

# UC Berkeley

## UC Berkeley Electronic Theses and Dissertations

### Title

Neuroprosthetic Skill Learning and Stimulation-Based Neuromodulation

### Permalink

<https://escholarship.org/uc/item/15s4r0f5>

### Author

Zippi, Ellen Luisa

### Publication Date

2022

Peer reviewed|Thesis/dissertation

Neuroprosthetic Skill Learning and Stimulation-Based Neuromodulation

by

Ellen Luisa Zippi

A dissertation submitted in partial satisfaction of the

requirements for the degree of

Doctor of Philosophy

in

Neuroscience

in the

Graduate Division

of the

University of California, Berkeley

Committee in charge:

Professor Jose M. Carmena, Co-Chair

Professor Joni Wallis, Co-Chair

Professor Robert T. Knight

Professor Richard Ivry

Summer 2022

© 2022 Copyright, Ellen Luisa Zippi

All Rights Reserved

## Abstract

### Neuroprosthetic Skill Learning and Stimulation-Based Neuromodulation

by

Ellen L. Zippi

Doctor of Philosophy in Neuroscience

University of California, Berkeley

Professor Jose M. Carmena, Co-Chair

Professor Joni Wallis, Co-Chair

Brain-machine interface (BMI) systems hold great promise for improving quality of life in patients with a number of motor and cognitive disabilities, but significant advancements must be made before these systems are clinically viable. Developing neuroprosthetic devices and adaptive deep brain stimulators that are neurobiologically-informed is essential for increasing the usability and scalability of these systems. To do so, we must first develop an understanding of how the brain interacts with these systems. Previous work has shown that learning to control a BMI is associated with neural adaptations that depend on large-scale networks across multiple brain regions. In this work, we aim to develop a better understanding of this large-scale learning process by analyzing how motor cortical population dynamics differ in groups of neurons used to directly control a BMI and the remaining recorded population. Then, we investigate the role of corticostriatal circuits in BMI control by simultaneously recording from dorsolateral prefrontal cortex, the caudate nucleus of the striatum, and motor cortex as a nonhuman primate controls a motor cortical BMI. Finally, we assess how high-frequency microstimulation administered in the striatum changes functional decision-making behavior and underlying neural representations of value in the caudate and anterior cingulate cortex with the goal of developing a better understanding of how stimulation-based therapies can be used for neuropsychiatric disorders. Altogether, this work aims to improve our understanding of the neurobiology of BMIs with the goal of developing neurobiologically informed neuroprosthetic devices.

## Acknowledgements

First and foremost, I'd like to thank my advisor, Jose Carmena. Anyone who knows Jose knows that he has an infectious enthusiasm that always leaves you feeling excited. He gave me remarkable freedom to shape my own research projects and pursue my interests, which was an important experience for my development as a researcher.

In addition, I'd like to express my deepest appreciation to Joni Wallis for co-chairing my thesis committee and for sharing her knowledge and experience with me over the past several years. Joni's guidance and assistance with my thesis work was invaluable.

Many additional faculty members at UC Berkeley played an important role in my graduate career. I'd like to thank Bob Knight and Richard Ivry for serving on my thesis committee and providing valuable feedback and support. I'm also greatly appreciative of my Qualifying Exam Committee: Michael Yartsev, David Foster, Na Ji, and Frederic Theunissen. Their input helped guide and shape my thesis projects.

I've been fortunate to work with many incredibly talented colleagues and I owe several people sincere thanks for their direct contributions to the work presented in this thesis. Albert You and I collaborated on work investigating the cortical dynamics of subpopulations of neurons with and without direct input to a BMI decoder (Chapter 2). Gabby Shvartsman played an instrumental role in the data collection and analysis of work investigating the role of prefrontal cortex and caudate in BMI control (Chapter 3). Finally, Samantha Santacruz and I collaborated on work investigating the effect of caudate microstimulation on neural representations of value and error (Chapter 4).

In addition to these scientific contributions, I'd also like to thank my colleagues in the Carmena lab for their support and kindness. First, I'd like to acknowledge Samantha. Samantha took me under her wing as soon as I arrived at Berkeley. I feel very fortunate to have had the opportunity to learn from her. It's been inspiring to watch her start her own lab and continue teaching others.

I'd also like to thank Nuria, who has a special talent for providing the most thoughtful and insightful feedback on projects. She also consistently brightens any room with her constant encouragement and is an absolute pleasure to work with. I also need to thank her for inspiring me to adopt my dog, Luna!

I also want to acknowledge one of my closest friends in the lab, Albert. He is incredibly fun to work with (and bike with, hike with, play with his chinchilla with, etc.). I will always remember our snack breaks and late-night spike sorting sessions fondly.

Our lab manager for most of my time in the lab, Maki, is unbelievably thoughtful and considerate. I'm grateful for her assistance with surgeries, recordings, and animal care as well as her grounding personality. NHP research can be overwhelming, but Maki always helped to make everything feel smooth and under control.

The work in this thesis wouldn't have been possible without the best research assistant one could ever hope for – Gabby. Gabby joined the lab as an ambitious undergraduate student who captained a tennis team, led an art & mind club, took a full course load, and still managed to spend every day in the basement with me learning how to perform NHP neurophysiology experiments. In addition to her help with data collection and analysis, I'd also like to express my gratitude towards her company during our long basement recording sessions. I'll always be thankful for our post recording yoga sessions!

Another colleague I am thankful for is Paul. Paul stuck it out to the end of the Carmena lab with me and is a great friend who made coming to the lab more fun. He was always ready to help me troubleshoot new problems or to grab a beer after work.

I'd also like to thank Emanuele for his friendship and advice, David for many fun conversations, and Tanner for being one of the most welcoming and thoughtful people I've ever met.

I look up to all of you as scientists and I am grateful for our shared experiences in the lab together. There were always people I could count on in the Carmena lab and I am incredibly thankful for that.

Major thanks must also go to the OLAC veterinary and husbandry staff. So much work goes into NHP research, and it wouldn't be possible without them. In particular, I'd like to thank Christie, Kelly, Jenn, Monica, and Ashley for their help with the work included in this thesis. I'd also like to acknowledge Joni, Samantha, and Charlie for performing the surgeries.

I'm also grateful to my cohort members, Hayley, Dariya, Sylvia, Molly, Lily, Celia, Ryan, and Christiane, for their friendship and moral support throughout my graduate career. Additionally, I was fortunate to have a strong, encouraging group of friends that I met in graduate school. I'd like to thank Amanda, Josel, Libby, and Juan for always being there to celebrate accomplishments or commiserate over setbacks. Your friendship has been instrumental in keeping my spirits and motivation alive while completing a PhD during a pandemic. I'd also like to thank my fiancé, Luis, who has been making me smile every day for the past seven years. And finally, I'd like to thank my parents, Pam and Pierre, and my brother, Josef, for their encouragement and support.

# Publications and presentations related to this work

## Journal articles and peer-reviewed conference proceedings

Zippi, E.L.\* , You, A.K.\* , Ganguly, K., & Carmena, J.M. “Selective modulation of population dynamics during neuroprosthetic skill learning”, *in Review at Scientific Reports*

You, A.K.\* , Zippi, E.L.\* , & Carmena, J.M. (2019), “Large-Scale Neural Consolidation in Brain-Machine Interface Learning.” *2019 9th International IEEE/EMBS Conference on Neural Engineering*

Santacruz, S.R., Zippi, E.L., & Carmena, J.M. (2018), “Modulation of Brain-Based Value Signals Using Electrical Stimulation.” *IEEE Engineering in Medicine and Biology Society*

## Presentations

Zippi E.L.† , & Carmena J.M. Neuroprosthetic skill learning and stimulation-based neuromodulation. Oral presentation, Helen Wills Neuroscience Thesis Symposium, Berkeley, CA, May 2022.

Zippi E.L.† , & Carmena J.M. Large-Scale Modification of Cortical & Subcortical Networks During Neuroprosthetic Skill Learning. Oral presentation, Helen Wills Neuroscience Institute Brain Lunch, Berkeley, CA, March 2021.

Santacruz S.R., Zippi E.L.† , & Carmena J.M. Modulation of value and error signals using Caudate electrical stimulation. Poster presentation, Center for Engineering and Prostheses Retreat, Berkeley, CA, December 2019.

Santacruz S.R., Zippi E.L.† , & Carmena J.M. Neurotherapeutic effects of closed-loop stimulation in a nonhuman primate model of anxiety. Poster presentation, Society for Neuroscience Annual Meeting, Chicago, IL, October 2019.

Zippi, E.L.† , Santacruz S.R., & Carmena J.M. Neurotherapeutic effects of closed-loop stimulation in a nonhuman primate model of anxiety. Oral presentation, Berkeley Neuroscience Retreat, Tahoe City, CA, September 2019.

You, A, Zippi, E.L.† , & Carmena, J.M. Large-Scale Neural Consolidation in Brain-Machine Interface Learning. Poster presentation, IEEE NER. San Francisco, CA, April 2019.

[\*authors contributed equally, †presenting author]

# Contents

<b>Chapter 1: Introduction</b> .....	<b>2</b>
1.1 Brain Machine Interface (BMI) Systems .....	2
1.2 Clinical Applications of BMIs .....	4
1.2.1 External Devices .....	5
1.2.2 Internal Devices .....	5
1.3 BMIs as Scientific Tools .....	6
1.4 Remaining Challenges .....	7
1.5 Chapter Previews .....	8
<b>Chapter 2: Selective Modulation of Motor Cortical Population Dynamics in Groups of Neurons Causally Linked to Behavior</b> .....	<b>9</b>
2.1 Introduction .....	9
2.2 Methods .....	11
2.2.1 Animal Subjects .....	11
2.2.2 Electrophysiology .....	12
2.2.3 Manual Control Training Before BMI .....	12
2.2.4 BMI Tasks .....	12
2.2.5 Preprocessing Pipeline .....	13
2.2.6 Factor Analysis .....	13
2.2.7 Quantification & Statistical Analyses .....	15
2.3 Results .....	15
2.3.1 Increases in Neural Variability with Learning .....	19
2.3.2 Increases in Coordinated Neural Activity with Learning .....	22
2.3.3 Changes in Neural Covariance with Learning .....	25
2.3.4 Changes in Neural Variability with a Novel Decoder .....	27
2.4 Discussion .....	28
<b>Chapter 3: Simultaneous Large-Scale Recordings During Neuroprosthetic Control in Nonhuman Primates</b> .....	<b>33</b>
3.1 Introduction .....	33
3.2 Methods .....	35
3.2.1 Segmenting MRIs for Anatomical Models .....	35
3.2.2 Large-Scale Semi-Chronic Microdrive System .....	35
3.2.3 Implanting the Chamber and Microdrive .....	35
3.2.4 Impedance Monitoring and Advancing Electrodes .....	36

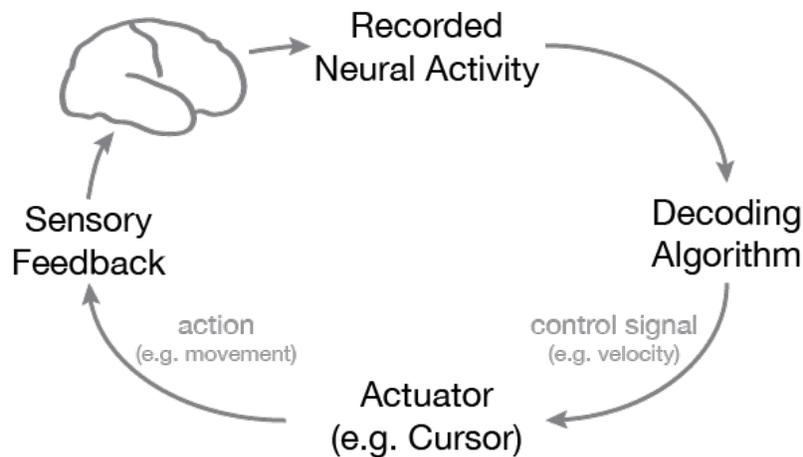
3.2.5 Intracortical Recording.....	37
3.2.6 Center-Out Task.....	37
3.2.7 Brain-Machine Interface.....	38
3.2.8 Neural Data Analysis.....	39
3.2.9 Quantification & Statistical Analyses.....	41
3.3 Results.....	41
3.3.1 Online BMI Performance.....	43
3.3.2 Control Type Classification with Power Features.....	44
3.3.3 Differences in Directed Functional Connectivity.....	50
3.4 Discussion.....	55
<b>Chapter 4: Modulation of Value Encoding in Caudate and Anterior Cingulate Cortex through Caudate Microstimulation.....</b>	<b>60</b>
4.1 Introduction.....	60
4.2 Methods.....	61
4.2.1 Surgery.....	61
4.2.2 Stimulation.....	62
4.2.3 Behavioral Task.....	62
4.2.4 Reinforcement Learning Models.....	63
4.2.5 Neural Correlates.....	63
4.2.6 Quantification & Statistical Analyses.....	64
4.3 Results.....	64
4.3.1 Caudate Microstimulation Biases Target Selection.....	66
4.3.2 Value Encoding in Caudate and ACC.....	67
4.3.3 Recruitment of Value Encoding Units with Caudate Microstimulation.....	69
4.4 Discussion.....	72
<b>Chapter 5: Conclusions and Open Questions.....</b>	<b>74</b>
5.1 Summary of Contributions.....	74
5.2. Open Questions and Future Directions.....	75
5.2.1 Neuroprosthetic Skill Learning.....	76
5.2.2 Stimulation-Based Neuromodulation.....	77
5.3 Conclusions.....	79
<b>Chapter 6: Bibliography.....</b>	<b>80</b>

# Chapter 1: Introduction

Advances in neural recording technology and machine learning over the past several decades have enabled a direct connection of brains with machines. These brain-machine interfaces (BMIs) allow for the real-time transformation of neural activity into control signals for external devices, such as computer cursors or robotic arms, or for internal devices, such as deep brain stimulators (DBS). These systems create new ways for users to engage with their environments, enabling numerous clinical and scientific applications.

## 1.1 Brain Machine Interface (BMI) Systems

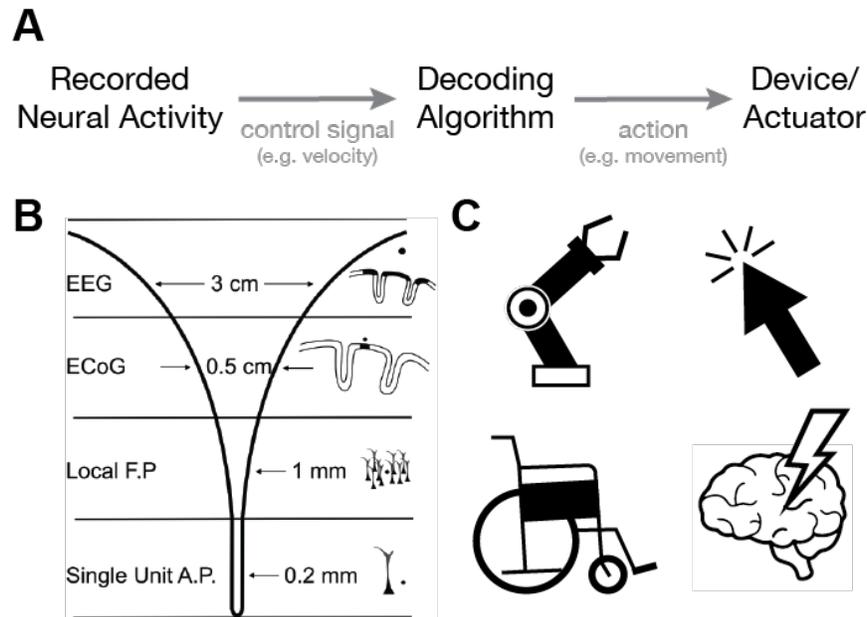
A typical BMI consists of three main components: the neural activity signal, the decoding algorithm that transforms the neural activity into a control signal, and the actuator, or machine, that is being controlled (Figure 1.1). Recorded neural activity is input to the online decoding algorithm and converted into a control signal. This control signal then sends instructions to a device. In many systems, there is an additional feedback component in which the user receives sensory feedback, such as visual or auditory feedback, to aid in learning to control the BMI.



**Figure 1.1 Schematic representation of the general architecture of a closed-loop BMI system.** Recorded neural activity is mapped into control signals for an actuator via a decoding algorithm. These systems typically provide sensory feedback (e.g., visual feedback) to the user, creating a closed-loop system.

There are a vast number of manifestations of BMI systems. For instance, many types of neural signals have been used to control a BMI, including electroencephalography (EEG), electrocorticography (ECoG), local-field potentials (LFPs), and single- and multi-unit action potentials (Figure 1.2B). Noninvasive BMIs, such as those using EEG

signals, offer poor spatial and temporal resolution relative to invasive recording techniques, limiting the complexity of control signals that these recordings can support. On the other hand, invasive techniques, which offer superior spatial and temporal resolution, require surgery before they can be used and are susceptible to signal degradation over time. While all of these recording techniques have been demonstrated to provide effective control of a BMI, array recordings, including LFP and single- and multi-unit action potentials, continue to be the most commonly used for BMIs in both a clinical and research setting.



**Figure 1.2 Examples of different manifestations of BMI systems.** (A) Schematic representation of the general architecture of a BMI system. (B) Figure from Schwartz et al., 2006. Representation of different neural recording modalities and comparisons of their spatial scale of recording. (C) Representation of possible brain-controlled actuators, including robotic arm, computer cursor, wheelchair, and brain stimulator.

A variety of decoding algorithms have also been used to create a mapping between the selected neural activity signal and the BMI actuator. Depending on the goal of the BMI, this decoder may be used to estimate a subject’s motor intent for translation into a desired action or to detect abnormal brain activity for translation into stimulation parameters for adaptive DBS. These algorithms map the recorded neural activity to the desired control signal via regression or classification. Within motor BMIs, decoders can either be biomimetic or non-biomimetic (Ganguly & Carmena, 2010; Shenoy & Carmena, 2014). Biomimetic decoders aim to mimic the natural biological mapping between neural activity and movement as closely as possible. Non-biomimetic

decoders require subjects to learn arbitrary mappings between neural activity and movement, requiring some amount of neural adaptation for proficient control. While biomimetic decoders may seem superior for clinical use cases, they can be more difficult to implement due to our current lack of understanding of the brain's natural mapping between neural activity and movement. Furthermore, the neural adaptation that occurs when learning to control a non-biomimetic decoder can be useful when using BMIs as a scientific tool to study learning. Additionally, neural adaptation during BMI learning may enable more robust BMIs (Orsborn & Carmena, 2013). Finally, decoding algorithms can be fixed or adaptive. While fixed decoders remain constant, closed-loop decoder adaptation (CLDA) is an approach to improve BMI performance by modifying the decoder during operation of the BMI (Dangi et al., 2013; Gilja et al., 2012; Orsborn et al., 2012, 2014). Using knowledge of task goals and assumptions about a user's intent (e.g., assuming a user intends to reach in a straight line towards a presented target at a known location), supervised algorithms can update the decoder to better reflect the user's intentions.

Finally, a variety of devices have been successfully controlled in real time using combinations of these neural activity signals and decoding algorithms (Figure 1.2C). Overall these BMI systems can be used to transform neural activity into a variety of control signals to be used for a myriad of devices including virtual objects (Carmena et al., 2003; Degenhart et al., 2020; Ganguly & Carmena, 2009; Gilja et al., 2012; Hochberg et al., 2006; Jarosiewicz et al., 2008; Leuthardt et al., 2004; Olsen et al., 2021; Schalk et al., 2008; Serruya et al., 2002; Silversmith et al., 2021; Taylor et al., 2002; Wander et al., 2013; Willett et al., 2021; Wolpaw & McFarland, 2004), robotic devices (Carmena et al., 2003; Chapin et al., 1999; Collinger et al., 2013; Flesher et al., 2021; Hochberg et al., 2012; Taylor et al., 2002; Velliste et al., 2008), wheelchairs (Carlson & del R. Millan, 2013; Iturrate et al., 2009; K.-T. Kim et al., 2018; Millán, 2010), muscle stimulators (Ajiboye et al., 2017; Collinger et al., 2018; Ethier et al., 2012; Moritz et al., 2008), and deep brain stimulators (Hoang et al., 2017; Little et al., 2013, 2016; Rosa et al., 2017; Swann et al., 2018).

## 1.2 Clinical Applications of BMIs

Neurological disorders including stroke, dementia, mood disorders, brain injuries, Parkinson's disease, and more affect up to one billion people worldwide (Feigin et al., 2019). Together, they make up the largest global disease burden as measured by disability-adjusted life years, a time-based measure that accounts for years of life lost due to premature mortality and adjusted years of life lost due to time lived in a state of less than full health (Feigin et al., 2019, 2020). Furthermore, this global burden has continued to increase in recent years as the aging population continues to grow

worldwide (Carroll, 2019). Thus, neurological disorders pose an important public health challenge that affects people across the globe, irrespective of age, gender, or income.

### *1.2.1 External Devices*

BMIs have useful clinical applications for a number of neurological disorders. Closed-loop BMIs can be used to control neuroprosthetic systems in which users with motor control deficits can manipulate robotic arms or computer cursors, enabling them to interact with their environments in new ways. These devices have allowed patients to perform reach-to-grasp motions with a robotic arm to feed themselves (Collinger et al., 2013; Flesher et al., 2021; Hochberg et al., 2012), to surf the internet or operate point-and-click computer applications using a brain-controlled computer cursor (Kim et al., 2011; Nuyujukian et al., 2018), and to control wheelchair movements (Iturrate et al., 2009; Kim et al., 2018). Additionally, these closed-loop systems, in combination with predictive text completion, have enabled locked-in users suffering from amyotrophic lateral sclerosis to operate speller or typing systems (Jarosiewicz et al., 2015; Pandarinath et al., 2017). Other systems allow users to communicate by decoding neural signals to synthesize speech or handwriting (Anumanchipalli et al., 2019; Willett et al., 2021; Wilson et al., 2020). These applications are important examples of ways in which BMIs can be used to improve the quality of life for those suffering from various neurological disorders.

### *1.2.2 Internal Devices*

In addition to controlling external devices, BMIs can be used to control internal devices such as deep brain stimulators. Traditional DBS is a device-based therapy that sends electrical signals to specific regions of the brain to treat certain neurological diseases. For example, stimulation of the subthalamic nucleus or the globus pallidus internus is a common treatment for Parkinson's Disease in patients whose symptoms are not adequately controlled with medication (Benabid, 2003; Benabid et al., 2009; Deuschl et al., 2006; Rodriguez-Oroz et al., 2005). In these traditional DBS systems, neurologists manually adjust stimulation parameters every 3–12 months after implantation and a consistent electrical current is delivered between programming sessions. Recently, closed-loop stimulation has been shown to result in fewer side effects, equal or better effectiveness, and less power consumption (Parastarfeizabadi & Kouzani, 2017). These closed-loop DBS systems are referred to as adaptive DBS. In these systems, stimulation is delivered when the brain is detected to be in an abnormal state or stimulation parameters are dynamically adjusted based on variations in a recorded biosignal over time. While these devices have only been approved by the

United States Food and Drug Administration for Epilepsy and Parkinson's Diseases (Hoang et al., 2017; Little et al., 2013; Rosin et al., 2011; Sun et al., 2008; Wu et al., 2021), recent work has suggested that closed-loop systems may improve the efficacy of stimulation-based therapies for mood disorders as well (Scangos et al., 2021; Shanechi, 2019; Widge et al., 2018). Identifying reliable biomarkers of disease and understanding how stimulation affects the neurophysiology of brain circuits involved in neurological disorders will further the possible clinical applications for these BMIs.

### 1.3 BMIs as Scientific Tools

In addition to offering promising assistive technologies and clinical therapies for a number of neurological disorders, BMIs are also a valuable scientific tool. Previous work has demonstrated that motor cortical BMIs are not simply an artificial replacement for dysfunctional neural circuits, but rather a tool that users learn to control through the modulation of neural signals (Orsborn & Carmena, 2013). Thus, BMIs provide a unique means of studying learning. Because the neural activity controlling the BMI output is selected by the experimenter and the mapping between this activity and control is known, BMIs provide a well-defined, simplified system without simplifying the actual learning process (Golub et al., 2016; Moxon & Foffani, 2015; Orsborn & Carmena, 2013; Orsborn & Pesaran, 2017). This is in contrast to the natural motor system, in which there are more motor-related neurons than can be recorded and the mapping between their activity and behavioral output is unknown.

In the case of motor BMI systems, the process of learning proficient control has been well studied. A key feature of this neuroprosthetic skill learning is that users can learn to volitionally modulate neural activity (Fetz, 2007; Green & Kalaska, 2011). To proficiently control a BMI, users must learn a neural representation of the mapping between neural activity and the actuator that facilitates successful control. Much like the representation of natural motor skills, this neuroprosthetic representation has been shown to be stable across days, readily recallable, and resistant to interference from learning a second mapping (Ganguly & Carmena, 2009). Experimenters can manipulate the mapping between neural activity and output control to study the neuroplasticity underlying learning to control a novel actuator. These manipulations have shed light onto the cognitive processes underlying sensorimotor control (Athalye et al., 2017, 2018; Clancy et al., 2014; Ganguly et al., 2011; Ganguly & Carmena, 2009; Golub et al., 2016, 2018; Jarosiewicz et al., 2008; Koralek et al., 2012, 2013; Neely et al., 2018; Oby et al., 2019; Orsborn et al., 2014; Sadtler et al., 2014).

## 1.4 Remaining Challenges

Whether by translating thought into action for the purpose of movement restoration or assistive communication or by detecting and responding to biomarkers of disease via adaptive deep brain stimulation, BMIs hold a lot of clinical promise. However, this emerging technology is still limited in its usefulness as a clinical tool by both technological and biological challenges.

Technological challenges include improving the biocompatibility and longevity of implants, as well as improving algorithms for decoding neural activity and artifact removal (Rapeaux & Constandinou, 2021). Clinical applications of BMIs typically rely on wired, invasive neural probes such as Utah arrays, which cause damage to brain tissue and can result in a chronic inflammatory response, cell death, and encapsulation of the electrodes by microglia leading to degradation of signal quality (Biran et al., 2005; Polikov et al., 2005; Turner et al., 1999). Furthermore, these wired systems are more susceptible to infection and are bulky and inconvenient to use. Work on miniaturizing these systems, making them wireless, and improving biocompatibility of materials is critical for the advancement of BMIs as a clinical tool. Non-invasive BMI systems, such as those that rely on EEG signals, do not face these same limitations but are limited to recording lower-frequency signals due to the attenuation of voltage signals through the skull (Fahimi Hnazaee et al., 2020). Improving decoding algorithms and artifact detection and removal for both invasive and non-invasive neural signals will be important for improving BMIs as a clinical tool.

Biological challenges facing BMIs include improving our understanding of how information about intention, mental state, and disease is encoded in the brain, as well as how the brain interacts with and learns to control novel actuators (Kawala-Sterniuk et al., 2021). State-of-the-art motor BMIs currently operate using relatively few degrees-of-freedom relative to the natural motor system and can be slow and unintuitive to use. It can take users days of practice to master control of a motor BMI and this process of learning has been reported to be frustrating to some. Several studies have demonstrated that BMIs create new systems that engage learning and adaptation (Carmena et al., 2003; Fetz, 2007; Taylor et al., 2002). Understanding these systems may be important for developing successful, intuitive neuroprosthetic devices. Additionally, advances in understanding how information about action intention and language is encoded in the brain is important for developing BMIs that are more natural and intuitive to use and identifying reliable biomarkers of mental state and disease is critical for improving BMIs for adaptive DBS.

The focus of this thesis will be on questions pertaining to the second of these two challenges and aims to improve our understanding of how the brain learns to control a BMI and how brain stimulation can affect the neural encoding of specific information. By furthering our understanding of the neurobiology of BMIs, we hope to develop more neurobiologically-informed neuroprosthetic devices that are easier to learn and scale.

## 1.5 Chapter Previews

In this thesis, we will cover work exploring neuroprosthetic skill learning and control as well as stimulation-based neuromodulation. Specifically, we aim to further our understanding of the neurobiology of BMIs by gaining insight into how distributed networks in the brain learn to control a virtual cursor and how microstimulation affects neural representations of value.

In Chapter 2, we compare the changes in coordinated spatiotemporal activity that occur in subpopulations of neurons whose activity was directly used as input to a BMI decoder and the remaining recorded motor cortical subpopulation. This project provides insight into how motor cortex explores population-level activity to produce coordinated neural dynamics in populations of neurons relevant for control, shedding light onto how novel motor skills are learned.

In Chapter 3, we present a novel technique for simultaneous recording of cortical and subcortical brain areas during motor BMI control. We identified task-relevant changes in spectral power in dorsolateral prefrontal cortex, caudate, and motor cortex as well as changes in directed functional connectivity between these regions. This work is an important step in understanding the large-scale networks involved in reinforcing neural patterns of activity that are necessary for proficient BMI control.

Finally, in Chapter 4 we explore stimulation-based neuromodulation during a decision-making task. We present work investigating how caudate microstimulation capable of biasing choice behavior affects underlying neural representations of task-relevant stimulus values in the caudate and anterior cingulate cortex. In this study, we aim to develop a better understanding of how stimulation-based therapies may be used to regulate valuation of choices in neuropsychiatric patients.

## Chapter 2: Selective Modulation of Motor Cortical Population Dynamics in Groups of Neurons Causally Linked to Behavior

Brain-machine interfaces (BMIs) provide a framework for studying how cortical population dynamics evolve over learning in a task in which the mapping between neural activity and behavior is precisely defined. Learning to control a BMI is associated with the emergence of coordinated neural dynamics in populations of neurons whose activity serves as direct input to the BMI decoder (direct subpopulation). While previous work shows differential modification of firing rate modulation in this population relative to a population whose activity was not directly input to the BMI decoder (indirect subpopulation), little is known about how learning-related changes in cortical population dynamics within these groups compare. To investigate this, we monitored both direct and indirect subpopulations as two macaque monkeys learned to control a BMI. We found that while the combined population increased coordinated neural dynamics, this increase in coordination was primarily driven by changes in the direct subpopulation. These findings suggest that motor cortex refines cortical dynamics by increasing neural variance throughout the entire population during learning, with a more pronounced coordination of firing activity in subpopulations that are causally linked to behavior. This work was done in collaboration with Albert K. You, Karunesh Ganguly, and Jose M. Carmena. It is currently in review at Scientific Reports and an earlier version of the manuscript is available as a preprint on BioRxiv (Zippi, You et al., 2021).

### 2.1 Introduction

Learned behaviors are reinforced through mechanisms involving both cortical and subcortical structures (Donchin et al., 2012; Krakauer et al., 2004; Sing & Smith, 2010; Sutton & Barto, 1998). Just as behavioral actions are reinforced, so is the cortical population activity required to efficiently produce these actions (Athalye et al., 2017; Costa, 2011; Tumer & Brainard, 2007). Studying mechanisms of cortical reinforcement underlying behavioral reinforcement can be challenging as the exact neural population controlling the desired behavior is unknown. Early studies of biofeedback demonstrated that activity in motor cortex can be reinforced and volitionally controlled using reward and sensory feedback of the firing rate (Fetz, 1969; Fetz & Baker, 1973). Later, initial research on brain-machine interfaces (BMIs) showed that subjects could learn to control external devices (e.g. computer cursors or robotic arms) by learning to modulate the activity of a population of neurons and that the neural encoding of these prosthetic movements changed over time and decreased in variability (Carmena et al., 2003; Ganguly & Carmena, 2009; Musallam et al., 2004; Taylor et al., 2002;

Zacksenhouse et al., 2007). These BMIs allow for precisely defined mappings between recorded neural activity and behavior (Carmena et al., 2003; Shenoy & Carmena, 2014; Taylor et al., 2002). Studies leveraging BMIs to study learning-related changes in cortical activity have demonstrated that neuroprosthetic skill learning can require the production of novel cortical dynamics to obtain skillful control (Ganguly et al., 2011; Ganguly & Carmena, 2009; Orsborn et al., 2014).

While classical approaches examine individual neurons to understand fundamentally how motor cortical activity is reinforced, more recent methods looking at population-level activity have uncovered how dynamic processes may govern movement planning and execution (Ames & Churchland, 2019; Churchland et al., 2010, 2012; Elsayed et al., 2016; Heming et al., 2019; Kao et al., 2015; Kaufman et al., 2014; Pandarinath et al., 2018; Russo et al., 2018; Shenoy et al., 2013; Suresh et al., 2020; Zhou et al., 2019), as well as learning (Golub et al., 2018; Hennig et al., 2021; Oby et al., 2019; Perich et al., 2018; Sadtler et al., 2014; X. Sun et al., 2022; Vyas et al., 2018; Vyas, Golub, et al., 2020; Vyas, O'Shea, et al., 2020; Zhou et al., 2019). Population-level activity is often characterized by low-dimensional dynamics that capture patterns of co-activation across neurons within a population (Gallego et al., 2017). These population-level dynamics arise from input connectivity and within-population connectivity. Two parallel mechanisms have been proposed to reinforce specific cortical population dynamics; fast reinforcement of dynamics that naturally produce a desired behavior and slower reinforcement that refines them to result in more reliable production of neural activity patterns (Athalye et al., 2020; Dayan & Cohen, 2011).

Previous studies have shown that neural populations are constrained to generate activity patterns within a pre-existing covariance structure within short timescales (Golub et al., 2018; Perich et al., 2018; Sadtler et al., 2014; X. Sun et al., 2022), suggesting that it is faster to learn to control and repurpose pre-existing cortical population dynamics than it is to modify them. When decoder perturbations that change the behavioral output associated with specific neural activity were introduced after subjects had already achieved proficient control using a BMI there was an immediate deficit in performance. However, over training, subjects were able to recover performance of cursor control and furthermore, experienced a washout when the perturbation was removed, but only when the perturbation did not require alteration of the natural covariance pattern among the recorded neurons (Sadtler et al., 2014). Other work has demonstrated that animals can learn to control BMIs that require neural patterns outside of the pre-existing covariance structure over the course of multiple days (Athalye et al., 2017; Oby et al., 2019). This eventual modification of cortical dynamics suggests that learning novel skills requires the production of new underlying population activity that develops over longer timescales.

With motor cortical BMIs, a small subset of neurons of all possible neurons in motor cortex is selected to use as input to the decoder (direct neurons). These neurons exist within a large network of other motor cortical neurons (indirect neurons). While the selection of direct neurons from all recorded units in our experiments was arbitrary and there was initially no functional difference between the those selected to be direct and those that are not, previous work has shown that differences in the neural activity of these two groups emerge with learning (Clancy et al., 2014; Ganguly et al., 2011; Gulati et al., 2014; Koralek et al., 2012, 2013; Neely et al., 2018; So et al., 2012). For example, it has been demonstrated that the task-relevant modulation of indirect neurons gradually reduces relative to direct neurons over learning (Ganguly et al., 2011). Additionally, it has been shown in rodents that coherence develops between dorsal striatum and direct neurons, but not indirect neurons (Koralek et al., 2012, 2013; Neely et al., 2018). In a study using 2-photon calcium imaging to record neural activity, mice initially modulated activity of both direct and indirect neurons, but predominantly modulated direct activity after learning (Clancy et al., 2014). Thus, it is likely that the initial cortical dynamics that produce desirable outcomes involve both direct neurons and the surrounding cortical network. Over time, as these cortical dynamics are refined, they may adapt to exclude neurons that do not directly drive behavior.

If this hypothesis is true, we expect differences in how cortical population-level dynamics within direct and indirect subpopulations change over time as well. As cortical dynamics are modified for more efficient control, the direct subpopulation would be expected to undergo further modification than the indirect subpopulation as additional modifications to indirect activity would not directly result in desirable outcomes. Here, we investigate this idea by studying recorded ensembles of motor cortical neurons while only a subset was assigned to have a causal role during BMI control and characterize the differential changes in coordinated neural dynamics between direct and indirect subpopulations.

## 2.2 Methods

### 2.2.1 *Animal Subjects*

All procedures were conducted in compliance with the NIH Guide for the Care and Use of Laboratory Animals and were approved by the University of California at Berkeley Institutional Animal Care and Use Committee.

Two adult male rhesus monkeys (*Macaca mulatta*) were chronically implanted in the brain with arrays of 64 microelectrodes (Innovative Neurophysiology, Durham NC) (Ganguly & Carmena, 2009). Monkey P was implanted in the left hemisphere in the

arm area of both primary motor cortex (M1) and dorsal premotor cortex (PMd), and in the right hemisphere in the arm area of M1, with a total of 192 microwires across three implants. Monkey R was implanted bilaterally in the arm area of M1 and PMd (256 microwires across four implants). Only activity from M1 was included in the direct ensembles (Monkey P: right M1; Monkey R: left M1) and only activity from the same hemisphere was included in the indirect ensemble. Array implants were targeted for pyramidal tract neurons in layer 5. Localization of target areas was performed using stereotactic coordinates from a neuroanatomical atlas of the rhesus brain (Paxinos et al., 2000).

### *2.2.2 Electrophysiology*

Neural activity was recorded using the MAP system (Plexon, Dallas TX). Stable units, to be part of the direct ensemble, were selected based on waveform shape, amplitude, relationship to other units on the same channel, interspike interval distribution, and the presence of an absolute refractory period. Only units from primary motor cortex were used which had a clearly identified waveform with signal-to-noise ratio of at least 4:1. Activity was sorted prior to recording sessions using an online spike-sorting application (Sort Client; Plexon). Stability of waveforms was confirmed by analyzing the stability of PCA projections over days (Wavetracker; Plexon).

Direct units are defined as the units being used to control the BMI. Indirect units consisted of the remaining recorded units. For analyses including only stable units from the same hemisphere, stability in the indirect ensemble was assessed using pairwise cross-correlograms, autocorrelograms, waveform shapes, and mean firing rates (Fraser & Schwartz, 2012).

### *2.2.3 Manual Control Training Before BMI*

Before starting the BMI learning experiments, subjects were overtrained on the task performed with arm movements using a Kinarm (BKIN Technologies, Kingston ON) exoskeleton which restricted shoulder and elbow movements to the horizontal plane.

### *2.2.4 BMI Tasks*

Data from Ganguly & Carmena, 2009, in which subjects performed a self-initiated, eight-target, center-out reaching task, was analyzed. In these experiments, a cursor on a screen was continuously controlled by neural activity. Subjects self-initiated trials by moving the cursor to a center target. One of the eight peripheral targets was randomly selected each trial. Self-initiated trials consisted of those in which the animal moved the cursor to the center target and held for 250-300 ms. Successful trials required the

animal to move the cursor to the peripheral target within 15 s of initiating the trial and hold the cursor at the target for 250-300 ms. Successful trials resulted in a juice reward; failed trials were repeated. During BMI control, both arms were lightly restrained to restrict arm movement during BMI control. During selected sessions, video and surface electromyogram (EMG) recordings from proximal muscle groups were performed to confirm minimal arm movements occurred during BMI control.

After the initial 19-days of performing the BMI task with a fixed-decoder, Monkey P learned a second decoder over the course of four days. Within each of these four days, Monkey P performed one training block of the new decoder, followed by one training block of the old decoder. Both of these decoders were fixed and used the same direct ensemble as input.

### 2.2.5 Preprocessing Pipeline

For all analyses, neural data was binned into 100 ms bins to match the decoder timescale. Additionally, learning was analyzed over “training epochs,” where each epoch consisted of 150 self-initiated trials. Learning took place over the first 15 training epochs (2250 self-initiated trials). We chose to analyze the data across training epochs, rather than days, to eliminate the effect of variable numbers of trials each day. Only the first 15 training epochs (2250 self-initiated trials) were analyzed; we defined early and late learning as the first and last seven of these training epochs, respectively. Monkey P initiated a total of 3589 trials. Monkey R initiated a total of 2357 trials.

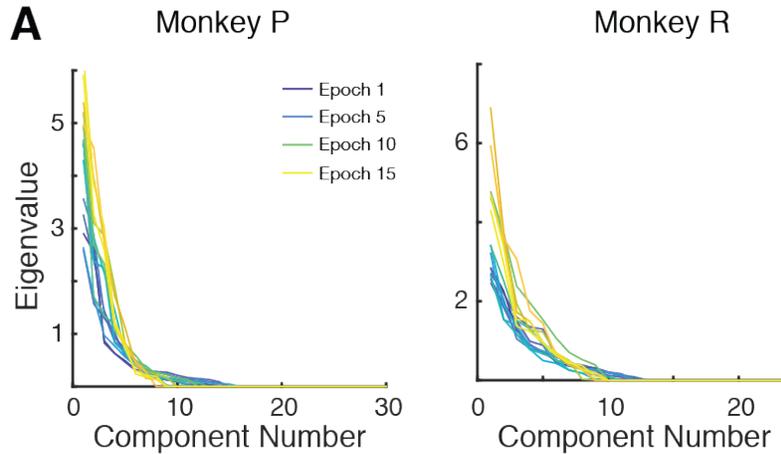
### 2.2.6 Factor Analysis

#### *Shared-Over-Total Variance Ratio*

Factor analysis (FA) was conducted on the neural population for each epoch to observe underlying correlated neural activity. FA decomposes population signals into correlated and uncorrelated components. For a given neuron  $i$ , correlated activity is represented by the shared variance  $\Sigma_{ii}^{shared}$ , and the degree to which the activity was correlated over learning was represented by the ratio of shared-over-total variance of the neural population (SOT). We calculated the SOT ratio according to the methods described in Althaye et al., 2017.

$$SOT = \frac{trace(\Sigma^{shared})}{trace(\Sigma^{total})}$$

Since FA decomposes the neural activity into correlated and uncorrelated components, we can reduce the dimensionality of the neural data by examining how much of the neural variance is captured by the shared (correlated) components. The number of dimensions was selected such that 90% (or more) of the shared variance was captured. A scree plot quantifying the variance captured by each factor for the combined population is shown in Figure 2.1.



**Figure 2.1 Scree plots for direct and indirect combined population.** Factor analysis was performed on neural activity from the combined direct and indirect population for each training epoch individually. A scree plot is shown for each training epoch indicated by color.

### Total Variance

We also considered how the total variance changed between early and late learning. This was the sum of the private and shared variances  $\Sigma_{ii}^{total} = \Sigma_{ii}^{shared} + \Sigma_{ii}^{private}$ .

### Partial Shared-Over-Total Variance Ratio

We quantified the respective contributions of subpopulations to the SOT ratio using the partial shared-over-total variance (pSOT) ratio. Here, we compared the sum of the shared variance for each subpopulation over the total variance for the entire population. A relative measure was used to account for the fact that the direct and indirect ensembles were different sizes. That is,

$$pSOT_{direct} = \frac{\text{trace}(\Sigma_{direct}^{shared})}{\text{trace}(\Sigma_{all}^{private}) + \text{trace}(\Sigma_{all}^{shared})}$$

$$pSOT_{indirect} = \frac{\text{trace}(\Sigma_{indirect}^{shared})}{\text{trace}(\Sigma_{all}^{private}) + \text{trace}(\Sigma_{all}^{shared})}$$

where  $\Sigma_i$  denotes the covariance matrix for the (sub)population  $i$ .

### *Shared Space Alignment*

We used the “shared space alignment” to measure the similarity between the shared variance of two different training epochs. The shared space alignment is the fraction of shared variance from one epoch captured in the shared space of a second epoch and thus ranges from 0 to 1. We calculated shared alignment according to the methods described in Athalye et al., 2017. Given two epochs, A and B, we first compute the projection matrix into Epoch B’s shared space,  $col(U^B)$ . We then project  $\Sigma^{A,shared}$  onto B’s shared space,  $P_{UB}\Sigma^{A,shared}P_{UB}^T$ . Finally, the alignment is calculated,

$$Shared\ Alignment = \frac{trace(P_{UB}\Sigma^{A,shared}P_{UB}^T)}{trace(\Sigma^{A,shared})}$$

### *2.2.7 Quantification & Statistical Analyses*

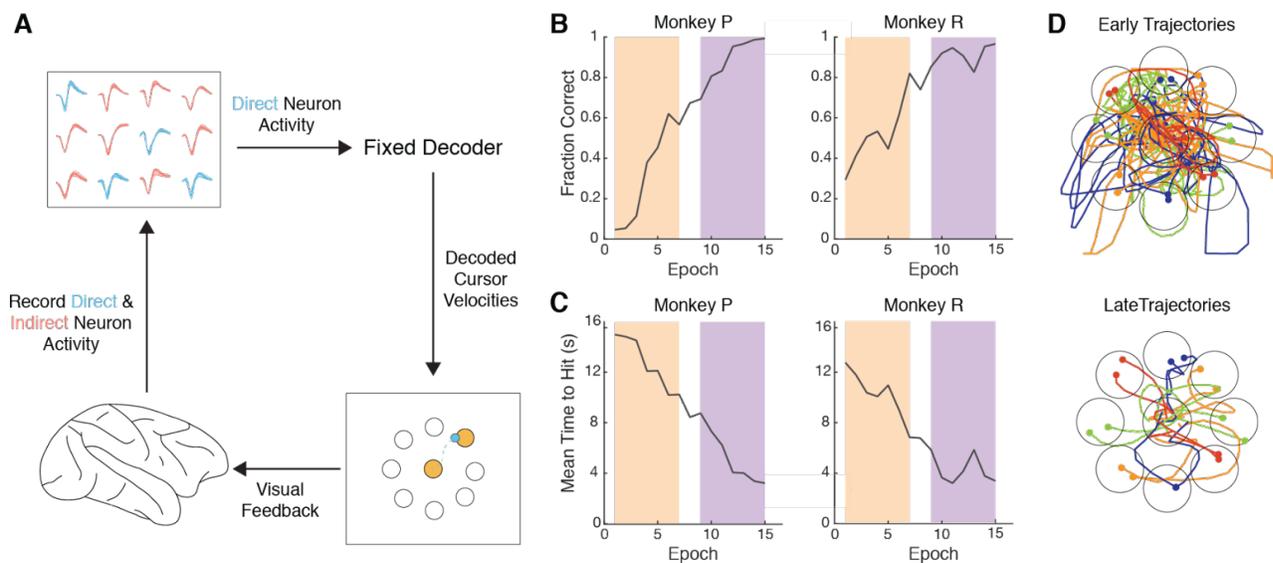
All analyses were performed on each training epoch separately. Trends were analyzed for significance with linear regressions. Epochs 1-7 and Epochs 9-15 were grouped into “early” and “late,” respectively. Epoch 8 was omitted so that there were equal numbers of epochs for both early and late learning. Groupings of early and late epochs were compared using an unpaired two-sample t-test when comparing unstable populations of neurons and a paired one-sample t-test when comparing stable populations. Significance was considered as  $p < 0.05$ . A single asterisk was used to denote significance of at least  $p < 0.05$ , a double asterisk was used to denote  $p < 0.01$ , and a triple asterisk was used to denote  $p < 0.001$ .

For t-tests resulting in a significant change in one group but not the other (e.g., direct but not indirect; near but not far), a random permutation test in which the group labels were shuffled 500 times was calculated to ensure that the results were not a consequence of a small sample size. In all cases, the t-statistic for the group with the significant change was outside 95% of the distribution of t-statistics obtained from the shuffled labels tests indicating that the significant group was not a chance combination of neurons.

## *2.3 Results*

Two rhesus macaques (P and R) were chronically implanted with bilateral microelectrode arrays in primary motor and dorsal premotor cortices, with electrodes from a single hemisphere used for BMI control and subsequent analyses (see

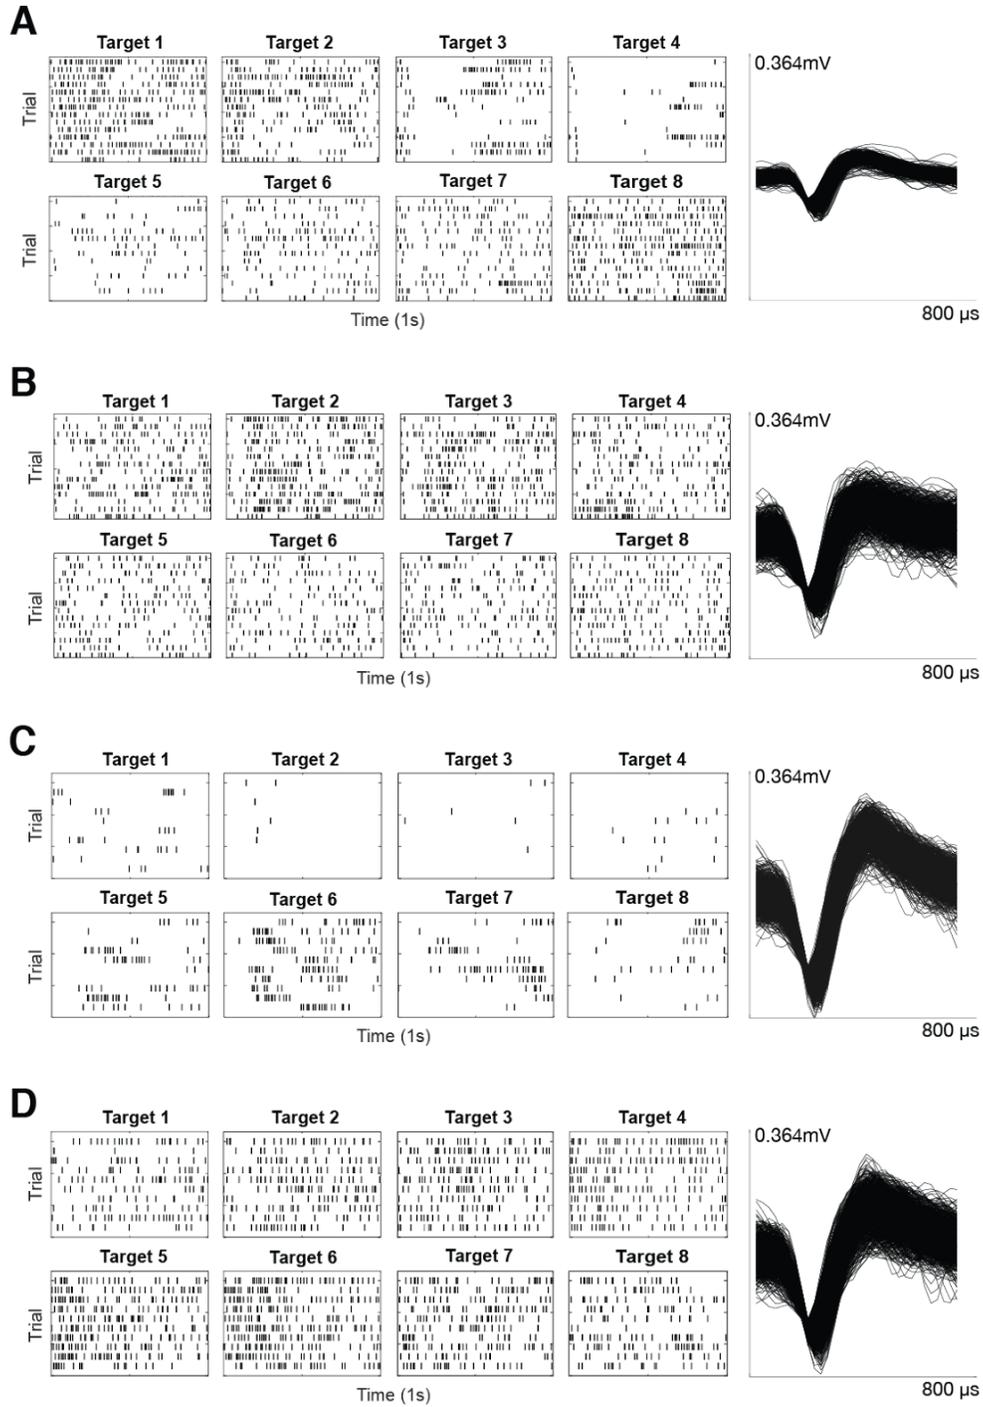
Methods). The monkeys learned to perform a two-dimensional, self-initiated, center-out BMI task, in which they drove a cursor under neural control to one of eight randomly instructed peripheral targets for a juice reward (Figure 2.2A). The next trial was initiated by driving the cursor back to the center target. Trials from all days of the experiment were concatenated then separated into 150-trial epochs since the number of successful trials was lower in early days of learning. Both animals increased the fraction of successful trials (Figure 2.2B) and decreased the time to move the cursor from the center target to the peripheral target (Figure 2.2C) over the course of the first 15 epochs. Example cursor trajectories from early and late learning are shown in Figure 2.2D. To capture the correlates of learning before performance saturated, only the first 15 epochs were used for each animal.



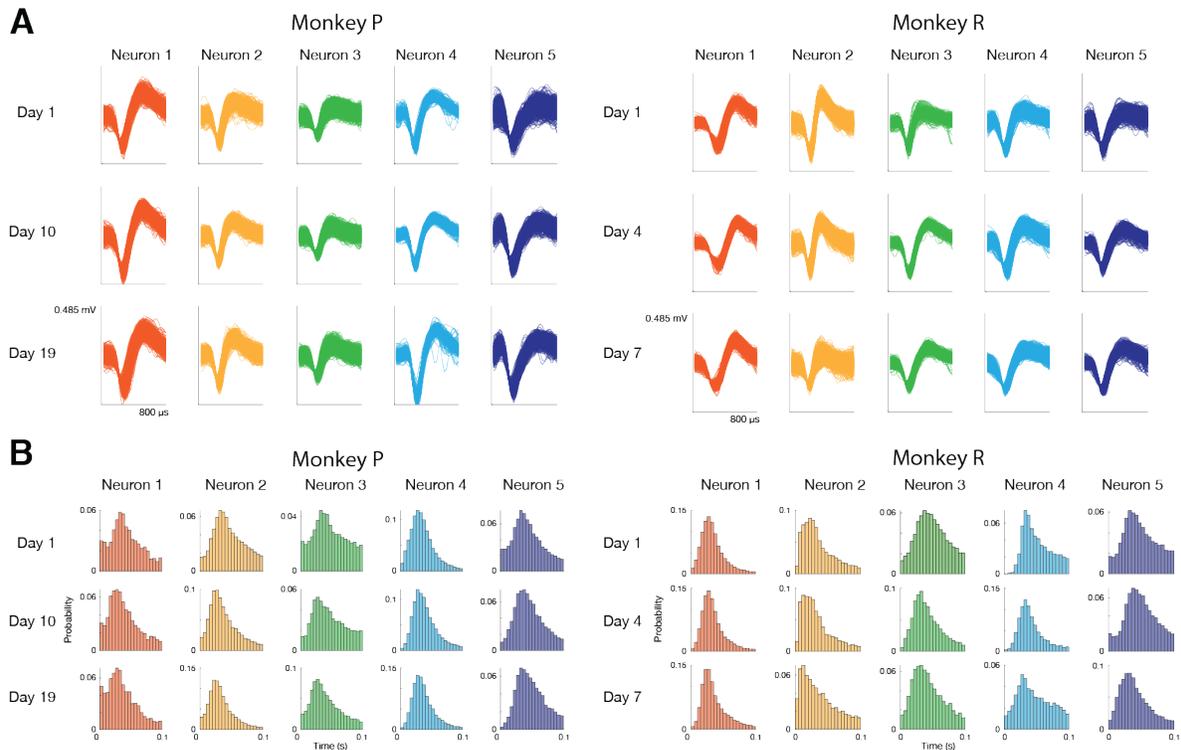
**Figure 2.2 Experimental setup and behavioral performance.** (A) Experimental setup from Ganguly & Carmena, 2009. Activity recorded from direct neurons (blue) in M1/PMd was input into a fixed linear decoder and used to drive a computer cursor to perform a center-out task (see Methods for details). Activity from indirect neurons (red) was simultaneously recorded but was not input into the decoder. (B) Performance improves over the first 15 training epochs for Monkey P and Monkey R. Each training epoch consists of 150 initiated trials. For some analyses, epochs were divided into groups of early (orange) and late (purple). Fraction of initiated trials that were successful increased over training epochs. (C) The time to reach a target decreased over training epochs. (D) Representative examples of single-trial cursor trajectories during the early and late learning.

Because the BMI decoder used in the experiment was novel to the subjects, they had to initially explore the neural population activity space. Over time, the subjects learn

from the behavioral consequences of explored activity patterns and select target-achieving, rewarded patterns of activity. To characterize the neural dynamics associated with this neuroprosthetic learning, we examined recorded ensembles of motor cortical neurons from which only a subset was assigned a causal role during control. We define these neurons, whose activity was used as a direct input to the BMI decoder, as “direct neurons” (Monkey P, N = 15; Monkey R, N = 10). The remaining recorded motor cortical neurons recorded using the same two 4x4 mm 64-channel microelectrode arrays (interelectrode distance 500  $\mu\text{m}$ ), whose activity was not used as direct input to the BMI, we define as “indirect neurons” (Monkey P, N = 29-69; Monkey R, N = 87-187). Spiking activity and waveforms for a representative direct and indirect unit are shown in Figure 2.3. For some analyses in which it is important to consider the same population across epochs we refer to indirect neurons that were stably recorded across all 15 epochs as “stable indirect neurons” (Monkey P, N = 17; Monkey R, N = 14). Stability of the indirect neurons was assessed using the methods described in Fraser & Schwartz, 2012. This method uses pairwise cross-correlograms, the autocorrelogram, waveform shape, and mean firing rate to classify neurons and has previously been used on recordings obtained from chronically implanted microelectrode arrays to assess the stability of neurons across days (Downey et al., 2018; Golub et al., 2014; Vasileva & Bondar, 2021; Yoo et al., 2021). Example waveforms from representative stable indirect units for each animal on the first day, middle day, and last day of recording are shown in Figure 2.4.



**Figure 2.3 Spiking activity of representative units during BMI control.** Raster plots separated by target (left) and 1000 waveforms (right). The height and width the waveform axes are 0.364 mV and 800  $\mu$ s, respectively. (A) Spiking activity of a direct neuron from Monkey P. (B) Spiking activity of an indirect neuron from Monkey P. (C) Spiking activity of a direct neuron from Monkey R. (D) Spiking activity of an indirect neuron from Monkey P.

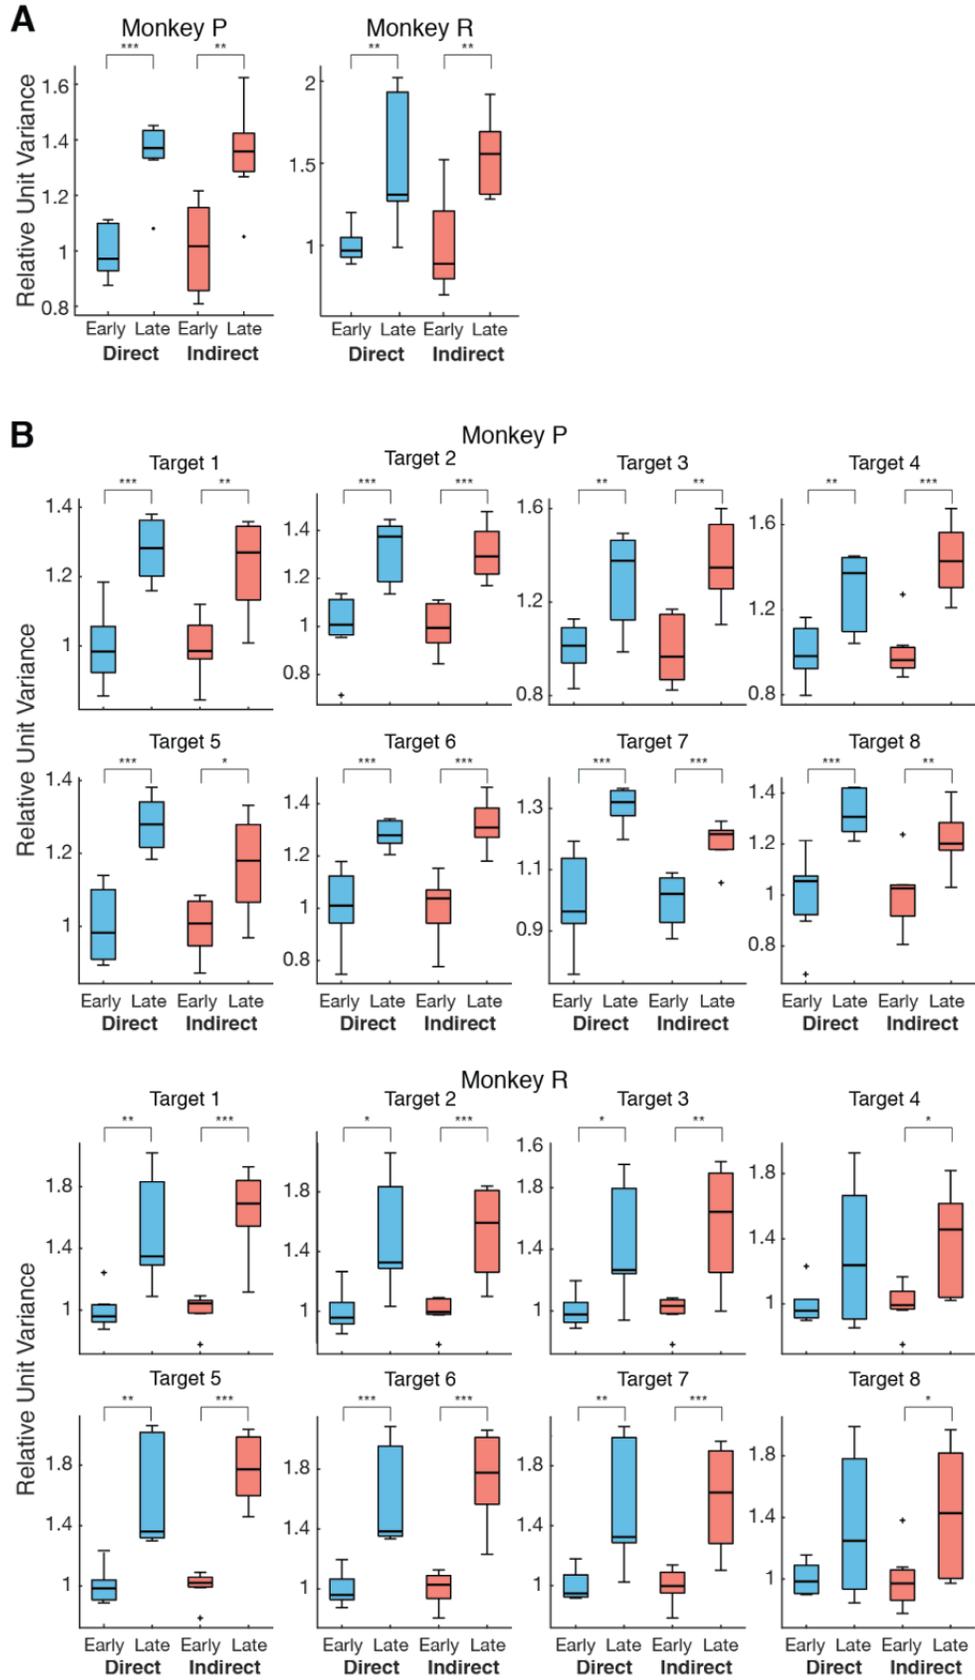


**Figure 2.4 Stability of indirect neuron recordings.** (A) Waveforms from five representative putative stable units as determined by the methods in Fraser & Schwartz, 2012 on the first day, middle day, and last day of recording for each animal. The height and width of each box are 0.485 mV and 800  $\mu$ s, respectively. (B) Inter-spike interval (ISI) distributions are shown for the first day, middle day, and last day of recording for each animal for the same five representative stable units.

### 2.3.1 Increases in Neural Variability with Learning

First, we examined how the neural firing rate variance changed in each subpopulation over learning. Changes in neural variance are often used as a proxy for neural exploration, as increasing the variance in firing rate allows for neurons to form different coordinated patterns of firing (Arduin et al., 2013; Athalye et al., 2017; Dhawale et al., 2017; Hennig et al., 2021; Mandelblat-Cerf et al., 2009; Sternad, 2018). Past work has shown neural activity fires in more coordinated patterns as behavior stabilizes, thereby decreasing the dimensionality in neural space over learning (Athalye et al., 2017; Golub et al., 2018; Oby et al., 2019; Sadtler et al., 2014). We commonly refer to these low-dimensional spaces as manifolds or neural subspaces. In order to observe changes in these neural subspaces, epochs were separated between early and late for each animal (Epochs 1-7 and Epochs 9-15, respectively) to track differences as behavioral performance improved. The firing rate for each neuron was binned in 100 ms intervals.

The binned firing rate variance for each neuron was then averaged across all neurons for each epoch. In early learning, the mean firing rate variance was significantly higher in the direct subpopulation than the indirect subpopulation for one subject and there was no significant difference for the other. We normalized the firing rate variance for each subpopulation based on the mean variance in early learning to assess relative change in variance from early to late learning within each subpopulation. In both animals, we observed an increase in relative unit variance between early and late learning for both direct and indirect subpopulations (Figure 2.5A). This increase in variance suggests a concerted effort of neural exploration that exists in a broader network that includes both direct and indirect neurons. To ensure that this increase in variance was not due to a change in the distribution of time spent at each target, we repeated this analysis within each target (Figure 2.5B). Trials to each target were evenly divided across 15 epochs and relative unit variance during early (first 7 epochs) and late learning (last 7 epochs) was considered for each target individually. The relative unit variance increased within both direct and indirect subpopulations for Monkey P and for six of the eight targets for Monkey R (Table 2.1), indicating that this result was not due to a change in the distribution of time spent at each target. Because the change in relative unit variance was consistent across targets, subsequent analyses included trials to all targets to increase the statistical power associated with more trials. Furthermore, because the task required two-dimensional control to achieve success at all targets, grouping trials across all targets provided better insight into how learning occurs in a generalized two-dimensional space rather than for target-specific activity.



**Figure 2.5 Neural variance increases with learning.** Variance was calculated for each neuron and then averaged across neurons. Relative variance was calculated by normalizing to the mean early variance within subpopulation. (a) Both direct and indirect subpopulations increased relative neural variance from early to late learning (Unpaired t-test; Monkey P: direct  $p = 8.63e-5$ , indirect  $p = 0.002$ ; Monkey R: direct  $p = 0.007$ , indirect  $p = 0.003$ ). (b) Analysis was repeated within target. Both direct and indirect subpopulations increased relative neural variance from early to late learning within target (Unpaired t-tests results reported in Table 2.1.)

Target	Monkey P		Monkey R	
	Direct	Indirect	Direct	Indirect
1	1.74e-04	0.003	0.006	8.64e-05
2	9.01e-04	1.72e-04	0.010	7.72e-04
3	0.004	0.001	0.015	0.002
4	0.006	1.35e-04	0.117 (n.s.)	0.018
5	7.24e-05	0.012	0.001	2.12e-06
6	3.78e-04	1.25e-04	5.71e-04	4.51e-05
7	3.17e-04	5.25e-04	0.005	9.75e-04
8	6.07e-04	0.006	0.093 (n.s.)	0.037

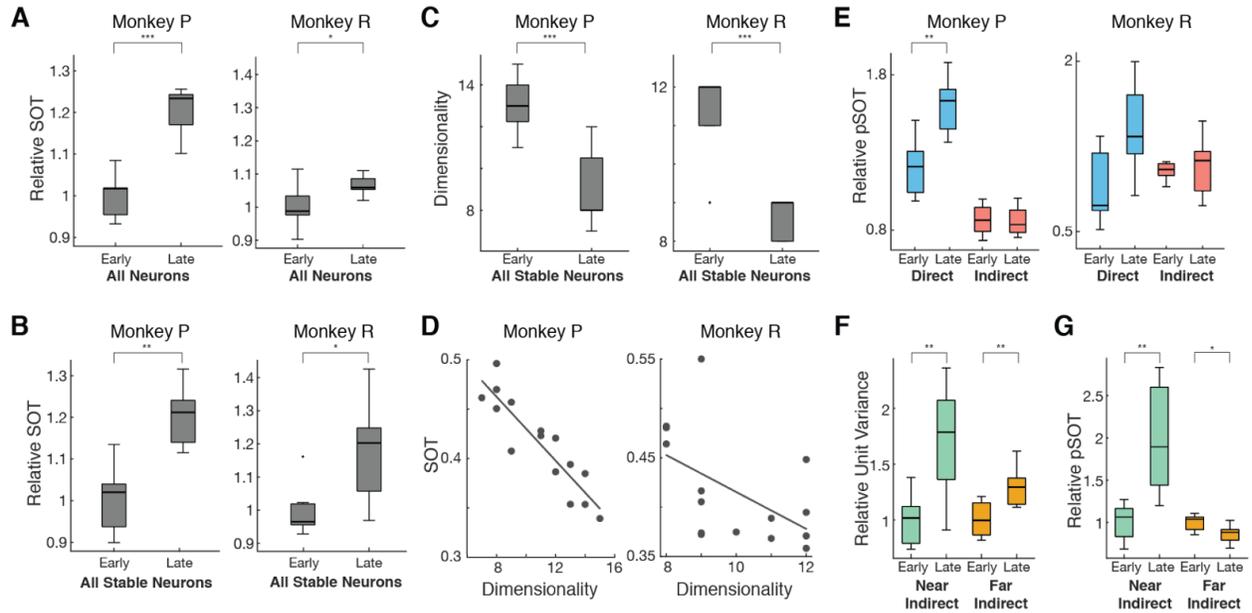
**Table 2.1. Relative variance within target.** P-values from unpaired t-test comparing relative neural variance in early and late learning calculated within each target separately.

### 2.3.2 Increases in Coordinated Neural Activity with Learning

Previous work has shown that neurons fire in increasingly coordinated patterns as performance improves (Athalye et al., 2017). We consider these changes in coordinated firing as a proxy for consolidation of neural population dynamics since the neural variance is stabilizing onto low-dimensional subspaces. We use factor analysis (FA) to separate the neural variance in the population into two components - private and shared variances (Everitt, 1984). The shared variance is the variance between neurons in the population and can be thought of as the underlying correlated firing pattern in the recorded population (Churchland et al., 2010; Yu et al., 2009). Conversely, the private variance denotes the amount of variance each neuron has that

is independent from the rest of the population. Past studies have explored the roles of these components in the direct neuron population, showing private variance as a proxy for exploration while an increase in shared variance is correlated to skill consolidation (Athalye et al., 2017, 2018). Previous work has quantified the amount of coordinated neural activity as a measure of the balance between shared and private variance: the ratio of shared variance over total variance (SOT) (Athalye et al., 2017, 2018; You et al., 2019). Here, we compared the proportion of the total neural variance that is captured in shared spaces for the combined direct and indirect population in each epoch as an estimate of coordination within the recorded population across learning.

Since the indirect population consisted of different units each epoch, we normalized the SOT to the mean SOT for early epochs (Figure 2.6A). In both animals, we found that the relative SOT increased between early and late learning for the entire recorded population including both direct and indirect neurons. Together, with the increase in variance over learning, our results indicated a high level of increased coordination that occurs within the entire recorded population driven by increased exploration as BMI performance improves. To assess that the effect was not due to day-to-day differences in the population, we conducted the same analysis on neurons that were stably recorded across all 15 epochs, which yielded consistent results (Figure 2.6B). Along with an increase in SOT, previous work has shown learning-related decreases in dimensionality of the shared neural subspace for the direct subpopulation (Athalye et al., 2017). We found a similar decrease in dimensionality for the entire stably recorded population (Figure 2.6C) and furthermore found that the dimensionality of the neural subspace in each epoch was significantly correlated with the SOT of each epoch (Figure 2.6D).



**Figure 2.6. Increases in coordinated neural activity over learning are primarily driven by direct neurons.** (A) The relative shared-over-total variance (SOT) ratio

was calculated with respect to the mean early SOT across the entire recorded population. Relative SOT increased between early and late learning, indicating an overall increase in coordination of neural activity in the entire recorded population (Unpaired t-test; Monkey P,  $p = 1.31e-5$ ; Monkey R,  $p = 0.033$ ). (B) Relative SOT also increased between early and late learning for a stably recorded population consisting of the same units each epoch (Paired t-test; Monkey P,  $p = 0.002$ ; Monkey R,  $p = 0.027$ ). (C) Dimensionality of the neural subspace for the stably recorded population decreased from early to late learning (Unpaired t-test; Monkey P:  $p = 4.10e-4$ ; Monkey R:  $p = 8.24e-5$ ). (D) The dimensionality of the neural subspace is correlated with SOT (Linear regression; Monkey P:  $R^2 = 0.813$ ,  $p = 4.41e-6$ ; Monkey R:  $R^2 = 0.277$ ,  $p = 0.044$ ). (E) Respective contributions of each sub-populations to the SOT ratio (pSOT, see Methods for details) were calculated in early and late learning relative to contributions in early learning. Only direct pSOT relative to early learning increased from early to late learning (Paired t-test; Monkey P, direct  $p = 0.003$ , indirect  $p = 0.805$ ; Monkey R, direct  $p = 0.084$ , indirect  $p = 0.675$ ). To test that the change in direct was not due to a chance grouping of neurons, we repeated the t-test 500 times while shuffling direct and indirect labels and compared the true direct t-statistic to a distribution of t-statistics from the shuffled populations (Permutation test; Monkey P, direct  $p = 0.014$ ; Monkey R, direct  $p = 0.078$ ). (F) Both near and far indirect neurons exhibited significant increases in neural variance (Unpaired t-test; near  $p = 0.006$ , far  $p = 0.007$ ). (G) Only near indirect neurons exhibited a significant increase in pSOT relative to early learning (Unpaired t-test;  $p = 0.003$ ). Far indirect neurons exhibited a significant decrease in pSOT (Unpaired t-test;  $p = 0.022$ ). To test that changes in pSOT are not due to a chance grouping of neurons, we repeated the

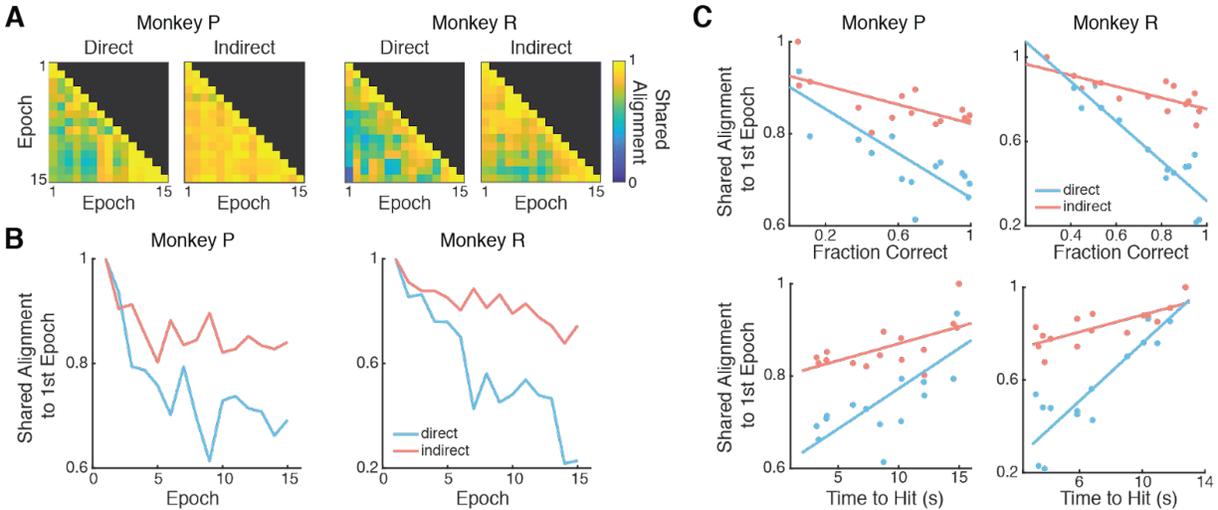
t-tests 500 times while shuffling near and far labels (Permutation test; near,  $p = 0.008$ ; far,  $p = 0.002$ ).

While the increase in SOT indicates more coordination within the entire population, it does not explain whether these changes are driven by a specific subpopulation. To answer this question, we considered the partial shared-over-total variance (pSOT) ratio of the stably recorded population to see how the same population of neurons change coordinated firing activity over learning (see Methods). Intuitively, the pSOT ratio asks how much of the overall change in coordinated activity was driven by one subpopulation versus the other. We see that, relative to early learning, there was an increase in  $pSOT_{\text{direct}}$  in late learning but not in  $pSOT_{\text{indirect}}$  for both animals (Figure 2.6E), indicating that the increase in coordination of population activity seen across the stably recorded population was driven by the increase in coordination of population activity within the direct subpopulation. While this result was only statistically significant in Monkey P, Monkey R (who had fewer direct neurons) exhibited the same trend. The larger increase in  $pSOT_{\text{direct}}$  suggests that the increased neural exploration in the network was primarily a consequence of changes in coordinated patterns specific to the direct neurons. To further characterize these changes in the indirect neurons, we separated all of the indirect neurons in Monkey P into “far” and “near” indirect neurons. “Far” indirect neurons ( $N = 29-69$ ) were those recorded on electrodes not containing direct neurons. In contrast, “near” indirect neurons ( $N = 7-14$ ) were indirect neurons that existed on the same electrode shanks as direct neurons. Monkey R was excluded from these analyses due to recording too few near indirect neurons during several epochs ( $N = 0-10$ ). We found that neural variance increased for both far and near indirect neurons between early and late learning (Figure 2.6F). However, the pSOT only increased for the near indirect subpopulation and significantly decreased for the far indirect subpopulation (Figure 2.6G). Together, these results suggest that while neural exploration exists in broader networks consisting of both direct and indirect neurons, activity in neurons closer in proximity to direct neurons becomes more coordinated than activity in neurons farther away from direct neurons. Differences in near and far indirect neurons could be due to synaptic projections from other brain regions into M1. Because neurons in closer spatial proximity may be more likely to share inputs, the increase in pSOT seen in the near indirect neurons may be the result of an increase in neural variability across an interconnected network.

### 2.3.3 Changes in Neural Covariance with Learning

To characterize how neural exploration modified the direct and indirect neural subspaces differently, we quantified these changes by calculating the shared alignment pairwise between each training epoch’s shared covariance matrix for each

subpopulation according to the methods described in Athalye et al., 2017 (Figure 2.7A). The shared alignment measures the similarity of covariance planes to compare how much of the shared space of one epoch projects onto the shared space of another epoch. Intuitively, given both two-dimensional shared subspaces, the shared alignment compares the angle between the two planes. Orthogonal planes, or subspaces, would result in a shared alignment of 0 and perfectly aligned planes would result in a shared alignment of 1. If the shared subspace consolidates with learning, as has been shown in direct subpopulations (Athalye et al., 2017), we would expect the shared subspace to rotate away from the initial subspace over learning. If the shared subspace remains fixed over learning, we would predict that the alignment between the first epoch and later epochs remains high, indicating little change in the coordinated activity of the population. Since we are interested in how the subspaces pertaining to specific populations change over time, we analyzed only the neurons that were stable across learning. We found that the shared alignment decreased from the first epoch for both subpopulations (Figure 2.7B). This indicates that both subpopulations rotated their low-dimensional subspaces, suggesting that neurons may adapt on a network level that includes both direct and indirect neurons. Furthermore, this rotation of the low-dimensional subspaces is correlated with behavior (Figure 2.7C). As the shared subspace diverges from where it began in the first epoch, the fraction of correct trials significantly increases and the time it takes for the cursor to reach the target significantly decreases. While this is true for both the direct and indirect subpopulations, the extent of rotation as measured by the shared alignment with the first epoch and the proportion of the variation in the shared alignment that is predictable from the behavior were higher for the direct subpopulation than the indirect subpopulation.

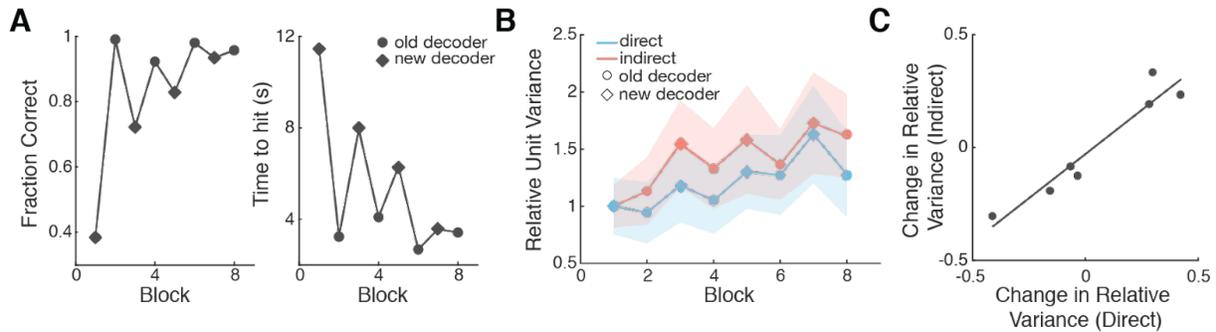


**Figure 2.7. Rotation of low-dimensional neural subspace.** (A) Shared alignment was calculated pairwise between epochs for each subpopulation. (B) Alignment diverges from the first epoch in both subpopulations (Linear regression; Monkey P: direct,  $R^2 = 0.530$ ,  $p = 0.001$ , indirect,  $R^2 = 0.352$ ,  $p = 0.012$ ; Monkey R: direct  $R^2 = 0.857$ ,  $p = 4.63e-7$ , indirect  $R^2 = 0.760$ ,  $p = 1.41e-5$ ). The slopes for direct and indirect shared alignment across epochs were significantly different (One-way ANCOVA; Monkey P:  $p = 0.024$ ; Monkey R:  $p = 5.84e-6$ ). (C) Alignment is correlated with fraction correct (Top, linear regression; Monkey P: direct,  $R^2 = 0.657$ ,  $p = 2.0e-4$ , indirect  $R^2 = 0.481$ ,  $p = 4.2e-4$ ; Monkey R: direct  $R^2 = 0.870$ ,  $p < 1.0e-6$ , indirect  $R^2 = 0.589$ ,  $p = 8.0e-4$ ) and time to hit (Bottom, linear regression; Monkey P: direct,  $R^2 = 0.535$ ,  $p = 0.002$ , indirect  $R^2 = 0.380$ ,  $p = 0.014$ ; Monkey R: direct  $R^2 = 0.815$ ,  $p < 1.0e-6$ , indirect  $R^2 = 0.592$ ,  $p = 8.0e-4$ ).

### 2.3.4 Changes in Neural Variability with a Novel Decoder

To further explore how indirect and direct neural activity may adapt together, we analyzed data from a second experiment in which Monkey P learned to perform the same BMI task with a new decoder following proficient control with the original learned decoder (Figure 2.8A). The new decoder used the same direct neurons as the original decoder, but the decoder weight assigned to each direct neuron was changed so that the same activity patterns result in different cursor movements when using the different decoders. Eight experimental blocks were performed over the course of four days, alternating between control with the new and previously learned decoder each day. The neural variance for direct and indirect neurons was calculated within each of these eight blocks. Note that only stable indirect neurons were used for this analysis since we wanted to explicitly track how the variance changed as a function of block number. We found that both subpopulations increased and decreased neural variance together over blocks, with similar changes in variance between blocks occurring in

both direct and indirect neurons (Figure 2.8B). Thus, increases in neural variability with changing decoders over shorter timescales involved increased exploration not only by the direct neurons, but also by the supporting indirect neurons. Because there was an increase in neural variability in both subpopulations with each decoder swap, we were also able to assess whether the amount change in neural variance was similar between the two subpopulations. We found that the changes in the firing rate variance of the direct subpopulation were correlated to changes in the indirect subpopulation (Figure 2.8C).



**Figure 2.8. Neural variance modulates concomitantly between subpopulations.**

(a) Monkey P learned to perform BMI with a new decoder following proficient control with the old decoder. 8 experimental blocks were performed, alternating between a new decoder (diamond) and the previously learned decoder (circle). Fraction of initiated trials that were successful increased over training blocks (left) and the time to reach a target decreased over training blocks (right). (b) Both subpopulations increase their neural variance over blocks (Linear regression; Direct  $R^2 = 0.623$ ,  $p = 0.020$ , Indirect  $R^2 = 0.655$ ,  $p = 0.015$ ). The relative variances across blocks are correlated between subpopulations (Pearson's  $r$ ,  $r = 0.856$ ,  $p = 0.007$ ). (c) Each point represents the change in relative variance between two consecutive blocks. The changes in relative variance within the direct and indirect subpopulations are correlated (Pearson's  $r$ ,  $r = 0.820$ ,  $p = 0.024$ ).

## 2.4 Discussion

In this study, we explored how changes in cortical population dynamics underlie skill learning. Specifically, we examined how cortical dynamics of subpopulations of neurons change over learning when the mapping of neural activity to behavior is precisely defined using a BMI. That is, how does the adaptation of the subpopulations used as inputs into a BMI decoder (direct neurons) compare to that of a subpopulation not used for decoding (indirect neurons)?

Our results revealed that learning-related neural state space exploration included neurons from both the direct and indirect subpopulations. Both subpopulations increased variance in their firing rates (Figure 2.5) and there was an increase in coordination of neural firing patterns across the combined population that was correlated with a decrease in dimensionality of the neural space (Figure 2.6). While we could not directly compare these changes in neural variance and coordinated activity to undetected movements made by the animals, both animals were observed to have minimal movements during BMI control (Ganguly & Carmena, 2009). Previous literature has suggested that volitional modulation of neural activity in animals may be related to movements, cognitive imagery, or shifts in attention (Fetz, 2007; Hwang et al., 2013). Thus, it is possible that these learning-related changes in neural variance are the result of an underlying behavioral strategy through which the animals learn to modulate their neural activity.

An increase in neural variability with learning, as observed in our results, has also been seen in previous studies. One explanation for this increase in variability is that it allows the brain to explore new activity patterns that may improve behavior (Athalye et al., 2017). Another study finding the same increase in neural variability in early learning proposed that this increase may be the result of changing internal states or increased neural engagement (Hennig et al., 2021). Thus, this increase in variance of the firing rates of both direct and indirect neurons may be a result of an increased exploratory drive or a change in the animals' internal state, which may be indicative of its arousal or uncertainty about its environment.

Previous work observing increases in coordinated population activity suggested that these changes are characteristic of more stereotyped behavior over learning and this increased neural covariance has been associated with subjects making straighter, more direct paths to the targets (Athalye et al., 2017). Furthermore, changes in covariance structure have been shown to relate to synaptic connectivity (Okun et al., 2015). Thus, when considering the entire recorded population, it appears as though the whole population adapts together to facilitate learning. However, past work has shown differences in adaptation between direct and indirect neurons (Ganguly et al., 2011; Gulati et al., 2014; Koralek et al., 2013). When we considered the relative contribution of each subpopulation to the overall increase in coordination, we found that the indirect subpopulation contributed very little to the increase in coordinated patterns (Figure 2.6E). That is, while we witnessed an increase in coordination in the entire recorded population, there was less within-group coordination in indirect neurons compared to direct neurons. Our metrics of coordination (e.g., SOT) rely on averaging the amount of correlated activity between pairs of neurons. Coordination of firing activity occurring more heavily in one subpopulation would nevertheless increase the

SOT in the entire population. Altogether, while both subpopulations exhibit similar levels of exploration over learning as demonstrated by an increase in neural variance, the exploration by the indirect subpopulation results in less of an increase in coordination of neural firing patterns than that of the direct subpopulation.

Changes in coordinated neural activity resulted in rotations of the neural space over learning in both direct and indirect subpopulations (Figure 2.7). The neural space can be intuitively thought of as the lower-dimensional space in which co-activations of neurons exist and rotating this neural space corresponds to adjusting which neurons are more active given the state of other neurons in the population. Following this interpretation, both direct and indirect subpopulations similarly adapted their coordinated firing patterns within subpopulation over learning. This adaptation in coordinated firing patterns was also correlated with behavioral improvements in the task for both the direct and indirect subpopulations, however, the relationship was stronger for the direct subpopulation than the indirect subpopulation for both subjects. This result, along with the increase in pSOT observed in the direct subpopulation, suggests that as the direct subpopulation increases coordination of neural activity its covariance exhibits greater changes than that of the indirect subpopulation over learning.

We also examined how the neural firing rate variance of both the direct and indirect subpopulations changed in an experiment where two decoders were swapped each day, requiring changes in neural activity over short timescales. We found that changes in neural variance occurred with each decoder swap and were proportional between the two subpopulations (Figure 2.8). This suggests that increases in neural variability occur over both short timescales, as seen in this experiment, and over longer timescales, as seen in the initial 15 epochs, involve both neurons within the supporting cortical network and direct neurons. Furthermore, the changes in neural variance between the two subpopulations were correlated. Consequently, both the direct and indirect subpopulations may be adapting via the same mechanism but to different extents. However, this experiment was limited to only four days of switching between decoders. It is unclear whether or not these parallel changes in neural variance between the direct and indirect subpopulations would continue if the animal was given more extensive practice.

Our findings that both direct and indirect subpopulations increase neural variability similarly but exhibit differential changes in coordination could be explained by existing hypotheses on how the brain learns to refine coordinated neural dynamics (Athalye et al., 2020). Specifically, small networks of cortical neurons may be driven by upstream cortical and subcortical inputs. We found that indirect neurons adapted in similar ways

as direct neurons, suggesting that some indirect neurons may in fact be adapting with or alongside direct neurons. When we investigated how distance from direct neurons influenced these results, we found that indirect neurons in closer spatial proximity to direct neurons increased coordination more than indirect neurons that were farther away (Figure 2.6G). While the upstream projections from cortical and subcortical structures are not necessarily spatially organized, our results are consistent with the hypothesis that upstream structures may be driving changes in smaller groups of interconnected neurons (Vyas, Golub, et al., 2020). Furthermore, it has previously been shown that when disparities are present between the control space and neural space (i.e. how well the decoders aligned to the natural firing patterns of the neurons), neurons with larger disparities adapt more over learning compared to neurons with smaller disparities (Athalye et al., 2017; Chase et al., 2012; Jarosiewicz et al., 2008; Orsborn et al., 2014). A recent study has shown that task-related neurons, consisting of direct neurons as well as task-modulated indirect neurons, increase coherency to slow-wave activity (SWA) during sleep which has been linked with consolidation (Gulati et al., 2014). This suggests that in addition to online task practice, neural reactivations during sleep can aid in exploring the contributions of direct and indirect neural population relative to successful outcomes and reward. Notably, indirect neurons that were closely tied to reward were preserved and resembled direct neurons; this might explain why some indirect neurons were modified during neuroprosthetic skill acquisition. This also provides further evidence that mechanisms of reinforcement learning may underlie our observed phenomena. Thus, it is quite plausible that adaptation of neural activity over BMI learning is attributed to finding the clusters of neurons with a direct effect on behavior, which may include both direct and indirect neurons, depending on their specific network connectivity and temporal association with successful outcomes.

In this study, we used factor analysis (FA) to find correlations in the neural activity of the recorded population. Underlying this model are latent factors, which are variables that coarsely group neurons together based on coordinated activity patterns. Importantly, the activity of a single neuron can be associated with multiple latent factors. While we were agnostic to what the latent factors in FA may correspond to in this study, they may be physiologically analogous to upstream connections from cortical or subcortical structures that drive changes in small clusters of neurons containing direct neurons. The idea that neural reinforcement is dependent on cortico-cortical and cortico-striatal circuits, similar to behavioral reinforcement, has been previously supported by studies using BMIs. For example, as rodents learned to produce specific patterns of cortical activity, coherence between these neurons and dorsal striatum emerged and neurons in dorsal striatum developed target-predictive

modulation of firing activity (Koralek et al., 2012, 2013; Neely et al., 2018). Furthermore, mice without functional NMDA receptors in striatal projection neurons could not learn to re-enter a cortical pattern that led to reward. Thus, cortico-striatal plasticity is necessary for learning to efficiently produce the cortical activity patterns required to obtain rewards. These findings along with the results from our study further support the hypothesis that smaller clusters of neurons, which may include both direct and indirect neurons, are adapted over learning more than clusters of neurons that do not drive behavior.

Overall, our results demonstrate that the brain learns to modify cortical population dynamics in subpopulations relevant for behavioral control. When using a BMI, we find that neurons with direct input to the decoder as well as neurons in the surrounding cortical network increase exploration and consolidate their firing activity onto a low-dimensional neural space. The degree of coordination among the population is dependent on the relationship of the neural activity to the behavioral output. Thus, the brain may not be reinforcing the activity of single neurons, but rather reinforcing cortical population dynamics that are relevant to producing a desired behavior. These findings indicate that the brain learns to control a BMI by refining cortical population-level dynamics, suggesting that BMI decoders extracting information based on population-level statistics, such as the covariance structure of the population, may be more effective compared to traditional decoding methods based on the statistics of individual neurons. Dimensionality reduction techniques such as FA allow us to pull out the correlated activity in a population of neurons. That is, the shared variance obtained using FA represents the concerted activity of a population and if we assume that the uncorrelated activity is largely noise, then using these population-level statistics effectively increases the signal-to-noise ratio of the neural activity. Building decoders based on these smoothed neural signals may translate to smoother output signals (e.g., cursor movements).

In conclusion, this study demonstrates an emergence of coordinated population dynamics within both populations of neurons whose activity is directly used as input for BMI control as well as within the surrounding network. The extent to which these subpopulations modify their coordinated activity varies, with the direct subpopulation exhibiting larger changes. Understanding the role of modifications of adjacent indirect activity in obtaining precise control of a BMI may help us understand the neural adaptation that is required for achieving long-term, stable control of a BMI.

## Chapter 3: Simultaneous Large-Scale Recordings During Neuroprosthetic Control in Nonhuman Primates

Brain-machine interfaces (BMIs) allow for the real-time transformation of neural activity into control signals for external devices. Learning to control these devices engages a wide array of learning mechanisms that rely on distributed cortical and subcortical areas. Many of these mechanisms are associated with both the caudate nucleus of the striatum (Cd) and the dorsolateral prefrontal cortex (DLPFC). Previous work in rodent BMI has demonstrated involvement of the striatum in neuroprosthetic skill learning, however DLPFC has been largely ignored in motor BMI studies. In this chapter, I present a novel method for recording motor cortical units from a nonhuman primate (NHP) to input to a BMI while simultaneously recording from multiple cortical and subcortical brain regions. This technique enables the investigation of the role of distributed brain regions and their interactions during BMI control to create a more detailed, mechanistic understanding of how neuroprosthetic control is implemented in the brain. To demonstrate this, we investigated changes in local field potentials (LFP) from Cd and DLPFC that occur during BMI control. We found changes in spectral power that were predictive of control type and directed functional connectivity with motor cortex that differed from baseline. These findings provide further evidence that a distributed network of cortical and subcortical areas is involved in neuroprosthetic skill learning and control. This work was done in collaboration with Gabrielle Shvartsman, Joni Wallis, and Jose M. Carmena.

### 3.1 Introduction

The ability to volitionally modulate the activity of single neurons (Chapin et al., 1999; Fetz, 1969, 2007; Fetz & Baker, 1973; Kennedy & Bakay, 1998) or populations of neurons (Khanna et al., 2013, 2017; Leuthardt et al., 2004; Wolpaw et al., 1991, 2002) is fundamental to the operation of closed-loop motor BMIs. With a BMI, subjects must learn to produce specific patterns of activity to control an external device. Previous work has revealed that this ability is a learned skill by demonstrating that performance increases with practice (Carmena et al., 2003; Ganguly & Carmena, 2009; Moritz et al., 2008; Schalk et al., 2008; Wolpaw & McFarland, 2004). The process of learning to volitionally modulate neural activity for the purpose of controlling an external device is referred to as neuroprosthetic skill learning.

Neuroprosthetic tasks share characteristics with both concrete motor tasks and abstract cognitive tasks (Green & Kalaska, 2011; Wander et al., 2013). Specifically, an abstract neuroprosthetic task requires users to produce certain patterns of neural activity to manipulate their physical environment without physical movement. The

process of learning that specific patterns of neural activity result in a desired outcome is essential for neuroprosthetic skill learning and can be thought of as operant conditioning of cortical activity patterns. The prefrontal cortex and basal ganglia play key roles in learning arbitrary stimulus-response associations common in abstract tasks (Antzoulatos & Miller, 2011, 2014; Badre et al., 2010; Boettiger, 2005; Genovesio et al., 2005; Muhammad et al., 2006; Pasupathy & Miller, 2005). Learning to associate a specific pattern of neural activity with a particular outcome, such as a specific cursor movement, may also involve the prefrontal cortex and basal ganglia. Specifically, both dorsolateral prefrontal cortex (DLPFC) and the caudate nucleus of the striatum (Cd) have been implicated in learning abstract associations (Antzoulatos & Miller, 2011, 2014; Pasupathy & Miller, 2005). These two regions are extensively interconnected with one another, as well as with sensory, motor, and higher-level associational areas (Haber, 2003, 2016; Petrides & Pandya, 2006), making them prime candidates for involvement in neuroprosthetic skill learning. Furthermore, cortico-striatal plasticity has been shown to be necessary for both motor learning (Costa et al., 2004; Hikosaka et al., 1999) and neuroprosthetic learning (Koralek et al., 2012, 2013; Neely et al., 2018).

Previous work investigating role of cortico-cortical and cortico-striatal interactions involved in neuroprosthetic learning and control have relied on simplified, one-dimensional tasks in rodents (Koralek et al., 2012, 2013; Neely et al., 2018) or patients with electrocorticography arrays (Wander et al., 2013, 2016). Furthermore, interactions between DLPFC and Cd during a neuroprosthetic task have not been studied despite their role in learning abstract associations. This is partially due to the technical difficulties of simultaneous recordings from cortical and subcortical regions in a nonhuman primate (NHP) during an abstract neuroprosthetic task. Proficient control in abstract neuroprosthetic tasks takes several days to learn (Athalye et al., 2017; Ganguly & Carmena, 2009; Oby et al., 2019), requiring a stable recording system such as chronically implanted microelectrode arrays. However, most chronic recording techniques are not suited for recording from subcortical regions. To overcome these challenges, we used a large-scale semi-chronic microdrive to simultaneously record neural activity from cortical and subcortical structures as NHPs learned to control a two-dimensional abstract neuroprosthetic task. This recording technique enables the study of large-scale networks underlying neuroprosthetic skill learning.

In this chapter, we present this novel recording technique for a neuroprosthetic task and investigate the role of Cd, DLPFC, motor cortex, and their interactions in an abstract neuroprosthetic task. We demonstrate that animals can learn to control a BMI with motor cortical recordings from a semi-chronic microdrive and that there are task-

relevant changes in spectral power within and interactions between Cd, DLPFC, and motor cortex.

## 3.2 Methods

### 3.2.1 *Segmenting MRIs for Anatomical Models*

For each subject, anatomical images with resolution of 1 x 0.84 x 0.84 mm were obtained using a 3T Siemens TIM/Trio MRI scanner with a two-channel receive-only head coil. 3D Slicer was used to construct models of the cranium and the neural targets (Fedorov et al., 2012). Boundaries of the skull and neural recording targets were defined using the segmentation and model functions in 3D Slicer. Regions of interest such as primary motor cortex (M1), dorsal premotor cortex (PMd), dorsolateral prefrontal cortex (DLPFC), caudate (Cd), and putamen (Pu) were manually traced in 3D Slicer using the Paxinos primate atlas as a reference (Paxinos et al., 2000). The resulting neuroanatomical and cranial models were used to custom-fit a titanium chamber to each subject and decide on stereotaxic coordinates for implantation.

### 3.2.2 *Large-Scale Semi-Chronic Microdrive System*

The large-scale semi-chronic microdrive system used in these experiments has been previously described in detail (Dotson et al., 2017; Qiao et al., 2016). It consists of a guide array, an actuator block, a printed circuit board (PCB), and a screw guide. The actuator block houses a set of linear actuators ( $n = 124$ ), each consisting of a miniature stainless steel lead screw, a threaded brass shuttle, and a compression spring. Each actuator provided 32 mm of electrode travel at a resolution of 8 turns/mm. The actuators were spaced at 1.5 mm intervals. Electrodes consisted of both glass-coated Tungsten electrodes (AlphaOmega) and Platinum-Iridium electrodes (MicroProbes).

### 3.2.3 *Implanting the Chamber and Microdrive*

The large-scale semi-chronic microdrive system was implanted in three stages: chamber implantation, craniotomy, and microdrive implantation. Each stage was followed by a period of recovery and testing to ensure the animals were healthy before proceeding to the next step. Fluid from inside the chamber tested negative for infection before advancing to the next procedure at all stages to assure the sterility of the hermetically sealed chamber. All procedures were conducted in compliance with the National Institute of Health (NIH) Guide for the Care and Use of Laboratory Animals and were approved by the University of California at Berkeley Institutional Animal Care and Use Committee.

To implant the chamber in the first stage, the skin, fascia, muscle, and periosteum are retracted from the cranial bone over an area that is slightly larger than the footprint of the chamber. The chamber is then placed in position and sealed around its perimeter with a thin bead of C&B-Metabond (Parkell) cement. The coordinates for placing the chamber were determined using an anatomical model from an MRI. The chamber is then anchored to the skull using Titanium bone screws and acrylic bone cement. At the end of this stage, the cranial bone is left intact, and the chamber is hermetically sealed with silicone, a rubber gasket, and a short plug screwed into the titanium chamber. The chamber is then covered with a protective cap and the animal was given a minimum of two weeks to recover from this procedure before proceeding to the next stage.

The second surgical step involved making a craniotomy inside of the chamber. After removing the cap, plug, and gasket, the chamber is rinsed, and the cranial bone surface debrided of connective tissue. The skull can then be removed within the chamber. The chamber, now with the cranial bone removed inside, is hermetically sealed again. This time, a tall plug is used to minimize bone and tissue regrowth in between surgical procedures. The animal was given a minimum of 48 hours to recover from this procedure before moving to the next stage.

To implant the microdrive in the final surgical stage, the dorm-fitting plug is removed, and the inside of the chamber is cleaned of connective tissue before replacing rubber gasket and silicone and mounting the microdrive. The microdrive is secured to the chamber using machine screws that pass through the actuator block. The titanium chamber served as the reference and ground connection to the animal using a machine screw linked to a trace on the PCB.

### *3.2.4 Impedance Monitoring and Advancing Electrodes*

After the microdrive had been mounted, the animals were given at least 48 hours to recover and receive a negative test for infection before advancing electrodes. Once the animal has been confirmed negative for infection, the electrodes can be advanced into the brain gradually over a period of 4-6 weeks. For the microdrive model used in these experiments, each leadscrew has a thread pitch of 125  $\mu\text{m}$  per turn. The electrodes are advanced by counterclockwise rotation of the leadscrew. This was done in an incremental manner by advancing a subset of 10–30 electrodes each day until unit activity was detected. Adjacent electrodes were never lowered in a single day to avoid excess tissue compression.

The impedance of the electrodes was monitored using a TDT NanoZ while advancing them into the brain. Initially, the impedance values will be high ( $>2.5 \text{ MOhm}$ ), indicating an open circuit while the electrodes are retracted inside of the microdrive. As

the electrodes are advanced outside of the microdrive and into the fluid between the brain and the bottom of the chamber, the impedance should drop to the electrode impedance. The impedance will begin to rise by about 1 MOhm as it is pressed into the dura. Then, when the electrode breaks through the dura, it will drop back to the current electrode impedance. If the impedance drops much lower ( $<100$  kOhm), this usually indicates a broken electrode tip.

Once the electrodes have passed the dura and are in the brain, they should be advanced by no more than 10 turns (1.25 mm) each day to avoid tissue compression. The electrode should be advanced until the neural target is reached. The number of turns can be estimated using the neuroanatomical models built in 3D Slicer (number of mm  $\times$  8 turns). Once the neural target is reached, the electrode position can be maintained. If desired, the electrode can be advanced further to search for new neural activity.

### *3.2.5 Intracortical Recording*

Neural data were recorded using the OmniPlex Neural Recording Data Acquisition System (Plexon Inc, Dallas, TX). Single- and multi-unit activity was sorted prior to beginning recording sessions using an online sorting application (Sort Client, Plexon). Wideband activity was recorded at 5 kHz. LFP activity was obtained by low-pass filtering at 250 Hz, notch filtering at 60 Hz and 120 Hz, and down-sampling to 1 kHz.

LFP activity was common median referenced by first z-scoring activity within each channel by subtracting the mean and dividing by the standard deviation in each recording session and then subtracting the median value across all channels at each time point. After subtracting the median to remove common signals, likely low frequency noise, across channels, the LFP activity was multiplied by the standard deviation and the mean was added to restore the LFP to its original scaling. Common median referencing was used instead of common average referencing because it is less susceptible to influence from outliers. The common-median-referenced LFP activity was then z-scored using the mean and standard deviation within channel across all recording sessions within a day. Unless otherwise stated, LFP analyses were performed using the average LFP across channels within each region of interest.

### *3.2.6 Center-Out Task*

Subjects performed a self-paced delayed center-out reaching task to eight targets. Trials were initiated by moving to the central targets. A successful trial required a short hold at the center, moving to the peripheral target within a specified time, and a brief

hold at the target. Successful trials resulted in a liquid reward; failed trials were repeated. Target directions were presented in a pseudo-randomized order.

Subjects were overtrained in the center-out task performed with arm movements before starting BMI. In this manual control (MC) version of the task, the subject's arm moved in a KINARM exoskeleton (BKIN Technologies) that restricted movements to the horizontal plane. Neural activity recorded during MC was used to train a BMI decoder. Using this BMI decoder, the animals performed the same task under BMI control (BC). During BC, the animals' arms were restricted in a fixed position within the exoskeleton and the animals were required to move the cursor to the target by modulation of motor cortex activity.

### 3.2.7 Brain-Machine Interface

#### *Brain-Machine Interface Algorithms*

Subjects learned to control a two-dimensional BMI cursor in real-time using a fixed velocity Kalman Filter decoder (Gilja et al., 2012; S.-P. Kim et al., 2008; W. Wu et al., n.d.). The Kalman Filter assumes two linear models:

$$x_{t+1} = Ax_t + w_t$$

$$y_{t+1} = Cx_t + q_t$$

where  $x_t$  and  $y_t$  are the cursor state and neural activity at time  $t$ , respectively. The first equation represents the state-transition model, which describes the state of the cursor over time. It is specified by the state-transition matrix  $A$  and additive Gaussian noise term  $w_t \sim N(0, W)$ . Equation 2 represents the observation model and describes the relationship between neural activity and cursor state. It is parameterized by the observation matrix  $C$  and additive Gaussian noise  $q_t \sim N(0, Q)$ . Neural activity was input as a vector of spike counts in 100 ms bins from the selected direct units.

#### *BMI Decoder Training*

Decoder parameters were initialized from neural and cursor kinematic data collected during the MC version of the task at the beginning of each recording day. Maximum likelihood estimation methods were used to fit initial parameters. Neural data from 10-20 single- and multi-units recorded from motor cortex were selected for BMI control each day (i.e., direct units). The population of direct units was highly overlapping from day to day, but there was some variability as units dropped out or new ones appeared. For Monkey Y, closed-loop decoder adaptation (CLDA) was performed using the SmoothBatch algorithm before fixing the BMI decoder. This algorithm uses knowledge

of task goals (i.e., reaching targets) to infer a subject's intent. The intended kinematics and observed neural activity during closed-loop BMI were used to re-estimate KF parameters. The SmoothBatch algorithm re-estimates the observation model of the KF (matrices C and Q), and updates were constrained to enforce smoothness. The SmoothBatch has been described in detail previously (Dangi et al., 2013; Gowda et al., 2014; Orsborn et al., 2012). CLDA was typically run for 2-5 minutes to provide the subject with adequate performance to allow successful reaches to all targets. For Monkey H, the initial decoder trained from MC data was used.

### *3.2.8 Neural Data Analysis*

All analyses were performed in Python with custom-written routines utilizing publicly available software packages including scipy, numpy, sklearn, and statsmodels. Unless otherwise specified, analyses were performed on the average LFP signal across all channels within each region of interest.

#### *Linear Classifiers*

To classify control type using power features, linear discriminant analysis (LDA) models using singular value decomposition (SVD) were trained and tested using the sklearn toolbox (Pedregosa et al., 2011) in Python. These models used power estimates obtained using Welch's method from segments of LFP recorded during the BMI task, the manual task, or the baseline period from a single recording day as features for classifying control type. To avoid negative effects on imbalanced data between classes, the number of trials per task type (either manual or BMI control) was matched to the class with the fewest number of trials per day by selecting a subset of trials from the larger class of equal size to the number of trials in the smaller class. The corresponding number of windows were selected from the baseline period. Typical frequency band cut-offs were used. Theta is defined as 4-8 Hz, alpha as 8-13 Hz, beta as 13-35 Hz, gamma as 35-75 Hz, and high gamma as 75-150 Hz.

Results for each model were validated using a 10-fold cross-validation within each recording day. Classification accuracies were compared to a chance accuracy obtained by shuffling task labels 1000 times. All frequency bands and regions of interest (ROIs) were trained and tested independently of one another first. Then, classifiers using power features from all frequency bands from different combinations of ROIs were compared. This process was performed separately for 500 ms segments of LFP recorded after the go cue and for 500 ms segments of LFP recorded prior to target achievement.

## Granger Causality

Granger causality was used to estimate the directional functional connectivity among all pairs of regions of interest. Granger causality relies on an autoregressive (AR) modeling framework, in which future values of a time series are modeled as a weighted combination of past values of time series. The quality of an AR-model is assessed by quantifying the variance of the model's residuals. If the variance of the AR-model's residuals is reduced by the inclusion of past measurements from a second time series, then the second time series is said to Granger-cause or G-cause the first. Thus, determining the Granger causality of one signal on another requires two AR-models:

$$X_t = A \cdot X_{t-1} + \varepsilon_t$$

$$X_t = A \cdot X_{t-1} + B \cdot Y_{t-1} + \varepsilon'_t$$

The first model is a univariate AR-model in which values of time series  $X$  are predicted as a weighted combination of past values of time series  $X$ . The second model is a bivariate AR-model in which values of time series  $X$  are predicted as a weighted combination of past values of time series  $Y$  in addition to past values of time series  $X$ .

A significant reduction of the variance of the residuals when comparing the univariate AR-model to the bivariate AR-model implies that inclusion of information about the past values of time series  $Y$  in the prediction of time series  $X$  improves prediction of time series  $X$  beyond inclusion of only past values of time series  $X$ . More explicitly, Granger causality is defined as the natural logarithm of a ratio of residual variances obtained from two different AR-models:

$$G_{Y \rightarrow X} = \ln \left( \frac{\text{var}(\varepsilon_t)}{\text{var}(\varepsilon'_t)} \right)$$

where  $\varepsilon_t$  and  $\varepsilon'_t$  represent the residuals from the univariate and bivariate AR-models, respectively. In cases where one signal G-causes another, the ratio of the variances is larger than one, leading to a G-causality value that is larger than zero (Bressler & Seth, 2011; Ding et al., 2006; Granger, 1969). Applying this logic, we obtained G-causality estimates using simultaneously recorded LFP signals.

We compared Bayesian information criterion (BIC) of AR-models of different orders,  $p$ , for each combination of LFP signals obtained from all pairs of regions during each trial from all recording sessions. We selected  $p=15$  to minimize BIC in the average case,

allowing for the signals to be sufficiently long enough to capture the data structure without over-parameterization.

To account for any spurious interactions and assess the statistical significance of our interactions of interest, we estimated the G-causality between time series that were not recorded at the same time or during the same task. LFP signals recorded during baseline were used to predict LFP signals recorded during BMI control and vice versa to get a distribution of values obtained from random chance. Only estimates in G-causality that were significantly greater than these values were considered statistically significant.

To determine whether interactions were task-relevant, we computed the gain in Granger causality from baseline as follows:

$$Gain_{Y \rightarrow X} = \frac{G_{Y \rightarrow X, BMI} - G_{Y \rightarrow X, Baseline}}{G_{Y \rightarrow X, Baseline}}$$

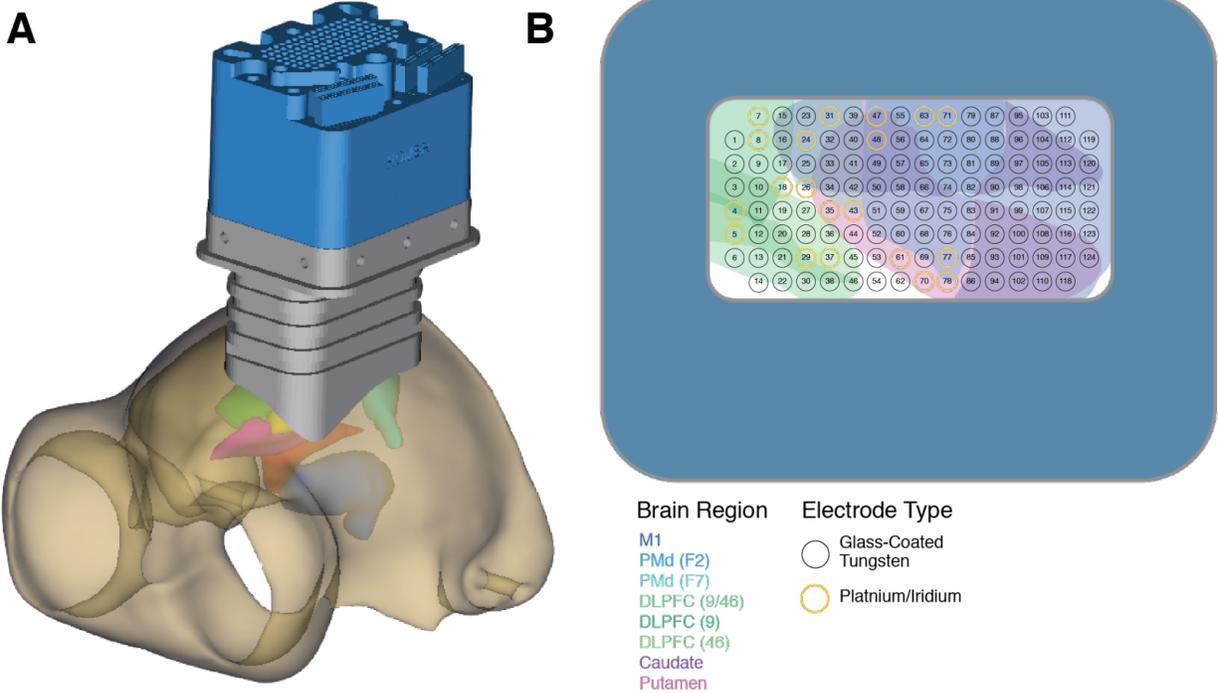
Values greater than zero indicate that the interaction is greater during the task than during baseline, while values less than zero indicate that the interaction during the task is below that during baseline.

### 3.2.9 Quantification & Statistical Analyses

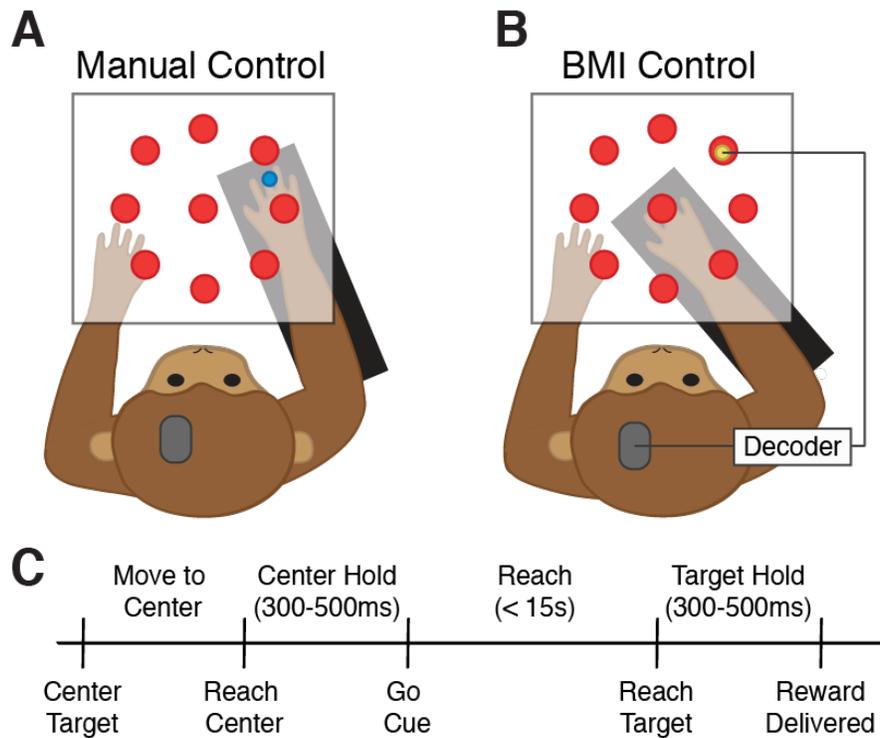
All analyses were performed within a single recording day and error is depicted across days. For analyses comparing two distributions or comparing a single distribution to a value, two-sample or one-sample t-tests were used, respectively. Bonferroni correction was used post hoc to correct for multiple comparisons. Significance is reported after correction for multiple comparisons.

## 3.3 Results

Two rhesus macaques (Monkey H and Monkey Y) were implanted with a custom-fit large-scale semi-chronic microdrive array on the left hemisphere (Figure 3.1). Single- and multi-unit recordings from motor cortex were used as input to a BMI decoder, while local-field potentials (LFP) were simultaneously recorded from primary motor cortex (M1), dorsolateral prefrontal cortex (DLPFC), and caudate (Cd). Each day, the animals performed a two-dimensional, self-initiated, center-out task. First, they performed the task under manual control followed by a four-minute baseline period. Then, they performed the same center-out task under BMI control, where they drove a cursor by modulating motor cortical activity to one of eight pseudo-randomly instructed peripheral targets for a juice reward (Figure 3.2).



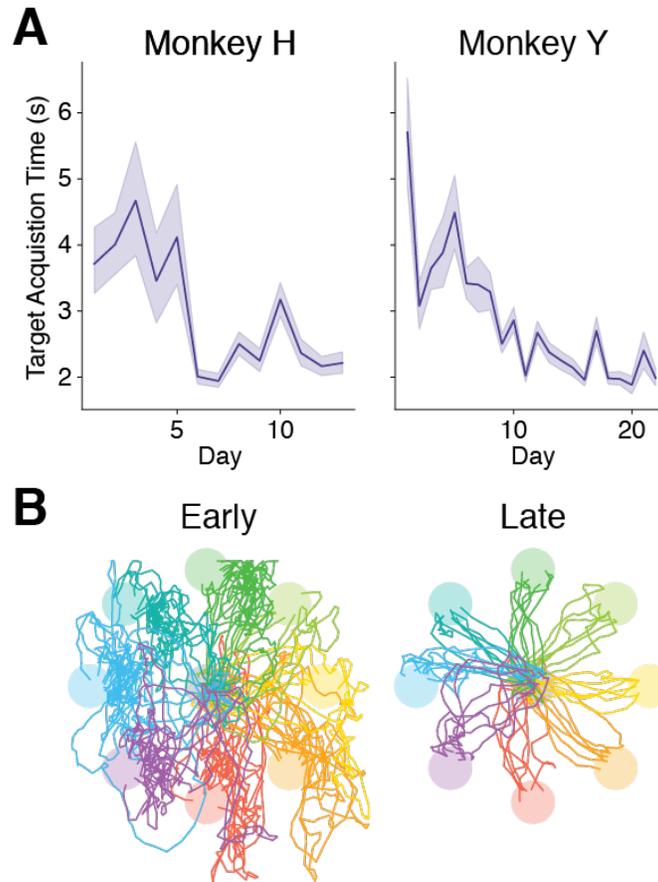
**Figure 3.1. Large-scale semi-chronic microdrive array for chronic, high-throughput, neural recordings.** (A) Three-dimensional model of the large-scale semi-chronic microdrive, neural targets, and skull of Monkey H obtained using *3D Slicer*. (B) Arrangement of electrodes overlaid with neural targets. Different brain regions are shaded in different colors as indicated by the legend. Electrodes were either glass-coated tungsten (black) or platinum-iridium (yellow).



**Figure 3.2. Task schematic representation of manual and BMI reaching.** Monkeys performed a two-dimensional, self-initiated, center-out movement task in both manual control (A) and BMI control (B). In manual control, the subject's arm was restricted to a two-dimensional plane within an exoskeleton. In BMI control, the subject's arm was locked in a fixed position. (C) Timeline of the center-out task. See Methods for details.

### 3.3.1 Online BMI Performance

Both monkeys successfully learned to perform the eight-target center-out task using a cursor under BMI control. Task performance showed clear improvements across days, despite using different direct units and decoders each day. The monkeys learned to produce faster, straighter cursor trajectories, resulting in a decreased target acquisition time over days (Figure 3.3A). Example single-trial cursor trajectories from the center target to each peripheral target for successful trials in early learning and late learning are shown in Figure 3.3B. This improved performance demonstrates that the animals were able to learn an abstract neuroprosthetic skill using neural activity recorded from a semi-chronic microdrive.

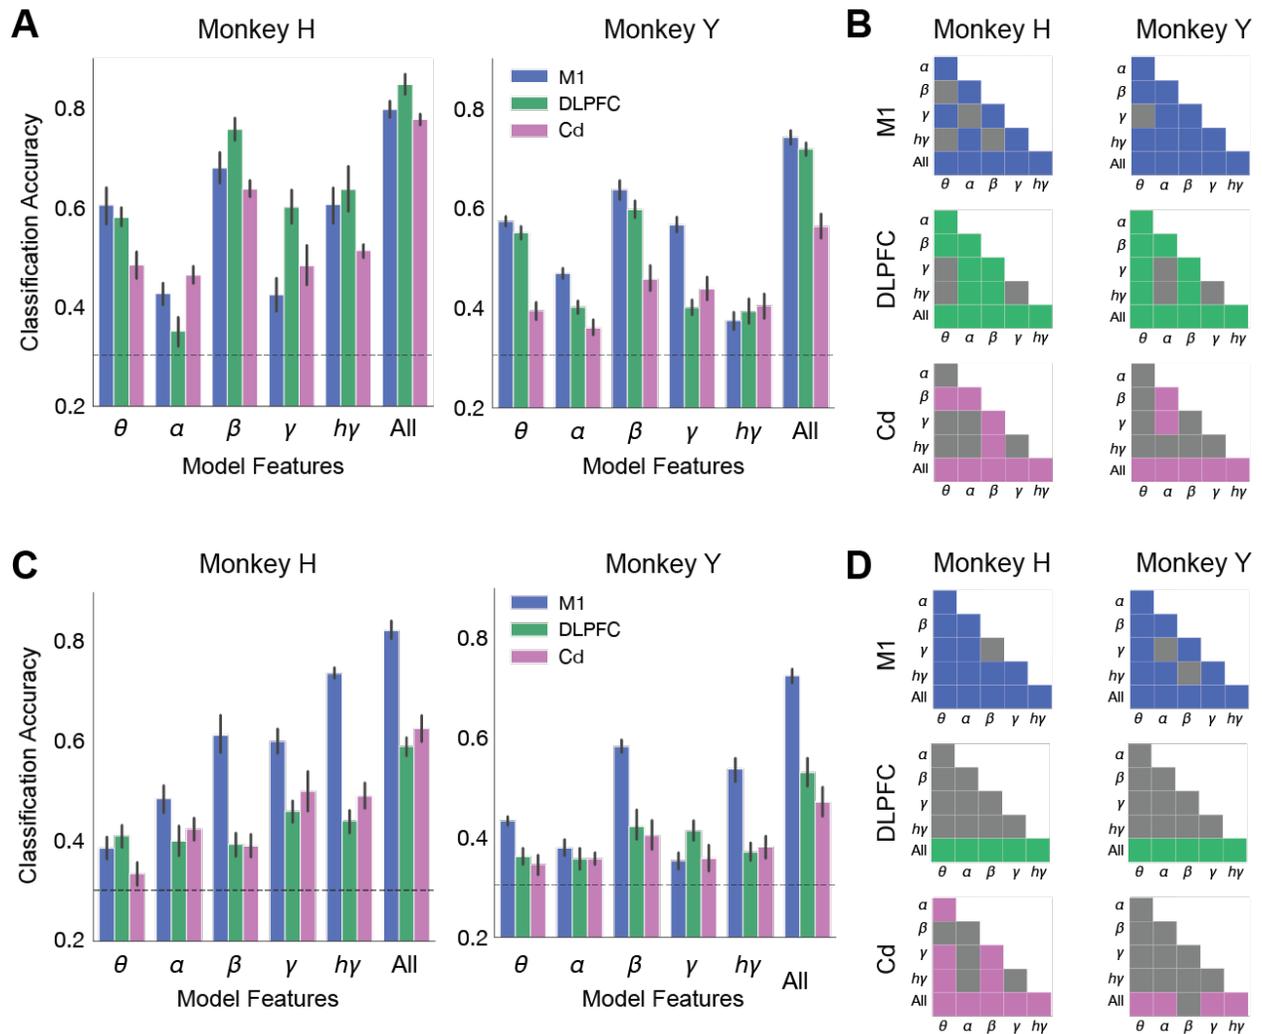


**Figure 3.3. Behavioral performance under BMI control.** (A) The average time to reach a target decreased over training days (Monkey H,  $n = 13$  days; Monkey Y,  $n = 22$  days). (B) Representative examples of single-trial cursor trajectories during the first and last days of recording from Monkey Y.

### 3.3.2 Control Type Classification with Power Features

Using the LFP recorded from multiple electrodes in M1, DLPFC, and Cd during BMI task performance, manual task performance, and a baseline period, we computed the power in various frequency bands using Welch's method of spectral estimation in 500 ms windows. Both monkeys' performance varied across days and across trials, but the 500 ms segment following the go cue of successfully initiated trials was more stereotyped and a period in which the monkeys were reliably engaged in the task. A second period in which the monkeys were reliably engaged in neuroprosthetic control was the 500 ms period immediately before target achievement on successful trials. Thus, 500 ms windows following the go cue and preceding target acquisition on successful trials were used for the BMI task and manual task. The mean spectral power across channels within each region of interest (ROI) was computed within these windows and within 500 ms baseline windows. These power estimates were used as

features in a linear discriminant analysis (LDA) classifier to classify control type: BMI, manual, or baseline. All frequency bands and ROIs were trained and tested independently of one another, and a chance classification estimate was obtained by shuffling control type labels within each recording day. An additional model including all power features from all frequency bands was also included for each ROI.



**Figure 3.4. Classification accuracy using spectral power features from different frequency bands.** LDA was done separately in each region of interest and accuracies for M1, DLPFC, and Cd are indicated by color (blue, green, and pink, respectively). (A) Average 10-fold cross-validated classification accuracy and error across days for individual frequency bands and all combined bands at the go cue for Monkey H (left) and Monkey Y (right). Chance accuracy shown as a dashed line. (B) Comparison of classification accuracies from classifiers using different power features as inputs at the go cue. Classification accuracies were compared within region (rows) and subject (columns). Blue, green, or pink boxes

indicate a significant difference via two-sample independent t-test after Bonferroni correction within M1, DLPFC, or Cd, respectively. (C, D) Same as (A, B) for power features obtained at target acquisition rather than at the go cue.

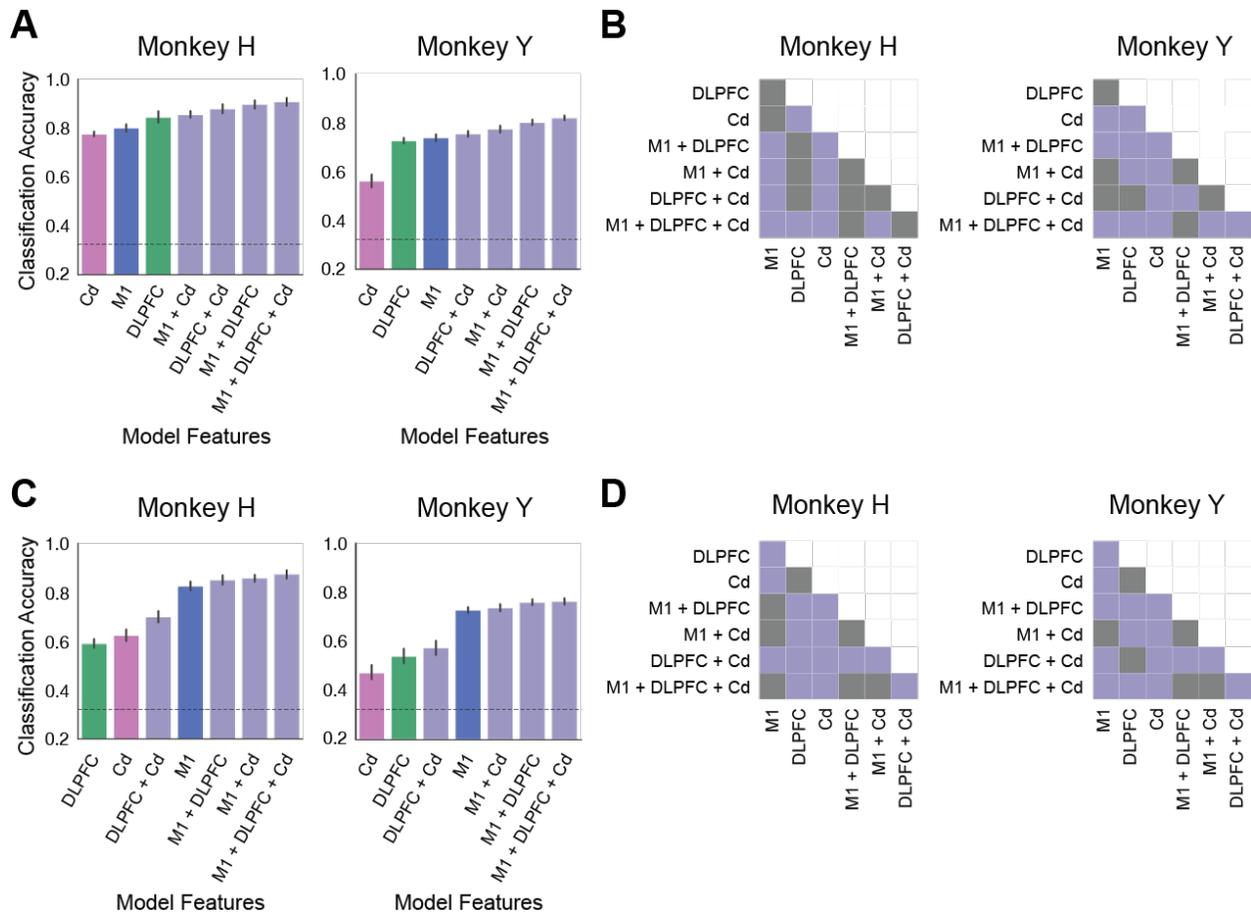
### *Comparing Power Features from Different Frequency Bands*

All of the linear classifiers except for the one using alpha power from DLPFC for Monkey H predicted control type significantly above chance at the go cue (Figure 3.4A). Mean classification accuracies using the different power features were compared within region and subject (Figure 3.4B). Across all regions, beta power features tended to demonstrate the highest classification performance of any individual frequency bands. The classifier using power from all frequency bands as input features performed better than any individual frequency band.

These analyses were repeated using the mean power estimates within each region prior to target acquisition. With power features from this task event, all linear classifiers predicted control type significantly above chance except for the one using theta power from Cd theta for Monkey H (Figure 3.4C). The mean classification accuracies using different power features as inputs were also compared for these models (Figure 3.4D). Similar to the go cue, the classifier using all frequency bands has the best performance. No single frequency band provided a consistently higher prediction of control type at target acquisition.

### *Comparing Power Features from Different Regions of Interest*

Power features from all three regions provided significant information about control type at both the go cue and target acquisition. To determine whether combining information across regions increases control type predictability, we compared classification accuracy using power features from all frequency bands within individual regions and from all possible combinations of regions (Figure 3.5).



**Figure 3.5. Classification accuracy using power across all frequency bands from individual and combinations of regions of interest.** (A) Average 10-fold cross-validated classification accuracy and error across days for individual frequency bands and all combined bands at the go cue for Monkey H (left) and Monkey Y (right). Classifiers are presented in order of ascending accuracy. Chance accuracy shown as a dashed line. (B) Comparison of classification accuracies from classifiers using power features from different combinations of regions of interest as inputs at the go cue. Purple boxes indicate a significant difference via two-sample independent t-test after Bonferroni correction. (C, D) Same as (A, B) but for power estimates obtained at target acquisition.

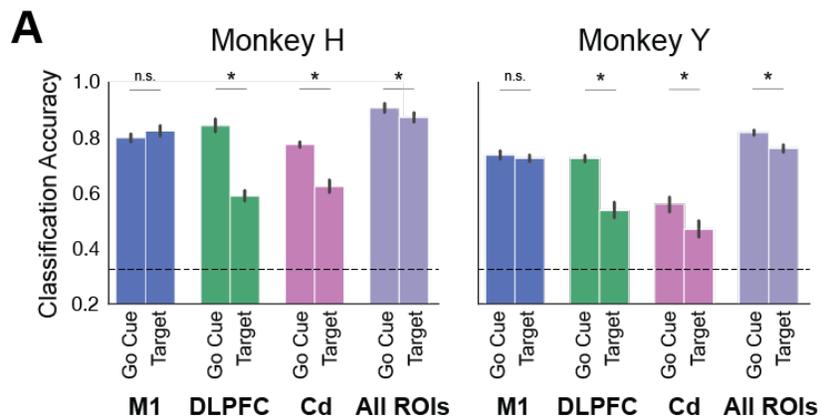
At the go cue, power features from DLPFC and M1 were significantly better at classifying control type than power features from Cd (Figure 3.5A, B). Furthermore, combining power features from DLPFC and M1 provided significantly higher classification performance than classifiers using power features from M1 only for both subjects and power features from DLPFC only for Monkey Y. While the classifier using power features from all three regions resulted in the highest classification accuracy for both subjects, there was no significant difference between classification accuracy when

using power features from all three regions and from M1 and DLPFC only for either subject. This suggests that most of the information about control type is contained in M1 and DLPFC at the go cue.

At target acquisition, the classifier using power features from M1 demonstrated significantly higher performance than the classifiers using power features from DLPFC or Cd (Figure 3.5C, D). Adding power features from DLPFC to M1 significantly improved performance for Monkey Y, but not for Monkey H. Similar to the go cue, classification accuracy was highest for the model using power features from all three regions but the difference between this model and the model using power features from only M1 and DLPFC was not significant. Overall, these results suggest that most of the information about control type is present in M1 power features at target acquisition.

### Comparing Task Events

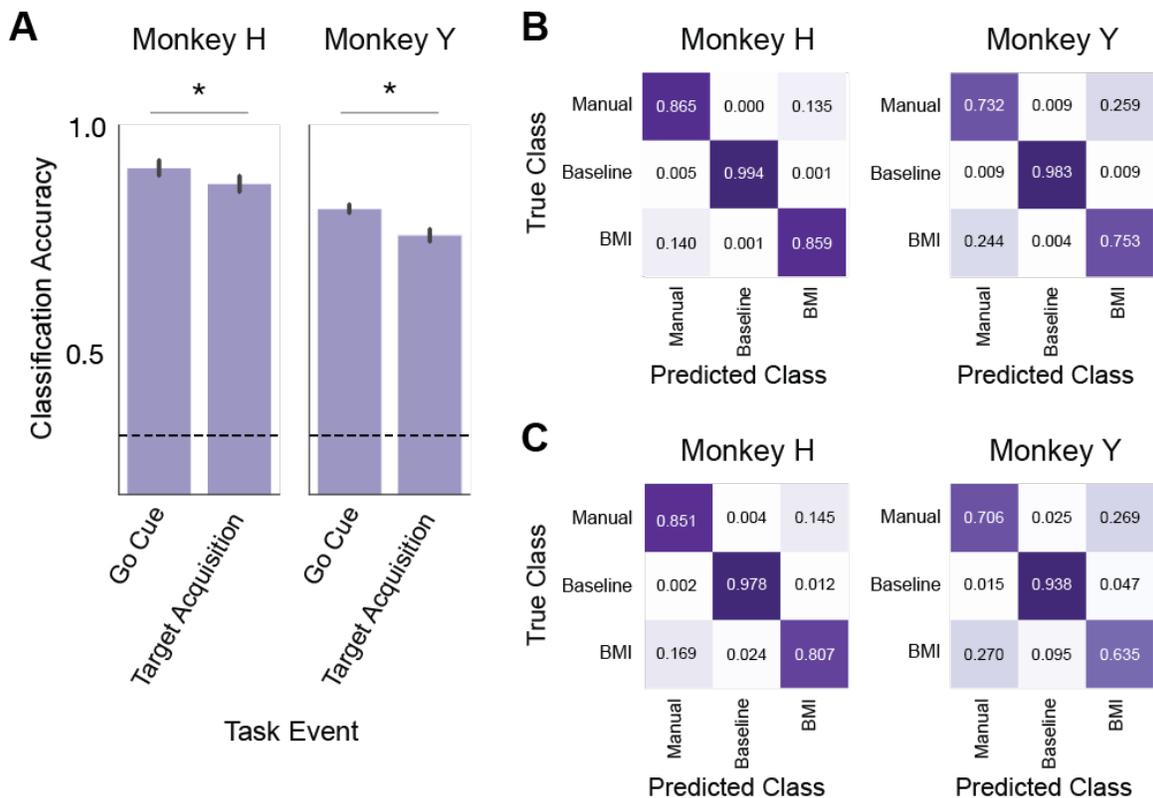
To compare classification performance using power features from different task events, we compared the classification accuracy from models using power features from M1, DLPFC, Cd, and all combined ROIs at the go cue and at target acquisition (Figure 3.6). Classification accuracy was significantly higher at the go cue than at the target for models using DLPFC, Cd, or all combined power features. Interestingly, there was no significant difference in classification accuracy between the two task events for the model using power features from M1.



**Figure 3.6. Classification performance using all frequency bands from M1, DLPFC, Cd, and all combined ROIs.** Classification accuracy is significantly higher at the go cue than at target acquisition for DLPFC (Monkey H,  $p = 3.4e-9$ ; Monkey Y,  $p = 3.1e-11$ ), Cd (Monkey H,  $p = 6.7e-9$ ; Monkey Y,  $p = 5.6e-9$ ), and all combined ROIs (Monkey H,  $p = 5.0e-3$ ; Monkey Y,  $p = 3.5e-5$ ), but not for M1 (Monkey H,  $p = 0.060$ ; Monkey Y,  $p = 0.352$ ).

## Misclassifications

The model with the highest classification accuracy for both subjects at both events included power features from all regions of interest. Classification accuracy was significantly higher at the go cue than at target acquisition when using power features from all regions (Figure 3.7A). The normalized confusion matrices of control type classification at the go cue and at target acquisition are depicted in Figure 3.7B and Figure 3.7C, respectively. At both task events, the main misclassifications are related to manual control and BMI control, indicating that power features during BMI control are more similar to manual control than to baseline. Thus, there is more distinction between BMI and manual power features at the go cue than at target acquisition.

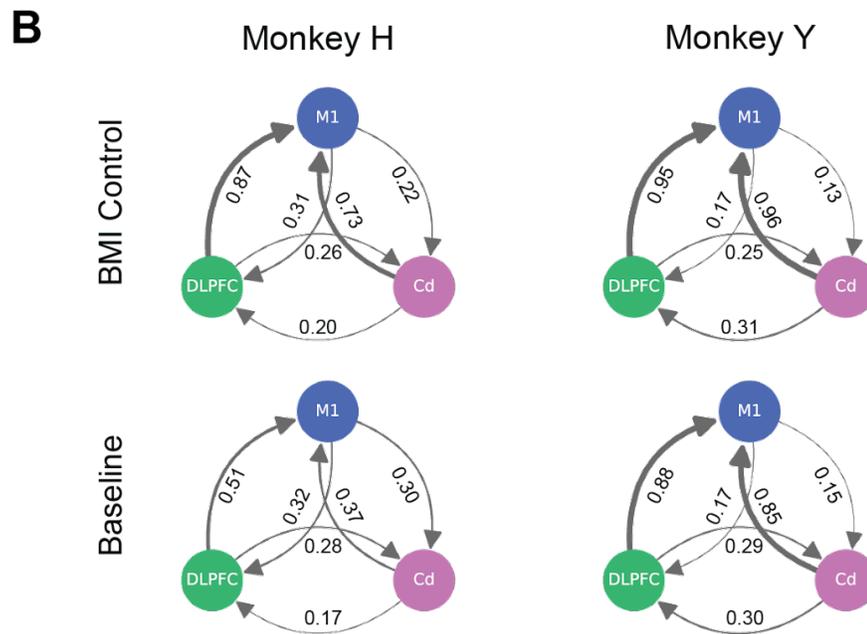
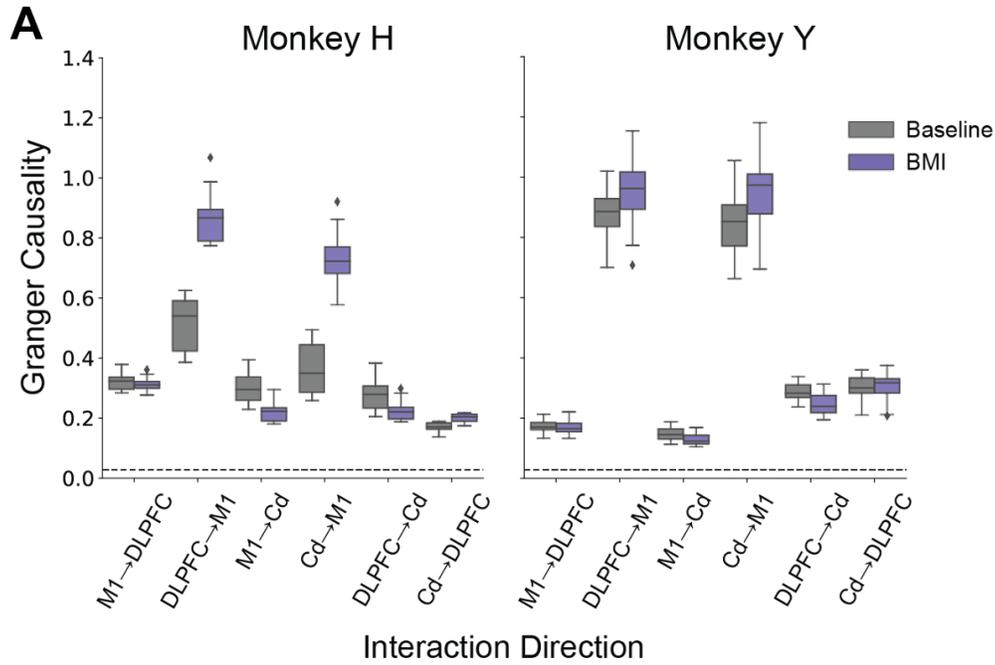


**Figure 3.7. Classification performance using all frequency bands from all ROIs.**

(A) Classification accuracy is higher at go cue than at target acquisition (Monkey H,  $p = 0.005$ ; Monkey Y,  $p = 3.4e5$ , two-sample independent t-test). (B, C) Average normalized confusion matrices of classifying control type across all days at (B) go cue and (C) target acquisition. Confusion matrix cells are normalized by true class.

### *3.3.3 Differences in Directed Functional Connectivity*

To investigate how these regions interact during neuroprosthetic control, we evaluated directed functional connectivity between the average LFP recorded within each region using Granger causality. Granger causality is a statistical test for determining whether one time series is useful in predicting another (see Methods for details). To obtain an estimate of spurious interactions, we first determined Granger causality estimates between the average LFP recorded during 500 ms segments of a baseline period and the average LFP recorded during 500 ms segments following the go cue in successfully initiated trials under BMI control. Because these signals were recorded at separate times in different tasks, any interactions that occur between regions would be artifactual. True interactions between regions within baseline and BMI control were considered significant if they were statistically different from this null distribution. To measure the Granger causality between pairs of regions in baseline and BMI control, we used the 500 ms LFP segments recorded during the baseline period and the 500 ms following the go cue of successfully initiated trial under BMI control, respectively. The average Granger causality value across trials within each day was determined and compared to that of the null distribution. All pairs of regions showed significant reciprocal interactions during both baseline and BMI control (Figure 3.8A). The direction and strength of these interactions is depicted by the direction and width of the arrows in Figure 3.8B.

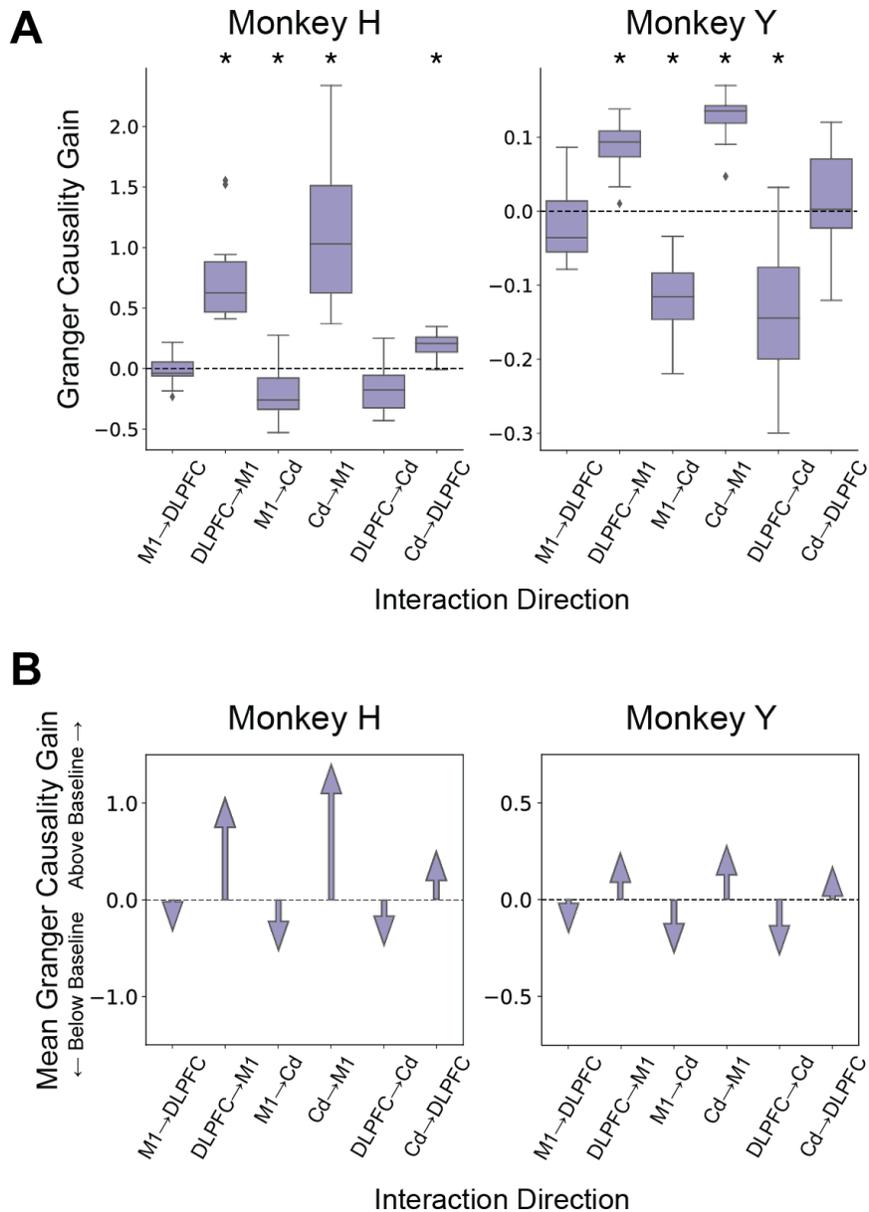


**Figure 3.8. Directed functional connectivity during baseline and BMI control.** Baseline estimates are based on 500 ms segments of LFP from baseline sessions. BMI estimates are based on the 500 ms following the go cue for successful trials. (A) Distribution of Granger causality estimates across baseline (gray) and BMI (purple) recording days. All interactions are significantly above the estimate of spurious interactions obtained by comparing LFP across session types (two-sample t-test, Bonferroni corrected). (B) Directed functional connectivity graphs for BMI control (top) and baseline (bottom) in which the direction and width of

arrows represents the mean direction and magnitude of significant interactions across days. The mean magnitude for each interaction across all recording sessions is also indicated next to the corresponding arrow.

### *Gain in Directed Functional Connectivity*

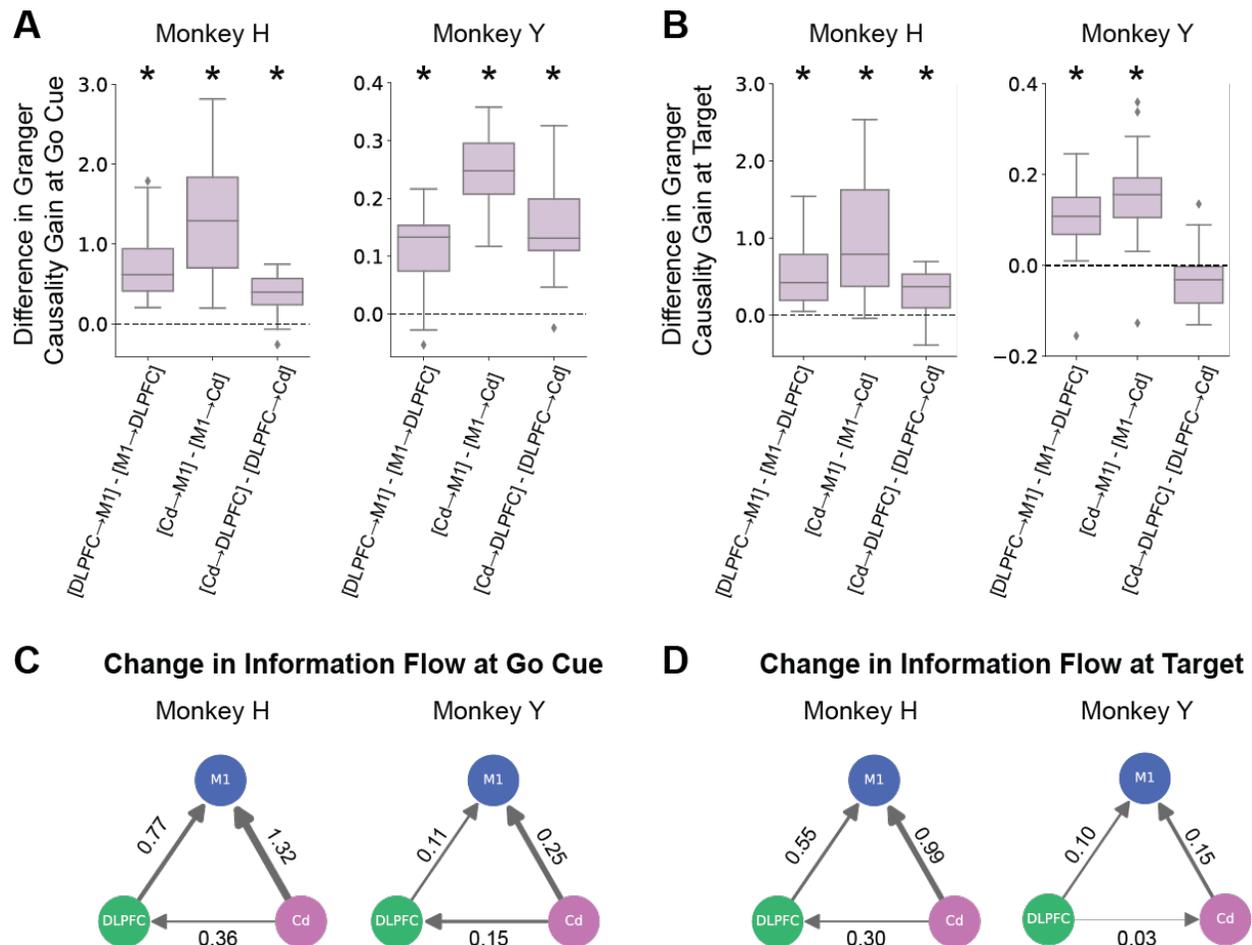
Cortico-cortical and cortico-striatal interactions involving M1, DLPFC, and Cd are involved in several planning, control, and learning processes (Haber, 2003, 2016). Reciprocal interactions between all pairs of regions were significantly above chance in both baseline and BMI control (Figure 3.8). To assess directed functional connectivity that is relevant to BMI control, we calculated the gain in Granger causality from baseline (see Methods for details) within each recording day (Figure 3.9A). Calculating the gain in Granger causality compares the interactions during BMI control to the interaction during baseline and accounts for variability across days. A Granger causality gain greater than zero indicates an increase from baseline, while a Granger causality gain less than zero indicates a decrease from baseline. The average Granger causality gain across days is depicted by the magnitude and direction of the arrows in Figure 3.9B. There is an increase in directed functional connectivity from DLPFC to M1 and Cd to M1, while directed functional connectivity decreases from M1 to Cd in both monkeys. There was also an increase from Cd to DLPFC for Monkey H and a decrease from DLPFC to Cd for Monkey Y.



**Figure 3.9. Gain in directed functional connectivity from baseline.** Baseline estimates are based on 500 ms segments of LFP from baseline sessions. BMI estimates are based on the 500 ms following the go cue for successfully initiated trials. (A) Distribution of the gain in Granger causality estimates above baseline across recording days. Granger causality gain estimates that are significantly different from zero are indicated with an asterisk (One-sample t-test, Bonferroni corrected). (B) Mean gain in directed functional connectivity from baseline. Magnitude of arrows represents the mean magnitude gain in interactions from baseline to BMI across days.

## Direction of Information Flow

Because reciprocal interactions all demonstrated opposite changes from baseline as indicated by the sign of the Granger causality gain, we took the difference between them to measure the magnitude of the change in direction of information flow from baseline to BMI (Figure 3.10). We subtracted the Granger causality gain in  $M1 \rightarrow DLPC$  from  $DLPC \rightarrow M1$ ,  $M1 \rightarrow Cd$  from  $Cd \rightarrow M1$ , and  $DLPC \rightarrow Cd$  from  $Cd \rightarrow DLPC$ . We found that information flow significantly increased from DLPC to M1, Cd to M1, and Cd to DLPC in the 500 ms period after the go cue of successfully initiated trials during BMI control (Figure 3.10A, C).



**Figure 3.10. Difference in gain in directed functional connectivity from baseline for reciprocal interactions.** Baseline estimates are based on 500 ms segments of LFP from baseline sessions. BMI estimates are based on the 500 ms following the go cue for successful trials (A, C) or the 500 ms before target acquisition (B, D). (A) Distribution of the difference in gain in Granger causality estimates above baseline between reciprocal interactions across recording days. Differences that are significantly different from zero are indicated with an asterisk (One-sample t-

test, Bonferroni corrected). (B) Same as (A) but for 500 ms prior to successful target acquisition during BMI control. (C) Direction and width of arrows represents the mean direction and magnitude of gain in information flow from baseline to BMI control 500 ms after the go cue on successfully initiated trials. The mean magnitude for each interaction across all recording sessions is also indicated next to the corresponding arrow scaled within each monkey. (D) Same as (C) but for 500 ms prior to successful target acquisition during BMI control.

The previously described Granger causality analyses were all performed using the 500 ms period following the go cue of successfully initiated trials. We repeated the analyses using the average LFP in the 500 ms preceding target achievement on successful BMI trials and found similar results to the period after the go cue (Figure 3.10B, D). There was an increase in information flow above baseline from DLPFC to M1 and Cd to M1 for both monkeys, and an increase in information flow from Cd to DLPFC in Monkey H. The increase in information flow from Cd to DLPFC was not present at the period prior to successful target acquisition for Monkey Y.

Overall, we find evidence of changes in directed functional connectivity between Cd, DLPFC, and M1 during neuroprosthetic control relative to baseline. Information flow increases from Cd to M1 and from DLPFC to M1 during both periods when the monkeys are engaged in neuroprosthetic control relative to a baseline period, suggesting that Cd and DLPFC also play a role in neuroprosthetic control.

### 3.4 Discussion

In this chapter, we demonstrate that a large-scale semi-chronic microdrive can be used to simultaneously record neural activity from cortical and subcortical regions during the real-time control of a motor cortical BMI. We leveraged these simultaneous recordings to characterize spectral power and directed functional connectivity within DLPFC, Cd, and M1 that are associated with BMI control.

Typical motor BMIs rely on chronically implanted microelectrode arrays, such as the Utah array. These recording devices have been used extensively for more than twenty years in nonhuman primates and people (Capogrosso et al., 2016; Collinger et al., 2013; Hochberg et al., 2012; Pandarinath et al., 2015; Sussillo et al., 2016; Velliste et al., 2008). Microelectrode arrays typically have a large number of electrodes distributed evenly across a few square millimeters of brain tissue at a fixed depth of only 1-2 mm, enabling the recording from hundreds of neurons with minimal set-up time post-implantation. While these arrays are excellent for obtaining stable, chronic

recordings for control of a BMI, they are not useful for recording from deep structures such as the striatum. Furthermore, if a unit is lost, the electrode placement is fixed and cannot be moved to find a new unit.

Linear multi-electrode probes, such as the Plexon V-Probe, are a popular alternative for studies requiring simultaneous recordings from large populations of neurons (Chiang et al., 2022; Knudsen & Wallis, 2021; Lara et al., 2009). These probes contain multiple channels along a single shank and can sample cells at different depths within the same region. The depth can be adjusted such that they can be used to record from surface areas or deeper structures. However, these probes are not ideal for BMI, as they are fragile and require several hours for insertion and tissue settling prior to every experiment. If the tissue is not given enough time to settle, the cells can drift from channel to channel, posing a difficult problem for a BMI decoder. Additionally, it is nearly impossible to record from the same cells from day to day.

In this study, we used a large-scale semi-chronic microdrive to obtain simultaneous recordings from cortical and subcortical structures while an animal controlled a BMI using a largely stable population of units recorded from motor cortex. These microdrives have been used to study of large-scale networks (Dotson et al., 2017; Qiao et al., 2016) but had not been used for BMI control. While these microdrive systems enabled semi-chronic simultaneous recordings from structures at different depths, they also have several limitations. In these experiments, a large fraction of electrodes broke when penetrating the dura (Monkey H,  $n = 0.36$ ; Monkey Y,  $n = 0.48$ ). While a durotomy may have improved this issue, it also leads to a higher risk of infection. We used a combination of platinum-iridium (Pt-Ir) and glass-coated tungsten (W) electrodes. An important note for future studies is that our Pt-Ir electrodes had a higher success rate in terms of penetrating the dura, however, the signal quality of the recordings obtained from these electrodes when compared to the W electrodes was significantly lower for both monkeys (Both Monkeys  $p < 1e-5$ , two-sample t-test). Due to the lack of single- and multi- unit recordings obtained in our regions of interest, particularly in Cd where most electrodes were Pt-Ir, we focused our analyses on LFP recordings.

Using the simultaneous recordings obtained during a BMI control task, manual control task, and baseline period, we identified control type relevant differences in spectral power across multiple frequency bands and multiple regions of interest. We trained and tested linear classifiers to distinguish power features from BMI control, manual control, and baseline. These analyses were repeated at the go cue and at target acquisition for successful BMI and manual trials. At the go cue, beta power tended to predict control type more accurately than any other individual frequency band (Figure

3.4A, B). In contrast, no single frequency band tended to outperform the others at target acquisition (Figure 3.5C, D). Beta oscillations have been implicated in top-down control (Buschman & Miller, 2007; Siegel et al., 2012), which may be evoked during the action selection processes that occur at the go cue. Furthermore, cortical beta power decreases during movement preparation and initiation (Leventhal et al., 2012; Pfurtscheller et al., 2003; Zhang et al., 2008). Thus, the higher classification accuracy of beta power features relative to other frequency bands at the go cue may be indicative of action selection and initiation strategies that are more dissociable between control types. At both the go cue and target acquisition, combining power features from all frequency bands led to higher classification accuracies (Figure 3.4). This result suggests that different frequency bands may carry non-overlapping or complementary information specific to the different control types.

When comparing classification accuracies between regions at the go cue, we found that power features from M1 and DLPFC were better predictors of control type than Cd (Figure 3.5). Furthermore, combining power features from M1 and DLPFC led to higher classification accuracy than either region individually. The classification accuracy obtained from this model was not significantly different from the model using power features from all three regions. DLPFC has been implicated in the implementation of behavioral rules, goal-directed behavior, and action planning (Boussaoud & Wise, 1993; Buschman & Miller, 2007; di Pellegrino & Wise, 1993; Fuster, 2000; Goldman-Rakic, 1996; Hoshi, 2006; Krämer et al., 2013; Miller & Cohen, 2001; Saito et al., 2005; Tanji et al., 2007; Tanji & Hoshi, 2008). These executive function processes are involved in the initiation and selection of movement that occurs at the go cue in center-out task under both manual and BMI control. The information contained in DLPFC power features was predictive of control type, suggesting that there are differential top-down cognitive strategies between BMI control, manual control, and a baseline rest period. In contrast, the classifier using power features from M1 performed significantly better than either the classifier using power features from DLPFC or Cd at target acquisition (Figure 3.5). Thus, there is more information about control type in M1 than either other region at this event. After the initial action selection and initiation, DLPFC may become less involved in BMI control or the activity in DLPFC may be less distinguishable between the different control types.

The main misclassifications occurred between BMI and manual power features, suggesting that power features during BMI and manual control are more similar to one another than either are to baseline (Figure 3.7). Furthermore, classification performance was greater at the go cue than at target acquisition for models using power features from DLPFC, Cd, and all three regions combined (Figure 3.6). Thus, there is less confusion between BMI and manual at the go cue than at target

acquisition which supports the hypothesis that the DLPFC and Cd are involved in action planning and selection and that neural representations of these processes are distinct during BMI control and manual control. Interestingly, this difference in classification accuracy at different task events was not true for the classifier using power features from M1 alone (Figure 3.6). This suggests that the control execution strategy implemented by the motor cortex is distinct at all task events, whereas the top-down action planning and selection strategies implemented by DLPFC and Cd is more distinct at the go cue than at target acquisition.

We also investigated how M1, DLPFC, and Cd interact during BMI control and baseline. All three regions showed significant reciprocal interactions with one another when compared to chance (Figure 3.8). To determine which interactions were important for BMI control, we found the gain in each interaction in BMI control relative to baseline and identified changes in directed functional connectivity. Reciprocal interactions tended to differ from baseline in opposite directions, so we measured the difference between them to quantify the magnitude of the change in the direction of information flow. We found that information flow increases from Cd to M1 and from DLPFC to M1 during neuroprosthetic control relative to baseline. Information flow also increased from Cd to DLPFC at the go cue for both monkeys and at target acquisition for Monkey H (Figure 3.10).

An important note is that granger causality reflects directed functional, and not anatomical, connectivity. Thus, these results indicate that there is increased information flow from Cd to M1, DLPFC to M1, and Cd to DLPFC during BMI control relative to baseline, and we make no claims about the mechanism of this increase in information flow. Despite both M1 and DLPFC projecting directly to Cd (Haber 2003), we see that information flows in the opposite direction. The cortex and basal ganglia are interconnected in cortico-striatal 'loops' (Delong et al., 1984; Kimura & Graybiel, 1995; Middleton & Strick, 2002). These cortico-striatal loops have been implicated in several aspects of goal-directed behaviors and control of volitional movements, including movement initiation, action selection, strategic planning, and motivation (Haber, 2003, 2016). Our findings that information flows from Cd to cortex align with results from previous studies demonstrating that Cd trains the cortex (Antzoulatos & Miller, 2014; Pasupathy & Miller, 2005). These studies found that task-relevant information was detected in Cd prior to cortex and that the direction of functional connectivity was from Cd to cortex. Cd activity has been implicated in action selection and control of volitional movement (Graybiel et al., 1994; Graybiel, 1998; Graybiel & Grafton, 2015; Hikosaka et al., 2000; Packard & Knowlton, 2002). Thus, it is likely that this information flow from Cd to cortex represents task-relevant information. Information flow also increased between cortical regions, from DLPFC to M1. This

further supports the hypothesis that DLPFC is involved in the top-down selection of neuroprosthetic actions.

Our studies focused on changes in LFP related to BMI control. Future work investigating how these differences evolve over learning must be done to understand how this neuroprosthetic control is learned over time. Neuroprosthetic skill learning is thought to rely on multiple parallel learning mechanisms: (1) the fast reinforcement of cortical activity that permits the re-entrance of particular cortical population dynamics that naturally produce desired outcomes, and (2) slower reinforcement that leads to the refinement of cortical population dynamics and a more reliable production of neural trajectories to drive skillful behavior on-demand (Athalye et al., 2020; Zippi, You et al., 2021). Previous work has demonstrated differential cortico-striatal plasticity during fast and slow motor skill learning (Costa et al., 2004) and different regions may play important roles in these neural reinforcement processes. Furthermore, M1, DLPFC, and Cd exist within larger cortico-basal ganglia-thalamic loops (Parent & Hazrati, 1995). Thus, these processes likely involve different nodes within this larger circuitry. Understanding how these regions interact together to learn and implement neuroprosthetic control will be important for understanding how neural activity facilitates and constrains learning and may be useful for developing neurobiologically informed neuroprosthetic devices.

## Chapter 4: Modulation of Value Encoding in Caudate and Anterior Cingulate Cortex through Caudate Microstimulation

The ability to evaluate options and use these values to inform choices is indispensable for adaptive, healthy decision-making behavior. Dysfunctional decision-making and reward processing are associated with a multitude of neuropsychiatric disorders, including anxiety, depression, and addiction. Multiple areas of the brain are involved in the neural processing of information related to decision-making, including the caudate (dorsomedial striatum) and the anterior cingulate cortex (ACC). The head of the caudate (Cd) is a region that contains flexible representations of value and receives value-related information from multiple cortical structures, as well as modulatory dopaminergic input from ventral striatum. The ACC also contains value and error information and is particularly implicated in flexible decision-making. In this chapter, we record neural activity in both brain areas during a two-armed bandit task and investigate how neural representations of task variables are modulated when Cd microstimulation is administered. Microstimulation delivered during the deliberation period of forced-choice trials to a particular stimulus significantly increases the likelihood of selecting that option in free-choice trials, and we find that the neural correlates of value signals reflect this bias. These results suggest that stimulation-based therapies may be used to regulate valuation of choices in neuropsychiatric patients. This work was done in collaboration with Samantha R. Santacruz, Joni Wallis, and Jose M. Carmena.

### 4.1 Introduction

Value-based decision-making involves an assessment of value associated with available items and the actions required to obtain them. The inability to appropriately evaluate one's options and use these values to inform their decisions is associated with a number of neuropsychiatric disorders (Hartley & Phelps, 2012; Paulus & Yu, 2012; Shepherd, 2013). Difficulties in finding effective behavioral or pharmaceutical treatments for these disorders has led to the exploration of alternative methods, such as deep brain stimulation (Holtzheimer et al., 2017; Holtzheimer & Mayberg, 2011; Insel, 2012; Kessler et al., 2005; Mayberg et al., 2005). Many of these stimulation-based approaches target the frontolimbic cortex, whose dysfunction has been associated with neuropsychiatric disorders (Fernando & Robbins, 2011).

Several studies have demonstrated that the caudate nucleus of the striatum (Cd) and the anterior cingulate cortex (ACC) encode task-relevant values for decision-making and action-selection (Kennerley et al., 2009; Kennerley & Wallis, 2009; Kennerley & Walton, 2011; Lau & Glimcher, 2008; Samejima, 2005). These brain regions are

essential in mediating how sensory information, experience, and motivation guide choice-behavior. In particular, Cd encodes values of alternative choices prior to selection (Balewski et al., 2022; H. F. Kim & Hikosaka, 2013; Lau & Glimcher, 2008; Samejima, 2005) and has been implicated in orienting value-guided fixations (Hikosaka et al., 2006; Watanabe & Hikosaka, 2005). ACC has been shown to play an important role in reinforcement-guided behavior by monitoring the history of action outcomes (Kennerley et al., 2006; Monosov, 2017; Walton et al., 2007). Previous work has shown that Cd microstimulation paired with a particular stimulus causally increases the likelihood of selecting that stimulus during a value-based decision-making task (Santacruz et al., 2017). This effect was shown to be stimulus-dependent and action-independent, indicating that Cd microstimulation can change the value of an associated stimulus. Electrical microstimulation is known to induce neural plasticity (Jackson et al., 2006; Madhavan et al., 2007). Thus, it is possible that electrical stimulation of Cd alters the neural representation of value.

In this work, we investigate the neural representation of value information during flexible decision-making behavior. Due to the prominent roles that Cd and ACC have in decision-making, we record neural activity in these nuclei while nonhuman primate (NHP) subjects perform a two-arm bandit decision-making task. The work leverages previous work in which it was demonstrated that electrical stimulation in the Cd of NHP subjects can preferentially increase the value of an associated choice (Santacruz et al., 2017). Here, we investigate the neural correlates of this functional change in behavior. We find that Cd microstimulation recruits more neurons to represent task-relevant stimulus values. These results further our understanding of the neural circuitry underlying value-based decision-making and support potential future applications of microstimulation to correct maladaptive plasticity underlying dysfunctional decision-making.

## 4.2 Methods

### 4.2.1 Surgery

Two rhesus macaques were implanted unilaterally with custom-machined recording chambers enabling access to caudate (Cd) and anterior cingulate cortex (ACC). Chamber positions were calculated based on images obtained from 1.5-T magnetic resonance imaging (MRI) scans of each subject's brain. For Monkey L, we used multi-channel electrodes (V-Probes, Plexon, Dallas, TX) along with standard methods for acute neurophysiology for neural recording. Single-channel stimulation electrodes were lowered into the neural tissue for Monkey L following the same techniques as multi-channel electrodes. For Monkey M, we used a custom semi-chronic microdrive

array to record and stimulate from moveable single microelectrodes that were chronically implanted (Gray Matter Research, Bozeman, MT). Two Platinum-Iridium microelectrodes (Alpha Omega; Microprobes) were used to administer stimulation in a bipolar manner. For Monkey L, on each experimental day electrodes were lowered manually using custom-built microdrives to a target depth in the head of the caudate. For Monkey M, electrodes were positioned in the caudate chronically.

#### *4.2.2 Stimulation*

The stimulation parameters used in this study were the same as those described in Santacruz et al., 2017 and are consistent with previous studies using electrical stimulation in nonhuman primates (Ditterich et al., 2003; Hanks et al., 2006; Nakamura & Hikosaka, 2006; Santacruz et al., 2017; Williams & Eskandar, 2006). Microstimulation pulse trains consisted of a series of charged-balanced biphasic pulses with no inter-pulse interval and a cathodal leading phase. Each phase was 200  $\mu$ s in duration and the pulse frequency was 200 Hz. Stimulation current amplitude in the range of 100 – 250  $\mu$ A. Stimulation trains lasted the duration of the center-hold period, 1000, during forced-choices trials as described in the behavioral task. This design ensured that the subject had negligible movement during the stimulation epoch.

#### *4.2.3 Behavioral Task*

The behavioral task has been described in detail elsewhere (Santacruz et al., 2017). Briefly, two subjects (Monkey L and Monkey M) used a joystick to control a computer cursor and select colored circular targets on a computer screen positioned in front of them. The color of each target was associated with the probability of reward. The colors changed each session and could appear on either side of the screen.

This task consisted of two trial types: (1) free-choice trials and (2) instructed trials. In free-choice trials, during the center-hold period the subjects held the cursor in the center target for 1000 ms. Two peripheral, colored targets were shown during this time. At the end of a successful hold, the subject freely moved the cursor to the target of choice and then completed a target-hold to indicate their selection for another 1000 ms. In instructed trials, only one peripheral target was presented during the center-hold period. After a successful hold, the subject moved the cursor to that target and completed a target-hold. A trial was considered successful if the subject completed the 1000 ms center-hold followed by holding at a peripheral target for 1000 ms within a 10 s period. The same trial was repeated up to 10 times until it was successfully completed and the subject advanced to the next trial.

Monkey L was trained in a two-target version of the task with the low-value target being assigned 40% likelihood of reward and high-value target being assigned 80% likelihood of reward. Monkey M was trained in a three-target version of the task in which the low-value target was assigned 35% likelihood of reward, the medium-value target was assigned 60% likelihood of reward, and the high-value target was assigned 85% likelihood of reward. In both versions, only two target options were displayed at once and the animals learned the relative probabilities over the course of the experiment.

#### 4.2.4 Reinforcement Learning Models

The subjects' free-choice behavior was modeled using Q-learning, a model-free reinforcement learning (RL) algorithm. The learning rate,  $\alpha$ , estimates how much the value of a choice is updated by new information at each time step. The inverse temperature,  $\beta$ , estimates how much expected rewards affect the probability of selecting a stimulus. The standard Q-learning algorithm (Sutton & Barto, 1998) consists of the following value update equations:

$$Q(t) = Q(t + 1) + \alpha\delta(t)$$

$$\delta(t) = r(t) - Q(t - 1)$$

Using a soft-max decision rule, the probability of selecting the low-value ( $a_{LV}$ ) target over the high-value ( $a_{HV}$ ) target is:

$$P(a_{LV})(t) | Q_{LV}(t), Q_{HV}(t) = \frac{1}{1 + \exp(\beta[Q_{HV}(t) - Q_{LV}(t)])}$$

where the variables  $Q_{LV}(t)$  and  $Q_{HV}(t)$  represent the values of the low-value and high-value targets at time  $t$ , respectively. The value update equations were updated on both instructed and free-choice trials, but the decision rule was only simulated for free-choice trials. The values obtained from the value update equations were used to assess neural correlates (see next section).

#### 4.2.5 Neural Correlates

Single-unit neuron activity was recorded throughout the decision-making task from Cd and ACC. This neural activity was recorded in naïve conditions, as well as with high-frequency electrical stimulation intervention where stimulation was delivered to Cd. We used multiple linear regression to determine how the activity of all recorded, well-isolated units co-varied with value, as well as movement variables (i.e., movement

time and reaction time), reward, and choice. The responses of individual neurons were fit using the following multiple linear regression:

$$y = \beta_1 R + \beta_2 C + \beta_3 MT + \beta_4 RT + \beta_5 Q_{LV} + \beta_6 Q_{MV} + \beta_7 Q_{HV}$$

where  $y$  is the firing rate in the window [0,400) ms following target presentation,  $R$  is for the presence or absence of reward,  $C$  is the chosen target color,  $MT$  is the movement time, and  $RT$  is the reaction time. The variables  $Q_{LV}$ ,  $Q_{MV}$ , and  $Q_{HV}$  represent the dynamic stimulus values estimates for the low-value, medium-value, and high-value target colors as determined by Q-learning value update equations. For Monkey L, there was no medium-value target and the  $Q_{MV}$  variable was excluded from the above equation. Statistical significance of regressors was determined using incremental F-statistic with a significance level of 0.05. Units were classified as value-coding if their activity co-varied with any Q value, regardless of whether they also co-varied with other regressors. Reward, choice, movement time, and reaction time neurons were categorized by the associated regressor with the largest slope. Units that did not significantly co-varying with any of these regressors were labeled as non-coding.

#### 4.2.6 Quantification & Statistical Analyses

All analyses were performed in Python using custom-written routines that utilize publicly available Python libraries. Bar charts report averages and standard error of the mean across sessions. Pie charts report the fraction of units within the recorded population averaged across sessions. Two-sample t-tests were used to compare distributions across stimulation sessions (Monkey L,  $n = 8$  sessions; Monkey M,  $n = 10$  sessions) and sham sessions (Monkey L,  $n = 12$  sessions; Monkey M,  $n = 11$  sessions).

### 4.3 Results

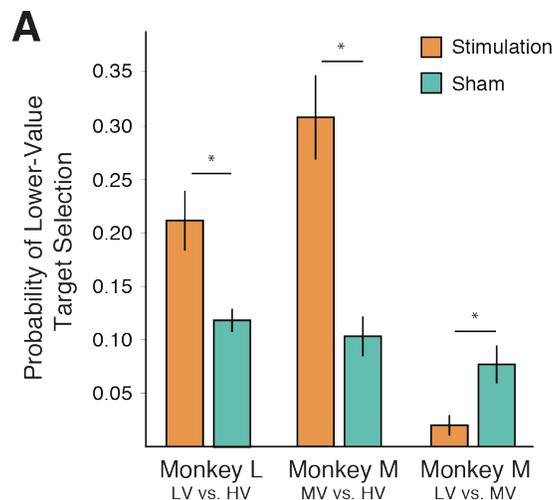
Two rhesus macaques (Monkey L and Monkey M) performed a probabilistic reward decision-making task in which they learned to choose between colored targets associated with unique reward probabilities (Figure 4.1) while neural activity was recorded from Cd and ACC. New target colors were selected and arbitrarily assigned reward probabilities in each session (see Methods for details). In each trial, each target could be presented on the left or right of the screen, requiring subjects to learn to associate the target color and not the action with a specific reward probability. Both subjects quickly learned to select the higher-value target with greater frequency than the lower-value target on free-choice trials; however, neither subject exclusively selected the higher-value target color after initial learning.



probabilities associated with the two versions of the task. (D) Description of task structure within the three blocks.

### 4.3.1 Caudate Microstimulation Biases Target Selection

To assess whether electrical stimulation significantly biased the decision-making policy towards the option paired with stimulation, we compared the probability of selecting the lower-value target on free-choice trials in Block A' (Figure 4.2). For Monkey L, stimulation was always paired with the low-value target. The likelihood of selecting this target on trials after stimulation was significantly higher than on trials after a sham control. For Monkey M, stimulation was paired with the middle-value target. On trials where the animal was presented with the middle-value and high-value targets, the middle-value target is the lower-value choice. Electrical stimulation increased the likelihood of selecting the middle-value target on these trials. In contrast, on trials where the animal was presented with the middle-value and low-value targets, the middle-value target represented the higher-value choice and electrical stimulation decreased the likelihood of selecting the middle value target, or lower-value option, on these trials. Overall, when given two options with unique reward probabilities, stimulation paired with a particular option increased the likelihood of selecting that option when compared to a sham control.



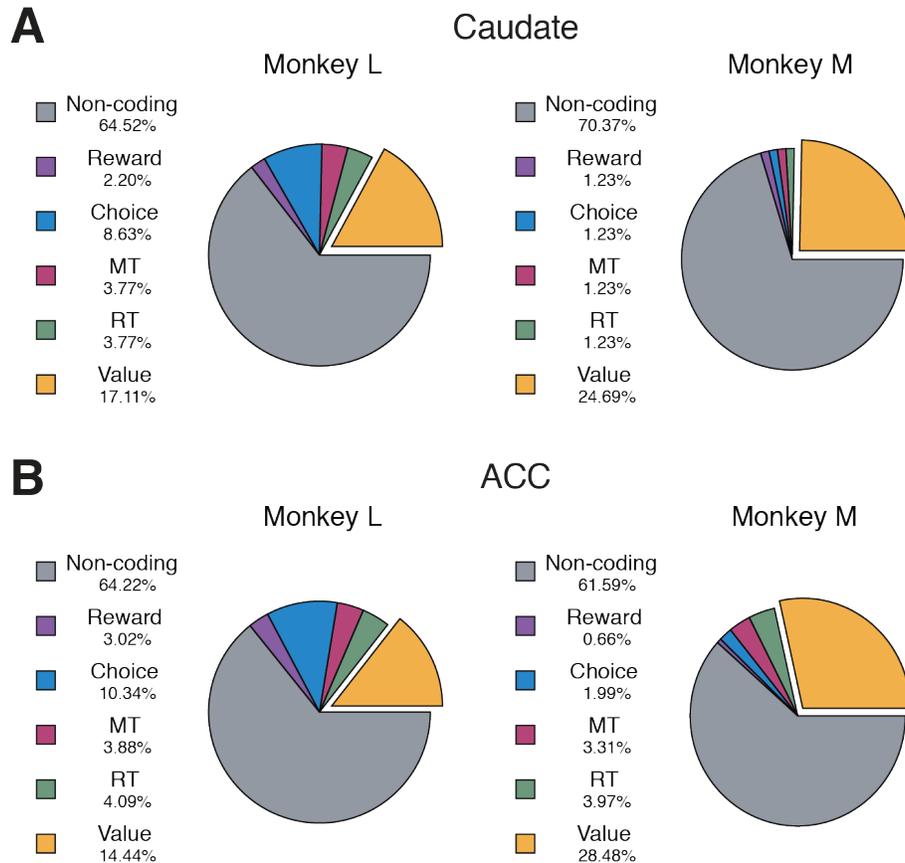
**Figure 4.2. Caudate microstimulation paired with a particular option increases the likelihood of selecting that option.** (A) The probability of selecting the lower-value target on free-choice trials after stimulation (orange) when stimulation is paired with the lower-value target or a sham control (teal). For Monkey L, stimulation was paired with the low-value (LV) target. The probability of selecting the lower-value target represents the probability that he selected the LV target over the high-value (HV) target on free-choice trials. For Monkey M, stimulation was paired with the middle-value (MV) target. On trials between the

MV and HV target and between the LV and MV target, the probability of selecting the lower-value target represents the probability of selecting the MV and LV target, respectively.

#### 4.3.2 Value Encoding in Caudate and ACC

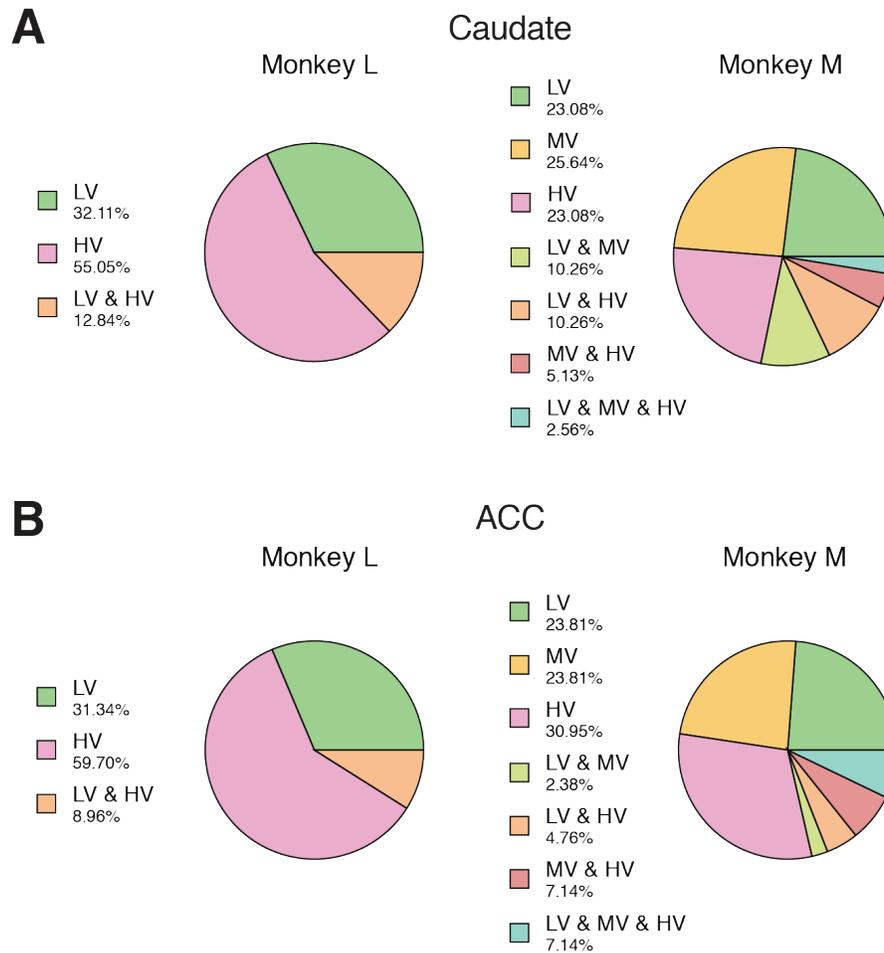
We recorded neural activity in Cd and ACC during the choice task and isolated individual units (Monkey L:  $n = 637$  Cd units,  $n = 464$  ACC units; Monkey M:  $n = 162$  Cd units,  $n = 151$  ACC units). Because Cd microstimulation paired with a particular stimulus resulted in an increased likelihood of selecting that stimulus, we hypothesized that this stimulation affects the neural computations of value. To test this theory, we first estimated the trial-by-trial value associated with each target using a Q-learning algorithm fit to the subject's decision-making behavior (see Methods for details). To ensure that the values had been learned, the last two-thirds of trials from Block A and trials from Block B (i.e., 133 trials for Monkey L and 200 trials for Monkey M) were used to fit the value update equations and estimate internal representations of value.

Linear regression of neural firing rates at the time of the presentation of the targets from each unit was performed using the values derived for each target from the Q-learning model as regressors. Additionally, we included reward (i.e., presence or absence), choice (i.e., chosen value), and motor covariates, including reaction time (RT) and movement time (MT), in the regression. We considered a unit to be value-coding if its firing rate co-varied with any target value, regardless of whether it also co-varied with other regressors. Reward-coding, choice-coding, MT-coding, and RT-coding units were categorized according to the largest significant coefficient. Units were considered non-coding if their firing rate activity did not significantly co-vary with any of the regressors. Statistical significance of regressors was determined using incremental F-statistic. We found that neural activity in both ACC and Cd encodes all stimulus values (Figure 4.3). We identified that 17.11% of Cd units and 14.44% of ACC units Monkey L and 24.69% of Cd units and 28.48% of ACC units in Monkey M encoded value, supporting the idea that these regions are involved in representing task-relevant values.



**Figure 4.3. Neural correlates of task-relevant variables in Cd and ACC.** (A) Percentage of units in Cd whose firing rate is significantly modulated by reward, choice, movement time (MT), reaction time (RT), or value based on linear regression analysis for Monkey L (left, N = 637) and Monkey M (right, N = 162). (B) Percentage of units in ACC whose firing rate is significantly modulated by each variable for Monkey L (left, N = 464) and Monkey M (right, N = 151). Statistical significance of regressors was determined using incremental F-statistic with a significance level of 0.05.

Within the value-coding units, units could encode individual target values or multiple target values. The distribution of value-coding neurons (i.e., the dark yellow section in Figure 4.3) is depicted in Figure 4.4. Though many units' firing rates co-varied with the values of multiple stimuli, it was more common for them to be correlated with a single target value. Furthermore, each individual target-value is represented by more than a quarter of value-coding neurons.

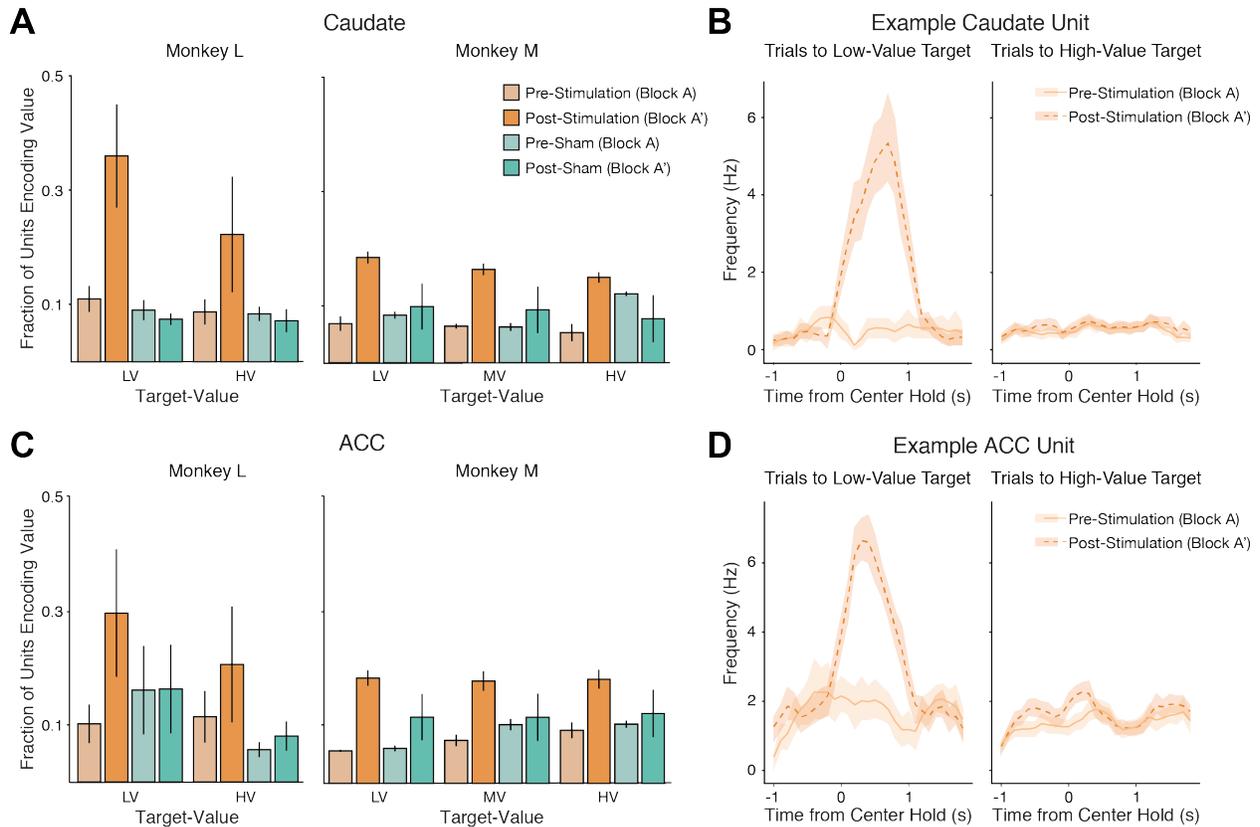


**Figure 4.4. Distributions of specific target-value associated units from value-coding units in Cd and ACC.** These pie charts expand upon the value-coding portion of the pie charts from Figure 4.3 in which neurons could be responsive to any individual or combination of target values. (A) Percentage of value-coding units in Cd whose firing rate co-varies with  $Q_{low}$ ,  $Q_{med}$ , and  $Q_{high}$ , and combinations thereof for Monkey L (left) and Monkey M (right). (B) Same as A but for value-coding units in ACC.

#### 4.3.3 Recruitment of Value Encoding Units with Caudate Microstimulation

To test the hypothesis that value-coding during the deliberation hold time is modulated by Cd microstimulation, we compared the fraction of units encoding value before and after stimulation and before and after the sham control (Figure 4.5). The same multiple linear regression analysis described previously was performed separately for Block A and Block A' from each session. If a unit's firing rate significantly co-varied with the value of a particular target, it was considered to encode that value.

Thus, units with firing rates that co-varied with multiple target values were included in the fraction of value-coding units for all the corresponding targets.



**Figure 4.5. Recruitment of value-coding units.** (A) Fraction of units in Cd whose firing rate is significantly predicted by stimulus-value for Monkey L (left) and Monkey M (right) in Block A and Block A'. Stimulation sessions are represented in orange and sham sessions are represented in teal. Block A is represented by lighter colors and Block A' is represented by darker colors. (B) Peri-stimulus time histogram (PSTH) from an example Cd unit from Monkey L whose firing rate was not modulated by either target before stimulation but became modulated by the LV target after stimulation. (C) Fraction of units in ACC whose firing rate is significantly predicted by stimulus-value for Monkey L (left) and Monkey M (right) in Block A and Block A'. (D) PSTH from example ACC unit in same conditions described in (B).

We compared the fraction of value-coding units associated with each of the possible target values within Cd and ACC before and after Cd stimulation and before and after a sham control (Figure 4.4). While there tends to be an increase in the fraction of value-coding units in Block A' compared to Block A after Cd stimulation is administered during the hold period of Block B, the increase from Block A to Block A'

was only statistically significant for the Cd units encoding the low-value target across stimulation sessions in Monkey L. The mean and standard error of the fraction of units encoding value for each target value in each block, as well as the results comparing Block A and Block A' are listed in Table 4.1 for Cd and Table 4.2 for ACC for both Monkey L and Monkey M.

		Monkey L		Monkey M		
		LV	HV	LV	MV	HV
<b>Stim</b>	<i>Block A</i>	0.11 ± 0.02	0.08 ± 0.02	0.08 ± 0.03	0.06 ± 0.01	0.12 ± 0.03
	<i>Block A'</i>	0.36 ± 0.07	0.22 ± 0.08	0.10 ± 0.02	0.09 ± 0.02	0.08 ± 0.02
	<i>T-Statistic</i>	-3.01 (p=0.020)	-1.13 (p=0.295)	-0.35 (p=0.733)	-1.11 (p=0.300)	1.01 (p=0.339)
<b>Sham</b>	<i>Block A</i>	0.09 ± 0.02	0.08 ± 0.01	0.07 ± 0.01	0.06 ± 0.01	0.05 ± 0.01
	<i>Block A'</i>	0.08 ± 0.01	0.07 ± 0.02	0.18 ± 0.08	0.16 ± 0.08	0.149 ± 0.08
	<i>T-Statistic</i>	0.70 (p=0.500)	0.76 (p=0.514)	-1.32 (p=0.215)	-1.18 (p=0.265)	-1.11 (p=0.293)

**Table 4.1. Comparison of the fraction of Cd units encoding value before and after stimulation or sham control.** Values represent the mean and standard error rounded to the nearest hundredth. Comparisons made using a two-sample paired t-test between Block A and Block A' for each target.

		Monkey L		Monkey M		
		LV	HV	LV	MV	HV
<b>Stim</b>	<i>Block A</i>	0.10 ± 0.03	0.11 ± 0.04	0.06 ± 0.00	0.10 ± 0.02	0.10 ± 0.03
	<i>Block A'</i>	0.30 ± 0.09	0.21 ± 0.08	0.12 ± 0.03	0.11 ± 0.03	0.122 ± 0.03
	<i>T-Statistic</i>	-1.95 (p=0.092)	-0.71 (p=0.502)	-1.79 (p=0.107)	-0.27 (p=0.794)	-0.45 (p=0.662)
<b>Sham</b>	<i>Block A</i>	0.16 ± 0.08	0.06 ± 0.01	0.06 ± 0.01	0.08 ± 0.02	0.09 ± 0.01
	<i>Block A'</i>	0.16 ± 0.08	0.08 ± 0.03	0.18 ± 0.08	0.18 ± 0.08	0.18 ± 0.08
	<i>T-Statistic</i>	-0.28 (p=0.783)	-0.76 (p=0.464)	-1.47 (p=0.172)	-1.1 (p=0.296)	-1.02 (p=0.330)

**Table 4.2. Comparison of the fraction of ACC units encoding value before and after stimulation or sham control.** Values represent the mean and standard error rounded to the nearest hundredth. Comparisons made using a two-sample paired t-test between Block A and Block A' for each target.

On average, there was an increase in the fraction of value-coding units across all target-values after Cd stimulation. However, this effect was only statistically

significant for Cd units encoding the LV target stimulus value in Monkey L. Example peri-stimulus time histograms for units whose activity was not modulated at the center hold prior to stimulation and became modulated after stimulation are shown in Figure 4.4B and Figure 4.4D for a Cd and ACC unit, respectively.

#### 4.4 Discussion

In this chapter, we demonstrate that Cd microstimulation paired with a particular stimulus in a flexible decision-making task altered value-based choice behavior. In one animal, there was a significant increase in the fraction of the neural population recorded from Cd that encoded the subjective stimulus value of the stimulus paired with this microstimulation. Across both animals, we observed an increase in the fraction of units in both Cd and ACC encoding task-relevant stimulus values, though these changes were not statistically significant. This work leveraged large-scale ensemble electrophysiology, electrical microstimulation, and behavior with the goal of elucidating neural mechanisms involved in the representation of stimulus value for decision-making behavior.

Value computations used to guide flexible decision-making behavior can be modulated by a number of factors, including risk and uncertainty (Rangel et al., 2008). Past work has found that increases in ACC activity are associated with increased levels of uncertainty, task difficulty, or probability of making an error (Behrens et al., 2007; Johnston et al., 2007). Similarly, an increased activation of both Cd and ACC was found during task-switching (Premereur et al., 2018), which requires flexibly shifting from one set of rules to another in response to changing environmental contingencies and may also contribute to an increased probability of making an error. Thus, it is possible that this increase in activity is the result of a larger proportion of the neural population in both Cd and ACC being modulated by task-relevant value information under periods of increased uncertainty or task difficulty.

Value has also been demonstrated to play a critical role in guiding attention (Anderson, 2016). When stimuli are associated with a promising reward outcome, they are given attentional priority. For example, if visual aspects of specific stimuli, such as color or orientation, are sufficient to differentiate value, robust value-driven attentional biases towards the reward-relevant stimulus occur (Anderson et al., 2011; Anderson & Yantis, 2012; Laurent et al., 2015). The basal ganglia has been implicated in guiding saccades to rewarded targets (Hikosaka et al., 2006) and reconfiguring attention to enhance it towards reward-relevant stimuli (Boroujeni et al., 2020). A recent study also found that activity in Cd aligns with value-guided saccades and is predictive of value-based choice behavior (Balewski et al., 2022). Thus, Cd microstimulation may

modify value computations used to guide decision-making by manipulating value-driven attentional biases.

In addition to the valuation of stimuli, value-based decision-making involves the process of selecting between these values (Rangel et al., 2008). A study investigating value representations in orbitofrontal cortex found activity in this region dynamically represents choice options, with neural representations alternating between states associated with each available option during a deliberation period (Rich & Wallis, 2016). Future work investigating the effect of Cd microstimulation on deliberation time and the dynamic representation of available options could provide further insight into how value information is integrated within the corticostriatal circuitry and used to guide decision-making behavior. Additionally, future work investigating the temporal specificity of value coding in Cd and ACC and how this may change as a result of electrical microstimulation may also provide further insight into the mechanism by which Cd microstimulation alters choice behavior.

Overall, we demonstrated that high-frequency stimulation delivered to the Cd can modulate decision-making processes. This change is possibly the result of changing underlying value representations in both Cd and ACC through the recruitment of value-coding neurons. Dysfunctional corticostriatal circuitry has been implicated in a number of neuropsychiatric disorders and has increasingly been the target of stimulation-based neurotherapies (Creed et al., 2015; Scangos et al., 2021; Shanechi, 2019). An inability to appropriately evaluate stimuli and use these values to inform decisions lies at the core of many of these neuropsychiatric disorders (Hartley & Phelps, 2012; Paulus & Yu, 2012; Shepherd, 2013). Our results suggest that electrical stimulation may offer a therapeutic approach to help regulate valuation in patients with impaired decision-making abilities.

## Chapter 5: Conclusions and Open Questions

BIMs are an emerging technology with great promise for future clinical therapies. The primary motivation for this work was to develop a better understanding of the neurobiology of BIMs. Knowledge of how the brain interacts with these systems, including both the neural adaptations that occur as the user learns to control a BIM and the change in neural representations that occur with neural stimulation, is essential for improving the performance of BIM systems as clinical tools. In this chapter, I will first summarize the contributions made by the work presented in this thesis. Then, I will propose future directions for work in neuroprosthetic skill learning and stimulation-based neuromodulation.

### 5.1 Summary of Contributions

This work in this thesis sheds light on how the brain adapts at multiple scales to control a motor cortical BIM and how microstimulation modulates neural circuits involved in flexible decision-making. The findings presented in this work will be important for advancing the development of neurobiologically informed neuroprosthetic devices.

In Chapter 2, we used factor analysis to compare coordinated spatiotemporal dynamics of subpopulations of neurons with direct input to a BIM decoder and the remaining recorded population. We partitioned the neural variance within the entire recorded population into variance arising from common, shared signals and variance arising from independent, private signals. We found that while the total variance increased in all groups of neurons and there was an overall increase in the proportion of shared variance, this proportional increase in shared variance was driven primarily by neurons with direct input to the BIM. This increase in the proportion of shared variance indicates an increase in coordination amongst the output-relevant neurons. We also observed changes in the covariance structure of both the direct and indirect subpopulations, with larger changes occurring in the direct subpopulation. These changes in the low-dimensional population structure of neural firing activity were correlated with changes in behavioral performance. Overall, these findings shed light onto how the motor cortex refines cortical population dynamics for control of a novel actuator.

In Chapter 3, we demonstrated simultaneous recordings from multiple cortical and subcortical brain regions as nonhuman primates (NHPs) learned to control a motor cortical BIM. To our knowledge, this is the first time a large-scale semi-chronic microdrive has been used to record neural activity for real-time BIM control. We used this method to investigate task-relevant activity within and outside of motor cortex.

We identified task-relevant changes in spectral power in both caudate and prefrontal cortex, as well as in motor cortex. Spectral power from all three regions can be used to classify BMI control, manual control, and baseline. The majority of misclassifications occurred between BMI control and manual control, suggesting that BMI control is more similar to manual control than to baseline. Furthermore, these misclassifications were less common at the go cue than at target acquisition. This may suggest larger differences in top-down strategies between BMI and manual control at the beginning of successful trials. Additionally, we identified task-relevant changes in directed functional connectivity between motor cortex, prefrontal cortex, and caudate. While there is a significant amount of information flow between all three regions during baseline and BMI control, BMI control was associated with an increase in information flow from caudate to motor cortex and prefrontal cortex to motor cortex above baseline. This work provides evidence that prefrontal cortex and caudate may play an important role in BMI control, especially at the go cue when top-down cognitive processes may guide action selection and initiation. These findings further demonstrate that neuroprosthetic control engages multiple cortical and subcortical brain regions.

In Chapter 4, we investigated how behaviorally biasing microstimulation administered in the caudate modulates underlying value representations in the caudate and anterior cingulate cortex. We demonstrated that changes in choice behavior resulting from high frequency microstimulation paired with a particular stimulus administered in the caudate was accompanied by an increase in the fraction of units encoding value in both caudate and anterior cingulate cortex. However, the increase in the fraction of units encoding value was only significant for units in the caudate encoding the stimulus paired with electrical stimulation for one animal. In addition to increasing our understanding of neural value computations underlying flexible decision-making, this work provides further evidence that stimulation-based therapies may be used to regulate valuation of choices in neuropsychiatric patients.

## 5.2. Open Questions and Future Directions

The work in this thesis contributes to our understanding of how the brain adapts to and interacts with a BMI system; however, there remain many unanswered questions. In this section, I present a number of open questions and future directions that may help to elucidate mechanisms of neuroprosthetic skill learning and aid in the development of the next generation of neurobiologically informed neuroprosthetics.

### *5.2.1 Neuroprosthetic Skill Learning*

Closed-loop control of BMI systems has been shown to engage neural adaptation, a process referred to as neuroprosthetic learning. While many studies have investigated this learning process and how it relates to natural motor learning, there remain many unanswered questions surrounding the neural adaptations that accompany the acquisition of proficient BMI control. Understanding this neuroplasticity is important for facilitating the development of robust BMIs that integrate with the brain and are intuitively controlled.

#### *Network Credit Assignment*

Previous literature, as well as Chapter 2 of this thesis, has highlighted differences in neural activity associated with neuroprosthetic learning between groups of neurons whose activity is directly input to a BMI decoder and the remaining motor cortical population (Ganguly et al., 2011; Gulati et al., 2014; Koralek et al., 2013; Zippi, You et al., 2021). Further work is needed to understand how the brain assigns credit to neurons whose activity causally drives output or reinforces specific patterns of neural activity that produce desirable outcomes. Identifying upstream regions within the cortico-striatal circuits known to be involved in BMI learning and control may be an important first step.

#### *Distributed Neural Circuits Underlying Neuroprosthetic Skill Learning*

In Chapter 3 of this thesis, we presented a new method for simultaneously recording from multiple cortical and subcortical structures while NHP learns to control a real-time BMI. Our analyses focused on changes in the spectral power within prefrontal cortex, caudate, and motor cortex and interactions between the LFP in these regions. Further characterization of these interactions with single- and multi-unit recordings in all these regions of interest will enable better spatial and temporal resolution. Additionally, the cortico-cortical and cortico-striatal interactions presented in Chapter 3, as well as the cortico-striatal interactions presented in previous work (Koralek et al., 2012, 2013; Neely et al., 2018) are all part of a much larger cortico-basal ganglia-thalamic loop, with many nodes participating in multiple large-scale interactions. Examining the role of other potentially task-relevant structures, such as the thalamus, will be important for characterizing the networks involved in neuroprosthetic learning and control. Furthermore, understanding which regions of the brain are involved in each aspect of neuroprosthetic learning and how they interact with one another to facilitate neuroprosthetic learning and control may enable us to develop BMIs that are more intuitive and easier to learn. Studies comparing these interactions in early and late learning, with new and learned decoders, and with intuitive and non-intuitive

mappings between neural activity and control will be useful in improving our understanding of their role in neuroprosthetic learning and control.

### *5.2.2 Stimulation-Based Neuromodulation*

BMI systems have the ability to restore lost function to patients with neurological and neuropsychiatric disorders by creating novel control pathways from the brain to external or internal devices. Internal devices that can be controlled by a BMI include stimulators that modulate neural activity through the injection of electrical current. While the majority of work in the field of BMIs to date has focused on designing systems with the potential to restore lost motor function in paralyzed patients, recent work has demonstrated that closed-loop BMIs for mood disorders may be a promising clinical tool for restoring dysfunctional mood-regulation or decision-making (Sani et al., 2018; Shanechi, 2019). Additionally, these systems will allow for researchers to probe the neural mechanisms involved in mood regulation and decision-making, increasing our understanding of the etiology of neuropsychiatric disorders. There remains much work to be done to improve our ability to control or regulate internal brain states via BMI systems.

#### *Characterization of Stimulation-Based Neuromodulation*

Deep brain stimulation (DBS) has set the stage for real-time causal intervention therapies for a number of neurological disorders. Electrical stimulation delivered via microelectrodes placed in specific regions has been shown to reduce subsets of symptoms in diseases ranging from Parkinson's Disease to major depression (Goodman & Alterman, 2012; Holtzheimer & Mayberg, 2011). However, the mechanisms by which DBS elicits its effects remain unclear. For example, electrical stimulation of certain regions can elicit mixed patterns of excitation and inhibition in diverse cell types (Maks et al., 2009). Furthermore, several studies have demonstrated that the stimulation of white matter tracts can also elicit beneficial effects (Holtzheimer & Mayberg, 2011). Understanding how electrical stimulation of specific regions modulates the circuitry involved in neurological disorders will be important for our understanding both the etiology of the disease and the mechanism of the therapy. An understanding of how stimulation-based neuromodulation influences neural circuitry, including neural representations of emotion and mood, will allow for the development of better neurotherapies for neuropsychiatric disorders. In chapter 4 of this thesis, we investigated how caudate microstimulation modulated behavior and value encoding. Using a similar approach to understand how stimulation of other regions within the cortico-striatal and cortico-limbic networks involved in decision-making and mood-

regulation will be important for advancing our ability to flexibly modulate internal brain states related to mood disorders and dysfunctional decision-making.

### *Reliable Biomarkers for Closed-Loop Stimulation*

Closed-loop DBS, or adaptive DBS, depends on the identification of reliable biomarkers of disease states. Being able to reliably detect abnormal patterns of neural activity that are related to specific disease states is a crucial step in developing closed-loop stimulation-based neurotherapies for neuropsychiatric disorders, in which stimulation parameters are dynamically adjusted according to a decoded signal (Provenza et al., 2019). Thus, developing closed-loop stimulation therapies for neuropsychiatric disorders, such as anxiety and depression, will rely on the identification of reliable biomarkers for these abnormal brain states. While some studies have shown that decoding mental states such as mood is possible (Sani et al., 2018), it is challenging because mood representations involve multiple distributed brain regions, whose functional organization is not well understood. Furthermore, mood disorders may manifest dissimilarly in different people. Studies identifying commonalities in affective representations across individuals or methods for developing personalized mood decoding will be necessary for the clinical use of these systems.

### *Neural Adaptations to Closed-Loop Stimulation*

Once reliable biomarkers have been identified and there is an understanding of how stimulation influences neural circuitry associated with various adaptive and maladaptive behaviors, closed-loop stimulation via BMI systems will allow for the observation and modulation of neural pathways involved in these behaviors with high spatio-temporal resolution. Neuropsychiatric diseases often involve specific maladaptive associations, such as ‘trigger’ stimuli, that elicit the symptoms (Belin et al., 2013). Traditional psychological treatments, such as cognitive-behavioral therapy, often focus on these ‘triggers’ with the goal of unlearning maladaptive associations. BMI systems may be well suited to aid in this process by quickly disrupting the neural circuitry underlying the maladaptive association whenever it is triggered. For example, previous work has demonstrated that electrical microstimulation can reverse addiction in rodent models (Creed et al., 2015). Thus, disruptive electrical stimulation may aid in the weakening of these associations over time until they ultimately no longer exist. Studies investigating the neural adaptations that accompany closed-loop electrical stimulation will be important for characterizing the effectiveness of BMIs as a clinical tool for neuropsychiatric disorders.

### 5.3 Conclusions

Directly interfacing brains with machines has opened a myriad of promising clinical and scientific applications. BMIs have tremendous potential to improve the quality of life for those with a number of cognitive and motor disabilities but have also revealed many challenges with understanding and interfacing with the brain. Further work investigating the neuroplasticity that occurs with introducing these systems to the brain and in neuroengineering will be critical for unlocking the full potential of BMI therapies. Large-scale simultaneous recording strategies provide promising new paths for studying the neurobiology of BMIs and may further enable us to develop neurobiologically informed neuroprosthetic systems.

## Chapter 6: Bibliography

- Ajiboye, A. B., Willett, F. R., Young, D. R., Memberg, W. D., Murphy, B. A., Miller, J. P., Walter, B. L., Sweet, J. A., Hoyen, H. A., Keith, M. W., Peckham, P. H., Simeral, J. D., Donoghue, J. P., Hochberg, L. R., & Kirsch, R. F. (2017). Restoration of reaching and grasping movements through brain-controlled muscle stimulation in a person with tetraplegia: A proof-of-concept demonstration. *The Lancet*, *389*(10081), 1821–1830. [https://doi.org/10.1016/S0140-6736\(17\)30601-3](https://doi.org/10.1016/S0140-6736(17)30601-3)
- Ames, K. C., & Churchland, M. M. (2019). Motor cortex signals for each arm are mixed across hemispheres and neurons yet partitioned within the population response. *ELife*, *8*, e46159. <https://doi.org/10.7554/eLife.46159>
- Anderson, B. A. (2016). The attention habit: How reward learning shapes attentional selection: The attention habit. *Annals of the New York Academy of Sciences*, *1369*(1), 24–39. <https://doi.org/10.1111/nyas.12957>
- Anderson, B. A., Laurent, P. A., & Yantis, S. (2011). Value-driven attentional capture. *Proceedings of the National Academy of Sciences*, *108*(25), 10367–10371. <https://doi.org/10.1073/pnas.1104047108>
- Anderson, B. A., & Yantis, S. (2012). Value-driven attentional and oculomotor capture during goal-directed, unconstrained viewing. *Attention, Perception, & Psychophysics*, *74*(8), 1644–1653. <https://doi.org/10.3758/s13414-012-0348-2>
- Antzoulatos, E. G., & Miller, E. K. (2011). Differences between Neural Activity in Prefrontal Cortex and Striatum during Learning of Novel Abstract Categories. *Neuron*, *71*(2), 243–249. <https://doi.org/10.1016/j.neuron.2011.05.040>
- Antzoulatos, E. G., & Miller, E. K. (2014). Increases in Functional Connectivity between Prefrontal Cortex and Striatum during Category Learning. *Neuron*, *83*(1), 216–225. <https://doi.org/10.1016/j.neuron.2014.05.005>
- Anumanchipalli, G. K., Chartier, J., & Chang, E. F. (2019). Speech synthesis from neural decoding of spoken sentences. *Nature*, *568*(7753), 493–498. <https://doi.org/10.1038/s41586-019-1119-1>
- Arduin, P.-J., Fregnac, Y., Shulz, D. E., & Ego-Stengel, V. (2013). “Master” Neurons Induced by Operant Conditioning in Rat Motor Cortex during a Brain-Machine Interface Task. *Journal of Neuroscience*, *33*(19), 8308–8320. <https://doi.org/10.1523/JNEUROSCI.2744-12.2013>
- Athalye, V. R., Carmena, J. M., & Costa, R. M. (2020). Neural reinforcement: Re-entering and refining neural dynamics leading to desirable outcomes. *Current Opinion in Neurobiology*, *60*, 145–154. <https://doi.org/10.1016/j.conb.2019.11.023>

- Athalye, V. R., Ganguly, K., Costa, R. M., & Carmena, J. M. (2017). Emergence of Coordinated Neural Dynamics Underlies Neuroprosthetic Learning and Skillful Control. *Neuron*, 93(4), 955-970.e5. <https://doi.org/10.1016/j.neuron.2017.01.016>
- Athalye, V. R., Santos, F. J., Carmena, J. M., & Costa, R. M. (2018). Evidence for a neural law of effect. *Science*, 359(6379), 1024–1029. <https://doi.org/10.1126/science.aao6058>
- Badre, D., Kayser, A. S., & D'Esposito, M. (2010). Frontal Cortex and the Discovery of Abstract Action Rules. *Neuron*, 66(2), 315–326. <https://doi.org/10.1016/j.neuron.2010.03.025>
- Balewski, Z. Z., Knudsen, E. B., & Wallis, J. D. (2022). Fast and slow contributions to decision-making in corticostriatal circuits. *Neuron*, 110(13), 2170-2182.e4. <https://doi.org/10.1016/j.neuron.2022.04.005>
- Banaie Boroujeni, K., Oemisch, M., Hassani, S. A., & Womelsdorf, T. (2020). Fast spiking interneuron activity in primate striatum tracks learning of attention cues. *Proceedings of the National Academy of Sciences*, 117(30), 18049–18058. <https://doi.org/10.1073/pnas.2001348117>
- Behrens, T. E. J., Woolrich, M. W., Walton, M. E., & Rushworth, M. F. S. (2007). Learning the value of information in an uncertain world. *Nature Neuroscience*, 10(9), 1214–1221. <https://doi.org/10.1038/nn1954>
- Belin, D., Belin-Rauscent, A., Murray, J. E., & Everitt, B. J. (2013). Addiction: Failure of control over maladaptive incentive habits. *Current Opinion in Neurobiology*, 23(4), 564–572. <https://doi.org/10.1016/j.conb.2013.01.025>
- Benabid, A. L. (2003). Deep brain stimulation for Parkinson's disease. *Current Opinion in Neurobiology*, 13(6), 696–706. <https://doi.org/10.1016/j.conb.2003.11.001>
- Benabid, A. L., Chabardes, S., Mitrofanis, J., & Pollak, P. (2009). *Deep brain stimulation of the subthalamic nucleus for the treatment of Parkinson's disease*. 8, 15.
- Biran, R., Martin, D. C., & Tresco, P. A. (2005). Neuronal cell loss accompanies the brain tissue response to chronically implanted silicon microelectrode arrays. *Experimental Neurology*, 195(1), 115–126. <https://doi.org/10.1016/j.expneurol.2005.04.020>
- Boettiger, C. A. (2005). Frontal Networks for Learning and Executing Arbitrary Stimulus-Response Associations. *Journal of Neuroscience*, 25(10), 2723–2732. <https://doi.org/10.1523/JNEUROSCI.3697-04.2005>
- Boussaoud, D., & Wise, S. P. (1993). Primate frontal cortex: Neuronal activity following attentional versus intentional cues. *Experimental Brain Research*, 95(1). <https://doi.org/10.1007/BF00229650>

- Bressler, S. L., & Seth, A. K. (2011). Wiener–Granger Causality: A well established methodology. *NeuroImage*, 58(2), 323–329.  
<https://doi.org/10.1016/j.neuroimage.2010.02.059>
- Buschman, T. J., & Miller, E. K. (2007). Top-Down Versus Bottom-Up Control of Attention in the Prefrontal and Posterior Parietal Cortices. *Science*, 315(5820), 1860–1862.  
<https://doi.org/10.1126/science.1138071>
- Capogrosso, M., Milekovic, T., Borton, D., Wagner, F., Moraud, E. M., Mignardot, J.-B., Buse, N., Gandar, J., Barraud, Q., Xing, D., Rey, E., Duis, S., Jianzhong, Y., Ko, W. K. D., Li, Q., Detemple, P., Denison, T., Micera, S., Bezdard, E., ... Courtine, G. (2016). A brain–spine interface alleviating gait deficits after spinal cord injury in primates. *Nature*, 539(7628), 284–288. <https://doi.org/10.1038/nature20118>
- Carlson, T., & del R. Millan, J. (2013). Brain-Controlled Wheelchairs: A Robotic Architecture. *IEEE Robotics & Automation Magazine*, 20(1), 65–73.  
<https://doi.org/10.1109/MRA.2012.2229936>
- Carmena, J. M., Lebedev, M. A., Crist, R. E., O’Doherty, J. E., Santucci, D. M., Dimitrov, D. F., Patil, P. G., Henriquez, C. S., & Nicolelis, M. A. L. (2003). Learning to Control a Brain–Machine Interface for Reaching and Grasping by Primates. *PLoS Biology*, 1(2), e42.  
<https://doi.org/10.1371/journal.pbio.0000042>
- Carroll, W. M. (2019). The global burden of neurological disorders. *The Lancet Neurology*, 18(5), 418–419. [https://doi.org/10.1016/S1474-4422\(19\)30029-8](https://doi.org/10.1016/S1474-4422(19)30029-8)
- Chapin, J. K., Moxon, K. A., Markowitz, R. S., & Nicolelis, M. A. L. (1999). Real-time control of a robot arm using simultaneously recorded neurons in the motor cortex. *Nature Neuroscience*, 2(7), 664–670. <https://doi.org/10.1038/10223>
- Chase, S. M., Kass, R. E., & Schwartz, A. B. (2012). Behavioral and neural correlates of visuomotor adaptation observed through a brain-computer interface in primary motor cortex. *Journal of Neurophysiology*, 108(2), 624–644.  
<https://doi.org/10.1152/jn.00371.2011>
- Chiang, F.-K., Wallis, J. D., & Rich, E. L. (2022). Cognitive strategies shift information from single neurons to populations in prefrontal cortex. *Neuron*, 110(4), 709–721.e4.  
<https://doi.org/10.1016/j.neuron.2021.11.021>
- Churchland, M. M., Cunningham, J. P., Kaufman, M. T., Foster, J. D., Nuyujukian, P., Ryu, S. I., & Shenoy, K. V. (2012). Neural population dynamics during reaching. *Nature*, 487(7405), 51–56. <https://doi.org/10.1038/nature11129>
- Churchland, M. M., Cunningham, J. P., Kaufman, M. T., Ryu, S. I., & Shenoy, K. V. (2010). Cortical Preparatory Activity: Representation of Movement or First Cog in a Dynamical Machine? *Neuron*, 68(3), 387–400. <https://doi.org/10.1016/j.neuron.2010.09.015>

- Clancy, K. B., Koralek, A. C., Costa, R. M., Feldman, D. E., & Carmena, J. M. (2014). Volitional modulation of optically recorded calcium signals during neuroprosthetic learning. *Nature Neuroscience*, *17*(6), 807–809. <https://doi.org/10.1038/nn.3712>
- Collinger, J. L., Gaunt, R. A., & Schwartz, A. B. (2018). Progress towards restoring upper limb movement and sensation through intracortical brain-computer interfaces. *Current Opinion in Biomedical Engineering*, *8*, 84–92. <https://doi.org/10.1016/j.cobme.2018.11.005>
- Collinger, J. L., Wodlinger, B., Downey, J. E., Wang, W., Tyler-Kabara, E. C., Weber, D. J., McMorland, A. J., Velliste, M., Boninger, M. L., & Schwartz, A. B. (2013). High-performance neuroprosthetic control by an individual with tetraplegia. *The Lancet*, *381*(9866), 557–564. [https://doi.org/10.1016/S0140-6736\(12\)61816-9](https://doi.org/10.1016/S0140-6736(12)61816-9)
- Costa, R. M. (2011). A selectionist account of de novo action learning. *Current Opinion in Neurobiology*, *21*(4), 579–586. <https://doi.org/10.1016/j.conb.2011.05.004>
- Costa, R. M., Cohen, D., & Nicolelis, M. A. L. (2004). Differential Corticostriatal Plasticity during Fast and Slow Motor Skill Learning in Mice. *Current Biology*, *14*(13), 1124–1134. <https://doi.org/10.1016/j.cub.2004.06.053>
- Creed, M., Pascoli, V. J., & Lüscher, C. (2015). Refining deep brain stimulation to emulate optogenetic treatment of synaptic pathology. *Science*, *347*(6222), 659–664. <https://doi.org/10.1126/science.1260776>
- Dangi, S., Orsborn, A. L., Moorman, H. G., & Carmena, J. M. (2013). Design and Analysis of Closed-Loop Decoder Adaptation Algorithms for Brain-Machine Interfaces. *Neural Computation*, *25*(7), 1693–1731. [https://doi.org/10.1162/NECO\\_a\\_00460](https://doi.org/10.1162/NECO_a_00460)
- Dayan, E., & Cohen, L. G. (2011). Neuroplasticity Subservient Motor Skill Learning. *Neuron*, *72*(3), 443–454. <https://doi.org/10.1016/j.neuron.2011.10.008>
- Degenhart, A. D., Bishop, W. E., Oby, E. R., Tyler-Kabara, E. C., Chase, S. M., Batista, A. P., & Yu, B. M. (2020). Stabilization of a brain-computer interface via the alignment of low-dimensional spaces of neural activity. *Nature Biomedical Engineering*, *4*(7), 672–685. <https://doi.org/10.1038/s41551-020-0542-9>
- Delong, M. R., Georgopoulos, A. P., Crutcher, M. D., Mitchell, S. J., Richardson, R. T., & Alexander, G. E. (1984). Functional Organization of the Basal Ganglia: Contributions of Single-Cell Recording Studies. In D. Evered & M. O'Connor (Eds.), *Novartis Foundation Symposia* (pp. 64–82). John Wiley & Sons, Ltd. <https://doi.org/10.1002/9780470720882.ch5>
- Deuschl, G., Krack, P., Bötzel, K., Dillmann, U., Gruber, D., Hilker, R., Koy, J., Lorenz, D., Mehdorn, H. M., Oertel, W., Reichmann, H., Schneider, G.-H., Steude, U., Tronnier, V., Wolf, E., & Voges, J. (2006). A Randomized Trial of Deep-Brain Stimulation for Parkinson's Disease. *N Engl J Med*, *13*.

- Dhawale, A. K., Smith, M. A., & Ölveczky, B. P. (2017). The Role of Variability in Motor Learning. *Annual Review of Neuroscience*, 40(1), 479–498.  
<https://doi.org/10.1146/annurev-neuro-072116-031548>
- di Pellegrino, G., & Wise, S. (1993). Visuospatial versus visuomotor activity in the premotor and prefrontal cortex of a primate. *The Journal of Neuroscience*, 13(3), 1227–1243.  
<https://doi.org/10.1523/JNEUROSCI.13-03-01227.1993>
- Ding, M., Chen, Y., & Bressler, S. L. (2006). *Granger Causality: Basic Theory and Application to Neuroscience* (arXiv:q-bio/0608035). arXiv. <http://arxiv.org/abs/q-bio/0608035>
- Ditterich, J., Mazurek, M. E., & Shadlen, M. N. (2003). Microstimulation of visual cortex affects the speed of perceptual decisions. *Nature Neuroscience*, 6(8), 891–898.  
<https://doi.org/10.1038/nn1094>
- Donchin, O., Rabe, K., Diedrichsen, J., Lally, N., Schoch, B., Gizewski, E. R., & Timmann, D. (2012). Cerebellar regions involved in adaptation to force field and visuomotor perturbation. *J Neurophysiol*, 14.
- Dotson, N. M., Hoffman, S. J., Goodell, B., & Gray, C. M. (2017). A Large-Scale Semi-Chronic Microdrive Recording System for Non-Human Primates. *Neuron*, 96(4), 769–782.e2.  
<https://doi.org/10.1016/j.neuron.2017.09.050>
- Downey, J. E., Schwed, N., Chase, S. M., Schwartz, A. B., & Collinger, J. L. (2018). Intracortical recording stability in human brain–computer interface users. *Journal of Neural Engineering*, 15(4), 046016. <https://doi.org/10.1088/1741-2552/aab7a0>
- Elsayed, G. F., Lara, A. H., Kaufman, M. T., Churchland, M. M., & Cunningham, J. P. (2016). Reorganization between preparatory and movement population responses in motor cortex. *Nature Communications*, 7(1), 13239. <https://doi.org/10.1038/ncomms13239>
- Ethier, C., Oby, E. R., Bauman, M. J., & Miller, L. E. (2012). Restoration of grasp following paralysis through brain-controlled stimulation of muscles. *Nature*, 485(7398), 368–371.  
<https://doi.org/10.1038/nature10987>
- Everitt, B. S. (1984). *An Introduction to Latent Variable Models*. Springer Netherlands.  
<https://doi.org/10.1007/978-94-009-5564-6>
- Fahimi Hnazaee, M., Wittevrongel, B., Khachatryan, E., Libert, A., Carrette, E., Dauwe, I., Meurs, A., Boon, P., Van Roost, D., & Van Hulle, M. M. (2020). Localization of deep brain activity with scalp and subdural EEG. *NeuroImage*, 223, 117344.  
<https://doi.org/10.1016/j.neuroimage.2020.117344>
- Fedorov, A., Beichel, R., Kalpathy-Cramer, J., Finet, J., Fillion-Robin, J.-C., Pujol, S., Bauer, C., Jennings, D., Fennessy, F., Sonka, M., Buatti, J., Aylward, S., Miller, J. V., Pieper, S., & Kikinis, R. (2012). 3D Slicer as an image computing platform for the Quantitative Imaging

- Network. *Magnetic Resonance Imaging*, 30(9), 1323–1341.  
<https://doi.org/10.1016/j.mri.2012.05.001>
- Feigin, V. L., Nichols, E., Alam, T., Bannick, M. S., Beghi, E., Blake, N., Culpepper, W. J., Dorsey, E. R., Elbaz, A., Ellenbogen, R. G., Fisher, J. L., Fitzmaurice, C., Giussani, G., Glennie, L., James, S. L., Johnson, C. O., Kassebaum, N. J., Logroscino, G., Marin, B., ... Vos, T. (2019). Global, regional, and national burden of neurological disorders, 1990–2016: A systematic analysis for the Global Burden of Disease Study 2016. *The Lancet Neurology*, 18(5), 459–480. [https://doi.org/10.1016/S1474-4422\(18\)30499-X](https://doi.org/10.1016/S1474-4422(18)30499-X)
- Feigin, V. L., Vos, T., Nichols, E., Owolabi, M. O., Carroll, W. M., Dichgans, M., Deuschl, G., Parmar, P., Brainin, M., & Murray, C. (2020). The global burden of neurological disorders: Translating evidence into policy. *The Lancet Neurology*, 19(3), 255–265.  
[https://doi.org/10.1016/S1474-4422\(19\)30411-9](https://doi.org/10.1016/S1474-4422(19)30411-9)
- Fernando, A. B. P., & Robbins, T. W. (2011). Animal Models of Neuropsychiatric Disorders. *Annual Review of Clinical Psychology*, 7(1), 39–61. <https://doi.org/10.1146/annurev-clinpsy-032210-104454>
- Fetz, E. E. (1969). *Operant Conditioning of Cortical Unit Activity* (No. 3870). 163(3870), 955–958. <https://doi.org/doi:10.1126/science.163.3870.955>
- Fetz, E. E. (2007). Volitional control of neural activity: Implications for brain-computer interfaces: Volitional control of neural activity. *The Journal of Physiology*, 579(3), 571–579. <https://doi.org/10.1113/jphysiol.2006.127142>
- Fetz, E. E., & Baker, M. A. (1973). Operantly conditioned patterns on precentral unit activity and correlated responses in adjacent cells and contralateral muscles. *Journal of Neurophysiology*, 36(2), 179–204. <https://doi.org/10.1152/jn.1973.36.2.179>
- Flesher, S. N., Downey, J. E., Weiss, J. M., Hughes, C. L., Herrera, A. J., Tyler-Kabara, E. C., Boninger, M. L., Collinger, J. L., & Gaunt, R. A. (2021). A brain-computer interface that evokes tactile sensations improves robotic arm control. *Science*, 372(6544), 831–836. <https://doi.org/10.1126/science.abd0380>
- Fraser, G. W., & Schwartz, A. B. (2012). Recording from the same neurons chronically in motor cortex. *Journal of Neurophysiology*, 107(7), 1970–1978.  
<https://doi.org/10.1152/jn.01012.2010>
- Fuster, J. M. (2000). Prefrontal neurons in networks of executive memory. *Brain Research Bulletin*, 52(5), 331–336. [https://doi.org/10.1016/S0361-9230\(99\)00258-0](https://doi.org/10.1016/S0361-9230(99)00258-0)
- Gallego, J. A., Perich, M. G., Miller, L. E., & Solla, S. A. (2017). Neural Manifolds for the Control of Movement. *Neuron*, 94(5), 978–984. <https://doi.org/10.1016/j.neuron.2017.05.025>
- Ganguly, K., & Carmena, J. M. (2009). Emergence of a Stable Cortical Map for Neuroprosthetic Control. *PLoS Biology*, 7(7), e1000153. <https://doi.org/10.1371/journal.pbio.1000153>

- Ganguly, K., & Carmena, J. M. (2010). Neural Correlates of Skill Acquisition with a Cortical Brain–Machine Interface. *Journal of Motor Behavior*, 42(6), 355–360. <https://doi.org/10.1080/00222895.2010.526457>
- Ganguly, K., Dimitrov, D. F., Wallis, J. D., & Carmena, J. M. (2011). Reversible large-scale modification of cortical networks during neuroprosthetic control. *Nature Neuroscience*, 14(5), 662–667. <https://doi.org/10.1038/nn.2797>
- Genovesio, A., Brasted, P. J., Mitz, A. R., & Wise, S. P. (2005). Prefrontal Cortex Activity Related to Abstract Response Strategies. *Neuron*, 47(2), 307–320. <https://doi.org/10.1016/j.neuron.2005.06.006>
- Gilja, V., Nuyujukian, P., Chestek, C. A., Cunningham, J. P., Yu, B. M., Fan, J. M., Churchland, M. M., Kaufman, M. T., Kao, J. C., Ryu, S. I., & Shenoy, K. V. (2012). A high-performance neural prosthesis enabled by control algorithm design. *Nature Neuroscience*, 15(12), 1752–1757. <https://doi.org/10.1038/nn.3265>
- Goldman-Rakic, P. S. (1996). *The prefrontal landscape: Implications of functional architecture for understanding human mentation and the central executive*. 9.
- Golub, M. D., Chase, S. M., Batista, A. P., & Yu, B. M. (2016). Brain–computer interfaces for dissecting cognitive processes underlying sensorimotor control. *Current Opinion in Neurobiology*, 37, 53–58. <https://doi.org/10.1016/j.conb.2015.12.005>
- Golub, M. D., Sadtler, P. T., Oby, E. R., Quick, K. M., Ryu, S. I., Tyler-Kabara, E. C., Batista, A. P., Chase, S. M., & Yu, B. M. (2018). Learning by neural reassociation. *Nature Neuroscience*, 21(4), 607–616. <https://doi.org/10.1038/s41593-018-0095-3>
- Golub, M. D., Yu, B. M., Schwartz, A. B., & Chase, S. M. (2014). Motor cortical control of movement speed with implications for brain-machine interface control. *Journal of Neurophysiology*, 112(2), 411–429. <https://doi.org/10.1152/jn.00391.2013>
- Goodman, W. K., & Alterman, R. L. (2012). Deep Brain Stimulation for Intractable Psychiatric Disorders. *Annual Review of Medicine*, 63(1), 511–524. <https://doi.org/10.1146/annurev-med-052209-100401>
- Gowda, S., Orsborn, A. L., Overduin, S. A., Moorman, H. G., & Carmena, J. M. (2014). Designing Dynamical Properties of Brain–Machine Interfaces to Optimize Task-Specific Performance. *IEEE Transactions on Neural Systems and Rehabilitation Engineering*, 22(5), 911–920. <https://doi.org/10.1109/TNSRE.2014.2309673>
- Granger, C. W. T. (1969). Investigating Causal Relations by Econometric Models and Cross-spectral Methods. *Econometrica*, 37(3), 424–438.
- Graybiel, A., Aosaki, T., Flaherty, A., & Kimura, M. (1994). The basal ganglia and adaptive motor control. *Science*, 265(5180), 1826–1831. <https://doi.org/10.1126/science.8091209>

- Graybiel, A. M. (1998). The Basal Ganglia and Chunking of Action Repertoires. *Neurobiology of Learning and Memory*, 70(1–2), 119–136. <https://doi.org/10.1006/nlme.1998.3843>
- Graybiel, A. M., & Grafton, S. T. (2015). The Striatum: Where Skills and Habits Meet. *Cold Spring Harbor Perspectives in Biology*, 7(8), a021691. <https://doi.org/10.1101/cshperspect.a021691>
- Green, A. M., & Kalaska, J. F. (2011). Learning to move machines with the mind. *Trends in Neurosciences*, 34(2), 61–75. <https://doi.org/10.1016/j.tins.2010.11.003>
- Gulati, T., Ramanathan, D. S., Wong, C. C., & Ganguly, K. (2014). Reactivation of emergent task-related ensembles during slow-wave sleep after neuroprosthetic learning. *Nature Neuroscience*, 17(8), 1107–1113. <https://doi.org/10.1038/nn.3759>
- Haber, S. N. (2003). The primate basal ganglia: Parallel and integrative networks. *Journal of Chemical Neuroanatomy*, 26(4), 317–330. <https://doi.org/10.1016/j.jchemneu.2003.10.003>
- Haber, S. N. (2016). Corticostriatal circuitry. *Dialogues in Clinical Neuroscience*, 18(1), 15.
- Hanks, T. D., Ditterich, J., & Shadlen, M. N. (2006). Microstimulation of macaque area LIP affects decision-making in a motion discrimination task. *Nature Neuroscience*, 9(5), 682–689. <https://doi.org/10.1038/nn1683>
- Hartley, C. A., & Phelps, E. A. (2012). Anxiety and Decision-Making. *Biological Psychiatry*, 72(2), 113–118. <https://doi.org/10.1016/j.biopsych.2011.12.027>
- Heming, E. A., Cross, K. P., Takei, T., Cook, D. J., & Scott, S. H. (2019). Independent representations of ipsilateral and contralateral limbs in primary motor cortex. *ELife*, 8, e48190. <https://doi.org/10.7554/eLife.48190>
- Hennig, J. A., Oby, E. R., Golub, M. D., Bahureksa, L. A., Sadtler, P. T., Quick, K. M., Ryu, S. I., Tyler-Kabara, E. C., Batista, A. P., Chase, S. M., & Yu, B. M. (2021). Learning is shaped by abrupt changes in neural engagement. *Nature Neuroscience*, 24(5), 727–736. <https://doi.org/10.1038/s41593-021-00822-8>
- Hikosaka, O., Nakahara, H., Rand, M. K., Sakai, K., Lu, X., Nakamura, K., Miyachi, S., & Doya, K. (1999). Parallel neural networks for learning sequential procedures. *Trends in Neurosciences*, 22(10), 464–471. [https://doi.org/10.1016/S0166-2236\(99\)01439-3](https://doi.org/10.1016/S0166-2236(99)01439-3)
- Hikosaka, O., Nakamura, K., & Nakahara, H. (2006). Basal Ganglia Orient Eyes to Reward. *Journal of Neurophysiology*, 95(2), 567–584. <https://doi.org/10.1152/jn.00458.2005>
- Hikosaka, O., Takikawa, Y., & Kawagoe, R. (2000). Role of the Basal Ganglia in the Control of Purposive Saccadic Eye Movements. *Physiological Reviews*, 80(3), 953–978. <https://doi.org/10.1152/physrev.2000.80.3.953>

- Hoang, K. B., Cassar, I. R., Grill, W. M., & Turner, D. A. (2017). Biomarkers and Stimulation Algorithms for Adaptive Brain Stimulation. *Frontiers in Neuroscience*, *11*, 564. <https://doi.org/10.3389/fnins.2017.00564>
- Hochberg, L. R., Bacher, D., Jarosiewicz, B., Masse, N. Y., Simeral, J. D., Vogel, J., Haddadin, S., Liu, J., Cash, S. S., van der Smagt, P., & Donoghue, J. P. (2012). Reach and grasp by people with tetraplegia using a neurally controlled robotic arm. *Nature*, *485*(7398), 372–375. <https://doi.org/10.1038/nature11076>
- Hochberg, L. R., Serruya, M. D., Friehs, G. M., Mukand, J. A., Saleh, M., Caplan, A. H., Branner, A., Chen, D., Penn, R. D., & Donoghue, J. P. (2006). Neuronal ensemble control of prosthetic devices by a human with tetraplegia. *Nature*, *442*(7099), 164–171. <https://doi.org/10.1038/nature04970>
- Holtzheimer, P. E., Husain, M. M., Lisanby, S. H., Taylor, S. F., Whitworth, L. A., McClintock, S., Slavin, K. V., Berman, J., McKhann, G. M., Patil, P. G., Rittberg, B. R., Abosch, A., Pandurangi, A. K., Holloway, K. L., Lam, R. W., Honey, C. R., Neimat, J. S., Henderson, J. M., DeBattista, C., ... Mayberg, H. S. (2017). Subcallosal cingulate deep brain stimulation for treatment-resistant depression: A multisite, randomised, sham-controlled trial. *The Lancet Psychiatry*, *4*(11), 839–849. [https://doi.org/10.1016/S2215-0366\(17\)30371-1](https://doi.org/10.1016/S2215-0366(17)30371-1)
- Holtzheimer, P. E., & Mayberg, H. S. (2011). Deep Brain Stimulation for Psychiatric Disorders. *Annual Review of Neuroscience*, *34*(1), 289–307. <https://doi.org/10.1146/annurev-neuro-061010-113638>
- Hoshi, E. (2006). Functional specialization within the dorsolateral prefrontal cortex: A review of anatomical and physiological studies of non-human primates. *Neuroscience Research*, *54*(2), 73–84. <https://doi.org/10.1016/j.neures.2005.10.013>
- Hwang, E. J., Bailey, P. M., & Andersen, R. A. (2013). Volitional Control of Neural Activity Relies on the Natural Motor Repertoire. *Current Biology*, *23*(5), 353–361. <https://doi.org/10.1016/j.cub.2013.01.027>
- Insel, T. R. (2012). Next-Generation Treatments for Mental Disorders. *Science Translational Medicine*, *4*(155). <https://doi.org/10.1126/scitranslmed.3004873>
- Iturrate, I., Antelis, J. M., Kubler, A., & Minguez, J. (2009). A Noninvasive Brain-Actuated Wheelchair Based on a P300 Neurophysiological Protocol and Automated Navigation. *IEEE Transactions on Robotics*, *25*(3), 614–627. <https://doi.org/10.1109/TRO.2009.2020347>
- Jackson, A., Mavoori, J., & Fetz, E. E. (2006). Long-term motor cortex plasticity induced by an electronic neural implant. *Nature*, *444*(7115), 56–60. <https://doi.org/10.1038/nature05226>
- Jarosiewicz, B., Chase, S. M., Fraser, G. W., Velliste, M., Kass, R. E., & Schwartz, A. B. (2008). Functional network reorganization during learning in a brain-computer interface paradigm.

- Proceedings of the National Academy of Sciences*, 105(49), 19486–19491.  
<https://doi.org/10.1073/pnas.0808113105>
- Jarosiewicz, B., Sarma, A. A., Bacher, D., Masse, N. Y., Simeral, J. D., Sorice, B., Oakley, E. M., Blabe, C., Pandarinath, C., Gilja, V., Cash, S. S., Eskandar, E. N., Friehs, G., Henderson, J. M., Shenoy, K. V., Donoghue, J. P., & Hochberg, L. R. (2015). Virtual typing by people with tetraplegia using a self-calibrating intracortical brain-computer interface. *Science Translational Medicine*, 7(313). <https://doi.org/10.1126/scitranslmed.aac7328>
- Johnston, K., Levin, H. M., Koval, M. J., & Everling, S. (2007). Top-Down Control-Signal Dynamics in Anterior Cingulate and Prefrontal Cortex Neurons following Task Switching. *Neuron*, 53(3), 453–462. <https://doi.org/10.1016/j.neuron.2006.12.023>
- Kao, J. C., Nuyujukian, P., Ryu, S. I., Churchland, M. M., Cunningham, J. P., & Shenoy, K. V. (2015). Single-trial dynamics of motor cortex and their applications to brain-machine interfaces. *Nature Communications*, 6(1), 7759. <https://doi.org/10.1038/ncomms8759>
- Kaufman, M. T., Churchland, M. M., Ryu, S. I., & Shenoy, K. V. (2014). Cortical activity in the null space: Permitting preparation without movement. *Nature Neuroscience*, 17(3), 440–448. <https://doi.org/10.1038/nn.3643>
- Kawala-Sterniuk, A., Browarska, N., Al-Bakri, A., Pelc, M., Zygarlicki, J., Sidikova, M., Martinek, R., & Gorzelanczyk, E. J. (2021). Summary of over Fifty Years with Brain-Computer Interfaces—A Review. *Brain Sciences*, 11(1), 43. <https://doi.org/10.3390/brainsci11010043>
- Kennedy, P. R., & Bakay, R. A. E. (1998). Restoration of neural output from a paralyzed patient by a direct brain connection: *NeuroReport*, 9(8), 1707–1711. <https://doi.org/10.1097/00001756-199806010-00007>
- Kennerley, S. W., Dahmubed, A. F., Lara, A. H., & Wallis, J. D. (2009). Neurons in the Frontal Lobe Encode the Value of Multiple Decision Variables. *Journal of Cognitive Neuroscience*, 21(6), 1162–1178. <https://doi.org/10.1162/jocn.2009.21100>
- Kennerley, S. W., & Wallis, J. D. (2009). Reward-Dependent Modulation of Working Memory in Lateral Prefrontal Cortex. *Journal of Neuroscience*, 29(10), 3259–3270. <https://doi.org/10.1523/JNEUROSCI.5353-08.2009>
- Kennerley, S. W., & Walton, M. E. (2011). Decision making and reward in frontal cortex: Complementary evidence from neurophysiological and neuropsychological studies. *Behavioral Neuroscience*, 125(3), 297–317. <https://doi.org/10.1037/a0023575>
- Kennerley, S. W., Walton, M. E., Behrens, T. E. J., Buckley, M. J., & Rushworth, M. F. S. (2006). Optimal decision making and the anterior cingulate cortex. *Nature Neuroscience*, 9(7), 940–947. <https://doi.org/10.1038/nn1724>

- Kessler, R. C., Demler, O., Frank, R. G., Olfson, M., Pincus, H. A., Walters, E. E., Wang, P., Wells, K. B., & Zaslavsky, A. M. (2005). Prevalence and Treatment of Mental Disorders, 1990 to 2003. *New England Journal of Medicine*, 352(24), 2515–2523.  
<https://doi.org/10.1056/NEJMsa043266>
- Khanna, P., So, K., & Carmena, J. M. (2013). Volitional phase control of neural oscillations using a brain-machine interface. *2013 6th International IEEE/EMBS Conference on Neural Engineering (NER)*, 1076–1079. <https://doi.org/10.1109/NER.2013.6696123>
- Khanna, P., Swann, N. C., de Hemptinne, C., Miocinovic, S., Miller, A., Starr, P. A., & Carmena, J. M. (2017). Neurofeedback Control in Parkinsonian Patients Using Electrocorticography Signals Accessed Wirelessly With a Chronic, Fully Implanted Device. *IEEE Transactions on Neural Systems and Rehabilitation Engineering*, 25(10), 1715–1724.  
<https://doi.org/10.1109/TNSRE.2016.2597243>
- Kim, H. F., & Hikosaka, O. (2013). Distinct Basal Ganglia Circuits Controlling Behaviors Guided by Flexible and Stable Values. *Neuron*, 79(5), 1001–1010.  
<https://doi.org/10.1016/j.neuron.2013.06.044>
- Kim, K.-T., Suk, H.-I., & Lee, S.-W. (2018). Commanding a Brain-Controlled Wheelchair Using Steady-State Somatosensory Evoked Potentials. *IEEE Transactions on Neural Systems and Rehabilitation Engineering*, 26(3), 654–665.  
<https://doi.org/10.1109/TNSRE.2016.2597854>
- Kim, S.-P., Simeral, J. D., Hochberg, L. R., Donoghue, J. P., & Black, M. J. (2008). Neural control of computer cursor velocity by decoding motor cortical spiking activity in humans with tetraplegia. *Journal of Neural Engineering*, 5(4), 455–476. <https://doi.org/10.1088/1741-2560/5/4/010>
- Kim, S.-P., Simeral, J. D., Hochberg, L. R., Donoghue, J. P., Friehs, G. M., & Black, M. J. (2011). Point-and-Click Cursor Control With an Intracortical Neural Interface System by Humans With Tetraplegia. *IEEE Transactions on Neural Systems and Rehabilitation Engineering*, 19(2), 193–203. <https://doi.org/10.1109/TNSRE.2011.2107750>
- Kimura, M., & Graybiel, A. M. (Eds.). (1995). *Functions of the Cortico-Basal Ganglia Loop*. Springer Japan. <https://doi.org/10.1007/978-4-431-68547-0>
- Knudsen, E. B., & Wallis, J. D. (2021). Hippocampal neurons construct a map of an abstract value space. *Cell*, 184(18), 4640–4650.e10. <https://doi.org/10.1016/j.cell.2021.07.010>
- 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.  
<https://doi.org/10.1016/j.neuron.2013.06.047>

- Koralek, A. C., Jin, X., Long II, J. D., Costa, R. M., & Carmena, J. M. (2012). Corticostriatal plasticity is necessary for learning intentional neuroprosthetic skills. *Nature*, *483*(7389), 331–335. <https://doi.org/10.1038/nature10845>
- Krakauer, J. W., Ghilardi, M.-F., Mentis, M., Barnes, A., Veytsman, M., Eidelberg, D., & Ghez, C. (2004). Differential Cortical and Subcortical Activations in Learning Rotations and Gains for Reaching: A PET Study. *J Neurophysiol*, *91*, 10.
- Krämer, U. M., Solbakk, A.-K., Funderud, I., Løvstad, M., Endestad, T., & Knight, R. T. (2013). The role of the lateral prefrontal cortex in inhibitory motor control. *Cortex*, *49*(3), 837–849. <https://doi.org/10.1016/j.cortex.2012.05.003>
- Lara, A. H., Kennerley, S. W., & Wallis, J. D. (2009). Encoding of Gustatory Working Memory by Orbitofrontal Neurons. *Journal of Neuroscience*, *29*(3), 765–774. <https://doi.org/10.1523/JNEUROSCI.4637-08.2009>
- Lau, B., & Glimcher, P. W. (2008). Value Representations in the Primate Striatum during Matching Behavior. *Neuron*, *58*(3), 451–463. <https://doi.org/10.1016/j.neuron.2008.02.021>
- Laurent, P. A., Hall, M. G., Anderson, B. A., & Yantis, S. (2015). Valuable orientations capture attention. *Visual Cognition*, *23*(1–2), 133–146. <https://doi.org/10.1080/13506285.2014.965242>
- Leuthardt, E. C., Schalk, G., Wolpaw, J. R., Ojemann, J. G., & Moran, D. W. (2004). A brain–computer interface using electrocorticographic signals in humans. *Journal of Neural Engineering*, *1*(2), 63–71. <https://doi.org/10.1088/1741-2560/1/2/001>
- Leventhal, D. K., Gage, G. J., Schmidt, R., Pettibone, J. R., Case, A. C., & Berke, J. D. (2012). Basal Ganglia Beta Oscillations Accompany Cue Utilization. *Neuron*, *73*(3), 523–536. <https://doi.org/10.1016/j.neuron.2011.11.032>
- Little, S., Beudel, M., Zrinzo, L., Foltynie, T., Limousin, P., Hariz, M., Neal, S., Cheeran, B., Cagnan, H., Gratwicke, J., Aziz, T. Z., Pogosyan, A., & Brown, P. (2016). Bilateral adaptive deep brain stimulation is effective in Parkinson’s disease. *Journal of Neurology, Neurosurgery & Psychiatry*, *87*(7), 717–721. <https://doi.org/10.1136/jnnp-2015-310972>
- Little, S., Pogosyan, A., Neal, S., Zavala, B., Zrinzo, L., Hariz, M., Foltynie, T., Limousin, P., Ashkan, K., FitzGerald, J., Green, A. L., Aziz, T. Z., & Brown, P. (2013). Adaptive deep brain stimulation in advanced Parkinson disease. *Annals of Neurology*, *74*(3), 449–457. <https://doi.org/10.1002/ana.23951>
- Madhavan, R., Chao, Z. C., & Potter, S. M. (2007). Plasticity of recurring spatiotemporal activity patterns in cortical networks. *Physical Biology*, *4*(3), 181–193. <https://doi.org/10.1088/1478-3975/4/3/005>
- Maks, C. B., Butson, C. R., Walter, B. L., Vitek, J. L., & McIntyre, C. C. (2009). Deep brain stimulation activation volumes and their association with neurophysiological mapping and

- therapeutic outcomes. *Journal of Neurology, Neurosurgery & Psychiatry*, 80(6), 659–666. <https://doi.org/10.1136/jnnp.2007.126219>
- Mandelblat-Cerf, Y., Paz, R., & Vaadia, E. (2009). Trial-to-Trial Variability of Single Cells in Motor Cortices Is Dynamically Modified during Visuomotor Adaptation. *Journal of Neuroscience*, 29(48), 15053–15062. <https://doi.org/10.1523/JNEUROSCI.3011-09.2009>
- Mayberg, H. S., Lozano, A. M., Voon, V., McNeely, H. E., Seminowicz, D., Hamani, C., Schwalb, J. M., & Kennedy, S. H. (2005). Deep Brain Stimulation for Treatment-Resistant Depression. *Neuron*, 45(5), 651–660. <https://doi.org/10.1016/j.neuron.2005.02.014>
- Middleton, F. A., & Strick, P. L. (2002). Basal-ganglia “Projections” to the Prefrontal Cortex of the Primate. *Cerebral Cortex*, 12(9), 926–935. <https://doi.org/10.1093/cercor/12.9.926>
- Millán, J. D. R. (2010). Combining brain-computer interfaces and assistive technologies: State-of-the-art and challenges. *Frontiers in Neuroscience*, 1. <https://doi.org/10.3389/fnins.2010.00161>
- Miller, E. K., & Cohen, J. D. (2001). An Integrative Theory of Prefrontal Cortex Function. *Annual Review of Neuroscience*, 24(1), 167–202. <https://doi.org/10.1146/annurev.neuro.24.1.167>
- Monosov, I. E. (2017). Anterior cingulate is a source of valence-specific information about value and uncertainty. *Nature Communications*, 8(1), 134. <https://doi.org/10.1038/s41467-017-00072-y>
- Moritz, C. T., Perlmutter, S. I., & Fetz, E. E. (2008). Direct control of paralysed muscles by cortical neurons. *Nature*, 456(7222), 639–642. <https://doi.org/10.1038/nature07418>
- Moxon, K. A., & Foffani, G. (2015). Brain-Machine Interfaces beyond Neuroprosthetics. *Neuron*, 86(1), 55–67. <https://doi.org/10.1016/j.neuron.2015.03.036>
- Muhammad, R., Wallis, J. D., & Miller, E. K. (2006). A Comparison of Abstract Rules in the Prefrontal Cortex, Premotor Cortex, Inferior Temporal Cortex, and Striatum. *Journal of Cognitive Neuroscience*, 18(6), 974–989. <https://doi.org/10.1162/jocn.2006.18.6.974>
- Musallam, S., Corneil, B. D., Greger, B., Scherberger, H., & Andersen, R. A. (2004). Cognitive Control Signals for Neural Prosthetics. *Science*, 305(5681), 258–262. <https://doi.org/10.1126/science.1097938>
- Nakamura, K., & Hikosaka, O. (2006). Facilitation of Saccadic Eye Movements by Postsaccadic Electrical Stimulation in the Primate Caudate. *Journal of Neuroscience*, 26(50), 12885–12895. <https://doi.org/10.1523/JNEUROSCI.3688-06.2006>
- Neely, R. M., Koralek, A. C., Athalye, V. R., Costa, R. M., & Carmena, J. M. (2018). Volitional Modulation of Primary Visual Cortex Activity Requires the Basal Ganglia. *Neuron*, 97(6), 1356–1368.e4. <https://doi.org/10.1016/j.neuron.2018.01.051>

- Nuyujukian, P., Albites Sanabria, J., Saab, J., Pandarinath, C., Jarosiewicz, B., Blabe, C. H., Franco, B., Mernoff, S. T., Eskandar, E. N., Simeral, J. D., Hochberg, L. R., Shenoy, K. V., & Henderson, J. M. (2018). Cortical control of a tablet computer by people with paralysis. *PLOS ONE*, *13*(11), e0204566. <https://doi.org/10.1371/journal.pone.0204566>
- Oby, E. R., Golub, M. D., Hennig, J. A., Degenhart, A. D., Tyler-Kabara, E. C., Yu, B. M., Chase, S. M., & Batista, A. P. (2019). New neural activity patterns emerge with long-term learning. *Proceedings of the National Academy of Sciences*, *116*(30), 15210–15215. <https://doi.org/10.1073/pnas.1820296116>
- Okun, M., Steinmetz, N. A., Cossell, L., Iacaruso, M. F., Ko, H., Barthó, P., Moore, T., Hofer, S. B., Mrsic-Flogel, T. D., Carandini, M., & Harris, K. D. (2015). Diverse coupling of neurons to populations in sensory cortex. *Nature*, *521*(7553), 511–515. <https://doi.org/10.1038/nature14273>
- Olsen, S., Zhang, J., Liang, K.-F., Lam, M., Riaz, U., & Kao, J. C. (2021). An artificial intelligence that increases simulated brain–computer interface performance. *Journal of Neural Engineering*, *18*(4), 046053. <https://doi.org/10.1088/1741-2552/abfaaa>
- Orsborn, A. L., & Carmena, J. M. (2013). Creating new functional circuits for action via brain–machine interfaces. *Frontiers in Computational Neuroscience*, *7*. <https://doi.org/10.3389/fncom.2013.00157>
- Orsborn, A. L., Dangi, S., Moorman, H. G., & Carmena, J. M. (2012). Closed-Loop Decoder Adaptation on Intermediate Time-Scales Facilitates Rapid BMI Performance Improvements Independent of Decoder Initialization Conditions. *IEEE Transactions on Neural Systems and Rehabilitation Engineering*, *20*(4), 468–477. <https://doi.org/10.1109/TNSRE.2012.2185066>
- Orsborn, A. L., Moorman, H. G., Overduin, S. A., Shanechi, M. M., Dimitrov, D. F., & Carmena, J. M. (2014). Closed-Loop Decoder Adaptation Shapes Neural Plasticity for Skillful Neuroprosthetic Control. *Neuron*, *82*(6), 1380–1393. <https://doi.org/10.1016/j.neuron.2014.04.048>
- Orsborn, A. L., & Pesaran, B. (2017). Parsing learning in networks using brain–machine interfaces. *Current Opinion in Neurobiology*, *46*, 76–83. <https://doi.org/10.1016/j.conb.2017.08.002>
- Packard, M. G., & Knowlton, B. J. (2002). Learning and Memory Functions of the Basal Ganglia. *Annual Review of Neuroscience*, *25*(1), 563–593. <https://doi.org/10.1146/annurev.neuro.25.112701.142937>
- Pandarinath, C., Ames, K. C., Russo, A. A., Farshchian, A., Miller, L. E., Dyer, E. L., & Kao, J. C. (2018). Latent Factors and Dynamics in Motor Cortex and Their Application to Brain–Machine Interfaces. *The Journal of Neuroscience*, *38*(44), 9390–9401. <https://doi.org/10.1523/JNEUROSCI.1669-18.2018>

- Pandarinath, C., Gilja, V., Blabe, C. H., Nuyujukian, P., Sarma, A. A., Sorice, B. L., Eskandar, E. N., Hochberg, L. R., Henderson, J. M., & Shenoy, K. V. (2015). Neural population dynamics in human motor cortex during movements in people with ALS. *ELife*, 4, e07436. <https://doi.org/10.7554/eLife.07436>
- Pandarinath, C., Nuyujukian, P., Blabe, C. H., Sorice, B. L., Saab, J., Willett, F. R., Hochberg, L. R., Shenoy, K. V., & Henderson, J. M. (2017). High performance communication by people with paralysis using an intracortical brain-computer interface. *ELife*, 6, e18554. <https://doi.org/10.7554/eLife.18554>
- Parastarfeizabadi, M., & Kouzani, A. Z. (2017). Advances in closed-loop deep brain stimulation devices. *Journal of NeuroEngineering and Rehabilitation*, 14(1), 79. <https://doi.org/10.1186/s12984-017-0295-1>
- Parent, A., & Hazrati, L.-N. (1995). *Functional anatomy of the basal ganglia. II. The place of subthalamic nucleus and external pallidum in basal ganglia circuitry*. 27.
- Pasupathy, A., & Miller, E. K. (2005). Different time courses of learning-related activity in the prefrontal cortex and striatum. *Nature*, 433(7028), 873–876. <https://doi.org/10.1038/nature03287>
- Paulus, M. P., & Yu, A. J. (2012). Emotion and decision-making: Affect-driven belief systems in anxiety and depression. *Trends in Cognitive Sciences*, 16(9), 476–483. <https://doi.org/10.1016/j.tics.2012.07.009>
- Paxinos, G., Huang, X.-F., & Toga, A. W. (2000). *The Rhesus Monkey Brain in Stereotaxic Coordinates*.
- Pedregosa, F., Varoquaux, G., Gramfort, A., Michel, V., Thirion, B., Grisel, O., Blondel, M., Prettenhofer, P., Weiss, R., Dubourg, V., Vanderplas, J., Passos, A., & Cournapeau, D. (2011). Scikit-learn: Machine Learning in Python. *MACHINE LEARNING IN PYTHON*, 6.
- Perich, M. G., Gallego, J. A., & Miller, L. E. (2018). A Neural Population Mechanism for Rapid Learning. *Neuron*, 100(4), 964–976.e7. <https://doi.org/10.1016/j.neuron.2018.09.030>
- Petrides, M., & Pandya, D. N. (2006). Efferent association pathways originating in the caudal prefrontal cortex in the macaque monkey. *The Journal of Comparative Neurology*, 498(2), 227–251. <https://doi.org/10.1002/cne.21048>
- Pfurtscheller, G., Graitmann, B., Huggins, J. E., Levine, S. P., & Schuh, L. A. (2003). Spatiotemporal patterns of beta desynchronization and gamma synchronization in corticographic data during self-paced movement. *Clinical Neurophysiology*, 114(7), 1226–1236. [https://doi.org/10.1016/S1388-2457\(03\)00067-1](https://doi.org/10.1016/S1388-2457(03)00067-1)
- Polikov, V. S., Tresco, P. A., & Reichert, W. M. (2005). Response of brain tissue to chronically implanted neural electrodes. *Journal of Neuroscience Methods*, 148(1), 1–18. <https://doi.org/10.1016/j.jneumeth.2005.08.015>

- Premereur, E., Janssen, P., & Vanduffel, W. (2018). Functional MRI in Macaque Monkeys during Task Switching. *The Journal of Neuroscience*, 38(50), 10619–10630. <https://doi.org/10.1523/JNEUROSCI.1539-18.2018>
- Provenza, N. R., Paulk, A. C., Peled, N., Restrepo, M. I., Cash, S. S., Dougherty, D. D., Eskandar, E. N., Borton, D. A., & Widge, A. S. (2019). Decoding task engagement from distributed network electrophysiology in humans. *Journal of Neural Engineering*, 16(5), 056015. <https://doi.org/10.1088/1741-2552/ab2c58>
- Qiao, S., Brown, K. A., Orsborn, A. L., Ferrentino, B., & Pesaran, B. (2016). Development of semi-chronic microdrive system for large-scale circuit mapping in macaque mesolimbic and basal ganglia systems. *2016 38th Annual International Conference of the IEEE Engineering in Medicine and Biology Society (EMBC)*, 5825–5828. <https://doi.org/10.1109/EMBC.2016.7592052>
- Rangel, A., Camerer, C., & Montague, P. R. (2008). A framework for studying the neurobiology of value-based decision making. *Nature Reviews Neuroscience*, 9(7), 545–556. <https://doi.org/10.1038/nrn2357>
- Rapeaux, A. B., & Constandinou, T. G. (2021). Implantable brain machine interfaces: First-in-human studies, technology challenges and trends. *Current Opinion in Biotechnology*, 72, 102–111. <https://doi.org/10.1016/j.copbio.2021.10.001>
- Rich, E. L., & Wallis, J. D. (2016). Decoding subjective decisions from orbitofrontal cortex. *Nature Neuroscience*, 19(7), 973–980. <https://doi.org/10.1038/nn.4320>
- Rodriguez-Oroz, M. C., Obeso, J. A., Lang, A. E., Houeto, J.-L., Pollak, P., Rehncrona, S., Kulisevsky, J., Albanese, A., Volkmann, J., Hariz, M. I., Quinn, N. P., Speelman, J. D., Guridi, J., Zamarbide, I., Gironell, A., Molet, J., Pascual-Sedano, B., Pidoux, B., Bonnet, A. M., ... Van Blercom, N. (2005). Bilateral deep brain stimulation in Parkinson's disease: A multicentre study with 4 years follow-up. *Brain*, 128(10), 2240–2249. <https://doi.org/10.1093/brain/awh571>
- Rosa, M., Arlotti, M., Marceglia, S., Cogiamanian, F., Ardolino, G., Fonzo, A. D., Lopiano, L., Scelzo, E., Merola, A., Locatelli, M., Rampini, P. M., & Priori, A. (2017). Adaptive deep brain stimulation controls levodopa-induced side effects in Parkinsonian patients: DBS Controls Levodopa-Induced Side Effects. *Movement Disorders*, 32(4), 628–629. <https://doi.org/10.1002/mds.26953>
- Rosin, B., Slovik, M., Mitelman, R., Rivlin-Etzion, M., Haber, S. N., Israel, Z., Vaadia, E., & Bergman, H. (2011). Closed-Loop Deep Brain Stimulation Is Superior in Ameliorating Parkinsonism. *Neuron*, 72(2), 370–384. <https://doi.org/10.1016/j.neuron.2011.08.023>
- Russo, A. A., Bittner, S. R., Perkins, S. M., Seely, J. S., London, B. M., Lara, A. H., Miri, A., Marshall, N. J., Kohn, A., Jessell, T. M., Abbott, L. F., Cunningham, J. P., & Churchland, M. M.

- (2018). Motor Cortex Embeds Muscle-like Commands in an Untangled Population Response. *Neuron*, 97(4), 953–966.e8. <https://doi.org/10.1016/j.neuron.2018.01.004>
- Sadtler, P. T., Quick, K. M., Golub, M. D., Chase, S. M., Ryu, S. I., Tyler-Kabara, E. C., Yu, B. M., & Batista, A. P. (2014). Neural constraints on learning. *Nature*, 512(7515), 423–426. <https://doi.org/10.1038/nature13665>
- Saito, N., Mushiake, H., Sakamoto, K., Itoyama, Y., & Tanji, J. (2005). Representation of Immediate and Final Behavioral Goals in the Monkey Prefrontal Cortex during an Instructed Delay Period. *Cerebral Cortex*, 15(10), 1535–1546. <https://doi.org/10.1093/cercor/bhi032>
- Samejima, K. (2005). Representation of Action-Specific Reward Values in the Striatum. *Science*, 310(5752), 1337–1340. <https://doi.org/10.1126/science.1115270>
- Sani, O. G., Yang, Y., Lee, M. B., Dawes, H. E., Chang, E. F., & Shanechi, M. M. (2018). Mood variations decoded from multi-site intracranial human brain activity. *Nature Biotechnology*, 36(10), 954–961. <https://doi.org/10.1038/nbt.4200>
- Santacruz, S. R., Rich, E. L., Wallis, J. D., & Carmena, J. M. (2017). Caudate Microstimulation Increases Value of Specific Choices. *Current Biology*, 27(21), 3375–3383.e3. <https://doi.org/10.1016/j.cub.2017.09.051>
- Scangos, K. W., Khambhati, A. N., Daly, P. M., Makhoul, G. S., Sugrue, L. P., Zamanian, H., Liu, T. X., Rao, V. R., Sellers, K. K., Dawes, H. E., Starr, P. A., Krystal, A. D., & Chang, E. F. (2021). Closed-loop neuromodulation in an individual with treatment-resistant depression. *Nature Medicine*, 27(10), 1696–1700. <https://doi.org/10.1038/s41591-021-01480-w>
- Schalk, G., Leuthardt, E. C., Brunner, P., Ojemann, J. G., Gerhardt, L. A., & Wolpaw, J. R. (2008). Real-time detection of event-related brain activity. *NeuroImage*, 43(2), 245–249. <https://doi.org/10.1016/j.neuroimage.2008.07.037>
- Serruya, M. D., Hastopoulos, N. G., Paninski, L., Fellows, M. R., & Donoghue, J. P. (2002). Instant neural control of a movement signal. *Nature*, 416(6877), 141–142. <https://doi.org/10.1038/416142a>
- Shanechi, M. M. (2019). Brain–machine interfaces from motor to mood. *Nature Neuroscience*, 22(10), 1554–1564. <https://doi.org/10.1038/s41593-019-0488-y>
- Shenoy, K. V., & Carmena, J. M. (2014). Combining Decoder Design and Neural Adaptation in Brain-Machine Interfaces. *Neuron*, 84(4), 665–680. <https://doi.org/10.1016/j.neuron.2014.08.038>
- Shenoy, K. V., Sahani, M., & Churchland, M. M. (2013). Cortical Control of Arm Movements: A Dynamical Systems Perspective. *Annual Review of Neuroscience*, 36(1), 337–359. <https://doi.org/10.1146/annurev-neuro-062111-150509>

- Shepherd, G. M. G. (2013). Corticostriatal connectivity and its role in disease. *Nature Reviews Neuroscience*, 14(4), 278–291. <https://doi.org/10.1038/nrn3469>
- Siegel, M., Donner, T. H., & Engel, A. K. (2012). Spectral fingerprints of large-scale neuronal interactions. *Nature Reviews Neuroscience*, 13(2), 121–134. <https://doi.org/10.1038/nrn3137>
- Silversmith, D. B., Abiri, R., Hardy, N. F., Natraj, N., Tu-Chan, A., Chang, E. F., & Ganguly, K. (2021). Plug-and-play control of a brain–computer interface through neural map stabilization. *Nature Biotechnology*, 39(3), 326–335. <https://doi.org/10.1038/s41587-020-0662-5>
- Sing, G. C., & Smith, M. A. (2010). Reduction in Learning Rates Associated with Anterograde Interference Results from Interactions between Different Timescales in Motor Adaptation. *PLoS Computational Biology*, 6(8), e1000893. <https://doi.org/10.1371/journal.pcbi.1000893>
- 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. <https://doi.org/10.1088/1741-2560/9/2/026004>
- Sternad, D. (2018). It's not (only) the mean that matters: Variability, noise and exploration in skill learning. *Current Opinion in Behavioral Sciences*, 20, 183–195. <https://doi.org/10.1016/j.cobeha.2018.01.004>
- Sun, F. T., Morrell, M. J., & Wharen, R. E. (2008). Responsive cortical stimulation for the treatment of epilepsy. *Neurotherapeutics*, 5(1), 68–74. <https://doi.org/10.1016/j.nurt.2007.10.069>
- Sun, X., O'Shea, D. J., Golub, M. D., Trautmann, E. M., Vyas, S., Ryu, S. I., & Shenoy, K. V. (2022). Cortical preparatory activity indexes learned motor memories. *Nature*, 602(7896), 274–279. <https://doi.org/10.1038/s41586-021-04329-x>
- Suresh, A. K., Goodman, J. M., Okorokova, E. V., Kaufman, M., Hatsopoulos, N. G., & Bensmaia, S. J. (2020). Neural population dynamics in motor cortex are different for reach and grasp. *ELife*, 9, e58848. <https://doi.org/10.7554/eLife.58848>
- Sussillo, D., Stavisky, S. D., Kao, J. C., Ryu, S. I., & Shenoy, K. V. (2016). Making brain–machine interfaces robust to future neural variability. *Nature Communications*, 7(1), 13749. <https://doi.org/10.1038/ncomms13749>
- Sutton, R. S., & Barto, A. G. (1998). *Reinforcement Learning: An Introduction*. The MIT Press.
- Swann, N. C., de Hemptinne, C., Thompson, M. C., Miocinovic, S., Miller, A. M., Gilron, R., Ostrem, J. L., Chizeck, H. J., & Starr, P. A. (2018). Adaptive deep brain stimulation for Parkinson's disease using motor cortex sensing. *Journal of Neural Engineering*, 15(4), 046006. <https://doi.org/10.1088/1741-2552/aabc9b>

- Tanji, J., & Hoshi, E. (2008). Role of the Lateral Prefrontal Cortex in Executive Behavioral Control. *Physiological Reviews*, 88(1), 37–57.  
<https://doi.org/10.1152/physrev.00014.2007>
- Tanji, J., Shima, K., & Mushiake, H. (2007). Concept-based behavioral planning and the lateral prefrontal cortex. *Trends in Cognitive Sciences*, 11(12), 528–534.  
<https://doi.org/10.1016/j.tics.2007.09.007>
- Taylor, D. M., Tillery, S. I. H., & Schwartz, A. B. (2002). Direct Cortical Control of 3D Neuroprosthetic Devices. *Science*, 296(5574), 1829–1832.  
<https://doi.org/10.1126/science.1070291>
- Tumer, E. C., & Brainard, M. S. (2007). *Performance variability enables adaptive plasticity of 'crystallized' adult birdsong*. 450, 6.
- Turner, J. N., Shain, W., Szarowski, D. H., Andersen, M., Martins, S., Isaacson, M., & Craighead, H. (1999). Cerebral Astrocyte Response to Micromachined Silicon Implants. *Experimental Neurology*, 156(1), 33–49. <https://doi.org/10.1006/exnr.1998.6983>
- Vasileva, L. N., & Bondar, I. V. (2021). Long-Term Stable Recording of Single-Neuron Spike Activity in the Amygdala in Conscious Rabbits. *Neuroscience and Behavioral Physiology*, 51(3), 322–331. <https://doi.org/10.1007/s11055-021-01075-5>
- Velliste, M., Perel, S., Spalding, M. C., Whitford, A. S., & Schwartz, A. B. (2008). Cortical control of a prosthetic arm for self-feeding. *Nature*, 453(7198), 1098–1101.  
<https://doi.org/10.1038/nature06996>
- Vyas, S., Even-Chen, N., Stavisky, S. D., Ryu, S. I., Nuyujukian, P., & Shenoy, K. V. (2018). Neural Population Dynamics Underlying Motor Learning Transfer. *Neuron*, 97(5), 1177–1186.e3.  
<https://doi.org/10.1016/j.neuron.2018.01.040>
- Vyas, S., Golub, M. D., Sussillo, D., & Shenoy, K. V. (2020). Computation Through Neural Population Dynamics. *Annual Review of Neuroscience*, 43(1), 249–275.  
<https://doi.org/10.1146/annurev-neuro-092619-094115>
- Vyas, S., O'Shea, D. J., Ryu, S. I., & Shenoy, K. V. (2020). Causal Role of Motor Preparation during Error-Driven Learning. *Neuron*, 106(2), 329–339.e4.  
<https://doi.org/10.1016/j.neuron.2020.01.019>
- Walton, M. E., Croxson, P. L., Behrens, T. E. J., Kennerley, S. W., & Rushworth, M. F. S. (2007). Adaptive decision making and value in the anterior cingulate cortex. *NeuroImage*, 36, T142–T154. <https://doi.org/10.1016/j.neuroimage.2007.03.029>
- Wander, J. D., Blakely, T., Miller, K. J., Weaver, K. E., Johnson, L. A., Olson, J. D., Fetz, E. E., Rao, R. P. N., & Ojemann, J. G. (2013). Distributed cortical adaptation during learning of a brain-computer interface task. *Proceedings of the National Academy of Sciences*, 110(26), 10818–10823. <https://doi.org/10.1073/pnas.1221127110>

- Wander, J. D., Sarma, D., Johnson, L. A., Fetz, E. E., Rao, R. P. N., Ojemann, J. G., & Darvas, F. (2016). Cortico-Cortical Interactions during Acquisition and Use of a Neuroprosthetic Skill. *PLoS Computational Biology*, *12*(8), e1004931. <https://doi.org/10.1371/journal.pcbi.1004931>
- Watanabe, K., & Hikosaka, O. (2005). Immediate Changes in Anticipatory Activity of Caudate Neurons Associated With Reversal of Position-Reward Contingency. *Journal of Neurophysiology*, *94*(3), 1879–1887. <https://doi.org/10.1152/jn.00012.2005>
- Widge, A. S., Malone, D. A., & Dougherty, D. D. (2018). Closing the Loop on Deep Brain Stimulation for Treatment-Resistant Depression. *Frontiers in Neuroscience*, *12*, 175. <https://doi.org/10.3389/fnins.2018.00175>
- Willett, F. R., Avansino, D. T., Hochberg, L. R., Henderson, J. M., & Shenoy, K. V. (2021). High-performance brain-to-text communication via handwriting. *Nature*, *593*(7858), 249–254. <https://doi.org/10.1038/s41586-021-03506-2>
- Williams, Z. M., & Eskandar, E. N. (2006). Selective enhancement of associative learning by microstimulation of the anterior caudate. *Nature Neuroscience*, *9*(4), 562–568. <https://doi.org/10.1038/nn1662>
- Wilson, G. H., Stavisky, S. D., Willett, F. R., Avansino, D. T., Kelemen, J. N., Hochberg, L. R., Henderson, J. M., Druckmann, S., & Shenoy, K. V. (2020). Decoding spoken English from intracortical electrode arrays in dorsal precentral gyrus. *Journal of Neural Engineering*, *17*(6), 066007. <https://doi.org/10.1088/1741-2552/abbfef>
- Wolpaw, J. R., Birbaumer, N., McFarland, D. J., Pfurtscheller, G., & Vaughan, T. M. (2002). Brain-computer interfaces for communication and control. *Clinical Neurophysiology*, *25*.
- Wolpaw, J. R., & McFarland, D. J. (2004). Control of a two-dimensional movement signal by a noninvasive brain-computer interface in humans. *Proceedings of the National Academy of Sciences*, *101*(51), 17849–17854. <https://doi.org/10.1073/pnas.0403504101>
- Wolpaw, J. R., McFarland, D. J., Neat, G. W., & Forneris, C. A. (1991). An EEG-based brain-computer interface for cursor control. *Electroencephalography and Clinical Neurophysiology*, *78*(3), 252–259. [https://doi.org/10.1016/0013-4694\(91\)90040-B](https://doi.org/10.1016/0013-4694(91)90040-B)
- Wu, W., Black, M. J., Gao, Y., Serruya, M., Shaikhouni, A., Donoghue, J. P., & Bienenstock, E. (n.d.). *Neural Decoding of Cursor Motion Using a Kalman Filter*. 8.
- Wu, Y.-C., Liao, Y.-S., Yeh, W.-H., Liang, S.-F., & Shaw, F.-Z. (2021). Directions of Deep Brain Stimulation for Epilepsy and Parkinson's Disease. *Frontiers in Neuroscience*, *15*, 680938. <https://doi.org/10.3389/fnins.2021.680938>
- Yoo, S. B. M., Tu, J. C., & Hayden, B. Y. (2021). Multicentric tracking of multiple agents by anterior cingulate cortex during pursuit and evasion. *Nature Communications*, *12*(1), 1985. <https://doi.org/10.1038/s41467-021-22195-z>

- You, A., Singhal, A., Moorman, H., Gowda, S., & Carmena, J. M. (2019). Neural Correlates of Control of a Kinetically Redundant Brain-Machine Interface. *2019 9th International IEEE/EMBS Conference on Neural Engineering (NER)*, 554–557. <https://doi.org/10.1109/NER.2019.8717010>
- Yu, B. M., Cunningham, J. P., Santhanam, G., Ryu, S. I., Shenoy, K. V., & Sahani, M. (2009). Gaussian-Process Factor Analysis for Low-Dimensional Single-Trial Analysis of Neural Population Activity. *Journal of Neurophysiology*, *102*(1), 614–635. <https://doi.org/10.1152/jn.90941.2008>
- Zacksenhouse, M., Lebedev, M. A., Carmena, J. M., O'Doherty, J. E., Henriquez, C., & Nicolelis, M. A. L. (2007). Cortical Modulations Increase in Early Sessions with Brain-Machine Interface. *PLoS ONE*, *2*(7), e619. <https://doi.org/10.1371/journal.pone.0000619>
- Zhang, Y., Chen, Y., Bressler, S. L., & Ding, M. (2008). Response preparation and inhibition: The role of the cortical sensorimotor beta rhythm. *Neuroscience*, *156*(1), 238–246. <https://doi.org/10.1016/j.neuroscience.2008.06.061>
- Zhou, X., Tien, R. N., Ravikumar, S., & Chase, S. M. (2019). Distinct types of neural reorganization during long-term learning. *Journal of Neurophysiology*, *121*(4), 1329–1341. <https://doi.org/10.1152/jn.00466.2018>
- Zippi, E. L., You, A. K., Ganguly, K., & Carmena, J. M. (2021). *Selective modulation of population dynamics during neuroprosthetic skill learning* [Preprint]. Neuroscience. <https://doi.org/10.1101/2021.01.08.425917>