunes EJ, Randall PA, Pardo M (2012) The behavioral pharmacology of effort-related choice behavior: dopamine, adenosine and beyond. J Exp Anal Behav 97:125-146.

Schneider GE (2014) Brain Structure and Its Origins: MIT Press.

Schwabe L, Wolf OT (2011) Stress-induced modulation of instrumental behavior: from goal-directed to habitual control of action. Behavioural brain research 219:321-328.

Sesack SR, Deutch AY, Roth RH, Bunney BS (1989) Topographical organization of the efferent projections of the medial prefrontal cortex in the rat: an anterograde tract-tracing study with Phaseolus vulgaris leucoagglutinin. J Comp Neurol 290:213-242.

Sgambato V, Abo V, Rogard M, Besson MJ, Deniau JM (1997) Effect of electrical stimulation of the cerebral cortex on the expression of the Fos protein in the basal ganglia. Neuroscience 81:93-112.

Shepherd (2003) The Synaptic Organization of the Brain: Oxford University Press, USA.

Siegel M, Donner TH, Engel AK (2012) Spectral fingerprints of large-scale neuronal interactions. Nat Rev Neurosci 13:121-134.

Smith KS, Graybiel AM (2016) Habit Formation Coincides with Shifts in Reinforcement Representations in the Sensorimotor Striatum. Journal of neurophysiology jn 00925 02015.

Stephenson-Jones M, Samuelsson E, Ericsson J, Robertson B, Grillner S (2011) Evolutionary conservation of the basal ganglia as a common vertebrate mechanism for action selection. Current biology : CB 21:1081-1091.

Striedter (2005) Principles of Brain Evolution: Sinauer Associates.

Tang C, Pawlak AP, Prokopenko V, West MO (2007) Changes in activity of the striatum during formation of a motor habit. Eur J Neurosci 25:1212-1227.

Tang CC, Root DH, Duke DC, Zhu Y, Teixeria K, Ma S, Barker DJ, West MO (2009) Decreased firing of striatal neurons related to licking during acquisition and overtraining of a licking task. J Neurosci 29:13952-13961.

Thorn CA, Atallah H, Howe M, Graybiel AM (2010) Differential dynamics of activity changes in dorsolateral and dorsomedial striatal loops during learning. Neuron 66:781-795.

Trytek ES, White IM, Schroeder DM, Heidenreich BA, Rebec GV (1996) Localization of motor- and nonmotor-related neurons within the matrix-striosome organization of rat striatum. Brain Res 707:221-227.

Tye KM, Prakash R, Kim SY, Fenno LE, Grosenick L, Zarabi H, Thompson KR, Gradinaru V, Ramakrishnan C, Deisseroth K (2011) Amygdala circuitry mediating reversible and bidirectional control of anxiety. Nature 471:358-362.

Van den Bercken JH, Cools AR (1982) Evidence for a role of the caudate nucleus in the sequential organization of behavior. Behavioural brain research 4:319-327.

van der Meer MA, Johnson A, Schmitzer-Torbert NC, Redish AD (2010) Triple dissociation of information processing in dorsal striatum, ventral striatum, and hippocampus on a learned spatial decision task. Neuron 67:25-32.

van der Meer MA, Redish AD (2011) Ventral striatum: a critical look at models of learning and evaluation. Current opinion in neurobiology 21:387-392.

VandenBerg PM, Hogg TM, Kleim JA, Whishaw IQ (2002) Long-Evans rats have a larger cortical topographic representation of movement than Fischer-344 rats: a microstimulation study of motor cortex in naive and skilled reaching-trained rats. Brain Res Bull 59:197-203.

Venkatraman S, Jin X, Costa RM, Carmena JM (2010) Investigating neural correlates of behavior in freely behaving rodents using inertial sensors. Journal of neurophysiology 104:569-575.

Verwey WB (2001) Concatenating familiar movement sequences: the versatile cognitive processor. Acta Psychol (Amst) 106:69-95.

Verwey WB, Abrahamse EL, de Kleine E (2010) Cognitive processing in new and practiced discrete keying sequences. Front Psychol 1:32.

Voorn P, Vanderschuren LJ, Groenewegen HJ, Robbins TW, Pennartz CM (2004) Putting a spin on the dorsal-ventral divide of the striatum. Trends in neurosciences 27:468-474.

Wang XJ (2010) Neurophysiological and computational principles of cortical rhythms in cognition. Physiol Rev 90:1195-1268.

West MO, Carelli RM, Pomerantz M, Cohen SM, Gardner JP, Chapin JK, Woodward DJ (1990) A region in the dorsolateral striatum of the rat exhibiting single-unit correlations with specific locomotor limb movements. Journal of neurophysiology 64:1233-1246.

Wu JH, Corwin JV, Reep RL (2009) Organization of the corticostriatal projection from rat medial agranular cortex to far dorsolateral striatum. Brain Res 1280:69-76.

Wymbs NF, Bassett DS, Mucha PJ, Porter MA, Grafton ST (2012) Differential recruitment of the sensorimotor putamen and frontoparietal cortex during motor chunking in humans. Neuron 74:936-946.

Yin HH, Knowlton BJ, Balleine BW (2004) Lesions of dorsolateral striatum preserve outcome expectancy but disrupt habit formation in instrumental learning. Eur J Neurosci 19:181-189.

Yin HH, Knowlton BJ, Balleine BW (2005a) Blockade of NMDA receptors in the dorsomedial striatum prevents action-outcome learning in instrumental conditioning. Eur J Neurosci 22:505-512.

Yin HH, Knowlton BJ, Balleine BW (2006) Inactivation of dorsolateral striatum enhances sensitivity to changes in the action-outcome contingency in instrumental conditioning. Behavioural brain research 166:189-196.

Yin HH, Mulcare SP, Hilario MR, Clouse E, Holloway T, Davis MI, Hansson AC, Lovinger DM, Costa RM (2009) Dynamic reorganization of striatal circuits during the acquisition and consolidation of a skill. Nature neuroscience 12:333-341.

Yin HH, Ostlund SB, Knowlton BJ, Balleine BW (2005b) The role of the dorsomedial striatum in instrumental conditioning. Eur J Neurosci 22:513-523.