

UC Berkeley

UC Berkeley Electronic Theses and Dissertations

Title

Large-Scale Neuronal Network Changes Underlying Neuroprosthetic Learning

Permalink

<https://escholarship.org/uc/item/4cr5v24j>

Author

Koralek, Aaron Christopher

Publication Date

2014

Peer reviewed|Thesis/dissertation

Large-Scale Neuronal Network Changes Underlying Neuroprosthetic Learning

By

Aaron Christopher Koralek

A dissertation submitted in partial satisfaction of the

requirements for the degree of

Doctor of Philosophy

in

Neuroscience

and the Designated Emphasis

in

Computational Science and Engineering

in the

Graduate Division

of the

University of California, Berkeley

Committee in Charge:

Professor Jose M. Carmena, Chair

Professor Yang Dan

Professor Shaowen Bao

Professor Michel M. Maharbiz

Spring 2014

© 2014 Copyright, Aaron C. Koralek
All Rights Reserved

Abstract

Large-Scale Neuronal Network Changes Underlying Neuroprosthetic Learning

By

Aaron Christopher Koralek

Doctor of Philosophy in Neuroscience

Designated Emphasis in Computational Science and Engineering

University of California, Berkeley

Professor Jose M. Carmena, Chair

Research on brain-machine interface (BMI) systems has flourished in recent years, with motor BMIs showing great promise as a therapeutic option for patients suffering from limb loss or immobility. In addition to this direct clinical application, BMI tasks also serve as a powerful research tool, in that they enable the researcher to directly define which cells are relevant for behavioral output and the ways in which activity in these cells affects the external world. Neuroprosthetic tasks also serve as a completely novel motor-like learning paradigm for subjects, as they invoke the motor system but do not involve natural body movements or muscle activity. Intriguingly, a large body of work has nevertheless suggested striking similarities between natural motor learning and neuroprosthetic learning, implying that the two forms of learning may share common neural mechanisms.

We developed a novel rodent paradigm to study neuroprosthetic learning in which rodents controlled a one-dimensional auditory cursor by modulating activity in primary motor cortex (M1) in order to hit one of two targets. We first use this paradigm to explore corticostriatal representations of neuroprosthetic skills, the ways in which these representations change over the course of learning, and the necessity of plasticity in the corticostriatal network for abstract learning. We then investigate fine-scale temporal coordination between M1 and the dorsal striatum over the course of neuroprosthetic learning, demonstrating the development of coherent interactions between M1 spikes and the striatal local field potential with high temporal precision. Importantly, we found these interactions to be specifically present in output-relevant neurons, despite close proximity to other neuronal populations. Finally, we modified this behavioral paradigm for use in conjunction with two-photon calcium imaging in head-fixed mice to examine the fine-scale spatial characteristics of network adaptations during neuroprosthetic learning. We demonstrate the development of coordinated network activity and the sparsening of task-relevant modulations over the course of learning. This novel imaging-based BMI paradigm also allows for a number of new techniques to now be

applied to research on neuroprosthetic learning. Together, these data suggest striking similarities between BMI and natural motor learning, demonstrating an important computational role for corticostriatal plasticity, neuronal coherence, and sparsening of cortical representations over the course of neuroprosthetic learning.

Acknowledgements

When I first started grad school, my interests were spread throughout neuroscience, with past experiences ranging from protein signaling pathways to human neuroimaging. As I began my rotations, I was fairly sure I would eventually return to neuroimaging, but I was excited by the opportunity to learn techniques from vastly different sectors of the field. However, upon working in Jose Carmena's lab, I immediately knew that systems level research was the right area for me. It was exciting to finally be able to directly measure neurons firing action potentials, and to then use those recordings for incredibly intriguing and clinically important research questions. I was finally able to get a much clearer view of neural activity in a field that had immediately relevant implications. I am incredibly grateful for the twists and turns of history that led me to my current area of research.

Above all else, I want to thank my advisor, Jose. From frenetic impromptu meetings to emails covered with exclamation marks, Jose always kept the momentum going and the enthusiasm high. His passion for science is contagious and has made the whole process of research not only rewarding, but also exciting and fun. In addition, Jose has given us all great freedom in pursuing our research questions and enough space to explore and learn what is and is not effective. It has been really wonderful, and I feel incredibly lucky to have had such an enthusiastic, ambitious, and fun advisor. Thank you, Jose. Next round is on me in Lisbon!

In addition, I want to thank Rui Costa for being somewhat of a co-advisor for much of the process. All of the studies presented here were done in collaboration with Rui, and he has been incredibly helpful throughout the entire process. Like Jose, Rui has also kept enthusiasm high and I'm not only grateful for his help during grad school, but also very excited to work with him soon as a postdoc and continue this great line of research.

I would also like to thank the other faculty members who have helped me greatly along the way. To start, I would like to thank my committee: Yang Dan, Shaowen Bao, and Michel Maharbiz. Our annual meetings and all of the input and support you have given me has been very important in keeping my thesis work on target. In addition, I would like to thank Dan Feldman for being a wonderful collaborator/mentor for the two-photon project. Dan was very open to the proposal of a collaboration that was fairly unrelated to his lab's research questions, and he has been incredibly helpful and supportive throughout the entire process.

I have also been fortunate enough to work with wonderful lab mates, colleagues, and collaborators. John and Subbu (my scientific big brothers), thank you for helping me get settled in the lab and for teaching me all of the techniques I know. Amy, I couldn't have made it through grad school without our much needed trips to the Kingfish and daily "no subject" emails about fly mating. Helene, thank you for our crazy last minute camping and kayaking trips (especially the ones that we snuck in between experiments). The whole Carmena lab in general has been very helpful and supportive throughout grad school, and I thank you all for being great coworkers: Kelvin, Ryan C., Simon, Sid,

Suraj, Vivek, Ryan N., Preeya, Ale, and Maryam. In addition, I have been lucky to have a number of fun and fruitful collaborations during grad school, and I'm very thankful for my great collaborators: Rikky, Simone, Phil, Peter, Dan, and Will.

A special thanks to Kelly Clancy for being both collaborator and close friend. I still can't believe we actually did the two-photon project, all starting from that random comment after cortex club. I'll forever cherish our hours in the cold, dark laser room together...and our celebratory trip to Guatemala!

And, of course, I am incredibly grateful to my friends for helping me stay sane through the stressful times. Julia, I couldn't have done it without the baked goods and ruv. Lissy, Adria, and James, our hippie hoedowns made it all bearable. Brian, the last minute trips to Joshua tree and ridiculous Halloween parties kept me smiling the whole time. Dina and BGP, thank you for all the "meet-in-the-middle" sierra camping adventures and heart-to-hearts on our various couches. Liz, thank you so much for all the "special time". And thank you to all my other friends who have made it so joyful along the way and continue to make me smile all day long.

Last but not least, thank you to my parents, Richard and Gayle, and the rest of my family: Jake, Kate, Annie, Kevin, Leif, and Ari. I am so lucky to have you all in my life for support, and I love you all.

Contents

1. Introduction	1
1.1. Brain-machine interface systems	1
1.1.1. BMI as a research tool	3
1.2. Similarities with natural motor learning	4
1.3. Chapter previews	4
2. Corticostriatal plasticity is necessary for learning neuroprosthetic skills	6
2.1. Introduction	6
2.2. Methods	7
2.2.1. Animals	7
2.2.2. Surgery	7
2.2.3. Electrophysiology	8
2.2.4. Behavioral task	8
2.2.5. Action-outcome manipulations	9
2.2.6. Data analysis	9
2.3. Results	11
2.3.1. Development of a rodent BMI task	11
2.3.2. Task performance does not rely on natural movements	12
2.3.3. Probing intentionality	13
2.3.4. Corticostriatal plasticity accompanies BMI learning	15
2.3.5. Corticostriatal plasticity is necessary for BMI learning	17
2.4. Discussion	19
3. Temporal precision and cell-specificity in corticostriatal coherence	20
3.1. Introduction	20
3.2. Methods	22
3.2.1. Surgery	22
3.2.2. Electrophysiology	22
3.2.3. Behavioral task	22
3.2.4. Intracortical microstimulation	24
3.2.5. Data analysis	24
3.3. Results	26
3.3.1. Acquisition of a neuroprosthetic skill	26
3.3.2. Corticostriatal coherence develops during neuroprosthetic learning	27
3.3.3. Coherence is specific to output-relevant neurons and time periods	29
3.3.4. Precise timing of corticostriatal activity	31
3.3.5. Network activity drives the 6-14 Hz LFP oscillation	34
3.4. Discussion	35
4. Fine-scale spatial characteristics of network reorganization	37

4.1. Introduction	37
4.2. Methods	37
4.2.1. Surgery	38
4.2.2. Two-photon imaging	38
4.2.3. Behavioral task	38
4.2.4. Data analysis	39
4.3. Results	42
4.3.1. Volitional modulation of calcium dynamics.....	42
4.3.2. Probing intentionality	42
4.3.3. Neuronal changes during neuroprosthetic learning	43
4.3.4. Network changes during neuroprosthetic learning	44
4.3.5. Fine-scale changes in spatial activity patterns during learning	45
4.4. Discussion	46
5. Conclusions and open questions.....	47
5.1. Summary of contributions.....	47
5.1.1. Additional engineering collaborations	49
5.2. Open questions and future directions	49
5.2.1. The role of the striatum in tasks that do not involve the motor system	49
5.2.2. Frontal and thalamic contributions to corticostriatal signaling	50
5.2.3. Structural plasticity during neuroprosthetic learning	50
5.2.4. Local circuit mechanisms of neuroprosthetic learning	50
5.3. Conclusion	50

Publications related to this work

Journal articles and peer-reviewed conference proceedings

- Clancy, K.B.*, Koralek, A.C.*, Costa, R.M., Feldman, D.E., & Carmena, J.M. (in review). Volitional modulation of optically recorded calcium dynamics for neuroprosthetic control. *Authors contributed equally.
- Koralek, A.C., Costa, R.M., & Carmena, J.M. (2013). Temporally precise cell-specific coherence develops in corticostriatal networks during learning. *Neuron*, 79(5):865-872.
- Biederman, W., Yeager, D., Narevsky, N., Koralek, A., Carmena, J., Alon, E., & Rabaey, J. (2013). A fully-integrated, miniaturized (0.125 mm²) 10.5 μ W wireless neural sensor. *IEEE Journal of Solid State Circuits*, 48(4):960-970.
- Koralek, A.C.*, Jin, X.*, Long, J.D., Costa, R.M., & Carmena, J.M. (2012). Corticostriatal plasticity is necessary for learning intentional neuroprosthetic skills. *Nature*, 483(7389):331-335. *Authors contributed equally.
- So, K., Koralek, A.C., Ganguly, K., Gastpar, M.C., & Carmena, J.M. (2012). Assessing functional connectivity of neural ensembles using directed information. *Journal of Neural Engineering*, 9(2):026004.
- Muller, R., Le, H.-P., Li, W., Ledochowitsch, P., Gambini, S., Bjorninen, T., Koralek, A., Carmena, J.M., Maharbiz, M.M., Alon, E., Rabaey, J.M. (in press). A miniaturized 64-channel, 225 μ W electrocorticographic neural sensor. *Proceedings of the International Solid State Circuits Conference*.
- Ledochowitsch, P., Koralek, A.C., Moses, D., Carmena, J.M., & Maharbiz M.M. (2013). Sub-mm functional decoupling of electrocortical signals through closed-loop BMI learning. *Proceedings of EMBC 2013, the 35th Annual International Conference of the IEEE Engineering in Medicine and Biology Society*, Osaka, Japan.
- Chu, P., Muller, R., Koralek, A., Carmena, J.M., Rabaey, J.M., & Gambini, S. (2013). Equalization for intracortical microstimulation artifact reduction. *Proceedings of EMBC 2013, the 35th Annual International Conference of the IEEE Engineering in Medicine and Biology Society*, Osaka, Japan.
- Koralek, A.C., Long, J.D., Costa, R.M., & Carmena, J.M. (2010). Corticostriatal dynamics during learning and performance of a neuroprosthetic task. *Proceedings of EMBC 2010, the 32nd Annual International Conference of the IEEE Engineering in Medicine and Biology Society*, Buenos Aires, Argentina.

Conference abstracts

- Koralek, A.C. & Carmena, J.M. (2012). Neuroprosthetic learning with de-efferented neuronal ensembles. *Annual meeting of the Society for Neuroscience*, New Orleans, LA.

- Ledochowitsch, P., Koralek, A., Tiefenauer, R.F., Carmena, J.M., & Maharbiz, M.M. (2012). Long-term implantation and lifetime testing of micro-ecog array for BMI. *Annual meeting of the Society for Neuroscience*, New Orleans, LA.
- Koralek, A.C. & Carmena, J.M. (2012). BMI learning results in highly precise cell-specific coherence in corticostriatal networks during learning. *Annual meeting on Computational and Systems Neuroscience*, Salt Lake City, UT.
- Koralek, A.C. (2012). Corticostriatal plasticity is necessary for learning intentional neuroprosthetic skills. *Presentation for the University of Utah Neural Engineering Research Group*.
- Koralek, A.C. & Carmena, J.M. (2011). The development of highly precise temporal dynamics in corticostriatal networks during learning. *Annual meeting of the Society for Neuroscience*, Washington, DC.
- Ganguly, K., Orsborn, A., Koralek, A., Long, J.D., Dangi, S., Schulman, J., Gowda, S., Moorman, H.G., & Carmena, J.M. (2011). Learning mechanisms underlying neuroprosthetic control. *Proceedings of EMBC 2011, the 33rd Annual International Conference of the IEEE Engineering in Medicine and Biology Society*, Boston, MA, USA.
- Koralek, A.C., Long, J.D., Costa, R.M., & Carmena, J.M. (2010). Corticostriatal dynamics during learning and performance of a neuroprosthetic task. *EMBC 2010, the 32nd Annual International Conference of the IEEE Engineering in Medicine and Biology Society*, Buenos Aires, Argentina.
- Koralek, A., Jin, X., Long, J.D., Costa, R.M., & Carmena, J.M. (2010). The role of corticostriatal dynamics in learning to operate a brain-machine interface. *Annual meeting of the Society for Neuroscience*, San Diego, CA.

Chapter 1:

Introduction

In recent years, it has become possible to directly connect machines with the brain. Such brain-machine interfaces (BMIs) manifest in a number of forms, with many BMIs aimed at restoring lost motor function and others aimed at restoring lost sensory function. For example, electrical stimulation is used in the case of cochlear implants or deep brain stimulation to restore functionally relevant activity in damaged neural pathways (Chopra et al., 2013; van Schoonhoven et al., 2013). In the case of motor neuroprostheses, recorded neural activity is entered into a decoding algorithm and used to directly control disembodied actuators in real-time, helping to restore motor function in patients suffering from limb loss or immobility (Fetz, 2007; Serruya et al., 2002; Taylor et al., 2002; Carmena et al., 2003). Recent work has demonstrated the utility of neuroprostheses for restoring movement in paralyzed human patients (Hochberg et al., 2006; Colinger et al., 2012), providing exciting prospects for clinical applications of this technology in humans in the future.

Importantly, a large body of work has recently demonstrated that BMI is not simply based on replacing dysfunctional neural circuits with an exact artificial replica, but rather involves a profound learning and adaptation process with striking similarities to natural learning (Green & Kalaska, 2011; Ganguly & Carmena, 2009). In addition to important clinical implications, such a view of BMI also suggests that neuroprosthetic tasks can be used as a research tool to study general neural mechanisms underlying learning from a unique perspective (Ganguly et al., 2011).

This dissertation focuses on BMIs aimed at restoring motor function in patients suffering from limb loss or immobility, and particularly on large-scale neuronal network changes that underlie the learning of neuroprosthetic skills. We begin this chapter with a review of motor BMIs, the ways in which such systems can be used as a unique research tool, and the similarities of BMI systems with the natural motor system. We then provide an overview of the chapters to come.

1.1 Motor brain-machine interface systems

Even when restricted to the motor realm, BMI systems span a wide range of manifestations. In general, most motor BMIs share a common closed-loop architecture (Figure 1.1). Recorded neural activity is entered into an online decoding algorithm that transforms that activity directly into movement of a disembodied actuator, which can take a number of forms. Sensory feedback is then supplied to subjects to facilitate learning and improve performance with the actuator. Apart from this general control loop, however, motor BMIs vary greatly.

One way in which motor BMIs differ is the choice of the control signal. Many potential signals exist for use in BMI, with an initial distinction between invasive and noninvasive methods. Noninvasive BMIs have been demonstrated using electroencephalography (Pascual et al., 2012) and even functional magnetic resonance imaging (Shen et al., 2014). However, noninvasive recording methods offer poor spatial and temporal resolution relative to invasive recording techniques, and there is therefore a limit on the complexity of behavioral performance that these control signals can support. Invasive methods, on the other hand, offer superior spatiotemporal resolution, but require surgery before they can be used. Invasive BMIs have been demonstrated using electrocorticography (Ganguly et al., 2009; Ledochowitsch et al., 2013), but the vast majority of BMI work has been performed with chronically-implanted arrays of penetrating electrodes (Serruya et al., 2002; Taylor et al., 2002; Carmena et al., 2003; Ganguly & Carmena, 2009). Penetrating electrodes can record single-unit activity, multi-unit activity, or local field potentials (LFPs), all of which can potentially be used for neuroprosthetic control (Carmena et al., 2003; Orsborn et al., 2012; So et al., 2014). All of these methods have been demonstrated to provide some level of skillful control when used in BMI, but array recordings continue to be the most commonly used technique. In addition, recordings can theoretically be taken from any region of the brain, although they are most often performed in primary motor cortex (M1).

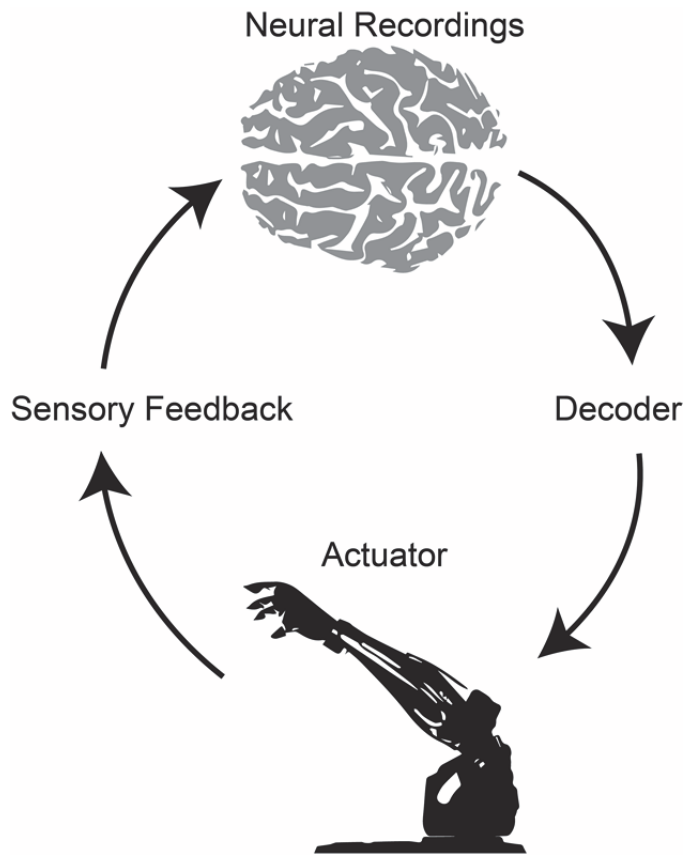


Figure 1.1 Schematic demonstrating the general architecture for BMI tasks. Neural activity is recorded and entered into a decoding algorithm that directly translates that activity into movement of an actuator. Sensory feedback is supplied to subjects to facilitate learning. A wide range of recording methods and neural signals can be used as the control signal.

Motor BMIs also differ in the choice of decoding algorithm. A major distinction here is between biomimetic and non-biomimetic decoders (Koyama et al., 2010; Ganguly &

Carmena, 2010; Cunningham et al., 2011). Biomimetic decoders attempt to mimic the brain's normal motor coding strategy. Subjects are therefore often asked to perform a natural motor task with their body while the researchers record neural activity, and the weights of the decoder are then trained based on these natural movements. Non-biomimetic decoders, on the other hand, do not attempt to have any relation to natural movements and instead require subjects to learn arbitrary mappings between neural activity and behavioral output. While biomimetic decoders intuitively seem far superior, our current lack of understanding of the brain's natural coding strategy makes them fairly difficult to implement effectively and they therefore have not been demonstrated to currently offer superior control relative to non-biomimetic decoders (Orsborn et al., 2012).

Finally, motor BMIs often differ in the choice of external actuator. Researchers have used neural activity to control the motion of virtual objects (Serruya et al., 2002; Taylor et al., 2002; Carmena et al., 2003; Dangi et al., 2013), auditory tones (Koralek et al., 2013), wheelchairs (Millan et al., 2009), robotic limbs (Hochberg et al., 2006; Collinger et al., 2012), and even the user's own body with closed-loop muscular stimulation (Ethier et al., 2012). In all cases, modulations in neural activity are directly translated into a consistent movement of the external device, but the specifics of that external device are determined by the specific research goal.

1.1.1 BMI as a research tool

In addition to the direct neurological applications discussed above, BMI has proven to be a powerful tool for investigating general neuronal mechanisms of learning (Green & Kalaska, 2011; Ganguly & Carmena, 2009; Koralek et al., 2012). During BMI learning, the researcher can directly define which cells are entered into the decoding algorithm, and therefore which cells are relevant for behavioral output. Thus, the task-relevant neuronal population is clearly defined. In the case of natural motor tasks, vast networks of neurons play a direct role in performance, and the entire network cannot be identified, nor can the entire network be recorded with current neural recording technologies. Importantly, the definition of the relationship between neural activity and behavioral output in BMI tasks often results in neighboring neurons displaying vastly different functional roles in the network, thereby enabling researchers to determine neural changes that are strictly related to task learning and performance. In the natural system, neighboring neurons often have similar functional roles, making network effects far more difficult to disentangle.

In addition, through even minor changes in the decoder weights, the researcher can define the precise ways in which neural activity in this output-relevant neuronal population affects behavioral output. This allows researchers to investigate how neuronal networks adapt in response to changes in their functional properties in ways that would be completely impossible with natural motor tasks.

BMI work to date has used such manipulations to demonstrate striking network reorganization while animals adapt to perturbations in the decoder weights (Jarosiewicz et al., 2008). Past work has also used neuroprosthetic tasks to demonstrate highly-specific functional changes in output-relevant neuronal populations (Ganguly et al., 2011). Together, this body of work has capitalized on neuroprosthetic tasks to elucidate network mechanisms that underlie behavioral improvement over the course of skill learning.

1.2 Similarities with natural motor learning

Intriguingly, a growing body of work has demonstrated striking similarities between natural motor learning and neuroprosthetic learning (Green & Kalaska, 2011). Pivotal work in primates demonstrated clear functional reorganization of spike patterns in M1 during the early stages of BMI learning, followed by stabilization of these firing patterns as animals become skilled at neuroprosthetic control (Ganguly & Carmena, 2009). In this work, the decoder was calibrated based on natural arm movements, and therefore the a priori optimal strategy for task performance would be to simply recapitulate the neural patterns that occurred during natural movement. However, this is not what the researchers observed. Instead, the brain appears to adapt to the dynamics of the task that are imposed by the decoder and produce novel firing patterns, just as the brain adapts to novel dynamics in natural motor tasks. As the animal becomes skilled at the task, neuronal firing patterns then stabilize and become reproducible from trial to trial for the remaining days of training. This is analogous to work on natural motor learning showing network plasticity in the early stages of skill acquisition followed by stabilization in later learning (Xu et al., 2009), as well as work showing neuronal adaptations following output perturbations with natural limb movements (Gandolfo et al., 2000). Together, these studies suggest that BMI may be seen as an extension of natural skill learning and may rely on similar neural mechanisms for storage and retrieval of learned skills.

1.3 Chapter previews

In Chapter 2, we present a novel rodent model of neuroprosthetic control, variants of which will be used throughout this dissertation. We probe the learned neuroprosthetic skills for intentionality, explore representations of neuroprosthetic skills in the corticostriatal system of rats, and use the task with mutant mice that lack long-term potentiation specifically in the dorsal striatum to determine the necessity of corticostriatal plasticity for neuroprosthetic skill acquisition and performance.

In Chapter 3, we explore the development of coordinated activity between M1 and the dorsal striatum during neuroprosthetic learning. We quantify this coordination using spike-field coherence, and further demonstrate cell-specificity and highly precise timing of these coherent interactions.

In Chapter 4, we modify the behavioral paradigm for use with two-photon calcium imaging in order to explore the fine-scale spatial characteristics of network

reorganization across the learning process. We demonstrate a sparsening of the cortical representation as the animals become more skilled and efficient at performing the task.

Chapter 2:

Corticostriatal plasticity is necessary for learning neuroprosthetic skills

We first investigate similarities between natural and neuroprosthetic learning by exploring neuronal representations of neuroprosthetic skills in basal ganglia circuitry. The dorsal striatum (DS) is the primary input to basal ganglia circuitry and mediates the transfer of information between the neocortex and subcortical loops (Graybiel et al., 1994; Hoover & Strick, 1991). Although the striatum has long been seen as the neural substrate for motor skill learning and performance (Yin et al., 2009; Barnes et al., 2005; Graybiel, 1998; Atallah et al., 2007), it receives input from vast expanses of the neocortex (Stern et al., 1997; Stern et al., 1998) and has therefore been proposed to play a more general role in the flow of information in large-scale neuronal networks, as well as participate in a wide range of cognitive processes (Garcia-Munoz et al., 2010; Grahn et al., 2008; Pennartz et al., 2009; Chudasama et al., 2006). Despite this possibility, it has remained difficult to dissociate broad, higher-level functions from basic motor control. Here we utilized a novel task in rodents designed to precisely examine these cognitive processes irrespective of natural motor behavior. Rodents learned to control the pitch of an auditory cursor to reach one of two targets by modulating activity in M1 in the absence of physical movements. Degradation of the association between action and outcome, as well as sensory-specific satiety tests, demonstrated that these cognitive actions were intentional and goal-directed rather than habitual. Striatal response profiles were altered as rodents learned to perform the task and strong functional interactions developed in corticostriatal networks across learning, suggesting that the DS serves a similar role in abstract learning as in natural motor learning. In addition, genetically disrupting the molecular mechanisms for plasticity in the DS resulted in marked learning deficits. Together, these results suggest that striatal plasticity is necessary for high-level cognitive actions that are dissociated from motor behavior, and further, that the striatum plays a more general role in cortical network dynamics than previously appreciated.

2.1 Introduction

The ability to learn new actions and perfect them with practice allows us to master amazing skills like playing the piano or riding a bicycle. Learning these skills usually implies moving faster, more accurately, and less variably (Brashers-Krug et al., 1996). However, mastering other types of skills, like playing chess or doing math calculations, often does not directly involve changes in movement (VanLehn, 1996). Cortico-basal ganglia circuits have been implicated in the learning, selection and execution of physical skills (Yin et al., 2009; Barnes et al., 2005; Graybiel, 2008; Hikosaka et al., 1999; Brasted & Wise, 2004; Kubota et al., 2009; Kimchi & Laubach, 2009). In particular, plasticity in the motor cortices and the striatum, the major input region of the basal ganglia, has been shown to accompany the learning of physical skills (Yin et al., 2009;

Rioult-Pedotti, 2000). The motor cortex and frontal cortices have also been implicated in the learning of cognitive skills (Georgopoulos et al., 1993; Gandolfo et al., 2000; Fincham & Andersen, 2006; Badre et al., 2010), and in learning to control neuroprosthetic devices in the absence of movement (Ganguly & Carmena, 2009; Ganguly et al., 2011). Some studies suggest that not only cortical areas, but also the striatum, are involved in learning cognitive skills (Beauchamp et al., 2003; Poldrack et al., 1999; Pasupathy & Miller, 2005). However, it is still unclear if the striatum is required for cognitive skill learning, and if corticostriatal circuits undergo plasticity during the learning of cognitive skills as they do during the learning of physical skills. Here, we use a novel behavioral paradigm in conjunction with electrophysiology and genetic manipulations in both rats and mice to investigate the role of corticostriatal circuits and corticostriatal plasticity in the learning of intentional cognitive actions.

2.2 Methods

2.2.1 Animals

All experiments were performed in compliance with the regulations of the Animal Care and Use Committees at the University of California, Berkeley and at the NIAAA, and according to NIH guidelines. Six male Long-Evans rats weighing roughly 250 grams were used for the experiments. In addition, striatal-specific NMDAR1-knockout mice and littermate controls (7 and 8, respectively) were generated by crossing *RGS9-cre* mice with *NMDAR1-loxP* mice, as described previously (Dang et al., 2006). Behavioral experiments were performed on *RGS9-cre+ / NMDAR1-loxP* homozygous mice, and control mice were littermates consisting of *RGS9-cre+*, *RGS9-cre+ / NMDAR1-loxP* heterozygous, and *NMDAR1-loxP* homozygous mice.

2.2.2 Surgery

Rodents were chronically implanted with microwire arrays ipsilaterally in M1 and the DS. In rats, 2 arrays were independently implanted: each array contained 16 tungsten microelectrodes (35 μ m diameter, 250 μ m electrode spacing, 8x2 configuration; Innovative Neurophysiology, Durham, NC). Stereotactic coordinates relative to bregma were used to center the arrays. In rats, these coordinates were: anteroposterior +2.0 mm, mediolateral +2.0 mm, and dorsoventral +1.5 mm for M1; and anteroposterior +0.5 mm, mediolateral +4.0 mm, dorsoventral +4.5 mm for DLS. In mice, a customized array contained 32 tungsten microelectrodes (8 columns x 4 rows configuration, 35 μ m diameter, 150 μ m electrode spacing, 200-1000-200 μ m electrode row spacing, 1.5 mm shorter for medial than lateral two rows; Innovative Neurophysiology, Durham, NC) was implanted unilaterally, with the medial two rows of electrodes targeting M1 and the lateral two rows of electrodes targeting DS; coordinates centered at anteroposterior +0.5 mm and mediolateral +2.0 mm, lowering dorsoventrally from brain surface 1.0-1.1 mm for M1 and 2.2 mm for DS). In all cases, M1 implants were targeted to record from layer 5 pyramidal neurons. Rats were anesthetized with Ketamine (50 mg/kg) and Xylazine (5 mg/kg) with supplemental isoflurane gas as needed. Mice were anesthetized with isoflurane. Craniotomies were sealed with cyanoacrylate and rodents

were allowed to recover for ten days after implantation before behavioral training. Rats were given dexamethasone treatment (0.5 mg/kg) for one week following surgery to minimize tissue damage around the implant (Zhong et al., 2007).

2.2.3 Electrophysiology

Single unit activity and local field potentials were simultaneously recorded with a Multichannel Acquisition Processor (MAP; Plexon Inc., Dallas, TX). Activity was sorted using an online sorting application (Plexon Inc., Dallas, TX) prior to each daily recording session. Only units with a clearly identified waveform and high signal-to-noise ratio were used. Sorting templates were further refined using an offline sorting application (Plexon Inc., Dallas, TX). Behavioral timestamps were sent to the MAP recording system through Matlab (Mathworks, Natick, MA) and synchronized to the neural data for later analyses.

2.2.4 Behavioral task

After rodents recovered from surgery, two ensembles of 2-4 well-isolated M1 units each were chosen for testing based on waveforms, interspike-interval histograms, and refractory periods. Activity from these ensembles was binned in 200 millisecond bins and entered into an online transform algorithm that related neural activity to the pitch of an auditory cursor, and by modulating activity in these ensembles, rodents controlled the pitch of the cursor. The specific transform used was:

$$f(t) = A_1 e^{(\sum_{i \in N_1} r_i(t))} - A_2 e^{(\sum_{j \in N_2} r_j(t))} + B$$

where f is the cursor frequency, $r_i(t)$ is the firing rate for neuron i at time bin t , N_1 and N_2 denote the units in ensembles one and two, respectively, and A_1 , A_2 , and B are coefficients that are set based on a daily baseline recording session of 250 cycles. This transform caused increased activity in the first ensemble to produce increases in the cursor pitch, while increased activity in the other ensemble produced decreases in the cursor pitch. Linear changes in firing rate resulted in exponential changes in cursor frequency, and frequency changes were binned in quarter-octave intervals to match rodent psychophysical discrimination thresholds (Han et al., 2007).

The rodents had to then precisely modulate these neuronal ensembles to move the cursor to one of two target pitches, one which was associated with a 20% (10% in mice) sucrose solution reward and one which was associated with a 45 mg (20 mg in mice) food pellet reward. Rodents were free to choose either reward at any time, although M1 activity levels had to return to baseline levels for a new trial to begin. A trial was marked incorrect if neither of these target states were achieved within 30 seconds of trial initiation.

To control for physical movements, rats also performed the task after being given injections of lidocaine in the whisker pad to locally inactivate sensory and motor nerve endings. Array placement in the vibrissa region of M1 was confirmed by applying trains

of 60 μ A biphasic pulses through the recording array and observing subsequent twitches of the whisker pad. Pulses lasted 200 μ s and were delivered at 350 Hz. 0.5-1 mL of a 2 mg/mL lidocaine solution was then injected into the whisker pad immediately before the behavioral session.

Behavioral sessions took place in a rodent operant box equipped with liquid and pellet dispensers (For rats: Lafayette Instrument Company, Lafayette, IL; For mice: Med Associates, Inc., St. Albans, VT). Recorded neural data was entered in real time to custom routines in Matlab that then translated those activity levels into the appropriate feedback pitch and played the pitch on speakers mounted on 2 sides of the operant box. Frequencies used for auditory feedback ranged from 1-24 kHz in quarter-octave increments. When a target was hit, a Data Acquisition board (National Instruments, Austin, TX) controlled by Matlab triggered the operant box to supply the appropriate reward to rodents.

2.2.5 Action-outcome manipulations

After initial training, the action-outcome contingency was degraded for 2 days. M1 ensemble activity still determined the pitch of the cursor, but reward was now given on a variable time schedule with equal probability of getting a reward after target achievement and no target achievement. Contingency was then reinstated for 2 days; performance had to return to previous levels before further task manipulations were performed.

A sensory-specific satiety test was also performed to manipulate expected reward value – Devaluation test. Animals were given free access to one of the task rewards in a cage different than their home cage for one hour before a choice session, thus reducing the subjective value of that reward. Animals were then placed into the operant training box and allowed to perform the task with free choice over which target to hit at any time during the session.

Finally, an omission test was performed to demonstrate that rodents were able to intentionally inhibit the learned cognitive actions. For this test, animals were no longer rewarded after achieving one of the behavioral targets and the reward associated with that target was now delivered when animals successfully inhibited the action for the duration of the trial (30 seconds). During this manipulation, the other target continued to be rewarded normally. This manipulation was performed separately for both targets to assess the intentional inhibition of the cognitive action.

2.2.6 Data analysis

All analyses were performed in Matlab (Mathworks, Natick, MA) with custom-written programs. Unit data were first binned in 1 millisecond time bins and digitized. As no difference was seen between the two targets in the basic task, data from the two targets were pooled. Peri-stimulus time histograms (PSTHs) were calculated in relation to target achievement and firing rate analyses were performed from 4 seconds before until

4 seconds after this event, as well as from 6 seconds before until 1 second after this event. Firing rates and PSTHs were smoothed with a moving window of 100 milliseconds. Representative days for early and late learning were chosen for each animal from days 2-4 and 8-11, respectively, based on behavioral performance, motivation, and signal quality. For firing rate z-score analyses, firing rates for individual units were binned in 100 millisecond time bins, averaged across events, and smoothed with a moving average. Mean firing rates for each unit were then z-scored and plotted during both early and late learning sessions. Cross-correlograms were also calculated with spiking activity in either M1 or DS being correlated to the occurrence of action potentials in the other region.

For coherence analyses, a multi-taper method was used to compute spectral estimates of spiking activity in both regions (Jarvis & Mitra, 2001; Thomson, 1982; Mitra & Pesaran, 1999). A total of 5 tapers were used and estimates were computed every 50 milliseconds with a window size of 500 milliseconds. Coherence between spiking activity in the two regions was calculated and defined as:

$$C_{xy} = \frac{|R_{xy}|}{\sqrt{|R_{xx}|} \sqrt{|R_{yy}|}}$$

where R_{xx} and R_{yy} are the power spectra and R_{xy} is the cross-spectrum. Power and spike-spike coherence estimates were calculated relative to the delivery of reward and averaged across trials and animals. Mean coherence in the theta band was calculated and defined as 4-8 Hz (4.5-9 Hz in mice).

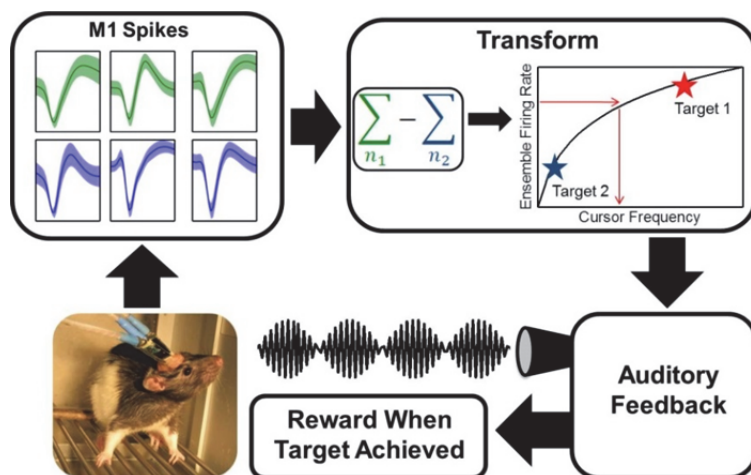


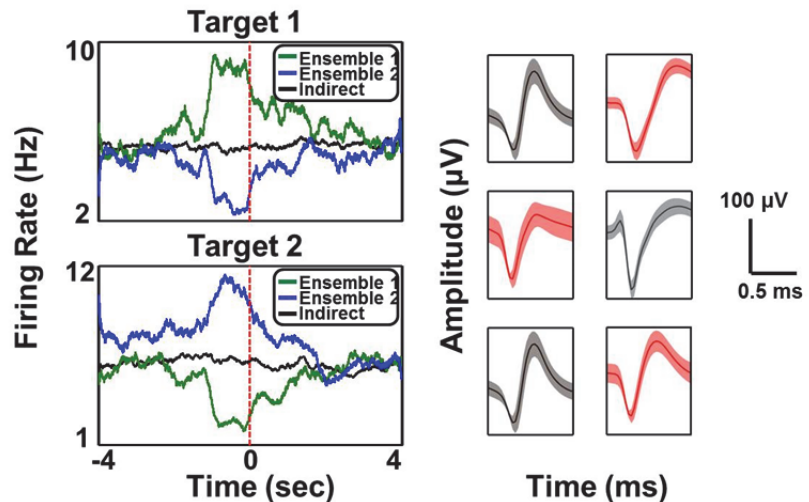
Figure 2.1 Task schematic. M1 unit activity was entered into an online transform algorithm that related ensemble activity to the pitch of an auditory cursor. Two opposing ensembles were chosen, with activity of one ensemble increasing the cursor pitch and activity of the other ensemble decreasing the cursor pitch. Constant auditory feedback about cursor location was supplied to rodents, and distinct rewards were supplied when rodents brought M1 activity into one of two target states.

2.3 Results

2.3.1 Development of a rodent BMI task

We developed a novel operant brain-machine interface task in which rodents were required to modulate activity in M1, rather than do a physical movement, to obtain reward (Figure 2.1). Modulation of M1 ensemble activity resulted in changes in the pitch of an auditory cursor, which provided constant auditory feedback to rodents about task performance. Reward was delivered when rodents precisely modulated M1 activity to move this auditory cursor to one of two target tones, and a trial was marked incorrect if no target had been hit within a set time limit (30 seconds). One of these targets was associated with a reward of sucrose solution, while the other target was associated with a pellet reward (see Methods). Two neural ensembles consisting of 2-4 well-isolated units each were used to control the auditory cursor. The action of these two ensembles opposed each other, such that increased activity in one ensemble produced increases in cursor pitch, while increased activity in the other ensemble caused decreases in cursor pitch. Thus, in order to achieve a high-pitched target, rodents had to increase activity in the first ensemble and decrease activity in the second, while the opposite modulations were necessary to hit a low-pitched target (Figure 2.2). Importantly, the opposing effect of activity in these ensembles ensures that rodents cannot simply increase activity broadly throughout the network, but rather must precisely modulate activity in specific neuronal populations. These firing rate modulations had to be maintained for several time bins (200-600 milliseconds) for a target to be hit. Hence, in this operant task, rodents had to bring the two M1 ensembles into a desired state irrespective of motor output.

Figure 2.2 Mean M1 ensemble firing rates for units in ensemble 1 (green), ensemble 2 (blue), and M1 units not used in the transform (black) in relation to the achievement of target 1 (top) or target 2 (bottom). Representative waveforms recorded from M1 in rats are shown on the right, with shaded regions denoting the standard deviation.



We trained six male Long-Evans rats on the task, and verified that they exhibited marked improvement in the percentage of correct trials over the course of 11 days (Figure 2.3a), following a characteristic learning curve. We could clearly see a phase of rapid improvement followed by a phase of slower learning, representing early (days 2-4)

and late (days 8-11) phases of learning. The percentage of correct trials increased significantly from early to late in learning (Figure 2.3b; $p < 0.001$), demonstrating that rats robustly learned to perform the task. This learning resulted in performance well above chance (Figure 2.3c; $p < 0.001$). Chance levels of target achievement were assessed by collecting daily baseline recordings in which neural activity was entered into the transform algorithm but rodents were not engaged in the task and received no auditory feedback or reward to guide learning and performance, thus measuring how often spontaneous fluctuations of neural activity would lead to target achievement. Analyses of M1 firing rates further showed that rats were producing the desired neuronal ensemble rate modulations during task performance (Figure 2.2).




Figure 2.3 a. Mean percentage of correct responses for all rats across days 1-11 of learning. Shaded regions denote the range of days from which the early and late learning analyses were performed. **b.** The percentage of correct responses for all rats increased significantly from early (light blue) to late (dark blue) in learning. **c.** Performance in late learning was significantly above our assessment of chance levels of target achievement.

2.3.2 Task performance does not rely on natural movements

We next investigated if animals were performing physical movements that would modulate the activity of those particular M1 ensembles. First, we monitored overall rodent movement with a three-axis accelerometer mounted on the recording headstage, which allowed us to measure if the animals produced any gross body or head movement during target achievement. Accelerometer traces exhibited no changes before and during target reaching, but did show prominent deflections after target reaching as the animals retrieved the reward (Figure 2.4a), demonstrating that rodents were not relying on gross motor behavior to perform the task. We also monitored movements of the vibrissae with electromyographic (EMG) recordings of the mystacial pad (electrodes targeted M1 areas controlling vibrissae movement, see Methods), and observed no significant EMG signals before target achievement, although there were clear EMG signals afterwards as animals retrieved and consumed the reward (Figure 2.4b). Importantly, there was no correlation between EMG activity and the spiking of the M1 neurons controlling the auditory cursor - the correlation coefficient for all trials in a behavioral session was 0.092 ± 0.003 (mean \pm s.e.m.), and the distribution of correlation coefficients across a session was not significantly different from zero ($p =$

0.57). This was observed across all training days, including during early learning. These data suggest that rats do not rely on physical movements to learn the task, although it is difficult to exclude the possibility that animals use some movement to generate neural activity to drive the auditory cursor during exploratory phases of the task in early learning. Nonetheless, the data shows that animals eventually learn to perform the task in the absence of overt movement. To further demonstrate that rats did not require vibrissae movements to control M1 activity, we injected lidocaine into the whisker pad to locally inactivate sensory and motor nerve endings during a session in late learning (see Methods). There was no significant change in performance during the temporary inactivation (Figure 2.4c; $p > 0.9$), with rats achieving 78.1 ± 2.2 % correct with lidocaine (mean \pm sem) versus 78.8 ± 6.5 % without lidocaine. Taken together, these data indicated that rodents were able to learn to operantly control M1 activity irrespective of any overt movement.

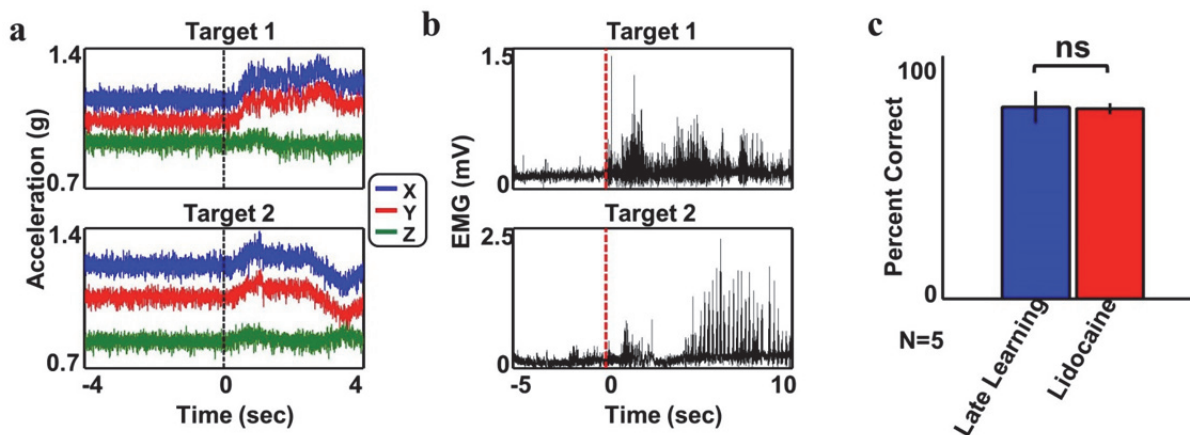


Figure 2.4 **a.** Representative accelerometer traces in relation to time of reward delivery (time zero) show no gross motor behavior leading to target achievement, but clear deflections as animals initiate movement to retrieve reward. **b.** Representative EMG traces time-locked to the time of reward delivery show no muscle activity in the mystacial pad before target achievement, but clear deflections as animals retrieve and consume reward. **c.** Mean performance in all rats when lidocaine was injected into the whisker pad before a behavioral session late in learning (red) compared to performance during a no-lidocaine session (dark blue). Lidocaine injection did not significantly decrease task performance.

2.3.3 Probing intentionality

Operant actions can be goal-directed or intentional if they are performed because of their consequences, or habitual if they are elicited in a particular situation based on past reinforcement history (Balleine & Dickinson, 1998; Yin et al., 2006; Hilario & Costa, 2008). Goal-directed actions are therefore sensitive to changes in the relation between performing the action and obtaining a reward (contingency) and to changes in the expected value of the reward, while habits are not. We asked if these cognitive actions, i.e. actions performed based on the modulation of specific neural activity but in the absence of physical movement, were performed intentionally because the animal

volitionally controlled M1 activity to get the outcome (goal-directed), or habitually due to reinforcement history. To test this, we first degraded the contingency between executing the action and obtaining the outcome: i.e. the auditory cursor was still under control of M1 ensemble activity, but the probability of obtaining reward was similar irrespective of target achievement, which had no effect on the rate of reward. Following two days of contingency degradation, rats markedly diminished their responding and the percentage of correct trials decreased significantly (Figure 2.5a; $p < 0.001$). When contingency was reinstated, rats resumed responding and the percentage of correct trials returned to plateau levels seen in late learning.

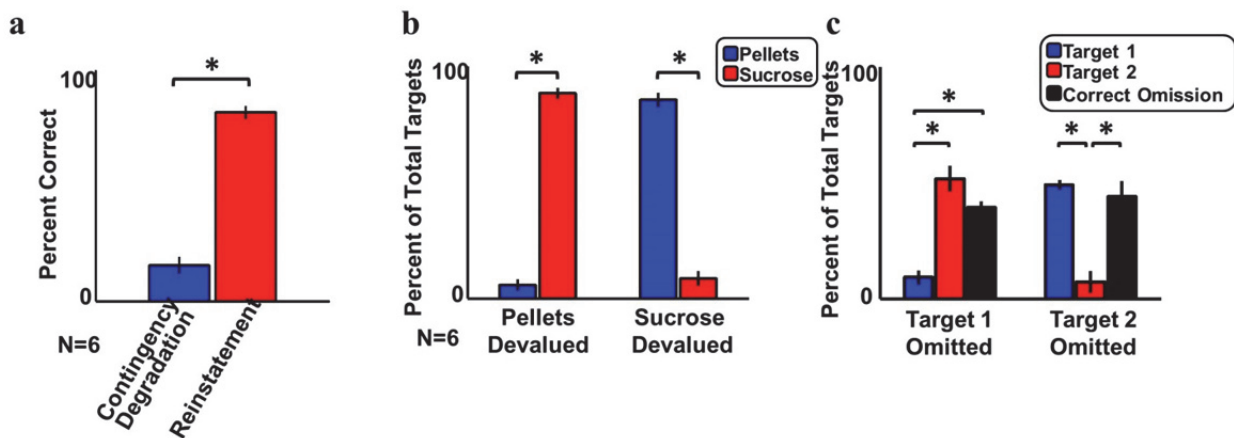


Figure 2.5 a. Significant reduction in response rate when the causal relation between target achievement and reward delivery was degraded (dark blue). When contingency was reinstated, performance rapidly returned to pre-degradation levels (red). **b.** Percentage of total correct trials that were directed at the target associated with pellet reward (blue) or sucrose solution reward (red) during choice sessions where rats had free access to pellets before the session (left; “Pellets Devalued”), or rats had free access to sucrose solution before the session (right; “Sucrose Devalued”). **c.** Percentage of total trials that involved responses toward target 1 (blue), target 2 (red), or response omissions (black) when omission tests were performed for target 1 (left) or target 2 (right).

To further investigate the intentional nature of the task, we performed a test where each of the outcomes was devalued using sensory-specific satiety. Rats were given free access to either sucrose solution or pellets for one hour before the behavioral session, thereby reducing the expected value of that outcome (Hilario et al., 2007). After specific devaluation of each outcome/reward, rats chose the target leading to that reward much less than the target leading to the reward that was not devalued (Figure 2.5b; $p < 0.001$), indicating that they performed the action driven by the expected value of the outcome. Importantly, there were no significant differences in reward preference during normal task performance when neither of the outcomes was devalued ($p > 0.25$). Finally, we asked whether rats were able to intentionally inhibit the reaching of one of the two targets in order to obtain the specific reward associated with that target. To examine this we performed an omission test, where the reward previously associated with reaching a particular target was now only delivered when rats successfully inhibited

reaching that target throughout the duration of the trial. If the target was reached during the 30s of trial duration, then no reward would be delivered and a new trial would be initiated. Importantly, reaching the other target continued to lead to reward as during training. Animals behaved in a goal-directed manner in the omission test for both targets, since they reduced the number of target reaches for the target they had to omit versus the no-omission target, while increasing the number of correctly omitted responses (Figure 2.5c; $p < 0.001$ for both comparisons). Taken together, these data show that the neuroprosthetic actions in our task are sensitive to changes in the causal relation between performing the action and obtaining the reward (contingency degradation and omission test), and to changes in the expected value of the outcome (sensory-specific devaluation), indicating that they are intentional and goal-directed rather than habitual.

2.3.4 Corticostriatal plasticity accompanies BMI learning

We next examined if learning to operantly control M1 activity in the absence of overt movement involves striatal plasticity, akin to what is observed for natural motor learning (Costa et al., 2004; Jin & Costa, 2010; Miyachi et al., 2002). We verified that the improvement in behavioral performance seen across learning was accompanied by a significant increase in firing rates in the DS in late learning compared to early learning (Figure 2.6a; $p < 0.001$). In addition to this general increase in firing rates, we noticed that firing rates of DS neurons exhibited greatest modulation during target reaching compared to baseline control periods (Figure 2.6b, before time zero indicating reward delivery), as observed during natural motor learning (Costa et al., 2004). This modulation was markedly enhanced in late learning (Figure 2.6b, firing rates for individual DS units are z-scored across time in a window surrounding target achievement). Indeed, the percentage of DS units exhibiting target-related firing rate modulation was significantly greater in late learning than early learning (Figure 2.7a; $p < 0.05$), indicating that DS neurons changed their activity during the volitional control of M1 activity, and that this change increased with learning.




Figure 2.6 a. Mean normalized firing rates in the DS increased significantly from early (light blue) to late (dark blue) in learning, mirroring the improvement in behavioral performance. Representative waveforms recorded from the DS are shown on the right, with shaded regions denoting the standard deviation. **b.** Z-scored firing rates for individual DS units in relation to target achievement (time zero). In late learning, the response profiles primarily exhibit modulation before target achievement.

We next investigated if learning, and the observed changes in DS target-related activity, were accompanied by corticostriatal plasticity, i.e. changes in the functional interactions between M1 and DS neurons. We noticed that cross-correlation histograms between the two regions in late learning exhibited pronounced oscillatory spike coupling (Figure 2.7b). To quantify this interaction, we calculated the coherence between spiking activity in the two regions in both early and late learning (Methods). The resulting coherograms exhibited a clear increase in coherence at low frequency bands in late learning relative to early learning (Figure 2.7c), and these frequencies corresponded to the oscillatory frequency seen in the cross-correlation histograms. Furthermore, mean coherence in the theta band (4-8 Hz) was significantly greater in late learning than early learning (Figure 2.7d; $p < 0.001$). Thus, cognitive skill learning is accompanied by dynamic changes in functional interactions between M1 and the DS neurons, suggesting an important role for corticostriatal plasticity in this novel task.

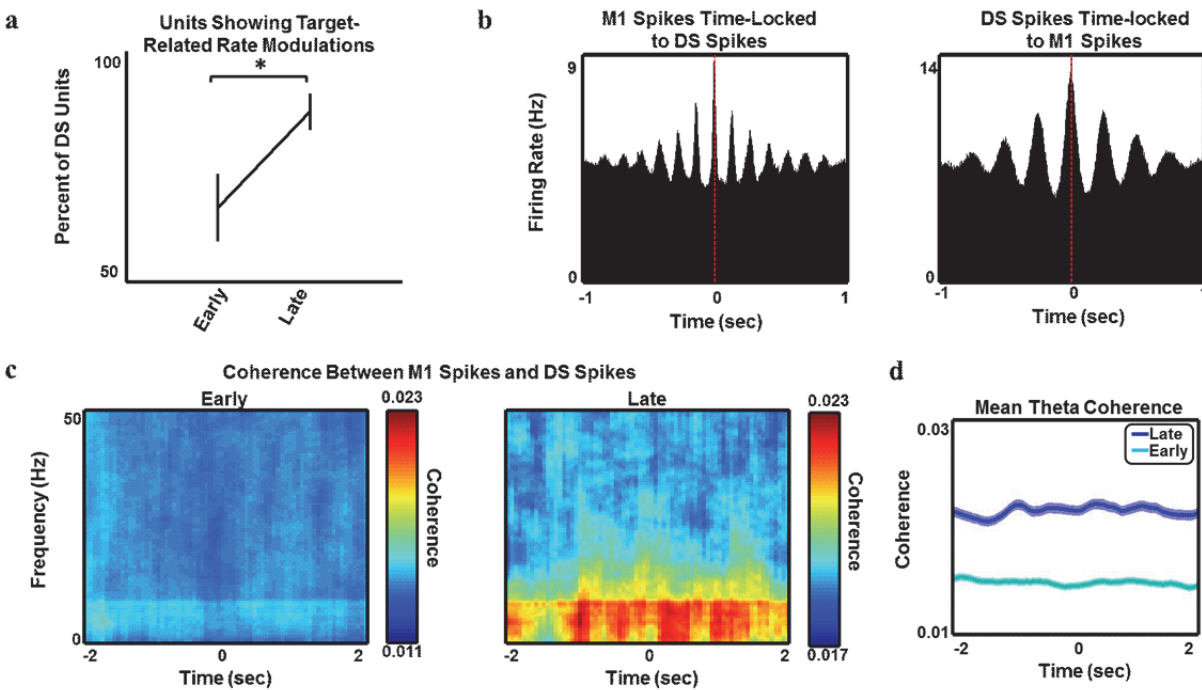


Figure 2.7 a. The percentage of DS units exhibiting target-related firing rate modulation increased significantly with learning. b. Cross-correlation histograms in late learning for M1 spiking activity in relation to DS spikes (left), and DS spiking activity in relation to M1 spikes (right), showing oscillatory coupling between the two regions. c. Coherence between M1 spikes and DS spikes in early (left) and late (right) learning shows a clear increase in low frequency coherence in late learning relative to early learning. d. Significant increase in mean coherence in the theta range in late (dark blue) versus early (light blue) learning. Shaded regions denote the standard error.

2.3.5 Corticostriatal plasticity is necessary for BMI learning

We therefore investigated if corticostriatal plasticity would be necessary for cognitive skill learning. N-Methyl-D-aspartic acid (NMDA) receptors in striatal medium spiny neurons are critical for corticostriatal long-term potentiation (Di Filippo, 2009). We generated mutant mice with a striatal-specific deletion of the NMDAR1 gene using the *Cre/loxP* system (*RGS9-cre+/NMDAR1-loxP^{-/-}*, referred to here as striatal NR1-knockout mice), and compared their learning to that of littermate controls (see Methods). Although control mice showed performance improvement across learning in the absence of physical movement as observed for rats (Figure 2.8a,b; $p < 0.001$), striatal NR1-knockout mice exhibited marked learning deficits on the task, with no significant increase in the percentage of correct trials from early to late in learning (Figure 2.8a; $p = 0.98$). Consistently, DS neurons in littermate controls exhibited a significant increase in firing rate across learning, but striatal NR1-knockout mice did not (Figure 2.8c; main effect of genotype $F_{1, 10} = 32.45$, $p < 0.001$; early vs. late $p < 0.05$ for CT and $p = 0.23$ for KO). Also, in control mice the proportion of DS neurons with significant target-related firing rate modulation increased with learning (Figure 2.8d,e; $p < 0.05$), but this was not observed in knockout mice (Figure 2.8e; $p = 0.28$). Finally, the development of functional corticostriatal interactions during learning was also abolished in striatal NR1-knockout mice, with no significant increase in the coherence between M1 and DS units observed during learning (Figure 2.8f, g, $F_{80, 10} = 0.65$, $p = 0.44$), despite littermate controls showing a clear increase as seen in rats (Figure 2.8f, g, $F_{80, 10} = 4.86$, $p < 0.05$). Taken together, these results demonstrate that the same striking plasticity in corticostriatal networks that was observed in rats during learning also occurs in control mice, but this plasticity is absent in mice lacking functional NMDA receptors in the striatum. These mutant mice do not show improvement with training, therefore indicating that corticostriatal plasticity is necessary for learning to intentionally modulate M1 states to obtain specific outcomes.

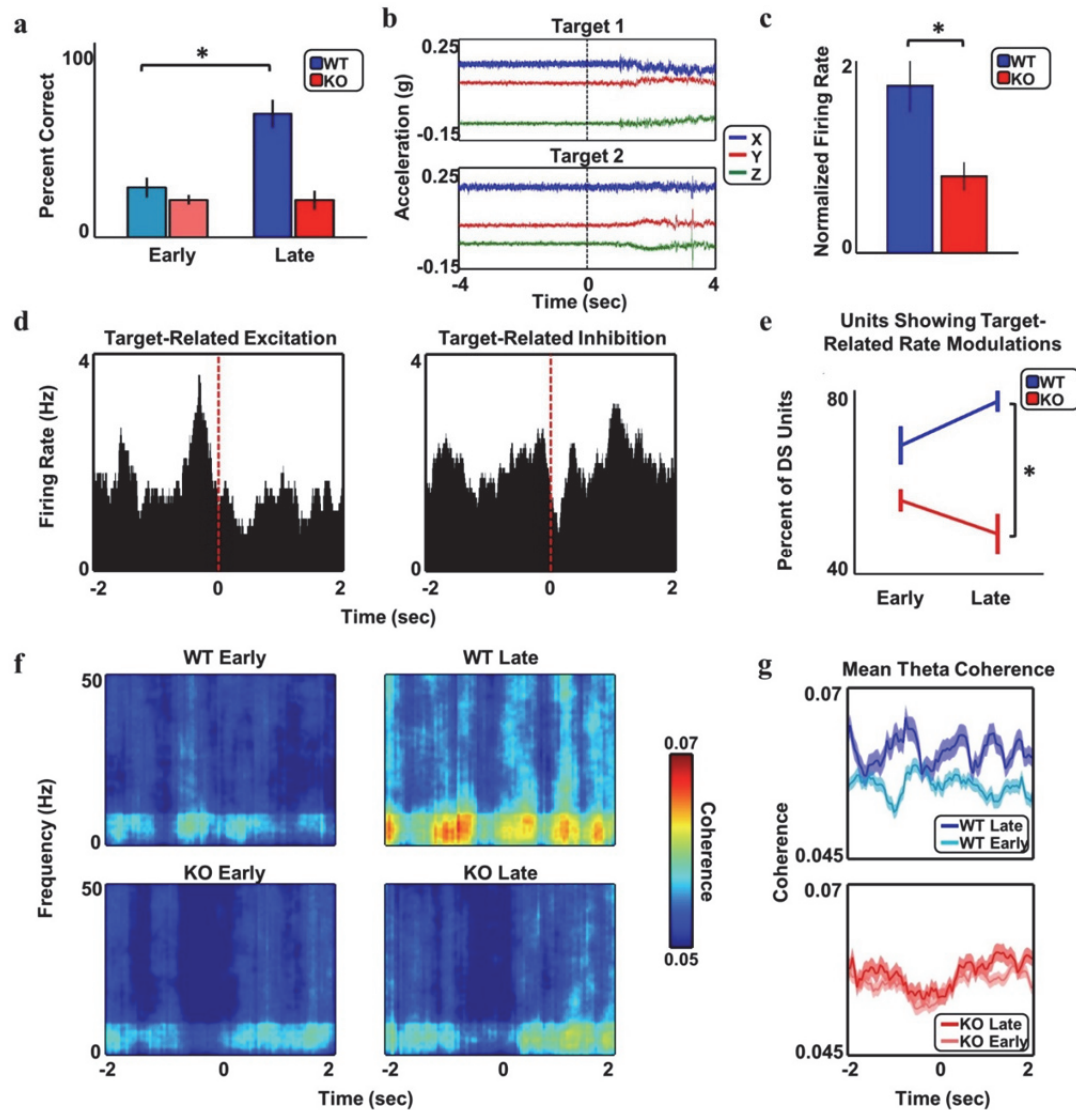


Figure 2.8 a. RGS9L-NR1-KO mice (red) exhibit no significant increase in the percentage of correct trials over the course of learning, despite clear performance improvement in littermate controls (blue). **b.** Accelerometer traces from control mice showing no clear oscillation during task performance, but clear deflections as mice retrieve reward. **c.** There is no significant increase in DS firing rates in knockout mice (red) from early to late in learning, although DS firing rates increase markedly in control mice (blue). **d.** DS units of littermate controls exhibit strong target-related firing rate modulations, including both excitation (left) and inhibition (right). **e.** The percentage of DS units showing significant target-related firing rate modulations increases significantly across learning in control mice (blue), but not in RGS9L-NR1-KO mice (red). **f.** Coherograms showing coherence between M1 spikes and DS spikes in early (left) and late (right) learning for both control mice (top) or knockout mice (bottom). Coherence in low frequency bands increases greatly from early to late learning in control mice, but there is no significant change in coherence in knockout mice. **g.** Mean coherence in the theta range for control (top) and knockout (bottom) mice. In control mice, there is a significant increase in coherence from early (light blue) to late (dark blue) learning, while there is very little change in knockout mice from early (light red) to late (dark red) learning.

2.4 Discussion

In summary, we used a novel task in rodents to demonstrate that corticostriatal networks exhibit profound plasticity during the learning of intentional cognitive skills, and further, that disrupting this plasticity impairs learning. This adds great support to claims that cortico-basal ganglia circuits, and more specifically the striatum, play a role in high-level cognitive processes. We observed that DS neurons strongly modulated their activity in relation to M1 activity even when M1 activity was dissociated from physical movements, suggesting that the striatum is important for learning and selecting cognitive actions that are controlled by cortical output. Ultimately, these data suggest that cortico-basal ganglia circuits are involved in learning thought actions and skills that do not require physical movement, indicating that they may have a broader role in intention and decision-making than previously acknowledged.

Our results also have important implications for the field of brain-machine interfaces. The cognitive actions investigated here form the basis for skillful neuroprosthetic control (Ganguly & Carmena, 2009; Fetz, 2007) and, as we have seen here, these actions recruit elements of the natural motor system outside of M1. Thus, our results suggest that neuroprosthetic movements capitalize on the neural circuitry for motor learning and therefore have great potential to feel naturalistic, generalize well to novel movements and environments, and benefit from our nervous system's highly-developed storage and retrieval mechanisms for skilled behavior.

Chapter 3:

Temporal precision and cell-specificity in corticostriatal coherence

After observing the development of striking oscillatory coupling in the corticostriatal system during neuroprosthetic learning, we next aimed to more precisely dissect these functional interactions, with a focus on the role these interactions might play in natural corticostriatal coding and the specificity of neuronal coherence. It has been postulated that the selective temporal coordination between neurons and the development of functional neuronal assemblies are fundamental for brain function and behavior. Still, there is little evidence that functionally relevant temporal coordination emerges preferentially in neuronal assemblies that directly control behavioral output. Here we investigate coherence between primary motor cortex and the dorsal striatum as rats learn an abstract operant task. Striking coherent interactions developed between these two regions as learning progressed. Interestingly, this coherence was selectively increased in cells relevant for behavioral output compared with other adjacent cells. Furthermore, the phase offset of these coherent interactions aligned closely with estimates of the corticostriatal conduction delay, demonstrating highly precise timing. Spikes from either region were followed by a consistent phase in the other, suggesting that network feedback reinforces the coherent oscillations. Together, these results demonstrate that temporally precise coherence develops during learning specifically between output-relevant neuronal populations and the striatum, and further suggest that correlations in oscillatory activity serve to synchronize large-scale brain networks to produce behavior.

3.1 Introduction

For any given task, the nervous system must coordinate the activity of large ensembles of individual neurons across distant brain regions. Even in seemingly trivial motor tasks, such as holding a cup of coffee, large ensembles of neurons must interact to properly control the musculature and monitor sensory feedback. Although the nervous system is equipped with dense anatomical connectivity to support interactions between cell groups, these interactions must be rapidly and flexibly altered as we move from one behavioral context to the next, and particularly as we learn a new task.

Brain-machine interface (BMI) tasks involve subjects learning to modulate neuronal activity in order to control a disembodied actuator (Fetz, 2007), and therefore provide a completely novel learning environment for subjects. As discussed above, past work has shown that neuroprosthetic skills rely on similar neural substrates as natural motor learning (Green & Kalaska, 2007) and therefore have similar computational requirements for rapid and flexible information transfer. Importantly, BMI tasks offer the unique advantage that the researcher can define which neuronal ensembles are directly

relevant for behavioral output, therefore allowing for an investigation of functional specificity within local neuronal populations.

Recent theories have proposed that alterations in the pattern of large-scale synchronous activity could serve as the substrate for the flexible neuronal associations necessary to coordinate activity in distant brain regions for performance of both natural and neuroprosthetic behaviors (Womelsdorf et al., 2007; Buschman et al., 2012; Canolty et al., 2010). Oscillatory local field potential (LFP) activity reflects rhythmic current flow across cell membranes in local ensembles and is hypothesized to alter the excitability of neuronal groups across different spatiotemporal scales (Buzsaki & Draguhn, 2004; Lakatos et al., 2005; Frohlich & McCormick, 2010). Therefore, precise temporal control in neuronal networks could enhance the efficiency of information transfer in specific populations (Wang et al., 2010; Tiesinga et al., 2001). It could also serve as a mechanism for synaptic gain control (Zeitler et al., 2008) and influence spike-timing-dependent plasticity (Huerta & Lisman, 1993; Harris et al., 2003), as spikes arriving at excitability peaks will have enhanced efficacy relative to poorly-timed spikes. The development of temporally coordinated activity in ensembles of neurons has been implicated in processes as diverse as perception (Rodriguez et al., 1999), expectation (von Stein et al., 2000), decision making (Pesaran et al., 2008), coordination (Dean et al., 2012), memory formation (Pesaran et al., 2002; Siegel et al., 2009; Fujisawa & Buzsaki, 2011), spatial cognition (Colgin et al., 2009), reward processing (van der Meer, 2011), and attentional shifting (Bollimunta et al., 2011; Lakatos et al., 2008; Fries et al., 2008; Steinmetz et al., 2000). In some cases, this synchrony manifests as spiking in one region becoming highly coordinated with LFP activity in a separate region (Pesaran et al., 2008). Importantly, many tasks evoke changes in the temporal pattern of spiking without concomitant changes in firing rate, suggesting that synchrony could serve as an additional information channel in neuronal circuits (Riehle et al., 1997). Alterations in synchrony and LFP dynamics have also been implicated in pathological states such as epilepsy (Bragin et al., 2010) and Parkinson's disease (Costa et al., 2006), highlighting their importance for normal brain functioning.

Despite increasing evidence that changes in synchronous LFP activity are related to changes in behavior during learning (DeCoteau et al., 2007), there is little evidence that temporal coordination during learning emerges selectively between the neurons that are controlling the behavior. For example, although previous work has demonstrated selectivity of corticomuscular coherence across hemispheres (Schoffelen et al., 2007), there is less evidence of selective coherence emerging in cells directly relevant for behavioral output, largely because the differential participation of neighboring neurons in behavior has been difficult to disentangle. In addition, investigating the progression of coherent interactions between distant brain regions across learning in individual animals has only recently become possible due to the development of chronically implantable multielectrode arrays. Corticostriatal networks exhibit plasticity during action learning (Costa et al., 2004; Hikosaka et al., 1999) which involves changes in coherence between distal regions (Koralek et al., 2012), and they therefore serve as an important model system for investigating changing interactions across learning. Here, we examine the dynamics and the specificity of the temporal interactions between distal

nodes of corticostriatal circuits during learning using a brain-machine interface paradigm that permits the definition of output-relevant neurons.

3.2 Methods

3.2.1 Surgery

All experiments were performed in compliance with the regulations of the Animal Care and Use Committee at the University of California, Berkeley. A total of eight male Long-Evans rats weighing roughly 250 grams were used for the experiments. Rats were chronically implanted with microwire arrays in both M1 and the DS. Each array contained 16 tungsten microelectrodes (35 μm diameter, 250 μm electrode spacing, 8x2 configuration; Innovative Neurophysiology, Durham, NC). Stereotactic coordinates relative to bregma were used to center the arrays (anteroposterior 2 mm, mediolateral 2 mm, and dorsoventral 1.5 mm for M1; and anteroposterior 0.5 mm, mediolateral 4 mm, dorsoventral 4.5 mm for DS). M1 implants were targeted for layer 5 neurons based on insertion depth, and this was verified histologically at the end of experiments (Koralek et al., 2012). Rodents were anesthetized with Ketamine (50 mg/kg) and Xylazine (5 mg/kg) with supplemental isoflurane gas as needed. Craniotomies were sealed with cyanoacrylate and rodents were allowed to recover for ten days after implantation before behavioral training. Rats were given dexamethasone treatment (0.5 mg/kg) for one week following surgery to minimize tissue damage around the implant (Zhong & Bellamkonda, 2007). Stimulation experiments were performed in the same animals used for training, with one additional animal that underwent the same surgical procedures but was not trained on the task.

3.2.2 Electrophysiology

Single unit activity and local field potentials were simultaneously recorded with a Multichannel Acquisition Processor (MAP; Plexon Inc., Dallas, TX). Activity was sorted using an online sorting application (Plexon Inc., Dallas, TX) prior to each daily recording session. Only units with a clearly identified waveform and signal-to-noise ratio greater than 2 were used. Sorting templates were further refined using an offline sorting application (Plexon Inc., Dallas, TX). Behavioral timestamps were sent to the MAP recording system through Matlab (Mathworks, Natick, MA) and synchronized to the neural data for later analyses. Recording arrays were grounded to a screw in the occipital bone, and both arrays were also referenced locally using the online program Ref2 (Plexon Inc., Dallas, TX) to eliminate effects of volume conduction. For referencing, an electrode on each array was chosen to be subtracted from all other electrodes on that array. This was done independently for both M1 and DS.

3.2.3 Behavioral task

After rodents recovered from surgery, two ensembles of 2-4 well-isolated M1 units each were chosen for inclusion in the “output” population based on waveforms, interspike-interval histograms, and refractory periods. No other selection criteria were used to

partition the recorded cells into each ensemble. Although these output units were consistently well-isolated, we also ensured that many well-isolated units were included in the indirect population to enable a proper comparison, and we further verified that there was no difference in baseline firing rate between the two populations. The units assigned to the output population remained relatively constant throughout training using the stability of spike waveforms, sorting templates, and interspike intervals across sessions as a guide. When recordings from an output unit were lost, the unit was replaced by a unit on the same electrode or a directly neighboring electrode, thus keeping constant the general cortical volume from which output cells were selected.

Activity from these ensembles of output units was binned in 200 millisecond bins and entered into an online transform algorithm that related neural activity to the pitch of an auditory cursor, and by modulating activity in these ensembles, rodents controlled the pitch of the cursor. The specific transform used was:

$$f(t) = A_1 e^{(\sum_{i \in N_1} r_i(t))} - A_2 e^{(\sum_{j \in N_2} r_j(t))} + B$$

where f is the cursor frequency, $r_i(t)$ is the firing rate for neuron i at time bin t , N_1 and N_2 denote the units in ensembles one and two, respectively, and A_1 , A_2 , and B are coefficients that are set based on a daily baseline recording session of 250 time bins. Ensemble firing rates were smoothed by a moving average of the past 3 time bins. This transform caused increased activity in the first ensemble to produce increases in the cursor pitch, while increased activity in the other ensemble produced decreases in the cursor pitch. Linear changes in firing rate resulted in exponential changes in cursor frequency, and frequency changes were binned in quarter-octave intervals to match rodent psychophysical discrimination thresholds (Han et al., 2007).

The rodents had to then precisely modulate these neuronal ensembles to move the cursor to one of two target pitches, one which was associated with a 20% sucrose solution reward and one which was associated with a 45 mg food pellet reward. Rodents were free to choose either reward at any time, although M1 activity levels had to return to baseline levels for a new trial to begin. Rats developed no clear preference for either reward during training ($p < 0.001$). A trial was marked incorrect if neither of these target states were achieved within 30 seconds of trial initiation, although in late learning animals only took 11.8 ± 1.1 seconds to hit targets (mean \pm sem).

Behavioral sessions took place in a rodent operant box equipped with liquid and pellet dispensers (Lafayette Instrument Company, Lafayette, IL). Recorded neural data was entered in real time to custom routines in Matlab that then translated those activity levels into the appropriate feedback pitch and played the pitch on speakers mounted on 2 sides of the operant box. Frequencies used for auditory feedback ranged from 1-24 kHz in quarter-octave increments. When a target was hit, a Data Acquisition board (National Instruments, Austin, TX) controlled by Matlab triggered the operant box to supply the appropriate reward to rodents.

3.2.4 Intracortical microstimulation

Intracortical microstimulation was applied to M1 while recording from the DS in order to estimate the corticostriatal conduction delay. Single monophasic cathodal pulses lasting 250 μsec with an amplitude of 300 μA were applied to M1 using a Model 2100 Isolated Pulse Stimulator (A-M Systems, Sequim, WA). Neuronal activity was simultaneously recorded from the DS using a separate recording headstage and preamplifier.

3.2.5 Data Analysis

Analyses were performed in Matlab (Mathworks, Natick, MA) with custom-written routines. Unit data were first binned in 1 millisecond time bins and digitized. As no difference was seen between the two targets in the basic task, data from the two targets were pooled. Peri-stimulus time histograms (PSTHs) were calculated in relation to target achievement and firing rate analyses were performed from 4 seconds before until 4 seconds after this event, as well as from 2 seconds before and after the event. Firing rates and PSTHs were smoothed with a moving window of 100 milliseconds. Neuronal modulation depths were calculated as the absolute value of the difference in firing rate for a given cell when hitting target 1 versus target 2. For a number of analyses, representative days for early and late learning were chosen for each animal from days 2-4 and 8-11, respectively, based on behavioral performance, motivation, and signal quality. Representative days contained a minimum of 100 completed behavioral trials and 8 well-isolated M1 units.

Coherence analyses were performed using algorithms from the Chronux toolbox (<http://chronux.org>) in conjunction with custom routines in Matlab. A multi-taper method was used to compute spectral estimates of spiking activity in both regions (Jarvis & Mitra, 2001; Thomson, 1982). A total of 5 tapers were used with a time-bandwidth product of 3, and estimates were computed every 50 milliseconds with a window size of 500 milliseconds. Coherence between spiking activity in the two regions was calculated and defined as:

$$C_{xy} = \frac{|R_{xy}|}{\sqrt{|R_{xx}|} \sqrt{|R_{yy}|}}$$

where R_{xx} and R_{yy} are the power spectra and R_{xy} is the cross-spectrum. Spectral analyses were calculated relative to the delivery of reward and averaged across trials and animals. M1 and DS were referenced independently, and therefore coherence was never calculated between signals sharing a common local reference. The mean event-related potential and time-varying firing rate for each LFP channel or recorded unit, respectively, was subtracted from each trial before calculation of coherence values. However, it has been shown that this ERP subtraction can also produce artifacts rather than removing them, and the correct procedure for controlling for evoked responses

remains an open question (Truccolo et al., 2002). The coherence analysis was therefore also performed without subtraction of the mean ERP and time-varying firing rate. Mean coherence in the alpha band was calculated and defined as 6-14 Hz. Significance testing on coherence estimates was performed on mean alpha band estimates across spike-field pairs using independent samples t-tests with the Bonferroni correction for multiple comparisons across time.

Coherence estimates can be affected by firing rate (Lepage et al., 2011) and we therefore performed a thinning procedure to equate firing rates between conditions in which rates differed (Gregoriou et al., 2009). Trial-averaged spike trains in the neuronal populations were smoothed with a moving average of 10 msec. The difference in firing rate between the populations normalized by the maximum firing rate at a given time point determined the probability that a spike would need to be removed from the population with a higher firing rate. Spikes were then removed from the population with a higher firing rate based on this probability in order to eliminate any possible influence of firing rate on coherence estimates.

Coherence phase offsets were then converted into temporal offsets for comparison with the corticostriatal conduction delay. Phase offset values were converted to temporal offsets by calculating the derivative of phase with respect to frequency across the 6-14 Hz band (Riddle & Baker, 2005) at every time point from 2 seconds before until 2 seconds after target reaching. The resulting distribution of estimates of the temporal delay was then compared against the distribution of conduction delay estimates obtained using ICMS. Unless otherwise noted, all histograms were normalized to their maximum value.

The STPC analysis is identical to calculations of intertrial phase coherence (Kim et al., 2007; Zervakis et al., 2011), except the analysis is time-locked to spikes rather than to behavioral events. For these analyses, LFP data were first bandpass filtered from 6-14 Hz and Hilbert transformed. The resulting complex values were normalized by their absolute value to extract phase and then averaged surrounding spiking activity. The amplitude of this measure therefore reflects the consistency of the LFP phase from spike to spike with a range of temporal lags, and differs from the other coherence calculations only in that the analysis is time-locked to the occurrence of spikes. However, by comparing the time course of coherence values surrounding spikes, the results of such an analysis are suggestive, though not conclusive, of the direction of influence between spiking and LFP activity. STPC values vary from 0 (no phase consistency) to 1 (perfect phase consistency). To test for significance, surrogate data sets were created by shuffling the timing of recorded spikes. 1000 of these surrogate data sets were created for each animal (6000 total), STPC was calculated for each, and the experimental values of STPC were compared against this distribution of surrogate STPC values. For calculations of STPC surrounding intracortical microstimulation, the stimulation artifact was removed by spline interpolation from the time of microstimulation until 50 msec after.

The signal-to-noise ratio for each recorded waveform was quantified as:

$$SNR = \frac{A}{2 * SD_{noise}}$$

where A is the peak-to-peak voltage of the mean waveform and SD_{noise} is the standard deviation of the residuals from each waveform after the mean waveform has been subtracted (Suner et al., 2005). Units included in the analysis had a minimum signal-to-noise ratio of 2.

3.3 Results

3.3.1 Acquisition of a neuroprosthetic skill

We developed a brain-machine interface (BMI) task in which rats were required to modulate activity in primary motor cortex (M1) irrespective of physical movements. Modulation of M1 ensemble activity produced changes in the pitch of an auditory cursor, which provided constant auditory feedback to rodents about task performance. Reward was delivered when rodents precisely modulated M1 activity to move this auditory cursor to one of two target tones, and a trial was marked incorrect if no target had been hit within a 30 second time limit. Two neural ensembles consisting of 2-4 well-isolated units each were randomly chosen and used to control the auditory cursor (see Methods). The action of these two ensembles opposed each other, such that increased activity in one ensemble produced increases in cursor pitch, while increased activity in the other ensemble caused decreases in cursor pitch. Thus, in order to achieve a high-pitched target, rodents had to increase activity in the first ensemble and decrease activity in the second, while the opposite modulations were necessary to hit a low-pitched target (Figure 3.1a). Firing rates were smoothed with a moving average of the past three 200-msec time bins, and rate modulations therefore had to be maintained for a target to be hit. In this sense, the task required rodents to volitionally bring M1 into a desired state irrespective of motor output. Importantly, this task allows us to directly define which cells are relevant for behavioral output, and therefore infer the causal link between neuronal activity in these cells and behavior.

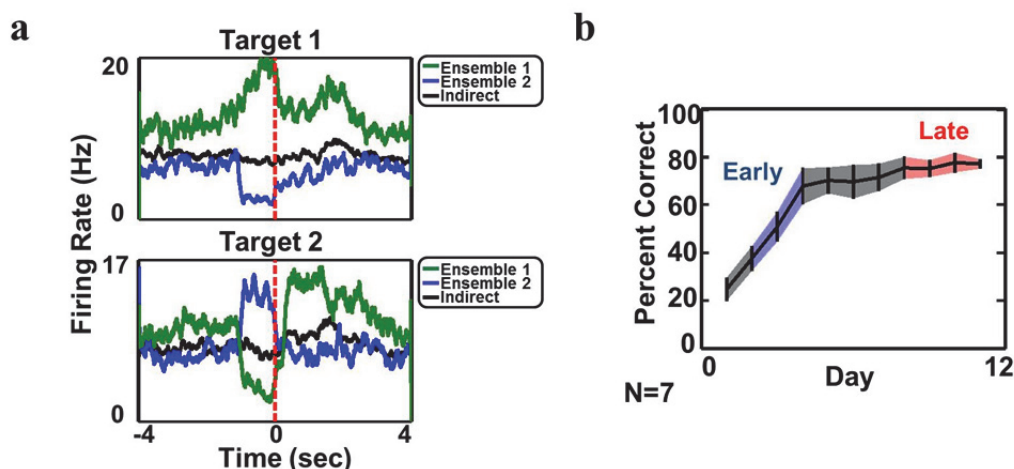


Figure 3.1 a. Mean M1 ensemble firing rates for units in ensemble 1 (green), ensemble 2 (blue), and M1 units not used in the transform (black) in relation to the achievement of target 1 (top) or target 2 (bottom). Time zero indicates target achievement (red dashed line). **b.** Mean percentage of correct responses for all rats across days 1-11 of learning. Shaded regions denote the s.e.m. and colored regions denote the range of days from which the early (blue) and late (red) analyses were performed.

We chronically implanted a group of rats ($N = 8$) with microelectrode arrays allowing us to simultaneously record activity in both M1 and the dorsal striatum (DS) throughout learning, and trained them in this paradigm. A subset of these rats were used in a previous study but underwent additional experimental manipulations for the present work, and two additional rats were used exclusively for this study. The mean percentage of correct trials increased greatly over the course of learning, and this measure followed a standard learning curve (Figure 3.1b). We could clearly see an initial phase of rapid improvement followed by a phase of slower learning, representing early (days 2-4) and late (days 8-11) phases of learning. The percentage of correct trials increased significantly from early to late in learning ($p < 0.001$) demonstrating that rats were able to properly learn to perform the task. Analyses of M1 firing rates further showed that rats were producing the desired ensemble rate modulations during task performance (Figure 3.1a).

3.3.2 Corticostriatal coherence develops during neuroprosthetic learning

We first investigated the relationship between spiking activity and the local field potential (LFP) oscillations recorded in each region during task engagement. We performed spike-triggered averaging of the LFP oscillation in late learning time-locked to spikes occurring either in the same region or in the other region. If spiking activity were independent of the LFP phase, then fluctuations would cancel out and produce a flat average LFP. We observed clear average LFP oscillations in both regions around action potentials from both regions; this oscillatory activity had a strong component between 6-14 Hz (Figure 3.2a). This is consistent with past work showing that oscillations in this range are particularly prominent in corticostriatal circuits when performing well-learned tasks (Berke et al., 2004), as well as work suggesting that motor cortex is predisposed

to operate in this frequency range (Castro-Alamancos, 2013). We therefore filtered the raw LFP from 6-14 Hz and calculated the predominant phase at which spikes occurred. Again, we observed clear phase-locking of spikes to the ongoing 6-14 Hz LFP in both regions (Figure 3.2b). Although the relationship between LFP and spiking activity is certainly complex and neurons may spike at several preferred LFP phases, there was nevertheless a dominant phase preference across both regions. Interestingly, both DS and M1 spikes occurred preferentially at the peak of the striatal 6-14 Hz LFP oscillation, suggesting that DS neuron firing is maximal at the peak of the striatal LFP.

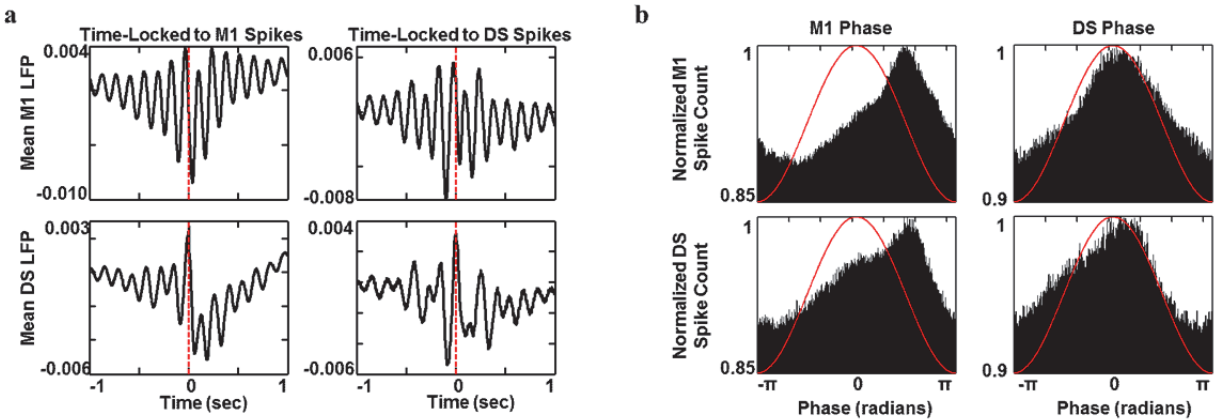


Figure 3.2 a. The mean M1 LFP (top row) or DS LFP (bottom row) time-locked to occurrences of spikes from M1 (left column) or DS (right column). All four average traces exhibit clear oscillatory activity with a strong component at roughly 8 Hz, showing that the phase at this frequency influences spiking activity. **b.** Spikes from M1 (top row) or DS (bottom row) fire at a preferred phase of the 6-14 Hz band in the M1 LFP (left column) or DS LFP (right column).

To further quantify these corticostriatal interactions and the ways in which they evolve during learning, we calculated coherence between spiking activity in M1 and LFP oscillations in DS. A total of 1936 spike-field pairs were analyzed (121 M1 units and 16 DS LFP channels). To avoid effects of evoked potentials on coherence estimates, the mean DS event-related potential (ERP) and M1 time-varying firing rate for each cell or LFP channel, respectively, were subtracted from individual trials before calculating coherence. We saw a profound increase in spike-field coherence across a range of low frequencies in late learning, when the rats were skillfully performing the task, relative to early learning (Figure 3.3a). This effect was most pronounced at frequencies between 6-14 Hz, and there was a significant increase in the mean coherence at these frequencies from early to late in learning (Figure 3.3b; $p < 0.001$, Bonferroni corrected). However, while subtracting the mean ERP often reduces the effect of evoked potentials on estimates of spike-field coherence, it has also been shown that such a procedure can produce artifacts (Truccolo, 2002). We therefore repeated the coherence analysis without subtracting the mean ERP, and again found a profound increase in 6-14 Hz coherence from early to late in learning. This change in coherence was not due to differences in trial number between early and late learning, as we matched the conditions for number of trials included in the analysis. Importantly, coherence was highest during target-reaching and decreased after trial completion at time 0 when the animals initiated movements towards the reward. Before trial completion, coherence

was significantly higher on correct relative to incorrect trials. In addition, coherence between the M1 LFP and DS LFP also increased from early to late in learning (Figure 3.3c), and this effect was also most pronounced between 6-14 Hz (Figure 3.3d). We therefore focused further analyses on this frequency band. These data suggest that corticostriatal neuronal ensembles became tightly coordinated over the course of learning.

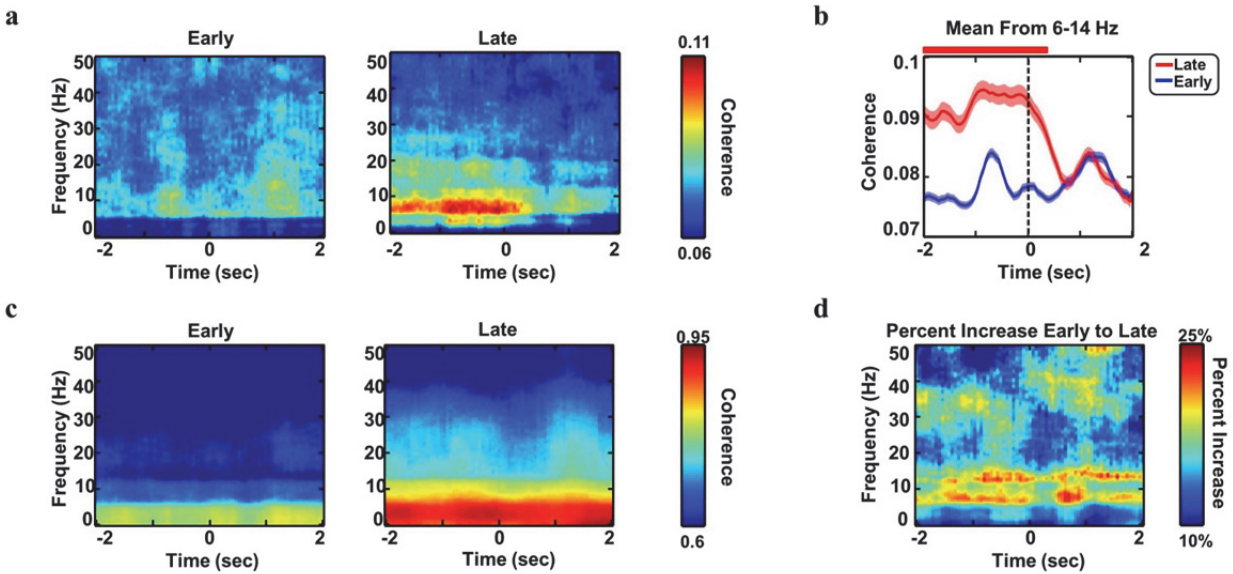


Figure 3.3 a. Coherograms illustrating the grand average of coherence between M1 spikes and DS LFP in early (left) and late (right) learning time-locked to target achievement. There is a clear increase from early to late in learning, with particularly pronounced activity in the 6-14 Hz band. **b.** The mean coherence from 6-14 Hz in early (blue) and late (red) learning time-locked to target achievement. Shaded regions denote s.e.m. Colored bars above plot designate time points with significant differences. **c.** Coherograms illustrating the grand average of coherence between M1 LFP and DS LFP in early (left) and late (right) learning time-locked to target achievement. There is a clear increase in low frequency coherence during learning. **d.** The percent increase in coherence from early to late in learning shows that this effect is most pronounced in the 6-14 Hz band.

3.3.3 Coherence is specific to output-relevant neurons and time periods

We then asked whether this increase in coherence between M1 spikes and DS LFP was observed for all M1 cells recorded or was specific for task-relevant cells. The operant neuroprosthetic task used here offers the unique advantage that the cells which are directly controlling the output of the BMI (hereafter “output cells”; N=31) are explicitly defined. Because past work has demonstrated enhanced neuronal modulations in output cells relative to cells not entered into the BMI (Ganguly et al., 2011) (hereafter “indirect cells”; N=89), we first examined the firing rate modulations that rats were producing during task performance. Although the indirect cells do not directly impact cursor movement, they are embedded in the same network as the output cells and modulation of their activity could therefore still play an indirect role in target

achievement. However, in late learning, rats modulated output cells significantly more than indirect cells surrounding target achievement (Figure 3.4a; $p < 0.001$), suggesting that the indirect cells were indeed being treated as less task-relevant than output cells. Importantly, we found that the M1-DS coherence that emerged during learning was highly specific to output cells (Figure 3.4b), even when they were recorded on the same electrode as indirect cells and separated from this population by less than 100 μm . This effect again appeared to be more pronounced in the 6-14 Hz range, with significantly larger coherence for output relative to indirect cells (Figure 3.4c; $p < 0.01$, Bonferroni corrected). We ensured that well-isolated units were included in both the output and indirect populations, and further verified that these populations did not differ in baseline firing rate (see Methods). Nevertheless, spikefield coherence estimates can be affected by firing rate (Lepage et al., 2011) and the task structure required differences in firing rates in the two populations during target achievement. We therefore performed a thinning procedure to equate firing rates between the two populations (Gregoriou et al., 2009) (see Methods). Despite differences in firing rate being removed, there remained a significant difference in spikefield coherence between output cells and indirect cells ($p < 0.001$, Bonferroni corrected), demonstrating that this effect was not driven by firing rate differences. To further ensure that our results were not affected by firing rate, we separated our analysis by both cell and trial type to examine trials in which output cells were required to increase their firing rate and trials in which output cells were required to decrease their firing rate to achieve the target. There was still a significant difference in coherence between output cells that were decreasing their firing rate relative to indirect cells ($p < 0.05$, Bonferroni corrected), despite no significant difference in firing rate between these same two populations. Finally, we performed the coherence analysis after removing cells with low signal-to-noise ratio from the indirect population and coherence remained higher in output cells than indirect cells, demonstrating that the effect was not driven by differences in waveform quality ($p < 0.05$, Bonferroni corrected). These coherent interactions were greatly diminished during the inter-trial intervals when rats were not actively engaged in the task (Figure 3.4d). Furthermore, during these periods, the difference in coherence between output and indirect populations was abolished (Figure 3.4e,f). These results suggest that the corticostriatal coherence that emerged during learning was highly specific for neurons that are directly relevant to behavioral output, even when they are closely intermingled with other cells, and that these precise interactions are flexible and appear rapidly during task performance.

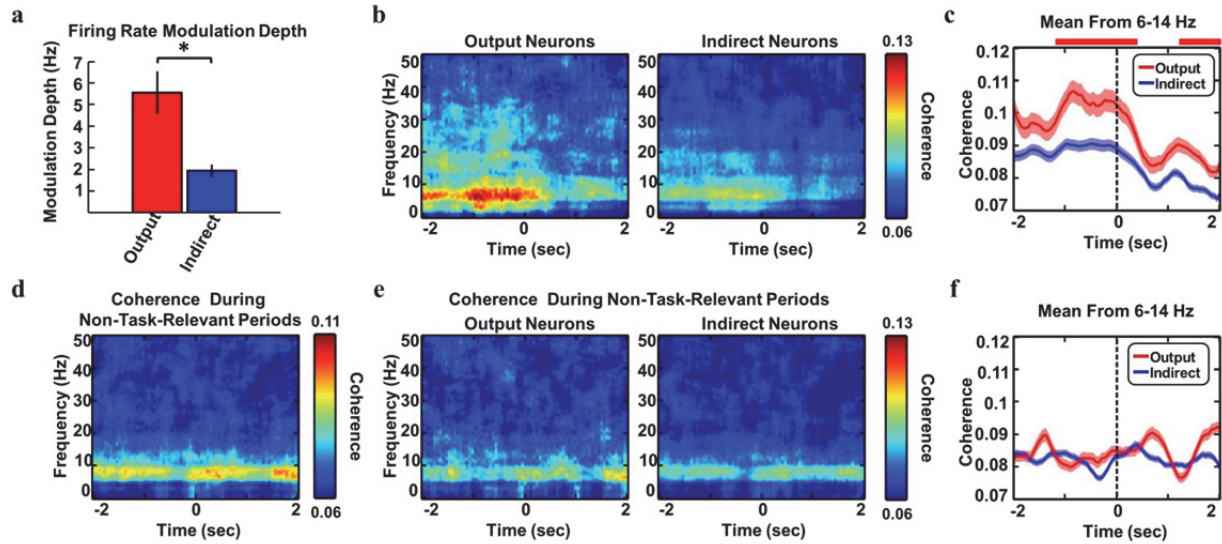


Figure 3.4 a. Firing rate modulation depth for output (red) and indirect (blue) cells in late learning. There is significantly greater modulation of output cells relative to indirect cells. **b.** Coherograms in late learning showing the grand average for output cells (left) and indirect cells (right) time-locked to target achievement. Coherence is markedly stronger in output than indirect cells. **c.** Mean coherence from 6-14 Hz in late learning for output cells (red) and indirect cells (blue) time-locked to target achievement. Shaded regions denote s.e.m. Colored bars above plot designate time points with significant differences. **d.** Coherence in late learning is greatly reduced during periods when rats are not actively engaged in the task. Plot shows the grand average across animals. **e.** When rats are not actively engaged in the task, there is no difference in coherence for output (left) and indirect (right) cells. **f.** Mean coherence from 6-14 Hz in late learning when rats are not actively engaged in the task again shows no difference between output (red) and indirect (blue) cells.

3.3.4 Precise timing of corticostriatal activity

Because we had found that M1 spikes were occurring preferentially at the peak of the DS LFP oscillations (Figure 3.2b), we next investigated the phase offset of the spike-field coherence. From the mean phase heat map, we see that there is a consistent negative phase offset in the 6-14 Hz range (Figure 3.5a). By convention, this suggests that M1 spikes precede the peak of the DS LFP in the 6-14 Hz band. Indeed, the phase offset at 6-14 Hz was commonly negative, as can clearly be seen in the distribution of phase offset values for every cell and every frequency in the alpha band (Figure 3.5b). When the phase offset values are used to estimate a temporal delay between M1 spikes and the DS LFP (see Methods), we see a clear preference for M1 cells to fire at an offset of roughly -5 to -7 msec relative to the DS LFP, as reflected in the mode of this distribution (Figure 3.5c; s.e.m. = 0.03 msec). This preference developed over the course of training and was not present in early stages of learning (Figure 3.5d) or in late learning during time periods when rats were not actively engaged in the task (Figure 3.5e), suggesting that it is not innately apparent in corticostriatal circuits. In addition, this temporal offset was specific to the 6-14 Hz band. These results show that M1 is on average spiking 6 msec before the peak of the LFP oscillations in DS when the animals are performing a well-learned task.

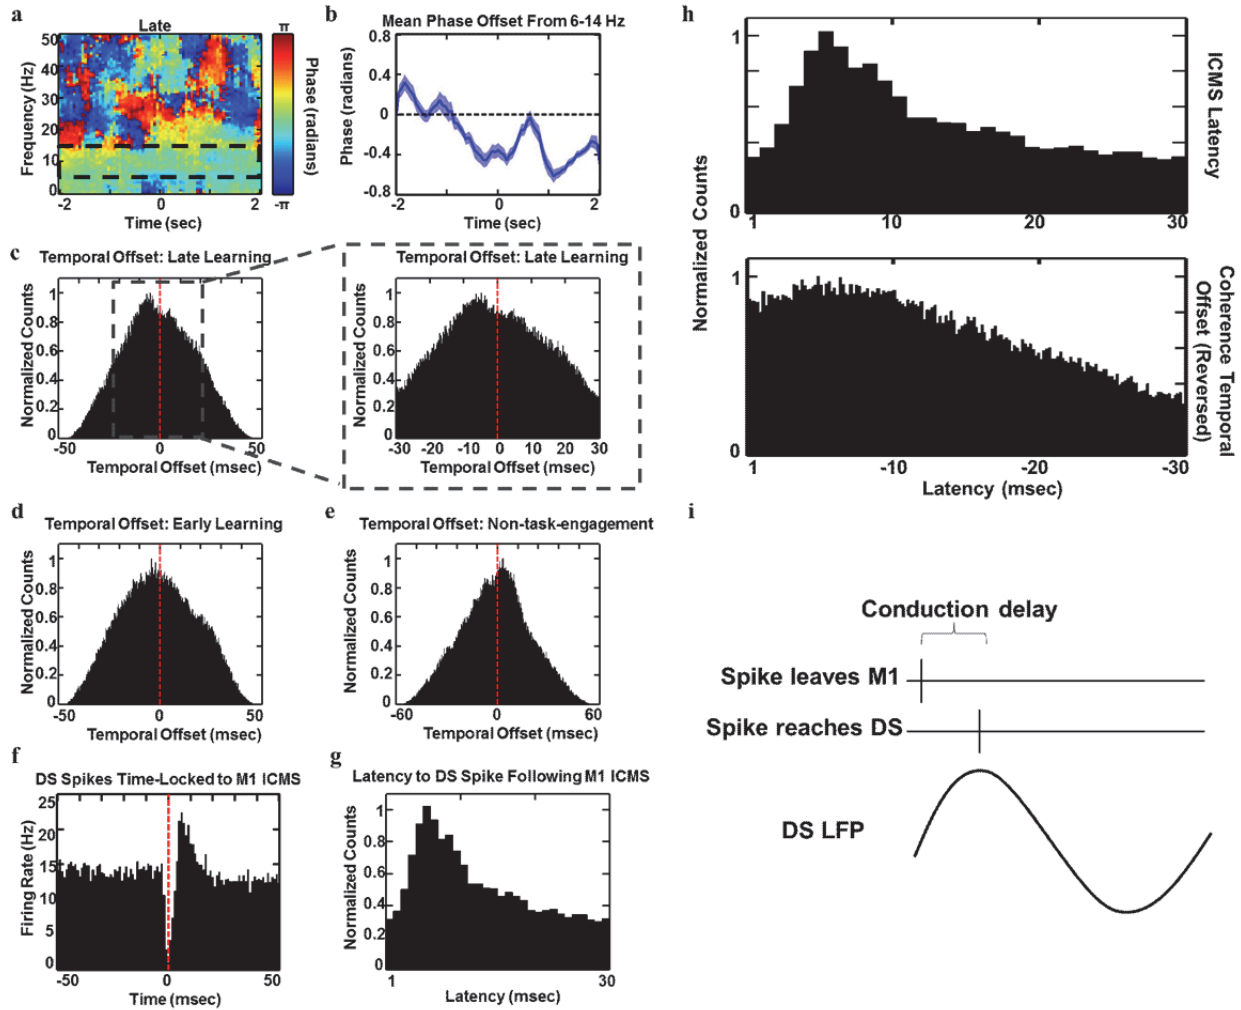


Figure 3.5 a. Phase values in late learning show a negative phase in the 6-14 Hz band relative to other frequency ranges. By convention, this suggests that M1 spikes lead the DS LFP. **b.** The mean phase from 6-14 Hz surrounding target achievement exhibiting consistently negative phase. **c.** The distribution of temporal offset estimates obtained from the coherence phase data for every trial and every M1 cell-DS LFP channel pair (see Methods) shows that M1 spikes most often occur 5-7 msec before the peak of the DS 6-14 Hz LFP. **d.** The distribution of temporal offset estimates in early learning does not show the same phase preference seen in late learning. **e.** The distribution of temporal offset estimates in late learning during time periods when the rats were not actively engaged in the task does not show the same phase preference seen during task engagement. **f.** Mean spiking response in the DS time-locked to application of ICMS to M1. **g.** Histogram of the latency to DS spikes following application of ICMS to M1 as a measure of the corticostriatal conduction delay. There is a clear peak 5-7 msec after application of ICMS. **h.** The ICMS-based estimate of the conduction delay (top) aligns remarkably well with temporal offset estimates from the spike-field coherence analysis (bottom). Temporal offsets are plotted with a reversed x-axis to correspond with the ICMS results. **i.** A working model for our results. M1 spikes precede the DS LFP 6-14 Hz band peak by 5-7 msec, which is on scale with the corticostriatal conduction delay. Thus, after accounting for this delay, M1 spikes arrive at the DS during peak excitability.

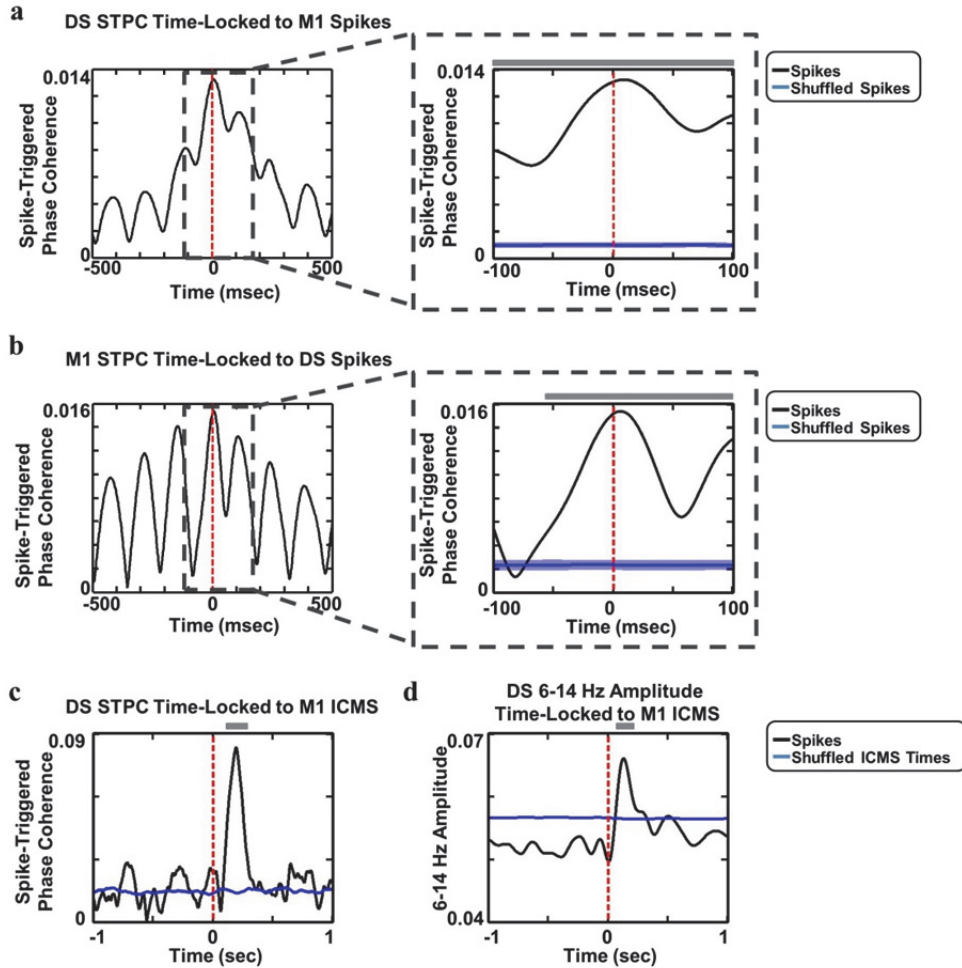


Figure 3.6 **a.** DS 6-14 Hz STPC time-locked to M1 spikes exhibits a marked peak immediately following M1 activity. DS STPC (black) is significantly greater than the distribution of 6000 STPC values obtained by shuffling the timing of recorded spikes (blue). Colored bars above plot designate time points with significant differences. **b.** M1 6-14 Hz STPC time-locked to DS spikes also exhibits a clear peak following DS activity. M1 STPC (black) is significantly greater than surrogate STPC values (blue). Colored bars above plot designate time points with significant differences. **c.** DS 6-14 Hz STPC (black) time-locked to application of ICMS to M1. ICMS in M1 is followed by a consistent phase in the DS. This is significantly greater than surrogate STPC values (blue). Colored bars above plot designate time points with significant differences. **d.** DS 6-14 Hz amplitude time-locked to application of ICMS to M1 (black). ICMS in M1 is followed by an increase in 6-14 Hz amplitude in the DS. This peak is significantly greater than values obtained on a surrogate dataset (blue). Colored bars above plot designate time points with significant differences.

The DS receives strong input from M1, and this temporal offset is concordant with past estimates of the conduction delay between the two regions (Cowan & Wilson, 1994), suggesting that M1 input may be driving DS firing (which occurs preferentially at the peak of striatal LFP oscillations; see Figure 3.2b). We therefore applied intracortical microstimulation (ICMS) to M1 while recording responses in the DS and estimated the delay between M1 stimulation and DS response. Brief cathodal pulses were applied to

M1 and produced a consistent spiking response in the DS (see Methods). A total of 3245 ICMS trials were performed in 7 animals over the course of several sessions. The mean peristimulus time histogram (PSTH) time-locked to stimulation shows a marked peak in DS spiking activity following application of ICMS to M1 (Figure 3.5f). For every cell, we then estimated the corticostriatal conduction delay by calculating the latency from M1 stimulation until the first DS spike occurred (Figure 3.5g). This distribution of latencies had a mode at 6 ms (s.e.m. = 0.1 msec), which is on scale with past estimates of the conduction delay obtained with intracellular recording (Van der Maelen & Kitai, 1980). There was striking alignment between this estimate of the conduction delay and the temporal offset values determined above, and these distributions were not significantly different from each other (Figure 3.5h, $p = 0.45$). Together these results suggest that M1 spikes in late learning are precisely timed to drive DS during task performance (Figure 3.5i).

3.3.5 Network activity drives the 6-14 Hz LFP oscillation

Our finding of a consistent nonzero phase lag concordant with the conduction delay between the two regions suggests that the two regions may interact directly rather than being coordinated by a third region. To further investigate a mechanism for these precise dynamics, we calculated spike-triggered phase coherence (STPC) in the 6-14 Hz band of both regions time-locked to spikes from either region (see Methods). STPC measures phase consistency from spike to spike. This measure will be 1 if, at a given time point, the phase is the same surrounding every spike, and it will be 0 if the phase is random. By investigating the time course of coherence fluctuations surrounding a spike, the STPC measure is suggestive of the direction of influence between spikes and LFP, although it cannot conclusively rule out the influence of a third region. Importantly, the DS STPC exhibited a pronounced peak after spikes from M1 are fired, showing that M1 spikes are followed by a highly consistent phase in the DS (Figure 3.6a; $p < 0.001$, Bonferroni corrected). Interestingly, we found a similar effect for the reverse calculation, with STPC in M1 significantly enhanced following spikes from the DS (Figure 3.6b; $p < 0.001$, Bonferroni corrected). This shows that activity in the DS is followed by a consistent phase in M1, and underscores that corticostriatal circuits function as re-entrant loops. To investigate whether M1 activity caused the 6-14 Hz activity or simply coordinated ongoing activity, we also calculated STPC in the DS surrounding application of ICMS to M1 (Figure 3.6c) and found that STPC was significantly enhanced following M1 ICMS ($p < 0.001$, Bonferroni corrected), suggesting that strong ICMS-induced activity in M1 produces entrainment that drives the DS 6-14 Hz oscillation. These peaks in STPC are significantly greater than values obtained with surrogate datasets in which the spike or event times were shuffled (see Methods). Importantly, M1 ICMS is also followed by an enhancement of 6-14 Hz amplitude in the DS (Figure 3.6d; $p < 0.001$, Bonferroni corrected), suggesting that strong M1 activity drives the 6-14 Hz activity in the DS rather than coordinating ongoing activity. Interestingly, the peak in 6-14 Hz amplitude following ICMS precedes the peak in STPC. This amplitude peak is again greater than values obtained with surrogate datasets, but as expected, surrogate values for the amplitude calculation are midway between the values seen in the actual data. They do not remain near zero like surrogate STPC

values because amplitude calculations are not as dependent on precise timing as phase calculations. Together these data suggest that, after learning, spiking in M1 or the DS produces a consistent LFP phase in the other region, resulting in reinforcement of coherent dynamics with precise timing throughout the network.

3.4 Discussion

In summary, we have shown that coherence develops in corticostriatal networks during learning with high temporal precision, and importantly, specifically involving the cells that control behavioral output, even when these cells are intermingled with other neuronal populations. This specificity suggests that coherence can serve to enhance communication between task-relevant populations and bias local competitive interactions in their favor. This, in turn, allows for rapid modulation of the functional connectivity between local ensembles and distant brain structures, and for flexible routing of specific signals throughout the brain as these signals become immediately relevant for behavior. Interestingly, this cell-specific coherence occurred predominantly in the alpha band, between 6-14 Hz. This is consistent with recent work showing low-frequency coherence between M1 spikes and DS spikes (Koralek et al., 2012). The slight shift in frequency between spike-spike coherence and spike-field coherence in the same task may reflect that spike-spike coherence measures similarity between the output spike trains of two regions, while spike-field coherence is thought to measure similarity between the output of one region and synchronous input to another (Zeitler et al., 2006). Differences between these measures in the dominant frequency of coherence could therefore reflect individual neurons not spiking on every cycle of the population rhythm or performing temporal integration of inputs.

A number of distinct rhythms have been previously observed in this frequency range. While some of these rhythms, such as high voltage spindles or mu rhythms, are thought to be generated in thalamocortical circuits (Hughes & Crunelli, 2005), other forms of 6-14 Hz LFP activity in M1 are thought to be generated via local circuit mechanisms (Castro-Alamancos, 2013). Importantly, sleep spindles in this frequency range have been associated with memory consolidation (Steriade & Timofeev, 2003). In addition, alpha band activity in the visual system (Kandel & Buzsaki, 1997) and mu rhythms in the sensorimotor system (Nicoletis et al., 1995), both centered roughly at 6-14 Hz, have frequently been associated with disengagement from external stimuli. Thus, our finding of enhanced phase-locking of M1 spikes to the DS alpha band LFP in late learning could reflect the rodents learning to disengage the corticostriatal system from the musculature in order to perform our neuroprosthetic task.

In addition, the precise timing of neuronal inputs that we observed could have consequences for network dynamics and plasticity throughout the brain. A large body of work has shown that temporal precision modulates the induction and direction of long-lasting synaptic plasticity (Dan & Poo, 2004). Indeed, computational models have demonstrated the importance of timing for spike-timing-dependent plasticity and information transfer through neuronal networks (Wang et al., 2010). Input timing is particularly important for the regulation of dendritic calcium levels in striatal cells and, in

turn, synaptic plasticity (Kerr & Plenz, 2004). Thus, the precise temporal dynamics demonstrated here may have important functional consequences for corticostriatal plasticity and its role in learning.

Our results also suggest the intriguing possibility that these precise temporal interactions can be maintained by activity within the network reinforcing the synchronous LFP oscillations. Cortico-basal ganglia circuits are organized as closed feedback loops (Hikosaka et al., 1999), with activity in any node influencing the flow of information throughout the system. Our finding of enhanced STPC following spikes in either M1 or DS therefore suggests that this flow of feedback through re-entrant corticostriatal loops maintains the orderliness and strength of coherent oscillations in the system. Indeed, while past work has suggested that oscillations spanning a range of frequencies may be produced in the thalamus, removal of corticothalamic feedback by decortication results in highly disordered oscillations (Contreras et al., 1996), highlighting the importance of network feedback mechanisms in the control and organization of coherent activity.

In summary, our data support coherence as an effective means by which functional cell assemblies can quickly form and disband to meet task demands, as well as demonstrating ways in which such neuronal interactions can be learned and adapted to support a lifetime of flexible, skilled behavior.

Chapter 4:

Fine-scale spatial characteristics of network reorganization

While we have seen striking reorganization in neuronal networks over the course of neuroprosthetic learning, microelectrode arrays suffer from a number of limitations. In particular, array recordings offer somewhat poor spatial resolution relative to several other techniques. To address this issue and dissect the spatial characteristics of network changes during learning, we developed a task involving operant conditioning of spike-related calcium signals recorded with high spatial resolution using two-photon imaging. Mice learned to modulate activity in both motor and somatosensory cortices, exhibiting learning across- and within-sessions, and this learning was accompanied by striking modifications to spatially localized networks. Crucially, the use of two-photon imaging allowed dissection of neuronal changes associated with neuroprosthetic learning with unprecedented spatial resolution.

4.1 Introduction

Brain-machine interfaces (BMIs) have gained great momentum in recent years as a therapeutic option for patients suffering from limb loss or immobility (Fetz, 2007; Carmena et al., 2003; Hochberg et al., 2012; Collinger et al., 2013). In addition to this direct clinical application, BMI tasks also provide a unique approach to studying sensorimotor learning, as they enable arbitrary mapping between neuronal activity, behavioral output, and reward (Green & Kalaska, 2011). Recent work has used BMI to demonstrate network adaptations in response to output perturbations (Jarosiewicz et al., 2008), as well as highly specific functional changes in output-relevant neurons (Ganguly et al., 2011; Koralek et al., 2013). However, past BMI work has been based on spatially sparse electrode recordings and lacks fine-scale spatial information about local network modifications. To address this issue, we developed a BMI task in awake, head-restrained mice using 2-photon calcium imaging in local neural ensembles, in which activity of every neuron in a small field of view (150 by 150 microns) is imaged. We used this novel calcium-based BMI paradigm (CaBMI) to probe the fine-scale spatial characteristics of network reorganization in cortical layer (L) 2/3 of both motor and somatosensory cortices during BMI learning.

4.2 Methods

All animal procedures were performed in accordance with University of California Berkeley Animal Care and Use Committee regulations. 6 C57BL/6J and 4 CD1 male wild-type mice were used in these experiments, ranging in age from postnatal day 30–45. Mice were housed with a 12-h dark, 12-h light reversed light cycle. All behavioral tests were performed in the same cohort of mice.

4.2.1 Surgery

Mice were anesthetized using 2% isoflurane (vol/vol) and placed in a stereotaxic apparatus. Body temperature was maintained at 37 °C using a feedback-controlled heating pad (FHC, 40-90-8D) and a small incision was made in the scalp. The skull was cleaned and a steel headplate was affixed over M1 (1 mm rostral, 1 mm lateral to Bregma) or S1 (1 mm caudal, 3 mm lateral to Bregma) using Metabond dental cement (Parkell, S380). A 3-mm craniotomy was opened over M1 or S1, and 200 nl of AAV2.9 Syn.GCamp6f.WPRE.SV40 (Chen et al., 2013) (University of Pennsylvania Vector Core) was injected 250 μm below the pia using a Nanoliter 2000 injector (World Precision Instruments). The tracer was delivered using a pulled glass pipette (tip diameter = 40–60 μm) at a rate of 40 nl min⁻¹. The pipette was left in the brain for 10 min after completion of the injection to prevent backflow. After the pipette was removed, the brain was covered with silicone oil (Sigma product # 181138) and a glass coverslip was affixed to the skull with dental cement, as previously described (Holtmaat et al., 2012). We allowed 2 weeks for recovery and gCaMP6f expression.

4.2.2 Two-photon imaging

In vivo imaging was performed with a Moveable Objective Microscope (Sutter) using a Chameleon Ultra Ti:Sapphire mode-locked laser (Coherent) tuned to 900 nm. Photons were collected with a Hamamatsu photomultiplier tube (H10770PA-40) using a Nikon objective (16 \times , 0.8 NA). Animals were head-fixed on a custom-made spring mounted imaging platform and placed under the two-photon microscope. This setup allowed them to run freely, and their movements were recorded by an accelerometer fixed to the underside of the platform. Frames of 128 \times 512 pixels (\sim 160 \times 160 μm) were collected at 7.23 Hz using ScanImage software (Pologruto et al., 2003) at 130–180 μm below the pia. The same imaging fields were used every day, localized by landmarks in the surface blood vessels. Imaged fields were stable over the course of training, and because the cortex was stabilized by a snugly fitting coverslip, only severe movements caused motion artifacts. Motion correction for slow drift in the imaging field was performed manually. Any period of gross movement during the task that caused cells to move out of their regions of interest (ROIs) resulted in poor task performance, as $\Delta F/F$ of E1 was reduced. In this sense, mice were punished for excessive movement and seem to have learned to remain still during the task.

4.2.3 Behavioral task

Two ensembles of 1–11 single cells each were chosen for inclusion in the output population. Cells with bright nuclei, indicating overexpression, were excluded, as were cells with many, poorly separable calcium events, an activity pattern indicative of fast-spiking interneurons. No other selection criteria were used to partition the recorded cells into each ensemble. We also ensured that many cells with good signal were included in the indirect population to enable a proper comparison. The cells assigned to the output population were changed on some days.

Ensemble activity was measured as mean $\Delta F/F$ for all component neurons. Fluorescence values from these ensembles were binned in 200-ms bins and entered into an online transform algorithm that related neural activity to the pitch of an auditory cursor. By modulating activity in these ensembles, rodents controlled the pitch of the cursor. The modulations that we required of the mice were calibrated daily based on a baseline recording session of roughly 2 min. Next, 10–15 min of spontaneous baseline activity was recorded to assess chance levels of performance. Fluorescence values were smoothed by a moving average of the past three time points. Changes in the frequency of the auditory cursor were binned in quarter-octave intervals to match rodent psychophysical discrimination thresholds (Han et al., 2007). Mice then had to modulate calcium dynamics in these neuronal ensembles to move the cursor to a high-pitched target tone that was associated with a 10% sucrose (wt/vol) solution reward. A trial was marked incorrect if a target was not achieved within 30 s of trial initiation. A trial was self-initiated when E1 and E2 activity returned to baseline levels (either by decreased activity in E1 or increased activity in E2), which reset the tone to its starting pitch.

Chance levels of performance (Figure 4.1b) were determined by running the animal on the task without reward or auditory feedback. Hits resulted when spontaneous fluctuations were large enough push the decoder to the target frequency. Failure to hit a target in 30 seconds resulted in a miss. The chance region represents the mean chance performance and s.e.m., pooled over all animals and all days.

ROIs were extracted from recorded neural data in real time. These ROIs were entered into custom routines in MATLAB (Mathworks) that translated fluorescence levels into the appropriate feedback pitch and played the pitch on speakers mounted on two sides of the imaging platform. Frequencies used for auditory feedback ranged from 1–24 kHz in quarter-octave increments. When a target was hit, a MATLAB-controlled Data Acquisition board (National Instruments) triggered the operant box to supply the appropriate reward to rodents. Each daily training session lasted 48 ± 2 min (71 ± 4 trials).

4.2.4 Data analysis

All analyses were performed with custom written routines in MATLAB. Recorded movies were spatially aligned using the `dfregistration` routine in MATLAB (Guizar-Sicairos et al., 2007). ROIs were manually selected to include the soma of neurons that appeared consistently throughout all recorded movies. Fluorescence traces were extracted from each ROI and data is presented as the relative change in fluorescence, $\Delta F/F$ (Chen et al., 2011).

No statistical methods were used to pre-determine sample sizes, but our sample sizes are similar to those generally employed in the field. For analyses of behavioral performance during the contingency degradation (Figure 4.1f), the first ten trials of a session were removed before calculating performance to exclude the transition period and reflect the animal's performance once the animal had fully learned the new reward contingency. For all sliding window analyses, sessions were divided into an equal

number of bins to determine the window size, and the step size was a fraction of this window size.

For the data plotted in Figure 4.2g, mean z score values during task engagement were binned by distance from E1 or E2 centroid. The first bin included all cells from 0–50 μm from the centroid of the output ensemble (close cells), the second bin included all cells 50–100 μm from the output ensemble, and the final bin including all cells 100+ μm from the output ensemble. ‘Distant’ cells included all cells at a distance of greater than 50 μm from the E1 centroid. We include data from 3 d late in training from 5 mice where 20 or more indirect cells were apparent in the field.

For the cross-correlation histograms, fluorescence traces from output cells were z-scored and values above 3 s.d. were considered an event. The first time point in which fluorescence values crossed this threshold during each event was used for time-locking. Fluorescence values in other populations of cells were then averaged around these indices.

In all cases, multiple comparisons were controlled for using the Bonferroni correction. Differences between groups were tested with *t* tests. To evaluate trends over time, we tested whether the slope of a fitted linear regression was significantly different from zero. All statistical tests were two-tailed.

For testing the activity modulations for low versus more active cell groups in Figure 2b, the high active group included cells with spontaneous event rate greater than the median spontaneous event rate, the low active group included cells with spontaneous event rate less than the median.

Data distributions were assumed to be normal, but this was not formally tested. Data collection and analysis were not performed blind to the experimental conditions. Randomization was not performed, as the experiment primarily involved within-animal comparisons and there were no multiple experimental cohorts.

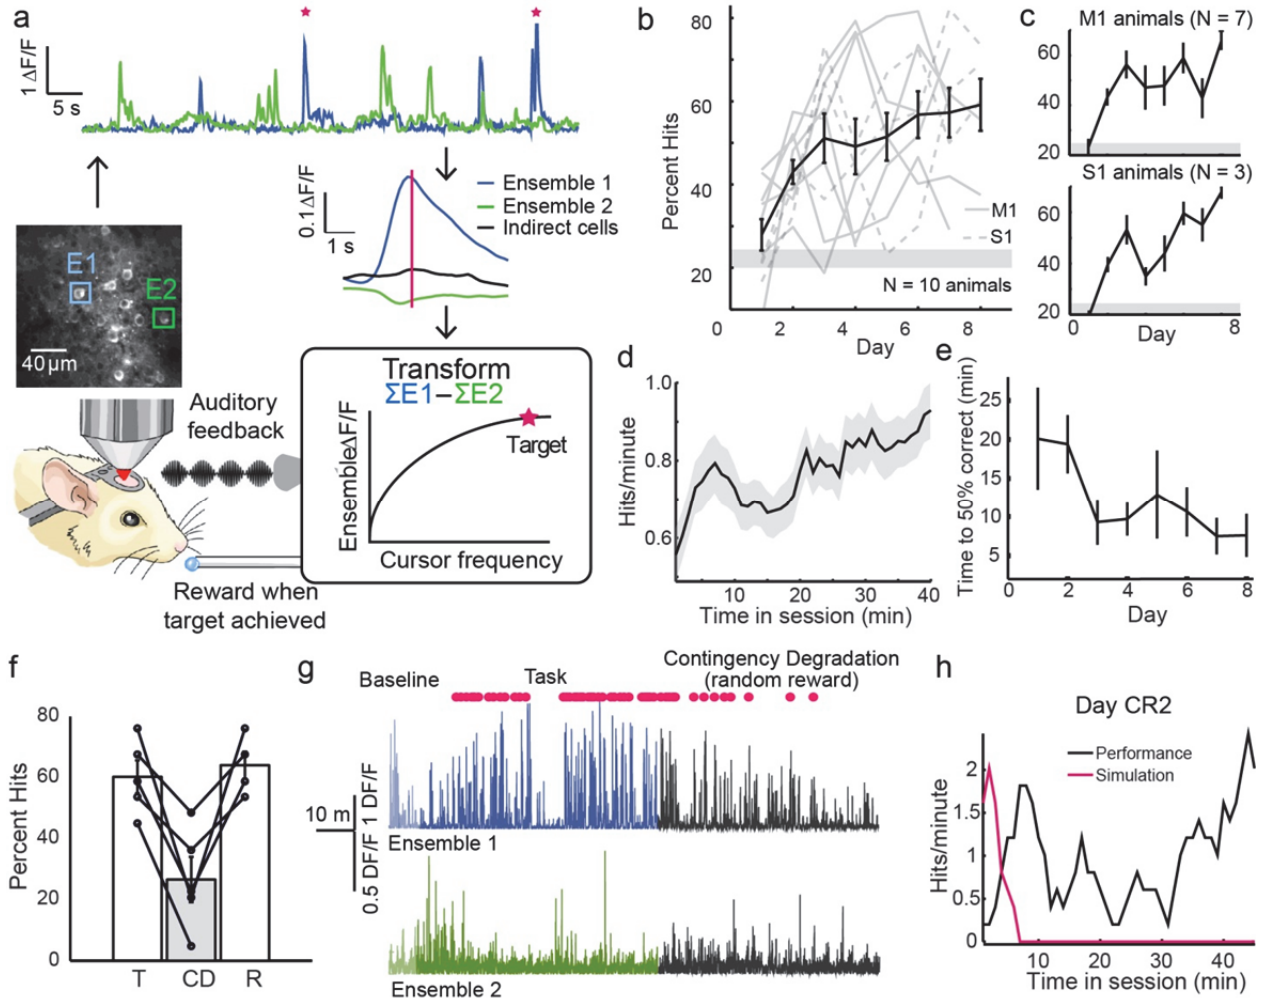


Figure 4.1 **a.** Task schematic. **b.** Behavioral performance over 8 days of training. Mean performance is shown in black, and individual animals are shown in gray. Error bars denote s.e.m. Shaded region denotes chance levels of performance. **c.** Behavioral performance for animals trained to use M1 (top) and S1 (bottom). **d.** Hit rate increases over the course of individual sessions. Shaded region denotes s.e.m. **e.** Animals achieve 50% performance faster over the course of training. **f.** Performance rapidly dropped compared to normal task levels (T) when the reward was randomized (contingency degradation, CD). Performance returned to previous levels when contingency was reinstated (R). **g.** $\Delta F/F$ in E1 increases during the task and rapidly decreases during contingency degradation. Likewise, target hits (red) increase in frequency over training, and decrease during contingency degradation. **h.** At the beginning of day 2 of contingency reversal, the animal initially performs as if the previous day's transform algorithm were still in use, but quickly learns the new transform.

4.3 Results

4.3.1 Volitional modulation of calcium dynamics

We trained ten mice expressing the genetically-encoded calcium indicator gCaMP6f in L2/3 of either primary motor (M1) or primary somatosensory (S1) cortex to modulate neural activity in response to auditory feedback (Figure 4.1a; see Methods). This task was adapted from one used previously with electrode-based recordings (Koralek et al., 2012). Each day, two ensembles containing 1-11 neurons each were chosen to control the task. Ensemble activity was measured as mean $\Delta F/F$ for all component neurons. The ensembles opposed each other, such that increased activity in one ensemble (“E1”) above its baseline (measured in a 10-15 min spontaneous activity period prior to task onset) increased the pitch of the auditory feedback, while increased activity in the other ensemble (“E2”) decreased the pitch. Tone pitch was updated every 200 ms as a function of the relative activity (E1-E2), and a sugar water reward was delivered when a specific high-pitched target was reached within 30 seconds of trial initiation (hit trial). Incorrect trials (target not hit within 30 seconds) were signaled with white noise. A trial was self-initiated when E1 and E2 activity returned to baseline levels, which reset the tone to its starting pitch. Ensemble cells were occasionally changed when calcium signals had significantly degraded.

Mice learned the task rapidly (Figure 4.1b), with initial rapid improvement (1-3 days) followed by slower improvement (4-8 days). After just 1 day of training, mice performed above chance level (Figure 4.1b, shaded region, 10 mice, $p = 0.0036$ on day 2, $t(8) = 4.07$). Similar learning occurred using either M1 or, more surprisingly, S1 (Figure 4.1c). Hit rate increased significantly within each daily session (Figure 1d, $p = 2.6 \times 10^{-5}$, $t(43) = 4.7$, $R^2 = 0.34$). Mice reached a criterion performance level (50% hits) faster across sequential days of training (Figure 4.1e, $p = 0.0247$, $t(6) = 2.98$, $R^2 = 0.596$), suggesting that within-session learning occurs faster as between-session learning progresses. As seen previously (Koralek et al., 2012), performance was not impaired by lidocaine injection into the contralateral mystacial pad ($p = 0.876$, $t(3) = 0.17$), and gross limb movements were absent preceding target hits, indicating that performance does not rely on natural movement and that neural activity, particularly in S1, is not driven by whisker reafference (data not shown).

4.3.2 Probing intentionality

We next asked whether these modulations were sensitive to the contingency between action and outcome (Hilario et al., 2007). We performed a contingency degradation, whereby after an animal successfully learned the task, we ceased rewarding target hits and instead delivered rewards under a variable interval schedule. Mice quickly ceased responding (Figure 4.1f-g; $p = 0.0089$, $t(4) = 4.76$). When reward was reinstated the next day using the same E1 and E2 ensembles as the previous day, mice again performed at normal levels (Figure 4.1f; $p = 0.791$, $t(3) = 0.289$). Thus, the animal’s performance was sensitive to reward contingency. Post-hoc analysis of imaging data showed that E1 activity increased during task performance, and decreased during

degradation (Figure 4.1g). On a separate day, we performed a contingency reversal ($N = 3$ mice) in which E1 and E2 identities were reversed between one day (termed day CR1) to the next (day CR2), requiring mice to reverse E1 and E2 neural activity patterns to obtain reward. At the beginning of day CR2, E2 in one example mouse showed clear bursting activity (consistent with its identity as E1 on day CR1), and E1 showed little activity (consistent with its identity as E2 on CR1). This pattern quickly reversed as the mouse learned the new contingency. We compared the behavioral hit rate on day CR2 to a simulated hit rate based on the E1/E2 identity and transform algorithm from day CR1. The simulation showed initially high performance that then dropped to zero, indicating that this mouse initially performed according to the learned CR1 transform, but then quickly adapted to the new CR2 transform (Figure 4.1h). Across all mice, the ratio of E1/E2 activity increased over the course of day CR2 ($N = 3$, $p < 0.01$), demonstrating that the animals learn to flexibly up-modulate E1 over E2. Together, these data indicate that mice can modulate calcium signals in a contingency-dependent manner, and that these modulations can be flexibly applied to arbitrarily chosen cells.

4.3.3 Neuronal changes during neuroprosthetic learning

We next investigated neural changes that accompany the learning process. We used a sliding window analysis to examine how activity changed over the course of a session in E1 and E2 cells (“output cells”). Mean $\Delta F/F$ increased for E1 cells over the course of individual sessions (Figure 4.2a; $p = 1.17 \times 10^{-11}$, $t(18) = 15.09$, $R^2 = 0.927$), decreased during subsequent contingency degradation ($p = 0.05$, $t(18) = 2.08$), mirroring the changes in hit rate. In contrast, mean $\Delta F/F$ did not significantly change for E2 cells during the task or contingency degradation, indicating that task learning was primarily driven by modulation of E1 $\Delta F/F$. This may reflect a bias toward volitional increases, rather than decreases, of mean L2/3 neuron activity over baseline.

Calcium imaging detects activity even in neurons that are rarely active, which are numerous in L2/3 (O’Connor et al., 2010; Barth et al., 2012). Although detectable with the extracellular recordings used in most BMI studies, these cells are undersampled by extracellular recordings and are often neglected in BMI studies (Shoham, 2006). Thus, their role in BMI learning is unclear. We observed a 30-fold range of baseline spontaneous activity across L2/3 cells (measured in the pre-task period) (Figure 4.2b). E1 cells with low spontaneous activity increased frequency of calcium events during the task, whereas cells with high spontaneous activity tended to slightly suppress burst frequency during task. E2 cells tended to lower their frequency of calcium events during task engagement, independent of spontaneous responsiveness (Figure 4.2b), even though mean fluorescence did not change over the session. Thus, task learning tended to preferentially recruit low-active E1 neurons to become more active. “Silent” L2/3 neurons clearly contributed to learning, because behavioral learning occurred normally even when all E1 cells had very low or zero baseline activity suggesting a role for “silent” L2/3 neurons in learning (Barth et al., 2012). Although the task structure encourages correlated network activity, this is not strictly enforced and mice can achieve targets with uncoordinated activity modulations. Within multi-cell E1 ensembles, multiple cells increased fluorescence around hits, and not just single neurons.

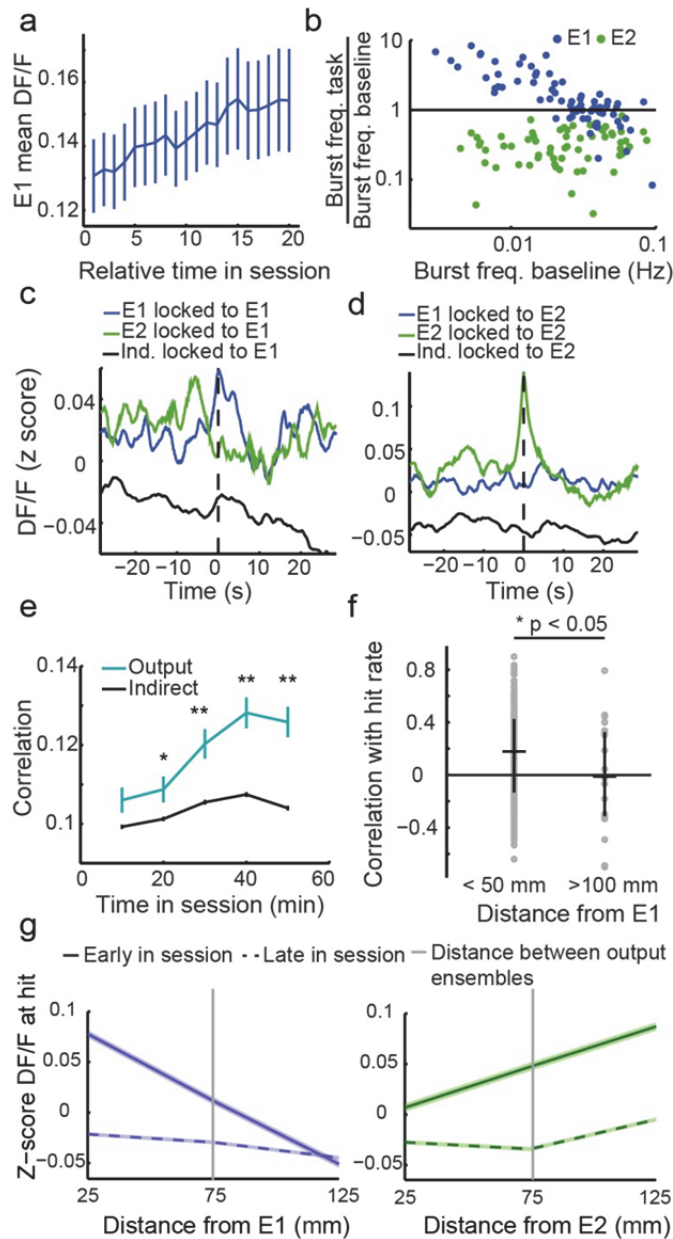


Figure 4.2 **a.** Mean fluorescence increases in E1 cells over the course of a behavioral session. **b.** E1 cells with low levels of spontaneous activity increase their activity more during the task than cells with high levels of spontaneous activity. E2 cells suppress their activity evenly. Plotted on a logarithmic axis. **c.** Activity in E1, E2 and indirect cells time-locked to large events in E1 cells. **d.** Activity in E1, E2 and indirect cells time-locked to large events in E2 cells. **e.** Correlations increase between output cells (blue-green) over the course of the session, with no similar increase in correlations between indirect cells (black). **f.** Indirect cells located near output cells have more task-related activity than those located far from output cells. **g.** Early in a session (solid lines), target-related rate modulations in indirect cells decrease with distance from cells in E1 (blue) and increase with distance from cells in E2 (green). Later in the session (dashed lines) there are no significant modulations in indirect cells, regardless of distance from output cells.

4.3.4 Network changes during neuroprosthetic learning

To examine higher-level network dynamics during learning, we first calculated the mean cross-correlation histograms time-locked to the occurrence of large fluorescence bursts in either E1 or E2 (see Methods). E1 and E2 cells developed coordinated, synchronous activity with other cells in the same ensemble (Figure 4.2c,d). E2 cells also developed a tendency to spike before E1 cells (Figure 4.2c,d), likely reflecting a strategy of bursting E2 in order to reset the cursor for trial initiation, followed by bursting of E1 to drive the cursor to target. This coordinated activity was not present in non-E1, non-E2 cells that were simultaneously imaged (Figure 4.2c,d; “indirect cells”). This prompted us to

investigate activity correlations between cells over the course of individual sessions. We found that correlations between output cells in the same ensemble increased significantly over the course of the session ($p = 0.0198$, $t(3) = 4.55$, $R^2 = 0.874$), while correlations between indirect cells did not (Figure 4.2e). By the second time point, output cells in the same ensemble were significantly more correlated than indirect cells ($p = 0.0056$, $t(60,366) = 2.77$). Within each session, correlations between output cells were initially similar to those between indirect cells, but output cells became more correlated than indirect cells as the session progressed. This enhanced correlation between output cells in the same ensemble was observable in individual animals. Output cells also became more correlated over days of training, even though neural composition of ensembles changed. This is analogous to increased correlations of functionally similar cells during natural motor learning (Komiyama et al., 2010) and could reflect tight coupling with millisecond precision that has been demonstrated with penetrating electrodes (Engelhard et al., 2013). These data suggest that output cells become more coordinated as animals learn the task, and this effect is not present in the indirect population as a whole.

4.3.5 Fine-scale changes in spatial activity patterns during learning

We next examined the spatial organization of learning-related activity at a fine spatial scale (~10-100 μm) that is not possible with electrode-based BMI. Learning performance did not vary systematically with distance between output ensembles (measured by E1 and E2 centroids), indicating that the proximity of output cells does not confer an advantage in learning ($p = 0.906$). Performance did depend on the number of cells in an ensemble: animals were more successful at learning the task with fewer neurons, suggesting that it was difficult to maintain coordinated control over multiple randomly selected neurons. Additionally, high baseline correlations between E1 and E2 cells predicted worse correlations. Next, we analyzed the spatial profile of learning within local networks surrounding the E1 and E2 ensembles. For each indirect cell, we calculated the correlation between its mean fluorescence and a moving average of the animal's instantaneous hit rate in temporal windows over the course of the session. We found that activity in indirect cells close to E1 (<50 μm away from E1 centroid) was significantly more correlated with hits than activity in distant indirect cells (>100 μm away from E1; Figure 4.2f, $p = 0.048$, $t(249) = 1.98$). Finally, we calculated mean target-related activity modulations in indirect cells for early and late epochs within daily learning sessions (Figure 4.2g). Early in sessions, indirect cells near the E1 centroid exhibited increased $\Delta F/F$ around hits, while more distant indirect cells did not. In contrast, indirect cells near the E2 centroid showed modest activity suppression early in sessions in indirect cells near E2, with increased activity evident in cells distant from E2 and close to E1. This shows that early in the session mice modulate activity in a relatively large local network surrounding output cells (Ganguly et al., 2011), but as the session progresses this target-related activity modulation in indirect cells disappears, such that only direct E1 and E2 cells exhibited task-related modulations. This suggests that as learning progresses mice are able to hone in on the individual output cells and precisely modulate the activity of these cells for efficient target achievement. However, indirect neurons that were more highly spontaneously correlated with E1 cells, and

therefore more likely to be embedded in the same local sub-network (Harris et al., 2013), were more likely to continue modulating their activity during the task ($p=5e-5$). Given the rapid falloff of spontaneous correlations with distance, such fine-scale effects cannot be measured with electrode-based recording methods. Together, these data demonstrate that mice initially modulate activity in a larger local network of cells that falls off with distance from the output cells, but over the course of learning are able to restrict activity modulations to the cells that are directly relevant to behavioral output, or preferentially functionally coupled to these output cells. This restriction in spatial activity is similar to sparsening of cortical representation observed during classical conditioning (Gdalyahu et al., 2012).

4.4 Discussion

To our knowledge, this CaBMI task is the first demonstration that mice can volitionally modulate calcium dynamics in L2/3 of M1 and S1, and the use of imaging has allowed for the dissection of learning-related network modifications during BMI with unprecedented spatial resolution. Mice rapidly learn the task and exhibit learning both within- and across-sessions. We found the learned neuronal modulations in this task to be very sensitive to changes in reward contingency, suggesting that they are goal-directed (Dias-Ferreira et al., 2009). This learning is accompanied by gradually increasing correlations between direct output cells. Furthermore, this learning involves a transient increase in activity modulation in neighboring indirect cells that disappears as learning progresses and circuits are refined so that primarily output cells are modulated. These findings demonstrate that the cortex can identify and select specific cells and neural patterns that are relevant for obtaining specific outcomes (Costa et al., 2011). Importantly, this novel paradigm provides a powerful tool for investigating the spatial extent of functional and structural plasticity during neuroprosthetic learning, and can be easily combined with other techniques, such as whole-cell recording and tracer injections, to precisely dissect learning strategies employed by the cortex.

Chapter 5:

Conclusions and open questions

The ability to learn new skills and perfect them with practice is fundamental to our daily lives. Whether it be refining the precise motor patterns used for playing a musical instrument, or learning the series of movements necessary to cook a great meal, or even developing an intuition for more abstract skills like chess and mathematics, our ability to learn from experience is what has led to the nearly endless repertoire of behaviors that human beings are capable of producing. However, the neuronal mechanisms of this form of learning and the seemingly limitless capacity of the procedural memory system remain poorly understood.

This work has investigated neuronal mechanisms of skill learning from the novel perspective of neuroprosthetic skills. By exploring these questions through a unique vantage point, we have been able to explore neural systems for skill learning that do not directly involve the musculature, and we have seen striking functional adaptations over the course of learning that are specific on the single cell level. In addition, we have seen that BMI tasks rely on similar neural substrates as natural skill learning, highlighting the intriguing ways in which these two areas of active research exist synergistically.

5.1 Summary of contributions

A primary motivation of this work was to investigate similarities between natural motor learning and neuroprosthetic learning, and the ways in which the two sub-disciplines can inform one another and capitalize on the respective benefits of each. Throughout this work, we have seen important implications for both BMI and basic science that could have only been elucidated through this novel perspective. In addition, a growing literature has supported the similarities between these two forms of learning (Green & Kalaska, 2011), and utilized this paradigm to gain important insights into the neural mechanisms at play (Jarosiewicz et al., 2008).

In Chapter 2, we presented a novel rodent paradigm for neuroprosthetic control that required rodents to learn an arbitrary mapping between neuronal activity and behavioral output. Learning followed a typical learning curve across slightly over a week of training, with rapid improvement in the first few days followed by slower improvement up to a plateau level in subsequent days. Using a three-axis accelerometer and EMG recordings in the mystacial pad, we verified that the animals were producing the desired neural modulations irrespective of gross motor output or more subtle muscular activity. We further verified, through three distinct manipulations on outcome value, that the animals were performing these modulations in a goal-directed rather than habitual manner.

In the second half of Chapter 2, we found striking changes in corticostriatal firing patterns in rats as task learning progressed. We found that striatal neurons fired significantly more strongly in late learning and a significantly larger proportion of striatal neurons exhibited target-locked firing rate modulations in late learning than in early learning, similar to what has been seen in natural motor learning (Yin et al., 2009). We also found changes in the interactions between M1 and the striatum, with significantly greater coherence between spikes in the two regions as learning progressed. Finally, we ran the task on both wildtype mice and mutant mice lacking functional NMDA receptors in the DS. Wildtype mice learned the task well and exhibited the same neuronal changes observed in rats, but mutant mice were unable to perform the task and exhibited none of the neural changes discussed above. Together, this work demonstrates that functional corticostriatal plasticity is necessary for animals to learn abstract skills that do not involve the musculature, and further suggests that BMI learning capitalizes on the neural systems for learning that are already in place in the nervous system.

In Chapter 3, we examined these coherent corticostriatal interactions in greater depth, in line with the theory of “neuronal communication through neuronal coherence” (Fries, 2005). We found that coherence increases between M1 spikes and the striatal LFP over the course of learning. These coherent interactions were highly specific to the population of output cells relative to neighboring indirect cells, even when these cells were recorded on the same electrode. In addition, this coherence was only present surrounding target achievement, and not during the inter-trial interval when rats were not actively performing the task. We then measured the corticostriatal conduction delay and found it to match the coherence phase offset, demonstrating that activity in the network becomes precisely timed to effectively drive other downstream regions. Finally, we found that activity in any node of the network is followed by reinforced phase consistency in other nodes, suggesting a form of positive feedback in maintaining these precise, coherent dynamics.

In Chapter 4, we adapted the behavioral paradigm for use with two-photon imaging. Mice again learned an arbitrary mapping between neuronal activity patterns and behavioral output, but now they were head-fixed during task performance and they were modulating calcium signals from cells in L2/3 rather than spikes from cells in L5. Their performance improved both within single sessions, as well as across days of training. We again found that the neuronal modulations were being performed in a goal-directed manner, and we found that mice could flexibly apply these learned modulations to arbitrarily chosen cell populations during the contingency reversal manipulation. Over the course of learning, output cells developed striking coordinated activity. Finally, using the fine-scale spatial information available with two-photon imaging, we found that at the beginning of training sessions, mice were modulating activity in a local network of indirect cells surrounding output cells, and these task-relevant modulations fell off with distance from output cells. Over the course of learning, however, mice honed in on the behaviorally-relevant neuronal populations and eventually modulated activity only in the output cells.

5.1.1 Additional engineering collaborations

In addition to the findings discussed in this dissertation, this work resulted in a number of interesting collaborations with other graduate students developing novel methods for the recording and analysis of neural signals. These projects will not be discussed in depth, but will be briefly mentioned here.

First, the data and findings from this work were used to validate the use of directed information as a measure of functional connectivity in neuronal networks (So et al., 2012). In addition, the experiments presented here on intracortical microstimulation were used in collaborative work to develop new methods for the reduction of stimulation artifacts when performing simultaneous stimulation and recording (Chu et al., 2013). This work together has forwarded the development of novel algorithms for acquiring and analyzing high-quality neural data.

In addition, several collaborations have aimed to develop and validate novel devices for neural recording. For example, we implanted micro-electrocorticography grids chronically in rats and then trained the rats to use the recorded neural signals for neuroprosthetic control (Ledochowitsch et al., 2013). In addition, we tested a new circuit for completely wireless neural recording, including wireless powering and data transfer, using the same micro-electrocorticography grid (Muller et al., 2014). Together, these collaborations furthered the development of novel technologies for neural recording.

5.2 Open questions and future directions

As we have seen, the fields of natural motor learning and neuroprosthetic learning mutually benefit from interactions between the disciplines. There is now ample evidence to support the view that the two processes share many common neural substrates, and there are similarly many opportunities to utilize these commonalities to answer a number of important open research questions.

5.2.1 The role of the striatum in tasks that do not involve the motor system

By utilizing neuroprosthetic tasks that do not involve the musculature, we have been able to demonstrate a role for the striatum in abstract skill learning, or the learning of skills that are irrespective of physical movements. However, the neuroprosthetic tasks used here still utilized neuronal activity from motor and somatosensory cortices, which are still part of the broader motor system. To more definitively demonstrate a role for corticostriatal plasticity in tasks that are completely divorced from the motor system, control signals must be acquired from brain regions that are far removed from the motor system. Recent work in our lab has begun to use activity from primary visual cortex towards this aim, and has showed promising initial results.

5.2.2 Basal ganglia and thalamic contributions to corticostriatal signaling

The work presented in Chapter 3 demonstrated intriguing mechanisms between M1 and the striatum to facilitate temporal precision and synchrony throughout the network. However, these brain regions are part of much larger cortico-basal ganglia-thalamic loops, with many nodes also participating in larger-scale interactions. To more precisely dissect the mechanisms of neuronal communication and network formation, recordings must be taken from other regions in the system, including all other nuclei in the basal ganglia, as well as functionally relevant nuclei in the thalamus that participate in such large-scale neuronal coordination.

5.2.3 Structural plasticity during neuroprosthetic learning

With the development of a neuroprosthetic task involving two-photon calcium imaging, researchers are now able to visually identify the relevant network involved in neuroprosthetic learning. Importantly, this allows for the investigation of structural changes that accompany neuroprosthetic learning. For example, initial stages of motor learning are characterized by the rapid turn-over of dendritic spines, and these spines stabilize as performance improves (Xu et al., 2009). Similar mechanisms could be at work during neuroprosthetic learning, but this can only now be examined with structural imaging of the network. In addition, neuroprosthetic tasks can be used to investigate whether any observed structural changes are specific to output-relevant neuronal populations, or whether learning is accompanied by broader structural plasticity in cortex.

5.2.4 Local circuit mechanisms of neuroprosthetic learning

Finally, with the use of the novel two-photon BMI paradigm, we can begin to answer more mechanistic questions about functional changes that occur during neuroprosthetic learning. For example, because the network can now be visualized, researchers can now easily insert intracellular electrodes into output cells to examine how the balance between inhibitory and excitatory inputs changes over the course of learning. In addition, a number of tracers can be injected into relevant neuronal populations to follow changes in connectivity that occur during learning. All of these important research questions have only recently become possible with the development of a BMI task using two-photon imaging.

5.3 Conclusion

Directly interfacing machines with the mind has not only provided clinically important therapies for patients suffering from limb loss or immobility, but has also provided an exciting and completely novel vantage point for studying neural mechanisms of complex, abstract skill learning. By capitalizing on the benefits of neuroprosthetic tasks, we have been able to demonstrate high selectivity in fundamental neuronal processes that researchers have previously only been able to describe on much larger spatial scales. Although a number of open questions remain, the development of an imaging-

based neuroprosthetic task allows for a greater repertoire of experimental tools that can be applied towards elucidating these remaining unknown processes. Through further work, research at the intersection of natural and neuroprosthetic learning will continue to shed light on the fundamental neural mechanisms that allow us to learn a lifetime of flexible, skilled behaviors.

Bibliography

- Atallah, H.E., Lopez-Paniagua, D., Rudy, J.W., & O'Reilly, R.C. Separate neural substrates for skill learning and performance in the ventral and dorsal striatum. *Nature Neurosci.*, 10, 126-131 (2007).
- Badre, D., Kayser, A.S., & D'Esposito, M. Frontal cortex and the discovery of abstract action rules. *Neuron*, 66, 315-326 (2010).
- Balleine, B.W. & Dickinson, A. Goal-directed instrumental action: contingency and incentive learning and their cortical substrates. *Neuropharmacology*, 37, 407-419 (1998).
- Barnes, T.D., Kubota, Y., Hu, D., Jin, D.Z., & Graybiel, A.M. Activity of striatal neurons reflects dynamic encoding and recoding of procedural memories. *Nature*, 437, 1158-1161 (2005).
- Barth, A.L. & Poulet, J.F.A. Experimental evidence for sparse firing in the neocortex. *Trends in Neurosciences*, 35, 345-355 (2012).
- Beauchamp, M.H., Dagher, A., Aston, J.A., & Doyon, J. Dynamic functional changes associated with cognitive skill learning of an adapted version of the Tower of London task. *Neuroimage*, 20, 1649-1660 (2003).
- Berke, J.D., Okatan, M., Skurski, J., & Eichenbaum, H.B. Oscillatory entrainment of striatal neurons in freely moving rats. *Neuron*, 43, 883-896 (2004).
- Bollimunta, A., Mo, J., Schroeder, C.E., & Ding, M. Neuronal mechanisms and attentional modulation of corticothalamic alpha oscillations. *J. Neurosci.*, 31, 4935-4943 (2011).
- Bragin, A., Engel, J. Jr., & Straba, R.J. High-frequency oscillations in epileptic brain. *Curr. Opin. Neurol.*, 23, 151-156 (2010).
- Brashers-Krug, T., Shadmehr, R., & Bizzi, E. Consolidation in human motor memory. *Nature*, 382, 252-255 (1996).
- Brasted, P.J. & Wise, S.P. Comparison of learning-related neuronal activity in the dorsal premotor cortex and striatum. *Eur. J. Neurosci.*, 19, 721-740 (2004).
- Buschman, T.J., Denovellis, E.L., Diogo, C., Bullock, D., & Miller, E.K. Synchronous oscillatory neural ensembles for rules in the prefrontal cortex. *Neuron*, 76, 838-846 (2012).
- Buzsaki, G. & Draguhn, A. Neuronal oscillations in cortical networks. *Science*, 304, 1926-1929 (2004).
- Canolty, R.T., Ganguly, K., Kennerley, S.W., Cadieu, C.F., Koepsell, K., Wallis, J.D., & Carmena, J.M. Oscillatory phase coupling coordinates anatomically dispersed functional cell assemblies. *PNAS*, 107, 17356-17361 (2010).
- Carmena, J.M. et al. Learning to control a brain-machine interface for reaching and grasping by primates. *PLoS Biology*, 1, e2 (2003).
- Castro-Alamancos, M.A. The motor cortex: a network tuned to 7-14 Hz. *Front. Neur. Circuits*, 7, doi:10.3389/fncir.2013.00021 (2013).
- Chen, T.-W. et al. Ultrasensitive fluorescent proteins for imaging neuronal activity. *Nature*, 499, 295-300 (2013).
- Chopra, A., Klassen, B.T., & Stead, M. Current clinical application of deep-brain stimulation for essential tremor. *Neuropsychiatr. Dis. Treat.* 9, 1859-1865 (2013).
- Chu, P., Muller, R., Koralek, A., Carmena, J.M., Rabaey, J.M., & Gambini, S.

- Equalization for intracortical microstimulation artifact reduction. Proceedings of the 35th Annual International Conference of the IEEE EMBS, pp. 245–248 (2013).
- Chudasama, Y. & Robbins, T.W. Functions of frontostriatal systems in cognition: comparative neuropsychopharmacological studies in rats, monkeys, and humans. *Biol. Psych.*, 73, 19-38 (2006).
- Colgin, L.L., Denninger, T., Fyhn, M., Hafting, T., Bonnevie, T., Jensen, O., Moser, M.B., & Moser, E.I. Frequency of gamma oscillations routes flow of information in the hippocampus. *Nature*, 462, 323-327 (2009).
- Collinger, J.L. et al. High-performance neuroprosthetic control by an individual with tetraplegia. *The Lancet*, 381, 557–564 (2013).
- Contreras, D., Destexhe, A., Sejnowski, T.J., & Steriade, M. Control of spatiotemporal coherence of a thalamic oscillation by corticothalamic feedback. *Science*, 274, 771-774 (1996).
- Costa, R.M., Cohen, D., & Nicolelis, M.A. Differential corticostriatal plasticity during fast and slow motor skill learning in mice. *Curr. Biol.*, 14, 1124-1134 (2004).
- Costa, R.M., Lin, S.C., Sotnikova, T.D., Cyr, M., Gainetdinov, R.R., Caron, M.G., & Nicolelis, M.A. Rapid alterations in corticostriatal ensemble coordination during acute dopamine-dependent motor dysfunction. *Neuron*, 52, 359-369 (2006).
- Costa, R.M. A selectionist account of de novo action learning. *Current Opinion in Neurobiology*, 21, 579–586 (2011).
- Cowan, R.L. & Wilson, C.J. Spontaneous firing patterns and axonal projections of single corticostriatal neurons in the rat medial agranular cortex. *J. Neurophys.*, 71, 17-32 (1994).
- Cunningham, J. P., Nuyujukian, P., Gilja, V., Chestek, C. A., Ryu, S. I., and Shenoy, K. V. A closed-loop human simulator for investigating the role of feedback control in brain-machine interfaces. *Journal of Neurophysiology* 105, 1932 – 1949 (2011).
- Dan, Y. & Poo, M.-M. Spike timing-dependent plasticity of neural circuits. *Neuron*, 44, 23-30 (2004).
- 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.*, 103, 15254-15259 (2006).
- Dangi, S., So, K., Orsborn, A.L., Gastpar, M.C., and Carmena, J.M. Brain-machine interface control using broadband spectral power from local field potentials. In Proceedings of the 35th Annual International Conference of the IEEE EMBS, pp. 285–288 (2013).
- Dean, H.L., Hagan, M.A., & Pesaran, B. Only coherent spiking in posterior parietal cortex coordinates looking and reaching. *Neuron*, 73, 829-841 (2012).
- DeCoteau, W.E., Thorn, C., Gibson, D.J., Courtemanche, R., Mitra, P., Kubota, Y., & Graybiel, A.M. Learning-related coordination of striatal and hippocampal theta rhythms during acquisition of a procedural maze task. *PNAS*, 104, 5644-5649 (2007).
- Dias-Ferreira, E. et al. Chronic stress causes frontostriatal reorganization and affects decision-making. *Science*, 325, 621–625 (2009).
- Di Filippo, M. et al. Short-term and long-term plasticity at corticostriatal synapses: implications for learning and memory. *Behav. Brain Res.*, 199, 108-118 (2009).

- Engelhard B, Ozeri, N., Israel, Z., Bergman, H. & Vaadia, E. Inducing gamma oscillations and precise spike synchrony by operant conditioning via brain-machine interface. *Neuron*, 77, 361–375 (2013).
- Ethier, C., Oby, E. R., Bauman, M. J., and Miller, L. E. Restoration of grasp following paralysis through brain-controlled stimulation of muscles. *Nature* 485, 368–371 (2012).
- Fetz, E.E. Volitional control of neural activity: implications for brain-computer interfaces. *J. Physiol.*, 579, 571-579 (2007).
- Fincham, J.M. & Anderson, J.R. Distinct roles of the anterior cingulate and prefrontal cortex in the acquisition and performance of a cognitive skill. *Proc. Natl. Acad. Sci. USA*, 103, 12941-12946 (2006).
- Fries, P. A mechanism for cognitive dynamics: neuronal communication through neuronal coherence. *Trends Cogn Sci*, 9, 474-480 (2005).
- Fries, P., Wolmesdorf, T., Oostenveld, R., & Desimone, R. The effects of visual stimulation and selective visual attention on rhythmic neuronal synchronization in macaque area V4. *J. Neurosci.*, 28, 4823-4835 (2008).
- Frohlich, F. & McCormick, D.A. Endogenous electric fields may guide neocortical network activity. *Neuron*, 67, 129-143 (2010).
- Fujisawa, S. & Buzsaki, G. A 4 Hz oscillation adaptively synchronizes prefrontal, VTA, and hippocampal activities. *Neuron*, 72, 153-165 (2011).
- Gandolfo, F., Li, C., Benda, B.J., Schioppa, C.P., & Bizzi, E. Cortical correlates of learning in monkeys adapting to a new dynamical environment. *Proc. Natl. Acad. Sci. USA*, 97, 2259-2263 (2000).
- Ganguly, K. & Carmena, J.M. Emergence of a stable cortical map for neuroprosthetic control. *PLoS Biol.*, 7, e1000153 (2009).
- Ganguly, K., Secundo, L., Ranade, G., Orsborn, A., Chang, E.F., Dimitrov, D.F., Wallis, J.D., Barbaro, N.M., Knight, R.T., & Carmena, J.M. Cortical representation of ipsilateral arm movements in monkey and man. *J Neurosci*, 29, 12948-12956 (2009).
- Ganguly, K., and Carmena, J. M. Neural correlates of skill acquisition with a cortical brain-machine interface. *Journal of Motor Behavior* 42, 355–360 (2010).
- Ganguly, K., Dimitrov, D.F., Wallis, J.D. & Carmena, J.M. Reversible large-scale modification of cortical networks during neuroprosthetic control. *Nat. Neurosci* doi:10.1038/nn.2797 (2011).
- Garcia-Munoz, M., Carrillo-Reid, L., & Arbutnott, G.W. Functional anatomy: dynamic states in basal ganglia circuits. *Front. Neuroanat.*, 4, 1-7 (2010).
- Gdalyahu, A. et al. Associative fear learning enhances sparse network coding in primary sensory cortex. *Neuron*, 75, 121–132 (2012).
- Georgopoulos, A.P., Taira, M., & Lukashin, A. Cognitive neurophysiology of the motor cortex. *Science*, 260, 47-52 (1993).
- Grahn, J.A., Parkinson, J.A., & Owen, A.M. The cognitive functions of the caudate nucleus. *Prog. Neurobiol.*, 86, 141-155 (2008).
- Graybiel, A.M., Aosaki, T., Flaherty, A.W., & Kimura, M. The basal ganglia and adaptive motor control. *Science*, 265, 1826-1831 (1994).
- Graybiel, A.M. The basal ganglia and chunking of action repertoires. *Neurobiol. Learn. Mem.*, 70, 119-136 (1998).

- Graybiel, A.M. Habits, rituals, and the evaluative brain. *Ann. Rev. Neurosci.*, 31, 359-387 (2008).
- Green, A.M. & Kalaska, J.F. Learning to move machines with the mind. *Trends Neurosci.*, 34, 61-75 (2011).
- Gregoriou, G.G., Gotts, S.J., Zhou, H., & Desimone, R. High-frequency, long-range coupling between prefrontal and visual cortex during attention. *Science*, 324, 1207-1210 (2009).
- Guizar-Sicairos, M., Thurman, S.T. & Fienup, J.R. Efficient subpixel image registration algorithms. *Optics Letters*, 33, 156–158 (2008).
- Han, Y.K., Köver, H., Insanally, M.N., Semerdjian, J.H., & Bao, S. Early experience impairs perceptual discrimination. *Nature Neurosci.*, 10, 1191-1197 (2007).
- Harris, K.D., Csicsvari, J., Hirase, H., Dragoi, G., & Buzsaki, G. Organization of cell assemblies in the hippocampus. *Nature*, 424, 552-556 (2003).
- Harris, K.D. & Mrsic-Flogel, T.D. Cortical connectivity and sensory coding. *Nature*, 503, 51–58 (2013).
- Hikosaka, O., Nakahara, H., Rand, M.K., Sakai, K., Lu, X., Nakamura, K., Miyachi, S., & Doya, K. Parallel neural networks for learning sequential procedures. *Trends Neurosci.*, 22, 464-471 (1999).
- Hilario, M.R., Clouse, E., Yin, H.H., & Costa, R.M. Endocannabinoid signaling is critical for habit formation. *Front. Integr. Neurosci.*, 1, 1-12 (2007).
- Hilario, M.R. & Costa, R.M. High on habits. *Front. Neurosci.*, 2, 208-217 (2008).
- Hochberg, L.R. et al. Reach and grasp by people with tetraplegia using a neutrally controlled robotic arm. *Nature*, 485, 372–375 (2012).
- Holtmaat, A. et al. Imaging neocortical neurons through a chronic cranial window. *Cold Spring Harbor Protocols*, 2012, 694–701 (2012).
- Hoover, J.E. & Strick, P.L. Multiple output channels in the basal ganglia. *Science*, 259, 819-821 (1991).
- Huerta, P.T. & Lisman, J.E. Heightened synaptic plasticity of hippocampal CA1 neurons during a cholinergically induced rhythmic state. *Nature*, 364, 723-725 (1993).
- Hughes, S.W. & Crunelli, V. Thalamic mechanisms of EEG alpha rhythms and their pathological implications. *Neuroscientist*, 11, 357-372 (2005).
- Jarosiewicz, B., Chase, S.M., Fraser, G.W., Velliste, M., Kass, R.E., & Schwartz, A.B. Functional network reorganization during learning in a brain-computer interface paradigm. *PNAS*, 105, 19486-19491 (2008).
- Jarvis, M. & Mitra, P. Sampling properties of the spectrum and coherency of sequences of action potentials. *Neural Comput.*, 13, 717-749 (2001).
- Jin, X. & Costa, R.M. Start/stop signals emerge in nigrostriatal circuits during sequence learning. *Nature*, 466, 457-462 (2010).
- Kandel, A. & Buzsaki, G. Cellular-synaptic generation of sleep spindles, spike-and-wave discharges, and evoked thalamocortical responses in the neocortex of the rat. *J. Neurosci.*, 17, 6783-6797 (1997).
- Kerr, J.N.D. & Plenz, D. Action potential timing determines dendritic calcium during striatal up-states. *J. Neurosci.*, 24, 877-885 (2004).
- Kim, Y.J., Grabowecky, M., Paller, K.A., Muthu, K., & Suzuki, S. Attention induces synchronization-based response gain in steady-state visual evoked potentials. *Nat. Neurosci.*, 10, 117-125 (2007).

- Kimchi, E.Y. & Laubach, M. Dynamic encoding of action selection by the medial striatum. *J. Neurosci.*, 29, 3148-3159 (2009).
- Komiyama, T. et al. Learning-related fine-scale specificity imaged in motor cortex circuits of behaving mice. *Nature*, 464, 1182–1186 (2010).
- Koralek, A.C., Jin, X., Long, J.D., Costa, R.M., & Carmena, J.M. Corticostriatal plasticity is necessary for learning intentional neuroprosthetic skills. *Nature*, 483, 331-335 (2012).
- 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).
- Koyama, S., Chase, S. M., Whitford, A. S., Velliste, M., Schwartz, A. B., and Kass, R. E. Comparison of brain–computer interface decoding algorithms in open-loop and closedloop control. *Journal of Computational Neuroscience* 29, 73–87 (2010).
- Kubota, Y., Liu, J., Hu, D., DeCoteau, W.E., Eden, U.T., Smith, A.C., & Graybiel, A.M. Stable encoding of task structure coexists with flexible codeing of task events in sensorimotor striatum. *J. Neurophys.*, 102, 2142-2160 (2009).
- Lakatos, P., Shah, A.S., Knuth, K.H., Ulbert, I., Karmos, G., & Schroeder, C.E. An oscillatory hierarchy controlling neuronal excitability and stimulus processing in the auditory cortex. *J. Neurophys.*, 94, 1904-1911 (2005).
- Lakatos, P., Karmos, G., Mehta, A.D., Ulbert, I., & Schroeder, C.E. Entrainment of neuronal oscillations as a mechanism of attentional selection. *Science*, 320, 110-113 (2008).
- Ledochowitsch, P., Koralek, A.C., Moses, D., Carmena, J.M., & Maharbiz, M.M. Sub-mm functional decoupling of electrocortical signals through closed-loop BMI learning. *Proceedings of the 35th Annual International Conference of the IEEE EMBS*, pp. 5622–5625 (2013).
- Lepage, K.Q., Kramer, M.A., & Eden, U.T. The dependence of spike field coherence on expected intensity. *Neural Computation*, 23, 2209-2241 (2011).
- Millán, J. d. R., Galan, F., Vanhooydonck, D., Lew, E., Philips, J., and Nuttin, M. Asynchronous non-invasive brain-actuated control of an intelligent wheelchair. In *Proceedings of the 31st Annual International Conference of the IEEE EMBS*, pp. 3361–3364 (2009).
- Mitra, P. & Pesaran, B. Analysis of dynamic brain imaging data. *Biophys. J.*, 76, 691-708 (1999).
- Miyachi, S., Hikosaka, O., & Lu, X. Differential activation of monkey striatal neurons in the early and late stages of procedural learning. *Exp. Brain Res.*, 146, 122-126 (2002).
- Muller, R., Le, H.-P., Li, W., Ledochowitsch, P., Gambini, S., Bjorninen, T., Koralek, A., Carmena, J.M., Maharbiz, M.M., Alon, E., & Rabaey, J.M. A miniaturized 64-channel 225 μ W wireless electrocorticographic neural sensor. Presented at the annual meeting of the International Solid-State Circuits Conference, San Francisco (2014).
- Nicolelis, M.A., Baccala, L.A., Lin, R.C., & Chapin, J.K. Sensorimotor encoding by synchronous neural ensemble activity at multiple levels of the somatosensory system. *Science*, 268, 1353-1358 (1995).
- O'Connor, D.H., Peron, S.P., Huber, D. & Svoboda, K. Neural activity in barrel cortex underlying vibrissa-based object localization in mice. *Neuron*, 67, 1048–1061

- (2010).
- Orsborn, A.L., Dangi, S., Moorman, H.G., & Carmena, J.M. Closed-loop decoder adaptation on intermediate time-scales facilitates rapid BMI performance improvements independent of decoder initialization conditions. *IEEE TNSRE*, 20, 468-477 (2012).
- Pascual, J., Velasco-Alvarez, F., Muller, K.R., & Vidaurre, C. First study towards linear control of an upper-limb neuroprosthesis with an EEG-based Brain-Computer Interface. In *Proceedings of the Annual International Conference of the IEEE EMBS*, pp. 3269-73 (2012).
- Pasupathy, A. & Miller, E.K. Different time courses of learning-related activity in the prefrontal cortex and striatum. *Nature*, 433, 873-876 (2005).
- Pennartz, C.M.A., Berke, J.D., Graybiel, A.M., Ito, R., Lansink, C.S., van der Meer, M., Redish, A.D., Smith, K.S., & Voorn, P. Corticostriatal interactions during learning, memory processing, and decision making. *J. Neurosci.*, 29, 12831-12838 (2009).
- Pesaran, B., Pezaris, J.S., Sahani, M., Mitra, P.P., & Andersen, R.A. Temporal structure in neuronal activity during working memory in macaque parietal cortex. *Nat. Neurosci.*, 5, 805-811 (2002).
- Pesaran, B., Nelson, M.J., & Andersen, R.A. Free choice activates a decision circuit between frontal and parietal cortex. *Nature*, 453, 406-409 (2008).
- Poldrack, R.A., Prabhakaran, V., Seger, C.A., & Gabrieli, J.D. Striatal activation during acquisition of a cognitive skill. *Neuropsychology*, 13, 564-574 (1999).
- Pologruto, T.A., Sabatini, B.L. & Svoboda, K. ScanImage: flexible software for operating laser scanning microscopes. *BioMed Eng OnLine*, 2, 13 (2003).
- Riddle, C.N. & Baker, S.N. Manipulation of peripheral neural feedback loops alters human corticomuscular coherence. *J. Physiol.*, 566, 625-639 (2005).
- Riehle, A., Grun, S., Diesmann, M., & Aertsen, A. Spike synchronization and rate modulation differentially involved in motor cortical function. *Science*, 278, 1950-1953 (1997).
- Riout-Pedotti, M.S., Friedman, D., & Donghue, J.P. Learning-induced LTP in neocortex. *Science*, 290, 533-536 (2000).
- Rodriguez, E., George, N., Lachaux, J.P., Martinerie, J., Renault, B., & Varela, F.J. Perception's shadow: long-distance synchronization of human brain activity. *Nature*, 397, 430-433 (1999).
- Schoffelen, J.-M., Poort, J., Oostenveld, R., & Fries, P. Selective movement preparation is subserved by selective increases in corticomuscular coherence. *J. Neurosci.*, 31, 6750-6758 (2011).
- Serruya, M.D., Hatsopoulos, N.G., Paninski, L., Fellows, M.R., & Donoghue, J.P. Instant neural control of a movement signal. *Nature*, 416, 141-142 (2002).
- Shen, G., Zhang, J., Wang, M., Lei, D., Yang, G., Zhang, S., & Du, X. Decoding the individual finger movements from single-trial functional magnetic resonance imaging recordings of human brain activity. *Eur J Neurosci.*, doi: 10.1111/ejn.12547. (2014).
- Shoham, S., O'Connor, D.H. & Segev, R. How silent is the brain: is there a "dark matter" problem in neuroscience? *J Comp Physiol A* 192, 777-784 (2006).
- Siegel, M., Warden, M.R., & Miller, E.K. Phase-dependent neuronal coding of objects in short-term memory. *PNAS*, 106, 21341-21346 (2009).

- So, K., Koralek, A.C., Ganguly, K., Gastpar, M.C., & Carmena, J.M. Assessing functional connectivity of neural ensembles using directed information. *J Neural Eng*, 9, doi: 10.1088/1741-2560/9/2/026004 (2012).
- So, K., Dangi, S., Orsborn, A.L., Gastpar, M.C., & Carmena, J.M. Subject-specific modulation of local field potential spectral power during brain-machine interface control in primates. *J Neural Eng*, 11:026002 (2014).
- Steinmetz, P.N., Roy, A., Fitzgerald, J., Hsiao, S.S., Johnson, K.O., & Niebur, E. Attention modulates synchronized neuronal firing in primate somatosensory cortex. *Nature*, 404, 187-190 (2000).
- Steriade, M. & Timofeev, I. Neuronal plasticity in thalamocortical networks during sleep and waking oscillations. *Neuron*, 37, 563-576 (2003).
- Stern, E.A., Kincaid, A.E., & Wilson, C.J. Spontaneous subthreshold membrane potential fluctuations and action potential variability of rat corticostriatal and striatal neurons in vivo. *J. Physiol.*, 77, 1697-1715 (1997).
- Stern, E.A., Jaeger, D., & Wilson, C.J. Membrane potential synchrony of simultaneously recorded striatal spiny neurons in vivo. *Nature*, 394, 475-478 (1998).
- Suner, S., Fellows, M.R., Vargas-Irwin, C., Nakata, G.K., & Donoghue, J.P. Reliability of signals from a chronically implanted, silicon-based electrode array in non-human primate primary motor cortex. *IEEE TNSRE*, 13, 524-541 (2005).
- Taylor, D.M., Tillery, S.I., & Schwartz, A.B. Direct cortical control of 3D neuroprosthetic devices. *Science*, 296, 1829-1832 (2002).
- Thomson, D. Spectrum estimation and harmonic analysis. *Proc IEEE*, 70, 1055-1096 (1982).
- Tiesinga, P.H.E., Fellous, J.-M., Jose, J.V., & Sejnowski, T.J. Optimal information transfer in synchronized neocortical neurons. *Neurocomputing*, 38-40, 397-402 (2001).
- Truccolo, W.A., Ding, M., Knuth, K.H., Nakamura, R., & Bressler, S.L. Trial-to-trial variability of cortical evoked responses: implications for the analysis of functional connectivity. *Clin. Neurophysiol.*, 113, 206-226 (2002).
- Van der Maelen, C.P. & Kitai, S.T. Intracellular analysis of synaptic potentials in rat neostriatum following stimulation of the cerebral cortex, thalamus, and substantia nigra. *Brain Res. Bull.*, 5, 725-733 (1980).
- van der Meer, M.A.A., & Redish, A.D. Theta phase precession in rat ventral striatum links place and reward information. *J. Neurosci.*, 31, 2843-2854 (2011).
- VanLehn, K. Cognitive skill acquisition. *Ann. Rev. Neurosci.*, 47, 513-539 (1996).
- van Schoonhoven, J., Sparreboom, M., van Zanten, B.G., Scholten, R.J., Mylanus, E.A., Dreschler, W.A., Grolman, W., & Maat, B. The effectiveness of bilateral cochlear implants for severe-to-profound deafness in adults: a systematic review. *Otol Neurotol*, 34, 190-198 (2013).
- Venkatraman, S., Jin, X., Costa, R.M., & Carmena, J.M. Investigating neural correlates of behavior in freely behaving rodents using inertial sensors. *J. Neurophys.*, 104, 569-575 (2010).
- von Stein, A., Chiang, C., & Konig, P. Top-down processing mediated by interareal synchronization. *PNAS*, 97, 14748-14753 (2000).

- Wang, H.-P., Spencer, D., Fellous, J.-M., & Sejnowski, T.J. Synchrony of thalamocortical inputs maximizes cortical reliability. *Science*, 328, 106-109 (2010).
- Womelsdorf, T., Schoffelen, J.-M., Oostenveld, R., Singer, W., Desimone, R., Engel, A.K., & Fries, P. Modulation of neuronal interactions through neuronal synchronization. *Science*, 316, 1609-1612 (2007).
- Xu, T., Yu, X., Perlik, A.J., Tobin, W.F., Zweig, J.A., Tennant, K., Jones, T., & Zuo, Y. Rapid formation and selective stabilization of synapses for enduring motor memories. *Nature*, 462, 915-919 (2009).
- Yin, H.H., Knowlton, B.J., & Balleine, B.W. Inactivation of dorsolateral striatum enhances sensitivity to changes in the action-outcome contingency in instrumental conditioning. *Behav. Brain Res.*, 166, 189-196 (2006).
- Yin, H.H. et al. Dynamic reorganization of striatal circuits during the acquisition and consolidation of a skill. *Nature Neurosci.*, 12, 333-341 (2009).
- Yokota, T., Saito, Y., & Miyatake, T. Conduction slowing without conduction block of compound muscle and nerve action potentials due to sodium channel block. *J. Neurol. Sci.*, 124, 220-224 (1994).
- Zeitler, M., Fries, P., & Gielen, S. Assessing neuronal coherence with single-unit, multi-unit, and local field potentials. *Neural Computation*, 18, 2256-2281 (2006).
- Zeitler, M., Fries, P., & Gielen, S. Biased competition through variations in amplitude of gamma oscillations. *J. Comput. Neurosci.*, 25, 89-107 (2008).
- Zervakis, M., Michalopoulos, K., Iordanidou, V., & Sakkalis, V. Intertrial coherence and causal interaction among independent EEG components. *J. Neurosci. Methods*, 197, 302-314 (2011).
- Zhong, Y. & Bellamkonda, R.V. Dexamethasone-coated neural probes elicit attenuated inflammatory response and neuronal loss compared to uncoated neural probes. *Brain Res.*, 1148, 15-27 (2007).