

UC San Diego

UC San Diego Electronic Theses and Dissertations

Title

Learning, execution, and bias of movements in the mammalian cortex

Permalink

<https://escholarship.org/uc/item/0zn1b6w9>

Author

Dahlen, Jeffrey Edward

Publication Date

2017

Peer reviewed|Thesis/dissertation

UNIVERSITY OF CALIFORNIA, SAN DIEGO

Learning, execution, and bias of movements in the mammalian cortex

A dissertation submitted in partial satisfaction of the requirements for the degree

Doctor of Philosophy

in

Neurosciences with a Specialization in Computational Neuroscience

by

Jeffrey Edward Dahlen

Committee in charge:

Professor Takaki Komiyama, Chair
Professor Douglas Nitz
Professor Samuel Pfaff
Professor Massimo Scanziani
Professor Mark Tuszynski

2017

The Dissertation of Jeffrey Edward Dahlen is approved, and is acceptable in quality and form for publication on microfilm and electronically:

Chair

University of California, San Diego

2017

DEDICATION

For Mom, Dad, and Jason.

TABLE OF CONTENTS

Signature Page	iii
Dedication	iv
Table of Contents	v
List of Figures	vi
Acknowledgements.....	viii
Vita	xii
Abstract of the Dissertation.....	xiv
Chapter 1. Introduction	1
Chapter 2. Motor cortex-independent movement execution induced by long-term learning	10
2.1 Summary	10
2.2 Introduction	10
2.3 Results.....	12
2.4 Discussion	21
2.5 Methods.....	25
Chapter 3. History-based action selection bias in posterior parietal cortex	51
3.1 Summary	51
3.2 Introduction	51
3.3 Results.....	53
3.4 Discussion	64
3.5 Methods.....	68
Chapter 4. General Discussion	102
References	110

LIST OF FIGURES

Figure 2.1 Two-direction forelimb movement task.	37
Figure 2.2 Expert mice exhibit bias for more stereotyped movements.	38
Figure 2.3 M1 activity in expert mice is tuned for non-biased movements.....	39
Figure 2.4 M1 inactivation does not inhibit initiation of expert movements.....	41
Figure 2.5 M1 inactivation enhances existing bias in expert mice	42
Figure 2.6 M1 inactivation during early learning impairs movement initiation and execution	43
Figure 2.7 Biased movements are more stereotyped than non-biased movements	44
Figure 2.8 Optogenetic inactivation of PV inhibitory neurons inactivates cortex.....	45
Figure 2.9 Control sessions of expert inactivation experiments	47
Figure 2.10 Kinematics of movements in early-learning are not reliably different between the two directions	48
Figure 2.11 Control cohort of early learning inactivation experiments	49
Figure 3.1 Task.	83
Figure 3.2 Choice-outcome history biases future decision, driving choice variability.	84
Figure 3.3 Pre-stimulus activity during the inter-trial interval in posterior parietal cortex (PPC) predicts the upcoming choice.....	86
Figure 3.4 Pre-stimulus choice-selectivity in PPC encodes a mixture of history and bias information.	87
Figure 3.5 History tuning of PPC neurons reflects the history dependency of bias.	88
Figure 3.6 Pre-stimulus population activity in PPC reflects the history-dependent bias and previous trial history.	89
Figure 3.7 Inactivating pre-stimulus activity in PPC alters internal bias.....	91
Figure 3.8 Inactivating pre-stimulus activity in PPC weakens the history dependency of choice.	93
Figure 3.9 Two-alternative forced-choice tasks using the joystick apparatus.....	94
Figure 3.10 Stimulus dependency increases, while history dependency decreases in non-memory task compared to memory task.	95
Figure 3.11 ITI choice selectivity in PPC is not a reflection of ongoing limb movements.....	96
Figure 3.12 ITI choice tuning persists in early and late execution trials.	97
Figure 3.13 The preferred directions of choice selectivity are equally distributed.....	98

Figure 3.14 PPC ITI inactivation effect on internal model accuracy, task performance, reaction time, and movement time.	99
Figure 3.15 Control sessions.	100

ACKNOWLEDGEMENTS

I have been incredibly fortunate to have parents that have provided my brother and I with every possible tool to succeed. Their foresight and selfless dedication to our success moved them from their friends, a home they had built with their bare hands, and a reasonable commute, to the incredible city of San Juan Capistrano – a place that I will forever call Home – in order to provide us with a stable community and quality education. They fostered my curiosity by allowing me disassemble nearly every appliance in the house (despite varying levels of successful reassembly), and they encouraged me to explore all of my interests from rubber band planes, to model trains, to remote control automobiles.

In addition to the incredible humans that are my parents I have in my younger, bigger, better-looking brother, a hero that I have and will always look up to. My brother, like my parents, is a remarkable and selfless human who has always been a source of encouragement, a voice of reason, and partner in crime for the seemingly boundless energy that we share.

Beyond my immediate family, I have a very close extended family. When I was lost in science for months at a time they never questioned my absences from family get-togethers, and when I felt discouraged their unwavering faith provided me with motivation and confidence. Without the unconditional and steadfast support of my family I would be nowhere, and without their incredible examples of what it is to be a good human I would be no one.

All of my academic successes are a result of the integrated effort of the brilliant, patient, and passionate teachers and mentors I had throughout my primary schooling. I am eternally grateful for the joy and excitement you brought into the classroom, your collective efforts contributed to my desire to never stop learning.

Scientifically I am forever indebted to Dr. Nathan Urban who accepted me into his lab with open arms as a young and naïve undergraduate at Carnegie Mellon University. Beyond allowing me the freedom to explore exactly what a scientist is, and what a scientist does, Nathan encouraged me to pursue my graduate degree and provided me with skills that have been invaluable throughout my graduate career. Nathan's lab, the Great Hall of Brain Science, and the rest of the Carnegie Mellon science community – including James Fitzpatrick, Krishnan Padmanabhan, Anne-Marie Oswald, and Shreejoy Tripathy (among many, many others) – will always hold a special place in my heart.

At UC San Diego I was fortunate to be one of the first graduate students in Dr. Takaki Komiyama's lab. In Takaki's lab I got to do incredibly exciting, inspiring, and challenging work at the forefront of neuroscience research. Takaki gave me the freedom to explore and take ownership of my project, for which I am incredibly thankful and proud. I am also very thankful to my thesis committee members – Doug Nitz, Sam Pfaff, Massimo Scanziani and Mark Tuszynski – for providing me with thoughtful and constructive feedback throughout my graduate work.

Beyond science, Takaki has cultivated a collaborative and fun lab environment, full of people that I am proud to consider close friends and colleagues. A special thanks goes to Andy Bo Bandy Peters, Monica Chu, Wankun Lovie Li, Aki Mitani, Madan Mukundan, Jun Lee, Simon Chen, EunJung Hwang, Nathan Hendricksonley, Haixin Liu, Anna Kim, Keelin Oneil, as well as the rest of the Komiyama lab past and present, for the stimulating conversation, endless support, and most importantly friendship.

As an undergraduate, completing my graduate work at UC San Diego, a top neuroscience program in the country, was a pipe dream. Living the dream over the past six years of my life has been an incredible gift that I have thoroughly enjoyed. However

this experience would not have been the same without my brilliant classmates. A world of thanks go to my classmates – Higa, Sirawaj Itthipuripat, Ashley Juavinett, Jason Keller, my fellow Carnegie Mellon Alum Stacy Kurnikova – and the rest of the Neurosciences Graduate Program, for always lending an ear and a beer when I needed it most, and for standing alongside me for some of my fondest memories of graduate school. I am thankful to have met you all, and I am excited to see where your careers take you.

As my friends know, the academic side of my life must be balanced with a physical outlet. In graduate school I was welcomed into the UC San Diego Triathlon and Cycling Teams with open arms, who became like a second family to me. Thanks go to all of my teammates and training partners for distracting me from the lab, and keeping me sane for the past 6 years.

Living in San Diego has been wonderful, but my roommates throughout graduate school made it even more wonderful. Thanks go to Pedro Pezzin, Brooks Taylor, and Kyle Begovich for tolerating all of my surfboards, bicycles, and crazy schedule. You are Saints for being so patient, caring, and kind to me throughout my graduate studies.

Lastly, but not leastly, I would be remiss if I didn't give thanks to my fellow Carnegie Mellon swim teammates – Scott Moorby and Ming Hsu – for staying in touch and keeping it light for all these years later. Daily chats with you guys helped me keep things in perspective, and for that, I thank you.

Chapter 2 currently being prepared for submission and will include Dr. EunJung Hwang, Madan Mukundan and Dr. Takaki Komiyama as authors. This material was included with the generous consent of all authors. The dissertation author is a primary researcher on this work.

Chapter 3 currently being prepared for submission and will include Dr. EunJung Hwang, Madan Mukundan and Dr. Takaki Komiyama as authors. This material was included with the generous consent of all authors. The dissertation author is a primary researcher on this work.

VITA

- 2010 Bachelor of Science in Biological Sciences
Carnegie Mellon University, Pittsburgh, PA
- 2014 Master of Science in Neuroscience
University of California, San Diego, La Jolla, CA
- 2017 Doctor of Philosophy in Neurosciences
with a Specialization in Computational Neuroscience
University of California, San Diego, La Jolla, CA

PUBLICATIONS

Hovis KR, Ramnath R, **Dahlen JE**, LaRocca GM, Bier MR, Urban NN. Activity Regulates Functional Connectivity from the Vomeronasal Organ to the Accessory Olfactory Bulb. *Journal of Neuroscience* 32(23):7907-16 (2012).

Dahlen JE, Jimenez DA, Gerkin RC, Urban NN. Morphological Analysis of Activity-Reduced Adult-Born Neurons in the Mouse Olfactory Bulb. *Frontiers in Neuroscience* 5:66 (2011).

CONFERENCE PRESENTATIONS

Dahlen JE, Hwang EJ, Mukundan M, Komiyama T. "The role of the primary motor cortex in the execution of skilled movements." Wiring and Functional Principles of Neural Circuits (2016), San Diego CA

Dahlen JE, Hwang EJ, Mukundan M, Komiyama T. "Planning and execution of movements revealed by in vivo calcium imaging in mice." Society for Neuroscience 46th Annual Meeting (2016), San Diego CA

Hwang EJ, **Dahlen JE**, Mukundan M, Komiyama T. "Internally-generated action-selection bias in the posterior parietal cortex." Society for Neuroscience 46th Annual Meeting (2016), San Diego CA

Dahlen JE, Hwang EJ, Komiyama T. "Movement preparation and execution in the mouse motor cortex during a novel head-fixed sensorimotor discrimination task." Society for Neural Control of Movement 25th Annual Meeting (2015), Charleston SC

Dahlen JE, Hwang EJ, Komiyama T. "A novel head-fixed sensorimotor discrimination task to investigate movement preparation and execution in the mouse motor cortex." Society for Neuroscience 44th Annual Meeting (2014), Washington DC

Dahlen JE, Jimenez DA, Urban NN. "Morphological Analysis of Activity Reduced Adult Born Neurons in the Mouse Main Olfactory Bulb." Society for Neuroscience 40th Annual Meeting (2009), Chicago IL

AWARDS

- 2015 National Institutes of Health, Ruth L. Kirschstein National Research Service Award Predoctoral Fellowship
- 2013 Society for Neuroscience, Neuroscience Scholars Program
- 2012 National Science Foundation, Alliance for Graduate Education and the Professoriate Fellowship
- 2006 Carnegie Mellon University, Presidential Scholarship

ABSTRACT OF THE DISSERTATION

Learning, execution, and bias of movements in the mammalian cortex

by

Jeffrey Edward Dahlen

Doctor of Philosophy in Neurosciences,
with a Specialization in Computational Neuroscience

University of California, San Diego, 2017

Professor Takaki Komiyama, Chair

Producing movements is a fundamental output of the brain. Over many millennia the brain has undergone significant evolutionary changes to accommodate the increasingly complex ways that we can move and interact with our environment. The mammalian cortex is one such structure that has evolved to provide animals with the capacity to perform highly skilled movements. Among the many cortical areas, the primary motor cortex (M1) has been recognized as a necessary structure for learning to produce such skilled movements. Alternatively, it is also known that M1 is not necessary for executing all movements. The consensus of the field is that whether or not a movement is dependent on M1 is determined entirely by the relative complexity of the movements. Therefore movements are divided into two categories: complex skilled movements that are M1 dependent, and simple movements that can be executed

independent of M1. However there has been no work attempting to understand how the state of learning affects a movement's dependency on M1 for execution. For example, is it possible for movements that are initially M1 dependent to become M1 independent after learning? The second chapter of this dissertation investigates this variable role of M1 in movement execution by combining a novel two-direction forelimb movement task with in vivo two-photon calcium imaging and optogenetic perturbation. Briefly, we found that once movements become sufficiently learned and reproducible, they are indeed no longer dependent on M1 for execution. In conjunction with this result, we found that neurons in M1 are more active and their activity was more consistent for the less-learned, less-stereotyped movements.

Following the acquisition of skills and movements, how do we choose which movements to make given the environmental conditions and understanding of possible outcomes? The third chapter of this thesis investigates this question using a novel memory guided sensorimotor discrimination task, in vivo two-photon calcium imaging and optogenetic perturbation to identify the posterior parietal cortex as a possible locus assembling choice-outcome history prior to decision output.

Chapter 1. Introduction

The necessity to interact with the environment has driven the evolution of the brain in all animals, including humans. In the earliest single-celled organisms, the development of flagellum and chemosensory receptors provided organisms the ability to sense and react their surroundings. Though seemingly trivial, this evolutionary step represented a profoundly important event that allowed organisms to become autonomous beings more capable of survival than their ancestors who were subservient to their environment. Millennia of evolution have since, not surprisingly, expanded the ways in which modern animals and humans now move within and manipulate their environments. The complexities of these movements have thus necessitated the evolution of multifaceted control systems spread across many brain areas to govern the intricacies of the movements we are constantly producing. In addition to movement control systems, the brain has evolved an equally sophisticated decision-making system to select which of the many possible actions to execute.

This dissertation will focus on two cortical areas involved in movement generation and selection, the primary motor cortex (M1) and the posterior parietal cortex (PPC). The remainder of this first chapter will provide a brief introduction to both M1 and PPC. The second chapter will aim to provide a better understanding of the role of M1 in learning and executing complex, skilled movements. Lastly, the third chapter will investigate the role of PPC in decision-making during action selection.

Neuroscience before electricity

Prior to the discovery of electricity, studies of the brain involved understanding how mechanical manipulations of the brain affected an individual. In addition to studies surrounding how different head injuries and lesions affected individuals, the state of the art of neuroscience research was direct mechanical stimulation to stimulate the brain,

such as poking with a scalpel or forceps. Mechanical stimulation experiments showed that subcortical structures and white matter elicited movements and expressions of pain, while stimulations of cortex elicited no such responses (Gross et al., 2007). Based entirely on these results, it was generally accepted that subcortical structures were the primary locus of movement generation. Intriguingly, this conclusion ignored many case studies of head injuries noting loss of motor function following cortical lesions. This led to a variety of theories of the functional role of the cortex, ranging from serving as a support structure to provide nourishment to subcortical tissues below to the study of phrenology which attempted to explain higher cognitive roles of cortex in terms of the shape of the skull (Taylor and Gross, 2003).

The discovery of the motor cortex

An astronomical leap forward in neuroscience research came when Gustav Fritsch and Eduard Hitzig utilized electricity to stimulate cortical tissue (Fritsch and Hitzig, 1870). Lacking proper animal research facilities Fritsch and Hitzig carried out their electrical stimulation experiments on anesthetized and awake dogs in Fritsch's home. Given the experimental conditions and the technologies of the time, the experiments carried out in Fritsch's home are quite remarkable in several respects. First, Fritsch and Hitzig made the first observation that stimulating cortex elicited movements on the contralateral side of the body. This observation was followed shortly after with degeneration studies that reconciled these observations with anatomical data of the pyramidal tract, which serves as the direct output of the motor cortex to the spinal cord (Asanuma and Sakata, 1967a).

Second, they reported a somatotopic representation of body within the cortex. That is, they noted that as they moved their stimulating electrode down what we now know as M1, they were able to elicit movements in different body parts. Soon after their

initial experiments this observation was confirmed in a variety of species, suggesting that the somatotopic representation of body parts within M1 was an evolutionarily conserved property of the cortex (Taylor and Gross, 2003). Interestingly, several decades following the discovery somatotopic representation of M1 it was discovered that the size of the representations of individual body parts were correlated with the relative complexity of movements that could be elicited by them (Penfield and Boldrey, 1937).

Lastly, Fritsch and Hitzig noted that lesioning regions representing specific body parts resulted in a corresponding loss of control of the body part. For example, they observed that following identification of the right foreleg by electrical stimulation, lesion of the area caused partial, though not complete, loss of function of the right foreleg. They hypothesized that because they did not see complete loss of function that there may be other motor centers not affected by the lesion that were capable of partial control of the limb.

Organization of the motor cortices

Following the initial discoveries of Fritsch and Hitzig, additional motor cortical areas, such as the premotor and supplementary motor cortices (among others), were discovered through electrical stimulation and histological studies (Dum and Strick, 2002). With this discovery came the identification of multiple independent somatotopic representations of the body across the different motor cortical areas. Interestingly, while low threshold stimulations of M1 produced simple movements, premotor and supplementary motor areas required higher threshold stimulation and evoked more complex movements (Fulton, 1935; Graziano, 2009). In addition to larger stimulation thresholds, histological studies showed that premotor and supplementary motor areas project to M1 (Pandya and Kuypers, 1969). Together these data suggest a hierarchical

organization of the cortical motor centers, with M1 containing the most basic functional units of movement.

Beyond its interconnectivity with cortical movement areas, M1 receives inputs from contralateral M1, the somatosensory cortex, and thalamus. This extensive connectivity with movement-related brain areas provides additional support for its central role in movement execution as well as motor learning. For example, signals from contralateral M1 are believed to be involved in coordinating bilateral movements (Rouiller et al., 1994), while somatosensory inputs to M1 are believed to couple motor and proprioceptive sensory signals thus facilitating the learning of new movements (Mao et al., 2011).

Apart from its connectivity with other sensorimotor cortical areas, thalamic projections provide extensive inputs to M1 from the basal ganglia and cerebellum. Each of these three movement related areas (the cerebellum, basal ganglia, and M1) serve different, equally important roles in normal movement execution. With connections to both M1 and sensory inputs to the spinal cord, the circuitry of the cerebellum provide online predictive error of movements during execution, which are to be crucial to motor learning in M1 (Shadmehr et al., 2010; Shmuelof and Krakauer, 2011).

Classically, the basal ganglia has been thought to be involved in action selection (Doya, 1999). However based on its extensive connectivity with the motor cortex through the thalamus, it is likely that the basal ganglia plays a more critical role in movement execution. Consistent with this idea, recent work has proposed that the basal ganglia may be contain information about movement kinematics similar to M1 as discussed below (Yttri and Dudman, 2016). In addition to its ability to code kinematic parameters of movement, the basal ganglia is known to be active later in learning, following movement acquisition (Costa et al., 2004). Together these data provide the intriguing possibility of

the basal ganglia as a location for the long-term storage of movements once they are sufficiently learned, and thus no longer require intervention by M1.

The different roles of each of these movement-related brain areas throughout learning, as well as their interconnectedness highlights the incredibly complex nature of how movements are learned and encoded by the brain. Despite relative homology between subcortical movement areas among many species, including mice and rats, there has been little progress in understanding their functions in movement acquisition and control with respect to their interactions with M1 and motor learning. Alternatively the functional properties of M1 are more clearly understood (Kaas, 2008), though remain far from comprehensive.

Movement coding in M1

There has been a large body of work surrounding the role of M1 in learning and executing skilled movements. Modern work on the functional role of M1 came with the advent of extracellular recording techniques, pioneered by Hubel and Wiesel, who recorded from the visual system of cats (Hubel and Wiesel, 1959). Following their work, Edward Evarts became the first to use extracellular recording techniques to record from M1 in awake cats (Evarts, 1965) and later primates performing trained movements (Evarts, 1966). With the adoption of these recording techniques, Evarts himself then established a framework for how M1 encodes movements, by attempting to understand the tuning of M1 neurons with respect to forces exerted during movements (Evarts, 1968).

The idea that neurons in M1 can be tuned to different parameters of movements has been a topic of debate in movement research following Evarts initial experiments. After Evarts' description of force-tuned M1 neurons, subsequent work has described M1 neurons being tuned to direction (Georgopoulos et al., 1982a), amplitude (Fu et al.,

1993; Messier and Kalaska, 2000), speed (Moran and Schwartz, 1999), or velocity and acceleration (Ashe and Georgopoulos; Reina et al., 2001) among many other parameters of movement. Reconciling the many different coding schemes of M1 neurons has thus become an interesting topic of research (Todorov, 2000).

Despite the ability of M1 neurons to correlate with a variety of kinematic parameters, M1 neurons ultimately work together within a population to achieve a concerted output. Similar to studies observing individual tuning of neurons, there have been many attempts to understand the population coding of M1, beginning with relatively simple movement parameters such as direction (Georgopoulos et al., 1986) to more complex descriptions of M1 populations in terms of dynamical systems (Churchland et al., 2012; Shenoy et al., 2013).

Motor learning

Beyond simply understanding the functional principles of M1, a considerable effort has been made to understand how the kinematic parameters of new movements are encoded by M1 during learning. Learning to perform a new movement, or motor learning, has been defined as the acquisition and refinement of a new movement over discrete phases (Luft and Buitrago, 2005), and is known to require M1 (Kawai et al., 2015; Peters et al., 2014).

Throughout each phase of motor learning M1 experiences significant plasticity across multiple temporal and spatial scales (Peters et al., 2017). For example, during acquisition of a novel movement there is an expansion of the somatotopic map revealed by stimulation studies (Kleim et al., 1998; Nudo et al., 1996), as well as significant synaptic changes following learning (Chen et al., 2015; Cichon and Gan, 2015; Rioult-Pedotti et al., 2000; Xu et al., 2009). Activity from single neurons recorded over individual sessions has demonstrated that there are substantial changes to their tuning

properties throughout learning (Costa et al., 2004; Gandolfo et al., 2000; Paz et al., 2003; Rokni et al., 2007). Further, the relatively recent development of genetically encoded calcium indicators (Chen et al., 2013) along with two-photon calcium imaging (Svoboda and Yasuda, 2006) has provided researchers with the ability to track hundreds to thousands of neurons over many days. These new technologies have provided incredible insight to how populations of M1 neurons change throughout skill acquisition (Huber et al., 2012; Masamizu et al., 2014; Peters et al., 2014), further highlighting the plasticity of M1 throughout learning.

Movement execution in M1

Despite a significant body of work aimed at understanding of the different schemes used by M1 to acquire and encode movements, the necessity of M1 in generating movements remains unclear. For example, early experiments demonstrated that lesions of M1 severely impaired the ability of rats to retrieve food pellets (Castro, 1972). Alternatively, recent work has demonstrated that M1 lesions did not affect the ability of rats to perform a serial lever-pressing task. How can we reconcile these seemingly contradictory results?

One explanation may be conjured by looking at these results with respect to the relative difficulties of the movements performed. For example, it has been well established through M1 lesions in both primates as well as rodents has been demonstrated to cause deficits in generating dexterous movements (Alaverdashvili and Whishaw, 2008; Passingham et al., 1983; Travis and Woolsey, 1956). Based on these results, the prevailing view has thus become one which divides movements into two categories: dexterous and non-dexterous movements. As noted, dexterous movements, such as food pellet retrieval which requires a high degree of digit control, require M1 for

execution. Alternatively non-dexterous movements, such as an unrestrained lever-pressing task, do not require M1 for execution (Kawai et al., 2015).

However beyond the complexity of movements, the state of learning of a movement presents an additional variable that may affect a movement's dependency on M1 for execution. For example, in addition to the expansion of movement representations in M1 during motor learning (Kleim et al., 1998; Nudo et al., 1996), additional work has demonstrated that following asymptotic performance of a new movement that these representations then contract (Molina-Luna et al., 2008; Pascual-Leone et al., 1994). This result is interesting considering the previously mentioned work noting that M1 is crucial for learning, and undergoes significant plasticity and reorganization during the acquisition of new movements and skills (Kawai et al., 2015; Peters et al., 2014). If the contraction of representations following asymptotic performance is a continuation of learning, are these later changes in M1 the result of M1 becoming less necessary for the execution of the movement? Chapter 2 of this thesis will investigate this question using a novel two-direction forelimb based movement task combined with in vivo two-photon calcium imaging and optogenetic perturbation.

The posterior parietal cortex in action selection

Beyond the many areas involved in movement execution, a number of other motor and non-motor related brain areas contribute to the selection of actions to be executed. These areas include the parietal, prefrontal and premotor areas as well as subcortical structures (Constantinidis and Procyk, 2004). Among these areas, the posterior parietal cortex (PPC) is anatomically located between sensory and motor areas and is thus well situated to integrate the external environment with the goals of the animal. Further, its extensive connectivity with areas known to be involved with working memory, such as the prefrontal cortex, suggests a potential role in working memory.

Indeed lesion studies of PPC have confirmed working memory deficits (Berryhill and Olson, 2008).

One manifestation of working memory comes in the form of choice and outcome information from completed actions (Seo et al., 2009a). For example, when the rules of a task are unknown, the history of decisions and outcomes contribute towards the selection of future actions. Such history-dependent action selection has been implicated in PPC among other cortical and subcortical areas (Seo et al., 2009a). However the way in which the neural representations of these signals evolve during action selection remains unknown. Further, traditional approaches to assess causal roles of such neural representations, such as pharmacological manipulation or lesion studies, preclude the ability to understand the temporal dynamics surrounding these representations and thus their role in subsequent action selection. Chapter 3 of this thesis utilizes a novel visually-guided action selection task with behavioral modeling, in vivo two-photon calcium imaging, and optogenetic perturbation to investigate the neural substrates of history signals in the PPC and assess their causal role in action selection.

Chapter 2. Motor cortex-independent movement execution induced by long-term learning

2.1 Summary

The primary motor cortex (M1) is important for the execution of some but not all movements. To investigate the variable role of M1 in movement execution, we developed a forelimb-based two-direction movement task performed by head-fixed mice. Mice were trained for many weeks before achieving high levels of performance. At the expert stage, each mouse exhibited an idiosyncratic bias toward one of the two directions, and movements to the biased direction were more stereotyped than movements to the non-biased direction. Two-photon calcium imaging revealed that M1 neurons were more active and their activity pattern was more consistent during the non-biased, less stereotyped movements than during the biased, stereotyped movements, suggesting a preferential contribution of M1 to less stereotyped movements. Supporting this model, optogenetic inactivation of M1 substantially amplified the existing bias (i.e., movements to the less stereotyped direction became rarer). In contrast to expert mice, M1 inactivation early in learning severely impaired the ability to initiate and execute movements in both directions. These results reveal that the involvement of M1 in movement execution changes throughout learning, and that well-practiced movements can become independent of M1 with learning.

2.2 Introduction

A fundamental function of the brain is to generate movements, allowing for interaction with others and the environment. Movement execution is controlled by the interplay among many interconnected brain regions. Among these areas, a central locus

in movement control is the primary motor cortex (M1). M1 is an evolutionarily conserved structure (Kaas, 2008) with direct projections to spinal cord motor circuits (Asanuma and Sakata, 1967b), and low current electrical stimulation of M1 can drive body movements (Fritsch and Hitzig, 1870; Graziano et al., 2002; Neafsey and Sievert, 1982; Tennant et al., 2011). Single-unit recordings in monkeys well-trained with forelimb-based reaching tasks have revealed that M1 activity is reliably related to various features of the executed movements at the levels of both single neurons (Georgopoulos et al., 1982b; Scott, 2008) and neuronal populations (Churchland et al., 2012; Georgopoulos et al., 1986; Shenoy et al., 2013). These results have painted the view that M1 is a critical region for movement execution.

However, M1 is not necessary for execution of all movements, as even complete removal of the cortex does not lead to full paralysis (Travis and Woolsey, 1956). While there are likely differences in M1 functions across species (Lemon and Griffiths, 2005), M1 lesions in general seem to affect dexterous movements more strongly, leaving non-dexterous and innate movements relatively intact (Castro, 1972; Kawai et al., 2015). These results suggest that M1 is preferentially involved in executing dexterous movements. Importantly, however, it remains unclear whether the movements that require M1 at one point in time always remain under the control of M1. In support of the idea that M1's role in movement execution is dynamic, M1 circuits are profoundly flexible, exhibiting many forms of changes across multiple spatial scales (Peters et al., 2017). For example, motor learning can change the somatotopic map revealed by microstimulation (Kleim et al., 1998; Nudo et al., 1996) as well as induce synaptic changes within M1 (Chen et al., 2015; Cichon and Gan, 2015; Rioult-Pedotti et al., 2000; Xu et al., 2009). Single-unit activity from single-session electrophysiological recordings (Costa et al., 2004; Gandolfo et al., 2000; Paz et al., 2003; Rokni et al., 2007), and more

recently longitudinal population imaging (Huber et al., 2012; Masamizu et al., 2014; Peters et al., 2014), has established that M1 ensemble activity changes during motor skill learning. These results preclude the possibility that motor skill learning is achieved simply by learning to activate the M1 neurons that relate stably to the desired movement. Rather, it has become clear that the relationship between M1 activity and movements itself is shaped with learning. Given the dynamic nature of the relationship between M1 activity and movement, an intriguing question is whether the degree of M1 dependence of movement execution may also change throughout skill acquisition. In fact, years of motor training reduce metabolic activity in M1 associated with the movement (Picard et al., 2013).

To explore whether the movements requiring M1 early in learning can become M1 independent, we trained mice to perform novel forelimb movements in two directions over months. At the expert stage, mice exhibited a bias for one movement over the other. The biased movement was more stereotyped compared to the non-biased movement. Surprisingly, M1 activity as measured by two-photon calcium imaging was stronger and more consistent during non-biased, less stereotyped movements. Furthermore, inactivating M1 amplified the bias for the stereotyped movement. These results suggest that the biased, more stereotyped movements are less dependent on M1. In contrast to the expert stage, M1 inactivation at the beginning stage of learning impaired movements in both directions, indicating that both movements initially require M1. Our results demonstrate that long-term training can make originally M1-dependent movements become less dependent on M1, and suggest a role for M1 in supporting the execution of newly acquired, less skilled movements.

2.3 Results

A novel two-direction forelimb movement task

To investigate the involvement of motor cortex in movement execution, we developed a forelimb-based two-direction motor task in which head-fixed mice manipulated a joystick with their left paw (Figure 2.1A-C). In each trial, one of two visual stimuli (full-field, full-contrast, vertical or horizontal gratings moving forward or down, respectively) was presented. Two seconds after stimulus onset, an auditory 'go' cue designated the response period. During the response period, mice were required to move the joystick beyond a set threshold (~10 mm) to one of two invisible targets in the direction corresponding to the visual stimulus (Figure 2.1B) to receive a water reward (Figure 2.1C). The visual stimulus remained on during the response period, thus there was no requirement for working memory. Unsuccessful trials (no movements, movements initiated prior to the go cue, movements that did not reach a target, and movements to the wrong target) were not rewarded.

We developed a training procedure with multiple blocks that incrementally shaped the structure of the task (Figure 2.1D-F, Experimental Procedures). On average mice required 49 ± 14 sessions (mean $\pm$ S.D.) of training (one session per day) before reaching expertise (arbitrarily defined as >55 % of correct withholding until the go cue and >70 % of correct discrimination in withheld trials). Once mice reached the expert level of performance, they continued on to either imaging or inactivation experiments (described below). Behavioral performance during imaging sessions and control trials of inactivation sessions remained at expert levels: completing 265 ± 55 trials per session, answering 98.1 ± 2.7 % trials, reliably withholding movements until the go cue for 81.3 ± 9.6 % of trials, and correctly discriminating 73.4 ± 7.5 % of trials (Figure 2.1G-J).

Expert mice exhibit bias for more stereotyped movements

Similar to previously described discrimination tasks (Busse et al., 2009; Guo et al., 2014a), mice biased their responses, answering with one movement more frequently than the other. We quantified the bias as:

$$Bias = probability(Move_{Forward} | Stim_{Forward}) - probability(Move_{Down} | Stim_{Down})$$

With this quantification, a bias of +1 means that mice always respond with a forward movement (regardless of visual stimuli presented) and a bias of -1 means that mice always respond with a downward movement. Many expert mice showed some bias (Figure 2.2A). The biases were evenly distributed to both directions across mice (bias across all mice: 0.0015 ± 0.1546 , $p = 0.96$, Student's t-test, Figure 2.2A), with 8 mice showing significant (outside 95% C.I., binomial test) bias for forward and 6 mice showing significant bias for down, indicating that bias was not the result of common mechanical constraints of the behavioral setup, or common innate differences in two movements.

Based on the significant bias of individual mice, we considered the possibility that the movements in the biased and non-biased directions were qualitatively different. To examine this possibility, we examined the trajectories of joystick movement in trials with movements to the two targets. An example of this can be seen in Figure 2.2B, which shows 10 randomly selected movements to each of the forward and down directions from 5 consecutive sessions for a mouse that shows a bias for the down movement. For this mouse, the movements to the down target were highly stereotyped, showing a high degree of overlap across trials (purple trajectories, Figure 2.2B). To quantify movement stereotypy, we calculated the variability of the endpoint angle and mean squared error from the average trajectory (M.S.E., see Experimental Procedures). Variability of endpoint angle represents the stereotypy of the direction of movements, and M.S.E. represents the similarity of movement trajectories compared to the average movement (to each direction). We observed that both the variability of movement endpoint angles

and mean M.S.E. were significantly smaller for biased movements compared to non-biased movements (variability of endpoint angle: $p < 0.05$; M.S.E.: $p < 0.05$, Figure 2.7A-B), and the relative difference in movement stereotypy was stable across sessions in expert mice (Figures 2C and 7C). Furthermore, the degree of bias significantly correlated with the normalized difference in motor stereotypy between biased and non-biased targets (Experimental Procedures), such that a larger difference in motor stereotypy coincided with a stronger bias for the stereotyped direction (Figure 2.2D&E). These results indicate that the movements in the biased direction of each mouse were more stereotyped than the movements in the non-biased direction.

M1 activity is stronger and more consistent for less stereotyped movements

Given the observed differences in stereotypy between biased and non-biased movements, we hypothesized that M1 would exhibit differential activity characteristics for biased versus non-biased movements. To explore this possibility, we recorded activity of M1 neurons using two-photon calcium imaging. Once mice reached behavioral expertise, we injected AAV-syn-GCaMP6f in the forelimb area of the right motor cortex. Two to five weeks after the AAV injection, we started imaging the activity of neurons in layer 2/3 expressing GCaMP6f during task performance in 29 sessions across 8 mice. Of the imaged neurons, 1,216 (mean: 46 / session, range: 19-98) showed significant movement-related activity and were included in analysis (Experimental Procedures). Movement-related neurons were split into two groups; 841 (69.2%, e.g. Cells 1-3, Figure 2.3A) showed significantly different activity in the two movements ('movement-tuned cells'), while 375 (30.8%, e.g. Cell 4, Figure 2.3A) showed indiscriminate activity for both movements.

Examining the movement-tuned cells, we observed that in most imaging sessions the fraction of cells showing stronger activity for the non-biased movement was

higher than the fraction of cells showing stronger activity for the biased movement ($p < 0.01$, Student's t-test, Figure 2.3B-C). Accordingly, the magnitude of bias significantly correlated with the abundance of cells that were more active during the non-biased movement ($R = 0.55$, $p < 0.01$; Figure 2.3D). Additionally, we observed significant correlations when comparing the fraction of cells more active for the non-biased movement to the relative difference in stereotypy between the two movements (Figure 2.3E-F). In other words, the less stereotyped one direction was, the more often cells were tuned to that direction. Thus, we conclude that M1 is more active during the non-biased, less stereotyped movement than during the biased, more stereotyped movement.

In addition to the difference in the amount of activity, we investigated whether consistency of M1 activity differs between biased and non-biased movements. To quantify the consistency of population activity across trials, we measured the correlation of population activity on a trial-by-trial basis. The trial-by-trial correlation was compared between the populations of neurons tuned to biased and non-biased movements during biased and non-biased movements respectively. We observed that the correlation of population activity of neurons tuned to non-biased movements was significantly higher than neurons tuned to biased movements ($p < 0.001$, Figure 2.3G, Student's t-test). This result indicates that M1 activity was more consistent during the non-biased movements, which is surprising given that non-biased movements are more variable than the biased movements (Figure 2.2D and 2E). This prompted us to investigate the relationship between M1 activity and movement trajectories on a trial-by-trial basis. Plotting the trial-by-trial correlation of population activity against the trial-by-trial correlation of movement trajectories, we found that the slope was significantly steeper for non-biased movements than biased movements ($p < 0.05$, for both ordinary and Deming linear regressions,

bootstrap, Figure 2.3H). Together these analyses indicate that M1 population activity is more consistent and more closely related to movements during the non-biased movement compared to the biased movement, suggesting a more important role of M1 in controlling the execution of non-biased, less stereotyped movements.

Inactivation of M1 does not inhibit movement initiation in expert mice

Given the differences in M1 activity during biased and non-biased movements, we next investigated the necessity of M1 in performing these movements by acutely inactivating M1 in expert animals while performing the task. Inactivation of M1 was accomplished by optogenetically activating parvalbumin (PV)-positive inhibitory neurons (Lien and Scanziani, 2013; Zhao et al., 2011). We expressed Channelrhodopsin-2 (ChR2) selectively in PV inhibitory neurons by injecting a Cre-dependent AAV encoding ChR2 in PV-Cre transgenic mice, or by crossing PV-Cre mice with Ai32 mice (containing lox-stop-lox-ChR2 in the ROSA locus, (Madisen et al., 2012)). Extracellular recordings in primary visual cortex showed that blue light activation of ChR2-infected PV inhibitory neurons acutely reduced firing rates of putative excitatory neurons by $93.4 \pm 5.8 \%$ (Figure 2.8). Using this technique, we bilaterally inactivated M1 during the visual stimulus and response periods in a subset (~15 %) of randomly interleaved trials (n = 14 mice, Figure 2.4A). To control for potential non-specific behavioral effects of the blue light, we also performed control sessions in which the blue light was pointed away from the brain (Experimental Procedures). We did not observe non-specific behavioral effects of the blue light in control sessions (Figure 2.9).

When we compared the fraction of trials with no movements between control (light-off) and inactivation (light-on) trials, we saw little effect of M1 inactivation. Only in 2 animals was there an increase in no-movement trials, and the effect was not statistically significant across mice (Figure 2.4B). Furthermore, we saw no significant effect of M1

inactivation on the reaction time (Figure 2.4C), nor did M1 inactivation impair the ability of mice to reach a target once movements were initiated (Figure 2.4D). Therefore, M1 inactivation does not significantly affect the ability of the expert mice to initiate movements and reach a target. Interestingly, there was a significant decrease in the ability of mice to withhold movements until the go cue ($p < 0.001$, Figure 2.4E), consistent with previous work suggesting the role of M1 in withholding movements (Stoltz et al., 1999; Zagha et al., 2015).

M1 inactivation in expert mice amplifies the bias for the more stereotyped movements

Despite limited effects on the initiation of movements, analysis of behavioral performance revealed that M1 inactivation resulted in a significant decrease in the discrimination performance (fraction of trials with correct target out of trials in which the mouse reached a target, $p < 0.001$, Figure 2.4F). The decrease in performance was not random, but was accompanied by an increase in bias ($p < 0.001$, Figure 2.4G).

The bias during M1 inactivation was related to the bias exhibited in control trials. For example, Figure 2.5A shows 75 trials during one inactivation session (session 87 from Figure 2.2B) of a mouse that has a bias for down response. Notice that during most inactivation trials the mouse responded with a downward movement regardless of which stimulus was presented. Across inactivation sessions of this mouse, M1 inactivation enhanced the bias exhibited during control trials (Figure 2.5B). Across mice, we observed a significantly positive correlation between the bias during control trials and inactivation trials ($R = 0.29$, $p < 0.05$; Figure 2.5C), with the slope significantly larger than one (slope > 1 defined by Deming regression, $p < 0.05$, bootstrap). In agreement with this finding, the bias exhibited during M1 inactivation significantly correlated with the relative stereotypy of the two directions in control trials, such that M1 inactivation

increased the trials with movements in the more stereotyped direction (Figure 2.5D-E). Together these results indicate that the bias induced through inactivation of M1 is an enhancement of the bias in control trials. With M1 inactivation, expert mice strongly prefer to answer with the more stereotyped movement.

M1 inactivation early in learning impairs movements

The results so far indicate that M1 is not essential for the execution of the biased, more stereotyped movement in the expert mice. These results are surprising given the wide-spread belief that M1 is essential for movement control described by our lab and others (Guo et al., 2015; Peters et al., 2014). A potential explanation for this discrepancy surrounds the fact that the mice in our experiments have been over-trained to perform these movements prior to inactivation and imaging experiments. In our study, imaging and inactivation experiments began after 49 ± 14 sessions of training, substantially more than the training animals underwent in the previously mentioned studies. Therefore, we postulated that overtraining may have led the biased movements to no longer require M1 for execution. To investigate this hypothesis, we tested whether M1 is required for the execution of the movements early in learning, by inactivating M1 in a separate cohort of mice ($n = 9$) during the first block of training.

Inactivation was performed on 3-5 sessions within the first 15 days of behavioral training (first inactivation session number 3.6 ± 0.5 , mean $\pm$ S.D., $n = 35$ sessions). In these early sessions, there was no significant correlation between bias and movement stereotypy as was described in expert animals (Figure 2.10A-B). In contrast to expert inactivation, inactivation during early-learning caused a significant impairment in movement initiation, exemplified by an increase in the fraction of no-movement trials ($p < 0.05$, Figure 2.6A). In trials where mice managed to initiate a movement, early inactivation significantly increased the reaction time ($p < 0.01$, Figure 2.6B) as well as

the fraction of trials in which movements failed to reach a target ($p < 0.05$, Figure 2.6C). Additionally, we did not see the same enhancement of bias observed during M1 inactivation in experts (Figure 2.10C). The observed behavioral phenotypes were not due to non-specific effects of the blue light (Figure 2.11).

In addition to these deficits in movement initiation and execution, we noticed by visual inspection that M1 inactivation in early learning often caused overt deficits in limb control. While mice in unperturbed trials grabbed the joystick with their paw and seldom removed it from the joystick, mice in inactivation trials often appeared unable to grip the joystick or maintain control of their forelimb (Supplemental Movies 1 and 2). During these trials in which animals took their hands off the joystick, the joystick position was no longer a faithful representation of the forelimb movements of the mouse. To quantify this effect, we recorded video in a subset of the inactivation sessions and manually scored the periods when the paw was off the joystick. Scorers were blind to the identity of the mouse, which experimental group mice belonged to (early control, early inactivation, expert inactivation), as well as the condition of individual trials (i.e. frames were masked so scorers could not tell which trials were light-on trials). Each mouse performed 3 sessions with video recording, and each movie was scored by 3 independent scorers (scores were generally consistent across scorers therefore we report the mean of the 3 scores for each session). This analysis demonstrated that M1 inactivation during early learning resulted in the release of joystick in a substantial fraction of trials ($p < 0.05$, Figure 2.6D). Again, this effect was a result of inactivation and not due to non-specific effects of light (Figure 2.11D). In contrast to early inactivation, M1 inactivation in expert mice did not cause this effect (Figure 2.6E), consistent with the results presented earlier. Taken together, in nearly half of trials during early learning, M1 inactivation caused movement impairment (no movements, failed movements, or joystick release) ($p < 0.05$,

Figure 2.6F). Movement impairment was observed with stimulus in both biased and non-biased directions ($p < 0.05$ in both directions, Figure 2.6F). This effect is in contrast to experts who did not demonstrate a significant effect of M1 inactivation on movement (Figure 2.6G). These results demonstrate a crucial role of M1 in the initiation and controlled execution of less skilled movements during learning. However, with long-term training, M1 can become dispensable especially for highly stereotyped movements.

2.4 Discussion

Asymmetric expertise of two learned movements

Here we investigated the role of the primary motor cortex in movement control, by developing a novel motor task requiring mice to make two distinct movements. Reaching expert levels of performance required many weeks of daily training. However, even at expert levels, mice showed asymmetric levels of expertise reflected in differing degrees of stereotypy between the two directions. The direction of more stereotyped movement was variable across animals. We currently do not know what underlies the idiosyncratic choice of stereotyped movement direction. Regardless of the origin, our behavioral task provided an opportunity to probe the involvement of M1 in executing movements with different levels of skillfulness. We found that in expert mice, at least one of the two directions of movements did not require M1 for execution, and the direction that was spared in M1 inactivation was the direction in which the movement was more stereotyped.

Neurons in M1 showed stronger and more consistent activity during the less stereotyped movement that required M1 for execution. Conversely, during the more stereotyped movements, M1 showed less activity that was more variable, and appeared to be dispensable with respect to movement execution. The variable activity pattern and

lack of M1 dependency could be explained if the more stereotyped movement direction somehow bypassed M1 for execution throughout the entire duration of training.

However, inactivation experiments early in learning, which impair movements in both directions, refute this possibility. Rather, these results suggest that as the mice became more skilled for a particular movement, M1 gradually became less necessary for its execution. Our lab previously showed that reproducible population activity patterns emerged in M1 during the first two weeks of motor skill learning (Peters et al., 2014). Our current results suggest that further overtraining can make the movement more independent of M1, which is accompanied by a reduction in the reproducibility of population activity.

Loss-of-function methodologies

We studied the involvement of M1 in movement execution by suppressing the function of M1. Such loss-of-function experiments fall in two categories: chronic (lesion) and acute (e.g. optogenetics, pharmacology and pharmacogenetics) manipulations. Acute manipulations can result in stronger deficits compared to chronic lesion (Otchy et al., 2015), as the circuits can adapt to the lesion over time by compensatory plasticity (Castro-Alamancos and Borrell, 1995). This is akin to developmental compensation that can result in a lack of phenotype in certain gene knock-out experiments, while acute knock-out can still reveal that the gene is necessary for the process under investigation. In this study, we opted to use optogenetic inactivation of M1 on a trial-by-trial basis, which should circumvent compensation as much as possible. Behavioral effects from acute manipulations are evidence for the involvement of the area. It is important to remember, however, that such results from acute manipulations do not distinguish the type of involvement; the area could provide instructive information for the behavior or only play a permissive role for other instructive areas to function (Otchy et al., 2015).

While our current study does not distinguish these possibilities, our main finding is significant and striking: a given movement that initially depended on M1 for execution became independent of M1 as a function of learning.

Involvement of M1 in different types of movements

It is well established that not all types of movements require M1. This is clearly shown by experiments with M1 lesion and transection of corticospinal tract, which result in permanent loss of some but not all movements (Castro, 1972; Travis and Woolsey, 1956). A key feature of movements that affects M1 dependence appears to be the level of dexterity (Alaverdashvili and Whishaw, 2008; Castro, 1972; Guo et al., 2015). Here, dexterous movements are broadly defined as those that involve fine coordination of distal muscles, especially those that are not species-typical. Non-dexterous movements, including species-typical movements such as locomotion and grooming, tend to be more preserved following lesion to cortical movement-control pathways (Berridge and Whishaw, 1992). These distinctions may arise from hard-wired circuits that have evolved in subcortical structures for species-typical and innate movements (Berridge and Whishaw, 1992). Within this context, our task falls in the dexterous category. In our task, mice needed to grab the joystick handle with one arm and move it in two directions while under head-fixation. Furthermore, the joystick handle is made of a thin brass rod, the grasping of which required similar levels of dexterity as manipulating a piece of dry pasta that has been used as a paradigm for dexterous motor skill learning (Allred et al., 2008; Whishaw and Coles, 1996). Because of these features, our task requires a set of motor coordination that is unlikely to occur within the normal repertoire of mouse behaviors. Consistent with the idea that our task requires dexterous movements, M1 inactivation impaired the execution and control of movements in our task during early training.

A recent and influential study proposed a model regarding the role of M1 with a series of well-controlled experiments (Kawai et al., 2015). In this study, rats were trained to learn to sequence together a set of non-dexterous movements. While rats with M1 lesion initially performed similarly to control animals, they improved much more slowly than their non-lesion counterparts. However, when M1 was lesioned in expert rats, they were able to retain the learned sequence after a post-lesion recovery period of 10 days. Based on these observations, the authors proposed that M1 is required for learning to temporally assemble pre-existing, non-dexterous movement components, functioning as a tutor for subcortical structures to learn the sequence. Apart from this role in learning, M1 appears to be dispensable for the execution of non-dexterous movements.

Our results here are distinct from – but nonetheless compatible with – this model. Unlike the movements studied by Kawai et al., the movements in our task involve dexterous control of the forelimb and critically depend on M1 for execution early in learning. However, with many weeks of training, these movements became less dependent on M1, suggesting that the role in execution of these movements became offloaded from M1, likely to subcortical structures. The subcortical structures that control movement execution in expert animals may include the brainstem, cerebellum and basal ganglia, all of which contain pathways to control motor circuits in the spinal cord (Hikosaka et al., 2002; Yin and Knowlton, 2006). Our results do not exclude the possibility that, in addition to its role in movement execution early in learning, M1 also plays an instructive role in this offloading process by acting as a tutor for the subcortical structures, as proposed for non-dexterous movement sequence learning (Kawai et al., 2015). In fact, it has been proposed that behaviorally-relevant circuits could have an additional function to entrain shortcut circuits through Hebbian plasticity (Makino et al., 2016).

It is widely believed that M1 is essential for the execution of dexterous movements. Our study extends this notion and shows that the degree of learning is another major factor affecting the involvement of M1 in executing particular movements. Training over weeks can cause a movement that is initially M1-dependent to become M1-independent. These results reinforce the idea that the same movements can be generated by different neurons and brain areas at different stages of learning. While the potential advantage of having multiple movement control systems for different stages of learning is unclear, an attractive hypothesis is that each system has a unique level of stability and flexibility. M1, with its high levels of plasticity, may be best suited to encode newly acquired skills which may maintain the need to be modified. When the skill is overlearned and the need for further modification is reduced, it may be beneficial to offload the movement to a more stable system, which may both allow stable storage and also make M1 available to learn other, novel movements. This sequential learning process is reminiscent of other systems of learning. For example, episodic memory is believed to be initially stored in hippocampus and later transferred to cortex (Frankland and Bontempi, 2005). Furthermore, reward-guided operant learning initially depends on dorsomedial striatum but later the dependence shifts to dorsolateral striatum as the behavior becomes habitual (Thorn and Graybiel, 2014; Yin et al., 2009). In both contexts, the dynamic areas initially important for learning may offload the information to other areas for stable long-term storage. Our results underscore the dynamic and fluid nature of learning circuits.

2.5 Methods

Animals

All procedures were in accordance with protocols approved by the UCSD Institutional Animal Care and Use Committee and the guidelines of the National Institute of Health. Mice (calcium imaging: cross between *Gad2-IRES-Cre* [JAX 010802] and *Rosa26-CAG-LSL-tdTomato* [JAD007914]; optogenetic inactivation: *PV-Cre* [JAX 008069], or cross between *PV-Cre* and Ai32 [JAX 024109]) were housed in a room with a reversed light cycle (12 h-12 h). Experiments were performed during the dark period.

Behavioral apparatus

The behavioral apparatus was housed in a soundproof box (40 × 40 × 40 cm) and training sessions were performed in the dark. The components of the task included a joystick (M11L061P; CHProducts), a 17-inch computer monitor (for visual stimulus presentation, placed ~15 cm from the right eye of the mouse; E1713Sb; Dell), and a water port (with photodiodes to sense licking). The joystick handle was custom machined and fitted with a 1/16 inch thick brass rod that mice manipulated with their left forepaw (Figure 2.1A). An electromagnet (EM050-3-222; APW) mechanically immobilized the joystick at the origin during inter trial intervals and visual stimulus period in block 1 training (Figure 2.1D), but only during inter trial intervals in block 2 and 3 (Figure 2.1E and F). The joystick had a dynamic range of 56° in each angular direction. The 2D angular position of the joystick was continuously recorded at 1 kHz using a data acquisition card (USB6008; National Instruments) and custom Matlab software. The task-sequence execution, stimulus selection, auditory cue presentation and reward dispensation were coordinated (and recorded) by an open source real-time Linux/Matlab software package BControl (<http://brodywiki.princeton.edu/bcontrol/>). The presentation of visual stimuli (100% contrast, full-field, square wave drifting gratings 0.04 cycles/degree and 3 cycles/sec) was implemented using Psychtoolbox (an open source Matlab toolbox; <http://psychtoolbox.org/>).

Behavioral training

Adult mice (six weeks or older, male and female) were implanted with a custom head-fixation plate that was attached to the skull. Following a minimum 3 days of recovery, daily water consumption was limited to a controlled volume (typically 1 mL/day). After 3-10 days of water restriction mice began behavioral training.

In the two-direction forelimb-based movement task (Figure 2.1A), one of two orthogonal visual stimuli (forward or down drifting gratings, see *Visual stimulus presentation* below) was presented on the computer monitor. Two seconds after visual stimulus onset, an auditory cue (6 kHz pure tone) marked the beginning of the response period (up to 10 sec) during which the joystick entering the correct target (hereafter referred to simply as ‘target’; Figure 2.1B) in the same direction as the gratings triggered a water reward. The visual stimulus and auditory cue remained on throughout the response period (Figure 2.1C). Errors (i.e. entering the incorrect target and movements before the ‘go’ cue) triggered a white noise sound and led to an immediate trial termination. Following reward, trial-termination, or no response, the return of the joystick to the origin ended the trial and initiated an inter-trial interval (ITI, 4 sec). During the ITI the joystick was immobilized at the origin by the electromagnet. At the end of the ITI (simultaneous with the beginning of the visual stimulus onset), the electromagnet was disengaged and the joystick was free to move.

Mice were trained under head-fixation in the behavioral setup approximately 1 hour per day over a period of 1-3 months. The task was shaped to reach the final version described above through 3 training blocks (Figure 2.1D-F). In the first training block, mice received a water reward as long as they moved the joystick to the correct target within a 30 second response period (regardless of whether or not they hit the incorrect target first). The nature of the visual stimulus presentation (see *Visual stimulus*

presentation below) required mice to move the joystick in both directions during the first training block. As mice become proficient at moving the joystick in both directions, we increased the target distance from 6.7° (~6 mm) to 11.1° (~10 mm), and we decreased the response time to 10 seconds until they reached the targets during the response period in more than 80% of trials (Figure 2.1D).

In block 2 mice were trained to refrain from pushing or leaning on the joystick prior to the go cue. During block 2, the joystick was released from electromagnet immobilization concurrent with the onset of the visual stimulus, and mice were rewarded only if they initiated joystick movements during the response period (i.e. successfully withheld movements until the go cue). Trials in which mice responded before the go cue were considered errors and immediately terminated, resulting in a white noise error sound. Similar to block 1, mice were rewarded if they reached the correct target, regardless of whether or not they hit the incorrect target first. Block 2 training continued until mice achieved withholding performance above 55% (arbitrarily defined). Withholding performance continued to improve during block 3 training, reaching $81.3 \pm 9.6\%$ withholding during imaging and control trials of optogenetic experiments (Figure 2.1I).

In block 3, in addition to withholding movement until the go cue, mice were trained to move the joystick only to the correct target. In this training block, trials were considered errors and immediately terminated if mice reached the incorrect target, or moved before the go cue (as in block 2). Training continued in block 3 until discrimination performance reached 70% (Figure 2.1F). Discrimination performance was computed for all trials that reached a target regardless of whether or not trials were successfully withheld. Once this performance criterion was achieved, the ITI length was increased up to 4 seconds (from 1 second). After achieving expertise during training

(>55 % of correct withholding until the go cue and >70 % of correct discrimination) mice underwent either a second surgery to prepare for either imaging experiments, or began optogenetic inactivation of M1.

Visual stimulus presentation

Early versions of the task employed random presentation of the visual stimulus. In these versions, many mice developed a bias for one target that was so strong they were unable to learn the task (data not shown). Therefore, we modified the stimulus selection statistics to encourage mice to make movements to both directions. In blocks 1 & 2 of training, visual stimulus alternated between forward and down with the added rule that after error trials the same stimulus was repeated. In block 3 the visual stimulus was randomly selected with following rules: 1) after three consecutive rewarded trials in one direction, the stimulus switched to the other direction, and 2) after error trials the same stimulus was repeated (similar to training blocks 1 & 2).

Imaging neural activity

After achieving expertise in the behavioral task, we temporarily suspended training and allowed mice unlimited access to water for at least 2 days prior to craniotomy and virus injections as previously described (Peters et al., 2014). The craniotomy spanned the caudal forelimb area of primary motor cortex (M1; stereotaxic coordinates relative to bregma: 1.5 mm lateral, 0.3 mm anterior) in the right hemisphere. Virus (AAV2-1-hSyn-GCaMP6f diluted in saline 1:8; UPenn Vector Core) was injected at 5 sites (~20 nL per site) in M1 at a depth of approximately 250 μ m beneath the dura in layer 2/3. After injections, the craniotomy was covered with an optical window fixed in place with dental cement.

Following surgery and recovery, mice resumed training until performance recovered to pre-surgery levels and we were able to visualize sufficient GCaMP6f

expression (2-5 weeks after surgery). We imaged cortical activity in layer 2/3 at the depth of $\sim 200\ \mu\text{m}$ with excitation at 925 nm from a Ti-Sa laser (Spectra-physics) using a two-photon resonance scanning microscope (B-scope; Thorlabs). Each imaging field was 512×512 pixels covering $472 \times 508\ \mu\text{m}$ and imaging was performed at ~ 28.4 Hz. The duration of each behavior-imaging session was limited by when mice were disengaged from the task, completed 170 rewarded trials, or reach 1 hour (whichever came first). Mice completed $\sim 162 \pm 27$ rewarded trials in each behavior-imaging session. We imaged on average 3.2 ± 2.8 fields across 8 mice.

Neural activity analysis

Single cell activity: Using custom Matlab software, fluorescence images were aligned frame by frame to compensate for lateral motions. Regions of interest (ROIs) corresponding to individual neurons were manually drawn on the motion-corrected fluorescence images by circumscribing the cell bodies based on their GCaMP6f fluorescence intensity distinguishable from the background (mean: 160 / session, range: 77 - 289). Pixels inside each ROI were considered as a single cell, whereas pixels extending radially outward from the cell boundary by 2-6 pixels were considered background. In the case the background included other cells' ROIs, those pixels were excluded. To estimate the activity of a single cell, 70% of the average pixel intensity in its background was subtracted from the average pixel intensity inside the cell. The time series of the background-adjusted intensity was transformed to dF/F by dynamically estimating the baseline intensity (i.e. the 8th percentile of the intensity distribution in the 20 second window centered at each time point).

Movement-related neurons: We identified movement-related neurons that showed significant activity modulation during the peri-movement period (-0.5 s prior to movement onset to 1.0 s after movement onset). The mean activity trace of each neuron

was calculated by aligning dF/F traces of all rewarded trials to the movement onset and calculating the Z-scored activity. Neurons were selected as task related if they contained 3 consecutive frames with a Z-score ≥ 2.57 (corresponding to the 99th percentile of activity).

Movement-tuned neurons: To avoid potential confounds of differences of activity being attributed to differences in the duration of biased and non-biased movements, we analyzed movement tuning using successful trials with movement durations < 1 second. Neurons were considered movement-tuned if the mean peri-movement activity of correct trials was significantly greater ($p < 0.01$, Wilcoxon Rank-Sum test) for one target over another.

Trial-by-trial population activity correlation: The population activity correlation between two trials was the correlation coefficient between the two concatenated activity time series (-0.5 to 1 sec from movement onset) of a population of neurons. When computing the population activity correlation for neurons showing stronger activity for the biased (or non-biased) movements, we used trials in which movements were in the biased (or non-biased) direction.

Analysis of movement–activity correlation: To examine the relationship between movement and activity, we measured movement similarity and activity similarity for each pair of trials, and then performed a linear regression between the two similarity measures. The movement similarity was the correlation coefficient between the two joystick traces (the concatenated x and y position time series for 1 sec from movement onset). The activity similarity was the same as the population activity correlation described above. Because neither movement nor activity similarity is free of measurement noise, we applied Deming linear regression when estimating the slope relating the two measures. However, ordinary linear regression produced qualitatively

similar results when extreme values outside the range estimated by Deming regression were excluded.

Optogenetic silencing of M1 in expert mice

Mice used for optogenetic silencing of M1 (either *PV-Cre*, $n = 10$; or *PV-Cre* :: *Ai32*, $n = 7$) were implanted with a head-fixation bar. *PV-Cre* mice were bilaterally injected with a virus carrying Channelrhodopsin-2 (AAV2-1-EF1A-DIO-hChR2, UPenn Vector Core) through a thinned skull window over the forelimb region of the primary motor cortex (coordinates relative to bregma: ± 1.5 mm lateral, $+0.3$ mm anterior). Approximately 100 nL of virus was injected at one location at each of two depths 200 μ m and 600 μ m from the dura. *PV-Cre* :: *Ai32* mice received bilateral cranial windows over the forelimb region of the primary motor cortex. After surgery mice followed the same training protocol outlined above prior to completing M1 inactivation experiments.

Once mice reached expertise we inspected the previously thinned skull or cranial windows for clarity, re-thinning or replacing the cranial window if necessary. Then we acclimatized mice to trials with the blue light over 1-7 training sessions. During acclimation training sessions 470 nm blue light (transmitted through bifurcated LED fibers, 7.5mW per fiber; Doric) was shone ~ 1 mm above the head-fixation bar (4-5 mm away from the windows) on a randomly selected 15% of trials. All mice adapted to the blue light and achieved performance levels similar to those prior to introduction of blue light within several training sessions.

Each M1 inactivation experiment was performed across 10-12 daily sessions. Control sessions were alternated with M1 inactivation session-by-session, or in block format (e.g. 5 inactivation sessions followed by 5 control sessions). In control sessions, the LED fibers were pointed above the head-fixation bar (4-5 mm away from the windows), whereas during M1 inactivation sessions LED fibers were placed directly

above M1 on both hemispheres. Aside from LED fiber placement, all procedures were identical between control and M1 inactivation sessions.

For both control and M1 inactivation sessions, light-on trials were pseudo-randomly selected following the restriction that there be at least 5 light-off trials between any two adjacent light-on trials. This restriction was introduced to avoid potential behavioral adaptation to M1 inactivation due to consecutive and/or frequent exposures. Light stimulation was applied to 15% of trials following this rule.

Optogenetic silencing of M1 during early learning

For optogenetic inactivation of M1 during early learning, *PV-Cre* :: Ai32 mice (n = 9 for inactivation; n = 7 for control) were prepared as described above. Following recovery from surgery mice began block 1 of task training. After a brief period of task familiarization (3-7 days), optogenetic inactivation of M1 was performed over 3 sessions identical to expert mice (described above). A separate cohort of control mice underwent identical preparation (i.e. surgery, water restriction, initial task familiarization) however the LED fiber was placed away from the windows above the head bar (similar to control sessions described above) to control for non-specific behavioral effects of blue light.

Extracellular electrophysiology

Extracellular recordings were performed similar to those previously described (Liu et al., 2016). Adult mice (*PV-Cre* :: Ai32, n = 2), 6 weeks or older were anesthetized with urethane (1.2 g/kg, intraperitoneal) and given the sedative chlorprothixene (0.05 mL of 4 mg/mL, intramuscular) and implanted with a T-shaped head bar for head fixation. Body temperature was maintained at 37°C using a feedback-controlled heating pad (40-90-8D; FHC). A uniform layer of silicone oil was applied to the eyes to prevent drying. A craniotomy ~1 mm in diameter was made over the middle of V1 (~2.75 mm lateral to the midline, ~0 mm anterior to the lambda suture), and sterile saline was placed in the well

of the craniotomy to keep the brain moist. A 16-channel linear silicon probe (a1x16-5mm-25-177; NeuroNexus) mounted on a manipulator (Luigs & Neumann) was slowly advanced into the brain to a depth of $\sim 750\ \mu\text{m}$.

Recordings were started 20 min after insertion of the probe into V1. Signals were amplified 400-fold, band-pass filtered (0.3-5000 Hz, with the presence of a 60 Hz notch filter; A-M Systems 3600) and then digitized at 32 kHz (PCIe-6259; National Instruments) with custom Matlab software.

Visual stimulus was presented across three computer monitors (VX2450wm-LED, 60 Hz refresh rate, gamma-corrected; Viewsonic) mounted orthogonally to each other to form a square enclosure that covered $\sim 270^\circ$ of the visual field along the azimuth. The mouse head was immobilized at the center of the enclosure. Visual stimuli were generated using Psychtoolbox. The gratings drifted clockwise or counterclockwise in an oscillatory manner (amplitude $\pm 5^\circ$; grating spatial frequency 0.08 cpd; oscillation frequency 0.4 Hz; contrast 100%; mean luminance 40 cd/m^2). Trials were spaced by an inter-stimulation interval of at 8 seconds.

Optogenetic stimulation of V1 was accomplished by shining 470 nm blue light through an optical fiber pointed at V1. We recorded from V1 using three different blue light intensities: 3.5 mW, 7.0 mW, and 10.5 mW. Blue light intensities were varied in separate blocks of trials (i.e. 100 trials of 3.5 mW, followed by 100 trials of 7.0 mW). During optogenetic cortical inactivating trials, 10 seconds of blue light stimulation were applied in the middle of 12 seconds of visual stimulus. Trials of cortical inactivation (light-on) were interleaved with control trials (light-off).

Multi-unit analysis

Multiunit activity was isolated using spike-sorting software in Matlab as previously described (Liu et al., 2016). The raw extracellular signal was band-pass filtered between

0.5 and 10 kHz. Spiking events were detected with a threshold of 3.5 times the standard deviation of the filtered signal. Spike waveforms of four adjacent electrode sites were clustered using a *k-means* algorithm. Multi-unit spiking activity was defined as all spiking events exceeding the detection threshold after the removal of electrical noise or movement artifacts by the sorting algorithm. Individual spiking events were assigned to one of the 16 recording sites according to where they showed the largest amplitude.

Video recording and movement scoring

Movements were recorded for the duration of a behavioral session using a monochrome camera (DMK 23U618; The Imaging Source) at 30 Hz. Videos of mice performing the task were then stripped of identifiers (e.g. mouse name and date), randomly named, and masked such that the identity of the trial (e.g. light-on or light-off) was unknown by visual inspection. We then asked undergraduate volunteers to manually inspect the videos to mark frames when the left forelimb was removed from the joystick. Each video was scored by 3 independent scorers.

Movement analysis

Movement onset: defined as the first time at which the joystick velocity exceeded $22.2^\circ/\text{sec}$ ($\sim 20\text{ mm/sec}$) continuously for 20 msec and the joystick moved at least 1.3° ($\sim 1.1\text{ mm}$) from the origin. Reaction time: measured for all trials correctly withheld until the go cue as the time from the go cue to movement onset. Movement duration: measured for all rewarded trials as the time from movement onset to the time when the joystick entered any target region. Endpoint angle: The endpoint was calculated for all rewarded trials as the point when the derivative of the distance from the origin became negative. That is, we defined the endpoint of the movement as the point when the trajectory of the movement reached the maximum distance from the origin. We then took the 2 dimensional coordinates from the endpoint and measured the angle between the

movement and the target axis. The target axis was defined as the principle axis of the target (e.g. for the forward target the target axis was the x-axis, and for the down target the target axis was the y-axis). Mean squared errors (M.S.E.) of trajectories: The M.S.E. was calculated for all rewarded trials by taking the mean squared difference in joystick trajectory from a template movement. The template movement was defined for each direction as the average trajectory from movement onset to the median time to peak velocity. Relative differences in stereotypy metrics (variability of end point angle or M.S.E.) between forward and down movements (Figure 2.2C-E, Figure 2.3E-F, Figure 2.5D-E, Figure 2.7C, and Figure 2.10A-B) were calculated as:

$$\text{Difference in Stereotypy} = \frac{(Metric_{Down} - Metric_{Forward})}{(Metric_{Down} + Metric_{Forward})}$$

Statistical analysis

Unless otherwise stated conditions were paired within an animal (i.e. light-on vs light-off trials). Accordingly, when testing for a specific hypothesis that an experimental condition has a significant effect (e.g. M1 inactivation) we used a Wilcoxon Signed-Rank test.




Figure 2.1 Two-direction forelimb movement task.

(A) Task schematic.

(B) Stimulus-response rule.

(C) Task trial structure.

(D-F) Training block performance. Each grey line represents a single animal. Red circle indicates block 'graduation' date (N = 22 animals).

(G-J) Mean $\pm$ standard deviation of: number of trials per session (265 ± 55), fraction of trials answered ($98.1 \pm 2.7\%$), fraction of trials withheld until go-cue ($81.3 \pm 9.6\%$) and fraction of trials correctly discriminating ($73.4 \pm 7.5\%$) across all expert (imaging, inactivation and inactivation control) sessions (N = 174 sessions across 22 animals from **D-F**).




Figure 2.2 Expert mice exhibit bias for more stereotyped movements.

(A) Distribution of bias across all expert mice (22 animals in **Figure 2.1G-J**). Distribution was centered around 0 (0.0015 ± 0.1546 , $p = 0.96$, Student's t-test). Green and purple bars represent animals with significant bias for forward or down respectively ($p < 0.05$, binomial test).

(B) Ten randomly chosen movements to forward (green) and down (purple) targets in 5 consecutive sessions from animal 5 in C (red arrow).

(C) Normalized difference of movement stereotypy, measured as endpoint angle variability, between forward and down directions across the last 10 days of experiments for imaging and expert inactivation mice in A. Relative movement stereotypy is stable within each animal. Red arrow represents example mouse in B and **Figure 2.5A-B**.

(D-E) Bias correlates with the difference in movement stereotypy in each direction. Each dot represents one inactivation or imaging session (174 sessions across 22 mice).

Figure 2.3 M1 activity in expert mice is tuned for non-biased movements

(A) Calcium traces (mean $\pm$ standard deviation) to forward (green) and down (purple) movements aligned to the movement onset (grey vertical line) from 4 example neurons in the field of view in **B**.


(B) Spatial distribution of movement-related neurons in one imaging field. Masks represent individual movement-related neurons that were either tuned to the biased movement (forward, purple), tuned to the non-biased movement (down, green), or indiscriminately active (white). Note the relative abundance of neurons tuned to the non-biased movement.

(C) Distribution of fraction of movement-tuned neurons that are tuned to non-biased movements in all imaging sessions. Red solid line represents mean of all sessions (0.58 ± 0.14), significantly larger than 0.5 (black dashed line, $p < 0.01$, Student's t-test), indicating that more neurons are tuned for the non-biased movement. $N = 26$ imaging sessions from 8 mice.

(D-F) Fraction of cells preferring down movements increases as the forward bias and forward stereotypy increase. Each dot represents one imaging session.

(G) Trial-by-trial correlation of population activity is lower for biased (grey, $N = 60,416$ trial pairs across 26 imaging sessions) than non-biased (black, $N = 53,319$ trial pairs) movements. $p < 0.001$, Student's t-test.

(H) Trial-by-trial correlation of population activity is more related (greater slope) to the trial-by-trial correlation of non-biased movements than biased movements. $p < 0.05$ for both ordinary and Deming linear regressions, bootstrap.






Figure 2.4 M1 inactivation does not inhibit initiation of expert movements

(A) Inactivation trial task structure.

(B-G) Fraction of trials with no movements (**B**, $p = 0.21$), reaction time of successfully withheld trials (**C**, $p = 0.15$), fraction of trials with movements that failed to reach either target (**D**, $p = 0.07$), fraction of trials withheld until go-cue (**E**, $p < 0.001$), fraction of trials correctly discriminating (**F**, $p < 0.001$), and magnitude of bias (**G**, $p < 0.001$) for control (light off) and inactivation (light on) trials. Thin lines, mean across inactivation sessions of single mice ($N = 14$); thick lines, median across mice.




Figure 2.5 M1 inactivation enhances existing bias in expert mice

(A) Example performance on 75 trials from session 87 of the mouse in **Figure 2.2B**. Blue triangles, inactivation trials; black dots, correct trials; red x's, incorrect trials; blue squares, failure to reach either target. Top row down visual stimulus presented, bottom row forward visual stimulus presented. Observe that in most inactivation trials the animal responded down regardless of which stimulus was presented.

(B) Bias from all inactivation sessions of the mouse shown in **A** ($N = 5$ sessions). Note that the mouse had a bias for down movements in control trials, which was enhanced in the inactivation trials.

(C) Bias in control trials correlates with bias in inactivation trials. Dots represent individual inactivation sessions ($N = 70$), colors represent different mice ($N = 14$). Regression slope > 1 (Deming regression, $p < 0.05$, bootstrap) indicating that inactivation bias is enhancement of control bias.

(D-E) Inactivation bias correlates with the difference in movement stereotypy in the two directions. Dots and colors same as in **C**.




Figure 2.6 M1 inactivation during early learning impairs movement initiation and execution

(A-C) Fraction of trials with no movements (**A**, $p < 0.05$), reaction time (**B**, $p < 0.05$), and fraction of trials with movements that failed to reach either target (**C**, $p < 0.01$) for control (light off) and inactivation (light on) trials in early-learning inactivation. Thin lines, mean across inactivation sessions of single mice ($N = 9$); thick lines, median across mice.

(D-E) Fraction of trials in which the forelimb left the joystick in early-learning inactivation (**D**, $p < 0.05$, $N = 6$) and expert inactivation (**E**, $p = 0.18$, $N = 5$). Thin lines, mean across inactivation sessions of single mice; thick lines, median across mice.

(F) Fraction of trials with movement impairment, which include no movements, movements that failed to reach either target, and joystick release, in early-learning inactivation. Left, both stimulus conditions ($p < 0.05$); middle, stimulus in biased direction ($p < 0.05$); right, stimulus in non-biased directions ($p < 0.05$). Thin lines, mean across inactivation sessions of single mice; thick lines, median across mice.

(G) Same as F for expert inactivation. Left, both stimulus conditions ($p = 0.18$); middle, stimulus in biased direction ($p = 0.43$); right, stimulus in non-biased directions ($p = 0.18$).




Figure 2.7 Biased movements are more stereotyped than non-biased movements

(A) Variability of end point angle ($p < 0.001$) and **(B)** mean M.S.E. ($p < 0.001$) for biased and non-biased movements. Each line represents one inactivation or imaging session (174 sessions across 22 mice). **(C)** Normalized difference of movement stereotypy measured as MSE between forward and down directions across the last 10 days of experiments for imaging and expert inactivation mice (N = 22).

Figure 2.8 Optogenetic inactivation of PV inhibitory neurons inactivates cortex


(A) Spike raster of multiunit activity from a single electrode on a linear silicon probe for control (blue light off, left) and inactivation (blue light on, right) trials over a single recording block of blue light intensity 7.5 mW. V-Stim On: visual stimulus onset (indicated by grey vertical line); V-Stim Off: visual stimulus offset (indicated by grey vertical line); blue shaded region indicates when blue light was on.

(B) Firing rate normalized to baseline firing rate. Baseline firing rate was calculated as the mean firing rate in periods 2 sec before visual stimulus onset and 2 sec after visual stimulus offset. Binned in 200 msec time windows for control (blue light off, left) and inactivation (blue light on, right) trials in **A**.

(C) Ratio (inactivation / control) of normalized firing rates in **B**.

(D) Reduction of firing rate during blue light period (blue shaded region in **C**). Each point represents mean $\pm$ standard deviation of reduction of multiunit activity across all electrodes in two recording sessions.

(E) Number of multiunit recordings for different blue light intensities.






Figure 2.9 Control sessions of expert inactivation experiments

(A-F) Control sessions were performed identically to inactivation sessions (**Figure 2.4A**), but with blue light pointed away from brain. We did not observe non-specific effects of blue light in control sessions.




Figure 2.10 Kinematics of movements in early-learning are not reliably different between the two directions

(A-B) Bias is not correlated with the difference in variability of end point angle and M.S.E. (respectively) in early-learning. Each dot represents one inactivation session, different colors represent different mice (35 sessions from 9 mice).
(C) Control bias is not related to inactivation bias in early learning.




Figure 2.11 Control cohort of early learning inactivation experiments

(A-G) A separate cohort of mice were run identically to mice in **Figure 2.6**, but with blue light pointed away from the brain. We did not observe non-specific effects of blue light.

Chapter 2 currently being prepared for submission and will include Dr. EunJung Hwang, Madan Mukundan and Dr. Takaki Komiyama as authors. This material was included with the generous consent of all authors. The dissertation author is a primary researcher on this work.

Chapter 3. History-based action selection bias in posterior parietal cortex

3.1 Summary

Even identical sensory evidence can lead to different decisions over time and across individuals. Here we explore the neural mechanisms underlying decision variability. In a novel visually-guided action selection task that we developed, the response of mice to identical sensory stimuli exhibited high degrees of trial-to-trial variability. Behavioral modeling revealed that variability arose from the idiosyncratic use of recent choice-outcome history, which biased subsequent choices. Two-photon calcium imaging identified that pre-stimulus activity in the posterior parietal cortex (PPC) closely reflected these idiosyncratic biases and history. Chronic imaging revealed that history-tuning of PPC neurons changed over time to reflect the ongoing strategies. Optogenetic inactivation of the PPC pre-stimulus activity, but not trial activity, weakened the idiosyncratic biases, i.e., the relationship between choice-outcome history and decisions. Consequently, when maladaptive biases were disrupted by optogenetic inactivation, the behavioral performance improved. These results uncover that PPC mediates subjective uses of history in biasing action selection.

3.2 Introduction

Decisions in response to a given sensory input are variable across individuals, as they are subject to individuals' subjective and often subconscious internal biases (Busenitz and Barney, 1997; Greenwald et al., 1998; Kable and Glimcher, 2007; Powell and Ansic, 1997; Santesso et al., 2008). These internal biases can change on a moment-by-moment basis depending on recent personal history and subjective rules to

utilize this history (Abrahamyan et al., 2016; Barraclough et al., 2004; Brehm, 1956; Busse et al., 2011; Egan et al., 2007; Gold et al., 2008; O'Doherty et al., 2004; Sugrue et al., 2004; Sul et al., 2011). For example, imagine when you are deciding on a meal to order at your favorite restaurant. Some may tend to order the same meals repeatedly, while others may explore a variety of options over time. If you enjoyed the dish you ordered the last time you dined there, you may be more inclined to order it again. These biases, shaped by the choice-outcome history, can allow one to infer the rules of the environment and generate adaptive behavioral strategies (Barraclough et al., 2004; O'Doherty et al., 2004; Sugrue et al., 2004; Sul et al., 2011). History-dependent biases are not limited to adaptive contexts, but instead human and animal subjects show diverse idiosyncratic history-dependent biases (e.g. win-stay, lose-switch, or lose-stay) even when biases are maladaptive (Abrahamyan et al., 2016; Brehm, 1956; Busse et al., 2011; Egan et al., 2007; Fründ et al., 2014; Gold et al., 2008). The prevalent history-dependency of decisions suggests that tracking choice-outcome history to form subjective bias is a fundamental aspect of decision-making, yet the neural circuits mediating this process are largely unknown.

Responses of neurons in the sensorimotor pathway including areas implicated for decision-making show degrees of variability even for identical sensory inputs and motor outputs (Churchland et al., 2011, 2010). Such variability is often treated as random noise (Brunton et al., 2013; Faisal et al., 2008; Rolls and Deco, 2011). However, for future decisions to be biased by history, it is necessary that neural responses are modulated by history, accounting for some of the observed neural variability. Indeed, choice-outcome history has been shown to modulate the activity in a variety of brain areas, including parietal, prefrontal and premotor cortex and subcortical structures (Barraclough et al., 2004; Histed et al., 2009; Nakamura et al., 2012; Parker et

al., 2016; Seo et al., 2009b; Sul et al., 2011; Zalocusky et al., 2016). For example, neurons in the lateral intraparietal area of the monkey posterior parietal cortex (PPC) that has been implicated for the accumulation of sensory evidence are modulated by the choice and outcome information of recent trials (Seo et al., 2009b; Shadlen and Newsome, 2001). However, history affects action selection biases in flexible and complex ways that vary over time and across individuals (Barracough et al., 2004; Seo et al., 2009b). How the history signals in the brain may account for such a flexible relationship between history and decision bias is unclear. Furthermore, these brain areas often contain intermingled neurons with diverse temporal activity profiles (Harvey et al., 2012; Hwang and Andersen, 2011; Seo et al., 2009b), and the roles of specific temporal windows of history-related activity cannot be accessed with traditional lesion or pharmacological inactivation approaches.

Here we combined behavioral modeling, two-photon calcium imaging and temporally-precise inactivation to explore the mechanisms of the subjective, history-dependent decision bias in mice performing a novel visually-guided action selection task. Our behavioral model identified diverse, idiosyncratic relationships between choice-outcome history and action selection bias. These idiosyncratic biases were highly correlated with the pre-stimulus activity of a subset of neurons in PPC. Temporally-precise inactivation revealed a causal role of the pre-stimulus activity of PPC, but not trial activity, in action biases.

3.3 Results

Visually-guided action selection task for head-fixed mice

We developed a task in which head-fixed mice moved a joystick with their left forelimb in one of two directions in response to visual cues (Figure 3.1A and 3.9A). In

each trial, one of two visual stimuli (gratings moving forward or downward) was presented for one second. This stimulus period was followed by a two sec memory period. The end of the memory period triggered an auditory go cue, and mice were required to move the joystick in the remembered direction of the visual stimulus to receive a water reward. Trials with no movements, movements before the go cue, and movements in the wrong direction were not rewarded. Mice performed one session per day, 281 ± 55 trials per session.

After 2-4 months of incremental training (Figure 3.9B and Methods), mice achieved a plateau level of performance which was significantly above chance (Wilcoxon one-sided signed rank test, $p < 6 \times 10^{-5}$), but far from perfect (Figure 3.1B). The behavioral performance was not limited by a difficulty in visual discrimination, as mice performed significantly better in another version of the task without the memory period (Figure 3.1C and 3.10A). The suboptimal performance in the memory task, characterized by variable responses to the same visual stimulus in individual trials, is essential for uncovering internal biases underlying choice variability as shown below.

History-based action selection bias revealed by behavioral modeling

We hypothesized that the choice variability in individual trials reflects a systematic fluctuation of hidden internal biases that are shaped by the recent history of the mice, rather than a random fluctuation due to neural noise. To test this hypothesis, we built a logistic regression model of the behavior. A similar model has been previously described (Busse et al., 2011; Lau and Glimcher, 2005). Briefly, this model predicts the choice of each mouse on individual trials by utilizing the sensory stimulus of the current trial, the choice-outcome history from previous trials, and a constant choice preference (Equation 1 and Methods). Accordingly, the portion of the equation excluding the current stimulus corresponds to the estimate of the internal bias on each trial. Regression was

performed in each session independently to identify the weight of each term and the time constants of history terms that best fit the behavior. To avoid overfitting, the accuracy of the model was quantified in a cross validated manner in which the model was built using a fraction of the trials in the session ('training set') and evaluated for the accuracy on the remaining trials ('test set').

$$\log \frac{\text{probability}\{\text{choice}(N)=\text{forward}\}}{\text{probability}\{\text{choice}(N)=\text{downward}\}} = w_s \cdot \text{stimulus}(N) + w_o \cdot \sum_{k=1}^{N-1} \text{outcome}(k) \cdot e^{-\frac{N-1-k}{\tau_o}} + w_c \cdot \sum_{k=1}^{N-1} \text{choice}(k) \cdot e^{-\frac{N-1-k}{\tau_c}} + w_{oc} \cdot \sum_{k=1}^{N-1} \text{outcome}(k) \cdot \text{choice}(k) \cdot e^{-\frac{N-1-k}{\tau_{oc}}} + \text{constant} \quad (1)$$

This model with choice-outcome history predicted the behavior significantly better than the stimulus alone (Figure 3.2A-B), indicating that the choice variability observed in this task is indeed not random. Instead, a significant part of the variability arises from a systematic influence of choice-outcome history that biases the decision on a trial-by-trial basis. To assess the contributions of distinct components of history information to the internal bias, we built partial models using only a subset of history information at a time (Figure 3.2C). On average, the outcomes (reward or error) of previous trials carried the largest predictive power, followed by the choices (forward or downward) of previous trials. The weight of each variable in the full model (Eq.1) was examined in sessions in which the corresponding variable exerted significant influence on the choice (Figure 3.2D). We found that the weight for the visual stimulus was significant and positive in a majority of sessions (9/17), demonstrating that animals properly used the stimulus information despite a low behavioral performance. The weight for the previous trial outcomes was significant in 9/17 sessions with short (< 1) time constants, indicating that mice tended to choose one direction after reward and the other direction after error (e.g., forward following reward trials and downward following error trials, irrespective of the

choices of previous trials). This strategy differs from the so-called ‘win-stay/lose-switch’ in which choice depends on both the outcomes and choices of previous trials (e.g., forward following rewarded forward trials, but downward following rewarded downward trials). Furthermore, the choices of previous trials significantly contributed to the prediction in 8/17 sessions, all with negative weights and longer (~10 trials) time constants, showing that mice had a tendency to equalize the frequencies of both forward and downward choices over time (Figure 3.2D). In the non-memory version of the task in which mice performed better (Figure 3.1C), the weight for visual stimulus was significantly larger, and the outcome history weight was significantly smaller than in the memory task (Figure 3.10B), indicating an enhanced use of the stimulus information and reduced importance of history information.

Mice employed somewhat common strategies in the memory task as described above, but the weights and time constants of individual terms were highly variable across individual sessions (Figure 3.2D). Consistent with the variability of weights and time constants, model accuracy was significantly worse when a model built for one session was applied to another session of the same mouse (Figure 3.2E). Thus, the rules that each mouse employed to use choice-outcome history were variable over days, similar to those previously reported in human subjects (Abrahamyan et al., 2016). The detriment in model accuracy was even larger when a model from one mouse was applied to another mouse, demonstrating an idiosyncratic nature of the strategies (Figure 3.2E). Taken together, the imperfect behavioral performance in conjunction with the behavioral model gives us an opportunity to estimate the hidden internal biases underlying decision variability on individual trials that are not directly measurable.

We note that these strategies can be somewhat adaptive or maladaptive, due to our stimulus selection algorithm. Specifically, the same stimulus was repeated after error

trials, and the stimulus was changed to the other stimulus after three consecutive rewarded trials with the same stimulus. In all other trials, stimulus was randomly selected. These rules were introduced to discourage mice from selecting the same choice (forward or downward) in every trial. Thus, the two common strategies described above, biasing choice based on the outcome history of the immediately preceding (N-1) trial, and equalizing choice frequencies, were both adaptive and helped mice perform better than chance. In contrast, a strong constant preference of one choice was an example of maladaptive strategies, as it results in repetition of the same error for multiple trials. The varying degrees of ‘adaptiveness’ of models in different sessions were quantified by assessing the success rate of the internal bias model (excluding the stimulus term in the full model) of each session in simulation in which the stimulus was selected according to the same rules. We found that 53% of sessions showed significantly adaptive strategies, while 24% were significantly maladaptive (Figure 3.2F).

Internal biases are represented in the pre-stimulus activity of posterior parietal cortex

To explore the neural basis of these subjective, history-based internal biases, we applied two-photon calcium imaging to record the neural ensemble activity in PPC while mice performed the task (Figure 3.3A). We chose PPC because it is widely implicated in decision making processes, and PPC is highly interconnected with visual and motor areas, placing it in an ideal location to bias the transformation of visual information to motor outputs (Harvey et al., 2012; Platt and Glimcher, 1999; Raposo et al., 2014; Shadlen and Newsome, 2001). We imaged neurons in layer 2/3 expressing GCaMP6f in the 17 sessions from 8 mice whose behavioral analyses were presented in the previous section. 991 unique neurons (mean: 58 / session, range: 15-123) showed significant task-related activity and were included in the analysis.

Consistent with previous recording studies in monkeys and rodents (Guo et al., 2014b; Hanks et al., 2006; Harvey et al., 2012; Shadlen and Newsome, 2001), many PPC neurons (66%, 652/991) exhibited choice-selective activity such that the activity in forward-pressing and downward-pressing trials was significantly different. Choice-selective neurons presented varying timing of choice selectivity, and largely distinct populations of neurons showed choice selectivity during the stimulus, memory and movement periods (Figure 3.3B-D). Surprisingly, a large fraction of choice-selective neurons (56%, 362/652) showed significant choice selectivity during the inter-trial interval (ITI) before the stimulus was presented (Figure 3.3C-E). In other words, these neurons had predictive information about the eventual choice of the mouse before the trial was initiated.

A trivial explanation for the ITI choice-selectivity is that the mice are already preparing or partially executing the movement during the ITI, and thus ITI choice-selective neurons are movement-related neurons. To address the possibility that mice are partially executing the movement during the ITI (e.g. leaning on the joystick), we performed additional experiments in which the joystick was unfixed in the ITI of a subset of trials. Choice selectivity of PPC neurons during the ITI in the trials when mice did not apply force on the unfixed joystick remained the same as in the joystick-fixed condition, suggesting that the ITI choice selectivity does not arise from partial execution of movements (Figure 3.11). Moreover, a majority of ITI choice-selective neurons (69%, 251/362) were non-movement neurons, distinct from movement neurons that showed maximal choice selectivity during the peri-movement period (-0.5 to 1.5 sec around movement onset; Figure 3.4A). Conversely, a majority of non-movement choice-selective neurons showed choice-selectivity during the ITI (85%, 251/292). The choice-selectivity of these non-movement choice-selective neurons appears to be related to

their sensitivity to the recent choice-outcome history. Of non-movement neurons, 92% (270/292) and 46% (135/292) were significantly tuned during the ITI to the outcome and choice of the immediately preceding (N-1) trial, respectively (Figure 3.4B). Accordingly, these N-1 outcome- and choice-tuned neurons are differentially modulated during ITI even when the choice in the upcoming N trial is the same, depending on the previous trial conditions (Figure 3.4C).

Notably, many non-movement choice-selective neurons (43%, 125/292) were tuned to both N-1 outcome and N-1 choice, suggesting that multiple types of distinct history information are mixed at the level of individual neurons. However, in theory, the concurrent tuning to multiple variables could result from the correlation of N trial choice with both N-1 outcome and N-1 choice as shown by our behavioral modeling. To disentangle this confounding relationship, we evaluated how the ITI activity of individual ITI choice-selective neurons was modulated by history information (N-1 trial outcome and N-1 trial choice) and N choice (used as a binary estimate of internal bias) independently, by focusing on trials in which 2 of the 3 variables were identical. For example, neuron 1 shown in Figure 3.4D was strongly modulated by N-1 outcome, exhibiting different levels of activity after rewarded and error trials. This modulation by N-1 outcome was clearly present even when we considered only the trials in which N-1 choice and N choice were fixed, indicating that the N-1 outcome modulation of this neuron was not a secondary effect of correlation between N-1 outcome and N-1 choice or N choice. In addition, this neuron was also modulated by both N-1 choice and N choice, independent from its modulation by N-1 outcome. Overall, large fractions of ITI choice-selective neurons exhibited independent tuning for N-1 outcome, N-1 choice, and N choice (Figure 3.4E), indicating that distinct history and bias information is encoded in overlapping but distinct populations of individual PPC neurons.

Importantly, such tuning of individual neurons was not fixed but sensitive to the current strategies employed by the mice. We addressed this issue in a subset of experiments in which we imaged the activity of the same populations of PPC neurons across multiple sessions (386 neurons in 4 mice were imaged across 4 to 7 sessions). For example, the neuron in Figure 3.5A imaged over 5 sessions is tuned to N-1 outcome in all imaged sessions during the ITI. However, the strength of its N-1 outcome tuning varied across sessions, tracking the strength of the influence of previous outcome on the subsequent choice as estimated by the accuracy of the partial model using the previous outcome information only. That is, the neuron showed more pronounced N-1 outcome tuning in sessions in which the previous outcome influenced the upcoming choice more strongly. Such flexible modulation of N-1 outcome tuning was consistent across PPC neurons (Figure 3.5B). Similar effects were found for N-1 choice tuning (Figure 3.5C). The flexible sensitivity of PPC neurons to distinct history information may underlie the flexible, subjective use of history to generate bias revealed by our behavioral modeling (Figure 3.2E).

In line with this notion that PPC represents the subjective bias, the weighted sum of ITI activity of neurons simultaneously recorded from PPC was able to fit very closely the internal biases estimated by our behavioral model (Figure 3.6A-C). The excellent fit was specific to the internal biases, and ITI activity could not fit the trial-shuffled internal biases (Figure 3.6B). The PPC ITI activity closely tracked the fluctuations of the strengths of internal biases even in trials of the same choice (Figure 3.6B-C), suggesting that the PPC ITI activity reflected continuously varying internal biases rather than categorical choice. PPC ITI activity better predicted the subsequent choice when the bias direction estimated from PPC ITI activity matched that of the subsequent stimulus,

supporting the idea that the final choice is made by integrating the biases encoded in PPC and the subsequent stimulus (Figure 3.6D).

The earlier analyses of single neuron responses showed that individual PPC neuron encode history and choice information in a mixed and heterogeneous manner (Figure 3.4). Consistently, the PPC ensemble activity also encodes a mixture of history and bias information. When the ensemble ITI activity was fit separately to the internal biases and the outcome in the N-1 trial, we found that the population activity encoded both bias and N-1 outcome independently. That is, even for the same value of internal bias, population activity was still separable depending on the N-1 outcome (Figure 3.6E-F). Similar results were found between bias and N-1 choice (Figure 3.6E-F). Thus, PPC neuronal population encoded both history and bias information independently during the ITI.

Optogenetic inactivation of PPC ITI activity alters internal bias

The results so far indicate that PPC contains information about action selection biases during the pre-stimulus ITI. To address whether this information in PPC is indeed used to bias the subsequent actions, as opposed to the alternative possibility that actions are biased by activity elsewhere and PPC activity simply correlates with it, we used optogenetics to inactivate PPC during the task. We injected Cre-dependent AAV encoding Channelrhodopsin-2 (ChR2) in PV-Cre mice to express ChR2 in parvalbumin-positive inhibitory neurons in PPC and these mice were trained with the task (N=8). Once their performance reached a plateau, we started inactivation sessions in which blue light was applied bilaterally to inactivate PPC during the ITI of a small subset (~15%) of trials (Figure 3.7A). If PPC is indeed essential for the internal biases based on choice-outcome history, then PPC inactivation should weaken the idiosyncratic relationship between history and choice. We tested this idea by building a behavioral

model with a subset of the unperturbed light-off trials as the training set and testing the accuracy of choice prediction for the remaining unperturbed light-off and inactivated light-on trials. Consistent with our hypothesis, the model built with unperturbed trials ('light-off model') was significantly better at predicting the choice on other light-off trials than light-on trials (Figure 3.7B-D and 3.14A). This result shows that PPC inactivation altered the idiosyncratic relationship between choice-outcome history and the subsequent actions. Such an effect was not observed in control sessions of the same mice in which the light was directed at the head bar instead of PPC (Figure 3.15A). Consequently, the light-off model performed significantly worse in inactivation light-on trials than control light-on trials (Wilcoxon one-sided signed rank test, $p < 0.004$), indicating that the effect was due to PPC inactivation and not due to non-specific effects of light. The altered history dependency occurred without significant changes in reaction time or movement time (Figure 3.14).

To our surprise, the history-choice relationship was not altered when we inactivated PPC during the trial period from the beginning of the visual stimulus until the end of the trial ($N=7$, a subset of the PPC ITI inactivation mice; Figure 3.15E-F). Accordingly, the model performed significantly worse in ITI inactivation trials than trial inactivation trials (Wilcoxon one-sided rank-sum test, $p < 0.01$). These results suggest that the bias information encoded in PPC during the pre-stimulus ITI is subsequently maintained elsewhere to guide behavior independent of later PPC activity. Importantly, the consistent effect of ITI inactivation on internal bias was specific to PPC. When we inactivated the primary motor cortex (M1) during the ITI in a separate set of mice ($N=7$), the effect was variable across animals (Figure 3.7G), thus no significant difference was observed between unperturbed light-off and inactivated light-on trials. Therefore, the

altered relationship between history and subsequence choice is not a general effect of inactivation, but instead it is specific to PPC.

To delineate the nature of the altered relationship between history and choice by PPC ITI inactivation, we fit the light-on and light-off trials with separate full models (Eq. 1), and compared the weights of the two models. Because of a greater number of light-off trials, we built light-off models using randomly sampled light-off trials matching the number of light-on trials 100 times. The light-on model was compared to the distributions of the 100 light-off models of each session. We found that the light-on model predicts the choice significantly worse than the mean light-off model across 56 sessions (Wilcoxon signed rank test, $p < 0.016$; Figure 3.8A). The decrease of model accuracy was accompanied by significant decreases in the weights for the outcome, choice, and outcome-choice interaction history (bootstrap, $p < 0.001$, $p < 0.001$, and $p < 0.05$ respectively; Figure 3.8B). These weight changes were not observed in control sessions (Figure 3.15E), and thus the decreases in the three history weights were significantly larger than in control sessions (bootstrap, $p < 0.001$, $p < 0.001$, and $p < 0.05$, respectively for outcome, choice, and outcome-choice history). These results indicate that PPC inactivation during the ITI weakened the dependency of subsequent action choice on choice-outcome history.

Given that some sessions showed adaptive and maladaptive strategies (Figure 3.2F), we hypothesized that weakened history dependency would deteriorate behavioral performance in sessions with adaptive strategies, and improve performance in sessions with maladaptive strategies. Consistent with this prediction, the change in behavioral performance induced by PPC ITI inactivation was negatively correlated with the degree of adaptiveness of the strategies (Pearson correlation coefficient = -0.5 , $p < 0.0001$; Figure 3.8C). In contrast, the change in behavioral performance and the degree of

adaptiveness was not significantly correlated in control sessions (Pearson correlation coefficient = -0.14, $p=0.08$; Figure 3.15F). Accordingly, the slope of linear regression on performance change against adaptiveness was significantly steeper in the inactivation sessions compared to control sessions (bootstrap; $p<0.005$). These changes in behavioral performance provide additional evidence that PPC ITI activity is essential for the history-dependent biases, and the bidirectional effects suggest that PPC is responsible for the range of variable and idiosyncratic strategies to utilize the history information.

3.4 Discussion

Our decisions are fundamentally influenced by our hidden internal biases (Dolan and Sharot, 2011; De Martino et al., 2006; Tversky and Kahneman, 1981). How often do we regret the suboptimal decisions we made while ignoring compelling objective evidence? Here we combined a variety of cutting-edge approaches to uncover an important role of PPC in the subjective use of history in biasing action selection.

Choice-outcome history and action bias

Our decision biases often evolve continuously over time depending on our choice-outcome history (Brehm, 1956; O'Doherty et al., 2004). A classic example is the 'win-stay/lose-switch' strategy in which one tends to repeat the same choice after a good outcome and switch to another choice after a bad outcome. In many cases, however, the influence of history on decision biases is nuanced and the individual is often unaware of their own subjective rules governing the relationship between history and biases. These biases are inherently internal and difficult to measure experimentally, rendering their neural basis hard to study. Attempts have been made to estimate these internal calculations by mathematical modeling of behavior based on measurable

parameters (Busse et al., 2011; Gold et al., 2008; Sugrue et al., 2004). Our behavioral model built on this line of work and successfully extracted the complex contributions of choice-outcome history on action biases in our task. The relationship between choice-outcome history and bias was variable from session to session and animal to animal, resonating with the idiosyncratic rules humans employ in everyday and perceptual decision making (Abrahamyan et al., 2016; Busenitz and Barney, 1997; Drolet, 2002; Kable and Glimcher, 2007; Powell and Ansic, 1997).

The strategies of the mice that were revealed by the behavioral model included adaptive and maladaptive strategies. Adaptive strategies took advantage of the hidden algorithm for non-random stimulus presentation in a subset of trials, which mice must have learned during many weeks of training. However, the impact of these adaptive strategies was subtle, and on average mice performed only slightly better on the trials in which the stimulus was deterministic (Figure 3.9C). The changes in various adaptive and maladaptive strategies over time likely represent the continued search by the mice for a better strategy.

Causal link between PPC pre-stimulus activity and action biases

The pre-stimulus activity of PPC during the ITI closely reflected the subjective, internal bias estimated by our behavioral model and accordingly predicted the future choice. Choice-predicting pre-stimulus activity has been reported in various brain areas including the visual, parietal, premotor, and prefrontal cortex (Coe et al., 2002; Maoz et al., 2013; Padoa-Schioppa, 2013; Platt and Glimcher, 1999; Shadlen and Newsome, 2001; Wyart and Tallon-Baudry, 2009). However, in contrast to the current study, these previous studies did not systematically relate the pre-stimulus activity to decision variables such as history-dependent internal bias and could not distinguish it from

stochastic neural noise such as ongoing fluctuations of baseline. Moreover, the causal relationship between pre-stimulus activity and future choice has not been tested.

To our knowledge, we have demonstrated for the first time that the PPC pre-stimulus activity is essential for the influence of biases on subsequent actions. This was shown by optogenetic inactivation during the ITI that weakened the idiosyncratic relationship between choice-outcome history and subsequent choice. In agreement with the behavioral model that extracted adaptive and maladaptive strategies, the effect of PPC ITI perturbation on behavioral performance was strategy-specific; it impaired behavioral performance in mice that utilized adaptive strategies, while the performance improved in mice that used maladaptive strategies. These results support the general role of PPC in subjective, history-dependent biases, including those that are adaptive, neutral, and maladaptive.

Our temporally precise optogenetic inactivation revealed that the effect of PPC inactivation was specific to ITI, and perturbation after stimulus onset did not cause a measurable effect on behavior. This result implies that bias information encoded in PPC during ITI is unloaded to some other areas after the ITI and maintained in a PPC-independent manner. PPC neurons have projections to various brain areas (Harvey et al., 2012; Wang et al., 2012), and identifying these downstream areas that are responsible for the bias execution is an important topic of future research.

It is important to note that our result does not imply that the post-stimulus choice-selective activity in PPC has no functions. In fact, several studies reported altered behavioral performance in perceptual discrimination tasks following PPC perturbation (Goard et al., 2016; Hanks et al., 2006; Raposo et al., 2014). PPC also contributes to movement planning and execution, and inactivation in monkeys can affect movement end point control (Hwang et al., 2012). Such a role in fine motor control or

sensory evidence accumulation is distinct from the bias coding that we describe here and was not tested in this study.

Mixed coding of history and bias in posterior parietal cortex

We found that the PPC ITI activity contains both history and bias information mixed at the level of individual neurons. This observation clearly excludes two extreme possibilities; 1) PPC only contains history information and is upstream of bias computation, and 2) PPC only contains bias information and is downstream of bias computation. While the precise circuit mechanisms underlying the transformation of history information into bias are extremely difficult to uncover, based on the mixed representation of history and bias in PPC, we favor the view that PPC participates in the computation of subjective bias from history information. It is important to note that these PPC neurons that encode history and bias information are intermingled with other neurons that are selectively active during visual stimulation, delay, and movement periods. PPC thus likely contains multiplexed, parallel pathways dedicated to the processing of distinct forms of information.

An intriguing observation is that the history tuning of individual PPC neurons reflected the ongoing strategies employed by the mouse. Specific history information modulated PPC ensembles during ITI more strongly when it had a stronger impact on bias. We propose that such flexible tuning of PPC neurons to past events underlies the variable and idiosyncratic use of history that we describe here. The various ways by which identical events affect the future behaviors of different individuals are an important manifestation of individuals' personalities, and the variable tuning of PPC to history information may be a neural substrate for such a process.

Relation to other findings in rodent PPC

Although the functional homology between primate and rodent PPC is not fully established, several rodent PPC studies have found neural response properties analogous to primate PPC and started to provide further insights into PPC circuits and functions (Funamizu et al., 2016; Goard et al., 2016; Hanks et al., 2015; Harvey et al., 2012; Raposo et al., 2014). Additionally, we note that two recent findings in rodent PPC resonate with our current study: 1) PPC population activity exhibits slow dynamics that integrate recent events (Morcos and Harvey, 2016), and 2) PPC perturbation affects internally-guided decisions (Erlich et al., 2015). Our finding that PPC neurons encode a mixture of history and bias to influence action selection demonstrates an important functional consequence of the former observation. The representation of history-dependent internal bias in PPC presents a mechanism for PPC to affect internally-guided decisions. Furthermore, our finding that intermingled but distinct sets of neurons represent specific sets of information lays foundation for investigating functional diversity in PPC microcircuits, likely linked with projection target areas.

3.5 Methods

Animals

All procedures were in accordance with protocols approved by the UCSD Institutional Animal Care and Use Committee and guidelines of the National Institute of Health. Mice (calcium imaging: cross between *Gad2-IRES-Cre* [JAX 010802] (Taniguchi et al., 2011) and *Rosa26-CAG-LSL-tdTomato* [JAX 007914] (Madisen et al., 2010) or *Rosa26-CAG-LSL-tdTomato* or cross between *Camk2a-tTA* [JAX 003010] and tetO-GCaMP6s [JAX 024742]; optogenetic perturbation: *PV-Cre* [JAX 008069] (Hippenmeyer et al., 2005)) were housed in a room with a reversed light cycle (12 h–12 h). Experiments were performed during the dark period.

Behavioral task

Mice were trained under head-fixation in the behavioral apparatus, approximately 1 hour per day over a period of 2-4 months. The task was shaped to reach the final version through 8 training steps (Figure 3.9B, see **Behavioral Training**). The visual stimulus was randomly selected between forward or downward drifting gratings with the following constraints: 1) after three consecutive rewarded trials in one direction, the stimulus always switched to the other direction, and 2) after error trials, the same stimulus was repeated. These constraints were implemented to discourage the mice from choosing only one direction and settling at 50% discrimination accuracy. Despite the deterministic stimulus after an error or a third consecutive reward in one direction, the mice performed only slightly better in those trials than random trials (Wilcoxon one-side signed rank, $p < 0.04$; Figure 3.9C), indicating that mice did not fully utilize these hidden stimulus presentation rules to their advantage.

Because of the pseudo-random rules of stimulus presentation, the fraction of correctly discriminating trials (*# of trials hitting the correct target / # of trials hitting any target*) achieved by random choice would not be 50% if there was a constant choice preference. So we estimated the constant choice preference within a session and converted it to a probability to choose each choice using the following formulae:

$$\text{Probability of choice 1} = \frac{1}{2} \times \left(\frac{\# \text{ of trials}_{\text{choice 1} | \text{stimulus 1}}}{\# \text{ of trials}_{\text{stimulus 1}}} + \frac{\# \text{ of trials}_{\text{choice 1} | \text{stimulus 2}}}{\# \text{ of trials}_{\text{stimulus 2}}} \right)$$

$$\text{Probability of choice 2} = 1 - \text{Probability of choice 1}$$

Then, the chance level performance for the given session was computed by simulating random binary choice with the estimated probabilities under the same pseudo-random rules 1000 times (Figure 3.1B).

Behavioral training

Adult mice (six weeks or older, male and female) were implanted with a custom head-fixation plate on the skull. Following a minimum 3 days of recovery, daily water consumption was limited to a controlled amount (typically 1 mL/day). Behavioral training began following 3 to 10 days of water restriction.

A custom-built behavioral apparatus housed in a box (40 × 40 × 40 cm) included a joystick (M11L061P; CHProducts), a 17 inch computer monitor (for visual stimulus presentation; placed ~15 cm from the right eye of the mouse), and a water port with photodiodes to sense licking (Figure 3.1A). The stock joystick handle was custom machined and retrofitted with a 1/16 inch thick brass rod that mice manipulated with their left forepaw (Figure 3.9A). An electromagnet (EM050-3-222; APW) was situated so it could be used to mechanically immobilize the joystick at the origin. The joystick had a dynamic range of 56° in each angular direction forming a spherical endpoint space (Figure 3.9A). The 2D angular position of the joystick was continuously recorded at 1 kHz using a data acquisition card (USB6008; National Instruments) and custom Matlab software. The task-sequence execution, stimulus selection, auditory cue presentation, reward dispensation, and task time recording were coordinated by an open source real-time Linux/Matlab software package BControl (<http://brodywiki.princeton.edu/bcontrol/>). The presentation of visual stimuli (100% contrast, full-field, square wave drifting gratings 0.04 cycles/degree and 3 cycles/sec) was implemented using Psychtoolbox (an open source Matlab toolbox; <http://psychtoolbox.org/>).

In the two-alternative forced-choice task (Figure 3.1A), one of two orthogonal visual stimuli (forward or downward moving gratings) was presented for 1 sec, followed by a 2-sec memory period. After the memory period, an auditory cue (6 kHz pure tone) marked the response period (up to 10 sec) during which the joystick entering the correct

target area (hereafter referred to simply as 'target'; Figure 3.1A and Figure 3.9A) in the same direction as the gratings triggered a water reward. Errors (i.e., entering the incorrect target, and movements before the go cue) triggered a white noise sound and led to an immediate trial-abortion. Following reward, trial-abortion, or no response, the return of the joystick to the origin ended the trial and initiated an inter-trial interval (ITI, 4 or 8 sec, constant within each session). During the ITI the joystick was immobilized at the origin by an electromagnet. At the end of the ITI (simultaneous with the beginning of visual stimulus onset), the electromagnet was disengaged, and the joystick was free to move. Thus, if mice already pushed the joystick in any direction from the ITI, the joystick would have moved out of the origin as soon as the next trial stimulus period began and the trial was most likely aborted (see below for the withholding requirement). However, these movements immediately after the ITI (movement onset within 100 ms from stimulus onset) were rare (1%, 45/4747).

Mice were trained under head-fixation in the behavioral apparatus, approximately 1 hour per day over a period of 2-4 months. The task was shaped to reach the final version through 8 training steps (Figure 3.10B). In the first step, the mice received a water reward as long as they moved the joystick to the correct target within a 30-sec response period (regardless of whether or not they hit the incorrect target first). As they became more proficient with pressing the joystick in both directions, we increased the target distance from 6.7° (~6 mm) to 11.1° (~10 mm). In step 2, we decreased the response time to 10 sec until they reached the targets during the 10 sec response period in more than 80% of trials. In both steps 1 and 2, the joystick was mechanically fixed by the electromagnet until the auditory go cue.

In step 3, to prevent mice from pushing or leaning on the joystick before the go cue, we released the joystick from electromagnet immobilization simultaneously with

visual stimulus onset, and rewarded the mice only if they moved the joystick during the response period (i.e. withheld movements until after the go cue) and reached the correct target. Trials in which mice responded before the go cue were considered errors and immediately aborted, resulting in a white noise error sound. Step 3 training continued until mice achieved withholding performance above 80%:

$$\text{Withholding performance} = \frac{\# \text{ of responding trials after go cue}}{\# \text{ of all responding trials}}$$

In step 4, mice were trained to discriminate between the two distinct visual stimuli (forward and downward drifting gratings) and reach the correct target after the go cue. In this step, trials were considered errors and immediately aborted if mice reached the incorrect target, or moved before the go cue (as in Step 3). Step 4 continued until they achieved both withholding and discrimination performance above 80%:

$$\text{Discrimination performance} = \frac{\# \text{ of trials hitting the correct target}}{\# \text{ of trials hitting any target}}$$

Discrimination performance was computed for all trials that reached a target regardless of whether or not the trials were successfully withheld. Once this performance criterion was achieved, the ITI length was gradually increased to 4 or 8 sec (step 5). In step 6, we turned off the visual stimulus during the response period (i.e. visual stimulus was turned off simultaneously with the go cue). In step 7, the stimulus period was shortened to 1.8 sec and a 0.2 sec memory period was introduced. In the final step, the visual stimulus period was gradually decreased to 1 sec and the memory period was gradually increased to 2 sec. With the 2-sec memory period, the discrimination performance rarely improved above 60% (even after prolonged training). Thus, we trained each mouse until their discrimination performance in the 2-sec memory task reached 60% on average.

A subset of mice performed sessions containing randomly interleaved non-memory and memory trials (both with a 3-sec pre-movement period between the stimulus onset and the go cue; in non-memory trials, the visual stimulus stayed on until a target was reached) during training. In those sessions, the discrimination accuracy was consistently lower in memory than non-memory trials, indicating that the memory load, rather than the sensory discriminability, impaired performance in the memory task (Figure 3.1C). Further supporting the negative effect of memory load, performance declined with the length of memory period from non-memory, to 1-sec memory, to 2-sec, performed in separate sessions (Figure 3.9D).

Behavioral model

In our behavioral model, the choice on a given trial is predicted by a weighted sum of the current stimulus, the history of past trial outcome, choice, and their interaction, and a constant (equation 1 and 2). Past trials were temporally discounted in an exponentially decaying manner (i.e. stronger effect from more recent trials) with time constants fit independently for each history variable. Stimulus, outcome, and choice were all binary variables with the value of 1 or -1. However, in trials in which mice did not reach a target, choice was zero and outcome was -1 (error).

We repeated the following procedure for a fixed set of time constants (varying from 0.01 to 100 for each history variable), and selected the time constants and weights that produced the highest model accuracy as the best-fit regression parameters. For given time constants, we found best weights using logistic regression on a training set (equation 1), and then estimated the choice sequence in a designated test set using the best weights (equation 2). The two-step process was 10-fold cross-validated. That is, trials within a session were divided in 10 non-overlapping parts, where each part served as a test set once, and the other nine parts as a training set. The fit of the model (or

simply, model accuracy) was measured as the fraction of test trials in which the estimated choice matched the actual choice.

$$\widehat{choice}(N) = \begin{cases} 1, & \text{if } p > 0.5 \\ -1, & \text{otherwise} \end{cases} \quad (2)$$

In partial models, we used a subset of variables and performed the same regression procedure. For instance, when estimating the effect of inactivation on trial-history dependency of choice, we used a partial model without the stimulus term and compared the partial model accuracy between light-on and light-off trials.

To assess the statistical significance of history information in predicting future choices, we applied a likelihood ratio test between the full model and a partial model that contains only stimulus and constant terms. We used $p < 0.05$ as a significance threshold.

To determine whether the specific history-dependent strategy of a given session was adaptive or not (Figure 3.2.F), we generated a sequence of choices following the estimated history model (i.e., the partial model without the stimulus term) and the same stimulus rules described earlier. After simulating 100 sequences, if the fraction of correctly discriminating trials was greater than 0.5 in more than 95% of the iterations, the strategy was classified as adaptive. If the fraction of correctly discriminating trials was less than 0.5 in more than 95% of the iterations, it was maladaptive. In the other cases, the strategy was neutral.

Imaging neural activity

After mice reached the discrimination threshold of 60% in the 2-sec memory task, we paused training and allowed unlimited water access at least for 2 days prior to craniotomy and virus injections. The craniotomy spanned both the posterior parietal cortex (PPC; stereotaxic coordinates relative to bregma: 1.7 mm lateral, 2.0 mm

posterior) and the forelimb region of the primary motor cortex (M1; stereotaxic coordinates relative to bregma: 1.5 mm lateral, 0.3 mm anterior) in the right hemisphere. Viruses (AAV2-1-hSyn-GCaMP6f diluted in saline 1:7, or AAV2-9-hSyn-GCaMP6f diluted in saline 1:7; UPenn Vector Core) were injected at 5 sites (~20 nL per site) in PPC and M1 at a depth of approximately 250 μm beneath the dura, in layer 2-3. After the injections, the craniotomy (~2 mm $\times$ 3.5 mm) was covered with an optical window fixed in place with dental cement. Two of the three mice in the free-joystick task condition (Figure 3.10) were generated by crossing Camk2-tTA and tetO-CGCaMP6s and received the same procedures for craniotomy and optical window implant without virus injections.

Following surgery and recovery (14-35 days after the surgery), we imaged cortical activity in layer 2-3 at the depth of ~200 μm with excitation at 925 nm from a Ti-Sa laser (Spectra-physics) using a two-photon microscope (B-scope, Thorlabs). Each imaging field was 512 $\times$ 512 pixels covering 472 $\times$ 508 μm and imaging was performed at ~28.4 Hz. The duration of each behavior-imaging session limited to 1.5 hours ended when the mouse was disengaged from the task, or completed 170 rewarded trials. Mice completed ~135 (range: 88-172) rewarded trials in each imaging session.

For each mouse, 1-4 different imaging fields were studied within PPC (one field per session). For some mice (N=4), the same fields were imaged repeatedly over 4-7 sessions. Of the repeatedly imaged fields, except for the analysis tracking selectivity for immediately preceding trial outcome and choice across sessions (Figure 3.6), we included only the session with the most number of active cells.

Selection of task related cells

Active cells: To detect calcium transients, we used a zero-mean dF/F trace in which the mean dF/F was subtracted from the original dF/F. Using Matlab function

findpeaks, we first identified tentative transient peaks. If the amplitude of a detected peak was at least 0.5 and greater than 3.3 times the standard deviation of dF/F velocity per frame, the peak was counted as a calcium transient. To focus our analysis on stable and reliable cells, we only included the cells that showed calcium transients at a rate greater than 1 transient/minute in both the first and second half of a session and the average peak amplitude of all transients is greater than 5 times the standard deviation of dF/F. By these criteria, the average number of analyzed cells (or, active cells) in a single PPC field was 73 (range: 22-140).

Task-related cells: Of the active cells, we identified task-related cells that showed significant activity modulation during the task as following. The mean activity trace of each neuron was calculated by aligning dF/F traces to behavioral events and averaging across all correct trials. Three different behavioral events were used to align dF/F traces: stimulus onset (-6 to 3 s), movement onset (-2 to 7 s), and reward onset (-4 to 5 s). A cell was considered to be task-related if its mean activity fell outside the 99.9th percentile of its dF/F distribution in three consecutive frames in any of the three alignments. For this criterion, the false positive rate estimated on temporally-shifted dF/F traces, by a random amount for each trial, was 4.4%.

PPC inactivation experiment

Mice for PPC inactivation experiments (PV-Cre; N=8) were implanted with a head-fixation bar and bilaterally injected with virus carrying Channelrhodopsin-2 (AAV2-1-EF1A-DIO-hChR2, undiluted; UPenn Vector Core) through a thinned skull over PPC. Approximately 100 nL of virus was injected in one location at each of two depths, 200 μ m and 600 μ m from the dura. After the surgery, following the same training protocol as the head-plate implanted mice described above, we trained them to perform the task over a period of 2-4 months.

Once a mouse reached the 60% discrimination criterion, we inspected the previously thinned-skull area and performed re-thinning if necessary. Then, we conducted 1-7 light acclimation sessions to minimize non-specific light effects on behaviors. In the acclimation sessions, bifurcated blue LED fibers (470 nm, 11-20 mW for inactivation in each fiber, Doric) were placed ~2 mm above the head-fixation bar, away from the cortical region expressing Channelrhodopsin, and lights were turned on during the ITI of randomly selected 15% of trials. Most mice recovered their previous task performance within 1-2 days.

Each inactivation experiment was performed across 14 - 16 daily sessions. Control and inactivation sessions alternated day-by-day for all but 5 mice. The 5 mice (3 trial inactivation and 2 ITI inactivation) performed control and perturbation sessions sequentially in 7-day blocks. In control sessions, the LED lights were directed above the head-fixation bar, whereas in perturbation sessions they were placed directly above PPC on both hemispheres. Except for this difference, all procedures were identical between control and inactivation sessions.

In both control and inactivation sessions, light-on trials were pseudo-randomly selected with a constraint that there be at least 5 light-off trials between any two adjacent light-on trials to avoid potential behavioral adaptation to cortical perturbation due to consecutive and/or frequent exposures. Under this restriction, light stimulation was applied to approximately 15% of trials.

M1 inactivation experiment

The procedures were identical to the PPC inactivation experiment described above except for the following difference. In three of the seven mice, ChR2 was expressed in PV positive neurons by crossing PV-Cre mice with Ai32 mice containing lox-stop-lox-ChR2 in the ROSA locus, and an optical window was placed over M1. We

did not observe behavioral differences during the inactivation experiment between these three mice and the rest.

Trial selection

In behavioral model analyses, choice was predicted only for trials in which mice reached any of the two targets after the stimulus onset, ~248 trials (range: 137-344; 91%) per session. In neural data analyses, we included trials in which mice reached any target within 1 sec after the go cue. Error trials in which mice moved the joystick before the go cue (~15% of trials) were excluded due to the possibility that neural activity during stimulus and memory period in those trials might be contaminated with immediate movement planning and execution. Slow trials (target acquisition taking longer than 1sec from the go cue; ~16%) were also excluded to reduce neural variability associated with highly dissimilar movement kinematics within the same categorical choice. By these criteria, ~185 trials per session (range: 92-293) were analyzed. The early and late trials excluded from the neural analysis showed similar choice tuning to the regular trials (Figure 3.12).

Neural activity analysis

Single cell activity: Using custom Matlab program, fluorescence images were aligned frame by frame to compensate for lateral motions. Regions of interest (ROIs) were manually drawn on the motion-corrected fluorescence images, by circumscribing the cell bodies based on their GCaMP fluorescence intensity distinguishable from the background. Pixels inside each ROI were considered as a single cell, whereas pixels extending radially outward from the cell boundary by 2-6 pixels were considered background. In case the background included other cells' ROIs, those pixels were excluded. To estimate the activity of a single cell, 70% of the average pixel intensity in its background was subtracted from the average pixel intensity inside the cell(Chen et al.,

2013). The time series of the background-adjusted intensity was transformed to dF/F by dynamically estimating the baseline intensity (i.e., the 8th percentile of the intensity distribution in the 20 second window centered at each time point)(Harvey et al., 2012). For GCaMP6s signals recorded to compare ITI tuning between free-joystick and fixed-joystick conditions, dF/F was further transformed into an estimate of spike rates using the spike-triggered mixture model (<https://github.com/lucastheis/c2s>)(Theis et al., 2016).

Choice selectivity in trial epochs: For the task-related cells, their choice selectivity was examined in 9 non-overlapping 1-sec epochs (Figure 3.3): the first 6 epochs aligned to stimulus onset (-4 to 2 s), and the latter 3 epochs aligned to movement onset (-1.5 to 1.5 s). The ITI started ~1.3 s after movement onset on average. To obtain both selectivity strength and significance, we performed receiver operating characteristic (ROC) analysis on the time-averaged activity in each epoch, using the binary choice as label and the activity as score. For a given area under the ROC curve (AUROCC), double the distance from 0.5 (i.e., $2 \times |AUROCC - 0.5|$) was taken as the selectivity strength. For a significance test, we used the 99.9th percentile of the null distribution of selectivity strength ($p < 0.01$ with Bonferroni correction for multiple comparisons) estimated by choice label shuffling per cell and epoch, 1000 times. The preferred directions of choice selective neurons were nearly equally distributed (Figure 3.15).

Selectivity for other behavioral variables: Selectivity or tuning for other behavioral variables (e.g. N-1 trial outcome and choice) was computed in the same way as choice selectivity, but with those binary variables as the label or score (Figure 3.5D and 3.6).

Fitting trial-by-trial internal bias with PPC ITI activity: We used a linear regression to fit the trial-by-trial fluctuation of internal bias with the trial-by-trial ITI population activity, following 10-fold cross-validation method (Figure 3.4 and 3.5). To avoid overfitting, features that significantly contribute to the linear regression of the internal

bias were selected using Matlab function *stepwisefit* on all trials, before applying linear regression. The fit was measured as r^2 achieved on test sets.

Linear classifier: To compute the prediction power of the neural activity on binary behavioral variables such as (N-1) trial outcome, (N-1) choice, and N choice, we computed the classifying accuracy of a linear classifier (Figure 3.5B). The weights and constant of each classifier were estimated by a logistic regression represented in equation 3 on a training set, and its performance was evaluated on a test set using equation 4, following the standard 10-fold cross-validation method. The classifier performance was defined as the fraction of test trials in which the prediction matched the actual variable. ITI population neural activity is high dimensional (dimension = N cells $\times$ 4 ITI epochs/cell = 4N). Thus to avoid overfitting, only the features (i.e., selective epochs of selective cells) that significantly contributed to the regression were selected using Matlab function *stepwisefit* using all trials before applying the classification analysis.

$$\log \frac{\text{probability}\{\text{behavior}=1\}}{1-\text{probability}\{\text{behavior}=1\}} = w \cdot \text{predictor}(N) + \text{const} \quad (3)$$

$$\widehat{\text{behavior}}(N) = \begin{cases} 1, & \text{if } p > 0.5 \\ -1, & \text{otherwise} \end{cases} \quad (4)$$

Neural distance from the decision boundary: We computed the signed Euclidian distance of the ITI population activity from the linear decision boundary of the N trial choice classifier for each trial (Figure 3.4D). That is, the distance for the activity on one side of the boundary was positive, and the other side negative. Given the strong correlation between the PPC ITI activity and the internal bias estimated from our behavior model, the neural distance serves as a proxy for the strength of internal bias.

Decoding N-1 trial history information independent of internal bias: To estimate the amount of history information independent of internal bias in the population ITI activity, we built two classifiers. The first classifier decoded N-1 outcome information

from the population ITI activity. The second classifier decoded N-1 outcome from the internal bias related activity (i.e., the weighted sum of the population ITI activity that best fit the internal bias estimated from our model). Then, to compute the outcome information independent of the internal bias, we subtracted the accuracy of the second classifier from the first (Figure 3.5B). The independent N-1 choice information was computed similarly.

Choice selectivity independent of immediately preceding trial outcome and choice information: Because the immediately preceding (N-1) trial outcome and choice information had predictable power for the upcoming trial choice, we inspected whether the ITI activity remains choice selective even for the trials that followed the same outcome and choice conditions in the N-1 trial. Thus, the ITI choice selectivity was examined in each of the four possible conditions of the N-1 trial: 1) post-reward and post-downward, 2) post-reward and post-forward, 3) post-error and post-downward, and 4) post-error and post-forward (Figure 3.5D). The ITI activity would not be N choice selective in any of these 4 conditions if its activity purely encoded N-1 outcome or choice because the N-1 outcome and choice were the same within each condition. To ensure statistical power, we examined the choice selectivity of cells only for the conditions with at least 18 trials per choice direction (118 neurons with one condition and 65 neurons with two conditions). In each of these conditions, choice selectivity of each cell was assessed using ROC analysis. The fraction of conditions in which the cells are choice selective was compared to the null distribution estimated by shuffling choice labels in each condition.

History information tuning independent of N trial choice: Similarly to the choice selectivity independent of history information described above, tuning for N-1 trial outcome was assessed in 4 different conditions, in which N-1 choice and N choice were

fixed. Likewise, tuning for N-1 choice was assessed in 4 conditions in which N-1 outcome and N choice were fixed.

Movement analysis

Movement onset was defined as the first time at which the joystick velocity exceeded $22.2^\circ/\text{sec}$ ($\sim 20 \text{ mm/sec}$) continuously for 20 msec and the joystick moved at least 1.3° ($\sim 1.1 \text{ mm}$) from the origin. The reaction time was measured as the time from the go cue and movement onset, and the movement duration was measured as the time from movement onset to the time when the joystick entered any target region (Figure 3.14).

Statistical Analysis

Throughout the paper, we performed non-parametric tests to avoid normality assumption. When simply assessing whether the medians of paired samples acquired from the same subjects are different, we used Wilcoxon signed rank test. When testing for unpaired samples acquired from different subjects, we used Wilcoxon rank sum test. When testing a specific hypothesis that the median of one set of samples is greater (or smaller) than the median of the other set acquired from the same subjects, we used Wilcoxon one-sided signed rank test. For two sets of samples acquired from different subjects, we used Wilcoxon one-sided rank-sum test for testing specific hypotheses. For statistical tests for means, we used bootstrapping.




Figure 3.1 Task.

(A) Top: task schematic. Middle: task trial structure. Bottom: stimulus-action-outcome rule.

(B) Fraction of correctly discriminating trials is significantly greater than chance. Black, mean $\pm$ s.e.m. across mice; grey, individual mice. Wilcoxon one-sided signed rank test.

(C) Fraction of correctly discriminating trials is significantly greater in non-memory trials (visual stimulus stays on throughout the trial) when randomly interleaved with memory trials. Black, mean $\pm$ s.e.m. across mice; grey, individual mice. Wilcoxon one-sided signed rank test.

Figure 3.2 Choice-outcome history biases future decision, driving choice variability.

(A) Top: example trial-by-trial sequence of an animal's choice (black) with partial model fit using stimulus information only (cyan). Bottom: choice sequence from top plot (black) with full model fit using stimulus, trial history, and a constant (green).


(B) Full model including both stimulus and history information predicts choice more accurately than partial model with stimulus only. Black, mean $\pm$ s.e.m. across sessions; grey, individual sessions. Wilcoxon one-sided signed rank test.

(C) Partial model accuracy using only one history variable (outcome, choice, or outcome-choice interaction) at a time, indicating that outcome history and choice history have larger contributions to model accuracy. Black, mean $\pm$ s.e.m. across sessions; grey, individual sessions.

(D) Time constant of each history variable. Sessions are included only if the corresponding partial model (with one history variable and a constant) is significantly better than the model with only a constant (likelihood-ratio test between the partial and constant models, $p < 0.05$). Black, mean $\pm$ s.e.m. across sessions; grey, individual sessions.

(E) Choice prediction accuracy is highest when models are estimated from the current session than different sessions of the same animal, or different animals, indicating that history-dependent strategies vary across sessions and animals. Black, mean $\pm$ s.e.m. across sessions; grey, individual sessions. Wilcoxon one-sided signed rank test.

(F) History-dependent strategies are classified as maladaptive (cyan), neutral (blue), or adaptive (magenta), depending on the fraction of correctly discriminating trials achieved by the estimated history models (using only history but not stimulus) relative to chance ($p < 0.05$, Experimental Procedures). Sessions with adaptive strategies tend to be associated with better behavioral performance (Pearson correlation coefficient = 0.64, $p < 0.005$).






Figure 3.3 Pre-stimulus activity during the inter-trial interval in posterior parietal cortex (PPC) predicts the upcoming choice.

(A) Two-photon calcium image of PPC neurons. Four example neurons shown in **B** are demarcated with colored lines.

(B) Activity of the four example neurons from a single imaging session shown in **A** (mean $\pm$ s.e.m.). Each neuron is choice selective during the inter-trial interval (ITI, Cell 1), stimulus period (Cell 2), memory period (Cell 3), or movement period (Cell 4).

(C) Trial-average activity of PPC choice-selective neurons for preferred choice trials. Each row represents the normalized activity of a single neuron sorted by its peak activity timing.

(D) Example PPC neurons that are choice selective during the ITI.

(E) Fraction of neurons (mean $\pm$ s.e.m.) with significant choice selectivity ($p < 0.01$ with Bonferroni correction) in 9 non-overlapping 1-sec epochs out of all choice-selective neurons.




Figure 3.4 Pre-stimulus choice-selectivity in PPC encodes a mixture of history and bias information.

(A) Break-down of choice-selective cells. The majority of cells that show significant choice-selectivity during ITI ('ITI choice cells') are different from those that show maximal choice selectivity during peri-movement period ('Move cells', peri-movement period defined as -0.5 to 1.5 sec around movement onset).

(B) Neurons are tuned to the previous trial (N-1) choice and/or outcome. Tuning was measured as the area under the ROC curve, ranging from 0 to 1. The value of 0.5 indicates no tuning. Dots are all choice-selective neurons (N=652). Red dots are neurons with stronger activity during the ITI following reward (N=146; 'Reward cells'), and pink dots are neurons with stronger activity following errors (N=124; 'Error cells'). Blue dots are neurons with stronger activity following forward choice (N=63; 'Forward cells'), cyan dots are neurons with stronger activity following downward choice (N=72; 'Downward cells'), and grey dots are choice-selective neurons that do not show significant tuning to N-1 choice or outcome.

(C) Trial-average activity of Reward, Error, Downward, Forward, and Move cells in their preferred choice direction trials under 4 different N-1 trial conditions. Each row represents the normalized activity of a single neuron sorted by its peak activity timing.

(D) Previous trial (N-1) outcome, previous trial (N-1) choice, and upcoming trial (N) choice tuning of 3 example PPC neurons. Tuning for each variable is assessed in 4 different conditions in each of which the other two variables are fixed as indicated in the tables (top). In each condition, the mean traces of dF/F are plotted for two values of the variable of interest (red vs. pink, blue vs. cyan, or green vs. lime). PPC cells are influenced by all three variables in complex ways (e.g., cell 1 encodes N-1 outcome predominantly, yet N-1 choice and N choice also modulate its ITI activity). Grey shaded region indicates the ITI, and the triangle indicates the average movement onset time.

(E) Fraction of cells whose ITI activity is significantly tuned for each of (N-1) outcome, (N-1) choice, or N choice, independent of the other two.




Figure 3.5 History tuning of PPC neurons reflects the history dependency of bias.

(A) Top: a single neuron imaged in 5 separate sessions, sorted in the descending order of the (N-1) trial outcome tuning strength. Bottom: pre-stimulus activity (mean $\pm$ s.e.m.; -6 to 0 s from stimulus onset) of the example neuron following reward and error trials. Each pair of numbers represents the tuning strength of the neuron for (N-1) trial outcome and outcome dependency of the internal bias (estimated as the accuracy of the partial model using only outcome information), respectively. Note that the outcome tuning is stronger in sessions in which the past outcome information has stronger contributions to the bias.

(B) Between-session differences in (N-1) outcome tuning of the ITI activity of single neurons plotted against between-session differences in choice dependency on outcome information (Pearson correlation coefficient=0.14, $p<7.1e-6$). Neurons with significant outcome tuning in at least one of the two sessions are included.

(C) The same as B but for the (N-1) trial choice.

Figure 3.6 Pre-stimulus population activity in PPC reflects the history-dependent bias and previous trial history.

- (A) Sequences of internal bias (thin black lines, estimated from the behavioral model) and weighted sum of ITI activity (thick yellow lines) across simultaneously imaged PPC neurons in three example mice. The weights for individual neurons were determined using a linear regression between the ITI activity and the internal bias.
- (B) Goodness of fit of PPC ITI activity to the internal bias across (from left to right): all trials, only forward choice trials, only downward trials, and shuffled trials. Only the sessions in which the history-dependent bias is significant (likelihood-ratio test between the full model and a partial model that contains only stimulus and constant terms, $p < 0.05$, 12 sessions; Experimental Procedures) are included. Black, mean $\pm$ s.e.m. across sessions; grey, individual sessions. Wilcoxon one-sided signed rank test.
- (C) Data from a single session showing the weighted sum of PPC ITI activity across simultaneously imaged neurons against the internal bias estimated from the behavioral model. Each dot represents a single trial.
- (D) Fraction of forward choice as a function of the internal bias estimated from the ITI activity of simultaneously imaged PPC neurons based on a linear regression (mean $\pm$ s.e.m. across 12 sessions). Black, forward stimulus trials; red, downward stimulus trials. The estimated internal bias in each session was normalized such that it represents the signed distance from the decision boundary (Methods) that divides the forward and downward choice trials. To achieve the same fraction of forward choice, the ITI activity needs to be further towards the forward bias domain in downward stimulus trials than in forward stimulus trials, indicating that the final choice is made by integrating bias and stimulus.
- (E) Example population ITI activity from one session projected on a 2-dimensional plane. Left: the x-axis shows the activity encoding internal bias (i.e., estimated internal bias from the population activity, as in a) and the y-axis shows the activity encoding N-1 trial outcome (i.e., estimated N-1 trial outcome from the population activity). Each dot represents a single trial. Population activity in trials with similar internal bias still encodes N-1 trial outcome information (reward or error). Right: x-axis same as the left and y-axis shows the activity encoding N-1 trial choice.
- (F) History information independent of internal bias in the ITI activity, computed as (prediction accuracy of history information decoded from the ITI activity – prediction accuracy of history information decoded from the ITI activity projected to the internal bias axis). Each dot represents a single imaging session in which the behavioral model predicts the choice sequence significantly ($p < 0.001$). Black, mean $\pm$ s.e.m. across mice. Wilcoxon one-sided signed rank test.




Figure 3.7 Inactivating pre-stimulus activity in PPC alters internal bias.

(A) Schematic of inactivation experiment. Control (blue light directed away from PPC) and inactivation (the light directed to PPC) sessions alternated day-to-day (for 14-16 days). Continuous blue light was applied during the ITI in randomly selected trials (15%; light-on trials) in both control and inactivation sessions.

(B) Choice sequence (black) and behavioral model fit (orange) in an example inactivation session. In this example, the mouse tended to alternate choice (i.e., the mouse most heavily weighted the previous choice history) in light-off trials, but this tendency was reduced in light-on trials.


(C) The effect of PPC ITI inactivation on the model fit in 7 separate inactivation sessions in a single mouse. Black, mean $\pm$ s.e.m. across sessions; grey, individual sessions. Wilcoxon signed rank test. The light-off model was built on a subset of light-off trials and its accuracy was assessed on the remaining light-off or light-on trials.

(D) Average light-off model accuracy in light-off and light-on trials in inactivation sessions. Black, mean $\pm$ s.e.m. across mice; grey, individual mice. Wilcoxon one-sided signed rank test.

(E) In trial inactivation sessions, blue light was applied from stimulus onset to the end of randomly selected trials (15%).

(F) The same as D, but for trial inactivation in PPC.

(G) The same as D, but for ITI inactivation in M1.






Figure 3.8 Inactivating pre-stimulus activity in PPC weakens the history dependency of choice.

(A) The light-on model accuracy versus light-off model accuracy. Light-on model was built by fitting the full model on light-on trials. Light-off model was built on randomly selected light-off trials matching the number of light-on trials. The x-axis represents the mean of 100 light-off models in each session, and the y-axis represents the light-on model for the corresponding session. The light-on model accuracy is significantly lower than the light-off model accuracy (Wilcoxon signed rank test, $p < 0.02$).

(B) Changes in the magnitude of model weights by PPC ITI inactivation. The weight magnitude of light-on model was translated into a z-score relative to the distribution of weight magnitudes of light-off models in each session. The weight magnitude for outcome history, choice history, and outcome-choice interaction history significantly decreased by PPC ITI inactivation (bootstrap, $p < 0.001$, $p < 0.001$, and $p < 0.05$ respectively). Black, mean $\pm$ s.d. across sessions.

(C) Change in behavioral performance during ITI inactivation (light-on – light-off) as a function of the degree of adaptiveness of the strategies (cyan: maladaptive, blue: neutral, and magenta: adaptive strategy sessions, respectively). The negative correlation between the performance change and the degree of adaptiveness indicates that ITI perturbation improves performance in sessions with maladaptive strategies, whereas it deteriorates performance in sessions with adaptive strategies. Black, linear regression.




Figure 3.9 Two-alternative forced-choice tasks using the joystick apparatus.

(A) The two target areas in the joystick endpoint space.
 (B) Behavioral training procedure.
 (C) Discrimination performance in random stimulus versus deterministic stimulus trials. Each pair of connected circles represents the average of approximately 20 sessions performed by a single mouse, and the black thick line represents mean $\pm$ s.e.m. across 14 mice. Wilcoxon one-sided signed rank test.




Figure 3.10 Stimulus dependency increases, while history dependency decreases in non-memory task compared to memory task.

(A) Discrimination performance in the non-memory (task described in step 5) and 2-sec memory trials. Each task was performed in separate sessions. Each circle represents a single session, and the black thick line represents mean $\pm$ s.e.m. Wilcoxon one-sided rank sum test.

(B) The magnitude of weights in the full model in the non-memory versus memory task (stimulus, outcome history, choice history, outcome-choice interaction history from left to right). The magnitude of stimulus weights are significantly larger for the non-memory than memory task. The magnitude of outcome history weight is significantly smaller for the non-memory than memory-task. Wilcoxon one-sided rank sum test.




Figure 3.11 ITI choice selectivity in PPC is not a reflection of ongoing limb movements.

(A) Nine example neurons that are choice selective during the ITI (yellow shade) in the original condition in which the joystick was fixed during the ITI. The average activity for forward (black) versus downward (red) choice trials in the fixed-joystick (solid) and free-joystick conditions (dashed). Fixed- and free-joystick conditions were randomly interleaved, and approximately 15% were free-joystick trials. The four dotted vertical lines mark stimulus onset, offset, movement onset, and the next trial ITI onset.

(B) The choice tuning strength during the ITI in the fixed- versus free-joystick condition for 434 cells that are significantly choice selective during the ITI in the fixed-joystick condition (Mean $\pm$ s.e.m.). The black dashed line represents the unity line.

(C) The same as b, but for 155 cells that are significantly choice selective during the ITI in the free-joystick condition (Mean $\pm$ s.e.m.). The black dashed line represents the unity line.




Figure 3.12 ITI choice tuning persists in early and late execution trials.

(A) The activity in preferred (red) versus non-preferred (blue) trials for example cells 1-2, and the mean activity of all choice cells in regular, early, and late trials.

(B) ITI tuning of regular versus early trials (Mean $\pm$ s.e.m.). The dashed line represents the unity line.

(C) ITI tuning of regular versus late trials.




Figure 3.13 The preferred directions of choice selectivity are equally distributed.

(A) Fractions of cells preferring forward versus downward choice across all choice-selective PPC cells in 11 non-overlapping 1-sec epochs.

(B) Fraction of cells preferring forward choice in individual imaging fields with at least 15 choice-selective cells in a given epoch. Dots are individual fields and the thick black line indicates mean $\pm$ s.e.m.




Figure 3.14 PPC ITI inactivation effect on internal model accuracy, task performance, reaction time, and movement time.

(A) Inactivation sessions. Each circle represents a single session. The x-axis is the mean of light-off trials in a given session, and the y-axis is the mean of light-on trials. The black lines represent the unity line. Wilcoxon signed rank test for difference between the light-off and light-on trials.

(B) Control sessions.




Figure 3.15 Control sessions.

(A-C) Average light-off model accuracy in light-off and light-on trials in control sessions. Black, mean $\pm$ s.e.m. across mice; colors, individual mice. Wilcoxon one-sided signed rank test. Compare these to **Figure 3.7D,F, and G**. There is no significant non-specific effects of blue lights.

(D) The light-on model accuracy versus light-off model accuracy in control sessions. PPC ITI and M1 ITI control sessions are combined ($N=106$). The light-on model accuracy is not significantly different from the mean light-off model accuracy in control sessions (Wilcoxon signed rank test, $p=0.08$). Compare this plot to **Figure 3.8A**.

(E) Changes in the magnitude of model weights in ITI control sessions. No weight shows significant change by blue light (bootstrap, $p>0.05$). Compare this plot to **Figure 3.8B**.

(F) Change in discrimination performance by blue light during ITI (light-on – light-off) as a function of the degree of adaptiveness of the strategies (cyan: maladaptive, blue: neutral, and magenta: adaptive strategy sessions, respectively). Compare this plot to **Figure 3.8C**.

Chapter 3 currently being prepared for submission and will include Dr. EunJung Hwang, Madan Mukundan and Dr. Takaki Komiyama as authors. This material was included with the generous consent of all authors. The dissertation author is a primary researcher on this work.

Chapter 4. General Discussion

This dissertation investigated two different cortical brain areas, the primary motor cortex (M1) as well as the posterior parietal cortex (PPC), and their role in learning, execution and bias of movements. First, we examined how the dependency on M1 for the initiation and execution changes with learning. To investigate this question we developed a novel two-direction forelimb movement task. In our behavioral task mice were required to manipulate a joystick in response to visual stimuli to acquire a water reward. Achieving expert performance required several months of daily training over three blocks of training that gradually shaped the behavioral task structure, progressing from a simple associative learning structure (e.g. move the joystick to receive a water reward) to a complex discrimination task. The complex nature of this behavioral task allowed us to make several observations, and identify several potential avenues for future work.

First, we demonstrated using optogenetic activation of inhibitory interneurons that M1 is necessary for the initiation and controlled execution of movements during the early stages of learning. During this early stage of learning many different behavioral changes are occurring simultaneously. For example, mice are making an association between joystick movement and water reward, as well as exploring the movement space of the joystick to identify movement trajectories that will result in a reward. This raises an important question as to whether, in our task, M1 is required for the initiation and execution of the movements, or rather M1 is simply necessary for the initial associative learning. In other words, by silencing M1 are mice simply not capable of integrating the sensory cues towards the initiation of movements, or is M1 required for the online controlled execution throughout the movements themselves? Interestingly, during M1 inactivation early in learning we observed significant impairment in the ability of mice to initiate movements

as well as an inability of mice to control their forelimb. This evidence suggests that M1 is necessary for both the initiation and controlled execution of movements early in learning. Additional evidence from our lab has indicated similar results demonstrating mice are capable of initiating movements (though severely impaired), but the movements themselves appear erratic and uncontrolled (Peters et al., 2014). However, as mentioned previously, at the earliest stages of learning of the behavioral task presented here, both associative learning as well as motor learning are occurring simultaneously. Therefore to definitively answer this question using our behavioral task would require segregating the associative learning phase from the motor learning phase of behavioral training, which may not be possible with this task design. Perhaps the ideal experiment would be performing inactivation throughout the first block of learning to observe if the dependency of movement initiation and execution changes as mice master the sensorimotor associations. Our model suggests that movements early during training depend on M1 for initiation and execution, therefore we would expect that throughout this initial phase of learning (or until the emergence of a stereotyped movement, see below) movements would be dependent on M1.

Following this initial phase of learning, during which movements are dependent on M1 for initiation and execution, we continued training mice to achieve high levels of performance on our two direction movement task. Interestingly, during training mice developed a bias for one target over the other. Further, the magnitude of the bias was correlated with the difference in stereotypy. Using optogenetic silencing of M1, identical to those experiments carried out early in learning, we observed that rather than impairing both movements (as was observed during M1 inactivation early in learning), we observed that M1 inactivation in expert mice impaired only the non-biased and less stereotyped movement, and had no effect on biased and stereotyped movement. These

results suggest that at this stage in learning the biased and stereotyped movement is no longer dependent on M1. This is surprising given the results observed from M1 inactivation early in learning, as well as previous work (Guo et al., 2015; Peters et al., 2014). We hypothesize that both the significant duration of training the mice undergo and the bias for one direction over another leads one movement (the biased movement) being practiced beyond the point of stereotypy, or over-trained. Once the movement becomes highly stereotyped, our model proposes that the movement is then transferred from being explicitly controlled by M1 to being controlled by subcortical movement areas.

Previously it was believed that movements were segregated into one of two categories, dexterous movements that require M1 for initiation and execution (Castro, 1972; Guo et al., 2015), and non-dexterous movements that can be executed independent of M1 (Kawai et al., 2015). As noted, our model hypothesizes that movements initially dependent on M1 for initiation and execution can become independent of M1 following stereotypy and over-training. At face value this is an attractive model for the brain: pushing stable and stereotyped movements from the highly plastic motor cortex to a (theoretically) more stable system. However, this begs the question as to whether or not all movements that are initially dependent on M1 can become independent following sufficient training.

It is the opinion of the author of this dissertation that indeed dexterous movements can become independent from M1, though these dexterous movements must be sufficiently constrained such that they require no modifications between movement trials. For example, in a food pellet retrieval task, small fluctuations in the starting position of the hand and shoulder require significant modifications to the movement program, as small changes in starting position drastically changes how the digits of the forelimb will grasp the food pellet. By placing the animal in a more

constrained environment, in a head-fixed position and requiring the forelimb to begin its movement from the same starting position, the movements will require neither online control nor appraisal of the environment prior to execution. After acquiring stereotypy and sufficient over-training, one would hypothesize that it would become independent of M1 similar to those biased movements described above.

This question can be extended to the behavioral task used in these experiments: given additional training, would the non-biased movements become independent of M1, similar to the biased movements? Following our proposed model we would expect these non-biased movements to eventually reach stereotypy, and subsequently become independent of M1 similar to the biased movements. However, it is likely that selection of a biased movement is a strategy employed by the mice to solve the behavioral task, and therefore would persist even after significant durations of training (i.e. constantly biasing movements is a maladaptive strategy adopted by the mice). Within our dataset we observed mice using exactly this strategy even after over a year of behavioral training.

In addition to demonstrating the role of M1 in movement execution at different learning stages, we recorded the activity of neurons in M1 during task execution. Briefly, we observed in expert mice that activity in M1 is more consistent and relevant for non-biased movements when compared to activity for biased movements. This result itself is intriguing, as it seems counterintuitive that more stereotyped movements have more variable activity, and less stereotyped movements have more consistent activity. One interpretation of this result is that once movements become sufficiently stereotyped and independent of M1, activity in M1 neurons represents movement irrelevant information. Alternatively, sparse activity in M1 could be sufficiently robust that it is encoding variability of movements along a dimension yet to be identified. Following this hypothesis it is possible that M1 maintains representations of these movements even after they can

be executed independent of M1. This would suggest that stereotyped movements become distributed throughout the brain rather than strictly offloaded from M1 to an individual subcortical area. Distributed coding of skilled movements might allow for more flexible utilization of these movements, under a more diverse set of environmental conditions. Should this be the case, then how would the brain select which motor pathway to utilize? In our experiments, similar to previous work (Ebbesen et al., 2017; Stoltz et al., 1999; Zagha et al., 2015), mice were unable to withhold movements when M1 was inactivated. This provides some evidence for M1 to coordinate movement execution, integrating decision and sensory information to select which pathway it will use and when it will use it to execute appropriate movements.

Regarding the coding of movements in M1, our lab has previously identified learning related changes including an initial expansion of movement-related activity followed by a refinement as movements become more stereotyped (Peters et al., 2014). These changes occur together with synaptic reorganization (Chen et al., 2015; Peters et al., 2014; Xu et al., 2009), which likely coincide with expansion (Kleim et al., 1998; Pascual-Leone et al., 1994) of behaviorally relevant regions of M1. Given our results, are these changes reflecting M1 encoding fine details of the movement, or alternatively, are these changes the result of M1 beginning the process of creating alternative circuits it will then use to offload the learned movement other brain areas? Evidence for the former (refinement of activity as stereotypy is achieved (Peters et al., 2014)) and latter (gradual encoding of movement details in subcortical networks (Costa et al., 2004)) suggest, not surprisingly, that both are intertwined throughout learning. Future experiments analyzing compartment specific spine changes (Chen et al., 2015) and projections from M1 to other movement generating areas will provide additional insight to how these processes work together to facilitate learning. Interestingly, once learned, expansions of movement

representations are followed by contractions of those same representations in M1 while preserving the fidelity of the learned movements (Molina-Luna et al., 2008; Pascual-Leone et al., 1994). Similar to those questions raised previously, could these reductions in the size of movement representations indicate the offloading of movements from M1? Understanding the details surrounding the changes in movement representation in M1 once a movement is mastered will provide an idea of the temporal timeline during which stereotyped movements are encoded subcortical areas.

Which subcortical areas are responsible for executing movements independent of M1? Both the basal ganglia and cerebellum participate in motor skill learning with differing roles (Doyon et al., 2003). Most notably, the basal ganglia has been observed to remain active following mastery suggesting a role in long-term retention of skills (Grafton et al., 1994) and is known to be interconnected with M1, making it a potential candidate for housing learned movements independent of M1. Experimentally subcortical movement centers could be identified through injection of a retrograde tracer, such as the modified rabies virus, into individual muscles. The modified rabies virus (Waller), is a viral tool that provides the ability to trace connected neurons by “jumping” across synapses in a retrograde manner. For example, the ventral root of the spinal cord carries motor neuron efferents to muscles. Thus injecting the modified rabies virus into the ventral root would infect spinal cord motor neuron, and subsequently cross the synapse to infect neurons projecting from the brain to the spinal cord. Using this technique one could identify potential candidate areas that are, at least, anatomically capable of driving movements.

Assuming this model of movement execution is true, namely that there are two different areas that can independently drive movements, what areas of the brain are responsible for selecting the actions to be executed, and subsequently activating the

appropriate motor programs? Among the many brain areas involved in action selection PPC has been implicated in the integration of history-dependent variables, such as choice and outcome, towards action selection (Constantinidis and Procyk, 2004). We investigated PPC's role in action selection by developing a novel visually-guided action selection task and combining it with in vivo two-photon calcium imaging and optogenetic perturbation. As noted, previously it has been demonstrated that the PPC contains choice-predicting activity (Platt and Glimcher, 1999). However performing manipulations to test the effects of PPC manipulation on subsequent action selection has remained challenging. Using similar optogenetic techniques as described above, we were able to perform trial-by-trial manipulations of PPC activity to observe how it influences subsequent actions. Briefly, we observed that perturbing PPC activity had variable effects on subsequent actions depending on the general strategy of the mice. For example, if mice employed maladaptive strategies, for example constantly biasing movements in one direction, inactivating PPC improved performance on the task. Alternatively, if mice employed adaptive strategies, inactivating PPC impaired performance on the task.

An avenue of future work will be to identify how action selection information from PPC is integrated and acted upon by downstream areas. For example, it is known that PPC projects to the secondary motor cortex, as well as the striatum (Yamawaki et al.). The secondary motor cortex then integrates information from PPC, sensory cortices, as well as inputs from other frontal areas (Reep et al., 1990; Zingg et al., 2014), and then subsequently sends projections to the striatum. These projection patterns suggest that the striatum may serve as a subcortical movement area responsible for M1 independent execution of stereotyped movements in the two direction forelimb based movement task described above. Interestingly, inactivation of PPC in the two direction forelimb based

movement task did not result in any behavioral effects (data not shown), likely because the previously described task has no working memory component. Nonetheless, the connectivity between PPC, M2, M1 and the striatum suggest the striatum is a potential candidate for executing movements independent of M1.

References

- Abrahamyan, A., Silva, L.L., Dakin, S.C., Carandini, M., and Gardner, J.L. (2016). Adaptable history biases in human perceptual decisions. *Proc. Natl. Acad. Sci. U. S. A.* 113, E3548-3557.
- Alaverdashvili, M., and Whishaw, I.Q. (2008). Motor cortex stroke impairs individual digit movement in skilled reaching by the rat. *Eur. J. Neurosci.* 28, 311–322.
- Allred, R.P., Adkins, D.L., Woodlee, M.T., Husbands, L.C., Maldonado, M.A., Kane, J.R., Schallert, T., and Jones, T.A. (2008). The Vermicelli Handling Test: A simple quantitative measure of dexterous forepaw function in rats. *J. Neurosci. Methods* 170, 229–244.
- Asanuma, H., and Sakata, H. (1967a). Functional organization of a cortical efferent system examined with focal depth stimulation in cats. *J. Neurophysiol.* 30, 35–54.
- Asanuma, H., and Sakata, H. (1967b). Functional organization of a cortical efferent system examined with focal depth stimulation in cats. *J. Neurophysiol.* 30, 35–54.
- Ashe, J., and Georgopoulos, A.P. Movement parameters and neural activity in motor cortex and area 5. *Cereb. Cortex* 4, 590–600.
- Barraclough, D.J., Conroy, M.L., and Lee, D. (2004). Prefrontal cortex and decision making in a mixed-strategy game. *Nat. Neurosci.* 7, 404–410.
- Berridge, K.C., and Whishaw, I.Q. (1992). Experimental Brain Research Cortex, striatum and cerebellum: control of serial order in a grooming sequence. *Exp Brain Res* 90, 275–290.
- Berryhill, M.E., and Olson, I.R. (2008). The right parietal lobe is critical for visual working memory. *Neuropsychologia* 46, 1767–1774.
- Brehm, J.W. (1956). Postdecision changes in the desirability of alternatives. *J. Abnorm. Psychol.* 52, 384–389.
- Brunton, B.W., Botvinick, M.M., and Brody, C.D. (2013). Rats and Humans Can Optimally Accumulate Evidence for Decision-Making. *Science* 340, 95–98.
- Busenitz, L.W., and Barney, J.B. (1997). Differences between entrepreneurs and managers in large organizations: Biases and heuristics in strategic decision-making. *J. Bus. Ventur.* 12, 9–30.
- Busse, L., Wade, A.R., and Carandini, M. (2009). Representation of concurrent stimuli by population activity in visual cortex. *Neuron* 64, 931–942.
- Busse, L., Ayaz, A., Dhruv, N.T., Katzner, S., Saleem, A.B., Schölvinck, M.L., Zaharia,

- A.D., and Carandini, M. (2011). The Detection of Visual Contrast in the Behaving Mouse. *J. Neurosci.* 31, 11351–11361.
- Castro, A.J. (1972). The effects of cortical ablations on tongue usage in the rat. *Brain Res.* 45, 251–253.
- Castro-Alamancos, M.A., and Borrell, J. (1995). Functional recovery of forelimb response capacity after forelimb primary motor cortex damage in the rat is due to the reorganization of adjacent areas of cortex. *Neuroscience* 68, 793–805.
- Chen, S.X., Kim, A.N., Peters, A.J., and Komiyama, T. (2015). Subtype-specific plasticity of inhibitory circuits in motor cortex during motor learning. *Nat. Neurosci.* 18, 1109–1115.
- Chen, T.-W.W., Wardill, T.J., Sun, Y., Pulver, S.R., Renninger, S.L., Baohan, A., Schreiter, E.R., Kerr, R. a, Orger, M.B., Jayaraman, V., et al. (2013). Ultrasensitive fluorescent proteins for imaging neuronal activity. *Nature* 499, 295–300.
- Churchland, A.K., Kiani, R., Chaudhuri, R., Wang, X.-J., Pouget, A., and Shadlen, M.N. (2011). Variance as a Signature of Neural Computations during Decision Making. *Neuron* 69, 818–831.
- Churchland, M.M., Yu, B.M., Cunningham, J.P., Sugrue, L.P., Cohen, M.R., Corrado, G.S., Newsome, W.T., Clark, A.M., Hosseini, P., Scott, B.B., et al. (2010). Stimulus onset quenches neural variability: a widespread cortical phenomenon. *Nat. Neurosci.* 13, 369–378.
- Churchland, M.M., Cunningham, J.P., Kaufman, M.T., Foster, J.D., Nuyujukian, P., Ryu, S.I., and Shenoy, K. V (2012). Neural population dynamics during reaching. *Nature* 487.
- Cichon, J., and Gan, W.-B. (2015). Branch-specific dendritic Ca²⁺ spikes cause persistent synaptic plasticity. *Nature* 520, 180–185.
- Coe, B., Tomihara, K., Matsuzawa, M., and Hikosaka, O. (2002). Visual and Anticipatory Bias in Three Cortical Eye Fields of the Monkey during an Adaptive Decision-Making Task. *J. Neurosci.* 22, 5081–5090.
- Constantinidis, C., and Procyk, E. (2004). The primate working memory networks. *Cogn. Affect. Behav. Neurosci.* 4, 444–465.
- Costa, R.M., Cohen, D., and Nicolelis, M.A.L. (2004). Differential corticostriatal plasticity during fast and slow motor skill learning in mice. *Curr. Biol.* 14, 1124–1134.
- Dolan, R.J., and Sharot, T. (2011). *Neuroscience of Preference and Choice: Cognitive and Neural Mechanisms* (Academic Press).

- Doya, K. (1999). What are the computations of the cerebellum, the basal ganglia and the cerebral cortex? *Neural Networks* 12, 961–974.
- Doyon, J., Penhune, V., and Ungerleider, L.G. (2003). Distinct contribution of the cortico-striatal and cortico-cerebellar systems to motor skill learning. *Neuropsychologia* 41, 252–262.
- Drolet, A. (2002). Inherent Rule Variability in Consumer Choice: Changing Rules for Change's Sake. *J. Consum. Res.* 29, 293–305.
- Dum, R.P., and Strick, P.L. (2002). Motor areas in the frontal lobe of the primate. *Physiol. Behav.* 77, 677–682.
- Ebbesen, C., Doron, G., Lenschow, C., and Brecht, M. (2017). Vibrissa motor cortex activity suppresses contralateral whisking behavior. *Nat. Neuro* 20, 82–89.
- Egan, L.C., Santos, L.R., and Bloom, P. (2007). The Origins of Cognitive Dissonance Evidence From Children and Monkeys. *Psychol. Sci.* 18, 978–983.
- Erich, J.C., Brunton, B.W., Duan, C.A., Hanks, T.D., and Brody, C.D. (2015). Distinct effects of prefrontal and parietal cortex inactivations on an accumulation of evidence task in the rat. *Elife* 4, e05457.
- Evarts, E. V (1965). Relation of Discharge Frequency To Conduction Velocity in Pyramidal Tract Neurons. *J. Neurophysiol.* 28, 216–228.
- Evarts, E. V (1966). Pyramidal tract activity associated with a conditioned hand movement in the monkey. *J. Neurophysiol.* 29, 1011–1027.
- Evarts, E. V (1968). Relation of pyramidal tract activity to force exerted during voluntary movement. *J. Neurophysiol.* 31, 14–27.
- Faisal, A.A., Selen, L.P.J., and Wolpert, D.M. (2008). Noise in the nervous system. *Nat. Rev. Neurosci.* 9, 292–303.
- Frankland, P.W., and Bontempi, B. (2005). The organization of recent and remote memories. *Nat. Rev. Neurosci.* 6, 119–130.
- Fritsch, G., and Hitzig, E. (1870). Ueber die elektrische Erregbarkeit des Grosshirns. *Arch. Für Anat. Physiol. Und Wissenschaftliche Med.* 300–332.
- Fründ, I., Wichmann, F.A., and Macke, J.H. (2014). Quantifying the effect of intertrial dependence on perceptual decisions. *J. Vis.* 14, 1–16.
- Fu, Q.G., Suarez, J.I., and Ebner, T.J. (1993). Neuronal specification of direction and distance during reaching movements in the superior precentral premotor area and primary motor cortex of monkeys. *J. Neurophysiol.* 70, 2097–2116.

- Fulton, J. (1935). A note on the definition of the “motor” and “premotor” areas. *Brain* 58, 311–316.
- Funamizu, A., Kuhn, B., and Doya, K. (2016). Neural substrate of dynamic Bayesian inference in the cerebral cortex. *Nat. Neurosci. advance on*.
- Gandolfo, F., Li, C.-S.S., Benda, B.J., Schioppa, C.P., and Bizzi, E. (2000). Cortical correlates of learning in monkeys adapting to a new dynamical environment. *Proc. Natl. Acad. Sci. U. S. A.* 97, 2259–2263.
- Georgopoulos, A.P., Kalaska, J.F., Caminiti, R., and Massey, J.T. (1982a). On the relations between the direction of two-dimensional arm movements and cell discharge in primate motor cortex. *J. Neurosci.* 2(11), 1527–1537.
- Georgopoulos, A.P., Schwartz, A.B., and Kettner, R.E. (1986). Neuronal population coding of movement direction. *Science* 233, 1416–1419.
- Georgopoulos, a P., Kalaska, J.F., Caminiti, R., and Massey, J.T. (1982b). On the relations between the direction of two-dimensional arm movements and cell discharge in primate motor cortex. *J. Neurosci.* 2(11), 1527–1537.
- Goard, M.J., Pho, G.N., Woodson, J., and Sur, M. (2016). Distinct roles of visual, parietal, and frontal motor cortices in memory-guided sensorimotor decisions. *Elife* 5, 1–30.
- Gold, J.I., Law, C.-T., Connolly, P., and Bennur, S. (2008). The relative influences of priors and sensory evidence on an oculomotor decision variable during perceptual learning. *J. Neurophysiol.* 100, 2653–2668.
- Grafton, S.T., Woods, R.P., and Tyszka, M. (1994). Functional imaging of procedural motor learning: Relating cerebral blood flow with individual subject performance. *Hum. Brain Mapp.* 1, 221–234.
- Graziano, M. (2009). *The Intelligent Movement Machine* (Oxford University Press).
- Graziano, M.S.A.M.S.A., Taylor, C.S.R.C.S.R., Moore, T., and Cooke, D.F. (2002). The Cortical Control of Movement Revisited. *Neuron* 36, 349–362.
- Greenwald, A.G., McGhee, D.E., and Schwartz, J.L.K. (1998). Measuring individual differences in implicit cognition: The implicit association test. *J. Pers. Soc. Psychol.* 74, 1464–1480.
- Gross, C.G., Gall, F.J., Jackson, J.H., and Broca, P. (2007). The Discovery of Motor Cortex and its Background. *J. Hist. Neurosci.* 16, 320–331.
- Guo, J.Z., Graves, A.R., Guo, W.W., Zheng, J., Lee, A., Rodríguez-González, J., Li, N., Macklin, J.J., Phillips, J.W., Mensh, B.D., et al. (2015). Cortex commands the performance of skilled movement. *Elife* 4, 1–18.

- Guo, Z., Li, N., Huber, D., Ophir, E., Gutnisky, D., Ting, J., Feng, G., and Svoboda, K. (2014a). Flow of cortical activity underlying a tactile decision in mice. *Neuron* 81, 179–194.
- Guo, Z.V., Li, N., Huber, D., Ophir, E., Gutnisky, D., Ting, J.T., Feng, G., and Svoboda, K. (2014b). Flow of Cortical Activity Underlying a Tactile Decision in Mice. *Neuron* 81, 179–194.
- Hanks, T.D., Ditterich, J., and Shadlen, M.N. (2006). Microstimulation of macaque area LIP affects decision-making in a motion discrimination task. *Nat. Neurosci.* 9, 682–689.
- Hanks, T.D., Kopec, C.D., Brunton, B.W., Duan, C.A., Erlich, J.C., and Brody, C.D. (2015). Distinct relationships of parietal and prefrontal cortices to evidence accumulation. *Nature* 520, 220–223.
- Harvey, C.D., Coen, P., and Tank, D.W. (2012). Choice-specific sequences in parietal cortex during a virtual-navigation decision task. *Nature* 484, 62–68.
- Hikosaka, O., Nakamura, K., Sakai, K., and Nakahara, H. (2002). Central mechanisms of motor skill learning. *Curr. Opin. Neurobiol.* 12, 217–222.
- Hippenmeyer, S., Vrieseling, E., Sigrist, M., Portmann, T., Laengle, C., Ladle, D.R., and Arber, S. (2005). A Developmental Switch in the Response of DRG Neurons to ETS Transcription Factor Signaling. *PLoS Biol* 3, e159.
- Histed, M.H., Pasupathy, A., and Miller, E.K. (2009). Learning Substrates in the Primate Prefrontal Cortex and Striatum: Sustained Activity Related to Successful Actions. *Neuron* 63, 244–253.
- Hubel, D.H., and Wiesel, T.N. (1959). Receptive fields of single neurones in the cat's striate cortex. *J. Physiol.* 148, 574–591.
- Huber, D., Gutnisky, D.A., Peron, S., O'Connor, D.H., Wiegert, J.S., Tian, L., Oertner, T.G., Looger, L.L., Svoboda, K., O'Connor, D.H., et al. (2012). Multiple dynamic representations in the motor cortex during sensorimotor learning. *Nature* 484, 473–478.
- Hwang, E.J., and Andersen, R.A. (2011). Effects of visual stimulation on LFPs, spikes, and LFP-spike relations in PRR. *J. Neurophysiol.* 105, 1850–1860.
- Hwang, E.J., Hauschild, M., Wilke, M., and Andersen, R.A. (2012). Inactivation of the parietal reach region causes optic ataxia, impairing reaches but not saccades. *Neuron* 76, 1021–1029.
- Kaas, J.H. (2008). The evolution of the complex sensory and motor systems of the human brain. *Brain Res. Bull.* 75, 384–390.

- Kable, J.W., and Glimcher, P.W. (2007). The neural correlates of subjective value during intertemporal choice. *Nat. Neurosci.* 10, 1625–1633.
- Kawai, R., Markman, T., Poddar, R., Ko, R., Fantana, A.L., Dhawale, A.K., Kampff, A.R., and Ölveczky, B.P. (2015). Motor Cortex Is Required for Learning but Not for Executing a Motor Skill. *Neuron* 86, 800–812.
- Kleim, J.A., Barbay, S., and Nudo, R.J. (1998). Functional reorganization of the rat motor cortex following motor skill learning. *J Neurophysiol* 80, 3321–3325.
- Lau, B., and Glimcher, P.W. (2005). Dynamic Response-by-Response Models of Matching Behavior in Rhesus Monkeys. *J. Exp. Anal. Behav.* 84, 555–579.
- Lemon, R.N., and Griffiths, J. (2005). Comparing the function of the corticospinal system in different species: Organizational differences for motor specialization? *Muscle Nerve* 32, 261–279.
- Lien, A.D., and Scanziani, M. (2013). Tuned thalamic excitation is amplified by visual cortical circuits. *Nat. Neurosci.* 16, 1315–1323.
- Liu, B.-H., Huberman, A.D., and Scanziani, M. (2016). Cortico-fugal output from visual cortex promotes plasticity of innate motor behaviour. *Nature* 538, 383–387.
- Luft, A.R., and Buitrago, M.M. (2005). Stages of motor skill learning. *Mol Neurobiol* 32, 205–216.
- Madisen, L., Zwingman, T.A., Sunkin, S.M., Oh, S.W., Zariwala, H.A., Gu, H., Ng, L.L., Palmiter, R.D., Hawrylycz, M.J., Jones, A.R., et al. (2010). A robust and high-throughput Cre reporting and characterization system for the whole mouse brain. *Nat. Neurosci.* 13, 133–140.
- Madisen, L., Mao, T., Koch, H., Zhuo, J., Berenyi, A., Fujisawa, S., Hsu, Y.-W.A., Garcia, A.J., Gu, X., Zanella, S., et al. (2012). A toolbox of Cre-dependent optogenetic transgenic mice for light-induced activation and silencing. *Nat. Neurosci.* 15, 793–802.
- Makino, H., Hwang, E.J., Hedrick, N.G., and Komiyama, T. (2016). Circuit Mechanisms of Sensorimotor Learning. *Neuron* 92, 705–721.
- Mao, T., Kusefoglu, D., Hooks, B.M., Huber, D., Petreanu, L., and Svoboda, K. (2011). Long-Range Neuronal Circuits Underlying the Interaction between Sensory and Motor Cortex. *Neuron* 72, 111–123.
- Maoz, U., Rutishauser, U., Kim, S., Cai, X., Lee, D., and Koch, C. (2013). Predeliberation activity in prefrontal cortex and striatum and the prediction of subsequent value judgment. *Front. Neurosci.* 7.
- De Martino, B., Kumaran, D., Seymour, B., and Dolan, R.J. (2006). Frames, biases, and

rational decision-making in the human brain. *Science* 313, 684–687.

- Masamizu, Y., Tanaka, Y.R., Tanaka, Y.H., Hira, R., Ohkubo, F., Kitamura, K., Isomura, Y., Okada, T., and Matsuzaki, M. (2014). Two distinct layer-specific dynamics of cortical ensembles during learning of a motor task. *Nat Neurosci* 17, 987–994.
- Messier, J., and Kalaska, J.F. (2000). Covariation of primate dorsal premotor cell activity with direction and amplitude during a memorized-delay reaching task. *J. Neurophysiol.* 84, 152–165.
- Molina-Luna, K., Hertler, B., Buitrago, M.M., and Luft, A.R. (2008). Motor learning transiently changes cortical somatotopy. *Neuroimage* 40, 1748–1754.
- Moran, D.W., and Schwartz, A.B. (1999). Motor cortical representation of speed and direction during reaching. *J. Neurophysiol.* 82, 2676–2692.
- Morcos, A.S., and Harvey, C.D. (2016). History-dependent variability in population dynamics during evidence accumulation in cortex. *Nat. Neurosci.* 19, 1672–1681.
- Nakamura, K., Santos, G.S., Matsuzaki, R., and Nakahara, H. (2012). Differential reward coding in the subdivisions of the primate caudate during an oculomotor task. *J. Neurosci. Off. J. Soc. Neurosci.* 32, 15963–15982.
- Neafsey, E.J., and Sievert, C. (1982). A second forelimb motor area exists in rat frontal cortex. *Brain Res.* 232, 151–156.
- Nudo, R.J., Milliken, G.W., Jenkins, W.M., and Merzenich, M.M. (1996). Use-dependent alterations of movement representations in primary motor cortex of adult squirrel monkeys. *J. Neurosci.* 16, 785–807.
- O'Doherty, J., Dayan, P., Schultz, J., Deichmann, R., Friston, K., and Dolan, R.J. (2004). Dissociable Roles of Ventral and Dorsal Striatum in Instrumental Conditioning. *Science* 304, 452–454.
- Otchy, T.M., Wolff, S.B.E., Rhee, J.Y., Pehlevan, C., Kawai, R., Kempf, A., Gobes, S.M.H., and Ölveczky, B.P. (2015). Acute off-target effects of neural circuit manipulations. *Nature* 528, 358–363.
- Padoa-Schioppa, C. (2013). Neuronal Origins of Choice Variability in Economic Decisions. *Neuron* 80, 1322–1336.
- Pandya, D.N., and Kuypers, H.G.J.M. (1969). Cortico-cortical connections in the rhesus monkey. *Brain Res.* 13, 13–36.
- Parker, N.F., Cameron, C.M., Taliaferro, J.P., Lee, J., Choi, J.Y., Davidson, T.J., Daw, N.D., and Witten, I.B. (2016). Reward and choice encoding in terminals of midbrain dopamine neurons depends on striatal target. *Nat. Neurosci.* 19.

- Pascual-Leone, A., Grafman, J., and Hallett, M. (1994). Modulation of Cortical Motor Output Maps During Development of Implicit and Explicit. *Source Sci. New Ser.* 263, 1287–1289.
- Passingham, R.E., Perry, V.H., and Wilkinson, F. (1983). The long-term effects of removal of sensorimotor cortex in infant and adult rhesus monkeys. *Brain* 106 (Pt 3), 675–705.
- Paz, R., Boraud, T., Natan, C., Bergman, H., and Vaadia, E. (2003). Preparatory activity in motor cortex reflects learning of local visuomotor skills. *Nat. Neurosci.* 6, 882–890.
- Penfield, W., and Boldrey, E. (1937). Somatic motor and sensory representation in the cerebral cortex of man as studied by electrical stimulation. *Brain* 60, 389–443.
- Peters, A.J., Chen, S.X., and Komiyama, T. (2014). Emergence of reproducible spatiotemporal activity during motor learning. *Nature* 510, 263–267.
- Peters, A.J., Liu, H., and Komiyama, T. (2017). Learning in the Rodent Motor Cortex. *Annu. Rev. Neurosci.* 40, 77–97.
- Picard, N., Matsuzaka, Y., and Strick, P.L. (2013). Extended practice of a motor skill is associated with reduced metabolic activity in M1. *Nat. Neurosci.* 16, 1340–1347.
- Platt, M.L., and Glimcher, P.W. (1999). Neural correlates of decision variables in parietal cortex. *Nature* 400, 233–238.
- Powell, M., and Ansic, D. (1997). Gender differences in risk behaviour in financial decision-making: An experimental analysis. *J. Econ. Psychol.* 18, 605–628.
- Raposo, D., Kaufman, M.T., and Churchland, A.K. (2014). A category-free neural population supports evolving demands during decision-making. *Nat. Neurosci.* 17, 1784–1792.
- Reep, R.L., Goodwin, G.S., and Corwin, J. V. (1990). Topographic organization in the corticocortical connections of medial agranular cortex in rats. *J. Comp. Neurol.* 294, 262–280.
- Reina, G.A., Moran, D.W., and Schwartz, A.B. (2001). On the relationship between joint angular velocity and motor cortical discharge during reaching. *J. Neurophysiol.* 85, 2576–2589.
- Riout-Pedotti, M., Friedman, D., and Donoghue, J. (2000). Learning-induced LTP in neocortex. *Science* 290, 2335–2336.
- Rokni, U., Richardson, A.G., Bizzi, E., and Seung, H.S. (2007). Motor learning with unstable neural representations. *Neuron* 54, 653–666.

- Rolls, E.T., and Deco, G. (2011). Prediction of Decisions from Noise in the Brain before the Evidence is Provided. *Front. Neurosci.* 5, 33.
- Rouiller, E.M., Babalian, A., Kazennikov, O., Moret, V., Yu, X.H., and Wiesendanger, M. (1994). Transcallosal connections of the distal forelimb representations of the primary and supplementary motor cortical areas in macaque monkeys. *Exp. Brain Res.* 102, 227–243.
- Santesso, D.L., Dillon, D.G., Birk, J.L., Holmes, A.J., Goetz, E., Bogdan, R., and Pizzagalli, D.A. (2008). Individual Differences in Reinforcement Learning: Behavioral, Electrophysiological, and Neuroimaging Correlates. *Neuroimage* 42, 807–816.
- Scott, S.H. (2008). Inconvenient truths about neural processing in primary motor cortex. *J. Physiol.* 586, 1217–1224.
- Seo, H., Barraclough, D.J., and Lee, D. (2009a). Lateral Intraparietal Cortex and Reinforcement Learning during a Mixed-Strategy Game. *J. Neurosci.* 29, 7278–7289.
- Seo, H., Barraclough, D.J., and Lee, D. (2009b). Lateral Intraparietal Cortex and Reinforcement Learning during a Mixed-Strategy Game. *J. Neurosci.* 29, 7278–7289.
- Shadlen, M.N., and Newsome, W.T. (2001). Neural Basis of a Perceptual Decision in the Parietal Cortex (Area LIP) of the Rhesus Monkey. *J. Neurophysiol.* 86, 1916–1936.
- Shadmehr, R., Smith, M.A., and Krakauer, J.W. (2010). Error Correction, Sensory Prediction, and Adaptation in Motor Control. *Annu. Rev. Neurosci.* 33, 89–108.
- Shenoy, K. V, Sahani, M., and Churchland, M.M. (2013). Cortical control of arm movements: a dynamical systems perspective. *Annu. Rev. Neurosci.* 36, 337–359.
- Shmuelof, L., and Krakauer, J.W.W. (2011). Are we ready for a natural history of motor learning? *Neuron* 72, 469–476.
- Stoltz, S., Leigh Humm, J., and Schallert, T. (1999). Cortical injury impairs contralateral forelimb immobility during swimming: A simple test for loss of inhibitory motor control. *Behav. Brain Res.* 106, 127–132.
- Sugrue, L.P., Corrado, G.S., and Newsome, W.T. (2004). Matching Behavior and the Representation of Value in the Parietal Cortex. *Science* 304, 1782–1787.
- Sul, J.J.H., Jo, S., Lee, D., and Jung, M.M.W. (2011). Role of rodent secondary motor cortex in value-based action selection. *Nat. Neurosci.* 14, 1202–1208.

- Svoboda, K., and Yasuda, R. (2006). Principles of Two-Photon Excitation Microscopy and Its Applications to Neuroscience. *Neuron* 50, 823–839.
- Taniguchi, H., He, M., Wu, P., Kim, S., Paik, R., Sugino, K., Kvitsani, D., Fu, Y., Lu, J., Lin, Y., et al. (2011). A Resource of Cre Driver Lines for Genetic Targeting of GABAergic Neurons in Cerebral Cortex. *Neuron* 71, 995–1013.
- Taylor, C.S.R., and Gross, C.G. (2003). Twitches versus movements: a story of motor cortex. *Neurosci.* 9, 332–342.
- Tennant, K. a, Adkins, D.L., Donlan, N. a, Asay, A.L., Thomas, N., Kleim, J. a, and Jones, T. a (2011). The organization of the forelimb representation of the C57BL/6 mouse motor cortex as defined by intracortical microstimulation and cytoarchitecture. *Cereb. Cortex* 21, 865–876.
- Theis, L., Berens, P., Froudarakis, E., Reimer, J., Román Rosón, M., Baden, T., Euler, T., Tolia, A.S., and Bethge, M. (2016). Benchmarking Spike Rate Inference in Population Calcium Imaging. *Neuron* 90, 471–482.
- Thorn, C.A., and Graybiel, A.M. (2014). Differential Entrainment and Learning-Related Dynamics of Spike and Local Field Potential Activity in the Sensorimotor and Associative Striatum. *J. Neurosci.* 34, 2845–2859.
- Todorov, E. (2000). Direct cortical control of muscle activation in voluntary arm movements: a model. *Nat. Neurosci.* 3, 391–398.
- Travis, A.M., and Woolsey, C.N. (1956). Motor performance of monkeys after bilateral partial and total cerebral decortications. *Am. Journal Phys. Med.* 35, 273–310.
- Tversky, A., and Kahneman, D. (1981). The framing of decisions and the psychology of choice. *Science* 211, 453–458.
- Wang, Q., Sporns, O., and Burkhalter, A. (2012). Network analysis of corticocortical connections reveals ventral and dorsal processing streams in mouse visual cortex. *J. Neurosci.* 32, 4386–4399.
- Whishaw, I., and Coles, B. (1996). Varieties of paw and digit movement during spontaneous food handling in rats: Postures, bimanual coordination, preferences, and the effect of forelimb cortex lesions. *Behav. Neurosci.* 77, 135–148.
- Wyart, V., and Tallon-Baudry, C. (2009). How Ongoing Fluctuations in Human Visual Cortex Predict Perceptual Awareness: Baseline Shift versus Decision Bias. *J. Neurosci.* 29, 8715–8725.
- Xu, T., Yu, X., Perlik, A.J., Tobin, W.F., Zweig, J.A., Tennant, K., Jones, T., and Zuo, Y. (2009). Rapid formation and selective stabilization of synapses for enduring motor memories. *Nature* 462.

- Yamawaki, N., Radulovic, X.J., Gordon, X., and Shepherd, M.G. Systems/Circuits A Corticocortical Circuit Directly Links Retrosplenial Cortex to M2 in the Mouse.
- Yin, H.H., and Knowlton, B.J. (2006). The role of the basal ganglia in habit formation. *Nat. Rev. Neurosci.* 7, 464–476.
- Yin, H.H., Mulcare, S.P., Hilário, M.R.F., Clouse, E., Holloway, T., Davis, M.I., Hansson, A.C., Lovinger, D.M., and Costa, R.M. (2009). Dynamic reorganization of striatal circuits during the acquisition and consolidation of a skill. *Nat. Neurosci.* 12, 333–341.
- Yttri, E.A., and Dudman, J.T. (2016). Opponent and bidirectional control of movement velocity in the basal ganglia. *Nature* 533, 402–406.
- Zagha, E., Ge, X., and McCormick, D.A. (2015). Competing Neural Ensembles in Motor Cortex Gate Goal-Directed Motor Output. *Neuron* 88, 565–577.
- Zalocusky, K.A., Ramakrishnan, C., Lerner, T.N., Davidson, T.J., Knutson, B., and Deisseroth, K. (2016). Nucleus accumbens D2R cells signal prior outcomes and control risky decision-making. *Nature* 531, 642–646.
- Zhao, S., Ting, J.T., Atallah, H.E., Qiu, L., Tan, J., Gloss, B., Augustine, G.J., Deisseroth, K., Luo, M., Graybiel, A.M., et al. (2011). Cell type-specific channelrhodopsin-2 transgenic mice for optogenetic dissection of neural circuitry function. *Nat. Methods* 8, 745–752.
- Zingg, B., Hintiryan, H., Gou, L., Song, M.Y., Bay, M., Bienkowski, M.S., Foster, N.N., Yamashita, S., Bowman, I., Toga, A.W., et al. (2014). Neural Networks of the Mouse Neocortex. *Cell* 156, 1096–1111.