NORTHWESTERN UNIVERSITY

The Roles of Primary Motor Cortex and Dorsal Premotor Cortex in Planning and Executing Movements under Uncertainty

A DISSERTATION

SUBMITTED TO THE GRADUATE SCHOOL
IN PARTIAL FULFILLMENT OF THE REQUIREMENTS

For the degree

DOCTOR OF PHILOSOPHY

Field of Biomedical Engineering

By

Brian Michael Dekleva

EVANSTON, ILLINOIS

March 2018

ABSTRACT

Each movement we make represents the final output of complex processes in the nervous system. Studies of motor control often attempt to minimize further complication by using controlled environments to generate repeated movements. However, in natural situations, the motor system faces the much more complicated task of interacting with an uncertain environment. Often this uncertainty arises from noise. For example, when attempting to swat a fly, we must rely on an erratic visual cue (the fly's flight path) to estimate its future location and then develop the appropriate motor command. In this case, the goal of the sensorimotor system is to translate noisy sensory information into the single most appropriate action. At other times, uncertainty arises from the presence multiple discrete options. For example, when picking out an apple at the store, we examine the color and size of each one first and then reach to the one we deem best. In this situation, the uncertainty does not pertain to the executed movement itself—after all, the locations of the apples are static—but rather to the expected consequences for each option. The goal of the sensorimotor system in this case is to use sensory cues (and/or previous experiences) to decide between the multiple possible actions. Regardless of the source and type of uncertainty encountered during motor control, the sensorimotor system should also use the result (missing the fly or picking a rotten apple) to inform future situations. In this work, I examine how primary motor cortex (M1) and dorsal premotor cortex (PMd) plan and execute reaching movements when faced with noisy sensory information or multiple potential reach targets. I show that when faced with noisy information about a target location, PMd develops a low-fidelity reach plan, reflecting a subjective sense of uncertainty about the decision. I then show that when faced with two potential reach targets, neither PMd nor M1 contains multiple

plans for both targets. Instead, motor cortex appears to decide on one option very quickly and later switches if required. Finally, I show that both PMd and M1 contain differential responses to the consequences (success or failure) of executed movements, potentially revealing a mechanism for driving motor learning within the network. These results together provide insight into how motor cortices cope with uncertainty when planning and executing movements.

TABLE OF CONTENTS

LIST OF FIGURES

LIST OF TABLES

CHAPTER 1: INTRODUCTION

Any interaction with the environment includes some degree of uncertainty. Uncertainty can arise at multiple stages during movement preparation and execution, and needs to be considered by the sensorimotor system if movements are to be effective. In the following sections, I will outline the various sources of noise and ambiguity as well as their behavioral consequences. I will then review the literature that aims to explain these behavioral consequences through hypothesized and/or observed neural mechanisms in the sensorimotor system. Finally, I will present a brief overview of reward-based motor learning, its application to uncertain motor tasks, and the potential implications for motor cortex.

TYPES OF UNCERTAINTY

Overview

Fittingly, the meaning of the term *uncertainty* is not precise. It has been used in different motor control contexts to indicate completely different phenomena. Various sources of noise and uncertainty impact decision-making, movement planning, and movement execution, each with different implications for motor control. Given the wide range of potential meanings, it is necessary to clarify the types of uncertainty that factor into the type of motor planning that I will explore later. In this section I will provide a brief overview of the types of uncertainty relevant to the task of movement planning, as well as published observations that indicate their impact on behavior.

Environmental and sensor noise

Noise first arises in the environment itself, which is often dynamic and stochastic. Using the example from the abstract, swatting a fly is challenging largely because of the unpredictability of the fly's path. The consequence of this kind of environmental noise from a motor control perspective is the lack of an explicit endpoint goal. The brain must plan an action using only an internally-generated estimate of the fly's future location based on its current location and movement, and incorporating transmission delays in the visual and motor systems. Since the endpoint goal cannot be generated from an external reference, any errors or uncertainty introduced during the process of estimation will add to the total uncertainty contained in the movement plan itself.

The challenge of estimating an endpoint goal is compounded by the imperfect nature of sensory systems. Even vision—though a high acuity sense—does not perfectly reflect the state of the world. This is especially evident through optical illusions, where context can bias the perception of color, size, etc (Franz, Gegenfurtner, Bülthoff, & Fahle, 2000; Weiss, Simoncelli, & Adelson, 2002). More common in experimental paradigms is the use of confusing stimuli that push the limits of perception. This is evident in discrimination tasks across multiple sensory modalities, including vision (A. K. Churchland, Kiani, & Shadlen, 2008; Coallier, Michelet, & Kalaska, 2015; Newsome, Britten, & Movshon, 1989), audition (Ritter, Simson, & Vaughan, 1972), and somatosensation (Hernandez et al., 2010; Mountcastle, Steinmetz, & Romo, 1990; Romo, Brody, Hernández, & Lemus, 1999; Romo, Hernández, & Zainos, 2004). In a classic study by Newsome, et al (Newsome et al., 1989), subjects observed a large cluster of dots—

some moving to left and some to the right—and were tasked with determining which movement direction was most prevalent. In this type of situation, the line between environmental noise and sensory noise is blurred. While the dot movement is stochastic (suggesting an environmental noise source), an ideal observer would be perfectly able to complete the task (i.e.,, any difficulty in doing so must arise from sensory noise). Since motor control is a practical, rather than theoretical endeavor, there seems little reason to consider these as two separate cases. Therefore, I will use the term *sensory uncertainty* to refer to the uncertainty resulting from either pure environmental noise, limitations of the sensory system, or a combination of both.

While vision may be the most common modality used to modify uncertainty in experimental settings, it is certainly not the only sense necessary for movement and subject to noise. Proprioception, the sense of body position, is essential for effective natural control. In the absence of proprioception (equivalent to infinite proprioceptive uncertainty), movements require significant cognitive effort and a heavy reliance on vision (Abbott, 2006). Even in healthy individuals, proprioception is subject to significant drift over time (Bowditch & Southard, 1882; Wann & Ibrahim, 1992). This proprioceptive noise induces uncertainty in the state of the effector (e.g., hand position and limb orientation in the case of reaching), which makes it difficult to plan the necessary muscle activations even if the endpoint goal is perfectly described. Under these limitations in sensation (namely vision and proprioception), the task faced by the motor system amounts to planning a movement from one poorly defined location to another.

Motor and neural noise

The motor system, like the sensory system, also contains noise. When subjects are asked to make repeated reaches to a clear visual target, the hand trajectories and endpoints are variable (M. M. Churchland, Afshar, & Shenoy, 2006; Robert J Van Beers, Haggard, & Wolpert, 2004). Some of this variability might be attributable to noise in proprioception (and to a lesser extent, vision), but under ideal experimental conditions the impact of sensory noise should be small. Ultimately, all movements rely on signals transmitted from motor brain areas to muscles. The contractions of those muscles are not perfectly reliable, meaning that equivalent neural commands are not guaranteed to result in equivalent movements. Noise at the neuromuscular junction (Hamilton, Jones, & Wolpert, 2004; Jones, Hamilton, & Wolpert, 2002; Osu et al., 2004), as well as factors like fatigue (Cortes, Onate, & Morrison, 2014), diminish the reliability of the command-to-output relationship and add uncertainty at the execution level.

Even with a perfectly deterministic relationship between neural signals and muscle contractions, movement accuracy would still be limited by variability in the neural signals themselves. Neural firing is inherently stochastic, with action potentials approximating a Poisson process (Ma, Beck, Latham, & Pouget, 2006; Shadlen & Newsome, 1998; Tolhurst, Movshon, & Dean, 1983). The brain can limit the impact of this variability in individual neurons by combining signals from large neural populations. However, recent studies examining the behavior of neural populations during movement planning suggest that a great deal of variability in motor output arises from noise in planning-related neural activity (Afshar et al., 2011; M. M. Churchland, A. Afshar, et al., 2006). Thus, even if the brain had perfect knowledge of the environment, the movement endpoint goal, and the state of the limb, the stochastic nature of

neural signals and muscle activation would inject variability and uncertainty into the motor output.

In addition to purely internal motor noise (e.g.,muscle contractions), external factors can also impact uncertainty in movement execution. For example, using a tool or reaching for an object adds uncertainty at the execution level by changing the dynamics of the arm. Natural movements cannot depend purely on a feedback control strategy due to the delays in both the sensory and motor systems. Instead, the brain likely generates an internal model of the body's dynamics for feedforward control through the formation and calibration of a so called "internal inverse model" (Reza Shadmehr & Mussa-Ivaldi, 1994; Daniel M Wolpert, Ghahramani, & Jordan, 1995). Accuracy of movements therefore depend on the accuracy of the internal model. If the dynamics of the effector (e.g.,arm holding a hammer) are unknown, they must be estimated, and any such estimation necessarily induces uncertainty in the motor output. However, while uncertainty in the internal model is certainly exacerbated by external forces (grasped objects, etc.), a similar but smaller effect likely exists for all movements. We can safely assume that part of the uncertainty and noise in movement execution arises from an imperfect knowledge of the dynamical state of the body.

Noise versus ambiguity

The sources of uncertainty outlined above all fall under the general category of noise. That is, uncertainty in motor control arises from the inability of the sensorimotor system to perfectly control the movement from point A to point B. But how does the brain arrive at the decision to move to point B in the first place? Rarely do we encounter situations where there is

only one possible movement. We live in a world full of choices, and must constantly decide between multiple potential actions. This type of uncertainty is separate from the stochastic noise present in the sensory and motor systems, and instead reflects ambiguity that must be addressed through higher-order decision-making processes. In the simple example presented in the abstract (choosing an apple at the store), the brain must consider the array of movement options and then decide which one to perform. Often there are multiple factors involved in that decision, including the expected reward (this apple is bigger), energy expenditure (that apple is farther away), and previous experience (those apples were rotten last week).

The creation of a movement plan in the brain represents the output of a decision-making process. For this reason, experiments studying decision-making are often performed within motor control contexts such as target-based reaching tasks. Task designs usually combine aspects of environmental and sensory noise and provide partial evidence for two possible choices. In the classic example by Newsome, et al (Newsome et al., 1989)—described in the above section regarding environmental and sensory noise—subjects observed a stochastic dot motion stimulus that (noisily) indicated whether to reach to the right or to the left. Variants of this task—noisy stimuli with a binary choice—are common in studies of decision-making in the context of motor control . While this approach certainly induces uncertainty at the higher-order decision-making level, it does not cause any significant low-level or execution-related uncertainty. That is, while there may be significant uncertainty about which target to choose, the targets themselves are well defined. Regardless of how noisy the cue may be, once a decision is made the subsequent movement is essentially trivial. This is not the case for tasks in which the target (Izawa & Shadmehr, 2008; Tassinari, Hudson, & Landy, 2006) or feedback of the effector (Konrad P Kording & Wolpert, 2004; Wei & Kording, 2010) is obscured with noise. Both types of tasks

require compensation for environmental and/or sensory uncertainty, but at different points in the pathway that processes sensation into movement execution.

BEHAVIORAL EFFECTS OF UNCERTAINTY

Overview

Uncertainty is ubiquitous in motor control, and its effects on sensation, decision-making, and movement execution directly influence task behavior. Much of what is known or hypothesized about uncertainty processing in the brain has come not from direct neural recordings, but rather from examining these behavioral outcomes. The specific neural mechanisms of uncertainty processing are still up for debate, but there is little question that the brain does indeed take uncertainty into account when planning and controlling movement. In this section I will provide an overview of the existing literature describing the effects of different types of uncertainty on behavior and the hypothesized governing principles that may explain them.

Bayesian integration: combining multiple cues

As described earlier in this introduction, the stochastic nature of the environment and noisy, imperfect sensory systems create significant obstacles for motor control. One general method for improving accuracy in the face of noisy signals is to combine information from multiple modalities. For example, accurate estimates of both arm position and target location are necessary for motor control. However, as discussed previously, the estimates provided by both

vision and proprioception are corrupted by noise. Rather than rely on either sense alone, performance can be improved by combining estimates from both.

The optimal combination of uncertain cues was formalized centuries ago by Thomas Bayes (Bayes, 1763). Assuming Gaussian noise, the final combined estimate should weight the two individual estimates in inverse proportion to their variances. Intuitively, this means that an estimate that is reliable (small variance) should be weighted more strongly than one that is unreliable (large variance). The resulting estimate has smaller variance than either of the individual estimates.

Evidence for Bayesian integration has been found in a number of sensory tasks, suggesting that it is (or least approximates) a basic aspect of sensation. When attempting to locate the source of an object, subjects combine visual and auditory information in a Bayesian manner (Alais & Burr, 2004; Battaglia, Jacobs, & Aslin, 2003). Similarly, both vision and touch are used when estimating features (e.g.,orientation, location, etc.) of an object (Ernst & Banks, 2002). More applicable to the case of motor control, van Beers (R. J. van Beers, Baraduc, & Wolpert, 2002) showed that subjects used a Bayesian combination of vision and proprioception when estimating hand location. While these examples show evidence for the combination of two separate sensory systems (i.e., vision/audition, vision/touch, vision/proprioception), Bayesian behavior has also been observed when using multiple cues from a single modality. For example, perception of surface slant depends on both binocular information and surface texture (Hillis, Watt, Landy, & Banks, 2004; Knill & Saunders, 2003), which are both visual in nature. From these experiments and others, it appears that humans compensate for uncertainty caused by limitations in sensory acuity (or environmental noise) by combining estimates from multiple modalities.

While Bayesian integration defines the optimal linear combination of two estimates, the specifics of the desired behavior can change what is actually "optimal". For example, a study by Greenwald et al. (Greenwald & Knill, 2009) explored how subjects relied on two types of visual cues (again, binocular information and texture) when grasping an object or when placing it on a surface. They found that subjects relied more on binocular information when grasping, since depth information was more relevant in that task. Likewise, a paper by Sober and Sabes (Sober & Sabes, 2005) explored the weighting of proprioception and vision during a reaching task. Executing a reach involves planning the action in an external coordinate frame ("I want to reach to that cup") and then translating that plan into the appropriate muscle contractions. Sober and Sabes found that while external goals relied heavily on vision, the translation into movement relied on proprioception. In this case, the "optimal" plan differed even between the stages of a single reach. Another version of task-specific optimality can be seen through the influence of environmental factors. When subjects are asked to reach to a target while avoiding some nearby obstacle, they will bias their movements accordingly (Greenwald & Knill, 2009; Trommershauser, Gepshtein, Maloney, Landy, & Banks, 2005; Trommershäuser, Maloney, & Landy, 2003). Similarly, a recent paper by Acerbi et al (Acerbi, Vijayakumar, & Wolpert, 2017) examined the online correction of reaches to uncertain targets. They found that subjects' compensatory corrections to a perturbation in cursor position changed as a function of target uncertainty. When the target was more uncertain, subjects corrected less while still maintaining a high level of task success. Thus, evidence from sensorimotor behavior indicates Bayesian integration is employed by the sensorimotor system to optimize task-specific performance.

Bayesian integration: incorporating prior experience

The cue-combination studies referenced above indicate that estimation of a single quantity (e.g.,object location, orientation, etc.) can be improved by combining information from multiple senses. But often we must rely on information from only a single sensory modality, as when estimating the location of a silent object using only vision. It may appear at first that the reliability (certainty) of such an estimate is simply equivalent to that of the immediate sensory inputs. However, there is another source of information available that can improve performance—experience. In the Bayesian framework, this is referred to as the *prior*, while current sensory information is termed the *likelihood*. As an example of how incorporating a prior can improve behavior, consider the act of reaching for a doorknob in a dark room. It may be too dark see anything but the outline of the door (i.e.,, a weak likelihood estimate), but experience indicates that the knob will be located on one side about three feet from the floor (i.e.,, strong prior estimate). The fact that nobody would attempt to find a doorknob by first reaching over their head is a testament to the influence of prior experience on sensorimotor behavior.

There is experimental evidence to indicate that sensorimotor behavior is strongly influenced by prior knowledge—sometimes with detrimental effects. A 2002 study by Weiss et al (Weiss et al., 2002) investigated motion perception of simple visual stimuli and found that errors could be explained by an underlying expectation (prior) of low velocities. A study of the perception of object orientation found that perception was influenced by subjects' expectations about the location of the light source (Mamassian & Landy, 2001). Evidence of prior weighting is not only found in visual perception tasks. In the somatosensory system, two temporally and spatially offset taps on the arm can elicit the sensation of additional taps that "jump" along the length of the arm—an illusion called the "cutaneous rabbit effect". Goldreich et al (Goldreich,

2007) showed that this illusion likely arises from the expectation that something moving along the arm will do so with limited speed. In these examples, conscious sensation itself was affected by prior belief and expectation.

While some priors are likely developed over the course of a lifetime (e.g., the expectation that light sources tend to be overhead), others develop on a much faster timescale. In a 2004 paper by Kording and Wolpert (Konrad P Kording & Wolpert, 2004), subjects performed reaches to a target with only brief, noisy visual feedback about their cursor location. Additionally, on each trial the cursor was offset from the actual hand location by a random amount drawn from a Gaussian distribution with non-zero mean. As predicted by Bayesian models, the subjects' reach trajectories indicated that they weighted the visual feedback and average cursor offset according to their relative uncertainty. Importantly, this integration occurred despite the subjects reporting no conscious knowledge that the cursor was offset. This result demonstrates that priors can be quickly integrated into behavior, and do so at an unconscious level within the sensorimotor system.

The concept of Bayesian integration in the sensorimotor system can also be extended to the continuous, moment-by-moment control of movement. Due to errors at the level of movement execution (see section *motor and neural noise*), the motor output must constantly be compared with the desired action and subsequently corrected. A 2008 study by Izawa and Shadmehr (Izawa & Shadmehr, 2008) investigated the effect of uncertainty on movement correction through reaches to blurry targets. On some trials, the target jumped (imperceptibly) to a new location midway through the reach, forcing a correction of the movement trajectory. The speed and magnitude of this corrective movement depended on the blurriness (uncertainty) of the initial target; the more uncertain the initial target, the faster and larger the subject's corrective

movement. This result is predicted by a Bayesian approach to online correction in which estimates (priors) are continually updated using current sensory information. Thus, whether at the scale of a lifetime or a couple hundred milliseconds, sensorimotor behavior appears to be largely influenced by experience and expectation.

Behavior with multiple available actions

The term "uncertainty" can refer to either stochastic noise or ambiguity in available choices. The two-alternative forced choice (2AFC) framework gauges perception of a stimulus (with variable noise-based uncertainty) through a subject's choice of target (target A or target B). The process of completing such a task includes sensory processing, high level decision-making, and motor-level execution. However, in most cases, the entirety of task-relevant uncertainty is contained within the sensory and decision-making processing stages. The noisy cue induces uncertainty in the visual system, which in turn creates uncertainty in the decision about which target to choose. However, given the binary nature of the task, the uncertainty does not bleed through into the output. No matter the uncertainty in the cue, the mapping from decision to action is determined entirely by the premise of the task (e.g., rightward dot motion indicates a rightward reach), a transformation which appears to occur rapidly (Gallivan, Stewart, Baugh, Wolpert, & Flanagan, 2017). The only type of uncertainty that persists throughout the movement in these types of tasks is the uncertainty about choice outcome. This is likely important for motor learning (explored in later sections), but does not have a direct impact on low-level control of the chosen movement itself.

A variant of the 2AFC task that more directly ties ambiguity-related uncertainty to the actual movement is the go-before-you-know paradigm. In this scenario, subjects are again provided with two potential targets, but are only provided with information about which is correct once a movement has been initiated. Unlike the standard 2AFC task, where movement reflects the rule-based mapping from decision-making to motor output, a go-before-you-know task integrates the two processes.

One consistent observation in ambiguous target tasks is the presence of intermediate movements (Chapman et al., 2010; Gallivan, Logan, Wolpert, & Flanagan, 2016; Georgopoulos, Kalaska, & Massey, 1981; Haith, Huberdeau, & Krakauer, 2015; He & Kowler, 1989; Henis & Flash, 1995; Stewart, Baugh, Gallivan, & Flanagan, 2013). Given two possible targets of equal likelihood, subjects tend to direct their initial reach trajectory somewhere between the targets. This phenomenon has been interpreted as either (1) the result of averaging two simultaneous reach plans, or (2) an optimal strategy used to improve task performance. The spatial averaging viewpoint suggests that when confronted with multiple possible actions, the brain constructs simultaneous, parallel reach plans (Chapman et al., 2010; Gallivan et al., 2016; Pastor-Bernier & Cisek, 2011). When movement is initiated before the choice can be made, the two plans are executed at once, combining to form a single, intermediate movement. The alternative view proposes that intermediate movements arise from a single reach plan that is directed between the two choices in a way that offers a task-relevant advantage (Haith et al., 2015; Nashed, Diamond, Gallivan, Wolpert, & Flanagan, 2017; A. L. Wong & Haith, 2017). Supporting this view is the observation that intermediate movements are not typically observed for fast reaches (A. L. Wong & Haith, 2017). When reaches are relatively slow, intermediate movements can bring the hand closer to both targets while the appropriate corrective action is planned. However, fast reaches do

not allow enough time to plan a corrective movement following reach initiation, so intermediate movements are not beneficial. Similarly, the presence of barriers between the two targets also eliminates the presence of intermediate movements (Burk, Ingram, Franklin, Shadlen, & Wolpert, 2014; Haith et al., 2015). These results suggest that intermediate movements are intentional and functional elements of motor control, rather than errors resulting from inadequate or simultaneous reach specification. In the following section I will review the implication of each interpretation on the neural basis of decision-making and movement planning.

MOTOR CORTEX

Overview

To translate action intent into voluntary movement, the nervous system must ultimately output signals to control the activation of muscles. The transformation from high-level action goals into muscle contraction likely involves multiple regions throughout the central nervous system, including frontal brain areas and motor cortices, as well as the cerebellum, brainstem, and spinal cord. Of particular interest are the primary motor and premotor cortices, which contain direct outputs to the spinal cord. In the following sections, I will review these two cortical areas and their hypothesized involvement in movement planning and execution.

Primary motor cortex

Primary motor cortex—also referred to as Brodmann Area 4 or M1—lies on the most posterior aspect of the frontal lobe, and extends from the anterior wall of the central sulcus to the precentral gyrus. The main cytoarchitectonic feature that distinguishes the primary motor cortex

from other nearby motor areas is the presence of large pyramidal neurons in layer 5, called Betz cells. Betz cells comprise a portion of the output pyramidal neurons originating in M1 that project directly to the spinal cord through the corticospinal pathway. In addition to corticospinal projections, M1 also contains neurons that project to brainstem nuclei through the corticobulbar pathway. Though considered primarily a cortical output structure, M1 is extensively interconnected with other cortical and subcortical areas, including premotor areas in the frontal lobe, parietal cortex, thalamus, the cerebellum, and basal ganglia (Dum & Strick, 2002; Hoover & Strick, 1999; Kelly & Strick, 2003; Purves, Augustine, Fitzpatrick, & Katz, 1997).

A major feature of the primary motor cortex—and indeed what first led to its discovery—is the ease with which electrical stimulation causes movement (Fritsch & Hitzig, 1960; Penfield & Boldrey, 1937). Experiments exploiting the low stimulation thresholds found that electrically stimulating different sub regions of M1 elicited different types of movement. This led to the characterization of the topographical structure within M1. Stimulation of the most medial aspect caused movements in the lower extremities. Stimulations more lateral elicited movements in the trunk and proximal arm, followed by lower arm and hands. The most lateral stimulations caused movement in the face and mouth. This organization generally tracks the somatotopy in the sensory cortex, which is located just posterior to the motor cortex across the central sulcus.

While these early stimulation experiments characterized the general mediolateral organization in M1, later electrophysiological studies also identified functional differences along the anteroposterior dimension. Recordings from single neurons in the arm region of M1 indicated that neurons located closer to or within the central sulcus more directly related to muscle activity than did those more anterior in M1 (D. J. Crammond & Kalaska, 1996; Johnson, Ferraina, Bianchi, & Caminiti, 1996). A more recent study using retrograde labeling provided anatomical

support for these results, finding that a large percentage of sulcal M1 neurons projected directly onto motor neurons (Rathelot & Strick, 2009). In contrast, M1 output neurons originating from the gyrus projected onto interneuronal circuits within the spinal cord. The distinction between these two output types led the authors to conclude that the region of M1 that lies within the central sulcus is an evolutionarily newer structure ("new M1"). The development of direct control of lower motor neurons likely allows for more precise and dexterous control over hand and finger movements. More general and gross movement—such as proximal limb movement—is therefore hypothesized to reside within "old M1" on the precentral gyrus.

Although the primary motor cortex is clearly involved with the generation of voluntary movement, the exact mechanism of control has long been a subject of debate. In the late 1960s, electrophysiological recordings from awake, behaving monkeys showed a correlation between the discharge rate of individual pyramidal tract neurons (originating from M1) and force generation at the wrist (Evarts, 1968). Since then, several other studies have shown a direct relationship between neural activity in M1 and the strength of force production or muscle contraction (Holdefer & Miller, 2002; Kakei, Hoffman, & Strick, 1999; Morrow, Jordan, & Miller, 2007; Pohlmeyer, Solla, Perreault, & Miller, 2007). This might be expected, given that much of the output from M1 projects on the spinal cord and motor neurons. However, even though the outputs from M1 must at some level control muscle activation, the exact "language" of M1 activity is still an open question.

One popular view of the neural "code" employed by M1 arose in the 1980s with the work of Georgopoulos (Georgopoulos, Kalaska, Caminiti, & Massey, 1982). He recorded the activity of individual neurons in M1 and found that the rate of discharge modulated with the direction of hand motion during reaching. Furthermore, the relationship to reach direction could be well

approximated by a sinusoid (tuning curve) with parameters fit to the baseline firing rate, modulation depth, and preferred direction. The preferred direction (often shortened to PD) refers to the direction of hand movement corresponding to the highest firing rate for a given neuron. This kinematic framework in which individual M1 neurons are represented as encoding high-level features of movement found a strong foothold, especially in the recent field of brain machine interfaces (Carmena et al., 2003; Chapin, Moxon, Markowitz, & Nicolelis, 1999; Serruya, Hatsopoulos, Paninski, Fellows, & Donoghue, 2002; Taylor, Tillery, & Schwartz, 2002; Wessberg, Stambaugh, Kralik, & Beck, 2000). While the approach provides a simple signal for control purposes, it does not fully capture the nature of M1 activity, which can also correlate with other factors such as force or muscle activity and even posture (Scott & Kalaska, 1995; Scott, Sergio, & Kalaska, 1997). This might indicate a highly integrated, multimodal representation in M1. It might also imply that the mechanism of movement generation is not based on any parameterization. This concept will be explored at greater length in a later section, *Analyzing motor cortical signals*.

Dorsal premotor cortex

Directly anterior to the primary motor cortex on the precentral gyrus is the premotor cortex, also known as Brodmann area 6. The area can be distinguished from the primary motor cortex by the lack of layer 5 Betz cells, and from the prefrontal cortex by the lack of layer 4 granulation (Purves et al., 1997; Rajkowska & Goldman-Rakic, 1995). The scope of "premotor cortex" encompasses several distinct areas, including the supplementary motor area (SMA) along the most medial aspect, and dorsal and ventral premotor areas (PMd and PMv), split by the spur

of the arcuate sulcus (Lu, Preston, & Strick, 1994; Picard & Strick, 2001). Activity within these three main divisions has been shown to correlate with aspects of voluntary motor control. SMA appears to be largely involved with movement sequencing (Mushiake, Inase, & Tanji, 1991; Tanji, Shima, & Mushiake, 1996). PMv displays a wide range of movement-related effects, including visual and somatosensory integration and the coordination of reaching and grasping (Eiji Hoshi & Tanji, 2002; E. Hoshi & Tanji, 2006; Jeannerod, Arbib, Rizzolatti, & Sakata, 1995; Murata et al., 1997; Romo et al., 2004). PMd, which is the focus of my investigations presented later, can be further divided—around the level of the genu of the arcuate sulcus—into rostral (PMDr) and caudal (PMDc) aspects (Picard & Strick, 2001). These rough divisions (despite no clear anatomical boundary) segregate the two main functional roles observed in PMd. The caudal aspect (PMDc) seems closely tied to movement planning and execution (D. J. Crammond & Kalaska, 2000; Weinrich & Wise, 1982), not unlike the adjacent primary motor cortex. The rostral aspect (PMDr) more closely resembles prefrontal areas, with responses that appear to reflect task-related abstractions like rule-based transformations from visual cues to movement (Mirabella, Pani, & Ferraina, 2011; Ohbayashi, Ohki, & Miyashita, 2003). Intracortical electrical stimulation of PMDr can even elicit eye movement (Fujii, Mushiake, & Tanji, 2000), also suggesting integration with the visual system.

The premotor areas are strongly connected to the primary motor cortex, as well as subcortical structures (e.g., basal ganglia) and the parietal cortex (Dum & Strick, 2002; Kurata, 1991; Stetson & Andersen, 2014). Traditionally the connection between premotor cortex and primary motor cortex was viewed as hierarchical, with premotor areas thought to perform high-level processing and primary motor cortex the control of low-level motor output. However, the finding that premotor and supplementary motor cortices significantly contribute to the

corticospinal pathway challenges this view (Dum & Strick, 1991; Martino & Strick, 1987). While premotor areas may indeed play an important role in processing the signals that arrive in M1, they also have a means of directly controlling movement. Electrical stimulation of premotor cortex has been shown to elicit movement, though at higher current thresholds than primary motor cortex (Weinrich & Wise, 1982).

Electrophysiological studies of the dorsal premotor cortex have mainly focused on the role that the caudal region (adjacent to M1) has in movement planning. Analysis methods of PMd activity largely mirror those used on primary motor cortex activity, using directional tuning curves to describe discharge as a function of reach angle. This approach first led to the observation that activity in PMd could indicate the direction of a visually cued reach long before the movement was initiated (D. J. Crammond & Kalaska, 2000; Godschalk, Lemon, Kuypers, & Van der Steen, 1985; Riehle & Requin, 1989). Such kinematic representations persisted even when the visual cue was removed, suggesting that they were indeed related to movement planning, and not just the presence of a visual cue (P. Cisek & Kalaska, 2005; D. J. Crammond & Kalaska, 2000). PMd appears to operate using a diverse set of movement plan representations, and displays a mix of externally-driven (i.e.,, visual location of a target) and motor-related (i.e.,, associated reach direction) activity (Batista et al., 2007; M. M. Churchland, Santhanam, & Shenoy, 2006; Gail, Klaes, & Westendorff, 2009; Shraga Hocherman & Wise, 1990; McGuire & Sabes, 2009; Messier & Kalaska, 2000; Pesaran, Nelson, & Andersen, 2006; Schaffelhofer & Scherberger, 2016; L. Shen & Alexander, 1997). This mixed selectivity suggests that part of PMd's role in planning may be to transform visual cues to the appropriate movement plan.

Regardless of the specific coordinate frame, PMd is clearly involved in the planning of goal-directed movement. A major question, then, is how planning processes proceed in the face

of goal ambiguity. The first description of PMd activity under incomplete goal specification came in a 1989 paper by Riehle and Requin (Riehle & Requin, 1989). Monkeys were cued to two different magnitudes of wrist extension and wrist flexion. On most trials, the monkeys were given simultaneous information about the required direction (extension/flexion) and extent. However, on other trials, the delivery of information was sequential, creating periods of partial goal specification. PMd activity on these trials suggested that it could begin to plan aspects of the upcoming movement even in the absence of a fully specified target. This observation was extended through work by Bastian et al. in 1998 (Bastian, Riehle, Erlhagen, & Schöner, 1998), who found that early planning-related activity in PMd appeared to represent a single reach plan whose width changed to encompass the full range of potential reach targets. In 2005, Cisek and Kalaska (P. Cisek & Kalaska, 2005) studied PMd planning activity when presented with two opposing reach targets. They showed that when the two targets first appeared, PMd simultaneously developed two corresponding reach plans. Following a visual cue identifying the correct target, the correct representation was strengthened, and the incorrect representation was suppressed. This study led to widescale acceptance that PMd can simultaneously represent multiple discrete reach plans. However, this result has only been shown using trial-averaged, single neuron analyses. It is possible that the apparent simultaneous representation is in fact an artifact of trial-averaging, and that PMd only ever plans one reach at a time (and switches that plan if necessary). In Chapter 3, I provide evidence of this alternative hypothesis, using population recordings to track moment-by-moment reach planning.

Analyzing motor cortical signals

Since the origins of *in vivo* neural recordings, our understanding of brain function has relied heavily on single neuron recordings. In the motor cortices, individual neurons were found to correlate with both kinematic (Georgopoulos et al., 1982; S Hocherman & Wise, 1991; Messier & Kalaska, 2000) and kinetic (Evarts, 1968; Georgopoulos, Ashe, Smyrnis, & Taira, 1992; Kakei et al., 1999) features of limb control control. This led to a search for the true "language" of motor cortex. Does motor cortex represent movement kinematics or kinetics? Are representations developed in a global coordinate frame, a visual coordinate frame, or an intrinsic (e.g., joint) coordinate frame? Unfortunately, experiments designed to differentiate these possibilities often came to the unsatisfying conclusion that motor cortex contains a mixed representation encompassing virtually any and all features related to movement (Batista et al., 2007; Pesaran et al., 2006). Regardless, most approaches continued with the general assumption that activity from single neurons reflected the representation (or encoding) of some unknown—and probably complex—movement parameter(s).

The parameter-based view of motor cortical activity did prove useful for decoding purposes (Serruya et al., 2002; Taylor et al., 2002), and still forms the basis of many current analytical approaches. However, there are features of motor cortical activity that seem inconsistent with this approach. Temporal firing rate profiles are highly variable, both across neurons and across different movements (M. M. Churchland & Shenoy, 2007). This dynamic complexity can result in a single neuron appearing to display different tuning properties (e.g.,preferred direction) at different times within a single movement (Fu, Flament, Coltz, & Ebner, 1995; Mason, Johnson, Fu, Gomez, & Ebner, 1998), even rotating one hundred and eighty degrees between the planning and execution phases of a movement. Kinematic tuning

properties in M1 also shift when reaching against external forces (Sergio & Kalaska, 1998). These observations make little sense if neurons are assumed to consistently represent the values of specific movement parameters. Both the need for multivariable coordinate frames (i.e., combined representations of visual coordinates, force generation, kinematics, etc.) and inconsistencies even within those coordinate frames (rotating PDs) draws the parameter-based view of motor cortex into question.

While motor cortical activity does unquestionably correlate with external movement parameters like reach direction, the goal of motor cortex is to generate movement, not represent it (Fetz, 1992; Kalaska & Crammond, 1992). There is no obvious reason why single neuron activity must reflect representation of any movement variable. The motor cortex—a deeply interconnected network of neurons—must ultimately output signals to control coordinated spinal cord circuits and muscle contractions. Rather than attempt to interpret the activity of single neurons with respect to extrinsic movement variables directly, it may be more appropriate to view them in relation to the function of the entire network. This population-based approach offers a mechanism for movement generation that can also explain the single-neuron phenomena (e.g.,dynamic preferred directions) seemingly at odds with a representational view (M. M. Churchland, G. Santhanam, et al., 2006; Michaels, Dann, & Scherberger, 2016).

Population-based analysis methods are particularly useful when applied to simultaneously recorded neuron data. Early electrophysiological experiments were limited to recordings from one or a small number of electrodes driven into the brain. However, with the development of chronic multi-electrode arrays came the ability to record from tens to hundreds of neurons simultaneously. An increasingly popular approach for analyzing array recordings is to use dimensionality reduction methods. Dimensionality reduction begins with the assumption that

although a brain area may contain millions or billions of individual neurons, the network functions at a much more limited scale. The goal of dimensionality reduction is to use the correlations between individual neurons to reveal a small number (on the order of tens) of large-scale patterns in the population activity. Interpretation of the activity within the context of the low-dimensional space is often referred to as the "neural state". These methods have been used on data combined across sessions from highly stereotypic movements to provide general insight into motor cortical network function (M. M. Churchland et al., 2012; Elsayed, Lara, Kaufman, Churchland, & Cunningham, 2016). However, the most compelling application is to simultaneously recorded activity, which can allow the neural state to explain individual movement features. An instantaneous readout of the neural state can prediction reaction time (Afshar et al., 2011; Michaels, Dann, Intveld, & Scherberger, 2015) and even hesitation and change of mind (Kaufman, Churchland, Ryu, & Shenoy, 2015). The utility of low-dimensional neural states in single-trial analysis will be discussed more in Chapter 3.

NEURAL MECHANISMS OF UNCERTAINTY PROCESSING

Overview

As described in previous sections, behavior on motor tasks is affected by environmental, sensory, and motor uncertainty. Importantly, injecting noise and uncertainty into, say, a sensory cue, does not simply increase variance in the motor output. Instead, behavior may be biased in a way that implies knowledge of the uncertainty itself. For the brain to incorporate uncertainty necessitates a neural mechanism (or multiple mechanisms) for representing uncertainty. Theoretical neuroscientists have posited several potential means by which cortex might achieve

uncertainty representation. Here I will review a few of the most common hypotheses, and discuss the potential implications for motor cortex function. I will also discuss hypothesized neural mechanisms for dealing with choice ambiguity.

Models of noise compensation

The brain might represent uncertainty in an arbitrary signal in two basic ways: (1) it could encode separate estimates for the value itself (e.g.,its mean) and its uncertainty (e.g.,variance), or (2) it could represent the full distribution of possible values. Most modeling approaches assume the latter case, and suggest that cortical signals do not simply reflect single parameter values, but distributions. These approaches can be loosely grouped together as population code models.

The concept of the population code is based on neural tuning, which is described in the above sections discussing motor cortex. Neurons in sensory and motor areas tend to display broad tuning characteristics with respect to external variables like visual motion or reach direction. In many cases this tuning is unimodal, with neurons responding maximally (via firing rate) for one specific stimulus or parameter value. The result of these tuning characteristics is that a single stimulus value will elicit graded responses throughout the neural population. Due to the inherent variability in neural firing, this population response can be readily interpreted as a probability distribution (J. M. Beck et al., 2008; Deneve, Latham, & Pouget, 2001; Ma et al., 2006; Salinas & Abbott, 1994; Sanger, 2003), forming the basis of the probabilistic population code (PPC) hypothesis.

PPC models achieve probability representation through gain modulation. That is, more uncertain stimuli elicit weaker responses. This model of cortical function has favorable properties, in part due to the Poisson-like nature of neural firing. The Poisson distribution is a member of the exponential family of functions, and is therefore particularly useful for Bayesian integration since distributions can be combined through addition rather than multiplication. Thus, multisensory integration in a PPC framework requires only the summing of neural responses (Jazayeri & Movshon, 2006; Ma et al., 2006). The model offers a simple and plausible mechanism by which the inherent properties of individual neurons can underlie optimal macro-level behavior. However, it is unclear how easily PPC—which relies on mean rates—might encompass other aspects of Bayesian inference. For example, it does not naturally include marginalization, which describes the dependence of a given probability on the considered set of possible outcomes (Lochmann & Deneve, 2011). It also considers all neurons as independent samples of a stimulus, which is likely a poor assumption considering the interconnected nature of cortex.

A separate model suggests that neural activity might represent a sampling of internal representations of the environment (Berkes, Orbán, Lengyel, & Fiser, 2011; Hinton & Sejnowski, 1983; Hoyer & Hyvärinen, 2003). Under this model, cortex performs Bayesian inference to update the posterior distribution over all environmental states that may have led to the observed sensory input. This view attributes the variability observed in neural spiking patters to a different source than the PPC model. The PPC model assumes that a single input elicits stochastic neural responses, and thus neural variance is simply a result of intrinsic properties of neurons themselves. Alternatively, the sampling models suggest that neural variance reflects the unavoidable uncertainty in performing Bayesian inference. That is, each sample of neural

activity reflects the (essentially noiseless) representation of a single world state that may have led to the observed sensory input (Hoyer & Hyvärinen, 2003). Neural variability thus arises from the distribution of inferred world states. The sampling model predicts that increased uncertainty in a sensory cue will increase temporal variability, interneuron variability, or both. This contrasts with the PPC model, which predicts lower overall firing rates (and thus lower variability, per Poisson statistics) with increased uncertainty.

A third popular model of uncertainty representation suggests that uncertainty is represented through the temporal dynamics of individual neurons (Deneve, 2008a, 2008b). Like the sampling model, this temporal model assumes that neural activity reflects an inferred world state. However, under this model, the activity of a single neuron reflects an update to the evidence it receives regarding a single environmental parameter. More specifically, each neuron spikes when its synaptic input cannot be readily predicted from its own recent firing history. This model accurately recreates many features of neural firing, including Poisson firing statistics. Uncertainty representation is directly incorporated into the model, as noise affects the probability that a neuron will receive novel supporting evidence.

Each of the three models described above has found limited physiological support, which likely indicates that none fully encapsulates the nature of neural uncertainty representation. Additionally, the models have only been used to explicitly model sensory areas like the middle temporal visual area (MT). Uncertainty representation within motor areas thus remains largely speculative. As mentioned by Ma (2006), it is unclear if uncertainty representation even makes sense within motor areas, since movement generation likely necessitates a single motor representation. The results provided in Chapter 2 aim to address this open question of uncertainty representation within motor areas.

Models of decision-making between multiple options

Modeling the processing of ambiguity-related uncertainty is equivalent to modeling decision-making. Given two or more potential options, how does the brain use the available information (if any differentiating information even exists) to decide between them? On this front, the most notable and widespread modeling approach has been the "bounded integrator" (Bogacz & Gurney, 2007; Carpenter & Williams, 1995; Grossberg & Pilly, 2008; Mazurek, 2003; Ratcliff, 1978; Reddi, Asrress, & Carpenter, 2003; Smith & Ratcliff, 2004; Stone, 1960; Usher & McClelland, 2001; Wang, 2002; K. F. Wong & Wang, 2006). In this model (assuming two potential targets), sensory evidence for each target accumulates over time. The difference in evidence thus represents the decision variable, and must surpass a predefined threshold to reach a final decision. This model accurately predicts behavioral phenomena, such as the effect of sensory evidence on reaction time and the presence of incorrect choices (Carpenter & Williams, 1995; Ratcliff, 1978; Reddi et al., 2003; Smith & Ratcliff, 2004). As with noise-based uncertainty models, the bounded integrator model has been applied mostly in the context of vision and eye movement (Carpenter & Williams, 1995; Grossberg & Pilly, 2008; Mazurek, 2003; Reddi et al., 2003; Smith & Ratcliff, 2004). Indeed, neural activity in multiple eye-related brain areas has been found to closely mimic an integrator during decision-making (Gold & Shadlen, 2000; Leon & Shadlen, 2003; Munoz & Wurtz, 1995; K. Shen & Pare, 2007).

In the reaching system, a model of decision-making based on "biased competition" between action choices (P. Cisek, 2006, 2007) closely resembles a bounded integrator. This model frames the simultaneous encoding result observed in PMd (see *Dorsal premotor cortex*) as

the fundamental mechanism by which movement decisions are made. In a standard bounded integrator model, decisions depend on the *difference* in evidence between multiple options. The biased competition model instead assumes a distinct representation for each action possibility. The strength of each representation is consequently affected by sensory evidence (which "biases" the competition between options) as well as inhibitory connections between the corresponding sub-populations. Single neuron results appear to support the biased competition model in PMd (Coallier et al., 2015; Pastor-Bernier & Cisek, 2011), although alternative interpretations will be discussed in Chapter 5.

A problematic feature of the bounded integrator model is its sluggish response to changes in evidence. That is, if evidence were to quickly switch from supporting target A to target B, the new evidence would only slowly be incorporated into the output. Neural recordings in PMd suggest a more instantaneous response to evidence, which does not fit with the integrator model of decision making (Thura & Cisek, 2014). An extension of the biased competition model, called the urgency gating model (Thura, Beauregard-Racine, Fradet, & Cisek, 2012), was developed to deal with this issue. The urgency gating model posits that decision-making areas (principally PMd) don't integrate sensory evidence, but rather track it on a moment-by-moment basis. The evidence is then combined with a time-varying urgency signal, which dictates the certainty required to for a decision. This model captures the tradeoff that exists between making quick decisions with little evidence and waiting for more information.

REINFORCEMENT AND MOTOR LEARNING

Overview

The elements of movement planning and execution described thus far consider only individual movements performed on short timescales. However, motor control is not limited to the task of generating useful movement, but also includes the process of learning what "useful" means in different contexts. This is especially true in conditions of high uncertainty, when initial performance is likely to be poor. Learning is an integral aspect of controlling movement, and requires incorporating previous movement outcomes (e.g.,success or failure) into the processes underlying future motor plans. In the following sections I will provide a brief overview of motor learning, with an emphasis on long-term, reward-based learning.

Both internal and external feedback drives motor learning

The motor system is inherently adaptive, as evidenced by the ability to learn and improve motor performance in novel contexts. Two illustrative examples of this are curl field reaching experiments and visual perturbation experiments. During a classic curl field reaching experiment by Shadmehr and Mussa-Ivaldi (Reza Shadmehr & Mussa-Ivaldi, 1994), subjects were asked to make targeted reaching movements using a manipulandum that applied forces to distort the dynamics of movement. Initially, trajectories displayed strong deviations from normal reaches, but over time subjects learned to properly account for the new dynamics to produce straight reaches. Importantly, when the forces were removed, subjects temporarily made errors in the opposite direction. These errors under normal conditions indicated that the previous adaptation did not represent a cognitive strategy that could be applied or ignored, but rather a fundamental

restructuring of the motor system itself. Similar results were found for visual perturbation experiments, in which subjects gradually adapted to shifted visual feedback (Martin, Keating, Goodkin, Bastian, & Thach, 1996). Again, returning to the normal condition elicited oppositely-directed errors. These two experimental paradigms demonstrate that the sensorimotor transformation from sensory input to movement execution is a highly adaptive process.

The motor learning processes involved in dynamic or visual perturbation experiments—and indeed motor learning in general—were traditionally thought to be wholly internal (Doya, 2000; Mazzoni & Krakauer, 2006; R. Shadmehr & Krakauer, 2008). That is, motor adaptation was thought to depend only on the comparison between desired motor output and observed sensory feedback. Under this view, external factors like reward are extraneous to the task of motor adaptation. This interpretation was bolstered by a result from Mazzoni and Krakauer (2006), which showed that motor adaptation can at times be maladaptive. Subjects made reaching movements under a visuomotor rotation that offset cursor movement from hand movement by forty-five degrees. They were instructed at the start of the experiment to compensate for the perturbation by aiming for an adjacent target. This cognitive strategy led to successful compensation at first, but performance degraded over time as implicit learning within the sensorimotor system gradually "corrected" for the imposed rotation. This result suggests that correspondence between motor intent and sensory feedback (i.e., sensory prediction error) is the driving force behind motor adaptation.

While sensory prediction error may drive motor adaptation, it cannot account for the entirety of motor learning. As discussed in previous sections, movement execution requires both an extrinsic action plan and subsequent transformation of that plan into motor (muscle) signals. Motor adaptation through sensory prediction error describes the process by which the goal-to-

output transformation can be continually updated. However, in the absence of sensory feedback, task-related feedback can still be used to drive motor learning. In their 2011 paper, Izawa and Shadmehr showed that subjects can successfully adapt to a visuomotor rotation when only provided with feedback of their reaches' success or failure (Izawa & Shadmehr, 2011). Result-based feedback can also modulate the extent of motor adaptation, with reward and punishment differentially affecting aspects of learning (Abe et al., 2011; Galea, Mallia, Rothwell, & Diedrichsen, 2015). A punishment signal (loss of reward) significantly increased the rate of motor adaptation compared to sensory feedback alone. A separate study showed that reward-related feedback enhanced adaptation for both gradual and abrupt perturbations (Nikooyan & Ahmed, 2015). These results and others (Dayan & Balleine, 2002; Hollerman & Schultz, 1998; O'Doherty, Dayan, Friston, Critchley, & Dolan, 2003; W. Schultz, 2006) indicate the strong influence of external reward-related feedback on motor learning.

Neural substrates of motor learning

Both sensory feedback (i.e., sensory prediction error) and reward-related feedback contribute to motor learning, but are likely governed by separate neural systems. The primary candidate area driving motor adaptation from sensory prediction error is the cerebellum (R. Shadmehr & Krakauer, 2008; R. Shadmehr, Smith, & Krakauer, 2010; Tseng, Diedrichsen, Krakauer, Shadmehr, & Bastian, 2007). Supporting this hypothesis are studies showing that individuals with cerebellar deficits have difficulty in learning motor skills (Nicolson et al., 1999; Sanes, Dimitrov, & Hallett, 1990). The strong interconnection between cerebellum and the premotor and primary motor cortices (Dum & Strick, 2002) indicates a potential means by which

the cerebellum could directly influence motor plans throughout learning. Facilitation of the motor cortex using transcranial direct current stimulation (TDCS) improves motor learning (Nitsche et al., 2003), which strengthens the argument that the cerebellum facilitates learning through direct modification of cortical structures.

The effect of reward-related feedback on motor learning is thought to operate through dopaminergic signaling arising in the basal ganglia (Frank, Woroch, & Curran, 2005; Hollerman & Schultz, 1998; Wolfram Schultz, 1998, 2000, 2002). Dopaminergic neurons have widespread influence, and as a result, reward-related signals have been characterized in a number of different brain areas, including the midbrain, prefrontal cortex, and motor cortices (Matthew R Roesch & Olson, 2003, 2004; Wolfram Schultz, 2000; Wallis & Kennerley, 2010). The anterior cingulate cortex is also strongly implicated in reward-based learning (Cohen & Ranganath, 2007; Kennerley, Walton, Behrens, Buckley, & Rushworth, 2006). Event-related potentials (ERP) measured through electroencephalography (EEG) are thought to arise in the anterior cingulate cortex, and appear to signal reward-prediction error (Cohen & Ranganath, 2007). The anterior cingulate cortex shares connections with motor cortices (Paus, 2001), providing yet another potential pathway by which reward might influence motor cortical function. Given the widespread cortical effect of dopaminergic signaling, it is quite reasonable to assume that motor cortex itself might hold signatures of reward-related feedback. In Chapter 4, I present evidence that reward (or more accurately, failure) signals do indeed exist within motor cortex.

SUMMARY

In this chapter, I provided an overview of the different types of uncertainty that influence sensorimotor function, and their effects on behavioral outcomes. I also reviewed the potential neural mechanisms (mainly within motor cortex) that might address the problem of moving under uncertainty. Since motor control involves both movement generation as well as movement optimization, I summarized the effect of movement outcome (i.e., reward) on motor learning and the brain areas that might provide reward-based influence on motor cortex. In the subsequent chapters, I will present evidence that motor cortices (dorsal premotor cortex and primary motor cortex) are intimately engaged in uncertainty processing and modified by reward-based feedback. Chapter 2 investigates how environmental uncertainty in a reach goal affects movement representations during planning and execution. Chapter 3 addresses movement planning with multiple potential targets, and shows single-trial results that contradict current models of decision-making in PMd. In Chapter 4, I move from movement planning to movement outcome, and provide evidence of reward-related signaling within motor cortex. In Chapter 5, I discuss the findings from Chapters 2 through 4 and the implications for the functional roles of PMd and M1 during movement planning, decision-making, and motor learning.

CHAPTER 2: UNCERTAINTY LEADS TO PERSISTENT EFFECTS ON REACH REPRESENTATIONS IN DORSAL PREMOTOR CORTEX

Brian M Dekleva, Pavan Ramkumar, Paul A Wanda, Konrad P Kording, Lee E Miller

FOREWORD

This chapter is an adapted version of a paper published by eLife (2016). In it, we aim to address the question, "How does noise-related uncertainty in a reach goal affect planning- and execution-related cortical activity?". We find that monkeys' uncertainty about where to reach is reflected—throughout the entirety of planning and execution—in the activity in dorsal premotor cortex, but not primary motor cortex. This result not only indicates that PMd is sensitive to task uncertainty, but also suggests potential properties of decision-making within the motor system. As discussed in Chapter 1, decisions are often thought to reflect the output of processes within PMd. However, here we show that during targeted reaching, the quality of the decision represented in PMd might be limited by subjective sense of uncertainty. Implications of this finding on decision-making, error correction, and learning will be discussed in Chapter 5.

ABSTRACT

Every movement we make represents one of many possible actions. In reaching tasks with multiple targets, dorsal premotor cortex (PMd) appears to represent all these possible actions simultaneously. However, in many situations we are not presented with explicit choices. Instead, we must estimate the best action based on noisy information and execute it while still uncertain of our choice. Here we asked how both M1 and PMd represented reach direction during a task in which a monkey made reaches based on noisy, uncertain target information. We found that with increased uncertainty, neurons in PMd actually enhanced their representation of unlikely movements throughout both planning and execution. The magnitude of this effect was highly variable across sessions, and was correlated with a measure of the monkeys' behavioral uncertainty. These effects were not present in M1. Our findings suggest that PMd represents and maintains a full distribution of potentially correct actions.

INTRODUCTION

Each motor action we perform reflects only one of the many available or considered actions. In some situations, the full set of potential actions comprises a set of discrete choices (e.g., which of these three apples should I pick?). In these cases, the task for the sensorimotor system is to evaluate each option and decide which will lead to the most favorable outcome. However, these "target selection" situations represent only one type of motor related decision-making. In many other scenarios the sensorimotor system cannot simply *select* between multiple explicit options, but instead must *estimate* the best action based on continuous – and often noisy

– sensory information and learned experience. Reaching toward a familiar object seen only in the peripheral vision, or under poor illumination is one such example.

Though target selection represents only one type of sensorimotor task, it dominates the current literature on neural correlates of motor-related decision making. This is true for both the visuomotor system (Basso & Wurtz, 1997; Britten, Newsome, Shadlen, Celebrini, & Movshon, 1996; Fetsch, Pouget, DeAngelis, & Angelaki, 2011; Newsome et al., 1989; Shadlen & Newsome, 2001) and the reaching system (Bastian, Schoner, & Riehle, 2003; P. Cisek & Kalaska, 2005; Coallier et al., 2015; Messier & Kalaska, 2000; Thura & Cisek, 2014). These studies do vary significantly in the method by which they provide cues to elicit a motor response. These cues may indicate different parameters of the action, such as the direction or extent of the movement (Bastian et al., 2003; D. Crammond & Kalaska, 1994; Gail et al., 2009; Messier & Kalaska, 2000; Welsh & Elliott, 2005). They can be discrete (Meegan & Tipper, 1998; Thura & Cisek, 2014; Wood et al., 2011) or continuous (Gold & Shadlen, 2001; Hernandez et al., 2010; Resulaj, Kiani, Wolpert, & Shadlen, 2009), and can even span different sensory modalities (Hernandez et al., 2010; Romo et al., 2004). However, all share a common characteristic: the action is directed towards one of multiple mutually exclusive targets. This mutual exclusivity is a constraint specific to the task of target selection and does not exist in target estimation, since no explicit options are presented. It is therefore not obvious how the results from target selection tasks may or may not extend to the case of target estimation.

In both target selection and estimation, there is some degree of uncertainty in the decision making process as well as the final decision itself. This uncertainty largely depends on the ambiguity of the available cues. If the task includes a completely unambiguous cue indicating the correct choice, the decision will contain practically no uncertainty whatsoever. For example, one

standard multiple target selection task used in non-human primate reaching studies (Bastian et al., 2003; P. Cisek & Kalaska, 2005) briefly presents a monkey with two or more potential reach targets before indicating the correct one. In this situation the animal may be initially uncertain about which target is correct, but that uncertainty vanishes with the disambiguating cue. Variants of this task provide more ambiguous cues and allow the animal to choose one of two targets while still unsure about the correct choice (Coallier et al., 2015; Thura & Cisek, 2014). These kinds of tasks result in decisions that are made despite a lingering uncertainty in the decision.

Studies of reach-related brain areas during target selection tasks have suggested that the dorsal premotor cortex (PMd) plays a significant role in sensorimotor decision-making. Historically, PMd has been viewed as a movement planning area, displaying activity consistent with a representation of upcoming movements to visual targets (Paul Cisek, Crammond, & Kalaska, 2003; L. Shen & Alexander, 1997; Weinrich & Wise, 1982). Later studies showed that these pre-movement representations can include multiple simultaneous potential targets (P. Cisek & Kalaska, 2005) and reflect motor plans even in the absence of visual targets (Klaes, Westendorff, Chakrabarti, & Gail, 2011). Furthermore, the representations during multiple-target tasks are modulated by decision-related variables (Coallier et al., 2015; Pastor-Bernier & Cisek, 2011). These more recent results are consistent with an interpretation that activity in dorsal premotor cortex modulates with the complexity (or uncertainty) of a motor decision.

In general, sensorimotor decision-making should take into account the uncertainty present in all task-relevant information sources – namely the current sensation and prior experience. When sensation provides a highly reliable action cue (e.g., when reaching toward a well-lit, foveated object), it can be used exclusively to plan and execute the appropriate motor output. However, as uncertainty in sensation increases, it becomes more beneficial to combine

sensory information with information learned through prior experience. The optimal method for integrating sensory and prior information was formulated centuries ago as Bayes' theorem (Bayes, 1763). A direct application of Bayes' theorem states that cues should be weighted in inverse proportion to their variance (Knill & Saunders, 2003; K. P. Kording & Wolpert, 2006). The Bayes optimal decision will lead to better results than either cue alone, but will still contain a certain degree of uncertainty.

Bayesian models have been used to describe human behavior in a wide array of psychophysical studies (Cheng, Shettleworth, Huttenlocher, & Rieser, 2007), including visual (Knill & Saunders, 2003; Mamassian & Landy, 2001; Weiss et al., 2002), auditory (Battaglia et al., 2003), somatosensory (Goldreich, 2007), cross-modal (Alais & Burr, 2004; Ernst & Banks, 2002; Gu, Angelaki, & Deangelis, 2008; Rowland, Stanford, & Stein, 2007), and sensorimotor (Greenwald & Knill, 2009; Konrad P Kording & Wolpert, 2004; Trommershauser, Maloney, & Landy, 2008; R. J. van Beers et al., 2002) applications. In these tasks, behavior generally matched the predictions of various Bayesian models of optimal performance, which has been taken as evidence that the brain does indeed incorporate information about the relative uncertainty of various cues when planning and executing movements.

To probe the effect of target estimation uncertainty on M1 and PMd, we designed a task in which monkeys estimated the location of reach targets using knowledge of the average target location (learned through experience) and noisy visual cues. Although M1 activity appeared to reflect only the direction of the executed reach, we found that the monkeys' uncertainty about where to reach correlated with changes in PMd activity during both movement planning and execution. The magnitude of these uncertainty-related effects in PMd was spatially tuned. Neurons whose strongest response direction (their preferred direction, or PD) was aligned with

the planned reach direction remained largely unchanged, while neurons with PDs opposite the reach direction experienced a significant increase in activity with increased uncertainty. Neurons with intermediate PDs displayed somewhat smaller uncertainty-related effects. The uncertainty-related change in this off-direction neural activity varied considerably across sessions, not only because of experimentally altered prior and likelihood uncertainty, but also apparently because of the monkeys' own subjective uncertainty in their final action decisions. We found that the magnitude of these cross-session differences correlated with estimates of the monkey's decision-related uncertainty.

RESULTS

Task Performance During Reaching to Certain and Uncertain Targets

Our goal in this study was to understand the effect of uncertainty on movement representations in the motor system. To this end, we designed a behavioral task in which monkeys (two Rhesus macaques) made decisions about where to reach using a planar robotic manipulandum, based on the learned history of target distributions and uncertain visual cues. During the first block of trials, the monkeys made center-out reaches with an instructed delay to well-specified (zero uncertainty) targets that were randomly distributed across eight locations (Fig. 2.1A, top). In the second block of trials, the target locations were randomly drawn from a circular normal (von Mises) *prior* distribution centered on a single direction that remained constant for the remainder of the session. Additionally, the monkey did not receive veridical feedback about the location of the target, but instead saw a noisy distribution of five (monkey M) or ten (monkey T) lines (Fig. 2.1A, bottom). These lines were drawn from a *likelihood*

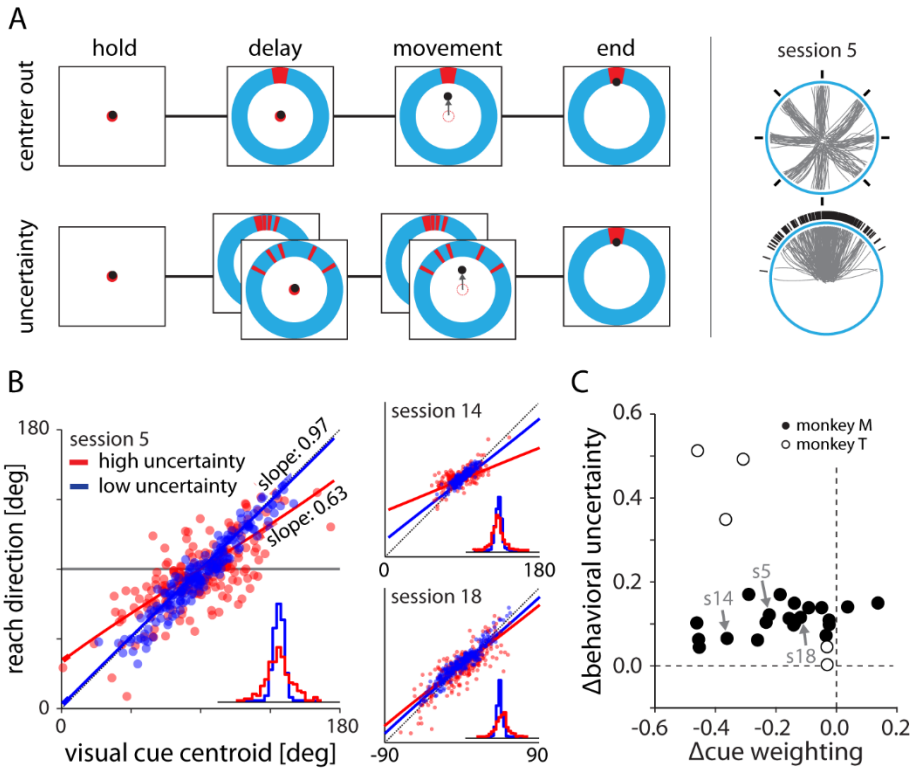

Figure 2.1 | Experimental setup and behavior
Experimental setup and behavior. (A) Monkeys made planar center-out reaches with instructed delay to visual targets. Illustrations on right show target locations (black) and reach trajectories (gray) for trials in the center out and uncertainty blocks for an example session. In the center out block, targets were distributed uniformly across eight directions and were cued with no uncertainty. In the uncertainty block, targets were sampled from a von Mises distribution and cued with stochastically sampled lines with either low or high variance. (B) Scatter plots of cue centroid versus reach direction for three sessions, with each dot representing a single trial. Under high uncertainty, the endpoints reflected an increased bias toward the average target location – indicated by a reduction in slope – and increased variability surrounding the fit line. (C) With the exception of two datasets from monkey M, fits to the behavioral scatter plots reveal reduced slope (negative Δcue weighting) for higher uncertainty targets. All datasets show greater residual variance with greater uncertainty.

distribution – also von Mises – centered on the correct target location and provided the monkey with noisy information about the target location. Each session contained at least two likelihood distributions of low and high variance, randomly interleaved across trials.

Therefore, during uncertainty trials, the monkey had two pieces of information available to estimate the target location: (1) the noisy visual cue and (2) a learned estimate of the distribution of previous target locations. According to Bayes' rule, optimal performance on the task would require the monkey to use the centroid of the displayed line segments (likelihood estimate) and the average target location (prior estimate), weighted according to the inverse of their variances. In general, this means that using an appropriately weighted sum of both the likelihood and prior estimates will on average result in fewer errors than either cue alone.

Fits to the scatter plot between the centroid of the visual cue and the reach direction reveal the monkey's relative weighting of the visual cue (the likelihood) and its estimate of the average target location (the prior; see methods for more information). A fitted line with a slope of zero would indicate complete reliance on the prior, while a slope of one would indicate reliance only on the likelihood. Panel B of Figure 2.1 shows several representative sessions. In each, the monkey relied more on the visual cue when its uncertainty was low (blue symbols) than when it was high (red symbols). We summarize the difference in visual cue weighting between the uncertainty conditions (Δcue weighting) for each session by subtracting the slopes of the fitted lines. The negative values of Δcue weighting in Figure 2.1C reveal that both monkeys almost always relied less on the visual cue during high uncertainty trials. This indicates that the monkeys combined information from both the displayed lines and the average target location in a Bayesian-like manner to estimate the location of the required reach target.

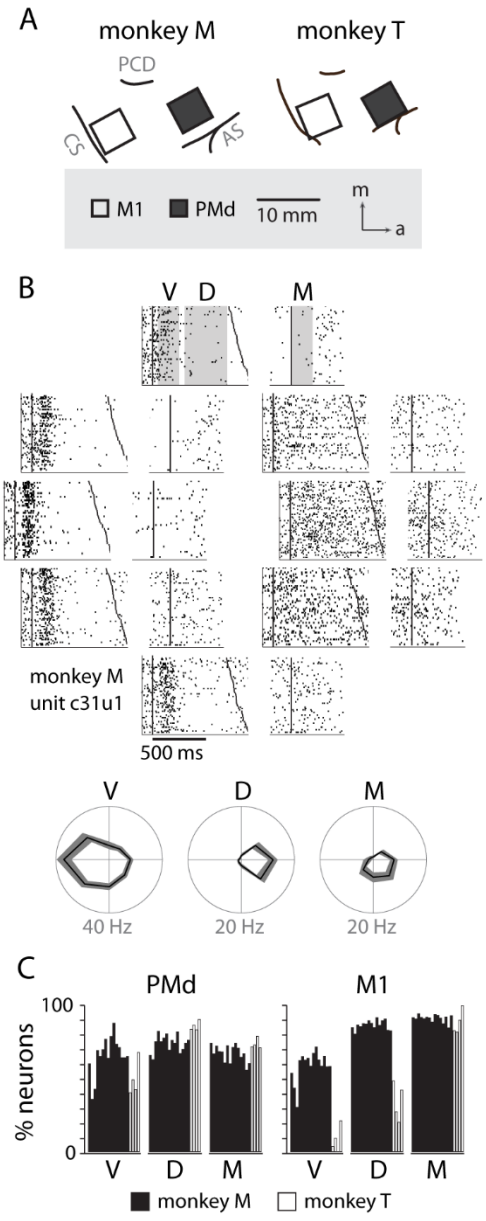

Figure 2.2 | Neural recordings and directional tuning
(A) Each monkey was implanted with two 96-channel microelectrode arrays, targeting the primary motor cortex (M1) and dorsal premotor cortex (PMd). (B) An example raster of a neuron in PMd displaying directional tuning, summarized below in three temporal periods: visual (V), delay (D) and movement (M). (C) Percentage of neurons from each session with significant tuning in each of the temporal periods.

Although there was a general tendency towards lesser weighting of the visual cue when it was more uncertain, there was a great deal of variability in that trend across sessions. In some instances, fits to the two uncertainty conditions revealed large differences in visual cue weighting (Fig. 2.1B red and blue fitted slopes, session 14) while in others the relative weight was nearly identical (Fig. 2.1B, session 18). Similarly, the behavioral uncertainty (as measured using the variance of the fit residuals) was sometimes very different between two conditions (Fig. 2.1 inset distributions, session 5) and sometimes nearly identical (Fig. 2.1, session 14). We characterized the total difference in uncertainty between the two conditions (Δbehavioral uncertainty) for each session by subtracting the circular standard deviation of the fit residuals. These two within-session metrics (Δcue weighting and Δbehavioral uncertainty) were very weakly correlated for monkey M and negatively correlated for monkey T (Fig. 2.1C). This variability provided a diverse set of uncertainty-related behavioral effects on which to examine neural activity.

Neural Activity

During the center out block of trials (zero uncertainty, eight discrete targets) many neurons in PMd displayed a robust burst of activity in the visual time period, followed by a more moderate, tonic response for the remainder of the delay period (e.g., Fig. 2.2B). We more formally described the population trends by calculating the percentage of neurons tuned in the visual (V), delay (D), and movement (M) time periods. The results for each session are shown in Figure 2.2C. We performed the same analysis for M1 neurons (Fig. 2.2C, right). In general, M1 displayed a bias toward delay and movement period tuning while PMd showed about equal percentages of tuned neurons for each time period.

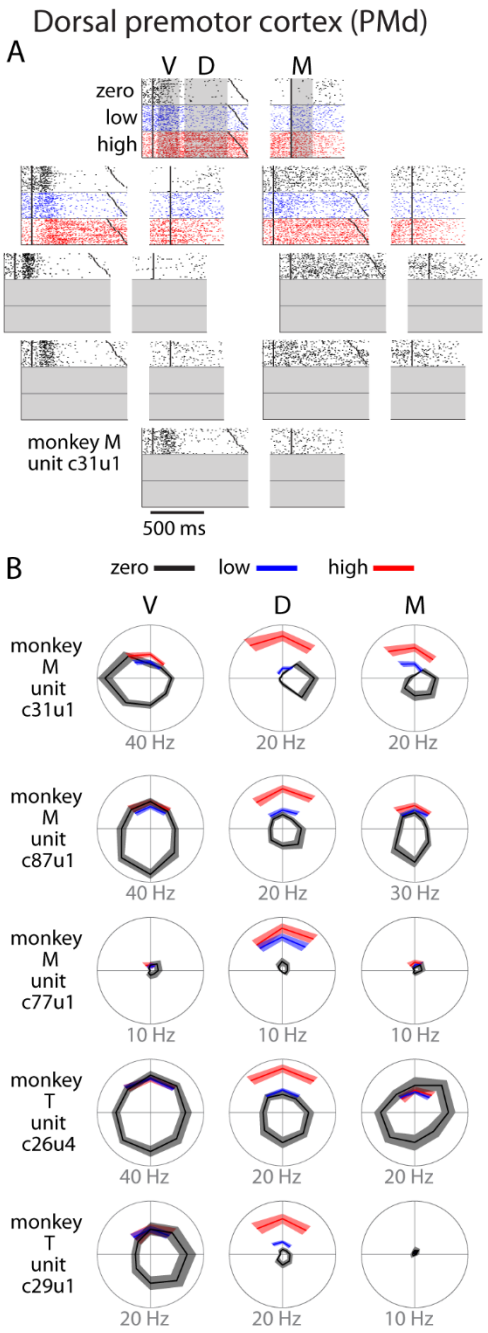

Figure 2.3 | Single unit activity in PMd
(A) Raster plot for an example neuron. Activity is aligned to either the visual cue appearance (left) or movement onset (M; right). Colors indicate zero (black), low (blue), and high (red) uncertainty conditions. (B) Directional tuning for other example neurons. Due to the nature of the task, reaches made during uncertain conditions with a non-uniform prior did not span all directions. Many neurons showed an increase in delay (D) or movement (M) activity as a function of uncertainty. Bounds on the tuning plots represent bootstrapped 95% confidence of the mean estimate.

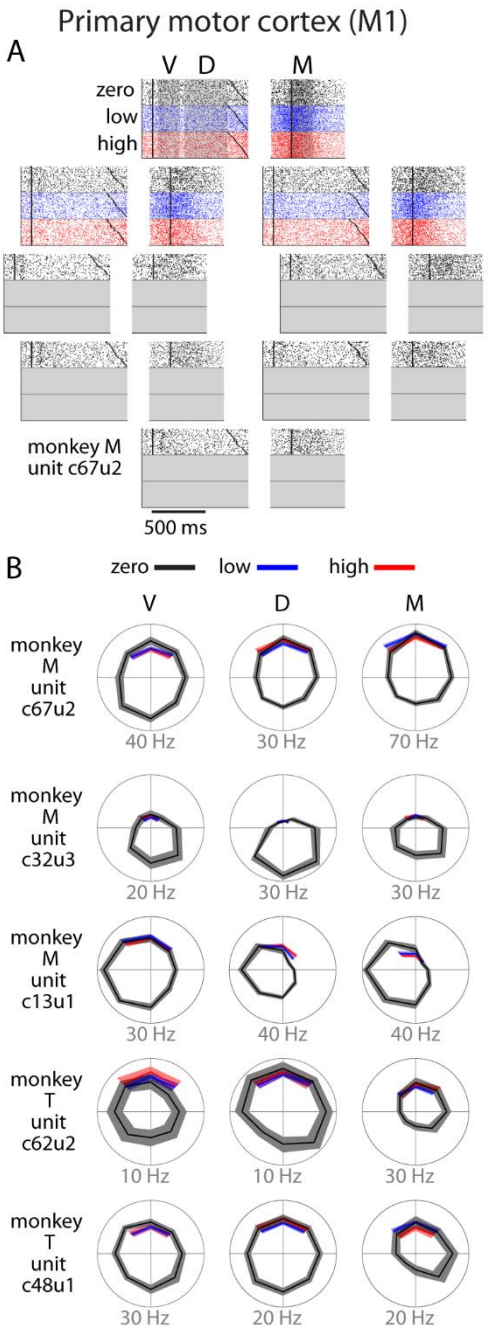

Figure 2.4 | Single unit activity in M1
(A) Raster plot for an example neuron with same conventions as Figure 2.3. (B) Directional tuning for other example neurons. In general, M1 activity was well-modulated by reach direction, but appeared to be largely unaffected by the uncertainty condition. Bounds on the tuning plots represent bootstrapped 95% confidence of the mean estimate.

During the remaining experimental blocks consisting of uncertain targets, we found many neurons in PMd to be more active during high uncertainty trials than low uncertainty trials (red vs. blue in Figure 2.3). This effect was most prominent during the delay (D) period, with some carryover into movement (M). Some neurons that had been essentially inactive during the block of zero-uncertainty reaches became strongly activated during the delay period of high-uncertainty trials (e.g.,c77u1 and c29u1, Fig. 2.3). We also noted that there was a greater tendency for increased activity in those neurons having PDs that were not aligned to the direction of movement (e.g.,c31u1 and c87u1, Fig. 2.3). Importantly, we found that greater uncertainty only ever led to increased activity.

M1 neurons did not display nearly the same degree of modulation with uncertainty as PMd neurons (Fig. 2.4). We observed neurons with strong directional tuning in all time periods, but this tuning was consistent across all uncertainty conditions. In general, analysis of single unit behavior suggested that M1 activity reflected only the reach direction and was largely unaffected by uncertainty.

Quantifying effects of uncertainty on firing rates

The anecdotal observations in Figures 3 and 4 strongly suggest that higher uncertainty leads to increased neural discharge in PMd but not in M1. Additionally, the magnitude of the uncertainty-related effect in individual PMd neurons was dependent on the neurons' tuning characteristics. A neuron experienced the greatest uncertainty-related activity increase when the reach was directed away from its preferred direction. To further examine this relationship

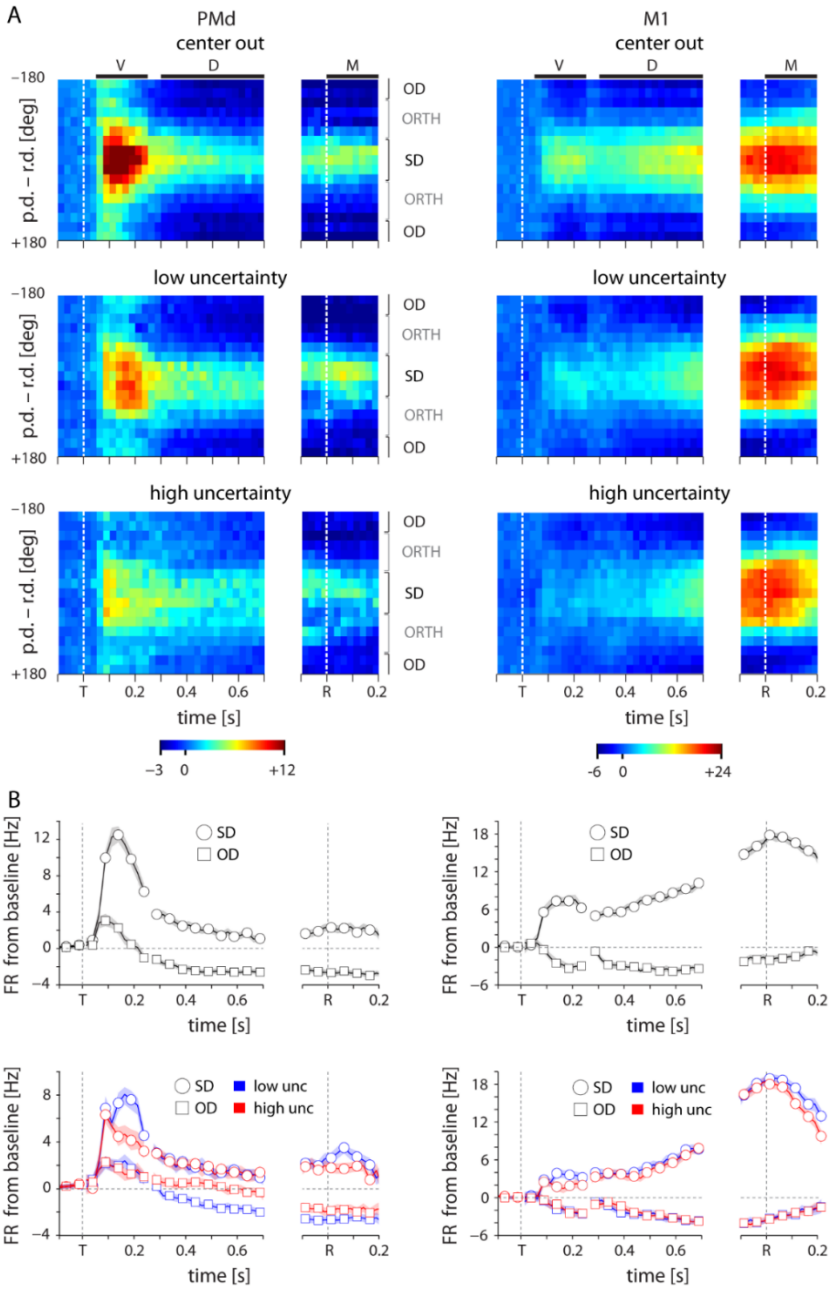

Figure 2.5 | Tuning-related changes in activity with uncertainty
Tuning-related changes in activity with uncertainty. (A) Spatiotemporal activity maps for PMd and M1. Neurons were binned on each trial by the distance between their preferred directions and the reach direction. Color indicates average change in firing rate from baseline in spikes per second. Left and right plots in each panel are alighend to taget onset (T) and reach onset (R) respectively. (B) Average change from baseline for SD and OD neurons in the initial center-out block (zero uncertainty; top) and subsequent blocks with low (bue) or high (red) uncertainty targets (bottom). High uncertainty trials resulted in reduced early activity for both SD and OD neurons, but an increase in OD activity for the remainder of the delay and movement phases. ORTH neurons were omitted for visibility. Error bars represent bootstrapped 95% confidence bounds on the mean estimate. For all plots, PDs were calculated separately for visual, delay, and movement epochs.

between tuning and uncertainty-related activity changes, we created spatiotemporal activity maps for both cortical areas in the manner of Cisek and Kalaska (2005) (Fig. 2.5). We binned each neuron's responses based on the angle between its PD and the reach direction. We then averaged across trials, resulting in population activity profiles centered on the reach direction.

In the zero-uncertainty condition, many PMd neurons displayed a burst of activity directly following cue appearance. This quickly resolved into a clear, maintained representation of the upcoming reach direction throughout the remainder of the delay and movement periods (Fig. 2.5A, top left). In contrast, M1 activity built more slowly as the trial evolved, ultimately producing a strong spatial representation of the executed reach direction (Fig. 2.5A, top right). However, while the recruitment of M1 neurons during low and high uncertainty conditions was similar (Fig. 2.5A, right), the representation in PMd differed significantly across these conditions. During high uncertainty trials, the representation of the reach direction was present but significantly less distinct, most notably due to increased activity in neurons with PDs far away from the reach direction (Fig. 2.5A, bottom left). To summarize this tuning-related effect, we partitioned the neurons into three groups for each trial: same direction (SD; preferred direction within 45 degrees of the reach direction), opposite direction (OD; preferred direction within 45 degrees of the anti-reach direction), and orthogonal direction (ORTH; preferred direction within 45 degrees orthogonalthe reach direction). After averaging the activity of these populations, it became clear that while both SD and OD neurons in PMd were less active immediately after high-uncertainty target appearance, the OD neurons increased their activity in the subsequent D and M periods. Thus the main effect of higher target uncertainty was an

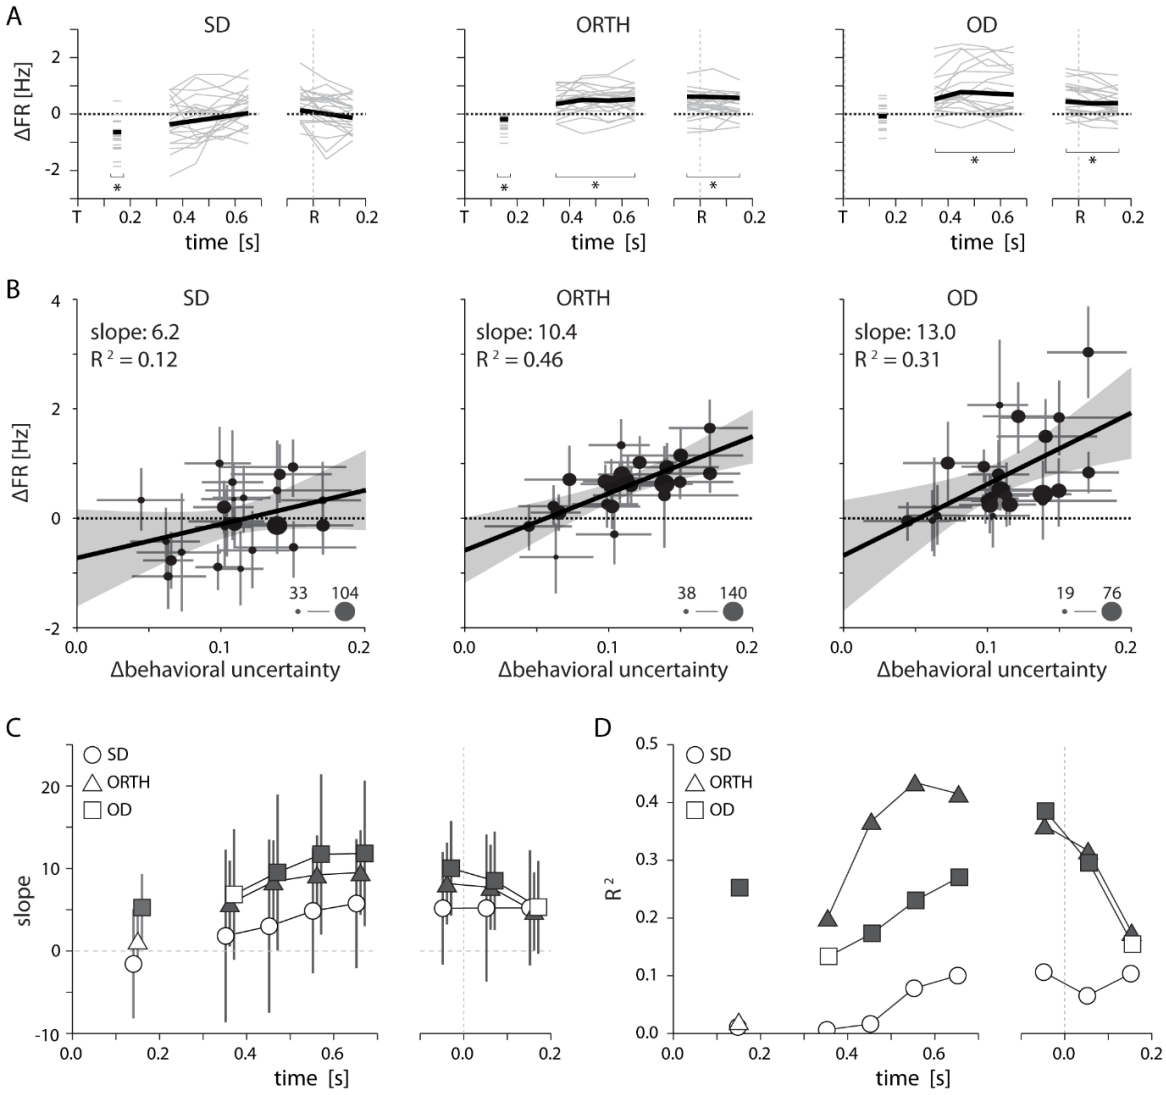

Figure 2.6 | Relationship between PMd activity and behavioral uncertainty
(A) Thin lines indicate the average difference in firing rate between high and low uncertainty trials for individual sessions. Heavy lines mark the mean across sessions. While SD neurons displayed an average change near zero, activity for ORTH and OD neurons was consistently higher for high uncertainty trials (B) Differences in firing rate between high and low uncertainty conditions as a function of the difference in behavioral uncertainty for a single time window 500-700 ms after target appearance. The correlation was weak for same direction neurons, but strongly positive for orthogonal and opposite direction neurons. Thus, the greater the difference in behavioral uncertainty, the larger the difference in activity for ORTH and OD neurons. Marker size indicates the number of contributing neurons for each session (C) The slopes from *B* calculated during the visual period (50-250ms after target appearance (left) and for 100ms time windows throughout the delay (middle) and movement (right) periods. The larger effect of behavioral uncertainty on OD and ORTH activity compared to SD activity persisted throughout planning and execution. (D) R^2 values for the linear fits in *C*. Filled symbols in *C* and *D* represent significant correlations, $p < 0.05$. All error bars represent bootstrapped 95% confidence bounds on the mean estimates.

increase in the PMd activity in neurons with preferred directions further away from the reach direction.

To summarize this uncertainty effect over sessions, we calculated the difference in average firing rates between low and high uncertainty conditions for SD, ORTH, and OD neurons. In most sessions, ORTH and OD activity during the delay and movement periods was significantly greater in the high uncertainty condition, while SD activity showed little change (Fig. 2.6A – monkey M; Fig. 2.7A – monkey T). However, the increase in OD activity varied considerably across sessions. We reasoned that the sessions with the greatest OD activity differences might correspond to the sessions with the greatest differences in the monkeys' uncertainty. To test this, we calculated the difference in behavioral uncertainty (Δbehavioral uncertainty) between uncertainty conditions for each session (see Methods: *behavioral task*). By plotting the activity differences as a function of Δbehavioral uncertainty, we found strong positive correlations for OD activity, but none for SD (Fig. 2.6B – monkey M; Fig. 2.7B – monkey T). For monkey M, the slope of the relation increased from SD to ORTH to OD neurons (Fig. 2.6B), consistent with the single-session example shown in Figure 2.5. We found very similar effects of uncertainty among OD neurons for monkey T (Fig. 2.7B). These findings suggest that as the monkeys became less certain about their decision of where to reach, the representations of less likely reach directions increased.

We also found that the tuning-related effect of uncertainty persisted throughout the entirety of movement planning and even after the initiation of the reach. We applied the analysis in Figure 2.6B to different time periods throughout the trial and plotted the slopes (Fig. 2.6C) and R^2 (Fig. 2.6D) relating Δbehavioral uncertainty to changes in SD, ORTH, and OD activity. For both monkeys, the difference in OD activity first displayed a significant correlation with

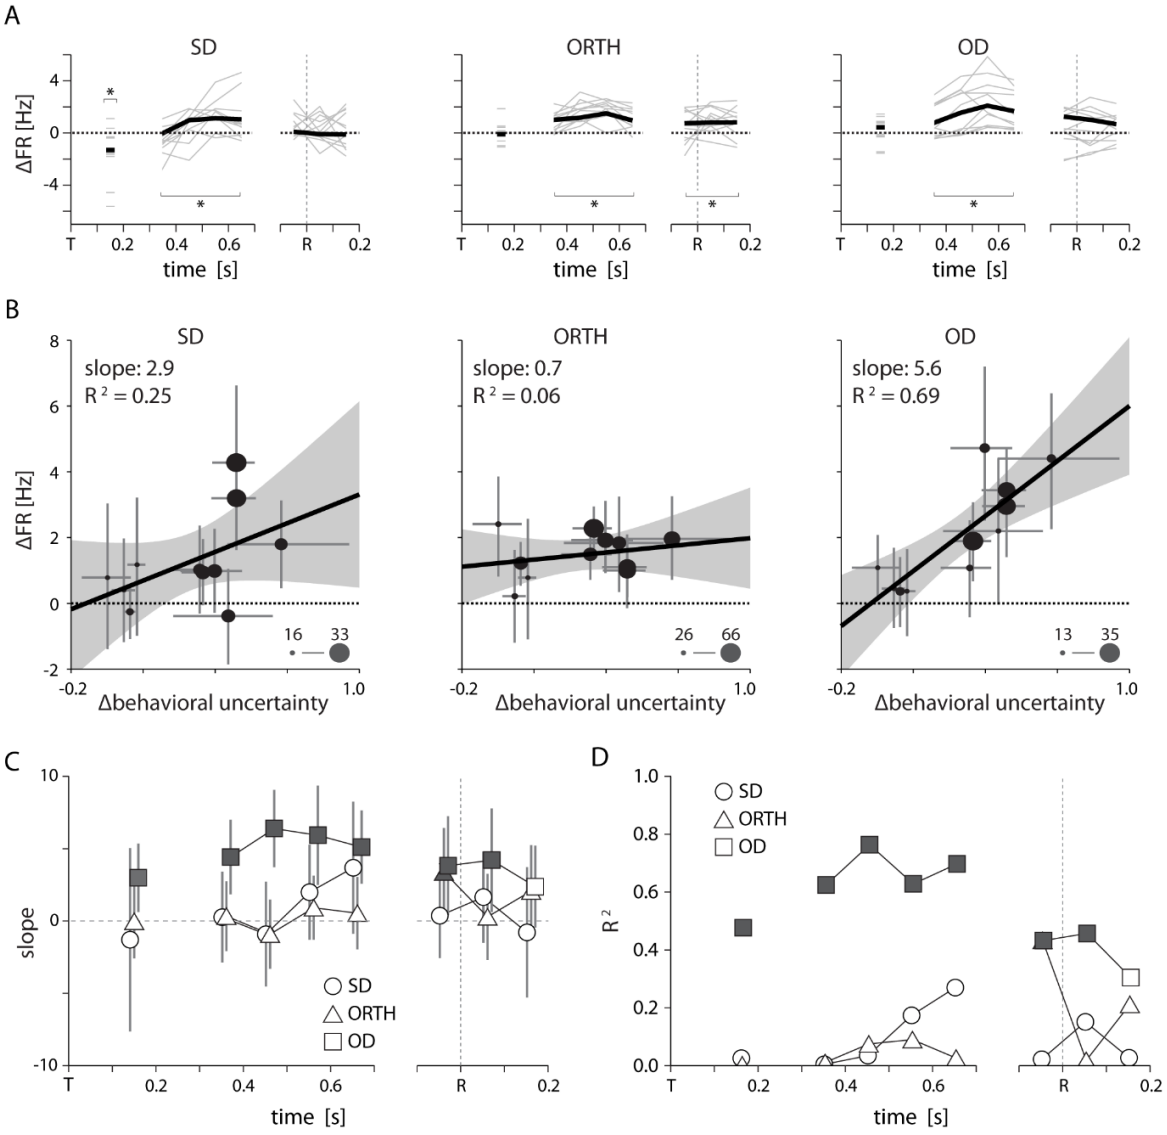

Figure 2.7 | Summary of uncertainty related activity in PMd for Monkey T
All conventions as in Fig. 2.6. Although we had only five sessions for monkey T, by splitting larger sessions into multiple blocks we obained 11 total data points. Specifics are given in Table 2.2.

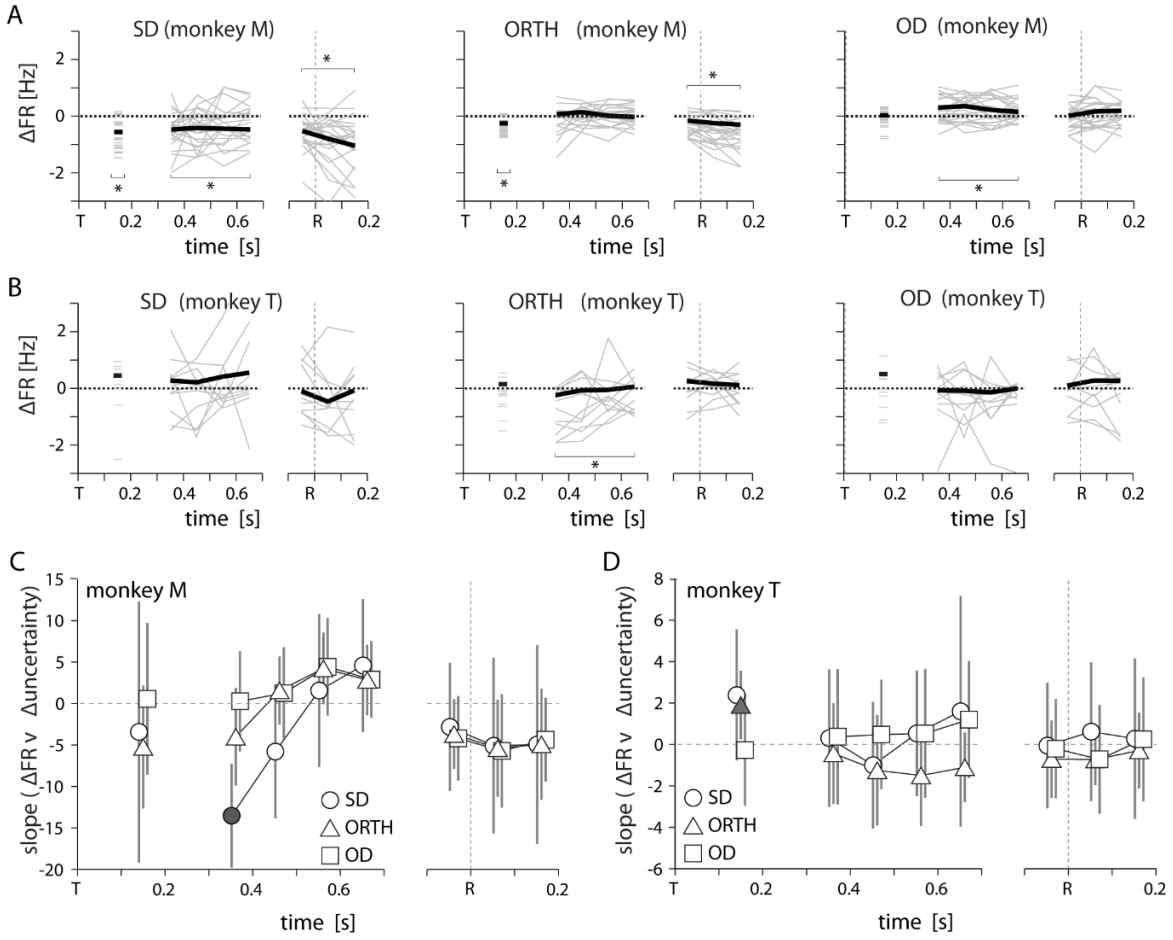

Figure 2.8 | Summary of uncertainty-related activity in M1 for both monkeys
All conventions as in Figs. 2.6, 2.7. Specifics of how we obtained datapoints for monkey T in panels B and D are given in Table 2.2.

Δbehavioral uncertainty during the visual period (figured 6 and 7, panels C and D). This effect persisted throughout the remainder of the delay period and the initiation of movement. ORTH activity displayed a similar trend but with a consistently shallower slope, indicating a weaker effect of uncertainty. SD neurons never displayed any significant correlation with uncertainty. For monkey T, only OD activity was consistently correlated with uncertainty throughout the delay and movement periods (Fig. 2.7C, D). Thus it appears that movement representations in PMd remain affected by decision-related uncertainty leading up to and throughout execution of a movement.

There was also substantial cross-session variability in the M1 firing rate measured between high and low uncertainty. For monkey M, SD activity was generally lower for high uncertainty trials and OD activity was slightly higher (Fig. 2.8A). However, there was rarely any correlation between the firing rate difference and the difference in behavioral uncertainty. For monkey M, SD activity was negatively correlated with uncertainty at the beginning of the delay period (300-400ms following target appearance; Fig. 2.8C). This effect dissipated quickly and was never observed for monkey T. As a result, we conclude that behavioral uncertainty had no significant effect on M1 activity during movement planning or execution.

Although the correlations between behavioral uncertainty and OD activity in PMd were significant, we considered the possibility that the neural effects were actually driven directly by the monkeys' relative weighting of the visual cue. To disassociate these two possibilities, we examined the independent correlations of OD activity with each of the two metrics in selectively subsampled groups of sessions. When we chose sessions that caused Δbehavioral uncertainty and Δcue weighting to be highly correlated (further exaggerating their normal relation), both metrics explained the change in OD activity (Fig. 2.9A). However, for subsampled groups of sessions

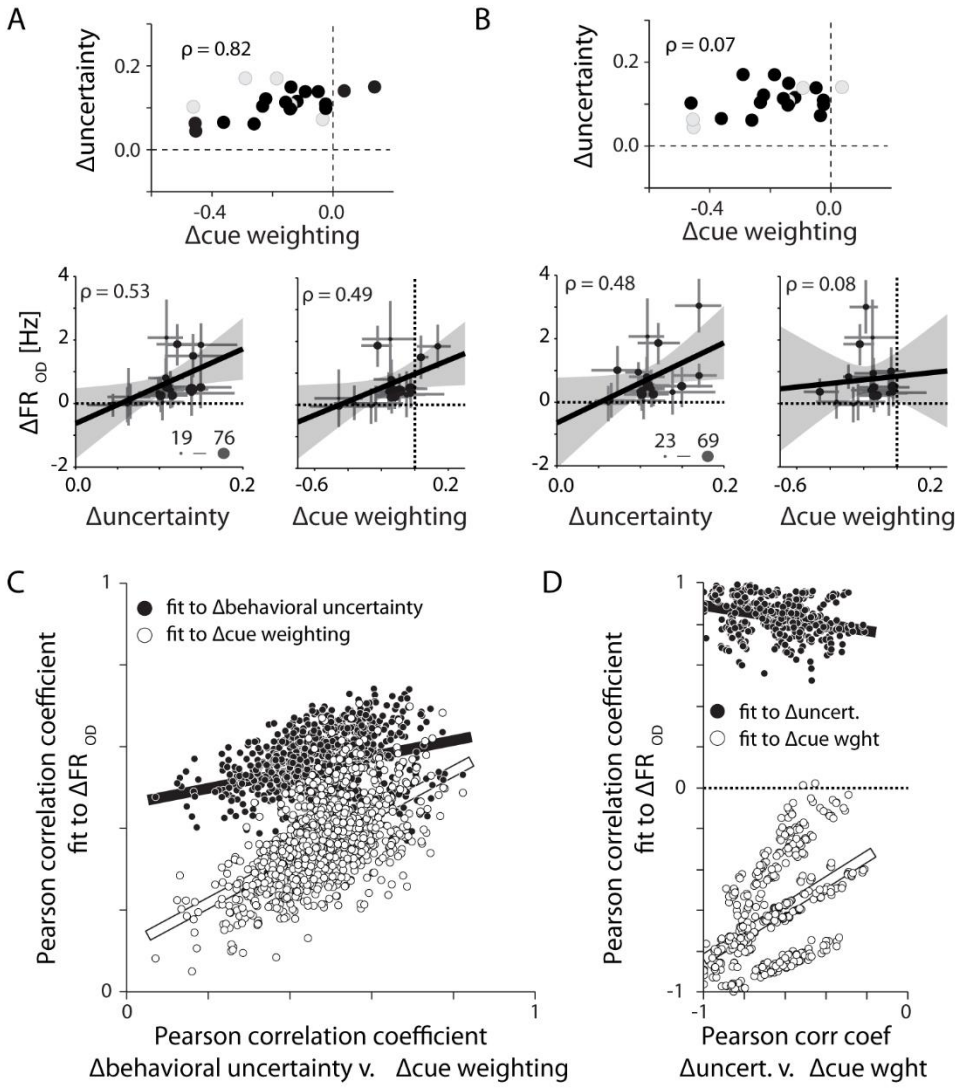

Figure 2.9 | Differences in PMd activity correlate with differences in behavioral uncertainty
(A) Eighteen sessions (filled symbols) selected for monkey M in order to increase the correlation between Δbehavioral uncertainty and Δcue weighting (top). Across these select sessions both metrics could explain the observed differences in OD activity (bottom). (B) Alternate subsampling that *minimized* the correlation between the two behavioral metrics (top). This resampling did not change the correlation between changes in OD activity and Δbehavioral uncertainty (lower left). However, it eliminated the correlation between Δcue weighting and OD activity (lower right). (C) Correlations of OD differences with Δbehavioral uncertainty (filled) and Δcue weighting (open) for 1000 unique 18-session subsamples. Each is plotted against the correlation between Δbehavioral uncertainty and Δcue weighting. The correlation with Δbehavioral uncertainty was consistently stronger than with Δcue weighting, while the correlation with Δcue weighting was only strong when Δcue weighting and Δbehavioral uncertainty were well correlated with each other. (D) Same as in C, but for monkey T. Each subsample contains six trial blocks. Unlike monkey M, Δcue weighting and Δbehavioral uncertainty were negatively correlated across sessions. Regardless, OD activity in PMd was still positively correlated with Δbehavioral uncertainty.

with poor correlation between the two metrics, only Δbehavioral uncertainty remained well correlated with OD activity (Fig. 2.9B). In fact, differences in OD activity correlated better with Δbehavioral uncertainty than with Δcue weighting for almost any randomly subsampled group of sessions for either monkey (Fig. 2.9C, D). This suggests that the firing rate changes in PMd actually reflect differences in the monkeys' overall decision-related uncertainty, not merely the uncertainty associated with the visual cue.

Another way of examining the evolution of target-related information in M1 and PMd is to use it to predict the monkey's choice of reach direction. For a representative session, Figure 2.10A (left) shows that, although it was possible to predict the monkey's reach direction from PMd activity, the predictions were consistently less accurate for high uncertainty trials than for low uncertainty trials. Accuracy rather rapidly reached these levels within about 200ms of target appearance, but then increased more slowly throughout the remainder of the trial. On the other hand, the ability to decode reach direction from M1 improved steadily through the delay period (Fig. 2.10A, right). This was true for both high and low uncertainty trials, with only slightly higher delay-period decoding accuracy for low uncertainty trials. At the time of movement initiation, the M1 decoder was equally accurate for both conditions.

Across all sessions, we observed results similar to the single session example. The PMd decoder nearly always performed better during low uncertainty trials than high uncertainty trials (Fig. 2.10B), especially during the visual and delay periods. PMd decoding generally did improve at the time of movement, however the difference in decoder performance between low and high uncertainty conditions remained significant. T-Tests on the performance difference between low and high uncertainty revealed significantly better low-uncertainty performance in

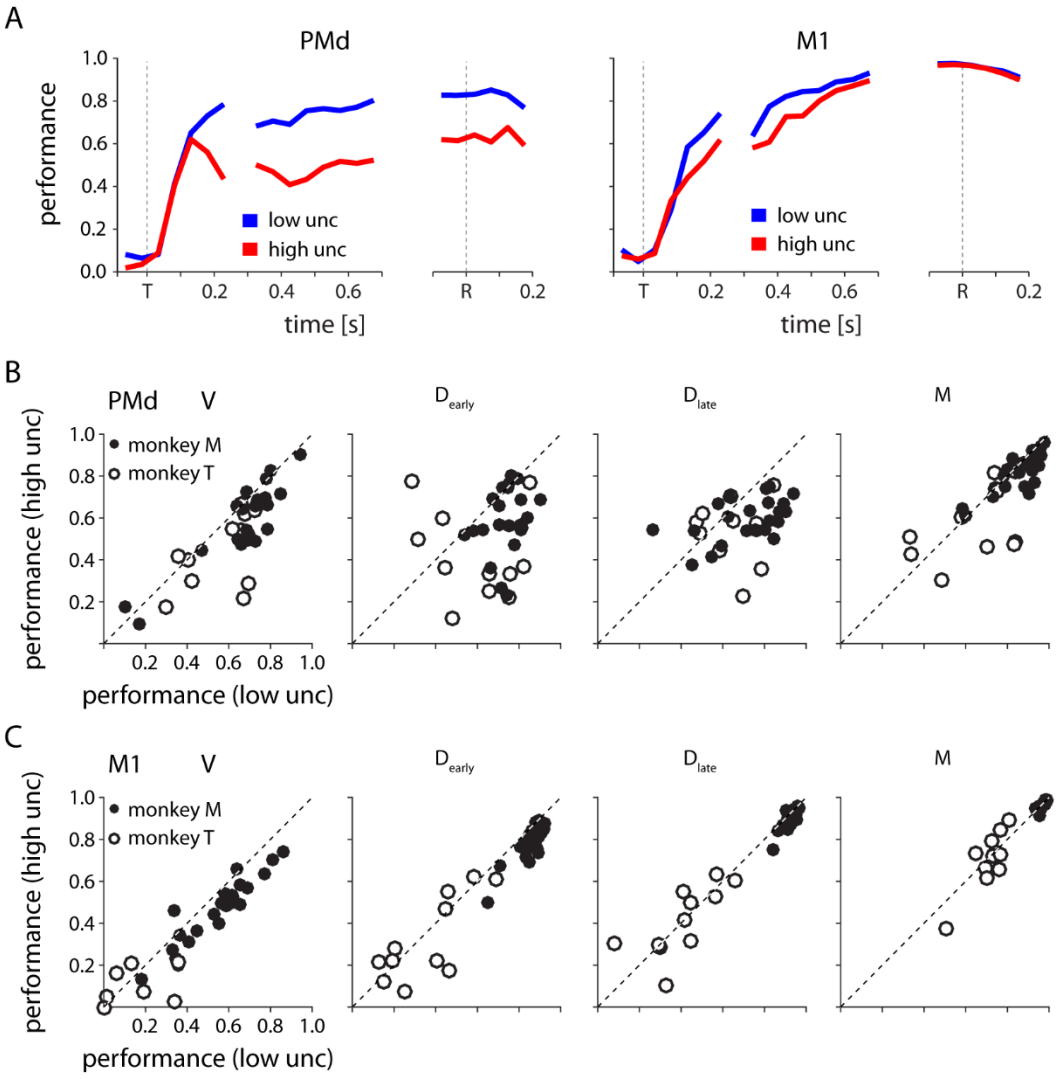

Figure 2.10 | Decoding reach direction from neural activity on single trials
Based on PD computed during center-out (zero-uncertainty) reaches. (A) The performance of PMd (left) and M1 (right) decoders as a function of time for one example session. Performance is defined as one minus the circular variance of the decoder error (B) PMd decoder performance in low v. high uncertainty conditions for four 200ms time windows spanning target appearance to movement in all sessions for both monkeys. Each point represents a single session from monkey M (closed) or monkey T (open) (C) Same as in B, but for M1.

all behavioral periods for monkey M ($p < 0.05$) and all except the movement time period for monkey T (discounting sessions with overall poor decoding, see Methods). In M1, decoding performance was also slightly better for low uncertainty trials during the visual and delay periods (t-Test, $p < 0.05$), although only for monkey M (Fig. 2.10C). This effect of uncertainty was much smaller than that observed in PMd. At the time of movement there was no bias in performance between low and high uncertainty trials. In general, we found decoding from M1 to be more accurate than from PMd, and less affected by uncertainty – especially at the time of movement.

Controls

We considered several alternate explanations for the effects of uncertainty, including differences in the visual stimuli, inhomogeneous distribution of the target prior over sessions, and variations in the kinematics of reaching.

To test for possible visual effects, we performed three control sessions with a single monkey (monkey M) in which half of the high-uncertainty trials contained an additional, different colored line segment at the correct target location (Fig. 2.11A). These sham trials had almost exactly the same visual properties as high uncertainty trials, but did not actually induce any uncertainty. The monkey learned to rely entirely on the new cue line (Fig. 2.11B). Comparing the difference in activity between actual high uncertainty and sham uncertainty trials, we found that OD (and to some extent ORTH) activity was greater under the actual uncertainty conditions (Fig. 2.11C). This suggests that our main finding of uncertainty-related changes in ORTH and OD activity cannot be explained simply as the result of differences in the visual information.

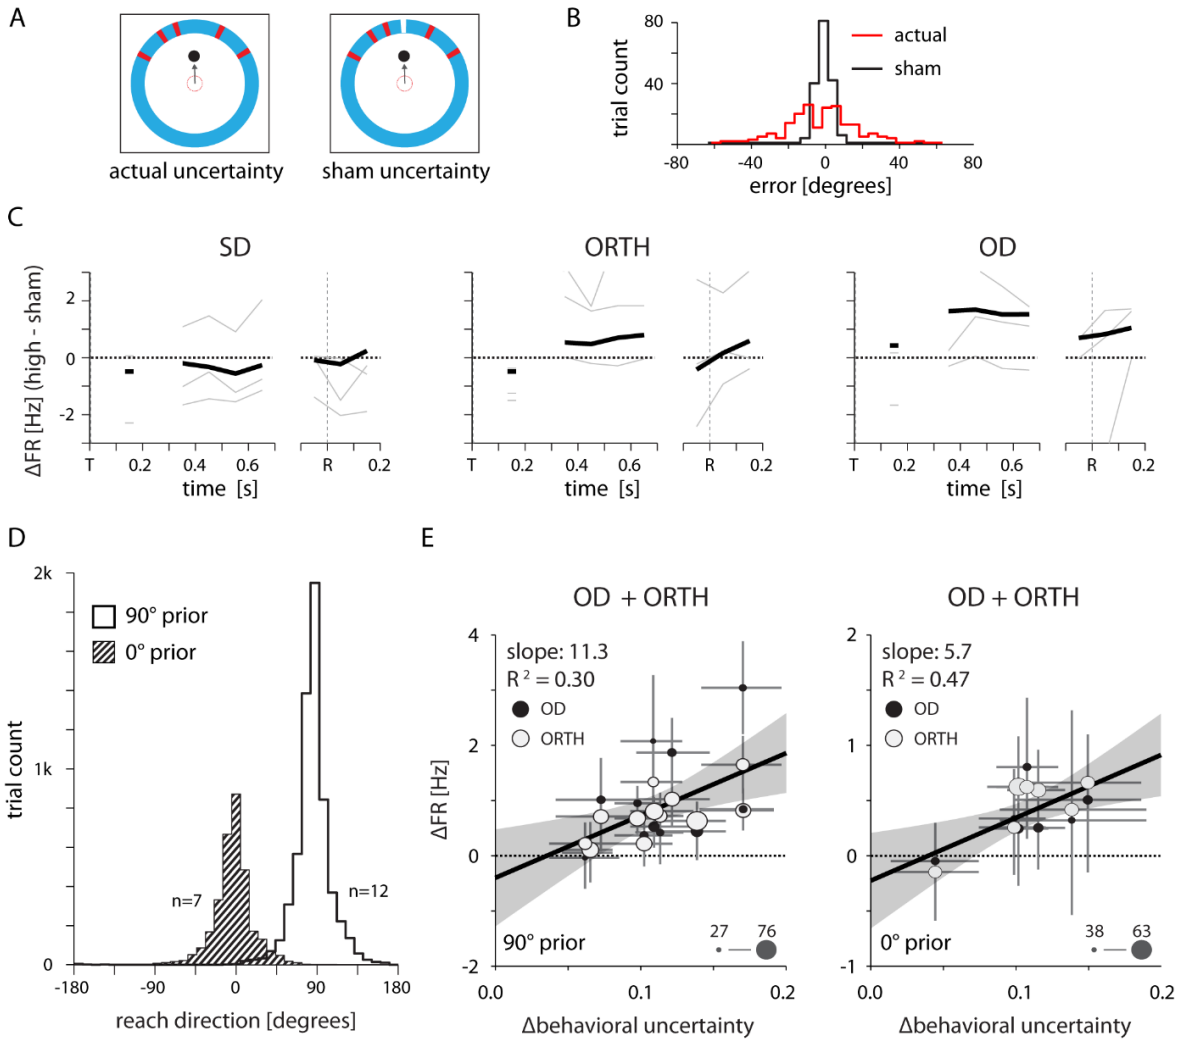

Figure 2.11 | Controls for visual properties and average reach directions
(A) Design of a control experiment to test whether the uncertainty-related effect could be explained solely by differences in the visual stimuli between conditions. Half of the trials contained a high-uncertainty cue (top left) and the other half contained sham high-uncertainty trials that included an additional line of a different color to indicate the veridical target location (top right). (B) Reaching errors were much smaller for the sham trials, indicating that the monkey learned to rely on the veridical cue. (C) Thin lines indicate the average difference in firing rate between actual and sham uncertainty trials for individual sessions. Heavy lines mark the mean across sessions. OD activity was higher during high uncertainty trials, despite the nearly equivalent visual properties (D) Control to test whether the neural effects could be explained by differences in the average target location across sessions. We selected two groups of sessions that each contained a consistent average reach direction. (E) Correlations between changes in OD and ORTH activity and Δbehavioral uncertainty for the two groups of sessions, 500-700 ms after target appearance. OD and ORTH activity within each group of sessions still correlated with Δbehavioral uncertainty.

We also considered the possibility that the effects on neural activity resulted from changes in the average target location (and subsequent reach direction) across sessions. We tested this possible explanation by analyzing only groups of sessions that shared a single average target direction. Figure 2.11D shows the distribution of reach directions for two groups of sessions for Monkey M in which the average target location was at either 0 or 90 degrees. Analyzing these two sets of sessions separately revealed a correlation between changes in OD/ORTH activity and Δbehavioral uncertainty (Fig. 2.11E) that was very similar to the full data set (Figs 7 and 8).

We anticipated that both the reaction time and peak speed might be affected by the target uncertainty, and might indirectly give rise to the firing rate changes we observed in PMd. In fact, these differences were rather small, but to test this possibility, we resampled the trials in a given session to reverse the sign of the uncertainty effect on either reaction time or peak speed (Figure S2.1). These manipulations had no effect on the correlation between PMd activity and Δbehavioral uncertainty, indicating that the difference was not simply driven by kinematics.

DISCUSSION

Summary

In this study, we set out to examine the neural effects within the motor system of uncertainty during a target estimation task. We showed that when visual cues of target location were made less informative, monkeys biased their reach direction toward the average target location that they had learned over the course of previous trials (their *prior* estimate) in a Bayesian-like manner. Activity in dorsal premotor cortex (PMd) changed systematically as a function of the resulting uncertainty in the monkeys' final estimate of target location, with higher

uncertainty leading to higher activity in PMd neurons. This effect was not present in primary motor cortex (M1). The extent to which uncertainty affected the activity of PMd neurons depended on their directional tuning properties. Neurons with preferred directions aligned to the ultimate reach direction showed no correlation with uncertainty, while those with orthogonal or opposite direction tuning displayed significant increases in activity with increased uncertainty. This can be interpreted as an increase in uncertainty causing in increase in the representation of less likely movements directions.

Representation of the process of target selection versus estimation

The uncertainty-related effect in PMd was present not only during movement planning, but also during execution – a result not readily predicted from previous studies. Several studies have recorded from PMd neurons as monkeys chose between multiple potential reach options (P. Cisek & Kalaska, 2005; Coallier et al., 2015; Klaes et al., 2011; Pastor-Bernier & Cisek, 2011; Thura & Cisek, 2014). Some even included ambiguous cues (Coallier et al., 2015; Thura & Cisek, 2014), which we might expect to induce uncertainty in the monkeys' decisions. The resulting representations of potential actions in PMd reflected in some sense the uncertainty in the choice prior to movement execution. However, in no studies before ours did this ambiguity in reach representation within PMd persist throughout movement execution. One study that used gradually accumulating evidence to trigger movement choice (Thura & Cisek, 2014) found that greater ambiguity in the evidence resulted in increased strength of the ultimately non-selected target representation prior to movement. About 300ms prior to movement, the activity corresponding to the selected target reached a peak that was consistent regardless of the level of

evidence. By the time of movement initiation, there was no evidence-related difference in activity for neurons tuned either to the selected or non-selected target. This is at odds with our finding of a persistent effect of uncertainty on the representation in PMd throughout the execution of movement.

That we did not observe a resolution in the reach representation prior to movement execution may reflect a difference in the decision-making processes associated with target *estimation* and target *selection*. Inherent to target selection is the knowledge that the correct action will only be one of several mutually exclusive options. That is information that can be integrated into the decision-making processes within sensorimotor areas like PMd. Since reaching anywhere that is not an explicit target will lead to failure, it is reasonable for the system to enforce a policy that the representation at the time of movement can only reflect one of the explicit target options. However, in target estimation tasks there are no such constraints on the executed action.

Differences in the roles of PMd and M1

This difference between target selection and estimation has important implications for the assumed roles of M1 and PMd within the sensorimotor system. Results from target selection tasks suggest that movement decisions result from the output of a biased neural competition between potential actions within PMd or related brain areas (P. Cisek, 2007; Gallivan et al., 2016; Pastor-Bernier & Cisek, 2011). This interpretation – that PMd ultimately decides on an action which is then transmitted to M1 – is especially convincing given the findings that decision-related variables (e.g.,difficulty, uncertainty, etc.) have no effect on PMd activity at the

time of movement initiation (Thura & Cisek, 2014). However, that an uncertain representation persisted throughout execution in our task may indicate that PMd is not necessarily the final step in the motor decision-making process. Instead, PMd may just reflect an updated estimate of the distribution of potentially useful actions, which is then translated into action by a downstream area such as M1. Perhaps the difference between a representation of remaining potential movements and a single selected action is less evident in target selection tasks.

Our results suggest that PMd contains a representation of something similar to the probability distribution of potentially useful reach directions. Furthermore, this representation appears relatively static, and does not resolve into a single unambiguous reach representation at any point in the delay or movement phases. M1, on the other hand, seems relatively unaffected by uncertainty and simply reflects the direction of the executed reach. These findings imply that PMd is not solely responsible for "deciding" which movement to execute, but instead contains only a noisy representation of potential reach directions that must be interpreted in some way by downstream areas like M1. Thus, we suggest that the processing that occurs in the connections between PMd and M1 "denoises" the PMd representation to provide a single, unambiguous movement decision.

The reach decoding results (Fig. 2.10) support the interpretation that PMd activity does not resolve into an unambiguous representation of reach direction prior to movement. At the time of movement initiation, decoding performance significantly improved from that of the delay-period levels, especially for high uncertainty trials. However, we still observed a bias towards better performance in decoding low uncertainty trials. M1, on the other hand, showed a steady increase in reach-related information leading up to movement initiation, that was slightly skewed towards better accuracy under low uncertainty. At the time of movement initiation, M1 was able

to decode reach direction with high accuracy, regardless of uncertainty condition. From these observations, we speculate that the decision about where to reach is not explicitly determined in PMd, but rather in the connections between PMd and M1. The noisy representation of potential actions in PMd appears to be pruned by M1, ultimately producing a single unambiguous motor command.

A downstream selection process could potentially occur through a *maximum a posteriori* (MAP) readout of the PMd representation. This kind of mechanism is not only consistent with the results of the current and previous studies, but could potentially explain the neural basis of sensorimotor learning. For example, we would expect that in very high uncertainty conditions (e.g., a novel behavioral task), PMd might contain nearly equal representations of all possible movements. As a consequence, small fluctuations due to noise within PMd would cause large variability in a downstream readout, driving exploration of the environment. As learning progressed and uncertainty decreased, the motor output would begin to converge on the optimal movement decision.

PMd reflects uncertainty in the decision, not the visual cue

Our task varied the monkeys' uncertainty in target estimation by manipulating both the history of target distribution and the noise in visual cues. We found that PMd activity changed not as a function of the weighting of those two pieces of information, but rather in proportion to the total uncertainty in the monkeys' final decision. This suggests that PMd contains uncertainty-related information pertaining to the final action decision that incorporates a range of information beyond that of the immediate target uncertainty. If uncertainty in visual information

were the sole driving force of changes in PMd planning- and execution-related activity, we would have observed very little difference in activity across sessions, since the visual cue properties were largely equivalent for all sessions. Instead, we found that activity modulated with the total behavioral uncertainty, which is a combination of visual uncertainty and prior expectation. This suggests that PMd likely reflects the combined uncertainty of all information sources relevant to a movement decision.

Comparison with existing theoretical models of uncertainty

There exist a number of theoretical models that address the potential neural representation of uncertainty (Deneve, 2008a; Hinton & Sejnowski, 1983; Hoyer & Hyvärinen, 2003; Ma et al., 2006; Zemel, Dayan, & Pouget, 1998). The predictions from these models encompass a wide range of neural behaviors, including temporal dynamics (Deneve, 2008a) and variability in spike timing (Deneve, 2008a; Hoyer & Hyvärinen, 2003). Unfortunately, our experimental design prevents us from performing fair and comprehensive tests of these model predictions. For example, our use of a static visual cue and instructed delay limits the potential interpretations regarding dynamic uncertainty codes. For these reasons, we hesitate to make any strong statements about the validity of any given models.

Despite the limitations of our experimental design, our results do bear some resemblance to admittedly simplistic interpretations of a few theoretical models. A probabilistic population code (PPC) model predicts that firing rates across a population should reflect the probability distribution – high uncertainty should therefore result in lower peak activity and higher non-peak activity (Ma et al., 2006). We did indeed observe an increase in non-peak activity with increased

uncertainty, and the spatiotemporal activity plots in Figure 2.5 do convincingly resemble probability distributions of reach direction. However, we did not see any consistent decrease in the peak activity with increasing uncertainty, which prevents us from interpreting the population activity as representing a true probability distribution. These findings also argue against the concept of divisive normalization, in which the total activity remains equivalent when representing multiple potential targets (P. Cisek & Kalaska, 2005; Pastor-Bernier & Cisek, 2011), at least in the context of target estimation.

Conclusions

Our results provide new insight into the behavior of PMd during movement planning. It is already well established that PMd can simultaneously represent all potential actions when faced with multiple, mutually exclusive visual targets (Bastian et al., 2003; P. Cisek & Kalaska, 2005). Our results provide the additional observation that PMd also represents and retains a distribution of potential motor plans that are not explicitly presented, but arise as possibilities during uncertain target estimation. The question of why this representation is maintained for the problem of target estimation but not target selection is an interesting one. One possibility is that it is simply an unavoidable result of the maintained, noisy inputs to PMd. That is, in the absence of explicit reach targets, the fidelity of the representation in PMd may be limited by the quality of available information. On the other hand, maintaining heightened representations of alternative movements in high uncertainty conditions be useful to the sensorimotor system for more rapid error correction or to drive subsequent motor learning. Experiments designed to test

these alternatives could help to further our understanding of the role of PMd in movement planning.

METHODS

All surgical and experimental procedures were fully consistent with the guide for the care and use of laboratory animals and approved by the institutional animal care and use committee of Northwestern University under protocol #IS00000367.

Behavioral task

The monkeys were seated in front of a vertical monitor and controlled an on-screen cursor using a planar robotic manipulandum. The behavioral task involved two or more blocks of trials. In the first block, monkeys performed a basic center-out reaching task with an instructed delay period. The monkey held the cursor within a central target for a random length center-hold period (700 – 1000 ms), after which a target (15 degrees wide) appeared in one of eight well-defined locations, distributed equally around an outer ring (Fig. 2.1A, top). Following an additional random delay period (700 – 1000 ms) the center target disappeared and the monkey received an auditory signal cueing him to reach to the outer target. Upon reaching the outer ring, the cursor froze. If the cursor was within the target, the monkey heard a success tone and received a small amount of juice. Otherwise, the monkey heard a failure tone and received no juice reward.

In the remaining (uncertainty) trial blocks, the target locations θ were not distributed uniformly among eight locations as before, but were instead selected randomly from a von Mises (circular normal) *prior*

$$f(\theta) = \frac{e^{\kappa \cos(\theta - \mu)}}{2\pi I_0(\kappa)} \quad \text{Eq. 2.1}$$

The mean of this prior distribution (μ) was always fixed for the duration of a session, but could vary in width (κ) across trial blocks. Additionally, during uncertainty trials the monkeys did not receive veridical visual cues about the target until the end of the trial. Instead, during planning and execution they were only shown several small lines (five for monkey M, ten for monkey T) sampled from a *likelihood* distribution (also von Mises) centered on the target location (Fig. 2.1A, bottom). These lines gave the monkey information about the target location, but with different levels of uncertainty depending on the variance of the distribution. Each session contained two different likelihood distributions, which were randomly interspersed across trials. The exact parameters used for each session are provided in Table 2.1. Upon reaching to the outer ring, the cursor froze and the ambiguous cue lines were replaced with the actual target (15 degrees, all conditions). The monkey subsequently received (or did not receive) reward as in the center out trial block.

Although we directly specified the variance (and therefore uncertainty) in the target distribution and the visual cue, the monkeys' subjective estimates of those parameters could deviate considerably from their true values. We therefore used the monkeys' actual responses throughout the session to estimate two values: the monkeys' weighting of the current visual cue, and the total uncertainty remaining in the monkeys' final estimate of the required reach direction. To do this, we assumed a Bayesian-like model of cue integration in which the final estimate was

the product of likelihood (visual cue) and prior (distribution of target locations) probability distributions. We modeled both of these as von Mises distributions. The product of two von Mises distributions can be approximated by a third, with mean

$$\mu_3 = \mu_1 + tan^{-1}\left(\frac{\sin(\mu_2 - \mu_1)}{\frac{k_1}{k_2} + \cos(\mu_2 - \mu_1)}\right) \quad \text{Eq. 2.2}$$

To obtain an estimate of the relative weighting of the visual cue for each uncertainty condition, we substituted the true target centroid location for μ_2, the true average target location for μ_1, and then fit $\left(\frac{k_1}{k_2}\right)$ to minimize the sum of the squared residuals between the model outputs and the monkeys' actual reach directions. The resulting equation for μ_3 describes the general function relating the centroid of the visual cue and the reach direction. (red and blue lines; Fig. 2.1B). Except for cases in which $|\mu_2 - \mu_1|$ is very large, this can be suitably approximated by the linear function

$$\mu_3 \cong \mu_1 + \frac{k_2}{k_1 + k_2}(\mu_2 - \mu_1) \quad \text{Eq. 2.3}$$

In all further analysis, we use the slope term $\left(\frac{k_2}{k_1 + k_2}\right)$ as a proxy for our estimate of the monkeys' relative weighting of the visual cue with respect to the summed prior and likelihood uncertainty. Slopes close to one represent high reliance on the visual cue, while slopes close to zero represent high reliance on the average prior target location.

The slope metric described above reveals only the monkeys' *relative* uncertainties in the likelihood and prior. It does not contain any information about the total magnitude of uncertainty present in the monkeys' decisions. We estimated this total uncertainty from the monkeys' behavior, by calculating the circular standard deviation of the residuals from each behavioral fit like those shown in Figure 2.1B. It is important to note that the behavioral uncertainty can be

affected by uncertainty in the estimate of average target location, uncertainty in the visual cue, and potentially other internal variables affecting the monkeys' behavior that we did not control (e.g.,motivation, attention).

Neural Recordings and Analysis

Throughout the experiments we recorded from neurons in M1 and PMd (Fig. 2.2A) using chronically implanted 96-channel microelectrode arrays (Blackrock Microsystems). Monkeys received appropriate pre- and post-operative antibiotics and analgesics. We identified single neurons from each session using offline sorter by isolating clusters within a principle component space projected from the waveform shapes of putative neurons (Plexon Inc., Dallas TX). There was likely significant overlap between sessions of the populations of recorded neurons, but we made no effort to track the identity of neurons across sessions. On each session, we used the activity from the center-out block of trials (zero uncertainty, eight target locations) to characterize the directional tuning characteristics of all neurons. Since many neurons (especially those in PMd) can have complex temporal profiles, we calculated preferred directions (PDs) in three distinct time periods: *visual* (50-250ms after target appearance), *delay* (300-700ms post-target), and *movement* (0-200ms after initiation of the reach) using a generalized linear model with Poisson noise:

$$\lambda = \exp[\alpha + \beta \cos(\theta - \theta^*)], \qquad \text{Eq. 2.4}$$

where λ is a vector of firing rates across trials, θ is a vector of reach directions, θ^* is the preferred direction, and $\alpha,\ \beta$ are scaling parameters. For each neuron, we also obtained

confidence bounds on the fit parameters through bootstrapping. A neuron was only considered to be significantly tuned if 95% of the bootstrapped estimates of θ^* were within forty-five degrees of the mean estimate. Due to the lower neuron count for monkey T, we relaxed this constraint to accept neurons with bootstrapped PDs within ninety degrees of the mean. For all analyses, we used only the preferred directions calculated within the appropriate time period (for example, delay-period tuning for all delay-period analyses). When analyzing a given time period, we excluded neurons without significant tuning in that period. Full details on the numbers of tuned neurons for all sessions is provided in Table 1.1.

Single Trial Decoding Analysis

We used a simple decoding approach based on each neuron's PD, computed from data collected during the center-out (zero-uncertainty) task. We first divided neurons according to their PDs, creating sixteen bins of 22.5 degrees each. We then averaged the activity of all neurons within each bin and fit a cosine to the resulting activity profile. The peak of this cosine defined the decoded reach direction. We characterized decoder performance for each uncertainty condition as one minus the circular variance of decoder error. Circular variance is bounded by 0 and 1. Therefore, a performance of 0 represents that the decoder did no better than random guessing, and a performance of 1 represents perfect decoding of the reach direction. This metric is similar to VAF, except it is not normalized by the total variance of the reach distribution. This is important in our dataset which contained very non-uniform distributions of reach directions. It provides a fair comparison of decoder performance regardless of differences in distributions between sessions or uncertainty conditions.

We assessed the effect of uncertainty condition on decoding performance by performing t-tests on the distributions of differences between low and high conditions for each monkey and time period. For monkey T, low neuron counts made decoding on a trial-by-trial basis much less accurate. Therefore, when assessing biases, we only included sessions in which the decoder performance on low uncertainty trials was greater than 0.5.

FIGURE SUPPLEMENTS

Table 2.1 | Experimental details for all sessions

In some instances we obtained multiple sessions from the same day (sessions 3-4, 5-7, 8-10, 11-12, 13-14, 16-17, and 26-27). In these cases, the sessions shared the same sorted neurons and center out trials. Uncertain trial blocks could differ in either target distribution or visual cue properties.

		center out	*uncertain target distribution*		*low uncertainty visual cue*		*high uncertainty visual cue*		*number of tuned neurons (PMd)*			*number of tuned neurons (M1)*		
session	monkey	# trials	mean	kappa	k	# trials	k	# trials	V	D	M	V	D	M
1	M	326	90	25	50	337	5	330	95	104	117	110	174	188
2	M	199	90	5	50	316	5	323	53	93	99	80	148	167
3	M	294	90	5	50	209	5	220	55	97	93	43	123	134
4				50	50	149	5	131						
5	M	223	90	5	50	196	5	189	105	125	104	129	177	190
6				50	50	236	5	243						
7				5	50	98	5	89						
8	M	283	0	5	50	244	5	230	84	97	88	136	183	190
9				50	50	197	5	224						
10				5	50	71	5	75						
11	M	290	90	50	50	312	5	312	115	128	100	132	184	194
12				5	50	373	5	403						
13	M	309	90	5	50	279	5	256	153	160	144	129	181	185
14				50	50	226	5	225						
15	M	275	90	25	50	228	5	253	104	113	98	124	186	188
16	M	339	180	5	50	184	5	158	115	111	104	129	163	180
17				50	50	171	5	145						
18	M	416	0	10	20	342	4	303	115	97	97	139	178	181
19	M	350	0	10	20	291	4	362	138	154	133	116	160	169
20	M	256	0	10	20	261	4	247	106	100	96	107	164	160
21	M	238	0	25	50	114	5	100	111	121	116	49	68	69
22	M	164	45	25	50	157	5	151	84	97	73	72	103	105
23	M	190	45	25	50	161	5	159	82	96	76	67	95	107
24	T	174	90	150	100	285	1	128	47	66	63	11	26	35
25	T	178	0	150	100	142	1	107	47	59	51	12	12	30
26	T	198	110	150	100	86	10	74	48	66	60	5	15	25
27				150	100	86	1	76						
28	T	357	110	5	50	162	10	179	28	30	25	6	16	20
Visual controls														
					sham cue									
VC 1	M	206	90	5	--	168	5	174	19	33	24	--	--	--
VC 2	M	251	90	5	--	205	5	228	7	17	13	--	--	--
VC 3	m	260	90	5	--	155	5	141	5	13	11	--	--	--

Table 2.2 | Subsampling of sessions for monkey T
Due to low sample size for monkey T, we subdivided larger sessions to obtain separate blocks of 100+ trials each. Here we show the trials contributing to each trial block and the subsequent numbers of low- and high-uncertainty trials.

		low uncertainty visual cue	*high uncertainty visual cue*
session	Trials contributing to sub-block	# trials	# trials
24a	1-100	69	31
24b	101-200	64	36
24c	201-300	72	28
24d	301-413	80	33
25a	1-100	59	41
25b	101-249	83	66
26	No subsampling	86	74
27	No subsampling	86	76
28a	1-100	44	56
28b	101-200	52	48
28c	201-341	66	75

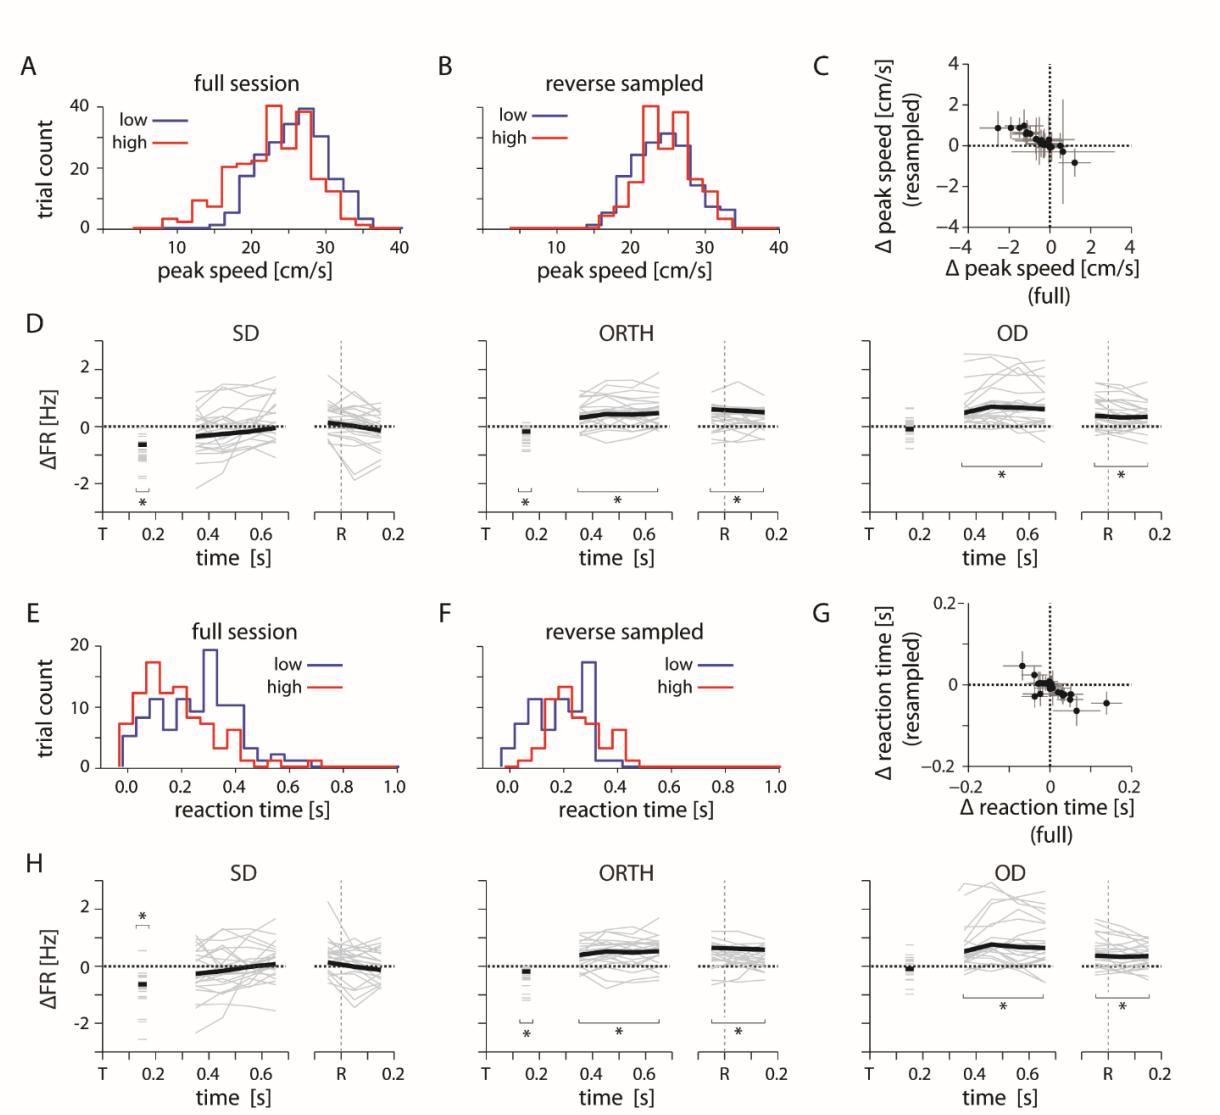

Figure S 2.1 | Kinematic controls
(A) Example distribution of peak reach speed for high and low uncertainty conditions. (B) Distribution of peak speeds for same session as in *A*, subsampled to reverse the condition-dependent difference. (C) Plot showing the difference in average peak speed between low and high uncertainty conditions for the full and reverse-sampled datasets. (D) Thin lines indicate the average difference in firing rate between high and low uncertainty conditions for an individual session, subsampled to reverse the trend of peak speed. Heavy lines mark the mean across sessions. (E) Example distribution of reaction time for high and low uncertainty conditions. (F) Distribution of reaction times for same session as *F*, subsampled to reverse the condition-dependent difference. (G) Plot showing the difference in average reaction time between low and high uncertainty conditions for the full and reverse-sampled datasets. (H) Same as in *D*. Average differences in peak speed and reaction time cannot explain the neural effect.

CHAPTER 3: SINGLE REACH PLANS IN DORSAL PREMOTOR CORTEX

Brian M Dekleva, Konrad P Kording, Lee E Miller

FOREWORD

This chapter contains a version of a paper recently submitted for publication. In it, I address the question of how PMd responds when faced with multiple potential reach targets. This extends the work in Chapter 2, which focused on planning related to a single—albeit noisy—goal. As discussed in Chapter 1, prominent hypotheses of motor decision making are centered on the concept of simultaneous representations of multiple options. These models are supported by physiological studies in motor areas, but the evidence has come only through trial-averaged, single-neuron recordings. Here I use a large population of simultaneously recorded neurons in PMd to search for simultaneous reach plans on individual trials. I show that in fact, the evidence for simultaneous planning is suspect, and can be explained by specific behavioral artifacts. This surprising result has significant implications for our interpretation of decision-making in PMd, which will be discussed in Chapter 5.

ABSTRACT

In many situations, we are faced with multiple potential actions, but must wait for more information before knowing which to perform. Movement scientists have long asked whether in these delayed-response situations the brain plans both potential movements simultaneously, or if it simply chooses one and then switches later if necessary. To answer this question, we used simultaneously-recorded activity from populations of neurons in macaque dorsal premotor cortex to track moment-by-moment deliberation between two potential reach targets. We found that the neural activity only ever indicated a single reach plan (with some targets favored more than others), and that initial plans often predicted the monkeys' behavior on both free-choice trials and incorrect trials. Our results suggest that cortex plans only one option at a time, and that decisions are strongly influenced by the initial response to the available set of movement options

INTRODUCTION

We often prepare for a movement by surveying the possible actions available and then waiting for more information before deciding. For example, a tennis player waiting to receive a serve can anticipate that he will need to perform a forehand or backhand return, but needs to first observe the trajectory of the ball before deciding between those two options. Likewise, we can plan to grasp an object without yet knowing which hand posture will be most useful. Similar parallels can be drawn for almost all types of movements, indicating that initial deliberation between potential actions is a ubiquitous aspect of motor control.

Although the initial deliberation between multiple potential actions is a widespread phenomenon, the neural processes underlying it are largely unknown. The brain could respond in

two ways: (1) by simultaneously representing several potential options, or (2) by initially representing only one and then later switching if necessary. Several behavioral studies have attempted to disentangle these two possibilities, most notably through the go-before-you-know paradigm in which a subject is given multiple reach targets and is then forced to move before knowing which is correct. Early movement trajectories are often directed in between the two presented options, which some have interpreted as a spatial averaging of two simultaneous plans (Chapman et al., 2010; Gallivan et al., 2016; Gallivan et al., 2017; Ghez et al., 1997; Stewart et al., 2013; Wood et al., 2011). However, an intermediate movement is not necessarily indicative of multiple simultaneous plans, and might instead reflect a single plan optimized for the task (Haith et al., 2015; Nashed et al., 2017; A. L. Wong & Haith, 2017). Thus, the neural processes that underlie the deliberation between potential movements cannot be readily interpreted from the movements themselves.

A few studies have directly recorded from individual neurons in motor cortex to probe planning-related activity in the face of multiple discrete movement options. Invariably, the results from these studies have supported a "simultaneous representation" hypothesis (P. Cisek & Kalaska, 2005; Coallier et al., 2015; Klaes et al., 2011; Pastor-Bernier & Cisek, 2011; Platt & Glimcher, 1997). However, these studies rely on single electrode recordings and trial-averaged data, with the implicit assumption that all trials reflect a single consistent process. Since delay-period planning activity has no measurable external signature, it is unclear how well this assumption actually holds. To accurately examine the wholly internal neural processes at play during movement deliberation requires analysis on the timescale of individual trials.

An alternative to averaging across trials is to combine the information obtained from many simultaneously recorded neurons. This population-based approach has increasingly been

adopted in motor areas, where it provides a reliable estimate of limb movement (Afshar et al., 2011; Carmena et al., 2003; Mante, Sussillo, Shenoy, & Newsome, 2013; Michaels et al., 2015; Pohlmeyer et al., 2007; Serruya et al., 2002; Taylor et al., 2002; Wessberg et al., 2000). An important, less exploited advantage of simultaneous recordings over trial-averaging is the ability to probe neural processes that have no measurable behavioral outcome. Multiple studies have used population recordings from cortex to identify changes of mind on single trials of a multiple potential target reaching task (Kaufman et al., 2015; Kiani, Cueva, Reppas, & Newsome, 2014). This ability to interpret activity on a short timescale in the absence of behavioral correlates thus provides a means to investigate deliberation between movement options.

Here we used simultaneous recordings from dorsal premotor cortex (PMd) of the macaque monkey to monitor the development of movement plans in the face of two potential movement options. Using dimensionality reduction methods on the population activity, we tracked the moment-by-moment planning in these areas over the course of single trials. We found that when presented with two options, PMd quickly developed a movement plan for one of them. This initial plan was strongly predictive of the monkey's eventual behavior (including freely chosen movements, reaction times, and task errors). Our results show no evidence of simultaneous representation in premotor cortex during the deliberation between multiple potential options, and indicate that ultimate movement decisions are strongly influenced by initial responses to the possible choices.

RESULTS

We trained two rhesus macaques on a center-out reaching task in which they controlled an on-screen cursor using a planar manipulandum (Fig 3.1a). Each session consisted of both 1-target and 2-target trials, randomly interspersed. On 2-target trials the monkey first positioned the cursor in the central target, after which two targets appeared spaced 180 degrees apart (Fig 3.1c, left). The targets remained on screen for between 750 and 1000 milliseconds (Target On), and then disappeared for 250-500 milliseconds (Target Blank) before reappearing for another 250-500 milliseconds (Cue). Finally, the central target disappeared and a tone cued the monkeys to reach to the target (Go). 1-target trials followed the same basic structure, except that the monkey was only shown one outer target throughout the trial (Fig 3.1d, left). On roughly ten percent of 2-target trials ("free-choice" trials) we omitted the Cue epoch, and instead instructed the monkey to reach immediately after the Target Blank epoch (Fig 3.1d, bottom). In all conditions, the monkeys made approximately straight reaches to the targets (Fig 3.1c,d, right).

The information from a single neuron is too noisy to provide reliable decoding of single-trial intent. Hence, our goal was to take advantage of the simultaneity of our recordings made from chronically implanted 96-electrode arrays (Blackrock microsystems) to obtain an estimate of the time-evolution of the recorded population's low-dimensional "neural state" (Afshar et al., 2011; Ames, Ryu, & Shenoy, 2014; Elsayed et al., 2016; Kao et al., 2015; Kaufman, Churchland, Ryu, & Shenoy, 2014; Kaufman et al., 2015; Michaels et al., 2015). For this dimensionality reduction approach to be valid, the correlations between neurons must remain fixed. However, if multiple targets are represented simultaneously in cortex, this may not be true. Consider three neurons with 1-target tuning functions such that neurons one and two are highly correlated with

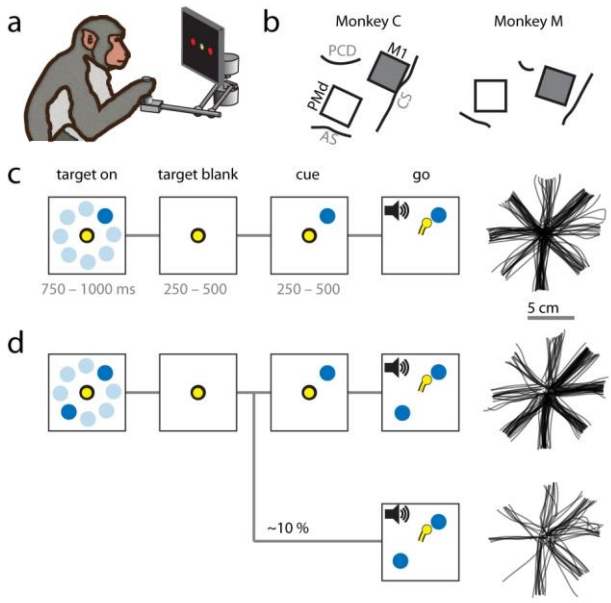

Figure 3.1 | Experimental setup
(a) Monkeys used a planar manipulandum to control a cursor. (b) Array placement in dorsal premotor (PMd) and primary motor (M1) cortex (CS = central sulcus, AS = Arcuate sulcus, PCD = precentral dimple). (c) Task events during single-target trials and resulting cursor trajectories for a single dataset. The monkeys reached to one of eight possible targets following an instructed delay. (d) Task events during two-target trials and resulting cursor trajectories. The monkeys were initially presented with two opposing targets. On the majority of trials, the correct target was displayed during the Cue epoch; roughly 10% did not contain a Cue epoch. Instead, the monkeys were forced to select one of the two targets with no information about which would lead to reward ("free-choice" trials).

each other, but anti-correlated with neuron three (Fig 3.2a). If the population contained only one movement plan at a time (the stay-or-switch hypothesis), all correlations would remain consistent (Fig 3.2c). If instead the population planned both targets simultaneously, then all neurons would display the same response to a given pair of targets, regardless of the eventual reach direction. This behavior would cause some neuron pairs (e.g., neurons one and three) to be negatively correlated on 1-target trials, but positively correlated on 2-target trials (Fig 3.2b). Such nonstationarity would invalidate the dimensionality reduction-based approach.

To examine the question of whether neural correlations remained fixed, we estimated the covariance between all pairs of neurons during the Target Blank epoch of both 1- and 2-target trials (single session example shown in Figure 3.2d). Over all sessions, we found that pairwise covariances did not change systematically between trial types (>90% of neuron pairs maintained the same sign; Fig 3.2e). These results most closely matched simulated results from the postulated stay-or-switch encoding hypothesis, not simultaneous encoding. This suggested that a dimensionality reduction approach would be valid to examine single-trial neural encoding.

We used principal components analysis (PCA) to reduce the population activity to a denoised, 10-dimensional "neural state". We then used the neural state to calculate at each point in time the likelihood that the monkey was planning a reach to each of the eight possible targets (Fig 3.3a). Figure 3.3b shows the time course of the likelihoods for all eight directions throughout an example 1-target (rightward) trial (corresponding to the neural state trace shown in Figure 3.3a). As expected, the greatest likelihood, evident shortly after target appearance, was for a rightward reach. This representation persisted until the end of the trial. To display single-trial results more compactly, we will from here on present only the difference in likelihood (Δlikelihood) between the target direction and the anti-target direction (Figure 3.3c). This

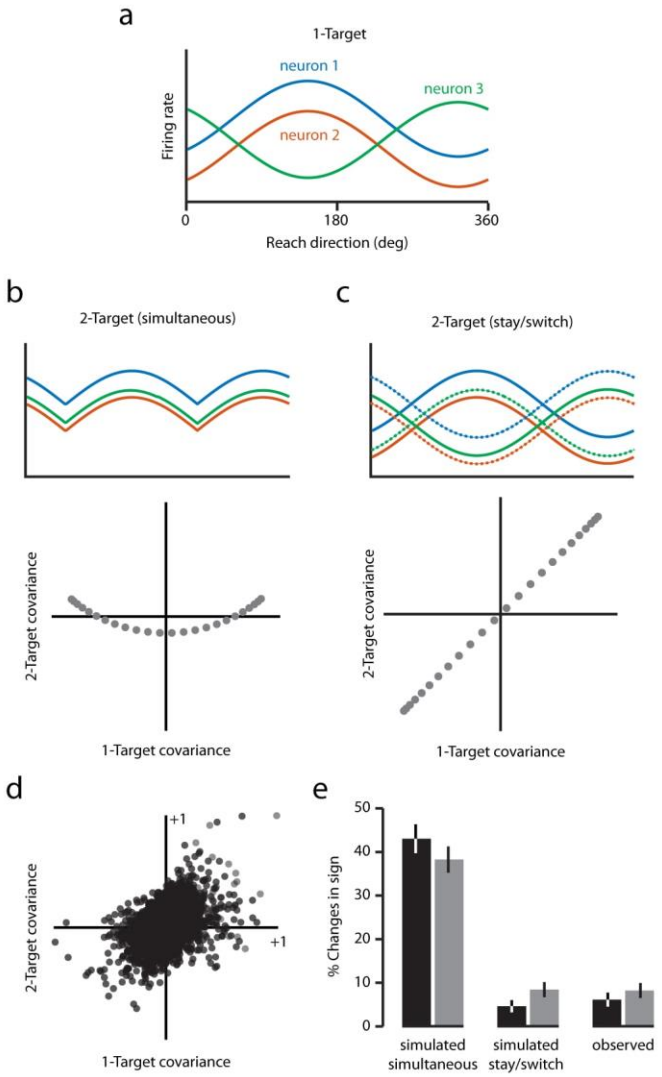

Figure 3.2 | Hypothesized neural responses to multiple simultaneously encoded targets

(a) Activity of three hypothetical neurons as a function of reach direction during single-target trials. All three are cosine tuned, but neurons 1 and 2 are most active for the same reach direction (about 140 degrees), while neuron 3 is most active for a reach in the opposite direction. (b) top: Hypothesized tuning of the neurons from *a* during two-target trials, given a simultaneous encoding scheme, in which all three neurons should have the same tuning form. bottom: Plot of the pair-wise covariances during 1-target vs. those during 2-target trials for 100 simulated "simultaneous encoding" neurons. (c) Same as in b, but under the stay-or-switch hypothesis in which only one of the two possible reach directions is encoded at a time. (d) Pairwise covariances for 1- and 2-target trials from a single session. In general, there was no systematic change in covariance between the two conditions. (e) Percentage of all neuron pairs with significantly non-zero covariances that differed in sign between single-target and two-target trials. For both monkeys (black – monkey C, grey – monkey M), the observed data closely resemble the simulated stay-or-switch response.

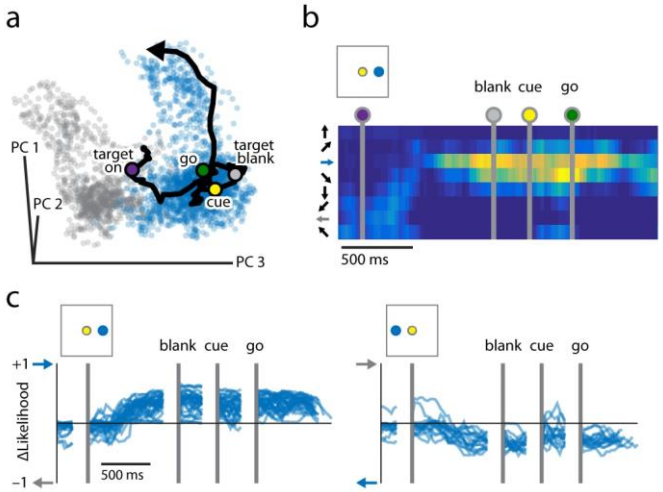

Figure 3.3 | Reach representations in PMd on single-target trials

(a) Low-dimensional representation of the neural activity throughout a single trial (black). Blue and grey clusters reflect all neural states that occurred during reaches to rightward and leftward targets, respectively. (b) The likelihood that the neural state reflected a reach to each of the eight directions calculated throughout the course of a trial. (c) Differences in likelihood between the displayed target direction and the opposite direction for two target directions. Each trace reflects a single trial.

approach proved capable of accurately decoding the target direction from planning activity (Target Blank period) on 1-target trials (97% - monkey C, 98% - monkey M, cross-validated).

The results from the covariance analysis (Fig 3.2) suggested that even on 2-target trials, the monkeys planned only one reach at a time. To track this evolving reach plan, we decoded the neural activity on 2-target trials in the same way as we did for 1-target trials. Figure 3.4a shows single-trial target decodes from all left/right 2-target trials in a single session (Fig 3.4a). As expected, reaches to the left were clearly distinguishable from those to the right in the neural state at the time of movement execution (after Go cue). However, activity early in the trials overwhelmingly resembled rightward reach plans (88%). Both monkeys had a preference for one target over the other for most target pairs (Fig 3.4b). However, an individual monkey's preferences were not consistent across sessions, nor were they governed by a clear spatial bias. Monkey M, for example, preferred upward and rightward reach plans on the first session, but downward and leftward plans on the third session. Monkey C seemed to treat each pair independently, with no apparent patterns either within or across sessions.

To validate the observed target preferences during initial planning, we compared each monkey's decoded preferences to their actual reaches on free-choice trials. Free-choice did not contain a Cue epoch, so the monkeys were forced to pick one of the two targets without any information about which was correct. Over all sixteen left/right free-choice trials from the example session in Figure 3.4, the monkey reached to the rightward target eleven times (69%, Fig 3.5a). This bias mirrors the rightward preference that we observed in initial plans on cued 2-target trials. Across all target pairs, we found a strong correspondence between the target preferences decoded from cued 2-target neural activity and the actual reach preferences

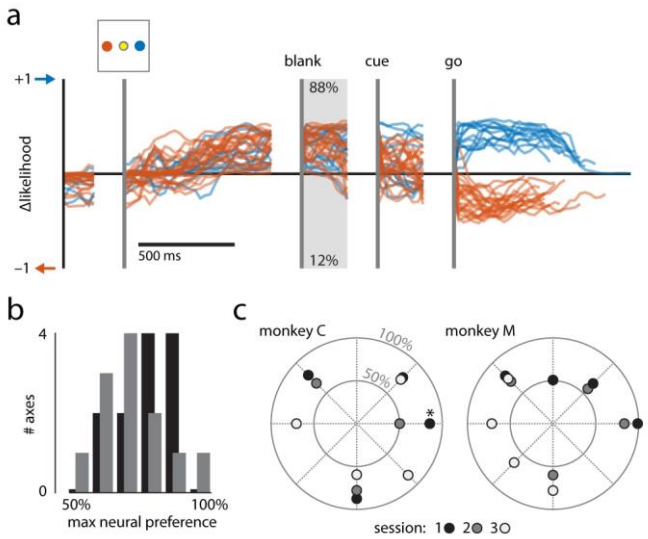

Figure 3.4 | Preferential reach representations in PMd during two-target trials

(a) The difference in likelihoods (as in Figure 1c) for all left/right trials in a single session. Blue traces indicate those for which the monkey was cued to the right. Red traces indicate a cue to the right. In the epoch prior to cue appearance, 88% of activity indicated a rightward reach plan (b) Distribution of neural preferences across target axes from all sessions (black – monkey C, grey – monkey M). (c) Neural preferences were not stable across sessions. Each point reflects the neural preference for the given axis on a particular session (example from a is indicated with an asterisk).

displayed on free-choice trials (Fig 3.5b). This correlation with behavior indicates that our decoding of the neural state accurately reflected the monkey's plan, even early in the trial.

While the monkeys tended to display a strong preference for one target in each pair (~75:25% on average; Fig 3.4b), the preference was almost never absolute (100:0%). Thus at times, the monkeys began planning to move to a target that they generally did not favor. While the results from Figure 3.5b show that initial plans were largely predictive of free-choice responses, it was unclear to what extent that correlation depended on the favorability of the initially planned target. Were initial plans made to favored targets more likely to be carried through to execution than those to non-favored targets? To answer this question, we first characterized the favored target for each target pair by identifying the most common reach plan (from activity in the Target Blank period) across trials. On each free-choice trial, we then decoded the initial reach plan and determined whether it matched the eventual reach direction. We found no difference between initial plans to favored and non-favored targets; both were equally predictive of the reach direction (Fig 3.5c). This suggests that while there appears to be a general long-term preference for some targets over others, decisions made on individual trials are largely determined by the initial plan.

While in general our findings suggest the existence of only single-reach plans, the strengths of those plans varied widely across trials. Consider the left/right responses in Figure 3.4a. While most early trial activity indicated rightward reach plans, we also observed trials with leftward (negative Δlikelihood) or weak-to-nonexistent plans (Δlikelihood near zero). Since we define plan strength as the difference between two likelihood calculations, a weak plan could in theory result from two strong, simultaneous plans. However, our results from Figure 3.2 argue against this possibility. We predicted that the magnitude of the reach plan on individual trials

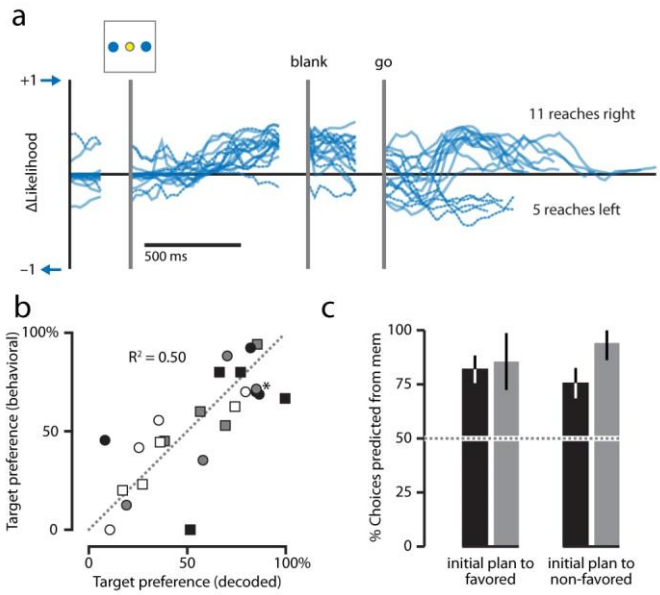

Figure 3.5 | Neural preferences match behavioral choice preferences

(a) Δlikelihood plots for all left/right free-choice trials in the example session from Figure 4a. The monkey chose the rightward target on eleven of the sixteen trials (b) Across all sessions, the monkeys' decisions on guess trials for each axis correlated with the calculated neural preference. Circles correspond to monkey C, squares to monkey M. White, grey, and black fills correspond to the different sessions, in order (d) Bar plot indicating the degree to which Target Blank period activity successfully predicted the monkeys' choices on guess trials for each session and target axis. Black bars correspond to monkey C, grey to monkey M.

would correlate with some aspect of the kinematics of the executed movement. To test this, we calculated the magnitude of Δlikelihood in a 100ms window preceding the Go cue on each two-target trial and compared it to the subsequent reaction time. Figure 3.6a shows this relationship for all leftward reaches on an example session. Reaction times were markedly shorter when Δlikelihood indicated a strong leftward plan at the time of the Go cue, and longer when it indicated a rightward plan. This negative correlation between decoded plan strength and reaction time occurred for nearly all reach directions for both monkeys (Fig 3.6b). A linear mixed effects model confirmed the negative relationship after accounting for differences across reach directions and sessions (monkey C: coeff=–165, $p = 0.00011$; monkey M: coeff=–174, $p \approx 0$). We found a similar negative correlation between Δlikelihood and reaction time on free-choice trials (Fig 3.6c,d; monkey C: coeff=–238, $p = 0.00013$; monkey M: coeff=–126, $p=0.0046$). These results indicate that switching a reach plan or developing a new plan delayed the initiation of movement.

Although both monkeys understood the two-target task, at times they chose the incorrect target (monkey C: 16%, monkey M: 27%). Using the same decoding approach as before, we examined the progression of reach plans during these error trials to understand the source of incorrect movement choices. Using the reach plans decoded early (Target Blank) and late (Cue) in the trial, we characterized three types of errors that encompassed over 97% of all errors. During type 1 error trials, the monkey maintained a consistent reach plan throughout, and did not deviate from it even after receiving a cue to the opposite target (Fig 3.7a). During type 2 errors, the monkey switched to the correct plan after cue presentation, but then reverted to the initial, incorrect plan upon receiving the Go cue (Fig 3.7b). Finally, type 3 errors represent last-second switches to the incorrect target, despite having planned the correct reach throughout the trial (Fig

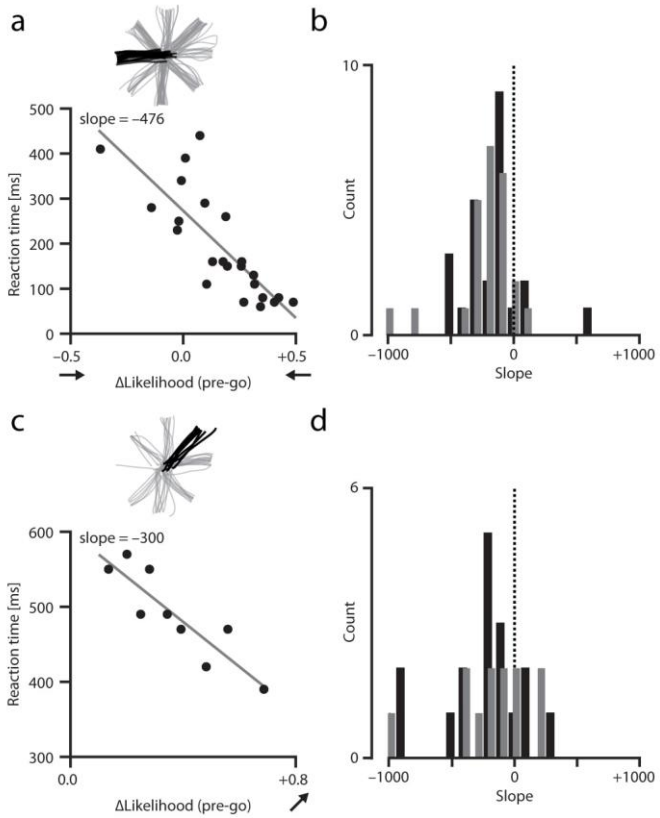

Figure 3.6 | Effect of pre-Go representations on reaction time during two-target trials

(a) Reaction time as a function of the strength of the reach representation in PMd calculated 100ms prior to the Go cue for all leftward reaches on two-target trials during a single session. Reaction times decreased as the calculated strength of the representation increased. (b) Histogram of the slopes (as calculated in a) for all target axes from all sessions. (monkey C – black, monkey M – grey). (c) As in a for all reaches to the upper left target during free-choice trials. Reaction time decreased with the strength of the decoded representation. (d) as in b for free-choice trials.

3.7c). The two monkeys were not equally predisposed to each error type (Fig 3.7d). Monkey C made fewer errors over all, and committed types 1 and 2 with about equal frequency. Monkey M made more errors that were predominantly type 1. However, although monkey M appeared more likely to ignore the cue altogether than did monkey C, the high incidence of type 1 and type 2 errors suggests that errors overwhelmingly arose due to overconfidence in the initial plan. This was true for both initial plans made to favored targets and those to non-favored target (Fig 3.7d).

DISCUSSION

We found that when faced with two potential reach targets, monkeys formed only a single motor plan. This finding was unexpected, given that several previous studies support a simultaneous encoding mechanism during movement planning (Christopoulos, Bonaiuto, & Andersen, 2015; P. Cisek, 2007; P. Cisek & Kalaska, 2005; Coallier et al., 2015; Gallivan et al., 2016; Gallivan et al., 2017; Klaes et al., 2011; McKinstry, Dale, & Spivey, 2008; Pastor-Bernier & Cisek, 2011). It may be that premotor cortex is incapable of planning multiple movements simultaneously, but that the trial averaging approaches used by previous studies were unable to differentiate between simultaneous encoding and single encoding with strong planning preferences. When we replicated the trial-averaged guess-and-switch control analysis from Cisek and Kalaska (P. Cisek & Kalaska, 2005), our results initially appeared to support the simultaneous encoding interpretation (Fig 3.S1). However, additional simulations showed that this was a spurious result caused by strong target preferences (Fig 3.S2). A stricter version of the test—modified to remove the effect of target preferences—revealed no evidence of simultaneous encoding (Fig 3.S1).

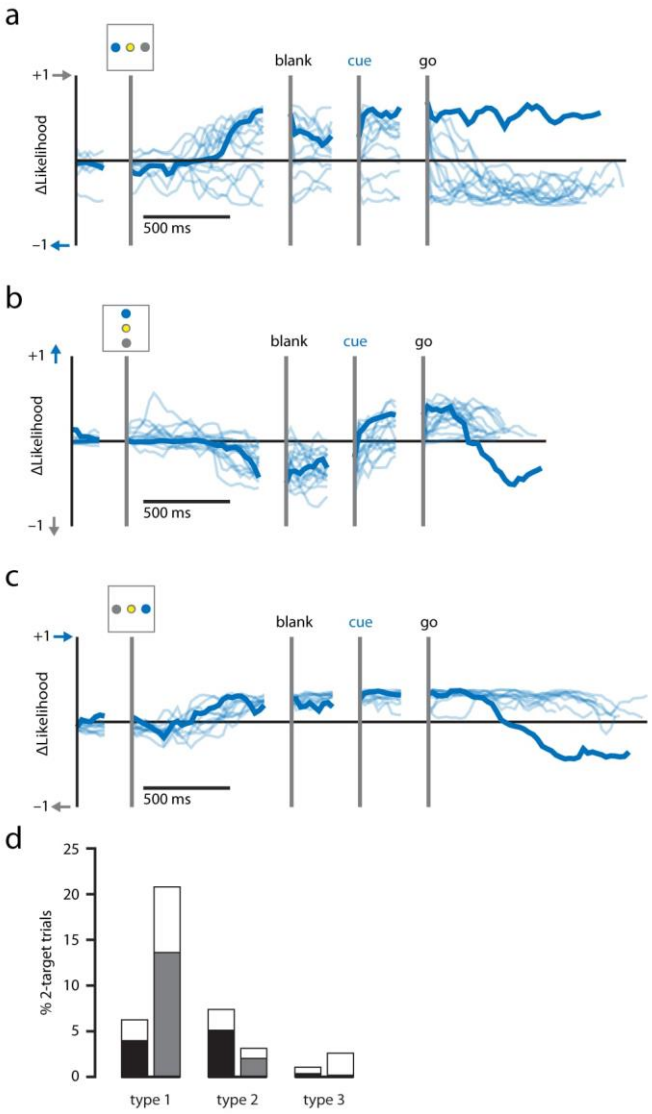

Figure 3.7 | Heterogeneity in types of errors
(a) An example error (dark) made to the rightward target when cued to the left (monkey M, session 2). This type of error can be characterized as an early plan that is not changed even after receiving a contradictory cue (error type 1). Light traces indicate all correct trials with the same leftward cue (b) An incorrect reach to the bottom target (monkey C, session 1). This type of error can be characterized as a late reversion to the initial plan (type 2). (c) An incorrect reach the leftward target (Monkey M, session 1). This error represents a late switch away from the correct target plan (type 3). (d) Histogram showing the frequency of each error type for both monkeys (monkey C – black, monkey M – grey). Filled and open segments indicate whether the errors were made to a favored or non-favored target, respectively.

It is possible that the monkeys in previous studies did not have target preferences during early planning. However, choice biases appear to be quite common in two target tasks (Kaufman et al., 2015; Klaes, Schneegans, Schoner, & Gail, 2012; Klaes et al., 2011). Some studies have attempted to minimize this effect by explicitly adjusting reward values (Klaes et al., 2011) or other aspects of the task (Kaufman et al., 2015) to encourage reaches to non-favored targets. Carefully executed, these approaches can equalize the number of movements to each target. However, eliminating biases in the final choice does not necessarily eliminate biases during planning. We found that while early planning activity did largely predict free choices, it was not without exception. There were some trials in which PMd activity suggested one reach direction throughout the trial, only to switch just moments before movement execution (Fig 5a). This type of last-second switching was relatively rare in our experiment, but a task designed to balance choices may simply encourage more frequent plan switching. If so, analysis based on the assumption that all targets are treated equally may not be valid during early planning phases of the trial. Trial-averaging methods simply cannot provide the temporal resolution necessary for determining instantaneous states of mind in high-order brain areas.

The monkeys in our study displayed quite strong preferences, but the source of those preferences is not obvious. They were not consistent across sessions, so it is unlikely that they resulted from the minimization of a cost function based on physiology or task performance (e.g., effort, energy expenditure, etc.) (Cos, Belanger, & Cisek, 2011; Morel, Ulbrich, & Gail, 2017; R. Shadmehr, Huang, & Ahmed, 2016). Rather, it seems that they reflected the monkeys' attempts at strategies that they believed could increase the likelihood of reward. Since the reward structure on free-choice trials was completely random, no such successful strategy was possible. The apparent randomness of the target preferences may therefore reflect the monkeys'

overinterpretation of brief patterns in the target presentations or even unsubstantiated guesses. Determining the source of target preferences and their effect on planning would require further experiments with a dynamic reward structure.

The high degree to which the initial reach plan predicted free-choice behavior suggests that early responses in PMd strongly influence the eventual movement decision. Supporting this idea is the observation that early reach plans on error trials almost always matched the movement direction (these type 1 and type 2 errors accounted for over 90% of all errors; Fig7d). That is, the monkeys made errors either because they unwaveringly stuck with their initial plan or switched back to it after briefly considering the other option. Changes of mind have been observed before in cortex (Bollimunta, Totten, & Ditterich, 2012; Kaufman et al., 2015; Kiani et al., 2014; Pastor-Bernier, Tremblay, & Cisek, 2012; Thura & Cisek, 2014), but not in direct contradiction to an explicit visual cue. In these cases, it appears that the monkey's first reaction to the pair of targets carried more weight than the subsequent—and completely informative—visual cue. This may indicate influence from another brain area overriding sensory inputs, or biases within PMd that resist changes away from an initial plan.

Recording from a large neural population allowed us to decode planning-related activity on a short timescale. Since the plans decoded early in the trial were predictive of both reach direction and reaction time on free-choice trials, we can be confident in the validity of the decoded neural states throughout the trial. This high temporal resolution was especially necessary when examining the link between instantaneous reach plans at the time of the Go cue and subsequent reaction time. Similar to previous studies (Ames et al., 2014; Coallier et al., 2015; Riehle & Requin, 1993), we found that the strength of the decoded reach plan was negatively correlated with reaction time on both 2-target and free-choice trials. This extended to

trials with incorrect plans at the time of the Go cue, where switching came at the cost of a significantly longer reaction time (Fig 3.6a). On the other end of the spectrum, 2-target trials with strong plans to the correct target sometimes showed very rapid reaction times (Fig 3.6a). Movement planning and execution are thought to involve separate neural processes (Haith, Pakpoor, & Krakauer, 2016), and our results suggest that the timing of movement execution is highly variable. The fastest reaction times indicate that the monkeys initiated movements in anticipation of the Go cue (Mark M Churchland, Byron, Ryu, Santhanam, & Shenoy, 2006; Rosenbaum, 1980). However, on trials with strong incorrect plans they were still able to withhold movement initiation to re-plan for the correct reach. This may suggest that the monkeys knew when they were planning the wrong reach, and that they delayed movement initiation to allow for the time needed to correct the plan.

The ability to interpret a quickly changing neural state with high temporal resolution is essential when studying high order brain functions like decision-making. While trial-averaged analysis methods have contributed a great deal to our understanding of movement planning and decision-making, they may at times suggest an oversimplified view of neural processes. The heterogeneity in single-trial responses observed in this study and others suggest that decision-making cannot be fully explained through a simple model of weighted sensory cues. Large-scale recordings (likely in multiple brain areas simultaneously) are almost certainly necessary to fully characterize the processes leading to a decision.

METHODS

Subjects

We trained two male rhesus macaque monkeys (*Macaca mulatta*) to make reaches using a planar robotic manipulandum for water or juice reward. All procedures were approved by the Northwestern University Institutional Animal Care and Use Committee (Protocol number #IS00000367) and were consistent with the Guide for the Care and Use of Laboratory Animals.

Behavioral task

The monkeys sat in a chair facing a vertical monitor and used a planar robotic manipulandum to control an on-screen cursor. Each 2-target trial (40%) began once the monkey held the cursor within a central target (1.5cm radius) for 500 ms, after which two outer targets appeared (target on; 750—1000ms), always 180 degrees apart. The target locations were restricted to eight different locations, equally spaced at a radius of 7cm. After the target on period, the outer targets briefly disappeared (target off; 250—500ms). This provided a time window in which we could analyze the neural response to the previously presented targets with little chance of visual confounds. After this target off period, one of the targets reappeared (cue on; 250—500ms), providing the monkey with complete information of the correct. The missing target then reappeared, returning to the initial 2-target presentation. The center target was extinguished, and a tone signaled the monkey to move (go; <5s). The trial ended when the cursor reached one of the outer targets. If the target was correct (as indicated by the target presented during the cue on period), the monkey received a success tone coincident with delivery of a

liquid reward. If incorrect, the monkey received a failure tone and the screen displayed the location of the correct target.

One target trials (40%) followed the same structure as above, except only the correct target was ever shown. Thus the monkey had complete information about the correct target from the beginning of the trial. On twenty percent of all trials, we omitted the cue period altogether. This resulted in 2-target free-choice trials (10%), where the monkey was forced to choose one of the two targets without any information as to which was correct. We maintained the same reward structure during these trials, so there was always a 50% chance that the monkey would receive a reward. The 1-target trials without cue period were not used in any analyses.

Neural recordings and preprocessing

Both monkeys were implanted with two chronic, 96-electrode arrays (Blackrock Microsystems, Salt Lake City UT) positioned over the arm area of primary motor cortex (M1; 1.5mm electrode length) and straddling the rostral/caudal division of dorsal premotor cortex (PMd; 1.0mm electrode length). We discriminated single neurons offline by isolating clusters within a principal components space calculated from the waveform shapes of putative neurons (Plexon Inc., Dallas TX). For monkey C this yielded unit counts of 143, 143, and 108. For monkey M: 154, 132, and 114.

We used the spiking events from each recorded unit to calculate a continuous estimate of firing rate by first convolving with a half-Gaussian kernel (s.d. 150ms) and then downsampling to 200Hz. We chose a half-Gaussian kernel to ensure a causal relationship between spiking

events and the estimated firing rate. We then applied a square root transform to all firing rate estimates.

Functional correlations: estimation and simulation

For the neural correlation hypothesis testing shown in Figure 3.2d,e, we wished to compare only the components related to reach direction. We assumed that each pairwise covariance $cov_{total}(h_i, h_j)$ resulted from the linear sum of the desired reach direction tuning-related covariance $cov_{tuning}(h_i, h_j)$ and noise covariance $cov_{noise}(h_i, h_j)$. We obtained $cov_{tuning,1T}(h_i, h_j)$—that is, the tuning-related covariance across all 1-target trials—directly by averaging the firing rates (square root transformed) to each target direction. For 2-target trials, we could not calculate a mean firing rate per target direction, since the monkey could have been planning a reach to either target (or both or neither). Instead, we assumed that the noise covariance would not change drastically between 1-target and 2-target trials and then estimated the tuning covariance:

$$cov_{tuning,2T}(h_i, h_j) = cov_{total,2T}(h_i, h_j) - cov_{noise,1T}(h_i, h_j), \qquad \text{Eq. 3.1}$$

To simulate the covariance changes expected given a stay-or-switch response (Figure 3.2e), we first calculated non-parametric tuning curves $f_i(\theta)$ and noise covariance matrix $\Sigma(H)$ from 1-target trials. For each simulated 2-target trial, we randomly selected between the two possible reach directions according to the choice preferences exhibited on 2-target free-choice trials. We then modeled the activity of all neurons by taking the direction-related firing rates (from the

tuning curves) and adding noise drawn from a multivariate Gaussian distribution defined by $\Sigma(H)$. Finally, we applied the same analysis as described above using the simulated firing rates.

To simulate the simultaneous encoding response, we followed the same procedure as for the stay-or-switch simulation, except using modified tuning curves. We assumed a simple model of simultaneous tuning in which each neuron's firing rate reflected the larger of the two firing rates expected for either target alone. That is, given the tuning curve $f_i(\theta)$, where θ is the reach direction, we defined each simultaneous encoding tuning curve to be

$$g_i(\theta) = \max[f(\theta), f(\theta + 180°)] \qquad \text{Eq. 3.2}$$

We calculated the percentages in Figure 3.2e from neuron pairs for which the bootstrapped 95% bounds of $cov_{tuning,1T}$ and $cov_{tuning,2T}$ did not span zero. This restriction prevented any potential effects from being diluted by neuron pairs for which we could not adequately estimate the tuning-related covariance.

Dimensionality reduction and reach plan likelihood estimation

On each session, we grouped all 1-target firing rates into the matrix $M \in \mathbb{R}^{N \times T}$, where N is the number of neurons and T is the number of time points obtained by concatenating all 1-target trials. We then performed PCA on the matrix M and projected all firing rates (from 1-target, 2-target, and 2-target free-choice) onto the top ten principal axes. This provided a 10-dimensional neural state at every point in time.

We used the neural state readout S to interpret, at every time point, the likelihood that the population activity corresponded to a reach plan to each of the eight possible target directions. For each 1-target reach direction i, we assembled the neural states (from 500ms post-target appearance until trial completion) observed across all trials into the set $\{C_i\}$. We then used the Mahalanobis distance D_M between the neural state and the training set to determine the likelihood that the neural state corresponded to direction i:

$$P\left(i\middle|D_M(S,\{C_i\})\right) = \frac{P\left(D_M(S,\{C_i\})\middle|i\right)}{\sum_{j=1}^{8} P\left(D_M(S,\{C_i\})\middle|j\right)} \quad \text{Eq. 3.3}$$

Note that our calculation of $P(i)$ is dependent only on $D_M(S,\{C_i\})$ and not $D_M\left(S,\left\{C_{j\neq i}\right\}\right)$. That is, the likelihood of each direction depends only on the distance of the neural state to the training set of that same direction. Thus, the likelihood values for all states do not sum to one. This was to allow for the possibility of dual or intermediate reach plans during 2-target trials, which could lead to neural states convincingly similar to multiple training sets.

The Δlikelihood value used throughout the results reflects the difference in likelihood values between either the two presented target directions (2-target trials) or the single presented target and the opposite direction (1-target trials).

Reaction time correlation

We defined reaction time as the time from the go cue to the time when the cursor had moved 1cm along a line projected from center to outer target. Using other definitions of reaction time based on velocity did not change the results.

We used a linear mixed effects model to test for a general correlation between reaction time and Δlikelihood across all reach directions and sessions. We used the *fitlme* function in Matlab (Mathworks) with reaction time as the response variable, Δlikelihood as the predictor variable, and reach direction (separated by session) as the grouping variable, with uncorrelated random effects for intercept and Δlikelihood.

Supplementary Materials

Our results suggest that monkeys plan a single reach on any given 2-target trial, contradicting previous reports of simultaneous reach representations in PMd. We speculate that this different result might have been due to our use of many neurons recorded simultaneously, rather than single unit recordings aggregated across sessions. However, it is also possible that the monkeys in our study simply adopted a strategy that differed from that of the earlier studies. To differentiate between these two possible explanations, we replicated several single-neuron analyses from previous studies. We reasoned that if our monkeys had used a different strategy, then our single-neuron results should not resemble those from studies finding simultaneous representation.

In Fig 3.S1a,b, and c, we show the results when we applied the three major analyses used by Cisek and Kalaska (2005) to our data. We used only neurons recorded during single sessions to avoid the possibility of duplicates. On 2-target trials, directional tuning appeared to be narrower and overall activity lower than on 1-target trials, matching the results from the referenced paper. Additionally, the analysis in Fig 3.S1c indicates that most neurons were more active during 2-target trials than would be expected under a guess-and-switch approach. This

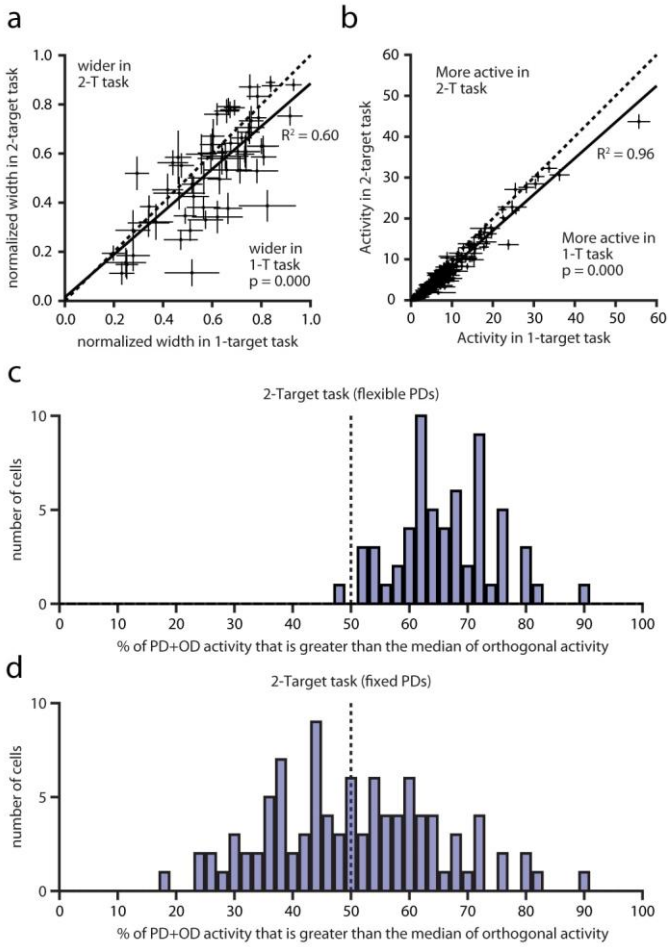

Figure S 3.1 | Single neuron analysis and guess-and-switch control analysis

(a) Plot of normalized tuning width during 1-target and 2-target trials (See Cisek and Kalaska, 2005) (b) Firing rates of single neurons for targets in their preferred direction on 1- and 2-target trials (see Cisek and Kalaska, 2005) (c) Control used by Cisek and Kalaska (2005) to discount the possibility of a guess-and-switch approach. Skew toward 100% suggests simultaneous representation (d) Same analysis as in c, but using only 1-target PDs.

result therefore appears to support the simultaneous representation hypothesis, in contradiction to our main result. However, we found this to be the case only if we calculated neuronal preferred directions for the 2-target trials separately from the 1-target trials, as was done in the original studies. Since the simultaneous representation hypothesis assumes there is no change in PD between 1- and 2-target trials, using a single PD calculated only from 1-target trials should not affect the results. Nonetheless, we found that this simple constraint eliminated all evidence of simultaneous representation (Fig 3.S1d).

We hypothesized that the inconsistency of the results between Supplementary Figures c and d was due to target preferences during planning that influenced the 2-target tuning curve calculations. To test this, we ran a simulation of stay-or-switch planning on a 2-target task using a population of 500 artificial neurons. Each neuron's directional tuning followed an exponentiated sine wave (amplitude between 5 and 10Hz), with equally distributed preferred directions and Poisson noise. We simulated 4000 1-target trials and 4000 2-target trials, imposing on the 2-target trials either a weak or strong target preference (Fig 3.S2a). The imposed target preferences meant that for the 0/180 degree target pair, we simulated a 0 degree target plan for either 60% (weak preference) or 90% (strong preference) of trials. From the resulting simulated firing rates, we recalculated separate preferred directions for 1-target and 2-target trials. The differences between these two calculations were usually quite small (Fig 3.S2b), which at first glance seems to indicate that the two can be used interchangeably. However, the control analyses (Fig 3.S2c,d) showed that this was not the case. Using separate PD calculations when target preferences were strong caused a spurious result that appeared to strongly support the simultaneous representation hypothesis (Fig 3.S2c, red distribution highly skewed toward 100%). Restricting the PDs to those calculated from 1-target trials corrected this problem, and

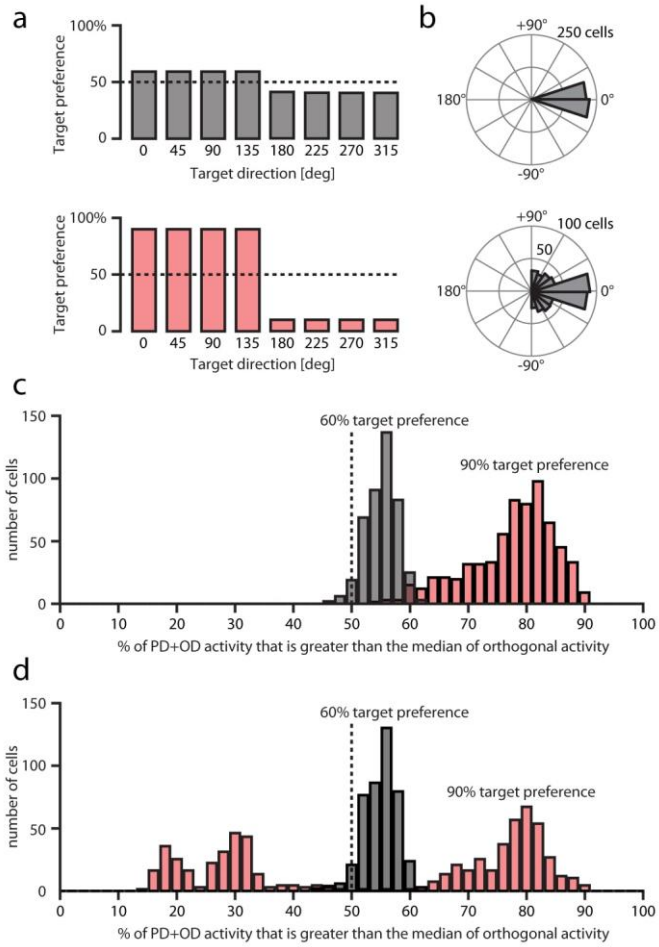

Figure S 3.2 | Stay-or-switch with biases can mimic simultaneous tuning

(a) Weak (top; 60%) and strong (bottom; 90%) target preferences used to simulate stay-or-switch population activity. (b) Rose plots showing the difference in preferred directions between 1-Target and 2-Target trials for the simulated population. Top corresponds to the session with weak target preferences (60%), bottom to strong (90%). (c) Analysis (from Cisek and Kalaska, 2005) used to test for the possibility of a stay-or-switch approach. Rightward skew is meant to support the simultaneous encoding hypothesis. (d) Same analysis as in c, but only using PDs calculated during 1-Target trials.

correctly discounted the simultaneous hypothesis (Fig 3.S2d, bimodal red distribution spanning 50%). When target preferences were weak, the analysis successfully ruled out the possibility of simultaneous encoding (Fig 3.S2c,d, grey bars; slight skew above 50% was due to non-symmetric tuning profiles and completely overlapped the 1-target distribution). These results show how seemingly valid assumptions (e.g., that a small average difference between two PD calculations would be inconsequential) can lead to misleading interpretations.

CHAPTER 4: PREMOTOR AND MOTOR CORTICES ENCODE REWARD

Pavan Ramkumar, Brian Dekleva, Sam Cooler, Lee Miller, Konrad Kording

FOREWORD

This chapter is a version of an article published by PLOS ONE in 2016. The idea for this paper began with the peculiar observation (made by Sam Cooler and me) that neurons in PMd seemed to display bursts of activity at the end of failed trials. Typical motor cortex experiments are not particularly suited to discovering this this type of effect for two main reasons: (1) the behavioral tasks do not contain a significant number of failures and (2) analyses are restricted to the neural activity occurring strictly within each trial. The task we used in Chapter 2 is somewhat unique in that frequent failure was unavoidable. This made any related neural effect much more noticeable. Still, it seemed possible—and indeed even likely—that the result-dependent activity differences we saw reflected some movement-related confound. For example, perhaps after failing the monkey made a rapid movement back to the start target so he could begin a new trial. To address these kinds of possibilities, Pavan Ramkumar modeled single-neuron activity by incorporating every conceivable confounding variable. To our surprise, nothing but trial outcome could account for the differences in post-trial activity. We thus concluded that motor cortex may in fact receive transient reward- or failure-dependent signaling. The absence of previous results on this front made the finding unexpected. However, the great extent of learning that occurs within motor cortex itself (see Chapter 1) provides an explanation for why such signaling could—and indeed should—exist. This concept will be discussed further in Chapter 5.

ABSTRACT

Rewards associated with actions are critical for motivation and learning about the consequences of one's actions on the world. The motor cortices are involved in planning and executing movements, but it is unclear whether they encode reward over and above limb kinematics and dynamics. Here, we report a categorical reward signal in dorsal premotor (PMd) and primary motor (M1) neurons that corresponds to an increase in firing rates when a trial was not rewarded regardless of whether or not a reward was expected. We show that this signal is unrelated to error magnitude, reward prediction error, or other task confounds such as reward consumption, return reach plan, or kinematic differences across rewarded and unrewarded trials. The availability of reward information in motor cortex is crucial for theories of reward-based learning and motivational influences on actions.

INTRODUCTION

How the brain learns based on action outcomes is a central question in neuroscience. Theories of motor learning have usually focused on rapid, error-based learning mediated by the cerebellum, and slower, reward-based learning mediated by the basal ganglia (Shmuelof & Krakauer, 2011). Different combinations of reward and sensory feedback result in different learning rates. For instance, positive and negative rewards influence motor learning differently (Abe et al., 2011; Galea et al., 2015). When reward is combined with sensory feedback, it can accelerate motor learning (Nikooyan & Ahmed, 2015). Reward is thus a fundamental aspect of learning (Dayan & Balleine, 2002; Hollerman & Schultz, 1998; O'Doherty et al., 2003; W. Schultz, 2006). Various reward signals have been characterized in the midbrain, prefrontal and limbic cortices (Matthew R Roesch & Olson, 2003, 2004; M. R. Roesch & Olson, 2005;

Wolfram Schultz, 2000; Wallis & Kennerley, 2010). Yet, we do not know how neurons in the motor system obtain the reward information that could be useful for planning subsequent movements.

The dorsal premotor cortex (PMd) and the primary motor cortex (M1) are known to be involved in planning and executing movements. We know this because movement goals (e.g., direction of upcoming movement), kinematics (e.g., position, velocity and acceleration) and dynamics (e.g.,forces, torques, and muscle activity) are reflected in the firing rates of motor cortical neurons (Ashe & Georgopoulos, 1994; Cheney & Fetz, 1980; Evarts, 1968; Georgopoulos et al., 1992; Holdefer & Miller, 2002; Morrow et al., 2007; Scott & Kalaska, 1995). If movement plans need to be modified based on previous actions, then information about their outcomes must reach motor cortices. In many real-world settings, task outcomes typically manifest in the form of reward.

Recently, Marsh et al. (Marsh, Tarigoppula, Chen, & Francis, 2015) have shown a robust modulation of M1 activity by reward expectation both during movement and observation of movement. To further investigate the nature of this potential reward signal, we trained monkeys to reach to targets based on noisy spatial cues and rewarded them for correct reaches. We induced different reward expectation on a trial-by-trial basis and quantified the representation of reward in PMd and M1. We observed that ~28% of PMd neurons and ~12% of M1 neurons significantly modulated their firing rates following trials that were not rewarded. The effect could not be explained simply by kinematic variables such as velocity or acceleration, reward consumption behavior, or upcoming movement plans, nor by task variables that may bias successful task performance, such as the noise in the target cue, the reward history, or the

precision of the reach. This effect might constitute an important piece in the larger puzzle of how motor plans are modified based on reward.

RESULTS

Our goal in this study was to investigate whether the motor system — in addition to planning and executing actions — also encodes responses to reward, which are key for learning about the environment and modifying motor plans. To this end, we trained two macaque monkeys to make center–out reaches to uncertain targets and rewarded them for successful reaches. The monkeys made reaching movements while grasping the handle of a planar manipulandum, their hand position represented by an on-screen cursor (Fig 4.1A). During this task, we recorded from two 96-channel microelectrode arrays (Fig 4.1B, Blackrock Microsystems), chronically implanted in the primary motor cortex (M1) and the dorsal premotor cortex (PMd).

A trial began when the monkey moved the cursor to the central target (Fig 4.1C). The true location of the target was not shown. Instead, a noisy cue was presented at a 7-cm radial distance from the center to indicate the approximate location of the outer target. The cue comprised a cluster of line segments generated from a distribution centered on the true target. Monkeys were trained to reach only after a combined visual/auditory go cue, which was delivered after a variable (0.8–1.0 s) delay period following cue appearance. At the end of the reach, the actual circular target (15° diameter) was displayed. If the monkey had successfully reached the target, an appetitive auditory cue announced the subsequent delivery of a juice reward. If the reach ended outside the target, an aversive auditory cue announced the failure of the reach and no reward was delivered. After the end of the trial, the monkey was cued to return

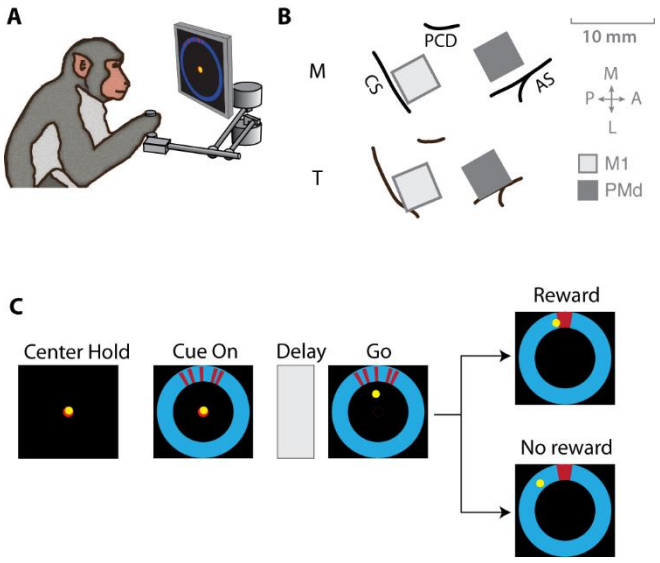

Figure 4.1 | Reaching task to uncertain targets

(A) Monkeys made center–out reaches using a planar manipulandum that controlled an on-screen cursor. (B) We recorded from chronic microelectrode arrays implanted in dorsal premotor cortex (PMd) and primary motor cortex (M1). CS: central sulcus, AS: arcuate sulcus, PCD: precentral dimple (C) Monkeys reached towards a target that was cued using a set of 5 line segments, whose dispersion varied from trial to trial. In each trial, the true target location was sampled from a von Mises "prior" distribution centered at 90° (clockwise from the rightmost point on the annulus) with one of two concentration parameters specifying a broad or narrow prior. The line segments making up the cue were then sampled from a "likelihood" von Mises distribution centered on the target location, with one of two concentration parameters ($\kappa = 5$ or 50 for Monkey M and $\kappa = 1$ or 100 for Monkey T) specifying a broad or narrow spread. Adapted from Dekleva et al., 2016.

to the center target in order to begin the next trial. The median inter-trial interval across animals and sessions was 2.79 ± 0.25 seconds. On any given trial, since the actual target was not shown, the monkey had to infer its location, potentially by combining the information in the noisy target cue with prior knowledge accumulated about the target location in previous trials (Dekleva, Ramkumar, Wanda, Kording, & Miller, 2016).

We assume that the monkeys calibrated their expected reward based on the cue uncertainty. To manipulate their reward expectation, we varied cue uncertainty from trial to trial. Specifically, we determined the dispersion of the line segment cluster on each trial by drawing the location of each line segment from either a narrow or a broad distribution (see Fig.1 for details). A narrower spread of line segments indicated the target location with lesser uncertainty than a broader spread. To verify that animals indeed change their reward expectation, we looked at the latency of movement onset after the go cue. We found that animals indeed started their reach later on average when they were more uncertain about the target location (35 ± 28 ms for narrow spreads; 111 ± 16 ms for broad spreads; mean ± SEM across animals and sessions). Thus, manipulating the dispersion on each trial is likely to have induced trial-by-trial changes in reward expectation.

Neural coding of reward

We asked if the firing rates of PMd and M1 neurons indicated whether a reward was obtained in the trial by comparing the peristimulus time histograms (PSTHs) aligned to the end of the trial timestamp (corresponding to the auditory cue that indicated whether a reward will be delivered) for rewarded and unrewarded trials. We matched the kinematics of the trials across the two conditions to control for trivial firing rate consequences of behavioral differences (see Fig

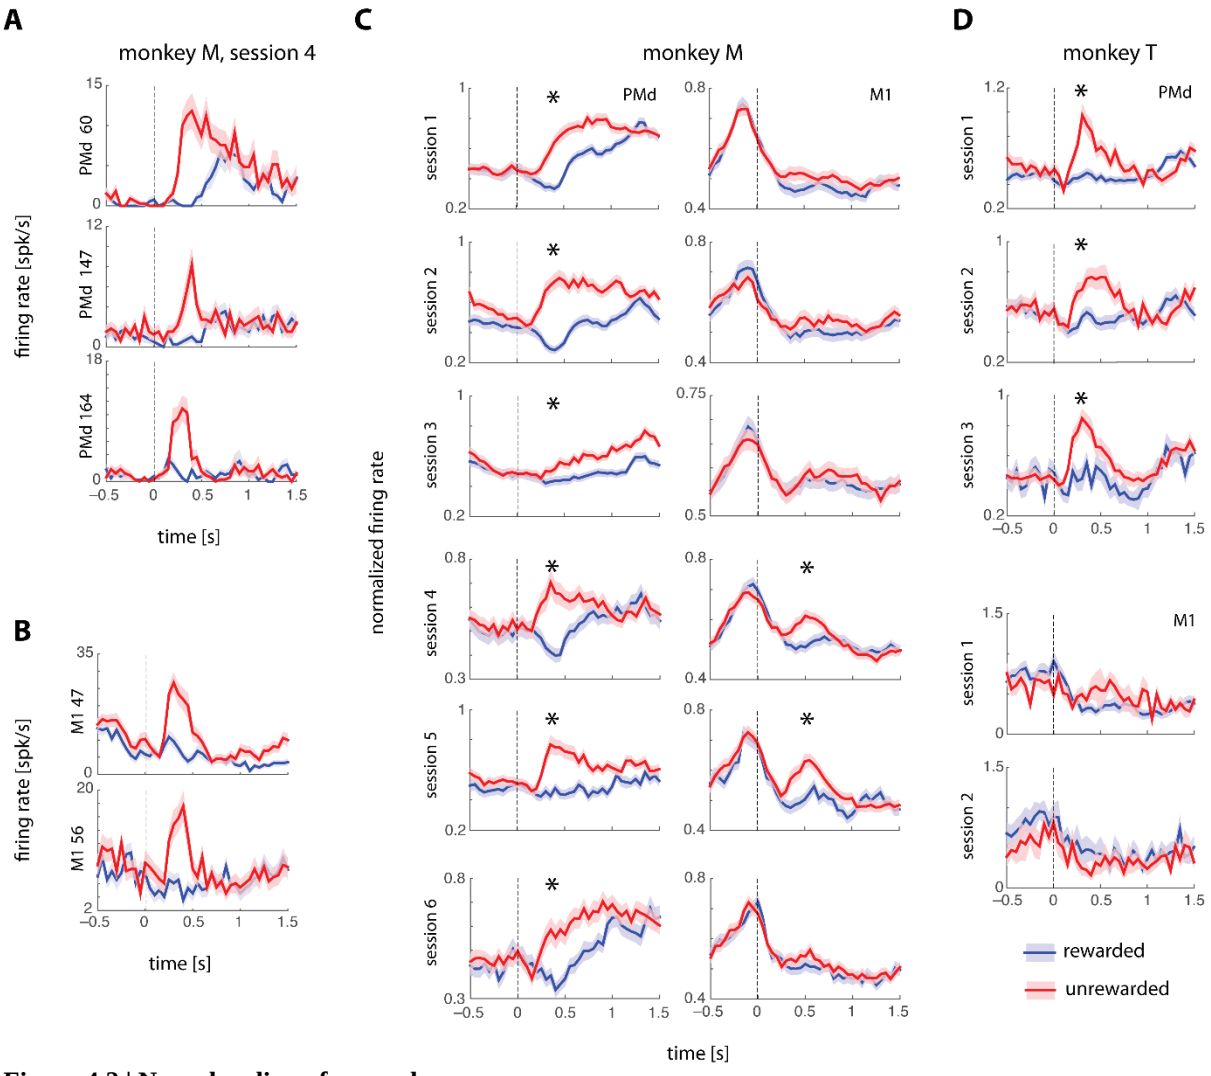

Figure 4.2 | Neural coding of reward

(A, B) Example neurons from PMd and M1 showing reward modulation that persisted after controlling for kinematic differences between rewarded and unrewarded trials (see S1 File, Fig. A). Error bars show standard errors (SEMs) across trials. Vertical dashed line at zero indicates the time of reward. (C, D) Trial-averaged, population-averaged normalized firing rates (mean ± SEMs across neurons) for PMd and M1 from two monkeys. For each session, significant differences between rewarded and unrewarded peak PSTH amplitudes are indicated using an asterisk. PMd shows a clear increase after unrewarded trials compared to kinematically-matched rewarded trials.

4.S1 and Table A for details about matching kinematics). We found that ~28% of PMd neurons and ~12% of M1 neurons modulated their firing rates in response to reward or lack thereof. Nearly 25% of all PMd neurons recorded increased their firing rates following unrewarded trials compared with rewarded trials (Figs. 2A, 2B show example PMd and M1 neurons from one session). In comparison, only ~3% of PMd neurons increased their firing rates after rewarded trials. For M1, these numbers were ~8% and ~4%, respectively.

We then determined the extent to which this effect was visible across the entire population. To do so, we normalized single neuron PSTHs computed from kinematically-matched trials by setting the peak of each PSTH to 1, and computing separate PTSHs for M1 and PMd. There was a significant increase in population-wide normalized firing rate following unrewarded trials in both monkeys (Fig 4.2C, D). PMd had a significant effect in all 9 sessions (6 from Monkey M and 3 from Monkey T), whereas the effect in M1 was significant in only 2 out of 9 sessions.

The firing rate effect of unrewarded trials in PMd and to a lesser extent, M1, is completely confounded by the fact that only successful trials were rewarded. Thus, increased firing rate for unrewarded trials could potentially be an intrinsic signal of success or failure, or might indicate some other correlate of the outcome of a goal-directed movement. To eliminate this confound, we ran a separate experiment on one monkey (Monkey M) in which we withheld reward in a subset of successful reaches. We found that firing rates increased even for these successful but unrewarded trials (Fig 4.S2), suggesting that the increased activity following unsuccessful trials is related to lack of extrinsic reward, not an intrinsic measure of task outcome.

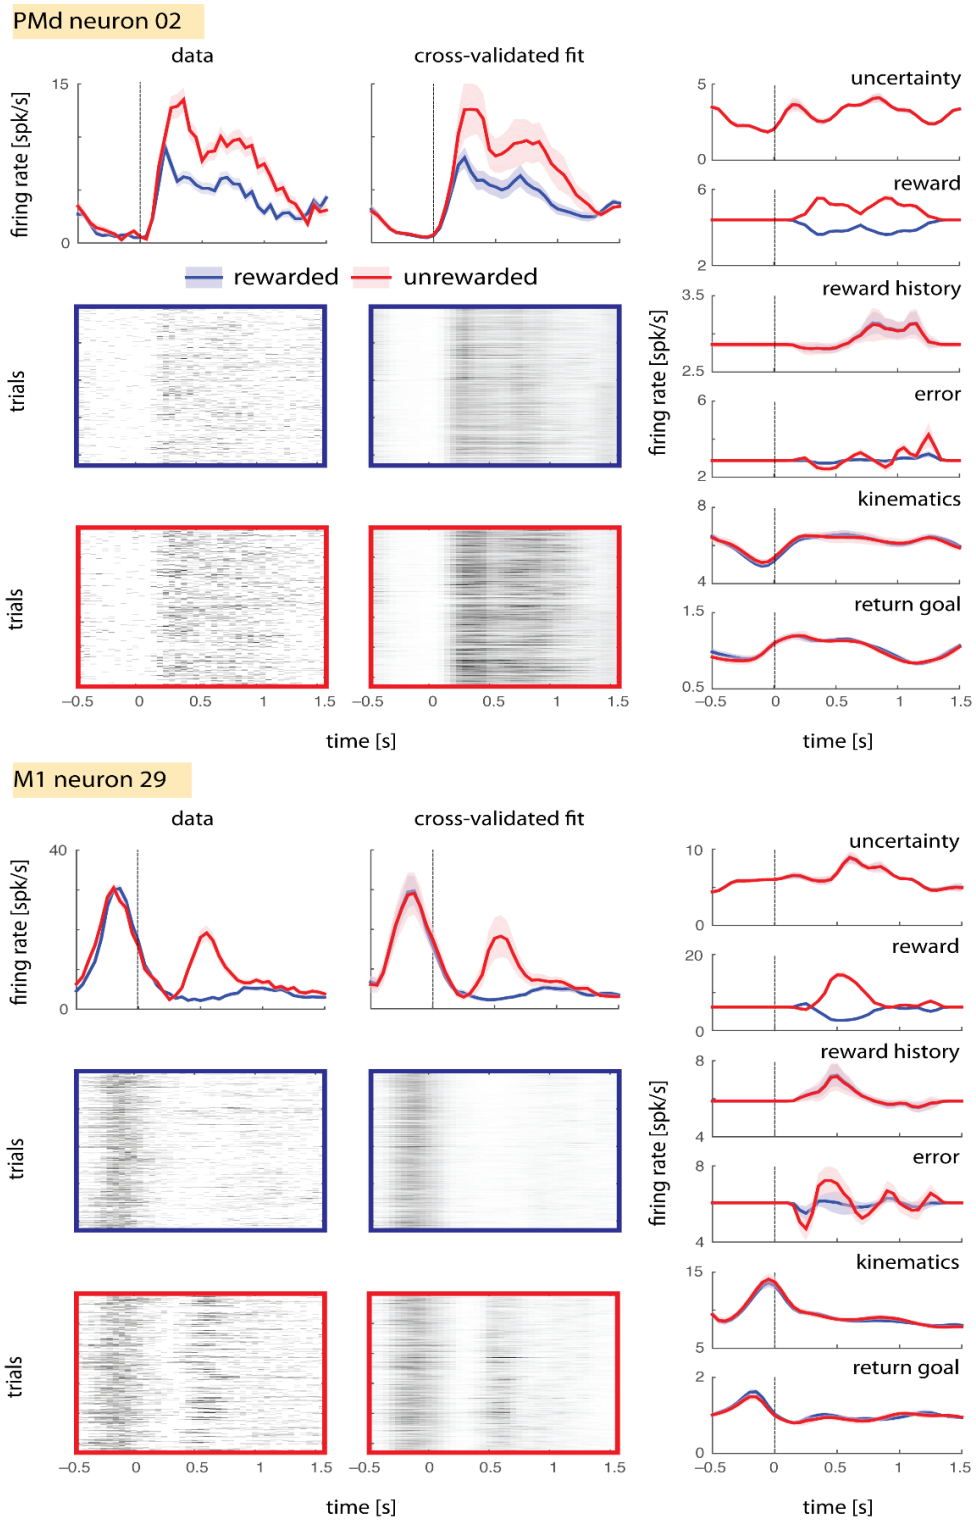

Figure 4.3 | Generalized linear modeling of reward coding.
Model predictions for two example neurons are shown. Left: PSTHs for rewarded (blue) and unrewarded (red) trial subsets are shown for the test set, along with corresponding single-trial rasters for both data and model predictions on the test set. Right: Component predictions corresponding to each covariate.

Putative reward signal is not explained away by task confounds

Several other variables could potentially confound this putative reward signal as well. As we did for kinematic differences (Fig 4.S1), examining groups of rewarded and unrewarded trials that are matched for these confounding variables is a potential means of disambiguating the source of the effect. Yet, adequately matching all possible confounding variables is impossible because of the trade-off between precise matching of numerous potential confounds, and adequate remaining sample size. Instead, we controlled for potential confounds by using multiple linear regression models of trial-by-trial firing rates. Specifically, we modeled single neuron spike trains using Poisson generalized linear models (GLMs) (Fernandes, Stevenson, Phillips, Segraves, & Kording, 2014; Park, Meister, Huk, & Pillow, 2014; Ramkumar et al., 2016). See Methods for more information.

To construct the GLM, we modeled neural spike counts during a 2-second epoch (–0.5 to 1.5 seconds, in 10-ms bins) around the reward onset. Spike counts were modeled as a function of the reward, which we represented as a binary variable (+1 for rewarded trials, –1 for unrewarded trials), aligned to the reward onset. In addition, we included the following confounding variables in the multiple regression.

1. *Kinematics*. PMd and M1 neurons are known to encode kinematic variables during movement planning and execution (Ashe & Georgopoulos, 1994; M. M. Churchland, G. Santhanam, et al., 2006; Georgopoulos et al., 1992). Therefore, we included instantaneous velocity and acceleration time series, binned in 10-ms time bins.

2. *Uncertainty*. Previous work has suggested that PMd can encode plans for more than one potential target (P. Cisek & Kalaska, 2005) and we have recently shown that PMd encodes uncertainty about the reach target location (Dekleva et al., 2016). Further, cue uncertainty influences the likelihood of a successful outcome, since monkeys are more successful in low-uncertainty trials. Although we did not find an effect of uncertainty on PSTHs aligned to reward time (Fig 4.S3), we included a measure of trial-specific uncertainty as a confounding variable. Specifically, we used the dispersion of the target cue line segments, where dispersion is the largest circular distance between all possible pairs of line segments.

3. *Reward history*. The outcome of the previous trial (and more generally, the history of reward) can influence the level of satiety, and thus the motivation and perceived value of a potential reward (Matthew R Roesch & Olson, 2004). To control for this possibility, we included the previous trial's outcome as a binary covariate (+1 for success, –1 for failure).

4. *Error*. The reward-related signal might be useful for reinforcement learning (temporal difference learning) if it encoded some information about the discrepancy between the reach direction and the true target direction (presented visually at the end of the trial). To test whether PMd/M1 neurons encode error magnitude, we included the unsigned reaching error (reach precision) as a covariate.

5. *Return goal*. Another potential confound is that the movement plan for the return reach to the center target may be modulated by recently obtained reward. Although a separate control analysis (Fig 4.S4) suggested there were no systematic differences between return reach

planning for rewarded and unrewarded trials, here we controlled for this possibility by including a covariate that specified the return reach direction; in particular, we used two covariates specifying the direction cosines (cosine and sine) of the return reach direction.

Environmental covariates — uncertainty, error, and reward — that are potential causes of firing rate changes, invariably *lead* spikes in M1 and PMd. By contrast, movements are the consequence of motor cortical activity. Therefore, kinematic variables are likely to *lag* spikes. To model these latency differences for the different covariates, we used temporal basis functions (see Methods for details).

Our model accurately captured the reward-related activity of many neurons. Comparing the data and cross-validated fit (see Methods) panels of Fig 4.3, we see that the trial-averaged data PSTHs and trial-averaged model-predicted firing rates are extremely similar. Across 70 PMd and 191 M1 neurons in a representative session (Monkey M, session 4), the model explained almost all the variance in the trial-averaged data (mean ± standard deviation of R^2 = 0.96 ± 0.05 and 0.93 ± 0.08 for PSTHs averaged across successful and unsuccessful trials, respectively). These high R^2s suggest that the model includes almost all potential sources of predictable variance. Therefore, if the reward covariate cannot be explained away by the confounding covariates, it is likely that the neurons represent reward.

To understand whether the reward covariate explains a significant fraction of the variance, we visualized the predictions of individual model covariates (reward, kinematics, uncertainty, reward history, error, and return goal), for rewarded and unrewarded trial subsets. Among all model covariates, only the reward covariate made different predictions for firing rates in rewarded and unrewarded trials (Fig 4.3, right panels); the predictions of other covariates were

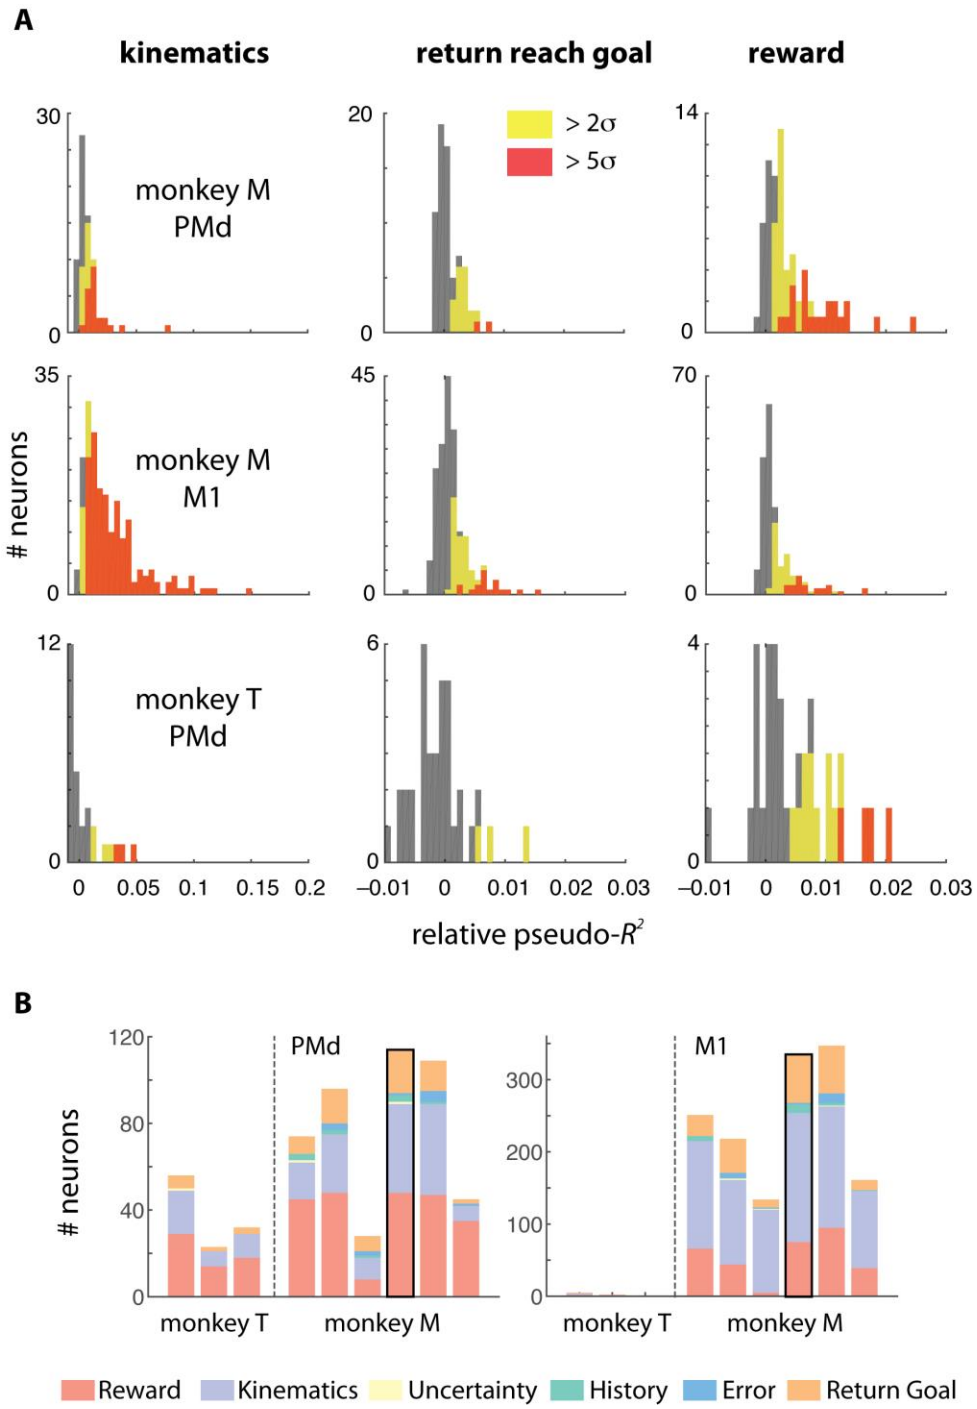

Figure 4.4 | Across-session summary of reward encoding

(**A**) The distribution of effect sizes across the population from one session in each monkey (Monkey M, Session 4 and Monkey T, Session 1) as measured by mean relative pseudo-R^2s are shown for kinematics, return reach goal, and reward. A large fraction of neurons in both PMd and M1 remained statistically significantly modulated by reward after controlling for cue uncertainty, reward in the previous trial, error magnitude, instantaneous kinematics, and planning of the return reach. (**B**) A stacked bar shows the number of significant neurons at the 2-sigma level for each session. The representative session for which histograms of effect sizes are shown above is indicated with a black border.

similar across these conditions. This preliminary analysis of example neurons seemed to suggest that reward was indeed the predominant driver of firing rate variance.

To quantify the marginal effects of reward, kinematics and other confounding variables, we also built partial models leaving each covariate out and then comparing these respective partial models against the full model using the relative pseudo-R^2 metric as a measure of effect size (see Methods for definitions and details). Briefly, we used two-fold cross-validation to quantify error estimates on pseudo-R^2s. We used each half of the data as a training set to fit the full and partial models and computed the pseudo-R^2s on the other half (the test set). We obtained 95% confidence intervals (CIs) on the cross-validated test set pseudo-R^2 by bootstrapping on the test sets and used these to determine which neurons were significantly predicted by the covariate of interest at the 2σ and 5σ significance levels (see Methods).

Three of the covariates accounted for a large fraction of the variance in PMd and M1 firing rates. As expected, a large number of neurons (41/70 in PMd, and 179/191 in M1; 2σ significance criterion; Fig 4.4A, left panel) were significantly modulated by reach kinematics (instantaneous velocity and acceleration) in a representative session (Monkey M, session 4). Further, many neurons (20/70 PMd and 67/191 M1; Fig 4.4A, middle panel) encoded the direction of the upcoming return reach. However, a large fraction of PMd (48/70), and M1 (75/191) neurons also encoded reward (Fig 4.4A, right panel). By comparison, a negligible number of neurons in either PMd or M1 encoded cue uncertainty (1), error magnitude (3) or reward history (15) — these were likely false positives that did not survive a multiple-comparison correction. These results were very similar across multiple sessions in both monkeys (Fig 4.4B). For Monkey T, the quality of the M1 array had degraded at the time of these

experiments, and the spike-sorted neurons had extremely low firing rates, insufficient to fit reliable multivariate GLMs. Therefore we were only able to quantify the effects in PMd.

In the generalized linear models, we could not directly control for actual reward consumption, which involves mouth and neck movements, and might therefore affect firing rates in PMd and M1. However, we mainly observed *increases* in neural activity when reward was not received. This would only be possible if the monkey had made more vigorous mouth or neck movements in the absence of reward, than while actually consuming reward (e.g.,by potentially sucking harder on the tube when no reward was delivered). We could not quantitatively control for this possibility because we did not measure kinematics or EMG signals from the face. However, we made doubly sure that the monkey correctly interpreted the auditory cue signaling lack of reward. To do this, we filmed one monkey during a separate session (Monkey M; see Fig 4.S5) and observed that it simply sat still during trials when no reward was delivered, without making any oral contact with the reward delivery tube. Therefore, the robust reward signal cannot be explained by reward consumption. Taken together, our results suggest that a large fraction of PMd and M1 neurons encode the outcome of the task independently of uncertainty, error magnitude, kinematics, reward history, reward consumption, and the return reach plan.

DISCUSSION

We asked if the premotor and motor cortices, implicated in planning and executing movements, might also represent the reward associated with those movements. We found a strong representation of reward in PMd firing rates, with a lesser effect in M1. The increase in firing rates was observed in response to the absence, and to a lesser extent, the occurrence of extrinsic reward, but not the intrinsic success or failure of the trial. We then asked if the reward

signal encoded motivation or satiety (modeled by reward history), prediction error (modeled by cue uncertainty), or movement precision (modeled by error magnitude), but found no evidence for any such signals. We also confirmed that although kinematics and return movement planning could explain firing rate variance, neither of them could explain away the reward signal. Although the motor cortex has traditionally been thought of as a brain area that sends control signals to the spinal cord and muscles, recent studies (Marsh et al., 2015) including ours, establish the important additional effect of reward on motor cortex activity.

One weakness of this experiment is the lack of statistical power to ask if there were trial-by-trial reward dependent learning effects. The target location and the dispersion of the cue lines were drawn at random on each trial; hence, there was very little opportunity to transfer knowledge from one trial to the next. As trial-by-trial learning is generally relatively slow (and further slowed by high feedback uncertainty (Wei & Kording, 2010) we expected only a small trial-by-trial effect. Not surprisingly then, we did not find that the reward signal was directly tied to the behavioral performance of subsequent trials. If and how the reward signal does influence trial-by-trial learning should be investigated with further experiments. A second weakness of our design, which is typically common across many animal experiments, is that it does not rule out the possibility of covert motor rehearsal following error or lack of reward. Such rehearsal might activate premotor and motor cortices without resulting in overt behavior. Although we rule out any direction-specific effects of the reported reward outcome signal, it is impossible to definitively rule out the influence of covert rehearsal. This is an important constraint that future experiments must contend with.

Reward is a central feedback mechanism that regulates motivation, valuation, and learning (Wolfram Schultz, 2000; W. Schultz, 2006). Existing computational theories of this

phenomenon, such as reinforcement learning and temporal difference learning (Hollerman & Schultz, 1998; O'Doherty et al., 2003), have been successful in explaining the dopaminergic prediction error signal, but reward coding in the brain is far more heterogeneous and pervasive than just that (Dayan & Balleine, 2002; Wallis & Kennerley, 2010). The large majority of brain areas implicated in reward processing, such as basal ganglia, ventral striatum, ventral tegmental area, and the orbitofrontal cortex, are predictive in nature, with predictions including reward probability, reward expectation, and expected time of future reward (Wolfram Schultz, 2000). Dopaminergic neurons in the ventral striatum also encode the mismatch between predicted and obtained rewards, combining reward prediction with reward feedback. Thus far, only the lateral prefrontal cortex has been shown to encode reward feedback without any predictive component. To our knowledge, the previous studies examining reward-related signaling in the premotor and motor cortices (Marsh et al., 2015; Matthew R Roesch & Olson, 2003, 2004) reported a predictive code for reward magnitude and reward expectation but not for reward outcome or feedback. Previous studies have implicated the motor cortex in error-related signaling (Diedrichsen, Hashambhoy, Rane, & Shadmehr, 2005; Nir Even-Chen, Stavisky, Kao, Ryu, & Shenoy, 2015). Here, we show for the first time that single neurons in premotor and motor cortices encode reward-related feedback. Our finding adds another piece to the heterogeneity of reward representation in the midbrain and cortex, which will help extend future theories of reward-based learning.

The latency of the reward signal in PMd and M1 is on the order of 400–600 milliseconds. This latency is much slower than the rapid (~100 ms) reward prediction error signal observed in dopaminergic neurons in the midbrain (Dommett et al., 2005; Wolfram Schultz, 1998). Thus, the pathway to reward outcome representation in the motor cortex is likely to be mediated by the

basal ganglia-thalamo-cortical loop. In particular, we know the striatum, which receives projections from reward-sensitive dopaminergic neurons, feeds back to the cortex through other basal ganglia structures and the thalamus (Haber, 2011). The anterior cingulate cortex is also implicated in decision-making based on past actions and outcomes (Kennerley et al., 2006). This is an alternate possibility for the origin of the signal that we observe in premotor and motor cortices. Thus, it is likely that the motor cortex, along with prefrontal cortex and other areas, reflects rather than generates the reward signal.

At present, the function of this reward outcome signal in the motor cortices is unclear. A recent EEG-fMRI study (Fouragnan, Retzler, Mullinger, & Philiastides, 2015) suggests that two distinct value systems shape reward-related learning in the brain. In particular, they found that an earlier system responding preferentially to negative outcomes engaged the arousal-related and motor-preparatory brain structures, which could be useful for switching actions if needed. Therefore, the reward signal in PMd and M1 could potentially induce the cortical connectivity changes required for correcting subsequent motor plans based on mistakes. Further investigations of our finding might thus potentially reveal the mechanisms by which the brain acquires new motor skills.

Behavioral studies of motor control are at the advanced stage of describing trial-to-trial learning and generalization to novel contexts using sophisticated Bayesian decision theory and optimal control models (Galea et al., 2015; K. P. Kording & Wolpert, 2006; Krakauer & Mazzoni, 2011; R. Shadmehr et al., 2010; Daniel M Wolpert & Ghahramani, 2004). Yet, we are only beginning to understand how different neural systems work together to achieve these behaviors. We have shown a robust reward signal in premotor and motor cortex that is not simply the result of movement kinematics or planning. Establishing a link between this reward

signal and motor learning could potentially open up a new area of research within computational motor control.

METHODS

Single-neuron and population PSTHs

We calculated peri-stimulus time histograms (PSTHs) of firing rates (spikes/s), in 25-ms windows aligned to the reward timestamp and averaged them across trials. Error bars were computed as standard errors of mean across trials. To test whether neurons were significantly modulated by reward, we compared mean firing rate in a [0, 1.5] second interval after the reward timestamp across rewarded and unrewarded trials using a one-sided *t*-test, with a significance level of $\alpha = 0.05$, Bonferroni-corrected for the number of neurons recorded in a single session. To calculate population-averaged PSTHs, we took the mean trial-averaged PSTHs, normalized them to have a peak firing-rate of 1, and then averaged these across neurons. Error bars were computed as standard errors of mean across trials.

Generalized Linear Modeling: Temporal basis functions

We used raised-cosine temporal basis functions to model the latencies between environmental events, firing rates, and kinematics. We used 4 basis functions with equal widths of 400 ms, and equispaced from each other with centers separated by 200 ms. We convolved each covariate time series with its respective basis set and then used these to predict firing rates. To prevent discontinuities between trial epochs, we zero-padded each trial with 500 ms (i.e.,, 50 time bins of 10 milliseconds, each), concatenated them, convolved the zero-padded time series, and then removed the zero-padding.

Generalized Linear Modeling: Model fitting

We fit models using the Matlab glmnet package which solves the convex maximum-likelihood optimization problem using coordinate descent (Friedman, Hastie, & Tibshirani, 2010; Hastie, Tibshirani, & Friedman, 2009). To prevent overfitting, we regularized model fits using elastic-net regularization (Qian, Hastie, Friedman, Tibshirani, & Simon, 2013). We did not optimize the hyperparameters, but we found that a choice of $\lambda = 0.1$ (which determines the weight of the regularization term) and $\alpha = 0.1$ (which weights the relative extent of L_1 and L_2 regularization) resulted in comparable training and test-set errors, and therefore did not inordinately over-fit or under-fit the data. We also cross-validated the model by fitting it to one random half of the trials and evaluating it on the other half. To evaluate model goodness of fit, we computed the pseudo-R^2, which is related to the likelihood ratio. The idea of the pseudo-R^2 metric is to map the likelihood ratio into a [0, 1] range, thus extending the idea of the linear R^2 metric to non-Gaussian target variables. We used McFadden's definition of pseudo-R^2 (Fernandes et al., 2014; McFadden, 1973; Ramkumar et al., 2016). For each neuron, we computed bootstrapped 95% confidence intervals of the pseudo-R^2s.

Generalized Linear Modeling: Model comparison

To quantify whether individual covariates explain unique firing rate variance, we used partial models, leaving out the covariate of interest and comparing this partial model against the full model. To quantify this nested model comparison, we also used the relative pseudo-R^2 metric. We obtained 95% confidence intervals on this metric using bootstrapping, for each cross-

validation fold. We then treated the minimum of the lower bounds and the maximum of the upper bounds across cross-validation folds as confidence intervals. From these CIs, we approximated 2σ and 5σ significance levels, by calculating appropriate lower bounds for each significance level and comparing these lower bounds against zero.

SUPPLEMENTARY MATERIALS

Matching kinematics does not explain away reward signal

Motor cortical activity strongly predicts the kinematics of movement, such as its velocity and direction (Galea et al., 2015; Nikooyan & Ahmed, 2015). More confident or highly motivated reaches may be initiated earlier or performed faster, and may thus result in neural firing rate differences between successful and unsuccessful trials that are actually unrelated to reward. Two aspects of our task design could have potentially influenced variables such as confidence or motivation. First, even before movement onset, cue uncertainty can influence movement-planning confidence — the target location can be inferred with higher confidence from low-uncertainty cues than high-uncertainty cues. Second, after the trial completion, failure to secure reward might increase the urgency or desire to secure a reward in the next trial. Therefore, we asked whether kinematics— including latency of the return movement onset, instantaneous velocity, and acceleration —were systematically different between rewarded and unrewarded trials. Indeed, we found that monkeys started their return movement earlier for unrewarded trials, presumably because drinking the juice reward costs time and attention. The trial-averaged velocity and acceleration traces, when aligned to the time of reward (Fig 4.S1A, left panel), revealed small but significant differences between successful and unsuccessful trials. Therefore, any differences in neural firing rates related to reward might actually be attributed to

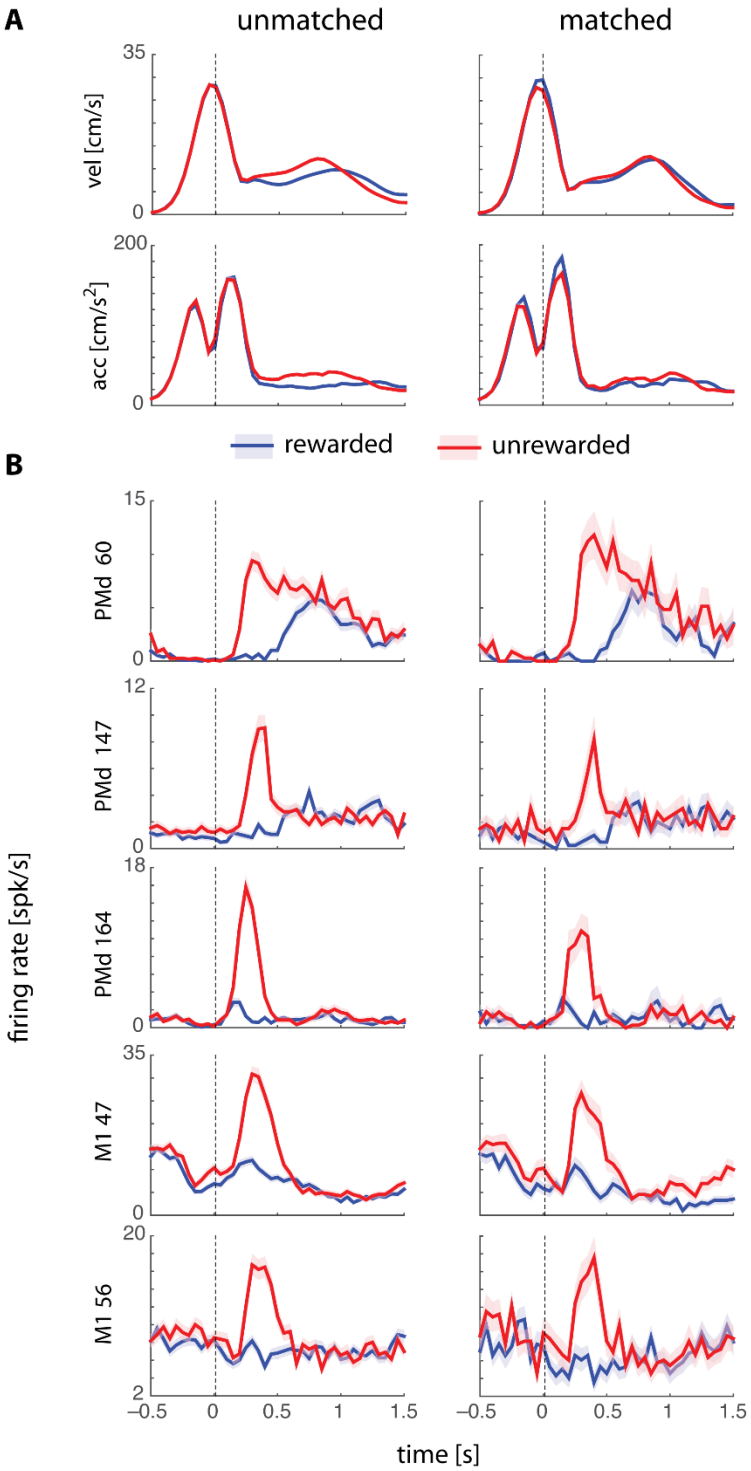

Figure S 4.1 | Velocity and acceleration control
Traces for rewarded (blue) and unrewarded (red) conditions averaged across all trials (unmatched) and across trials selected to have similar latencies and amplitudes for one representative session (Monkey M, session 4). The dashed line at zero represents the reward onset. (B) PSTHs of example neurons (mean ± SEM) from the same representative session for unmatched and matched conditions.session. The representative session for which histograms of effect sizes are shown above is indicated with a black border.

Table 4.1 | Number of trials before and after matching for kinematics

Monkey	Session	Before matching		After matching	
		n_rewarded	n_unrewarded	n_rewarded	n_unrewarded
M	1	368	217	43	102
M	2	344	260	39	144
M	3	183	122	53	50
M	4	426	265	112	98
M	5	412	244	50	80
M	6	194	145	55	76
T	1	254	104	40	68
T	2	122	83	4	25
T	3	121	84	7	55

kinematic differences. To control for these differences before analyzing the neural data, we matched the kinematics across conditions from 200 ms onwards after the reward (Fig 4.S1A, right panel) by selectively subsampling trials with similar return-movement peak velocities and latencies. For Monkey M, we included all trials where the peak velocity was between 11 and 16 cm/s, and the peak time was between 550 and 950 milliseconds. For Monkey T, we included all trials where the peak time was between 550 and 950 milliseconds, with no restriction on peak velocity.

Putative reward signal is not related to intrinsic success

Because all successful trials were rewarded, it was not possible to tell whether the motor system actually encoded a failure in accurate movement completion, or the absence of extrinsic reward. Therefore, in a separate control experiment, we attempted to make this distinction by including a set of catch trials during the low uncertainty condition. In these catch trials the monkey was not rewarded, despite being successful. By comparing PSTHs for regular trials with PSTHs for these catch trials, we found that firing rates increased for any unrewarded trial, whether or not the animal had been successful (Fig 4.S2), Therefore, PMd and M1 neurons encode the presence or absence of extrinsic reward, not intrinsic success or task outcome.

Reward signal was not related to prediction error

Dopaminergic neurons in the midbrain and prefrontal cortex are known to encode the mismatch between the magnitude of expected and obtained reward (Hollerman & Schultz, 1998). This difference is known as the reward prediction error, and is a useful learning signal in computational models of reinforcement learning (Abe et al., 2011; Shmuelof & Krakauer, 2011).

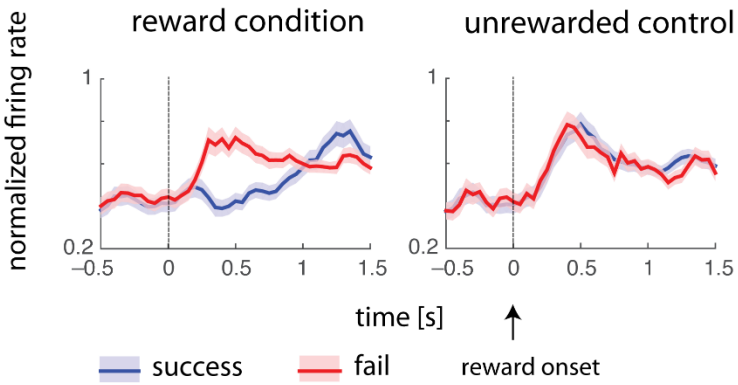

Figure S 4.2 | Tuning to extrinsic reward, not intrinsic success
PMd neurons increased their firing rates (trial-averaged, normalized, mean ± SEM) when no reward was received, for successful as well as unsuccessful trials. The dashed line at zero indicates the reward onset time

Therefore, we asked if the firing-rate increases in PMd and M1 resulting from the absence of reward might encode reward-prediction errors. Expected reward likelihood was indirectly related to the visual uncertainty of the target cue, which was displayed unpredictably in each trial by sampling from one of two distributions — a narrow one (low uncertainty) and a broad one (high uncertainty). Monkeys were more successful in low uncertainty trials and thus may have associated them with greater expectation of reward. A violation of this expectation would result in a reward-prediction error; a neural code for reward-prediction error magnitude would predict higher firing rates for successful high-uncertainty and failed low-uncertainty trials. To this end, we compared reward-aligned PSTHs for rewarded and unrewarded trials, separately for high and low uncertainty conditions. We found that unrewarded trials always had higher firing rates than did rewarded trials regardless of cue uncertainty (example neurons in Fig 4.S3). Thus, unlike midbrain dopaminergic neurons, motor cortical neurons do not encode reward-prediction errors.

Reward signal is distinct from return movement plan

The post-reward-related increase in firing rates could potentially be explained by the fact that consuming the reward costs time and attention. Since consuming the reward takes attention away from planning the return movement, it could be that the return reach plan is better attended to in the unrewarded trials. If this were the case, the apparent reward-related activity that we observed should be explained by the return movement plan. One test for this confound is to examine whether the putative reward activity is spatially tuned to the return reach direction. We tested for this possibility by constructing a spatio-temporal activity plot over the entire

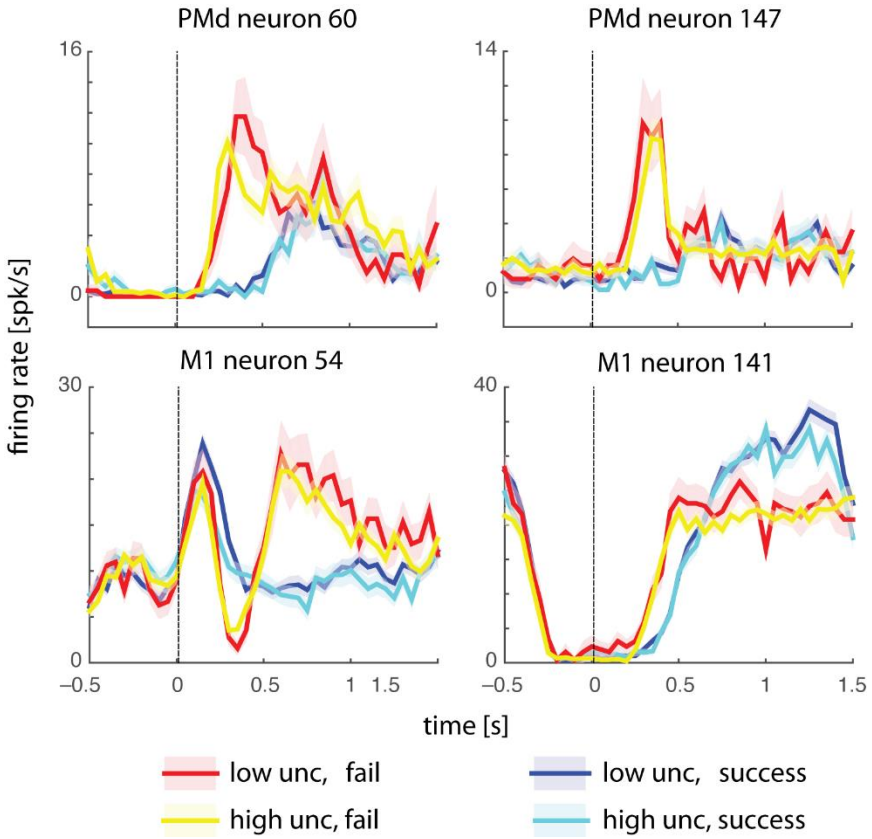

Figure S 4.3 | No effect of uncertainty on reward encoding
PSTHs aligned to the time of reward (dashed line at zero) for example neurons from PMd and M1. The averages are separated by both trial outcome (rewarded or unrewarded) and cue uncertainty (low vs. high). Firing rates were strongly modulated by reward but not uncertainty, suggesting that the code is not related to reward-prediction error.

population, in which activity was averaged across both neurons and trials, separated into bins by the difference between each neuron's preferred direction and the return-reach direction of the trial (Fig 4.S4). We used 64 uniformly spaced directional bin centers between $-\pi$ and π, and time bins of 25 milliseconds. Although we observed a slight increase in the activity profile (Fig 4.S4, top and middle panels: 200–500 ms after reward in PMd and 500–800 ms in M1) for return reaches toward the preferred direction, there was no significant difference between rewarded and unrewarded trials (Fig 4.S4, bottom panels). These results suggest that the return reach plan is independent of the reward signal.

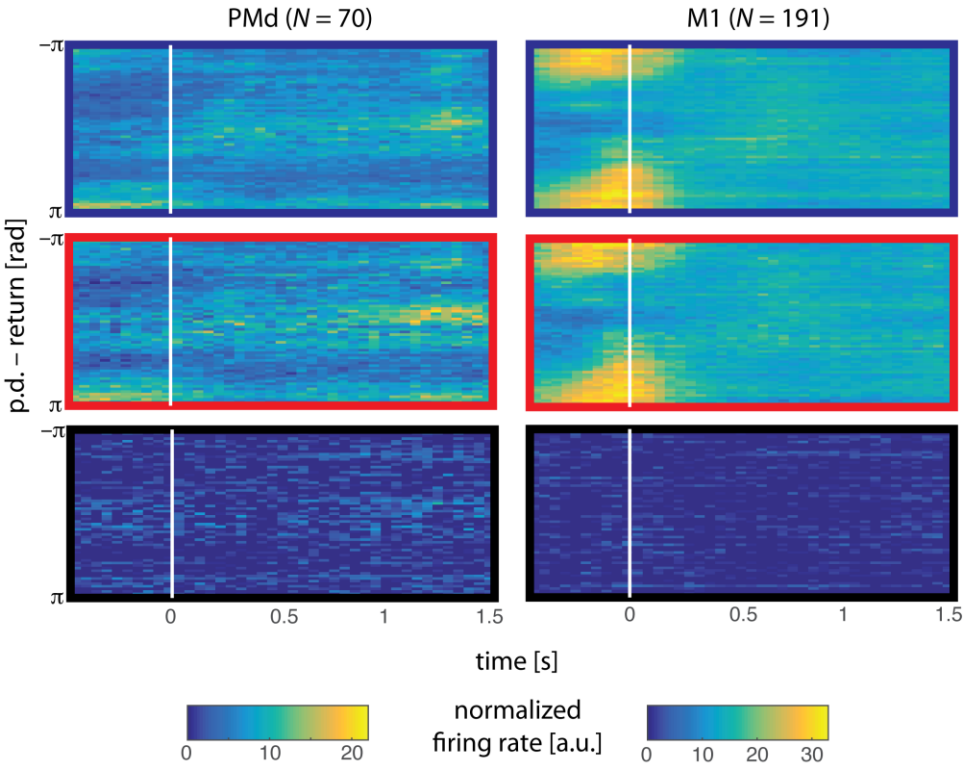

Figure S 4.4 | Reward signal was not spatially tuned
Single session, monkey M: A spatiotemporal activity profile over the entire population, separated by the circular distance between the upcoming return reach direction and each neuron's preferred direction, averaged across trials. Blue border: rewarded trials, Red border: unrewarded trials, Black border: difference. An increase in activity profiles for upcoming return reaches in a given neuron's preferred direction for both classes of trials is indicative of the spatial tuning for the return reach plan. No significant differences were found between rewarded and unrewarded trials, suggesting that the spatially-tuned return reach plan was independent of reward encoding.

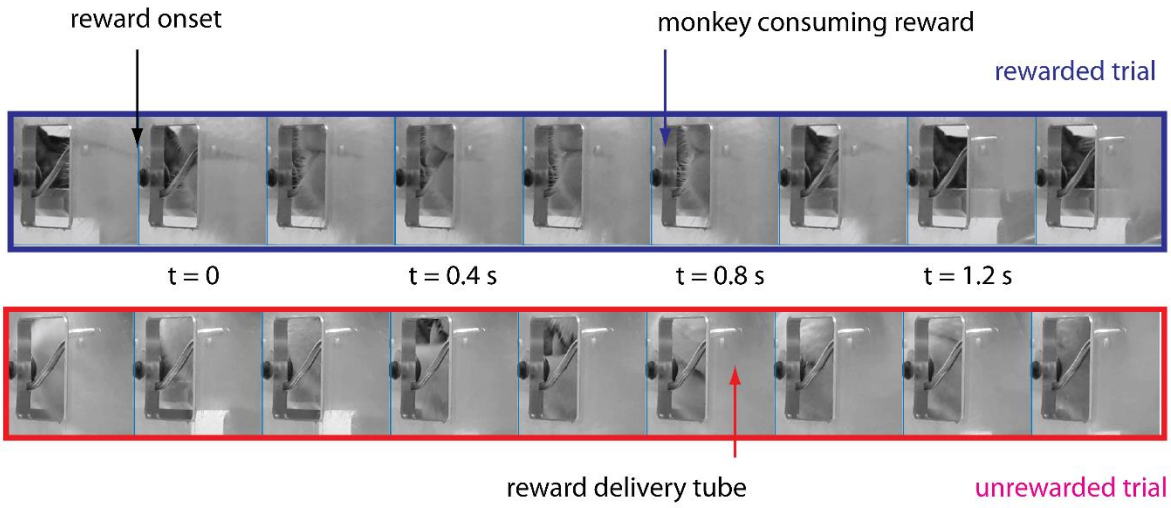

Figure S 4.5 | No mouth movements in unrewarded trials
The increase in firing rate after lack of reward could potentially be explained by increased mouth or neck movement activity, e.g.,if the monkey would suck harder at the tube. We filmed the monkey's behavior after rewarded (example above) and unrewarded trials (example below) but found that in general, the monkey did not make contact with the tube after unrewarded trials.

CHAPTER 5: DISCUSSION

SUMMARY

In Chapter 1, I provided an overview of the elements of motor control related to uncertainty, namely the execution of movement under noisy and/or ambiguous conditions, as well as how outcomes of those movement can drive motor learning. In Chapters 2 and 3, I presented work that aimed to characterize the roles of dorsal premotor cortex (PMd) and primary motor cortex (M1) in the planning and execution of reaching movements when faced with two main types of uncertainty: (1) noise, and (2) ambiguity. I found that when faced with noisy information about the location of a reach target, monkeys, like humans (Konrad P Kording & Wolpert, 2004; Tassinari et al., 2006), integrate both sensory and prior information to obtain nearly optimal task performance (see Chapter 2). During planning and execution, target representations within PMd were strongly modulated by the magnitude of the monkeys' subjective feelings of uncertainty about their choices (as estimated from behavior). Surprisingly, this uncertainty-related component of neural activity did not resolve prior to action initiation. Rather, after emerging quickly upon trial onset, it persisted throughout the movement. M1 activity, on the other hand, did not change with uncertainty, suggesting separate roles of PMd and M1 in the transition from action selection to action execution.

From these results, I then moved to the question of target ambiguity (Chapter 3). Specifically, does uncertainty in a decision between two discrete targets induce the same neural effects in PMd as does uncertainty about a single noisy target? To address this question, I trained monkeys to perform a reaching task in which I first presented two opposing targets before indicating which was correct. According to a well-respected model (P. Cisek, 2007), the initial

presentation of two targets should elicit simultaneous representations in PMd. However, I found that population activity on single trials only ever indicated one movement plan at a time. Furthermore, I showed how preferences for certain targets—a commonly observed phenomenon—can lead to a spurious finding of simultaneous representation when using the trial-averaged analysis methods of the prior study. These results show that PMd does not reflect the uncertainty caused by choice ambiguity, and likely has a different role in motor decision-making than previously thought.

Finally, in Chapter 4, I presented evidence of activity in PMd and M1 that reflected the signaling of task failure (i.e., no reward). These areas are normally—and understandably—studied during movement planning and execution. However, the motor cortex is also highly adaptive, and can quickly learn to control movement as circumstances change (e.g., as a result of altered limb dynamics or modified sensory feedback). This adaptability indicates a direct influence of task-related feedback on the motor system. In support of this, we found that neurons in both PMd and M1 displayed a transient response in activity following a task failure cue. Although we can only speculate on the corresponding mechanism, our result shows evidence that motor cortex receives input regarding task-specific performance. In the following sections, I will discuss the implications of results from Chapters 2, 3, and 4 both with respect to existing literature and in relation to each other.

PMD AND M1 DURING DECISION-MAKING

Overview

As stated in the title of a 2012 paper by Wolpert and Landy (D. M. Wolpert & Landy, 2012), "motor control is decision-making". Each movement reflects the output of complex processes incorporating multisensory cues, anticipated risks and rewards, and prior experiences—all of which are plagued by noise and uncertainty. Knowledge of how the motor system arrives at a single executed movement plan under such uncertainty is central to understanding motor control in general. In Chapters 2 and 3 I presented results from PMd and M1 during the planning and execution of reaches under two different types of uncertainty (environmental/sensory noise and ambiguous target choice). M1 displayed no uncertainty-related changes in activity, and appeared to encode the executed reach reliably under all uncertainty conditions. In contrast, PMd was strongly affected by sensory noise, but not target ambiguity, for the duration of movement planning and execution. In the following sections, I will discuss how we might incorporate these results into a cohesive view of PMd and M1 as elements of a motor decision-making neural process.

Planning movements under noise-related uncertainty

In order to hypothesize about the neural mechanisms underlying movement generation in the presence of uncertainty, we first need to understand how motor cortex actually "represents" movement at the most basic level. This is still an open question in motor neuroscience. Traditionally, motor cortex activity has been analyzed from a parameter-centric viewpoint, in which the firing rates of individual neurons correspond to values of movement-related variables (e.g., hand velocity, muscle activation). Even in low-uncertainty conditions, tuning properties

tend to be broad, meaning that neurons exhibit gradated activity with movement parameters (e.g., reach direction). A neuron that is strongly active during rightward reaching will also be active—although less so—for neighboring reach directions. As a result, the entire neural population can be interpreted as containing weighted representations of all possible movements. The same concept naturally extends to more uncertain conditions: higher uncertainty leads to a broader distribution of represented movements.

If motor cortex—particularly PMd—represents the full distribution of potential movements, what then determines the singular, executed motor decision? A prominent model of motor cortex places PMd as the primary motor decision-making area (P. Cisek, 2006, 2007; Thura & Cisek, 2014). In this model, a movement plan develops in M1 only after PMd arrives at a final decision. However, the results from Chapter 2 contradict this hypothesis, as PMd activity reflected uncertainty in the reach goal throughout planning and execution. The fact that PMd did not arrive at a final, de-noised movement representation prior to (or even after) movement onset suggests that motor decision-making does not take place exclusively within PMd.

Although uncertainty leads to persistent effects in PMd, it does not appear to affect M1 directly at any point during movement planning or execution (Chapter 2). These two observations together suggest that movement decision-making—that is, the final specification of a single movement in cortex—occurs in the connections between PMd and M1. This could perhaps result from differing levels of lateral inhibition in the two areas. Lateral inhibition describes the inhibitory effect that an excited neuron can have on neighboring neurons. This allows for a "sharper", more specific representation by preventing incidental activation of nearby but functionally dissimilar neurons. Although more commonly described in sensory brain areas, lateral inhibition is present in motor cortex as well (Asanuma & Rosén, 1973). The strength of

lateral inhibition in motor cortex is modulated via input from the basal ganglia, with separate regions projecting to primary motor and premotor cortices (S. Beck & Hallett, 2011). This independent control of the allowable specificity in movement representation raises the possibility that movement plans become fully defined during the transition from PMd to M1 via increased lateral inhibition in the latter.

A population-based interpretation of uncertain movement representations

More recently, some neuroscientists have rejected the traditional parameter-based views in favor of a "population dynamics" interpretation of cortex (Ames et al., 2014; M. M. Churchland et al., 2012; M. M. Churchland, Cunningham, Kaufman, Ryu, & Shenoy, 2010; Michaels et al., 2016). This view stresses the influence of recurrent connections within neural populations, that result in constrained, population-wide patterns of correlated activity which can be driven by external inputs. Under this interpretation, there is no explicit relationship between individual neurons and specific movement parameters. Movement planning instead involves setting the initial state of the population that will, following a largely constrained dynamical evolution, generate the temporal patterns of cortical activity required to control muscles (and/or spinal circuits). How this translation from planning to execution might occur under uncertain conditions is not obvious. With the parameter-based approach, uncertainty could be described as causing a broadening of (i.e.,, greater uncertainty in) the distribution of represented movements. Under the population dynamics hypothesis, the same intuition does not hold, since movements are not "represented" in the first place. However, we can reframe "representation" in the dynamical context to mean the initial population state that ultimately produces a given

movement. In low uncertainty situations, the population achieves an initial state that is highly predictive of the movement, and so can also be said to accurately "represent" that movement. Under high uncertainty, the population might not achieve such well-specified initial states and instead settle on intermediate states that unfold somewhat unpredictably upon movement execution. The resulting variability in motor output might even contribute to motor learning in high uncertainty situations (Herzfeld & Shadmehr, 2014).

As with the representation-based interpretation, finalization of the motor plan from a dynamic system viewpoint likely occurs only during the transition from planning to execution. Dimensionality reduction-based analysis methods of population activity suggest that movement preparation and movement execution correspond to different low-dimensional, orthogonal subspaces (Elsayed et al., 2016; Stavisky, Kao, Ryu, & Shenoy, 2017). That is, activity constrained to one correlation structure is used during planning, while a separate but functionally linked correlation structure (within the same neural population) is used during execution. Upon movement initiation, the neural state (that component set within the planning subspace) unfolds into the execution subspace to produce movement. Adding uncertainty to a movement plan might affect activity within the planning subspace such that the transformation to the execution subspace is less predictable. Even so, these uncertain initial states still elicit precise motor-related dynamics, as evidenced by the lack of uncertainty-related effects on M1 or movement kinematics (Chapter 2). Regardless of whether one adopts a parameter- or dynamics-based perspective of motor cortex, the production of a single executed movement from an uncertain plan likely results from the inherent clarification that occurs in the transition between planning (i.e., PMd-dominated) and execution (i.e., M1-dominated).

Motor plan transformation from PMd to M1

In the previous sections I suggested that final movement decisions are not completed within a planning phase, but rather emerge from a forced transition between planning and movement domains. The main role of PMd might then be to simply transform task-relevant inputs (e.g.,multisensory cues, prior knowledge) into a motor-planning context. This view is consistent with the apparent mixed selectivity of PMd neurons (Batista et al., 2007; M. M. Churchland, G. Santhanam, et al., 2006; Gail et al., 2009; McGuire & Sabes, 2009; Messier & Kalaska, 2000; Pesaran et al., 2006; Schaffelhofer & Scherberger, 2016; L. Shen & Alexander, 1997), which includes both external goal coordinates and internal, execution-level movement representations. Further evidence can be found in the time-varying nature of PMd neural responses, which suggests some change in coordinate system even within a single reach plan. For example, cells exhibit strong, short-latency bursts of activity that appear to represent the locations of visual stimuli, but quickly shift to represent the associated movements. As a more extreme example, as monkeys learn to associate certain spatial cues with oppositely directed reaches (e.g.,an upward target that indicates a downward reach), PMd first appears to encode the cue location, then later the reach direction (Klaes et al., 2011). This is consistent with the view that PMd is involved with translating behaviorally relevant cues into an appropriate motor context. If the main function of PMd is indeed as a simple "translator" of the motor system, then it is expected that uncertainty present in sensory inputs will exist in the motor plan as well.

PERSISTENCE OF UNCERTAIN REACH PLANS IN PMD

Overview

The previous discussion of action finalization addressed the obstacle that persistent uncertainty-related effects in PMd may pose for movement execution. I proposed that the uncertainty in a movement plan might be filtered out during the transition from planning to execution. This explains how movements can be executed in the absence of a well-developed plan, but does not address why uncertain reach plans persist throughout movement in PMd. It is possible that this persistence serves no function, but is simply a side effect of separate mechanisms controlling movement planning and movement initiation (Haith et al., 2016; A. L. Wong, Goldsmith, Forrence, Haith, & Krakauer, 2017). It is also possible that maintaining an uncertain plan contributes to motor control in some way. In the following sections, I present two possible contributions: error correction and/or motor learning.

Error correction

Many of the reaching tasks used to study motor cortex (e.g., the center-out paradigm) involve ballistic movements. Ballistic reaches—especially when following an instructed delay period—are preferable experimentally because they can be completely separated into movement planning and open-loop movement execution components. This simplified motor behavior limits the number of potential confounds for analysis purposes, but is not representative of motor control in general. In most situations, we do not simply plan movements and then execute them in a fully feedforward manner. Rather, we make series of complex, continuous movements and use feedback to correct them.

The reach plans observed in PMd might persist throughout movement to aid in correcting the ongoing movement. As discussed above and in Chapter 1, PMd appears to have a role in transforming visual cues into motor commands. It strongly connects with M1 and contains significant corticospinal projections, suggesting both indirect and direct influences on movement execution. These properties make it a likely candidate for aiding in online movement correction. In support of this view, a study in human subjects showed that disrupting PMd via transcranial magnetic stimulation (TMS) greatly reduced error correction during a visuomotor adaptation task (Lee & van Donkelaar, 2006). Thus, the persistence of reach plans throughout movement may not reflect "persistence" of a previous motor plan, but rather an ongoing process for modifying the action if necessary.

If PMd does indeed contribute to error correction, what might the results from Chapter 2 suggest about the impact of uncertainty on this process? Unfortunately, the experimental approach used in Chapter 2 was not designed to address error correction directly, so the results are ill-suited for developing a specific model. However, the general observation that uncertain stimuli induced uncertain (i.e., broad or poorly-specified) representations in PMd provides some insight into previous behavioral observations. For example, subjects reaching to an uncertain target location can quickly adjust their reaches if the target suddenly jumps to a new location (Izawa & Shadmehr, 2008). The same corrective movements are delayed and smaller when the initial reach is toward a more certain target. This behavioral phenomenon might be a direct result of uncertainty-related effects in PMd. The broadened representations in PMd may reflect something akin to contingency plans which can be quickly executed if needed.

Motor learning

Another possible role of persistent uncertainty-related movement plans in PMd is in guiding motor learning. While M1 is largely studied for its role in movement execution, it is also highly plastic (Sanes & Donoghue, 2000). Motor learning generally requires knowledge of both the intended movement and its outcome. Reach plans in PMd might therefore persist throughout movement in order to interact with subsequent result-dependent signaling. I will discuss this possibility—and the potential relationship with the failure-based signaling described in Chapter 4—in greater depth in a following section.

If PMd does indeed maintain movement representations for the purpose of motor learning, then we should expect uncertainty in those representations to affect learning in some way. The experimental setup in Chapter 2 was not designed to track learning, so the results do not provide direct evidence that uncertainty in PMd affects the rate of motor learning. However, the results do show that PMd contains information about the subjective level of uncertainty in the executed movement. There is also ample behavioral evidence that sensory uncertainty slows the rate of motor learning (Wei & Kording, 2010). I suggest that these two observations might be linked; that PMd activity at the time of movement completion directly contributes to functional changes within motor cortex.

EFFECTS OF NOISE BUT NOT AMBIGUITY IN PMD

The results from Chapter 3 show that when monkeys were confronted with two potential reach targets, PMd activity only ever reflected a motor plan to one of them. This directly contradicts previous studies describing a parallel motor encoding mechanism in PMd during

similar tasks (P. Cisek & Kalaska, 2005; Coallier et al., 2015; Klaes et al., 2011; Pastor-Bernier & Cisek, 2011). As mentioned in Chapter 3 it may be the case that the single-neuron and trial-average based analyses used in these studies were ill equipped to discriminate between parallel encoding and single reach encoding when the monkeys had intrinsic target selection biases. Although it is impossible to rule out that the difference was a result of task or monkey differences between the current and former studies, the results in Chapter 3 show no evidence of simultaneous representation.

The apparent lack of simultaneous reach representations in PMd was surprising, given not only the previous studies, but also the results from Chapter 2. When faced with the task of estimating target location from noisy information, the level of uncertainty clearly affected the representation in PMd. Specifically, we concluded that under high uncertainty, PMd represents a broad distribution of potential movements (or in the dynamics-based interpretation, a flexible and ill-defined initial state). It seems only logical, then, that when faced with the ambiguity of multiple targets, PMd might contain a multimodal distribution reflecting the potential movements. The fact that we did not observe this may help to explain the type of motor processing performed in PMd—i.e.,, its position in the sensorimotor pipeline.

The term "uncertainty" is commonly used to describe both noise and ambiguity. However, these have different implications for the development of a motor plan. In the case of noise-related uncertainty (as in Chapter 2), there is no single sensory cue that can be used as a reference to specify the movement goal. Instead, the sensorimotor system must use the available information (sensory cues, prior knowledge, idiosyncratic preferences, etc.) to generate the motor plan from the ground up. Alternatively, planning under target ambiguity involves a more top-down type of processing. Although the choice between options may be unclear, the

alternative movements themselves are well-defined. Once a choice is made, the motor planning aspect is drastically simplified. Thus, the two cases differ mainly in the quality of final action specification: noise forces the motor system to develop an internally-generated movement plan, while ambiguity ultimately transposes into a goal-directed movement. Brain areas responsible for low-level movement execution are unlikely to respond differently during single and multi-target tasks.

Following from the above reasoning, it is possible that we did not observe simultaneous representation in Chapter 3 because PMd is further downstream in the decision-making process than is commonly believed. The current view—which was developed from the simultaneous encoding result—is that movement-related decisions are made directly within motor structures rather than high level, cognitive brain areas (Christopoulos et al., 2015; P. Cisek, 2006, 2007; Pastor-Bernier & Cisek, 2011; Thura & Cisek, 2014). However, this might overstate the scope of PMd. PMd is ultimately an output-level brain area. While its response to visual targets and its role in movement planning do suggest some degree of higher-order processing, it seems unlikely that PMd is responsible for every type of movement related decision. The lack of evidence for simultaneous reach representations instead suggests that PMd is involved with the translation of high-level goals into movement, and not the decision-making processes that define those high-level goals.

Another possible reason why PMd did not exhibit simultaneous encoding might be that the monkeys' behavioral strategies simply did not require it. The simultaneous encoding hypothesis relies heavily on an assumption—albeit a plausible one—of how one *should* complete a two-target task. To experimenters, simultaneous encoding is an attractively simple mechanism: the brain equally represents all possible actions, and then those representations compete (biased

by the available information) until only one is left. However, consider the basic structure of the task used by Cisek and Kalaska (2005). Two targets appear onscreen, followed by a cue to indicate which one is correct. This task can be performed with only a single reach plan (to one of the targets) combined with a matching rule to dictate whether to switch (i.e., does the color of the cue match the color of the planned target?). The monkeys' behavior in our experiments suggests such a strategy. Admittedly, the task in Chapter 3 didn't explicitly require the monkeys to memorize the locations and features of the two targets. However, if this memorization component is essential for eliciting simultaneous plan representations in PMd, then we cannot conclude that such simultaneous representation is a general neural mechanism of decision-making. No matter the reasoning behind the difference in results (task demands, monkey preferences, etc.), the results from Chapter 3 argue strongly against the interpretation that simultaneous representation reflects a ubiquitous aspect of ambiguity-related decision-making—if it exists at all.

TASK OUTCOME AND MOTOR CORTEX

External versus internal outcome signaling

The results from Chapter 4 provide evidence for transient, result-based (no reward) signaling in both PMd and M1. Reward and failure can alone drive motor learning in the absence of continuous visual feedback or even endpoint error (Izawa & Shadmehr, 2011; Nikooyan & Ahmed, 2015). The presence of some type of outcome signal in motor cortex is therefore not entirely surprising. However, it is not clear why the effect was limited to just the absence of external reward. A recent study also found outcome-dependent changes in motor cortex while

monkeys used a brain-computer interface (BCI) to move a cursor to different targets on a screen (N. Even-Chen, Stavisky, Kao, Ryu, & Shenoy, 2017). However, they concluded that the increases in activity were driven by error rather than task failure or lack of external reward. These two different results suggest that outcome-related signaling is a flexible mechanism that can be modified for different task constraints.

The main question surrounding the result from Chapter 4 and that shown by Even-Chen et al. is whether the outcome-related signals are based on external rewards or an internal evaluation of movement. Chapter 4 suggests an external signal source (reward) while Even-Chen et al. suggest an internal source (error in cursor location). This can potentially be explained by a difference in the sources of errors for the two tasks. In the task from Chapter 4, monkeys made simple, straight reaching movements that were well-practiced over multiple sessions and over years of natural arm control. Thus, is unlikely that there was ever significant deviation between their intended movements and their executed movements (i.e., internal reach errors). We can therefore assume that failure on the task largely resulted from the monkeys' misunderstanding of the task, rather than an inability to accurately control movement. Conversely, the BCI control paradigm used by Even-Chen et al. naturally involves more "motor" error, since motor intent cannot be perfectly decoded from neural activity. The main source of failure was therefore an inability to accurately control movement, rather than confusion about the task itself. This difference in the source of error and task performance might explain the two different types of outcome signals. When poor task performance is driven by a misunderstanding of the task, the outcome-based signal in motor cortex depends on external reward (or lack thereof). When poor task performance is caused by inaccurate motor control, it depends on internal errors. This

hypothesis provides a physiological basis for the claim that the estimated source of task error impacts motor learning and generalization (Berniker & Kording, 2008).

Potential utility of outcome-based signaling

One reason why motor cortex might respond strongly to failure is because of the potential need for a corrective movement. The increased activity we observed in motor cortex following failed trials may reflect heightened alertness meant to facilitate the formation of a new movement plan. This view is supported by an imaging study in humans (Fouragnan et al., 2015), which provided evidence for a fast, failure-dependent increase in brain areas tied to both movement planning and arousal. Due to the design of the behavioral task used in Chapter 4 (which was the same as Chapter 2), there was no opportunity to correct or re-attempt a failed movement, yet we observed failure-based responses nonetheless. It seems likely then that these signals reflect a subconscious, automatic mechanism of the reward-motor circuitry, and not an explicit or conscious strategy.

Alternatively, increased motor cortical activity following trial failure might indicate a mechanism for driving motor learning. The ultimate purpose of the sensorimotor system is to make useful and beneficial interactions with the environment. Failure indicates that the motor system should modify its behavior to improve the chance of reward in future situations. This motivational component might rely on the connections between motor cortex and the reward-related basal ganglia. Recordings from the substantia nigra of monkeys indicate that the strength of responses from dopaminergic neurons following reward is larger when the reward is unexpected (Hollerman & Schultz, 1998). This reward-related feedback is also tied to greater

retention of a learned motor behavior (Fouragnan et al., 2015). Alternatively, punishment results in faster learning rates during motor adaptation (Galea et al., 2015), indicating that failure-related feedback can facilitate changes within motor areas. It is beyond the scope of this thesis to hypothesize about the specific mechanisms of motor learning at the level of individual neurons. However, I propose that motor cortex (PMd, M1, and potentially other motor areas) might receive failure-based signaling from the basal ganglia to increase plasticity and enhance motor learning through functional reorganization.

Interaction between uncertainty and task outcome

The modeling approach used in Chapter 4 to characterize neural responses did not reveal any significant effect of uncertainty (as estimated from cue dispersion; see Chapter 4) on the magnitude of failure-related responses. This is somewhat surprising, since failure should be more unexpected on trials with lower uncertainty and therefore likely to induce a greater response in motor cortex. This expectation follows from current knowledge of reinforcement learning, which suggest that the magnitude of result prediction error (that is, the unexpectedness of success or failure) influences learning rate (Galea et al., 2015; Nikooyan & Ahmed, 2015; Schonberg, Daw, Joel, & O'Doherty, 2007). Learning is disincentivized when uncertainty is high to prevent overcorrection or over-adaptation to uninformative noise.

It is possible that we found no effect of uncertainty on the failure-based response because the metric did not accurately capture the true level of uncertainty. We defined "uncertainty" as the dispersion in the visual cue (refer to Chapter 2). However, this estimate was likely a poor proxy for the monkeys' actual subjective assessment of uncertainty. The results from Chapter 2

show that the amount of uncertainty in a reach plan cannot be gauged from task parameters alone. To gain a more accurate estimate would require either evaluation of behavioral responses (such as the endpoint variance metric used in Chapter 2) or inference directly from planning-related neural activity.

Even if the magnitude of the failure-based response does not change with movement uncertainty, it is still possible that there exists some interactive effect. There is evidence that failure-based signaling in premotor areas inhibits the consolidation of motor memories (Fouragnan et al., 2015), presumably to avoid similar errors in the future. High uncertainty reach plans in PMd might cause similar effects, even for successful movements. Chapter 2 shows that high uncertainty principally causes an increase in PMd activity, which persists throughout movement. While likely originating from separate sources—failure signaling suggests input from serotinergic midbrain structures (Fouragnan et al., 2015)—both failure and high uncertainty appear to have similar effects on PMd. Thus, even if a highly uncertain movement is successful, the corresponding increase in activity caused by the uncertainty might act to interrupt dopaminergic reinforcement mechanisms. While speculative, this kind of mechanism could explain the slow learning rates observed under high uncertainty conditions (Burge, Ernst, & Banks, 2008; Wei & Kording, 2010).

LIMITATIONS AND SUGGESTED FUTURE DIRECTIONS

Overview

While the results presented here provide new insight into the contributions of PMd and M1 to motor control and motor learning, there are a number of limitations that should be

addressed by future work. In the following sections, I identify those limitations and discuss what I believe are deficits in both the experimental and analytical approaches used. I also offer suggestions for possible directions of future work that might improve and/or extend the scope and impact of these initial findings.

Noise-related uncertainty and PMd – Chapter 2

The examination of noise-related uncertainty during reach planning (Chapter 2) was founded on a traditional, parameter-based view of cortex. These approaches are not ideal for large, simultaneously recorded populations of neurons. The tuning curves and preferred directions associated with these analyses can (and do) change throughout the course of a trial, and many neurons must be discarded because they do not have the required properties (e.g.,non-cosine tuning). For these reasons, and perhaps others, we were unable to identify any single-trial effects of uncertainty. Future work should reexamine the data collected from Chapter 2 using modern, population-based analyses. It is possible that a dimensionality reduction technique similar to the one used in Chapter 3 (Byron et al., 2009; Kaufman et al., 2015) might provide an accurate moment-by-moment estimate of uncertainty. This would open the possibility of tracking the effect of learning on uncertainty, or characterizing how uncertain movement decisions develop on individual trials.

In addition to these analytical considerations, the results from Chapter 2 could also be clarified and extended through supplementary experiments. One of the most compelling findings was that uncertainty-related effects in PMd persisted throughout movement. However, the interpretation of this result is unclear. As noted, it might indicate a mechanism for movement correction or for aiding motor learning. A future experiment incorporating a target jump (a la

Izawa and Shadmehr) could reveal a correlation between reach plan uncertainty (determined directly from PMd) and properties of the corrective movement. Similarly, a motor learning task could be used to examine how the uncertainty of a single movement might affect the extent of adaptation on a subsequent trial.

Ambiguity and decision-making – Chapter 3

The results from Chapter 3 argue against the view that PMd represents all potential movement plans simultaneously. This contradicts the results from previous experiments (Christopoulos et al., 2015; P. Cisek & Kalaska, 2005; Coallier et al., 2015; Klaes et al., 2011; Thura & Cisek, 2014) with similar—but not the same—task designs. Future experiments should attempt to fully replicate the structure of those tasks, but incorporate simultaneous population recordings in PMd to track the moment-by-moment reach plans.

Beyond the representation of single versus multiple reaches, the results also indicate a strong influence of target preferences (biases) on decision-making in ambiguous situations. However, future experiments are needed to illuminate the origins and strengths of those biases. The task could be altered to provide unequal rewards to manipulate the monkey's target bias, and to determine how such a reward structure and altered bias might change the neural response. Providing higher rewards for initially non-preferred targets would undoubtedly increase the proportion of reaches to those targets. However, would the neural responses immediately following target appearance change to reflect the new target preferences? It is possible that the monkey would simply learn to switch to those plans mid-trial, without any change to the short-latency neural responses thus revealing... Similar experiments could also help quantify the subjective values attributed to preferred targets. For example, free-choice responses (or reach

plans decoded from neural activity) could be used to modify the reward structure continuously until behavior suggested equal preference for all targets. The associated rewards would indicate the strength of the initial preference.

Outcome signaling – Chapter 4

As in Chapter 2, the results from Chapter 4 were based on the behavior of single neurons. Population-wide analytical methods may be able to illuminate trial-by-trial differences in failure signaling that were not visible in individual neurons. Also, the potential interaction between uncertainty and the failure signal should be reevaluated using an improved estimate of uncertainty based on behavior and/or neural markers.

The main limitation of Chapter 4, however, is that the task was not designed to test specifically for failure signaling. Additional experiments should be performed to directly address the various hypothetical implications of a failure signal: (1) facilitation of corrective movement planning, (2) increased plasticity and (3) blocked memory consolidation. The effect on corrective planning could easily be tested by presenting a second, unexpected target on some trials. Correlation between the size of the failure signal and some aspect of the secondary movement (e.g., reaction time or velocity) would provide convincing evidence that it facilitates motor plan switching. Alternatively, a standard motor adaptation task (e.g., visual rotation) with a condition providing only outcome-related feedback (Izawa & Shadmehr, 2011) could be used to address whether the failure signal observed in cortex directly affects motor learning.

CONCLUSION

Planning and executing movements is rarely a trivial matter. The sensorimotor system must account for various sources of sensory and intrinsic noise, consider multiple potential actions, and continuously update its understanding of how to interact with a stochastic and dynamic environment. The work here provides initial steps for understanding how two components of the motor system—PMd and M1—contribute to those tasks. Movement plans in PMd (but not M1) are sensitive to noise-related uncertainty in the reach goal, an effect that persists throughout movement execution. Upon movement termination, both PMd and M1 respond to failure with bursts of activity. Together these results suggest an important role of PMd and M1 in planning and executing uncertain movements, and learning from task outcomes.

REFERENCES

Abbott, A. (2006). Neuroprosthetics: In search of the sixth sense. *Nature, 442*(7099), 125-127.

Abe, M., Schambra, H., Wassermann, E. M., Luckenbaugh, D., Schweighofer, N., & Cohen, L. G. (2011). Reward improves long-term retention of a motor memory through induction of offline memory gains. *Curr Biol, 21*(7), 557-562. doi:10.1016/j.cub.2011.02.030

Acerbi, L., Vijayakumar, S., & Wolpert, D. M. (2017). Target Uncertainty Mediates Sensorimotor Error Correction. *PLoS One, 12*(1), e0170466. doi:10.1371/journal.pone.0170466

Afshar, A., Santhanam, G., Yu, B. M., Ryu, S. I., Sahani, M., & Shenoy, K. V. (2011). Single-trial neural correlates of arm movement preparation. *Neuron, 71*(3), 555-564. doi:10.1016/j.neuron.2011.05.047

Alais, D., & Burr, D. (2004). The ventriloquist effect results from near-optimal bimodal integration. *Curr Biol, 14*(3), 257-262. doi:10.1016/j.cub.2004.01.029

Ames, K. C., Ryu, S. I., & Shenoy, K. V. (2014). Neural dynamics of reaching following incorrect or absent motor preparation. *Neuron, 81*(2), 438-451. doi:10.1016/j.neuron.2013.11.003

Asanuma, H., & Rosén, I. (1973). Spread of mono-and polysynaptic connections within cat's motor cortex. *Experimental brain research, 16*(5), 507-520.

Ashe, J., & Georgopoulos, A. P. (1994). Movement parameters and neural activity in motor cortex and area 5. *Cerebral cortex, 4*(6), 590-600.

Basso, M. A., & Wurtz, R. H. (1997). Modulation of neuronal activity by target uncertainty. *Nature, 389*(6646), 66.

Bastian, A., Riehle, A., Erlhagen, W., & Schöner, G. (1998). Prior information preshapes the population representation of movement direction in motor cortex. *Neuroreport, 9*(2), 315-319.

Bastian, A., Schoner, G., & Riehle, A. (2003). Preshaping and continuous evolution of motor cortical representations during movement preparation. *European Journal of Neuroscience, 18*(7), 2047-2058. doi:10.1046/j.1460-9568.2003.02906.x

Batista, A. P., Santhanam, G., Yu, B. M., Ryu, S. I., Afshar, A., & Shenoy, K. V. (2007). Reference frames for reach planning in macaque dorsal premotor cortex. *J Neurophysiol, 98*(2), 966-983. doi:10.1152/jn.00421.2006

Battaglia, P. W., Jacobs, R. A., & Aslin, R. N. (2003). Bayesian integration of visual and auditory signals for spatial localization. *Josa a, 20*(7), 1391-1397.

Bayes, T. (1763). An essay towards solving a problem in the doctrine of chances.

Beck, J. M., Ma, W. J., Kiani, R., Hanks, T., Churchland, A. K., Roitman, J., . . . Pouget, A. (2008). Probabilistic population codes for Bayesian decision making. *Neuron, 60*(6), 1142-1152. doi:10.1016/j.neuron.2008.09.021

Beck, S., & Hallett, M. (2011). Surround inhibition in the motor system. *Experimental brain research, 210*(2), 165-172.

Berkes, P., Orbán, G., Lengyel, M., & Fiser, J. (2011). Spontaneous cortical activity reveals hallmarks of an optimal internal model of the environment. *Science, 331*(6013), 83-87.

Berniker, M., & Kording, K. (2008). Estimating the sources of motor errors for adaptation and generalization. *Nat Neurosci, 11*(12), 1454-1461. doi:10.1038/nn.2229

Bogacz, R., & Gurney, K. (2007). The basal ganglia and cortex implement optimal decision making between alternative actions. *Neural computation, 19*(2), 442-477.

Bollimunta, A., Totten, D., & Ditterich, J. (2012). Neural dynamics of choice: single-trial analysis of decision-related activity in parietal cortex. *J Neurosci, 32*(37), 12684-12701. doi:10.1523/JNEUROSCI.5752-11.2012

Bowditch, H. P., & Southard, W. F. (1882). A comparison of sight and touch. *The Journal of physiology, 3*(3-4), 232-245.

Britten, K. H., Newsome, W. T., Shadlen, M. N., Celebrini, S., & Movshon, J. A. (1996). A relationship between behavioral choice and the visual responses of neurons in macaque MT. *Visual neuroscience, 13*(1), 87-100.

Burge, J., Ernst, M. O., & Banks, M. S. (2008). The statistical determinants of adaptation rate in human reaching. *J Vis, 8*(4), 20 21-19. doi:10.1167/8.4.20

Burk, D., Ingram, J. N., Franklin, D. W., Shadlen, M. N., & Wolpert, D. M. (2014). Motor effort alters changes of mind in sensorimotor decision making. *PLoS One, 9*(3), e92681. doi:10.1371/journal.pone.0092681

Byron, M. Y., Cunningham, J. P., Santhanam, G., Ryu, S. I., Shenoy, K. V., & Sahani, M. (2009). *Gaussian-process factor analysis for low-dimensional single-trial analysis of neural population activity.* Paper presented at the Advances in neural information processing systems.

Carmena, J. M., Lebedev, M. A., Crist, R. E., O'Doherty, J. E., Santucci, D. M., Dimitrov, D. F., . . . Nicolelis, M. A. (2003). Learning to control a brain-machine interface for reaching and grasping by primates. *PLoS Biol, 1*(2), E42. doi:10.1371/journal.pbio.0000042

Carpenter, R. H., & Williams, M. L. (1995). Neural computation of log likelihood in control of saccadic eye movements. *Nature, 377*(6544), 59-62. doi:10.1038/377059a0

Chapin, J. K., Moxon, K. A., Markowitz, R. S., & Nicolelis, M. A. (1999). Real-time control of a robot arm using simultaneously recorded neurons in the motor cortex. *Nature neuroscience, 2*(7).

Chapman, C. S., Gallivan, J. P., Wood, D. K., Milne, J. L., Culham, J. C., & Goodale, M. A. (2010). Reaching for the unknown: multiple target encoding and real-time decision-making in a rapid reach task. *Cognition, 116*(2), 168-176. doi:10.1016/j.cognition.2010.04.008

Cheney, P. D., & Fetz, E. E. (1980). Functional classes of primate corticomotoneuronal cells and their relation to active force. *Journal of neurophysiology, 44*(4), 773-791.

Cheng, K., Shettleworth, S. J., Huttenlocher, J., & Rieser, J. J. (2007). Bayesian integration of spatial information. *Psychol Bull, 133*(4), 625-637. doi:10.1037/0033-2909.133.4.625

Christopoulos, V., Bonaiuto, J., & Andersen, R. A. (2015). A biologically plausible computational theory for value integration and action selection in decisions with competing alternatives. *PLoS Comput Biol, 11*(3), e1004104. doi:10.1371/journal.pcbi.1004104

Churchland, A. K., Kiani, R., & Shadlen, M. N. (2008). Decision-making with multiple alternatives. *Nat Neurosci, 11*(6), 693-702. doi:10.1038/nn.2123

Churchland, M. M., Afshar, A., & Shenoy, K. V. (2006). A central source of movement variability. *Neuron, 52*(6), 1085-1096. doi:10.1016/j.neuron.2006.10.034

Churchland, M. M., Byron, M. Y., Ryu, S. I., Santhanam, G., & Shenoy, K. V. (2006). Neural variability in premotor cortex provides a signature of motor preparation. *Journal of Neuroscience, 26*(14), 3697-3712.

Churchland, M. M., Cunningham, J. P., Kaufman, M. T., Foster, J. D., Nuyujukian, P., Ryu, S. I., & Shenoy, K. V. (2012). Neural population dynamics during reaching. *Nature, 487*(7405), 51-56. doi:10.1038/nature11129

Churchland, M. M., Cunningham, J. P., Kaufman, M. T., Ryu, S. I., & Shenoy, K. V. (2010). Cortical preparatory activity: representation of movement or first cog in a dynamical machine? *Neuron, 68*(3), 387-400. doi:10.1016/j.neuron.2010.09.015

Churchland, M. M., Santhanam, G., & Shenoy, K. V. (2006). Preparatory activity in premotor and motor cortex reflects the speed of the upcoming reach. *J Neurophysiol, 96*(6), 3130-3146. doi:10.1152/jn.00307.2006

Churchland, M. M., & Shenoy, K. V. (2007). Temporal complexity and heterogeneity of single-neuron activity in premotor and motor cortex. *J Neurophysiol, 97*(6), 4235-4257. doi:10.1152/jn.00095.2007

Cisek, P. (2006). Integrated neural processes for defining potential actions and deciding between them: a computational model. *J Neurosci, 26*(38), 9761-9770. doi:10.1523/JNEUROSCI.5605-05.2006

Cisek, P. (2007). Cortical mechanisms of action selection: the affordance competition hypothesis. *Philos Trans R Soc Lond B Biol Sci, 362*(1485), 1585-1599. doi:10.1098/rstb.2007.2054

Cisek, P., Crammond, D. J., & Kalaska, J. F. (2003). Neural activity in primary motor and dorsal premotor cortex in reaching tasks with the contralateral versus ipsilateral arm. *Journal of neurophysiology, 89*(2), 922-942.

Cisek, P., & Kalaska, J. F. (2005). Neural correlates of reaching decisions in dorsal premotor cortex: specification of multiple direction choices and final selection of action. *Neuron, 45*(5), 801-814. doi:10.1016/j.neuron.2005.01.027

Coallier, E., Michelet, T., & Kalaska, J. F. (2015). Dorsal premotor cortex: neural correlates of reach target decisions based on a color-location matching rule and conflicting sensory evidence. *J Neurophysiol, 113*(10), 3543-3573. doi:10.1152/jn.00166.2014

Cohen, M. X., & Ranganath, C. (2007). Reinforcement learning signals predict future decisions. *J Neurosci, 27*(2), 371-378. doi:10.1523/JNEUROSCI.4421-06.2007

Cortes, N., Onate, J., & Morrison, S. (2014). Differential effects of fatigue on movement variability. *Gait Posture, 39*(3), 888-893. doi:10.1016/j.gaitpost.2013.11.020

Cos, I., Belanger, N., & Cisek, P. (2011). The influence of predicted arm biomechanics on decision making. *J Neurophysiol, 105*(6), 3022-3033. doi:10.1152/jn.00975.2010

Crammond, D., & Kalaska, J. (1994). Modulation of preparatory neuronal activity in dorsal premotor cortex due to stimulus-response compatibility. *Journal of neurophysiology, 71*(3), 1281-1284.

Crammond, D. J., & Kalaska, J. F. (1996). Differential relation of discharge in primary motor cortex and premotor cortex to movements versus actively maintained postures during a reaching task. *Experimental brain research, 108*(1), 45-61.

Crammond, D. J., & Kalaska, J. F. (2000). Prior information in motor and premotor cortex: activity during the delay period and effect on pre-movement activity. *Journal of neurophysiology, 84*(2), 986-1005.

Dayan, P., & Balleine, B. W. (2002). Reward, motivation, and reinforcement learning. *Neuron, 36*(2), 285-298.

Dekleva, B. M., Ramkumar, P., Wanda, P. A., Kording, K. P., & Miller, L. E. (2016). Uncertainty leads to persistent effects on reach representations in dorsal premotor cortex. *Elife, 5*. doi:10.7554/eLife.14316

Deneve, S. (2008a). Bayesian spiking neurons I: inference. *Neural computation, 20*(1), 91-117.

Deneve, S. (2008b). Bayesian spiking neurons II: learning. *Neural computation, 20*(1), 118-145.

Deneve, S., Latham, P. E., & Pouget, A. (2001). Efficient computation and cue integration with noisy population codes. *Nature neuroscience, 4*(8), 826-831.

Diedrichsen, J., Hashambhoy, Y., Rane, T., & Shadmehr, R. (2005). Neural correlates of reach errors. *J Neurosci, 25*(43), 9919-9931. doi:10.1523/JNEUROSCI.1874-05.2005

Dommett, E., Coizet, V., Blaha, C. D., Martindale, J., Lefebvre, V., Walton, N., . . . Redgrave, P. (2005). How visual stimuli activate dopaminergic neurons at short latency. *Science, 307*(5714), 1476-1479.

Doya, K. (2000). Complementary roles of basal ganglia and cerebellum in learning and motor control. *Current Opinion in Neurobiology, 10*(6), 732-739.

Dum, R. P., & Strick, P. L. (1991). The origin of corticospinal projections from the premotor areas in the frontal lobe. *Journal of Neuroscience, 11*(3), 667-689.

Dum, R. P., & Strick, P. L. (2002). Motor areas in the frontal lobe of the primate. *Physiology & behavior, 77*(4), 677-682.

Elsayed, G. F., Lara, A. H., Kaufman, M. T., Churchland, M. M., & Cunningham, J. P. (2016). Reorganization between preparatory and movement population responses in motor cortex. *Nat Commun, 7*, 13239. doi:10.1038/ncomms13239

Ernst, M. O., & Banks, M. S. (2002). Humans integrate visual and haptic information in a statistically optimal fashion. *Nature, 415*(6870), 429-433.

Evarts, E. V. (1968). Relation of pyramidal tract activity to force exerted during voluntary movement. *Journal of neurophysiology, 31*(1), 14-27.

Even-Chen, N., Stavisky, S. D., Kao, J. C., Ryu, S. I., & Shenoy, K. V. (2015). *Auto-deleting brain machine interface: Error detection using spiking neural activity in the motor cortex.* Paper presented at the Engineering in Medicine and Biology Society (EMBC), 2015 37th Annual International Conference of the IEEE.

Even-Chen, N., Stavisky, S. D., Kao, J. C., Ryu, S. I., & Shenoy, K. V. (2017). Augmenting intracortical brain-machine interface with neurally driven error detectors. *J Neural Eng, 14*(6), 066007. doi:10.1088/1741-2552/aa8dc1

Fernandes, H. L., Stevenson, I. H., Phillips, A. N., Segraves, M. A., & Kording, K. P. (2014). Saliency and saccade encoding in the frontal eye field during natural scene search. *Cereb Cortex, 24*(12), 3232-3245. doi:10.1093/cercor/bht179

Fetsch, C. R., Pouget, A., DeAngelis, G. C., & Angelaki, D. E. (2011). Neural correlates of reliability-based cue weighting during multisensory integration. *Nat Neurosci, 15*(1), 146-154. doi:10.1038/nn.2983

Fetz, E. E. (1992). Are movement parameters recognizably coded in the activity of single neurons? *Behavioral and brain sciences, 15*, 679-679.

Fouragnan, E., Retzler, C., Mullinger, K., & Philiastides, M. G. (2015). Two spatiotemporally distinct value systems shape reward-based learning in the human brain. *Nat Commun, 6*, 8107. doi:10.1038/ncomms9107

Frank, M. J., Woroch, B. S., & Curran, T. (2005). Error-related negativity predicts reinforcement learning and conflict biases. *Neuron, 47*(4), 495-501. doi:10.1016/j.neuron.2005.06.020

Franz, V. H., Gegenfurtner, K. R., Bülthoff, H. H., & Fahle, M. (2000). Grasping visual illusions: No evidence for a dissociation between perception and action. *Psychological Science, 11*(1), 20-25.

Friedman, J., Hastie, T., & Tibshirani, R. (2010). Regularization paths for generalized linear models via coordinate descent. *Journal of statistical software, 33*(1), 1.

Fritsch, G., & Hitzig, E. (1960). On the electrical excitability of the cerebrum. *Some Papers on the Cerebral Cortex. Charles C Thomas, Springfield, IL*, 73-96.

Fu, Q., Flament, D., Coltz, J., & Ebner, T. (1995). Temporal encoding of movement kinematics in the discharge of primate primary motor and premotor neurons. *Journal of neurophysiology, 73*(2), 836-854.

Fujii, N., Mushiake, H., & Tanji, J. (2000). Rostrocaudal distinction of the dorsal premotor area based on oculomotor involvement. *Journal of neurophysiology, 83*(3), 1764-1769.

Gail, A., Klaes, C., & Westendorff, S. (2009). Implementation of spatial transformation rules for goal-directed reaching via gain modulation in monkey parietal and premotor cortex. *J Neurosci, 29*(30), 9490-9499. doi:10.1523/JNEUROSCI.1095-09.2009

Galea, J. M., Mallia, E., Rothwell, J., & Diedrichsen, J. (2015). The dissociable effects of punishment and reward on motor learning. *Nat Neurosci, 18*(4), 597-602. doi:10.1038/nn.3956

Gallivan, J. P., Logan, L., Wolpert, D. M., & Flanagan, J. R. (2016). Parallel specification of competing sensorimotor control policies for alternative action options. *Nat Neurosci, 19*(2), 320-326. doi:10.1038/nn.4214

Gallivan, J. P., Stewart, B. M., Baugh, L. A., Wolpert, D. M., & Flanagan, J. R. (2017). Rapid Automatic Motor Encoding of Competing Reach Options. *Cell Rep, 18*(7), 1619-1626. doi:10.1016/j.celrep.2017.01.049

Georgopoulos, A. P., Ashe, J., Smyrnis, N., & Taira, M. (1992). The motor cortex and the coding of force. *Science, 256*(5064), 1692-1695.

Georgopoulos, A. P., Kalaska, J. F., Caminiti, R., & Massey, J. T. (1982). On the relations between the direction of two-dimensional arm movements and cell discharge in primate motor cortex. *Journal of Neuroscience, 2*(11), 1527-1537.

Georgopoulos, A. P., Kalaska, J. F., & Massey, J. T. (1981). Spatial trajectories and reaction times of aimed movements: effects of practice, uncertainty, and change in target location. *Journal of neurophysiology, 46*(4), 725-743.

Ghez, C., Favilla, M., Ghilardi, M., Gordon, J., Bermejo, R., & Pullman, S. (1997). Discrete and continuous planning of hand movements and isometric force trajectories. *Experimental brain research, 115*(2), 217-233.

Godschalk, M., Lemon, R., Kuypers, H., & Van der Steen, J. (1985). The involvement of monkey premotor cortex neurones in preparation of visually cued arm movements. *Behavioural brain research, 18*(2), 143-157.

Gold, J. I., & Shadlen, M. N. (2000). Representation of a perceptual decision in developing oculomotor commands. *Nature, 404*(6776), 390.

Gold, J. I., & Shadlen, M. N. (2001). Neural computations that underlie decisions about sensory stimuli. *Trends in cognitive sciences, 5*(1), 10-16.

Goldreich, D. (2007). A Bayesian perceptual model replicates the cutaneous rabbit and other tactile spatiotemporal illusions. *PLoS One, 2*(3), e333. doi:10.1371/journal.pone.0000333

Greenwald, H. S., & Knill, D. C. (2009). A comparison of visuomotor cue integration strategies for object placement and prehension. *Vis Neurosci, 26*(1), 63-72. doi:10.1017/S0952523808080668

Grossberg, S., & Pilly, P. K. (2008). Temporal dynamics of decision-making during motion perception in the visual cortex. *Vision Res, 48*(12), 1345-1373. doi:10.1016/j.visres.2008.02.019

Gu, Y., Angelaki, D. E., & Deangelis, G. C. (2008). Neural correlates of multisensory cue integration in macaque MSTd. *Nat Neurosci, 11*(10), 1201-1210. doi:10.1038/nn.2191

Haber, S. N. (2011). 11 Neuroanatomy of Reward: A View from the Ventral Striatum. *Neurobiology of sensation and reward*, 235.

Haith, A. M., Huberdeau, D. M., & Krakauer, J. W. (2015). Hedging your bets: intermediate movements as optimal behavior in the context of an incomplete decision. *PLoS Comput Biol, 11*(3), e1004171. doi:10.1371/journal.pcbi.1004171

Haith, A. M., Pakpoor, J., & Krakauer, J. W. (2016). Independence of Movement Preparation and Movement Initiation. *J Neurosci, 36*(10), 3007-3015. doi:10.1523/JNEUROSCI.3245-15.2016

Hamilton, A. F., Jones, K. E., & Wolpert, D. M. (2004). The scaling of motor noise with muscle strength and motor unit number in humans. *Exp Brain Res, 157*(4), 417-430. doi:10.1007/s00221-004-1856-7

Hastie, T., Tibshirani, R., & Friedman, J. (2009). Overview of Supervised Learning. In *The Elements of Statistical Learning* (pp. 1-33).

He, P., & Kowler, E. (1989). The role of location probability in the programming of saccades: Implications for "center-of-gravity" tendencies. *Vision Research, 29*(9), 1165-1181.

Henis, E. A., & Flash, T. (1995). Mechanisms underlying the generation of averaged modified trajectories. *Biological cybernetics, 72*(5), 407-419.

Hernandez, A., Nacher, V., Luna, R., Zainos, A., Lemus, L., Alvarez, M., . . . Romo, R. (2010). Decoding a perceptual decision process across cortex. *Neuron, 66*(2), 300-314. doi:10.1016/j.neuron.2010.03.031

Herzfeld, D. J., & Shadmehr, R. (2014). Motor variability is not noise, but grist for the learning mill. *Nat Neurosci, 17*(2), 149-150. doi:10.1038/nn.3633

Hillis, J. M., Watt, S. J., Landy, M. S., & Banks, M. S. (2004). Slant from texture and disparity cues: optimal cue combination. *J Vis, 4*(12), 967-992. doi:10.1167/4.12.1

Hinton, G. E., & Sejnowski, T. J. (1983). *Optimal perceptual inference.* Paper presented at the Proceedings of the IEEE conference on Computer Vision and Pattern Recognition.

Hocherman, S., & Wise, S. (1991). Effects of hand movement path on motor cortical activity in awake, behaving rhesus monkeys. *Experimental brain research, 83*(2), 285-302.

Hocherman, S., & Wise, S. P. (1990). Trajectory-selective neuronal activity in the motor cortex of rhesus monkeys (Macaca mulatta). *Behavioral neuroscience, 104*(3), 495.

Holdefer, R. N., & Miller, L. E. (2002). Primary motor cortical neurons encode functional muscle synergies. *Exp Brain Res, 146*(2), 233-243. doi:10.1007/s00221-002-1166-x

Hollerman, J. R., & Schultz, W. (1998). Dopamine neurons report an error in the temporal prediction of reward during learning. *Nature neuroscience, 1*(4), 304-309.

Hoover, J. E., & Strick, P. L. (1999). The organization of cerebellar and basal ganglia outputs to primary motor cortex as revealed by retrograde transneuronal transport of herpes simplex virus type 1. *Journal of Neuroscience, 19*(4), 1446-1463.

Hoshi, E., & Tanji, J. (2002). Contrasting neuronal activity in the dorsal and ventral premotor areas during preparation to reach. *Journal of neurophysiology, 87*(2), 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*(6), 3596-3616. doi:10.1152/jn.01126.2005

Hoyer, P. O., & Hyvärinen, A. (2003). *Interpreting neural response variability as Monte Carlo sampling of the posterior.* Paper presented at the Advances in neural information processing systems.

Izawa, J., & Shadmehr, R. (2008). On-line processing of uncertain information in visuomotor control. *J Neurosci, 28*(44), 11360-11368. doi:10.1523/JNEUROSCI.3063-08.2008

Izawa, J., & Shadmehr, R. (2011). Learning from sensory and reward prediction errors during motor adaptation. *PLoS Comput Biol, 7*(3), e1002012. doi:10.1371/journal.pcbi.1002012

Jazayeri, M., & Movshon, J. A. (2006). Optimal representation of sensory information by neural populations. *Nat Neurosci, 9*(5), 690-696. doi:10.1038/nn1691

Jeannerod, M., Arbib, M. A., Rizzolatti, G., & Sakata, H. (1995). Grasping objects: the cortical mechanisms of visuomotor transformation. *Trends in neurosciences, 18*(7), 314-320.

Johnson, P. B., Ferraina, S., Bianchi, L., & Caminiti, R. (1996). Cortical networks for visual reaching: physiological and anatomical organization of frontal and parietal lobe arm regions. *Cerebral cortex, 6*(2), 102-119.

Jones, K. E., Hamilton, A. F. d. C., & Wolpert, D. M. (2002). Sources of signal-dependent noise during isometric force production. *Journal of neurophysiology, 88*(3), 1533-1544.

Kakei, S., Hoffman, D. S., & Strick, P. L. (1999). Muscle and movement representations in the primary motor cortex. *Science, 285*(5436), 2136-2139.

Kalaska, J. F., & Crammond, D. J. (1992). Cerebral cortical mechanisms of reaching movements. *Science, 255*(5051), 1517-1523.

Kao, J. C., Nuyujukian, P., Ryu, S. I., Churchland, M. M., Cunningham, J. P., & Shenoy, K. V. (2015). Single-trial dynamics of motor cortex and their applications to brain-machine interfaces. *Nat Commun, 6*, 7759. doi:10.1038/ncomms8759

Kaufman, M. T., Churchland, M. M., Ryu, S. I., & Shenoy, K. V. (2014). Cortical activity in the null space: permitting preparation without movement. *Nat Neurosci, 17*(3), 440-448. doi:10.1038/nn.3643

Kaufman, M. T., Churchland, M. M., Ryu, S. I., & Shenoy, K. V. (2015). Vacillation, indecision and hesitation in moment-by-moment decoding of monkey motor cortex. *Elife, 4*, e04677. doi:10.7554/eLife.04677

Kelly, R. M., & Strick, P. L. (2003). Cerebellar loops with motor cortex and prefrontal cortex of a nonhuman primate. *Journal of Neuroscience, 23*(23), 8432-8444.

Kennerley, S. W., Walton, M. E., Behrens, T. E., Buckley, M. J., & Rushworth, M. F. (2006). Optimal decision making and the anterior cingulate cortex. *Nat Neurosci, 9*(7), 940-947. doi:10.1038/nn1724

Kiani, R., Cueva, C. J., Reppas, J. B., & Newsome, W. T. (2014). Dynamics of neural population responses in prefrontal cortex indicate changes of mind on single trials. *Curr Biol, 24*(13), 1542-1547. doi:10.1016/j.cub.2014.05.049

Klaes, C., Schneegans, S., Schoner, G., & Gail, A. (2012). Sensorimotor learning biases choice behavior: a learning neural field model for decision making. *PLoS Comput Biol, 8*(11), e1002774. doi:10.1371/journal.pcbi.1002774

Klaes, C., Westendorff, S., Chakrabarti, S., & Gail, A. (2011). Choosing goals, not rules: deciding among rule-based action plans. *Neuron, 70*(3), 536-548. doi:10.1016/j.neuron.2011.02.053

Knill, D. C., & Saunders, J. A. (2003). Do humans optimally integrate stereo and texture information for judgments of surface slant? *Vision Research, 43*(24), 2539-2558. doi:10.1016/s0042-6989(03)00458-9

Kording, K. P., & Wolpert, D. M. (2004). Bayesian integration in sensorimotor learning. *Nature, 427*(6971), 244.

Kording, K. P., & Wolpert, D. M. (2006). Bayesian decision theory in sensorimotor control. *Trends Cogn Sci, 10*(7), 319-326. doi:10.1016/j.tics.2006.05.003

Krakauer, J. W., & Mazzoni, P. (2011). Human sensorimotor learning: adaptation, skill, and beyond. *Curr Opin Neurobiol, 21*(4), 636-644. doi:10.1016/j.conb.2011.06.012

Kurata, K. (1991). Corticocortical inputs to the dorsal and ventral aspects of the premotor cortex of macaque monkeys. *Neuroscience research, 12*(1), 263-280.

Lee, J. H., & van Donkelaar, P. (2006). The human dorsal premotor cortex generates on-line error corrections during sensorimotor adaptation. *J Neurosci, 26*(12), 3330-3334. doi:10.1523/JNEUROSCI.3898-05.2006

Leon, M. I., & Shadlen, M. N. (2003). Representation of time by neurons in the posterior parietal cortex of the macaque. *Neuron, 38*(2), 317-327.

Lochmann, T., & Deneve, S. (2011). Neural processing as causal inference. *Curr Opin Neurobiol, 21*(5), 774-781. doi:10.1016/j.conb.2011.05.018

Lu, M. T., Preston, J. B., & Strick, P. L. (1994). Interconnections between the prefrontal cortex and the premotor areas in the frontal lobe. *Journal of Comparative Neurology, 341*(3), 375-392.

Ma, W. J., Beck, J. M., Latham, P. E., & Pouget, A. (2006). Bayesian inference with probabilistic population codes. *Nat Neurosci, 9*(11), 1432-1438. doi:10.1038/nn1790

Mamassian, P., & Landy, M. S. (2001). Interaction of visual prior constraints. *Vision Research, 41*(20), 2653-2668.

Mante, V., Sussillo, D., Shenoy, K. V., & Newsome, W. T. (2013). Context-dependent computation by recurrent dynamics in prefrontal cortex. *Nature, 503*(7474), 78-84. doi:10.1038/nature12742

Marsh, B. T., Tarigoppula, V. S., Chen, C., & Francis, J. T. (2015). Toward an autonomous brain machine interface: integrating sensorimotor reward modulation and reinforcement learning. *J Neurosci, 35*(19), 7374-7387. doi:10.1523/JNEUROSCI.1802-14.2015

Martin, T., Keating, J., Goodkin, H., Bastian, A., & Thach, W. (1996). Throwing while looking through prisms: I. Focal olivocerebellar lesions impair adaptation. *Brain, 119*(4), 1183-1198.

Martino, A. M., & Strick, P. L. (1987). Corticospinal projections originate from the arcuate premotor area. *Brain research, 404*(1), 307-312.

Mason, C. R., Johnson, M. T., Fu, Q.-G., Gomez, J. E., & Ebner, T. J. (1998). Temporal profile of the directional tuning of the discharge of dorsal premotor cortical cells. *Neuroreport, 9*(6), 989-995.

Mazurek, M. E. (2003). A Role for Neural Integrators in Perceptual Decision Making. *Cerebral cortex, 13*(11), 1257-1269. doi:10.1093/cercor/bhg097

Mazzoni, P., & Krakauer, J. W. (2006). An implicit plan overrides an explicit strategy during visuomotor adaptation. *J Neurosci, 26*(14), 3642-3645. doi:10.1523/JNEUROSCI.5317-05.2006

McFadden, D. (1973). Conditional logit analysis of qualitative choice behavior.

McGuire, L. M., & Sabes, P. N. (2009). Sensory transformations and the use of multiple reference frames for reach planning. *Nat Neurosci, 12*(8), 1056-1061. doi:10.1038/nn.2357

McKinstry, C., Dale, R., & Spivey, M. J. (2008). Action dynamics reveal parallel competition in decision making. *Psychological Science, 19*(1), 22-24.

Meegan, D. V., & Tipper, S. P. (1998). Reaching into cluttered visual environments: spatial and temporal influences of distracting objects. *The Quarterly Journal of Experimental Psychology A: Human Experimental Psychology*.

Messier, J., & Kalaska, J. F. (2000). Covariation of primate dorsal premotor cell activity with direction and amplitude during a memorized-delay reaching task. *Journal of neurophysiology, 84*(1), 152-165.

Michaels, J. A., Dann, B., Intveld, R. W., & Scherberger, H. (2015). Predicting Reaction Time from the Neural State Space of the Premotor and Parietal Grasping Network. *J Neurosci, 35*(32), 11415-11432. doi:10.1523/JNEUROSCI.1714-15.2015

Michaels, J. A., Dann, B., & Scherberger, H. (2016). Neural Population Dynamics during Reaching Are Better Explained by a Dynamical System than Representational Tuning. *PLoS Comput Biol, 12*(11), e1005175. doi:10.1371/journal.pcbi.1005175

Mirabella, G., Pani, P., & Ferraina, S. (2011). Neural correlates of cognitive control of reaching movements in the dorsal premotor cortex of rhesus monkeys. *J Neurophysiol, 106*(3), 1454-1466. doi:10.1152/jn.00995.2010

Morel, P., Ulbrich, P., & Gail, A. (2017). What makes a reach movement effortful? Physical effort discounting supports common minimization principles in decision making and motor control. *PLoS Biol, 15*(6), e2001323. doi:10.1371/journal.pbio.2001323

Morrow, M. M., Jordan, L. R., & Miller, L. E. (2007). Direct comparison of the task-dependent discharge of M1 in hand space and muscle space. *J Neurophysiol, 97*(2), 1786-1798. doi:10.1152/jn.00150.2006

Mountcastle, V. B., Steinmetz, M., & Romo, R. (1990). Frequency discrimination in the sense of flutter: psychophysical measurements correlated with postcentral events in behaving monkeys. *Journal of Neuroscience, 10*(9), 3032-3044.

Munoz, D. P., & Wurtz, R. H. (1995). Saccade-related activity in monkey superior colliculus. I. Characteristics of burst and buildup cells. *Journal of neurophysiology, 73*(6), 2313-2333.

Murata, A., Fadiga, L., Fogassi, L., Gallese, V., Raos, V., & Rizzolatti, G. (1997). Object representation in the ventral premotor cortex (area F5) of the monkey. *Journal of neurophysiology, 78*(4), 2226-2230.

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. *Journal of neurophysiology, 66*(3), 705-718.

Nashed, J. Y., Diamond, J. S., Gallivan, J. P., Wolpert, D. M., & Flanagan, J. R. (2017). Grip force when reaching with target uncertainty provides evidence for motor optimization over averaging. *Sci Rep, 7*(1), 11703. doi:10.1038/s41598-017-10996-6

Newsome, W. T., Britten, K. H., & Movshon, J. A. (1989). Neuronal correlates of a perceptual decision. *Nature, 341*(6237), 52-54.

Nicolson, R. I., Fawcett, A. J., Berry, E. L., Jenkins, I. H., Dean, P., & Brooks, D. J. (1999). Association of abnormal cerebellar activation with motor learning difficulties in dyslexic adults. *The Lancet, 353*(9165), 1662-1667. doi:10.1016/s0140-6736(98)09165-x

Nikooyan, A. A., & Ahmed, A. A. (2015). Reward feedback accelerates motor learning. *J Neurophysiol, 113*(2), 633-646. doi:10.1152/jn.00032.2014

Nitsche, M. A., Schauenburg, A., Lang, N., Liebetanz, D., Exner, C., Paulus, W., & Tergau, F. (2003). Facilitation of implicit motor learning by weak transcranial direct current stimulation of the primary motor cortex in the human. *Journal of cognitive neuroscience, 15*(4), 619-626.

O'Doherty, J. P., Dayan, P., Friston, K., Critchley, H., & Dolan, R. J. (2003). Temporal difference models and reward-related learning in the human brain. *Neuron, 38*(2), 329-337.

Ohbayashi, M., Ohki, K., & Miyashita, Y. (2003). Conversion of working memory to motor sequence in the monkey premotor cortex. *Science, 301*(5630), 233-236.

Osu, R., Kamimura, N., Iwasaki, H., Nakano, E., Harris, C. M., Wada, Y., & Kawato, M. (2004). Optimal impedance control for task achievement in the presence of signal-dependent noise. *Journal of neurophysiology, 92*(2), 1199-1215.

Park, I. M., Meister, M. L., Huk, A. C., & Pillow, J. W. (2014). Encoding and decoding in parietal cortex during sensorimotor decision-making. *Nat Neurosci, 17*(10), 1395-1403. doi:10.1038/nn.3800

Pastor-Bernier, A., & Cisek, P. (2011). Neural correlates of biased competition in premotor cortex. *J Neurosci, 31*(19), 7083-7088. doi:10.1523/JNEUROSCI.5681-10.2011

Pastor-Bernier, A., Tremblay, E., & Cisek, P. (2012). Dorsal premotor cortex is involved in switching motor plans. *Front Neuroeng, 5*, 5. doi:10.3389/fneng.2012.00005

Paus, T. (2001). Primate anterior cingulate cortex: where motor control, drive and cognition interface. *Nature reviews neuroscience, 2*(6), 417-424.

Penfield, W., & Boldrey, E. (1937). Somatic motor and sensory representation in the cerebral cortex of man as studied by electrical stimulation. *Brain: A journal of neurology*.

Pesaran, B., Nelson, M. J., & Andersen, R. A. (2006). Dorsal premotor neurons encode the relative position of the hand, eye, and goal during reach planning. *Neuron, 51*(1), 125-134. doi:10.1016/j.neuron.2006.05.025

Picard, N., & Strick, P. L. (2001). Imaging the premotor areas. *Current Opinion in Neurobiology, 11*(6), 663-672.

Platt, M. L., & Glimcher, P. W. (1997). Responses of intraparietal neurons to saccadic targets and visual distractors. *Journal of neurophysiology, 78*(3), 1574-1589.

Pohlmeyer, E. A., Solla, S. A., Perreault, E. J., & Miller, L. E. (2007). Prediction of upper limb muscle activity from motor cortical discharge during reaching. *J Neural Eng, 4*(4), 369-379. doi:10.1088/1741-2560/4/4/003

Purves, D. E., Augustine, G. J., Fitzpatrick, D. E., & Katz, L. C. (1997). *Neuroscience*: Sunderland, MA, US: Sinauer Associates.

Qian, J., Hastie, T., Friedman, J., Tibshirani, R., & Simon, N. (2013). Glmnet for matlab. In: Stanford, CA: Stanford Univ. http://www. stanford. edu/∼ hastie/glmnet_matlab.

Rajkowska, G., & Goldman-Rakic, P. S. (1995). Cytoarchitectonic definition of prefrontal areas in the normal human cortex: I. Remapping of areas 9 and 46 using quantitative criteria. *Cerebral cortex, 5*(4), 307-322.

Ramkumar, P., Lawlor, P. N., Glaser, J. I., Wood, D. K., Phillips, A. N., Segraves, M. A., & Kording, K. P. (2016). Feature-based attention and spatial selection in frontal eye fields during natural scene search. *J Neurophysiol, 116*(3), 1328-1343. doi:10.1152/jn.01044.2015

Ratcliff, R. (1978). A theory of memory retrieval. *Psychological Review, 85*(2), 59.

Rathelot, J.-A., & Strick, P. L. (2009). Subdivisions of primary motor cortex based on cortico-motoneuronal cells. *Proceedings of the National Academy of Sciences, 106*(3), 918-923.

Reddi, B., Asrress, K., & Carpenter, R. (2003). Accuracy, information, and response time in a saccadic decision task. *Journal of neurophysiology, 90*(5), 3538-3546.

Resulaj, A., Kiani, R., Wolpert, D. M., & Shadlen, M. N. (2009). Changes of mind in decision-making. *Nature, 461*(7261), 263-266. doi:10.1038/nature08275

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. *Journal of neurophysiology, 61*(3), 534-549.

Riehle, A., & Requin, J. (1993). The predictive value for performance speed of preparatory changes in neuronal activity of the monkey motor and premotor cortex. *Behavioural brain research, 53*(1), 35-49.

Ritter, W., Simson, R., & Vaughan, H. G. (1972). Association cortex potentials and reaction time in auditory discrimination. *Electroencephalography and clinical neurophysiology, 33*(6), 547-555.

Roesch, M. R., & Olson, C. R. (2003). Impact of expected reward on neuronal activity in prefrontal cortex, frontal and supplementary eye fields and premotor cortex. *Journal of neurophysiology, 90*(3), 1766-1789.

Roesch, M. R., & Olson, C. R. (2004). Neuronal activity related to reward value and motivation in primate frontal cortex. *Science, 304*(5668), 307-310.

Roesch, M. R., & Olson, C. R. (2005). Neuronal activity in primate orbitofrontal cortex reflects the value of time. *J Neurophysiol, 94*(4), 2457-2471. doi:10.1152/jn.00373.2005

Romo, R., Brody, C. D., Hernández, A., & Lemus, L. (1999). Neuronal correlates of parametric working memory in the prefrontal cortex. *Nature, 399*(6735), 470.

Romo, R., Hernández, A., & Zainos, A. (2004). Neuronal Correlates of a Perceptual Decision in Ventral Premotor Cortex. *Neuron, 41*(1), 165-173. doi:10.1016/s0896-6273(03)00817-1

Rosenbaum, D. A. (1980). Human movement initiation: specification of arm, direction, and extent. *Journal of Experimental Psychology: General, 109*(4), 444.

Rowland, B. A., Stanford, T. R., & Stein, B. E. (2007). A model of the neural mechanisms underlying multisensory integration in the superior colliculus. *Perception, 36*(10), 1431-1443. doi:10.1068/p5842

Salinas, E., & Abbott, L. (1994). Vector reconstruction from firing rates. *Journal of computational neuroscience, 1*(1), 89-107.

Sanes, J. N., Dimitrov, B., & Hallett, M. (1990). Motor learning in patients with cerebellar dysfunction. *Brain, 113*(1), 103-120.

Sanes, J. N., & Donoghue, J. P. (2000). Plasticity and primary motor cortex. *Annual review of neuroscience, 23*(1), 393-415.

Sanger, T. D. (2003). Neural population codes. *Current Opinion in Neurobiology, 13*(2), 238-249. doi:10.1016/s0959-4388(03)00034-5

Schaffelhofer, S., & Scherberger, H. (2016). Object vision to hand action in macaque parietal, premotor, and motor cortices. *Elife, 5*. doi:10.7554/eLife.15278

Schonberg, T., Daw, N. D., Joel, D., & O'Doherty, J. P. (2007). Reinforcement learning signals in the human striatum distinguish learners from nonlearners during reward-based decision making. *J Neurosci, 27*(47), 12860-12867. doi:10.1523/JNEUROSCI.2496-07.2007

Schultz, W. (1998). Predictive reward signal of dopamine neurons. *Journal of neurophysiology, 80*(1), 1-27.

Schultz, W. (2000). Multiple reward signals in the brain. *Nature reviews neuroscience, 1*(3), 199-207.

Schultz, W. (2002). Getting formal with dopamine and reward. *Neuron, 36*(2), 241-263.

Schultz, W. (2006). Behavioral theories and the neurophysiology of reward. *Annu Rev Psychol, 57*, 87-115. doi:10.1146/annurev.psych.56.091103.070229

Scott, S. H., & Kalaska, J. F. (1995). Changes in motor cortex activity during reaching movements with similar hand paths but different arm postures. *Journal of neurophysiology, 73*(6), 2563-2567.

Scott, S. H., Sergio, L. E., & Kalaska, J. F. (1997). Reaching movements with similar hand paths but different arm orientations. II. Activity of individual cells in dorsal premotor cortex and parietal area 5. *Journal of neurophysiology, 78*(5), 2413-2426.

Sergio, L. E., & Kalaska, J. F. (1998). Changes in the temporal pattern of primary motor cortex activity in a directional isometric force versus limb movement task. *Journal of neurophysiology, 80*(3), 1577-1583.

Serruya, M. D., Hatsopoulos, N. G., Paninski, L., Fellows, M. R., & Donoghue, J. P. (2002). Brain-machine interface: Instant neural control of a movement signal. *Nature, 416*(6877), 141-142.

Shadlen, M. N., & Newsome, W. T. (1998). The variable discharge of cortical neurons: implications for connectivity, computation, and information coding. *Journal of Neuroscience, 18*(10), 3870-3896.

Shadlen, M. N., & Newsome, W. T. (2001). Neural basis of a perceptual decision in the parietal cortex (area LIP) of the rhesus monkey. *Journal of neurophysiology, 86*(4), 1916-1936.

Shadmehr, R., Huang, H. J., & Ahmed, A. A. (2016). A Representation of Effort in Decision-Making and Motor Control. *Curr Biol, 26*(14), 1929-1934. doi:10.1016/j.cub.2016.05.065

Shadmehr, R., & Krakauer, J. W. (2008). A computational neuroanatomy for motor control. *Exp Brain Res, 185*(3), 359-381. doi:10.1007/s00221-008-1280-5

Shadmehr, R., & Mussa-Ivaldi, F. A. (1994). Adaptive representation of dynamics during learning of a motor task. *Journal of Neuroscience, 14*(5), 3208-3224.

Shadmehr, R., Smith, M. A., & Krakauer, J. W. (2010). Error correction, sensory prediction, and adaptation in motor control. *Annu Rev Neurosci, 33*, 89-108. doi:10.1146/annurev-neuro-060909-153135

Shen, K., & Pare, M. (2007). Neuronal activity in superior colliculus signals both stimulus identity and saccade goals during visual conjunction search. *J Vis, 7*(5), 15 11-13. doi:10.1167/7.5.15

Shen, L., & Alexander, G. E. (1997). Preferential representation of instructed target location versus limb trajectory in dorsal premotor area. *Journal of neurophysiology, 77*(3), 1195-1212.

Shmuelof, L., & Krakauer, J. W. (2011). Are we ready for a natural history of motor learning? *Neuron, 72*(3), 469-476. doi:10.1016/j.neuron.2011.10.017

Smith, P. L., & Ratcliff, R. (2004). Psychology and neurobiology of simple decisions. *Trends Neurosci, 27*(3), 161-168. doi:10.1016/j.tins.2004.01.006

Sober, S. J., & Sabes, P. N. (2005). Flexible strategies for sensory integration during motor planning. *Nat Neurosci, 8*(4), 490-497. doi:10.1038/nn1427

Stavisky, S. D., Kao, J. C., Ryu, S. I., & Shenoy, K. V. (2017). Motor Cortical Visuomotor Feedback Activity Is Initially Isolated from Downstream Targets in Output-Null Neural State Space Dimensions. *Neuron, 95*(1), 195-208 e199. doi:10.1016/j.neuron.2017.05.023

Stetson, C., & Andersen, R. A. (2014). The parietal reach region selectively anti-synchronizes with dorsal premotor cortex during planning. *J Neurosci, 34*(36), 11948-11958. doi:10.1523/JNEUROSCI.0097-14.2014

Stewart, B. M., Baugh, L. A., Gallivan, J. P., & Flanagan, J. R. (2013). Simultaneous encoding of the direction and orientation of potential targets during reach planning: evidence of multiple competing reach plans. *J Neurophysiol, 110*(4), 807-816. doi:10.1152/jn.00131.2013

Stone, M. (1960). Models for choice-reaction time. *Psychometrika, 25*(3), 251-260.

Tanji, J., Shima, K., & Mushiake, H. (1996). Multiple cortical motor areas and temporal sequencing of movements. *Cognitive Brain Research, 5*(1), 117-122.

Tassinari, H., Hudson, T. E., & Landy, M. S. (2006). Combining priors and noisy visual cues in a rapid pointing task. *J Neurosci, 26*(40), 10154-10163. doi:10.1523/JNEUROSCI.2779-06.2006

Taylor, D. M., Tillery, S. I. H., & Schwartz, A. B. (2002). Direct cortical control of 3D neuroprosthetic devices. *Science, 296*(5574), 1829-1832.

Thura, D., Beauregard-Racine, J., Fradet, C. W., & Cisek, P. (2012). Decision making by urgency gating: theory and experimental support. *J Neurophysiol, 108*(11), 2912-2930. doi:10.1152/jn.01071.2011

Thura, D., & Cisek, P. (2014). Deliberation and commitment in the premotor and primary motor cortex during dynamic decision making. *Neuron, 81*(6), 1401-1416. doi:10.1016/j.neuron.2014.01.031

Tolhurst, D. J., Movshon, J. A., & Dean, A. F. (1983). The statistical reliability of signals in single neurons in cat and monkey visual cortex. *Vision Research, 23*(8), 775-785.

Trommershauser, J., Gepshtein, S., Maloney, L. T., Landy, M. S., & Banks, M. S. (2005). Optimal compensation for changes in task-relevant movement variability. *J Neurosci, 25*(31), 7169-7178. doi:10.1523/JNEUROSCI.1906-05.2005

Trommershauser, J., Maloney, L. T., & Landy, M. S. (2008). Decision making, movement planning and statistical decision theory. *Trends Cogn Sci, 12*(8), 291-297. doi:10.1016/j.tics.2008.04.010

Trommershäuser, J., Maloney, L. T., & Landy, M. S. (2003). Statistical decision theory and trade-offs in the control of motor response. *Spatial vision, 16*(3), 255-275.

Tseng, Y. W., Diedrichsen, J., Krakauer, J. W., Shadmehr, R., & Bastian, A. J. (2007). Sensory prediction errors drive cerebellum-dependent adaptation of reaching. *J Neurophysiol, 98*(1), 54-62. doi:10.1152/jn.00266.2007

Usher, M., & McClelland, J. L. (2001). The time course of perceptual choice: the leaky, competing accumulator model. *Psychological Review, 108*(3), 550.

van Beers, R. J., Baraduc, P., & Wolpert, D. M. (2002). Role of uncertainty in sensorimotor control. *Philos Trans R Soc Lond B Biol Sci, 357*(1424), 1137-1145. doi:10.1098/rstb.2002.1101

Van Beers, R. J., Haggard, P., & Wolpert, D. M. (2004). The role of execution noise in movement variability. *Journal of neurophysiology, 91*(2), 1050-1063.

Wallis, J. D., & Kennerley, S. W. (2010). Heterogeneous reward signals in prefrontal cortex. *Curr Opin Neurobiol, 20*(2), 191-198. doi:10.1016/j.conb.2010.02.009

Wang, X.-J. (2002). Probabilistic decision making by slow reverberation in cortical circuits. *Neuron, 36*(5), 955-968.

Wann, J. P., & Ibrahim, S. F. (1992). Does limb proprioception drift? *Experimental brain research, 91*(1), 162-166.

Wei, K., & Kording, K. (2010). Uncertainty of feedback and state estimation determines the speed of motor adaptation. *Front Comput Neurosci, 4*, 11. doi:10.3389/fncom.2010.00011

Weinrich, M., & Wise, S. P. (1982). The premotor cortex of the monkey. *Journal of Neuroscience, 2*(9), 1329-1345.

Weiss, Y., Simoncelli, E. P., & Adelson, E. H. (2002). Motion illusions as optimal percepts. *Nat Neurosci, 5*(6), 598-604. doi:10.1038/nn858

Welsh, T. N., & Elliott, D. (2005). The effects of response priming on the planning and execution of goal-directed movements in the presence of a distracting stimulus. *Acta Psychol (Amst), 119*(2), 123-142. doi:10.1016/j.actpsy.2005.01.001

Wessberg, J., Stambaugh, C. R., Kralik, J. D., & Beck, P. D. (2000). Real-time prediction of hand trajectory by ensembles of cortical neurons in primates. *Nature, 408*(6810), 361.

Wolpert, D. M., & Ghahramani, Z. (2004). Computational motor control. *Science, 269*, 1880-1882.

Wolpert, D. M., Ghahramani, Z., & Jordan, M. I. (1995). An internal model for sensorimotor integration. *Science*, 1880-1882.

Wolpert, D. M., & Landy, M. S. (2012). Motor control is decision-making. *Curr Opin Neurobiol, 22*(6), 996-1003. doi:10.1016/j.conb.2012.05.003

Wong, A. L., Goldsmith, J., Forrence, A. D., Haith, A. M., & Krakauer, J. W. (2017). Reaction times can reflect habits rather than computations. *Elife, 6*. doi:10.7554/eLife.28075

Wong, A. L., & Haith, A. M. (2017). Motor planning flexibly optimizes performance under uncertainty about task goals. *Nat Commun, 8*, 14624. doi:10.1038/ncomms14624

Wong, K. F., & Wang, X. J. (2006). A recurrent network mechanism of time integration in perceptual decisions. *J Neurosci, 26*(4), 1314-1328. doi:10.1523/JNEUROSCI.3733-05.2006

Wood, D. K., Gallivan, J. P., Chapman, C. S., Milne, J. L., Culham, J. C., & Goodale, M. A. (2011). Visual salience dominates early visuomotor competition in reaching behavior. *J Vis, 11*(10). doi:10.1167/11.10.16

Zemel, R. S., Dayan, P., & Pouget, A. (1998). Probabilistic interpretation of population codes. *Neural computation, 10*(2), 403-430.