NORTHWESTERN UNIVERSITY

The Roles of Primary Motor and Dorsal Premotor Cortex in Motor Adaptation

A DISSERTATION

SUBMITTED TO THE GRADUATE SCHOOL

IN PARTIAL FULFILLMENT OF THE REQUIREMENTS

for the degree

DOCTOR OF PHILOSOPHY

Field of Biomedical Engineering

By

Matthew George Perich

EVANSTON, ILLINOIS

September 2017

ABSTRACT

The brain has a remarkable ability to rapidly adapt behavior. On the scale of development, the brain learns to control the complex dynamics of our limbs by forming and pruning synaptic connections. However, the motor system also learns on much shorter timescales, such as when learning to hit the bullseye of a dartboard, or using a new tool. This process, referred to as motor learning, may be mediated by interactions between the cerebellum and areas of the cerebral cortex including primary motor (M1) and dorsal premotor cortex (PMd). PMd is involved in movement planning with strong connectivity with M1, the main cortical output to the spinal cord. Both are intimately interconnected with the cerebellum. Although cortical reorganization is believed to underlie the long-term learning of motor skills, it is unlikely to account for the rapid adaptation that is observed experimentally. In this project, I studied how behavioral adaptation arises from the coordinated activity of neural populations in the motor cortex. I recorded from neurons in both M1 and PMd as monkeys learned to compensate for perturbations applied to their reaching movements. I show that throughout learning neurons in M1 maintain a fixed relationship with the dynamics of movement, suggesting that adaptation may not involve cortical reorganization. Instead, behavioral changes are mediated through the altered recruitment of M1 neurons. I then study the population-level activity patterns of M1 and PMd and show that PMd plays a direct role in learning how to control the dynamics of the limb by modifying the motor plans sent to M1, potentially by interactions with the cerebellum. These modified plans are executed without changing the functional interactions between neurons in either area, or from PMd to M1. These results provide new insight into the process of motor adaptation and the neural control of movement, and highlight a population-wide mechanism that could help to explain rapid learning processes through the brain.

TABLE OF CONTENTS

LIST OF FIGURES

CHAPTER 1

INTRODUCTION

MOTOR LEARNING

During development, the brain learns to control the complex dynamics of our limbs by forming and pruning synaptic connections. However, the motor system also learns throughout our lives, such as when learning to play an instrument or a sport, and on much shorter time scales, such as when learning to use a novel tool. For decades, neuroscientists and psychologists have attempted to understand what mechanisms exist in the brain to allow such adaptability. Understanding how the motor system can adapt is fundamental to the goal of understanding how the brain controls movement. In the following sections, I will review literature exploring the range and limits of the brain's ability to adapt its motor output, a process commonly called "motor learning". I will then review the anatomy and physiology of areas of the motor cortex, and how they may be involved in motor learning.

PSYCHOPHYSICS OF MOTOR LEARNING

Overview

Through practice, the motor system can learn to change its output in order to better achieve a desired behavior. This process of motor learning occurs on a variety of timescales (Krakauer and Shadmehr 2006), and involves the concerted effort of many brain regions (Grafton et al. 2008). Motor learning is commonly studied in the laboratory environment using systematic perturbations to our natural movements (Shadmehr and Mussa-Ivaldi 1994; Lackner and Dizio 1994; Wolpert et al. 1995). In these tasks, subjects are instructed to perform a simple action, such as to rapidly move a finger (Classen et al. 1998; Muellbacher et al. 2002), reach to a specific location (Shadmehr and Mussa-Ivaldi 1994; Lackner and Dizio 1994; Wolpert et al. 1995), or throw a dart (Martin et al. 1996a). The experimenters then apply the perturbation, inducing

errors. Over time, subjects use sensory feedback regarding these errors to adjust their motor output to achieve the desired goal (Wolpert et al. 1998). By studying the behavioral signatures of this adaptation, researchers hope to explain not only how the motor system can adapt, but also how it plans and executes natural movements.

Internal models of the limb

To help understand and explain the process of motor learning, we can draw inspiration from the control of robotics and view the motor system as a feedback controller (Slotine 1985; Wolpert et al. 1995). Reaching movements are believed to be planned in the coordinates of the hand to achieve a desired smooth kinematic trajectory (Morasso 1981; Flash and Hogan 1985). Yet, the motor command must ultimately be transformed into a pattern of muscle activations that actually drive the limb. This process may be mediated by an "inverse model" of the limb (Shadmehr and Mussa-Ivaldi 1994). While making a movement, sensory feedback is used to shape the motor output, but delays in the motor command output and subsequent sensory signals would make pure feedback control difficult. Thus, when the brain generates a motor plan, it does so using a "forward model" of the motor output (Wolpert et al. 1995). This forward model predicts the sensory consequences of the action. Potential errors, then, can be identified by comparing the actual motor output against the model's expectation (Wolpert et al. 2011).

Many studies have proposed that motor learning can be achieved by modifying the brain's forward and inverse models of the limb (Shadmehr and Mussa-Ivaldi 1994; Wolpert et al. 1995, 1998; Lalazar and Vaadia 2008). These modifications can occur rapidly, even within a single trial (Thoroughman and Shadmehr 2000). Learning can even occur without the subject's knowledge, indicating that the motor system can adapt these internal models subconsciously

(Kagerer et al. 1997; Mazzoni and Krakauer 2006). In the following sections, I will describe a variety of experiments that were designed to investigate what role these forward and inverse internal models play in coordinating movement.

Perturbations of limb dynamics

Experimentally, motor learning has been studied using a number of paradigms, each providing insight into distinct aspects of this system. One primary class of experiments modifies the dynamics of the motor output. A classic study by Lackner and Dizio placed subjects in a spinning room to induce Coriolis forces that perturbed their ability to reach outwards from the body (Lackner and Dizio 1994). In another study, Shadmehr and Mussa-Ivaldi modified the dynamics of planar reaching movements using a velocity-dependent force applied to the hand by robotic motors (Shadmehr and Mussa-Ivaldi 1994). A variant of this paradigm, referred to as the "curl field", is commonly used to study motor learning (Gandolfo et al. 1996; Brashers-Krug et al. 1996; Shadmehr and Brashers-Krug 1997; Thoroughman and Shadmehr 1999, 2000; Krakauer et al. 1999; Caithness et al. 2004; Mattar and Gribble 2005; Smith and Shadmehr 2005; Thoroughman and Taylor 2005; Smith et al. 2006; Hwang et al. 2006; Orban de Xivry et al. 2011). In all of these experiments, the subjects learned to compensate for the altered dynamics to make reaches with straight endpoint kinematic trajectories.

To achieve this type of learning, it is believed that subjects adapt their inverse internal model of the limb dynamics. In psychophysics, evidence of the adapted internal model comes from the presence of "after effects", or oppositely-directed errors produced when the perturbation is removed (Shadmehr and Mussa-Ivaldi 1994). After effects indicate that the novel dynamics have been internalized, forcing subjects to readapt to their normal movement execution dynamics

once the perturbation is removed. The learned dynamics in one posture generalize to novel postures (Shadmehr and Mussa-Ivaldi 1994; Gandolfo et al. 1996; Thoroughman and Taylor 2005). Furthermore, the specific generalization patterns suggest that the inverse internal model is represented in the intrinsic joint coordinates (Shadmehr and Mussa-Ivaldi 1994; Gandolfo et al. 1996), rather than in any external or kinematic coordinate frame. This generalization provides further evidence that learning arises from a modified inverse model of the limb rather than something like a look up table. In the following chapters, I will explore potential neural mechanisms in the motor cortex underlying this adapted internal model.

Learning with motor primitives

The rapid rate of motor adaptation belies the complexity of the dynamics of the limb. The brain can readily adjust its output to perform the desired behavioral goal, accounting for the complex limb mechanics and large number of degrees of freedom. One possible explanation is that the brain uses flexible combinations of discrete motor modules, referred to as motor primitives (Thoroughman and Shadmehr 2000; Mussa-Ivaldi and Bizzi 2000), to simplify the control problem. This framework can predict the behavioral error patterns during CF learning, as well as the limitations on the ability to learn (Thoroughman and Shadmehr 2000). In Chapter 3, I will discuss a conceptual model of cortical function which offers a possible explanation for these primitives.

Although we can rapidly engage these primitives to achieve a variety of learned behaviors, there are combinations of perturbations that appear to be incompatible, and cannot be learned at the same time. For example, the motor system does not seem to be able to learn inverse internal models for multiple conflicting curl fields simultaneously (Brashers-Krug et al. 1996; Wigmore

et al. 2002). Intriguingly, a more recent study showed that opposing curl fields can be learned simultaneously if they are presented such that each field requires a separate preparatory state, whether by perturbation (Krakauer et al. 1999) or by task design (Sheahan et al. 2016). This raises the intriguing possibility that a specific adapted motor plan is crucial to motor learning, an idea that will be explored in Chapter 4.

Visuomotor perturbations

A second common class of motor learning experiments perturbs the visual feedback that the subject receives. In the classic psychophysical studies, participants wore optical prism goggles that shifted the visual feedback of the world, inducing errors as subjects performed a motor behavior (Held and Schlank 1959; Held and Freedman 1963; Cohen 1967; Martin et al. 1996a, b). Over time, the subjects learned to compensate for the visual shift. As with the curl field, there were after effects when the prisms were removed. An analogous paradigm that is commonly used in motor learning studies is the visuomotor rotation (VR) (Krakauer et al. 1999, 2005, 2006; Mazzoni and Krakauer 2006; Schlerf et al. 2012). The VR is typically implemented in a virtual environment such as the cursor on a computer screen, made to track the motion of the hand. The cursor feedback is then rotated by a fixed angle, and the subjects must compensate in order to acquire the desired target. Throughout this process, the motor system is believed to modify its internal models (Wolpert et al. 1995) to reduce the sensory error.

Although the solution to the VR seems simple, the ability to learn a VR depends on a number of factors, including the starting posture (Baraduc and Wolpert 2002), the speed of presentation (Kagerer et al. 1997; Werner et al. 2014), or the context (Ingram et al. 2013). An intriguing study by Mazzoni et al. demonstrated that such adaptation occurs even when subjects

are instructed to pursue the explicit strategy to compensate for the VR (Mazzoni and Krakauer 2006). In this case, subjects were given a 45° rotation and instructed to reach to neighboring target 45° away in order to compensate. Thus, there was a mismatch between the explicit intention of the subject ("reach to the neighboring target") and the prediction of the forward model, which expected the cursor to go to the neighboring target rather than the original target. Movement errors gradually increased as the motor system adapted to this error, what they called implicit learning. This adaptation was at odds with the explicit intention of the subject. This result illustrates that sensory errors during VR learning engage an implicit adaptive process, and that learning likely cannot be reduced to a purely cognitive strategy. This observation is useful for interpreting the results I will present in Chapter 4 using a similar VR paradigm.

Multiple timescales of motor learning

The internal models described above can be adapted within a single session, yet most real-life skills are formed over months or years of practice. Evidence from psychophysical studies in humans suggests that motor learning occurs on at least two distinct timescales (Karni et al. 1998; Kleim et al. 2004; Costa et al. 2004; Smith et al. 2006). The early phase of learning occurs on immediate exposure to a perturbation and is characterized by rapid reduction in errors (Riek et al. 2012). A later phase is characterized by a slower adaptation rate as behavior begins to stabilize. There is compelling evidence that these two processes require different underlying neural mechanisms (Baraduc et al. 2004; Costa et al. 2004; Riek et al. 2012; Herzfeld et al. 2014).

Several studies have attempted to understand how memories of motor skills are stored and improved over time. Shadmehr and Brashers-Krug allowed subjects to adapt to a force field on one day, and re-tested them with the same field a day later. They observed that the subjects

performed significantly better in the early trials of the second day, with a faster rate of adaptation, than they did on the first day (Shadmehr and Brashers-Krug 1997). This phenomenon, which was also previously observed in the primate oculomotor system (Kojima et al. 2004), is commonly referred to as "savings" (Krakauer et al. 2005; Smith et al. 2006; Huang et al. 2011). Between sessions, the motor memory is "consolidated" in the brain for long-term storage (Brashers-Krug et al. 1996; Shadmehr and Brashers-Krug 1997; Shadmehr and Holcomb 1997; Muellbacher et al. 2002; Krakauer and Shadmehr 2006; Debas et al. 2010). After this time, the motor memory is seemingly immune to interference, and can be readily recalled. The subsequent savings have been shown to persist for as long as months (Shadmehr and Brashers-Krug 1997).

Consolidation is advantageous because it places the motor memory in a protected state. A professional tennis player, for example, can learn to play ping pong without fear of forgetting how to swing their racket. However, a group of three laboratories published a study attempting to replicate the effect of consolidation of motor memories in a variety of tasks. They found that their subjects consistently were not able to consolidate the motor memories (Caithness et al. 2004), and instead were susceptible to interference from conflicting perturbations. This was true even days after the initial learning, allowing plenty of time for consolidation of the initial motor memory. Furthermore, effective consolidation has proven elusive in many visuomotor rotation and prism goggle paradigms (Caithness et al. 2004; Krakauer and Shadmehr 2006). The authors proposed that motor memories do not get consolidated into a protected state, but instead merely become inactive. The memories can then be recalled and further modified when necessary, but may become susceptible to interference. These complications suggest that the necessary conditions for effective long-term memory storage and retrieval are complex, and likely depend

on the task demands (Goedert and Willingham 2002), as well as the characteristics of the neural processes that mediate the learning. In the experiments I present in the following chapters, I will focus exclusively on within-session learning, which is largely comprises the fast stage of learning before consolidation. However, I will extensively discuss the implications of my results for long-term learning in Chapter 5.

MOTOR CORTEX AND THE NEURAL CONTROL OF MOVEMENT

Overview

A goal-directed movement such as a reach ultimately arises from the activation of spinal motoneurons that control the muscles of the limb. However, there are a number of cortical and sub-cortical structures, including the spinal cord, brainstem motor nuclei, cerebellum, and the motor cortices, that play important roles in transforming the high-level goal into a low-level pattern of muscle activations (Kalaska and Crammond 1992). In this section, I will review the rich history of literature studying the anatomy and physiology of the motor and premotor cortices, two of the critical output areas of the cerebral cortex.

Primary motor cortex anatomy and physiology

The primary motor cortex (M1), or Brodmann's Area 4 (Brodmann 1909), was first identified by its low threshold for eliciting movements with electrical stimulation (Fritsch and Hitzig 1870; Ferrier 1873). M1 lies along the precentral gyrus and extends into the central sulcus. The area is characterized by the presence of large projection neurons in Layer 5, better known as Betz cells (Betz 1874). Electrical stimulation of M1 with intracortical microstimulation (ICMS) readily elicits muscle contractions or movement with low current thresholds (Ferrier

1873; Bucy 1933; Penfield and Boldrey 1937; Fetz and Baker 1969). M1 is the primary output of the cerebral cortex to the spinal cord, with direct projections to spinal circuits via the pyramidal tract (corticospinal neurons) (Rathelot and Strick 2009).

The spatial arrangement of M1 appears to have a functional organization. Experimenters have used electrical stimulation to map the surface of the gyrus and identify a coarse topographic map of motor output, with the medial aspect marked by leg and proximal arm movements, and the lateral aspect corresponding to distal arm and facial movements (Penfield and Boldrey 1937). Neurophysiological studies of single neurons suggested that cells deep within the central sulcus were more intimately related to the activation of muscles than those cells found on the gyrus (Johnson et al. 1996; Crammond and Kalaska 1996) A recent study by Rathelot and Strick found evidence that the sulcal M1 had direct projections to spinal motoneurons, and may be specialized for finger movements (Rathelot and Strick 2009). The authors used retrograde viral labeling to identify corticospinal projections and found that many cells in the sulcus had monosynaptic connections with motoneurons, suggesting that the hand regions in the sulcus evolved to achieve to dexterous control. In contrast, most of the corticospinal projections from the gyrus (which includes the arm region studied in the later chapters) synapse on spinal interneurons, and thus have an indirect effect on the spinal motoneurons that cause movement. They drew an evolutionary distinction between the gyrus, called “Old M1”, and the phylogenetically newer sulcal region, called “New M1”, proposing that over time primates evolved the direct connections to achieve the more complex demands of dexterous finger movements. In the experiments described in the following chapters, I focus on neural activity on the gyrus which is related to proximal limb movements for reaching.

The precise function of M1 during behavior has been the subject of considerable debate over the last several decades. In 1982, Georgopoulos found that M1 activity recorded during a reaching behavior correlated with the direction of reach (Georgopoulos et al. 1982). He averaged the neural firing rate over the course of the reach and fit cosine tuning curves to each neuron according to the model $D = b_0 + b_1 \bullet \cos(\theta + b_2)$, where D represents the discharge of the neuron and θ the direction of movement. There were three parameters to fit for these tuning curves, representing the mean firing rate (b_0), the depth of modulation (b_1), and the preferred direction (PD; b_2). The PD represents the direction of hand motion for which the cell fires maximally. A later study expanded this model to include the effect of movement speed (Moran and Schwartz 1999). A series of studies, discussed more extensively below, used similar kinematic tuning models to attempt to understand the role of motor cortical areas in motor learning (Gandolfo et al. 2000; Li et al. 2001; Padoa-Schioppa et al. 2002, 2004; Paz et al. 2003; Xiao et al. 2006; Richardson et al. 2008; Arce et al. 2010a, b; Mandelblat-Cerf et al. 2011). Additionally, this kinematic framework has been used extensively in Brain Computer Interfaces (BMIs) to control the movement of computer cursors or robotic arms (Paninski et al. 2002; Taylor et al. 2002; Carmena et al. 2003; Hochberg et al. 2006; Collinger et al. 2013).

Although the kinematic tuning curves employed by Georgopoulos are adequate descriptions of the basic neural firing statistics during reaching, they provide an incomplete view of the information contained within M1. Much earlier, in 1968, Evarts showed that pyramidal tract neurons modulated their activity with the amount of force applied about the wrist (Evarts 1968). In 1969, Fetz and Baker compared the activity of neurons in the precentral cortex to electromyograms (EMG) of muscles in the leg and showed that increased muscle activation typically coincided with bursts of neural activity (Fetz and Baker 1969). In the following years,

many studies have shown a direct relationship between M1 activity and muscle activity (Fetz et al. 1976; Thach 1978; Morrow and Miller 2003; Pohlmeyer et al. 2007; Cherian et al. 2011; Ethier et al. 2012), and its activity has been shown to depend on several other non-kinematic variables such as force (Humphrey et al. 1970; Cheney and Fetz 1980), joint torques (Fetz et al. 1986; Bauswein et al. 1991; Werner et al. 1991), and posture (Scott and Kalaska 1995, 1997). These studies make it clear that, the information contained within M1 has a complex and multi-variate relation to the dynamics of limb movement (Kalaska and Crammond 1992; Churchland and Shenoy 2007).

Premotor cortex anatomy and physiology

The premotor cortex, or Brodmann's Area 6, lies directly anterior to M1, and is distinguished from M1 by its cytoarchitecture, notably a lack of Betz cells (Bucy 1933, 1935). Early studies identified a number of additional features that distinguish premotor cortex from M1. The area has a higher threshold for eliciting movement from electrical stimulation (Bucy 1933; Weinrich and Wise 1982; Dum and Strick 2002). Anatomically, the premotor cortex receives diverse corticocortical inputs (Kurata 1991), and has strong bidirectional connectivity with M1 (Dum and Strick 2002, 2005) and the cerebellum (Dum and Strick 2003). Although premotor cortex is typically viewed as hierarchically "above" M1, it also has many projections to the spinal cord and other subcortical structures (Dum and Strick 1991). Thus, the premotor cortex may provide information about movement to these structures in parallel with M1.

In monkeys, the premotor cortex is subdivided by the spur of the arcuate sulcus into dorsal (PMd) and ventral (PMv) regions, each with distinct inputs, outputs, and functional roles (Kurata and Hoffman 1994; Hoshi and Tanji 2002, 2006). The activity of PMv is related to both motor

and cognitive processes (Rizzolatti et al. 2002), but is believed to play an important role in grasping (Hoshi and Tanji 2006). PMv receives information about both visual and somatosensory stimuli, and appears to encode movements in extrinsic space (Kakei et al. 2001). Rather than the multi-sensory grasp-related activity of PMv, PMd appears to have a more specific motor role during reaching (Hoshi and Tanji 2002). PMd can be further subdivided into rostral and caudal segments, each with different functional characteristics (Fujii et al. 2000). While the rostral aspect appears to play a role in eye movements, and can even elicit saccades with electrical stimulation (Fujii et al. 2000), the caudal aspect relates specifically to movements of the limb.

In the following chapters, I will focus exclusively on the caudal aspect of PMd, which plays an important role in preparing and executing reaching movements (Weinrich and Wise 1982; Shen and Alexander 1997; Hoshi and Tanji 2002; Cisek et al. 2003; Cisek and Kalaska 2005; Churchland et al. 2006a, 2010a). In 1980, Roland et al. studied the cerebral blood flow of healthy subjects during movements and found that premotor cortex was preferentially activated when a movement required a new "motor program", or when the motor program needed to be adjusted based on sensory feedback (Roland et al. 1980). PMd activity can be used to predict the intended action of a monkey long before movement begins (Cisek and Kalaska 2005; Santhanam et al. 2006; Thura and Cisek 2014). Furthermore, activity in PMd preferentially reflects the target or goal of a reaching movement (Vaadia et al. 1988; Shen and Alexander 1997), and even the probability of an upcoming reach(Dekleva et al. 2016; Glaser et al. 2017) . Based on these observations, PMd is believed to help facilitate the transformation from planned kinematics to executed dynamics (Shen and Alexander 1997; Batista et al. 2007).

A POPULATION-LEVEL VIEW OF CORTICAL ACTIVITY

Overview

For many years, since the work of Ramon y Cajal (y Cajal 1995) and the Neuron Doctrine, the neuron has been treated as the basic functional unit of the brain. Through the anatomical and neurophysiological studies described in the preceding sections, it is apparent that M1 is intimately related to the execution of movement and that PMd plays an important role in preparing the movement. However, the relationship between neural activity in both of these areas and behavior is complex. Many cells appear to represent multiple external covariates, while others represent no single covariate. Ultimately, since many of the available measurable behavioral outcomes, such as hand kinematics or joint torques, are correlated, it is difficult to determine what precisely is encoded by neurons in M1. Recent work has begun to question whether neurons in M1 must encode any particular behavioral variable. In 1992, Eb Fetz proposed that the purpose of neurons in M1 is to generate patterns of muscle activation, and as such, neurons do not necessarily need to represent any specific external covariation (Fetz 1992). More recently, Churchland et al. studied the activity of single neurons during reaching and found widely heterogeneous firing patterns, with a great deal of temporal complexity (Churchland and Shenoy 2007). These properties make it difficult to relate neural activity reliably to any single behavioral covariate. Ultimately, the neurons in the motor cortices must cause movement, not simply represent it (Fetz 1992; Kalaska and Crammond 1992), so it is unlikely that we will find such simple, lawful relationships at the single neuron level.

While researchers have been able to gain much insight into brain function by studying single neurons, each neuron is not independent. The activity of a given neuron is ultimately a reflection of its inputs from other cells (Sussillo et al. 2015). Furthermore, behavior ultimately arises from

the concerted activity of vast populations of neurons (Fetz 1992; Hatsopoulos et al. 1998). As such, several recent theories have proposed that, while each neuron has a unique input/output relationship, when it comes to behavior, the basic building blocks may better be viewed as the coordinated activity of neural populations (Churchland et al. 2012; Shenoy et al. 2013; Kaufman et al. 2014; Elsayed et al. 2016). In this section, I will review a growing body of work suggesting that there is substantial insight to be gained by studying the brain at the level of neural populations. I will then draw from these ideas for the analyses presented in Chapter 4 to study how population-level interactions can give rise to adapted behaviors.

Dimensionality reduction for neural data analysis

The majority of studies described in the previous sections used microelectrodes to record the activity of individual neurons. While modern recording technology allows us to record the activity of hundreds or even thousands of neurons simultaneously, this remains only a small fraction of the number of neurons used to generate even simple actions. Yet, the activity of these small populations of neurons can explain a large fraction of task-relevant variability (Cunningham and Yu 2014; Gao and Ganguli 2015), suggesting that the brain does not use all of the potential degrees of freedom provided by the neurons to generate behavior. This idea was recently formalized by Gao and Ganguli, who theorized that the "dimensionality", or the number of degrees of freedom in the neural population, is bounded by the complexity of a given behavior (Gao and Ganguli 2015). The dimensionality of the neural population can be thought of as the number of independent signals within the population, and is typically defined as the number of components needed to explain the majority of variance. The apparently low dimensionalities observed in M1, typically on the order of 10 (Santhanam et al. 2009; Sadtler et al. 2014), is a

consequence of the simple behavioral tasks used in neurophysiology experiments. Thus, in order to obtain a more complete view of motor cortical function, we may need to employ more complex behavioral tasks. The discussion contained in Chapter 3 explores these ideas in more depth.

Population analyses typically begin with a form of dimensionality reduction. Dimensionality reduction reduces a large number of correlated signals, in this case the spiking activity of individual neurons, into a smaller number of linearly independent signals (Cunningham and Yu 2014). The simplest and most common method is Principle Components Analysis (PCA) (Law and Jolliffe 1987). PCA can be used to identify the shared covariance patterns across a population of neurons and orders them according to the variance explained. The low-dimensional components define the "neural manifold" (Stopfer et al. 1997; Sadtler et al. 2014), which is composed of a relatively small number of components that capture the majority of population variance. The meaning and implications of such manifolds are discussed in depth in Chapter 3, and later used to study cortical activity during motor learning in Chapter 4.

Using population analyses to explain behavior

While traditional single neuron techniques typically require averaging over time and many repetitions of the same behavior to achieve adequate statistics, we can now leverage the large numbers of recorded neurons to study trial-to-trial variations in behavior. Much of the evidence that dimensionality reduction can help to explain how neural populations give rise to behavior comes from the work of Krishna Shenoy and his colleagues. One study used Factor Analysis, a dimensionality reduction technique that is analogous to PCA, to show that PMd activity during movement planning can be reduced to a simple and orderly structure in the neural state space

(Santhanam et al. 2009). Despite the relatively large numbers of recorded neurons, the behavioral output could be explained with only a few dimensions of activity. Another study used population activity to show that trial-to-trial neural variability is reduced following the presentation of a stimulus, and that longer reaction times occurred when more neural variability was present at the time of go cue, suggesting that the movement was not adequately prepared (Churchland et al. 2006b). A later study from the same group showed that the position in neural state space at the time of go cue could predict abnormal reaction times, illustrating a direct link between the dimensionality-reduced view of population activity and the subsequent behavior (Afshar et al. 2011).

One the most compelling examples of population-level analysis comes from Kaufman et al. (Kaufman et al. 2014), who used population activity to explain how the brain can plan a movement to be made at a future time without actually causing that movement. Their analysis exploited the observation that population activity has a larger dimensionality than the behavioral output. They used dimensionality reduction to identify a six-dimensional representation of neural activity, and a three-dimensional representation of muscle activity in a standard reaching task. The larger neural dimensionality necessitated the existence of an "output-null" space, which captured neural activity that produced no output in the muscles, as well as an "output-potent" space that captured the direct mapping from neural activity to muscles. They used a linear mapping between neurons and muscles to define the potent and null spaces, and showed that pre-movement preparatory activity existed in the null dimensions only. Thus, motor planning can occur without causing movement because the output-potent dimensions cancel out. Intriguingly, they found a lawful relationship between activity in the null space and the subsequent behavior. This result gave a meaningful interpretation of the null space activity, which served to set the

preparatory state for the system dynamics of the neural population necessary to cause the desired movement (Churchland et al. 2010a, 2012).

Together, these studies highlight the advantages of studying neural populations. In many of these cases, such as the null space of Kaufman et al., the analysis uncovered a population-level mechanism that was not observable with individual neurons. Besides the motor system, similar techniques have been used to understand odor processing in the locust (Broome et al. 2006), the process of decision making in prefrontal cortex (Mante et al. 2013), and processing in the visual (Cowley et al. 2016) and auditory cortices (Okun et al. 2015), suggesting a common framework linking the function of areas throughout the brain.

Population activity may be constrained by the network connectivity

There is evidence that the neural covariance patterns found through dimensionality reduction are more than mere mathematical descriptions of the data, but instead seem to capture underlying connectivity within the population. Okun and colleagues used population recordings from multiple cortical areas in three species to show that the covariance patterns found through PCA correlated with actual synaptic connectivity (Okun et al. 2015). First, they showed that the correlations between a given neuron and the surrounding population were the same during both spontaneous and stimulus-evoked activity. Second, they showed that this correlation predicted the magnitude of the response of each neuron to optogenetic stimulation of the surrounding population, suggesting that the correlations captured some causal link. Lastly, they showed that the neural covariance during behavior correlated with the number of synapses shared between neurons.

In another experiment, Sadtler et al., used a learning paradigm with a Brain Computer Interface (BCI) to show that the neural covariance patterns obtained through dimensionality reduction reflect actual constraints on the possible network activity (Sadtler et al. 2014). They used FA to identify a ten-dimensional neural manifold, which the mapped linearly to the two-dimensional velocity of a cursor. They asked the monkey to learn two classes of perturbations to this decoder: the "within-manifold" perturbation rotated the mapping between the manifold and the cursor; the "outside-manifold" perturbation rotated the mapping of each neuron to the manifold, breaking the normal covariance patterns. While the monkeys could readily learn the within-manifold perturbation, they were unable to learn the outside-manifold perturbation in a single session. These results, together with those of Okun et al, provide strong evidence that the low-D covariance patterns reflect actual network connectivity.

MOTOR CORTEX AND CEREBELLUM DURING MOTOR LEARNING

Overview

During motor learning, the adapted behavior ultimately results from changes in the activity of M1, PMd, and other cortical areas, likely involving interactions with the cerebellum. In this section, I will provide a brief review of psychophysical and neurophysiological experiments that support this view.

Involvement of M1 in motor learning

Since M1 plays a major role in the execution of voluntary movement, it is likely that the adapted behavior observed in motor learning is ultimately caused by changes in activity of M1 neurons. While the precise role of M1 in developing, storing, and recalling motor memories is

not well understood, M1 appears to play a crucial role in the consolidation of a motor memory. Using whole-brain functional imaging, Shadmehr et al. showed that neural correlates of a motor memory emerged over time in the hours after learning a skill (Shadmehr and Holcomb 1997). This change in M1 activation could predict the subsequent performance. A later study showed that disruption of M1 with repetitive transcranial magnetic stimulation (rTMS) did not impair the within-session learning of a CF, but did eliminate consolidation of the memory (Richardson et al. 2006). A similar experiment showed that disruptive rTMS between sessions of practice with a ballistic finger movement task impaired the retention, again suggesting a critical role of M1 in consolidating the memory (Muellbacher et al. 2002). However, this experiment used a simple motor paradigm involving rapid and precise finger movements. A follow-up study using the CF paradigm could not replicate this result (Baraduc et al. 2004). The authors ultimately concluded that learning novel dynamics requires a different, more distributed neural mechanism than learning a skill that relies on the existing dynamics.

Although CF and VR learning share many common features, adaptation to these two perturbations likely require distinct neural processes within the motor system. Krakauer et al. demonstrated that learning a VR did not interfere with consolidation of a CF memory (Krakauer et al. 1999). Although the authors interpreted this as independent learning of kinematics and dynamics based on two independent sources of error, a later study clarified that the kinematic variable used for the perturbation determines if there will be interference (Tong et al. 2002). The standard CF is velocity-dependent, while the VR is static or position dependent. Tong et al. showed that a position-dependent force field interfered with the VR, suggesting that the interference cannot be entirely attributed to the altered effector dynamics.

Regardless of this distinction, it appears that the learning of dynamic and static perturbations, such as the standard CF and VR, may involve distinct neural circuits (Krakauer 2003). Diedrichsen et al. used fMRI to identify regions of the brain that were active during CF and VR learning and found distinct patterns of activation in the two paradigms (Diedrichsen et al. 2005). Although both perturbations involved both the cerebellum and motor cortex, only the VR also involved interactions with parietal cortex. Regions of parietal cortex are important for reaching (Snyder et al. 1997), possibly by mediating coordinate transformations between the goal in visual coordinates and the desired trajectory in body-centered coordinates (Buneo et al. 2002). Since the VR effectively perturbs how the visual feedback maps onto the hand trajectory, a computational study has suggested that VR learning may involve interactions between motor cortex and the posterior parietal cortex (Tanaka et al. 2009).

Involvement of the cerebellum in motor learning

The cerebellum is known to play a necessary role in motor learning. Patients with cerebellar degeneration are unable to learn to compensate for a CF, suggesting that the cerebellum is necessary to adapt the internal model (Smith and Shadmehr 2005). The circuitry of the cerebellum is organized in a manner that seems to enable the storage and recall of memories (Albus 1971). Early studies implicated the activity of Purkinje cells in the cerebellar cortex in driving motor learning (Gilbert and Thach 1977), and numerous theories have proposed how the cerebellum could implement error-based learning (Kawato and Gomi 1992; Wolpert et al. 1998). Increasing cerebellar excitability with direct current stimulation improves the ability of subjects to adapt to a CF (Galea et al. 2011; Herzfeld et al. 2014), and functional imaging shows activation of the cerebellum during both CF and VR learning (Diedrichsen et al. 2005). Although

the experiments in the following chapters will focus on the motor cortices, the strong connectivity of these areas with the cerebellar nuclei (Dum and Strick 2003) suggests that the cerebellum is closely involved in shaping the activity of cortical neurons during learning.

Long-term learning mechanisms and plasticity

There is considerable evidence that extended training results in persistent changes in the synaptic connectivity within the cortex. A study of the hand representations of expert violinists and other string players found a larger cortical representation in the motor cortex contralateral to the fretting hand (Elbert et al. 1995). They interpreted these results as long-term plasticity that enables the highly dexterous finger movements required to play the instrument. In monkeys, Nudo et al. observed cortical reorganization during skill learning directly. They trained monkeys on a behavioral task and used electrical stimulation to map the forelimb cortex before and after learning and found changes in the identified motor maps (Nudo et al. 1996). This process likely requires synaptogenesis (Kleim et al. 2002), and may operate via NMDA receptor activation (Bütefisch et al. 2000).

Motor learning over days has also been shown to change the horizontal connections within rat M1 (Rioult-Pedotti et al. 1998). In mice, Peters et al. used calcium imaging techniques to monitor the activity of large populations of neurons in M1 during motor skill learning (Peters et al. 2014). They found that, across days, the population activity patterns became increasingly stable, with corresponding changes in the dendritic spines linking neurons. Interfering with the formation of these spines prevented the animals from learning the skill (Chen et al. 2015). Neural plasticity underlying long-term learning has been observed in a number of other species and learning models, including sensory associations in birds (Jeanne et al. 2013) and behavioral

habituation in *Aplysia Californica* (Bailey and Chen 1983). Intriguingly, in the *Aplysia* experiment, structural plasticity was not observed in the hours immediately following learning of the response, suggesting that short-term learning may require a fundamentally different mechanism (Bailey and Chen 1988). A similar finding was observed in rats: synaptogenesis only began to occur after extensive training (Kleim et al. 2004).

Activity of primary motor cortical neurons during curl field learning

A series of studies from the Bizzi group investigated the learning-related changes in neural activity during curl field adaptation (Gandolfo et al. 2000; Li et al. 2001; Padoa-Schioppa et al. 2002, 2004; Xiao et al. 2006; Richardson et al. 2008, 2012). In the earliest studies (Gandolfo et al. 2000; Li et al. 2001), the authors recorded from single neurons in M1 and characterized the spatial tuning (Georgopoulos et al. 1982) of each neuron before (Baseline), during (Force), and after (Washout) CF learning. They found that the preferred direction (PD) of the majority of M1 cells rotated in the direction of the curl field. Using a statistical test on the PDs, they assessed whether there were significant changes in tuning between each of Baseline, Force, and Washout epochs. They identified five classes of cells. "Kinematic" cells showed no change in tuning throughout the session, while "Dynamic" cells changed with the curl field, then rotated back to their original tuning in Washout. Another class, labeled "Other", had different tuning in all epochs. However, the most intriguing cell types, called "Memory I" cells, showed a change in tuning during Force that persisted into Washout. This class was complemented by "Memory II" cells, which did not change during Force but changed in the opposite direction in Washout. Together, these two classes of Memory cells balanced each other such that the average change between Baseline and Washout was near zero. A later study used chronically-implanted

electrode arrays to study the activity of the same M1 neurons over five consecutive sessions with a CF (Richardson et al. 2012). This allowed the authors to assess each individual cell across days for changes in neural activity that could explain the across-session effect of savings (Krakauer et al. 2005). They identified persistent changes in neural tuning between the first and second sessions that they interpreted as an effect of internal model consolidation.

The authors of the original Bizzi studies interpreted the PD changes, and appearance of Memory cells, as evidence of an adapted kinematics-to-dynamics transformation (Shadmehr and Mussa-Ivaldi 1994). The Memory cells, in this framework, represented a learning-related reorganization of M1. This concept was later expanded into a general framework for learning using unstable representations (Rokni et al. 2007), though another study found reaching representations to be quite stable over days (Chestek et al. 2007). A study by Stevenson et al. offers a potential explanation for this difference by showing that poorly estimated tuning curves can appear to change as a result of noise, even when the underlying neural behavior is known to be unchanged (Stevenson et al. 2011). This observation raises the possibility that the appearance of Memory cells following curl field learning may be the result of behavioral or measurement noise, rather than a sign of learning-related plasticity. This possibility will be addressed in-depth in Chapter 2.

A study by Cherian et al. proposed an alternative interpretation for the changes in M1 tuning during CF adaptation (Cherian et al. 2013). The authors simultaneously recorded the activity of neurons in M1 as well as EMG from numerous muscles of the proximal limb. They found that both neurons and muscles immediately changed their spatial tuning when the curl field was imposed, with no progressive change in PDs associated with the behavioral adaptation. Thus, the authors concluded that the PD changes observed by the Bizzi group simply reflected the altered

movement dynamics of the CF. Furthermore, they found no evidence for Memory cells, which they instead attributed to randomly distributed classification errors (Stevenson et al. 2011).

Another series of studies from the Vaadia group also used a variant of the CF paradigm to study the adaptive responses of M1 neurons (Arce et al. 2010a, b; Mandelblat-Cerf et al. 2011). These experiments were unique in that they perturbed only one of the eight target directions of a center out task in order to study the generalization pattern of the model for the learned target to other, unperturbed targets. Arce et al. repeatedly presented only the perturbed target, allowing the monkeys to rapidly adapt, before testing on all eight targets (Arce et al. 2010b). One limitation of this approach is that the directional tuning of the recorded neurons could not be assessed during learning, since the monkeys only reached in a single direction. The authors studied the adaptive responses of cells that were active for the perturbed direction during reaches in that direction. Most of these cells changed their activity to reflect the altered dynamics, and some even appeared to retain this new representation when tested on the unperturbed targets, a phenomenon that the authors concluded was related to Memory cells.

While the Memory effects observed by Arce et al. seem to support the conclusions of the Bizzi studies (Arce et al. 2010b), a follow-up study from the Vaadia group appeared to reach a different conclusion (Mandelblat-Cerf et al. 2011). Like Arce et al, only a single target was perturbed, though all eight targets were presented throughout learning. The authors could thus assess directional tuning throughout the adaptation period. They found that, throughout learning, M1 activity related consistently to the movement that the monkey intended to make. They concluded that the observed changes in M1 activity during CF learning reflect the change in the monkey's intention, or "motor plan", not a change in the functional properties those cells.

The Bizzi group performed several similar curl field experiments while recording from numerous other motor-related regions of the brain. In PMd and the Supplementary Motor Area (SMA), they found evidence of a kinematics-to-dynamics transformation during motor planning (Padoa-Schioppa et al. 2002, 2004; Xiao et al. 2006), while PMv appeared to reflect only the altered dynamics (Xiao et al. 2006). In the cingulate motor area, neurons appeared to alter their discharge during force field learning, but did so in a non-directional manner (Richardson et al. 2008). Together, these studies suggest that learning is likely a consequence of the coordinated activity of a variety of interconnected motor and premotor regions (Diedrichsen et al. 2005). The studies of M1 by Mandelblat-Cerf et al. and Cherian et al. described previously do not support the hypothesis that the observed neural tuning changes in M1 are a consequence of plastic changes occurring within M1. Instead, they strongly suggest that compensation for the CF occurs above the inputs to M1, and that the new patterns of muscle activation that drive the adapted behavior are a consequence of altered recruitment of M1 neurons. This possibility was acknowledged by Richardson et al., who said that their Memory cells may result from a persistent change in input to M1 rather than plastic reorganization within M1 (Richardson et al. 2012). Given the strong bidirectional connectivity between PMd and M1 (Dum and Strick 2002, 2005), and the role of premotor cortex in motor planning (Roland et al. 1980; Weinrich and Wise 1982; Shen and Alexander 1997; Cisek and Kalaska 2005), PMd is an appealing candidate to look for the source of this altered recruitment. In Chapter 4, I will describe an experiment designed to test this hypothesis directly.

SUMMARY

In this chapter, I reviewed a wide range of literature that suggests that the motor and premotor cortices play an important role in motor learning. The following chapters will address one of the key questions in understanding the neural basis of motor learning: what mechanisms are employed by the motor cortices to adapt behavior rapidly? In Chapter 2, I will resolve the apparent discrepancies between the previous studies of M1 activity during CF learning. In Chapter 3, I will describe a conceptual framework for studying and interpreting the activity of neurons which focuses on population-level interactions. In Chapter 4, I will describe an experiment that uses simultaneous population recordings from M1 and PMd to shed light on how these areas are coordinated during CF and VR learning. I will provide evidence that PMd plays an important role in CF adaptation by adjusting a motor plan to compensate for the new limb dynamics. In the final chapter, I will discuss the implications of my work, and speculate on possible mechanisms that may underlie short-term and long-term motor learning.

CHAPTER 2

ALTERED TUNING IN PRIMARY MOTOR CORTEX DOES NOT ACCOUNT FOR BEHAVIORAL ADAPTATION DURING FORCE FIELD LEARNING

Matthew G. Perich, Lee E. Miller

FOREWORD

This chapter consists of a manuscript that has been published in the journal Experimental Brain Research (Perich and Miller 2017). The purpose of this paper was to resolve a discrepancy between the "Memory cell" result of the classic studies of the Bizzi group, and a more recent experiment from Lee Miller's lab. The latter study proposed that the changes in spatial tuning observed in M1 neurons merely reflect the altered dynamics of the motor output, though a number of procedural differences made direct comparison difficult. I recorded from neurons in the primary motor cortex of monkeys in a curl field task. I looked at the spatial tuning properties of each cell to look for evidence of cortical plasticity, but ultimately did not find any changes in neural properties with a time course that matched that of behavioral adaptation. I ultimately show that the neurons consistently relate to the dynamics of the motor output throughout learning, and use a basic musculoskeletal model to illustrate this concept in a simulated population of neurons.

ABSTRACT

Although primary motor cortex (M1) is intimately involved in the dynamics of limb movement, its inputs may be more closely related to higher order aspects of movement and multi-modal sensory feedback. Motor learning is thought to result from the adaption of internal models that compute transformations between these representations. While the psychophysics of motor learning has been studied in many experiments, the particular role of M1 in the process remains the subject of debate. Studies of learning-related changes in the spatial tuning of M1 neurons have yielded conflicting results. To resolve the discrepancies, we recorded from M1 during curl field adaptation in a reaching task. Our results suggest that aside from the addition of the load itself, the relation of M1 to movement dynamics remains unchanged as monkeys adapt behaviorally. Accordingly, we implemented a musculoskeletal model to generate synthetic neural activity having a fixed dynamical relation to movement and showed that these simulated neurons reproduced the observed behavior of the recorded M1 neurons. The stable representation of movement dynamics in M1 suggests that behavioral changes are mediated through progressively altered recruitment of M1 neurons, while the output effect of those neurons remain largely unchanged.

INTRODUCTION

Despite the complexity of the limb's dynamics, we make precise reaching movements with smooth hand kinematics, long thought to suggest that we plan reaches in a hand-centered reference frame (Morasso 1981; Flash and Hogan 1985). More recently, models based on optimal feedback control have been proposed that do not require explicit trajectory planning and transformation (Todorov 2004; Scott 2012). In either case, control signals must ultimately be

expressed in the intrinsic, musculoskeletal coordinates of motor execution. Both approaches posit the use of internal models of the limb (Wolpert et al. 1995), whether to transform a desired kinematic trajectory into kinetic commands (Kalaska and Crammond 1992; Shadmehr and Mussa-Ivaldi 1994), or for optimal state estimation (Shimansky et al. 2004; Todorov 2004; Scott 2012). These internal models are learned during development, but it is also possible to modify them on a shorter time scale, for example, when using a heavy tool that requires altered forces to produce a given movement.

The psychophysics of this short-term motor adaptation process have been studied by using Coriolis forces to perturb reaches made in a rotating room (Lackner and Dizio 1994) and by perturbing reaches with forces imposed by a robotic manipulandum (Shadmehr and Mussa-Ivaldi 1994). Over time, subjects learn to alter their patterns of muscle activity to compensate for these external forces and restore normal kinematic trajectories. When normal movement dynamics are restored, the transient appearance of behavioral "aftereffects", expressed as oppositely directed errors during re-adaptation to normal movement dynamics, is taken as evidence of internal model adaptation.

The pattern of generalization found in the experiments described above suggests that the learned internal model represents the dynamics of movement in terms of intrinsic coordinates of the limb (Shadmehr and Mussa-Ivaldi 1994). Given the considerable evidence that primary motor cortex (M1) encodes low-level details of motor execution (Evarts 1968; Fetz et al. 1986; Kalaska et al. 1989; Scott and Kalaska 1997; Morrow and Miller 2003; Sergio et al. 2005), it is reasonable to expect that this short-term adaptive process could be implemented by neurons in M1. Experimental evidence for the short-term adaptation of internal models within M1 comes from several studies of altered neural tuning as monkeys adapted to a curl-field (CF) perturbation

(Gandolfo et al. 2000; Li et al. 2001; Richardson et al. 2012). These studies compared the spatial tuning properties of individual M1 neurons before, during, and after CF adaptation. They reported that a subset of M1 neurons (called "Dynamic" cells) changed their tuning in response to the altered dynamics. Intriguingly, they also found a number of cells whose change in tuning persisted even when the CF was removed, together with a second, complementary group of cells that had no initial change, but changed subsequently when the CF was removed. These "memory cells", were thought to be evidence of the modification of internal models within M1 during CF adaptation. A follow-up study found no evidence of memory cells, with neural tuning changes simply mimicking the behavior of muscles (Cherian et al. 2013), although differences in task design between the studies made it difficult to compare the differing results directly.

The goal of the present study was to resolve the discrepancies between these studies and characterize the behavior of neurons in M1 prior to, in the presence of, and following a curl force field. Imposing the curl field led to apparent tuning changes in M1 that occurred as early as we were able to test, which were consistent with the altered dynamics of movement and directly proportional to the magnitude of the applied perturbation. Furthermore, simulated neurons with fixed dynamical relationships generated from a simple musculoskeletal model reproduced the critical features observed in the recorded units. These new results suggest that M1 neurons do not exhibit short-term learning-related changes reminiscent of internal model adaptation, but instead have a largely unchanging relation to the dynamics of movement that must have been formed over longer, perhaps developmental periods of time. Short-term behavioral adaptation thus appears to occur at a hierarchical level above the inputs to these cells, likely involving upstream areas that provide input to M1.

METHODS

Behavioral Task

Two monkeys (male, *macaca mulatta*) were seated in a primate chair and grasped the handle of a 2-D planar manipulandum that controlled a computer cursor on a screen in a standard center-out (CO) task (Figure 2.1a). A trial began when the monkey moved to the center target. After a hold period of 500-1500 ms, a 2cm target was randomly displayed in one of eight regularly spaced outer positions at a radial distance of 8 cm, followed by a variable delay period of 500-1500 ms before an auditory go cue (Figure 2.1b). Monkey C was initially trained without the instructed delay (which had been used in the prior studies), and the first six curl field sessions with this monkey omitted it. As we were interested in the movement-related activity of M1 it is unlikely that the lack of a delay impacted our conclusions, though to ensure that there were not differences we later collected an additional seven sessions from Monkey C with the delay. The monkeys were required to reach to the target within one second and hold for 500 ms to receive a liquid reward.

Each experimental session had three task epochs: a short Baseline epoch in which the monkey made unperturbed movements, a longer Adaptation epoch when the monkey adapted to reaching movements made with the curl field, and a Washout epoch when the CF was removed and the monkey readapted to normal movement dynamics. During CF trials, motors in the manipulandum applied forces at the endpoint proportional to the velocity of the hand according to Equation 1.

$$\vec{F} = \begin{bmatrix} F_x \\ F_y \end{bmatrix} = k \begin{bmatrix} \cos\theta_c & -\sin\theta_c \\ \sin\theta_c & \cos\theta_c \end{bmatrix} \begin{bmatrix} \dot{p}_x \\ \dot{p}_y \end{bmatrix} \quad (1)$$

where F is the applied force, $\dot{p}$ is hand velocity, and k is a constant (0.15 N•s/cm). The forces were exerted at a direction θ_c of 85° relative to the direction of hand movement to avoid instabilities that occurred when the force was applied at 90°. The CF was enabled for the duration of the Adaptation epoch including the return movement and intervals between trials.

Both monkeys required several experimental sessions to learn to tolerate the perturbation, initially completing only a small number of reaches in a given curl field session. Data analysis began with the first session in which a given monkey completed at least 25 reaches to each target during adaptation and had sufficient time in Washout for deadaptation. The data used for analysis began with the third and seventh interaction with the curl field for Monkey C and Monkey M, respectively. The experimental sessions reported here were typically not consecutive, but instead, had intervening sessions with adaptation to visual rotation. Consequently, we focus here on within-session learning, since the experiments were not designed to investigate long-term savings. The monkeys also performed a small set of control sessions of similar duration to the force field sessions, but with no applied force field. In these sessions, all other task parameters were identical to the curl field sessions. These control sessions allowed us to ascertain the baseline variability in our analyses in order to better understand the effect of the force field.

Implantation of Microelectrode Arrays

We implanted 100-electrode arrays with 1.5mm shaft length (Blackrock Microsystems, Salt Lake City) in the arm area of M1 of two monkeys. We placed the monkeys under isoflurane anesthesia and opened a craniotomy above the motor cortex. M1 was localized using visual landmarks and the arm area was identified using bipolar cortical surface stimulation to evoke

twitches of proximal muscles. The arrays were then inserted pneumatically. Figure 2.1c shows array implant locations for the two monkeys and neighboring cortical surface landmarks, based on photographs taken intraoperatively.

Analysis of Behavioral Adaptation

We computed the "takeoff angle" error between one vector drawn from the position of the hand at the start of movement to the position of the hand at the time of peak speed, and a second vector pointing from the hand directly to the target. This metric was designed to focus on the ballistic phase of movement and ignored force-induced error corrections later in the reach.

Neural Data Acquisition

Neural data were amplified, band-pass filtered (250 to 5000 Hz) and digitized using a Cerebus system (Blackrock Microsystems, Salt Lake City, UT). We identified threshold crossings of 5.5 times the root-mean square (RMS) amplitude of the signal on each of the channels and recorded spike times and brief waveform snippets. Additionally, we recorded kinematic data from the robot handle and endpoint force data using a 6-axis strain gauge in the handle to measure the net forces acting on the hand. After each session, we used Offline Sorter (Plexon, Inc, Dallas, TX) to sort all the waveforms that crossed a detection threshold. Importantly, we sorted the waveforms for all three epochs together, to ensure that we did not inadvertently introduce sorting differences. Since we sought to study well-isolated neurons, they were included only if they had a waveform signal to noise ratio greater than three (calculated as the average waveform peak-to-peak value divided by two times the standard deviation of the waveform shapes). We ensured that each single unit was reliably held throughout the session by

comparing the spikes from each epoch using a statistical test that incorporated the waveform shapes and inter-spike interval distribution (Tolias et al. 2007), an approach also used by a previous study in our lab (Rebesco et al. 2010). We excluded units with an average firing rate during movements <5 Hz to achieve adequate model fits.

Neural Tuning Analysis

For each unit, we fit directional tuning curves relating the firing rate (f) to the direction of movement kinematics (θ) (Georgopoulos et al. 1982) according to Equation 2.

$$f = f_m + DOT\cos(\theta - PD) \quad (2)$$

where f_m is the mean firing rate across directions, DOT represents the depth of tuning, and PD represents the preferred direction. We grouped movements into 45-degree bins according to their mean directions, and fit the model to the average firing rates in the eight bins.

We obtained confidence bounds for the PD estimate using a bootstrapping procedure with 1000 iterations. We excluded from analysis, poorly-tuned units whose bootstrapped PD confidence bounds exceeded ±20 degrees and whose lower 95% confidence bound for R^2 was below 0.5. In practice, this yielded results similar to a requirement of mean $R^2 \geq 0.7$, but ensured that every bootstrap fit used for the PD confidence interval was of acceptable quality.

The typical method for determining PDs uses target-based tuning, where the average neural activity throughout the movement is regressed to the target direction. We compared this target-centric tuning to similar tuning computed in a hand movement coordinate frame, relating average neural firing rate to the actual direction of hand motion. We selected a window of time beginning with the onset of movement (as determined by a velocity threshold crossing of 1.5 cm/s) until the time of peak speed. This window was chosen to focus on the ballistic, planned aspects of reach

and to avoid the later error correction phases of movement, particularly in the Adaptation epoch. To compensate for typical neuronal transmission delays, we shifted neural activity forward by 100 ms to better align with the kinematics.

Neural Tuning Comparison

We tested for changes in PD estimates for each single unit during the Baseline, Adaptation, and Washout epochs. We determined the significance of PD changes by computing confidence bounds on the differences between the 1000 reference Baseline PD estimates obtained with the previously described bootstrapping procedure, and the 1000 bootstrapped PD estimates in each tested epoch. A change was considered significant if the 95% confidence bounds of the differences did not include zero.

We classified cells into five categories based on their tuning behavior in the three epochs (Li et al. 2001). "Kinematic" cells had no change in PD. "Dynamic" cells had a significant change from Baseline PD during Adaptation but returned to Baseline tuning in the Washout. As in prior studies, we also defined two types of "Memory" cells. Type I had a change in PD during Adaptation that persisted in the Washout, while Type II changed only between Adaptation and Washout. A final class, termed "Other", had different tuning in all three epochs. To further characterize the behavior of these cells, we computed a "Memory cell" index, shown in Equation 3, to determine if Other cells more closely resembled Dynamic or Memory cells (Cherian et al., 2013).

$$Mem\ Ind = \frac{|PD_{WO} - PD_{BL}|}{min[|PD_{AD} - PD_{BL}|, |PD_{AD} - PD_{WO}|]} \quad (3)$$

Values greater than 1.0 indicated that the PD in Washout was closer to its value in Adaptation than in Baseline, and thus the cells were more like Memory cells, while values less than 1.0 were more like Dynamic cells.

Musculoskeletal model

We implemented a simple musculoskeletal arm model in order to generate simulated neurons with known relationships to behavior. The monkey's limb was modeled as a two-link planar manipulator fixed at the shoulder (Figure 2.7a). For each session of behavioral data, we used the real hand position and inverse kinematics to compute the joint angles (θ_e *and* θ_s) for the model system according to Equations 4 and 5.

$$\theta_e(t) = \cos^{-1}\left[\frac{p_x(t)^2+p_y(t)^2-L_1^2-L_2^2}{2L_1L_2}\right] \tag{4}$$

$$\theta_s(t) = \frac{[L_1+L_2\cos\{\theta_e(t)\}]p_y(t)-L_2\sin[\theta_e(t)]p_x(t)}{L_2\sin[\theta_e(t)]p_y(t)} \tag{5}$$

where L is the length of each limb segment. We then computed joint torques at the elbow (τ_e) and shoulder (τ_s) using inverse dynamics (Equation 6) derived in (Murray et al. 1994). The equations are reproduced with our notation here:

$$\begin{bmatrix}\tau_s(t)\\ \tau_e(t)\end{bmatrix} = \begin{bmatrix}A+2B\cos[\theta_e(t)] & B\cos[\theta_e(t)]+C\\ B\cos[\theta_e(t)]+C & C\end{bmatrix}\begin{bmatrix}\ddot{\theta}_s(t)\\ \ddot{\theta_e}(t)\end{bmatrix} +$$

$$\begin{bmatrix}-B\sin[\theta_e(t)]\dot{\theta}_e(t) & -B\sin[\theta_e(t)][\dot{\theta}_s(t)+\dot{\theta}_e(t)]\\ B\sin[\theta_e(t)]\dot{\theta}_s(t) & 0\end{bmatrix}\begin{bmatrix}\dot{\theta_e}(t)\\ \dot{\theta_e}(t)\end{bmatrix} \tag{6}$$

$$A = I_s + I_e + m_s r_s^2 + m_e(L_s^2 + r_e^2)$$

$$B = m_e L_s r_e$$

$$C = I_e + m_e r_e^2$$

where for each limb segment, m is the mass, r is the distance to the center of mass (assumed to be half of L), and $I = \frac{1}{3}mL^2$. Note that we did not make precise measurements of the mass of each limb segment, but instead used approximate values based on morphometry studies (Cheng and Scott 2000) and the total body mass of each monkey (Monkey C: 12 kg, Monkey M: 10 kg). We measured approximate segment lengths for each monkey (Monkey C: L_s = 25 cm, L_e = 20 cm; Monkey M: L_s = 20 cm, L_e = 24 cm), though qualitatively our results were consistent even with relatively large changes in these parameters.

Based on the torque computed from actual trajectories, we calculated patterns of muscle activation to drive the modeled limb using shoulder, elbow, and biarticular muscles for both flexion and extension. The moment arms of each muscle were given by the matrix [2 -2 0 0 1.5 -2; 0 0 2 -2 2 -1.5] (Lillicrap and Scott 2013). Given the redundancy in these muscles, we performed static optimization at each time step to find the activations for the desired joint torques. We normalized the required muscle force by the physiological cross-sectional area (PCSA) of each muscle. We used PCSA values from (Cheng and Scott 2000), assuming that each of our modeled muscles represented the sum of the corresponding groups of synergistic muscles. We minimized the total squared activation across the six muscles, with the additional constraint that activations must not be negative. During curl field trials, we modeled the effect of the CF using the parameters and equations described previously. These forces caused an additional torque about each joint that altered the optimal muscle activations needed to drive the limb with the observed kinematics.

We simulated a population of M1 neurons with a firing rate generated using randomly weighted combinations of the activations α_i of the six muscles. For each neuron, we generated a weight ω_i pseudo-randomly for each muscle from a Gaussian distribution with a mean of zero

and a standard deviation of one, modeling balanced excitatory and inhibitory effects on muscle activation (Kujirai et al. 1993; Mariño et al. 2005). We added the resulting sum to a random baseline spiking probability ω_0. These weighted sums were scaled to give physiologically realistic firing rates and treated as the mean rate of a Poisson process, λ, which generated a random spike count *n* for each time bin according to Equations 7 and 8.

$$\lambda = \omega_0 + \sum_i \omega_i \alpha_i;\ -1 \leq \omega_i \leq 1;\ 0 \leq \omega_0 \leq 0.1 \quad (7)$$

$$n|\lambda \sim \text{Poisson}(\lambda) \quad (8)$$

We included a threshold nonlinearity in the form of saturation at a λ value of 1 to avoid physiologically unrealistic firing rates. Neural activity was made to lead the intensity function by a latency drawn for each neuron from a uniform distribution between 70 and 130 ms to simulate the anticipated lag between neural activity and movement used for the recorded units. Importantly, we kept the weights between neurons and muscles constant throughout the session. Thus, the model provided time-varying patterns of neural activity sufficient to drive the limb in the presence of the curl field. We then performed the same kinematic analysis on these simulated neurons as we did on the recorded single units.

RESULTS

Behavioral adaptation to the force field

We analyzed thirteen sessions of data with curl field learning from Monkey C and nine curl field sessions from Monkey M. Of the 22 sessions, 14 used a counter-clockwise (CCW) curl field and 8 used a clockwise (CW) curl field. We saw no clear difference in behavioral or neural data between the two directions, aside from the sign of the effects. The experimental sessions for both monkeys occurred over several months, and each learning session was treated

independently. Additionally, we collected 11 control sessions that had a large number of trials with no imposed fields (five with Monkey C and six with Monkey M). We attempted to maximize the number of Adaptation trials that the monkey would perform in each session to provide more nearly complete adaptation. Consequently, the number of trials in each epoch varied from session to session depending on the monkey's motivation. The monkeys averaged 185 ± 21 (mean ± st.d.) movements in Baseline, 266 ± 35 movements during Adaptation, and 196 ± 100 movements for the Washout. Figure 2.1d shows representative kinematics and the discharge of 30 single units, sorted according to each cell's PD, recorded from Monkey M during the Baseline epoch. Note that the first movement (indicated by the upward arrow at the top of the panel, which was actually a movement away from the monkey) was accompanied by a burst from several units with upward PDs. These same units shut off during the subsequent return movement, and burst again (more weakly) for the later diagonal movements that had an upward component. Though somewhat noisy, similar observations can be made for the downward and leftward PD units.

To quantify behavioral adaptation, we used the "takeoff angle" error metric (see Methods), which compares the direction of initial hand movement to the location of the target. The average speed traces for both monkeys had a characteristic biphasic profile during Adaptation that appeared to represent an initial open-loop movement, followed by a correction. This effect decreased during learning, yet was present to some degree even in the later, more-adapted trials (Figure 2.2a). We thus evaluated the takeoff angle error metric in the brief window between movement onset and the time of peak speed, prior to the apparent error correction. For the CW curl field sessions, we flipped the sign of the error to align it with the CCW-field errors and to simplify analysis and presentation of the results. The CF-induced errors decreased significantly

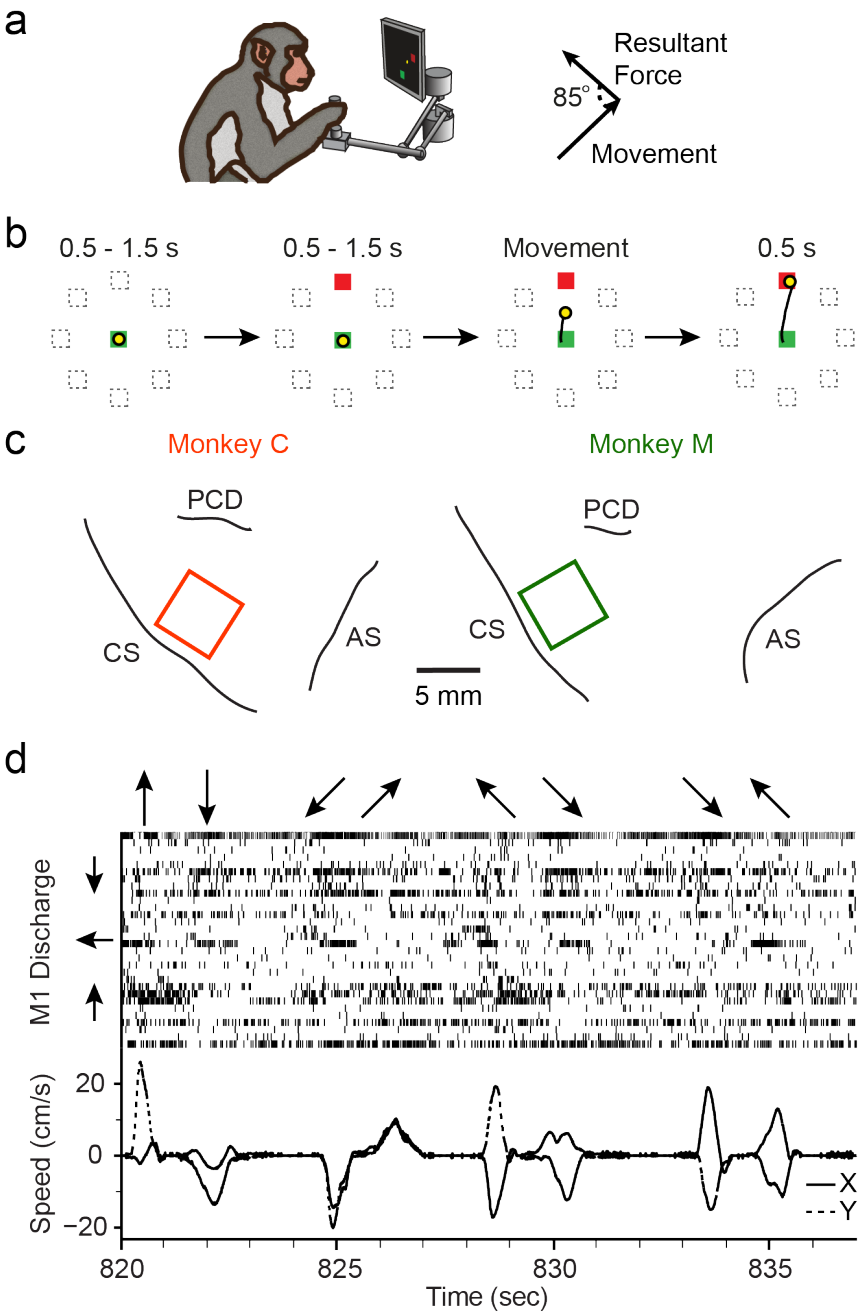

Figure 2.1 | Behavioral tasks and neural recordings. a) Macaque monkeys were seated in a chair and controlled a cursor on a computer screen using a 2-D planar manipulandum that covered a 20cm x 20cm workspace (top left). During the Adaptation epoch, the monkeys made reaches in a velocity-dependent force field applied to the endpoint of the hand, approximately orthogonal to its direction of motion (top right). b) The center-out task required that the monkey reach from the center of the workspace to one of eight outer targets (see Methods). c) We implanted electrode arrays, indicated by the squares, in the arm region of M1 of two monkeys. CS: Central Sulcus; AS: Arcuate Sulcus; PCD: Pre-central Dimple. d) Upper panel shows normalized neural spike rasters for 30 single units recorded simultaneously from Monkey M during a series of reaches. The arrows along the top illustrate the direction of each reach. The units have been sorted by their PDs, indicated approximately by the arrows to the left of the panel. The bottom panel illustrates X (solid) and Y (dashed) components of hand velocity.

for both monkeys with adaptation (Figure 2.2). At the beginning of Washout, both monkeys made significant errors due to the after-effects of adaptation.

Neural responses to altered dynamics

We identified a total of 523 single units across the 22 curl field sessions (Monkey C: 49 ± 7 units per session, Monkey M: 25 ± 10 units per session, mean ± st.d.). Since the recordings were made with chronically implanted arrays, these were certainly not all unique. However, groups of sessions were separated by as much as 18 months, increasing the diversity of the neural populations they represented. Many electrode channels yielded well-isolated single units (Across all channels: Monkey C: 28-49%; Monkey M: 12-35%; of these, rough one-quarter to one-half yielded more than one well discriminated, single unit).

We first examined the change in neural responses of M1 cells from Baseline to the Adaptation and Washout epochs. We computed the mean firing rate of each cell in the window of time from movement onset to peak speed, as described above. We determined the PD of each cell by fitting tuning curves with respect to the direction of hand motion in that window, rather than the target location. This approach did not require the assumption that the monkey reached straight to the target, as would otherwise be necessary. We divided the Adaptation and Washout epochs into three blocks, each one-third of the total movements. We looked initially at spatial tuning in the final block of each epoch when behavior was most likely to be stable.

For the curl field sessions, 36% of cells (190/523) met our requirements for spatial tuning (see Methods). In the control sessions, a higher percentage of cells (66%, 144/226) met the tuning

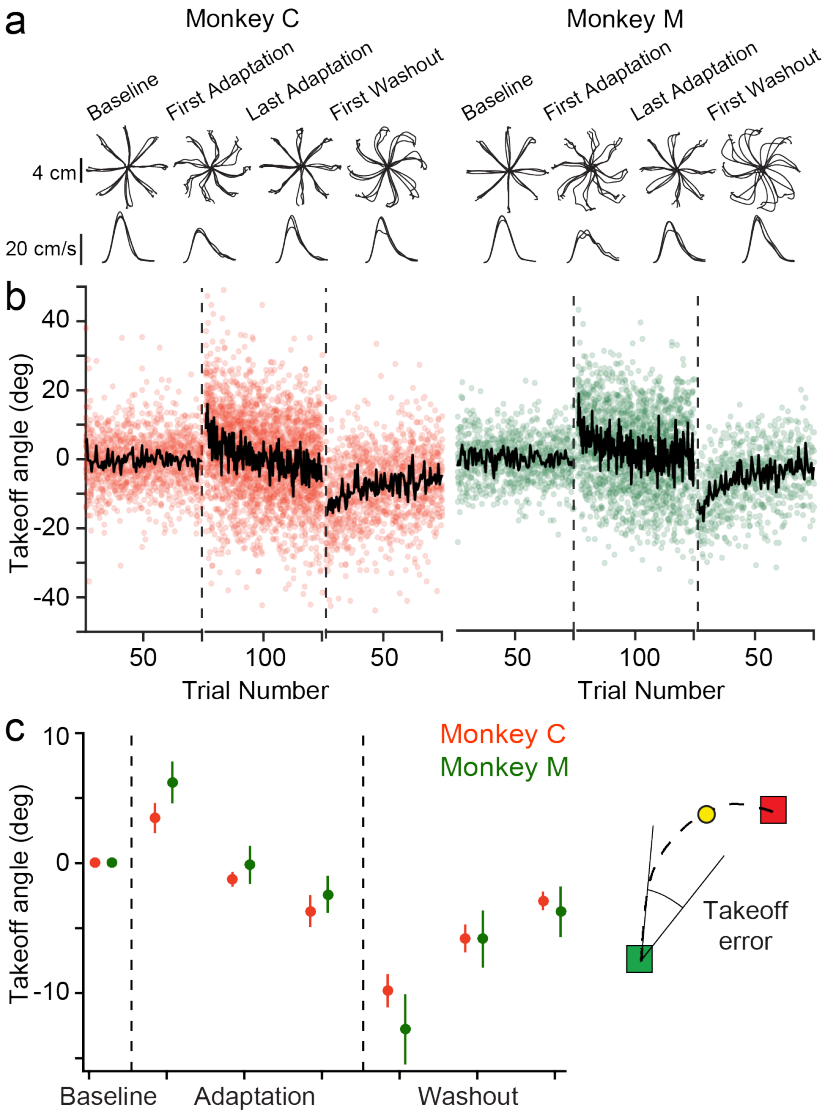

Figure 2.2 | Behavioral adaptation to the force field. a) The top row shows reach position traces for three representative sessions with clockwise curl field perturbations for Monkey C (left) and Monkey M (right). For each monkey, we plotted the final reach to each target during Baseline (left), the first reach and final reach to each target during Adaptation (middle), and the first reach to each target during Washout (right) for each of the three sessions. Trajectories during Baseline and late Adaptation were generally straight, while the early Adaptation and Washout trajectories were markedly curved. The bottom row shows speed profiles, averaged over the trajectories plotted in the top row for each session. b) We used the average takeoff angle (see Methods) at movement onset to characterize behavioral adaptation. The left panel shows these angles for the first 100 Baseline reaches, 200 Adaptation reaches, and 100 Washout reaches for each session with Monkey C. The black line indicates the mean across sessions for each trial (13 with Monkey C and 9 with Monkey M). Takeoff angle increased when the force field was applied and gradually decreased during Adaptation. There were after-effects during Washout that returned gradually to Baseline values. Note that all three epochs contained additional trials beyond those shown here, that varied in number across sessions. c) Behavioral adaptation for Monkey C (red) and Monkey M (green). For the Adaptation and Washout epochs, we averaged the take-off angle in non-overlapping blocks of reaches, each containing one-third of the total number of successful trials. These blocks of trials are used again in the later neural analysis (see Figure 2.4). Plotted data represent mean ± SEM across all sessions.

criteria due to the more consistent trial-to-trial behavior compared to the perturbed trials. Figure 2.3a shows example waveforms and tuning curves throughout the three epochs for six representative cosine-tuned single units. The changes in PD (ΔPD) from Baseline to Adaptation and Baseline to Washout epochs for all cells are shown in Figure 2.3b. As we did with the behavioral adaptation data, we changed the signs of ΔPDs in the CW sessions to align them with the CCW metrics. Across the population, the mean PD shifted by 21° ± 19° (mean ± st.d.) between Baseline and Adaptation, but only 1.6° ± 12° between Baseline and Washout (Figure 2.3b, solid black lines). We divided these cells into five classes based on the significance of their changes across the three epochs. Combined across all sessions, 55% of cells (104/190) were classified as "Dynamic", having a significant change of tuning between the Baseline and Adaptation epochs that returned to the Baseline PD during Washout. All of these cells are located near the horizontal axis with positive PD changes in Figure 2.3b. Fifty of the remaining cells (26%) were clustered near the origin, with statistically unchanged tuning and were classified as "Kinematic" cells.

We also found 13 (7%) cells that fit the definition of Memory I type cells, having PDs that rotated during the Adaptation epoch and kept the new tuning during Washout (Figure 2.3b, red diamonds). In addition, there were nine (5%) Memory II cells (stable tuning across the Baseline and Adaptation epochs but altered tuning in Washout). These cells (Figure 2.3b, green diamonds) were clustered near zero PD change during Adaptation. Lastly, we identified 14 (7%) cells with significantly different tuning in all three epochs, referred to in prior literature as "Other" cells. For these, we computed a "Memory cell index" (see Methods) describing whether the changes more closely resembled Dynamic cell behavior (values less than one) or Memory I cell behavior

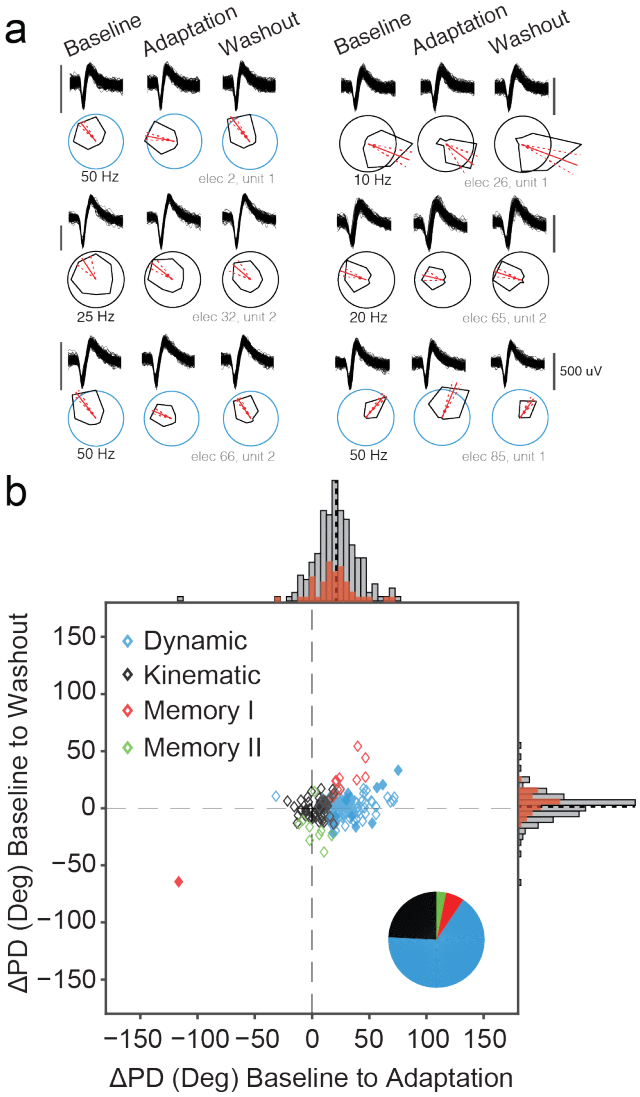

Figure 2.3 | Neural PD changes. a) Waveforms and tuning curves for six representative cosine-tuned units recorded from one session of Monkey C during the Baseline period and the final third of trials in the Adaptation and Washout periods. The black scale bar next to the waveforms indicates 500 uV. The red line indicates the PD, and the dashed red lines denote the 95% confidence bounds on the PD. The scaling circles are colored according to the cell classification. b) Change in preferred direction (ΔPD) from the Baseline to Adaptation epochs (abscissa) and from the Baseline to the Washout (ordinate). Histograms indicate marginal distributions, with gray being all tuned cells, and red being the best isolated cells (SNR > 6). Black dashed lines on the histograms represent the mean change from baseline for each distribution (21° in Adaptation and 1.6° in Washout). Most PDs changed significantly during Adaptation and reverted to their original tuning in Washout. Symbol colors and pie chart (inset) shows the proportion of each cell type for the population. The Dynamic and Memory I cells with solid fills represent Other cells that were reclassified using the Memory Cell index (see Methods).

(values greater than one). Thirteen of the fourteen cells had an index below one and if included in the Dynamic cell group would bring the total Dynamic population to 62% of all cells (117/190). We could find no tendency for the Kinematic and Dynamic cells to be located differentially on the arrays.

Our decision to regress neural activity against the motion of the hand differs from earlier studies (Gandolfo et al. 2000; Li et al. 2001), which averaged neural activity over the duration of the reach and regressed against the target direction. To ensure our results did not depend on the differences between these coordinate frames, we repeated these analyses using the more typical target-direction regressor. We found qualitatively similar PD changes during Adaptation, although the ΔPD distribution was slightly smaller on average and more variable (17° ± 22°). The choice of coordinate system had little impact on our cell classifications, with 29% Kinematic, 46% Dynamic, 13% Memory I or II, and 12% Other.

As a control, we analyzed data from several sessions in which there was no curl-field perturbation, dividing the large number of trials into three equal portions. Under these conditions, we would expect no changes in PDs, and for all single units to be classified as Kinematic. In fact, we found only 81% of cells (117/144) to be Kinematic. There was a small number of Dynamic (6%, 19/144) and Memory I (4%, 5/144) cells, and a slightly larger group of Memory II cells (9%, 13/144). No cells were classified as Other.

The cell classification in Figure 2.3 used a block of trials from the final third of the Adaptation and Washout epochs when adaptation or de-adaptation would have been most nearly complete. We repeated the classification using PD changes in the first and second blocks of trials to assess the time course of neural PD changes. If changing PDs were related to adaptation, we would expect to find smaller changes earlier in the Adaptation epoch. However, that was not the

case. Any cells that rotated did so fully in the first Adaptation block (Figure 2.4a). There was no progressive change in PD throughout the Adaptation or Washout epochs, even as behavioral performance continued to change (compare Figure 2.2c).

We then tested the stability of cell classifications using the first and second blocks of trials within Adaptation and Washout. If Memory cells result from motor learning, we would expect to find fewer early in Adaptation. At the population level, neuron classification rates were similar for all three blocks of trials (Figure 2.4b). However, at the single cell level, there were differences in the classifications of individual units. We compared each unit's classification during the middle and final blocks of the Adaptation and Washout epochs. Of the 190 cells that were tuned in the final block, 155 were tuned in all three blocks. The Dynamic population proved to be fairly robust; 76% (68/89) of Dynamic cells from block 3 were also classified as Dynamic in block 2, well above chance (approximately 20% given the five classes). A slightly lower proportion of Kinematic cells, (66%; 23/35), had consistent classification, while the classification of Memory (4/18) and Other cells (5/13) was close to chance. Together, these results show that short-term motor adaptation is not a product of progressive changes within M1, and that putative Memory cells are likely the result of statistical misclassification.

Magnitude of neural PD changes depends on the instantaneous force

The previous results used PDs computed within a brief window between movement onset and peak speed. However, the speed and amount of force exerted by the curl field varied even within this window. If M1 activity is related to movement dynamics, we should observe larger ΔPD when the monkey experienced larger curl forces. We fit PDs within six overlapping

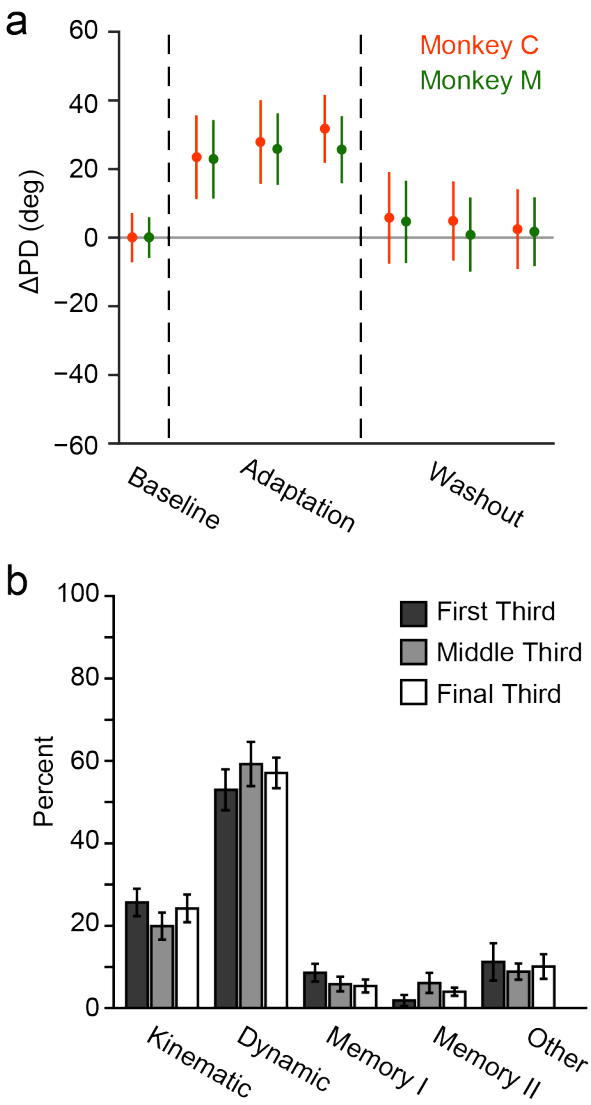

Figure 2.4 | PD changes did not correlate with behavioral adaptation. a) ΔPD (mean ± SEM) of the population as the monkeys adapted for three different blocks of data during Adaptation and Washout. The PDs for Monkey C (red) and Monkey M (green) rotated when the force field was applied, with no progressive changes across these blocks despite the changing behavior (Figure 2.2c). b) Neural population classifications for each of the blocks of Adaptation and Washout trials from Figure 2.4a. Error bars represent standard error of the mean across experimental sessions. There were similar proportions of each cell class within the three blocks.

windows (30% of the movement duration for each trial, typically ~150 ms) from the beginning of movement. The last of these ended approximately at the time of peak speed. The magnitude of the CF force varied throughout this period, which we compared to the ΔPD for all cells that had significant cosine tuning in a given window (Figure 2.5a). There was a strong correspondence between the two measures. However, Figure 2.5a shows single units from the same time windows pooled across multiple sessions and thus removes the correlations with behavior on individual sessions. Each point in Figure 2.5b represents, for a single time window from Figure 2.5a, the mean ΔPD for the population of recorded cells on one session and the corresponding RMS curl force experienced on that session. A linear fit across all data points from both monkeys was highly significant ($r=0.51$, $p\sim0$; black line of Figure 2.5b).

Neural tuning to the "motor plan" is stable during adaptation

The correlation between ΔPD and force suggests that endpoint kinematics may not be the appropriate reference frame with which to describe M1 activity. Therefore, we also expressed M1 activity in terms of the dynamics of the "motor plan", using an approach developed by Mandelblat-Cerf, et al (Mandelblat-Cerf et al. 2011). Briefly, during the window between movement onset and peak speed, we used endpoint velocity to compute the underlying CF force, which we subtracted from the force recorded by the force transducer. The resulting force vector provided an estimate of the net force at the hand that would have been produced by the monkey in the absence of the curl field and serves as a proxy for the movement that the monkeys intended to make, referred to as the "motor plan" (Figure 2.6a). In early CF trials, this motor plan should be directed towards the target, then rotate gradually to counter the CF over the course of adaptation. We computed the average motor plan force vector during the time between

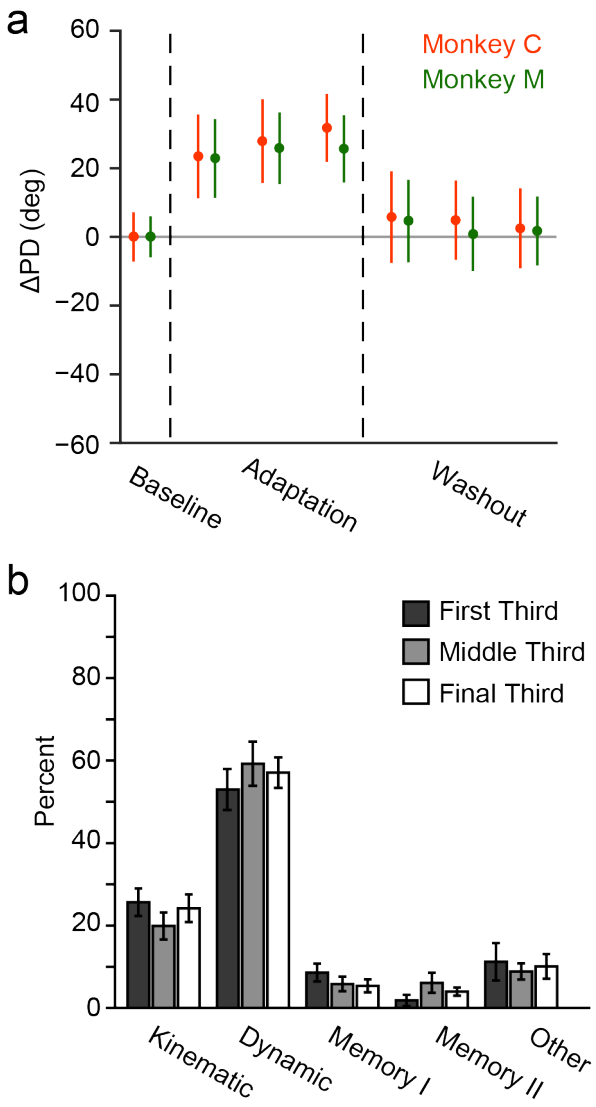

Figure 2.5 | Neural PD changes depend on force. a) We fit PDs for all single units in overlapping, sliding time windows beginning at movement onset and computed ΔPD between Baseline and Adaptation in each window. Values represent mean ± SEM across the cells that were tuned in any given window. This plot compares the time-varying change in RMS force measured at the handle (black lines, scale bar) with the population ΔPD (red for Monkey C, green for Monkey M, left axis) in each window. b) Each data point shows the per-session average of the ΔPD change observed in each of the time windows shown in Figure 2.5a plotted against the difference in endpoint force in that window. The fitted line gives the relationship ΔPD = 6.6° + 26° x ΔForce

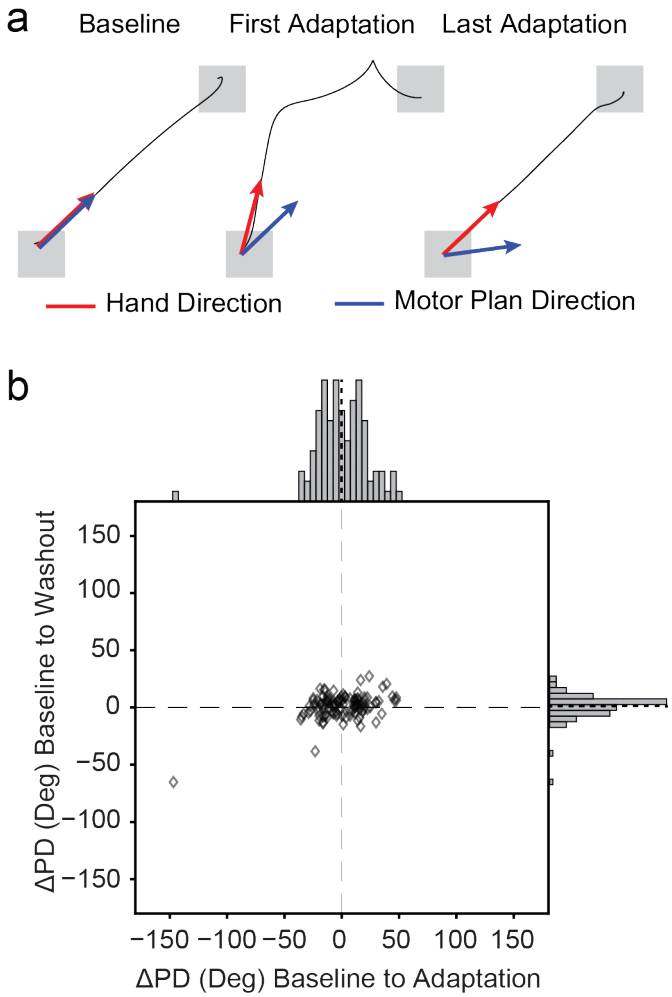

Figure 2.6 | M1 consistently encodes the desired motor plan. a) We identified the direction of the intended force application (blue arrows), or the "motor plan", by subtracting the effect of the curl field from the observed endpoint force vector (red arrows). In Baseline, these two coordinate frames were identical. Initially during Adaptation, the monkey intended to reach to the target, though the direction of the hand motion was altered by the curl field. In late Adaptation trials, the hand reached the target but the compensatory motor plan was rotated toward the direction of the perturbation. b) Despite the large behavioral effects, PD changes in the motor plan coordinate system, plotted as in Figure 2.3b, were centered near zero at the beginning of both adaptation and washout.

movement onset and peak speed. We then fit cosine tuning functions to the firing rate of our M1 units regressed to the direction of this vector. We assessed the change in the resulting PDs between Baseline and Adaptation trials. If this dynamic coordinate frame more accurately captures M1 modulation, the average ΔPD should be significantly smaller than that shown in Figure 2.3b. Indeed, neural tuning in this "motor plan" reference frame was quite stable during Adaptation, with a population ΔPD = 0.8° ± 22° (mean ± st.d.; p=0.95, one-sample t-test), suggesting that M1 activity more consistently relates to the net forces produced by the limb (Figure 2.6b).

Musculoskeletal Model

We repeated the tuning analyses using our simulated, muscle-based neurons and found results qualitatively similar to those of the recorded single units. We generated 500 simulated neurons (Figure 2.7b) for each of the 22 curl-field sessions from both monkeys to give a total of 11,000 neurons. The randomly weighted muscle force inputs allowed cells to have equal activation from antagonist muscle groups. As a result, not all cells produced significant spiking activity. As with the recorded units, we considered only cells that had a firing rate greater than five Hz during the movement period, which removed 1827 cells from further analysis. Of the remaining neurons, 46% (4279/9173) also met our cosine tuning criteria.

The PD changes during Adaptation were broadly distributed, with a mean significantly larger than that of the recorded units (68° ± 22°; Figure 2.7c, top marginal distribution), while the distribution of model neuron PD changes during Washout was a bit narrower (-1° ± 7°; Figure 2.7c, right marginal distribution). The PDs rotated immediately due to the effect of the force perturbation, which altered the relationship between endpoint kinematics and joint torques

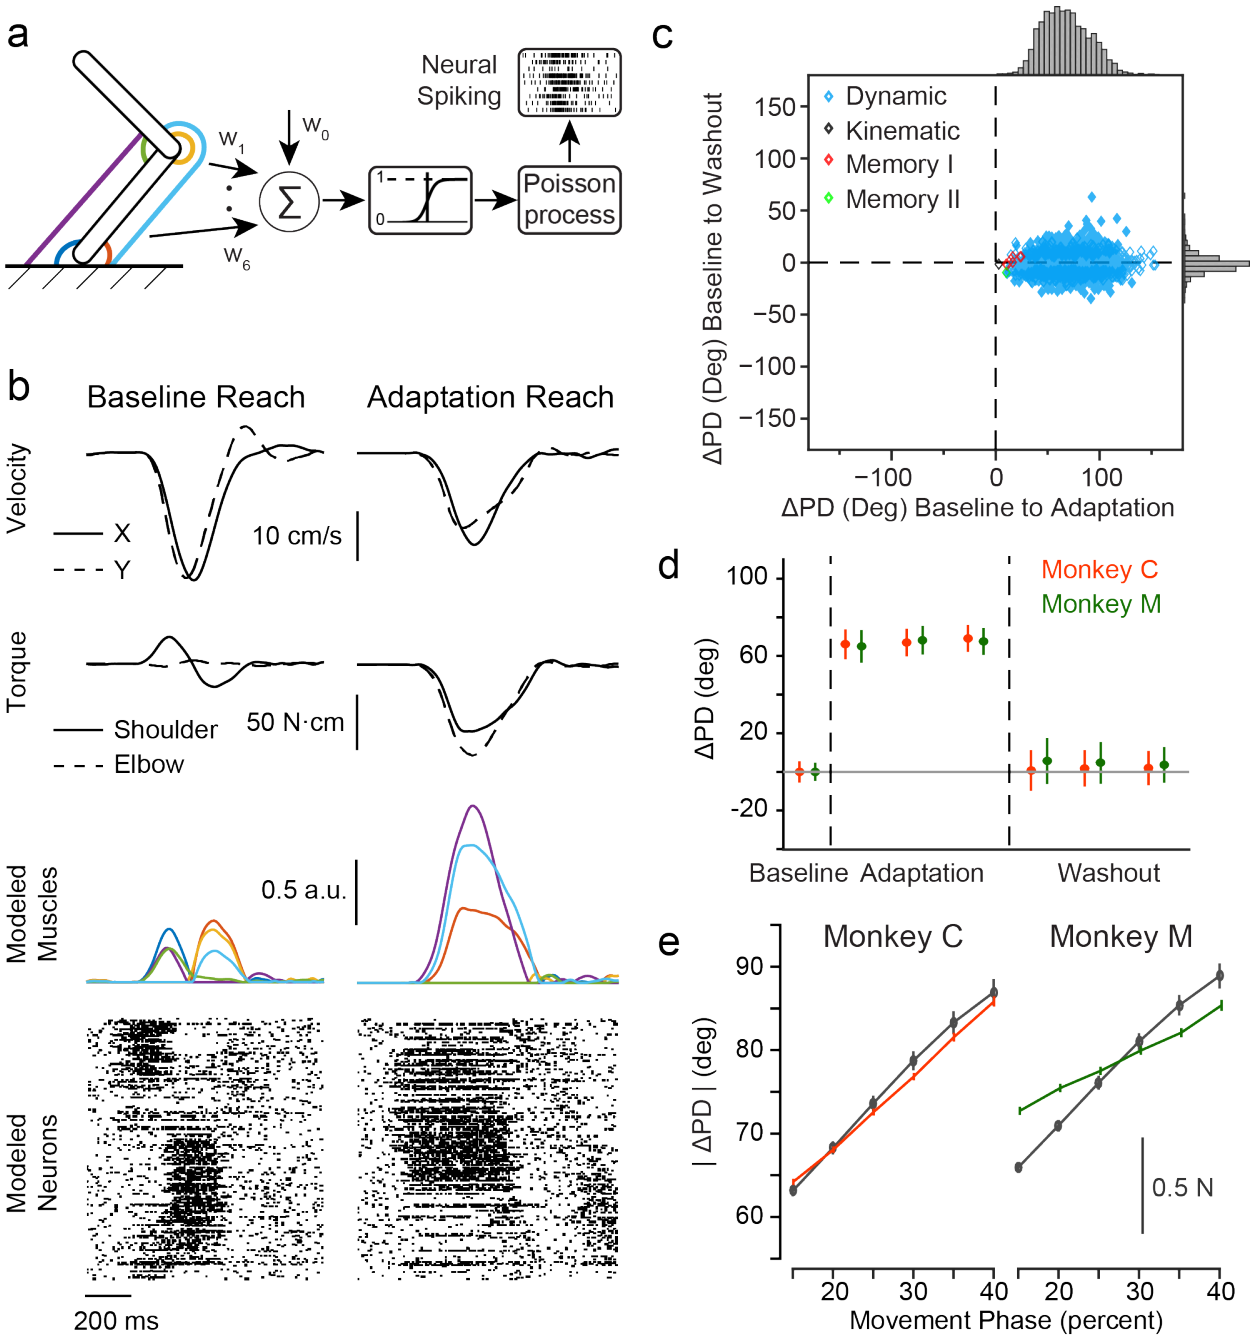

Figure 2.7 | Simulated dynamics neurons reproduce properties of recorded units. a) Schematic of biomechanical model and neural generative model. The limb was modeled as a two-link manipulandum controlled by six muscles (colored lines). Neural activity was generated from weighted sums of these muscle activities (see Methods). b) Biomechanical model examples from one session with Monkey M. The left column is a representative reach to the -135° target in Baseline, while the right column is a reach to the same target with the curl field. Top row: endpoint velocities in the x (solid) and y (dashed) directions. Second row: joint torques for shoulder (solid) and elbow (dashed) computed using the model. Third row: muscle activations from the biomechanical model. Colors correspond to the muscles in Figure 2.7a. Bottom row: raster plots showing the spiking activity of a subset of the simulated neurons sorted by PD. c) Same as Figure 2.3b for simulated neurons. d) Same as Figure 2.4a for simulated neurons. e) Same as Figure 2.5a for simulated neurons.

(Figure 2.7d). The ΔPD was thereafter constant during Adaptation, with no temporal correlation to behavior. Although the majority of cells were classified as Dynamic (87%, 3726/4279), there were many Other type cells (13%, 543/4279) and far smaller numbers of Kinematic (2/4279) and Memory (8/4279) cells. Since the neural firing rates were generated using fixed dynamical relationships, any change in spatial tuning in the Washout that lead to Other or Memory cells must have been due to a misclassification resulting from chance and/or different kinematic behavior rather than a change in underlying neural processes. Assuming all Memory and Other cells were the result of altered behavioral statistics in Washout, then 13% (551/4279) of simulated neurons were classified erroneously, somewhat less than the proportion observed in the recorded units (19%, 36/190). The PD changes of modeled neurons also had a dependence on time (and force) throughout the reach (Figure 2.7e), as did the real units (Figure 2.5a).

DISCUSSION

Summary

We investigated the spatial tuning of M1 neurons as monkeys made reaching movements against an externally imposed curl field that altered limb dynamics. We showed that the majority of single units were Dynamic, having PDs that changed when the field was imposed and returned to their Baseline tuning when it was removed. Importantly, the PDs of the units showed no progressive change even during ongoing behavioral adaptation, but instead, the size of ΔPD was well correlated with the magnitude of the CF force. This was true over the course of movements (Figure 2.5a) and across sessions (Figure 2.5b). When instead computed in terms of dynamic, "motor plan" coordinates, there was no change in PD. Lastly, we used a simple musculoskeletal model to generate a population of neurons with fixed relationships to joint torques. These

modeled neurons had a broad distribution of ΔPDs that was qualitatively similar to that of the recorded units. It has been proposed that learning results by adapting the brain's internal inverse model of limb dynamics (Shadmehr and Mussa-Ivaldi 1994). Overall, our results suggest that short-term adaptation of such an internal model likely occurs before the inputs to these M1 neurons, which themselves maintain a consistent relationship to the dynamics of movement.

Tuning changes do not suggest internal model adaptation within M1

The progressive learning of new motor skills must be encoded within the brain, presumably by changes in discharge that mimic the time course of the behavioral adaptation. It is reasonable to expect that M1 participates in this process. M1 tuning curve changes, especially those of Memory cells, have been thought to be an expression of this process (Gandolfo et al. 2000; Li et al. 2001; Arce et al. 2010b; Richardson et al. 2012). This theory suggests that changes in neural discharge are the result of altered functional properties either within M1 or in downstream neural circuits, embodying an adapted internal model. It is important to note that although these prior studies expressed neural activity within the classic kinematic coordinate frame, they did not assume that M1 acts as a kinematic controller. Indeed, they found that the activity of many M1 cells reflected movement dynamics, with a range of effects that was suggestive of a kinematic to dynamic transformation (Li et al. 2001). The critical difference between those studies and our own, however, is the suggestion that these properties are altered during the hour-long process of adaptation, for which we can find no compelling evidence.

In our experiment, there were progressive behavioral changes within sessions (Figure 2.2) and clear behavioral aftereffects during Washout. If changes within M1 mediate this short-term motor learning through internal model adaptation, we would predict that the magnitude of the

change in neural tuning or the proportion of Memory cells should gradually change as the monkey adapts. However, although the CF affected the great majority of the recorded and simulated neurons, we found no progressive change in their PDs throughout the Adaptation epoch (Figures 2.2 and 2.4). Thus, we interpret the PD changes as a direct consequence of the altered dynamics, rather than the subsequent process of adaptation. Although the PD changes appeared at the earliest point at which we were able to examine them, a limitation of the tuning analysis is that we must average over a relatively large number of trials to achieve adequate statistical power. Much of the behavioral adaptation occurred early, within the first block in the session. However, it is important to note that there were still significant behavioral changes even between blocks two and three (Figure 2.2c). If the tuning changes were the consequence of learning, there should be evidence of further tuning changes across all three blocks.

Furthermore, we saw little evidence of Memory cells. Although there was no progressive change in their number during learning, the classification of individual Memory cells did change unpredictably. Finally, we even found a small number of Memory cells in the control sessions, in which there was no applied curl field. A similar study that recorded the activity of single units in the supplementary motor area also detected Memory cells in control sessions approximately as often as during the learning sessions (Padoa-Schioppa et al. 2004). These observations raise the concern that Memory cells may be the result of measurement noise rather than a robust change in neural tuning. There are several potential explanations for the differences between our results and those in the prior studies. First, we used a stronger force field (0.15 compared to 0.07 Ns/cm). Since neural PDs rotate in proportion to the magnitude of force (Figure 2.5), the larger effect size likely made our classifications more robust. Additionally, our experiments contained more movements during Adaptation (266 on average, versus ~160-200 in prior studies), allowing

us to exclude more of the earlier, less adapted trials. More movements likely also allowed for more complete adaptation and stable behavior. If Memory cells arise from noisy estimates, then larger fields, more trials, and more stable behavior should decrease the chance of misclassification.

Although we classified cells as Dynamic and Kinematic for consistency with prior literature, the units were not distributed bimodally in a manner that would suggest two distinct classes. Rather, in both the recorded and simulated neurons there was a continuum of ΔPD magnitudes. Those recorded units with small rotations may represent muscles that were only minimally affected by the CF. Alternatively, this continuous distribution may well represent an internal model that transforms motor commands from kinematic to dynamic coordinates (Fetz 1992; Kalaska and Crammond 1992; Shadmehr and Mussa-Ivaldi 1994). In that case, however, it appears that these neurons are not the source of the short-term behavioral adaptation in our experiments. Instead, the internal model represented by the outputs of this neural circuit remained stable, suggesting that compensation for the modified external dynamics occurred at a higher level and was simply transmitted to M1.

Neural tuning is related to movement dynamics

If the tuning changes in M1 in the presence of a curl field do not correspond to behavioral adaptation, how are they to be explained? We propose that PD changes are primarily due to the direct effect of force on M1 firing rates (Evarts 1968; Scott and Kalaska 1995, 1997; Sergio and Kalaska 2003; Sergio et al. 2005; Gupta and Ashe 2009). Our evidence for this includes the very rapid PD rotation with the onset of the CF, the strong correlation between ΔPD and force (Figure 2.5), and the stable PD representation that occurred for the dynamic motor plan. Although the

discharge of these units changed during the course of learning to drive the adapted behavior, their spatial tuning throughout the Adaptation epoch was quite stable. Importantly, our conclusions do not need to assume that single M1 neurons necessarily encode specific variables (e.g., endpoint kinematics, endpoint force, joint angles, muscle activations) and indeed many studies of activity in M1 have suggested that individual neurons need not represent any particular movement covariate (Fetz 1992; Churchland and Shenoy 2007; Elsayed et al. 2016). We focused on a brief window, early in movement before online error correction, to simplify the interpretation of our results. Yet, there remains the possibility that some of the PD changes were the result of changes in short-latency proprioceptive feedback to M1 (Scott et al. 2015). Indeed, it has been proposed that proprioceptive feedback to M1 is necessary to learn curl field perturbations (Wolpert et al. 1995; Mathis et al. 2017), and that the adapted internal model is shared for both feedforward control and feedback control (Wagner and Smith 2008). Although afferent inputs likely influence the shape of at least some tuning curves, this influence is equivalent to that of premotor inputs with respect to our argument about the role of M1 in short-term adaptation. The rapid and sustained PD changes throughout adaptation, whether the result of altered premotor or afferent inputs to M1, suggest that the dynamical relation between M1 activity and movement was altered only by the addition of the load force, not subsequent adaptation. The very similar results in our musculoskeletal model, which lacked proprioceptive feedback, further support the conclusion that PD changes result primarily from the altered dynamics of the task.

It would be enlightening to pursue a more detailed model of the relation between M1 activity and muscle force, but that is well beyond the scope of this study. With their fixed generative relationship to dynamics, the simulated neurons exhibited tuning changes remarkably

similar to those of the recorded M1 units. There were, however, some differences. Notably, our model had a large percentage of Dynamic cells, many Other cells, and virtually no Kinematic or Memory cells. These differences are likely due in large part to two factors. First, there were larger model PD changes during Adaptation, and generally smaller changes during Washout compared to the real neurons, yielding more robust statistical changes underlying the Dynamic classification. Second, there was much less variability in the discharge of the simulated neurons compared to the real units, leading to narrower confidence intervals on the PD estimates. The activity of real neurons likely includes a number of sources of signal-dependent noise, or variability that is not strictly stochastic but is also not well-captured by our cosine tuning model. Such sources could include the posture or load dependence of neurons (Scott and Kalaska 1995; Sergio et al. 2005), the control of muscles not included in the model or not strongly modulated in the task, and other activity that is not related explicitly to the motor outputs (Fetz 1992; Churchland and Shenoy 2007). This combination of factors caused virtually all modeled neurons to be statistically different in Adaptation from both other epochs. This greatly decreased the likelihood of observing Kinematic cells. It also meant that most cells that happened to differ between Baseline and Washout became Other (all epochs different) rather than Memory. A smaller PD change in Adaptation would have increased the probability that the Other cells would instead be classified as Memory, since no change would be observed in Adaptation.

The broader distribution of modeled PD changes may have been the result of simplifications in our biomechanical model and our generative neural model. First, we assumed that motion was limited to two joints in the horizontal plane, while the actual posture also involved movements of most of the limb's seven degrees of freedom, many of which would not have contributed directly to the planar tuning curves. Second, our neurons were based on a small number of modeled

muscles, which are greatly simplified compared to the number of muscles that actually move the limb. This simple model did not include force-length and force-velocity characteristics, or posture-dependent pulling directions, and it minimized any co-contraction. Including these features would likely increase the number of single units with less than maximal responses to the curl field, thus providing a population that more closely resembles the recorded M1 cells.

It is important to note that the model was not intended to recreate the monkeys' behavior precisely; we used real behavioral kinematics simply to provide realistic movement speeds, trajectories, and effects of behavioral adaptation. We generated neural activity directly from the activity of the modeled muscles, maintaining a fixed dynamical relation between the two. Although the connectivity was random, there is much evidence that the real connectivity of corticomotoneuronal cells is more structured, with many cells favoring reciprocal activation of agonists and antagonists (Cheney et al. 1985). In practice, the majority of the modeled cells that were cosine-tuned had this structure. Although the firing of many M1 cells is closely related to forces at the hand and to patterns of muscle activity (Evarts 1968; Ashe 1998; Cabel et al. 2001; Sergio et al. 2005), a number of studies have shown that the discharge of M1 neurons correlates with many aspects of behavior beyond the activation of muscles (Fetz 1992; Churchland and Shenoy 2007). Our modeled neurons lack activity related to intrinsic network dynamics (Shenoy et al. 2013), movement planning and preparation (Alexander and Crutcher 1990), or afferent feedback (Asanuma et al. 1979; Cheney and Fetz 1984). As such, our modeled neurons cannot be expected to replicate the precise patterns of activity that would be observed in a real motor cortical population. Nevertheless, despite these simplifications, the generated neurons behaved in a qualitatively similar manner to the recorded units. The results from our experimental and

modeled data are consistent with the conclusion that M1 activity encodes information related to the dynamics of the intended movement that remains unchanged during adaptation.

Comparison with prior studies

Including the original study in 2000 (Gandolfo et al. 2000), several studies have proposed that short-term tuning changes in M1 during curl field learning reflect modification of an internal, inverse model of limb dynamics that drives the behavioral adaptation (Li et al. 2001; Arce et al. 2010b; Richardson et al. 2012). Another study reached a conclusion more like ours, that M1 activity was consistently correlated with the "motor plan" during adaptation (Mandelblat-Cerf et al. 2011). Their motor plan analysis, applied to our data, yielded similar results. However, that study did not include a Washout period, so they could not assess their neurons for any memory effects. Furthermore, they used a unique design, in which the CF was applied to a single reach direction (as did the 2010 Arce study). This difference makes comparison of both these studies with the earlier studies from Bizzi's group difficult.

Another important later study used chronically implanted electrode arrays, allowing some neurons to be tracked between sessions. That study replicated many of the earlier within-session findings, but also reported some neurons with persistent, cross-session PD shifts that were taken as evidence of long-term learning. However, they also acknowledged that the M1 tuning changes they described did not necessarily mean a change in the functional properties of M1 cells, but instead likely reflected altered recruitment by higher order brain areas, as we have concluded. Nonetheless, they interpreted the changes in PD and the presence of Memory cells as signatures of an adapted internal model, not simply the dynamics of the added load. Neither our results, nor those of the previous study from our group (Cherian et al. 2013), support this conclusion.

The discrepancies between Cherian et al. and the earlier studies have also been puzzling, though a number of differences in the design of the studies make direct comparison difficult. First, Cherian et al. compared movements between two oppositely directed curl fields. Interference between the two perturbations may have disrupted the learning process (Brashers-Krug et al. 1996), raising the possibility that the lack of Memory cells was simply a consequence of incomplete adaptation. Our experiment compared curl fields to null fields, as did the earlier studies. Second, Cherian et al. used a random-movement task with dynamics that were quite different from the standard center out movements used in prior literature. The lack of explicit planning and greater movement complexity may have led to different rates or signatures of learning (Smith et al. 2006; Sheahan et al. 2016). For the present study, we used the center-out task, which also allowed us to compare target-based and hand-movement reference frames. Finally, we added a biomechanical model, with known, fixed dynamics. Both our recording and modeling results remain at odds with the conclusions of the earlier studies (Gandolfo et al. 2000; Li et al. 2001; Arce et al. 2010b; Richardson et al. 2012).

Behavioral adaptation is mediated by altered recruitment of M1 neurons

Motor adaptation is a complex process that is dependent on many factors, including the type of perturbation (Krakauer et al. 1999), the perturbation schedule (Orban de Xivry et al. 2011), and even whether the subject has explicit knowledge of the perturbation (Mazzoni and Krakauer 2006). It is likely that motor learning involves the concerted efforts of multiple cortical and sub-cortical areas (Kawato 1999). Our results suggest that short-term adaptation to a curl field is not mediated by persistent plastic changes in the functional characteristics of the units we recorded in M1. There remains the possibility that adaptation occurs upstream of our recording but still

within M1. However, two studies investigating curl field learning by humans used repetitive transcranial magnetic stimulation to disrupt M1 and concluded that it is not a critical site for short term motor learning (Richardson et al. 2006; Overduin et al. 2009). The greater rate of adaptation for humans than monkeys and the possibility that humans may make more use of altered strategy than monkeys complicates the cross-species comparison. However, the inactivation results are consistent with the idea that the rapid adaptation to a novel dynamic environment results from changes in the recruitment of M1 neurons by higher cortical areas. Similar studies suggest that the supplementary motor area (SMA) helps to compensate for the altered dynamics, potentially by altering its inputs to M1. However, its responses to the CF varied quite broadly, during both learning and control conditions, suggesting a complex role for SMA (Padoa-Schioppa et al. 2002, 2004). Curl field adaptation is likely mediated as well by inputs from dorsal premotor cortex (Shadmehr and Holcomb 1997; Dum and Strick 2002, 2005), and undoubtedly involves cortical interactions with the cerebellum (Wolpert et al. 1998).

CHAPTER 3

NEURAL MANIFOLDS FOR THE CONTROL OF MOVEMENT

Juan A. Gallego, Matthew G. Perich, Lee E. Miller, Sara A. Solla

FOREWORD

The following chapter is an adapted version of a Perspective article published in the journal Neuron (Gallego et al. 2017). The article appears in a special issue titled "How the Brain Works". In the article, we discuss a new approach to understanding the neural control of movement that is based on population-level connectivity patterns. The article serves not only as a review of existing literature studying neural populations, but also as a novel conceptual contribution to the field. Throughout the text, we attempt to integrate existing knowledge and formalize a new framework for interpreting many experimental observations based on the idea of "neural modes". The ideas contained within started from a long series of intellectual discussions I had with Dr. Juan Gallego, a co-author on the manuscript. For the purposes of this dissertation, I have added an additional section analyzing population activity from M1 and PMd during reaching. With these data, I illustrate some important concepts about the neural manifolds, and the function of M1 and PMd. This chapter provides essential background for the analytical approach that I use in Chapter 4 to understand population-level activity during learning.

ABSTRACT

The analysis of neural dynamics in several brain cortices has consistently uncovered low-dimensional manifolds that capture a significant fraction of neural variability. These neural manifolds are spanned by specific patterns of correlated neural activity, the "neural modes." We discuss a model for neural control of movement in which the time-dependent activation of these neural modes is the generator of motor behavior. This manifold-based view of motor cortex may lead to a better understanding of how the brain controls movement.

INTRODUCTION

Since the work of Herbert Jasper (Jasper et al. 1958) and Ed Evarts (Evarts 1968), cortical function has been studied by recording single neuron activity while animals perform a variety of behaviors, including decision making (Newsome et al. 1989), sensation (Wurtz 1969), and movement (Humphrey et al. 1970; Georgopoulos et al. 1982).

In the motor system, the main focus of this article, single neuron studies typically involved recordings during repeated, stereotypical movements. Many of these experiments sought explicit representations that relate single-neuron activity to specific movement covariates, including but not limited to target position, endpoint and joint kinematics, endpoint forces, and muscle activity (Evarts 1968; Humphrey et al. 1970; Thach 1978; Georgopoulos et al. 1982; Morrow and Miller 2003). Although some of these efforts involved the decoding of population activity (Georgopoulos et al. 1982), they were restricted to models of non-interacting neurons whose individual activity was associated with specific movement covariates.

However, some of these studies also identified single neurons whose activity did not represent movement parameters (Fetz 1992; Churchland and Shenoy 2007; Scott 2008). If

neurons in primary motor cortex (M1) were to represent movement parameters, those representations ought to be most evident in corticomotoneuronal (CM) cells, which make direct connections onto spinal motoneurons (Fetz 1992). Yet, many of these CM cells do not represent any specific movement covariate (Fetz et al. 1989).

The ultimate role of M1 is to generate movement, not to represent it (Scott 2004; Cisek 2006; Churchland et al. 2012); thus, it is not surprising that many M1 neurons do not relate to any single movement covariate. The search for representations at the single-neuron level might actually divert us from understanding the neural control of movement. Early neural network simulations indicated that individual neurons need not explicitly encode movement covariates when the goal of M1 population activity is to generate realistic muscle activation patterns (Fetz 1992).

The role of neurons that do not explicitly represent any movement covariate can be explained by recent work based on optimal feedback control theory, which postulates that the goal of motor cortex is to produce a desired movement and force, taking into account the state of the muscles. This hypothesis avoids the need for explicit representation of movement covariates by single neurons, though some neurons may still represent movement covariates or high level task characteristics as a byproduct of the necessary control signals (Todorov 2000; Scott 2008).

Recent and accelerating technical developments provide the experimental tools for monitoring the activity of large numbers of neurons, as well as the statistical and modeling tools for analyzing how these neural populations perform the computations necessary to plan and execute movement (Gao and Ganguli 2015). The challenge of understanding the neural control of movement by analyzing neural population activity is formidable, as population activity in any specific area not only reflects its intrinsic dynamics, but must also respond to its inputs and

generate output projections based on the computations being performed (Sussillo et al. 2015). A simplification arises from the fact that neural computations are based on the joint activity of interconnected neurons (Fetz 1992; Hatsopoulos et al. 1998; Shenoy et al. 2013); the resulting population activity is thus likely constrained by the connectivity of the underlying network.

Here we argue that the underlying network connectivity constrains the possible patterns of population activity (Tsodyks et al. 1999; Sadtler et al. 2014; Okun et al. 2015), and that the possible patterns are confined to a low-dimensional manifold (Stopfer et al. 2003; Yu et al. 2009) spanned by a few independent patterns that we call *neural modes*. These modes capture a significant fraction of population covariance. It is the activation of these neural modes, rather than the activity of single neurons, that provides the basic building blocks of neural dynamics and function (Shenoy et al. 2013; Sadtler et al. 2014; Luczak et al. 2015).

We thus propose a generative model of the activity of individual neurons based on the activation of neural modes, and explain how the parameters of the model can be identified using dimensionality reduction methods. We then review work showing that these neural modes span task-specific neural manifolds in premotor and motor cortices. We propose that neural manifolds spanned by a surprisingly small number of neural modes are likely to simplify the neural control of movement, and speculate on the potential learning mechanisms underlying the emergence of this low-dimensional organization.

FROM SINGLE NEURONS TO NEURAL MANIFOLDS

Current multi-electrode arrays (MEAs) allow for the simultaneous recording of about a hundred neurons. This is many more than the small numbers recorded with single electrodes, but still a tiny fraction of the total number of neurons involved in movement generation. Despite this

limitation, Brain-Machine Interfaces (BMIs) based on these MEAs are able to predict reasonably well many behavioral variables (Paninski et al. 2002; Carmena et al. 2003; Ethier et al. 2012).

What is the underlying reason for this success? Intuitively, it is the high degree of correlation and redundancy across individual neural activity. This intuition has been recently made precise in elegant arguments on the low dimensionality of the stereotypical motor behaviors used in most motor control studies (Gao and Ganguli 2015). The relatively small number of independent signals needed to control behavior during the execution of such tasks only requires a small number of independent neural signals. These neural signals are the *latent variables* (Cunningham and Yu 2014) that describe the dynamics of the *neural modes*.

The participation of individual neurons in neural modes is illustrated in Figure 3.1a. Note that each neural mode includes a large fraction of the neurons in the population, and that a given neuron can participate in several neural modes. In this view, the time-dependent activity of individual neurons is simply a reflection of the latent variables (Figure 3.1b) (Macke et al. 2011; Kobak et al. 2016; Kaufman et al. 2016). Consider the neural space in Figure 3.1c; each axis represents the activity of one of the N recorded neurons (here, N=3). Assuming that network connectivity constrains the possible patterns of population activity (Tsodyks et al. 1999; Sadtler et al. 2014; Okun et al. 2015), the population dynamics will not explore the full high-dimensional neural space, but instead remain confined to a low-dimensional surface within the full space, the *neural manifold*. In the simplest linear case, the neural manifold is flat, as the hyperplane in Figure 3.1c, spanned by the two neural modes, $u_1$and u_2.

This geometrical picture illustrates a possible generative model for the dynamics of individual neurons: the activity $n_i(t)$ of the ith neuron, $1 \leq i \leq N$, results from a linear combination of latent variables $L_j(t)$ plus additive noise ε_i:

$$n_i(t) = \sum_j w_{ij} L_j(t) + \varepsilon_i \quad (1)$$

Here, $L_j(t)$ is the *j*th latent variable, the time-dependent activation of the *j*th neural mode. Each latent variable results from projecting the neural population activity onto the corresponding neural mode. The coefficient w_{ij} in the linear combination quantifies the contribution of the *j*th latent variable to the activity of the *i*th neuron. These *participation weights* relate to the internal connectivity of the network (Okun et al. 2015). The noise term ε_i represents intrinsic neural noise, and potentially other processes not accounted for in the model. By construction, neural population activity remains within the neural manifold except for small fluctuations (see how close the actual black trajectory is to the gray trajectory projected into the manifold in Figure 3.1c).

Dimensionality reduction techniques allow us to study neural population dynamics by finding a set of neural modes that span the neural manifold and identify relevant population features (Cunningham and Yu 2014). Common linear techniques for dimensionality reduction, such as principal component analysis (PCA) and factor analysis (FA), identify neural modes as dominant patterns of covariation across neurons and yield the parameters of the generative model (Eq. 1; Figure 3.1b). As an illustration, we show that neural data recorded during an isometric wrist task (Figure 3.2a,b) is largely accounted for by the latent variables in Figure 3.2c. The low-dimensionality of the neural manifold follows from the rapid increase of the explained variance with the number of neural modes (Figure 3.2d).

NEURAL MANIFOLDS: A FRAMEWORK TO STUDY NEURAL CONTROL OF MOVEMENT

The concept of the neural manifold and its associated latent variables has been used in a series of recent studies that replace the search for movement representation by single neurons to

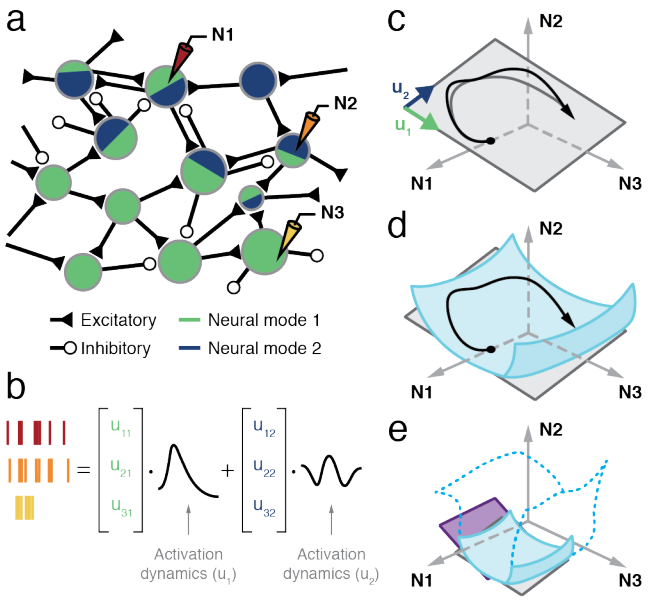

Figure 3.1 | The neural manifold hypothesis. a) Latent variables as a generative model for population activity. The relative area of the blue/green regions in each neuron represents the magnitude of the contribution of each latent variable to the neuron's activity. b) Spikes from three recorded neurons during task execution as a linear combination of two latent variables. c) Trajectory of time-varying population activity in the neural space of the three recorded neurons (black). The trajectory is mostly confined to the neural manifold, a plane shown in gray, spanned by the neural modes u_1 and u_2. d) A curved, nonlinear neural manifold, shown in blue. Linear methods would capture a flat, local approximation to a small task-specific region of the manifold. e) Linear manifolds for two different tasks shown as gray and purple planes. Are these two planes local linear approximations to different regions within a large, continuous manifold (transparent surface with blue contour), or are they distinct task-specific manifolds that may or not share neural modes?

consider instead movement planning and execution based on the activation of a few neural modes (Churchland and Shenoy 2007; Santhanam et al. 2009; Churchland et al. 2010b, a, 2012; Ahrens et al. 2012; Kaufman et al. 2014; Sadtler et al. 2014; Bruno et al. 2015; Overduin et al. 2015; Sussillo et al. 2015; Elsayed et al. 2016; Michaels et al. 2016).

One of the earliest findings of a neural manifold for movement control comes from Shenoy and colleagues (Santhanam et al. 2009), who analyzed population activity recorded with an MEA implanted in the arm area of dorsal premotor cortex (PMd) during a delayed center-out reach task. Single-neuron activity in PMd correlates with the direction toward the end point of an upcoming reach movement (Riehle and Requin 1989; Shen and Alexander 1997). Shenoy and colleagues used FA to obtain neural modes that accounted for the observed shared variance of individual neurons. They found that a three-dimensional manifold sufficed to identify target-specific clusters of latent activity during the delay period (Figure 3.3a).

A subsequent study (Churchland et al. 2010b) showed a systematic decrease in the trial-to-trial variability in the neural dynamics of both PMd and primary visual cortex (V1) following stimulus onset, as demonstrated in two-dimensional visualizations of the latent variables (Figure 3.3b). The low-dimensional manifold was characterized using Gaussian Process Factor Analysis (GPFA), a method that combines FA with temporal smoothing through a Gaussian kernel, to extract the low-dimensional trajectories defined by the latent variables during individual trials. The method was proposed and compared to static methods like PCA and FA in an earlier paper (Yu et al. 2009) that identified variability reduction following target presentation in PMd data.

The notion of a neural manifold and its associated latent variables was subsequently used by Churchland, Shenoy, and colleagues (Churchland et al. 2010a, 2012) to explain how neural activity in both PMd and M1 during movement planning (Riehle and Requin 1989) does not

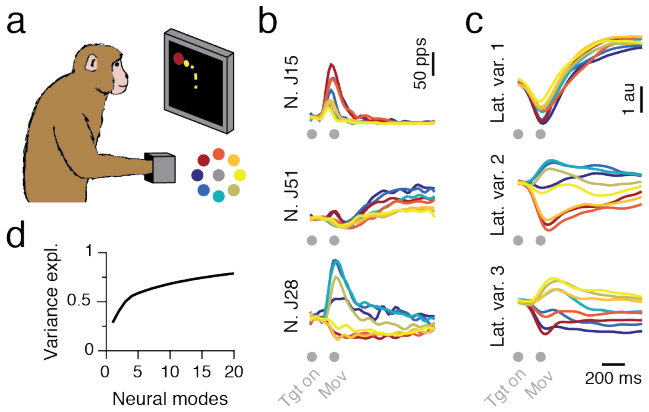

Figure 3.2 | Latent variables for an isometric center-out task. a) Monkey performing an isometric eight-target center-out wrist task. The targets are color coded (inset). The same target color code is used for the curves shown in panels b and c. b) Peristimulus time histogram for three out of the 68 neurons recorded in a single session. c) Latent variables track the activation of the three leading neural modes, identified with PCA. In panels b and c, target presentation ("Tgt on") is followed by movement onset ("Mov") after about 140 ms. d) Variance explained as function of the number of PCA neural modes. Ten modes sufficed to account for about 70% of the variance.

generate movement during the delay preceding a go signal (Cisek and Kalaska 2005). To explain how M1 could prepare movement without causing it, the same group (Kaufman et al. 2014) identified a six-dimensional neural manifold using PCA, then built a linear model that related these latent variables to three *muscle synergies* (d'Avella et al. 2003; Tresch and Jarc 2009), also identified by PCA. Based on this linear model, they divided the neural manifold into a *potent* space, whose activity controls muscle activity, and a *null* space, whose activity does not affect muscle activity (Figure 3.3c). They showed that preparatory activity lies in the null space; this condition-dependent activity provides an initialization from which the population dynamics evolve to generate the desired movement (Churchland et al. 2010a, 2012). In a recent follow-up study, the same group expanded this analysis to show that preparatory and movement activity lie in orthogonal spaces within the manifold, and that population dynamics evolve from one to the other (Elsayed et al. 2016) (Figure 3.3d).

The separation between potent and null spaces was also used by Slutzky and colleagues to investigate the long-term stability of BMIs (Flint et al. 2016). They found that the stability of all recorded neurons was not uniformly necessary to achieve stable BMI control, and showed that neural activity in the potent space was significantly more stable than neural activity in the null space. This finding provided evidence supporting the postulate that optimal feedback control allows the brain to control activity in the potent space, while activity in the null space can vary from trial to trial (Todorov and Jordan 2002).

Together, these studies strongly support the existence of low-dimensional manifolds in motor cortices. The notion that latent variables may constitute the building blocks of population activity allows us to consider the activity of individual neurons as one-dimensional samples of the manifold dynamics. These studies also suggest that the constraints embodied by the neural

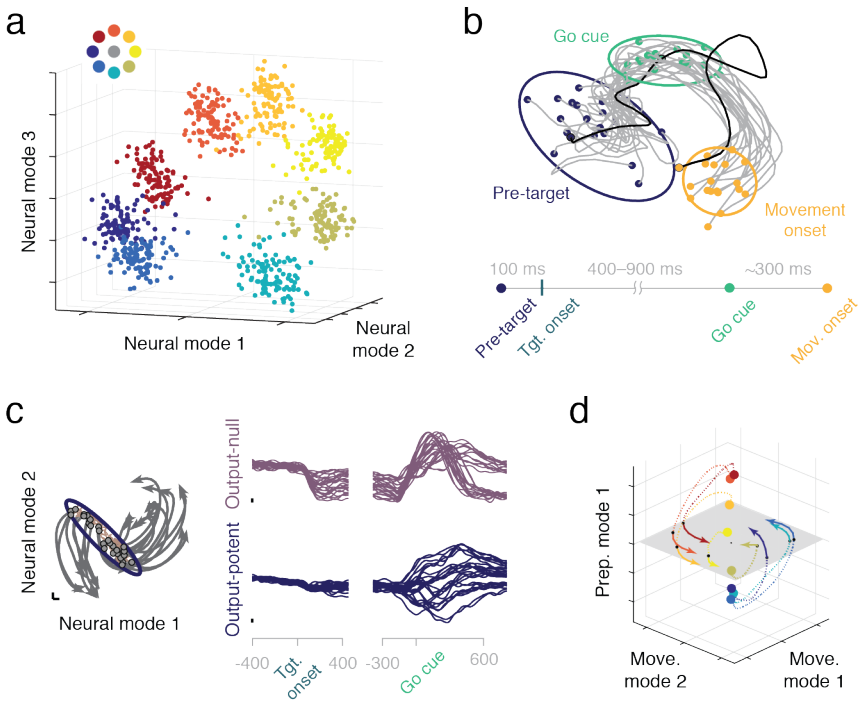

Figure 3.3 | Neural modes in motor cortices. a) Preparatory activity in PMd for an eight-target reach task. A neural manifold spanned by three neural modes reveals target-specific clusters. The targets are color coded (inset). Each point represents activity measured during the delay period for one trial. Adapted from (Santhanam et al. 2009). b) Population variability in PMd is reduced by stimulus presentation. For a delayed reach task (timeline at bottom of panel), population activity trajectories for individual trials are shown in gray in a two-dimensional manifold. The "Pre-target" confidence ellipsoid shrinks during "Go cue", as trajectories get closer together. Adapted from (Churchland et al. 2010b). c) The null space allows for movement preparation without execution. Population activity trajectories for each reach condition are shown in gray in a two-dimensional manifold. Preparatory activity sets the corresponding initial conditions (gray circles within purple ellipsoid). Null (purple, top) and potent (dark blue, bottom) latent variables, defined with respect to EMG activity, are shown on the right. Adapted from (Kaufman et al. 2014). d) Neural modes associated with movement preparation and execution span different manifolds. Two movement modes span a plane (in gray), while the orthogonal preparation mode spans a line that contains the initial conditions (colored circles). Projections of the full trajectories onto the movement manifold resemble the traces in panel c. Trajectories are color coded for each target, as per panel a. Adapted from (Elsayed et al. 2016).

manifolds simplify movement generation by providing a small number of signals that are independently controlled to achieve a desired behavior (Thoroughman and Shadmehr 2000; Mussa-Ivaldi and Solla 2004; Shenoy et al. 2013; Overduin et al. 2015; Song et al. 2015).

EMERGENCE OF NEURAL MANIFOLDS THROUGH LEARNING

Neural manifolds embody patterns of correlated activity; we hypothesize that these correlations reflect the connectivity of the underlying network of neurons. Since long-term learning can alter cortical connectivity (Rioult-Pedotti et al. 1998; Fu et al. 2012), we address potential connections between learning and the emergence of neural manifolds.

A revealing connection between neural manifolds and learning was discovered by Batista, Yu, and colleagues (Sadtler et al. 2014), who used a BMI paradigm in monkeys to address the question of why some motor skills are easier to learn than others. They used FA to identify a ten-dimensional neural manifold, and built a linear decoder from these latent variables into a two-dimensional center-out task. Once the monkeys had learned the task, the decoder was modified in one of two ways; see Figure 3.4a. When the modification required the use of existing neural modes (a within-manifold perturbation), the monkeys easily adapted in a single session (Figure 3.4b). In contrast, when the modification required the acquisition of new neural modes (an outside-manifold perturbation), the task proved significantly harder to learn (Figure 3.4c). The same group subsequently showed that monkeys could learn new neural modes (Oby et al. 2015) if they were guided to generate these new patterns progressively, over many days. These results suggest that short-term adaptation may be based on the generation of new combinations of preexisting neural modes, while long-term learning may require generating new neural modes. If neural modes arise from a neural circuitry that constrains spatiotemporal patterns of activity, it is

not surprising that modifying a neural manifold is more difficult than simply learning to use an existing neural manifold in novel ways.

A connection between the emergence of neural manifolds and learning also arises from a simulation (Sussillo et al. 2015) in which a Recurrent Neural Network (RNN) was trained to output the correct activity pattern for seven muscles recorded during a reach task (Churchland et al. 2012) (Figure 3.4d). In the RNN model, training refers to learning the connections between neurons to obtain a network able to perform the desired function, in analogy with the modification of synaptic connectivity in biological learning. Latent variables were then identified for the data generated by the trained RNN, and compared to those for the experimental data. The experimental data were best reproduced by the simplest RNN connectivity that allowed the network to output the measured muscle activity. A similar result was subsequently achieved with an RNN trained to output the x and y components of hand velocity for a delayed center-out reach task (Michaels et al. 2016).

In a complementary simulation study (Hennequin et al. 2014), Gerstner and colleagues proposed a neural network model with random excitatory recurrent connections and inhibitory connections that were learned to achieve stable population dynamics. This balanced network generated population activity patterns whose associated latent variables also exhibited the damped oscillations experimentally observed in the latent variables during reaching (Churchland et al. 2012). It is quite interesting that a network model not trained to produce a specific output – kinematics or muscle activity – but to stabilize its internal dynamics also exhibits oscillatory latent variables.

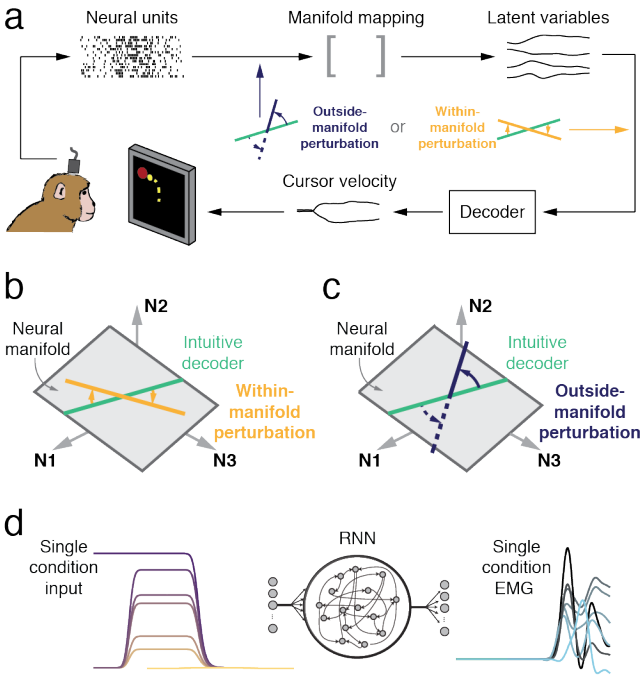

Figure 3.4 | Neural manifolds and learning. a) Existing neural modes facilitate the adaptation to variations of a learned task. In a BMI paradigm, ten latent variables representing the recorded activity of M1 neurons were mapped onto x and y components of cursor velocity. Once monkeys were proficient with the BMI, a b) within-manifold or c) outside-manifold perturbation was imposed on the decoder inputs. Adaptation to b) required a change in the relative activation of existing neural modes, while c) required the acquisition of new neural modes through changes in neural comodulation patterns. Adapted from (Sadtler et al. 2014). d) A recurrent neural network was trained to generate condition-specific EMG patterns selected through inputs representing preparatory activity. The latent variables that described the population activity of the trained recurrent network closely resembled those associated with the experimental data. Adapted from (Sussillo et al. 2015).

NEURAL MANIFOLDS FOR REACHING THROUGHOUT THE MOTOR CORTEX

The sensorimotor cortex is divided into a large number of functionally and cytoarchitecturally distinct regions (Brodmann 1909; Penfield and Boldrey 1937). An intriguing question is how the structure of the manifold evolves between these areas of the cortex. The manifold reflects the connectivity of local populations (Okun et al. 2015), such as those recorded by multielectrode arrays, but cells also send projections to other brain regions. Thus, there may be neural modes that are shared across the motor cortices. However, since each area performs distinct functions in coordinating movement, there must be area-specific modes as well. The observed dimensionality is a function of the complexity of a behavior (Gao and Ganguli 2015), yet it may also depend on the complexity of the activity in an area perhaps reflecting local computation. For the same behavioral output, we might expect high-order areas such as the pre-frontal cortex to require more neural modes than low-level output areas such as M1.

Here, we compared neural manifolds within M1 and PMd, two functionally distinct motor cortical areas. PMd appears to serve a more complex role in generating behavior than does M1. PMd has prominent activity during both motor planning related to sensory inputs and motor execution (Mushiake et al. 1991; Shen and Alexander 1997; Batista et al. 2007), while M1 is intimately involved in the execution of movement (Evarts 1968; Georgopoulos et al. 1982; Sergio et al. 2005). If neural modes are the fundamental building blocks of behavior, we predict that more modes will be needed to accomplish the diverse functions of PMd than for the motor execution function of M1. To test this hypothesis, we recorded simultaneously from populations of neurons in M1 and PMd using multielectrode arrays as monkeys performed a planar center-out reaching task. We adapted a method by Machens and colleagues to estimate the dimensionality, or the number of neural modes, of a recorded neural population (Machens et al.

2010). In brief, the method uses the trial-to-trial variability in each neuron's firing to estimate the amount of variance that can be explained simply by noise. This value is used to place a threshold on the PCA variance explained, thereby determining the number of non-noise dimensions.

As predicted, we consistently identified a larger dimensionality in PMd compared to M1 (Figure 3.5a). Since each neural mode captures population-wide activity, we expect that dimensionality within a local population should not depend on the particular neurons, or even the number of neurons that we recorded. We tested this by estimating the dimensionality of repeated random sub-samples of the recorded neurons. We found similar dimensionality even when removing 50% of the neurons, with PMd consistently higher-dimensional than M1 (Figure 3.5b).

These results illustrate that the more complex and higher-order functions performed by PMd involve a larger number of neural modes than M1. Given the strong connectivity between M1 and PMd (Dum and Strick 2002), some neural modes are likely to be shared. The shared modes may represent a direct means of information transfer between the populations, while the area-specific modes support the functions performed by that area (Sussillo et al. 2015). An interesting avenue for future research is to analyze the structure of the manifolds from simultaneously-recorded populations to understand whether these across-area modes exist, as well as their purpose in generating behavior.

NEURAL MANIFOLDS IN NON-MOTOR BRAIN CORTICES

Although we have focused on neural manifolds in motor cortices, it is important to remark that neural manifolds seem to be widely present across the brain. A number of studies have shown that the largely heterogeneous activity patterns of individual neurons in monkey

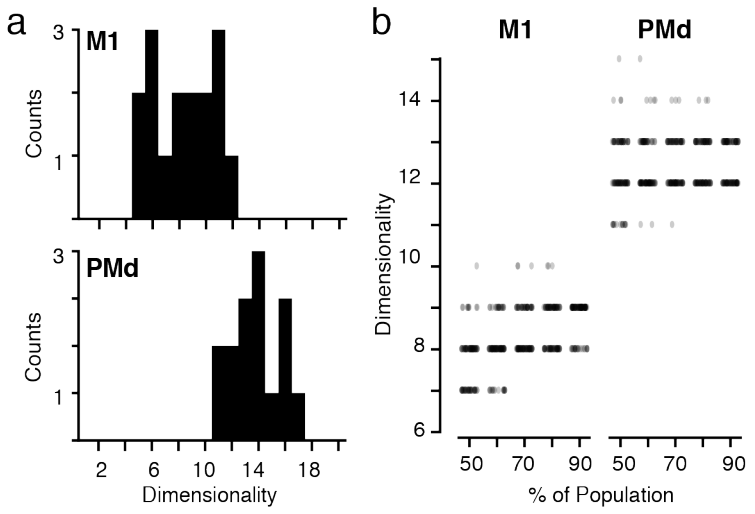

Figure 3.5 | PMd manifold has higher dimensionality than the M1 manifold. a) Summary histograms of dimensionality for M1 (top) and PMd (bottom). Each entry in the histogram represents a single session, taken from 22 sessions from two monkeys b) For one example session from Monkey C, the effect of population size on dimensionality. We randomly subsampled the neural populations 100 times at each percentage and repeated the dimensionality analysis. The result of each repetition is plotted as a single point with a random jitter on the horizontal axis to show the density. PMd (right) was consistently higher dimensional than M1 (left).

(Machens et al. 2010; Mante et al. 2013; Markowitz et al. 2015; Kobak et al. 2016) and rat (Durstewitz et al. 2010) prefrontal cortex, monkey (Churchland et al. 2010b) and rat (Forsberg et al. 2016) V1, rat olfactory cortex (Kobak et al. 2016), rat thalamus (Chapin and Nicolelis 1999), rat parietal cortex (Raposo et al. 2014), locust olfactory system (Stopfer et al. 2003), *aplysia* pedal ganglion (Bruno et al. 2015), and perhaps the entire zebrafish brain (Ahrens et al. 2012) can be explained as generated by a small set of latent variables associated with neural modes. In all these studies, neural modes and their time-varying activation helped describe previously unexplained mechanisms of neural function.

Studies of sensory cortices have provided clear evidence of network connectivity constraints on the activity of individual neurons. In the cat V1, the instantaneous activity of strongly tuned neurons is tightly linked to the population activity measured with optical imaging, both in response to stimulus presentation and during spontaneous activity (Tsodyks et al. 1999). Experiments in auditory and somatosensory cortices of awake and anesthetized rats further examined the relation between spontaneous and stimulus-evoked activity (Luczak et al. 2009) and found a surprising degree of conservation across these distinct regimes. The authors identified neural modes for spontaneous population activity, and found that neural modes in response to stimuli lay within the same neural manifold. Interestingly, the evoked responses sampled a smaller portion of the manifold than the spontaneous activity, an organization of population activity also found in monkey V1 (Cowley et al. 2016).

Recent experiments in both mouse and monkey V1, and in rat auditory cortex demonstrated that the correlation between single neuron activity and population activity is the same during both spontaneous and stimulus-evoked activity (Okun et al. 2015). These authors demonstrated that correlations between individual neurons and the population significantly predict all pairwise

correlations among the neurons. The degree of correlation between individual neurons and the population displays three interesting properties: 1) it predicts the responses of individual neurons during optogenetic stimulation of the surrounding population, suggesting that the relationship between population activity and that of a single neuron is causal; 2) it predicts the estimated number of synapses that a neuron receives from its neighbors, relating activity correlations to the underlying network structure; 3) it correlates strongly with the participation weight of the first FA factor on the activity of each neuron. Together, these results provide convincing if not direct evidence that network connectivity underlies the interactions among neurons captured by dimensionality reduction methods and the resulting neural modes.

OPEN QUESTIONS

The prevalence of neural manifolds across brain areas highlights their importance as a tool to understand brain function and reinforces the idea that neural processing may be built upon latent variables rather than on the activity of single neurons. The concept of neural manifolds is not restricted to flat surfaces; the manifold might well be a nonlinear surface within the neural space (Figure 3.1d). However, for stereotypical laboratory tasks as those discussed here, linear methods such as PCA or FA capture a large amount of the neural covariance. If the task-specific neural dynamics explore only a limited region within the manifold, a local linear approximation to a nonlinear neural manifold would work quite well (see the gray and purple planes in Figure 3.1e). For complex behaviors whose dynamics explore a larger region of neural space, linear methods may provide poor estimates of the neural manifold. Nonlinear methods for dimensionality reduction such as Locally Linear Embedding (LLE) (Roweis and Saul 2000),

Isomap (Tenenbaum et al. 2000), or Autoencoder Neural Networks (Hinton and Salakhutdinov 2006) might be needed.

These nonlinear methods have already been used in non-motor brain areas. Analysis of population data from the locust antennal lobe during an odor identification task using LLE revealed odor-specific neural manifolds that contain trajectories corresponding to different concentrations of the same odor (Stopfer et al. 2003). The analysis of the population activity of retinal ganglion cells using Isomap (Ganmor et al. 2015) identified activity clusters corresponding to similar visual stimuli. It remains an open question whether nonlinear methods might reveal mechanisms for the neural control of movement that linear methods have not revealed.

The studies discussed here focus on neural manifolds associated with specific tasks. Since organisms are able to execute a rich repertoire of motor tasks, how are the corresponding neural manifolds organized with respect to each other within the neural space? The neural space may contain distinct neural manifolds, each associated with a specific task. What is then the relationship among these distinct manifolds? Is each manifold spanned by its own unique neural modes, or do tasks in a class share some neural modes that represent common features? These are virtually unexplored questions; note however a recent suggestion that preparation and execution of a movement correspond to orthogonal but related manifolds (Elsayed et al. 2016). Whether a similar finding applies to manifolds corresponding to the execution of different tasks is yet unknown.

An alternative is that all motor behaviors might lie within a single, possibly nonlinear, universal neural manifold, with each task sampling a different region (Figure 3.1e). Because the tasks commonly studied are simple and stereotypical, the manifolds extracted from the recorded

neural data may capture only a small region of the universal manifold, a region spanned by task-specific neural modes. This picture agrees with theoretical arguments that limiting the complexity of behavior also constrains the dimensionality of the neural manifolds (Gao and Ganguli 2015). In this view, recording larger populations of neurons may not yield new insights; only by recording population activity during naturalistic, complex behaviors will we uncover the true underlying neural manifold.

We have argued that motor behaviors are generated by the time-varying activation of a small set of neural modes, population-wide activity patterns arising from network connectivity that define a low-dimensional manifold in neural space. We argue that the transition from a neuron-centric to a manifold-centric view of neural activity fosters a important advance in our understanding of brain function. Experiments involving longer and more complex motor tasks will require the development of increasingly sophisticated techniques for recording and data analysis; all these will be critical to advance our understanding of the relation between network connectivity, the resulting neural manifolds, and motor behaviors.

METHODS

Two monkeys (male, *mucaca mulatta*; Monkey C: 11.7 kg, Monkey M: 10.5 kg) were seated in a primate chair and made reaching movements with a custom 2-D planar manipulandum to control a cursor displayed on a computer screen. The monkeys performed a standard center-out reaching task with eight outer targets evenly distributed around a circle at a radius of 8cm. All targets were 2cm squares. Each trial began when the monkey moved to a center target. After a variable hold period (0.5 – 1.5 s), one of the eight outer targets appeared. The monkey had a variable instructed delay period (0.5 – 1.5 s). The monkeys then received an

auditory go cue, and the center target disappeared. The monkeys had one second to reach the target, where they had to hold for 0.5 s.

After extensive training, we surgically implanted chronic multi-electrode arrays (Blackrock Microsystems, Salt Lake City, UT) in M1 and PMd. From each array, we recorded 96 channels of neural activity using a Blackrock Cerebus system (Blackrock Microsystems, Salt Lake City, UT). The snippet data was manually processed offline using spike sorting software to identify single neurons (Offline Sorter v3, Plexon, Inc, Dallas, TX). Across 22 behavioral sessions, we isolated between 137 – 256 PMd and 55 – 93 M1 neurons for Monkey C, and 66 – 121 PMd and 26 – 51 M1 neurons for Monkey M. We excluded cells with a trial-averaged firing rates of less than 1 Hz.

We counted spikes in 10 ms bins and square root transformed the raw counts to stabilize the variance (Cunningham and Yu 2014). We then convolved the spike train of each neuron for each trial with a Gaussian kernel of width 100 ms to compute a smooth firing rate. We used Principal Component Analysis (PCA) to reduce the smoothed firing rates of the neurons in each session to a small number of components.

To estimate the dimensionality of a population, we adapted a method developed by Machens et al to estimate the dimensionality of our recorded populations (Machens et al. 2010). In brief, PCA provides an orthogonal basis set with the same dimensionality as the neural input. However, the variance captured by many of the higher dimensions (with the smallest eigenvalues) is typically quite small. We estimated the noise in the neural activity patterns using the trial-to-trial variation in the activity of each neuron. We sampled a random pair of trials for each reach direction and subtracted the activity of each neuron. This gave an estimate of the variance of each neuron across two different reaches to each target. We then ran PCA on the

neural “noise” space provided by this difference for all targets. We repeated this 1000 times, giving a distribution of eigenvalues for each of these noise dimensions. We used the 99% limit of these distributions to estimate the amount of noise variance explained for each dimension. This allowed us to put a ceiling on the amount of variance that could be explained by noise. The dimensionality was thus defined by the number of dimensions needed to explain 95% of the remaining variance.

CHAPTER 4

A NEURAL POPULATION MECHANISM FOR RAPID LEARNING

Matthew G. Perich, Juan A. Gallego, Lee E. Miller

FOREWORD

This chapter is adapted from a submitted manuscript (Perich et al. 2017). Although the content is for the most part the same as the submitted manuscript, I have rearranged the figures and expanded the text to elaborate on several important points. Using the manifold framework laid out in the previous chapter, I developed an analytical approach to study the population-level interactions between M1 and PMd during learning. This allowed me to test the hypothesis laid out at the end of Chapter 2 that PMd recruits M1 during CF learning. I show evidence that supports this hypothesis: there are specific changes in PMd activity patterns during CF learning. These changes were not observed during adaptation to a VR, suggesting a fundamental difference in how the brain learns to compensate for these two perturbations.

ABSTRACT

The brain's ability to adapt behavior is crucial to survival. Long-term learning of dexterous motor skills likely requires plastic changes in cortex (Nudo et al. 1996), but we can also learn even from errors in single movements (Thoroughman and Shadmehr 2000). Such trial-to-trial adaptation likely requires a faster mechanism. Here, we show how the brain can adapt behavior by redeploying existing population activity patterns without altering the functional structure of the cortex. We recorded from both primary motor cortex (M1) and dorsal premotor (PMd) cortex in macaque monkeys during motor learning. We trained computational models to predict single neuron spiking based on the activity of the surrounding neural population in order to study the functional relationships between neurons *within* the two areas (Truccolo et al. 2010). Intriguingly, the functional structure within each area was preserved throughout learning, suggesting that the underlying neural circuitry remained unaltered (Ahissar et al. 1992; Gerhard et al. 2013). To study the interaction *between* the areas, we separated the PMd activity (Cunningham and Yu 2014) into two sets of components: "potent" components that captured activity that mapped onto M1, and "null" components that captured activity patterns with effects only within PMd (Kaufman et al. 2014). As was true within each area, the activity of the potent components consistently predicted M1 spiking throughout its learning-related changes. In stark contrast, the mapping from the null components gradually changed with learning. These results show that, at a population level, PMd develops new motor plans (reflected in the null components) that are transmitted to M1 through an unchanged functional mapping (captured by the potent components) between the two areas. Use of the PMd-to-M1 null space as a neural scratch pad for the gradual development of new motor plans is a powerful mechanism that may explain a variety of rapid learning processes throughout the brain.

INTRODUCTION

In order to make skilled movements, sensory input is combined with internal state variables and transformed into a plan executed by the motor cortices (Kalaska et al. 1997). This transformation may be mediated in part by an "internal inverse model" that maps a desired motor action to the required low-level muscle commands. The dorsal premotor (PMd) and primary motor (M1) cortices, together with the cerebellum, are prime candidate locations for such an inverse model. PMd is involved in movement planning (Cisek and Kalaska 2005), with diverse inputs and strong connectivity with M1 (Dum and Strick 2002), while M1 is the main cortical output to the spinal cord (Rathelot and Strick 2009). Correction of movement errors, such as those caused by external perturbations (Shadmehr and Mussa-Ivaldi 1994), is thought to lead to alteration of the inverse model, leading to progressively more accurate movements even on a trial-by-trial basis (Thoroughman and Shadmehr 2000). The rapid speed of these changes seems incompatible with structural changes in synaptic connectivity (Bailey and Chen 1983, 1988). Furthermore, monkeys using brain-machine interfaces have great difficulty learning mappings between brain activity and cursor movement that require the normal pattern of covariation among recorded cortical neurons to be altered (Sadtler et al. 2014). Given that neural covariance patterns seem to be determined by synaptic connectivity (Okun et al. 2015), this result further suggests that changes in cortical connectivity may not be the primary mechanism for short-term learning. At the same time, the progressive change in performance over tens of minutes and the performance savings between sessions seem incompatible with a mechanism like the network reverberation that may underlie short-term working memory (Major and Tank 2004).

To reconcile these apparently contradictory observations, we recorded simultaneously from electrode arrays implanted in both M1 and PMd (Figure 4.1a) as monkeys learned to make

accurate reaching movements that were perturbed by a curl field (CF), a velocity-dependent force applied to the hand (Shadmehr and Mussa-Ivaldi 1994), or a visuomotor rotation (VR). The CF altered the dynamics such that straight reaches to each target required the monkeys to learn new muscle activation patterns (Thoroughman and Shadmehr 2000; Cherian et al. 2013), while the VR rotated the visual feedback that the monkeys received. We investigated whether changes in the functional relationships of the M1 and PMd populations could explain the adapted behavior.

RESULTS

Behavioral adaptation and single neuron activity

Two rhesus macaque monkeys to performed the standard center-out reaching task (Figure 4.1a). Each session began with reaches in a null field before the monkeys began to adapt to the CF (Shadmehr and Mussa-Ivaldi 1994; Li et al. 2001; Cherian et al. 2013), and progressively straighten their reaches (Figure 4.1b,c). Evidence of their learning was revealed by the occurrence of after-effects upon eventual return to the null field. As has been previously observed, neural activity during adaptation was strikingly heterogeneous across neurons (Churchland and Shenoy 2007), and nearly all cells showed large changes in firing rate (Figure 4.1d). Intriguingly however, the population correlation structure was surprisingly similar before and after learning, despite the large changes in neural firing in both areas (Figure 4.1e,f).

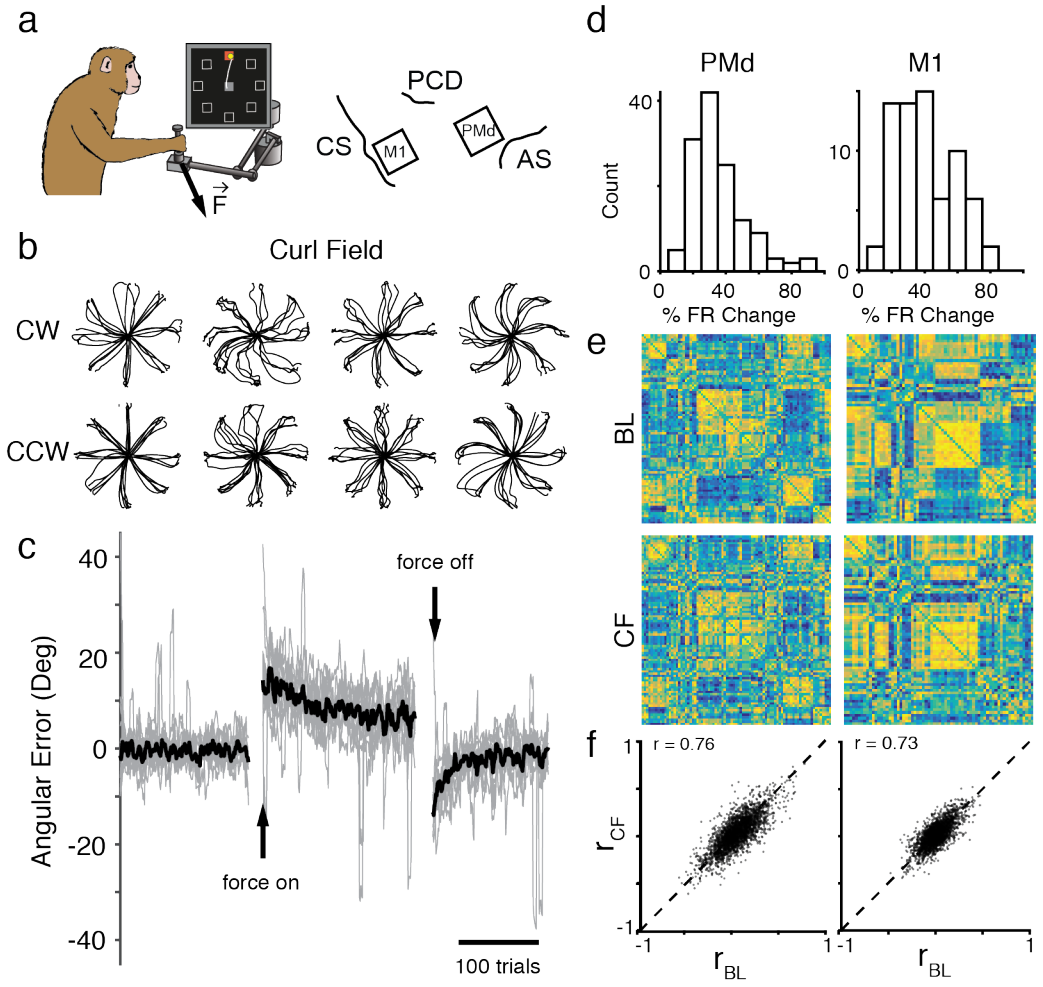

Figure 4.1 | Curl field task. a) Monkeys performed a standard center-out task with a variable instructed delay period following cue presentation (top). We recorded from single neurons in M1 and PMd (bottom; CS: central sulcus, PCD: pre-central dimple, AS: arcuate sulcus). b) Example position traces for the first reaches to each target from four sessions with a clockwise CF (top row) and five sessions with a counter-clockwise CF (bottom row). Sessions from both monkeys are included. Data from the three sessions with Monkey C with shorter reaches were plotted on a different scale to provide uniform length for visualization purposes. Curvature increased when the CF was imposed (top right), but straightened during learning (bottom left). We observed oppositely directed after-effects in Washout (bottom right). c) Angular error for all CF sessions across trials. Thin gray lines represent error on individual sessions, while thick black line is median across sessions. Error increased with the CF ("force on"), and after-effects can be seen upon removal ('force off"). d) Summary of percent of firing rate change for all cells recorded on a single session. e) Normalized pairwise correlations between all cells recorded on the same session as Panel c. Clustering was performed in Baseline (BL, left) as a means to visualize the correlation structure, and the same ordering was kept for late CF (right). f) Summary of pairwise correlations between BL and late CF for all combinations of neurons recorded in each of the nine CF sessions. A random subsample of 5,000 pairs is plotted. The value of 'r' for each plot indicates the Pearson's correlation coefficient to assess similarity between the null-field and CF conditions for all pairwise-correlations from all sessions.

Assessing population relationships during learning

We used Poisson Generalized Linear Models (GLMs) to predict the spiking of individual neurons based on the activity of the remaining neurons (see Methods) (Pillow et al. 2008; Truccolo et al. 2010). Using data from late in learning when behavior had stabilized, we trained three models (Figure 4.2a,b): one predicted M1 neurons from the M1 population activity (M1-M1), another predicted PMd neurons from the PMd population (PMd-PMd), and the third predicted M1 neurons from the PMd population (PMd-M1). We assessed model performance using a relative pseudo-R^2 (rpR^2) metric, which quantified the improvement in model performance due to the neural inputs above that of reach kinematics alone (see Methods). This removed the effect of shared variability due to behavior-related common inputs, and left the unique contributions of individual neurons. We tested generalization of each model from the late CF training data to the early CF trials. Good generalization would indicate that the relationships between neurons were unchanged during learning.

All three models performed similarly well when test on cross-validated training data (Figure 4.2c). We then tested whether the models could generalize to the Early CF trials. The within-area models (M1-M1, PMd-PMd) predicted the complex spiking changes remarkably well throughout learning (Figure 4.2d,e). However, early in learning, M1 spiking was poorly predicted from PMd (Figure 4.3a). Furthermore, the model's performance changed with a time course very much like that of behavioral adaptation (Figure 4.3b), suggesting that learning resulted from a change in the functional relationships between neurons in PMd and M1, even though the functional interactions within the two networks remained unchanged (Ahissar et al. 1992; Gerhard et al. 2013; Okun et al. 2015).

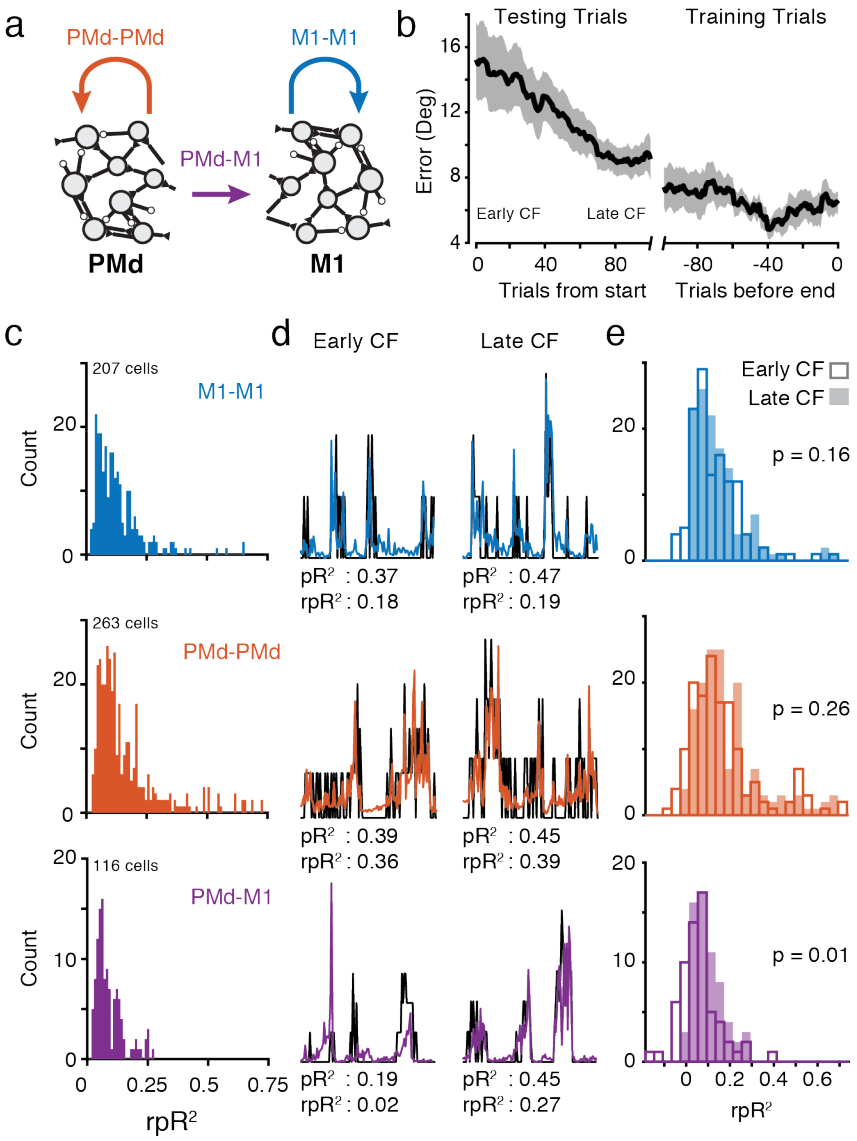

Figure 4.2 | GLM model performance during CF learning. a) We trained three models: two within an area (M1-M1, and PMd-PMd) and one between the areas (PMd-M1). b) Schematic of angular error (mean ± st.e. across sessions) during CF learning. We trained GLMs using data recorded late in learning, after behavior had stabilized (right) and tested them for generalization throughout the initial phase of learning, beginning at the first CF trial (left). We compared Early CF (highest error) and Late CF (lowest error) trials within the testing block. c) Cross-validated rpR^2 for all cells with significant fits for the three GLM models. Cells were pooled across two monkeys and nine total sessions. d) Spiking of three representative neurons (black) and model predictions (colors) during three early and three late learning trials. e) Summary histograms of rpR^2 values for predictions of a block of 5 trials in early CF (hollow) and late CF (solid). M1-M1 and PMd-PMd had similar distributions during early and late CF, but Early CF predictions by the PMd-M1 models were significantly lower than Late CF ($p = 0.01$, two-sample t-test).

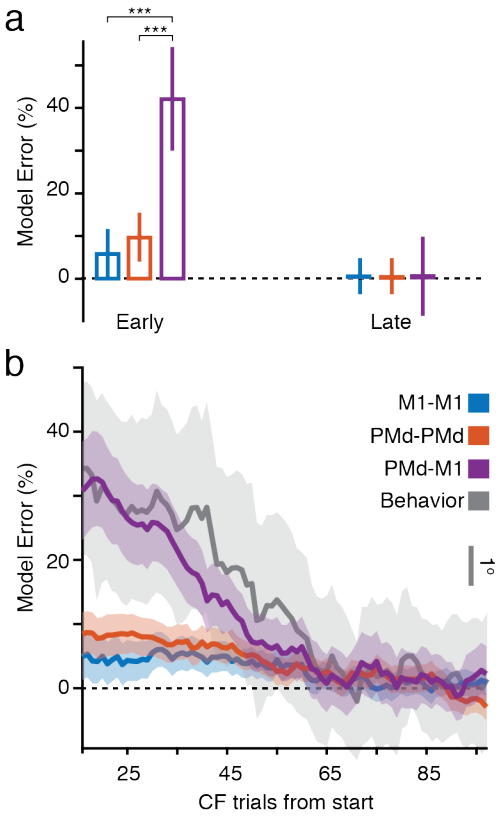

Figure 4.3 | Time course of GLM model performance changes. a) Percent error in model performance during early and late learning. Significant differences were observed between PMd-M1 during Early and Late, and PMd-M1 compared with M1-M1 and PMd-PMd during late, with a significance level defined at $p < 0.01$ (two-sample t-test). b) Time course of model performance changes. Predictions were made for individual trials, and then smoothed with a 30 trial moving average (see Methods). Plotted data are mean and standard error across predicted neurons for each model. Behavioral error processed with the same methods is overlaid in gray. The decrease in PMd-M1 model error followed the time course of behavior.

Separating PMd into potent and null components

We sought to explain how the relationships between PMd and M1 neurons could change while those within each area remained consistent. Intuitively, PMd population activity at once reflects its inputs, their subsequent processing, and the eventual outputs to M1 (Sussillo et al. 2015) (Figure 4.4a). We sought to separate these components of population activity by projecting PMd activity patterns onto output-null and output-potent spaces (Figure 4.4b, see Methods) (Kaufman et al. 2014). We used Principal Component Analysis (PCA) to represent the activity of the M1 and PMd populations as a small number of components that captured mutual covariance patterns across neurons (Cunningham and Yu 2014). PMd consistently contained more components than M1 (see Chapter 3), indicating the existence of a null space containing PMd activity that had no net effect on the M1 components. We hypothesized that these extra components and the resulting null space arise from planning-related computations performed within PMd that did not directly activate M1. Such null-space planning activity could account for the altered overall relationships between PMd and M1, while at the same time, allowing the potent space to maintain a stable mapping from PMd to M1.

We repeated the GLM analysis to predict the spiking of individual M1 neurons using either the PMd potent (Pot-M1) or null (Null-M1) components as inputs (Figure 4.4c). If the null components capture motor planning within PMd that changes with learning, the accuracy of the Null-M1 model should change with behavioral performance, much like the overall PMd-M1 model. However, if the updated motor plans are ultimately sent in a consistent manner to M1, Pot-M1 should remain unchanged. We compared these models against a third GLM which used all of the PMd components found by PCA. For both monkeys, Pot-M1 predicted M1 spiking consistently, while Null-M1 and PCA-M1 predictions changed (Figure 4.4d) with a time course

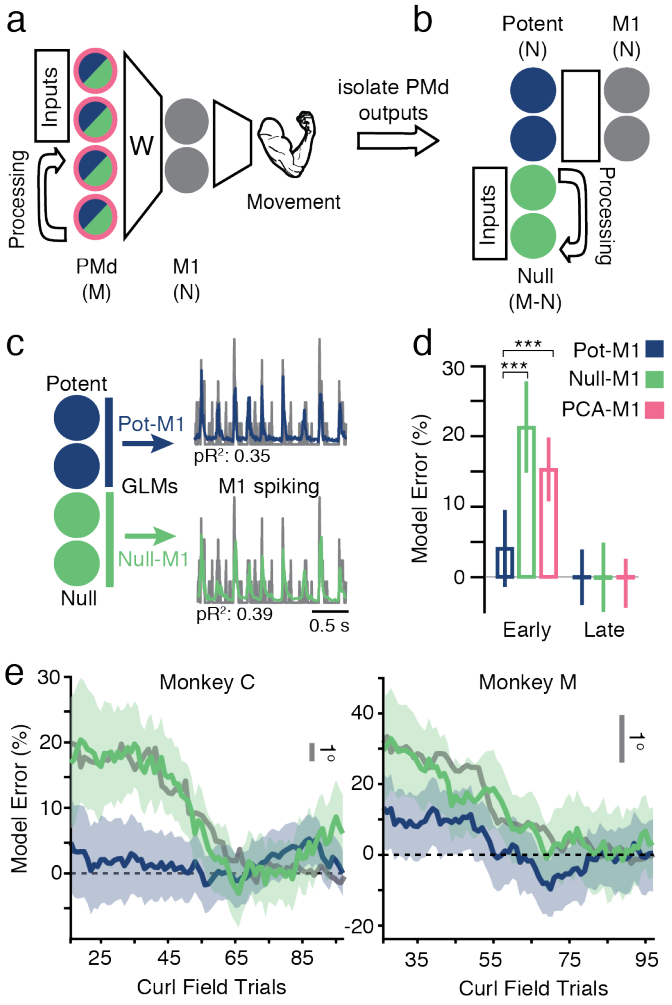

Figure 4.4 | Predictions from potent and null components. a) Hierarchical schematic of motor planning in PMd and M1. b) We devised an analysis to demix the PMd outputs to M1 from the other functions of the population. The former comprises the potent space, while the latter resides in the null space. c) The time-varying projections of PMd activity onto these potent and null axes were used as the inputs to a GLM model to predict M1 spiking. d) Bar plot comparing model error performance during early and late trials with potent (Pot-M1), null (Null-M1), and all PMd PCs (PCA-M1). Pot-M1 performed significantly better than the other models ($p < 0.01$, two-sample t-test). e) The time course of model performance for Pot-M1 and Null-M1 for all sessions with the two subjects. Gray line is the mean behavioral error corresponding to those trials.

like that of behavior (Figure 4.4e). For this analysis, we had to assume a dimensionality for both M1 and PMd. We verified that our results were not dependent on the precise values selected by repeating the GLM analysis for different dimensionalities, always assuming the dimensionality of PMd to be twice that of M1. We found that Null-M1 consistently performed worse during Early CF than Late CF for all dimensionalities tested (Figure 4.5). Pot-M1 performed significantly better than Null-M1, though for higher dimensionalities (where the potent space is overestimated), the GLMs began to perform worse. This is expected, since the components found when overestimating the true potent space can include noise, or even activity that would otherwise participate in the output-null dimensions.

The stability of Pot-M1 shows that, at a population level, there exists a direct mapping between PMd and M1 that persists throughout short-term motor adaptation. The stability of the potent mapping (along with the M1-M1 and PMd-PMd models) supports the conclusion that there were no structural changes in these cortical areas, since the potent and null spaces are defined simply by different weightings of the same neurons with the same connectivity (see Figure 4.8d). Learning, then, results from new activity patterns within PMd, which may be necessary to set a new preparatory state for M1 (Churchland et al. 2012; Kaufman et al. 2014).

GLM predictions during VR adaptation

We next asked if the observed changes within PMd are a necessary consequence of adapted behavior, or if they are indicative of a more specialized role for PMd in the CF task. On a separate set of sessions, the monkeys learned to reach in the presence of a static rotation of the reach-related visual feedback (visuomotor rotation; VR). Considerable evidence from behavioral studies in humans suggests that the brain areas involved in learning a static visual mapping differ

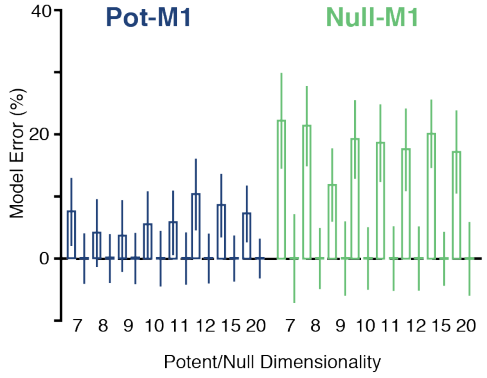

Figure 4.5 | Potent and null differences do not depend on selected dimensionality. Comparison of GLM performance error between early CF (left bars for each dimension) and late CF trials (right bars) for Pot-M1 and Null-M1 as a function of the selected dimensionality. Values for eight dimensions plotted here are those included in Figure 4.4d. Our primary effect that Pot-M1 generalizes to early CF trials better than Null-M1 was consistent for a range of dimensionalities.

from those required to learn novel effector dynamics (e.g., the CF) (Krakauer et al. 1999; Diedrichsen et al. 2005). Since VR learning appears to rely heavily on parietal cortex (Diedrichsen et al. 2005; Tanaka et al. 2009), hierarchically upstream yet of PMd, we hypothesized that VR adaptation would not result in a change in the functional relationship between M1 and PMd.

We repeated the above analyses using sessions where the monkeys adapted to a VR of 30 degrees (Figure 4.6a). There were a number of similarities with the CF sessions: behavioral errors were similar in magnitude and time course (Figure 4.6b,c), and there were highly varied, complex changes in neural activity patterns with a preserved correlation structure (Figure 4.6d-e). However, when we fit GLM models to predict single neuron spiking, all models, including PMd-M1, accurately generalized early and late in learning (Figure 4.7a), despite the clear behavioral adaptation (Figure 4.7b). Thus, there were no changes in the functional relationships between the PMd and M1 populations, despite diverse changes in single-neuron activity. This result highlights a fundamental difference in the neural adaptation to these two perturbations, and supports the view that VR adaptation occurs upstream of PMd, likely involving parietal cortex (Diedrichsen et al. 2005; Tanaka et al. 2009). It also strengthens our conclusions about the CF task: the poor generalization of the PMd-M1 GLM model is not a necessary consequence of changing behavior, but rather captures a previously undescribed mechanism by which the motor cortices drive sensorimotor adaptation through population-level activity patterns.

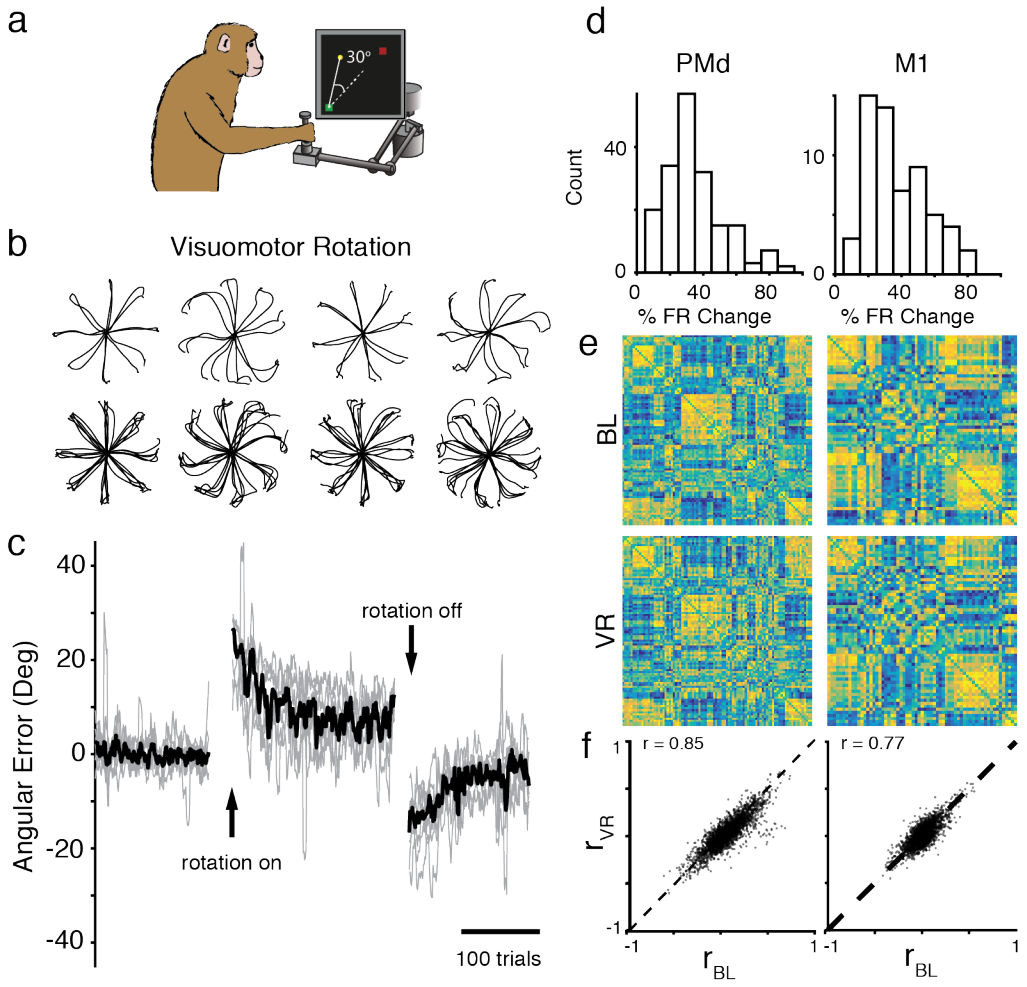

Figure 4.6 | Visuomotor rotation task. a) The monkeys also adapted to a visuomotor rotation (VR) using the same center-out task as Figure 4.1. The VR rotated the visual cursor feedback on the screen by 30 degrees. b) Position traces for the first (or last) reach to each target for four representative sessions with the VR, shown as in Figure 4.1b. c) Angular error for the VR sessions, as plotted in Figure 4.1c. The monkeys exhibited behavioral errors that were similar to those of the CF condition. d-f) Neural activity changes and pairwise correlations, as plotted in Figure 4.1.

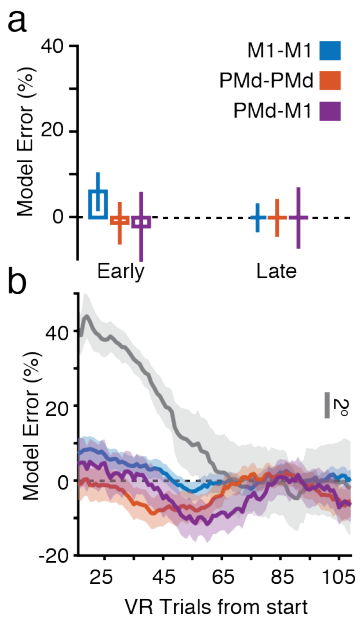

Figure 4.7 | GLM performance during VR learning. a) Model prediction error during early and late VR trials. There were no significant changes throughout learning. b) Same format as Figure 4.3b. All GLM models, including PMd-M1 generalized well from late VR to early VR trials, despite clear behavioral adaptation (gray).

Discussion

Long-term learning is known to alter connectivity in the motor cortex, resulting in increased horizontal connections (Rioult-Pedotti et al. 1998) and synaptogenesis(Kleim et al. 2002). Many have proposed that the brain uses similar plastic mechanisms to adapt behavior on shorter timescales (Classen et al. 1998; Li et al. 2001). However, structural changes would have impaired predictions of the GLM models (Ahissar et al. 1992; Gerhard et al. 2013). Hence our results suggest that, at least on the time scale of a single experimental session, there were no structural changes within PMd or M1. Instead, we show that we learn by exploring new activity patterns within the existing network structure. These new patterns offer a possible explanation for the changes in movement representations of single neurons previously reported during CF learning (Gandolfo et al. 2000; Li et al. 2001; Richardson et al. 2012). It is important to note that our GLM models and null/potent analysis finds functional relationships between the neural populations; it does not necessarily represent direct anatomical connectivity. However, we expect that activity reflecting any direct or indirect synaptic connectivity would be contained within these models.

Our lab has previously found that the relationship between M1 activity and the dynamics of the motor output remains unchanged during CF adaptation (Cherian et al. 2013), with no evidence for adaptive changes in either spatial tuning or firing rates that have a time course like that of learning. Therefore, we hypothesized that CF learning must be mediated by changes in recruitment of M1 by upstream areas, including PMd. Our current results directly support this interpretation: PMd exploits the null space to formulate new motor plans reflecting the modified task demands of the CF, which are then used to recruit M1 without changing the connectivity

within either area, or within the potent space from PMd to M1. The lack of any such change in null-space processing during the VR task suggests that VR adaptation occurs upstream of PMd.

Through lesion, computational, and recording studies, the cerebellum has been implicated in a variety of supervised, error-driven motor-learning problems, including both the curl field and visual rotation paradigms explored in this study (Diedrichsen et al. 2005; Galea et al. 2011; Izawa et al. 2012; Herzfeld et al. 2014). It is also considered to be a site at which both forward and inverse internal models may be learned (Wolpert et al. 1998; Imamizu et al. 2000). Many forms of cellular plasticity are present in the cerebellum, occurring at multiple sites and over several time scales (Zheng and Raman 2010; Yang and Lisberger 2013). The cerebellum also supports rapid, short term memory storage through the bistable properties of Purkinje cells(Loewenstein et al. 2005). The most direct evidence for its role in motor adaptation comes through Purkinje cell recordings during the adaptation of arm (Gilbert and Thach 1977), eye (Yang and Lisberger 2013), and head (Brooks et al. 2015) movements. Given the extensive interconnections with PMd (Dum and Strick 2003) the new motor plans during CF learning may arise from interactions between PMd and an evolving inverse internal model in the cerebellum (Wolpert et al. 1998; Thoroughman and Shadmehr 2000; Diedrichsen et al. 2005).

Other evidence suggests that while these internal models may depend on the cerebellum for their modification, they may actually be located elsewhere (Shadmehr and Holcomb 1997). Over a longer time period (hours to days), the motor memory is consolidated, possibly through structural changes in the cerebral cortex (Nudo et al. 1996; Peters et al. 2014). We propose that such structural changes could emerge to support the long-term refinement and recall of skills (Bailey and Kandel 1993; Peters et al. 2014), while rapid behavioral adaptation is mediated by modified population-wide activity patterns within the existing constrained network structure of

the motor cortices. Similar activity patterns have been found in prefrontal cortex for decision-making (Mante et al. 2013), working memory (Machens et al. 2010), and rule-learning (Durstewitz et al. 2010), in the motor cortex for movement planning (Kaufman et al. 2014), and in the parietal cortex for navigation (Harvey et al. 2012). These widespread observations suggest that the novel coordination mechanism between neuronal populations described here could be exploited throughout the brain for the rapid, flexible adaptation of behavior.

METHODS

Behavioral task

Two monkeys (male, *mucaca mulatta*; Monkey C: 11.7 kg, Monkey M: 10.5 kg) were seated in a primate chair and made reaching movements with a custom 2-D planar manipulandum to control a cursor displayed on a computer screen. We recorded the position of the handle at a sampling frequency of 1kHz using encoders. The monkeys performed a standard center-out reaching task with eight outer targets evenly distributed around a circle at a radius of 8cm. All targets were 2cm squares. The first three sessions with Monkey C used a radius of 6 cm. However, we observed no qualitative different in the behavioral or neural results for the shorter reach distance, and all sessions were thus treated equally. Each trial began when the monkey moved to a center target. After a variable hold period (0.5 – 1.5 s), one of the eight outer targets appeared. The monkey had a variable instructed delay period (0.5 – 1.5 s) which allowed us to study neural activity during explicit movement planning and preparation, in addition to movement execution. The monkeys then received an auditory go cue, and the center target disappeared. The monkeys had one second to reach the target, where they had to hold for 0.5 s.

In the curl field (CF) task, two motors applied torques to the elbow and shoulder joints of the manipulandum in order to achieve the desired endpoint force. The magnitude and direction of the force depended on the velocity of hand movement according to Equation 1, where $\vec{F}$ is the endpoint force, $\dot{p}$ is the derivative of the hand position $\vec{p}$, θ_c is the angle of curl field application (85°), and k is a constant (0.15 N•s/cm):

$$\vec{F} = \begin{bmatrix} F_x \\ F_y \end{bmatrix} = k \begin{bmatrix} \cos\theta_c & -\sin\theta_c \\ \sin\theta_c & \cos\theta_c \end{bmatrix} \begin{bmatrix} \dot{p}_x \\ \dot{p}_y \end{bmatrix} \quad (1)$$

In the visuomotor rotation (VR) task, hand position p was rotated by θ_r (here, chosen to be 30°) to provide altered cursor feedback $\vec{C}$ on the screen. The rotation was position-dependent so that the cursor would return to the center target with the return reach:

$$\vec{C} = \begin{bmatrix} C_x \\ C_y \end{bmatrix} = \begin{bmatrix} \cos\theta_r & -\sin\theta_r \\ \sin\theta_r & \cos\theta_r \end{bmatrix} \begin{bmatrix} p_x \\ p_y \end{bmatrix} \quad (2)$$

Both the CF and VR perturbations were imposed continuously throughout the block of learning trials, including the return to center and outer target hold periods.

Each session was of variable length since we allowed the monkeys to reach as long as possible to ensure that behavior had sufficient time to stabilize, and allow for large testing and training sets for the GLM. For the CF sessions, the monkeys performed a block of unperturbed Baseline trials (range across sessions: 170 – 225 rewards) followed by an Adaptation block with the CF perturbations (201 – 337 rewards). The session concluded with a Washout block, where the perturbation was removed and the monkeys readapted to making normal reaches (153 – 404 rewards). The curl field was applied in both clockwise (CW) and counter-clockwise (CCW) directions, though we saw no qualitative difference between the sessions. Monkey C had three CW sessions and two CCW sessions, while Monkey M had four CCW sessions. For the VR

sessions, the monkeys performed 154 – 217 successful trials in Baseline, 219 – 316 during VR (either CW or CCW), and then 162 – 348 in Washout. Monkey C performed two CW VR sessions and two CCW sessions, while monkey M performed three CCW sessions. There is considerable evidence that learning can be consolidated, resulting in savings across sessions (Huang et al. 2011). In this study, we minimized the effect of savings to focus on single-session learning. The monkeys typically: 1) received different perturbations day-to-day, as we alternated between CF and VR sessions, 2) received opposing directions of the perturbation on subsequent days, and 3) had multiple days between successive perturbation exposures.

Behavioral adaptation analysis

For a quantitative summary of behavioral adaptation, we used the errors in the angle of the initial hand trajectory. We measured the angular deviation of the hand from the true target direction 150 ms after movement onset. To account for the natural biases of the monkeys, we found the difference on each trial from the average deviation for that target in Baseline trials. Sessions with the CW and CCW perturbations were similar except for the sign of the effects. Thus, for the behavioral data in Figures 4.1 and 4.6, we pooled all perturbation directions together and simply flipped the sign of the CW errors. Figures 4.1b and 4.6b show the position traces for example CF and VR sessions, respectively. Since the target size was 2cm, there could be some deviation in the starting and ending positions, and subsequently some deviation in the total length of the reaches. For visualization purposes only, we normalized the length of each reach to begin in the center of the workspace and have a total linear distance of 8cm between the starting and ending points.

Neural recordings

After extensive training in the unperturbed center-out reaching task, we surgically implanted chronic multi-electrode arrays (Blackrock Microsystems, Salt Lake City, UT) in M1 and PMd. From each array, we recorded 96 channels of neural activity using a Blackrock Cerebus system (Blackrock Microsystems, Salt Lake City, UT). The snippet data was manually processed offline using spike sorting software to identify single neurons (Offline Sorter v3, Plexon, Inc, Dallas, TX). We sorted data from all three task epochs (Baseline, CF/VR learning, and Washout) simultaneously to ensure we reliably identified the same neurons throughout the sessions. With such array recordings, there is a small possibility that duplicate neurons can appear on different channels as a result of electrode shunting, which would influence our GLM models by providing perfectly correlated inputs for these cells. While such duplicate channels are often easily identifiable during recording, we took two precautionary steps to ensure our data included only independent channels. First, we used the electrode crosstalk utility in the Blackrock Cerebus system to identify and disable any potential candidates with high crosstalk. Second, offline we computed the percent of coincident spikes between any two channels, and compared this percentage against an empirical probability distribution from all sessions of data. We excluded any cells whose coincidence was above a 95% probability threshold (in practice, this was approximately 15-20% coincidence, which excluded no more than one or two low-firing cells per session).

Across all sessions, we isolated between 137 – 256 PMd and 55 – 93 M1 neurons for Monkey C, and 66 – 121 PMd and 26 – 51 M1 neurons for Monkey M. For the pairwise correlation analysis, we excluded cells with a trial-averaged firing rates of less than 1 Hz. Our GLM models were by necessity poorly fit for neurons with low firing rates. Thus, for the GLM

analyses, we only considered neurons with a trial-averaged mean firing rate greater than 5 Hz. Pooled across all monkeys and CF and VR sessions, this gave a population of 918 M1 and 2221 PMd neurons. Given the chronic nature of these recordings, it is certain that some individual neurons appeared in multiple sessions. However, our analyses primarily focus on the population-level relationships which we found to be robust to changes in the exact cells recorded, so we do not expect our results to biased by partial resampling.

Dimensionality reduction

We counted spikes in 10 ms bins and square root transformed the raw counts to stabilize the variance (Cunningham and Yu 2014). We then convolved the spike train of each neuron for each trial with a Gaussian kernel of width 100 ms to compute a smooth firing rate. We used Principal Component Analysis (PCA) to reduce the smoothed firing rates of the neurons in each session to a small number of components (Cunningham and Yu 2014). PCA finds the dominant covariation patterns in the population and provides a set of orthogonal basis vectors that captures most of the population variance. Importantly, the axes of PCA capture population-wide interactions, with nearly all neurons contributing to the dominant components.

For the null and potent space analysis described below, we needed to select dimensionalities for M1 and PMd. We adapted a method developed by Machens et al (Machens et al. 2010) to estimate the dimensionality of our recorded populations. In brief, PCA provides an orthogonal basis set with the same dimensionality as the neural input. However, the variance captured by many of the higher dimensions (with the smallest eigenvalues) is typically quite small. We estimated the noise in the neural activity patterns using the trial-to-trial variation in the activity of each neuron. We sampled a random pair of trials for each reach direction and subtracted the

activity of each neuron. This gave an estimate of the variance of each neuron across two different reaches to each target. We then ran PCA on the neural "noise" space provided by this difference for all targets. We repeated this 1000 times, giving a distribution of eigenvalues for each of these noise dimensions. We used the 99% limit of these distributions to estimate the amount of noise variance explained for each dimension. This allowed us to put a ceiling on the amount of variance that could be explained by noise. The dimensionality was thus defined by the number of dimensions needed to explain 95% of the remaining variance.

Potent and null space calculation

Using the above method, we estimated the dimensionality of the M1 and PMd populations on each session. Since we identified a larger dimensionality for PMd than M1, there existed a "null space" in PMd, which encompasses PMd activity that has no net effect on M1 (Kaufman et al. 2014). To identify the geometry of the null and potent spaces, we constructed multi-input multi-output (MIMO) linear models, W, relating the N-dimensional PMd space to the O-dimensional M1 space (with N > O):

$$M = WP \tag{3}$$

M (O x t) and P (N x t) are matrices whose rows contain the activity of each PC for M1 and PMd, respectively, and whose columns contain the time points (t). We evaluated the quality of fit for these linear mappings using R^2 (Figure 4.8b,c), and found that M1 activity was well-predicted for most of the leading dimensions. We then performed singular value decomposition (SVD) of the matrix W (O x N) that maps PMd onto M1:

$$W = USV' \tag{4}$$

SVD decomposes the rank-deficient rectangular matrix *W* into a set of orthonormal basis vectors that allows us to define the null and potent spaces. For our purposes, the matrix *V'* defines the vectors that define the potent and null spaces, with the first N rows corresponding to the potent space, and the remaining M - N rows defining the null space (Figure 4.8a; Equation 5):

$$V = \begin{bmatrix} v_{11} & \cdots & v_{1M} \\ \vdots & \ddots & \vdots \\ v_{N1} & \cdots & v_{NM} \end{bmatrix} \quad (5)$$

We used only trials from the Baseline period of each session to find the axes for PCA, as well as the null and potent spaces. The Baseline trials were independent of the CF/VR trials used for both testing and training the GLM models, ensuring that we did not bias our results to find any specific solutions. However, we obtained nearly identical results if we used all of the data, or data only from the CF/VR trials, indicating that the null and potent spaces identified through this analysis did not change throughout the session. It is also important to note that the null and potent spaces, as with the PCA axes, typically comprised population-wide activity patterns, rather than sub-groups of neurons (Figure 4.8d).

It is worth noting that although we defined the Null space as activity which produced zero output in the low-dimensional M1 components, we could still predict M1 spiking quite well from the Null space in the cross-validated training data. Although potentially unintuitive, it worked well for a number of reasons. First, we identified the potent and null spaces using population-wide components, and used these activity patterns to predict the spiking single neurons. Additionally, for a given reach direction within a condition, the stereotyped activity in the null space could be well-correlated with activity in the potent space (and subsequently M1) due to the

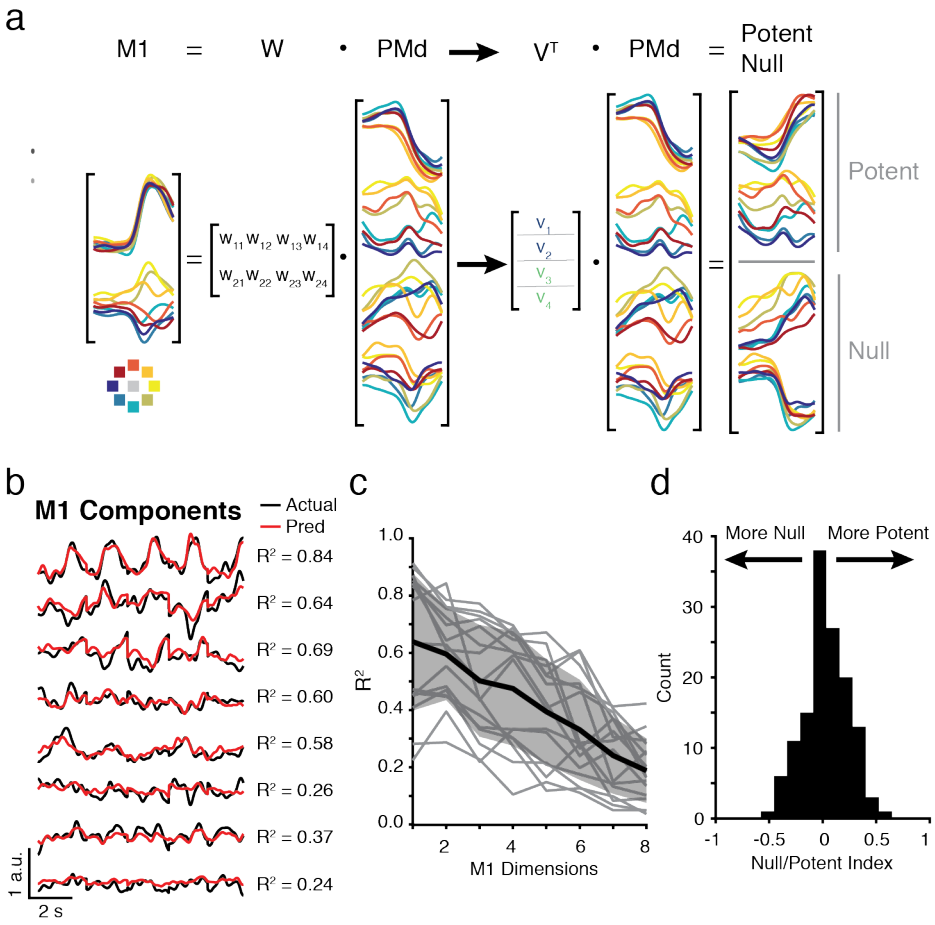

Figure 4.8 | Identifying output-potent and output-null spaces. a) Schematic representation of the method to identify output-potent and output-null spaces using trial-averaged data from a single session. Traces show the activity of components found by PCA on a representative session, with each color corresponding to one of the eight target directions. We used multi-linear regression to build a matrix *W* relating the activity of the M1 PCs (here, dimensionality of two) to the PMd PCs (here, dimensionality of four). Thus, *W* is a 2x4 matrix. Using Singular Value Decomposition (see Methods), we identified a matrix *V'*, the first two rows of which contained the basis vectors of the potent space, while the last two rows defined the null space. We multiplied the PMd PCs by this matrix to get the time-varying potent (top) and null (bottom) projections. b) Example predictions (red) of the first eight M1 PCs (black) from the first sixteen PMd PCs, with R^2 quantifying quality of fit for a single session. c) Summary of R^2 for M1 PC predictions across sessions (gray lines). Black line and gray shading indicate and mean and st.dev. across sessions. d) We attempted to identify potent or null subpopulations using an index that quantified the relative weights of each neuron onto the potent and null axes (see Methods). Values of 1 indicate the cell was exclusively potent, and values of -1 indicate the cell was exclusively null. The distribution of cells was centered around zero, indicating that the potent and null spaces captured population-wide activity patterns.

lawful relationship between them (Kaufman et al. 2014; Elsayed et al. 2016). It is under the condition of changing behavior that these correlations can begin to break.

Single neuron correlation analysis

We studied the correlations between individual neurons using the same smoothed firing rates we used for PCA. We then aligned each trial at the time of movement onset and isolated a window beginning 700 ms before and ending 800 ms after movement onset. We averaged across trials for each target direction, during both the pre-learning Baseline and the learning epochs. We excluded the first 50% of CF or VR trials to look at neural activity when adaptation was most complete. We then performed pairwise cross-correlations between all neurons recorded on each session during Baseline and late CF/VR. The coefficient of correlation values shown in Figures 4.1 and 4.6 were computed using the pairwise correlation values using all pairs of neurons from each session, as a means to quantify the similarity between the two conditions. For the heat maps shown in Figures 4.1 and 4.6, we normalized the range of each row to scale from -1 to 1 to enhance visualization. We then used a simple hierarchical clustering algorithm to sort the neurons in the Baseline condition. This same sorting order was used for the late CF heat map as a means of visually assessing the consistency in the correlation structure.

Generalized Linear Models

We trained Poisson Generalized Linear Models (Nelder and Baker 1972) (GLMs) to predict the spiking activity of individual neurons on a single-trial basis (Truccolo et al. 2010). GLMs extend Gaussian multilinear regression approaches for the Poisson statistics of neural spiking. We take weighted linear combinations of the desired covariates, x_i, such as limb kinematics:

$$\sum_i \theta_i x_i = X\Theta \quad (6)$$

The weighted covariates were passed through an exponential inverse link function. The exponential provides a non-negative conditional intensity function λ, analogous to the firing rate of the predicted neuron:

$$\lambda|X,\Theta = \exp(X\Theta) \quad (7)$$

The number of observed spikes, n, in any given time bin is assumed to a Poisson process with an instantaneous firing rate mean of λ:

$$n|\lambda \sim \text{Poisson}(\lambda\, dt) \quad (8)$$

Covariate inputs to the GLMs

In our analyses, we used GLMs to predict the spiking activity of single neurons based on the activity of the remaining population and kinematic signals. We binned the neural spikes at 50 ms intervals and downsampled the continuous kinematic signals to 20 Hz to match the binned spikes. We shifted the kinematic signals backwards in time by three spiking bins (150 ms) to account for transmission delays between cortical activity and the motor output. Previous studies have observed a broad range of delays (Moran and Schwartz 1999), so we convolved the kinematic signals with raised cosine basis functions centered at 0 ms and -100 ms, adapting the method of Pillow et al., where bases further back in history become wider (Pillow et al. 2008). By including these convolved signals as inputs to our GLM models, we allowed the neurons to have more flexible temporal relationships with the kinematics. Note that all GLM models included the same convolved endpoint position, velocity, and acceleration signals as covariates.

We trained two types of models: the Basic models included only kinematic covariates, while the Full models included both the kinematic covariates and the spiking activity of the single-

neuron populations (Figure 4.9a). For the GLMs with single neuron inputs, we trained three different types of Full models. M1-M1 models predicted the spiking activity of each M1 neuron from the activity of all other M1 neurons recorded on the same session, PMd-PMd models predicted the spiking of each PMd neuron from all other PMd neurons, and PMd-M1 models predicted M1 neurons using the activity of all PMd cells. For the GLM analysis with potent and null components (Figures 4.4 and 4.5), we used low-dimensional summaries of PMd population activity as inputs to the GLMs, rather than single neurons. For PCA-M1, we projected PMd activity into the PCA space (see above) and selected only the first 16 dimensions as input to the GLM. Since PCA captures population-wide covariance patterns, we expected that this approach would provide nearly identical results to the single neuron models of PMd-M1, and it was included primarily as a control. For Pot-M1 and Null-M1, we projected the time-varying PMd signals onto the basis vectors for the potent and null space, respectively (see above). We then used these time-varying signals as inputs to GLMs to predict the spiking of M1 neurons.

Training the GLMs

We trained the models using the last 50% of CF or VR trials when behavior was most stable, including only trials where the monkeys made successful reaches to acquire the outer target (reward trials). This allowed us to test the generalization of the GLMs during the early adaptation trials. For the CF, it was important to both train and test the GLMs using trials from the CF epoch to avoid extrapolating between the Null and CF conditions. When we imposed the CF, it changed the relationship between the kinematics and dynamics of limb movement. Thus, if we trained the GLM on Baseline trials, the relationship between kinematics and neural activity

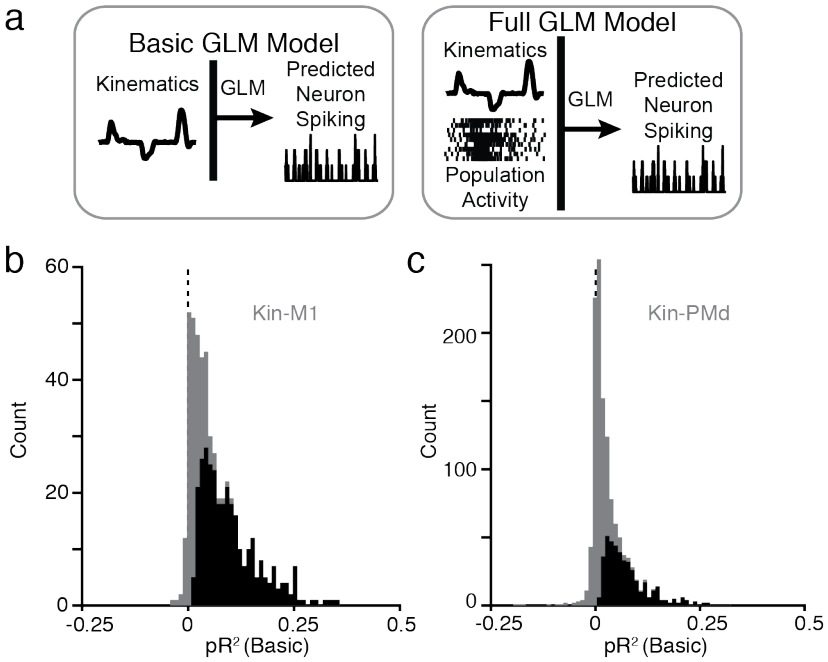

Figure 4.9 | Types of GLM models. a) Schematic representation of the GLM models. The Basic model included only kinematic covariates (see Methods), while the Full model included both kinematics and neural activity. The relative pseudo-R^2 metric was a comparison between these two models. b) Distribution of cross-validated pseudo-R^2 values for predictions of all M1 neurons from all sessions with the Basic model (gray). Black overlaid distribution shows cells with significant model fits (see Methods). c) Same as Panel b, but for predictions of PMd neurons.

changed immediately on CF trials (Cherian et al. 2013), leading to poor GLM generalization for all models. By both training and testing within the block of CF trials, we avoided the problem of extrapolating to new dynamics conditions. Although the VR sessions did not have this problem, we adopted this same approach the sake of consistency.

We trained the models using a maximum likelihood method (glmfit in Matlab, The Mathworks Inc). In the case of our full population spiking models, we had dozens to hundreds of covariate inputs for a single predicted output. Although we had very large numbers of training data points (typically on the order of 10,000 samples), there is the possibility our models were impaired by overfitting. We guarded against overfitting using ten-fold cross-validation of our training dataset. We also repeated our analyses using Lasso GLM for regularization and observed nearly identical results (data not shown). We thus chose to use the non-regularized GLM for simplicity and to reduce the computational load, since it did not impact our results.

Evaluating GLM performance

We evaluated GLM performance using a particular formulation of the pseudo-R^2 (pR^2). The pR^2 is analogous to the R^2 commonly used in model-fitting with Gaussian statistics, but it is generalized to incorporate the assumed Poisson statistics of the neural spiking data:

$$\mathrm{pR}^2 = 1 - \frac{\log L(n) - \log\, L(\hat{\lambda})}{\log L(n) - \log\, L(\bar{n})} \tag{9}$$

The pR^2 finds the difference in log-likelihood between the observed spiking data (n) and the model predictions ($\hat{\lambda}$). This value is compared against the difference in log-likelihood for the mean of the dataset ($\bar{n}$). We used the Likelihood (L) for Poisson data according to:

$$L = \prod_{t=1}^{T} \mathrm{Poisson}(n_t|\lambda_t) = \prod_{t=1}^{T} \frac{\lambda_t^{n_t} \exp(-\lambda_t)}{n_t!} \tag{10}$$

And thus, the log-likelihood (log L) across all time bins (t) of a given spike train is:

$$\log L = \sum_{t=1}^{T}(n_t \log \lambda_t - \lambda_t - \log n_t!) \quad (11)$$

Although the upper bound for pR^2 is one, poor model fits can be less than zero. A pR^2 of one indicates a perfect model fit, a value of zero indicates that the model prediction performs as well as finding the mean of the data, while values less than zero indicate that the model performed worse than merely fitting the mean. Typical pR^2 values are smaller in magnitude than those typically found with the Gaussian R^2. When evaluating GLM fits, we used a bootstrapping procedure with 1000 iterations to obtain 95% confidence bounds on the pR^2 value. We considered a model fit to be significant if this bootstrapped confidence interval was above zero, indicating that the model helped to explain the spiking activity. For many analyses, we used the relative pseudo-R^2 (rpR^2), which directly compares two separate GLM models. While pR^2 compared the log-likelihood of the model predictions to the mean of the data, the rpR^2 compares the predictions of a Full model to a Basic model with fewer covariates.

$$\mathrm{rpR^2(Basic, Full)} = 1 - \frac{\log L(n) - \log\, L(\hat{\lambda}_F)}{\log L(n) - \log\, L(\hat{\lambda}_B)} \quad (12)$$

Here, $\hat{\lambda}_F$, the Full model prediction, which includes both the kinematics and the population spiking, is compared to $\hat{\lambda}_B$, the prediction of the Basic model, which includes only kinematics. This metric thus quantifies the improvement in performance afforded by the additional neuronal inputs. Positive values indicate that the Full model performed better than the Basic model, while negative values indicate that predictions were better with kinematics alone. As with the pR^2, we obtained confidence bounds with a bootstrapping procedure and assessed significance by determining if the lower bound was above zero. This indicated that the addition of population

spiking added information over the kinematics alone, and thus could be capturing meaningful functional relationships between the population and the predicted cell.

For the plots intended to visualize the time course of GLM model changes, such as Figure 4.3b, we predicted neural spiking on individual trials. However, predictions could be quite noisy with such small numbers of data points. For example, if a cell fired very few spikes on a particular trial, the pR^2 may be quite low, even though the model generally performed quite well. To remove some of this variability, we smoothed the trial-to-trial predictions for each neuron (as well as the overlaid behavior) with a moving average. We chose a window of 30 trials, though we observed similar (but slightly more variable) traces even down to window sizes of 5-10 trials. Since there were rapid behavioral improvements in the early trials, we padded the beginning and end with NaNs, each of a length of half of the window size. This helped to prevent averaging out the changing behavioral effects, with the tradeoff of slightly increasing noise. In practice, our results were similar without this padding.

Selecting cells with significant population relationships

For most of our analyses, we studied cells that were well-predicted by our GLMs. We determined this by two main criteria using ten-fold cross-validation on the training data. First, we required that the Basic pR^2 was significantly above zero. This reduced the pool of candidate cells to 522/918 (57%) in M1 and 612/2221 (28%) in PMd, but was necessary so that the rpR^2 would be well defined. Qualitatively, we obtained similar results when we relaxed this criterion to include more cells. We also required that the rpR^2 was significantly above zero. We only included cells that were significantly above zero for all ten of the folds for all pR^2 and rpR^2. This

method was very conservative, but ensured that we only studied cells that were reliably predicted.

Statistical tests

For the GLM models, we assessed the significance of model fits empirically using a bootstrapping procedure on cross-validated data as described above. We additionally used two-sample Student's t-tests to compare the distributions of pseudo-R^2 changes in Early and late learning. For Figure 4.2e, this was done using the raw rpR^2 values. For Figures 4.3a, 4.4d, 4.5, and 4.7a, the t-test was done using the normalized change in rpR^2.

CHAPTER 5

DISCUSSION

SUMMARY OF FINDINGS

In the previous chapters, I presented the results of a series of experiments designed to investigate the role of the primary motor and dorsal premotor cortex in short-term motor adaptation. I trained monkeys to make movements with the altered movement dynamics of a curl field, or with the perturbed feedback of a static visuomotor rotation. Throughout these chapters, I studied how neural activity in M1 and PMd changed to drive behavioral adaptation. In Chapter 2, I presented a detailed analysis of the tuning properties of individual M1 neurons during CF adaptation. I provided evidence that CF learning does not result from changes in the functional properties of the recorded M1 neurons. I hypothesized that behavioral adaptation was mediated by altered recruitment of M1 neurons by PMd, rather than plastic reorganization within M1. In Chapter 3, I presented a new conceptual framework for interpreting neural population activity. This framework was used in Chapter 4 to directly investigate the relationship between PMd and M1 during CF and VR adaptation. I found that throughout both learning paradigms, functional relationships within M1 and PMd were unchanged. However, there was a specific, behavior-related change in the processing performed by PMd when preparing inputs for M1. In this chapter, I will discuss the implications of these findings, propose plausible mechanisms underlying these observations, and speculate on interesting future directions.

POSSIBLE MECHANISMS UNDERLYING MOTOR LEARNING

Short term motor learning may not require structural changes

The results of the preceding chapters suggest that, at least on the timescale of a single experimental session, plasticity does not occur in the motor cortices during motor learning. In Chapter 2, I studied the functional relationship between neurons in M1 and the motor output. I

found that the changes in kinematic tuning can be accounted for by the altered dynamics of the curl field, and there were no changes with a time course that could explain behavioral adaptation. Based on this observation, I concluded that there were no changes in the functional outputs of these cells, or in the downstream circuits such as the spinal cord. I interpret this observation to mean that M1 must be recruited differently by upstream premotor circuits. There remains the possibility that the altered recruitment of the recorded populations came from upstream plastic changes within M1, though this idea would be inconsistent with numerous psychophysical studies suggesting that M1 is not essential for acquisition of a motor skill (Richardson et al. 2006; Herzfeld et al. 2014). The GLM results of Chapter 3, discussed below, further support that short-term learning is not a result of reorganization within M1.

Transition to long-term learning

The results of Chapters 2 and 4 suggest that, on the timescale of seconds to minutes, behavior is adapted without engaging cortical plasticity. However, on the scale of hours to days or even years, it is known that the cortex does undergo plastic changes (Elbert et al. 1995; Nudo et al. 1996; Rioult-Pedotti et al. 1998; Kleim et al. 1998, 2004; Peters et al. 2014). This distinction of timescales has been shown elegantly in the *Aplysia californica*, a marine mollusk. A naive aplysia learns to avoid an approaching stimulus after just a single exposure. A study showed that several hours after this single exposure, there were significantly more, and larger, dendritic spines in the presynaptic terminals of sensory neurons that controlled the reflex, corresponding to the maintained memory (Bailey and Chen 1983). However, the changes in synaptic connectivity were not present in the minutes following learning, even though the subjects showed complete behavioral habituation (Bailey and Chen 1988).

Although the A*plysia* is a simple organism compared to primates, similar observations were made in a motor learning experiment with rats. No structural synaptic changes were observed during short-term behavioral improvements, while large-scale synaptogenesis and cortical reorganization occurred during long-term learning on the scale of days (Kleim et al. 2004). These observations are compelling. They suggest that brain has evolved mechanisms to adapt and maintain behavior on a short timescale without engaging the costly and time-consuming process of synaptic plasticity. When it is advantageous to remember a skill for an extended period of time, it can then be hardcoded into the synaptic circuitry. In the motor system, structural changes in M1 can be seen within days of practice (Peters et al. 2014). After the initial learning session, where the cerebellum plays a key role, the memory may move to the cortex to be consolidated (Galea et al. 2011). This can facilitate the long-term recall of the skill. There is psychophysical evidence for this theory, since the motor cortex does not appear to be crucial for the initial adaptation (Richardson et al. 2006), while a functional cerebellum is necessary (Smith and Shadmehr 2005). However, motor cortical disruption can impair the recall, suggesting a prominent role in consolidation.

The cerebellum as a candidate for short-term learning

What is the neural mechanism underlying the change in null-space processing I observed in the CF task? I propose that, during CF learning, PMd computes a new motor plan through its interactions with the cerebellum. The cerebellum is known to play an essential role in error-based learning (Diedrichsen et al. 2005), and is activated during learning of both curl fields and visuomotor rotations (Diedrichsen et al. 2005). The cerebellum is intimately connected with the parietal, premotor, and motor cortices (Dum and Strick 2003), and has a highly specialized

circuitry that makes it a prime candidate to engage with the motor areas of the cerebral cortex for trial-to-trial learning. The cerebellar Purkinje cells have unique physiological properties with mechanisms for substantially changing their output on a rapid timescale in response to such errors. Purkinje cells are bistable, with two activity states characterized by different intrinsic excitability and rates of tonic output (Loewenstein et al. 2005). Movement error signals are encoded by climbing fibers from the inferior olive, which synapse on Purkinje cells in the cerebellar cortex and cause complex spikes (Kawato and Gomi 1992). Activation of these climbing fibers can switch the Purkinje cells between the two states, causing persistent changes in cerebellar output in less than a second (Loewenstein et al. 2005).

The specialized circuitry and anatomical connectivity of the cerebellum provide plausible mechanisms for rapid behavioral adaptation. Errors experienced during the previous CF trial alter the output of the cerebellar Purkinje cells (Kawato and Gomi 1992), which may be relayed to PMd via the dentate nucleus (Dum and Strick 2003). PMd integrates this signal with visual and intention information from areas such as parietal cortex (Batista et al. 2007), allowing it formulate a new motor plan that can compensate for the altered dynamics. Given the anatomical connectivity between parietal cortex and cerebellum (Dum and Strick 2003), a similar mechanism could be employed in response to visual errors in the VR task, where the visual cue is transformed into a new endpoint of the planned movement.

This cerebellar-dependent framework may also be consistent with the slow and fast timescales of learning discussed in Chapter 1 (Karni et al. 1998). The fast phase of learning is characterized by rapid behavioral changes, but is highly susceptible to interference. The slow phase occurs over a longer timescale, but during this time the motor memory is less susceptible to interference. Within the framework of the CF learning experiments, I speculate that during a

session, the unique properties of Purkinje cells, including bistability, are exploited to rapidly change the output of the cerebellum in response to errors. This gives the rapid trial-to-trial learning. Over longer time periods, the cerebellum can undergo a more stable change using LTP or LTD (Zheng and Raman 2010). Ultimately, as described above, during long-term learning the cerebral cortex can begin to alter connectivity (Rioult-Pedotti et al. 1998; Kleim et al. 2004), accounting for the slower phase of learning.

INTERPRETATION OF THE RESULTS

Functional and synaptic connectivity

In the preceding chapters, I investigated whether motor learning causes reorganization of motor representations or functional connectivity in the motor cortex. The most probable mechanisms for such reorganization would be through synaptic plasticity. There are a number of means through which plasticity can be achieved. New synaptic connections can be formed between an axon and the dendrites of a second neuron. Although this process is believed to require a substantial amount of time (Bailey and Kandel 1993), the efficacy of existing synapses can be changed on a shorter timescale. In a common model, the coordinated activity of a pair of neurons triggers a biochemical process that increases or decreases the efficacy of that synapse (Markram and Tsodyks 1996; Nevian and Sakmann 2006; Delattre et al. 2015). Additionally, the intrinsic excitability of a cell can be regulated through leak currents or tonic inputs (Zheng and Raman 2010), such that the same synaptic input produces a different spiking rate. While other, more specialized, mechanisms exist, all of these processes will result in a change in the input/output relationship between neurons.

The evidence presented in the previous chapters suggests that short-term motor learning likely does not involve structural reorganization within the motor cortices. I developed analyses that identified the functional relationships between single neurons and behavior (Chapter 2), or between populations of neurons (Chapter 4). Since I used chronic extracellular array recordings, I cannot identify cell types, or trace synaptic connectivity between the recorded cells. Indeed, it is likely that most of the recorded cells only have indirect relationships (via one or more intermediate synapses, or via common inputs). However, I expect that any true synaptic connectivity would be captured within these functional models. Furthermore, I expect that any change in synaptic connectivity or efficacy would be reflected in the ability of the GLM models to predict spiking. This is expected for synaptically connected cells, where functional connectivity directly corresponds to anatomical connectivity (Gerhard et al. 2013). It is also true, however, at the level of the population, where connectivity between randomly selected cells is likely to be indirect (e.g. through one or more intermediate synapses). A plastic change between a subset of neurons in a network can be expected to change the functional interactions between those neurons and the surrounding population, even if they do not share direct, synaptic connections with those neurons (Ahissar et al. 1992). A change in functional connectivity does not necessarily imply a change in structural connectivity. However, if I observe no functional change, we can expect that there has been no structural change, or at least that any structural changes were too small to have a functional effect. The GLM measurements of functional connectivity can thus serve as reasonable proxies to evaluate whether there have been structural changes within the network.

Interpreting the GLM models

The GLM analysis was designed to study the interactions between M1 and PMd. The stability of the GLM models within each brain area (M1-M1 and PMd-PMd) suggests that there were no structural changes within the two populations. However, there are three potential explanations for the learning-related functional change between PMd and M1. 1) There may be a specific change within M1 that breaks the relationship with PMd. However, in this case, one would predict the M1-M1 model to show a similar change, which was not observed. 2) M1 may be directly recruited by other brain areas, for example the supplementary motor area (Padoa-Schioppa et al. 2002, 2004), without the influence of PMd. This explanation seems unlikely, since I observed adaptive changes in the neural activity within PMd. While compensation for the CF likely involves other brain areas (Baraduc et al. 2004), if adaptation were mediated entirely outside of PMd, one would not expect its activity to change. 3) PMd plays a specific role in recruiting M1 to compensate for the altered dynamics. The null and potent space analysis of Chapter 4 most strongly support this last explanation. Both the null and potent space GLM models would be equally impaired by changes within M1 (Explanation #1) or by influence of other areas (Explanation #2), yet only the null space model showed a learning-related change.

Motor planning in PMd is crucial for CF learning

The neural control of movement is distributed across a number of cortical and subcortical structures (Kalaska et al. 1997). My data cannot exclude the possibility that adaptation occurs in parallel with the circuitry of M1 and PMd, though the null space processing changes of Chapter 4 suggest a specific role of PMd in CF adaptation. Even if it is not the sole learning pathway, PMd is at least playing a crucial role in preparing the adapted movements with the CF. The

results can be interpreted in the context of the work of Churchland et al., where PMd serves to set the preparatory state for the movement activity of M1 (Churchland et al. 2010a, 2012; Kaufman et al. 2014). Part of the role of PMd may be to transform sensory information about the cue and intention into a motor plan (Shen and Alexander 1997; Batista et al. 2007), and PMd is known to modify its preparatory activity based on the behavioral outcome and context (Kurata and Wise 1988; Vaadia et al. 1988; di Pellegrino and Wise 1993; Cisek and Kalaska 2005). Evidence from psychophysical experiments suggests that a specific motor plan is important for CF learning (Sheahan et al. 2016). Subjects can learn opposing curl fields only if the task conditions are such that both curl fields require separate motor plans. The results of Chapter 4 suggest that PMd may be an important locus for this task-specific preparatory activity.

Visuomotor rotation learning is upstream of PMd

Learning the VR perturbation caused no plastic changes within or between the M1 and PMd populations. Intriguingly, I did not observe the change in null-space processing that occurred with the CF. It is important to note that in the VR task, PMd activity still changed during learning, and preparatory activity in PMd still played an important role in the reaching movements. The key difference, though, is that for a given hand trajectory with the VR, the processing performed by PMd was unchanged. I interpret this observation to mean that VR learning occurs upstream of PMd. The VR is known to engage areas of parietal cortex (Diedrichsen et al. 2005), and it has been proposed that VR learning can occur through interactions between motor cortex and areas of parietal cortex such as the parietal reach region (PRR) (Snyder et al. 1997), without changing the movement representation in M1 (Tanaka et al.

2009). Since I did not record from PRR in these experiments, I cannot provide direct evidence for this claim, though my results support this theory.

LIMITATIONS OF THE ANALYSES AND EXPERIMENTS

Practical considerations of the single-neuron cosine tuning model

In the analyses of Chapter 2, I focused exclusively on the preferred direction (PD) of each cell. However, the cosine tuning curves that I used were described by two other parameters: the mean firing rate and depth of modulation. In a series of analyses that were not included in the final paper, I studied the adaptive changes of these two parameters for all cells. Unlike the PDs, which showed consistent rotations in the direction of the CF, the changes in the other parameters were seemingly random. Cells were equally likely to increase or decrease their firing rates or modulation depth, and we observed no progressive shift at the population-level that could explain behavior. These results are consistent with the ultimate conclusion of Chapter 2 that adaptation is mediated by the altered recruitment of M1 neurons. The new recruitment patterns increase the activity of some cells and decrease the activity of others in a complex manner that differs across reach direction and over time.

A limitation of the cosine tuning analysis is that many cells are not well-described by a cosine function in hand coordinates. Thus, a large proportion of the neural population was excluded from analysis. It is possible that the adaptive changes could have occurred in the properties of these excluded cells, though this is not likely to be the case. The population-level analyses of Chapters 3 and 4 provide one avenue to address this limitation. Since I included every M1 neuron in the GLM and PCA analyses, I was not biased towards a particular

subpopulation. Yet, I reached a similar conclusion that M1 is not subject to reorganization during within-session learning.

Although extracellular recording methods may be biased towards recording the activity of larger excitatory cells, the neural population likely included a variety of cell types. Thus, the cells I recorded can be viewed as a random sample of cells that ultimately generate movement. I saw no evidence that the cosine tuning model, or the tuning changes with the CF, identified any subpopulations of neurons. In an unreported analysis, I tested for this specifically by separating the neural population into putative excitatory and inhibitory neurons based on the width of the recorded spike waveform (Kaufman et al. 2013). Although I saw no difference in tuning behavior during CF adaptation for these two populations, there is evidence that these cell types may contain different information about the control movement (Best et al. 2016). It may be enlightening in future experiments to perform a more careful analysis of neural sub-types.

Limitations of the population GLM approach

In Chapters 3 and 4, I included kinematic variables as covariates into the GLM. This was intended to remove the confounding effect of behavior-related common inputs to the population of recorded cells. By removing this confound, I hoped to study more closely the precise interactions between the neural populations. However, the inclusion of only kinematics may also be a limitation of the analysis. Given the considerable evidence that activity in M1 correlates with variables such as force (Evarts 1968) and muscle activation (Morrow and Miller 2003), kinematic variables such as endpoint position may not be the most appropriate inputs to the GLMs to predict M1 spiking. Ideally, the neural recordings would be complemented by recordings of EMG or joint torques. These variables could then be included along with the

kinematics to provide a more complete view of behavior, and allow me to study more accurately the neuron-to-neuron interactions. As a control, I included as inputs to the GLM, the readings of the force transducer in the handle (this setup was described in Chapter 2) and the results were unchanged. I ultimately used the kinematics as a high-level correlate of behavior, and did not intend to impose any precise model of M1 function. Thus, I do not anticipate that the results and conclusion are dependent on this detail.

A second limitation of the GLM approach is that I can only study cells that were well-predicted by my GLMs. In practice, I had to exclude many cells due to low pseudo-R^2 values. I do not expect that excluding these cells caused me to draw any incorrect conclusions, nor that these excluded cells represent a distinct subclass that might behave differently. The cells that were excluded primarily had low mean firing rates, and the conclusions did not depend on the choice of minimal firing rate. Furthermore, the exclusion only applied to those cells I chose to predict with the GLM output; I included the entire population as inputs. Any structural change within these excluded cells would have been apparent in the performance of my GLM models. Although I did not directly compare if these excluded cells are the same as those excluded by the cosine tuning analysis of Chapter 2, in both cases there was a relationship between quality of fit and firing rate, so there is likely to be overlap.

FUTURE DIRECTIONS

There are a number of compelling questions that could be addressed in subsequent studies. The results of Chapter 4 suggest that PMd helps to learn the altered dynamics of the curl field, potentially through interactions with the cerebellum. Since M1 also shares many connections with the cerebellum (Dum and Strick 2003), the learning effect may not be isolated to PMd

alone. Although the planning activity in M1 is weaker than PMd (Cisek and Kalaska 2005; Kaufman et al. 2010), neurons in M1 do have meaningful preparatory activity (Kaufman et al. 2013). In Chapter 4, I considered the activity of PMd with respect to M1 to determine the output-null space. An interesting follow-up experiment would be to record electromyograms (EMG) of proximal limb muscles during CF learning. As Kaufman et al. showed, null space planning can be seen in M1 with respect to muscles, allowing the brain to prepare to move without causing that movement (Kaufman et al. 2014). Some of this planning may be altered within M1 during CF learning, which could appear in the null space of muscle activity in a similar manner as the PMd to M1 null space.

Many other brain areas could be studied using this framework. Sensory feedback is necessary to get the error signal for CF learning (Wolpert et al. 1995), and primary somatosensory cortex (S1) is necessary to adapt to CF perturbations (Mathis et al. 2017). Simultaneous population recordings from M1/PMd and S1 could shed light on how they interact. Does corticocortical somatosensory feedback help PMd and M1 to develop the new motor plan, or is the S1 processing only used to compute errors that are sent to the cerebellum or other subcortical structures? Similarly, it would be interesting to test the hypothesis that parietal cortex learns to compensate for the VR. I hypothesize that if I recorded from populations of neurons from PRR and PMd simultaneously I would observe a similar change in null space processing from PRR to PMd, as PRR adjusts the visuomotor mapping from goal into a planned action (Buneo et al. 2002).

Another compelling, but more difficult, experiment would be to record from Purkinje cells in the cerebellar cortex simultaneously with PMd. This could provide a direct test of my interpretation of the experiments. I would predict that cerebellar Purkinje cells should show

functional changes with respect to the motor output, but have a fixed functional interaction with PMd. This would imply that the Purkinje cell changes do, indeed, cause the adapted behavior, and that the PMd effect arises from interactions with the cerebellum through the fixed architecture.

One further set of experiments, which I piloted during the course of these experiments, is to causally probe the motor system to identify any functional changes in motor or premotor cortex. In the first paradigm, I used intracortical microstimulation (ICMS) of M1 to evoke twitches in the proximal limb muscles. The twitches ultimately produced a force vector at the hand, which could be measured. This provided an input-output mapping between motor cortex and hand movement. By probing the system in this manner before, during, and after CF learning I hoped to test directly, whether plastic changes occurred downstream of the M1 cells that I recorded. If the magnitude and direction of the evoked force remains unchanged throughout learning, then it is unlikely the underlying circuitry has changed. I later developed a similar approach to validate the PMd to M1 connectivity of Chapter 4. In this paradigm, I used ICMS within PMd and attempted to evoke responses in the spiking of M1 cells as a direct test of functional connectivity. I could then compare these functional maps throughout learning to validate the results of the GLM. These experiments remain promising avenues for continued research.

In the experiments of Chapter 4, I imposed the VR perturbation suddenly, causing large errors on the first trial. At least in humans, this engages a cognitive strategy of voluntary error correction in order to compensate. There is compelling evidence from psychophysical studies that subjects can learn to adjust for gradual rotations, where the size of the perturbation is increased in small increments, without being aware that the perturbation has been applied or that they have adapted (Kagerer et al. 1997). After the rotation has gradually ramped up to the full

magnitude, the subjects show similar aftereffects to those observed with the sudden rotation, indicating that implicit learning was achieved in both cases. If adaptation to the sudden VR perturbation occurs in part as a result of an explicit cognitive strategy, or re-aiming the desired trajectory, then the gradual perturbation may more directly engage an implicit motor learning process (Mazzoni and Krakauer 2006). It is possible that during the sudden perturbation, the monkeys can form a new explicit motor plan and compensate before the inputs to PMd. An interesting experiment would be to repeat the analyses using a gradual onset perturbation to determine if it has an effect on PMd and M1 interactions.

There are also possible experiments related to the source of long-term learning and savings across sessions. In the current experiments, I focused on within-session learning. With a long-term study of the same skill, one could ask if the null-space processing I observed is necessary even for consolidated skills. It is possible that once a behavior is well-learned the memory can reside exclusively in the cortical circuitry, and that PMd will not need to perform the new null-space processing that I observed in the CF experiment. However, I would predict that null-space planning is necessary to adapt behavior, even for well-practiced movement. The effect of savings simply changes the rate at which you can achieve proficiency. Even expert athletes and musicians must warm up before they can achieve peak performance (Ajemian et al. 2010). Perhaps plasticity, then, only serves to make the online control faster and more efficient, not to completely replace the need to adapt online.

These experiments provide compelling evidence that studying population-level interactions can give insight that cannot be seen at the level of single neurons. However, while the analyses treated the recorded neural population as random samples of the underlying neural manifold (Gao and Ganguli 2015), there is a great degree of specialization in the cell types and

connectivity of the cortex. Another interesting future direction would be to explore how the population activity relates to the known cell-types and laminar organization of M1 and PMd. Although PMd has many corticospinal projections (Dum and Strick 1991), it lacks the large Betz cells that allow for high information transfer speeds (Bucy 1935). Although in these experiments I have treated the population as a random sample of underlying latent signals, there may be additional specialization that could be exploited to better understand the system. A limitation of the Utah arrays that I used is the inability to target particular cell types. Modern recording techniques such as two-photon calcium imaging allow large populations of neurons with particular genetic signatures or projection locations to be recorded, providing a plausible avenue to begin to answer these types of questions.

CONCLUSION

For most species, the ability to flexibly adapt movements is essential for survival. In this dissertation, I explored the mechanisms by which the dorsal premotor cortex can coordinate with the primary motor cortex, and presumably other brain areas, to change the behavioral output on a short timescale. Although there remains much work to be done to fully understand how this process occurs, these experiments are an intriguing step forward, with implications for understanding the mechanisms available to the brain to generate behavior using neural populations. These results also have strong implications for the field of brain-machine interfaces. By exploiting knowledge of the natural functions of different brain areas, as well as taking a principled approach to using cortical signals for brain decoders, we can develop more intuitive and adaptable interfaces with strong clinical potential.

REFERENCES

Afshar A, Santhanam G, Yu BM, et al (2011) Single-trial neural correlates of arm movement preparation. Neuron 71:555–564. doi: 10.1016/j.neuron.2011.05.047

Ahissar E, Vaadia E, Ahissar M, et al (1992) Dependence of cortical plasticity on correlated activity of single neurons and on behavioral context. Science (80-) 257:1412–5. doi: 10.1126/science.1529342

Ahrens MB, Li JM, Orger MB, et al (2012) Brain-wide neuronal dynamics during motor adaptation in zebrafish. Nature 485:471–477. doi: 10.1038/nature11057

Ajemian R, D'Ausilio A, Moorman H, Bizzi E (2010) Why professional athletes need a prolonged period of warm-up and other peculiarities of human motor learning. J Mot Behav 42:381–8. doi: 10.1080/00222895.2010.528262

Albus JS (1971) A theory of cerebellar function. Math Biosci 10:25–61. doi: 10.1016/0025-5564(71)90051-4

Alexander GE, Crutcher MD (1990) Preparation for movement: neural representations of intended direction in three motor areas of the monkey. J Neurophysiol 64:133–50.

Arce F, Novick I, Mandelblat-Cerf Y, et al (2010a) Combined adaptiveness of specific motor cortical ensembles underlies learning. J Neurosci 30:5415–25. doi: 10.1523/JNEUROSCI.0076-10.2010

Arce F, Novick I, Mandelblat-Cerf Y, Vaadia E (2010b) Neuronal correlates of memory formation in motor cortex after adaptation to force field. J Neurosci 30:9189–98. doi: 10.1523/JNEUROSCI.1603-10.2010

Asanuma H, Larsen KD, Yumiya H (1979) Direct sensory pathways to the motor cortex in the monkey: A basis of cortical reflexes. In: Integration in the Nervous System. Igaku-Shoin, New York, pp 223–238

Ashe J (1998) Force and the motor cortex. Behav Brain Res 87:253–270. doi: 10.1016/S0166-4328(96)00145-3

Bailey CH, Chen M (1983) Morphological basis of long-term habituation and sensitization in Aplysia. Science 220:91–3. doi: 10.1126/science.6828885

Bailey CH, Chen M (1988) Morphological basis of short-term habituation in Aplysia. J Neurosci 8:2452–2459.

Bailey CH, Kandel ER (1993) Structural Changes Accompanying Memory Storage. Annu Rev Physiol 55:397–426. doi: 10.1146/annurev.ph.55.030193.002145

Baraduc P, Lang N, Rothwell JC, Wolpert DM (2004) Consolidation of dynamic motor learning is not disrupted by rTMS of primary motor cortex. Curr Biol 14:252–6. doi: 10.1016/j.cub.2004.01.033

Baraduc P, Wolpert DM (2002) Adaptation to a visuomotor shift depends on the starting posture. J Neurophysiol 88:973–81.

Batista AP, Santhanam G, Yu BM, et al (2007) Reference frames for reach planning in macaque

dorsal premotor cortex. J Neurophysiol 98:966–83. doi: 10.1152/jn.00421.2006

Bauswein E, Fromm C, Werner W, Ziemann U (1991) Phasic and tonic responses of premotor and primary motor cortex neurons to torque changes. Exp Brain Res 86:303–310. doi: 10.1007/BF00228953

Best MD, Takahashi K, Suminski AJ, et al (2016) Comparing offline decoding performance in physiologically defined neuronal classes. J Neural Eng 13:26004. doi: 10.1088/1741-2560/13/2/026004

Betz V (1874) Anatomischer Nachweis zweier Gehirncentra. Cent für die medizinischen Wissenschaften 12:578–580, 595–599.

Brashers-Krug T, Shadmehr R, Bizzi E (1996) Consolidation in human motor memory. Nature 382:252–5. doi: 10.1038/382252a0

Brodmann K (1909) Vergleichende Lokalisationslehre der Grosshirnrinde. Johann Ambrosius Barth, Leipzig

Brooks JX, Carriot J, Cullen KE (2015) Learning to expect the unexpected: rapid updating in primate cerebellum during voluntary self-motion. Nat Neurosci 18:1–10. doi: 10.1038/nn.4077

Broome BM, Jayaraman V, Laurent G (2006) Encoding and decoding of overlapping odor sequences. Neuron 51:467–82. doi: 10.1016/j.neuron.2006.07.018

Bruno AM, Frost WN, Humphries MD (2015) Modular deconstruction reveals the dynamical and physical building blocks of a locomotion motor program. Neuron 86:304–318. doi: 10.1016/j.neuron.2015.03.005

Bucy PC (1933) Electrical Excitability and cyto-architecture of the premotor cortex in monkeys. Arch Neurol Psychiatry 30:1205–1225.

Bucy PC (1935) A comparative cytoarchitectonic study of the motor and premotor areas in the primate cortex. J Comp Neurol 62:293–331. doi: 10.1002/cne.900620203

Buneo C a, Jarvis MR, Batista AP, Andersen R a (2002) Direct visuomotor transformations for reaching. Nature 416:632–6. doi: 10.1038/416632a

Bütefisch CM, Davis BC, Wise SP, et al (2000) Mechanisms of use-dependent plasticity in the human motor cortex. Proc Natl Acad Sci 97:3661–5. doi: 10.1073/pnas.050350297

Cabel DW, Cisek P, Scott SH (2001) Neural activity in primary motor cortex related to mechanical loads applied to the shoulder and elbow during a postural task. J Neurophysiol 86:2102–8.

Caithness G, Osu R, Bays P, et al (2004) Failure to consolidate the consolidation theory of learning for sensorimotor adaptation tasks. J Neurosci 24:8662–71. doi: 10.1523/JNEUROSCI.2214-04.2004

Carmena JM, Lebedev M a, Crist RE, et al (2003) Learning to control a brain-machine interface for reaching and grasping by primates. PLoS Biol 1:E42. doi: 10.1371/journal.pbio.0000042

Chapin JK, Nicolelis MAL (1999) Principal component analysis of neuronal ensemble activity

reveals multidimensional somatosensory representations. J Neurosci Methods 94:121–140. doi: 10.1016/S0165-0270(99)00130-2

Chen SX, Kim AN, Peters AJ, Komiyama T (2015) Subtype-specific plasticity of inhibitory circuits in motor cortex during motor learning. Nat Neurosci 18:1109–1115. doi: 10.1038/nn.4049

Cheney PD, Fetz EE (1980) Functional classes of primate corticomotoneuronal cells and their relation to active force. J Neurophysiol 44:773–791.

Cheney PD, Fetz EE (1984) Corticomotoneuronal cells contribute to long-latency stretch reflexes in the rhesus monkey. J Physiol 349:249–72.

Cheney PD, Fetz EE, Palmer SS (1985) Patterns of facilitation and suppression of antagonist forelimb muscles from motor cortex sites in the awake monkey. J Neurophysiol 53:805–820.

Cheng EJ, Scott SH (2000) Morphometry of Macaca mulatta forelimb. I. Shoulder and elbow muscles and segment inertial parameters. J Morphol 245:206–224. doi: 10.1002/1097-4687(200009)245:3<206::AID-JMOR3>3.0.CO;2-U

Cherian a, Krucoff MO, Miller LE (2011) Motor cortical prediction of EMG: evidence that a kinetic brain-machine interface may be robust across altered movement dynamics. J Neurophysiol 106:564–75. doi: 10.1152/jn.00553.2010

Cherian A, Fernandes HL, Miller LE (2013) Primary motor cortical discharge during force field adaptation reflects muscle-like dynamics. J Neurophysiol 110:768–83. doi: 10.1152/jn.00109.2012

Chestek C a, Batista AP, Santhanam G, et al (2007) Single-neuron stability during repeated reaching in macaque premotor cortex. J Neurosci 27:10742–50. doi: 10.1523/JNEUROSCI.0959-07.2007

Churchland MM, Cunningham JP, Kaufman MT, et al (2012) Neural population dynamics during reaching. Nature. doi: 10.1038/nature11129

Churchland MM, Cunningham JP, Kaufman MT, et al (2010a) Cortical preparatory activity: representation of movement or first cog in a dynamical machine? Neuron 68:387–400. doi: 10.1016/j.neuron.2010.09.015

Churchland MM, Santhanam G, Shenoy K V (2006a) Preparatory activity in premotor and motor cortex reflects the speed of the upcoming reach. J Neurophysiol 96:3130–3146. doi: 10.1152/jn.00857.2006

Churchland MM, Shenoy K V (2007) Temporal complexity and heterogeneity of single-neuron activity in premotor and motor cortex. J Neurophysiol 97:4235–57. doi: 10.1152/jn.00095.2007

Churchland MM, Yu BM, Cunningham JP, et al (2010b) Stimulus onset quenches neural variability: a widespread cortical phenomenon. Nat Neurosci 13:369–378. doi: 10.1038/nn.2501

Churchland MM, Yu BM, Ryu SI, et al (2006b) Neural Variability in Premotor Cortex Provides

a Signature of Motor Preparation. J Neurosci 26:3697–3712. doi: 10.1523/JNEUROSCI.3762-05.2006

Cisek P (2006) Preparing for speed. Focus on "Preparatory activity in premotor and motor cortex reflects the speed of the upcoming reach". J Neurophysiol 96:2842–3. doi: 10.1152/jn.00857.2006

Cisek P, Crammond DJ, Kalaska JF (2003) Neural activity in primary motor and dorsal premotor cortex in reaching tasks with the contralateral versus ipsilateral arm. J Neurophysiol 89:922–42. doi: 10.1152/jn.00607.2002

Cisek P, Kalaska JF (2005) Neural correlates of reaching decisions in dorsal premotor cortex: specification of multiple direction choices and final selection of action. Neuron 45:801–14. doi: 10.1016/j.neuron.2005.01.027

Classen J, Liepert J, Wise SP, et al (1998) Rapid plasticity of human cortical movement representation induced by practice. J Neurophysiol 79:1117–23.

Cohen MM (1967) Continuous versus terminal visual feedback in prism aftereffects. Percept Mot Skills 24:1295–1302. doi: 10.2466/pms.1967.24.3c.1295

Collinger JL, Wodlinger B, Downey JE, et al (2013) High-performance neuroprosthetic control by an individual with tetraplegia. Lancet 381:557–64. doi: 10.1016/S0140-6736(12)61816-9

Costa RM, Cohen D, Nicolelis MAL (2004) Differential corticostriatal plasticity during fast and slow motor skill learning in mice. Curr Biol 14:1124–34. doi: 10.1016/j.cub.2004.06.053

Cowley BR, Smith MA, Kohn A, Yu BM (2016) Stimulus-Driven Population Activity Patterns in Macaque Primary Visual Cortex. 1–31. doi: 10.1371/journal.pcbi.1005185

Crammond DJ, Kalaska JF (1996) Differential relation of discharge in primary motor cortex and premotor cortex to movements versus actively maintained postures during a reaching task. Exp brain Res 108:45–61.

Cunningham JP, Yu BM (2014) Dimensionality reduction for large-scale neural recordings. Nat Neurosci. doi: 10.1038/nn.3776

d'Avella A, Saltiel P, Bizzi E (2003) Combinations of muscle synergies in the construction of a natural motor behavior. Nat Neurosci 6:300–308. doi: 10.1038/nn1010

Debas K, Carrier J, Orban P, et al (2010) Brain plasticity related to the consolidation of motor sequence learning and motor adaptation. Proc Natl Acad Sci U S A 107:17839–17844. doi: 10.1073/pnas.1013176107

Dekleva BM, Ramkumar P, Wanda PA, et al (2016) Uncertainty leads to persistent effects on reach representations in dorsal premotor cortex. Elife 5:1–24. doi: 10.7554/eLife.14316

Delattre V, Keller D, Perich M, et al (2015) Network-timing-dependent plasticity. Front Cell Neurosci 9:1–11. doi: 10.3389/fncel.2015.00220

di Pellegrino G, Wise SP (1993) Effects of attention on visuomotor activity in the premotor and prefrontal cortex of a primate. Somatosens Mot Res 10:245–262. doi: 10.3109/08990229309028835

Diedrichsen J, Hashambhoy Y, Rane T, Shadmehr R (2005) Neural correlates of reach errors. J Neurosci 25:9919–9931. doi: 10.1523/JNEUROSCI.1874-05.2005

Dum RP, Strick PL (2003) An unfolded map of the cerebellar dentate nucleus and its projections to the cerebral cortex. J Neurophysiol 89:634–639. doi: 10.1152/jn.00626.2002

Dum RP, Strick PL (1991) The origin of corticospinal projections from the premotor areas in the frontal lobe. J Neurosci 11:667–89.

Dum RP, Strick PL (2002) Motor areas in the frontal lobe of the primate. Physiol Behav 77:677–82.

Dum RP, Strick PL (2005) Frontal lobe inputs to the digit representations of the motor areas on the lateral surface of the hemisphere. J Neurosci 25:1375–86. doi: 10.1523/JNEUROSCI.3902-04.2005

Durstewitz D, Vittoz NM, Floresco SB, Seamans JK (2010) Abrupt transitions between prefrontal neural ensemble states accompany behavioral transitions during rule learning. Neuron 66:438–448. doi: 10.1016/j.neuron.2010.03.029

Elbert T, Pantev C, Wienbruch C, et al (1995) Increased cortical representation of the fingers of the left hand in string players. Science (80-) 270:305–7.

Elsayed GF, Lara AH, Kaufman MT, et al (2016) Reorganization between preparatory and movement population responses in motor cortex. Nat Commun 13239. doi: 10.1038/ncomms13239

Ethier C, Oby ER, Bauman MJ, Miller LE (2012) Restoration of grasp following paralysis through brain-controlled stimulation of muscles. Nature 485:368–71. doi: 10.1038/nature10987

Evarts E V (1968) Relation of pyramidal tract activity to force exerted during voluntary movement. J Neurophysiol 31:14–27.

Ferrier D (1873) Experimental Researches in Cerebral Physiology and Pathology. J Anat Physiol 8:152–155. doi: 10.1136/bmj.1.643.457

Fetz E, Baker M (1969) Response properties of precentral neurons in awake monkeys.

Fetz EE (1992) Are movement parameters recognizable coded in the activity of single neurons? Behav. Brain Sci. 15:679–690.

Fetz EE, Cheney PD, German DC (1976) Corticomotoneuronal connections of precentral cells detected by post-spike averages of EMG activity in behaving monkeys. Brain Res 114:505–510. doi: 10.1016/0006-8993(76)90973-2

Fetz EE, Cheney PD, Mewes K, Palmer S (1989) Control of forelimb muscle activity by populations of corticomotoneuronal and rubromotoneuronal cells. Prog Brain Res 80:437-49–30.

Fetz EE, Cheney PD, Palmer SS (1986) Activity of forelimb motor units and corticomotoneuronal cells during ramp-and-hold torque responses: comparisons with oculomotor cells. Prog Brain Res 64:133–141. doi: 10.1016/S0079-6123(08)63408-1

Flash T, Hogan N (1985) The coordination of arm movements: an experimentally confirmed mathematical model. J Neurosci 5:1688–703.

Flint RD, Scheid MR, Wright ZA, et al (2016) Long-Term Stability of Motor Cortical Activity: Implications for Brain Machine Interfaces and Optimal Feedback Control. J Neurosci 36:3623–3632. doi: 10.1523/JNEUROSCI.2339-15.2016

Forsberg LE, Bonde LH, Harvey MA, Roland PE (2016) The Second Spiking Threshold: Dynamics of Laminar Network Spiking in the Visual Cortex. Front Syst Neurosci 10:1–21. doi: 10.3389/fnsys.2016.00065

Fritsch G, Hitzig E (1870) Ueber dir elektrische Erregbarkeit des Grosshirns. Arch Anat Physiol Lpz 37:330–332.

Fu M, Yu X, Lu J, Zuo Y (2012) Repetitive motor learning induces coordinated formation of clustered dendritic spines in vivo. Nature 1:92–95. doi: 10.1038/nature10844

Fujii N, Mushiake H, Tanji J (2000) Rostrocaudal distinction of the dorsal premotor area based on oculomotor involvement. J Neurophysiol 83:1764–9.

Galea JM, Vazquez A, Pasricha N, et al (2011) Dissociating the roles of the cerebellum and motor cortex during adaptive learning: the motor cortex retains what the cerebellum learns. Cereb Cortex 21:1761–70. doi: 10.1093/cercor/bhq246

Gallego JÁ, Perich MG, Miller LE, Solla SA (2017) Neural manifolds for the control of movement. Neuron 94:978–984. doi: 10.1016/j.neuron.2017.05.025

Gandolfo F, Li C, Benda BJ, et al (2000) Cortical correlates of learning in monkeys adapting to a new dynamical environment. Proc Natl Acad Sci 97:2259–63. doi: 10.1073/pnas.040567097

Gandolfo F, Mussa-Ivaldi FA, Bizzi E (1996) Motor learning by field approximation. Proc Natl Acad Sci U S A 93:3843–6.

Ganmor E, Segev R, Schneidman E (2015) A thesaurus for a neural population code. Elife 4:1–19. doi: 10.7554/eLife.06134

Gao P, Ganguli S (2015) On simplicity and complexity in the brave new world of large-scale neuroscience. Curr Opin Neurobiol 32:148–155. doi: 10.1016/j.conb.2015.04.003

Georgopoulos a P, Kalaska JF, Caminiti R, Massey JT (1982) On the relations between the direction of two-dimensional arm movements and cell discharge in primate motor cortex. J Neurosci 2:1527–37.

Gerhard F, Kispersky T, Gutierrez GJ, et al (2013) Successful Reconstruction of a Physiological Circuit with Known Connectivity from Spiking Activity Alone. PLoS Comput Biol 9:32–34. doi: 10.1371/journal.pcbi.1003138

Gilbert PF, Thach WT (1977) Purkinje cell activity during motor learning. Brain Res 128:309–28.

Glaser JI, Perich MG, Ramkumar P, et al (2017) Population coding of conditional probability distributions in dorsal premotor cortex. bioRxiv 1–27. doi: https://doi.org/10.1101/137026

Goedert KM, Willingham DB (2002) Patterns of interference in sequence learning and prism adaptation inconsistent with the consolidation hypothesis. Learn Mem 9:279–292. doi: 10.1101/lm.50102.previously

Grafton ST, Schmitt P, Van Horn J, Diedrichsen J (2008) Neural substrates of visuomotor learning based on improved feedback control and prediction. Neuroimage 39:1383–1395. doi: 10.1016/j.neuroimage.2007.09.062

Gupta R, Ashe J (2009) Offline decoding of end-point forces using neural ensembles: Application to a brain machine interface. IEEE Trans Neural Syst Rehabil Eng 17:254–262. doi: 10.1109/TNSRE.2009.2023290

Harvey CD, Coen P, Tank DW (2012) Choice-specific sequences in parietal cortex during a virtual-navigation decision task. Nature 484:62–8. doi: 10.1038/nature10918

Hatsopoulos NG, Ojakangas CL, Paninski L, Donoghue JP (1998) Information about movement direction obtained from synchronous activity of motor cortical neurons. Proc Natl Acad Sci 95:15706–11.

Held R, Freedman SJ (1963) Plasticity in Human Sensorimotor Control. Science (80-) 142:455–462. doi: 10.1126/science.142.3591.455

Held R, Schlank M (1959) Adaptation to disarranged eye-hand coordination in the distance-dimension. Am J Psychol 72:603–605.

Hennequin G, Vogels TP, Gerstner W (2014) Optimal control of transient dynamics in balanced networks supports generation of complex movements. Neuron 82:1394–1406. doi: 10.1016/j.neuron.2014.04.045

Herzfeld DJ, Pastor D, Haith AM, et al (2014) Contributions of the cerebellum and the motor cortex to acquisition and retention of motor memories. Neuroimage. doi: 10.1016/j.neuroimage.2014.04.076

Hinton GE, Salakhutdinov RR (2006) Reducing the dimensionality of data with neural networks. Science 313:504–7. doi: 10.1126/science.1127647

Hochberg LR, Serruya MD, Friehs GM, et al (2006) Neuronal ensemble control of prosthetic devices by a human with tetraplegia. Nature 442:164–171. doi: 10.1038/nature04970

Hoshi E, Tanji J (2002) Contrasting neuronal activity in the dorsal and ventral premotor areas during preparation to reach. J Neurophysiol 87:1123–1128.

Hoshi E, Tanji J (2006) Differential involvement of neurons in the dorsal and ventral premotor cortex during processing of visual signals for action planning. J Neurophysiol 95:3596–616. doi: 10.1152/jn.01126.2005

Huang VS, Haith A, Mazzoni P, Krakauer JW (2011) Rethinking Motor Learning and Savings in Adaptation Paradigms: Model-Free Memory for Successful Actions Combines with Internal Models. Neuron 70:787–801. doi: 10.1016/j.neuron.2011.04.012

Humphrey DR, Schmidt EM, Thompson WD (1970) Predicting measures of motor performance from multiple cortical spike trains. Science 170:758–762. doi: 10.1126/science.170.3959.758

Hwang EJ, Smith M a, Shadmehr R (2006) Dissociable effects of the implicit and explicit memory systems on learning control of reaching. Exp Brain Res 173:425–37. doi: 10.1007/s00221-006-0391-0

Imamizu H, Miyauchi S, Tamada T, et al (2000) Human cerebellar activity reflecting an acquired internal model of a new tool. Nature 403:192–5. doi: 10.1038/35003194

Ingram JN, Flanagan JR, Wolpert DM (2013) Context-dependent decay of motor memories during skill acquisition. Curr Biol 23:1107–12. doi: 10.1016/j.cub.2013.04.079

Izawa J, Criscimagna-Hemminger SE, Shadmehr R (2012) Cerebellar contributions to reach adaptation and learning sensory consequences of action. J Neurosci 32:4230–9. doi: 10.1523/JNEUROSCI.6353-11.2012

Jasper H, Ricci GF, Doane B (1958) Ciba Foundation Symposium - Neurological Basis of Behaviour. In: Wolstenholme G, O'Connor C (eds). John Wiley & Sons, Ltd., Chichester, UK,

Jeanne JM, Sharpee TO, Gentner TQ (2013) Associative learning enhances population coding by inverting interneuronal correlation patterns. Neuron 78:352–363. doi: 10.1016/j.neuron.2013.02.023

Johnson PB, Ferraina S, Bianchi L, et al (1996) Cortical networks for visual reaching: Physiological and anatomical organization of frontal and parietal lobe arm regions; Representing spatial information for limb movement: Role of area 5 in the monkey. Cereb Cortex 6; 5:102; 391-119; 409.

Kagerer F a, Contreras-Vidal JL, Stelmach GE (1997) Adaptation to gradual as compared with sudden visuo-motor distortions. Exp Brain Res 115:557–61.

Kakei S, Hoffman DS, Strick PL (2001) Direction of action is represented in the ventral premotor cortex. Nat Neurosci 4:1020–5. doi: 10.1038/nn726

Kalaska J, Cohen D, Hyde M, Prud'homme M (1989) A comparison of movement direction-related versus load direction-related activity in primate motor cortex, using a two-dimensional reaching task. J Neurosci 2080–2102.

Kalaska JF, Crammond DJ (1992) Cerebral cortical mechanisms of reaching movements. Science 255:1517–23.

Kalaska JF, Scott SH, Cisek P, Sergio LE (1997) Cortical control of reaching movements. Curr Opin Neurobiol 7:849–859. doi: 10.1016/S0959-4388(97)80146-8

Karni a, Meyer G, Rey-Hipolito C, et al (1998) The acquisition of skilled motor performance: fast and slow experience-driven changes in primary motor cortex. Proc Natl Acad Sci 95:861–8.

Kaufman MT, Churchland MM, Ryu SI, Shenoy K V (2014) Cortical activity in the null space: permitting preparation without movement. Nat Neurosci 17:440–8. doi: 10.1038/nn.3643

Kaufman MT, Churchland MM, Santhanam G, et al (2010) Roles of monkey premotor neuron classes in movement preparation and execution. J Neurophysiol 104:799–810. doi: 10.1152/jn.00231.2009

Kaufman MT, Churchland MM, Shenoy K V (2013) The roles of monkey M1 neuron classes in movement preparation and execution. J Neurophysiol 110:817–825. doi: 10.1152/jn.00892.2011

Kaufman MT, Seely JS, Sussillo D, et al (2016) The Largest Response Component in the Motor Cortex Reflects Movement Timing but Not Movement Type. eNeuro. doi: 10.1523/ENEURO.0085-16.2016

Kawato M (1999) Internal models for motor control and trajectory planning. Curr Opin Neurobiol 9:718–27.

Kawato M, Gomi H (1992) A computational model of four regions of the cerebellum based on feedback-error learning. Biol Cybern 68:95–103.

Kleim J a, Barbay S, Nudo RJ (1998) Functional reorganization of the rat motor cortex following motor skill learning. J Neurophysiol 80:3321–5.

Kleim J a, Hogg TM, VandenBerg PM, et al (2004) Cortical synaptogenesis and motor map reorganization occur during late, but not early, phase of motor skill learning. J Neurosci 24:628–33. doi: 10.1523/JNEUROSCI.3440-03.2004

Kleim JA, Barbay S, Cooper NR, et al (2002) Motor Learning-Dependent Synaptogenesis Is Localized to Functionally Reorganized Motor Cortex. Neurobiol Learn Mem 77:63–77. doi: 10.1006/nlme.2000.4004

Kobak D, Brendel W, Constantinidis C, et al (2016) Demixed principal component analysis of neural population data. Elife 5:1–36. doi: 10.7554/eLife.10989

Kojima Y, Iwamoto Y, Yoshida K (2004) Memory of learning facilitates saccadic adaptation in the monkey. J Neurosci 24:7531–7539. doi: 10.1523/JNEUROSCI.1741-04.2004

Krakauer JW (2003) Differential Cortical and Subcortical Activations in Learning Rotations and Gains for Reaching: A PET Study. J Neurophysiol 91:924–933. doi: 10.1152/jn.00675.2003

Krakauer JW, Ghez C, Ghilardi MF (2005) Adaptation to visuomotor transformations: consolidation, interference, and forgetting. J Neurosci 25:473–478. doi: 10.1523/JNEUROSCI.4218-04.2005

Krakauer JW, Ghilardi MF, Ghez C (1999) Independent learning of internal models for kinematic and dynamic control of reaching. Nat Neurosci 2:1026–31. doi: 10.1038/14826

Krakauer JW, Mazzoni P, Ghazizadeh A, et al (2006) Generalization of motor learning depends on the history of prior action. PLoS Biol 4:1798–1808. doi: 10.1371/journal.pbio.0040316

Krakauer JW, Shadmehr R (2006) Consolidation of motor memory. Trends Neurosci 29:58–64. doi: 10.1016/j.tins.2005.10.003

Kujirai T, Caramia MD, Rothwell JC, et al (1993) Corticocortical inhibition in human motor cortex. J Physiol 471:501–19. doi: VL - 471

Kurata K (1991) Corticocortical inputs to the dorsal and ventral aspects of the premotor cortex of macaque monkeys. Neurosci Res 12:263–280. doi: 10.1016/0168-0102(91)90116-G

Kurata K, Hoffman DS (1994) Differential effects of muscimol microinjection into dorsal and

ventral aspects of the premotor cortex of monkeys. J Neurophysiol 71:1151–1164.

Kurata K, Wise SP (1988) Premotor cortex of rhesus monkeys: set-related activity during two conditional motor tasks. Exp Brain Res 69:327–343. doi: 10.1007/BF00247578

Lackner JR, Dizio P (1994) Rapid adaptation to Coriolis force perturbations of arm trajectory. J Neurophysiol 72:299–313.

Lalazar H, Vaadia E (2008) Neural basis of sensorimotor learning: modifying internal models. Curr Opin Neurobiol 18:573–581. doi: 10.1016/j.conb.2008.11.003

Law J, Jolliffe IT (1987) Principal Component Analysis. Stat 36:432. doi: 10.2307/2348864

Li CS, Padoa-Schioppa C, Bizzi E (2001) Neuronal correlates of motor performance and motor learning in the primary motor cortex of monkeys adapting to an external force field. Neuron 30:593–607.

Lillicrap TP, Scott SH (2013) Preference distributions of primary motor cortex neurons reflect control solutions optimized for limb biomechanics. Neuron 77:168–79. doi: 10.1016/j.neuron.2012.10.041

Loewenstein Y, Mahon S, Chadderton P, et al (2005) Bistability of cerebellar Purkinje cells modulated by sensory stimulation. Nat Neurosci 8:202–11. doi: 10.1038/nn1393

Luczak A, Barthó P, Harris KD (2009) Spontaneous Events Outline the Realm of Possible Sensory Responses in Neocortical Populations. Neuron 62:413–425. doi: 10.1016/j.neuron.2009.03.014

Luczak A, McNaughton BL, Harris KD (2015) Packet-based communication in the cortex. Nat Rev Neurosci 16:745–755. doi: 10.1038/nrn4026

Machens CK, Romo R, Brody CD (2010) Functional, but not anatomical, separation of "what" and "when" in prefrontal cortex. J Neurosci 30:350–60. doi: 10.1523/JNEUROSCI.3276-09.2010

Macke JH, Buesing L, Cunningham JP, et al (2011) Empirical models of spiking in neuronal populations. Adv Neural Inf Process Syst 24:1–9. doi: 10.1.1.230.7630

Major G, Tank D (2004) Persistent neural activity: Prevalence and mechanisms. Curr Opin Neurobiol 14:675–684. doi: 10.1016/j.conb.2004.10.017

Mandelblat-Cerf Y, Novick I, Paz R, et al (2011) The neuronal basis of long-term sensorimotor learning. J Neurosci 31:300–13. doi: 10.1523/JNEUROSCI.4055-10.2011

Mante V, Sussillo D, Shenoy K V, Newsome WT (2013) Context-dependent computation by recurrent dynamics in prefrontal cortex. Nature 503:78–84. doi: 10.1038/nature12742

Mariño J, Schummers J, Lyon DC, et al (2005) Invariant computations in local cortical networks with balanced excitation and inhibition. Nat Neurosci 8:194–201. doi: 10.1038/nn1391

Markowitz D a, Curtis CE, Pesaran B (2015) Multiple component networks support working memory in prefrontal cortex. Proc Natl Acad Sci 112:11084–11089. doi: 10.1073/pnas.1504172112

Markram H, Tsodyks M (1996) Redistribution of synaptic efficacy between neocortical

pyramidal neurons. Nature 382:807–810.

Martin T a, Keating JG, Goodkin HP, et al (1996a) Throwing while looking through prisms II. Specificity and storage of multiple gaze-throw calibrations. Brain 119:1199–1211. doi: 10.1093/brain/119.4.1183

Martin T a, Keating JG, Goodkin HP, et al (1996b) Throwing while looking through prisms. I. Focal olivocerebellar lesions impair adaptation. Brain 119 (Pt 4:1183–98.

Mathis MW, Mathis A, Uchida N (2017) Somatosensory Cortex Plays an Essential Role in Forelimb Motor Adaptation in Mice. Neuron 93:1493–1503.e6. doi: 10.1016/j.neuron.2017.02.049

Mattar A a G, Gribble PL (2005) Motor learning by observing. Neuron 46:153–60. doi: 10.1016/j.neuron.2005.02.009

Mazzoni P, Krakauer JW (2006) An implicit plan overrides an explicit strategy during visuomotor adaptation. J Neurosci 26:3642–5. doi: 10.1523/JNEUROSCI.5317-05.2006

Michaels JA, Dann B, Scherberger H (2016) Neural Population Dynamics during Reaching Are Better Explained by a Dynamical System than Representational Tuning. PLOS Comput Biol 12:e1005175. doi: 10.1371/journal.pcbi.1005175

Moran DW, Schwartz AB (1999) Motor cortical representation of speed and direction during reaching. J Neurophysiol 82:2676–92.

Morasso P (1981) Spatial control of arm movements. Exp Brain Res 42:223–7. doi: 10.1007/BF00236911

Morrow M, Miller L (2003) Prediction of muscle activity by populations of sequentially recorded primary motor cortex neurons. J Neurophysiol 89:2279–2288. doi: 10.1152/jn.00632.2002.Prediction

Muellbacher W, Ziemann U, Wissel J, et al (2002) Early consolidation in human primary motor cortex. Nature 415:640–4. doi: 10.1038/nature712

Murray RM, Li Z, Sastry SS (1994) A Mathematical Introduction to Robotic Manipulation.

Mushiake H, Inase M, Tanji J (1991) Neuronal activity in the primate premotor, supplementary, and precentral motor cortex during visually guided and internally determined sequential movements. J Neurophysiol 66:705–718.

Mussa-Ivaldi F a, Bizzi E (2000) Motor learning through the combination of primitives. Philos Trans R Soc Lond B Biol Sci 355:1755–69. doi: 10.1098/rstb.2000.0733

Mussa-Ivaldi FA, Solla SA (2004) Neural primitives for motion control. IEEE J Ocean Eng. doi: 10.1109/JOE.2004.833102

Nelder J, Baker R (1972) Generalized linear models. Encycl. Stat. Sci.

Nevian T, Sakmann B (2006) Spine Ca2+ signaling in spike-timing-dependent plasticity. J Neurosci 26:11001–13. doi: 10.1523/JNEUROSCI.1749-06.2006

Newsome WT, Britten KH, Movshon J a (1989) Neuronal correlates of a perceptual decision. Nature 341:52–54. doi: 10.1038/341052a0

Nudo RJ, Milliken GW, Jenkins WM, Merzenich MM (1996) Use-dependent alterations of movement representations in primary motor cortex of adult squirrel monkeys. J Neurosci 16:785–807.

Oby E, Degenhart A, Tyler-kabara E, et al (2015) Network constraints dictate the timescale of learning new brain-computer interfaces. In: Annual Meeting of the Society for Neuroscience. Chicago, IL,

Okun M, Steinmetz N a., Cossell L, et al (2015) Diverse coupling of neurons to populations in sensory cortex. Nature. doi: 10.1038/nature14273

Orban de Xivry J-J, Criscimagna-Hemminger SE, Shadmehr R (2011) Contributions of the motor cortex to adaptive control of reaching depend on the perturbation schedule. Cereb Cortex 21:1475–84. doi: 10.1093/cercor/bhq192

Overduin S a., D'Avella A, Roh J, et al (2015) Representation of Muscle Synergies in the Primate Brain. J Neurosci 35:12615–24. doi: 10.1523/JNEUROSCI.4302-14.2015

Overduin SA, Richardson AG, Bizzi E (2009) Cortical processing during dynamic motor adaptation. Adv Exp Med Biol 629:423–38. doi: 10.1007/978-0-387-77064-2_22

Padoa-Schioppa C, Li C-SR, Bizzi E (2004) Neuronal activity in the supplementary motor area of monkeys adapting to a new dynamic environment. J Neurophysiol 91:449–73. doi: 10.1152/jn.00876.2002

Padoa-Schioppa C, Li CSR, Bizzi E (2002) Neuronal correlates of kinematics-to-dynamics transformation in the supplementary motor area. Neuron 36:751–65.

Paninski L, Fellows MR, Donoghue JP, et al (2002) Instant neural control of a movement signal. Nature 416:141–2. doi: 10.1038/416141a

Paz R, Boraud T, Natan C, et al (2003) Preparatory activity in motor cortex reflects learning of local visuomotor skills. Nat Neurosci 6:882–90. doi: 10.1038/nn1097

Penfield W, Boldrey E (1937) Somatic motor and sensory representation in the cerebral cortex of man as studied by electrical stimulation.

Perich MG, Gallego JA, Miller LE (2017) A Neural Population Mechanism For Rapid Learning. bioRxiv 1–24. doi: https://doi.org/10.1101/138743

Perich MG, Miller LE (2017) Altered tuning in primary motor cortex does not account for behavioral adaptation during force field learning. Exp Brain Res. doi: 10.1007/s00221-017-4997-1

Peters AJ, Chen SX, Komiyama T (2014) Emergence of reproducible spatiotemporal activity during motor learning. Nature. doi: 10.1038/nature13235

Pillow JW, Shlens J, Paninski L, et al (2008) Spatio-temporal correlations and visual signalling in a complete neuronal population. Nature 454:995–9. doi: 10.1038/nature07140

Pohlmeyer E a, Solla S a, Perreault EJ, Miller LE (2007) Prediction of upper limb muscle activity from motor cortical discharge during reaching. J Neural Eng 4:369–79. doi: 10.1088/1741-2560/4/4/003

Raposo D, Kaufman MT, Churchland AK (2014) A category-free neural population supports evolving demands during decision-making. Nat Neurosci 17:1784–1792. doi: 10.1038/nn.3865

Rathelot J-A, Strick PL (2009) Subdivisions of primary motor cortex based on cortico-motoneuronal cells. Proc Natl Acad Sci U S A 106:918–23. doi: 10.1073/pnas.0808362106

Rebesco JM, Stevenson IH, Körding KP, et al (2010) Rewiring neural interactions by micro-stimulation. Front Syst Neurosci 4:1–15. doi: 10.3389/fnsys.2010.00039

Richardson AG, Borghi T, Bizzi E (2012) Activity of the same motor cortex neurons during repeated experience with perturbed movement dynamics. J Neurophysiol 107:3144–54. doi: 10.1152/jn.00477.2011

Richardson AG, Lassi-Tucci G, Padoa-Schioppa C, Bizzi E (2008) Neuronal activity in the cingulate motor areas during adaptation to a new dynamic environment. J Neurophysiol 99:1253–66. doi: 10.1152/jn.01096.2007

Richardson AG, Overduin S a, Valero-Cabré A, et al (2006) Disruption of primary motor cortex before learning impairs memory of movement dynamics. J Neurosci 26:12466–70. doi: 10.1523/JNEUROSCI.1139-06.2006

Riehle a, Requin J (1989) Monkey primary motor and premotor cortex: single-cell activity related to prior information about direction and extent of an intended movement. J Neurophysiol 61:534–549.

Riek S, Hinder MR, Carson RG (2012) Primary motor cortex involvement in initial learning during visuomotor adaptation. Neuropsychologia 50:2515–23. doi: 10.1016/j.neuropsychologia.2012.06.024

Rioult-Pedotti MS, Friedman D, Hess G, Donoghue JP (1998) Strengthening of horizontal cortical connections following skill learning. Nat Neurosci 1:230–4. doi: 10.1038/678

Rizzolatti G, Fogassi L, Gallese V (2002) Motor and cognitive functions of the ventral premotor cortex. Curr Opin Neurobiol 12:149–54.

Rokni U, Richardson AG, Bizzi E, Seung HS (2007) Motor learning with unstable neural representations. Neuron 54:653–66. doi: 10.1016/j.neuron.2007.04.030

Roland PE, Skinhøj E, Lassen N a, Larsen B (1980) Different Cortical Areas in Man in Organization of Voluntary Movements in Extrapersonal Space. J Neurophysiol 43:137–150.

Roweis ST, Saul LK (2000) Nonlinear dimensionality reduction by locally linear embedding. Science 290:2323–2326. doi: 10.1126/science.290.5500.2323

Sadtler PT, Quick KM, Golub MD, et al (2014) Neural constraints on learning. Nature 512:423–426. doi: 10.1038/nature13665

Santhanam G, Ryu SI, Yu BM, et al (2006) A high-performance brain-computer interface. Nature 442:195–8. doi: 10.1038/nature04968

Santhanam G, Yu BM, Gilja V, et al (2009) Factor-analysis methods for higher-performance neural prostheses. J Neurophysiol 102:1315–30. doi: 10.1152/jn.00097.2009

Schlerf JE, Galea JM, Bastian AJ, Celnik P a (2012) Dynamic modulation of cerebellar excitability for abrupt, but not gradual, visuomotor adaptation. J Neurosci 32:11610–7. doi: 10.1523/JNEUROSCI.1609-12.2012

Scott S (2004) Optimal Feedback Control and the Neural Basis of Volitional Motor Control. Nat Rev Neurosci 5:532–546. doi: doi: 10.1038/nrn1427

Scott SH (2008) Inconvenient truths about neural processing in primary motor cortex. J Physiol 586:1217–24. doi: 10.1113/jphysiol.2007.146068

Scott SH (2012) The computational and neural basis of voluntary motor control and planning. Trends Cogn Sci 16:541–549. doi: 10.1016/j.tics.2012.09.008

Scott SH, Cluff T, Lowrey CR, Takei T (2015) Feedback control during voluntary motor actions. Curr Opin Neurobiol 33:85–94. doi: 10.1016/j.conb.2015.03.006

Scott SH, Kalaska JF (1995) Changes in motor cortex activity during reaching movements with similar hand paths but different arm postures. J Neurophysiol 73:2563–7.

Scott SH, Kalaska JF (1997) Reaching Movements With Similar Hand Paths But Different Arm Orientations . I . Activity of Individual Cells in Motor Cortex. J Neurophysiol 77:826–852.

Sergio LE, Hamel-Pâquet C, Kalaska JF (2005) Motor cortex neural correlates of output kinematics and kinetics during isometric-force and arm-reaching tasks. J Neurophysiol 94:2353–78. doi: 10.1152/jn.00989.2004

Sergio LE, Kalaska JF (2003) Systematic changes in motor cortex cell activity with arm posture during directional isometric force generation. J Neurophysiol 89:212–28. doi: 10.1152/jn.00016.2002

Shadmehr R, Brashers-Krug T (1997) Functional stages in the formation of human long-term motor memory. J Neurosci 17:409–19.

Shadmehr R, Holcomb HH (1997) Neural correlates of motor memory consolidation. Science 277:821–5.

Shadmehr R, Mussa-Ivaldi FA (1994) Adaptive representation of dynamics during learning of a motor task. J Neurosci 14:3208–24.

Sheahan HR, Franklin DW, Wolpert DM (2016) Motor Planning, Not Execution, Separates Motor Memories. Neuron 92:773–779. doi: 10.1016/j.neuron.2016.10.017

Shen L, Alexander GE (1997) Preferential representation of instructed target location versus limb trajectory in dorsal premotor area. J Neurophysiol 77:1195–212.

Shenoy K V, Sahani M, Churchland MM (2013) Cortical control of arm movements: a dynamical systems perspective. Annu Rev Neurosci. doi: 10.1146/annurev-neuro-062111-150509

Shimansky YP, Kang T, He J (2004) A novel model of motor learning capable of developing an optimal movement control law online from scratch. Biol Cybern 90:133–145. doi: 10.1007/s00422-003-0452-4

Slotine J-JE (1985) The Robust Control of Robot Manipulators. Int J Rob Res 4:49–64. doi:

10.1177/027836498500400205

Smith M a, Ghazizadeh A, Shadmehr R (2006) Interacting adaptive processes with different timescales underlie short-term motor learning. PLoS Biol 4:e179. doi: 10.1371/journal.pbio.0040179

Smith M a, Shadmehr R (2005) Intact ability to learn internal models of arm dynamics in Huntington's disease but not cerebellar degeneration. J Neurophysiol 93:2809–21. doi: 10.1152/jn.00943.2004

Snyder LH, Batista AP, Andersen RA (1997) Coding of intention in the posterior parietal cortex. Nature 386:167–170. doi: 10.1038/386167a0

Song W, Cajigas I, Brown EN, Giszter SF (2015) Adaptation to elastic loads and BMI robot controls during rat locomotion examined with point-process GLMs. Front Syst Neurosci. doi: 10.3389/fnsys.2015.00062

Stevenson IH, Cherian A, London BM, et al (2011) Statistical assessment of the stability of neural movement representations. J Neurophysiol 106:764–74. doi: 10.1152/jn.00626.2010

Stopfer M, Bhagavan S, Smith BH, Laurent G (1997) Impaired odour discrimination on desynchronization of odour-encoding neural assemblies. Nature 390:70–4. doi: 10.1038/36335

Stopfer M, Jayaraman V, Laurent G (2003) Intensity versus identity coding in an olfactory system. Neuron 39:991–1004. doi: 10.1016/j.neuron.2003.08.011

Sussillo D, Churchland MM, Kaufman MT, Shenoy K V (2015) A neural network that finds a naturalistic solution for the production of muscle activity. Nat Neurosci 18:1025–33. doi: 10.1038/nn.4042

Tanaka H, Sejnowski TJ, Krakauer JW (2009) Adaptation to visuomotor rotation through interaction between posterior parietal and motor cortical areas. J Neurophysiol 102:2921–32. doi: 10.1152/jn.90834.2008

Taylor DM, Tillery SIH, Schwartz AB (2002) Direct cortical control of 3D neuroprosthetic devices. Science 296:1829–32. doi: 10.1126/science.1070291

Tenenbaum JB, de Silva V, Langford JC (2000) A global geometric framework for nonlinear dimensionality reduction. Science 290:2319–23. doi: 10.1126/science.290.5500.2319

Thach WT (1978) Correlation of neural discharge with pattern and force of muscular activity, joint position, and direction of intended next movement in motor cortex and cerebellum. J Neurophysiol 41:654–676.

Thoroughman K a, Shadmehr R (2000) Learning of action through adaptive combination of motor primitives. Nature 407:742–7. doi: 10.1038/35037588

Thoroughman K a, Shadmehr R (1999) Electromyographic correlates of learning an internal model of reaching movements. J Neurosci 19:8573–88.

Thoroughman K a, Taylor J a (2005) Rapid reshaping of human motor generalization. J Neurosci 25:8948–53. doi: 10.1523/JNEUROSCI.1771-05.2005

Thura D, Cisek P (2014) Deliberation and commitment in the premotor and primary motor cortex during dynamic decision making. Neuron 81:1401–16. doi: 10.1016/j.neuron.2014.01.031

Todorov E (2000) Direct cortical control of muscle activation in voluntary arm movements: a model. Nat Neurosci 3:391–398. doi: 10.1038/73964

Todorov E (2004) Optimality principles in sensorimotor control. Nat Neurosci 7:907–15. doi: 10.1038/nn1309

Todorov E, Jordan MI (2002) Supp Optimal feedback control as a theory of motor coordination. Nat Neurosci 5:1226–1235. doi: 10.1038/nn963

Tolias AS, Ecker AS, Siapas AG, et al (2007) Recording chronically from the same neurons in awake, behaving primates. J Neurophysiol 98:3780–3790. doi: 10.1152/jn.00260.2007

Tong C, Wolpert DM, Flanagan JR (2002) Kinematics and dynamics are not represented independently in motor working memory: evidence from an interference study. J Neurosci 22:1108–13.

Tresch MC, Jarc A (2009) The case for and against muscle synergies. Curr Opin Neurobiol 19:601–7. doi: 10.1016/j.conb.2009.09.002

Truccolo W, Hochberg LR, Donoghue JP (2010) Collective dynamics in human and monkey sensorimotor cortex: predicting single neuron spikes. Nat Neurosci 13:105–11. doi: 10.1038/nn.2455

Tsodyks M, Kenet T, Grinvald A, Arieli A (1999) Linking spontaneous activity of single cortical neurons and the underlying functional architecture. Science 286:1943–1946. doi: 10.1126/science.286.5446.1943

Vaadia E, Kurata K, Wise SP (1988) Neuronal activity preceding directional and nondirectional cues in the premotor cortex of rhesus monkeys. Somatosens Mot Res 6:207–230. doi: 10.3109/08990228809144674

Wagner MJ, Smith MA (2008) Shared internal models for feedforward and feedback control. J Neurosci 28:10663–10673. doi: 28/42/10663 [pii]\r10.1523/JNEUROSCI.5479-07.2008

Weinrich M, Wise SP (1982) The premotor cortex of the monkey. J Neurosci 2:1329–1345.

Werner S, Schorn CF, Bock O, et al (2014) Neural correlates of adaptation to gradual and to sudden visuomotor distortions in humans. Exp Brain Res. doi: 10.1007/s00221-014-3824-1

Werner W, Bauswein E, Fromm C (1991) Static firing rates of premotor and primary motor cortical neurons associated with torque and joint position. Exp Brain Res 86:293–302.

Wigmore V, Tong C, Flanagan JR (2002) Visuomotor rotations of varying size and direction compete for a single internal model in motor working memory. J Exp Psychol Hum Percept Perform 28:447–57.

Wolpert DM, Diedrichsen J, Flanagan JR (2011) Principles of sensorimotor learning. Nat Rev Neurosci 12:739–51. doi: 10.1038/nrn3112

Wolpert DM, Ghahramani Z, Jordan MI (1995) An internal model for sensorimotor integration.

Science (80-) 269:1880–2.

Wolpert DM, Miall RC, Kawato M (1998) Internal models in the cerebellum. Trends Cogn Sci 2:338–47.

Wurtz RH (1969) Visual receptive fields of striate cortex neurons in awake monkeys. J Neurophysiol 32:727–42.

Xiao J, Padoa-Schioppa C, Bizzi E (2006) Neuronal correlates of movement dynamics in the dorsal and ventral premotor area in the monkey. Exp brain Res 168:106–19. doi: 10.1007/s00221-005-0074-2

y Cajal SR (1995) Histology of the nervous system of man and vertebrates. Oxford University Press, USA

Yang Y, Lisberger SG (2013) Interaction of plasticity and circuit organization during the acquisition of cerebellum-dependent motor learning. Elife 2013:1–19. doi: 10.7554/eLife.01574

Yu BM, Cunningham JP, Santhanam G, et al (2009) Gaussian-Process Factor Analysis for Low-Dimensional Single-Trial Analysis of Neural Population Activity. J Neurophysiol 102:614–635. doi: 10.1152/jn.90941.2008

Zheng N, Raman IM (2010) Synaptic inhibition, excitation, and plasticity in neurons of the cerebellar nuclei. Cerebellum 9:56–66. doi: 10.1007/s12311-009-0140-6