

UCSF

UC San Francisco Electronic Theses and Dissertations

Title

Learning in the corticostriatal network

Permalink

<https://escholarship.org/uc/item/9pk520jp>

Author

Lemke, Stefan Mattia

Publication Date

2020

Peer reviewed|Thesis/dissertation

Learning in the corticostriatal network

by
Stefan Lemke

DISSERTATION

Submitted in partial satisfaction of the requirements for degree of
DOCTOR OF PHILOSOPHY

in

Neuroscience

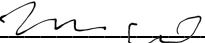
in the

GRADUATE DIVISION

of the

UNIVERSITY OF CALIFORNIA, SAN FRANCISCO

Approved:


DocuSigned by:

3A5C1850C2F1412... Michael Brainard
Chair

DocuSigned by:

DocuSigned by:407... Karunesh Ganguly

DocuSigned by:

DocuSigned by:4F0... Joshua Berke

DocuSigned by:

EF21CD5A6EE643A... Loren Frank

Committee Members

Copyright 2020

by

Stefan Lemke

To my parents, Lory and Pieranna, and my brother, David

Acknowledgements

There are many people I am indebted to for making my PhD work possible. I'd like to thank Karunesh Ganguly, my PhD advisor, for providing an endless source of excitement, motivation, and support. Your positivity and confidence in me have made it easy to look past what didn't work until it did.

I'd like to thank Dhakshin Ramanathan for his mentorship since my first year of graduate school. All the experiments we did together, both successful and not so successful, taught me how carry out good science and approach our work with positivity.

I'd like to thank all the members of the Ganguly lab that have made the lab an ideal place to be a graduate student: Tanuj Gulati, Jason Godlove, Seok Joon Won, Adelyn Tu-Chan, Nikhilesh Natraj, Doug Totten, Jaekyung Kim, Reza Abiri, Nick Hardy, Preeya Khanna, Harman Ghuman, Hoseok Choi, Kyungsoo Kim. I owe a lot of what I know to each of your unique expertise, generosity, and time.

I was very fortunate to have joined the Ganguly lab alongside two other students who are proficient scientists, Ling Guo and Tess Veuthey. Working together with both of you, solving both scientific and graduate school problems, made my years in the Ganguly lab more productive and a lot more fun. I'd also like to thank the other students we gained along the way: Kate Derosier, Daniel Silversmith, Sravani Kondapavulur, and David Darevsky. I have learned a lot from working with each of you. I'd also like to thank April Hishinuma, Chelsey Rodriguez, and Sapeeda Barati for their support as part of the Ganguly lab.

I'd like to thank Josh Berke and Daniel Egert for being wonderful collaborators. The opportunity to work together has allowed me to advance my work and learn about a new side of electrophysiology.

I'd like to thank Josh Berke, Michael Brainard, and Loren Frank for serving on my thesis committee. Your feedback has been indispensable over the past six years. I'd also like to thank Alexandra Nelson and Kira Poskanzer for shaping my thesis work through their time serving on my qualification exam committee.

I'd like to thank Pat Veitch and Lucita Nacionales for their help in navigating graduate school.

I'd like to thank my amazing mentors from St. Olaf College: Jay Demas, Kevin Crisp, and Shelly Dickinson. Without all of your time and encouragement I would not be at UCSF.

I want to thank the UCSF Neuroscience class of 2014, who made the transition to San Francisco and UCSF more fun and much more manageable. I especially want to thank Perry Spratt and Stephanie Holden for their friendship and support over the past six years.

I want to thank Trisha Vaidyanathan for giving me confidence and making everything more fun.

I want to thank my parents, Lory Lemke and Pieranna Garavaso, and my brother, David Lemke, for years and years of love and support.

Contributions

Chapter 2 is adapted from Lemke, S. M., Ramanathan, D. S., Guo, L., Won, S. J. & Ganguly, K. Emergent modular neural control drives coordinated motor actions. *Nat. Neurosci.* **22**, 1122–1131 (2019). doi:10.1038/s41593-019-0407-2

Chapter 3 is adapted from Lemke, S. M., Ramanathan, D. S., Darevsky, D., Egert, D., Berke, J.D., Ganguly, K. Sleep spindles coordinate corticostriatal reactivations during the emergence of automaticity. *bioRxiv* 2020.10.25.354282 (2020). doi:10.1101/2020.10.25.354282

Learning in the corticostriatal network

Stefan Lemke

Abstract

A fundamental question regarding movement control is how the brain regulates motor learning. In other words, how does initially variable control of a new behavior transition into stable control of a skilled behavior? Plasticity in the corticostriatal network, which includes the motor cortex and the striatum, an input structure of the basal ganglia innervated by the motor cortex, has been identified as a critical regulator of learning. In this thesis, I begin by presenting a historical perspective on the motor network, outlining why the corticostriatal network is uniquely positioned to regulate learning. I then present my work investigating the corticostriatal network during learning. In Chapter 2, I explore how the corticostriatal network drives learning of a complex coordinated action and show that distal vs. proximal movements are differentially encoded. In Chapter 3, I examine how and when corticostriatal plasticity occurs during learning and show that non-REM sleep and sleep spindles play central roles in shaping the corticostriatal network. I conclude by presenting a model for motor network organization in Chapter 4, proposing that corticostriatal plasticity may mediate learning by transitioning an initially cortically driven behavior to a subcortically driven behavior.

Table of Contents

Chapter 1: Introduction.....	1
The stability-plasticity dilemma.....	1
History of the distributed motor network.....	2
Brainstem / Cerebellum.....	6
Motor Cortex.....	8
Basal Ganglia.....	10
The corticostriatal network and learning.....	13
References.....	17
Chapter 2: What does the corticostriatal network encode?.....	26
Abstract.....	26
Introduction.....	27
Results.....	29
Discussion.....	55
Methods.....	60
References.....	73
Chapter 3: How is the corticostriatal network modified?.....	80
Abstract.....	80
Introduction.....	81
Results.....	83
Discussion.....	109

Methods	112
References	125
Chapter 4: Why corticostriatal plasticity matters: a model of motor network organization.....	131
A cortical-to-subcortical gradient in neural flexibility.....	133
Learning is mediated through plasticity in the corticostriatal network.....	137
Summary	142
References	143

List of Figures

Chapter 1:

Figure 1.1. Motor network model	1
Figure 1.2. Illustration of Galvani’s frog leg experiments	3
Figure 1.3. Hitzig and Frisch discovery of the motor cortex	4

Chapter 2:

Figure 2.1. Localization of electrodes	29
Figure 2.2. Corticostriatal projections	30
Figure 2.3. Refinement of skilled fine and gross movements is dissociable during reach-to-grasp skill learning	32
Figure 2.4. Fine and gross movements do not covary between days five through eight of reach-to-grasp skill learning	32
Figure 2.5. Coordinated movement-related activity emerges across M1 and DLS during skill learning	34
Figure 2.6. Non-zero phase lag between high coherence M1 and DLS LFP signals	35
Figure 2.7. LFP power and coherence increases for behaviorally-matched trials during learning	35
Figure 2.8. Peak frequency of M1-DLS LFP coherence covaries with movement duration on day eight	37
Figure 2.9. Coordinated M1 and DLS spiking activity emerges during skill learning	40
Figure 2.10. Details of M1 and DLS spiking activity	41

Figure 2.11. Coordinated M1 and DLS activity is specifically linked to skilled gross movements	43
Figure 2.12. M1 and DLS inactivation have differential effects on skilled fine and gross movements	46
Figure 2.13. Body posture during reaching before and after DLS inactivation	47
Figure 2.14. DLS inactivation decreases movement-related low-frequency M1 activity.....	49
Figure 2.15. Localization of muscimol infusion	49
Figure 2.16. Localization and behavioral effects of excitotoxic DLS lesions	50
Figure 2.17. Persistent disruption of skilled fine movements after M1 lesion.....	52
Figure 2.18. Localization of M1 photothrombotic lesion	52
Figure 2.19. Skilled fine movement representation in M1	54
Figure 2.20. Increased GPFA neural trajectory consistency during skill learning	54
Chapter 3:	
Figure 3.1. Offline striatal NMDA receptor activation is required to develop a consistent behavior.....	84
Figure 3.2. Animals do not adapt quickly to new pellet position after long-term reach-to-grasp training	86
Figure 3.3. Corticostriatal functional connectivity increases during offline periods	89
Figure 3.4. Offline increases in functional connectivity predict the emergence of low-dimensional cross-area neural dynamics during behavior	91
Figure 3.5. Corticostriatal transmission strength within offline periods is maximal during sleep spindles in NREM.....	94

Figure 3.6. Electrophysiology recordings from M1 and DLS.....	96
Figure 3.7. Corticostriatal transmission strength across behavioral states and NREM rhythms.....	97
Figure 3.8. Corticostriatal modulation across NREM rhythms.....	98
Figure 3.9. Striatal reactivations during sleep spindles reflect cortical input	101
Figure 3.10. Delta wave modulation change with training	102
Figure 3.11. Slow oscillation modulation change with training	103
Figure 3.12. Sleep spindle modulation predicts offline changes in corticostriatal coupling ...	105
Figure 3.13. Comparison of time spent in each behavioral state during pre- and post-sleep ..	106
Figure 3.14. Corticostriatal transmission strength changes for non-sleep spindle modulated pairs of M1 and DLS units	106
Figure 3.15. The interaction between sleep spindles and slow oscillations impact the role of sleep spindles within the corticostriatal network	108
Chapter 4:	
Figure 4.1. Proposed shared neural axis for different motor control schemes.....	132
Figure 4.2. Proposed cortical-to-subcortical gradient in neural flexibility-to-stability.....	134
Figure 4.3. M1 and DLS spiking properties.....	134
Figure 4.4. Comparison of M1 and DLS neural stability vs. flexibility	135
Figure 4.5. Comparison of M1 and DLS neural stability vs. flexibility for sub-selected spike networks.....	136
Figure 4.6 Pairwise correlation values of M1 and DLS networks	136
Figure 4.7. Proposed cortical-to-subcortical shift in movement control with learning.....	138

Chapter 1: Introduction

The stability-plasticity dilemma

The *stability-plasticity* dilemma is a critical constraint for artificial and biological neural systems. The dilemma is how a neural system can acquire new information (plasticity) without overriding older knowledge (stability). What this refers to, in the case of procedural learning, is the ability to learn new behaviors without forgetting others. For example, I can spend the next week learning as many new skills as possible: shooting a basketball, crocheting a blanket, playing the piano, yet at the end of the week I will not have forgotten how to ride my bike or sign my name. How does the brain achieve this?

A “simple” explanation to solve this dilemma is that the ability to explore new behaviors and the ability to execute well-established behaviors derive from distinct brain regions. For example, the stable control of *innate* behaviors such as locomotion has been linked to the brainstem^{1,2}, while the flexible control of exploratory movements associated with *new* behaviors, such learning to play the piano, has been linked to motor cortex (Figure 1.1a)^{3,4}.




Figure 1.1. Motor network model. a. Example distributed motor control. **b.** Potential neural controllers of a learned behavior.

However, what happens if I practice the piano until I can play with stability comparable to the control of locomotion (**Figure 1.1b**)? Does local plasticity allow the same brain region to shift from the control of flexible *new* to stable *learned* behaviors (dashed red arrow)? Does the brain region associated with stable *innate* behaviors also control stable *learned* behaviors (dashed blue arrow)? Does a new brain region assume control of stable *learned* behaviors (dashed black arrow)? These questions are the topic of this thesis. I built on a body of work implicating plasticity in the corticostriatal network as a key regulator in the transition from *new* to *learned* behaviors. This thesis aims to clarify what is (and is not) encoded by corticostriatal activity throughout learning of a complex coordinated action (Chapter 2) and explore how and when corticostriatal network plasticity occurs during learning (Chapter 3). In conclusion (Chapter 4), I propose a model of motor network organization to answer the question posed in **Figure 1.1b**. As an introduction, this chapter presents a historical perspective on our modern understanding of the motor network, outlining why the corticostriatal network is particularly relevant for learning and worthy of our focus.

History of the distributed motor network

The goal of this section is to provide a brief historical perspective that outlines why the motor network is considered *distributed*, that is, with *multiple regions able to control movement*.

Our modern understanding of motor control began with the first evidence that nerves activate muscles via electricity. By the early 18th century, anatomists had established that muscles were connected to the brain through nerves. However, the prevailing theories were that vibrations, ethereal spirits, or fluids flowed through nerves to activate muscles⁵. In 1714, Isaac Newton wrote in *General Scholium*⁶⁻⁸:

The members of animal bodies move at the command of the will, namely, by the vibrations of this Spirit, mutually propagated along the solid filaments of the nerves, from the outward organs of sense to the brain, and from the brain into the muscles. But these are things that cannot be explained in few words, nor are we furnished with that sufficiency of experiments which is required to an accurate determination and demonstration of the laws by which this electric and elastic spirit operates.

It wasn't until 1791 that Luigi Galvani established that nerve impulses are electrical in nature. In the *Commentary on the Effects of Electricity on Muscular Motion*^{5,9}, Galvani details ten years of experimental results, including famous work where he electrically stimulated muscles in frog legs, causing them to contract and demonstrating that nerves activate muscles via “animal electricity” (**Figure 1.2**). In the *Commentary*, Galvani theorizes that animal electricity is secreted by the brain and distributed by the nerves. He reports a clever experiment in which he extracts droplets of oil from nerves and concludes that nerves must have an oily covering that prevents leakage of electricity to surrounding tissues⁵. With the *Commentary*, Galvani uncovered the role of electricity as the means in which the nervous system communicates and thereby established the field of electrophysiology.




Figure 1.2. Illustration of Galvani's frog leg experiments. Muscular contraction caused by electrical stimulation, generated by connecting zinc (Z) and copper (C) wires. From *Galvani (1791)*.

Galvani's experiments themselves, however, did not provide direct evidence that the brain was a source of “animal electricity”. It was Galvani's nephew, Giovanni Aldini, who continued Galvani's work and, in 1802, showed it was possible to evoke muscle contractions by stimulating

the brain^{5,10}. Along with experimental work in animals, Aldini performed public demonstrations where he electrically stimulated the exposed brains of criminals who had been recently killed by decapitation. These gruesome experiments demonstrated that facial muscle contractions could be evoked through electrical stimulation of the brain^{5,11}. In these experiments, Aldini also reports the prescient result that stimulating the cerebral cortex of one hemisphere elicited facial muscle contractions of the opposite hemisphere, early evidence for what would come to be called the “motor cortex” nearly 70 years later^{10,12}.

Before the discovery of the motor cortex, however, the cortex was largely considered “unexcitable”, that is, that stimulation of the cortex would not evoke movement^{5,12}. It was instead believed that cortex was involved in “intellectual functions” while more primitive structures such as the *corpus striatum*, *cerebellum*, or *brainstem* (discussed below) were responsible for driving movement^{13,14}. This view was so prevalent that when, in 1870, Gustav Fritsch and Edvard Hitzig




Figure 1.3. Hitzig and Frisch discovery of the motor cortex. Illustration from *Fritsch and Hitzig (1870)* on the left and recreation emphasizing stimulation locations that caused facial (◇), neck (Δ), and leg (+ and #) movements from *Brazier (1988)* on the right.

published their discovery of motor cortex, outlining a specific region of the dog’s cortex that could elicit contralateral movement¹⁵ (**Figure 1.3**), prominent scientists questioned their conclusions. In 1875, Ludimar Herman, a well-known electrophysiologist of the time, wrote^{11,16}:

There had to be the suspicion that the authors obtained their positive results because of the strength of electrical current they were using, never before attempted, and that this high electrical current had set into action motor apparatus deeper within the brain. In fact, in an organ containing an incredibly fine net of apparatus in a compact mass, an electrical stimulus must appear most dangerous, since it cannot be limited to one particular area.

George Henry Lewes made a more succinct criticism, arguing that the observation that movements could be evoked by electrical stimulation of the cortex was not sufficient to classify cortex as a “motor center”^{11,17}:

We do not consider the centre of laughter to be located in the sole of the foot, because tickling the sole causes laughter.

These concerns, centered on the potential spread of electrical current to other brain regions^{13,14}, led to the practice of “punctate” mapping, where experimenters used the minimum current required to elicit movement. These experiments dispelled the worry of current spread and provided confirmatory evidence that electrical stimulation of cortex can evoke movements^{14,15,18}. Punctate mapping also allowed Charles Sherrington to establish the idea of motor cortex topography with his famous publication of “motor maps” in 1917¹⁹.

This work of Fritsch and Hitzig, confirmed by Sherrington and others, added motor cortex to the growing list of brain regions where movements had been observed with electrical stimulation, including the *corpus striatum* and *cerebellum*⁵. Anatomical studies also confirmed that multiple “descending tracts”, originating from *motor cortex* and *brainstem*, connected the brain to the spinal cord²⁰, altogether providing overwhelming evidence that multiple brain regions control movement. Over the next century, the study of each of these regions grew into largely independent fields of

study. In the following sections, I will briefly outline some of the main themes of this work. Of particular relevance to this thesis is the study of the motor cortex and basal ganglia, as modern work has shown that the interaction of these regions is critical for learning new behaviors. The final section in this chapter will focus on this network.

Brainstem / Cerebellum

Both the brainstem and cerebellum have long been considered central to movement due to the severity of occipital injuries¹². The proximity of the brainstem and cerebellum led some early investigators studying occipital injuries to mistakenly assign motor functions between these structures¹². However, in 1760, A.C. Lorry published results from a series of careful brainstem lesion experiments in cats and dogs establishing that the brainstem, rather than the cerebellum, controls innate movements such as respiration^{12,21}. The modern view of the cerebellum, on the other hand, can be traced back to the early 19th century, when Pierre Flourens proposed a role for the cerebellum in the *coordination* of movements:^{22,23}

All movements persist following ablation of the cerebellum: all that is missing is that they are not regular and coordinated. From this I have been induced to conclude that the production and the coordination of movements form two classes of essentially distinct phenomena and that they reside in two classes of organs also essentially distinct: with coordination in the cerebellum and production in the spinal cord and medulla oblongata.

Modern work on the brainstem and cerebellum has largely confirmed these different functions. Serial dissections confirmed the vital role for the brainstem in movement: animals in which the forebrain has been removed but with an intact brainstem retain their ability to perform a large

repertoire of innate movements compared to animals with transections at the brainstem-spinal cord junction²⁴.

More recent work has further elucidated that the brainstem contains a variety of nuclei that control disparate motor functions²⁵. Although many of these function relate to innate behaviors including orofacial movements, breathing, and locomotion, recent evidence has demonstrated that certain brainstem nuclei also control aspects of skilled forelimb control²⁵⁻²⁷. Of note, recent work showed that ablations of the medullary reticular formation ventral part, or MdV, disrupts grasping in a skilled reaching task but does not disrupt simple locomotor tasks²⁷. Similarly, lesions of the red nucleus have been shown to impair reaching and grasping but not locomotion²⁸. An important consequence of this work is our understanding of how forelimb control recovers after brain injury. It has been proposed that redundant controllers of forelimb movements in the cortex and brainstem provide substrates for functional recovery²⁹. Further work is required to understand whether cortex and brainstem regions interact or act as parallel controllers in skilled forelimb control.

Similarly, modern work has clarified the role of the cerebellum in movement coordination. The cerebellum is thought to compute the estimated sensory consequences of motor commands³⁰. These “forward models” allow for fast feedback corrections (that are faster than sensory feedback would allow) that are critical for coordinated movements. In addition to coordination, these forward models also allow for motor adaptation, that is, the ability to adapt the brain’s estimates is they are consistently biased³¹.

Despite the disparate functions assigned to the cerebellum and brainstem, neither region has been shown to play a central role in the learning of new behaviors (but see *ref. 28*). Although it is possible that future work uncovers a role for the cerebellum and brainstem in learning, this thesis focuses on two regions that have been shown to play central roles in learning: the motor cortex and basal ganglia.

Motor Cortex

Disputes over what aspects of movement the motor cortex controls began soon after the discovery of the motor cortex. While Fritsch and Hitzig's original discovery of the motor cortex suggested that motor cortex activated individual muscles, follow-up work by David Ferrier, who used longer electrical stimulation pulses, showed that activation of motor cortex can elicit complex movements that involve the coordination of several muscles¹⁴. These initial differences in results grew into a larger "muscles vs. movement" debate over the next century as neural recordings from motor cortex began with Edward Evarts' pioneering work³². Since then, neural activity in motor cortex has been linked to specific muscle activations, as well as a range of kinematic features such as movement amplitude, direction, speed, velocity, acceleration, and position³³. Altogether, this work presents a complex view of what motor cortex activity actually encodes. Indeed, Apostolos Georgopoulos, one of the pioneers of this work, reflects on the diversity of roles played by the motor cortex as follows³⁴:

It would be fair to say that motor cortex is involved in motor 'stuff', be it real movement, imagined movement, or motor cognition.

These complexities, combined with the lack of clear conclusions, have led to distinct conceptualizations of motor cortical activity. Krishna Shenoy, Mark Churchland, and colleagues,

recently presented the *Dynamical Systems* perspective of motor control as an alternative to conventional work focused on finding the aspects of movement or muscle activity that correlate best to motor cortical activity^{35,36}. The *Dynamical Systems* perspective proposes that neural activity during movement evolves according to laws of a dynamical system with an initial state set by motor preparatory activity, and that any correlations between motor cortical activity and specific muscle activations or kinematic features are incidental. Paul Cisek simply justified this view by stating: “*the role of the motor system is to produce movement, not to describe it*”³⁷. The power of the *Dynamical Systems* perspective has been demonstrated through accurate predictions of neural activity during movement, as well as new insights gained about how motor cortical activity may control movement^{38,39}.

The modern study of motor cortex has also continued to leverage lesion techniques. Such work has provided considerable evidence that a key role of motor cortex is to control dexterous movements, that is, the coordination of fine movements of the hands and fingers⁴⁰⁻⁴³. Intriguingly, however, recent work showed that motor cortex does not play the same central role in controlling non-dexterous movements⁴⁴. This work, by Bence Ölveczky and colleagues, provided compelling evidence that a complex sequence of non-dexterous movements learned with an intact motor cortex can be performed without disruption after a complete motor cortical lesion. The motor cortex is therefore only required to learn, but not execute, these movement sequences, presenting an intriguing dichotomy in the role of motor cortex, in the *learning vs. executing* non-dexterous movements.

Lastly, it is also important to note that “motor cortex” is often subdivided into multiple regions. For example, in rodents there exists both a primary and secondary motor cortical region⁴⁵. Each motor cortical region projects to many downstream brain regions in addition to the spinal cord and receives input from a range of brain regions. Therefore, motor cortical activity is both shaped “upstream” and “downstream” by the activity in other brain regions. Of particular relevance to this thesis are the projections from motor cortex to the striatum, a major input region of the basal ganglia, which have been shown to mediate learning of new behaviors^{46,47}.

Basal Ganglia

The term “basal ganglia” describes a set of ~7 subcortical regions bound together by reciprocal interconnections. Although the exact characterization of which brain regions make up the basal ganglia has (and continues to) vary between researchers, the modern conception of the basal ganglia is attributed to work by Walle Nauta and William Mehler published in 1965^{48,49}. Our contemporary view of the basal ganglia includes the caudate nucleus, putamen, globus pallidus (interna and externa), subthalamic nucleus, and substantia nigra (pars compacta and pars reticula)⁵⁰.

Far predating this view, however, was the initial characterization of the largest component of the basal ganglia, the *corpus striatum* (or *striatum*). The *striatum* is so called because of it is made up of striations of gray and white matter, an anatomical feature so distinct that Thomas Willis characterized it in 1664^{49,51}. Willis postulated that the corpus striatum may be involved in movement due to its placement in the brain, a striking hypothesis given that it occurred long before

“animal electricity” was even discovered (Willis subscribed to the theory that spirits traveled in nerves to activate muscles). In *Cerebri anatome* Willis wrote^{49,51}:

If one should enquire about their function, it immediately comes to one's mind that these bodies are placed between the cerebrum and its appendix and form a lodging place for both ... here, namely, as it were in a most frequented public place, the animal spirits concerned with the execution of willed action are directed into the appropriate nerves.

This prescient description has largely held true as functions of the basal ganglia have been investigated in the following centuries. Work in the first half of the 19th century demonstrated it was possible to elicit movement by stimulating the striatum⁵² and a more complete theory was put forth by David Ferrier in *The Functions of the Brain* in 1876¹⁸, where he differentiates the functions of the motor cortex and the striatum, suggesting the striatum is important for automatic movements:

It appears from these facts that the corpora striata proper are centres of innervation of the same movements as are differentiated in the cortical motor centres, but of a lower grade of specialization. The innervation of the limbs in all that relates to their employment as instruments of consciously discriminated acts is dependent on the cortical centres, while for all other purposes involving mere strength or automatism, primary or secondary, the corpora striata with the lower ganglia are sufficient.

In fact, modern work has placed the basal ganglia as a central controller of habits and automatic actions⁵³. Lesions in the basal ganglia are known to disrupt habitual, automatic, and learned behaviors⁵⁴. Additionally, disorders that impact the basal ganglia, such Parkinson's Disease, have provided further evidence for its central role in controlling movement^{49,55}. However, the precise role of the basal ganglia in controlling movement remain debated. I will describe two major proposed roles: selecting actions and controlling continuous movement kinematics.

The classical view of the basal ganglia is that the output nuclei, particularly the globus pallidus interna (GPi) and substantia nigra pars reticulata (SNr), keep motor regions in the cortex and brainstem under tonic inhibition to prevent unwanted movements^{56,57}. To generate a movement, therefore, this inhibitory drive must be released. This release from inhibition is proposed to occur through the activation of inhibitory striatal neurons, which make up the input nuclei of the basal ganglia and receive cortical, thalamic, and dopaminergic input^{56,57}. This influential conceptualization has led to the idea that the basal ganglia is critical for the initiation and “selection” of actions, but does not directly implicate the basal ganglia in the control of detailed kinematic features of movement. This latter idea is instead supported by considerable evidence that basal ganglia activity encodes kinematic features^{58,59} and that lesions of the basal ganglia disrupt the details of movement but do not prevent initiation of movements^{58,60}. An active field of study is investigating how these two proposed functions occur in the basal ganglia.

The goal of the preceding section was to demonstrate that the motor network is comprised of a variety of brain regions capable of driving movement. Even beyond the regions that have been mentioned, there are additional brain regions, although typically not considered “motor regions”, that influence movement. For example, the thalamus is responsible for connecting the basal ganglia back to the motor cortex^{61,62} and thalamic inputs to the basal ganglia are known to mediate basal ganglia function⁶³. Similarly, inputs from the amygdala to the basal ganglia have been shown to modulate basal ganglia function⁶⁴. Furthermore, with improved recording technologies, it is becoming increasingly clear that the study of any brain region in isolation provides a limited view of brain function⁶⁵ and that studying the interactions between regions is a critical step towards furthering our understanding of how the brain controls movement. In the next section, we will

begin to restrict our focus to the study of one such interaction, between the motor cortex and the striatum. This *corticostriatal network* has been shown to play a central role in learning new behavior, the focus of this thesis.

The corticostriatal network and learning

In the preceding sections I briefly described the major regions within the distributed motor network. In this section, my goal is to provide evidence that the *corticostriatal network* is a uniquely important hub within this network responsible for learning new behaviors.

We can think of learning in several ways. From a cognitive lens, we can think of learning as the association of an action and a reward. From the lens of motor control, we can think of learning as the combination and refinement of movements into a consistent action. This latter type of learning is typically referred to as *motor skill learning* and is classically conceptualized as “increasing spatial and temporal accuracy of movements with practice”⁶⁶. The central role such learning plays in our lives was expressed by Daniel Willingham in a 1998 review⁶⁶:

It would be a strange, cruel world without motor skill learning: Automobile drivers would get behind the wheel as if for the first time every day; there would be no virtuosic athletic and artistic performances to watch; and tying one's shoes in the morning would require minutes of intense concentration.

The corticostriatal network has been proposed to play a central role in motor skill learning. Below, I will describe two central factors that make the corticostriatal network particularly well suited to mediate learning (1) input/output *connectivity* and (2) evidence for *plasticity*.

We will first consider *connectivity*. The functional connectivity of the basal ganglia was briefly described in the preceding section: the output nuclei of the basal ganglia are thought to tonically inhibit movement through the inhibition of motor regions in the cortex and brainstem and excitatory inputs from cortex and thalamus to the striatum are thought to result in the reduction of this inhibitory output of the basal ganglia to facilitate movement^{56,57}. In the context of learning, plasticity in the inputs to the striatum would therefore allow for the specific, consistent activation of striatum to emerge with learning, leading to the consistent disinhibition of motor centers in the cortex and brainstem and stable, learned behavior. The inputs to the striatum are uniquely suited to undergo such plasticity as the striatum receives inputs from nearly all (if not all) regions of cortex⁶⁷, allowing the striatum to integrate widespread information about context and behavior. In fact, there is considerable evidence for convergence within the inputs to the striatum that allow a single site in the striatum to receive input from widespread regions of cortex⁶⁸. For example, it has been observed that projections from functionally linked areas of somatosensory and motor cortex, i.e., areas that represent the same body part, overlap in the striatum⁶⁹. Such convergence is enabled by the projection patterns of corticostriatal neurons. A typical corticostriatal projection neuron makes sparse connectivity with many striatal neurons, rather than strong connectivity with few striatal neurons⁶⁸, in an approximate 6:1 ratio of corticostriatal neurons to striatal projection neurons. It is also important to note that there are two broad classes of corticostriatal neurons: intratelencephalic (IT) neurons that project widely within the cortex and striatum and pyramidal tract (PT) neurons that project to the brainstem or spinal cord with collaterals to the striatum (as well as to the thalamus, subthalamic nucleus, zona incerta, globus pallidus internal segment, substantia nigra parts reticulata, superior colliculus, and pontine nuclei). The specific role of these two classes of neurons is an active field of study. The striatum also receives dopaminergic

innervation, which has been shown to both serve as a reward signal that influences plasticity and play a role in driving movement itself⁵⁷. Altogether, this unique connectivity places the striatum, and basal ganglia more broadly, in a unique position to mediate learning.

Next, we will consider *plasticity*. To mediate learning a brain network must have the ability to modify connectivity to represent learned actions. NMDA-dependent plasticity has been observed in the excitatory inputs to the striatum^{70,71}. Such excitatory inputs originate from cortex and thalamus and largely synapse onto spiny projection neurons within the striatum, which compose 95% of striatal cells. Blocking such NMDA-dependent plasticity in the striatum has been shown to prevent learning, suggesting that striatal input plasticity is a critical mediator of learning⁷²⁻⁷⁵. The focus of work investigating striatal input plasticity during learning has largely been on the cortical inputs to the striatum, both *in vitro*⁴⁷ as well as *in vivo*^{74,76-78}. However, recent work has also implicated plasticity in thalamic inputs^{79,80}. Additionally, 5% of striatal cells are interneurons of two classes: large aspiny acetylcholine interneurons and medium aspiny GABAergic interneurons⁶⁸. There is evidence that within the class of GABAergic interneurons, parvalbuminergic interneurons and calretinergic interneurons receive cortical input^{81,82}, suggesting another potential site for learning-related plasticity. Plasticity is likely regulated by dopamine, classically linked to reinforcement learning in the basal ganglia⁸³. Such demonstrations of plasticity in the striatum, particularly within inputs from cortex, suggests this is a fruitful area of study to further understand learning.

Prior to this thesis, it had been demonstrated that neural activity across the corticostriatal network becomes more coordinated and consistent as behaviors become more consistent with learning^{74,77}.

The goal of work in Chapter 2 was to understand, during learning of a complex coordinated action, what aspects of the behavior are controlled (or not controlled) by this emerging coordinated activity. We found that skilled proximal movements are encoded by coordinated corticostriatal activity, while skilled distal movements, such as grasping, are uniquely cortically represented. The goal of work in Chapter 3 was to understand how and when the corticostriatal network is shaped during learning, with a particular focus on the role of sleep which has been linked to procedural learning⁸⁴⁻⁹⁰. We found that non-REM sleep, generally, and sleep spindles, specifically, play a central role in shaping the corticostriatal network during learning and impact learned behavior. The goal of Chapter 4 is to present a model for the motor network and outlines a potential explanation for why corticostriatal plasticity is a critical regulator of learning.

References

1. Shik, M. L. & Orlovsky, G. N. Neurophysiology of locomotor automatism. *Physiological Reviews* **56**, 465–501 (1976).
2. Gatto, G. & Goulding, M. Locomotion Control: Brainstem Circuits Satisfy the Need for Speed. *Curr. Biol.* **28**, R256–R259 (2018).
3. Gremel, C. M. & Costa, R. M. Orbitofrontal and striatal circuits dynamically encode the shift between goal-directed and habitual actions. *Nat. Commun.* (2013). doi:10.1038/ncomms3264
4. Parsons, L. M., Sergent, J., Hodges, D. A. & Fox, P. T. The brain basis of piano performance. in *Neuropsychologia* (2005). doi:10.1016/j.neuropsychologia.2004.11.007
5. Finger, S. *Minds Behind the Brain: A history of the pioneers and their discoveries.* (2010). doi:10.1093/acprof:oso/9780195181821.001.0001
6. Newton, I. *Philosophiae naturalis principia mathematica. Philosophiae naturalis principia mathematica* (1687). doi:10.5479/sil.52126.39088015628399
7. Newton, I. GENERAL SCHOLIUM. in *The Principia: The Authoritative Translation and Guide* (2019). doi:10.1525/9780520964815-025
8. Brazier, M. *A History of Neurophysiology in the 17th and 18th Centuries.* (Raven Press, 1984).
9. Galvani, L. *Aloysii Galvani De viribus electricitatis in motu musculari commentarius. Aloysii Galvani De viribus electricitatis in motu musculari commentarius* (1791). doi:10.5479/sil.324681.39088000932442
10. Parent, A. Giovanni Aldini: From animal electricity to human brain stimulation. *Canadian*

- Journal of Neurological Sciences* (2004). doi:10.1017/S0317167100003851
11. Brazier, M. *A History of Neurophysiology in the 19th Century*. (Raven Press, 1988).
 12. Neuburger, M. & Clarke, E. *The Historical Development of Experimental Brain and Spinal Cord Physiology Before Flourens*. *Annals of Internal Medicine* **96**, (1982).
 13. Gross, C. G. The discovery of motor cortex and its background. *Journal of the History of the Neurosciences* (2007). doi:10.1080/09647040600630160
 14. Taylor, C. S. R. & Gross, C. G. Twitches versus movements: A story of motor cortex. *Neuroscientist* (2003). doi:10.1177/1073858403257037
 15. G, F. & E, H. 'Ueber die elektrische Erregbarkeit des Grosshirns'. *Arch. fur Anat. und Physiol. Wissenschaftliche Med.* **37**, 300–332 (1870).
 16. Hermann, L. *Ueber elektrische Reizversuche an der Grosshirnrinde*. (1875).
 17. LEWES, G. H. The Functions of the Brain. *Nature* **15**, 93–95 (1876).
 18. Ferrier, D. Experimental Researches in Cerebral Physiology and Pathology. *J. Anat. Physiol.* **8**, 152–155 (1873).
 19. Leyton, A. S. F. & Sherrington, C. S. OBSERVATIONS ON THE EXCITABLE CORTEX OF THE CHIMPANZEE, ORANG-UTAN, AND GORILLA. *Q. J. Exp. Physiol.* (1917). doi:10.1113/expphysiol.1917.sp000240
 20. NATHAN, P. W. & SMITH, M. C. LONG DESCENDING TRACTS IN MAN: I. REVIEW OF PRESENT KNOWLEDGE. *Brain* **78**, 248–303 (1955).
 21. Morcos, J. J. & Haines, S. J. History of brain stem surgery. *Neurosurg. Clin. N. Am.* (1993). doi:10.1016/s1042-3680(18)30565-5
 22. Glickstein, M., Strata, P. & Voogd, J. Cerebellum: history. *Neuroscience* **162**, 549–559 (2009).

23. Flourens, P. *Recherches expérimentales sur les propriétés et les fonctions du système nerveux dans les animaux vertébrés*. Baillière ed. (1842).
24. Roh, J., Cheung, V. C. K. & Bizzi, E. Modules in the brain stem and spinal cord underlying motor behaviors. *J. Neurophysiol.* (2011). doi:10.1152/jn.00842.2010
25. Ruder, L. & Arber, S. Brainstem Circuits Controlling Action Diversification. *Annual Review of Neuroscience* (2019). doi:10.1146/annurev-neuro-070918-050201
26. Azim, E., Jiang, J., Alstermark, B. & Jessell, T. M. Skilled reaching relies on a V2a propriospinal internal copy circuit. *Nature* **508**, 357–63 (2014).
27. Esposito, M. S., Capelli, P. & Arber, S. Brainstem nucleus MdV mediates skilled forelimb motor tasks. *Nature* **508**, 351–356 (2014).
28. Rizzi, G., Coban, M. & Tan, K. R. Excitatory rubral cells encode the acquisition of novel complex motor tasks. *Nat. Commun.* (2019). doi:10.1038/s41467-019-10223-y
29. Baker, S. N., Zaaime, B., Fisher, K. M., Edgley, S. A. & Soteropoulos, D. S. Pathways mediating functional recovery. in *Progress in Brain Research* (2015). doi:10.1016/bs.pbr.2014.12.010
30. Shmuelof, L. & Krakauer, J. W. Are we ready for a natural history of motor learning? *Neuron* **72**, 469–76 (2011).
31. Shadmehr, R., Smith, M. A. & Krakauer, J. W. Error correction, sensory prediction, and adaptation in motor control. *Annual Review of Neuroscience* (2010). doi:10.1146/annurev-neuro-060909-153135
32. Evarts, E. V. Pyramidal tract activity associated with a conditioned hand movement in the monkey. *J. Neurophysiol.* (1966). doi:10.1152/jn.1966.29.6.1011
33. Ebner, T. J., Hendrix, C. M. & Pasalar, S. Past, present, and emerging principles in the

- neural encoding of movement. *Adv. Exp. Med. Biol.* (2009). doi:10.1007/978-0-387-77064-2_7
34. Georgopoulos, A. P. & Carpenter, A. F. Coding of movements in the motor cortex. *Current Opinion in Neurobiology* (2015). doi:10.1016/j.conb.2015.01.012
 35. Churchland, M. M. *et al.* Neural population dynamics during reaching. *Nature* **487**, 51–6 (2012).
 36. Shenoy, K. V., Sahani, M. & Churchland, M. M. Cortical control of arm movements: A dynamical systems perspective. *Annual Review of Neuroscience* (2013). doi:10.1146/annurev-neuro-062111-150509
 37. Cisek, P. Preparing for speed. Focus on ‘Preparatory activity in premotor and motor cortex reflects the speed of the upcoming reach’. *Journal of Neurophysiology* (2006). doi:10.1152/jn.00857.2006
 38. Michaels, J. A., Dann, B. & Scherberger, H. Neural Population Dynamics during Reaching Are Better Explained by a Dynamical System than Representational Tuning. *PLoS Comput. Biol.* (2016). doi:10.1371/journal.pcbi.1005175
 39. Jackson, A. Untangling Neural Representations in the Motor Cortex. *Neuron* **97**, 736–738 (2018).
 40. Lawrence, D. G. & Kuypers, H. G. The functional organization of the motor system in the monkey. II. The effects of lesions of the descending brain-stem pathways. *Brain* **91**, 15–36 (1968).
 41. Lawrence, D. G. & Kuypers, H. G. The functional organization of the motor system in the monkey. I. The effects of bilateral pyramidal lesions. *Brain* **91**, 1–14 (1968).
 42. Peters, A. J., Liu, H. & Komiyama, T. Learning in the Rodent Motor Cortex. *Annu. Rev.*

- Neurosci.* **40**, 77–97 (2017).
43. Alaverdashvili, M. & Whishaw, I. Q. Motor cortex stroke impairs individual digit movement in skilled reaching by the rat. *Eur. J. Neurosci.* **28**, 311–322 (2008).
 44. Kawai, R. *et al.* Motor Cortex Is Required for Learning but Not for Executing a Motor Skill. *Neuron* **86**, 800–812 (2015).
 45. Barthas, F. & Kwan, A. C. Secondary Motor Cortex: Where ‘Sensory’ Meets ‘Motor’ in the Rodent Frontal Cortex. *Trends in Neurosciences* (2017). doi:10.1016/j.tins.2016.11.006
 46. Kupferschmidt, D. A., Juczewski, K., Cui, G., Johnson, K. A. & Lovinger, D. M. Parallel, but Dissociable, Processing in Discrete Corticostriatal Inputs Encodes Skill Learning. *Neuron* **96**, 476-489.e5 (2017).
 47. Yin, H. H. *et al.* Dynamic reorganization of striatal circuits during the acquisition and consolidation of a skill. *Nat. Neurosci.* **12**, 333–341 (2009).
 48. Nauta, W. J. H. & Mehler, W. R. Projections of the lentiform nucleus in the monkey. *Brain Research* (1966). doi:10.1016/0006-8993(66)90103-X
 49. Wilson, C. J. The History of the Basal Ganglia: Cells and Circuits. in *Handbook of Behavioral Neuroscience* (2016). doi:10.1016/B978-0-12-802206-1.00003-9
 50. Chakravarthy, V. S., Joseph, D. & Bapi, R. S. What do the basal ganglia do? A modeling perspective. *Biological Cybernetics* (2010). doi:10.1007/s00422-010-0401-y
 51. Willis, T. *Cerebri anatome: cui accessit nervorum descriptio et usus.* (1664).
 52. WILSON, S. A. K. AN EXPERIMENTAL RESEARCH INTO THE ANATOMY AND PHYSIOLOGY OF THE CORPUS STRIATUM. *Brain* (1914). doi:10.1093/brain/36.3-4.427
 53. Graybiel, A. M. Habits, Rituals, and the Evaluative Brain. *Annu. Rev. Neurosci.* **31**, 359–

- 387 (2008).
54. Graybiel, A. M. & Grafton, S. T. The striatum: where skills and habits meet. *Cold Spring Harb. Perspect. Biol.* **7**, a021691 (2015).
 55. Panigrahi, B. *et al.* Dopamine Is Required for the Neural Representation and Control of Movement Vigor. *Cell* **162**, 1418–1430 (2015).
 56. Grillner, S., Robertson, B. & Stephenson-Jones, M. The evolutionary origin of the vertebrate basal ganglia and its role in action selection. *Journal of Physiology* (2013). doi:10.1113/jphysiol.2012.246660
 57. Klaus, A., Alves Da Silva, J. & Costa, R. M. What, If, and When to Move: Basal Ganglia Circuits and Self-Paced Action Initiation. *Annual Review of Neuroscience* (2019). doi:10.1146/annurev-neuro-072116-031033
 58. Park, J., Coddington, L. T. & Dudman, J. T. Basal Ganglia Circuits for Action Specification. *Annual Review of Neuroscience* (2020). doi:10.1146/annurev-neuro-070918-050452
 59. Rueda-Orozco, P. E. & Robbe, D. The striatum multiplexes contextual and kinematic information to constrain motor habits execution. *Nat. Neurosci.* **18**, 453–460 (2015).
 60. Desmurget, M. & Turner, R. S. Testing Basal Ganglia Motor Functions Through Reversible Inactivations in the Posterior Internal Globus Pallidus. *J. Neurophysiol.* **99**, 1057–1076 (2008).
 61. Goldberg, J. H., Farries, M. A. & Fee, M. S. Basal ganglia output to the thalamus: still a paradox. *Trends Neurosci.* **36**, 695–705 (2013).
 62. Oldenburg, I. A. & Sabatini, B. L. Antagonistic but Not Symmetric Regulation of Primary Motor Cortex by Basal Ganglia Direct and Indirect Pathways. *Neuron* **86**, 1174–1181 (2015).

63. Lalive, A. L., Lien, A. D., Roseberry, T. K., Donahue, C. H. & Kreitzer, A. C. Motor thalamus supports striatum-driven reinforcement. *Elife* (2018). doi:10.7554/eLife.34032
64. Popescu, A. T., Saghyan, A. A. & Paré, D. NMDA-dependent facilitation of corticostriatal plasticity by the amygdala. *Proc. Natl. Acad. Sci. U. S. A.* (2007). doi:10.1073/pnas.0609831104
65. Stringer, C. *et al.* Spontaneous behaviors drive multidimensional, brainwide activity. *Science* (80-.). (2019). doi:10.1126/science.aav7893
66. Willingham, D. B. A Neuropsychological Theory of Motor Skill Learning. *Psychol. Rev.* (1998). doi:10.1037/0033-295X.105.3.558
67. Hunnicutt, B. J. *et al.* A comprehensive excitatory input map of the striatum reveals novel functional organization. *Elife* (2016). doi:10.7554/eLife.19103
68. Dudman, J. T. & Gerfen, C. R. The Basal Ganglia. *Rat Nerv. Syst.* 391–440 (2015). doi:10.1016/B978-0-12-374245-2.00017-6
69. Flaherty, A. W. & Graybiel, A. M. Two input systems for body representations in the primate striatal matrix: Experimental evidence in the squirrel monkey. *J. Neurosci.* (1993). doi:10.1523/jneurosci.13-03-01120.1993
70. Calabresi, P., Pisani, A., Mercuri, N. B. & Bernardi, G. Long-term Potentiation in the Striatum is Unmasked by Removing the Voltage-dependent Magnesium Block of NMDA Receptor Channels. *Eur. J. Neurosci.* (1992). doi:10.1111/j.1460-9568.1992.tb00119.x
71. Charpier, S. & Deniau, J. M. In vivo activity-dependent plasticity at cortico-striatal connections: Evidence for physiological long-term potentiation. *Proc. Natl. Acad. Sci. U. S. A.* (1997). doi:10.1073/pnas.94.13.7036
72. Dang, M. T. *et al.* Disrupted motor learning and long-term synaptic plasticity in mice

- lacking NMDAR1 in the striatum. *Proc. Natl. Acad. Sci. U. S. A.* (2006).
doi:10.1073/pnas.0601758103
73. Koralek, A. C., Jin, X., Long II, J. D., Costa, R. M. & Carmena, J. M. Corticostriatal plasticity is necessary for learning intentional neuroprosthetic skills. *Nature* **483**, 331–335 (2012).
 74. Santos, F. J., Oliveira, R. F., Jin, X. & Costa, R. M. Corticostriatal dynamics encode the refinement of specific behavioral variability during skill learning. *Elife* **4**, (2015).
 75. Xiong, Q., Znamenskiy, P. & Zador, A. M. Selective corticostriatal plasticity during acquisition of an auditory discrimination task. *Nature* (2015). doi:10.1038/nature14225
 76. Costa, R. M., Cohen, D. & Nicolelis, M. A. L. Differential Corticostriatal Plasticity during Fast and Slow Motor Skill Learning in Mice. *Curr. Biol.* **14**, 1124–1134 (2004).
 77. Koralek, A. C., Costa, R. M. & Carmena, J. M. Temporally Precise Cell-Specific Coherence Develops in Corticostriatal Networks during Learning. *Neuron* **79**, 865–872 (2013).
 78. Neely, R. M., Koralek, A. C., Athalye, V. R., Costa, R. M. & Carmena, J. M. Volitional Modulation of Primary Visual Cortex Activity Requires the Basal Ganglia. *Neuron* **0**, (2018).
 79. Kawai, R. *et al.* Motor Cortex Is Required for Learning but Not for Executing a Motor Skill. *Neuron* **86**, 800–812 (2015).
 80. Wolff, S. B. E., Ko, R. & Ölveczky, B. P. Distinct roles for motor cortical and thalamic inputs to striatum during motor learning and execution. *bioRxiv* 825810 (2019).
doi:10.1101/825810
 81. Dimova, R., Vuillet, J., Nieoullon, A. & Goff, L. K. Le. Ultrastructural features of the choline acetyltransferase-containing neurons and relationships with nigral dopaminergic

- and cortical afferent pathways in the rat striatum. *Neuroscience* (1993). doi:10.1016/0306-4522(93)90489-3
82. Tepper, J. M., Tecuapetla, F., Koós, T. & Ibáñez-Sandoval, O. Heterogeneity and diversity of striatal GABAergic interneurons. *Frontiers in Neuroanatomy* (2010). doi:10.3389/fnana.2010.00150
 83. Fisher, S. D. *et al.* Reinforcement determines the timing dependence of corticostriatal synaptic plasticity in vivo. *Nat. Commun.* (2017). doi:10.1038/s41467-017-00394-x
 84. Yang, G. *et al.* Sleep promotes branch-specific formation of dendritic spines after learning. *Science* (80-.). **344**, 1173–1178 (2014).
 85. Ramanathan, D. S., Gulati, T. & Ganguly, K. Sleep-Dependent Reactivation of Ensembles in Motor Cortex Promotes Skill Consolidation. *PLOS Biol.* **13**, e1002263 (2015).
 86. Gulati, T., Ramanathan, D. S., Wong, C. C. & Ganguly, K. Reactivation of emergent task-related ensembles during slow-wave sleep after neuroprosthetic learning. *Nat. Neurosci.* **17**, 1107–1113 (2014).
 87. Doyon, J. & Benali, H. Reorganization and plasticity in the adult brain during learning of motor skills. *Curr. Opin. Neurobiol.* **15**, 161–167 (2005).
 88. Rasch, B. & Born, J. About Sleep's Role in Memory. *Physiol. Rev.* **93**, 681–766 (2013).
 89. Fischer, S., Hallschmid, M., Elsner, A. L. & Born, J. Sleep forms memory for finger skills. *Proc. Natl. Acad. Sci.* **99**, 11987–11991 (2002).
 90. Walker, M. P., Brakefield, T., Morgan, A., Hobson, J. A. & Stickgold, R. Practice with Sleep Makes Perfect. *Neuron* **35**, 205–211 (2002).

Chapter 2: What does the corticostriatal network encode?

Chapter 2 is adapted from Lemke, S. M., Ramanathan, D. S., Guo, L., Won, S. J. & Ganguly, K. Emergent modular neural control drives coordinated motor actions. Nat. Neurosci. 22, 1122–1131 (2019). doi:10.1038/s41593-019-0407-2

Abstract

A remarkable feature of motor control is the ability to coordinate movements across independent body parts, or “effectors”, into a consistent, skilled action. To reach and grasp an object, effectors for “gross” arm and “fine” dexterous movements must be coordinated as a single action. How the nervous system achieves this coordination remains unknown. One possibility is that, with training, effectors are co-optimized to produce a coordinated action; alternatively, effectors may be modularly refined to function together. To address this question, we recorded neural activity in primary motor cortex (M1) and dorsolateral striatum (DLS) during reach-to-grasp skill learning in rats. With training, coordinated activity across M1 and DLS emerged, linked to the appearance of fast and consistent gross movements. Surprisingly, the emergence of skilled fine movements was independent of such activity, evolved over a longer timescale, and displayed a stronger cortical reliance. Consistent with these results, inactivation of M1 and DLS had distinct effects on skilled fine and gross movements. Our results indicate that skilled movement coordination can be achieved through emergent modular neural control.

Introduction

The deceptively simple act of reaching and grasping an object requires the precise coordination of both “gross” movements of the arm and “fine” movements of the fingers. Each of these distinct body parts, or “effectors”, plays a different role in the action and has distinct complexities in its control¹⁻³. For example, there are distinct degrees-of-freedom in movements of the arm and hand. How, then, does the nervous system coordinate such effectors to produce a unified skilled action? It has been suggested that such multi-effector coordination is achieved by globally optimizing movements with respect to biologically relevant task goals⁴⁻⁸. For example, in reaching and grasping, both fine and gross movements may be jointly optimized to achieve task success while minimizing parameters such as effort. Surprising little, however, is known about the emerging neural basis of such coordination during skill learning.

While many tasks have been used to study the neural basis of skill learning (e.g. reaching and grasping⁹⁻¹², lever pressing^{13,14}, accelerating rotarod¹⁵⁻¹⁷), learning is often measured by a global task parameter rather than changes in the actual movements involved (however, see Kawai, et al., 2015, Rueda-Orozco, et al., 2015). For example, while rodent reach-to-grasp skill learning requires the coordination of both fine and gross movements³, learning is commonly assessed using overall success rate rather than detailed analyses of the fine and gross movement components. Thus, a key goal of this study was to establish how changes in parameters such as success rate are achieved through changes in the underlying movements involved and, further, to determine the neural basis for the emergence of such skilled movements.

One possibility is that the emerging neural basis of multi-effector coordination reflects theories positing the global optimization of movements, i.e., changes in success rate are achieved through the co-optimization of both fine and gross movements. In this case, during reach-to-grasp skill learning, we would expect a pattern of neural activity linked to both fine and gross movements to emerge as fine and gross movements are refined concurrently. Alternatively, however, coordination may be achieved in a distributed fashion. In this case, we would expect modular patterns of neural activity to emerge independently that represent the control of fine or gross movements specifically. We hypothesized that monitoring neural activity across the motor network during learning of a multi-effector skill would allow us to distinguish between these possibilities.

Here, we report that effector-specific neural control emerged as a coordinated motor action was learned. We recorded neural activity in primary motor cortex (M1) and dorsolateral striatum (DLS), the primary striatal target of M1^{17,18}, throughout learning of a reach-to-grasp skill in rats. We observed that coordinated low-frequency activity emerged across M1 and DLS, linked to the emergence of fast and consistent gross movements. Intriguingly, the emergence of skilled fine movements was independent of this activity, evolved over a longer timescale, and displayed a stronger cortical reliance. Consistent with these results, inactivation of M1 and DLS had distinct effects on skilled fine and gross movements. Together, our results indicate that skilled movement coordination can be achieved through emergent modular neural control.

Results

We recorded neural signals, including single-unit activity and local field potentials (LFP) in M1 and DLS (Figure 2.1; Figure 2.2) as rats ($n=4$) learned a reach-to-grasp skill. Rats were trained for eight days using automated behavioral boxes, performing 75-150 trials each day¹⁹. Refinement of both “gross” movements of forearm, for an accurate reaching action, and “fine” movements of the digits, to successfully grasp the pellet, are required to learn this skill (Figure 2.3a). Consistent with past results^{9,10}, over eight days of learning, training resulted in faster and more consistent movements, as well as increased success rate (Figure 2.3b&c; reach duration: 874 ± 203 ms on day one to 262 ± 10 ms on day eight, mean \pm SEM across animals hereafter, mixed-effects models hereafter (see methods): $t(913) = -16.6$, $P = 3.6 \times 10^{-54}$; sub-movement timing variability: 281 ± 97 ms to 66 ± 34 ms, $t(913) = -4.4$, $P = 1.7 \times 10^{-5}$; forearm trajectory consistency: 0.86 ± 0.02 to 0.92 ± 0.02 mean correlation value, $t(516) = 4.4$, $P = 1.5 \times 10^{-5}$; success rate: $25.2 \pm 9.9\%$ to $51.4 \pm 10.3\%$, $t(913) = 9.4$, $P = 5.1 \times 10^{-20}$).

Refinement of skilled fine and gross movements is dissociable during skill learning

We first sought to determine how changes in success rate were related to changes in fine or gross movements. Intriguingly, we observed that success rate and changes in gross forearm movements,




Figure 2.1. Localization of electrodes. a. Illustration of M1 and DLS recording sites. b. Quantification of electrolytic lesion sites marking electrode locations for four learning animals.




Figure 2.2. Corticostriatal projections. **a.** Anterograde labeling of projections from M1 showing projections to the DLS. **b.** Retrograde labeling of projections to the DLS showing strong inputs from cortex, including M1.

measured by reach duration, sub-movement timing variability, and forearm trajectory consistency, seemed to evolve on different timescales. While measures of gross forearm movements reached a plateau within eight days, success rate remained variable (**Figure 2.3b**, D5-D8 in gray box). This dissociation suggested that the control of gross movements may stabilize while the control of fine movements of the digits remains variable, resulting in variable success rate. In fact, we observed that differences in forearm movements did not account for success on days five through eight as we found no significant differences between reach duration, sub-movement timing variability, or forelimb trajectory consistency for successful and unsuccessful trials on these days (**Figure 2.3d**; reach duration: $355 \pm 48\text{ms}$ for successful trials and $327 \pm 35\text{ms}$ for unsuccessful trials, $t(1310) = 0.3$, $P = 0.76$; sub-movement timing variability: $101.1 \pm 24.7\text{ms}$ and $100.6 \pm 34.9\text{ms}$, $t(114) = 1.1$, $P = 0.29$; forearm trajectory consistency: 0.92 ± 0.02 and 0.92 ± 0.01 mean correlation value, $t(301) = -0.001$, $P = 0.99$). Furthermore, success rate did not covary with measures of gross movements on days five through eight of training (**Figure 2.4**; reach duration: $p = 0.21$, $R = 0.11$; sub-movement timing variability: $p = 0.26$, $R = 0.10$).

Importantly, the control of skilled fine movements continued to evolve over a longer time scale after gross movements stabilized. In a separate “extended training” cohort (n=3), performing ~2500 trials over 4 weeks, average success rate reached a significantly higher rate than our “learning cohort” reached in eight days, while reach duration, sub-movement timing variability, and forearm trajectory consistency were not significantly different between cohorts (**Figure 2.3c**; reach duration: $262 \pm 10\text{ms}$ for learning cohort and $279 \pm 39\text{ms}$ for extended training cohort, $t(714) = 0.49$, $P = 0.62$; sub-movement timing variability: $66 \pm 34\text{ms}$ and $125 \pm 22.1\text{ms}$, $t(135) = 1.5$, $P = 0.12$; forearm trajectory consistency: 0.92 ± 0.02 and 0.91 ± 0.01 mean correlation value, $t(460) = 0.07$, $P = 0.94$; success rate: $51.4 \pm 10.3\%$ and $78.9 \pm 8.4\%$, $t(714) = 2.5$, $P = 0.01$). Altogether, this indicated that the refinement of skilled fine and gross movements is dissociable during reach-to-grasp skill learning.




Figure 2.3. Refinement of skilled fine and gross movements is dissociable during reach-to-grasp skill learning. **a.** Diagram of reach-to-grasp task and learning metrics: duration from movement onset (MO) to pellet touch (PT), duration from MO to retract onset (RO), forearm trajectory correlation, and success rate. **b.** Example learning curve (dots represent individual trials, lines are averaged over 30 trials; forearm trajectories are shown from day one and eight, individual trial trajectories in grey and mean trajectory in yellow). **c.** Difference in reach duration, sub-movement timing variability, forearm trajectory correlation, and success rate between day one of training (D1), day eight (D8), and performance from an extended training cohort (EXT). **d.** Differences in reach duration, sub-movement timing variability, and forearm trajectory correlation between successful (Suc.) and unsuccessful (Unsuc.) trials on days 5-8.




Figure 2.4. Fine and gross movements do not covary between days five through eight of reach-to-grasp skill learning. **a.** Scatterplots of mean movement duration, sub-movement timing variability and success rate for non-overlapping ten trial bins between days five and eight of training across animals. Values are z-scored within each animal.

Coordinated movement-related activity emerges across M1 and DLS during skill learning

We first explored the neural basis for the refinement of skilled gross movements. As gross movements were refined, consistently timed sub-movements and a multiphasic profile of forelimb velocity emerged (**Figure 2.5a**, top/middle). We observed rhythmic movement-related neural activity across M1 and DLS that reflected the precise timing of sub-movements. Specifically, we observed that coordinated low-frequency (~3-6Hz) activity emerged during movement across M1 and DLS that was closely related to the timing of sub-movements and forearm muscle activity, which also displayed a similar low-frequency component (**Figure 2.5a**, bottom).

Movement-related LFP signals displayed the emergence of coordinated low-frequency activity across M1 and DLS during skill learning. A clear increase in movement-related LFP power was observed from day one to eight in both M1 and DLS in the 3-6Hz frequency range (**Figure 2.5b**; M1: 1.0 ± 0.13 baseline normalized power on day one to 1.74 ± 0.1 on day eight, $t(146) = 9.1$, $P = 5.0 \times 10^{-16}$; DLS: 1.0 ± 0.7 to 1.67 ± 0.1 , $t(94) = 6.4$, $P = 5.1 \times 10^{-9}$). Movement-related LFP coherence between M1 and DLS LFP also significantly increased in the 3-6Hz frequency range (**Figure 2.5c**; 0.18 ± 0.03 coherence on day one to 0.24 ± 0.03 on day eight, $t(870) = 9.1$, $P = 9 \times 10^{-19}$). As striatal LFP signals can contain volume-conducted cortical signals with zero phase-lag coherence²⁰, we measured the phase-lag for high-coherence M1 and DLS channel pairs (i.e., coherence value > 0.2) and found phase lag values consistent with the connectivity between M1 and DLS and inconsistent with volume conducted signals (**Figure 2.6**). Additionally, increases in LFP power and coherence were not solely a result of faster and more consistent movements, as we observed significant increases in LFP power and coherence for behaviorally matched fast and consistent trials early and late in learning (**Figure 2.7**).




Figure 2.5. Coordinated movement-related activity emerges across M1 and DLS during skill learning. **a.** Diagram of reach-to-grasp task and learning metrics: duration from movement onset (MO) to pellet touch (PT), duration from MO to retract onset (RO), forearm trajectory correlation, and success rate. **b.** Example learning curve (dots represent individual trials, lines are averaged over 30 trials; forearm trajectories are shown from day one and eight, individual trial trajectories in grey and mean trajectory in yellow). **c.** Difference in reach duration, sub-movement timing variability, forearm trajectory correlation, and success rate between day one of training (D1), day eight (D8), and performance from an extended training cohort (EXT). **d.** Differences in reach duration, sub-movement timing variability, and forearm trajectory correlation between successful (Suc.) and unsuccessful (Unsuc.) trials on days 5-8.




Figure 2.6. Non-zero phase lag between high coherence M1 and DLS LFP signals. a. Example mean movement onset-locked LFP signals from M1 and DLS (top) and extracted phase (bottom). **b.** Histogram of phase difference between M1 and DLS LFP signals from 250ms before to 500ms after movement onset for example LFP pair (grey shaded region in a). **c.** Peak of phase difference histogram (as in b) for all high coherence (coherence value >0.2) pairs of M1-DLS channels on day eight.




Figure 2.7. LFP power and coherence increases for behaviorally-matched trials during learning. a. Differences in M1 3-6Hz LFP power, DLS 3-6Hz LFP power, and M1-DLS 3-6Hz LFP coherence between behaviorally-matched “fast” trials with a duration between 200-400ms on days one and two (Early Days) and days seven and eight (Late Days; grey lines represent individual animals and black lines represent mean and SEM across animals). **b.** Same as a for behaviorally-matched high forelimb trajectory correlation trials (individual trial trajectory with correlation value >0.9 to mean session trajectory).

We next characterized the relationship between emerging coordinated low-frequency activity across M1 and DLS and the timing of sub-movements. With training, sub-movement timing became precisely phase-locked to 3-6Hz LFP activity across both M1 and DLS, consistent with what we would expect if this activity was involved in generating sub-movements²¹⁻²⁴ (**Figure 2.5d**; significant increase in inter-trial coherence (ITC) of M1 LFP signals locked to movement onset, $t(102) = 3.8$, $P = 2 \times 10^{-4}$; M1 LFP signals locked to pellet touch, $t(102) = 4.7$, $P = 1 \times 10^{-7}$; M1 LFP signals locked to retract onset, $t(102) = 8.5$, $P = 2 \times 10^{-13}$; DLS LFP signals locked to movement onset, $t(96) = 9.6$, $P = 1 \times 10^{-15}$; DLS LFP signals locked to pellet touch, $t(96) = 6.6$, $P = 3 \times 10^{-9}$; DLS LFP signals locked to retract onset, $t(96) = 12.4$, $P = 1 \times 10^{-23}$). Given this close relationship between LFP activity and sub-movement timing, we next explored whether the frequency of LFP coherence between M1 and DLS covaried with movement duration. In fact, the peak frequency of LFP coherence on day eight was significantly higher for faster movements compared to slower movements (**Figure 2.8**). Together, these results indicated that coordinated low-frequency activity across M1 and DLS was closely linked to skilled gross movements.




Figure 2.8. Peak frequency of M1-DLS LFP coherence covaries with movement duration on day eight. **a.** Example mean M1 and DLS LFP signals for slowest third of trials (top) and fastest third of trials (bottom) on day eight for example animal. **b.** Difference in 3-6Hz M1-DLS LFP coherence between slowest third of trials (Slow Trials) and fastest third of trials (Fast Trials; normalized to mean coherence for slowest third of trials) on day one and day eight ($n = 4$ animals). Grey lines represent mean values from individual animals and black lines represent mean and SEM across animals. P values from mixed-effects models. **c.** Difference in peak M1-DLS frequency between 3-6Hz for slowest third of trials and fastest third of trials on day one and day eight ($n = 4$ animals). Grey lines represent mean values from individual animals and black lines represent mean and SEM across animals. P values from mixed-effects models.

Coordinated spiking activity emerges across M1 and DLS during skill learning

Coordinated low-frequency spiking activity also emerged across M1 and DLS during skill learning. Peri-event time histograms (PETHs) of M1 and DLS units displayed movement-related multiphasic activity locked to 3-6Hz LFP activity (**Figure 2.9a**). We investigated the relationship between movement-related spiking activity and LFP signals by quantifying the phase-locking of M1 and DLS spikes to low-frequency LFP signals. Phase-locking was quantified by generating polar histograms of the LFP phases at which each spike occurred for a single unit and LFP channel filtered in the 3-6Hz band in a one-second window around movement (**Figure 2.9b**). We compared all M1 and DLS units on day one and day eight to the same LFP channel in M1 and DLS. The non-uniformity of these histograms (indicating phase-locking) was quantified using a Raleigh test of circular non-uniformity that produced a z-statistic with a threshold for significance that allowed us to determine the percentage of units in M1 and DLS that were significantly phase-locked to LFP signals in either M1 or DLS. We observed an increase in the percentage of M1 and DLS units phase-locked to both M1 and DLS LFP signals with training (**Figure 2.9c**; black vertical dotted lines correspond to the $p = 0.05$ significance threshold of the natural log of the z-statistic, all unit-LFP pairs with z-statistics greater than this threshold were significantly phase-locked; M1 unit - M1 LFP pairs: 37.3% day one to 76% day eight, $p = 1 \times 10^{-6}$, Kolmogorov-Smirnov test; M1 unit - DLS LFP pairs: 35.0% to 56.8%, $p = 0.03$; DLS unit - M1 LFP pairs: 42.9% to 66.0%, $p = 0.03$; DLS unit - DLS LFP pairs: 24.2% to 66.3%, $p = 6 \times 10^{-8}$). We tested whether LFP phase-locked units were active during a specific phase of movement compared to other units and did not observe any clear timing differences in the average responses of phase-locked and not phase-locked units (**Figure 2.10**).

The percentage of units that displayed transient oscillatory activity in the 3-6Hz range during movement also increased with training (**Figure 2.9d**). This increase provided an LFP-independent measure of local changes in quasi-oscillatory activity. Of these quasi-oscillatory units, 47.1% of M1 units on day eight were also phase-locked to M1 LFP, compared to 42.1% on day one, and 53.9% of DLS units on day eight were also phase-locked to DLS LFP, compared to 50.0% on day one. Strikingly, the mean quasi-oscillatory unit cross-correlation across regions on day eight (average of all individual cross correlations of quasi-oscillatory M1 and DLS unit pairs) displayed a peak with a short time lag consistent with the connectivity of M1 and DLS, as well as secondary peaks corresponding to a 3-6Hz rhythm. This spiking relationship was not observed on day one (**Figure 2.9e**). Altogether, these results provided further evidence of emerging movement-related coordination across M1 and DLS during skill learning.




Figure 2.9. Coordinated M1 and DLS spiking activity emerges during skill learning. **a.** Example Peri-Event Time Histograms (PETH) from units in M1 (left) and DLS (right) displaying multi-phasic activity locked to 3-6Hz LFP activity. **b.** Diagram of spike-LFP phase locking. *Top:* Raster plot of example M1 unit spiking activity during movement aligned to movement onset (MO). *Middle:* Example unit PETH with M1 LFP activity overlaid and extracted 3-6Hz LFP phase. *Bottom:* Polar histogram of LFP phases at which spikes occurred. **c.** Cumulative density functions of z-statistics for every unit-LFP pair across and within each region (vertical dotted lines denote significance threshold of z-statistic at $p < 0.05$, % of respective unit-LFP pairs greater than threshold noted, lighter color is day one). **d.** *Left:* PETHs from example unit displaying multiphasic activity and example unit not displaying multiphasic activity and corresponding autocorrelations used for classifying quasi-oscillatory (QO; arrows denote “bumps” used for classification, see methods) and non-quasi-oscillatory units (Non-QO). *Right:* Percentage of units in M1 and DLS on day one and day eight that display quasi-oscillatory activity. **e.** Mean cross-correlation between all quasi-oscillatory M1 and DLS units on day one (grey) and day eight (blue).




Figure 2.10. Details of M1 and DLS spiking activity. **a.** Example waveforms from array spanning M1 and DLS for one session. **b.** Total recorded units and total task-related units across days in M1. **c.** Total recorded units and total task-related units across days in DLS. **d.** Histogram of firing rates for all units across all sessions, vertical lines denote median for each region. **e.** PETHs of all M1 units time-locked to movement onset on day one (left) and day eight (right). **f.** PETHs of all DLS units time-locked to movement onset on day one (left) and day eight (right). **g.** M1 unit activity during reaching. **h.** DLS unit activity during reaching. **i.** PETHs of all phase-locked and not phase-locked M1 units on day eight. **j.** PETHs of all phase-locked and not phase-locked DLS units on day eight.

Coordinated M1 and DLS activity is specifically linked to skilled gross movements

Given the close relationship between coordinated low-frequency activity across M1 and DLS and skilled gross movements, we expected their emergence to coincide during learning. In fact, we found that the emergence of movement-related M1-DLS 3-6Hz LFP coherence closely coincided with the transition to fast and consistent gross movements (**Figure 2.11a**). Across animals, we observed a significant correlation between each session's average movement-related 3-6Hz M1-DLS LFP coherence and the average reach duration, sub-movement timing variability, and forearm trajectory correlation of that session (**Figure 2.11b**; reach duration: $p = 3 \times 10^{-5}$, $R = -0.73$; sub-movement timing variability: $p = 2 \times 10^{-3}$, $R = -0.58$; forelimb trajectory correlation: $p = 6 \times 10^{-3}$, $R = 0.52$; Pearson Correlation, $n = 25$ sessions across 4 animals). As we observed variable success rate after the stabilization of gross movements (Fig. 1b, D5-D8 in gray box), we wondered whether coordinated M1-DLS activity was also related to this variability in success rate. We compared movement-related 3-6Hz M1-DLS LFP coherence between successful and unsuccessful trials on days five through eight and found no significant difference (**Figure 2.11a&c**; 0.20 ± 0.03 for successful trials and 0.21 ± 0.03 for unsuccessful trials, $t(2558) = 1.1$, $p = 0.28$). As we attribute whether trials were successful during this period to the control of skilled fine movements, this suggested that the production of skilled fine movements is independent of such activity. These results indicated that emerging coordinated activity across M1 and DLS is specifically linked to skilled gross movements.




Figure 2.11. Coordinated M1 and DLS activity is specifically linked to skilled gross movements. **a.** Time course of movement-related 3-6Hz LFP coherence from example M1-DLS channel pair over training period overlaid with timing of sub-movements and forearm trajectories from day one and day eight. **b.** Scatterplots of each session's mean movement-related 3-6Hz M1-DLS LFP coherence and mean reach duration, sub-movement timing variability, and forelimb trajectory correlation, each normalized per animal. **c.** Filtered LFP (3-6Hz) signals from example M1 and DLS channels for successful and unsuccessful trials on days 5-8, individual trials overlaid with mean signal (left) and difference in M1-DLS LFP coherence for successful and unsuccessful trials on days 5-8 (right; grey lines represent individual animals and black line represents mean and SEM across animals).

M1 and DLS inactivation have differential effects on skilled fine and gross movements

We next sought to causally test the role of M1 and DLS activity in producing skilled fine and gross movements. In a separate cohort of well-trained animals ($n=5$) implanted with infusion cannulas in both M1 and DLS, we acutely inactivated either M1 or DLS by muscimol infusion and observed the effects on skilled movements (**Figure 2.12a**). To dissect impairments of either skilled fine or gross movements we developed a novel reach-to-grasp task design in which the pellet position is either placed at the “far” position (same position used for training) or a “close” position (**Figure 2.12b**). A close position was added to generate a reaching condition in which the reliance on skilled gross movements for success was reduced, while skilled fine movements were still required to successfully grasp the pellet. Differential effects of region inactivation on success rate for the close and far position would therefore indicate differences in skilled fine and gross movement impairment.

Both acute M1 and DLS inactivation disrupted gross movements (**Figure 2.12c**; M1 inactivation/far position: reach duration: $662 \pm 63\text{ms}$ and $2.6 \times 10^3 \pm 667\text{ms}$, $t(311) = 11.2$, $P = 2 \times 10^{-24}$; sub-movement timing variability: $121 \pm 35\text{ms}$ and $586 \pm 228\text{ms}$, $t(53) = 5.7$, $P = 5 \times 10^{-7}$; M1 inactivation/close position: reach duration: $662 \pm 63\text{ms}$ and $2.7 \times 10^3 \pm 711\text{ms}$, $t(312) = 14.8$, $P = 6 \times 10^{-38}$; sub-movement timing variability: $121 \pm 35\text{ms}$ and $641 \pm 89\text{ms}$, $t(54) = 7.0$, $P = 3 \times 10^{-9}$; DLS inactivation/far position: reach duration: $448 \pm 63\text{ms}$ baseline and $889 \pm 279\text{ms}$ post-infusion, $t(725) = 7.5$, $P = 2 \times 10^{-13}$; sub-movement timing variability: $74 \pm 8\text{ms}$ and $237 \pm 77\text{ms}$, $t(136) = 3.7$, $P = 3 \times 10^{-4}$; DLS inactivation/close position: reach duration: $448 \pm 63\text{ms}$ and $995 \pm 284\text{ms}$, $t(735) = 10.0$, $P = 3 \times 10^{-22}$; sub-movement timing variability: $74 \pm 8\text{ms}$ and $195 \pm$

82ms, $t(138) = 4.2$, $P = 5 \times 10^{-5}$). This was consistent with our conclusion that coordinated activity across both M1 and DLS is closely linked to skilled gross movements.

Intriguingly, while both M1 and DLS inactivation also impaired success rate to the far position (**Figure 2.12c**; M1 infusion: $73.6 \pm 7.7\%$ baseline and $6.8 \pm 3\%$ post-infusion, $t(311) = -12.4$, $P = 8 \times 10^{-29}$; DLS infusion: $67.5 \pm 3.8\%$ and $34.7 \pm 8.6\%$, $t(725) = -8.95$, $P = 3 \times 10^{-18}$), only M1 inactivation impaired success rate at the close position (**Figure 2.12c**; M1 infusion: $73.6 \pm 7.7\%$ baseline and $10.8 \pm 6.6\%$ post-infusion, $t(312) = -12.1$, $P = 6 \times 10^{-28}$; DLS infusion: $67.5 \pm 3.8\%$ and $72.2 \pm 2.9\%$, $t(735) = -1.3$, $P = 0.20$). This suggested that while skilled gross movements rely on activity across both M1 and DLS, skilled fine movements display a stronger cortical reliance. No changes in behavior were observed with saline infusions (M1 and DLS saline infusion sessions combined; saline/far position: reach duration: 269 ± 43 ms baseline and 279 ± 53 ms post-infusion, $t(817) = 1.5$, $P = 0.13$; sub-movement timing variability: 135 ± 27 ms baseline and 140 ± 34 ms post-infusion, $t(817) = 0.93$, $P = 0.35$; success rate: $64.6 \pm 2\%$ baseline and $66.5 \pm 7.2\%$ post-infusion, $t(817) = 0.77$, $P = 0.44$).




Figure 2.12. M1 and DLS inactivation have differential effects on skilled fine and gross movements. **a.** Illustration of M1 and DLS muscimol inactivation. **b.** Illustration of two-position reach-to-grasp task design with a “far” pellet and “close” pellet position. **c.** Differences in reach duration, sub-movement timing variability, and success rate between trials before muscimol infusion (Baseline), trials after muscimol infusion reaching to the far pellet position (Far), and trials after muscimol infusion reaching to the close pellet position (Close) for M1 infusions (left) and DLS infusions (right; grey lines represent individual sessions and black lines represents mean and SEM of all sessions).

Behavioral and neural effects of DLS inactivation

We next sought to further dissect the effects of DLS inactivation. With DLS inactivation we observed decreased reach amplitude, compared to an increase in reach amplitude with M1 inactivation (DLS infusion: $1 \pm 1 \times 10^{-16}$ normalized reach amplitude baseline to $0.99 \pm 5 \times 10^{-$

3 normalized reach amplitude post-infusion, $t(593) = -4.1$, $P = 6 \times 10^{-5}$; M1 infusion: $1 \pm 1 \times 10^{-16}$ to 1.01 ± 0.02 , $t(322) = 6.1$, $P = 3 \times 10^{-9}$). This is consistent with work implicating the striatum in movement vigor^{25–30} and suggested that decreased reach amplitude may account for the decrease in success rate at the far position. In fact, when we compared DLS inactivated, post-infusion trials with “normal” reach amplitude (i.e., greater or equal to the average reach amplitude during pre-infusion baseline trials) to pre-infusion baseline trials, success rate was not significantly different ($74.4 \pm 4.4\%$ baseline and $75.33 \pm 11.8\%$ post-infusion, $t(670) = 1.5$, $P = 0.12$). This provided further evidence that DLS inactivation impairs gross movements involved in transporting the paw to the pellet, rather than the fine movements involved in grasping. Importantly, body posture at the time of movement onset was not significantly different after DLS inactivation compared to pre-infusion baseline trials (**Figure 2.13**).




Figure 2.13. Body posture during reaching before and after DLS inactivation. **a.** Example of top camera view and quantification of body axis/posture during reaching. **b.** Example body posture for all trials (posture is quantified at the time of movement onset) before DLS muscimol infusion (Baseline Session) and after DLS muscimol infusion (DLS Inactivation Session). **c.** Difference in posture variability and lateral bias for trials before DLS muscimol infusion and after DLS muscimol infusion (grey lines represent individual animals and black lines represents mean and SEM across animals).

We next sought to examine off-target effects of DLS inactivation. In a separate cohort of well-trained animals ($n=3$) implanted with infusion cannulas in DLS and electrodes in M1, we acutely inactivated DLS by muscimol infusion and observed the effects on M1 activity (**Figure 2.14; Figure 2.15**). With DLS inactivated, there was a significant decrease in movement-related 3-6Hz M1 LFP power compared to pre-infusion baseline (**Figure 2.14b**; $t(318) = 18.1$, $P = 5 \times 10^{-51}$). Intriguingly, this suggested that DLS is required for movement-related low-frequency activity in M1. Importantly, this change was not attributable to a general suppression of M1 activity as we found no significant decrease in movement-related firing rates in M1 with DLS inactivated (**Figure 2.14c**; $t(318) = 18.1$, $P = 0.56$). No changes in movement-related M1 LFP power or firing rate were observed after saline infusions compared to pre-infusion baseline (LFP power, $t(190) = 1.3$, $P = 0.20$; movement-related firing rate, $t(168) = 0.36$, $P = 0.72$). To test whether such acute off-target effects may be causing behavioral impairments, we performed excitotoxic lesions centered on DLS that lesioned DLS as well as small portions of surrounding cortex ($n=3$). Consistent with previous work³¹, we observed increased reach duration, sub-movement timing variability and decreased success rate two weeks post-lesion (**Figure 2.16**; reach duration: 428 ± 56 ms baseline and 631 ± 114 ms post-lesion, $t(620) = 4.7$, $P = 3 \times 10^{-6}$; sub-movement timing variability: 106 ± 13 ms and 296 ± 90 ms, $t(620) = 4.2$, $P = 4 \times 10^{-5}$; success rate: $63 \pm 4.3\%$ and $40.7 \pm 2.9\%$, $t(620) = -5.5$, $P = 5 \times 10^{-8}$). This suggests that acute off-target effects do not fully account for disruptions in skilled gross movements observed with acute DLS inactivation.




Figure 2.14. DLS inactivation decreases movement-related low-frequency M1 activity. **a.** Illustration of DLS muscimol infusion and M1 recording. **b. Left:** 3-6Hz filtered LFP aligned to movement onset from example M1 channel for trials before and after DLS inactivation, individual trials overlaid with mean signal. **Right:** Difference in movement-related 3-6Hz LFP power in M1 before and after DLS inactivation. **c. Left:** PETH from example M1 unit for trials before and after DLS inactivation. **Right:** Difference in movement-related firing rate before and after DLS inactivation.




Figure 2.15. Localization of muscimol infusion.
a. Localization of muscimol infusion into DLS.




Figure 2.16. Localization and behavioral effects of excitotoxic DLS lesions. **a.** DLS excitotoxic lesion localization (immunolabeling: DAPI in blue; NeuN in red). **b.** Differences in reach duration, sub-movement timing variability, and success rate between trials before DLS lesion (Pre Lesion) and trials two weeks after DLS excitotoxic lesion (Post Lesion; lines represent individual animals and black lines represent mean and SEM across animals).

Skilled reaching impairments with chronic M1 lesion

We next sought to further investigate impairments in skilled movements with chronic M1 lesions. It was recently shown that for a complex motor skill composed of non-dexterous movements, lesioning M1 after skill learning did not impair performance¹⁴. However, decreases in success rate have been reported with cortical lesions in the reach-to-grasp skill³². Therefore, we sought to determine if specific aspects of the reach-to-grasp skill are chronically impaired by a permanent M1 lesion. In a separate cohort of well-trained animals ($n=5$), we generated photothrombotic lesions centered on M1 (**Figure 2.17a**; **Figure 2.18**). In the first training session post-lesion (within 8 to 13 days post-lesion), reach duration and sub-movement timing variability were increased and success rate was decreased (**Figure 2.17b**; reach duration: 339 ± 63 ms baseline and

756 ± 130ms early post-lesion, $t(806) = 14.5$, $P = 3 \times 10^{-42}$; sub-movement timing variability: 183 ± 66ms and 365 ± 63ms, $t(195) = 6.7$, $P = 2 \times 10^{-10}$; success rate: 69.1 ± 2.5% and 19.4 ± 7.65%, $t(808) = -15.84$, $P = 2 \times 10^{-49}$). With training, gross movement metrics recovered while success rate remained disrupted. Comparing pre-lesion performance to performance once all behavioral performance measures had plateaued (within 15 to 73 days post-lesion), reach duration and sub-movement timing variability were not significantly different than pre-lesion performance (**Figure 2.17b**; reach duration: 339 ± 63ms baseline and 395 ± 117ms late post-lesion, $t(966) = 2.5$, $P = 0.1$; sub-movement timing variability: 183 ± 66ms and 254 ± 133ms, $t(234) = 2.4$, $P = 0.1$), but success rate remained significantly decreased (**Figure 2.17b**; 69.1 ± 2.5% baseline and 44.7 ± 2.75% late post-lesion, $t(968) = -8.3$, $P = 4 \times 10^{-16}$). We interpreted the differential recovery of skilled fine and gross movements as further evidence that skilled fine movements have a strong cortical reliance.




Figure 2.17. Persistent disruption of skilled fine movements after M1 lesion. **a.** Illustration of photothrombotic M1 lesion. **b.** Differences in reach duration, sub-movement timing variability, and success rate between trials before M1 lesion (Pre Lesion), trials during the first reaching session post-lesion (Early Lesion), and trials once a performance plateau had been reached (Late Lesion; grey lines represent individual animals and black lines represents mean and SEM across animals).




Figure 2.18. Localization of M1 photothrombotic lesion. **a.** Example photothrombotic M1 lesion.

Skilled fine movement representation in M1

Lastly, we explored the representation of skilled fine movements in M1 and DLS. We used gaussian-process factor analysis (GPFA) to find low-dimensional neural trajectory representations of population spiking activity in M1 and DLS on individual trials³³ (Figure 2.19a; Figure 2.20) and then compared trajectories for successful and unsuccessful trials during the period of training after gross movements had stabilized (e.g., Figure 2.3b, D5-D8). As we attribute whether trials were successful during this period to the control of skilled fine movements of the digits, we expected to find a difference in movement-related neural signals between successful and

unsuccessful trials if a region encodes the control of skilled fine movement. Alternatively, if a region does not encode the control of skilled fine movements, we did not expect to find a difference.

Strikingly, we observed a difference between trajectories for successful and unsuccessful trials in M1 but not DLS. To compare successful and unsuccessful trials we subtracted the mean neural trajectory for successful trials, i.e., the “successful template”, from each individual trial’s neural trajectory (**Figure 2.19b**) and calculated the mean absolute value of the deviation during each time point from 250ms before movement onset until pellet touch. We focused on this period as it includes the fine movements involved in shaping the digits for contact with the pellet but does not include differences in retraction or reward between successful and unsuccessful trials. As trials differed in the duration of this period, we interpolated trajectories such that they were all the same length (see methods). We found that M1 neural trajectories for unsuccessful trials had significantly higher deviation than successful trials starting after movement onset (**Figure 2.19c**, top; * = $p < 0.05$, mixed effects model w/Bonferroni correction for multiple comparisons). In DLS, however, deviation of successful and unsuccessful trials from the template did not differ (**Figure 2.19c**, bottom), providing evidence for a specifically cortical representation of skilled fine movements.




Figure 2.19. Skilled fine movement representation in M1. **a.** GPFA neural trajectories for trials on day eight for M1 (top) and DLS (bottom) from example animal. **b.** Illustration of method for calculating deviation from the mean successful template for successful and unsuccessful trials. **c.** Mean deviation (width depicts SEM) from successful template for successful and unsuccessful trials from 250ms before movement onset to pellet touch, across animals (* = $p < 0.05$, mixed-effects model w/Bonferroni correction for multiple comparisons).




Figure 2.20. Increased GPFA neural trajectory consistency during skill learning. **a.** GPFA neural trajectories for M1 (top) and DLS (bottom) on day one and day eight from example animal. **b.** Difference in consistency of GPFA trajectories between early days (days 1-4) and late days (days 5-8) in M1 (top) and DLS (bottom), across animals (grey dots represent individual days and black lines represents mean and SEM across days).

Discussion

In summary, we found that fine and gross movement refinement were behaviorally and neurally differentiable during reach-to-grasp skill learning in rats. Coordinated low-frequency activity emerged across M1 and DLS, linked to the emergence of skilled gross movements. The emergence of skilled fine movements was independent of this activity, evolved over a longer timescale, and displayed a stronger cortical reliance. Consistent with these results, inactivation of either M1 or DLS disrupted gross movements, while only inactivation of M1 disrupted fine movements. This work provides evidence that coordinated skills can emerge from the modular refinement of effectors. Importantly, this work was performed on effectors that vary greatly in degrees of freedom, further work is required to determine whether such modular control generalizes to other forms of coordination (e.g., arm/leg).

The role of M1 in skill learning and execution

Primary motor cortex has been ascribed multiple roles in learning and executing motor skills³⁴. A critical determinant for the role of M1 is whether the skill involves dexterous movements³⁴. Our results indicate that for a motor skill composed of both dexterous (fine) and non-dexterous (gross) movements, distinct patterns of M1 activity emerge independently during skill learning that are related to each component. This suggests that M1 plays dual, but separable, roles in learning a skill composed of both dexterous and non-dexterous components.

Motor cortex is known to play a critical role in the production of dexterous movements^{35,36}. Motor cortical projections to brainstem and spinal cord are thought to mediate this role³⁷⁻³⁹. Lesions to cortex disrupt this pathway and lead to chronic impairments in dexterous movements^{40,41}.

Interestingly, non-dexterous movements have been shown to recover after motor cortical lesion^{40,41}, suggesting a greater ability for gross movements to be generated subcortically. Our results are broadly consistent with this literature. We found a predominantly cortical representation of skilled fine movements (**Figure 2.19**). Moreover, permanent lesions of M1 were associated with chronic fine motor impairments while aspects of gross movements recovered (**Figure 2.17**). The observed deficits in both fine and gross movements with acute M1 inactivation (**Figure 2.12**) raise the possibility of off-target effects⁴².

The role of motor cortex in non-dexterous movements evolves with learning

A recent study found that bilateral lesions of premotor and motor cortex after learning did not impair performance of complex motor skill composed of non-dexterous movements. However, similar lesions before learning prevented skill acquisition¹⁴. This has led to the notion that non-dexterous skills are “transferred” to downstream structures and that motor cortex plays a “tutor” role, i.e., is required for learning but not execution of a motor skill¹⁴. What may allow non-dexterous control in the absence of M1? Skill learning is commonly associated with plasticity in corticostriatal projections^{13,16,17} and it has therefore been theorized that the striatum and thalamostriatal inputs may be critical downstream producers of these learned skills^{14,43,44}. Our results suggest that corticostriatal plasticity is indeed important to learn motor skills (**Figure 2.5**). Additionally, it is possible that the observed coordinated activity across M1 and DLS is a neurophysiological substrate for the motor cortical training or tutor signal. This coordinated low-frequency activity may provide a mechanism through which M1 activity patterns induce long-term plasticity in the DLS. Modeling has shown that temporally patterned inputs to the striatum can drive inter-striatal plasticity⁴⁴.

Importantly, in our work, recovery of skilled gross movements after motor cortical lesion occurred over time with training. This is different from Kawai, et al., 2015, in which full recovery was observed after a period of 10 days with no further training. One important difference that may explain this discrepancy is that the reach-to-grasp task contains both dexterous and non-dexterous components. Another important variable is the amount of training required for transfer. In Kawai, et al., 2015, the complex motor skill was learned over a long period of training (~20,000 training trials over ~30 days); it remains possible that an action containing dexterous movements may also be realized by downstream circuitry after cortical lesions with enough training.

The role of DLS in skill learning and execution

The striatum's role in producing skilled movements is also multifaceted. Receiving widespread cortical innervation, the striatum is positioned to play a central role in motor function as is evidenced by the motor dysfunction caused basal-ganglia disorders⁴⁵. Cortical innervation of the striatum projects back to cortex via basal-ganglia output nuclei and the thalamus, forming a corticostriatal "loop". The precise role of this loop in the production of motor skills remains unclear, although it has been shown that striatal activity can influence motor cortical activity with low-latency⁴⁶ (<200ms). It has also been shown in a Brain Machine Interface (BMI) task in which animals learn to modulate cortical activity to achieve reward that plasticity in corticostriatal projections is required, suggesting that cortex may require striatum to generate stereotyped patterns of activity⁴⁷. Consistent with this notion, we found disrupted movement-related LFP activity in M1 with DLS inactivation (**Figure 2.14**). Altogether, this suggests that the striatum may play a role in modulating cortical activity during skill learning. Additional work detailing the precise effect of

basal ganglia activity on cortical activity will be critical to understand the role of striatum in the production of skilled movements.

Recent work has also championed a role of the striatum in controlling movement vigor²⁵⁻³⁰. One such model focuses on the feed-forward convergent pathway of motor cortical neurons that project to “premotor” brainstem regions and their collaterals that project to the dorsal striatum, that in turn projects to the same “premotor” brainstem regions through basal ganglia output nuclei³⁰. This pathway through the striatum is theorized to play a role in controlling the gain of descending motor commands based on previous experience. One prediction of this “history-dependent gain” model (for details see Yttri & Dudman, 2018) posits that the striatum is active during the acceleration and deceleration phases of a movement, which was confirmed in a joystick²⁷ and locomotor²⁸ task. Similarly, we find that striatal units are active throughout the reaching action (**Figure 2.9; Figure 2.10**). Additionally, we found evidence that disrupting striatal activity effects movement vigor, as reach amplitude was decreased with DLS inactivation. Altogether, our results support the notion that DLS may play a role in both modulating cortical activity and controlling movement vigor for skilled gross movements. Our finding that DLS activity is linked to skilled gross movements in the rodent is consistent with work in the monkey^{48,49} and human^{50,51} basal ganglia, demonstrating a greater representation of proximal, compared to distal, portions of the limb.

What kind of “skill” is the reach-to-grasp task?

The term “motor skill” can describe a vast range of behaviors characterized by fast, accurate, and consistent movements. One group of skills has been formalized using the speed-accuracy tradeoff⁵². Learning such skills typically involves optimizing speed while maintaining accuracy.

For example, one can play a series of notes on the piano consistently and accurately on the first day of learning if the movements are performed slowly enough. Learning therefore involves optimizing speed. Other skills require optimizing accuracy without constraints on than speed (e.g., shooting a free throw). The rodent reach-to-grasp skill falls under this latter type of learning. Interestingly, while there were no explicit constraints on speed, we still observed a significant decrease in movement duration. This may be because the brain's motor network, in general, biases toward faster skills to increase rewards per unit time⁵³. Another possibility is that coordinated low-frequency activity between M1 and DLS is critical for consistent and accurate movements, and that speed is simply a byproduct of network dynamics. In fact, there is evidence that increases in speed are required to maximize efficiency or maintain consistency⁵⁴.

Coordination across the motor network

The study of the coordination between different regions in the motor network may allow us to dissect the details of each region's complex role. It has been suggested that oscillatory or rhythmic activity is central to multi-region communication⁵⁵⁻⁵⁷. Intriguingly, a growing body of work has also proposed that transient oscillatory activity is central to motor function^{21,58,59}. In fact, modeling has suggested that low-frequency activity may be an essential feature of neural activity that generates descending commands to muscles⁶⁰. While this work has focused on the role of cortex in such a process, the coordinated low-frequency activity we observe between M1 and DLS suggests that such activity is also present in other nodes in the motor network and that such rhythmic activity may allow for the coordination of activity across regions. Further work exploring multi-region interactions will be essential to understanding the interplay between cortex and striatum, and the greater motor network, during motor skill learning.

Methods

Animal Care and Surgery

All procedures were in accordance with protocols approved by the Institutional Animal Care and Use Committee at the San Francisco Veterans Affairs Medical Center. Male Long-Evans rats between 3-4 months old and 250-300 grams were used in this study. Animals were kept under controlled temperature and a 12-h light, 12-h dark cycle with lights on at 06:00 a.m.

All surgical procedures were performed using sterile technique under 2-4% isoflurane. For electrode and/or cannula implantation, surgery involved exposure and cleaning of the skull, preparation of the skull surface (using cyanoacrylate), and then implantation of skull screws for referencing and overall head-stage stability. Reference screws were implanted posterior to lambda, contralateral to the neural recordings. Ground screws were implanted posterior to lambda, ipsilateral to the neural recordings. Craniotomy and dural incision were performed, followed by implantation of neural probes or cannulas. Neural probes (32- or 64-channel Tucker-Davis Technologies (TDT) 33 μ m polyimide-coated tungsten microwire electrode arrays) or infusion cannulas (PlasticsOne) were implanted in the forelimb area of M1, centered at 3.5mm lateral and 0.5mm anterior to bregma and implanted in layer 5 at a depth of 1.5mm, and the dorsolateral striatum, centered at 4mm lateral and 0.5mm anterior to bregma and implanted at a depth of 4.5mm. Final location of electrodes was confirmed by electrolytic lesion (**Figure 2.1**). Muscimol localization was performed by infusing a fluorescent muscimol (**Figure 2.15**; BODIPY, TMR-X Conjugate) prior to perfusion and histology. The forearm was implanted with a pair of twisted electromyography (EMG) wires (0.007" single-stranded, teflon-coated, stainless steel wire; A-M Systems, Inc.) with a hardened epoxy ball (J-B Weld) at one end preceded by 1–2 mm of uncoated

wire under the ball. Wires were inserted into the muscle belly and pulled through until the ball came to rest on the belly. EMG wires were braided, tunneled under the skin to a scalp incision, and soldered into headstage connectors. Fascia and skin incisions were closed with a suture. The post-operative recovery regimen included administration of buprenorphine at 0.02 mg/kg and meloxicam at 0.2 mg/kg. Dexamethasone at 0.5 mg/kg and Trimethoprim sulfadiazine at 15 mg/kg were also administered post-operatively for 5 days. All animals recovered for 14 days prior to start of behavioral experiments.

For photothrombotic lesion surgery, the protocol was the same as for electrode and cannula implantation up to the craniotomy. After the craniotomy, rose bengal dye was injected into the femoral vein using an intravenous catheter. Next, the surface of the brain was illuminated with white light (KL-1500 LCD, Schott) using a fiber optic cable for 20 min. We used a 4-mm aperture for lesion induction (centered on the M1 coordinates reported above) and covered the remaining cortical area with an aluminum foil mask to prevent light penetration. After lesion induction the craniotomy was covered with a layer of silicone (Quiksil), followed by dental cement. The same post-operative recovery regimen as electrode and cannula implantation surgery was then implemented. After all experiments, rats were anesthetized and transcardially perfused with 0.9% sodium chloride, followed by 4% formaldehyde. The harvested brains were post-fixed for 24 h and immersed in 20% sucrose for 2 days. Coronal cryostat sections (40- μ m thickness) were then mounted and imaged for lesion and muscimol localization.

Behavior

For learning, rats naïve to any motor training were first tested for forelimb preference. This involved presenting approximately ten pellets to the animal and observing which forelimb was most often used to reach for the pellet. One-week later rats underwent surgery followed by a recovery period. Rats were then trained using an automated reach-box, controlled by custom MATLAB scripts and an Arduino micro-controller. This setup required minimal user intervention, as described previously¹⁹. Each trial consisted of a pellet dispensed on the pellet tray followed by an alerting beep indicating that the trial was beginning and then the door opening. Animals had to reach, grasp and retrieve the pellet. A real-time “pellet-detector” using an IR sensor centered over the pellet was used to determine when the pellet was moved, indicating the trial was over, and the door was closed. All trials were captured by video through a camera placed on the side of the behavioral box, which was synced with electrophysiology data using an Arduino digital output. Two types of cameras were used, Microsoft LifeCams which captured videos at 30Hz, and Basler Cameras which captured videos at 75Hz. Behavioral scoring, including timing of sub-movements and reconstruction of the forelimb trajectory, was performed manually. For inactivation experiments, a second camera was placed above the behavioral box to capture body posture during reaching. The learning paradigm consisted of 100-150 trial sessions performed each day for 8 consecutive days. Rats had 15 seconds in each trial to execute a reach before a 10 second inter-trial-interval in which the door was closed, which led to ~75-150 trials performed (i.e., trials where the pellet was displaced) each day. For the “extended training” cohort, a separate cohort of animals was trained more extensively using the same paradigm for 4 weeks, resulting in ~2500 trials performed.

Behavioral Analysis

Learning was assessed using four metrics (**Figure 2.3**): (1) reach duration defined as the time from the onset of movement (movement onset) to when the paw is fully retracted off of the pellet tray (retract onset), (2) sub-movement timing variability defined as the standard deviation across trials of the duration between paw touching the pellet (pellet touch) and when the paw is fully retracted off of the pellet tray (retract onset), (3) success rate defined as the percentage of reaches that resulted in retrieval of the pellet into the box, and (4) forelimb trajectory consistency defined as the average correlation between each individual trial's forelimb trajectory and the mean forelimb trajectory calculated over all trials in that session (computed separately in each of the two dimensions). These metrics were chosen as they measured changes in both gross movements of the forelimb involved in producing a consistent reach and fine movements of the fingers involved in successful grasping. For the scatter plots comparing changes in reach duration, sub-movement timing variability, and forelimb trajectory consistency across learning to changes in movement-related 3-6Hz M1-DLS LFP coherence (**Figure 2.11**), normalized values of reach duration, sub-movement timing variability, and forelimb trajectory consistency were computed by z-scoring the eight mean values corresponding to the eight days of training for each animal separately, then combining the normalized values across animals. To determine body posture with DLS inactivation we used a top camera and manually determined the body axis from nose to center of the body (**Figure 2.13**). We defined posture variability as the mean of the absolute value of the distance across all trials from the middle of the body to the average middle of the body position. We defined lateral bias as the mean distance in the x-axis across all trials from the middle of the body to the average middle of the body position.

Inactivation Experiments

We performed two sets of inactivation experiments. For both experiments, rats were first tested for forelimb preference, then underwent either dual cannula surgery (M1 and DLS cannula implantation) or cannula and electrode implantation surgery (DLS cannula and M1 electrode implantation). Following the recovery period, rats were trained for 10 days (100 trials/day). Following this training, inactivation experiments began. For each session, baseline performance was calculated from 100 trials performed before muscimol infusion. For M1/DLS inactivation experiments (**Figure 2.12**), infusion consisted of anesthetizing the rat (w/isoflurane) and infusion of 250nl of 1ug/ul muscimol (Tocris) in saline (0.9% sodium chloride) at a rate of 100nl/min in either M1 or DLS. After the two-and-a-half-minute infusion and a five-minute waiting period with the infusion cannula inserted, the rat was taken off anesthesia and allowed to recover for 2 hours. A 200 trials block was then performed alternating between 10 trials to the “close” position (5mm from pellet center to slot opening in behavioral box) and 10 trials to the “far” position (15mm from pellet center to slot opening in behavioral box). For DLS inactivation experiments (**Figure 2.14**), infusion protocol was the same except a volume of 1ul of muscimol was infused. After the ten-minute infusion and a five-minute waiting period, a 100-trial block was performed at the “far”/normal pellet position to examine effects of DLS inactivation on movement-related M1 activity.

Lesion Experiments

For photothrombotic lesion experiments (**Figure 2.17**), rats were first tested for forelimb preference, then trained for 10 days (100-150 trials/day). Pre-lesion performance was measured before animals underwent photothrombotic lesion surgery. Post-lesion, animals began performing

reaching trials at variable times, so “early” lesion performance was defined as the performance during the first session that animals were completing trials (within 8 to 13 days post-lesion). Animals underwent reach training until a performance plateau was reached, which was defined as “late” lesion performance (within 15 to 73 days). Photothrombotic lesion size was determined with immunohistochemistry (see below). For excitotoxic lesion experiments (**Figure 2.16**), rats were first tested for forelimb preference, then trained for 10 days (100 trials/day). Pre-lesion baseline performance was then measured. Excitotoxic lesions were then implemented with 500nl infusions of 10ug/ul ibotenic acid (7.4 pH; Abcam) at an infusion rate of 100nl/min at the same DLS coordinates as referenced above. The animals recovered for two weeks, then performance post-lesion was measured. Excitotoxic lesion position and size was determined with immunohistochemistry (see below).

In Vivo Electrophysiology

Units, LFP, and EMG activity were recorded using a TDT-RZ2 system (Tucker-Davies Technologies). Spike data were sampled at 24414 Hz and LFP/EMG data at 1017 Hz. ZIF-clip-based analog headstages with a unity gain and high impedance ($\sim 1 \text{ G}\Omega$) were used. Behavior-related timestamps (i.e., trial onset, trial completion) and video timestamps (i.e., frame times) were sent to the RZ2 analog input channel using an Arduino digital board and synchronized to neural data.

Neural Data Analysis

Analyses were conducted using a combination of custom-written scripts and functions in MATLAB 2015a/2017a (MathWorks), along with functions from the EEGLAB toolbox (<http://sccn.ucsd.edu/eeglab/>) and the Chronux toolbox (<http://chronux.org/>).

LFP analysis

Pre-processing steps for LFP analysis included: artifact rejection (manually removing noisy/broken channels), z-scoring entire recording session, and common-mode referencing using the median signal (at every time-point, the median signal across all channels in a region was calculated. This median signal was subtracted from every channel to decrease common noise and minimize volume conduction. Common-mode referencing was performed independently for the channels in each region, i.e., M1 and DLS.

In several instances we filtered LFP signals to isolate and display the low-frequency (3-6Hz) component of the signal (**Figure 2.5a&d**; **Figure 2.6a**; **Figure 2.8a**; **Figure 2.9a&b**; **Figure 2.11c**; **Figure 2.14b**). Filtering was performed using the EEGLAB function *eegfilt*. In addition to display purposes, we also used filtered LFP to characterize phase-locking of spiking activity specifically to low-frequency LFP signals. For this we used the Hilbert transform (MATLAB) to extract the phase information from low-frequency filtered LFP signals (**Figure 2.9b&c**). We also used filtered LFP and phase extraction to determine the movement-related phase lag between M1 and DLS LFP signals (**Figure 2.6**).

To quantify changes across frequencies in the amplitude of rhythmic activity in LFP signals we calculated movement-related LFP spectrograms and power spectrums within each region (**Figure 2.5; Figure 2.7; Figure 2.14b**). For learning comparisons, power was measured and compared for the same channels on day one and day eight across all channels (except those removed due to noise). This was carried out using wavelets with the EEGLab function *newtimef*⁶¹. To quantify phase-locking of LFP signals to specific sub-movements (movement onset, pellet touch, and retract onset) we calculated inter-trial coherence (ITC) of LFP signals across trials time-locked to these sub-movements (**Figure 2.5**). ITC was measured and compared for the same channels on day one and day eight across all channels (except those removed due to noise). ITC was computed using the EEGLab function *newtimef*⁶¹.

To characterize coordination of activity across regions we measured changes in movement-related spectral coherence between LFP channels in M1 and DLS (**Figure 2.5c; Figure 2.7; Figure 2.8; Figure 2.11**). For learning comparisons, coherence was measured for the same channels on day one and day eight, and specifically for channels with an increase in power of one baseline normalized unit from day one to day eight. Strong coherence in a specific frequency band indicates a constant phase relationship in that frequency between two signals and is theorized to indicate increased communication between regions^{56,57}. Spectral coherence was computed using chronux function *cohgramc*⁶². All comparisons of “movement-related” LFP power or coherence used power and coherence values generated from signals between 250ms before movement onset to 750ms after movement onset and trial averaging over relevant trials (e.g., all trials on day one or day eight).

To determine whether the emergence of coordinated low-frequency activity during training was attributable solely to faster and more consistent movements, we compared LFP power and coherence between “fast” trials (trials with a movement duration between 200 and 400ms) and trials with high forelimb trajectory correlation values (correlation values > 0.9) on days one and two vs. days seven and eight (**Figure 2.8**).

For the scatter plots comparing changes in reach duration, sub-movement timing variability, and forelimb trajectory consistency across learning to changes in movement-related 3-6Hz M1-DLS LFP coherence (**Figure 2.11**), normalized values of LFP coherence were computed by z-scoring the eight mean values corresponding to the eight days of training for each animal separately, then combining the normalized values across animals.

Spiking analysis

Thresholds for spiking activity were set online using a standard deviation of 4.5 (calculated over a one-minute baseline period using the TDT-RZ2 system), and waveforms and timestamps were stored for any event that crossed that threshold. Spike sorting was performed using Plexon OfflineSorter v4.3.0 (Plexon Inc.) with a PCA-based clustering method followed by manual inspection. We accepted units based on waveform shape, clear cluster boundaries in PC space, and 99.5% of detected events with an $ISI > 2\text{ms}$. All units were analyzed and not sorted into cell type based on waveform shape. Average firing rates were consistent with what has been previously reported for M1 and DLS units¹³ (**Figure 2.10**). Peri-event time histograms (PETHs) were generated by averaging spiking activity across trials in a session, locked to movement onset and binned at 25ms, then fitting a smoothing spline using MATLAB function *fit* (**Figure 2.9a&b&d**;

the smoothing spline was not applied to PETHs presented in **Figure 2.14c**). These PETHs were used to generate average movement-related spiking plots (**Figure 2.10e-j**). To investigate spiking activity relative to movement phase (**Figure 2.10g&h**), we interpolated spiking activity binned at 25ms during movement (from movement onset to retract onset) for each trial such that the resulting spiking activity from each trial was normalized to the same length and then averaged across trials. We also used PETHs to classify task-related units (**Figure 2.10**). We defined a unit as task-related if during movement its activity was greater than one standard deviation or less than one standard deviation away from its baseline activity.

To characterize low-frequency spiking activity, we generated histograms of the LFP phases at which each spike occurred for a single unit to a single LFP channel filtered in the 3-6Hz band in a one-second window around movement (-250ms before to 750ms after movement onset) across all trials of a session (**Figure 2.9b**). For learning comparisons, all units were compared to the same selected M1 and DLS LFP channel on day one and day eight. These histograms were generated for each unit-LFP channel pair both within and across regions. For every pair we then calculated the Rayleigh's z-statistic for circular non-uniformity. These z-statistics were then used to calculate the percentage of significantly non-uniform distributions across unit-LFP pairs with a significance threshold $p = 0.05$ (**Figure 2.9c**). A significantly non-uniform distribution signifies phase preference for spikes of a unit to an LFP signal.

To further characterize low-frequency spiking activity, we determined the percentage of units that displayed low-frequency (3-6Hz) quasi-oscillatory activity. To do this, we computed autocorrelations on each unit's PETH. If a unit's autocorrelation had a "peak" between 166-333ms

time lag (corresponding to 3-6Hz activity) the unit was considered quasi-oscillatory. A “peak” was defined as a higher average value between 166-333ms than between 100-166ms (**Figure 2.9d**).

To characterize spiking interactions between M1 and DLS we calculated the mean cross-correlation of movement-related spiking across regions for all M1 and DLS quasi-oscillatory units on day one and day eight (**Figure 2.9e**). To do this we concatenated spiking activity for all trials between -250ms to 750ms from movement onset for each M1 and DLS quasi-oscillatory unit and then computed the cross correlation for each M1 and DLS quasi-oscillatory unit pair using the MATLAB function *xcorr*.

To determine the effects of DLS inactivation on M1 spiking activity we compared movement-related firing rates from pre-infusion baseline trials and post-infusion trials. Movement-related firing rates were calculated by averaging the firing rate from -250ms before to 500ms after movement on each trial of the session (**Figure 2.14c**).

To characterize single-trial representations of population spiking activity we used Gaussian process factor analysis (GPFA)³³ to find low-dimensional neural trajectories for each trial (**Figure 2.19; Figure 2.20**). GPFA analyses were carried out using MATLAB based GUI DataHigh⁶³, 25ms time bins, and a dimensionality of 5. The first two factors were used for analysis as they accounted for >90% of shared variance explained in both M1 and DLS on each session. We found that the consistency of these trajectories, calculated by averaging the correlation of every trial’s neural trajectory to the mean neural trajectory of that session (performed in each dimension independently) provided a robust measure neural consistency as this measure increased in both M1

and DLS during learning (**Figure 2.20b**; M1: $t(1789) = 6.9$, $P = 7 \times 10^{-12}$; DLS: $t(766) = 4.8$, $P = 2 \times 10^{-6}$). We also determined the magnitude of deviation for each individual trial trajectory from the mean trajectory across all successful trials by taking the absolute value of the difference between the trajectory of each trial and the mean trajectory across all trials (**Figure 2.19b&c**; computed in each dimension independently). This was performed specifically for the time period between 250ms before movement onset until pellet touch. As this duration varied across trials, we interpolated each trial such that every trial was the same length (100 values) and then calculated the average deviation.

Statistics

Linear mixed-effects models were used to test the significance of differences across both behavioral and neural measures. Using these models accounts for the fact that units, channels or trials from the same animal are more correlated than those from different animals and is more stringent than computing statistical significance over all units, channels or trials⁶⁴. For example, to test for learning-related changes in reach duration across all trials from training day one to all trials from day eight, we implemented a linear mixed-effects model (using MATLAB *fitlme*) with random intercepts/effects for each rat ($n=4$) and reported the P values for the regression coefficients associated with day one and day eight. Similar models were used to test for changes in other behavioral or neural measures, including LFP power, LFP coherence, and LFP/sub-movement phase-locking.

Viral Injection

We performed two sets of viral injections in separate cohorts of rats to label anterograde projections from M1 and retrograde projections from DLS. To label anterograde projections from M1 we injected 750nl of AAV8-hsyn-JAWs-KGC-GFP-ER2 virus into two sites (1.5mm anterior, 2.7mm lateral to bregma, at a depth of 1.4mm and 0.5 posterior, 3.5mm lateral to bregma, at a depth of 1.4mm). To label retrograde projections from DLS we injected 750nl of retrogradeAAV-hsyn-JAWs-KGC-GFP-ER2 virus at one site (4mm lateral and 0.5 mm anterior to bregma at a depth of 4mm). Two weeks after injection rats were anesthetized and transcardially perfused with 0.9% sodium chloride, followed by 4% formaldehyde. The harvested brains were post-fixed for 24 h and immersed in 20% sucrose for 2 days. Coronal cryostat sections (40- μ m thickness) were then mounted and imaged with a fluorescent microscope.

Immunohistochemistry

Cryostat sections (40- μ m thickness) were pre-incubated with blocking buffer (2% goat serum, 0.1% bovine serum albumin and 0.3% Triton X-100 in 0.1 M PB) at room temperature, and then incubated with mouse anti-NeuN (1 mg/ml; Millipore, Temecula, CA) overnight. After washing, the sections were incubated with biotinylated anti-mouse IgG secondary antibody (5 mg/ml; Vector laboratories, Burlingame, CA) for 2 hrs. Sections were visualized by the DAB method, using ABC reagents using a Vector ABC kit (Vector laboratories) and peroxidase substrate solution (Vector laboratories). For fluorescence images, sections were incubated with Alexa Fluor 594-conjugated donkey anti-mouse IgG (Life Technologies Corporation, Grand Island, NY; 1:1000).

References

1. Schaffelhofer, S. & Scherberger, H. Object vision to hand action in macaque parietal, premotor, and motor cortices. *Elife* **5**, e15278 (2016).
2. Wang, X. *et al.* Deconstruction of Corticospinal Circuits for Goal-Directed Motor Skills. *Cell* **171**, 440–455.e14 (2017).
3. Whishaw, I. Q. An endpoint, descriptive, and kinematic comparison of skilled reaching in mice (*Mus musculus*) with rats (*Rattus norvegicus*). *Behav. Brain Res.* **78**, 101–11 (1996).
4. Diedrichsen, J., Shadmehr, R. & Ivry, R. B. The coordination of movement: optimal feedback control and beyond. *Trends Cogn. Sci.* **14**, 31–9 (2010).
5. Todorov, E. & Jordan, M. I. Optimal feedback control as a theory of motor coordination. *Nat. Neurosci.* **5**, 1226–1235 (2002).
6. Uno, Y., Kawato, M. & Suzuki, R. Formation and control of optimal trajectory in human multijoint arm movement. *Biol. Cybern.* **61**, 89–101 (1989).
7. Hoff, B. & Arbib, M. A. Models of Trajectory Formation and Temporal Interaction of Reach and Grasp. *J. Mot. Behav.* **25**, 175–192 (1993).
8. Flash, T. & Hogan, N. The coordination of arm movements: an experimentally confirmed mathematical model. *J. Neurosci.* **5**, 1688–703 (1985).
9. Kargo, W. J. & Nitz, D. A. Improvements in the Signal-to-Noise Ratio of Motor Cortex Cells Distinguish Early versus Late Phases of Motor Skill Learning. *J. Neurosci.* **24**, 5560–5569 (2004).
10. Li, Q. *et al.* Refinement of learned skilled movement representation in motor cortex deep output layer. *Nat. Commun.* **8**, 15834 (2017).

11. Hyland, B. Neural activity related to reaching and grasping in rostral and caudal regions of rat motor cortex. *Behav. Brain Res.* **94**, 255–69 (1998).
12. Ramanathan, D. S., Gulati, T. & Ganguly, K. Sleep-Dependent Reactivation of Ensembles in Motor Cortex Promotes Skill Consolidation. *PLOS Biol.* **13**, e1002263 (2015).
13. Santos, F. J., Oliveira, R. F., Jin, X. & Costa, R. M. Corticostriatal dynamics encode the refinement of specific behavioral variability during skill learning. *Elife* **4**, (2015).
14. Kawai, R. *et al.* Motor Cortex Is Required for Learning but Not for Executing a Motor Skill. *Neuron* **86**, 800–812 (2015).
15. Costa, R. M., Cohen, D. & Nicolelis, M. A. L. Differential Corticostriatal Plasticity during Fast and Slow Motor Skill Learning in Mice. *Curr. Biol.* **14**, 1124–1134 (2004).
16. Kupferschmidt, D. A., Juczewski, K., Cui, G., Johnson, K. A. & Lovinger, D. M. Parallel, but Dissociable, Processing in Discrete Corticostriatal Inputs Encodes Skill Learning. *Neuron* **96**, 476–489.e5 (2017).
17. Yin, H. H. *et al.* Dynamic reorganization of striatal circuits during the acquisition and consolidation of a skill. *Nat. Neurosci.* **12**, 333–341 (2009).
18. Hintiryan, H. *et al.* The mouse cortico-striatal projectome. *Nat. Neurosci.* **19**, 1100–14 (2016).
19. Wong, C. C., Ramanathan, D. S., Gulati, T., Won, S. J. & Ganguly, K. An automated behavioral box to assess forelimb function in rats. *J. Neurosci. Methods* **246**, 30–7 (2015).
20. Lalla, L., Rueda Orozco, P. E., Jurado-Parras, M.-T., Brovelli, A. & Robbe, D. Local or Not Local: Investigating the Nature of Striatal Theta Oscillations in Behaving Rats. *eneuro* **4**, ENEURO.0128-17.2017 (2017).

21. Hall, T. M., de Carvalho, F. & Jackson, A. A common structure underlies low-frequency cortical dynamics in movement, sleep, and sedation. *Neuron* **83**, 1185–99 (2014).
22. Riehle, A., Wirtsohn, S., Grün, S. & Brochier, T. Mapping the spatio-temporal structure of motor cortical LFP and spiking activities during reach-to-grasp movements. *Front. Neural Circuits* **7**, 48 (2013).
23. Jerbi, K. *et al.* Coherent neural representation of hand speed in humans revealed by MEG imaging. *Proc. Natl. Acad. Sci.* **104**, 7676–7681 (2007).
24. Dipietro, L., Poizner, H. & Krebs, H. I. EEG correlates of submovements. in *2011 Annual International Conference of the IEEE Engineering in Medicine and Biology Society* **2011**, 7429–7432 (IEEE, 2011).
25. Turner, R. S. & Desmurget, M. Basal ganglia contributions to motor control: a vigorous tutor. *Curr. Opin. Neurobiol.* **20**, 704–716 (2010).
26. Mazzoni, P., Hristova, A. & Krakauer, J. W. Why don't we move faster? Parkinson's disease, movement vigor, and implicit motivation. *J. Neurosci.* **27**, 7105–16 (2007).
27. Panigrahi, B. *et al.* Dopamine Is Required for the Neural Representation and Control of Movement Vigor. *Cell* **162**, 1418–1430 (2015).
28. Rueda-Orozco, P. E. & Robbe, D. The striatum multiplexes contextual and kinematic information to constrain motor habits execution. *Nat. Neurosci.* **18**, 453–460 (2015).
29. Dudman, J. T. & Krakauer, J. W. The basal ganglia: from motor commands to the control of vigor. *Curr. Opin. Neurobiol.* **37**, 158–166 (2016).
30. Yttri, E. A. & Dudman, J. T. A Proposed Circuit Computation in Basal Ganglia: History-Dependent Gain. *Mov. Disord.* **33**, 704–716 (2018).

31. Whishaw, I. Q., Zeeb, F., Erickson, C. & McDonald, R. J. Neurotoxic lesions of the caudate-putamen on a reaching for food task in the rat: acute sensorimotor neglect and chronic qualitative motor impairment follow lateral lesions and improved success follows medial lesions. *Neuroscience* **146**, 86–97 (2007).
32. Whishaw, I. Q., Alaverdashvili, M. & Kolb, B. The problem of relating plasticity and skilled reaching after motor cortex stroke in the rat. *Behav. Brain Res.* **192**, 124–136 (2008).
33. Yu, B. M. *et al.* Gaussian-Process Factor Analysis for Low-Dimensional Single-Trial Analysis of Neural Population Activity. *J. Neurophysiol.* **102**, 614–635 (2009).
34. Peters, A. J., Liu, H. & Komiyama, T. Learning in the Rodent Motor Cortex. *Annu. Rev. Neurosci.* **40**, 77–97 (2017).
35. Alaverdashvili, M. & Whishaw, I. Q. Motor cortex stroke impairs individual digit movement in skilled reaching by the rat. *Eur. J. Neurosci.* **28**, 311–322 (2008).
36. Guo, J.-Z. *et al.* Cortex commands the performance of skilled movement. *Elife* **4**, (2015).
37. Lemon, R. N. Descending Pathways in Motor Control. *Annu. Rev. Neurosci.* **31**, 195–218 (2008).
38. Miri, A. *et al.* Behaviorally Selective Engagement of Short-Latency Effector Pathways by Motor Cortex. *Neuron* **95**, 683–696.e11 (2017).
39. Ueno, M. *et al.* Corticospinal Circuits from the Sensory and Motor Cortices Differentially Regulate Skilled Movements through Distinct Spinal Interneurons. *Cell Rep.* **23**, 1286–1300.e7 (2018).
40. Lawrence, D. G. & Kuypers, H. G. The functional organization of the motor system in the monkey. I. The effects of bilateral pyramidal lesions. *Brain* **91**, 1–14 (1968).

41. Lawrence, D. G. & Kuypers, H. G. The functional organization of the motor system in the monkey. II. The effects of lesions of the descending brain-stem pathways. *Brain* **91**, 15–36 (1968).
42. Otchy, T. M. *et al.* Acute off-target effects of neural circuit manipulations. *Nature* **528**, 358–363 (2015).
43. Díaz-Hernández, E. *et al.* The Thalamostriatal Projections Contribute to the Initiation and Execution of a Sequence of Movements. *Neuron* **100**, 739–752.e5 (2018).
44. Murray, J. M. & Escola, G. S. Learning multiple variable-speed sequences in striatum via cortical tutoring. *Elife* **6**, e26084 (2017).
45. Dudman, J. T. & Gerfen, C. R. The Basal Ganglia. *Rat Nerv. Syst.* 391–440 (2015). doi:10.1016/B978-0-12-374245-2.00017-6
46. Oldenburg, I. A. & Sabatini, B. L. Antagonistic but Not Symmetric Regulation of Primary Motor Cortex by Basal Ganglia Direct and Indirect Pathways. *Neuron* **86**, 1174–1181 (2015).
47. Koralek, A. C., Jin, X., Long II, J. D., Costa, R. M. & Carmena, J. M. Corticostriatal plasticity is necessary for learning intentional neuroprosthetic skills. *Nature* **483**, 331–335 (2012).
48. DeLong, M. R., Crutcher, M. D. & Georgopoulos, A. P. Primate globus pallidus and subthalamic nucleus: functional organization. *J. Neurophysiol.* **53**, 530–543 (1985).
49. Filion, M., Tremblay, L. & Bédard, P. J. Abnormal influences of passive limb movement on the activity of globus pallidus neurons in parkinsonian monkeys. *Brain Res.* **444**, 165–76 (1988).

50. Taha, J. M., Favre, J., Baumann, T. K. & Burchiel, K. J. Characteristics and somatotopic organization of kinesthetic cells in the globus pallidus of patients with Parkinson's disease. *J. Neurosurg.* **85**, 1005–1012 (1996).
51. Baker, K. B. *et al.* Somatotopic organization in the internal segment of the globus pallidus in Parkinson's disease. *Exp. Neurol.* **222**, 219–25 (2010).
52. Shmuelof, L., Krakauer, J. W. & Mazzoni, P. How is a motor skill learned? Change and invariance at the levels of task success and trajectory control. *J. Neurophysiol.* **108**, 578–94 (2012).
53. Hikosaka, O., Yamamoto, S., Yasuda, M. & Kim, H. F. Why skill matters. *Trends Cogn. Sci.* **17**, 434–441 (2013).
54. Park, S.-W., Marino, H., Charles, S. K., Sternad, D. & Hogan, N. Moving slowly is hard for humans: limitations of dynamic primitives. *J. Neurophysiol.* **118**, 69–83 (2017).
55. Harris, A. Z. & Gordon, J. A. Long-Range Neural Synchrony in Behavior. *Annu. Rev. Neurosci.* **38**, 171–194 (2015).
56. Fries, P. A mechanism for cognitive dynamics: neuronal communication through neuronal coherence. *Trends Cogn. Sci.* **9**, 474–480 (2005).
57. Fries, P. Rhythms for Cognition: Communication through Coherence. *Neuron* **88**, 220–35 (2015).
58. Ramanathan, D. S. *et al.* Low-frequency cortical activity is a neuromodulatory target that tracks recovery after stroke. *Nat. Med.* **24**, 1257–1267 (2018).
59. Churchland, M. M. *et al.* Neural population dynamics during reaching. *Nature* **487**, 51–6 (2012).

60. Sussillo, D., Churchland, M. M., Kaufman, M. T. & Shenoy, K. V. A neural network that finds a naturalistic solution for the production of muscle activity. *Nat. Neurosci.* **18**, 1025–1033 (2015).
61. Delorme, A. & Makeig, S. EEGLAB: an open source toolbox for analysis of single-trial EEG dynamics including independent component analysis. *J. Neurosci. Methods* **134**, 9–21 (2004).
62. Bokil, H., Andrews, P., Kulkarni, J. E., Mehta, S. & Mitra, P. P. Chronux: a platform for analyzing neural signals. *J. Neurosci. Methods* **192**, 146–51 (2010).
63. Cowley, B. R. *et al.* DataHigh: graphical user interface for visualizing and interacting with high-dimensional neural activity. *J. Neural Eng.* **10**, 066012 (2013).
64. Aarts, E., Verhage, M., Veenvliet, J. V, Dolan, C. V & van der Sluis, S. A solution to dependency: using multilevel analysis to accommodate nested data. *Nat. Neurosci.* **17**, 491–496 (2014).

Chapter 3: How is the corticostriatal network modified?

Chapter 3 is adapted from Lemke, S. M., Ramanathan, D. S., Darevsky, D., Egert, D., Berke, J.D., Ganguly, K. Sleep spindles coordinate corticostriatal reactivations during the emergence of automaticity. bioRxiv 2020.10.25.354282 (2020). doi:10.1101/2020.10.25.354282

Abstract

Plasticity within the corticostriatal network is known to regulate the balance between behavioral flexibility and automaticity. Repeated training of an action has been shown to bias behavior towards automaticity, suggesting that training may trigger activity-dependent corticostriatal plasticity. However, surprisingly little is known about the natural activity patterns that may drive plasticity or when they occur during long-term training. Here we chronically monitored neural activity from primary motor cortex (M1) and the dorsolateral striatum (DLS) during both training and offline periods, i.e., time away from training including sleep, throughout the development of an automatic reaching action. We first show that blocking striatal NMDA receptors during offline periods prevents the emergence of behavioral consistency, a hallmark of automaticity. We then show that, throughout the development of an automatic reaching action, corticostriatal functional connectivity increases during offline periods. Such increases track the emergence of consistent behavior and predictable cross-area neural dynamics. We then identify sleep spindles during non-REM sleep (NREM) as uniquely poised to mediate corticostriatal plasticity during offline periods. We show that sleep spindles are periods of maximal corticostriatal transmission within offline periods, that sleep spindles in post-training NREM reactivate neurons across areas, and that sleep-

spindle modulation in post-training NREM is linked to observable changes in spiking relationships between individual pairs of M1 and DLS neurons. Our results indicate that offline periods, in general, and sleep spindles, specifically, play an important role in regulating behavioral flexibility through corticostriatal network plasticity.

Introduction

Automaticity allows animals to capitalize on invariance in the environment through the development of actions that, while inflexible to changes, are performed highly consistently in response to a specific stimulus¹⁻⁵. It has been demonstrated that the consistent production of an action emerges with repeated training over multiple days and is accompanied by coordinated neural activity across the corticostriatal network during action execution⁶⁻¹⁰. Importantly, the emergence of such actions has also been shown to require striatal NMDA receptor activation¹⁰⁻¹², suggesting that cortical activity patterns that modulate the striatum may be important drivers of activity-dependent plasticity^{13,14} and the emergence of coordinated corticostriatal activity. However, surprisingly little is known about the natural activity patterns related to repeated task training that underlie the emergence of automaticity.

One intriguing possibility is that time away from training - “offline” periods, including sleep - may play a role in modifying the corticostriatal network. This possibility is motivated by evidence that sleep-dependent reactivations of cortical neural ensembles active during task performance are essential for initial learning¹⁵⁻¹⁸. It is possible that the coordinated reactivations of both cortical and striatal ensembles modify the corticostriatal network and impact behavior during long-term training¹⁹⁻²². However, how cortical reactivation events engage downstream striatal ensembles

remains unclear. Moreover, how such cross-area activity may precisely modify the corticostriatal network and impact network activity during subsequent awake behavior is unexplored.

Currently, our understanding of how sleep impacts distributed brain networks is largely derived from the *systems consolidation theory*, where it has been shown that coordinated activity patterns across hippocampus and cortex lead to the formation of stable long-term memories in cortex that do not require the hippocampus^{23–25}. Notably, whether sleep impacts the connectivity across hippocampus and cortex has not been established. Therefore, one possibility is that, in the corticostriatal network, we similarly observe coordinated cross-area activity patterns during sleep but do not find evidence for the modification of corticostriatal connectivity during offline periods. Alternatively, it is possible that we find evidence that cross-area activity patterns during sleep modify the connectivity between cortex and striatum and impact network activity during subsequent behavior.

Here we establish that offline periods play an essential role in modifying the corticostriatal network during the emergence of automaticity and identify sleep spindles as uniquely poised to mediate such plasticity. We show that during post-training NREM, sleep-dependent reactivation events are coordinated across both cortex and striatum during sleep spindles and link such spindle-modulation to changes in functional connectivity across the corticostriatal network. These results suggest that sleep plays an important role in modifying cross-area connectivity within offline periods and that the modulation of activity patterns during sleep may offer novel therapeutic targets for unlearning maladaptive habits^{26,27}.

Results

To study how corticostriatal network activity evolves during long-term training, we implanted six adult rats with either microwire electrode arrays ($n = 4$) or custom built high-density silicon probes²⁸ ($n = 2$) in both primary motor cortex (M1) and the dorsolateral striatum (DLS), which receives the majority of M1 projections to the striatum²⁹ (**Figure 3.1a**). Neural activity across regions was monitored as rats underwent ~eight days of reach-to-grasp task training (range: 5-14 days, mean: 8.67 days). Each recording day consisted of a 2-3 hour pre-training block (“pre-sleep”), a 100-150 trial training block, and a second 2-3 hour post-training block (“post-sleep”; **Figure 3.1b**; pre-sleep length: 157.2 ± 5.8 minutes, post-sleep length: 166.6 ± 6.3 minutes, mean \pm SEM). The reach-to-grasp task requires rats to reach and grasp a food pellet through a small window present in their behavioral box. During pre- and post-sleep, behavioral states, i.e., wake, non-REM sleep (NREM), and REM sleep, were classified using standard methods based on cortical local field potential (LFP) power and movement measured from video or electromyography (EMG) activity³⁰.




Figure 3.1. Offline striatal NMDA receptor activation is required to develop a consistent behavior. **a.** Schematic displaying primary motor cortex (M1) and dorsolateral striatum (DLS) recording locations (left) and labeled M1 projections showing direct input to the DLS (right). **b.** Schematic showing each day's recording blocks during long-term training. **c.** Individual reach trajectories in grey overlaid with mean reach trajectory across trials in red for each day of training in example animal. **d.** Average reach velocity profile in x and y dimensions for each day of training in example animal. **e.** Reach velocity profile correlation for first eight days of training for individual animals in grey overlaid with mean \pm SEM across animals in black. **f.** Day-to-day evolution in reach velocity profile correlation with post-training DLS infusions of either AP5 or saline in example animal. **g.** Comparison of day-to-day changes in reach velocity profile correlation with post-training saline infusion, post-training AP5 infusion, or no infusion in learning cohort animals, showing specific decrease in day-to-day reach velocity profile correlation with post-training AP5 infusion. Individual day-to-day changes as grey dots overlaid with mean \pm SEM across all day-to-day changes in color.

Offline striatal NMDA receptor activation is required to develop a consistent behavior

With repeated training on the reach-to-grasp task, animals developed a consistent reaching trajectory (**Figure 3.1c**) and reaching velocity profile (**Figure 3.1d**). Measuring the correlation between the mean reaching velocity profile on each day of training and the final day of training revealed that a consistent day-to-day reaching action emerged within the first eight days of training (**Figure 3.1e**). Such day-to-day invariance in skilled reaching is consistent with the emergence of automaticity^{1-3,31}. To further test the automaticity of reaching after the emergence of invariant behavior, we moved the location of the food pellet such that reaches to the old pellet location would no longer be successful. If an animal were reaching flexibly, we would expect that reaching behavior would quickly adapt to the new position. Alternatively, if an animal were reaching automatically, we would expect that reaches would remain consistent despite decreased success. Consistent with automaticity, reach trajectories remained consistent and did not adapt to the new pellet position despite a large decrease in success rate (**Figure 3.2**). We also examined whether such automaticity emerged with long-term training or existed at the start of training by testing whether animals could reach flexibly during the first two days of exposure to the task ($n = 2$ rats with no neural implant). These animals were able to reach to each of the two different pellet positions with comparable success rates (animal 1: 43% and 72% success rate; animals 2: 44% and 50% success rate, ~200 trials in each animal), indicating that their reaching behavior was flexible prior to repeated training. Altogether this indicated that our training paradigm led to the formation of inflexible and automatic reach-to-grasp behavior.

To further test the link between offline plasticity in the corticostriatal network and increases in behavioral consistency, we trained a new cohort of animals ($n = 6$ rats) and infused $1\mu\text{l}$ of either NMDA receptor antagonist AP5 ($5\mu\text{g}/\mu\text{l}$) or saline into DLS immediately after training on each day (**Figure 3.1f**). This revealed that offline striatal NMDA activation was essential for the emergence of a consistent reaching behavior, as day-to-day changes in reach consistency were significantly decreased with AP5 infusions, compared to saline infusions or day-to-day improvements observed in the learning cohort (**Figure 3.1g**; $n = 24$ correlation change values with AP5 infusions, -0.03 ± 0.03 correlation value, $n = 24$ correlation change values with saline infusions, 0.07 ± 0.02 correlation value, $n = 40$ correlation change values in learning cohort, 0.04 ± 0.01 correlation value; AP5 infusions vs. saline infusions: $t(23) = 2.8$, $P = 8 \times 10^{-3}$, paired-sample




Figure 3.2 Animals do not adapt quickly to new pellet position after long-term reach-to-grasp training. **a.** Reach trajectories from example session following reach-to-grasp task training paradigm, showing that reach trajectories with the pellet in the learned position or a new position are largely overlapping. **b.** Histogram of single-trial correlation values for individual reach trajectories to the mean reach trajectory for trials with the pellet in the learned pellet position or a new pellet position in first example animal (left; two-sample Kolmogorov–Smirnov test between distributions from trials with pellet in learned position and pellet in new position) and success rate in pellet retrieval for learned and new pellet positions (right). **c.** same as **b** for second example animal.

t-test, AP5 infusions vs. learning cohort: $t(62) = 2.3$, $P = 0.03$, two-sample t-test, saline infusions vs. learning cohort: $t(62) = 1.3$, $P = 0.20$, two-sample t-test). Importantly, task engagement, as measured by reaction time from trial start to reach, did not differ for trials on subsequent days after AP5 or saline infusions (AP5: 218.2 ± 3.2 ms, saline: 222.4 ± 3.0 ms, $t(5198) = -0.96$, $P = 0.34$, two-sample t-test). Altogether, these results were consistent with the notion that offline plasticity in the corticostriatal network following training is critical for the emergence of automaticity.

Corticostriatal functional connectivity increases during offline periods

To measure long-term changes in corticostriatal functional connectivity during the emergence of automaticity, we measured LFP coherence across individual pairs of M1 and DLS electrodes. LFP signals can be stably recorded across multiple days allowing LFP coherence to provide a stable long-term measure of multi-region connectivity^{32,33}. Specifically, within the corticostriatal network, theta coherence (4-8Hz) has been previously shown to reflect coordinated population spiking activity^{8,9,34}. Therefore, we measured 4-8Hz LFP coherence during pre- and post-sleep on each day of training to determine when corticostriatal functional connectivity changed during long-term training (**Figure 3.3a**). LFP coherence was calculated specifically during NREM to control for any differences in the time spent in each behavioral state during pre- and post-sleep. Common-mode referencing was applied, separately in each region, to decrease common noise and minimize volume conduction⁸. We found that there was a significant correlation between each day's mean 4-8Hz LFP coherence across all channel pairs and reach velocity profile correlation ($r = 0.44$, $P = 7 \times 10^{-3}$, Pearson's r), indicating that offline LFP coherence reflects changes in corticostriatal functional connectivity that are related to the emergence of a consistent behavior.

We next sought to determine whether LFP coherence increased during training or offline periods. To do this, we specifically examined LFP channel pairs that increased in coherence from day one to day eight (33% of pairs increased, 17% decreased, and 50% did not change; increase or decrease defined as a change in coherence of at least 0.25). Remarkably, within the subset of channels that showed training-related increases in coherence over learning, increases occurred largely offline, i.e., between each day's post-sleep and the next day's pre-sleep, rather than online during training, i.e., between pre- and post-sleep on the same day (**Figure 3.3b&c**). More specifically, the distribution of online LFP coherence changes was not significantly different than zero, while the distribution of offline LFP coherence changes was skewed toward larger increases (**Figure 3.3d**; online LFP coherence changes: $t(422) = 1.2$, $P = 0.23$, offline LFP coherence changes: $t(422) = 18.8$, $P = 5 \times 10^{-57}$, one-sample t-test). Importantly, the subset of channels that showed training-related increases had a close relationship to the emergence of consistent behavior (**Figure 3.3e**; $r = 0.73$, $P = 4 \times 10^{-7}$, Pearson's r), providing evidence that offline increases in corticostriatal functional connectivity are relevant to the consistency of behavior during the emergence of automaticity.




Figure 3.3. Corticostriatal functional connectivity increases during offline periods. **a.** Schematic depicting M1 and DLS electrode pairs with high 4-8Hz LFP coherence (>0.6 coherence value measured in NREM) during pre- and post-sleep on one day of training and pre-sleep on the next day of training, showing an increase in the number of high LFP coherence pairs occurring offline rather than online, in example animal. **b.** LFP coherence spectrums (measured in NREM) across example M1 and DLS electrode pair for pre- and post-sleep periods represented in panel **a.** showing an increase in 4-8Hz LFP coherence largely occurring offline rather than online. **c.** LFP coherence (4-8Hz measured in NREM) for each pre- and post-sleep period throughout learning for example M1 and DLS electrode pair, showing increases in coherence largely occurring offline rather than online, overlaid with reach velocity profile correlation values for each day of training. **d.** Comparison of distributions of online (left) and offline (right) changes in LFP coherence (4-8Hz measured in NREM) averaged across training days for M1 and DLS electrode pairs across animals. **e.** Correlation between each day's mean LFP coherence (mean 4-8Hz measured in NREM during both pre- and post-sleep) and reach velocity profile correlation value.

Offline increases in functional connectivity predict the emergence of low-dimensional cross-area neural dynamics during behavior

We next examined how offline increases in corticostriatal functional connectivity may impact corticostriatal network activity during subsequent reach-to-grasp performance. We extracted low-dimensional neural trajectory representations of DLS spiking activity during reaching using principle components analysis (PCA). We then examined the evolution of how spiking activity in M1 could predict DLS neural trajectories over the course of training (**Figure 3.4a**). We found that the ability to predict DLS neural trajectories during reaching from M1 spiking activity increased with training, while the ability to predict the trajectory representations of DLS activity during a baseline, non-reaching, period did not significantly change (**Figure 3.4b**; reach activity: first two days of training: 0.15 ± 0.05 Pearson's r , last two days of training: 0.46 ± 0.05 Pearson's r , $t(30) = -4.4$, $P = 1 \times 10^{-4}$, two-sample t-test; baseline activity: first two days of training: 0.03 ± 0.02 Pearson's r , last two days of training: 0.01 ± 0.03 Pearson's r , $t(30) = 1.0$, $P = 0.30$, two-sample t-test). Notably, the ability to predict DLS neural trajectories during reaching from M1 spiking activity was significantly correlated to the mean 4-8Hz LFP coherence measured offline on each day of training (**Figure 3.4c**), indicating that offline increases in LFP coherence track the emergence of predictable cross-area dynamics during subsequent reach-to-grasp performance.




Figure 3.4. Offline increases in functional connectivity predict the emergence of low-dimensional cross-area neural dynamics during behavior. **a.** Trial-averaged neural trajectory (PC1 and PC2) of DLS activity during reaching (one second before to one second after pellet touch) on day one (left) and day eight (right) of training in example animal, overlaid with prediction of DLS neural trajectory from M1 spiking activity. **b.** Ability to predict DLS neural trajectory (PC1 and PC2) during reaching and during a baseline, non-reaching, period from M1 spiking activity on each day of training (mean \pm SEM across animals). **c.** Correlation between each day's mean LFP coherence (mean 4-8Hz measured in NREM during both pre- and post-sleep) and ability to predict DLS neural trajectory (PC1 and PC2) during reaching from M1 spiking activity averaged across training days for M1 and DLS electrode pairs across animals. **e.** Correlation between each day's mean LFP coherence (mean 4-8Hz measured in NREM during both pre- and post-sleep) and reach velocity profile correlation value.

Corticostriatal transmission strength within offline periods is maximal during sleep spindles in NREM

Given the evidence that offline periods are relevant for changes in corticostriatal functional connectivity, we next sought to identify the activity patterns that may be responsible for driving such plasticity across M1 and DLS. To do this, we first examined how corticostriatal transmission strength, i.e., the degree to which M1 neural activity drives DLS activity, differed across behavioral states during offline periods (**Figure 3.5a**). To measure this, we characterized putative

monosynaptically connected pairs of M1 and DLS units ($n = 1,100$ M1 and 579 DLS units) by determining whether there was a significant peak in the cross correlation of their spiking activity at the short-latency time lag consistent with the conduction and synaptic delays between M1 and DLS (~6ms time lag from M1 to DLS activity⁹; **Figure 3.5b**; 3,969/10,286 M1 and DLS unit pairs were classified as putatively connected; **Figure 3.6a&b**). We then compared the short-latency cross correlation magnitude (1-10ms time lag) for the population of putatively connected M1 and DLS pairs across behavioral states. To account for differences in firing rates across behavioral states (**Figure 3.6c&d**), we normalized each pair's cross correlation by the mean cross correlation value from 50-100ms time lag, where no consistent spiking relationship is expected between putatively connected pairs of M1 and DLS units. This revealed that corticostriatal transmission strength was maximal during NREM, compared to REM or wake (**Figure 3.5c & Figure 3.7a&b**).

Given the heterogeneous nature of NREM activity, we next explored the dynamics of corticostriatal transmission within NREM. We specifically detected NREM rhythms in M1 that have been previously related to activity-dependent plasticity in cortex, i.e., sleep spindles, slow oscillations, and delta waves^{18,35-37}, and examined whether activity in DLS was also modulated during these rhythms (**Figure 3.5d**). We found that both LFP signals and spiking in DLS were significantly modulated during slow oscillations, delta waves, and sleep spindles detected in M1 (**Figure 3.5e&f; Figure 3.8**). To compare corticostriatal transmission strength during these rhythms, we measured the short-latency cross correlation magnitude for the population of putatively connected M1 and DLS pairs using the spiking activity during each sleep rhythm. Importantly, we applied a previously established normalization method to isolate and subtract off the influence of firing rate changes or LFP phase-locking differences across NREM rhythms on

cross correlations³⁸. This revealed that sleep spindles were unique periods of boosted corticostriatal transmission strength, compared to slow oscillations or delta waves (**Figure 3.5g & Figure 3.7c&d**). Altogether, this indicated that sleep spindles during NREM may be particularly relevant periods for activity-dependent plasticity within the corticostriatal network, given the high transmission of activity from M1 to DLS.




Figure 3.5. Corticostriatal transmission strength within offline periods is maximal during sleep spindles in NREM. **a.** M1 local field potential (LFP) spectrogram and behavioral state detection from example session. **b.** Example M1 and DLS single unit sorting from high-density silicon probe (top) and cross correlation of spiking activity centered on DLS unit spiking for an example pair of M1 and DLS units showing a short-latency peak indicating putative monosynaptically connectivity (bottom). **c.** Comparison of normalized cross correlations of spiking activity from all putatively connected pairs of M1 and DLS units across behavioral states, showing that corticostriatal transmission strength is maximal in NREM (width of line represents mean \pm SEM). **d.** Snippet of LFP and single unit spiking activity from M1 and DLS during NREM overlaid with detected NREM rhythms in M1. **e.** Mean LFP and spiking activity during slow oscillations, delta waves, and sleep spindles in both M1 and DLS in example animal (top) and percentage of M1 and DLS units across animals significantly phase locked to M1 LFP during each NREM rhythm (significance threshold of $P = 0.05$, Rayleigh test of uniformity, bottom). **f.** Comparison of firing rate modulation distributions for M1 (left) and DLS (right) units across animals during slow oscillations, delta waves, sleep spindles, and a baseline NREM period. **g.** Comparison of normalized cross correlations of spiking activity for all putatively connected pairs of M1 and DLS units across NREM rhythms, showing that corticostriatal transmission strength is maximal during sleep spindles (width of line represents mean \pm SEM).




Figure 3.6. Electrophysiology recordings from M1 and DLS **a.** Illustration of target electrode locations in M1 and DLS and dimensions of microwire electrode and silicon probes. **b.** Cross correlations of spiking activity across all pairs of putatively connected and not putatively connected pairs of M1 and DLS units (width of line represents mean \pm SEM) and pie chart depicting percentage of all pairs that are classified as significantly connected. **c.** Comparison of firing rates in M1 across behavioral states, before and after training (1,100 M1 units; 3.4 ± 0.1 spikes/second in pre NREM vs. 3.7 ± 0.1 spikes/second in post NREM, mean \pm SEM, paired-sample t-test: $t(1099) = -6.8, P = 1 \times 10^{-11}$; 4.2 ± 0.1 spikes/second in pre REM vs. 4.4 ± 0.1 spikes/second in post REM, mean \pm SEM, paired-sample t-test: $t(1099) = -5.5, P = 6 \times 10^{-8}$; 4.0 ± 0.1 spikes/second in pre wake vs. 4.4 ± 0.1 spikes/second in post wake, mean \pm SEM, paired-sample t-test: $t(1099) = -9.5, P = 1 \times 10^{-20}$). **d.** Comparison of firing rates in DLS across behavioral states, before and after training (579 DLS units; 1.9 ± 0.1 spikes/second in pre NREM vs. 2.1 ± 0.1 spikes/second in post NREM, mean \pm SEM, paired-sample t-test: $t(578) = -6.5, P = 1 \times 10^{-10}$; 2.5 ± 0.1 spikes/second in pre REM vs. 2.7 ± 0.1 spikes/second in post REM, mean \pm SEM, paired-sample t-test: $t(578) = -5.5, P = 8 \times 10^{-8}$; 2.4 ± 0.1 spikes/second in pre wake vs. 2.8 ± 0.1 spikes/second in post wake, mean \pm SEM, paired-sample t-test: $t(578) = -6.8, P = 3 \times 10^{-11}$).




Figure 3.7. Corticostriatal transmission strength across behavioral states and NREM rhythms. **a.** Comparison of normalized cross correlations across behavioral states for all putatively connected pairs of M1 and DLS units in example animals implanted with either microwire electrode array (top) or high-density silicon probe (bottom), showing that corticostriatal transmission strength is maximal in NREM sleep (width of line represents mean \pm SEM) **b.** Comparison of mean short-latency correlation magnitude (1-10ms time lag) across behavioral states, values from days for animals implanted with microwires in grey, silicon probes in red, and mean \pm SEM across animals in black ($n = 52$ days across 7 rats; wake: $3.2 \times 10^{-4} \pm 0.3 \times 10^{-4}$ correlation value, REM: $2.7 \times 10^{-4} \pm 0.5 \times 10^{-4}$ correlation value, NREM: $5.1 \times 10^{-4} \pm 0.4 \times 10^{-4}$ correlation value; wake vs. REM: $t(51) = 1.8$, $P = 0.08$, paired-sample t-test, REM vs. NREM: $t(51) = -7.0$, $P = 6 \times 10^{-9}$, paired-sample t-test, wake vs. NREM: $t(51) = -7.7$, $P = 4 \times 10^{-10}$, paired-sample t-test). **c.** Comparison of normalized cross correlations across NREM rhythms for all putatively connected pairs of M1 and DLS units in example animals implanted with either microwire electrode array (top) or high-density silicon probe (bottom), showing that corticostriatal transmission strength is maximal during sleep spindles (width of line represents mean \pm SEM). **d.** Comparison of mean short-latency correlation magnitude (1-10ms time lag) across NREM rhythms, values from days for animals implanted with microwires in grey, silicon probes in red, and mean \pm SEM across animals in black ($n = 52$ days across 7 rats; sleep spindles: $9.9 \times 10^{-4} \pm 0.8 \times 10^{-4}$ correlation value, delta waves: $7.5 \times 10^{-4} \pm 0.6 \times 10^{-4}$ correlation value, slow oscillations: $5.7 \times 10^{-4} \pm 1.7 \times 10^{-4}$ correlation value; sleep spindles vs. slow oscillation: $t(51) = 4.0$, $P = 2 \times 10^{-4}$, paired-sample t-test, delta waves vs. slow oscillations: $t(51) = 1.2$, $P = 0.22$, paired-sample t-test, sleep spindles vs. delta waves: $t(51) = 2.4$, $P = 0.01$, paired-sample t-test).




Figure 3.8. Corticoatrial modulation across NREM rhythms. **a.** Comparison of firing rate modulation distributions for M1 (left) and DLS (right) units during slow oscillations and a baseline NREM period (M1: $P = 3 \times 10^{-28}$, DLS: $P = 6 \times 10^{-10}$, two-sample Kolmogorov–Smirnov test, followed by a shift test to assess how quartiles of the distribution differed). **b.** Comparison of firing rate modulation distributions for M1 (left) and DLS (right) units during delta waves and a baseline NREM period (M1: $P = 1 \times 10^{-16}$, DLS: $P = 2 \times 10^{-3}$, two-sample Kolmogorov–Smirnov test, followed by a shift test to assess how quartiles of the distribution differed). **c.** Comparison of firing rate modulation distributions for M1 (left) and DLS (right) units during sleep spindles and a baseline NREM period (M1: $P = 4 \times 10^{-24}$, DLS: $P = 2 \times 10^{-4}$, two-sample Kolmogorov–Smirnov test, followed by a shift test to assess how quartiles of the distribution differed).

Striatal reactivations during sleep spindles reflect cortical input

We next assessed whether sleep spindles, or other NREM rhythms, were significant predictors of day-to-day changes in behavioral consistency. We found that sleep spindle density (events/minute) during post-sleep, but not pre-sleep, was a significant predictor of day-to-day changes in reaching consistency (pre-sleep sleep spindles: $r = 0.07$, $P = 0.70$, post-sleep sleep spindles: $r = 0.38$, $P = 0.01$, Pearson’s r). Neither delta waves nor slow oscillations were significantly predictive of day-to-day changes in reaching consistency (pre-sleep delta waves: $r = 0.02$, $P = 0.87$, post-sleep delta waves: $r = 0.03$, $P = 0.83$, pre-sleep slow oscillations: $r = -0.11$, $P = 0.51$, post-sleep slow

oscillations: $r = -0.08$, $P = 0.60$, Pearson's r). Given the evidence that corticostriatal transmission is boosted during sleep spindles, a possible explanation for the unique relationship between post-training sleep spindle density and day-to-day increases in behavioral consistency is that sleep spindles drive activity-dependent corticostriatal plasticity that impacts behavior. If this were the case, we would expect relevant M1 and DLS neural populations to be preferentially engaged during sleep spindles after training. In fact, we found that reach modulated (RM) M1 units, characterized by a significant modulation of activity during the reaching action, were significantly more modulated during sleep spindles after training, while non-RM M1 units did not significantly change in modulation from pre- to post-sleep (**Figure 3.9a&b**; RM M1 units: $P = 0.02$, non-RM M1 units: $P = 0.75$, two-sample Kolmogorov–Smirnov test between distributions from pre- and post-sleep, followed by a shift test to assess how quartiles of the distributions differed; P values for the rest of Figure 3 reflect these statistical tests).

How does this then affect downstream neural activity during sleep spindles in DLS? Surprisingly, both RM and non-RM DLS unit populations were significantly more modulated during sleep spindles after training (**Figure 3.9c&d**; RM DLS units: $P = 8 \times 10^{-3}$, non-RM DLS units: $P = 2 \times 10^{-3}$). One possibility is that this occurs because DLS activity during sleep spindles is driven by M1 input, rather than strictly reflecting reach modulation, as is the case for M1 neurons. Consistent with this, we were able to separate DLS unit populations that increased or did not increase in sleep spindle modulation after training based on putative connectivity with M1 units (**Figure 3.9e&f**; RM DLS units with strong RM M1 input: $P = 1 \times 10^{-3}$, RM DLS units with weak or no RM M1 input: $P = 0.82$, non-RM DLS units with strong RM M1 input: $P = 3 \times 10^{-3}$, non-RM DLS units with weak or no RM M1 input: $P = 0.24$; strong RM M1 input was defined as putative connectivity

with three or more RM M1 units). In contrast, DLS units did not increase in modulation during either delta waves (**Figure 3.10**) or slow oscillations (**Figure 3.11**) after training. Altogether, this suggests that, while reactivations during sleep spindles reflect task modulation in M1, DLS reactivations reflect cortical input, suggesting a potential role for sleep spindles in reinforcing task-related corticostriatal connectivity, including novel connectivity such as projections from RM M1 units to previously non-RM DLS units.




Figure 3.9. Striatal reactivations during sleep spindles reflect cortical input. **a.** Trial-averaged spiking activity during reaching for all reach modulated (RM) and non-RM M1 units across days and animals. **b.** Comparison of distributions of sleep spindle modulation during pre- and post-sleep for RM (top) and non-RM (bottom) M1 units, showing increased modulation from pre- to post-sleep specifically in RM M1 units. **c.** Trial-averaged spiking activity during reaching for all reach modulated (RM) and non-RM DLS units across days and animals. **d.** Comparison of distributions of sleep spindle modulation during pre- and post-sleep for RM (top) and non-RM (bottom) DLS units, showing increased modulation from pre- to post-sleep in both RM and non-RM DLS units. **e.** Comparison of distributions of sleep spindle modulation during pre- and post-sleep for RM DLS units with strong RM M1 input (top) and weak or no RM M1 input (bottom), showing increased modulation from pre- to post-sleep specifically in RM DLS units with strong RM M1 input. **f.** Comparison of distributions of sleep spindle modulation during pre- and post-sleep for non-RM DLS units with strong RM M1 input (top) and weak or no RM M1 input (bottom), showing increased modulation from pre- to post-sleep specifically in non-RM DLS units with strong RM M1 input.




Figure 3.10. Delta wave modulation change with training. **a.** Trial-averaged spiking activity during reaching for all reach modulated (RM) and non-RM M1 units across days and animals. **b.** Comparison of distributions of delta wave modulation during pre- and post-sleep for RM (top) and non-RM (bottom) M1 units (RM M1 units: $P = 6 \times 10^{-4}$, non-RM M1 units: $P = 0.40$, two-sample Kolmogorov–Smirnov test between distributions from pre- and post-sleep, followed by a shift test to assess how quartiles of the distributions differed; P values for the rest of Figure 3.10 legend reflect these statistical tests). **c.** Trial-averaged spiking activity during reaching for all reach modulated (RM) and non-RM DLS units across days and animals. **d.** Comparison of distributions of delta wave modulation during pre- and post-sleep for RM (top) and non-RM (bottom) DLS units (RM DLS units: $P = 0.17$, non-RM M1 units: $P = 0.51$). **e.** Comparison of distributions of delta wave modulation during pre- and post-sleep for RM DLS units with strong RM M1 input (top) and weak or no RM M1 input (bottom; RM DLS units with strong RM M1 input: $P = 0.14$, RM DLS units with weak or no RM M1 input: $P = 0.72$). **f.** Comparison of distributions of delta wave modulation during pre- and post-sleep for non-RM DLS units with strong RM M1 input (top) and weak or no RM M1 input (bottom; non-RM DLS units with strong RM M1 input: $P = 0.85$, non-RM DLS units with weak or no RM M1 input: $P = 0.01$).




Figure 3.11. Slow oscillation modulation change with training. **a.** Trial-averaged spiking activity during reaching for all reach modulated (RM) and non-RM M1 units across days and animals. **b.** Comparison of distributions of slow oscillation modulation during pre- and post-sleep for RM (top) and non-RM (bottom) M1 units (RM M1 units: $P = 2 \times 10^{-4}$, non-RM M1 units: $P = 0.02$, two-sample Kolmogorov–Smirnov test between distributions from pre- and post-sleep, followed by a shift test to assess how quartiles of the distributions differed; P values for the rest of Figure 3.11 legend reflect these statistical tests). **c.** Trial-averaged spiking activity during reaching for all reach modulated (RM) and non-RM DLS units across days and animals. **d.** Comparison of distributions of slow oscillation modulation during pre- and post-sleep for RM (top) and non-RM (bottom) DLS units (RM DLS units: $P = 0.28$, non-RM M1 units: $P = 0.44$). **e.** Comparison of distributions of slow oscillation modulation during pre- and post-sleep for RM DLS units with strong RM M1 input (top) and weak or no RM M1 input (bottom; RM DLS units with strong RM M1 input: $P = 0.27$, RM DLS units with weak or no RM M1 input: $P = 0.85$). **f.** Comparison of distributions of slow oscillation modulation during pre- and post-sleep for non-RM DLS units with strong RM M1 input (top) and weak or no RM M1 input (bottom; non-RM DLS units with strong RM M1 input: $P = 0.85$, non-RM DLS units with weak or no RM M1 input: $P = 0.14$).

Sleep spindle modulation predicts offline changes in corticostriatal transmission strength

We next sought to directly examine whether sleep spindle modulation following training was related to modifications of the corticostriatal network within offline periods. To measure modifications of the corticostriatal network during offline periods, we calculated cross correlations of spiking activity across individual pairs of M1 and DLS units during the first and second half of each pre- and post-sleep period (**Figure 3.12a-c**). Cross correlations were generated specifically with spiking activity during NREM to control for any differences in time spent in each behavioral state. Consistent changes in the short-latency cross correlation magnitude from the first to second half of pre- or post-sleep would indicate a modification of corticostriatal transmission strength within the offline period. Given the evidence of reactivation during sleep spindles between RM M1 and putatively connected DLS units, we first specifically examined transmission strength changes within the pairs in this population that were significantly modulated to spindles (708/3,969 pairs in pre-sleep and 1,062/3,969 pairs in post-sleep). Strikingly, we observed an increase in corticostriatal transmission strength during post-sleep, but no significant change during pre-sleep (**Figure 3.12d&e**; pre-sleep: one-sample t-test: $t(707) = -0.9$, $P = 0.37$; post-sleep: one-sample t-test: $t(1061) = 10.8$, $P = 6 \times 10^{-26}$). Furthermore, increases in corticostriatal transmission strength across individual pairs of M1 and DLS units during post-sleep, but not pre-sleep, were correlated to the mean sleep spindle modulation of that pair (**Figure 3.12f&g**). Importantly, the amount of time spent in NREM was similar during pre- and post-sleep (**Figure 3.13**). In contrast, transmission strength across RM M1 and putatively connected DLS unit pairs that were not significantly modulated to sleep spindles (1,164/3,969 pairs in pre-sleep and 989/3,967 pairs in post-sleep) did not change during either pre- or post-sleep and, within this population, changes in corticostriatal transmission strength across individual pairs of M1 and DLS units were not

significantly correlated to the mean sleep spindle modulation of that pair (Figure 3.14). These results suggested that, following training, offline corticostriatal transmission strength changes are linked to sleep spindle modulation.




Figure 3.12. Sleep spindle modulation predicts offline changes in corticostriatal coupling.

a. Schematic of NREM spiking activity snippets from example M1 and DLS units depicting the evolution of M1 and DLS spiking relationships from the first to second half of pre- (left) and post-sleep (right). **b.** Cross correlations of spiking activity during NREM from example M1 and DLS unit pair during the first and second half of pre-sleep, showing no change in cross correlation magnitude. **c.** Same as **b** for post-sleep, showing an increase in short-latency cross correlation magnitude. **d.** Cross correlations of spiking activity during NREM for all pairs of RM M1 and putatively connected DLS units that are significantly modulated to sleep spindles during the first and second half of pre-sleep, showing no change in cross correlation magnitude (width of line represents mean \pm SEM). **e.** Same as **d** for post-sleep, showing an increase in short-latency cross correlation magnitude. **f.** Correlation between change in short-latency cross correlation magnitude and mean sleep spindle modulation for all pairs of RM M1 and putatively connected DLS units that are significantly modulated to sleep spindles during pre-sleep. **g.** Same as **f** for post-sleep.




Figure 3.13. Comparison of time spent in each behavioral state during pre- and post-sleep.
a. Comparison of durations spent in each behavioral state during pre- and post-sleep across days and animals (NREM: 41.2 ± 2.4 minutes of pre NREM vs. 44.3 ± 2.5 minutes of post NREM, mean \pm SEM, paired-sample t-test: $t(51) = -1.03$, $P = 0.31$; REM: 17.2 ± 1.4 minutes of pre REM vs. 16.3 ± 1.4 minutes of post REM, mean \pm SEM, paired-sample t-test: $t(51) = 0.56$, $P = 0.58$; wake: 98.8 ± 3.6 minutes of pre wake vs. 106.0 ± 4.3 minutes of post wake, mean \pm SEM, paired-sample t-test: $t(51) = -1.5$, $P = 0.14$).




Figure 3.14. Corticostriatal transmission strength changes for non-sleep spindle modulated pairs of M1 and DLS units. **a.** Cross correlations of spiking activity during NREM for all pairs of RM M1 and putatively connected DLS units that are not significantly modulated to sleep spindles during the first and second half of pre-sleep (left) and post-sleep (right), showing no changes in cross correlation magnitude (width of line represents mean \pm SEM) **b.** Correlation between change in short-latency cross correlation magnitude and mean sleep spindle modulation for all pairs of RM M1 and putatively connected DLS units that are not significantly modulated to sleep spindles during pre-sleep (left) and post-sleep (right).

The interaction between sleep spindles and slow oscillations impact the role of sleep spindles within the corticostriatal network

To understand why changes in corticostriatal transmission strength occurred specifically in post-sleep, but not pre-sleep, we examined the interaction between sleep spindles and slow oscillations, a relationship known to be relevant for sleep-dependent processing^{18,38,39}. We found that the distribution of temporal proximity to preceding slow oscillations in post-sleep significantly differed from the distribution in pre-sleep, with slow oscillations in closer proximity to sleep spindles during post-sleep (**Figure 3.15a**; $P=2\times 10^{-29}$, two-sample Kolmogorov–Smirnov test). We found that this close proximity of slow oscillations to sleep spindles increased firing rates during sleep spindles in both M1 and DLS (**Figure 3.15b&c**, M1 units: $P = 0.02$, DLS units: $P = 0.02$, two-sample Kolmogorov–Smirnov tests, followed by a shift test to assess how quartiles of the distribution differed). Notably, the rate of sleep spindles within 500ms after a slow oscillation was correlated to the mean change in corticostriatal transmission strength across all pairs of M1 and DLS units during post-sleep (**Figure 3.15d**). This suggested that proximity to slow oscillations may be an important factor in whether sleep spindles drive plasticity. Altogether, our results provide evidence for offline plasticity within the corticostriatal network after training and that sleep spindles, and their interactions with slow oscillations, are important mediators of such plasticity.




Figure 3.15. The interaction between sleep spindles and slow oscillations impact the role of sleep spindles within the corticostriatal network. **a.** Distributions of the temporal proximity to preceding slow oscillations for all sleep spindles during pre- and post-sleep across days and animals. **b.** Firing rate across M1 units during sleep spindles with close proximity to slow oscillations (<0.5 seconds) and all other sleep spindles (width of line represents mean \pm SEM). **c.** Comparison of distributions of peak firing rates for M1 (left) and DLS (right) units during sleep spindles with close proximity to slow oscillations (<0.5 seconds) and all other sleep spindles. **d.** Correlation between each days' post-sleep density of sleep spindles in close proximity to slow oscillations (<0.5 seconds; normalized by subtracting the rate in pre-sleep and normalized within each animal by z-scoring across days) and mean change in short-latency cross correlation across all pairs of RM M1 and putatively connected DLS units that are significantly modulated to sleep spindles (normalized by subtracting the change in pre-sleep and normalized within each animal by z-scoring across days).

Discussion

Skilled behaviors exist in a continuum between being flexible - adapting quickly to changes in the environment - and automatic - inflexible to changes but cognitively efficient¹⁻⁴. Plasticity within the corticostriatal network is thought to regulate the balance between flexibility and automaticity^{4,5,26}. In this study, we study long-term training that resulted in automaticity, as evidenced by day-to-day invariance in reaching behavior that persisted even when the food pellet was moved such that reaches were no longer successful in retrieving the pellet. We show that, during such long-term training, corticostriatal functional connectivity increased during offline periods and provide evidence that sleep spindles uniquely engage the corticostriatal network to mediate such plasticity.

Our results provide evidence that sleep plays an important role in modifying cross-area connectivity during learning. While coordinated cross-area reactivations during sleep have been reported across several brain networks^{19,40-45}, how such coordinated activity patterns precisely shape the connectivity across brain regions remains largely unexplored (but see ¹⁹). For example, studies informing the *systems consolidation theory* have posited that coordinated activity patterns across the hippocampus and cortex during sleep drive intra-cortical plasticity^{23,25,46}. However, it is not known whether sleep impacts hippocampal-cortical connectivity, despite evidence for a change in hippocampal-cortical coupling after learning⁴⁷. Our results thus suggest that the *systems consolidation theory* may need to be broadened to consider a role for sleep in increasing the coupling between connected regions that can impact subsequent wake network activity and behavior. As evidence for cross-area reactivations in different brain networks continues to grow⁴⁸,

it will be important to consider how such activity patterns may impact both local and cross-area plasticity.

Our measures of cross-area connectivity are based on the coordination of LFP signals and single unit spike timing across M1 and DLS. We observed a subset of both LFP electrodes and corticostriatal neuron pairs showing evidence of increased connectivity with training, indicating the selective strengthening of corticostriatal connectivity - but what is the neural basis for these changes? One possibility is that our functional measures of connectivity reflect changes in synaptic strength of M1 projections to the DLS. This is consistent with evidence for the strengthening of cortical inputs to the striatum with motor training⁴⁹. An alternative possibility is that coordinated inputs to both M1 and DLS drive increased functional connectivity. We believe our results are most consistent with a physical change in synaptic strength, as we observed evidence of increased cross-area connectivity in two distinct states, NREM, reflected as increased LFP coherence, and awake task performance, reflected in the emergence of predictable cross-area dynamics. Future work is required to determine whether our observations are consistent with structural changes in synaptic strength.

We also provide evidence that sleep spindles are uniquely poised to mediate the enhancement of corticostriatal coupling during the offline period following training. While sleep spindles have been previously suggested to be important for plasticity^{36,50}, the precise link between sleep spindles, plasticity, and behavior has remained unclear. Here we also show that striatal NMDA activation during the offline periods following training is required for increases in behavioral consistency. This suggests that sleep spindles may be important drivers of corticostriatal plasticity

through NMDA activation. This is consistent with work showing that corticostriatal plasticity is NMDA-dependent^{13,14}, as well as *in vitro* work examining how sleep spindle activity patterns might drive plasticity⁵⁰. Additionally, we provide evidence that the proximity of sleep spindles to preceding slow oscillations is an important regulator of plasticity, consistent with previous work^{18,39}. As slow oscillations have been linked to NMDA receptor activation⁵¹, one intriguing possibility is that slow oscillations gate sleep spindle plasticity through the activation of NMDA receptors.

Our results link offline corticostriatal plasticity to the emergence of predictable low-dimensional cross-area activity. It has been previously demonstrated that M1 exhibits consistent low-dimensional population neural dynamics during consistently produced motor actions⁵². There is also growing evidence that subcortical regions such as the DLS are important for stabilizing cortical activity patterns and the emergence of consistent behaviors^{8,12}. Consistent with this idea, task-related coordination of M1 and DLS activity emerges with skill acquisition⁸⁻¹⁰. Here we link offline increases in corticostriatal functional connectivity to the ability to predict low-dimensional population activity in DLS from M1 activity. This suggests a model in which consistent low-dimensional neural dynamics emerge across the motor network with training and that motor network plasticity during offline periods is important for the emergence of such consistent cross-area dynamics.

Lastly, we link offline corticostriatal plasticity in the corticostriatal network to the emergence of fast and consistent reaching behavior, reflected in the invariance of day-to-day reaching velocity profile. This is consistent with a range of studies demonstrating that sleep benefits speed and

consistency in motor tasks in humans^{53,54} and rodents^{37,55}, as well as rodent brain-machine interface (BMI) tasks^{17,18}. Therefore, our results suggest the possibility that a fundamental role of sleep is to modify the corticostriatal network to impact the consistency of behavior in a range of tasks. Further work is required to determine the precise role of offline corticostriatal plasticity in different contexts. One important avenue of research is to explore whether sleep can impact corticostriatal connectivity in the context of maladaptive automatic behaviors, such as addiction, that have been linked to the corticostriatal network^{26,56}. Notably, there is evidence that the reactivation of a stored memory can make the memory temporarily labile²⁷ and recent work has shown that the modulation of NREM rhythms can regulate modulate memory consolidation vs. forgetting¹⁸. Therefore, it will be informative to determine whether similar manipulations could be used in the context of maladaptive automatic behaviors to provide a therapeutic benefit.

Methods

Animal care and surgery

This study was performed in strict accordance with guidelines from the USDA Animal Welfare Act and United States Public Health Science Policy. Procedures were in accordance with protocols approved by the Institutional Animal Care and Use Committee at the San Francisco Veterans Affairs Medical Center. This study consists of experiments performed with fourteen male Long-Evans rats (approximately 12-16 weeks old), housed under controlled temperature and a 12-h light/12-h dark cycle with lights on at 6:00 a.m. Animal experiments were performed during the light period. All surgical procedures were performed using sterile techniques under 2–4% isoflurane. Six animals were implanted with either microwire electrodes ($n = 4$ animals; 32 or 64 channel 33 μ m diameter Tungsten microwire arrays with ZIF-clip adapter; Tucker-Davis

Technology) or high-density silicon probes ($n = 2$ animals; 256 channel custom-built silicon probes) targeted to both the forelimb area of M1, centered at 3.5mm lateral and 0.5mm anterior to bregma and implanted in layer V at a depth of 1.5mm, and the DLS, centered at 4mm lateral and 0.5mm anterior to bregma and implanted at a depth of 4mm. Six additional animals were implanted with infusion cannulas (PlasticsOne; 26Ga) targeted to the DLS. Surgery involved exposure and cleaning of the skull, preparation of the skull surface (using cyanoacrylate), and implantation of skull screws for overall headstage stability. In the animals implanted with neural probes, a reference screw was implanted posterior to lambda, contralateral to the neural recordings and a ground screw was implanted posterior to lambda, ipsilateral to the neural recordings. Craniotomy and durotomy were then performed, followed by implantation of neural probes or infusion cannulas and securing of the implant with C&B Metabond (Parkell, Product #S380) and Duralay dental acrylic (Darby, Product #8830630). In four of the animals implanted with neural probes, the forearm was also implanted with a pair of twisted electromyography (EMG) wires (0.007" single-stranded, Teflon-coated, stainless steel wire; A-M Systems) with a hardened epoxy ball (J-B Weld Company) at one end preceded by 1–2mm of uncoated wire under the ball. Wires were inserted into the muscle belly and pulled through until the ball came to rest on the belly. EMG wires were braided, tunneled under the skin to a scalp incision and soldered into an electrode interface board (ZCA-EIB32; Tucker-Davis Technology). The postoperative recovery regimen included administration of buprenorphine at 0.02mg/kg and meloxicam at 0.2mg/kg. Dexamethasone at 0.5mg/kg and trimethoprim/sulfadiazine at 15mg/kg were also administered postoperatively for 5 days. All animals recovered for at least one week before the start of behavioral training.

***In vivo* electrophysiology**

Units, local field potentials (LFP), and EMG activity were recorded using an RZ2 system (Tucker-Davis Technologies). For the microwire animals, spike data was sampled at 24,414Hz and LFP/EMG data at 1,017Hz. To record spiking data in these animals, thresholds for spiking activity were set online using a standard deviation of 4.5 (calculated over a 1-min baseline period using the RZ2 system). Waveforms and timestamps were stored for any event that crossed that threshold. Spike sorting was then performed using Offline Sorter v.4.3.0 (Plexon) with a principal component analysis-based clustering method followed by manual inspection. Spikes were sorted separately for each day, combining pre-sleep, training, and post-sleep sessions. We accepted units based on waveform shape, clear cluster boundaries in principal component space and 99.5% of detected events with an ISI>2ms. For silicon probe animals, signals were recorded at 24,414Hz. In these animals, spike times and waveforms were detected from the broadband signal using Offline Sorter v.4.3.0 (Plexon). Spike waveforms were then sorted using Kilosort2 (<https://github.com/MouseLand/Kilosort2>). We accepted units based on manual inspection using Phy (<https://github.com/cortex-lab/phy>) and 99.5% of detected events with an ISI>2ms.

Viral injection

To label anterograde projections in M1 we injected 750nl of AAV8-hsyn-JAWs-KGC-GFP-ER2 virus into two sites (1.5mm anterior, 2.7mm lateral to bregma, at a depth of 1.4mm and 0.5 posterior, 3.5mm lateral to bregma, at a depth of 1.4mm). Two weeks after injection rats were anesthetized and transcardially perfused with 0.9% sodium chloride, followed by 4% formaldehyde. The harvested brains were post-fixed for 24 h and immersed in 20% sucrose for

2 days. Coronal cryostat sections (40- μ m thickness) were then mounted and imaged with a fluorescent microscope.

Reach-to-grasp task

Rats naïve to any motor tasks were first tested for forelimb preference. This involved presenting approximately ten food pellets to the animal and observing which forelimb was most often used to reach for the pellet. Rats then underwent surgery for either neural probe or cannula implantation in the hemisphere contralateral to preferred paw. Following the one-week recovery period, rats were trained using an automated reach-box, controlled by custom MATLAB scripts and an Arduino microcontroller. This setup requires minimal user intervention, as described previously⁵⁵. Each trial consisted of a pellet dispensed on the pellet tray followed by an alerting beep indicating that the trial was beginning, then the door would open. Animals had to reach, grasp, and retrieve the pellet. A real-time ‘pellet detector’ using an infrared sensor centered over the pellet was used to determine when the pellet was moved, indicating the trial was over and then the door was closed. All trials were captured by a camera placed on the side of the behavioral box ($n = 2$ animals monitored with a Microsoft LifeCam at 30 frames/second; $n = 12$ animals monitored with a Basler ace acA640-750uc at 75 frames/second). For animals implanted with neural probes, each animal underwent five to fourteen days of training (~100–150 trials per day). For the infusion cannula implanted animals, each animal underwent ten days of training (100 trials per day). Rats had fifteen seconds to complete each trial, and trials were separated by a ten second inter-trial-interval. Reach trajectories were captured from video using DeepLabCut⁵⁶ to track the center of the rat’s paw as well as the food pellet. Reach trajectories consisted of the paw trajectory from 500ms before to 500ms after “pellet touch”, which was classified as the frame in which the paw was closest to the

pellet, before the pellet was displaced off the pellet holder. Only trials in which the pellet was displaced off the pellet holder were considered. We assessed behavioral consistency throughout training in both neural probe and cannula implanted animals by calculating the correlation between the mean velocity profile of reaches on each day of training and the mean velocity profile of reaches on the last day of training. These correlations were computed separately for the x and y dimensions and then averaged. At the end of training, we tested whether reaching behavior was automatic in two of the neural probe implanted animals by performing a 100 trial training session on the subsequent day with the pellet moved to a new location (~10mm lateral from original pellet position) and observing whether the animal's reaching behavior changed. We performed a similar experiment for two additional animals naïve to the task and without neural implant to test whether reaching was flexible or automatic at the start of training. These animals performed ~200 trials on two consecutive days. To calculate single-trial reach trajectory correlations, we first generated a mean trajectory in each dimension (x and y) for trials with the pellet in the learned position and trials with the pellet in the new position (mean trajectories were computed separately for each pellet position). Single trial trajectories were then correlated to the mean trajectory in each dimension and then averaged across the x and y dimension. To compare across pellet positions, we considered reach trajectories up to pellet touch (from 500ms before pellet touch to pellet touch), as automatic reaches with the pellet in the new position often missed the pellet and pellet holder completely.

DLS infusions

To test if blocking the activation of striatal NMDA receptors during the offline period after training disrupts increases in behavioral consistency, we infused either 1ul of saline or NMDA blocker

AP5 (5 μ g/ μ l) at an infusion rate 200nl/minute into the DLS immediately following training in six animals for ten consecutive days. In the first five days of training, we infused three rats with AP5 and three rats with saline, for the second five days, we switched the infusion, i.e., animals that received AP5 in the first five days, received saline for the second five days, and vice-versa.

Sleep classification

All neural data analyses were conducted using MATLAB 2019a (MathWorks) and functions from the EEGLAB (<http://sccn.ucsd.edu/eeglab/>) and Chronux (<http://chronux.org/>) toolboxes. Sleep was classified using cortical LFP signals and movement measured by video or EMG activity. LFP was preprocessed by artifact rejection, including manual rejection of noisy channels and z-scoring of each channel across the entire recording session. A mean LFP channel was then generated in M1 for sleep classification by averaging across all M1 channels. This mean M1 LFP channel was then segmented into non-overlapping 10 second windows. In each window the power spectral density was computed using the Chronux function *mtspecgramc* and then averaged over the delta (1–4Hz) and theta (5-10Hz/2-15Hz) frequency bands. Both LFP power bands were then normalized by z-scoring. Epochs with high delta power (>0 z-scored delta) and no movement were classified as NREM, epochs with high theta and low delta power (>0 z-scored theta and <0 z-scored delta) were classified as REM sleep, and other epochs were classified as wake (Watson, et al., 2016). All consecutive NREM or REM epochs that were less than 30 seconds long (3 consecutive epochs) were reclassified as wake.

Assessing corticostriatal functional connectivity using LFP coherence

To measure corticostriatal functional connectivity across days, we measured LFP coherence during NREM across all M1 and DLS electrode pairs on each pre- and post-sleep session using chronux function *cohgramc*. For these analyses, we first applied common-mode referencing using the median signal, i.e., at every time-point, the median signal across all channels in a region was calculated and subtracted from every channel to decrease common noise and minimize volume conduction. Common-mode referencing was performed independently for the channels in each region, i.e., M1 and DLS. We classified “high coherence LFP pairs” as electrodes with a mean 4-8Hz coherence >0.6 . To compare online changes in LFP coherence (from pre- to post-sleep on the same day) to offline changes in LFP coherence (from post-sleep on one day to pre-sleep on the next day), we computed a single value per pair for both online and offline coherence changes by averaging values across days of training.

Predicting cross-area activity

To assess cross-area dynamics, we first extracted low-dimensional representations of DLS activity by performing principal component analysis (PCA) on trial-averaged activity of DLS neurons time-locked to pellet touch and binned at 100ms, specifically for time bins from five seconds before to five seconds after pellet touch. Principal components were computed using MATLAB function *pca*. Spiking activity from five seconds before to five seconds after pellet touch and binned at 100ms was then projected onto each of the first two components to generate low-dimensional neural trajectory representations of population activity in DLS. We then fit a linear regression model to predict DLS reach-related neural trajectories from one second before to one second after pellet touch from single unit spiking activity in M1. A separate model was used to

predict each principle component, using MATLAB function *fitlm* and five-fold cross validation. For each time bin of the neural trajectory, the preceding 500ms of spiking activity for all M1 units, binned at 100ms, were used as predictors. A model was also fit on baseline, non-reaching, neural trajectories, calculated by projecting DLS spiking activity from five seconds to four second before pellet touch onto each of the first two computed principal components. The predictive ability of these models was assessed by calculating the correlation between the actual neural trajectories and the predicted trajectories.

NREM rhythm detection

The NREM rhythm detection applied here is based on an algorithm we have developed previously^{18,37}. A mean LFP channel was generated in M1 for NREM rhythm classification by averaging across all channels (same as used for sleep classification). To detect sleep spindles, this mean signal was filtered in the spindle band (10 – 16 Hz) using a zero-phase shifted, third order Butterworth filter. A smoothed envelope was calculated by computing the magnitude of the Hilbert transform of this signal then convolving it with a Gaussian window. Next, we determined two upper thresholds for spindle detection based on the mean and standard deviation (s.d.) of the spindle band envelope during NREM. Epochs in which the spindle envelope exceeded 2.5 s.d. above the mean for at least one sample and the spindle power exceeded 1.5 s.d. above the mean for at least 500ms were detected as spindles. Then, spindles that were sufficiently close in time (<300 ms) were combined. To detect slow oscillations and delta waves, the mean M1 signal was filtered in a low frequency band (2nd order, zero phase shifted, high pass Butterworth filter with a cutoff at 0.1Hz followed by a 5th order, zero phase shifted, low pass Butterworth filter with a cutoff at 4Hz). Next, all positive-to-negative zero crossings during NREM were identified, along with

the previous peaks, the following troughs, and the surrounding negative-to-positive zero crossings. Each identified epoch was considered a slow oscillation if the peak was in the top 15% of peaks, the trough was in the top 40% of troughs and the time between the negative-to-positive zero crossings was greater than 300ms but did not exceed 1 second. Each identified epoch was considered a delta wave if the peak was in the bottom 85% of peaks, the trough was in the top 40% of troughs and the time between the negative-to-positive zero crossings was greater than 250ms.

Characterizing putatively monosynaptically connected M1 and DLS

We characterized putatively monosynaptically connected pairs of M1 and DLS units by calculating the cross correlation of spiking activity binned at 1ms during the first five minutes of NREM during pre- and post-sleep concatenated together (10 minutes total) on each day of training for each pair of M1 and DLS units. We then measured the mean value of the short-latency cross correlation for each pair (1-10ms time lag centered on DLS spiking; consistent with the conduction and synaptic delay between M1 and DLS⁹) and compared this value to a shuffled distribution generated by shuffling DLS spike time bins and recalculating the cross correlation 1,000 times. If the non-shuffled short-latency correlation magnitude was greater than 95% of the shuffled distribution values, we classified the pair of units as putatively connected.

Comparing corticostriatal transmission strength across behavioral

To compare corticostriatal transmission strength across behavioral states, we generated a cross correlation of spiking activity binned at 1ms from each behavioral state (NREM, REM, and wake) for all putatively connected pairs of M1 and DLS units, during both pre- and post-sleep. To account

for firing rate differences across states, each pair's cross correlation was normalized by subtracting the mean cross correlation values from 100-150ms time lag.

Comparing corticostriatal transmission strength across NREM

To compare corticostriatal transmission strength across NREM rhythms, we generated a cross correlation of spiking activity binned at 1ms from each NREM rhythm (sleep spindles, delta waves, and slow oscillations) for all putatively connected pairs of M1 and DLS units. Spiking during sleep spindles consisted of spiking during the one second centered on sleep spindle peak (-500ms to 500ms). Spiking during slow oscillations and delta waves consisted of spiking during the one second around upstate peak (-500ms to 500ms). To account for the influence of firing rate differences or changes in LFP-phase locking across NREM rhythms, we applied a normalization step we previously developed³⁷. Briefly, we generated shuffled cross correlations between each M1 and DLS unit pair, with DLS spike times shuffled with respect to the NREM rhythm in which it fired. In this approach, both units maintain all their first-order relationships with the NREM rhythm; for example, the number of spikes, phase locking values, and phase preferences of individual units do not change after shuffling. However, the shuffling breaks the statistical relationship between the two neurons under examination. We repeated this shuffling 25 times and then subtracted the mean shuffled cross correlation from the unshuffled cross correlation.

NREM rhythm modulation

To determine the sleep spindle modulation of individual M1 and DLS units, spiking during each sleep spindle was time locked to the peak of the filtered LFP and binned at 10ms. Spiking was averaged across sleep spindles and modulation was calculated by taking the minimum to maximal

firing rate bin in the second around sleep spindle peak (-500ms to 500ms) divided by the minimum to maximal firing rate bin in a second long baseline period before each spindle (-1500ms to -500ms relative to spindle peak). To determine slow oscillation and delta wave modulation of individual M1 and DLS units, spiking during each slow oscillation or delta wave was time locked to the peak of the upstate and binned at 10ms. Spiking was averaged across slow oscillations or delta waves and modulation was calculated by taking the minimum to maximal firing rate bin in the second around upstate peak (-500ms to 500ms) divided by the minimum to maximal firing rate bin in a second long baseline period before each slow oscillation or delta wave (-1500ms to -500ms relative to upstate peak).

Characterizing reach modulated (RM) units

To characterize M1 and DLS reach modulated units, we generated trial-averaged peri-event time histograms (PETHs) of spiking activity for individual units during reaching locked to pellet touch in 25ms bins, from 5 seconds before to 5 seconds after pellet touch (400 total bins). Each unit's PETH was then z-scored and reach modulation was measured by taking the sum of the absolute value of the time bins from 1 second before pellet touch to 1 second after pellet touch (80 total bins). We then generated a distribution of shuffled modulations by shuffling all time bins and recalculating the modulation of the shuffled PETH and repeating this shuffling procedure one thousand times. Units with a non-shuffled modulation greater than the 99% percentile of the shuffled distribution were considered significantly reach modulated.

Characterizing DLS units with strong or weak M1 reach modulated input

To characterize DLS units with strong or weak M1 reach modulated input, we calculated the number of reach modulated M1 units that were putatively connected to each DLS unit. If a DLS unit was connected to 3 or more reach modulated M1 units, we classified that DLS unit as having strong M1 reach modulated input, if a DLS unit was connected to 2 or less M1 reach modulated units, we classified that DLS unit as having weak or no M1 reach modulated input.

Measuring corticostriatal transmission strength changes within pre- and post-sleep

To measure changes in corticostriatal transmission strength within pre- and post-sleep, we generated a cross correlation of spiking activity binned at 1ms from NREM activity during the first and second half of pre- and post-sleep. This was done for two populations of M1 and DLS unit pairs. The first population was all M1 and DLS unit pairs that contained a RM M1 unit, a DLS unit that was putatively connected to a RM M1 unit, and contained both M1 and DLS units that were significantly modulated to sleep spindles. The second population was all M1 and DLS unit pairs that contained a RM M1 unit, a DLS unit that was putatively connected to a RM M1 unit, and contained M1 and DLS units that were both not significantly modulated to sleep spindles. To determine which units were modulated to sleep spindles, we generated peri-event time histograms (PETHs) of sleep spindle activity locked to spindle peak in 10ms bins from 2 seconds before to 2 second after spindle peak (400 bins), averaged across all spindles. Sleep spindle modulation was then calculated by taking the minimum to maximal firing rate bin within the 1 second period centered on spindle peak (-500ms to 500ms). We then generated a distribution of shuffled modulations by shuffling the time bins and recalculating the modulation of this shuffled PETH. This shuffling procedure was repeated one thousand times to generate a distribution. Units with a

non-shuffled modulation greater than the 99% percentile of the shuffled distribution were considered significantly sleep spindle modulated.

Sleep spindle and slow oscillation proximity

Slow oscillation to sleep spindle proximity was determined by measuring the temporal proximity of the preceding slow oscillation zero-crossing (positive to negative LFP) to each sleep spindle peak. To determine the influence of slow oscillation proximity on sleep spindle modulation, we generated two PETHs locked to spindle peak and binned at 10ms for each unit. The first PETH was generated with sleep spindles that had a preceding slow oscillation within 500ms (“nested spindles”) and the second PETH was generated with sleep spindles that did not have a preceding slow oscillation within 500ms (“isolated spindles”). As there were more isolated spindles than nested spindles, the number of events used to generate each PETH was matched by randomly selecting isolated spindles to match the number of nested spindles. Modulation was then assessed by determining the peak firing rate bin in each PETH.

References

1. Isoda, M. & Hikosaka, O. Cortico-basal ganglia mechanisms for overcoming innate, habitual and motivational behaviors. *Eur. J. Neurosci.* **33**, 2058–2069 (2011).
2. Smith, K. S. & Graybiel, A. M. Habit formation. *Dialogues Clin. Neurosci.* **18**, 33–43 (2016).
3. Robbins, T. W. & Costa, R. M. Habits. *Current Biology* (2017). doi:10.1016/j.cub.2017.09.060
4. Graybiel, A. M. Habits, Rituals, and the Evaluative Brain. *Annu. Rev. Neurosci.* **31**, 359–387 (2008).
5. Yin, H. H. & Knowlton, B. J. The role of the basal ganglia in habit formation. *Nat. Rev. Neurosci.* **7**, 464–476 (2006).
6. Costa, R. M., Cohen, D. & Nicolelis, M. A. L. Differential Corticostriatal Plasticity during Fast and Slow Motor Skill Learning in Mice. *Curr. Biol.* **14**, 1124–1134 (2004).
7. Neely, R. M., Koralek, A. C., Athalye, V. R., Costa, R. M. & Carmena, J. M. Volitional Modulation of Primary Visual Cortex Activity Requires the Basal Ganglia. *Neuron* **0**, (2018).
8. Lemke, S. M., Ramanathan, D. S., Guo, L., Won, S. J. & Ganguly, K. Emergent modular neural control drives coordinated motor actions. *Nat. Neurosci.* **22**, 1122–1131 (2019).
9. Koralek, A. C., Costa, R. M. & Carmena, J. M. Temporally Precise Cell-Specific Coherence Develops in Corticostriatal Networks during Learning. *Neuron* **79**, 865–872 (2013).
10. Santos, F. J., Oliveira, R. F., Jin, X. & Costa, R. M. Corticostriatal dynamics encode the refinement of specific behavioral variability during skill learning. *Elife* **4**, (2015).
11. Dang, M. T. *et al.* Disrupted motor learning and long-term synaptic plasticity in mice

- lacking NMDAR1 in the striatum. *Proc. Natl. Acad. Sci. U. S. A.* (2006). doi:10.1073/pnas.0601758103
12. Koralek, A. C., Jin, X., Long II, J. D., Costa, R. M. & Carmena, J. M. Corticostriatal plasticity is necessary for learning intentional neuroprosthetic skills. *Nature* **483**, 331–335 (2012).
 13. Calabresi, P., Pisani, A., Mercuri, N. B. & Bernardi, G. Long-term Potentiation in the Striatum is Unmasked by Removing the Voltage-dependent Magnesium Block of NMDA Receptor Channels. *Eur. J. Neurosci.* (1992). doi:10.1111/j.1460-9568.1992.tb00119.x
 14. Charpier, S. & Deniau, J. M. In vivo activity-dependent plasticity at cortico-striatal connections: Evidence for physiological long-term potentiation. *Proc. Natl. Acad. Sci. U. S. A.* (1997). doi:10.1073/pnas.94.13.7036
 15. Yang, G. *et al.* Sleep promotes branch-specific formation of dendritic spines after learning. *Science (80-.)*. **344**, 1173–1178 (2014).
 16. Ramanathan, D. S., Gulati, T. & Ganguly, K. Sleep-Dependent Reactivation of Ensembles in Motor Cortex Promotes Skill Consolidation. *PLOS Biol.* **13**, e1002263 (2015).
 17. Gulati, T., Ramanathan, D. S., Wong, C. C. & Ganguly, K. Reactivation of emergent task-related ensembles during slow-wave sleep after neuroprosthetic learning. *Nat. Neurosci.* **17**, 1107–1113 (2014).
 18. Kim, J., Gulati, T. & Ganguly, K. Competing Roles of Slow Oscillations and Delta Waves in Memory Consolidation versus Forgetting. *Cell* **179**, 514-526.e13 (2019).
 19. Vahdat, S., Fogel, S., Benali, H. & Doyon, J. Network-wide reorganization of procedural memory during NREM sleep revealed by fMRI. *Elife* **6**, (2017).
 20. Boutin, A. *et al.* Transient synchronization of hippocampo-striato-thalamo-cortical

- networks during sleep spindle oscillations induces motor memory consolidation. *Neuroimage* **169**, 419–430 (2018).
21. Doyon, J., Gabbitov, E., Vahdat, S., Lungu, O. & Boutin, A. Current issues related to motor sequence learning in humans. *Curr. Opin. Behav. Sci.* **20**, 89–97 (2018).
 22. Doyon, J. & Benali, H. Reorganization and plasticity in the adult brain during learning of motor skills. *Curr. Opin. Neurobiol.* **15**, 161–167 (2005).
 23. Klinzing, J. G., Niethard, N. & Born, J. Mechanisms of systems memory consolidation during sleep. *Nat. Neurosci.* **22**, 1598–1610 (2019).
 24. Rasch, B. & Born, J. About Sleep’s Role in Memory. *Physiol. Rev.* **93**, 681–766 (2013).
 25. Squire, L. R., Genzel, L., Wixted, J. T. & Morris, R. G. Memory consolidation. *Cold Spring Harb. Perspect. Biol.* **7**, a021766 (2015).
 26. Lipton, D. M., Gonzales, B. J. & Citri, A. Dorsal Striatal Circuits for Habits, Compulsions and Addictions. *Front. Syst. Neurosci.* **13**, 28 (2019).
 27. Lee, J. L. C., Nader, K. & Schiller, D. An Update on Memory Reconsolidation Updating. *Trends Cogn. Sci.* **21**, 531–545 (2017).
 28. Egert, D. G. *et al.* Cellular-scale silicon probes for high-density, precisely-localized neurophysiology. *J. Neurophysiol.* (2020). doi:10.1152/jn.00352.2020
 29. Aoki, S. *et al.* An open cortico-basal ganglia loop allows limbic control over motor output via the nigrothalamic pathway. *Elife* **8**, (2019).
 30. Watson, B. O., Levenstein, D., Greene, J. P., Gelinás, J. N. & Buzsáki, G. Network Homeostasis and State Dynamics of Neocortical Sleep. *Neuron* **90**, 839–852 (2016).
 31. Rueda-Orozco, P. E. & Robbe, D. The striatum multiplexes contextual and kinematic information to constrain motor habits execution. *Nat. Neurosci.* **18**, 453–460 (2015).

32. Flint, R. D., Scheid, M. R., Wright, Z. A., Solla, S. A. & Slutzky, M. W. Long-Term Stability of Motor Cortical Activity: Implications for Brain Machine Interfaces and Optimal Feedback Control. *J. Neurosci.* **36**, 3623 (2016).
33. Yazdan-Shahmorad, A., Silversmith, D. B., Kharazia, V. & Sabes, P. N. Targeted cortical reorganization using optogenetics in non-human primates. *Elife* **7**, (2018).
34. Thorn, C. A. & Graybiel, A. M. Differential Entrainment and Learning-Related Dynamics of Spike and Local Field Potential Activity in the Sensorimotor and Associative Striatum. *J. Neurosci.* **34**, 2845–2859 (2014).
35. Huber, R., Felice Ghilardi, M., Massimini, M. & Tononi, G. Local sleep and learning. *Nature* **430**, 78–81 (2004).
36. Durkin, J. *et al.* Cortically coordinated NREM thalamocortical oscillations play an essential, instructive role in visual system plasticity. *Proc. Natl. Acad. Sci.* **114**, 10485–10490 (2017).
37. Silversmith, D. B., Lemke, S. M., Egert, D., Berke, J. D. & Ganguly, K. The Degree of Nesting between Spindles and Slow Oscillations Modulates Neural Synchrony. *J. Neurosci.* **40**, 4673–4684 (2020).
38. Niethard, N., Ngo, H.-V. V., Ehrlich, I. & Born, J. Cortical circuit activity underlying sleep slow oscillations and spindles. *Proc. Natl. Acad. Sci.* **115**, E9220–E9229 (2018).
39. Hoffman, K. L. & McNaughton, B. L. Coordinated reactivation of distributed memory traces in primate neocortex. *Science (80-.)*. **297**, 2070–2073 (2002).
40. Lansink, C. S., Goltstein, P. M., Lankelma, J. V., McNaughton, B. L. & Pennartz, C. M. A. Hippocampus Leads Ventral Striatum in Replay of Place-Reward Information. *PLoS Biol.* **7**, e1000173 (2009).
41. Sjulson, L., Peyrache, A., Cumpelik, A., Cassataro, D. & Buzsáki, G. Cocaine Place

- Conditioning Strengthens Location-Specific Hippocampal Coupling to the Nucleus Accumbens. *Neuron* **98**, 926-934.e5 (2018).
42. Ji, D. & Wilson, M. A. Coordinated memory replay in the visual cortex and hippocampus during sleep. *Nat. Neurosci.* **10**, 100–107 (2007).
 43. Girardeau, G., Inema, I. & Buzsáki, G. Reactivations of emotional memory in the hippocampus-amygdala system during sleep. *Nat. Neurosci.* **20**, 1634–1642 (2017).
 44. Ribeiro, S. *et al.* Long-Lasting Novelty-Induced Neuronal Reverberation during Slow-Wave Sleep in Multiple Forebrain Areas. *PLoS Biol.* **2**, e24 (2004).
 45. Yu, J. Y., Liu, D. F., Loback, A., Grossrubatscher, I. & Frank, L. M. Specific hippocampal representations are linked to generalized cortical representations in memory. *Nat. Commun.* (2018). doi:10.1038/s41467-018-04498-w
 46. Tingley, D. & Peyrache, A. On the methods for reactivation and replay analysis. *Philosophical Transactions of the Royal Society B: Biological Sciences* (2020). doi:10.1098/rstb.2019.0231
 47. Yin, H. H. *et al.* Dynamic reorganization of striatal circuits during the acquisition and consolidation of a skill. **12**, (2009).
 48. Rosanova, M. & Ulrich, D. Pattern-specific associative long-term potentiation induced by a sleep spindle-related spike train. *J. Neurosci.* **25**, 9398–9405 (2005).
 49. Chauvette, S., Seigneur, J. & Timofeev, I. Sleep oscillations in the thalamocortical system induce long-term neuronal plasticity. *Neuron* **75**, 1105 (2012).
 50. Churchland, M. M. *et al.* Neural population dynamics during reaching. *Nature* **487**, 51–6 (2012).
 51. Fischer, S., Hallschmid, M., Elsner, A. L. & Born, J. Sleep forms memory for finger skills.

- Proc. Natl. Acad. Sci.* **99**, 11987–11991 (2002).
52. Walker, M. P., Brakefield, T., Morgan, A., Hobson, J. A. & Stickgold, R. Practice with Sleep Makes Perfect: Sleep-Dependent Motor Skill Learning. *Neuron* **35**, 205–211 (2002).
 53. Nagai, H. *et al.* Sleep consolidates motor learning of complex movement sequences in mice. *Sleep* (2017). doi:10.1093/sleep/zsw059
 54. Gerdeman, G. L., Partridge, J. G., Lupica, C. R. & Lovinger, D. M. It could be habit forming: drugs of abuse and striatal synaptic plasticity. *Trends Neurosci.* **26**, 184–192 (2003).
 55. Wong, C. C., Ramanathan, D. S., Gulati, T., Won, S. J. & Ganguly, K. An automated behavioral box to assess forelimb function in rats. *J. Neurosci. Methods* **246**, 30–7 (2015).
 56. Mathis, A. *et al.* DeepLabCut: markerless pose estimation of user-defined body parts with deep learning. *Nat. Neurosci.* **21**, 1281–1289 (2018).
 57. Watson, B. O., Levenstein, D., Greene, J. P., Gelineas, J. N. & Buzsáki, G. Network Homeostasis and State Dynamics of Neocortical Sleep. *Neuron* **90**, 839–852 (2016).

Chapter 4: Why corticostriatal plasticity matters: a model of motor network organization

There is considerable evidence that the motor network is distributed, that is, activity in multiple brain regions in the motor network can drive movement (see Chapter 1). This raises the question, how are functions distributed across this network? One influential organizational scheme was put forth by Hanricus (Hans) Kuypers, who differentiated the function of brain regions based on connectivity to the spinal cord¹:

From these findings it has been concluded that the medially descending group A brain stem pathways, which are especially derived from centrally located brain stem structures, characteristically steer body and integrated limb and body movements as well as movement synergisms of the individual limbs involving their various parts. Group B brain stem pathways also exist; these descend laterally. Components of this group appear to add further resolution to brain stem control, and provide the capacity to execute relatively independent movements of the limbs, especially of their distal parts. The cortical pathways to the spinal cord and lower brain stem, which parallel the descending brain stem pathways, further amplify the brain stem control but, especially in primates, also provide the unique capacity to execute highly fractionated movements, exemplified by individual finger movements.

An alternative scheme, put forth by Okihide Hikosaka, differentiated function based on the “level” of control, rather than specific body parts controlled²:

I now propose that there are at least three levels in the control of action. The first level controls innate movements (Innate action). The second level controls movements that have been acquired by practice (Learned action). The third level controls movements that have not yet become automatic and require

attention and effort (New action). It should be stressed, however, that such non-automatic movements will eventually become automatic after long-term practice.

At first glance, these two motor network organization schemes seem quite distinct. On one hand, motor control is differentiated based on body part (I will call this “effector level”), while on the other hand, motor control is differentiated based on the level of control (“control level”) – how can we make sense of these two distinct schemes? In this chapter, I propose a model of motor network organization that aims to find a common axis between these schemes.

The proposed model is based on an axis of neural flexibility-to-stability. I propose that both the “effector level” and “control level” organizational schemes can be successfully mapped onto this axis (**Figure 4.1**). In the case of effector control, we can think of a gradient from proximal to distal movements. For example, distal effectors, such as the hand, have higher degrees-of-freedom (DOF) and thus require more flexible neural control than lower DOF proximal effectors such as the shoulder. Similarly, for the “control level”, performing and learning a new behavior requires more flexibility compared to executing a stable innate behavior. Therefore, we can map both “effector level” and “control level” schemes onto a shared neural axis. But how does this shared axis map onto the motor network? And how is corticostriatal plasticity involved? In this chapter I will argue for the following two claims:




Figure 4.1. Proposed shared neural axis for different motor control schemes.

- (1) There exists a cortical-to-subcortical gradient of neural flexibility-to-stability across the motor network.
- (2) Learning a new behavior involves transitioning from cortical to subcortical control of a behavior and is mediated through corticostriatal plasticity.

In this chapter, I will present evidence for the above claims, then discuss future work and conflicting evidence relating to the proposed model.

A cortical-to-subcortical gradient in neural flexibility

The first claim is that *there exists a cortical-to-subcortical gradient of neural flexibility-to-stability across the motor network* (**Figure 4.2**). This gradient is proposed to exist across the motor cortex, basal ganglia, and brainstem. I will begin by providing evidence for greater neural flexibility in cortex compared to the basal ganglia, and then discuss evidence that this gradient may exist across all three regions. To do this, we must first determine how to measure neural flexibility. Imagine we measure the activity in two networks, each made up of ten neurons. Given the size of the network and assuming we binarize neural activity such that, within each time bin, each neuron either has or does not have a spike (0 or 1), there are 1,024 possible neural activity patterns in this network (i.e., 1000000000, 0100000000, ... = 2^{10}). I propose a simple measure of neural flexibility vs. stability based on how often specific neural patterns are repeated. A stable network will repeat the same neural patterns more often than a flexible network.

To look for evidence of this gradient I compared the dynamics of neural activity from ten-neuron networks in primary motor cortex (M1) and the dorsolateral striatum (DLS). The data was obtained from rats with both M1 and DLS networks recorded simultaneously. Each recording was ~2 hours long and contained both wake and sleep, but no specific behavioral task. One consideration when comparing network dynamics between M1 and DLS is the difference




Figure 4.2. Proposed cortical-to-subcortical gradient in neural flexibility-to-stability.

in firing rate (**Figure 4.3a**). Firing rate is known to influence pairwise correlations³. Therefore, I generated circularly shuffled networks of M1 and DLS neurons in which each neuron retained its firing rate but pairwise correlations across the population were removed, creating an “independent




Figure 4.3. M1 and DLS spiking properties. **a.** Firing rate distributions in M1 and DLS. **b.** Circular shuffling method. **c.** Comparison of pairwise correlation between real and shuffled M1 (left) and DLS. (right)

network” (Figure 4.3b&c). Then, I compared how often the same neural patterns were produced in the real M1 and DLS networks, compared to the shuffled M1 and DLS networks. This approach aimed to reveal how removing the correlations in a network impacted neural flexibility. In DLS the increase in frequency of neural patterns between the real and independent networks was greater than in M1 (Figure 4.4), indicating that correlations in DLS drive neural stability to a greater degree than they do in M1, providing evidence towards a cortical-to-subcortical gradient in neural flexibility.




Figure 4.4. Comparison of M1 and DLS neural stability vs. flexibility. **a.** Scatterplot of pattern occurrences in real and shuffled network for M1 (left) and DLS (right). **b.** Comparison of distributions between ratio of real/shuffled pattern occurrence in M1 and DLS.

Another approach to normalize the firing rate differences between M1 and DLS is to randomly select spikes to match firing rates in each M1 and DLS unit. Using this method, I examined how the correlational structure in M1 and DLS impacted how often specific neural patterns are repeated without the influence of firing rate. Comparing the pattern occurrence between the real and shuffled networks, the real DLS network repeats specific patterns more often than the real M1 network, compared to each network’s shuffle counterparts, despite the same firing rate (Figure 4.5). I also observe higher pairwise correlation values in DLS vs. M1 with normalized firing rates

(Figure 4.6). Prior work has demonstrated how pairwise correlations have significant impact on network dynamics⁴, suggesting this difference may account for the stability difference we observe.




Figure 4.5. Comparison of M1 and DLS neural stability vs. flexibility for sub-selected spike networks. a. Scatterplot of pattern occurrences in sub-selected spike real and shuffled networks for M1 (left) and DLS (right). **b.** Comparison of distributions between ratio of real/shuffled pattern occurrence in M1 and DLS for sub-selected spike networks.




Figure 4.6. Pairwise correlation values of M1 and DLS networks. Comparison of pairwise correlation values for M1 and DLS networks with all spikes (left) and sub-selected spikes. (right)

Together, these analyses provide evidence for greater neural flexibility in motor cortex compared to the striatum. While further work is required to provide direct evidence for a gradient in neural flexibility vs. stability extending to the brainstem, I believe there are several pieces of evidence consistent with such a gradient. First, there is evidence that distal and proximal effectors are

differentially controlled by cortical and subcortical brain regions, respectively. For example, there is considerable evidence that cortex is uniquely required for dexterity⁵⁻⁸ while, on the other extreme, it is known that the brainstem controls movements of proximal body parts such as locomotion^{9,10}. Between these extremes, as I presented in Chapter 2, there also exists evidence for a difference in the control of proximal vs. distal effectors between cortex and the basal ganglia¹¹. Second, there is evidence that lesioning cortex prevents an animal from learning a sequence of non-dexterous movements. However, once learned, a complex non-dexterous movement sequence could be performed without disruption after complete motor cortical lesion¹². One interpretation of this work is that the exploration and flexibility involved with early learning require cortex, but the stable production of learned behaviors can be controlled subcortically in the basal ganglia and brainstem. Together these two pieces of evidence are consistent with a cortical-to-subcortical gradient in neural flexibility vs. stability that extends from the cortex to the brainstem.

Learning is mediated through plasticity in the corticostriatal network

The second claim is that *learning a new behavior involves transitioning from cortical to subcortical control of a behavior and is mediated through corticostriatal plasticity. (Figure 4.7).*

There is considerable evidence that corticostriatal plasticity is critical for learning. It has been shown that NMDA-dependent plasticity in the inputs to the striatum is required to develop a consistent behavior¹³⁻¹⁵. One obvious consequence of such plasticity is that motor cortical inputs to the striatum are potentiated and drive greater striatal activity. Consistent with this, across several learning paradigms, coordinated activity emerges across motor cortex and the dorsal striatum as a consistent behavior develops^{11,15,16}. However, there lacks an explanation for why corticostriatal plasticity leads to more consistent, learned behaviors.




Figure 4.7. Proposed cortical-to-subcortical shift in movement control with learning.

Here, I propose that corticostriatal plasticity drives learning by transitioning an initially cortically-driven novel behavior to a subcortically-driven learned behavior. To expand on this idea, I will differentiate between two different aspects of motor control: *action selection* and the *control of the kinematic aspects of movement*. I propose that, specifically, a cortical-to-subcortical shift in the kinematic control of movement allows variable kinematics to become more consistent with learning. In this model, early in learning, flexible neural activity in motor cortex, along with connected frontal and sensory cortices, encodes contextual information required for selecting actions and contains movement-related activity responsible for driving exploratory movement kinematics. With training and associated corticostriatal plasticity, the role of motor cortex shifts to providing a context-specific excitatory input to the striatum that serves to select actions, while the consistent patterns of activity driven in the striatum are responsible for stable and reliable kinematics. I will provide evidence for two claims supporting this idea below: (1) novel behaviors performed early in learning are driven by motor cortex and (2) consistent kinematics associated with learned behaviors are largely driven by the basal ganglia.

Consistent with the claim that *novel behaviors performed early in learning are driven by motor cortex* is evidence that, during early learning, motor cortical activity is highly modulated by movement¹⁷⁻¹⁹. It is possible that, for novel behaviors, motor cortex could be involved in both the selection of actions and the control of kinematic aspects of movement, as cortical activity contains motor information about movement kinematics, as well as sensory information about context relevant for action selection^{20,21}. However, despite evidence for this activity in cortex, it is a challenge to determine whether such movement-related modulation is truly “driving” movement. For example, it has been demonstrated that while motor cortical lesions prevent motor learning, these lesions do not prevent the execution of a variable behavior that fails to develop into a consistent learned behavior²². One interpretation of this work is that while behaviors, in general, can be executed without motor cortex, for a variable behavior to develop into a consistent learned behavior, the variable behavior must be driven by motor cortical activity during early learning. This requirement may exist because such activity in motor cortex is a prerequisite for plasticity in striatal-projecting motor cortical neurons that is known to be required for learning. For example, one possibility is that activity in striatal-projecting motor cortical neurons that drives movement in early learning primes these neurons for corticostriatal plasticity that occurs during post-training sleep, as suggested in Chapter 3. Notably, there are pyramidal tract (PT) motor cortical neurons that project to the brainstem/spinal cord with axon collaterals to the striatum²³, suggesting that the same motor cortical neurons could drive behavior during early learning, through projections to brainstem/spinal cord, and undergo corticostriatal plasticity involved in learning.

The second claim, that *consistent kinematics associated with learned behaviors are largely driven by the basal ganglia* is supported by the demonstration that, while motor cortex is required for

learning a sequence of non-dexterous movements, once learned, a kinematically complex non-dexterous behavior can be performed even after motor cortical lesion¹². Furthermore, follow up work has shown that lesioning the striatum disrupts this learned behavior, implicating the basal ganglia as the driver of such consistent kinematics²⁴. Notably, however, it has also been shown that a learned reaching movement can be interrupted though temporally precise inhibition of motor cortical activity²⁵. One interpretation of these studies is that the role of motor cortex transitions with learning from the “driver” of variable behavior to a provider of context-specific excitatory input to the striatum, where stable patterns of activity drive the kinematics associated with the learned behavior. One possibility is that this pattern of context-specific input serves the role of action selection. This change in the role of motor cortex suggests that alternative sources of subcortical input to the basal ganglia may similarly drive action selection. In fact, it has also been demonstrated that striatal-projecting thalamic neurons may play an important role in driving learned behavior²⁶. It is possible that, as training progresses, there is plasticity in the projections from the thalamus to the striatum, and such projections provide an excitatory input to the striatum that replaces the drive from motor cortex with sufficient training. The latter possibility is consistent with evidence that, with long-term learning, motor cortex “disengages” from learned behavior²⁷. Intriguingly sleep spindles originate in the thalamus²⁸, which is reciprocally interconnected to both the motor cortex and basal ganglia, suggesting that sleep may play a critical role in motor network plasticity as demonstrated between M1 and DLS in Chapter 3.

An important consideration for this model is the considerable evidence that dexterous movements require cortex, even once well learned^{7,8,11}. One possibility, consistent with the proposed model, is that, with learning, basal ganglia feedback to the motor cortex, relayed through thalamus,

becomes an important driver of cortical activity underlying dexterous movement. Different cortical populations, e.g., pyramidal tract vs. intratelencephalic²³, may underlie the multiple roles of cortex in (1) providing excitatory input to the striatum and (2) receiving feedback from the basal ganglia to drive dexterous movements. In this model, therefore, we can think of basal ganglia activity as the “driver” of both learned dexterous and non-dexterous movements, through its influence on motor cortex and brainstem motor regions, respectively. A benefit of this scheme is that dexterous and non-dexterous movements can be precisely coordinated by the basal ganglia in the case of learned movements that contain both types of movements, such as reaching and grasping actions. Consistent with the idea that learned cortical activity patterns require subcortical input from the basal ganglia is evidence that corticostriatal plasticity is required to learn an entirely cortically controlled brain-machine interface task¹⁴. Another important consideration, however, is whether skilled dexterous movement can be performed independently of basal ganglia feedback when not integrated into a skilled non-dexterous movement. For example, in Chapter 2, I presented evidence that the dexterous portion of a reach-to-grasp task is unaffected by striatal inactivation when the reward pellet is moved closer to the animal, removing the requirement for skilled non-dexterous movements. There is also evidence, from a different study investigating a cortically controlled brain-machine interface task, that striatal inhibition prevents learning of the task but does not impair performance once the task is learned²⁹. Therefore, the control of learned dexterous movements may vary based on the context in which they are performed.

To summarize, I have provided evidence that novel behaviors are driven by motor cortical activity and with learning become driven by striatal activity. I propose this transition is driven by corticostriatal plasticity, as there is considerable evidence that blocking plasticity in the inputs to

the striatum prevents learning¹³⁻¹⁵. It is important to note, however, that alternative models exist. For example, one such model proposes parallel corticostriatal loops: an associative loop between frontal cortex and the dorsomedial striatum and a sensorimotor loop between the motor cortex and the dorsolateral striatum³⁰. It has been proposed that the associative loop controls goal directed movement while the sensorimotor movements controls learned, habitual stimulus-response movements^{30,31}. Therefore, future work, including precise analyses to dissect what aspects of activity are, or are not, related to movement across the corticostriatal network, combined with careful manipulations of neural activity, will be required to understand the compatibility of distinct models and precisely understand how the “drivers” of movement shift with learning.

Summary

This thesis aims to further our understanding of how the brain regulates the transition from initially variable, new behaviors to stable, learned behaviors. In Chapter 2 of this thesis, I presented my work investigating how activity across the corticostriatal network differentially encodes learning of a complex coordinated action. In Chapter 3, I presented my work investigating the critical role that NREM sleep and sleep spindles play in corticostriatal plasticity during learning. In Chapter 4, I have sought to bring together this work, along with other evidence, to present a model for motor network organization and the role for corticostriatal plasticity in learning. Returning to the question presented at the start of this thesis: how does the brain regulate the transition from *new* to *learned* behaviors? I propose that plasticity between motor cortex and the basal ganglia allows an initially cortically driven behavior to be controlled subcortically.

References

1. Kuypers, H. G. J. M. Anatomy of the Descending Pathways. in *Comprehensive Physiology* (2011). doi:10.1002/cphy.cp010213
2. Hikosaka, O. Neural systems for control of voluntary action - A hypothesis. *Advances in Biophysics* (1998). doi:10.1016/S0065-227X(98)80004-X
3. De La Rocha, J., Doiron, B., Shea-Brown, E., Josić, K. & Reyes, A. Correlation between neural spike trains increases with firing rate. *Nature* (2007). doi:10.1038/nature06028
4. Schneidman, E., Berry, M. J., Segev, R. & Bialek, W. Weak pairwise correlations imply strongly correlated network states in a neural population. *Nature* (2006). doi:10.1038/nature04701
5. Lawrence, D. G. & Kuypers, H. G. The functional organization of the motor system in the monkey. I. The effects of bilateral pyramidal lesions. *Brain* **91**, 1–14 (1968).
6. Lawrence, D. G. & Kuypers, H. G. The functional organization of the motor system in the monkey. II. The effects of lesions of the descending brain-stem pathways. *Brain* **91**, 15–36 (1968).
7. Peters, A. J., Liu, H. & Komiyama, T. Learning in the Rodent Motor Cortex. *Annu. Rev. Neurosci.* **40**, 77–97 (2017).
8. Alaverdashvili, M. & Whishaw, I. Q. Motor cortex stroke impairs individual digit movement in skilled reaching by the rat. *Eur. J. Neurosci.* **28**, 311–322 (2008).
9. Shik, M. L. & Orlovsky, G. N. Neurophysiology of locomotor automatism. *Physiological Reviews* **56**, 465–501 (1976).
10. Gatto, G. & Goulding, M. Locomotion Control: Brainstem Circuits Satisfy the Need for

- Speed. *Curr. Biol.* **28**, R256–R259 (2018).
11. Lemke, S. M., Ramanathan, D. S., Guo, L., Won, S. J. & Ganguly, K. Emergent modular neural control drives coordinated motor actions. *Nat. Neurosci.* **22**, 1122–1131 (2019).
 12. Kawai, R. *et al.* Motor Cortex Is Required for Learning but Not for Executing a Motor Skill. *Neuron* **86**, 800–812 (2015).
 13. Dang, M. T. *et al.* Disrupted motor learning and long-term synaptic plasticity in mice lacking NMDAR1 in the striatum. *Proc. Natl. Acad. Sci. U. S. A.* (2006). doi:10.1073/pnas.0601758103
 14. Koralek, A. C., Jin, X., Long II, J. D., Costa, R. M. & Carmena, J. M. Corticostriatal plasticity is necessary for learning intentional neuroprosthetic skills. *Nature* **483**, 331–335 (2012).
 15. Santos, F. J., Oliveira, R. F., Jin, X. & Costa, R. M. Corticostriatal dynamics encode the refinement of specific behavioral variability during skill learning. *Elife* **4**, (2015).
 16. Koralek, A. C., Costa, R. M. & Carmena, J. M. Temporally Precise Cell-Specific Coherence Develops in Corticostriatal Networks during Learning. *Neuron* **79**, 865–872 (2013).
 17. Costa, R. M., Cohen, D. & Nicolelis, M. A. L. Differential Corticostriatal Plasticity during Fast and Slow Motor Skill Learning in Mice. *Curr. Biol.* **14**, 1124–1134 (2004).
 18. Kargo, W. J. & Nitz, D. A. Early skill learning is expressed through selection and tuning of cortically represented muscle synergies. *J. Neurosci.* **23**, 11255–69 (2003).
 19. Dayan, E. & Cohen, L. G. Neuroplasticity Subservicing Motor Skill Learning. *Neuron* **72**, 443–454 (2011).
 20. Gandolla, M. *et al.* Re-thinking the role of motor cortex: Context-sensitive motor outputs? *Neuroimage* (2014). doi:10.1016/j.neuroimage.2014.01.011

21. Hatsopoulos, N. G. & Suminski, A. J. Sensing with the motor cortex. *Neuron* (2011). doi:10.1016/j.neuron.2011.10.020
22. Kawai, R. *et al.* Motor Cortex Is Required for Learning but Not for Executing a Motor Skill. *Neuron* **86**, 800–812 (2015).
23. Dudman, J. T. & Gerfen, C. R. The Basal Ganglia. *Rat Nerv. Syst.* 391–440 (2015). doi:10.1016/B978-0-12-374245-2.00017-6
24. Dhawale, A. K., Wolff, S. B. E., Ko, R. & Ölveczky, B. P. The basal ganglia can control learned motor sequences independently of motor cortex. *bioRxiv* 827261 (2019). doi:10.1101/827261
25. Guo, J.-Z. *et al.* Cortex commands the performance of skilled movement. *Elife* **4**, (2015).
26. Wolff, S. B. E., Ko, R. & Ölveczky, B. P. Distinct roles for motor cortical and thalamic inputs to striatum during motor learning and execution. *bioRxiv* 825810 (2019). doi:10.1101/825810
27. Hwang, E. J. *et al.* Disengagement of motor cortex from movement control during long-term learning. *Sci. Adv.* (2019). doi:10.1126/sciadv.aay0001
28. Steriade, M., McCormick, D. A. & Sejnowski, T. J. Thalamocortical oscillations in the sleeping and aroused brain. *Science* (80-.). (1993). doi:10.1126/science.8235588
29. Neely, R. M., Koralek, A. C., Athalye, V. R., Costa, R. M. & Carmena, J. M. Volitional Modulation of Primary Visual Cortex Activity Requires the Basal Ganglia. *Neuron* **0**, (2018).
30. Kupferschmidt, D. A., Juczewski, K., Cui, G., Johnson, K. A. & Lovinger, D. M. Parallel, but Dissociable, Processing in Discrete Corticostriatal Inputs Encodes Skill Learning. *Neuron* **96**, 476-489.e5 (2017).

31. Graybiel, A. M. Habits, Rituals, and the Evaluative Brain. *Annu. Rev. Neurosci.* **31**, 359–387 (2008).

Publishing Agreement

It is the policy of the University to encourage open access and broad distribution of all theses, dissertations, and manuscripts. The Graduate Division will facilitate the distribution of UCSF theses, dissertations, and manuscripts to the UCSF Library for open access and distribution. UCSF will make such theses, dissertations, and manuscripts accessible to the public and will take reasonable steps to preserve these works in perpetuity.

I hereby grant the non-exclusive, perpetual right to The Regents of the University of California to reproduce, publicly display, distribute, preserve, and publish copies of my thesis, dissertation, or manuscript in any form or media, now existing or later derived, including access online for teaching, research, and public service purposes.

DocuSigned by:

Stefan Lemke

722F3F557C784C4...

Author Signature

12/16/2020

Date