THE UNIVERSITY OF CHICAGO

THE INVOLVEMENT OF PREMOTOR CORTEX IN EXECUTING REACH TO GRASP MOVEMENTS

A DISSERTATION SUBMITTED TO

THE FACULTY OF THE DIVISION OF THE BIOLOGICAL SCIENCES

AND THE PRITZKER SCHOOL OF MEDICINE

IN CANDIDACY FOR THE DEGREE OF

DOCTOR OF PHILOSOPHY

COMMITTEE ON COMPUTATIONAL NEUROSCIENCE

BY

MATTHEW DAVIS BEST

CHICAGO, ILLINOIS

MARCH 2016

TABLE OF CONTENTS

LIST OF FIGURES v

LIST OF TABLES viii

ACKNOWLEDGEMENTS ix

ABSTRACT x

INTRODUCTION 1

The kinematics of reach to grasp movements and the dual channels hypothesis 2

Premotor cortex as distinct from primary motor cortex 5

Anatomical distinction of PMd and PMv 6

Functional distinction of PMd and PMv 8

Representation of movement in single unit activity in PMd 9

Representation of movement in single unit activity in PMv 10

Revisiting the dual channels hypothesis 12

Modern evidence that further argues against the dual channels hypothesis 14

Beta oscillations as a lens into motor cortical function 17

Hypotheses to be tested 21

CHAPTER 1: Encoding of joint kinematics in premotor cortex during reach-to-grasp 24

Abstract 24

Introduction 24

Methods 26

Results 33

Discussion 56

CHAPTER 2: Beta oscillations in premotor cortex during reach-to-grasp 62
Abstract 62
Introduction.....62
Methods.....65
Results 72
Discussion 92
DISCUSSION.....97
Relationship of PM and MI 97
Separating representations of function and anatomy 101
Implications for neural prosthetics.....104
Future work 106
APPENDIX 1: Spatio-temporal patterning in primary motor cortex at movement onset 111
Abstract 111
Introduction.....111
Material and methods.....112
Results 119
Discussion 130
APPENDIX 2: Comparing offline decoding performance in physiologically defined neuronal classes.....137
Abstract 137
Introduction.....137

Material and methods 140

Results 147

Discussion 161

REFERENCES 165

LIST OF FIGURES

Figure 1.1 Methods 34

Figure 1.2 Grasp aperture as a function of experimental condition 35

Figure 1.3 Trial-averaged kinematics of the arm, wrist, and hand 37

Figure 1.4 Demixed principal components analysis 39

Figure 1.5 Trial averaged dPCA trajectories relative to go cue 41

Figure 1.6 Encoding models that include kinematics predict spiking activity better than set related models 43

Figure 1.7 Comparing encoding performance of models based on an anatomical division of the upper limb 45

Figure 1.8 Comparing encoding performance of models based on a functional division of the upper limb 46

Figure 1.9 Comparing anatomical and functional representations in premotor cortex ... 48

Figure 1.10 Comparing spike history with kinematics 50

Figure 1.11 Comparison of spike history effects at multiple timescales 52

Figure 1.12 Predicting responses in premotor cortex using either PMd or PMv activity. 54

Figure 2.1 Power spectra 69

Figure 2.2 Average power spectra across electrodes 73

Figure 2.3 Beta amplitude as a function of behavior state 74

Figure 2.4 Temporal evolution of beta amplitude 75

Figure 2.5 Logistic curve goodness of fit 76

Figure 2.6 Timing of beta attenuation in PMd and PMv in two animals 77

Figure 2.7 Electrode-Electrode BAM correlations 78

Figure 2.8 Spatio-temporal propagation of beta attenuation across the cortical surface 79

Figure 2.9 Relationship of body representation and beta attenuation 81

Figure 2.10 Comparing different methods of estimating beta attenuation 82

Figure 2.11 Beta attenuation varies with experimental condition 84

Figure 2.12 Frequency of wave activity 86

Figure 2.13 Duration of wave activity 87

Figure 2.14 Direction of wave activity 88

Figure 2.15 Wave direction is not informative about experimental condition 91

Figure A1.1 Experimental setup and beta attenuation analysis.....120

Figure A1.2 Array placement and beta attenuation maps 122

Figure A1.3 Unit spiking analysis 126

Figure A1.4 Relationship of beta attenuation to unit modulation 128

Figure A2.1 Quantifying spike waveform width.....148

Figure A2.2 Bimodal distribution of spike waveform widths.....149

Figure A2.3 Decoding performance using narrow and wide spiking neural ensembles 152

Figure A2.4 The number of taps does not explain the difference in decoding performance.....154

Figure A2.5 Response properties of narrow and wide spiking neural ensembles 156

Figure A2.6 Narrow spiking neurons outperform wide psiking populations even after controlling for differences in response properties 158
Figure A2.7 Underlying differences in response properties do not explain the difference in decoding performance 160
Figure A2.8 Decoding muscle activity using narrow and wide neural ensembles 161

LIST OF TABLES

Table 1: Summary of datasets used in narrow versus wide decoding.....150

ACKNOWLEDGEMENTS

Advisor and mentor

Dr. Nicholas G. Hatsopoulos

Thesis committee

Dr. Stephanie Palmer

Dr. Jason MacLean

Dr. Lee Miller

Other faculty

Dr. Phil Ulinski

Dr. Sliman Bensmaia

Dr. Wim van Drongelen

Dr. David Freedman

Dr. Melina Hale

Dr. Dan McGehee

ARC staff

Marek Niekrasz

Craig Wardrip

Alison Ostdiek

Jennifer McGrath

Karin Peterson

Maggie Bruner

Hatlab colleagues

Dr. Kazutaka Takahashi

Dr. Aaron Suminski

Dr. Karthikeyan Balasubramanian

Dr. Fritzie Arce-MacShane

Kevin Brown

Adil Tobaa

Noah Huh

Mukta Vaidya

Alex Rajan

Adam Dickey

Jeff Walker

Kai Qian

Josh Coles

Carrie Balcer

Administrative staff

Diane Hall

Sharon Montgomery

Elena Rizzo

Animal colleagues

Jaco

Lester

Mack

Athena

Rockstar

Roxie

Raju

Velma

Boo

ABSTRACT

Reaching to grasp is an ethologically relevant primate behavior supported by a large network of cortical areas. One classical hypothesis concerning the organization of this network posited that information about reaching and grasping is processed along separate, independent pathways that are temporally coordinated via a central clocking mechanism. According to this hypothesis, known as the dual channels hypothesis, the dorsal premotor cortex (PMd) should be exclusively involved with reaching, while the ventral premotor cortex (PMv) should be exclusively involved in grasping. As a corollary of this hypothesis, the premotor cortex and primary motor cortex should be hierarchically organized.

Here, we used multi-electrode arrays to record from both PMd and PMv simultaneously while animal subjects engaged in a reach-to-grasp task. We then developed encoding models that predicted the activity of neurons in these areas based on various kinematic features. We found that the kinematics of both reaching and grasping were predictive of activity in both PMd and PMv.

Additionally, we examined the spatio-temporal dynamics of attenuation in the beta frequency range (15-30 Hz) of the local field potential in both areas. We found that beta attenuation propagated spatially across the cortical surface in PMd in a similar fashion to beta attenuation in MI.

Collectively, these results are inconsistent with the expectations of the dual channels hypothesis, and instead, imply a more direct role of premotor cortex in movement execution.

Introduction

"Engineers are philosophers of the hand"

Saturday morning breakfast cereal #3971

One of the most fundamental behaviors we engage in with our upper limb is the reach to grasp. This behavior first emerged in our primate ancestors as they foraged for insects in trees and has been conserved amongst many species of primate (Kaas, Gharbawie, and Stepniewska 2013; Marzke, n.d.). Supporting this behavior is a large network of sub-cortical and cortical structures (Grafton 2013). In this work, we will review a long-standing theory about the cortical organization of this network and focus on the role of the premotor cortex within it. To understand the motivations of this theory, I will introduce two psychophysical observations about reach to grasp movements that informed the development of this theory. I will then review the neuroanatomical and neurophysiological data that were available at the time this theory was developed. Subsequently I will describe new data that are inconsistent with the expectations of this theory before presenting a series of experiments that directly test it.

In the chapters that follow, I articulate new evidence that challenges this classical theory on the involvement of premotor cortex in reach to grasp movements. Instead, I will suggest a more complex, direct role of the premotor cortex in executing these movements. In chapter 1, I show that single cells in premotor cortex encode complex kinematic features that span the entire upper limb. I demonstrate that both dorsal and ventral premotor cortex (PMd and PMv, respectively) contain anatomical and functional representations of both reaching and grasping. In chapter 2, I characterize the spatio-

temporal patterns of oscillatory neural activity that precede movement onset and highlight differences between PMd and PMv. I also examine how the phase gradient of these oscillations propagates as planar waves across the cortical surface during the peri-movement epoch. In the discussion, the findings of chapters 1 and 2 will be contrasted with the expectations of the classic theory and some future experiments will be proposed to further elucidate the role of premotor cortex in executing reach to grasp movements.

The kinematics of reach to grasp movements and the dual channels hypothesis

Marc Jeannerod and colleagues engaged in a series of studies spanning the 1980s and 1990s that carefully analyzed the psychophysical performance of human subjects engaging in reach to grasp movements. Based on these findings, Jeannerod argued that reach to grasp movements arose from two discrete, independent control processes. The first process, transport, involved translating visual information about the location of an object in space into a motor command that drove the hand to that object. The second process, manipulation, was concerned with the hand actively preshaping to conform to the grasped object before the hand actually made contact with the object. These two processes naturally mapped onto the concepts of reaching and grasping, and, according to Jeannerod, were temporally coordinated via a central clocking mechanism. There were two key psychophysical observations that motivated this perspective.

The first observation concerned the independence of transport and manipulation processes. It was observed that perturbations applied to one component of the

movement, either transport or manipulation, minimally impacted the other (Jeannerod 1988; Paulignan, Jeannerod, et al. 1991; Paulignan, MacKenzie, et al. 1991). When the size of the object was increased, grasp aperture, defined as the distance between the tip of the thumb and index finger, and used as a proxy for the manipulation process scaled accordingly. However, no changes were observed in wrist speed, a proxy for the transport component. Inversely, when the distance between the hand and the object was increased, wrist speed increased, but no change in grasp aperture was observed (Jeannerod 1981; Jeannerod 1984). In this respect, it was argued that reach to grasp movements were modular and could be decomposed into separate reaching and grasping sub movements.

Although reach to grasp movements were shown to be modular, the resulting components were nevertheless temporally coupled. This coupling was demonstrated by showing that the timing of peak wrist deceleration was highly correlated with the timing of maximum grasp aperture (Jeannerod 1981). This moment in time was argued to correspond to the transition between transport and manipulation phases of the movement and occurred when the movement was 75% complete. Here, movement completeness was defined temporally, and measured in terms of the duration between the start of movement and making contact with the object. This effect was present even when subjects' vision was blocked after movement initiation, but before object contact (Jeannerod 1984). The temporal coordination of reaching and grasping movements was achieved via a clocking mechanism that would align the timing of peak wrist deceleration with the timing of peak grasp aperture.

Following these two central observations regarding the nature of reach to grasp, Jeannerod introduced a hypothesis about how these movements were organized in the cerebral cortex (Jeannerod et al. 1995; Jeannerod 1984; Jeannerod 1981; Jeannerod 1988). The "dual channels" hypothesis posited that the control processes supporting transport and manipulation were manifested in discrete neural networks that became activated during reach to grasp. Temporal coordination between channels was achieved by a central clocking mechanism that synchronized information from the two channels. When Jeannerod posed this hypothesis, the functional architecture of the motor cortex appeared to be consistent with the behavioral segregation of reaching and grasping. At least two distinct cortical networks seemed to be involved in reach to grasp; a dorsal network appeared to be specialized for the control of reaching movements (Steven P. Wise et al. 1997; Grafton 2013), while a ventral network supported grasping (Jeannerod et al. 1995; Kurata and Tanji 1986; G Rizzolatti et al. 1988).

The primary literature on the function and organization of motor cortex that was available to Jeannerod at the time he formulated the dual channels hypothesis will be reviewed. Subsequently, I will introduce additional findings that indirectly challenge the dual channels hypothesis as an organizational principle of motor cortex.

Before continuing, it should be noted that the robustness of Jeannerod's psychophysical observations has been challenged. Several studies (some including Jeannerod as an author) have failed to replicate his observations on reach to grasp movements (Jakobson and Goodale 1991; Paulignan, MacKenzie, et al. 1991; Paulignan, Jeannerod, et al. 1991; Paulignan et al. 1997). However, one recent study

observed a weak, but significant correlation between the timing of peak wrist deceleration and grasp aperture in non-human primates (Vaidya et al. 2015).

Premotor cortex as distinct from primary motor cortex

Some of the earliest investigations into motor cortical organization argued that there were at least two motor fields in motor cortex (Fulton 1935; Weinrich and Wise 1982), with Fulton being the primary proponent of a distinct premotor cortex. Here, references to motor cortex referred to the precentral cortex that evokes movement when electrically stimulated, excluding the eye fields (Fulton 1935). In modern terms, the cortical territory that Fulton referred to as premotor cortex (PM) included both what would now be considered premotor cortex (PM), as well as supplementary motor area (SMA). Fulton found that lesions to human PM led to deficits in the ability to perform skilled movements, spasticity, forced grasping, and vasomotor disturbances (Bucy and Fulton 1933; Kennard, Viets, and Fulton 1934).

After Fulton, the idea of distinct premotor and primary motor cortices fell out of favor for many years, largely due to the work of Woolsey. While Woolsey acknowledged there being multiple motor fields in the precentral cortex (corresponding in modern terms to primary motor cortex (MI), and SMA), he concluded what would now be considered PM was not involved in motor function. He was led to this conclusion because electrical stimulation of the premotor cortex in anesthetized primates failed to evoke movements, and ablations of PM after lesions to either SMA or MI failed to cause additional motor deficits (Woolsey et al. 1952).

Beginning around 1980, there were a series of investigations into the putative existence of a premotor cortex, and it was finally accepted as being a distinct motor area that chiefly corresponded to Brodmann's area 6 (Weinrich, Wise, and Mauritz 1984; Godschalk et al. 1984). To distinguish PM from other motor areas, multiple anatomical and functional criteria were proposed (Weinrich and Wise 1982; S. P. Wise 1985). Anatomically, premotor cortex was found to contain a thin layer IV in contrast to the agranular MI (Barbas and Pandya 1987). Likewise, layer V of MI was known to contain a population of giant pyramidal cells, known as Betz cells (Meyer 1987), but these cells were not observed in PM (Weinrich and Wise 1982). Functionally, intra-cortical micro-stimulation (Stoney, Thompson, and Asanuma 1968) (ICMS) of PM was shown to evoke movements, however, at currents that were much higher than those needed to evoke movement in MI (Weinrich and Wise 1982). Additionally, single unit activity in PM was shown to respond to visuomotor signals in addition to pure motor action (Weinrich and Wise 1982).

Anatomical distinction of PMd and PMv

During the same era that premotor cortical was recognized as being distinct from primary motor cortex, anatomical evidence also suggested that the premotor cortex could be further subdivided into dorsal and ventral regions each containing a motor field (PMd and PMv, respectively).

One line of evidence supporting separate dorsal and ventral fields in PM came from careful study of the inputs to PM from parietal cortex. Using horseradish peroxidase (HRP) staining to retrogradely label connections, it was shown that PMv

received inputs primarily from area 7b (Godschalk et al. 1984) while PMd received its inputs chiefly from area 5 (Matelli, Luppino, and Rizzolatti 1985; Barbas and Pandya 1987). In studying the patterns of inputs to PM from frontal areas, it was found that PMd contained more spatially localized inputs from frontal areas as compared to PMv, and that PMd received input from the dorsal aspect of the dorsolateral prefrontal cortex (PFC) while PMv received input from the ventral aspect of the dorsolateral PFC (Barbas and Pandya 1987).

In addition to receiving different inputs, projections from the putative arm region of PMd to MI were shown to be more medial than projections from the putative hand/orofacial area of PMv (Kurata 1991). However, this result must be interpreted with caution as spatially restricted retrograde labeling of the upper limb region of MI demonstrated that multiple discrete regions of the PM provided input to the same site in MI (Muakkassa and Strick 1979; Matsumura and Kubota 1979). In addition to differential connectivity, the cellular composition of PMd and PMv were also found to be different with PMd containing slightly larger pyramidal cells in layer V and PMv containing a slightly larger layer IV (Matelli, Luppino, and Rizzolatti 1985; Barbas and Pandya 1987).

One potential explanation for the anatomical differences between PMd and PMv comes from the fact that they are believed to arise from different developmental sources (Sanides 1964; Sanides, n.d.). PMd is thought to come from a protogradation originating in the cingulate gyrus. And accordingly, PMd shares connections with the SMA, a phylogenetically older division of the same protogradation (Barbas and Pandya

1987). PMv is thought to come from a protogradation originating in the insular cortex, and accordingly, shares connections with its putative architectonic precursor areas (Sanides, n.d.; Barbas and Pandya 1987).

Functional distinction of PMd and PMv

Throughout the 1980s as anatomical evidence increasingly supported the distinction between PMd and PMv, researchers began to consider what functional differences may also distinguish these areas. Some insight was gained from lesion studies that showed differences between the two areas. Lesions to PMd resulted in deficits in the ability to perform conditional visuo-motor association tasks (Petrides 1986; Halsband and Passingham 1985). In contrast, lesions to PMv resulted in deficits at making visually guided grasping movements (G. Rizzolatti, Matelli, and Pavesi 1983).

In addition to lesion studies, much of the early work employed ICMS (Stoney, Thompson, and Asanuma 1968) to gain insight into the functional role of each area through understanding its output effects. It was shown that upper limb movements could be evoked by stimulating PMd; however, the current intensities required to evoke movements were substantially higher than those needed in MI (Weinrich and Wise 1982). While ICMS revealed stimulation effects spanning the entire upper limb, a preponderance of the observed effects were limited to the proximal aspect of the limb, i.e. either the shoulder or the elbow (Weinrich and Wise 1982). In contrast, early attempts at stimulating PMv evoked hand and orofacial movements almost exclusively (Kurata and Tanji 1986; Gentilucci et al. 1988). Like PMd, threshold currents in PMv were also higher than MI. In addition to stimulation, researchers also began to probe

the function of PMd and PMv by recording the activity of single units using extracellular electrodes while animals engaged in motoric tasks.

Representation of movement in single unit activity in PMd

Even the earliest recordings of single unit activity in PMd revealed a relationship between spiking activity and movement in most neurons. It was found that approximately 75% of all premotor cortical neurons modulated their firing rate during movement as compared to a baseline period between movements (Weinrich, Wise, and Mauritz 1984). Furthermore, it was noted that some of the cells exhibited directional selectivity in their firing. That is, cells only modulated when the animal moved its arm in a specific direction (Weinrich, Wise, and Mauritz 1984). In this respect, the activity of neurons in PMd seemed quite similar to that of MI.

Where these two areas differed, however, was their activity during movement planning. These early studies of PMd employed a behavioral paradigm that experimentally dissociated movement planning from movement execution (Weinrich, Wise, and Mauritz 1984; A. Riehle and Requin 1989). This type of task, known as an instructed delay task, involved presenting multiple cues to the animal. The first cue, the instruction cue, would prime the animal to make a specific type of movement to successfully complete the task. The second cue, the go cue, would signal to the animal when to begin moving. A temporal delay, typically at least 1 second, was present between instruction and go cues, and referred to as the instruction epoch. It was thought that motor planning occurred during the instruction epoch because reaction times, the time between the go cue and movement onset, were much faster when

animals performed instructed-delay tasks as compared to simple cued movement tasks lacking separate instruction and go cues (for a review see (Alexa Riehle 2005)).

When animals engaged in instructed delay tasks involving upper limb movements, a variety of neural activity patterns was observed in the spiking activity of single cells in PMd (Weinrich, Wise, and Mauritz 1984; A. Riehle and Requin 1989). Some cells exhibited modulation, defined as a changed firing rate relative to baseline, only during movement. Qualitatively, this activity pattern was quite similar to that of most cells in MI. However, it was not especially common in PMd with only 26% of cells active in this way (A. Riehle and Requin 1989).

In contrast, a substantial proportion of PMd units exhibited "set-related" activity; that is, units modulated during the instruction epoch, and this modulation continued until movement completion. These activity patterns were thought to encode abstract motor plans that were not directly tied to any particular kinematic reference frame and led to the suggestion that PMd was involved in both the planning and execution of voluntary movements (Weinrich, Wise, and Mauritz 1984). Other cells in PMd exhibited modulation only during the instruction epoch, and no modulation during movement itself. For both set-related and instruction-only activity patterns, spiking activity during the instruction epoch often exhibited directional tuning. Notably, the motor behavior associated with many of these instructed delay tasks involved making planar reaches, despite knowledge of the fact that ICMS could evoke wrist movements in PMd.

Representation of movement in single unit activity in PMv

Some of the first single unit recordings from PMv observed cells that were exquisitely sensitive to somatic and somatovisual stimulation of the mouth, and to a lesser extent, the hands in reversibly paralyzed, anesthetized monkeys (Giacomo Rizzolatti et al. 1981). A similar experiment conducted in awake, behaving animals demonstrated that many of these neurons also responded robustly to active orofacial movements in addition to sensory stimulation (G. Rizzolatti et al. 1981). Later work established a clear relationship between single unit activity in PMv and distal limb movements (Kurata and Tanji 1986), but did not employ the same instructed-delay reaching paradigm used to study PMd. Instead, animals performed a simple cued movement task, although the cue to initiate movement was varied across different sensory modalities, and it was shown that some neurons in PMv selectively modulated their activity to cues given by a specific modality (e.g. an auditory tone). The type of movements that modulated PMv neurons was qualitatively different from those in PMd. Whereas cells in PMd exhibited simple preferences for movement direction, neurons in PMv showed modulation for more elaborate upper limb movements described as "grasping with the hand and mouth," "holding," and "tearing." (G Rizzolatti et al. 1988) In addition to selectivity for movement type, the cells were also sensitive to the overall conformation of the hand during the motor task, with cells exhibiting selectivity for grip type (G Rizzolatti et al. 1988).

It was not until much later that PMv was explicitly implicated in the preparation of upper limb movements (Hocherman and Wise 1991). In particular, one study recorded from both PMd and PMv while an animal engaged in an instructed delay reaching task

that involved moving to one of three targets using one of three curved trajectories. It was shown that cells in both PMv and PMd were modulated according to the location of the object (Hocherman and Wise 1991). However, a much larger proportion of cells in PMd were also modulated by the trajectory taken to the object (Hocherman and Wise 1991). Both areas contained cells that were modulated during the instruction epoch, and in roughly equal proportion (Hocherman and Wise 1991).

Revisiting the dual channels hypothesis

Based on this body of anatomical and physiological evidence, the premotor cortex appeared to be divided into dorsal and ventral regions specialized for the planning and execution of reaching and grasping movements, respectively. The role of MI in this framework, then, was to temporally coordinate these two independent inputs, as it was the first area in the motor hierarchy to receive convergent input from both streams. This basic hierarchical organization of motor cortex around prehensile movements, articulated by Jeannerod in the dual channels hypothesis remained largely unchallenged until recently. Indeed, experimental emphasis was largely placed on understanding more precisely how sensory information was gradually transformed into motor action. Chiefly, this was accomplished by attempting to characterize the coordinate frame of computation for different cortical areas (for a summary see (Crawford, Henriques, and Medendorp 2011; Kalaska 2009)).

Even when the dual channels hypothesis was formulated, certain experimental findings suggested that it might provide an incomplete view of motor cortical organization. Anatomically, it was shown that PMd contained cortico-spinal projections

to both upper and lower cervical segments of the spinal cord, and that these projections were spatially segregated. This challenged the dual channels hypothesis in two ways. First, it challenged the putative hierarchical organization of PMd and MI by showing that PMd could also directly act upon the spinal cord. In contrast, PMv was not found to send projections to lower cervical segments of the spinal cord, although some projections to propriospinal neurons in the upper cervical segments were observed (He, Dum, and Strick 1993).

Second, it demonstrated that PMd sent projections to both upper and lower cervical spinal segments, and thus contained a representation of the entire upper limb, including the wrist and hand. This map was loosely somatotopically organized, although there was substantial overlap between projections to the arm and hand (He, Dum, and Strick 1993). This finding was consistent with the observation that stimulation of PMd could occasionally evoke distal limb movements (Weinrich and Wise 1982). A more recent systematic study of the effects of ICMS in PMd revealed a somatotopic map of stimulation effects in the upper limb, complementing the anatomical findings (Raos et al. 2003). Furthermore, single unit recordings have found a population of neurons in PMd that are modulated during a naturalistic grasping task (Raos et al. 2004; Hendrix, Mason, and Ebner 2009).

Another piece of evidence that challenged the dual channels hypothesis also concerned the limb representation in PMv. The caudal portion of PMv was found to contain a sparse population of neurons that responded to proximal arm movements (Gentilucci et al. 1988; Hocherman and Wise 1991). Thus, although PMd and PMv

preferentially represented arm and hand movements, respectively, the division between the two areas in terms of upper limb representation was not absolute, i.e. PMv was not implicated exclusively in hand control.

Modern evidence that further argues against the dual channels hypothesis

Recent technological advances in motion tracking systems have enabled researchers to simultaneously monitor movement in many degrees of freedom in the arm and hand while animals engage in unconstrained, 3D prehensile movements (Vargas-Irwin et al. 2010; Bansal et al. 2012; Saleh, Takahashi, and Hatsopoulos 2012; Mollazadeh et al. 2014; Aggarwal et al. 2013). At the same time, advances in neural recording devices have now made it routine to record the activity of tens to hundreds of cells in a cortical area simultaneously (Stevenson and Kording 2011). Subsequently, there has been a renewed interest in elucidating the contribution of premotor cortex to prehensile behavior. Particularly, these methodological advances have enabled researchers to directly probe the extent to which information about movement kinematics is represented in populations of neurons in these two areas. It has been shown that detailed hand kinematics can be decoded from both PMd and PMv suggesting that a full representation of the upper limb is present in PMd (Bansal et al. 2012; Aggarwal et al. 2013; Mollazadeh et al. 2014). While reaching kinematics have not been shown to be concretely represented in PMv, the position of the hand in space can be decoded from PMv activity suggesting indirectly that reaching may also be represented in PMv (Bansal et al. 2012; Aggarwal et al. 2013).

The dual channels hypothesis predicted that independent reaching and grasping signals would be temporally coordinated via a central clocking mechanism. The most likely candidate location for this clock would be the primary motor cortex as it receives input from PMd and PMv and also sends feedback projections to both areas (He, Dum, and Strick 1993). The mechanistic implementation of this clock, however, remained ambiguous in Jeannerod's statement of the dual channels hypothesis. Two potential neural implementations of the clocking mechanism will be discussed below. The first potential mechanism would involve single cells in MI integrating information from the two streams. There has been some debate, however, as to the extent to which the inputs from PM may be convergent in MI.

Anatomically, it has been shown that projections from the hand representation of PMv are not homotopic with MI. That is, the hand representation of PMv projects to arm, wrist, and hand representations in MI (Dancause et al. 2006). Additionally, although these connections span the entirety of the upper limb, the projections from PMv to MI are not homogenous throughout the entire upper limb region of MI, but rather form discrete functional clusters (Dancause et al. 2006). Presently, it remains unknown if these discrete functional clusters receive convergent input from PMd. However, it has been shown that some feedback projections from spatially localized sites in MI project to both PMd and PMv (Dum and Strick 2005). Functionally, it has been demonstrated that single MI neurons encode kinematics from multiple joints spanning both proximal and distal aspects of the limb (Saleh, Takahashi, and Hatsopoulos 2012).

An alternative model suggests that temporal coordination of reaching and grasping does not arise because of convergent input onto single cells, but rather is a product of dense horizontal connectivity within MI linking ensembles of neurons receiving inputs from reaching and grasping channels. In contrast to the previous study concerning single cell encoding, more recent work identified putative reaching and grasping populations of cells in MI, and argued that these two populations were temporally coordinated via a proportional derivative control law (Vaidya et al. 2015). Additionally, it has been shown that spatially expansive populations of MI neurons representing proximal and distal limb segments form large functionally connected networks during movement execution (Takahashi et al. 2015). This network is oriented consistently with large scale spatio-temporal patterns of neural activity that have been hypothesized to coordinate proximal and distal segments of the limb during reach to grasp (Alexa Riehle et al. 2013; Hatsopoulos, Olmedo, and Takahashi 2010). The precise propagation direction of spatio-temporal patterns would then be determined by the relative timing of information arriving from the two streams (Hatsopoulos, Olmedo, and Takahashi 2010).

Some recent studies, however, suggest that interpreting the previous findings as evidence of a clocking mechanism is incorrect. Generally, there is computational and physiological evidence to support the position that neural circuits exhibit characteristic temporal structure (Buonomano and Laje 2010; Goel and Buonomano 2014), and, a dynamical systems perspective has become the new standard approach to understanding population activity in motor cortex (Churchland et al. 2012; Sussillo et al.

2015; Shenoy et al. 2011). In the classical formulation of the dual channels hypothesis, there was an explicit representation of how long it would take to complete the reach, and the grasp; some mechanism would ensure that both movements were 75% complete (corresponding to the time of peak wrist deceleration and maximum grasp aperture) at the same time. To date, there has been little evidence to suggest that any single cells in MI function in this way. Moreover, in light of these new theories on time and cortical dynamics, it could be argued that no such explicit clocking mechanism may exist at all.

Instead, the temporal coordination of reaching and grasping may arise from the intrinsic dynamics of the motor cortical activity generating the movement. In this view, neurons do not explicitly represent time, or the temporal evolution of any movement parameter, but rather, they reflect the internal state of a complex, possibly chaotic (Laje and Buonomano 2013) dynamical system that generates movement (Shenoy et al. 2011). The preceding anatomical and physiological observations in MI are not incompatible with this new perspective. Specifically, ensemble spiking activity of neural populations in MI would move the system to a neural state from which coordinated reaching and grasping movements follow (Shenoy et al. 2011; Churchland et al. 2010). The goal of the dynamical systems perspective, then, is to discover the laws that govern the temporal evolution of neural activity in state space (Shenoy et al. 2011), and further understand how those laws give rise to coordinated reaching and grasping behavior.

Beta oscillations as a lens into motor cortical function

While the dual channels hypothesis is very much concerned with understanding the type of information that is represented in PMd and PMv, consideration of single unit tuning properties is not the only way to gain insight into motor cortical function. Indeed, electrophysiologists began recording motor cortical activity long before the discovery of single unit recording (Li and Jasper 1953). These early recordings were collected via large electrodes placed on either the scalp or cortical surface (P. D. H. Berger 1929; H. Berger 1933; Jasper and Penfield 1949). It was found that cortical activity could be decomposed into the sum of multiple time-varying sinusoidal curves. In particular, neural activity was often broken into two different frequency ranges: the alpha range (8-13 Hz) and the beta range (15-30 Hz).

Even the earliest studies were able to identify relationships between cortical rhythms and behavior (P. D. H. Berger 1929; Jasper and Penfield 1949). In particular, it was shown that the amplitude of beta oscillations attenuated when movements were performed (Jasper and Penfield 1949). Subsequent research has confirmed this phenomenon is a robust feature of motor cortical activity in human (Stancák Jr. and Pfurtscheller 1996; G. Pfurtscheller et al. 2003; Takahashi et al. 2011) and nonhuman primates (Sanes and Donoghue 1993; Donoghue et al. 1998; Rubino, Robbins, and Hatsopoulos 2006; Murthy and Fetz 1996a). Beta attenuation is not limited to upper limb movements; it has been documented during leg movements (Crone et al. 1998) as well as tongue movements (Crone et al. 1998; G. Pfurtscheller et al. 2003). Within a major body segment, beta attenuation appears to be widespread. In one study, it was found that simple finger flexion movements resulted in attenuation spanning the entire

upper limb region (Stancák and Pfurtscheller 1995). However, across body segments, beta amplitude actually seems to increase, that is, beta amplitude would increase in the leg representation during arm movements (Crone et al. 1998). It is also not limited to executed movements; in human subjects, beta attenuation occurs during motor imagery (Gert Pfurtscheller and Neuper 1997).

Despite its ubiquity as a cortical correlate of motor behavior, beta amplitude at single cortical sites does not appear to convey much information about movement. In the context of the upper limb, beta amplitude does not reflect the direction (Waldert et al. 2008; Rickert et al. 2005; Mehring et al. 2003), force (Tombini et al. 2009), or speed (Stancák and Pfurtscheller 1995; Stancák Jr. and Pfurtscheller 1996) of movements. This has led to some debate about the functional role of beta attenuation in motor cortex.

On one hand, some have argued that the beta oscillation represents an active braking process inhibiting movement production (Gilbertson et al. 2005; Pani et al. 2014). In support of this argument is the finding that beta amplitude is often high during epochs of postural stability (Gilbertson et al. 2005). Paradoxically, beta amplitude is also high during epochs of isometric force production (Murthy and Fetz 1992; Conway et al. 1995), which has led to the suggestion that high amplitude beta oscillations indicate the motor system is maintaining the status quo (Engel and Fries 2010). In this view, beta amplitude has been interpreted as an abstract indication of motor readiness or attention (Murthy and Fetz 1996a; Donoghue et al. 1998).

An alternative, albeit more limited perspective, argues that the specific phenomenon of beta attenuation serves as a mesoscopic reflection of the transition into movement execution-related activity (Jasper and Penfield 1949; Gert Pfurtscheller and Lopes da Silva 1999; Kilavik et al. 2013). There is an increasing body of experimental evidence supporting this interpretation of beta attenuation. Using an instructed-delay paradigm, it has been shown that beta oscillations are typically high in amplitude during motor planning (O'Leary and Hatsopoulos 2006). During epochs of high amplitude beta oscillations (not restricted to motor planning), single cells in motor cortex are known to entrain to the oscillation (Murthy and Fetz 1996b; Denker et al. 2007), but, this entrainment does not necessarily cause changes in the firing rate of these cells (Denker et al. 2007). In contrast, around the time of movement onset, the temporal dynamics of spiking in single cells become complex and heterogeneous across cells (Churchland and Shenoy 2007) as they begin to exhibit execution related activity. Correspondingly, spiking activity desynchronizes from the beta oscillation (Murthy and Fetz 1996b). Thus, beta attenuation may be interpreted as a mesoscopic reflection of the transition from phasic to complex, movement-related spiking patterns in ensembles of cells. In further support of this interpretation, it was recently shown that the spatio-temporal sequence of beta attenuation across the cortical surface in MI mirrors the spatio-temporal engagement order of motor cortical units in movement execution (Appendix 1).

This specific interpretation of beta attenuation does not prescribe a single, general role for beta oscillations in motor cortex. Nevertheless, it has some important implications for the genesis of beta oscillations. Specifically, it implies that beta

oscillations are, at least in part, a product of local cortical circuitry. Presently, the mechanistic origins of the beta oscillation remain elusive (for a recent review of three current hypotheses see (Khanna and Carmena 2015)).

Further complicating this interpretation of beta attenuation is that beta amplitude is high during sustained force production. During sustained force production, many cells in motor cortex exhibit modulated firing rates throughout the duration of force production (Arce-McShane et al. 2014; Evarts et al. 1983), though it remains unknown if the nature of this modulation is the same as modulation during transient changes in force levels. Another possible explanation is that sustained force production does not require the precise spatio-temporal sequencing of muscle activation that accompanies reaching movements (Scott 1997; Haggard and Wing 1995; Karst and Hasan 1991), but rather, large-scale simultaneous muscle activation (Khanna and Carmena 2015). In order to activate muscles in this fashion, there would have to be a large excitatory drive, a necessary condition for evoking large amplitude beta oscillations (Roopun et al. 2006).

Hypotheses to be tested

Presently, the dual channels hypothesis appears to be an increasingly inadequate theoretical model of motor cortical organization. In the sections that follow, we will explore the role of PMd and PMv in generating and executing reach to grasp movements. In chapter 1, we approach the question of representation in PM. By combining state of the art motion tracking technology with neural array recordings, we will ask if reaching and grasping are truly processed along separate, independent pathways. It will be demonstrated that information about reach to grasp movements is

not segregated along two pathways. We will show that whole upper limb kinematics are encoded in both PMd and PMv. Based on this finding, we will argue for a complete representation of the upper limb in each area. Furthermore, we will show that information about reaching and information about grasping is encoded by the spiking activity of cells in each area. Although we cannot argue representations in each area are discrete, there remain some differences between areas; chiefly, kinematics are more strongly represented in PMd than PMv. The dual channels hypothesis will be discussed in light of these two findings.

We will subsequently consider the role of premotor cortex in voluntary movement initiation in chapter 2. We will ask if PMd and PMv are differentially involved in movement initiation through consideration of the beta attenuation phenomenon in each area. We will explore the extent to which beta attenuation is spatio-temporally sequenced in each area and how it varies with experimental parameters associated with movement. It will be shown that the directionality and overall timing of beta attenuation are not constant across the various experimental conditions. Additionally, we will examine PMd and PMv LFP activity for evidence of planar wave propagation in the beta frequency range (Rubino, Robbins, and Hatsopoulos 2006). The existence of these waves during reach to grasp movements will be established. It will be confirmed that the directionality of these waves does not convey information about the experimental condition.

Following chapter 2 will be a discussion of the dual channels hypothesis informed by the findings of this thesis as well as a proposal for future work. The appendix contains other work that I completed during my time as a graduate student.

Chapter 1 - Encoding of joint kinematics in premotor cortex during reach-to-grasp

Abstract

Classically, it has been hypothesized that reach-to-grasp movements arise from two discrete cortical networks. As part of these networks, the dorsal premotor cortex has been implicated in the control of reaching movements, while the ventral premotor cortex has been implicated in the control of grasping movements. Here, we recorded from PMd and PMv simultaneously using multi-electrode arrays implanted in each area while rhesus macaques engaged in a reach-to-grasp task. Generalized linear models were used to predict the spiking activity of cells in both areas as a function of different kinematic parameters, as well as spike history. We found that single cells in both PMd and PMv encode the kinematics of the entire upper limb suggesting that this classical division of reach and grasp in PMd and PMv, respectively, does not accurately reflect the encoding preferences of cells in those areas.

Introduction

Reaching to grasp is a fundamental, ethologically relevant primate behavior (Kaas, Gharbawie, and Stepniewska 2013). Several cortical and sub-cortical structures are involved in the generation of these movements (Kaas, Gharbawie, and Stepniewska 2013; Grafton 2013). Recently, much emphasis has been placed on understanding the behavior of single cells in primary motor cortex (MI), one of the main sources of cortical output to the motor periphery, during reach to grasp movements (Saleh et al. 2010; Saleh, Takahashi, and Hatsopoulos 2012; Mollazadeh et al. 2011; Mollazadeh et al. 2014). The motor cortex, however, does not exhibit a clear hierarchical organization

with MI as the sole source of cortical output (He, Dum, and Strick 1993; Richard P. Dum and Strick 2005). Rather, MI is part of a network of cortical areas involved the production of reach to grasp movements. Two other cortical areas that are densely connected horizontally with MI are dorsal and ventral premotor cortex (PMd and PMv, respectively) with these two areas also providing a large proportion of cortical inputs to MI (Richard P. Dum and Strick 2005).

In the present study, we sought to characterize the encoding properties of neurons in PMd and PMv to better understand the functional organization of the motor cortex. Classically, it was thought that reach to grasp movements arose from the temporal coordination of activity in independent reaching and grasping brain networks (Jeannerod 1988). In this framework, PMd was part of the reaching network, and thus, was concerned with reaching movements and the proximal segment of the limb, while PMv was part of the grasping network and concerned with distal aspects of the limb, namely, the hand (Grafton 2013; Hoshi and Tanji 2007).

Here, we directly tested whether representations of reaching and grasping were strictly segregated in PMd and PMv. We applied a novel dimensionality reduction technique to dissociate anatomical representations (i.e. based on encoding specific joints) from functional representations (i.e. encoding reach-related or grasp-related kinematic synergies). We found no evidence of a strict segregation of reaching and grasping activity in PMd or PMv. Instead, we observed that both areas contained complete representations of the anatomy of the upper limb, and functional reach and grasp representations.

Methods

Neurophysiology

All surgical and experimental procedures were approved by the University of Chicago Animal Care and Use Committee and conformed to the principles outlined in the Guide for the Care and Use of Laboratory Animals (NIH publication no 86-23 revised 1985). Two rhesus macaques (*macaca mulatta*) were implanted with 96-electrode Utah arrays in the dorsal and ventral premotor cortex contralateral to their working arm. Electrodes were 1.5 mm in length except for the array in PMd of animal J (1.0 mm). Neural spiking activity from the electrodes were amplified with a gain of 5000, band-pass filtered between 0.3 Hz and 7.5 kHz, and recorded digitally (14 bit resolution) using a Cerebus acquisition system (Blackrock Microsystems, Salt Lake City, UT). A threshold was set above the noise floor on each channel (5.5 standard deviations) and every time this threshold was crossed a 1.6 ms sample (sampled at 30 kHz) was recorded as a putative spike waveform. These threshold crossings were subsequently sorted offline using a semi-manual clustering procedure (Offline Sorter, Plexon Inc., Dallas, TX). Only sorted waveforms with a signal-to-noise ratio (SNR) greater than 3 were used in subsequent analyses. Here, signal to noise ratio is defined as the peak minus the trough of the average waveform divided by the average standard deviation of the waveform across time.

Behavioral Task

Two male rhesus macaques were trained to perform to perform a reach to grasp task with their left hand. A robot (RV-1A-S11 6-axis robot; rixan.com) presented each

animal with objects to grasp at four different spatial locations in its peripersonal space. The set of objects consisted of geometric shapes that were designed to evoke a variety of different hand conformations when grasped. Object and location pairings were randomly varied on a trial-by-trial basis to minimize the amount of motor planning the animal could do. Additionally, the animal's vision was occluded between trials with a pneumatically controlled screen.

The precise set of objects that were presented to animal L consisted of 5 different objects (Cylinder, Small Disc, Key, Large Disc, Ring) in different orientations for a total of 11 different grasp conformations (Cyl Horizontal, Cyl Out, Cyl Vertical, S Disc O, Key, L Disc H, L Disc V, Ring H, Ring V, S Disc H, S Disc V). On average, the animal completed 21 repetitions of each object/location pairing, although some combinations were less frequent than others due to the random sampling of objects, with the minimum number of repetitions being 5. Animal J was presented with three objects (Cylinder, Small Disc, Ring) to evoke four different grasp conformations (Cyl O, S Disc O, Ring H, Ring V). On average, animal J completed 36 repetitions of each object/location pairing, and all combinations were sampled at least 15 times.

Motion capture and inverse kinematics

A 10-camera motion capture system (Vicon Motion Tracking System, Oxford, UK) detected and recorded the 3-D position of retroreflective markers affixed to the animal. A total of 30 markers were placed on the animal's dorsal hand and dorsolateral arm enabling the tracking of 21 degrees of freedom in the arm and hand (listed below). The motion capture system was also used to monitor the position of the object and vision

block screen. The time series of 3D marker positions were low pass filtered (bi-directional 4th order Butterworth filter 15 Hz cutoff). To measure grasp aperture, we computed the Euclidean distance between markers attached to the distal most aspects of the thumb, and index finger (D1 and D2, respectively).

To compute wrist speed, we numerically differentiated the positional data of three markers on the animal's wrist. After numerical differentiation, wrist velocities were again low pass filtered (using the same filter design as before). Wrist speed was defined as the average tangential velocity of the three wrist markers. Movement onset, derived from wrist speed, was defined as the first moment that wrist speed exceeded 10% of the max wrist speed. Reaction times were subsequently defined as the difference in time between movement onset, and release of the vision block.

We used open-source software (Delp et al. 2007) to infer joint angles from marker positional data. We estimated joint angles for 21 degrees of freedom in the arm and hand including: humerus tilt, humerus rotation, elbow flexion, wrist pronation, wrist adduction, wrist flexion, 1 carpal-metacarpal (CMC) flexion, 1 CMC adduction, 5 CMC extension, 1 metacarpal-phalangeal (MCP) flexion, 1 MCP adduction, 2 MCP adduction, 2 MCP flexion, 2 phalangeal-medial (PM) flexion, 3 MCP flexion, 3 MCP adduction, 4 MCP flexion, 4 MCP adduction 4 PM flexion, 5 MCP flexion, and 5 MCP adduction. Inverse kinematics were bidirectionally low pass filtered with a 6 Hz cutoff.

Demixed principal components analysis

We used a novel dimensionality reduction technique, demixed principal components analysis (dPCA), to identify kinematic synergies associated with reaching,

grasping, the interaction of reach and grasp, or synergies that were common across all experimental conditions (Kobak et al. 2014; Brendel, Romo, and Machens 2011). This analysis was motivated by two observations about the kinematics. First, on a given trial, there are often strong correlations between the kinematics of different joints suggesting that most of the variance in the kinematics may be described by a few kinematic synergies, often estimated using principal components analysis (Santello and Soechting 1998; Mason, Gomez, and Ebner 2001). Second, the correlation patterns across joints during reach-to-grasp movements depend upon the exact nature of the movement, thus, different experimental conditions may produce different correlation patterns (Todorov and Ghahramani 2004). Briefly, the dPCA algorithm attempts to find linear combinations of kinematic variables that explain the most variance, yet also vary only with a given experimental parameter.

We will now present a mathematical explanation of dPCA that closely mirrors the description in (Kobak et al. 2014). Suppose that every trial, e (of E total trials), has an associated location, l (out of L possible locations), and object o (out of O possible objects). Also, suppose that on every trial we have recorded the activity of J kinematic features at T different time points. We computed the trial-average kinematic trajectory of every kinematic feature, j, for each object and location pairing denoted $R_{jlo}(t)$. We then gathered all the trial-averaged kinematic trajectories into a single matrix, $\boldsymbol{X}$ of size J x LOT. We subsequently demeaned $\boldsymbol{X}$ such that the average of any kinematic feature over all locations, objects, and times was 0. It has been previously shown that for such

a matrix, X, it can be decomposed into independent parts called marginalizations that satisfy:

$$\boldsymbol{X} = \boldsymbol{X}_t + \boldsymbol{X}_{tl} + \boldsymbol{X}_{to} + \boldsymbol{X}_{lto} = \sum_{\phi} \boldsymbol{X}_{\phi}$$

Where $\boldsymbol{X}_t$ denotes the time varying, but location and object-invariant part of $\boldsymbol{X}$, obtained by averaging X over all locations and objects, i.e.

$$\boldsymbol{X}_t = \frac{1}{LO} \sum_{l=1}^{L} \sum_{o=1}^{O} \boldsymbol{X}_{\boldsymbol{tlo}}$$

The location dependent term, $\boldsymbol{X}_{tl}$, is obtained by averaging $\boldsymbol{X} - \boldsymbol{X}_t$ over all objects, and similarly, the object dependent term, $\boldsymbol{X}_{to}$ is obtained by averaging $\boldsymbol{X} - \boldsymbol{X}_t$ over all locations. Finally, the location-object interaction, $\boldsymbol{X}_{lto}$ is obtained by computing $\boldsymbol{X}_{lto} = \boldsymbol{X} - \boldsymbol{X}_t - \boldsymbol{X}_{tl} - \boldsymbol{X}_{to}$. It is important to note that every marginalization of $\boldsymbol{X}$ has the same dimension as $\boldsymbol{X}$, i.e. J x LOT.

Using the decomposition of $\boldsymbol{X}$ into its marginalizations, dPCA aims to find directions in $\mathbf{R}^J$ that explain as much variance as possible, with the additional constraint that this variance should come from only one marginalization. Unlike standard PCA, dPCA relaxes the assumption that all the directions must be strictly orthogonal. Formulated as a constrained optimization problem, dPCA seeks to minimize the following objective:

$$L = \sum_{\phi} \left\| \boldsymbol{X}_{\phi} - F_{\phi} D_{\phi} \boldsymbol{X} \right\|_F^2$$

where F_{ϕ} is an encoder matrix with q columns, D_{ϕ} is a decoder matrix with q rows, and $\|.\|^2_F$ denotes the Frobenius norm of a matrix. An efficient algorithm was recently proposed to solve this objective function (Kobak et al. 2014).

In order to compute the cumulative variance explained by the first *p* dPCs, we cannot simply add variances because the dPCs are not strictly orthogonal. Instead, the cumulative variance explained by the first *p* dPCs is given by:

$$\frac{\|\boldsymbol{X}\|_F^2 - \left\|\boldsymbol{X} - [f_{1:p}][d_{1:p}]'\boldsymbol{X}\right\|_F^2}{\|\boldsymbol{X}\|_F^2}$$

where $[f_{1:p}]$ is a matrix of the first *p* encoding vectors and $[d_{1:p}]$ is a matrix of the first *p* decoding vectors. We emphasize that the cumulative proportion of variance explained by each component is based off of the trial averaged data in the matrix, $\boldsymbol{X}$.

Generalized linear model analysis

Input features

We attempted to predict the instantaneous probability of spiking in a given neuron from many different combinations of extrinsic covariates (i.e. kinematics and experimental factors) and intrinsic covariates (i.e. neural signals).

Extrinsic covariates

We used the position and velocity of 21 joints in the arm and hand at various time lags as input features to our encoding model. We used 8 different time lags from -156 ms (i.e. spikes lag kinematics by 156 ms) to 208 ms (i.e. spikes lead kinematics by 208 ms) in 52 ms steps. In total, there were 168 (21 joints x 8 lags) positions and velocities for a total of 336 kinematic features. Other extrinsic features included kinematic synergies identified via demixed principal components analysis (described below). Additional categorical variables indicating the reach location and object on a given trial were defined; inter-trial intervals were coded as one category in these variables.

Intrinsic covariates

In addition to extrinsic features, intrinsic features such as a given neuron's own spike history and the spike history of other neurons may also contribute to a neuron's firing rate. To account for the fact that spike history may be relevant in different ways at different time scales, we filtered binary spike trains with raised cosine basis functions of the form:

$$y(t) = 0.5\cos(a\log(t+c) - \phi) + 0.5$$

for t such that $a\log(t+c) \in [\phi - \pi, \phi + \pi]$, and 0 elsewhere (Saleh, Takahashi, and Hatsopoulos 2012). We specified three different values of ϕ, the temporal peak of the cosine curve, to account for short (16 ms), medium (44 ms), and long (108 ms) time scale spike history effects. Each basis vector was convolved with the binary spike train of a given neuron, thus giving rise to 3 spike history vectors for every neuron.

Logistic regression

We used logistic regression to predict the probability that a neuron fired a spike in a small time window (4 ms) based on previously described input features. Mathematically, this model may be expressed as:

$$\log\left[\frac{p_n(t)}{1-p_n(t)}\right] = \beta_0 + \sum_{k=1}^{K}\beta_k X_k(t) + \sum_{j=1}^{J}\beta_j X_j(t)$$

where $p_n(t)$ is the probability that neuron n fires a spike at time t, β_0 represents the baseline probability that the cell will spike, $X_k(t)$ is the value of the k^{th} (of K) extrinsic feature at time t, $X_j(t)$ is the value of the j^{th} (of J) intrinsic feature at time t with corresponding weights β_k and β_j, respectively.

Data segmentation and assessing goodness of fit

Model goodness of fit (GoF) was quantified using the area under the receiver operating characteristic curve (AUROC) (Saleh et al. 2010; Saleh, Takahashi, and Hatsopoulos 2012; Hatsopoulos, Xu, and Amit 2007; Truccolo, Hochberg, and Donoghue 2010). We assessed goodness of fit using 10 folds of cross-validated test data, and, unless otherwise noted, all reported AUROC values are taken to be the median across cross-validation folds. Each training fold was comprised of at least 329 experimental trials and inter-trial activity between those trials. Because the folds were based on experimental trials rather than data points, the number of test and training data points varied across folds, however, at least 531,747 and 1,541,496 data points were in every fold in animals J and L, respectively. The biggest model we fit contained 418 features, so the ratio of data points to features was always greater than 1,000.

Results

The kinematics of unconstrained reaching to grasp

We used a marker-based infrared motion tracking system combined with an advanced biomechanical model of the primate upper limb (Delp et al. 2007) to infer the kinematics of 21 joints in the arm and hand as two rhesus macaques engaged in an unconstrained reach-to-grasp task. We presented multiple objects at different locations in the animal's workspace designed to elicit a diversity of prehensile movements that included both precision and power grips (Fig. 1.1).

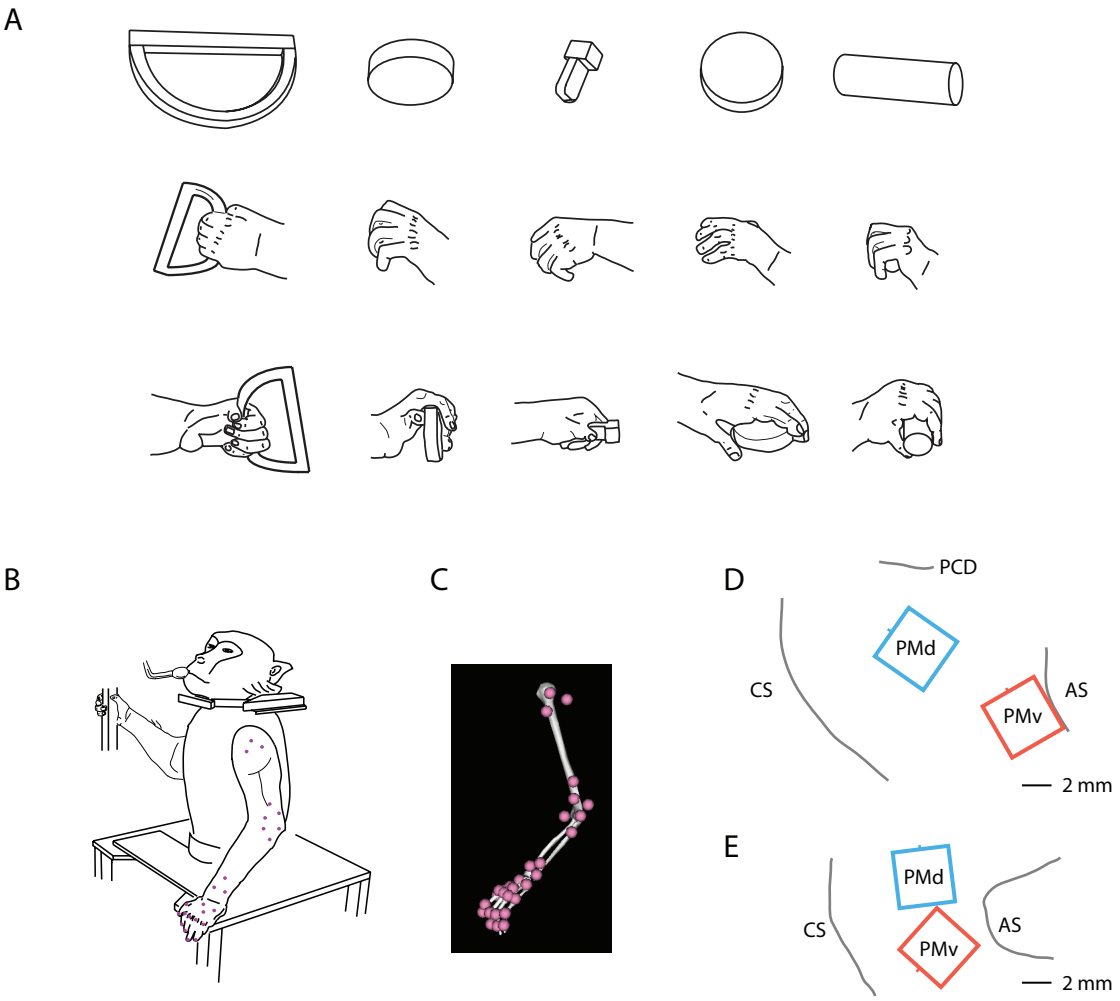

Figure 1.1 Methods. A. Animals were trained to grasp a variety of geometric shapes including (left to right) ring (vertical), small disc (out), key, small Disc (horizontal), cylinder (horizontal). These drawings depict static hand conformation during grasping of these objects. **B.** Drawing of animal J in rest position. Pink dots correspond to the approximate placement of the infrared markers. **C.** Depiction of the arm in the same position as **B.** in the Opensim software environment. **D-E.** Placement of electrode arrays in animals L and J, respectively.

We observed that the average temporal profiles of grasp aperture (defined as the distance between the distal-most markers on digits 1 and 2) were qualitatively different across experimental conditions (Fig. 1.2A). Quantitatively, the maximum grasp aperture was significantly different across the four grasping conditions shown in Fig. 1.2B (Kruskal Wallis (KW) test, animal J: X^2 with 3 d.f. = 330.84, $p < 0.00000001$; animal L: X^2 with 3 d.f. = 199.91, $p < 0.00000001$). Similarly, we observed differences in the peak

wrist speed across object locations (result not shown; KW test, animal J: X^2 with 3 d.f. = 87.16, $p < 0.00000001$; animal L: X^2 with 3 d.f. = 272.19, $p < 0.00000001$).

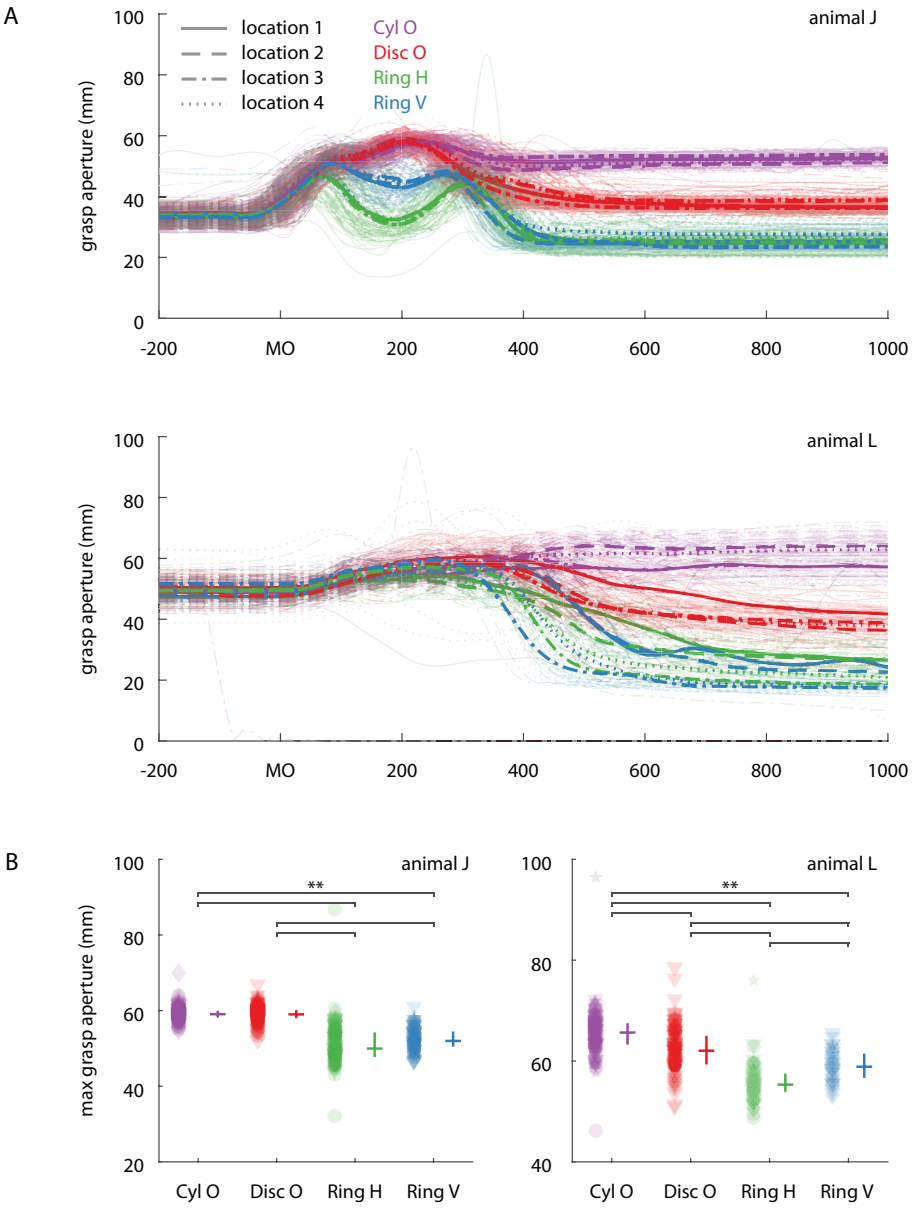

Figure 1.2 Grasp aperture as a function of experimental condition. A. We computed the grasp aperture, defined as the distance between the distal-most thumb and index finger markers and plotted it relative to movement onset (MO) in animal J (top) and animal L (bottom). Each semi-transparent line corresponds to the temporal profile of grasp aperture on a single trial, while the thick, opaque lines correspond to within-condition trial averages. Different line styles indicate the different object locations while color is used to indicate the type of object that was grasped. **B.** We found that grasp aperture varied significantly with the object to be grasped. Here, object location is indicated using the marker style (location 1: diamonds, location 2: circles, location 3:

triangles, location 4: squares). Raw data are shown in the left of each column, while the median and interquartile range are shown on the right. The black horizontal lines at the top of the panel indicate statistically significant differences at $p < 0.005$ (Bonferroni corrected *t*-test, Matlab function multcompare).

Differences in kinematics across experimental conditions were also evident at the level of individual joints (Fig. 1.3). From visual inspection of Fig. 1.3A, it can be seen that the different object locations elicited different patterns of arm movement and the different object types elicited different patterns of hand movement. Less obvious, however, is that the object to be grasped also influenced arm kinematics and, to a lesser extent, the position of the object in space influenced hand kinematics (Fig. 1.3B).

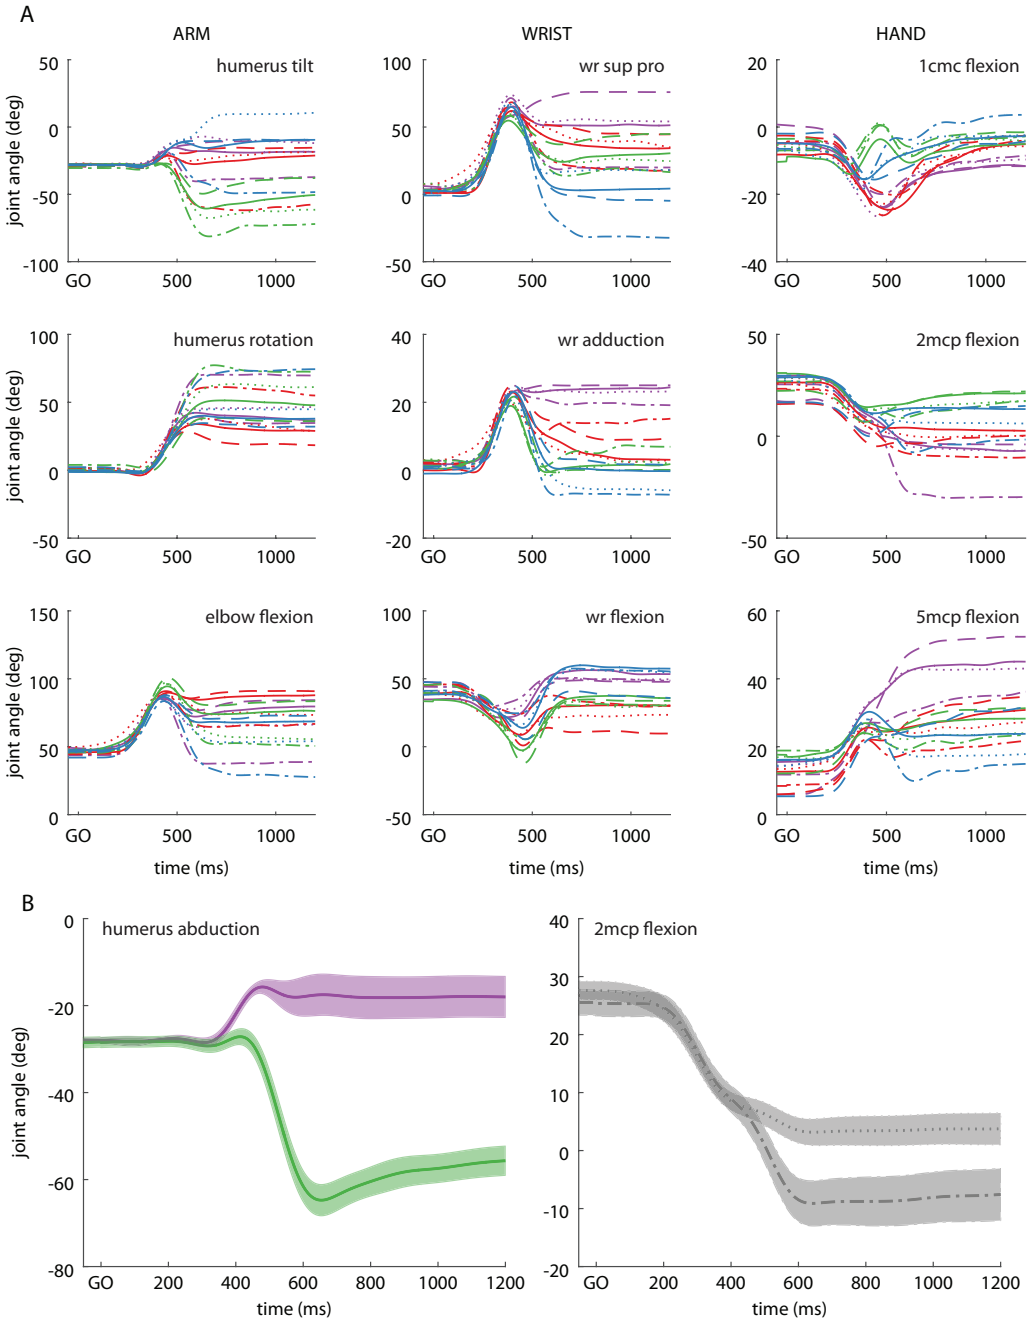

Figure 1.3 Trial-averaged kinematics of the arm, wrist and hand. A. Within-condition trial-averaged kinematics for three degrees of freedom in the arm (left column), wrist (center column) and hand (right column) of animal J. Line style indicates object location while color indicates object; conventions the same as Fig. 1.2. **B.** We found that the object can significantly affect arm kinematics (left). Here, we computed the average humerus abduction over all object locations when grasping Cyl o (purple trace) and ring h (green trace). Shaded area indicates +- 2 standard errors of the mean (S.E.M.) Similarly, reach location can affect hand kinematics (right). Flexion in the metacarpophalangeal joint of the second digit (2mcp flexion) was computed over all objects and is shown as a function of object location.

To control for the fact that both reaching and grasping affect kinematics in the entirety of the upper limb, we employed a novel dimensionality reduction technique, demixed principal components analysis (dPCA), to identify functionally defined kinematic synergies. Each kinematic synergy (i.e. each dPC) could be either reach-related, grasp-related, condition-independent (i.e. common to all experimental conditions), or reach-grasp interaction (i.e. specific to a particular experimental condition). As with other studies that have applied dimensionality reduction techniques to kinematic data, we found that a handful of functional kinematic synergies described the majority of the variability in the kinematics (Fig. 1.4A-D).

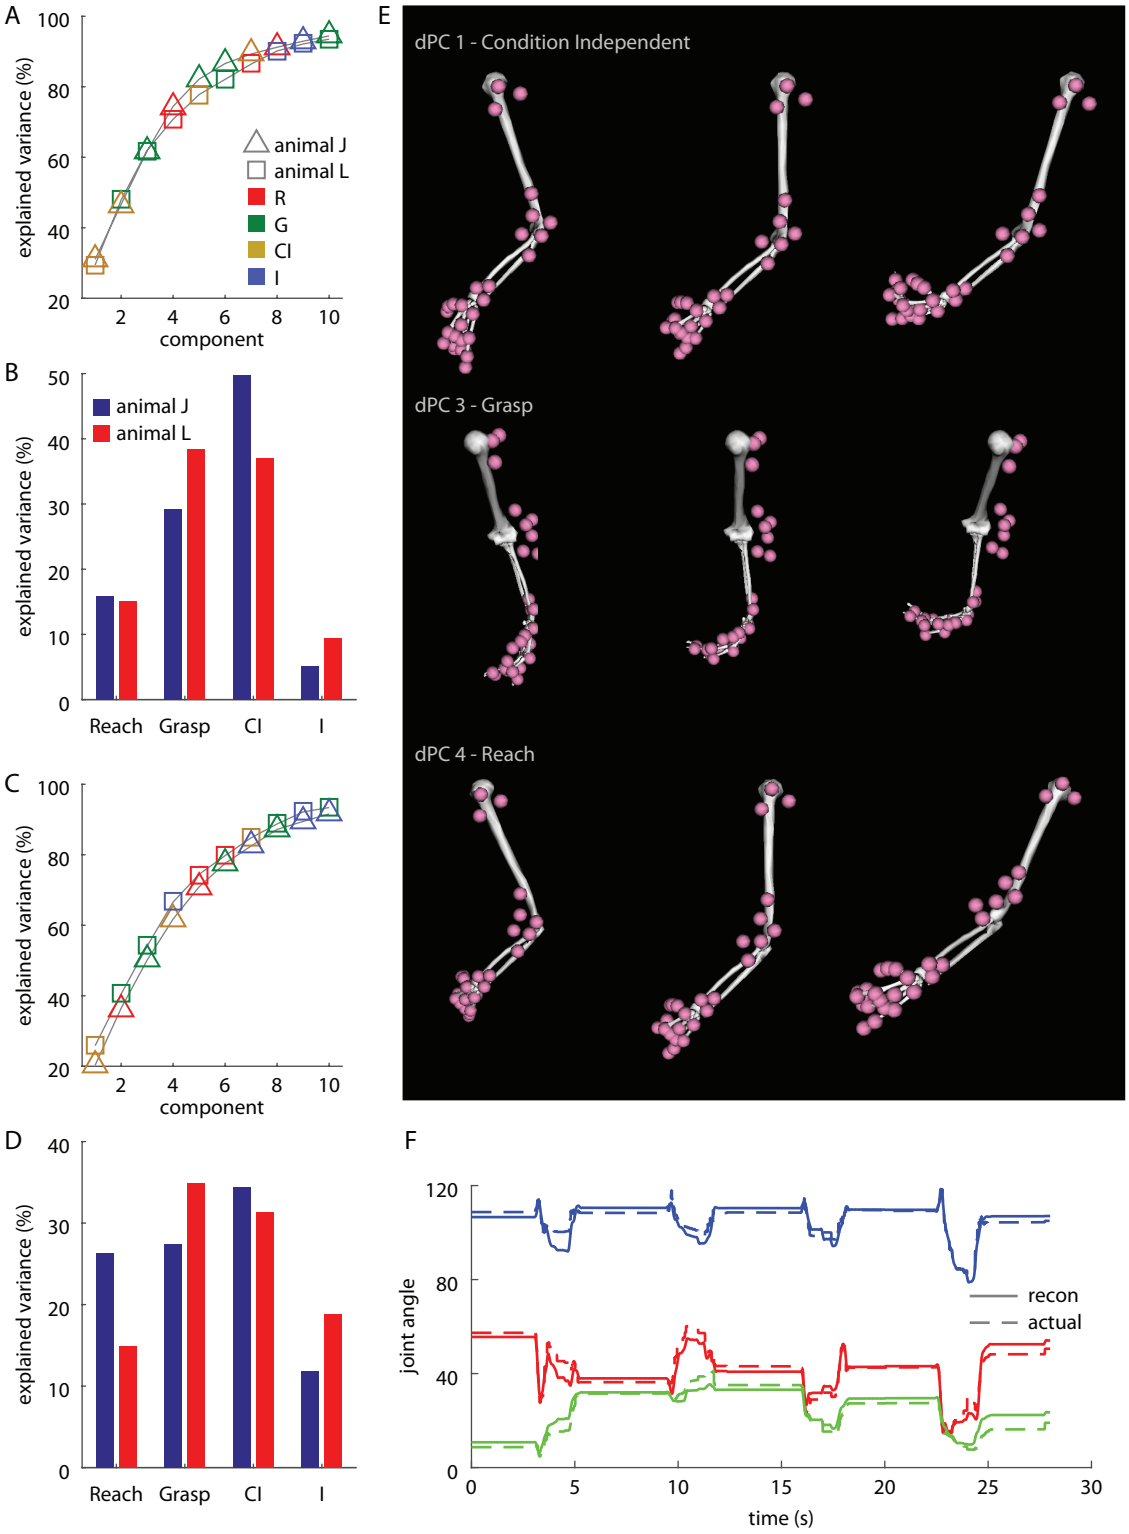

Figure 1.4 Demixed principal components analysis. A. Scree plot of demixed PCA. The cumulative percentage of explained variance for a given number of dPCs is shown for joint angles. Each individual component varies only along one task parameter. It can be either reach related (R, red points), grasp related (G, green points), condition independent (CI, gold points), or an interaction between reach and grasp condition (I,

blue points). Marker shape (either square or triangle) indicates animal. Note that the scree plots are highly similar across animals. **B.** Percentage of variance explained by the different task conditions in joint angle data. **C-D.** Same conventions as **A-B.** but applied to joint angular velocities. **E.** Visualization of kinematic synergies revealed by the demixed PCA. The posture of the upper limb is shown at various projections along the largest condition independent synergy (top row), the largest grasp synergy (middle row), and the largest reach synergy (bottom row). Data are from animal J. **F.** Reconstructed (solid lines) and actual (dashed lines) kinematic profiles of elbow flexion (blue), wrist flexion (red), and 2mcp flexion (green) during several seconds of behavior. Kinematics were reconstructed from the largest 8 dPCs. Data shown are from animal L.

In both animals, we found 8 components were sufficient to explain 90% of the variance in joint angles (Fig. 1.4A), while 10 and 9 components were needed to explain 90% of the variance in joint angular velocities in animals J and L, respectively (Fig. 1.4C). In both animals, we observed that a substantial proportion of the kinematic variance was common across all experimental conditions, and that grasping components explained more variance than reaching components (Fig. 1.4B,D). Examples of the kinematic synergies in animal L are shown in Fig. 1.4E. Finally, the original kinematics were recovered from the dPCA analysis (Fig. 1.4F). The temporal evolution of these kinematic synergies as a function of time relative to the go cue is shown in Fig. 1.5 for both animals.

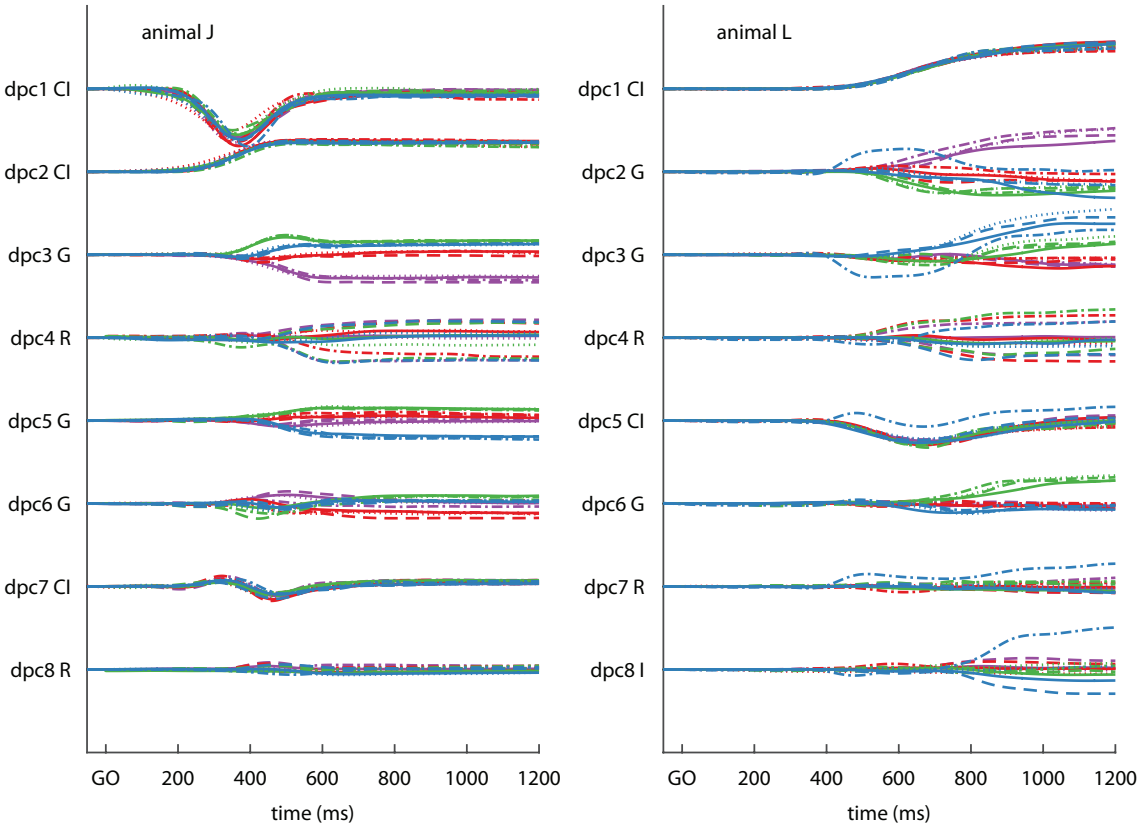

Figure 1.5 Trial averaged dPCA trajectories relative to go cue. The temporal profile of each dPC is shown for the various object locations and objects. Conventions are the same as Fig. 1.3. To the left of each dPC is an indication of whether it is a condition-independent (CI), grasp (G), reach (R), or interaction (I) component. The components are ordered from top to bottom in terms of variance explained; that is, dPC1 explained more variance than all others.

Encoding models

We used generalized linear models (GLM) to develop encoding models of premotor cortical activity.

Encoding of extrinsic features in premotor cortex

We compared how explicit moment-by-moment kinematic features are encoded in premotor cortex as compared to a more abstract, set-related representation of kinematics. In the first model, we used a detailed kinematic model containing the activity of 21 joints in the arm and hand at a variety of different lags (see Methods for

details). In contrast, the set-related model assumed a single firing rate for each experimental condition. We calculated the area under the receiver operating characteristic curve (AUROC) of each model and found that AUROC values across cells were significantly greater than 0.5, the AUROC that would be expected if there were no relationship between kinematics and spiking. (Fig. 1.6A; Wilcoxon signed rank test, all $p < 0.0006$). In both animals, we observed that kinematics better predicted activity in PMd as compared to PMv (Mann-Whitney *U*-test, animal J: $Z = 2.85$, $p < 0.004$; animal L: $Z = 4.12$, $p < 0.00005$). Additionally, for both PMd and PMv, we found that encoding models based on kinematics better predicted spiking activity than the set-related models (Fig. 1.6B-C; animal J PMd: $Z = 5.74$, $p < 0.0001$; PMv: $Z = 4.98$, $p < 0.0001$; animal L PMd: $Z = 3.45$, $p < 0.0001$; PMv: $Z = 4.11$, $p < 0.0001$).

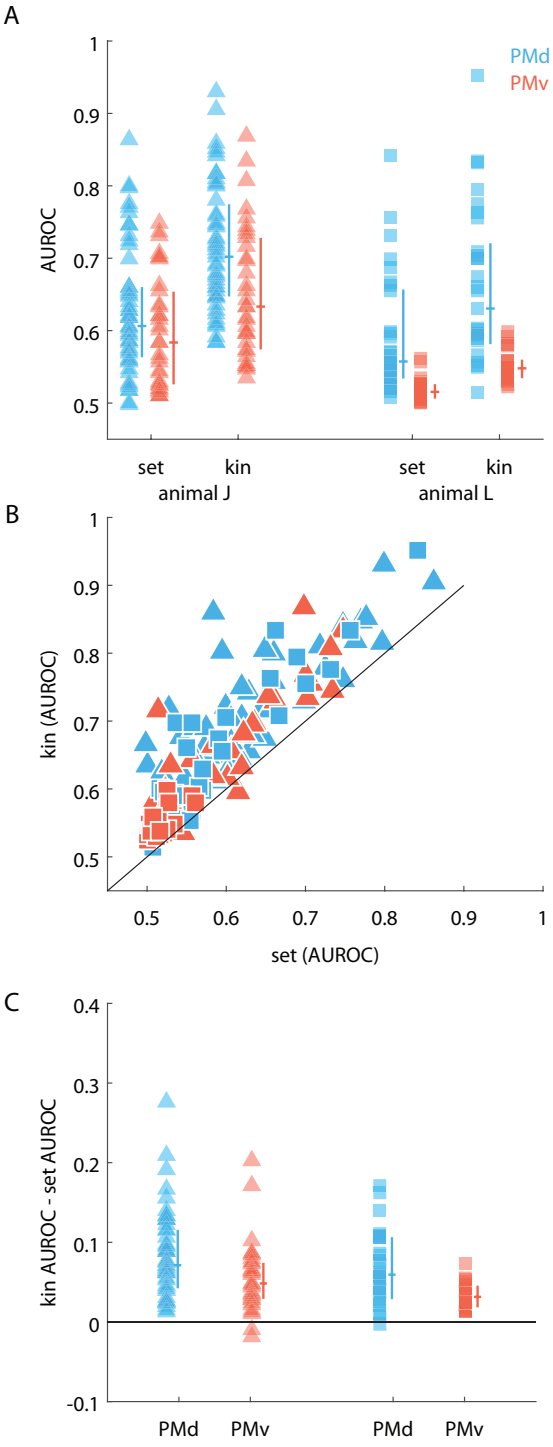

Figure 1.6 Encoding models that include kinematics predict spiking activity better than set-related models. A. We fit encoding models using either set-related activity or kinematics to each neuron in our sample and measured the area under the receiver operating characteristic curve (AUROC) across 10 folds of cross validated data. Here, each point corresponds to the median AUROC across folds as a function of animal (triangles or squares) and model type. Cells in PMd are blue, while cells in PMv are red. The vertical line next to each point-set indicates the interquartile range while

the median is indicated with the horizontal line. **B.** Plotting kinematic AUROC versus set AUROC for each of the cells in **A**. It can be seen that almost all cells fall above the unity line (black trace). **C.** Distribution of differences between kinematic and set-related models. The lines to the immediate right of each point-set follow the same conventions as **A.** Kinematic models predicted spiking activity better than set-related activity models in PMd and PMv in both animals (see main text for statistics).

We tested the extent to which anatomical segments of the limb were represented in both areas. Encoding models were fit using only arm kinematics, (i.e. shoulder and elbow kinematics), or only hand kinematics (including the wrist). We found evidence that both arm and hand kinematics were encoded in each area as the median AUROC across neurons was significantly greater than 0.5 (Fig. 1.7A; Wilcoxon signed rank test, all $p < 0.001$). To ensure that these effects were not because of correlations between arm and hand kinematics, we performed a control analysis in which we subtracted off either arm or hand AUROC from the AUROC of the full model containing both arm and hand kinematics (Fig. 1.7B). Suppose, for example, that we subtracted the AUROC of the hand kinematics model from the AUROC of the full model and that this difference was positive. Then, arm kinematics predicted spiking activity above and beyond the activity that was predicted by hand kinematics; this cell then would be classified as encoding arm kinematics. We found that a majority of cells encoded arm and hand kinematics (86% and 58%, respectively, pooling across animals and areas); however, the proportion of cells that encoded both arm and hand kinematics simultaneously was not significantly different than would be expected if they were independent (binomial test, $B_{0.497,\ 121} = 62$, $p > 0.05$).

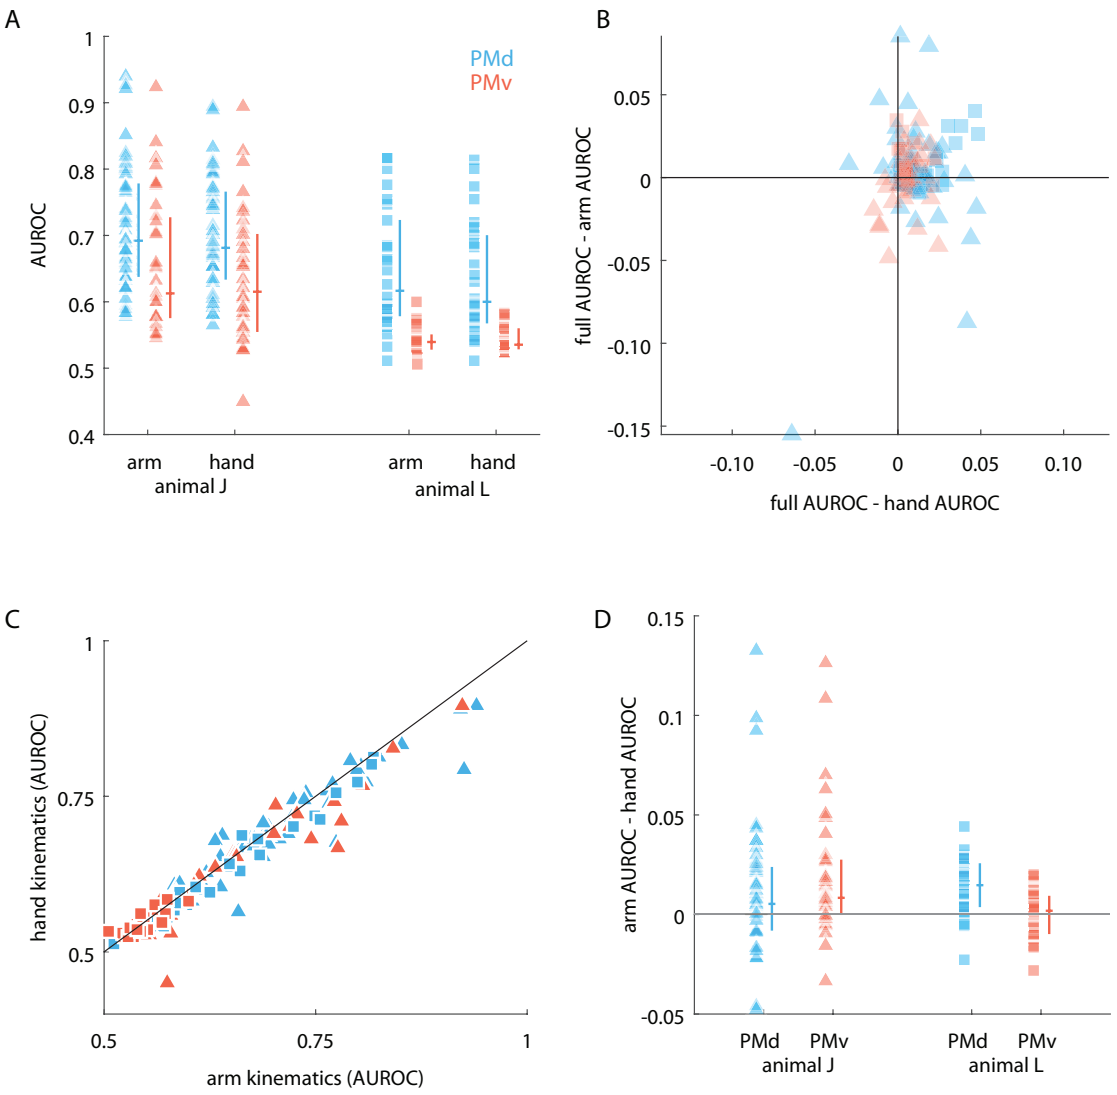

Figure 1.7 Comparing encoding performance of models based on an anatomical division of the upper limb. A. We fit GLMs using either arm or hand (including wrist) kinematics. For each model type and animal, the distribution of median AUROCs across folds is shown as raw data (left column; each point corresponds to a cell) and interquartile range + median (vertical and horizontal lines, respectively). **B.** Median AUROCs were compared between the full model that included all kinematic terms, and reduced models that contained only arm, or only hand kinematics. **C.** We observed that arm and hand AUROCs were highly correlated, consistent with the strong temporal correlations that were observed in the kinematics. **D**. Distribution of differences in arm and hand AUROCs. We found that arm kinematics were preferentially encoded by PMd in both animals and also in PMv in animal J because a significant proportion of cells preferred arm kinematics in those two areas (see text for statistics).

We subsequently tested whether the arm and hand were preferentially represented in PMd and PMv, respectively by comparing the AUROC of the arm-only

and hand-only models. We found that arm kinematics were preferentially represented in the PMd of both animals (Fig. 1.7C-D, animal J: $Z = 2.51$, $p < 0.02$; L: $Z = 3.35$, $p < 0.001$). Arm kinematics were also preferentially represented in the PMv of animal J ($Z = 2.57$, $p < 0.01$), but not in animal L ($Z = -0.41$, $p > 0.05$).

In addition to examining anatomical representation, we also compared functional representations of prehensile movements in PMd and PMv. We fit encoding models using only reaching components, or only grasping components of the dPCA analysis. We found both reach and grasp representations in both PMd and PMv (Fig. 1.8A; Wilcoxon signed rank test, all $p < 0.0001$).

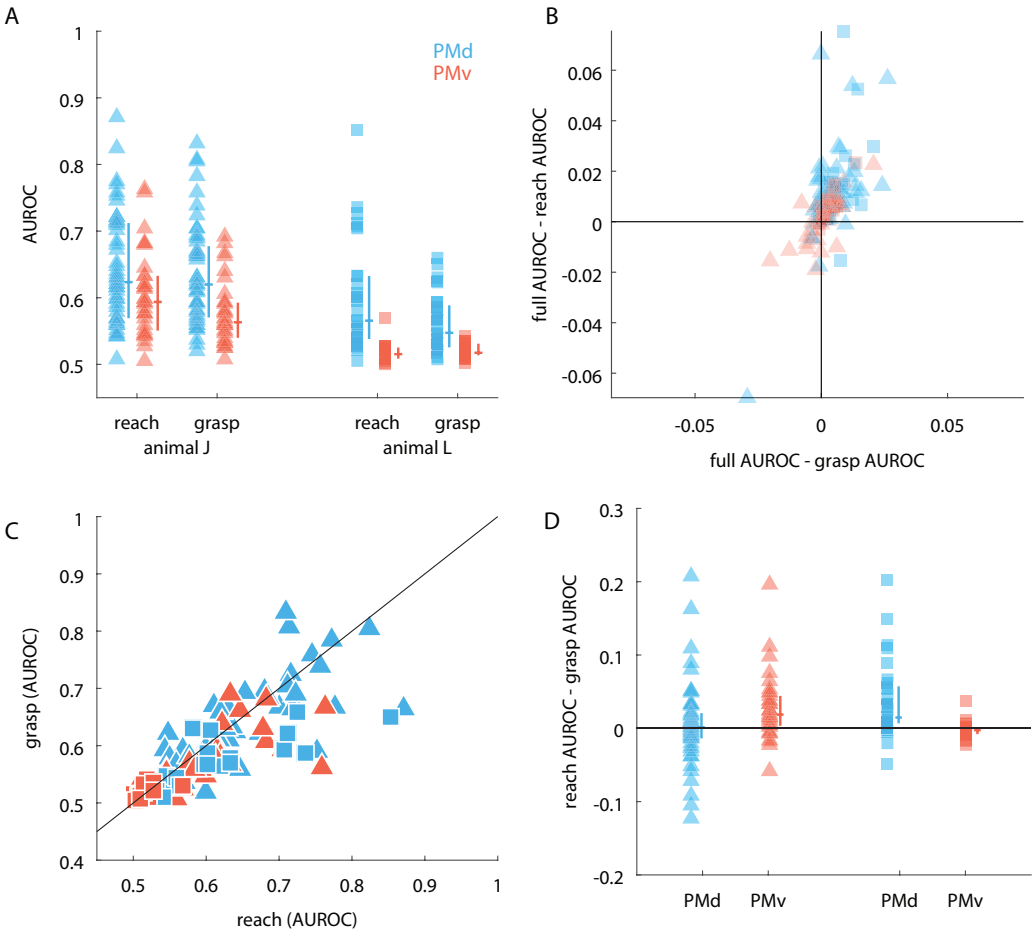

Figure 1.8 Comparing encoding performance of models based on a functional division of the upper limb. A. We fit GLMs using either reach or grasp kinematic

synergies. For each model type and animal, the distribution of median AUROCs across folds is shown as raw data (left column) and interquartile range + median (vertical and horizontal lines, respectively). **B.** Median AUROCs were compared between the full model that included both reach and grasp synergies, and reduced models that contained only reach or only grasp synergies. **C.** We observed that reach and grasp AUROCs were correlated, but less so than arm and hand model AUROCs. **D**. Distribution of differences in reach and grasp AUROCs. We found that reach kinematic synergies were preferentially encoded by PMv in animal J and PMd in animal L (see text for statistics).

Again, as a control analysis, we assessed the extent to which a full model containing both reach and grasp synergies outperformed a model containing only reach, or only grasp (Fig. 1.8B). Here, we observed that the proportion of cells that simultaneously encoded both reach and grasp was significantly more frequent than would be expected under chance if reach and grasp were independent (binomial test, $B_{0.628,\ 136} = 98$, $p < 0.009$). Although many cells encoded both reach and grasp, there was a preferential representation of reaching in the PMv of animal J and PMd of animal L because, as a population, the difference in AUC between the reach only and grasp only models was significantly different from 0 (Fig. 1.8C-D, Wilcoxon signed rank test, J: $Z = 3.58$, $p < 0.0005$; L: $Z = 3.75$, $p < 0.0005$).

Finally, we considered the relationship between the anatomical and functional kinematic features encoded by each cell. To this end, we compared the difference of arm and hand model AUROCs to the difference of reach and grasp model AUROCs. Under the classical model, we might assume that cells that preferentially encode the kinematics of the arm also prefer the kinematics of reaching, and correspondingly, cells that prefer the hand also prefer grasping (Fig. 1.9A).

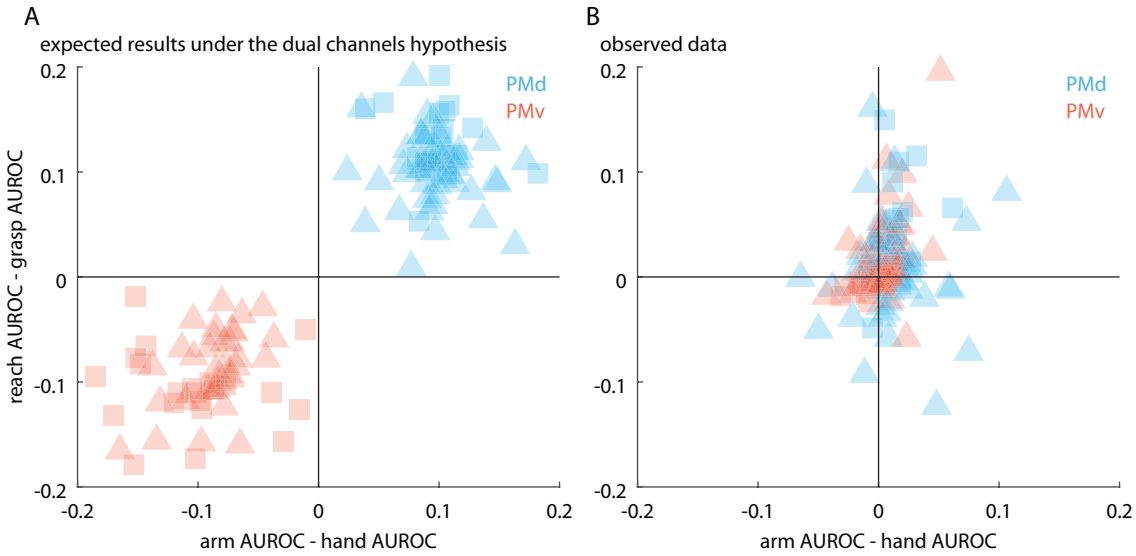

Figure 1.9 Comparing anatomical and functional representations in premotor cortex.
We compared the differences of arm and hand AUROCs to the differences of reach and grasp AUROCs to determine if there was a systematic relationship such that cells that preferentially encoded arm kinematics likewise encoded reaching synergies, and similarly, cells that preferentially encoded hand kinematics likewise encoded grasping synergies. **A.** Expected results based on the dual channels hypothesis. PMd would be concerned exclusively with arm kinematics and reaching, while PMv would be concerned with hand kinematics and grasping. There would be a segregation of both function and anatomy between PMd and PMv. **B.** Actual results based on observed data. A chi-squared test of independence revealed no significant relationship between functional and anatomical encoding preferences.

Such a simplistic relationship, however, does not appear to capture the patterns we observed in our data because we observed several cells that preferred both hand and reach or arm and grasp (Fig. 1.9B). Alternatively, another possibility is that anatomical encoding preferences are not related to functional encoding preferences. We tested this hypothesis quantitatively and found no evidence to suggest that function and anatomy are not independent (chi-squared test of independence, X^2 with 1 d.f. = 3.03, $p > 0.05$).

Encoding of intrinsic features in premotor cortex

In addition to extrinsic features, we also explored how intrinsic features are encoded in premotor cortex. We began by comparing the encoding performance of the full kinematic model to the full intrinsic, i.e. spike history model. The spike-history model contained information about the spike histories of all other simultaneously recorded neurons at a variety of temporal scales (Fig. 1.10; see methods for details).

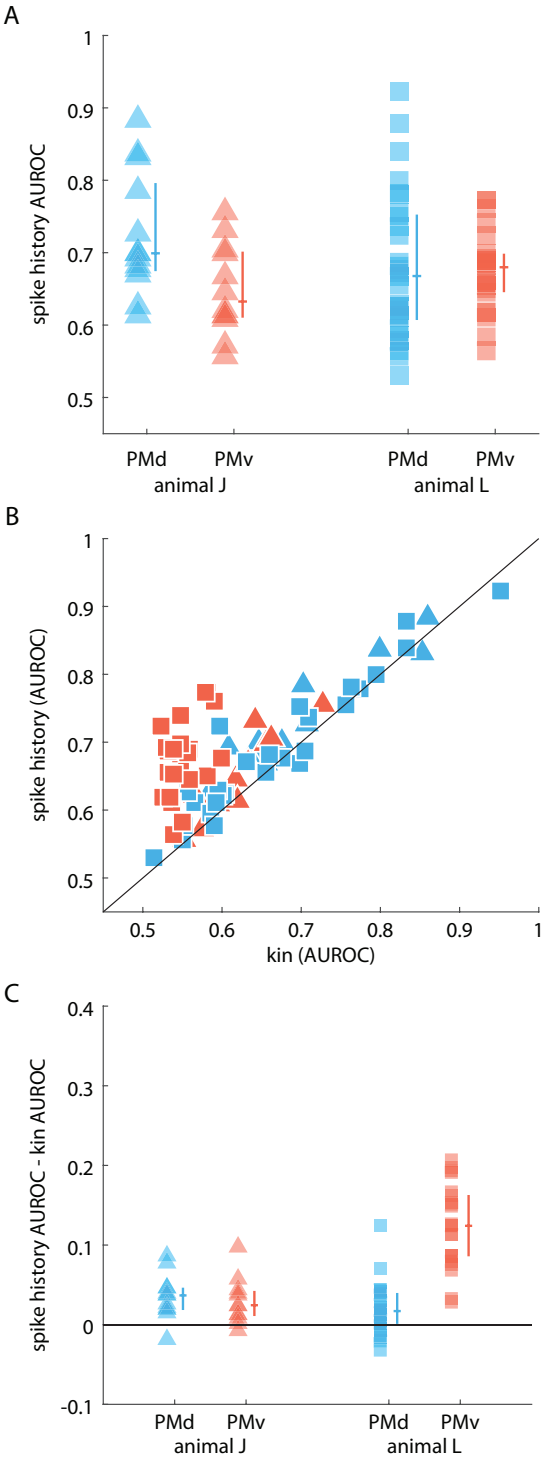

Figure 1.10 Comparing spike history with kinematics. A. We found that spike history from all cells at multiple temporal scales effectively predicted spiking activity (see main text for statistics). Each point corresponds to a cell from a given animal (shape) and cortical area (color). Lines next to each column of data indicate interquartile range and median. **B.** We compared the performance of the spike history model to the full kinematic model and observed that generally, spike history was a more effective

predictor of spiking activity than kinematics. **C.** The distribution of differences between the spike history and kinematic models shown as a function of animal (shape) and cortical area (color). Lines indicate interquartile range and median values.

We found that spike history was an effective predictor of spiking activity (Wilcoxon signed rank test, animal J PMd: $W = 91$, $p < 0.0005$; PMv: $W = 78$, $p < 0.0005$; animal L PMd: $Z = 4.7$, $p < 0.000005$ PMv: $Z = 4.1$, $p < 0.00005$). In animal J, we found spike history predicted PMd activity significantly better than PMv (Mann-Whitney *U*-test, $Z = 2.20$, $p < 0.03$), but there was no difference in animal L across areas ($Z = -0.33$, $p > 0.05$). Indeed, spike history better predicted spiking activity than the kinematics in all areas (Fig. 1.10B-C, all Bonferroni corrected $p < 0.014$).

We wanted to identify which temporal scale led to the best encoding performance. We fit models using each of the three different temporal scales (Fig. 1.11A) and found that the longest temporal scale (108 ms) led to the best encoding performance in both animals irrespective of cortical area (Fig. 1.11B).

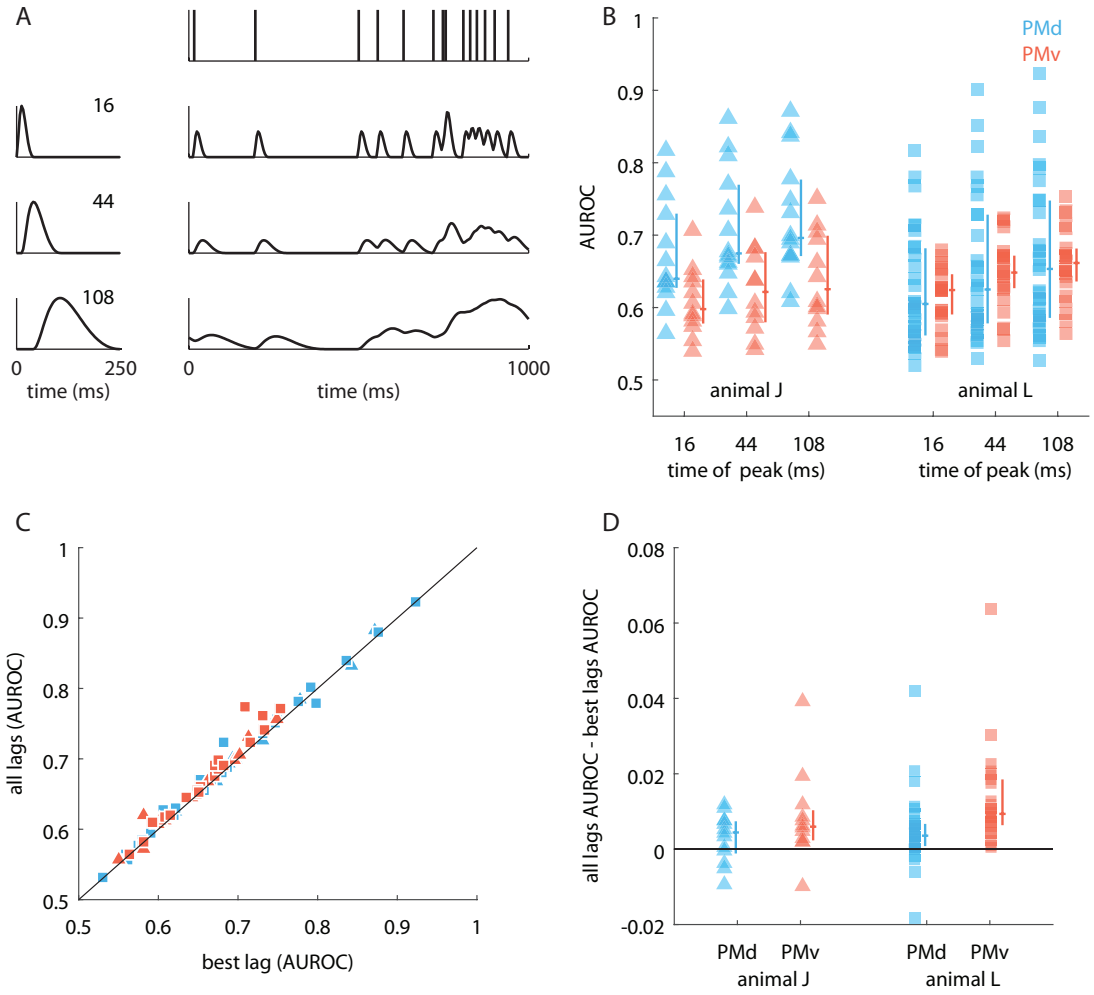

Figure 1.11 Comparison of spike history effects at multiple timescales. A. The three raised cosine basis functions are shown (left column). On the right is an example spike train (top) and the convolution of that spike train with each raised cosine basis function (below). The three different basis functions are designed to represent spike history effects over different timescales from short (16 ms) to long (108 ms). **B.** Median AUROC across cross validation folds as a function of basis function, cortical area, and animal. In both areas, the longest spike history term had the best encoding performance. **C.** Comparing the best single basis AUROC of each cell to the AUROC of its corresponding full spike history model. **D.** Distribution of the differences in AUROC between the full spike history model and a model based on the best single basis function (lag). To the right of each data column are the median and interquartile range of the data.

We subsequently wondered if the information at each temporal scale was totally redundant, or, if using spike history at multiple temporal scales led to superior predictive ability. Accordingly, we compared the AUROC at the best temporal scale for each cell

and compared it to the AUROC of the full spike history model including all temporal scales (Fig. 1.11C-D). In the PMd of animal J, we observed that there was no significant difference between the AUROC of the full spike history model and the model at the best lag (Wilcoxon rank sum test, PMd: $W = 69$, $p < 0.11$). In contrast, spike history at multiple spatial scales better predicted spiking activity in the PMv of both animals (J: $W = 69$, $p < 0.02$, L: $Z = 4.11$, $p < 0.00005$) and the PMd of animal L ($Z = 3.30$, $p < 0.001$).

Additionally, we examined if cells within a given cortical area were better predicted by the spike history of other cells in that area. In both animals, we estimated the correlation between every pair of neurons by computing the Pearson correlation coefficient of their spike history at the widest lag. In animal L, we observed two clusters in the correlation matrix; cells within an area exhibited correlated activity, however, there was almost no correlation between cells across areas (Fig. 1.12B).

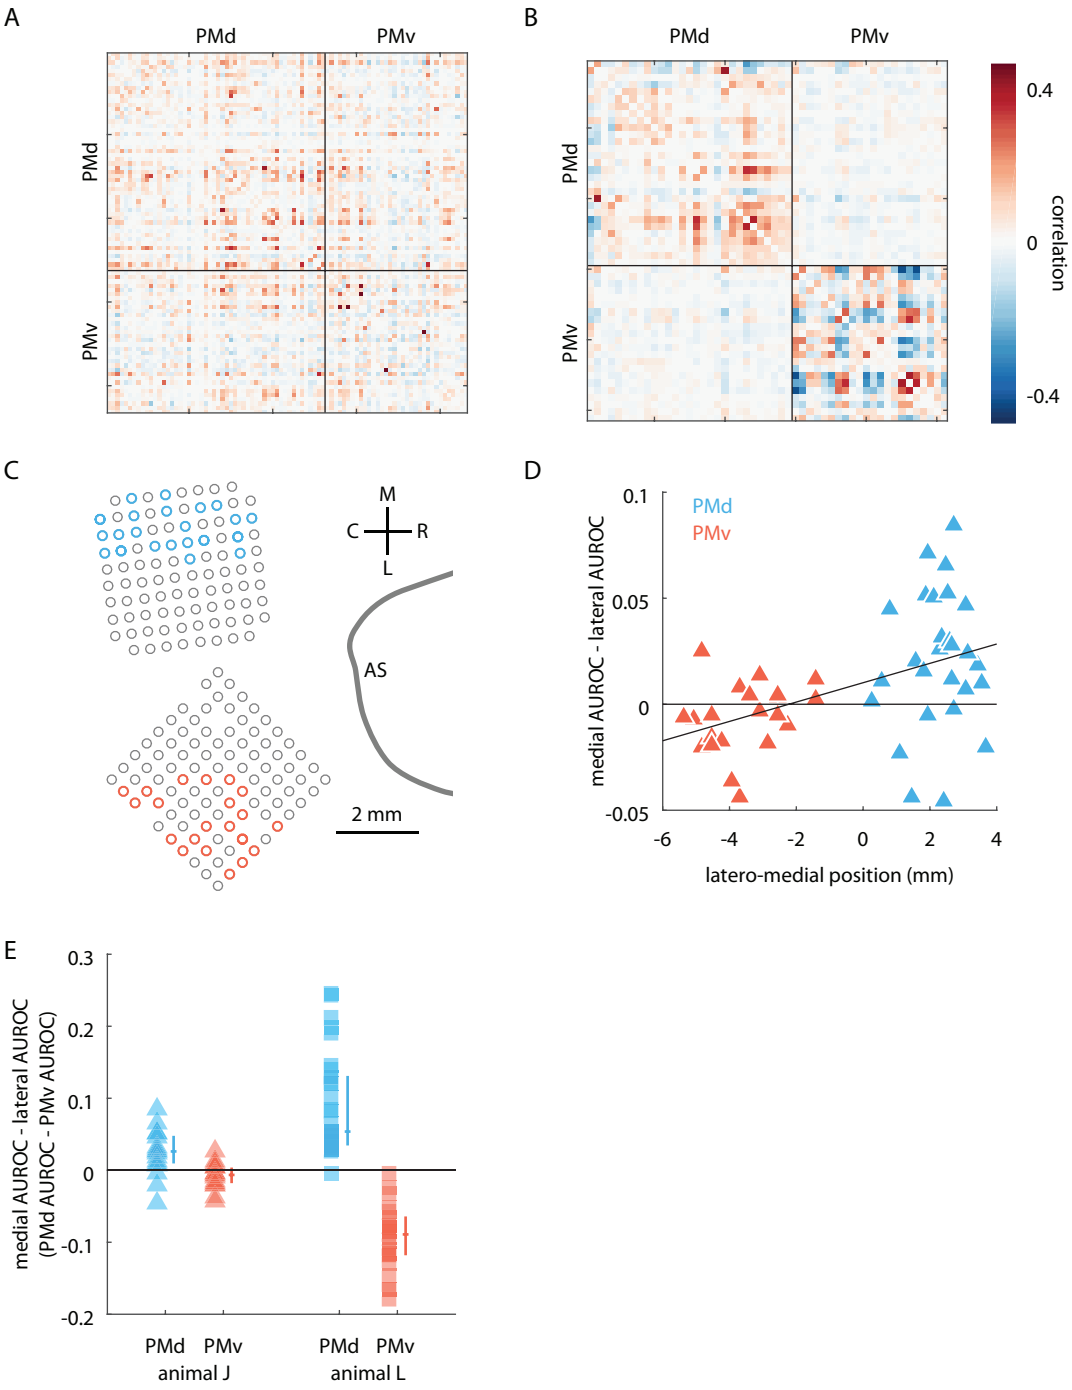

Figure 1.12 Predicting responses in premotor cortex using either PMd or PMv activity. A-B. Pairwise correlation matrices for all recorded cells in animal J and L, respectively. Color indicates the correlation coefficient between two cells such that highly correlated cells appear red, while highly anti-correlated cells appear blue (scale given on right). Note that the diagonal has been colored white. Qualitatively, there does not appear to be a clear division of PMd and PMv based on the correlation matrix in animal J. In contrast, in animal L, there appear to be two clear segregations based on anatomical location. **C**. For animal J, we fit encoding models using cells located on either the most medial (blue cells, corresponding to cells that are most likely in PMd), or

most lateral (red cells, corresponding to cells that are most likely in PMv) as predictors. The approximate location of the arcuate sulcus (AS) is indicated. **D.** We examined the differential performance of the medial and lateral encoding models for each cell as a function of its latero-medial position relative to the genu of the AS (position 0). Positions lateral to the genu are negative, while positions medial to it are positive. **E.** Distribution of differences between medial and lateral encoding models for animal J, or fully PMd and PMv encoding models for animal L as a function of the anatomical location of the cell (color). Lines indicate median and interquartile range. Generally, spiking activity of cells in PMd are better predicted by the activity of other cells in PMd while the spiking activity of cells in PMv are better predicted by other cells in PMv.

We did not observe the same structure in animal J (Fig. 1.12A); however, its two arrays were much closer together than those of animal L. We fit encoding models using the activity of cells in either PMd or PMv. The self-spike history of each neuron was excluded. In animal L, we found that spiking activity of cells in PMd was better predicted by the spike history of other PMd cells (Fig. 1.12E, Wilcoxon signed rank test, $Z = 4.68$, $p < 0.000003$) and similarly spiking activity of PMv cells were better predicted by the spike history of other PMv cells (Fig. 1.12E, $Z = -4.11$, $p < 0.00005$). Because animal J's were closer together, we only used cells that were located in the medial aspect of the PMd array, or the lateral aspect of the PMv array as input features to the encoding model (Fix. 12C). In animal J, we observed that spiking activity of cells on the PMd array were better described by the spike history of cells in the medial aspect of PMd (Wilcoxon signed rank test, $Z = 2.96$, $p < 0.004$). In PMv, we observed a qualitatively similar result to animal L, however, it was only weakly significant ($Z = 1.72$, $p < 0.09$). We further examined if a cell's latero-medial position on the cortical surface affected the relative AUROC of the medial (i.e. PMd) and lateral (i.e. PMv) encoding models. We performed an analysis of covariance (ANCOVA) predicting the differential performance of the medial and lateral models based each cell's latero-medial position,

and the Utah array from which it was recorded. We observed no significant interaction between latero-medial position and array conditions ($F_{1,44} = 0.001$, $p > 0.05$), and moreover, no significant difference between arrays ($F_{1,45} = 0.08$, $p > 0.05$). However, we did observe a significant main effect of latero-medial position (Fig. 1.12D $F_{1,46} = 13.24$, $p < 0.0007$). If the *x*-intercept from the previous model (shown in Fig. 1.12D) was used as the putative boundary between PMd and PMv; that is, cells to the right of the *x*-intercept putatively belong to PMd with cells to the left putatively belonging to PMv, then spiking activity of putative PMd neurons was better predicted by the history of other PMd cells (Fig. 1.12E, $Z = 3.06$, $p < 0.002$) and spiking activity of PMv neurons was better predicted by the history of other PMv cells (Fig. 1.12E, $Z = 2.01$, $p < 0.045$).

Discussion

Demixed principal components analysis

Here, we applied a novel dimensionality reduction technique, demixed principal components analysis (Kobak et al. 2014), to identify kinematic synergies associated with either reaching or grasping during an unconstrained prehensile movement task. In both animals, we observed that a few functionally defined kinematic synergies accounted for most of the variance in the data. Moreover, we found that reaching and grasping synergies each explained a large proportion of the kinematic variance suggesting that presenting multiple objects at multiple locations successfully elicited a diversity of upper limb movements. In contrast, comparatively little variance was due to interactions between reaching and grasping conditions, supporting the decomposition of

kinematics into reaching and grasping sub-movements (Jeannerod 1984; Haggard and Wing 1995).

Although this method provides a promising avenue for gaining insight into prehensile movements, there are some limitations to the approach. Chiefly, the kinematic decomposition was based on trial-averaged data, and failed to account for trial-to-trial variation within an experimental condition. An addition limitation was that it did not support the use of continuous variables as conditions. That is, object locations were not treated as positions in Euclidean space, but rather discrete, nominal covariates. Finally, on individual trials, the scores of the dPCs often exhibited strong moment-by-moment correlations suggesting that even though reach and grasp were decoupled across conditions, within a condition there may remain some strong correlation.

Implications for cortical control of reach to grasp

Here, we used GLMs to develop encoding models of spiking activity in premotor cortex based intrinsic and extrinsic features. We found evidence that challenges the dogmatic view of premotor cortical organization being organized around discrete reaching and grasping pathways. In particular, we demonstrated that there was a complete anatomical representation of the upper limb in PMd. This finding is consistent with previous studies that used intracortical-microstimulation to study motoric representations in PMd, and further work that showed firing rates of PMd neurons are modulated by changes in grasping condition (Raos et al. 2003; Raos et al. 2004; Fattori et al. 2010). Additional work has demonstrated that whole arm kinematics can be

decoded from PMd ensemble activity (Bansal et al. 2012). Our work complements that study by showing that spiking activity encodes kinematics, and, particularly, single cells encode complex, whole-arm kinematics. Similarly, in PMv, we also found a complete representation of the upper limb, although, in absolute terms, it was weaker than in PMd. This result is consistent with previous reports that ICMS of PMv occasionally evoked arm movements (Godschalk et al. 1981; Godschalk et al. 1985; Hocherman and Wise 1991).

Kinematically, we demonstrated that arm kinematics were affected by grasping condition and hand kinematics were affected by reaching condition, so, while we demonstrated that whole arm representations are present in both PMd, and PMv, that was insufficient to argue against our null hypothesis that reach and grasp were processed in independent pathways. We used a novel dimensionality reduction technique to dissociate reaching from grasping and found that both reaching and grasping kinematic synergies predicted the activity of cells in PMd and PMv. Moreover, across the entire population of recorded cells, we found no relationship between anatomical and functional representations. That is, the anatomical preferences of a cell (i.e. whether it was biased towards arm or hand kinematics) were not related to its functional preferences (i.e. whether it preferred reach or grasp synergies).

Differences between PMd and PMv

In this work, we compared the encoding properties of neurons simultaneously recorded in the dorsal and ventral premotor cortex of two rhesus macaques. By using

this simultaneous recording paradigm, our results were not confounded by potential behavioral differences across sessions.

In both animals, we observed that movement kinematics were more strongly represented in PMd as compared to PMv. One potential explanation for this finding is that PMd is seemingly more directly related to motor output. Pyramidal cells in layer V of PMd comprise approximately 10% of fibers from the frontal lobe in the cortico-spinal tract (R. P. Dum and Strick 1991), and intracortical microstimulation of PMd is known to elicit arm movements (Weinrich and Wise 1982; Weinrich, Wise, and Mauritz 1984). PMv too sends projections to spinal cord, but it only projects to propriospinal neurons located in the upper cervical spinal segments (He, Dum, and Strick 1993). Recent work using ICMS has shown that PMv acts on spinal cord primarily through facilitation of MI (Cerri et al. 2003; Shimazu et al. 2004).

In both animals, we observed that while single cells in PMd encode whole arm kinematics, they exhibit a preferential encoding of arm kinematics. This is consistent with the classical studies of PMd implicating it in the control of reaching movements (Weinrich and Wise 1982; Weinrich, Wise, and Mauritz 1984; Kurata and Tanji 1986; Riehle and Requin 1989; Pesaran, Nelson, and Andersen 2006). Yet, in PMv, we observed no consistent bias for either arm or hand kinematics.

Additionally, in PMv, we found that spike history at multiple time scales was a better predictor of spiking activity than a single timescale. This suggests that PMv may be more influenced by several internal processes operating at different timescales, and may also explain why kinematics were not strongly represented in PMv. One possibility

may be to use a more sophisticated state-space analysis (Shenoy, Sahani, and Churchland 2013; Aggarwal et al. 2013) to elucidate the relationship between population activity in PMv and movement. An alternative possibility is that these types of naturalistic movements are not the most effective drivers of activity in PMv, but rather a different class of movement altogether may be represented (Bonini et al. 2014; Lehmann and Scherberger 2013).

Inter-animal differences

We observed some differences in the kinematics of the two animals. Animal J had faster reaction times than animal L, and, generally, executed faster movements that were qualitatively more stereotyped. Although the movements were executed at different speeds and with different reaction times, we observed many similarities in the dPCA analysis across the two animals suggesting that the kinematic synergies we identified are robust across a naturalistic range of self-selected movement speeds. In future work, it could be tested if the same set of dPCA synergies predicts movement kinematics at a variety of speeds by having an animal engage in a reach-to-grasp task with an explicit speed cue.

An additional difference in the kinematics between the two animals was their resting posture. In animal L, resting grasp aperture was proportionately much larger than in animal J. Indeed, the grasp aperture was maximal at the rest position for some trials in animal L. In contrast, animal J maintained a more closed aperture during the rest epoch. This finding may explain why the temporal profiles of the largest dPC in each animal are different. The temporal profile of the largest CI dPC in animal J was a

bell-shaped curve, reflecting that the hand opened and closed during the reach to grasp movement. In contrast, the temporal profile of the largest CI dPC in animal L was a sigmoidal curve corresponding to the fact that his hand started opened and gradually conformed to the shape of the object to be grasped. Finally, the animals adopted qualitatively different strategies to successfully grasp the ring object. Grasp aperture of the ring object in animal J followed a biphasic profile that was absent in animal L.

At the neural level, there were also several differences between the two animals. We observed that arm kinematics were preferentially represented in the PMv of animal J, but not L. One potential explanation of this finding is based on array placement. In both animals, we saw a preferential representation of the arm in PMd. The PMv array of animal J is much closer to its PMd array than in animal L. Accordingly, arm representation in the PMv of J may be stronger because it is closer to PMd. We might also argue that array placement also explains our findings concerning spike history effects. In animal L, spiking activity in PMd was better predicted by the history of other cells in PMd, and vice versa for PMv. In animal J, this result, while still significant, was not as strong as in animal L potentially because of the closer spatial proximity of the two arrays.

Chapter 2 - Beta oscillations in premotor cortex during reach-to-grasp

Abstract

Movement modulates oscillations in the beta frequency range (15-30 Hz) of motor cortical local field potentials (LFPs). And yet, the relationship between the LFP and movement has remained poorly understood. One promising approach to elucidate this relationship involves consideration of spatio-temporal patterns of beta activity. In this chapter, we examine how the amplitude and phase of beta oscillations vary across the cortical surface during a reach to grasp task. We find that the amplitude of beta oscillations attenuates around the time of movement onset. This attenuation phenomenon propagates across the cortical surface as a linear wave in PMd, but exhibits no strong spatio-temporal structure in PMv. Based on this finding, it is argued that PMd is directly involved with movement initiation, but not PMv. The directionality of phase gradients in both premotor cortical areas is examined. We document the existence of planar wave activity in PMv, and show that wave activity in both areas does not provide information about experimental condition.

Introduction

Oscillatory activity in the beta frequency range of local field potentials (15-30 Hz) is a hallmark of motor cortical activity (Jasper and Penfield 1949; Sanes and Donoghue 1993; Murthy and Fetz 1992). First described in the 1930s by Jasper and Penfield, beta oscillations were immediately correlated with movement (Jasper and Penfield 1949) and have remained an active area of research in part because a deeper understanding of the relationship between beta activity and movement has remained elusive. Work in

both humans and nonhuman primates has shown that beta amplitude does not correlate with properties of movement including direction (Rickert et al. 2005), force(Tombini et al. 2009), and speed (Stancák Jr. and Pfurtscheller 1996). Accordingly, movement cannot be decoded from the amplitude of beta band activity (Spinks et al. 2008). While current evidence suggests that the amplitude of beta oscillations does not directly encode any parameters of movement, beta activity is still clearly related to movement at a more global level. Beta oscillations are high in amplitude during epochs of isometric force production (Baker et al. 1999) as well as postural stability (Sanes and Donoghue 1993; Gilbertson et al. 2005). These findings have led to theories positing that beta plays an active role in movement. Specifically, high amplitude beta oscillations are maintaining the motor status quo (Gilbertson et al. 2005; Engel and Fries 2010), or, are a "brake" that actively inhibits movement (Jenkinson and Brown 2011).

At the same time, a number of recent studies have documented that beta oscillations are spatially heterogeneous within the primary motor cortex (MI) (Rubino, Robbins, and Hatsopoulos 2006; Takahashi et al. 2015). During motor planning, the phase of beta oscillations is not constant across the cortical surface, but rather propagates as planar waves (Rubino, Robbins, and Hatsopoulos 2006). These waves have been related to the emergence of a transient network of effective connections between neurons with narrow spike waveforms (Quinn et al. 2011; Takahashi et al. 2015). Around movement onset, these planar waves are diminished and beta activity transitions from a high to low amplitude state in a process referred to as beta attenuation. The spatio-temporal distribution of beta attenuation is also not constant

across the cortical surface, but rather, propagates linearly (Appendix 1) along a similar direction as the planar waves. And yet, the motor cortex is a highly distributed network of cortical areas involved in the control of movement. Presently, it remains unknown if the spatio-temporal patterns of activity that are observed in MI are also present in other areas of motor cortex, and particularly, the premotor cortex.

The premotor cortex can be further subdivided into at least two regions: dorsal and ventral (PMd and PMv, respectively) with each contributing differentially to prehensile movements (Barbas and Pandya 1987). During both motor planning and execution, beta activity in PMd bears a strong resemblance to that of MI with waves occurring during planning and then amplitude attenuating around movement onset (Rubino, Robbins, and Hatsopoulos 2006). Much less has been documented about beta activity in PMv during reach-to-grasp, although beta amplitude has also been shown to attenuate at movement onset in this area (Spinks et al. 2008).

Here, we characterized spatio-temporal variations in the amplitude and phase of beta activity during reaching-to-grasp in dorsal and ventral premotor cortex using simultaneous recordings from these areas. Because we simultaneously recorded from both areas, we emphasized differential involvement of these areas to behavior in an effort to further elucidate the functional differences between PMd and PMv. We found that beta attenuation exhibited spatio-temporal patterning in both areas; however, there was a qualitative difference in the nature of attenuation between PMd and PMv. Additionally, we documented the existence of planar wave propagation in PMv. Unlike PMd and MI, waves in PMv did not have a consistent orientation; however, for both PMd

and PMv wave propagation did not vary significantly with movement features related to reaching or grasping.

Methods

Neurophysiology

All surgical and experimental procedures were approved by the University of Chicago Animal Care and Use Committee and conformed to the principles outlined in the Guide for the Care and Use of Laboratory Animals (NIH publication no 86-23 revised 1985). Two rhesus macaques (*macaca mulatta*) were implanted with 96-electrode Utah arrays in the dorsal and ventral premotor cortex contralateral to their working arm. Electrodes were 1.5 mm in length except for the PMd array in animal J in which the electrodes were 1.0 mm. Neural signals from the electrodes were amplified with a gain of 5000; band-pass filtered between 0.3 Hz and 7.5 kHz, and recorded digitally (14 bit resolution) using a Cerebus acquisition system (Blackrock Microsystems, Salt Lake City, UT). Local field potential recordings were sampled at 2 kHz.

Behavioral Task

Two male rhesus macaques were trained to perform to perform a reach to grasp task with their left hand. A robot (RV-1A-S11 6-axis robot; rixan.com) presented the animals with objects to grasp at four different spatial locations in their peripersonal space. The set of objects consisted of geometric shapes that were designed to evoke a variety of different hand conformations when grasped. Object and location pairings were randomly varied on a trial-by-trial basis to minimize the amount of motor planning

the animal could do. Additionally, animals' vision was occluded between trials with a pneumatically controlled screen.

The precise combination of objects that were presented varied between the two animals. Animal J grasped 3 shapes (Cyl-O, Disc-O, Ring H and V) for a total of 4 unique grasping conditions while animal L grasped 4 shapes (Cyl-O, Disc-O, key, and Ring H; Ring V did not have many successful trials so key grip was substituted) for a total of 4 unique grasping conditions. Animals completed at least 15 repetitions of each object-location pairing. These objects were presented to the same 4 spatial locations in the workspace for both animals.

Motion Tracking

A 10-camera motion capture system (Vicon Motion Tracking System, Oxford, UK) detected and recorded the 3-D position of retroreflective markers affixed to animal. A total of 30 markers were placed on the animal's dorsal hand and dorsolateral arm enabling the tracking of 27 degrees of freedom in the arm and hand. The motion capture system was also used to monitor the position of the object and vision block screen.

Three markers were placed on the animals' wrist, and were used to compute wrist speed. The time series of 3D marker positions for each of these three markers were low pass filtered (bi-directional 4^{th} order Butterworth filter 15 Hz cutoff) and then numerically differentiated to compute wrist speed. After numerical differentiation, wrist velocities were again low pass filtered (using the same filter design as before). Wrist speed was defined as the average tangential velocity of the wrist markers. Movement

onset was defined as the first moment that wrist speed exceeded 10% of the max wrist speed.

Intracortical Microstimulation

The Utah array was used to deliver intracortical microstimulation (ICMS) to premotor cortex. Cathodal-anodal pulse trains were generated by a Blackrock current stimulator (pulse width: 0.2 ms, train-frequency: 333 Hz, train-duration 75 ms). Current intensities varied between 20 and 100 μA.

We began stimulating each electrode at 80 μA. If, after several repetitions, we did not evoke a movement, we stimulated at 100 μA. If movements were still not evoked, we considered that electrode inactive. However, when evoked movements were observed, we attempted to isolate the functional specificity of that electrode. If complex, or whole arm movements were evoked at a given current intensity, we would drop the current in 10 μA increments until the movement was well isolated, or no longer visible. For each electrode, the evoked movements at the lowest current that still reliably evoked movement were recorded as the behavior for that electrode. The sequence of electrodes that we stimulated through was chosen pseudorandomly.

Movements evoked by microstimulation were characterized by two experimenters observing the animal. A movement was recorded only if it was visible to both experimenters and readily repeatable. We stimulated only when the animal appeared to be relaxed.

We characterized the behavior of each electrode as either proximal, distal, or hybrid. Hybrid effects, in turn were divided into mostly proximal and mostly distal

subgroups. Shoulder and elbow movements were classified as proximal, while forearm muscle twitches and wrist movements were classified as distal. If a hybrid evoked movement contained both proximal and distal components it was classified according to the qualitative strength of each component, e.g. a movement evoking deltoid twitches, biceps twitches, and wrist flexion would be classified as mostly proximal. In our sample, 25 out of the 85 electrodes with stimulation effects were hybrid, i.e. they contained both proximal and distal components.

Analysis

Power Spectra

Power spectra were estimated for each channel individually using the Chronux (http://chronux.org/) spectral analysis toolbox. Data from the entire recording session were used to estimate these spectra. We discarded channels that had no modulation in the beta frequency range based on visual inspection. These discarded channels represented a small portion of the total number of channels recorded (Fig. 2.1 for exemplar datasets from J and L).

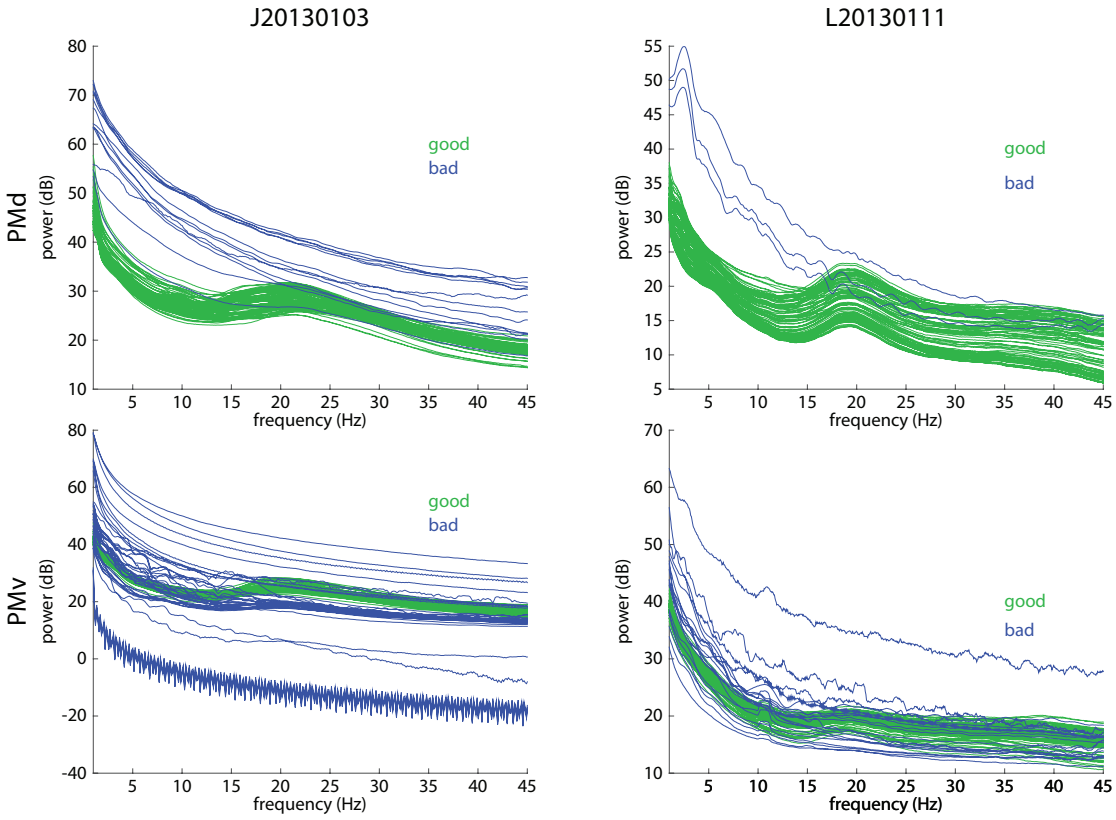

Figure 2.1 Power spectra. Power spectra from each recorded electrode were estimated using a multi-taper spectral density estimator (time-bandwidth product: 5, number of tapers: 9). Channels were visually classified as either good or bad based on power spectra. In animal J, 83 and 56 of the 96 channels were good in PMd and PMv respectively (93 and 77 for animal L). Only good channels were used in subsequent analyses.

LFP preprocessing

Several preprocessing steps were applied to local field potentials before the beta attenuation analysis. Raw local field recordings were first filtered into the beta frequency range using a bidirectional band-pass filter (4^{th} order Butterworth filter whose pass-band was equal to the peak frequency in beta +- 3 Hz for a total of a 6 Hz pass-band), and then their Hilbert amplitude was computed. The Hilbert amplitudes were used for the beta attenuation analysis while Hilbert phases were used to compute phase gradient waves. Beta amplitudes were trial registered to movement onset.

Sigmoidal curve fitting

The amplitude of beta oscillations is known to attenuate during active movement. We modeled this attenuation phenomenon with a logistic curve of the form:

$$y = \frac{h - l}{1 + \ln(-s * (x - m))} + l$$

Where the two parameters of the model, the slope, *s*, and the midpoint, *m*, were found using a non-linear least squares curve fitting algorithm (matlab function lsqcurvefit). The other two model parameters, *h* and *l*, corresponding to the upper and lower asymptotes of the sigmoid were set to be equal to the average beta amplitude in a window of time relative to movement onset. For the upper asymptote, *h*, it was the average beta amplitude 1500 to 1000 ms before movement onset, while, the lower asymptote, *l*, was the average beta amplitude 0 to 500 ms after movement onset. Model goodness of fit was assessed using the coefficient of determination, R^2, corresponding to the proportion of variance explained by the logistic model.

Bootstrap analysis

We compared properties of beta attenuation across experimental conditions using a bootstrap analysis. For every experimental condition, on a given iteration of the bootstrap (out of a total of 50), we randomly selected half of the trials from that condition, and fit a logistic curve to those data. Trials were sampled without replacement on each iteration of the bootstrap, but with replacement over iterations. That is, a given trial could be randomly selected on the first iteration of the bootstrap, and then again later on in another iteration, although the same trial couldn't be selected twice on a given iteration of the bootstrap. This bootstrap procedure was used to

estimate the variance of parameters in logistic model. Beta attenuation midpoints (BAMs) as reported in the results were typically derived by averaging the BAMs over every iteration of the bootstrap.

Comparing BAOs across experimental conditions

To compare BAOs across experimental conditions we used an *F*-test of nested linear models. Using Wilkinson's notation (Wilkinson and Rogers 1973), we fit a full model of the form:

$$BAM \sim (loc * obj) + (r+c) - loc:r - loc:c - obj:r - obj:c$$

Where loc is a categorical variable encoding the object's position, obj is a categorical variable encoding which object was presented, and r and c are integers indicating the row and column coordinates of the electrode on the Utah array. We fit a reduced model of the form:

$$BAM \sim (loc * obj) + r + c$$

Where abbreviations are the same as the full model. The difference between the full and reduced model is that in the full model, each experimental condition has its own unique BAO (determined by two three-way interaction terms: loc:obj:r and loc:obj:c), while the reduced model identifies one BAO, but let's the model intercept (i.e. the average BAM) vary across experimental conditions.

Phase gradient waves

Beta phases on each electrode were estimated using the Hilbert transform as described above. Then, local phase gradients around each electrode were computed numerically by taking differences along the rows and columns of the Utah array. The

extent to which these local phase gradients agreed with one another determined the phase gradient directionality (PGD) of the wave. Mathematically, PGD may be defined as:

$$PGD(t) = \frac{\|\bar{\phi}\|}{\overline{\|\phi\|}}$$

Where $\|\bar{\phi}\|$is the norm of the average direction, $\bar{\phi}$, and $\overline{\|\phi\|}$ is the average of the norm of the individual phase gradients, ϕ (see (Rubino, Robbins, and Hatsopoulos 2006) for additional information about the methodology used in wave calculations).

Results

Beta power in PMd and PMv

We estimated power spectra from multiple datasets in animals J and L (Fig. 2.2).

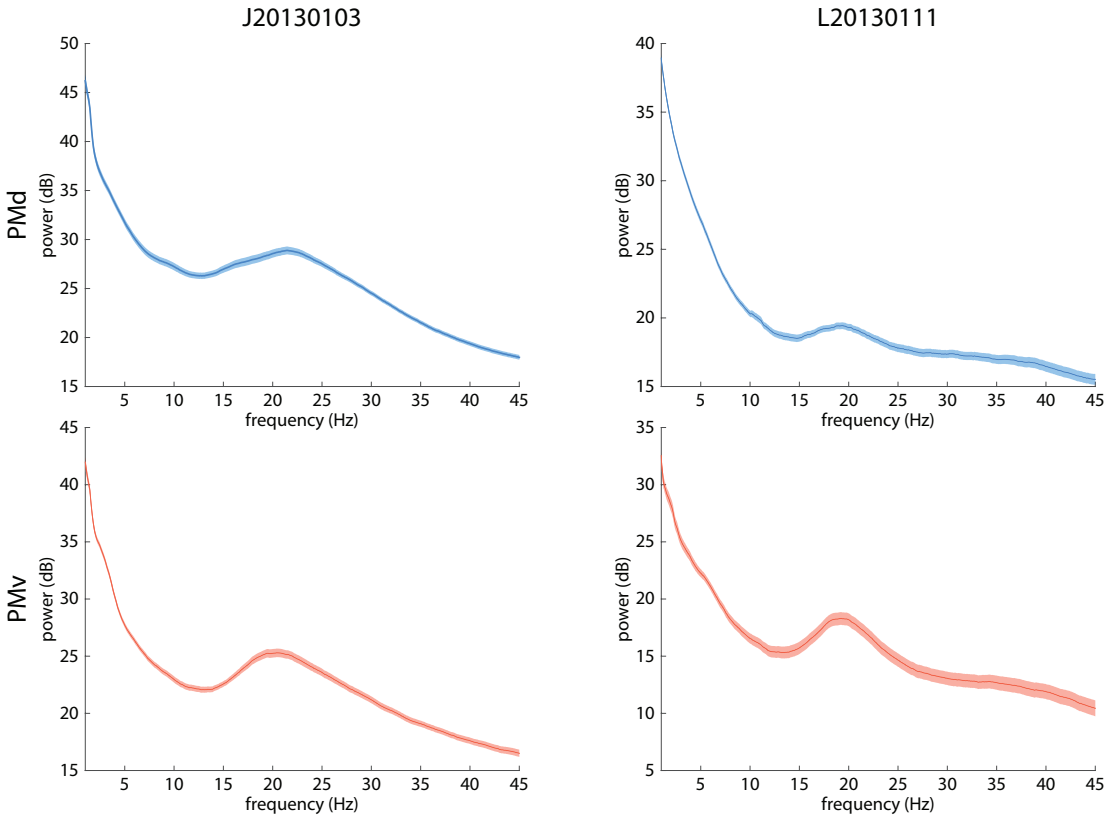

Figure 2.2 Average power spectra across electrodes (+- 2 standard errors of the mean (S.E.M.)). There is a substantial peak in the beta frequency range (15-30 Hz) in each dataset.

The average frequency of the beta peak was 19.94+-1.00 and 19.96+-0.41 Hz (mean + 2 S.E.M) in PMd and PMv respectively across datasets in animal J. In animal L, the average frequency of the beta peak was 19.10 +- 0.24 and 17.40 + 3.91 Hz in PMd and PMv, respectively. The average peak frequency was not significantly different across PMd and PMv (paired *t* test, $t_{18} = 0.5242$, $p < 0.61$). Pooling across areas, we found that the average frequency of peak beta was significantly different between animals (unpaired *t* test assuming equal variance, $t_{18} = 2.49$, $p < 0.023$).

Timing of beta attenuation

In primary motor cortex, low amplitude oscillations in the beta frequency range are believed to be a cortical correlate of active movement (Sanes and Donoghue 1993;

Engel and Fries 2010; Kilavik et al. 2013). In both PMd and PMv, we also observed that the amplitude of beta activity was low during active movement, as compared to stationary epochs. In particular, we found strong evidence that beta amplitude 500 ms after the start of movement was significantly lower than beta amplitude 1000 to 500 ms before the start of movement in both animals (Fig. 2.3).

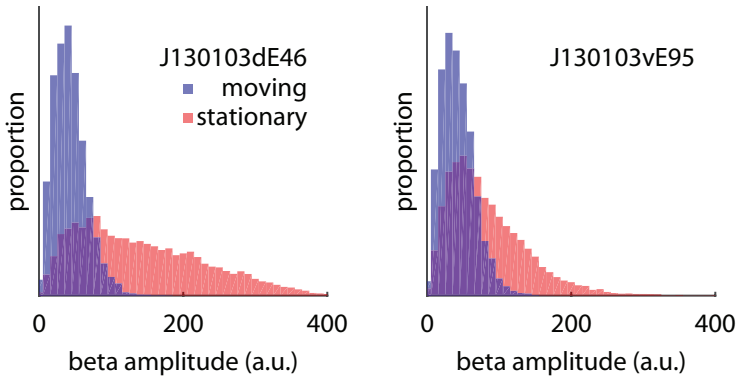

Figure 2.3 Beta amplitude as a function of behavioral state on a single electrode in PMd (left) and PMv (right). During epochs of active movement (blue), beta amplitude is lower, whereas during stationary epochs (red), beta power is higher.

We sought to characterize the transition between high and low amplitude states at a finer temporal resolution. In particular, we first asked if the transition from high to low amplitude, henceforth referred to as beta attenuation, was a reflection of sensory processing performed by the premotor cortex in response to changing visual information (i.e. the vision block being released), or, if it was a reflection of motor processing (i.e. related to the upcoming movement). To answer this question, we examined the timing of attenuation relative to movement onset in trials with both short and long reaction times. If beta attenuation were related to visual processing, then its timing should be earlier on long reaction time trials, whereas, if it were related to motoric processing, its timing should be constant relative to movement onset.

We modeled this attenuation phenomenon with a logistic curve (see Methods for

computational details, Fig. 2.4).

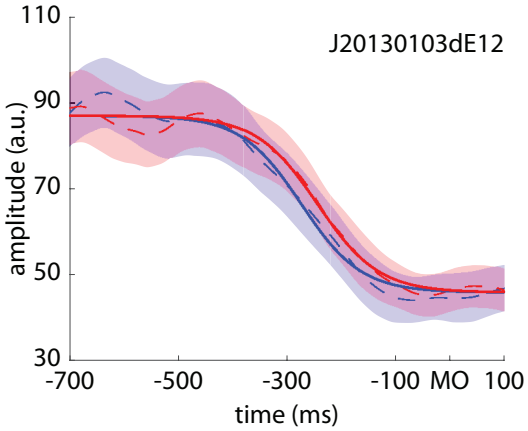

Figure 2.4 Temporal evolution of beta amplitude as a function of time relative to movement onset (MO) on a single electrode in the PMd of animal J. Here, we show trial averaged beta amplitude (dashed lines, shaded area indicates +- 2 S.E.M.) for early (blue) and late (red) reaction time trials. Each trace was fit with a logistic curve (solid line). The midpoints of these curves were not significantly different.

Specifically, we fit logistic curves to beta amplitude profiles for each recorded electrode. We then compared the midpoints of these curves, dubbed the beta attenuation midpoint (BAM), from each good electrode across the two reaction time conditions. To assess the statistical significance of this difference, we compared it to a null distribution of differences based on random partitions of trials into two groups without consideration of reaction time. In animal J, we found only two electrodes in PMd that had significantly different BAMs while no electrodes in PMv showed the same effect. In animal L, no electrodes in either PMd or PMv showed significantly different BAMs in the two reaction time conditions suggesting that reaction time is likely coupled with movement onset.

Effect of experimental condition on beta attenuation

In the previous analysis, we examined if variation in beta attenuation timing across trials was due to differences in reaction time and found no strong evidence of an

effect. We further considered what effect, if any, the specific object to be grasped and its location in the workspace had on beta attenuation timing. To this end, we fit logistic curves to data from each experimental condition (i.e. location-object pairs) and compared BAMs across conditions. To facilitate this comparison, we employed a bootstrap analysis in which random samples of trials were repeatedly drawn from each experimental condition. We measured the goodness of fit for each model using the coefficient of determination, R^2. We found that goodness of fit values were higher for PMd than PMv when averaged across experimental conditions (Fig. 2.5, two way ANOVA, $F_{1,306} = 262.42$, $p < 1e\text{-}8$), suggesting that the temporal dynamics of beta activity may be different between PMd and PMv. Similarly, goodness of fit values were higher in animal J than L (two way ANOVA, $F_{1,306} = 292.07$, $p < 1e\text{-}8$).

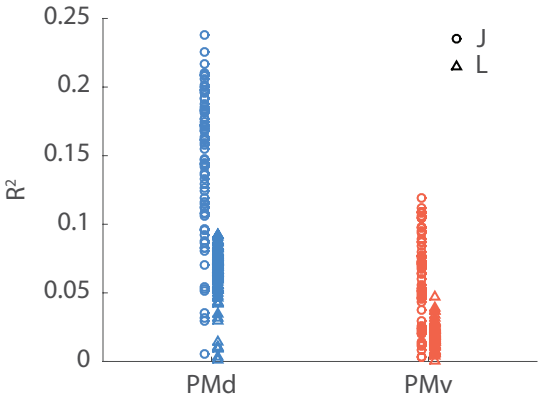

Figure 2.5 Logistic curve goodness of fit as a function of anatomical area and animal. Each point on the plot corresponds to an electrode from either PMd (blue) or PMv (red) in animal J (circles) or L (triangles). R^2 values were significantly higher in PMd than PMv, and significantly higher in animal J than L (see main text for statistics).

In the previous section, we found that the temporal dynamics of beta activity were differentially described by a logistic curve suggesting their activity profiles were different. We subsequently asked if the overall timing, as measured by the BAM, was significantly different. We computed the median BAM across electrodes for each experimental

condition. In animal J, we found that BAMs were significantly later in PMd than PMv (paired *t*-test, $t_{15} = 6.72$, $p < 0.000007$, Fig. 2.6) while in animal L, BAMs were later in PMv than PMd (paired *t*-test, $t_{15} = -4.62$, $p < 0.0004$). We attempted to relate this inconsistency in timing to differences in behavior, but were unable to find a strong relationship.

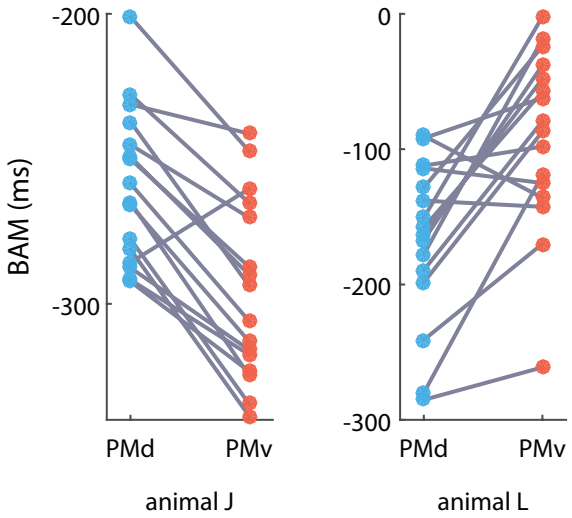

Figure 2.6 Timing of beta attenuation in PMd (blue) and PMv (red) in two animals. Note that overall attenuation timing is different for both animals, and that the relative timing of attenuation in PMd and PMv is different across animals.

Within each area, we examined the relationship between BAMs recorded at different recording sites. We observed that the BAMs associated with each experimental condition were highly correlated at nearby recording sites, and this

correlation decayed gradually with distance (Fig. 2.7).

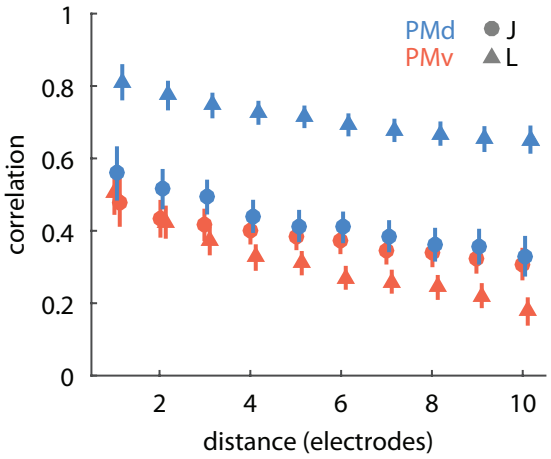

Figure 2.7 Electrode-Electrode BAM correlations as a function of distance between electrodes. Distance between electrodes was measured using the taxicab metric. Error bars indicate +- 2 S.E.M.

We modeled this decay with a linear relationship. BAM correlations between electrodes were higher in PMd than PMv ($F_{1,35}$ = 808.52, p < 1e-8), although both decayed at the same rate, i.e. decayed with the same linear slope ($F_{1,33}$ = 2.37, p < 0.13). Additionally, there was a highly significant difference overall in the correlation level of BAMs between animals J and L ($F_{1,35}$ = 189.26, p < 1e-8).

Spatiotemporal patterning of beta attenuation

Previous work has shown that the spatial propagation of beta attenuation across the cortical surface in MI is well approximated by a linear gradient (Appendix 1). In premotor cortex, we observed a similar spatial progression of beta attenuation when pooling data across all experimental conditions (Fig. 2.8).

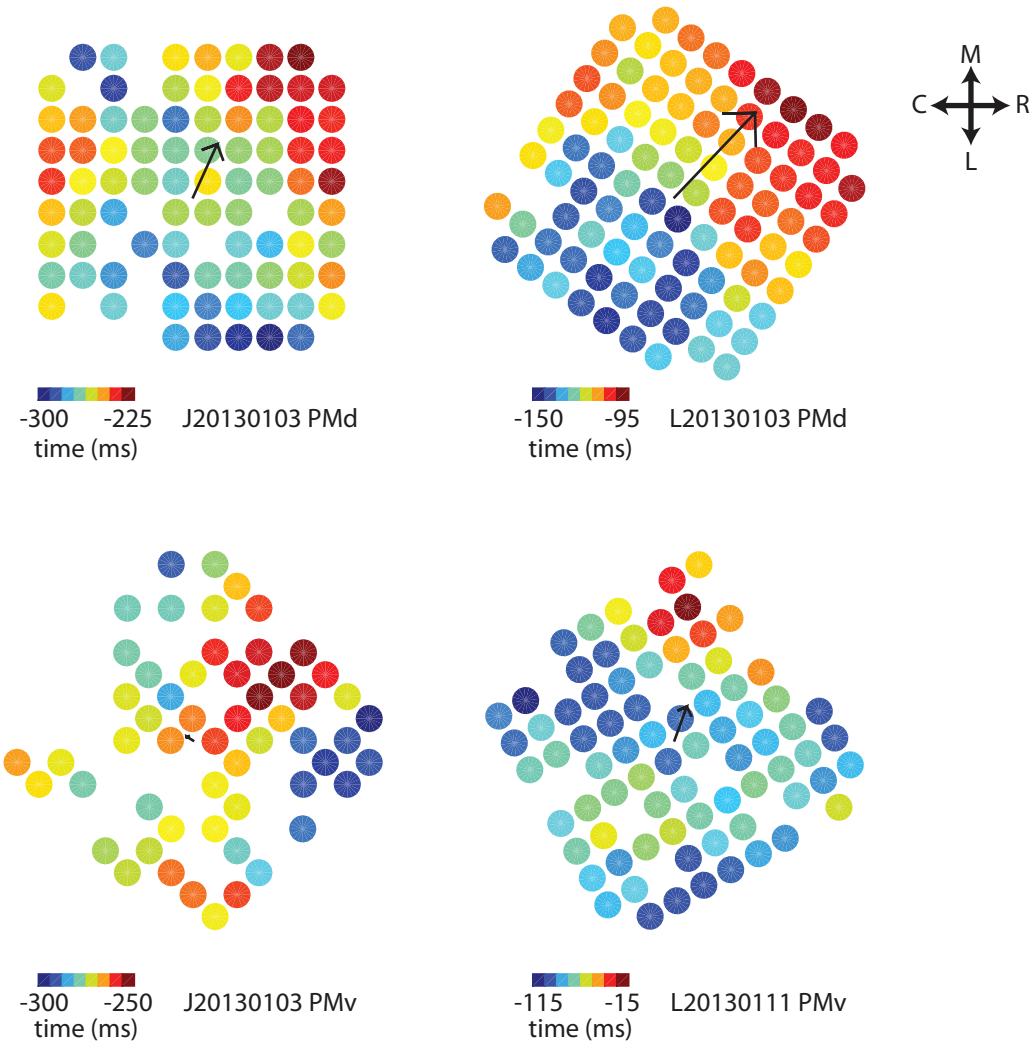

Figure 2.8 Spatio-temporal propagation of beta attenuation across the cortical surface. In both PMd and PMv, we found that beta attenuation did not happen simultaneously across the cortical surface, but rather was well approximated by a linear gradient. Here, we plotted BAMs (relative to movement onset) based on their spatial location on the electrode array and used linear regression to characterize their spatial propagation. We found that beta attenuation was organized along a lateral-to-medial axis in PMd of both animals such that early BAMs (blue) were more lateral to later BAMs (red). In PMv, there was no such consistent relationship across animals, although both animals did have significant linear gradients (see main text for statistics).

We visualized BAMs as a function of their spatial position on the electrode array. Linear regression was used to characterize the beta attenuation orientation (BAO) in each area. In PMd, we observed a significant linear relationship between BAMs and their spatial location (animal J: $F_{2,80} = 20.6$, $p < 1e\text{-}7$, $R^2 = 0.34$; L: $F_{2,90} = 96.2$, $p < 1e\text{-}8$,

$R^2 = 0.68$). In both animals, J and L, the BAO was situated along a lateral-to-medial vector.

In contrast to PMd, spatial patterning in PMv was more heterogeneous. In animal J, there was slight, but significant linear patterning ($F_{2,53} = 4.26$, $p < 0.019$, $R^2 = 0.14$). In animal L, linear spatial patterning was also present, albeit weaker than in PMd ($F_{2,74} = 9.87$, $p < 0.0002$, $R^2 = 0.21$). The BAOs in animals J and L were not consistently aligned and the strength of the linear relationship was much smaller than in PMd suggesting a qualitative difference in the behavior of beta attenuation in these two areas.

One current thought is that spatiotemporal variations in primary motor cortical activity are largely driven by the underlying representation of the limb in cortex (Murphy, Wong, and Kwan 1985; Riehle et al. 2013). Here, we compared BAMs to a somatotopic map of premotor cortical outputs based on an intra-cortical microstimulation (ICMS) experiment. In PMd of animal J, we found that the BAO was nearly orthogonal to the somatotopic map (Fig. 2.9,107.7°).

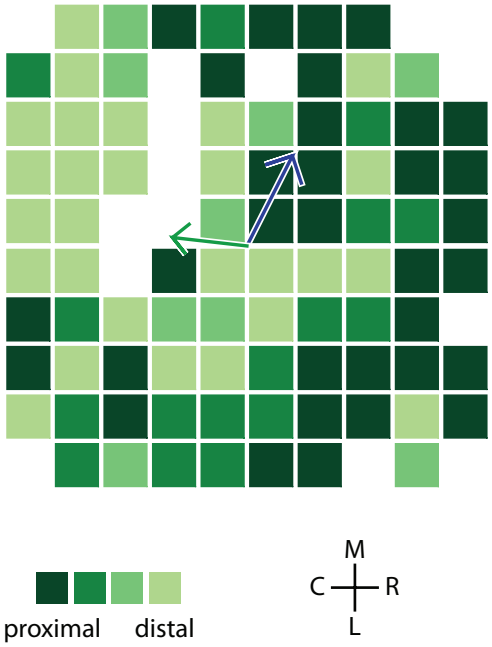

Figure 2.9 Relationship of body representation and beta attenuation. We used ICMS to characterize the somatotopic organization of PMd in animal J. We found that the proximal-to-distal gradient of limb representation (green arrow) was nearly orthogonal to the BAO (blue arrow)

We performed similar ICMS experiments in PMv, but were unable to reliably evoke movements. In PMd of animal L, we observed some evoked movements on a few electrodes, however, there were too few to reliably identify any somatotopy (17 electrodes).

Relating BAMs to BATs

Previous work on the spatiotemporal dynamics of beta attenuation identified a beta attenuation time (BAT) based on the moment when beta amplitude crossed an arbitrary threshold. This procedure did not involve any model fitting, unlike the approach used here to characterize beta attenuation. We defined pseudo-BATs (pBATs) based on the time when the fitted logistic curve first crossed the attenuation threshold used to calculate a BAT. We compared raw BAMs, pBATs and BATs as well as BAOs estimated from each of these three statistics in dataset J20130103 (Fig. 2.10).

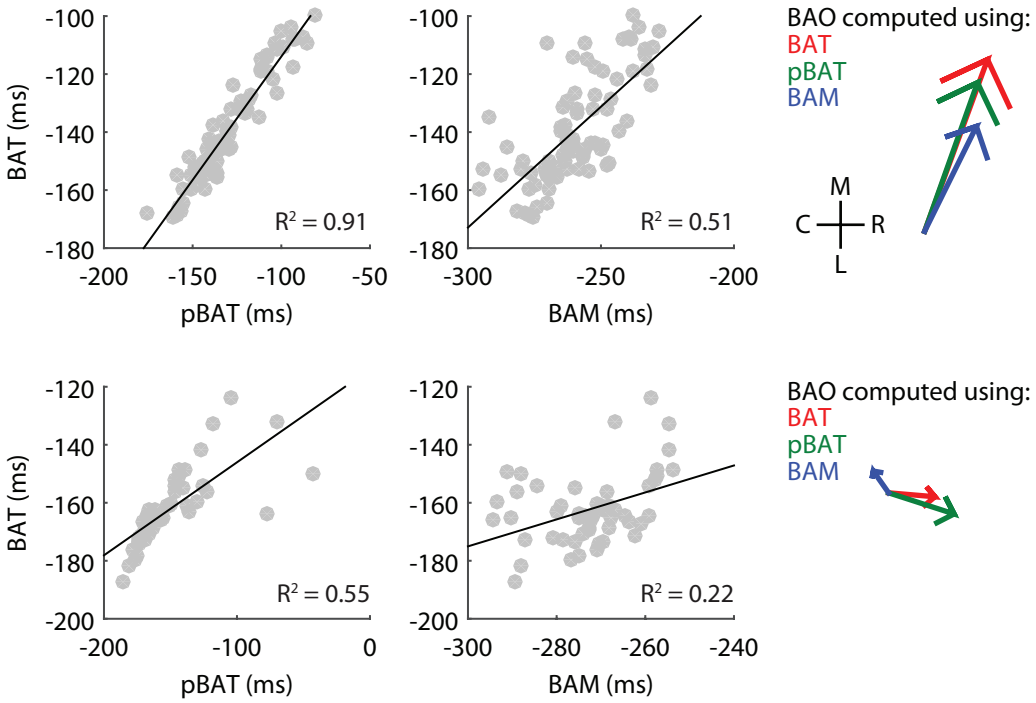

Figure 2.10 Comparing different methods of estimating beta attenuation We compared different measures of beta attenuation including BAMs, BATs, and pseudo-BATs (pBATs) computed from BAMs. In PMd (top row) we found a close correspondence between all measures. The BAOs from each measure were highly similar. In PMv, however, BATs and BAMs were not as closely related and the resulting BAOs had different orientations.

In PMd, we observed a strong correlation between all three measures of attenuation, with pBATs being extremely tightly coupled with actual BATs. The resulting BAOs estimated from each of these three summary statistics were highly similar, with the largest angle between any pair of BAOs being 6.7°. In PMv, however, there was a much weaker relationship between BAMs and BATs. The BAOs derived from each of these two summary statistics were 122.0° apart.

Effect of experimental condition on the BAO

On individual electrodes in every dataset, we observed that BAMs were significantly different across experimental conditions; yet, the correlation of BAMs

between neighboring electrodes was high across experimental conditions suggesting that the spatiotemporal pattern of BAMs may be similar across conditions. Here we considered if the spatiotemporal pattern of beta attenuation, expressed as the spatial distribution of BAMs, was similar across experimental conditions, and shifted only by an overall additive constant, or, if the spatial distribution of BAMs, i.e. the BAO, varied across experimental conditions.

We further operationalized this hypothesis by formulating it as a linear regression problem. We predicted BAMs of every electrode based on the location the object was presented, which object was presented, and the two spatial coordinates of the electrode. To test if the BAO is different across experimental conditions, we included interaction terms between electrodes' spatial location and experimental condition. We found that BAOs were significantly different in animal J (PMd: $F_{18,1244} = 8.54$, $p < 1e\text{-}8$; PMv: $F_{18,620} = 6.47$, $p < 1e\text{-}8$) and animal L (PMd: $F_{18,300} = 3.57$, $p < 0.000003$; PMv: $F_{18.284} = 3.22$, $p < 0.00002$) implying that different experimental conditions have different BAOs (Fig. 2.11 for one exemplar dataset).

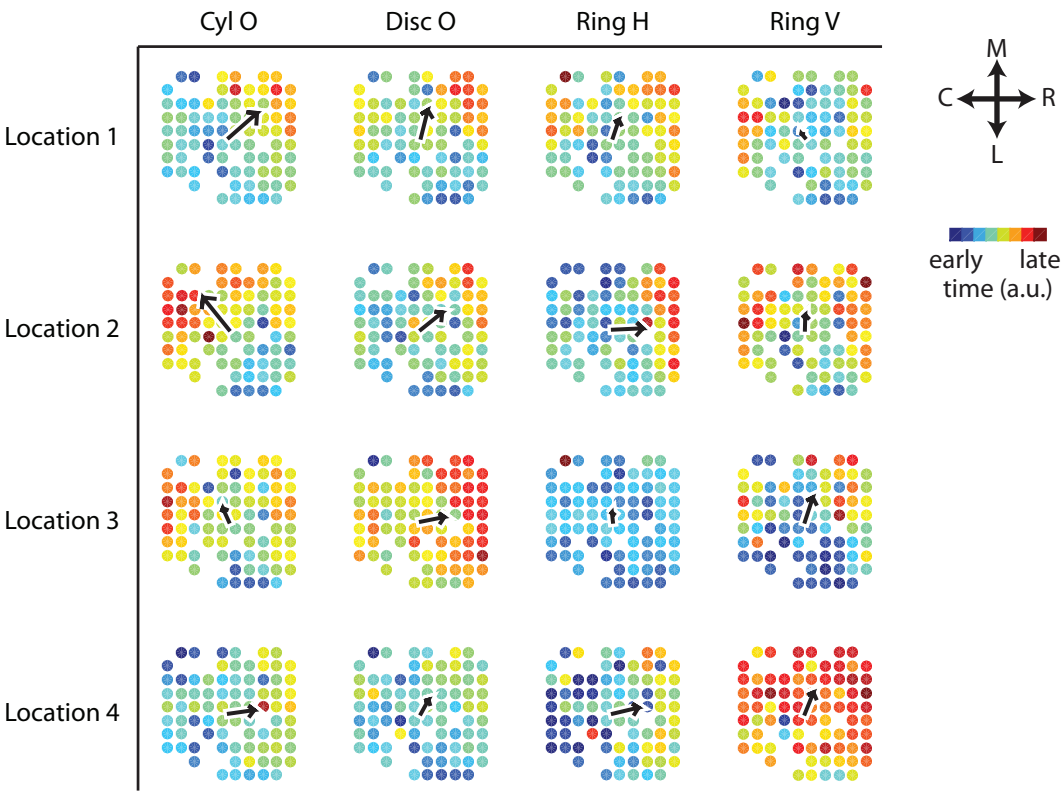

Figure 2.11 Beta attenuation varies with experimental condition. BAM maps as a function of experimental condition in PMd of animal J. Each subpanel depicts the spatio-temporal progression of BAMs as a function of their spatial location on the electrode array. The BAO for each condition is shown with a black arrow whose magnitude is proportional to the R^2 of the linear model characterizing the BAO. Note that the overall timing of beta attenuation varies across conditions (the same color corresponds to the same time across all panels, e.g. beta attenuation occurs earlier in the Ring H/location 3 condition than the Ring V/location 4 condition). BAOs are also significantly different across experimental conditions (i.e. the BAOs point in different directions for different conditions; see text for statistics).

Phase gradient waves

In both primary motor cortex and PMd, it has been shown that the instantaneous phase of LFP oscillations is not constant across the cortical surface, but rather exhibits linear wave activity (Rubino, Robbins, and Hatsopoulos 2006; Takahashi et al. 2015; Takahashi et al. 2011). These patterns have been implicated in processing of visual information related to upcoming movements though a direct link to movement has

remained elusive. We examined our data for evidence of planar wave activity in PMv and compared these results to PMd.

We computed the phase gradient directionality (PGD) at every moment in time on every trial. This statistic quantifies the alignment of phase gradients on different electrodes and can be thought of as a measurement of the strength of the wave. To infer if phase gradients were significantly aligned at a given moment in time, we employed a spatial shuffling analysis. We repeatedly recomputed PGD values on spatially shuffled data to identify moments in time that were significantly above chance. In both animals, significant wave activity in PMd was observed on every successful trial in a window spanning movement onset (1000 ms before MO to 500 ms after). Although waves were present on every trial, they were transient and consisted of only 25.7% of the total peri-movement time in animal J and 50.0% of the time in animal L. Wave activity was sparser in PMv. Again, it was observed on every trial, although it comprised only 15.4% and 45.9% of the total peri-movement window in animals J and L, respectively.

We subsequently considered the temporal properties of the wave. We first counted the number of wave transients occurring on each trial (Fig. 2.12).

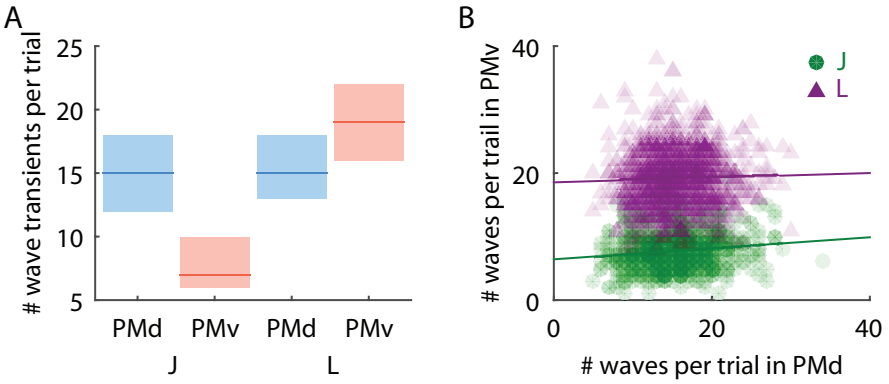

Figure 2.12 Frequency of wave activity. A. We observed that the number of wave transients per trial was different between PMd and PMv in both animals. **B.** In animal J (green circles), there was a significant correlation between the number of wave transients in PMd and PMv on a trial-by-trial basis, while no relationship was present in animal L (purple triangles; see main text for statistics).

Here, a wave transient was defined as a duration of time when the wave PGD value exceeded the maximum PGD of spatially shuffled data that were repeatedly reshuffled (10 times). In PMd, we found that there were 15 transients per trial on median in both animals, while in PMv, there were 7 and 19 in animals J and L, respectively, resulting in significantly different numbers of transients in each animal (J: Z = 25.7, p < 1e-8; L: Z = -18.01, p < 1e-8). We compared the number of wave transients on a trial-by-trial basis in PMd and PMv. In animal J, we observed a significant correlation between the number of wave events in PMd and PMv ($F_{1,613}$ = 11.8, p < 0.006) while no such correlation was present in animal L ($F_{1,1064}$ = 1.21, p < 0.272).

We additionally compared the distribution of wave transient durations (Fig. 2.13A).

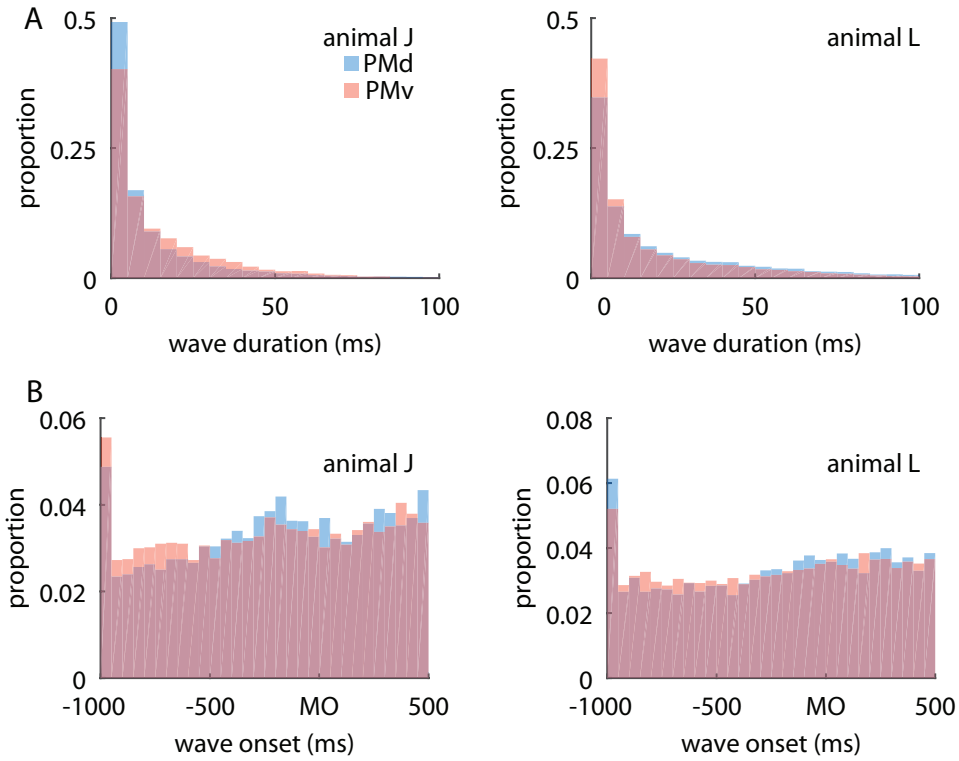

Figure 2.13 Duration of wave activity. Distribution of wave durations (A) and wave onsets (B) in animals J and L (left and right columns) in PMd (blue) and PMv (red) relative to movement onset (MO). The average wave duration was significantly different between PMd and PMv in both animals. Additionally, the distribution of wave onsets was significantly different across areas in both animals (see main text for statistics).

In animal J, waves were 12.5 and 14.9 ms on average in PMd and PMv, respectively. In animal L, waves were substantially longer at 24.2 and 18.0 ms. For both animals, the distribution of wave durations in each area was significantly different (KS-test, J: $Z = 0.10$, $p < 1e\text{-}8$; L: $Z = 0.09$; $p < 1e\text{-}8$). The distribution of wave timing was also significantly different across the two areas in both animals (Fig. 2.13B; KS-test J: $Z = 0.04$, $p < 0.00006$; L: $Z = 0.03$, $p < 0.0000006$).

Previous work has shown that wave propagation follows a bimodal distribution organized along the medio-lateral axis in PMd during planar reaching. During reach to grasp, wave propagation was unimodal and organized along a mediorostral vector in both animals (Fig. 2.14A). In PMv, less consistency was observed across animals.

Wave propagation was bimodal in J along a mediocaudal-laterorostral axis while it was unimodal along a medial direction in L (Fig. 2.14B).

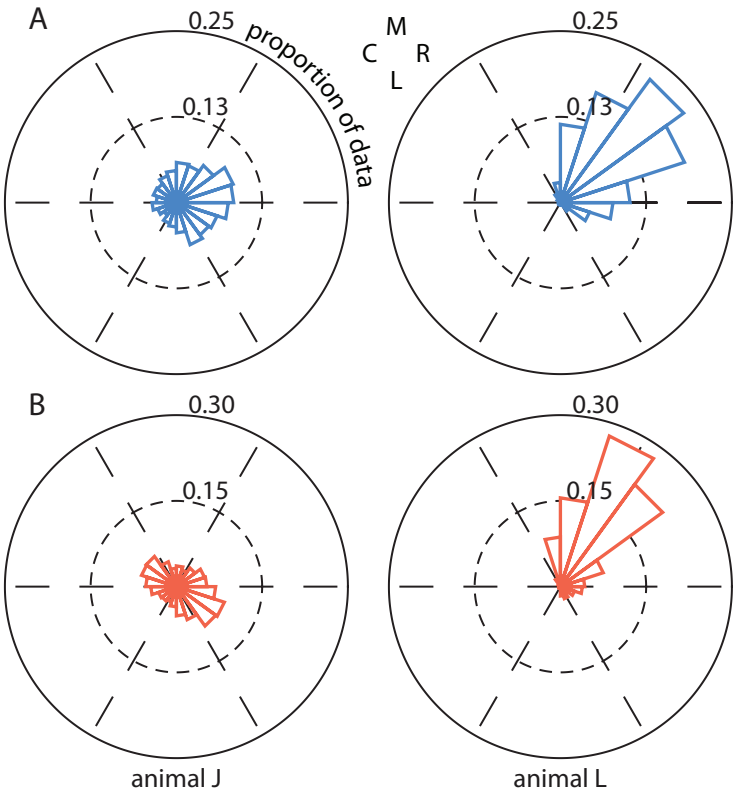

Figure 2.14 Direction of wave activity. Distribution of wave propagation direction in animals J and L (left and right) in PMd (blue) and PMv (red). Note that waves consistently propagate along a mediorostral direction in PMd for both animals, whereas there is no consistency in propagation direction or modality in PMv. Anatomical coordinates are indicated as medial (M), lateral (L), rostral (R), caudal (C).

Effect of experimental condition on wave activity

In both PMd and PMv of animal J and PMd of animal L, we found that experimental condition had no significant effect on the number of transient waves or their duration (Kruskal-Wallis test for equal medians). Animal L's PMv, however, showed weakly significant variation in the number of waves (χ^2 with 15 d.f. = 27.2, $p <$ 0.03) and their duration (χ^2 with 15 d.f. = 29.5, $p < 0.02$) with experimental condition. We performed an analog of the ANOVA for circular data, the Watson-Williams test, and

observed that the mean propagation direction of the waves significantly varied with experimental condition in both areas of both animals, however, in every dataset, the test's assumption of unimodality was violated, so its interpretation remains ambiguous.

As a complementary approach to elucidating the behavioral significance of the wave propagation direction, we performed a time-resolved information theoretic analysis that examined the mutual information (MuI) between binned wave propagation direction and experimental condition. To assess the statistical significance of the MuI values, a shuffle analysis was performed in which the experimental condition information for each trial was shuffled. We then subtracted off the 95^{th} percentile of observed shuffled MuI values from the observed values. After performing this correction, only MuI values greater than 0 were significant. We measured the frequency of significant MuI values in each animal and as a function of which experimental parameter was modified (either object or location). In animal L, we observed significant information between the wave direction and object identity 4% and 5% of the time in PMd and PMv, respectively. Significant information was observed between wave propagation direction and object location 4% and 5% of the time in PMd and PMv, respectively. For both object and location parameters, the frequency of significant MuI was entirely at chance levels suggesting that wave propagation direction does not provide information about either experimental parameter. In animal J, significant MuI between wave direction and object identity occurred 9% and 6% of the time in PMd and PMv, respectively (Fig. 2.15). Given that significant MuI occurred more often than would be expected in PMd, we further split the data into two epochs, pre and post movement onset (MO). In the post-

MO condition, we observed significant information 15% of the time with much of the information occurring between 200-600 ms after movement onset. Thus, there may be a very weak representation of object identity in the wave propagation direction in animal J following movement onset. We did not observe significant information between wave propagation direction and object location more frequently than chance would predict in either PMd or PMv of animal J (Fig. 2.15).

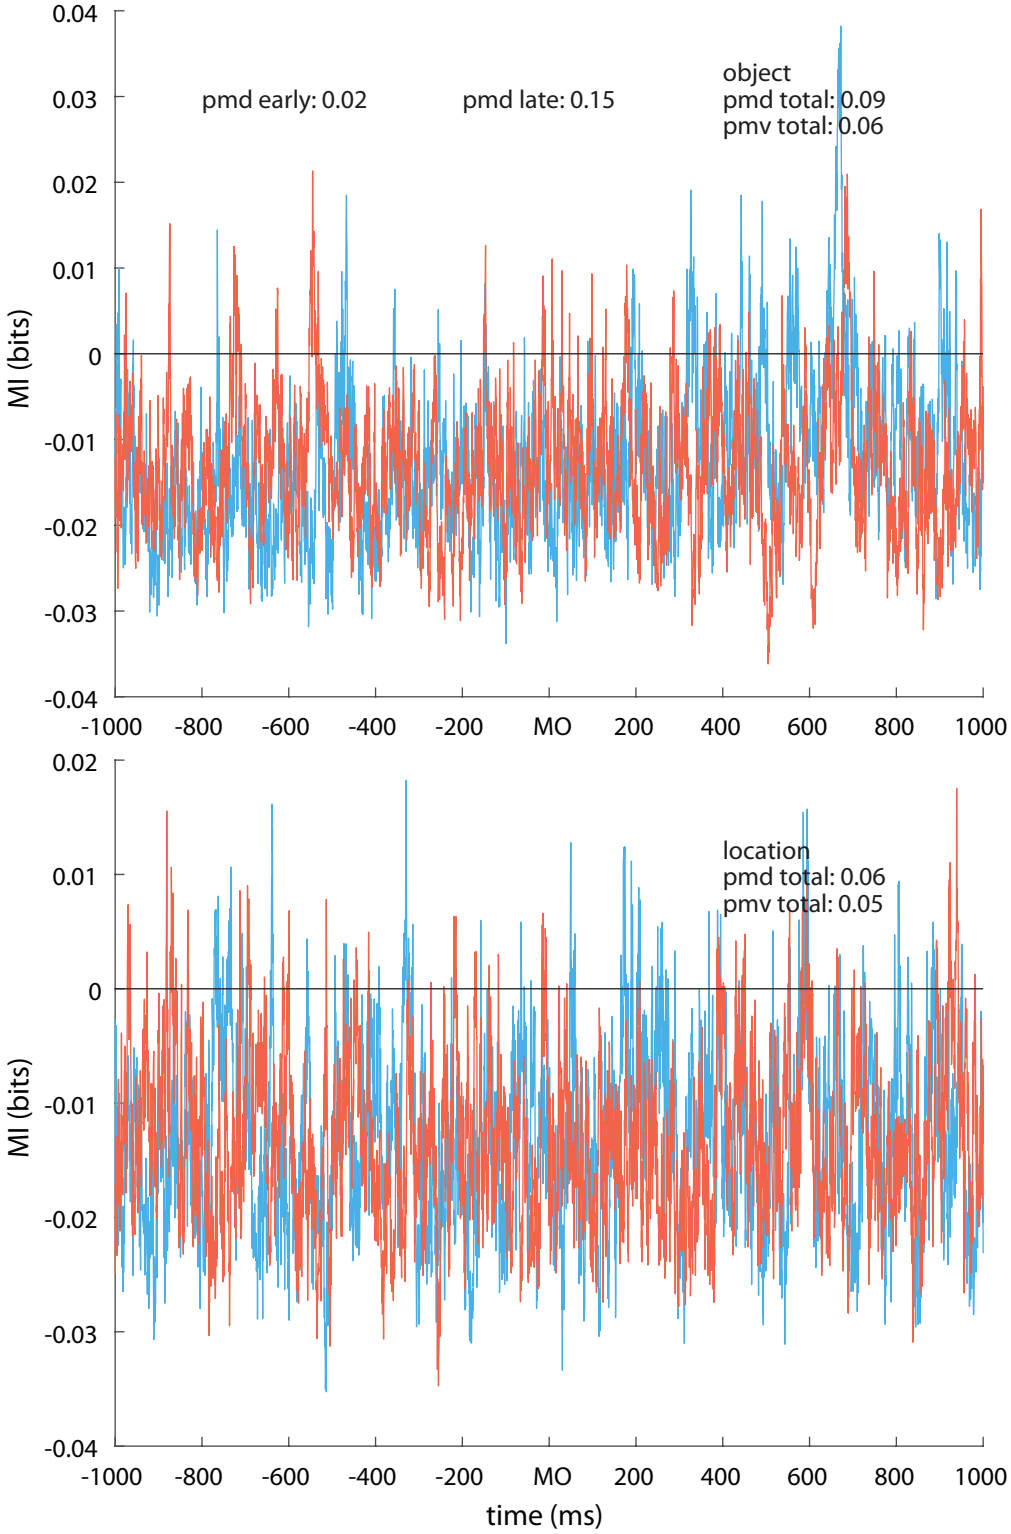

Figure 2.15. Wave direction is not informative about experimental condition. We computed the time-resolved mutual information between wave propagation direction and object or location experimental conditions. Data shown are from animal J. We observed significant mutual information between wave direction in PMd and object

identity slightly more often than would be expected under chance (top). Significant MuI occurred primarily in a window between 200-600 ms after movement onset. MuI between wave direction in PMv and object identity occurred at chance levels. In both PMd and PMv, we observed no appreciable relationship between wave propagation direction and object location (bottom).

Discussion

Beta oscillations are a ubiquitous feature of motor cortical activity and were one of the first documented cortical rhythms (Jasper and Penfield 1949). Here, we have considered how the amplitude and phase of these oscillations vary between PMd and PMv during a naturalistic reach-to-grasp task.

As a preliminary analysis, we examined the spectral content of LFP. We found that the peak frequency in the beta range was not significantly different between PMd and PMv suggesting that the beta oscillations in each area may be created by similar neural circuit architectures.

Beta attenuation

Beta attenuation has been previously interpreted as a mesoscopic reflection of the transition between preparatory and execution related cortical activity and has been shown to coincide with the sequential recruitment of primary motor neurons for movement (Appendix 1). Here, we showed that beta activity attenuates in both PMd and PMv prior to movement. We found that the relative timing of attenuation was inconsistent across the two animals used in this study. One interpretation of this finding, then, is that the relative sequence in which PMd and PMv were recruited to execute the movement was different for these two animals, which may possibly suggest that different neural strategies were used to accomplish the same behavior by the two

animals. Another possibility is that the animals each employed their own behavioral strategy to successfully complete the task. Notably, movements were self-initiated and paced at a speed chosen by the animal.

We found that a sigmoidal model better described beta attenuation in PMd than PMv. Beta attenuation has been previously interpreted as a mesoscopic reflection of the transition from preparatory to execution related activity in individual neurons (Pfurtscheller and Lopes da Silva 1999). During movement execution, many neurons exhibit directional selectivity in their firing rates, however, fewer neurons in PMv exhibit directional selectivity as compared to PMd (Suminski et al. 2015). Because fewer neurons conform to this stereotypical response type, the temporal dynamics of beta activity are not as well described by the sigmoidal model, which implicitly assumes that neurons become desynchronized during execution. This may reflect PMv's relative position in the motor hierarchy. In both MI and PMd, areas that send many direct projections to the spinal cord (He, Dum, and Strick 1993), the temporal dynamics of beta amplitude are well described by a sigmoidal function, whereas in PMv, which acts on spinal cord primarily through MI (Cerri et al. 2003; Shimazu et al. 2004), beta activity is not fit as well by that model.

In this work, we characterized beta attenuation by fitting a sigmoidal model in contrast to a previous model-free approach based on the time at which the beta amplitude crossed a threshold (Appendix 1). This new approach enabled us to quantitatively measure the extent to which beta activity was attenuating in addition to estimating when this attenuation occurred. In PMd, which exhibited clear attenuation,

there was a close correspondence between BAMs, the new measure of beta attenuation, with BATs, the old measure based on threshold crossings. In PMv, where the temporal dynamics of beta attenuation were more complex than a simple sigmoid, BAMs and BATs were less similar.

Functional significance of spatio-temporal patterns

Here, we have identified a novel spatio-temporal pattern of neural activity that precedes movement initiation in premotor cortex. The role of these patterns in movement initiation has remained elusive. However, two hypotheses have emerged based on data from MI. According to the first hypothesis, sites in motor cortex are recruited consistently with the sequential activation of upper limb muscles (Murphy, Wong, and Kwan 1985; Riehle et al. 2013). For many upper limb movements, including prehension, the dominant muscle activity pattern is proximal-to-distal (Karst and Hasan 1991; Scott 1997). Therefore, this theory would predict that sites in premotor cortex associated with the shoulder would be recruited before sites associated with the wrist. In our data, however, we found that the BAO was nearly orthogonal to the somatotopic proximal-to-distal gradient in PMd. Although we were unable to identify a somatotopic map in PMv, we speculate that there would be a similar inconsistency between the BAO and somatotopic gradient in PMv given that it does not interact directly with spinal cord, and is thus, unlikely to be chiefly concerned with low level details of movement execution.

Instead, we hypothesize that the computations performed by premotor cortex to initiate a movement are spatially distributed and these patterns reflect that activity. Our

theory makes two specific predictions: that inputs to premotor cortex from various other cortical and subcortical structures are spatially localized and integrated via a dense network of horizontal connectivity within PM, and, second, that the resulting spatio-temporal pattern may be organized to optimally drive downstream targets, in particular, the spinal cord.

Wave activity

Propagating waves in the phase gradients of beta oscillations were previously documented in MI and PMd during planar reaching, but never before had they been shown in PMv, nor in the context of naturalistic behavior. Here, we extended those previous results by documenting the existence of planar waves in PMd and PMv during unconstrained reach-to-grasp movements. In PMd, we found that waves propagated along a mediorostral vector in both animals. In contrast to planar reaching, the distribution of wave direction in PMd was largely unimodal, though the propagation vector was similarly aligned to the bimodal axis that had been reported during planar reaching.

One current theory posits that wave propagation direction (and BAO) is determined by anisotropies in the distribution of anatomical and functional connections between neurons (Takahashi et al. 2015). In MI and PMd, wave propagation has been documented along the rostro-caudal and medio-lateral axes respectively in multiple animals with various array placements within these regions. In PMv, however, we found inconsistent wave propagation directions in two animals. We speculate that this is because horizontal connectivity in PMv exhibits greater local variability, meaning that

wave activity in PMv is not dominated by a single axis, but may vary locally depending upon the precise placement of the electrode array.

As with previous studies, many properties of the wave did not bear a strong relationship to features of movement (Rubino, Robbins, and Hatsopoulos 2006; Takahashi et al. 2011). The number of waves and their duration was not significantly different across experimental conditions in either PMd or PMv in either animal. Though we observed that the average wave propagation direction differed across experimental conditions, the assumption of strong unimodality was violated in every instance obscuring the interpretation of our result. A complementary information theoretic analysis revealed that there was no information between the directionality of the waves and object location. In animal J, we observed very weak evidence that wave direction may be informative about object identity in a window of time following movement onset, but, no corresponding result was found in animal L. Taken as a whole, these results suggest that wave direction is largely unaffected by experimental condition. Combined with the observation that wave direction distributions in MI have not been shown to encode information about movement direction, we suggest that these waves may be involved in different aspects of motor cortical function other than movement encoding including possibly action selection or processing of salient visual information by the motor cortex.

Discussion

"The future is already here – it's just not very evenly distributed."

William Gibson

This thesis was concerned with elucidating the role of premotor cortex in the execution of reach to grasp movements. This class of movements was chosen as an object of study because of its profound ethological relevance to primates. A preexisting theory concerning the organization of the brain networks involved in the production of these movements was presented in the introductory chapter of this thesis. This theory argued that reach to grasp movements could be decomposed into separate reaching and grasping movements that were generated by discrete neural networks. In subsequent chapters, aspects of this theory were critically tested, and shown to be deficient in explaining the neural activity of premotor cortex.

Because each chapter included its own standalone discussion section, the remainder of this section will focus on seating the general conclusions of this thesis in the broader context of current motor cortical neurophysiology. As part of this discussion, the limitations of the dual channels theory will be reviewed, as well as its utility as an explanatory framework for future studies. Additional discussion will be devoted to furthering the lines of inquiry motivated by this thesis with special focus placed on elucidating the importance of spatial and spatio-temporal patterning in MI.

Relationship of PM and MI

The dual channels hypothesis posited a largely hierarchical relationship between premotor and primary motor cortices. Within this framework, premotor cortex was

thought to be substantially more involved in translating visual information into a motor plan, and then transmitting that plan to be executed by the primary motor cortex. Anatomically, the horizontal connectivity between these areas (Richard P. Dum and Strick 2005) as well as their patterns of descending output to spinal cord (He, Dum, and Strick 1993) do not support such a simplistic organization of the motor cortex. As with the anatomy, the results of this thesis, too, are largely inconsistent with a hierarchical organization of motor cortex.

In previous work, it was shown that neurons in MI encode complex temporal kinematic trajectories spanning multiple joints (Saleh, Takahashi, and Hatsopoulos 2012). In chapter 1 of this thesis, we observed several similar findings to those reported by Saleh and colleagues. As with MI, we found that many cells in PM were better predicted by temporally extensive joint trajectories as compared to static relationships at a single, fixed lag. Additionally, many cells in PM were well-predicted by functional kinematic synergies. Unlike Saleh and colleagues, we performed a kinematic decomposition that identified functionally relevant kinematic synergies, rather than raw synergies explaining maximal variance, so a direct comparison between synergy encoding in MI and PM is not feasible, yet it can be argued that both have been shown to encode synergies.

This does not imply, however, that there are no differences in encoding properties between PM and MI. In particular, when comparing the preference of single cells for intrinsic versus extrinsic features, it was found that spiking activity in MI cells are better predicted by extrinsic features (i.e. kinematics) (Saleh, Takahashi, and

Hatsopoulos 2012). In contrast, in both PMd and PMv, we found that intrinsic features (i.e. spike history) better predicted the spiking activity of cells in each area.

It should be noted that the spike history terms that were used in these two studies were similar, but not identical. In particular, Saleh and colleagues used more basis functions to model spike history effects at a wider variety of characteristic lags. Using more basis functions confers two principal advantages: an ability to examine spike history at longer time-scales, and at a denser temporal resolution. The downside of this approach is that a larger model must be fit, and with finite data, is more prone to overfitting. Thus, one possibility is that spike history is only seemingly more important in PM than MI because MI spike history models were overfit. But, this possibility seems unlikely because kinematic encoding models in MI had high AUC values and larger numbers of parameters than spike history models. Conversely, it could be argued that the kinematic models in PM were overfit, but again this is unlikely given that models were also fit using a small number of kinematic synergies and realized a same level of performance as the full kinematic models.

The phenomenon of beta attenuation in MI was the subject of a recent study (Appendix 1). In MI, beta attenuation was found to occur approximately 100 ms before movement onset (though there was some variability across animals). In both PMd and PMv, we found that beta attenuation also occurred approximately 100 ms before movement onset. Although the tasks performed by the animals in these two studies were different, these findings suggest that beta attenuation is largely concurrent in both premotor cortex and MI. Based on the interpretation of beta attenuation as a

mesoscopic reflection of units engaging in execution related activity, this implies that PM and MI both transition into a neural state associated with movement execution at approximately the same time. Recent work examining population dynamics of ensembles of cells in these areas further supports this finding (Suminski et al. 2015).

At a higher level, this finding is consistent with a new perspective that premotor cortex is considerably closer to the motor periphery than previously appreciated. It has been speculated that the spatio-temporal patterning of beta attenuation in MI is necessary for normal movement initiation, presumably because the spinal cord is preferentially driven by MI activity exhibiting that pattern (Appendix 1). Based on this interpretation, it would suggest that PMd is also intimately involved in transmitting information to spinal cord because beta attenuation is also strongly spatio-temporally patterned. This finding is compatible with anatomical evidence suggesting that PMd supplies much input to the corticospinal tract (R. P. Dum and Strick 1991; He, Dum, and Strick 1993). We did not observe strong spatio-temporal patterning of beta attenuation in PMv, which is consistent with its putative role as a facilitator of MI modulation (Shimazu et al. 2004; Cerri et al. 2003). Collectively, this evidence is inconsistent with the hierarchical organization of motor cortex proposed in the dual channels hypothesis. Specifically, it would not predict such strong spatio-temporally patterned beta attenuation in PMd.

Previous work examined wave propagation in both MI and PMd during a planar reaching task (Rubino, Robbins, and Hatsopoulos 2006). Our results are entirely consistent with those of the previous study. We found that waves propagate along a

latero-medial gradient in PMd during naturalistic movement. Previously, these waves were speculated to coordinate the proximal-to-distal (PD) sequence of muscle activation (Hatsopoulos, Olmedo, and Takahashi 2010), yet, here, we found the waves propagate along an axis that is inconsistent with the PD gradient revealed by ICMS. Here, during naturalistic reach to grasp, we observed that wave propagation revealed no information about either the spatial location of the object in the workspace, or which object was to be grasped. Thus, while waves were present in both PMd and MI, they may serve different roles in each area with waves in MI being more directly involved in movement execution.

Expanding upon these previous results, we documented the existence of wave-like activity in PMv. As with PMd, we observed no relationship between PMv wave propagation direction and experimental condition, and, the wave propagation direction was different across animals. This finding suggests that PMv is a heterogeneous area, and that perhaps our electrode arrays were not placed in precisely the same part of PMv.

Separating representations of function and anatomy

One of the most striking results in all of systems neuroscience is that of the sensory homunculus (Kandel, Schwartz, and Jessell 2013). And yet, there is no corresponding clear map of body representation in motor cortex. At a coarse scale, body segments are well separated; however, within a body segment, such as the upper limb, the representation is much more complex, and overlapping (Boudrias et al. 2010). This has led to an argument revolving around what information may actually be mapped

spatially across the cortical surface. On one hand, several (Riehle et al. 2013; Park, Belhaj-Saïf, and Cheney 2004; Rathelot and Strick 2006) have argued for a somatotopic organization; that is, each point on the cortical surface represents a muscle or group of muscles. An alternative perspective is that actions are mapped to points on the cortical surface, such that different behaviors are represented on different parts of the cortical surface (M. S. A. Graziano 2005; Michael S. A. Graziano 2016).

The support for both of these perspectives comes from ICMS of motor cortex. Stimulation with single pulses, or short pulse trains at low current intensities reveals a putative somatotopic map (Park, Belhaj-Saïf, and Cheney 2004). In contrast, long duration, high amplitude stimulation evokes naturalistic movements and has led to the suggestion of action maps (Michael S. A. Graziano 2016). The precise mechanism that causes these behaviors to emerge in this high amplitude stimulation regime remains poorly understood. However, it was the subject of one recent study (Griffin et al. 2014).

In this thesis, we performed ICMS that was largely consistent with the somatotopic approach. That is, we identified the effects of stimulation at threshold current intensities on each of the electrodes. We identified a coarse proximal-to-distal gradient of limb representation based on the ICMS study. Classically, it has been assumed that any putative spatio-temporal pattern of neural activity is determined solely by an interaction between the temporal sequence of limb activation and the spatial disparity in limb representation in cortex (for a longer statement of this hypothesis see (Riehle et al. 2013; Murphy, Wong, and Kwan 1985)). Here, we found that large-scale patterns of neural activity, manifested as beta attenuation, were not spatio-temporally

structured according to the putative proximal-to-distal gradient of limb representation revealed by ICMS. This finding suggests that spatio-temporal variability in motor cortical responses are not solely dependent on limb representation, but may be determined by other factors as well.

We also identified the anatomical tuning preferences (i.e. preference for arm or hand kinematics) for every single cell in our sample. Using a novel dimensionality reduction technique, we identified functional kinematic synergies that were distinct from anatomical representations of the limb. From this analysis, in single cells we found no relationship between their preference for arm or hand kinematics and their preference for reaching and grasping.

Collectively, our findings at the mesoscopic level of the LFP, and microscopic level of single unit activity suggest that limb representations may be multiplexed with behavioral representations in a complex way. These multiplexed representations in single cells would lead to overlaying representational maps distributed over the cortical surface. The notion that cells may exhibit this type of multiplexed representation is not unprecedented. Indeed, the study of mixed selectivity in prefrontal cortex has been an exciting area of recent research (Mante et al. 2013; Rigotti et al. 2013). Here, we argue that there may similarly be mixed representations of the underlying anatomy and the behavioral intent of the animal. Such mixed selectivity might possibly explain why identifying a single coordinate frame for motor cortical computation has proved challenging (Wu and Hatsopoulos 2006; Reimer and Hatsopoulos 2009).

This idea concerning a dissociation of function and anatomy strongly contradicts the expectation of the dual channels hypothesis. Particularly, it would be predicted that PMd should care exclusively about arm kinematics and reaching movements, and PMv should care exclusively about hand kinematics and grasping movements.

Implications for neural prosthetics

Pragmatically, the findings of this thesis may inform the development of a clinically viable brain machine interface (BMI). It has been previously shown that whole arm kinematics can be decoded from the activity of neural ensembles in PMd and PMv (Mollazadeh et al. 2014; Aggarwal et al. 2013; Bansal et al. 2012). While these studies demonstrate the feasibility of using PM cortical activity to drive a prosthetic limb, they do not provide insight as to why PM activity is effective at doing so. In studying the encoding properties of these cells, we discovered that both areas represent complex spatio-temporal kinematic trajectories of multiple joints, often spanning the entirety of the limb. Accordingly, using PM cortical activity to decode kinematic synergies may lead to better decoding performance than attempting to predict the activity of single joints, the standard paradigm employed in many decoding studies (Aggarwal et al. 2013; Carmena et al. 2003; M.D. Best et al. 2013; Matthew D. Best et al. 2014).

We found that extrinsic parameters, either raw kinematics, or kinematic synergies, both were more strongly predictive of activity in PMd as compared to PMv. Although kinematics could predict spiking activity in PMv neurons significantly better than chance, kinematics were nevertheless a fairly ineffective predictor of spiking activity as typical area under the receiver operator characteristic curve (AUROC) values

were less than 0.6. Recently, much interest has been placed on the development of a bi-directional brain machine interface; that is, one that provides feedback to the brain from sensors embedded in the prosthesis. Here, we suggest that PMv may be an effective locus for delivery of that feedback signal from the prosthetic. Given that kinematics cannot be reliably decoded from it, yet it is strongly reciprocally connected with MI, it would be an ideal location to provide sensory feedback.

Finally, it has been suggested throughout this thesis that there are fundamental differences between the neural activity that is associated with movement execution and during other times. Despite some physiological evidence to support this perspective (Churchland et al. 2012; Kaufman et al. 2014; Shenoy, Sahani, and Churchland 2013), many studies focused on decoding neural activity in a BMI context fail to account for these differences. Accordingly, neural activity is translated into movement, even in the absence of any such intention.

That movement is always translated into action is not particularly problematic in a laboratory setting in which an animal subject is encouraged to engage in a BMI task. But, for a human being with an upper limb neural prosthesis navigating the minutiae of everyday living, such constant movement could prove to be remarkably burdensome. Action observation produces robust responses in premotor cortical neurons (and even in primary motor cortex as well) (Hatsopoulos and Suminski 2011; Suminski, Tkach, and Hatsopoulos 2009). So, activities such as watching sports on television, or seeing a movie in a theatre could lead to wild, unintended movements of the prosthetic limb. One trivial solution to this problem would involve intentionally disabling the prosthesis

during “passive” activities, much like we may intentionally set our cell phones to silent mode. But, this approach fails to appreciate our natural ability to imagine movements without executing them.

One major reason that this problem has been largely ignored is that identifying the intention to move from spontaneous, baseline motor cortical activity in the absence of any external cues or task structure has proven a formidable challenge (Chi et al. 2007). In this thesis, we identified spatio-temporal activity patterns of trial-averaged PMd activity that characterize movement onset. The extent to which these patterns may be identified on a single trial basis remains unknown; however, they offer a promising avenue for distinguishing epochs of intended movement from baseline activity.

These patterns were characterized, in part, based on the spatial relationship between electrodes on the array. Almost all current BMI decoding algorithms fail to account for the spatial arrangement of their electrodes, and thus ignore any putative spatial inter-dependencies amongst the neural signals sampled by the electrodes. This is despite the fact that causal studies have indicated the brain may be able to discern spatial patterns of ICMS (O’Doherty et al. 2012). Thus, there is a growing body of evidence to suggest that next-generation BMI decoding algorithms ought to give consideration to the spatial properties of both the electrode array, and the neural tissue in which it is embedded.

Future work

This thesis has tested Jeannerod’s dual channels theory of motor cortical organization through consideration of the involvement of premotor cortex in reach to

grasp movements. In multiple instances throughout this thesis, experimental evidence was presented that was inconsistent with the expectations of this theory. Accordingly, it was argued to no longer be an effective explanatory framework for understanding premotor cortical function. Instead, it was suggested that premotor cortex and primary motor cortex work in concert to generate reach to grasp movements. A growing body of literature now appreciates the perspective that there is not a strict hierarchy between premotor and primary motor cortices (Suminski et al. 2015; Richard P. Dum and Strick 2005; He, Dum, and Strick 1993), or that reach and grasp are strictly segregated (Hocherman and Wise 1991; Raos et al. 2004; Bansal et al. 2012; Aggarwal et al. 2013). The quote that begins this chapter pays service to the fact that there was not a single clear moment in time at which this perspective became dominant, but rather, emerged from evidence accumulated over many years by multiple researchers.

Arguably the most significant contribution of this thesis is that premotor cortical activity exhibits spatio-temporal structure that is associated with specific behavioral events. Here, it was argued that this spatio-temporal patterning is not a mere reflection of the spatio-temporal dynamics in the motor periphery (Riehle et al. 2013; Murphy, Wong, and Kwan 1985), but rather reveals the structure of the cortical computations that support movement. In particular, it was hypothesized that the spatio-temporal patterning of beta attenuation indicates premotor cortical neurons engaging in movement-execution related spiking. We will now describe an experiment designed to causally test this hypothesis.

In this thesis and elsewhere, it has been argued that movement onset is characterized by specific spatio-temporal sequences of neural activity. Accordingly, it may be argued that these patterns are necessary for normal movement initiation, and that disruption of these patterns should delay movement onset. Here, we will describe an experiment using spatio-temporally patterned ICMS to test this hypothesis.

In this experiment, awake behaving rhesus macaques will receive ICMS to either PM or MI while they perform a reach-to-grasp task. Spatio-temporally varying cathodal-anodal pulse trains will be generated with a Blackrock current stimulator (pulse width: 200 μs, train frequency: 250 Hz, train duration 68 ms).

Three unique spatio-temporal patterns of stimulation will be used. All three patterns will stimulate the same set of electrodes, distributed along the beta attenuation orientation of a given array. One pattern of stimulation will be chosen to mimic the spatio-temporal pattern of beta attenuation. The stimulation will propagate across the Utah array at approximately the same speed as the beta attenuation phenomenon. This stimulation pattern will be referred to as the "WITH" condition. The second pattern, "AGAINST," is the exact reverse of the "WITH" condition. For the last condition, "RANDOM," we will stimulate without any obvious spatial structure along the BAO by randomly selecting sites to stimulate.

On each experimental session, and for each stimulation pattern, we will identify the current intensity threshold that evokes movement. One experimenter will observe and actively palpate the animal during stimulation to visually or tactilely identify muscle activity due to stimulation, while another experimenter performs stimulation. To ensure

that any putative behavioral effects are a result of disrupting intrinsic motor cortical dynamics and not directly because of stimulation, the experiment will proceed with currents equal to 80% of threshold for stimulation in MI, and 25% of threshold for stimulation in PM. In both cases, this should result in current intensities around 10 μA.

In order to probe the temporal specificity of any putative behavioral effect, the timing at which stimulation is administered will be systematically varied. Stimulation will be delivered before, during, and after the attenuation phenomenon to test for temporal specificity of stimulation effects. Stimulation timing will be varied across experimental sessions.

The specific pattern of stimulation that will be delivered on a trial-to-trial basis will vary, but will be balanced across conditions. In addition to active stimulation conditions, an equal proportion of trials will also lack stimulation.

Behaviorally, we will measure the animals' reaction time (RT) as well as detailed movement kinematics. We expect to find that stimulation that is inconsistent with the BAO will lead to delayed movement initiation as compared to stimulation that is consistent with the BAO. We expect that this effect will be temporally limited to occur only during the beta attenuation epoch and not elsewhere. We expect to observe this spatio-temporal effect of stimulation in PMd and MI, but not PMv, because no strong spatio-temporal pattern of beta attenuation was observed there.

At the same time, one critical aspect of the dual channels hypothesis remains untested by this proposed experiment: to what extent is the representation of reach to grasp movements redundant between PMd and PMv? To date, few studies have

performed simultaneous recordings in these areas despite the fact that there is known to be dense anatomical interconnectivity between them (Richard P. Dum and Strick 2005). While decoding studies have demonstrated that information about hand kinematics is present in population activity from both areas, they do not provide insight into how this information is encoded by these populations of cells. An open question remains: is the information about grasping in PMd present because of an interaction with PMv, or, is it propagated along as a separate, independent representation of grasping? And, to a lesser extent, is any putative reach representation in PMv due to interaction with PMd, or rather an independent representation? Resolving these two open questions will directly, and critically, test the dual channels hypothesis.

Appendix 1 - Spatio-temporal patterning in primary motor cortex at movement onset

This scientific content of this chapter was accepted for publication at Cerebral Cortex and was in press at the time this thesis was submitted.

Abstract

Voluntary movement initiation involves the engagement of large populations of motor cortical neurons around movement onset. Despite knowledge of the temporal dynamics that lead to movement, the spatial structure of these dynamics across the cortical surface remains unknown. In data from four rhesus macaques, we show that the timing of attenuation of beta frequency local field potential oscillations, a correlate of locally activated cortex, forms a spatial gradient across primary motor cortex (MI). We show that these spatio-temporal dynamics are recapitulated in the engagement order of ensembles of MI neurons. We demonstrate that these patterns are unique to movement onset, and suggest that movement initiation requires a precise spatio-temporal sequential activation of neurons in MI.

Introduction

Large ensembles of neurons in primary motor cortex (MI) are used to initiate movements. Even the simplest, single joint movement of the arm or hand recruits neurons distributed across the entire upper limb region of MI (Sanes and Donoghue 1993; M. H. Schieber and Hibbard 1993); however, it remains unknown if the order in which these units are engaged is related to their spatial position on the cortical sheet.

At the same time, the amplitude of oscillations in the beta band (15-30 Hz) of the local field potential (LFP) attenuates sharply around movement onset (Jasper and

Penfield 1949; Sanes and Donoghue 1993; Murthy and Fetz 1996a; Rubino, Robbins, and Hatsopoulos 2006). It is thought that this attenuation is a mesoscopic reflection of activated motor cortex, and coincides with an epoch of enhanced corticospinal excitability (Pfurtscheller and Lopes da Silva 1999). Although the spatio-temporal dynamics of beta attenuation are poorly understood, beta oscillations have been shown to propagate along the rostro-caudal dimension during motor preparation (Rubino, Robbins, and Hatsopoulos 2006). These planar waves may be determined by the underlying pattern of horizontal connectivity within MI which exhibits a preponderance of long range connections distributed along the rostro-caudal dimension (Gatter and Powell 1978; Huntley and Jones 1991). Using a Granger causality analysis, a recent study found that pairs of neurons are preferentially connected along this dimension, and that sequential spiking between these pairs of neurons contains more information about movement direction (Takahashi et al. 2015).

We reasoned that if spatial patterning in MI is related to its horizontal connectivity, spatial patterns may also emerge during movement initiation. Thus, we hypothesized that movement initiation is characterized by a specific spatio-temporal sequence of neural activity structured along the rostro-caudal axis. Furthermore, if such a pattern is indeed a unique signature of movement initiation, we would expect it to be absent during preparation, another epoch in which motor cortex is known to be transiently activated.

Materials and methods

Behavioral task

Four macaque monkeys (Macaca mulatta) of either sex were operantly trained to perform an instructed-delay center out reaching task (for a complete description of the task, see (Rubino, Robbins, and Hatsopoulos 2006)). Briefly, each animal used an exoskeletal robot (BKIN Technologies, Kingston, Ontario, CA) to make planar reaching movements that controlled the position of a cursor on a horizontal screen above the animal's arm. The animal had to hold the cursor on a center target for 500 milliseconds before an instruction cue appeared. The instruction cue signaled the location of a peripheral target to which the animal had to ultimately move. Peripheral targets were located at eight evenly spaced radial intervals around the center target with a distance of 5-7 cm between the center and peripheral targets. After receiving the instruction cue, the animal had to hold the cursor in the center target for a random amount of time (uniformly distributed between 1000 and 1500 milliseconds) until a go cue appeared. Upon presentation of the go cue, the animal could move towards the peripheral target, and, if it held the cursor at the cued peripheral target for 500 milliseconds, it received a juice reward (Fig. A1.1A).

Behavioral data selection

We computed the reaction time and movement duration of each trial. We defined reaction time as the difference between the moment the cursor left the center target (i.e. movement onset) and go cue appearance. We excluded trials with reaction times less than 100 ms and greater than 750 ms. Movement duration was defined as the time between movement onset and acquisition of the peripheral target. We imposed a threshold on movement durations such that they had to be at least 150 ms, but no more

than 1500 ms. Additionally, we observed a small percentage of trials that were considerably slower than the median, so we excluded the slowest 25% of trials from each movement direction to minimize variability arising from the behavior.

Neurophysiology

All surgical and experimental procedures were approved by the University of Chicago Animal Care and Use Committee and conform to the principles outlined in the Guide for the Care and Use of Laboratory Animals (NIH publication no 86-23, revised 1985). All monkeys were implanted with 96-electrode Utah arrays (1.0 mm length, 10 x 10 grid with 400 μm inter-electrode spacing, Blackrock Microsystems, Salt Lake City, UT) in MI contralateral to their working hand. During a recording session, signals from up to 96 electrodes were amplified with a gain of 5000, band-pass filtered between 0.3 Hz and 7.5 kHz and recorded digitally (14-bit) at 30 kHz using a Cerebus acquisition system (Blackrock Microsystems, Salt Lake City, UT). Local field potentials (LFPs) were sampled at 1 kHz and digitally band pass filtered between 0.3 and 500 Hz or 0.3 to 250 Hz using the Cerebus. Single units were sorted offline using a semi-manual procedure (Offline Sorter, Plexon Inc., Dallas, TX).

Beta attenuation analysis

We performed a number of preprocessing steps on the LFPs to estimate beta attenuation. To begin, we found the frequency of peak power in the beta range (15-30 Hz) for each dataset, averaged over all electrodes and time. LFPs from each electrode were then filtered bidirectionally with a bandpass filter (Butterworth, 4th order) whose passband was 3 Hz on either side (i.e. a total bandwidth of 6 Hz) of the peak beta

frequency power (animal Rs, 18 Hz; Rj, 18 Hz; Rx, 27 Hz; V, 21 Hz rounded to the nearest Hertz). We found that rounding the center of the filter to the nearest integer frequency had no effect on subsequent results. Power spectra from each animal were computed using the chronux (http://chronux.org) function mtspectrumc, log transformed into decibels, and then averaged over electrodes.

The Hilbert transform was applied to the band-pass filtered data to estimate the instantaneous beta amplitude of each LFP. For each electrode, the instantaneous amplitudes were trial-aligned to movement onset, and then averaged across trials, thereby estimating beta attenuation.

We found the maximum and minimum values of each electrode's beta profile in a window spanning +- 500 ms around movement onset. Beta profiles were normalized so that they ranged from 0 to 1. The time at which the normalized profile crossed an attenuation threshold (0.15) was computed for each electrode. This time, referred to as the beta attenuation time (BAT), is one of the primary summary statistics of this paper. Early analysis experimented with several values of the attenuation threshold, and we found comparable results over a moderate range of attenuation thresholds. For some electrodes, visual inspection of their power spectrum revealed no significant bump in the beta frequency band; these electrodes were discarded from future analysis (2 from V and Rs, 4 from Rj, and 7 from Rx).

Computing the beta attenuation orientation

We characterized the spatio-temporal progression of beta attenuation across the Utah array by using linear regression to estimate a vector that described the spatial

orientation of BATs from earliest to latest. This vector will henceforth be referred to as the beta attenuation orientation (BAO). For each electrode, we predicted the BAT, denoted y_{rc}, from the spatial location of that electrode. The subscripts *r* and *c* are used to indicate the row and column coordinates of the electrode on the array. Mathematically, this relationship can be expressed as

$$y_{rc} = b_r r + b_c c + a,$$

where b_r and b_c are the coefficients of the rows and columns, and *a* is a constant time offset. The arctangent of the ratio of b_r and b_c indicates the orientation of beta attenuation, i.e. the BAO.

Goodness of fit was quantified using the coefficient of determination, R^2, defined as

$$R^2 := 1 - \frac{\sum_{rc}(y_{rc} - \hat{y}_{rc})^2}{\sum_{rc}(y_{rc} - \bar{y})^2}$$

where y_{rc} and $\hat{y}_{rc}$ denote the observed and fitted beta attenuation times for electrode *rc*, and $\bar{y}$ corresponds to the average beta attenuation time across all electrodes.

We used two fundamentally different procedures to test for the significance of the BAO. The first involved performing an *F*-test for the significance of our regression model by comparing it to a model that contained only the offset term *a*. The second test for significance of the BAO was based on a spatial shuffling of our data (Riehle et al. 2013). After computing BATs for each electrode, we shuffled the spatial location of the BATs on the Utah array. We then estimated the BAO and corresponding R^2 of the shuffled data. This procedure was repeated one million times to estimate a null distribution of R^2 values in a population where spatial information had been explicitly

destroyed. We additionally tested the consistency of the BAO. We repeatedly (1000 times) randomly divided each dataset into two equal sized halves and estimated a BAO for each half. We then computed the angle between the two BAOs and compared it to a null distribution of BAOs based on spatially shuffled data (using the preceding shuffling procedure).

Circular dissimilarity test

To statistically quantify the similarity between two directional vectors (e.g. the BAOs estimated from two discrete subsets of trials), the angle between them was measured and compared to the distribution of angles if two uniform random vectors were drawn independently of one another. This null distribution is uniform between 0° and 180°. To compare multiple angles simultaneously, their sum was computed and compared to the sum of n uniformly distributed random variables (on the interval [0, 180]), which, is given by the Irwin-Hall distribution (rescaled to the interval [0 180]). For large values of n, the Irwin-Hall distribution is approximately normal with mean $n/2$ and variance $n/12$.

Unit spiking analysis

For every dataset, each unit's spiking activity was binned into 1 ms bins, and convolved with a Gaussian kernel (25 ms s.d. for monkeys Rs and V, and 50 ms s.d. for Rj due to lower trial counts) to smooth firing rate estimates. We trial-aligned the smoothed firing rates to either the instruction cue, or movement onset, and then performed an analysis to identify the time at which each cell was most informative about movement. On every trial, we predicted the instantaneous probability that the animal

was moving to each of the targets. We used the entropy of these probabilities to quantify the uncertainty about movement direction in that neuron. These probabilities were estimated from each neuron individually using a multinomial logistic regression model that will be subsequently described in greater mathematical detail.

We represented the firing rate of a given neuron, *i*, in a matrix, *X*, where each row corresponded to the *j*th of *J* trials, and each column corresponded to a time, *t*, relative to movement onset. At every moment in time and on every trial, i.e. for each *t* and *j*, we used the firing rate of neuron *i* as the input of our multinomial logistic regression model. This model may be expressed as:

$$\hat{p}_{kjt} = \frac{\exp\left[\alpha_{kt} + \beta_{kt} \cdot X_{jt}^{i}\right]}{\sum_{k=1}^{8} \exp[\alpha_{kt} + \beta_{kt} \cdot X_{jt}^{i}]}$$

where $\hat{p}_{kjt}$ represents the estimated probability that the animal moved to target *k* on trial *j* at time *t*, and α_{kt}, β_{kt} represents the regression coefficients associated with target *k* at time *t*. Due to low trial counts in monkey Rj, we pooled adjacent movement directions as a single target location (i.e. targets at 0° and 45° were pooled). These 8 (4 for Rj) probabilities comprise a probability distribution, Z_{jt}, and we estimated the entropy of that distribution, $H(Z_{jt})$, as follows:

$$H(Z_{jt}) = -\sum_{k=1}^{K} \hat{p}_{kjt} \log_2 \hat{p}_{kjt}$$

To disambiguate chance fluctuations in entropy due to stochastic spiking from behaviorally significant modulation, we performed a bootstrap analysis. We repeatedly shuffled target locations (200 times) and recomputed the entropy, $H(Z_{jt})$, for each neuron. For each moment in time, we subtracted the 5th percentile of bootstrapped

entropy values (corresponding to reductions in entropy due to stochastic spiking) from the unshuffled entropy values and added back 3 bits so that the corrected data were on approximately the same scale as uncorrected data. After performing this correction, only entropy values less than 3 bits (2 for Rj) were considered significantly task modulated.

We manually identified local minima in the entropy profiles during the instruction (50 to 350 ms after the instruction) and execution epochs (175 ms before to 50 ms after movement onset for Rs and V, 325 to 25 before movement onset for Rj because its beta attenuation range was earlier, and longer). The time of the local entropy minimum defined that unit's modulation time, i.e. its UMT. In additional analyses, we defined UMTs based on the timing of threshold crossings in the entropy profile s and observed no qualitative difference between minima and threshold crossings. The same linear regression procedure that was used to find the BAO was also used to find the unit modulation orientation (UMO). The UMO represents the spatial orientation along which motor cortical neurons are sequentially engaged from early to late.

Results

We recorded single unit spiking activity and local field potentials (LFPs) from multi-electrode arrays implanted in the primary motor cortex (MI) of four rhesus macaques while they engaged in an instructed-delay, center-out reaching task (Fig. A1.1A).

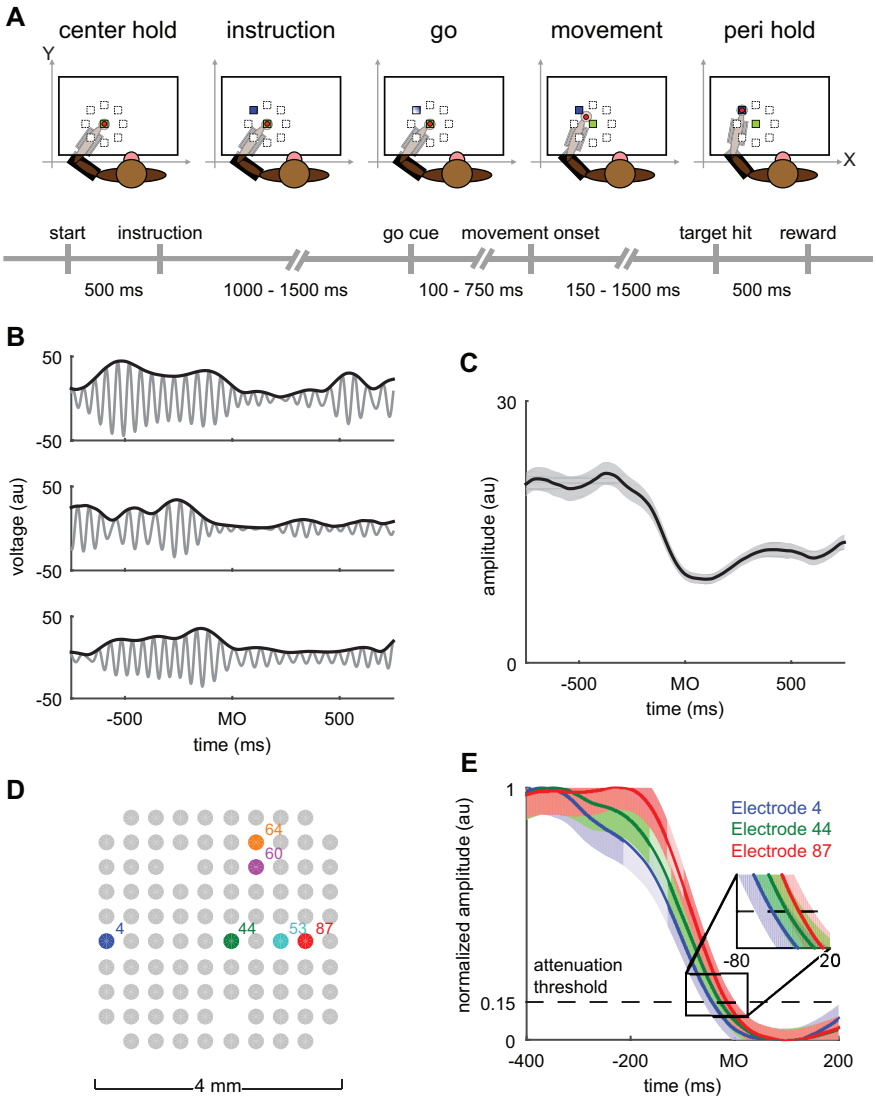

Figure A1.1 Experimental setup and beta attenuation analysis. A. Animals were trained to perform an instructed-delay center out reaching task. The animals used a 2-link robotic exoskeleton to control the position of a cursor on a screen projected above their arm. **B.** LFPs were filtered into the beta frequency range and aligned to movement onset (MO) for every trial. The beta oscillation of one electrode's LFP (gray trace) and its amplitude (black trace) in arbitrary units (au) are shown for three different example trials from animal Rs. **C.** Beta amplitudes from the electrode in **B** were averaged across all trials including all movement directions to estimate the trial-averaged beta amplitude aligned to movement onset (black trace indicates average amplitude, gray area indicates +- 2 standard errors of the mean (SEM)). **D.** LFPs were simultaneously recorded across multiple sites in MI during the experiment. **E.** Trial-averaged beta amplitude profiles for three different electrodes (color corresponding to the sites highlighted in **D**) are shown (mean +- 2 SEM). The time at which the beta amplitude passed an attenuation threshold (horizontal black dashed line) was estimated for every electrode. Note that beta activity on each electrode does not simultaneously pass the attenuation threshold (**E** inset). Data are from animal Rs.

The amplitude of beta oscillations attenuated around movement onset (Fig. A1.1B-C) as has been shown in many previous studies (Sanes and Donoghue 1993; Murthy and Fetz 1996a). This phenomenon is considered to be a mesoscopic reflection of local motor cortical activation and is often measured using macroscopic EEG electrodes in humans or intracortical electrodes in animals without reference to their spatial position within MI under the assumption that the entire cortical area is activated simultaneously. Instead, we considered whether MI is activated in a spatio-temporally organized fashion at movement onset by measuring the time at which the amplitude of beta band activity crossed an attenuation threshold across different sites in MI (Fig. A1.1D-E, see Methods for details). These trial-averaged beta attenuation times (BATs) varied systematically along a rostro-caudal axis, and were characterized using a linear regression model that predicted the BAT of an electrode from its spatial location on the array (Fig. A1.2A-B).

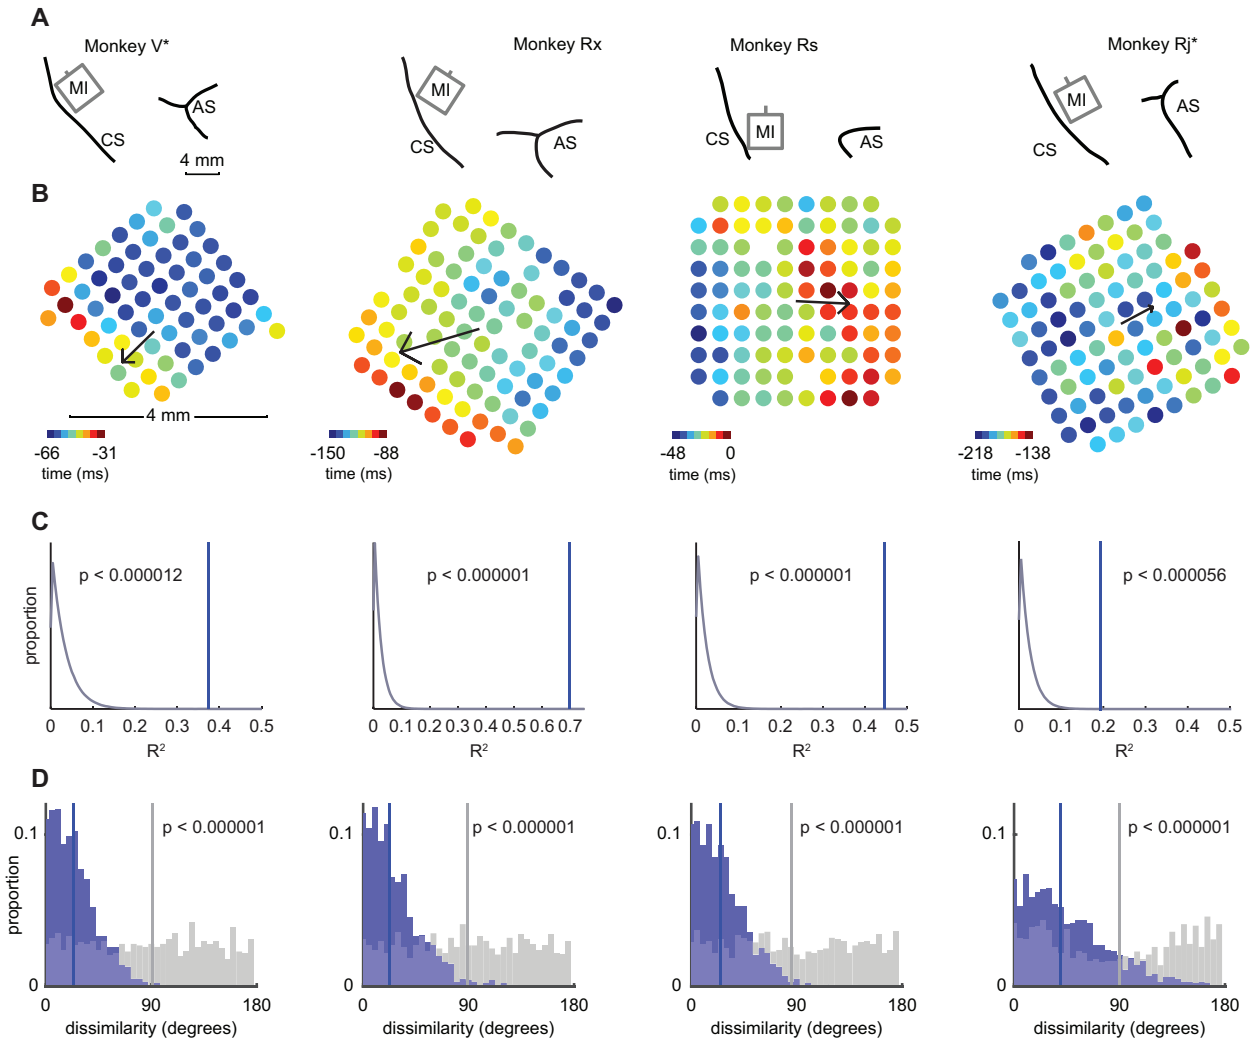

Figure A1.2 Array placement and beta attenuation maps. A. Line drawings depicting placement of the multielectrode arrays in the upper limb area of MI in four monkeys. The central and arcuate sulci are indicated (CS and AS, respectively). Data from monkeys V and Rj have been flipped vertically to normalize the location of the CS and AS across animals. **B.** Heat maps depicting the timings of attenuation in beta amplitude across all electrodes on each array relative to movement onset (time = 0). We fit a linear model to describe the spatial progression of beta attenuation times from earlier to later. This beta attenuation orientation (BAO) is indicated by a black arrow whose magnitude is proportional to the model's goodness of fit. **C.** We performed a spatial shuffle analysis to test for the significance of the BAO. The spatial location of BATs was shuffled one million times. For each shuffle, we computed the BAO of the shuffled data and its goodness of fit, R^2. The distribution of R^2 values for each monkey from shuffled data (gray curve) is shown, as well as the R^2 statistic observed in actual data (blue vertical line). The BAO was highly significant in all animals. **D.** The consistency of the BAO was assessed with a bootstrap analysis. For each iteration of the bootstrap (1000 total), we randomly partitioned each dataset into halves. A BAO was estimated from the data in each half, and the angle between them was measured. The distribution of angular differences (blue bars) was compared to a null distribution of

angular differences estimated from spatially shuffled data (gray bars) using the same shuffling procedure as **C**. We found strong evidence that the angular differences we observed between BAOs estimated on subsets of trials were significantly less than what would be expected under chance (Kolmogorov-Smirnov test, *p* indicated on figure).

BATs were well predicted by this model (R^2, Monkey V: 0.37, Rx: 0.70, Rs: 0.45, Rj: 0.30), and it was highly significant compared to a null model without spatial information (F-test, $p < 0.00001$ for all monkeys). The resultant vector of our regression model defined a beta attenuation orientation (BAO) corresponding to the gradient from earliest to latest BAT (arrows in Fig. A1.2B). As a further independent test for the significance of the BAO, we performed a shuffle analysis where we randomly shuffled the spatial position of each BAT and fit a planar regression model to the shuffled BATs. We repeated this shuffling procedure one million times and compared the model goodness of fit to the unshuffled data with that of the shuffled data (Fig. A1.2C).

The consistency of the BAO was established in three ways: across trials, across attenuation thresholds, and across frequencies. We assessed the consistency of the BAO across trials by randomly partitioning the data into two subsets and comparing the BAOs across the two halves (Fig. A1.2D). As a further test of the consistency of the BAO across trials, we computed movement-direction specific BAOs and found that they were oriented similarly (Fig. S1). We assessed the consistency of the BAO at different values of the attenuation threshold (ranging from 0.1 to 0.4 in steps of 0.05) using a circular dissimilarity test. In each animal, we found strong evidence that the BAOs were not significantly different across threshold values (V: $h_{21} = 0.10$; Rx: $h_{21} = 4.21$; Rs: $h_{21} = 0.23$; Rj: $h_{21} = 2.53$; all $p < 0.000001$). To address the consistency of the BAO across frequencies, we analyzed wide-band filtered beta band activity (15-30 Hz). We

computed the sample correlation, ρ, between BATs and wide-band BATs (wBATs). In each dataset a significant correlation was observed (V: ρ = 0.55, $p < 0.000003$; Rx: ρ = 0.42, $p < 0.00004$; Rs: ρ = 0.74, $p <$ 1e-17; Rj: ρ = 0.44, $p < 0.000007$). As a further test, we compared the BAO to the wide-band BAO (wBAO). We found that there was a significant correspondence between the BAO and wBAO using a circular dissimilarity test ($h_4 = 160.0$, $p < 0.026$).

In all animals, the BAO was organized along the rostro-caudal axis, although its precise directionality varied across animals. The BAO was rostro-caudal for monkeys V and Rx and caudo-rostral for monkeys Rs and Rj. One potential explanation for this difference in BAO is that the electrode arrays were placed at different positions medio-laterally along MI. For animals V and Rx, the array centers were 3 and 4 mm, respectively, medial to the genu of the arcuate sulcus, whereas for animals Rs and Rj, the array centers were exactly at the genu of the arcuate sulcus. Therefore, the BAO was oriented in the caudal direction for the arrays in a more medial position whereas the BAO was in the rostral direction for the arrays situated more laterally. As would be expected if attenuation reflects local activation of MI, BATs occurred before movement onset, and, their temporal range was approximately 50 ms from earliest to latest across the 4mm extent of the array, which we term the beta attenuation epoch.

If beta attenuation is truly a reflection of local network activation, then the spatial patterning in beta activity should also be present in the engagement order of single units across MI. The temporal dynamics of individual neurons, in MI, however, are quite heterogeneous across neurons prior to movement onset and firing rate modulation

times vary with different movement directions (Murphy, Wong, and Kwan 1985; Lecas et al. 1986; Scott 1997; Churchland and Shenoy 2007), so directly measuring activation times from firing rates is challenging (Fig. S2). Instead, we inferred modulation times of individual neurons by measuring when they were most informative about the upcoming movement direction (Fig. A1.3, see Methods for computational details).

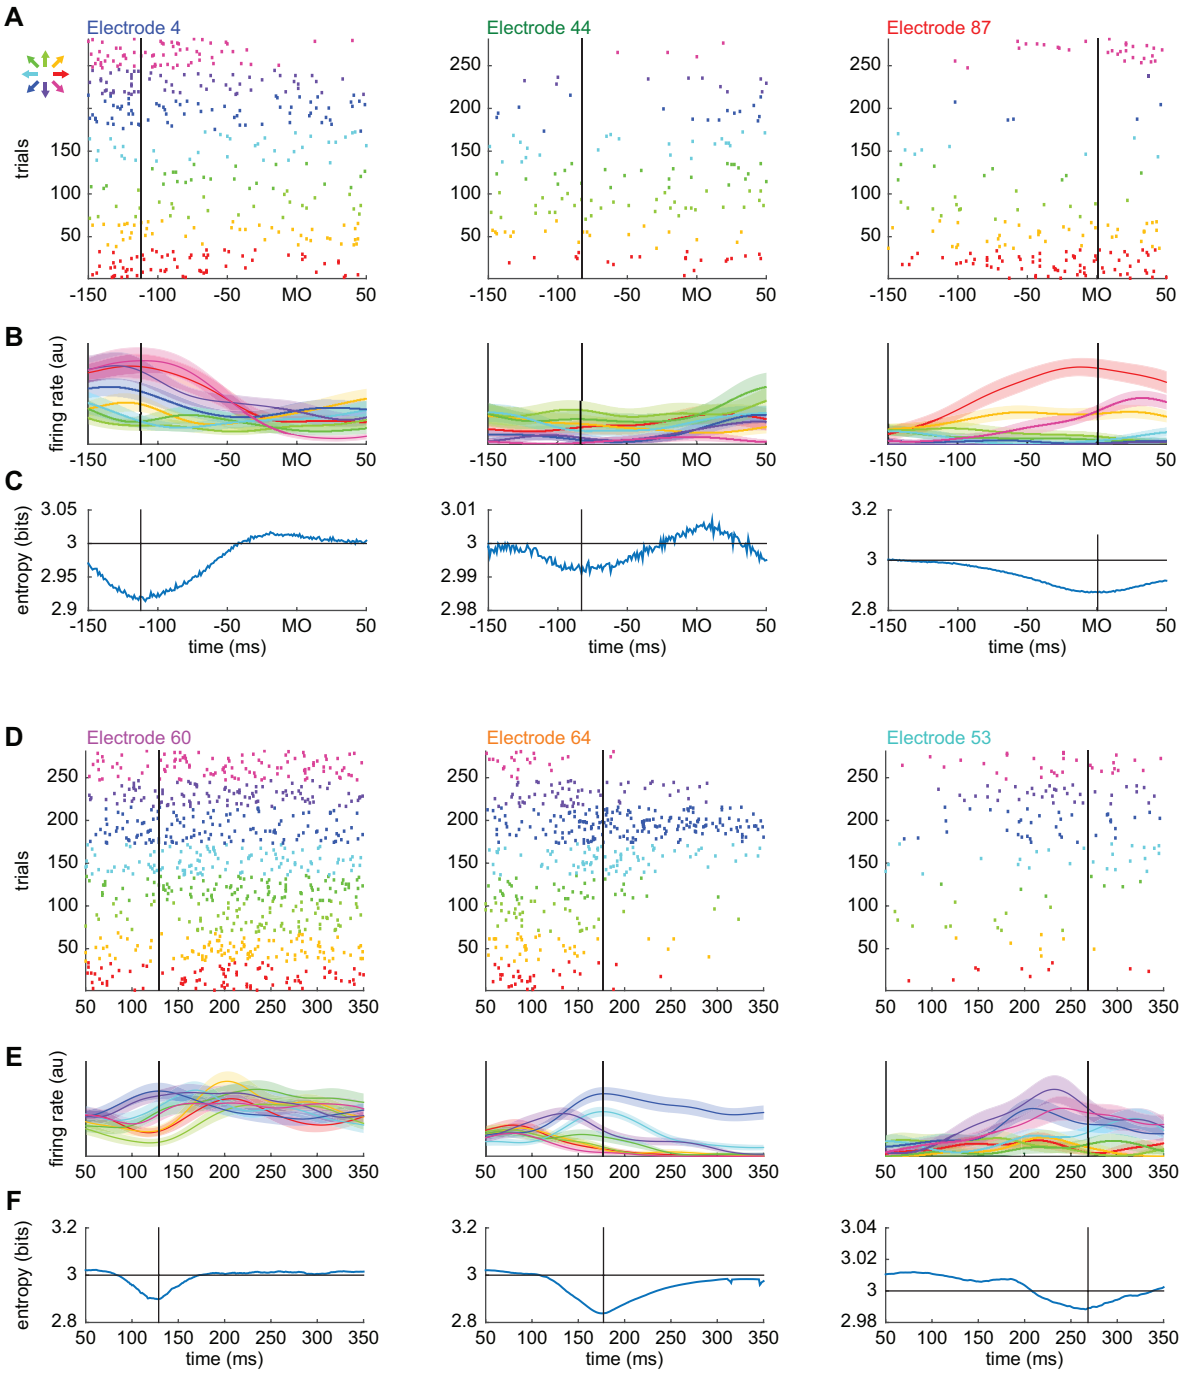

Figure A1.3 Unit spiking analysis. A. Raster plot of unit spiking activity on three electrodes corresponding to those in figure A1.1D. Each row in the raster indicates spiking activity on one trial aligned to movement onset. Color indicates the direction of movement. **B.** Unit spiking activity was convolved with a Gaussian kernel and averaged across trials. Trial averaged firing rates +- 1 SEM are shown for each movement direction. **C.** Time-resolved entropy of the target direction probability distribution conditioned on the firing rates in **B.** The dip in entropy reflects the directionally selective modulation of the neuron, and is used to estimate its modulation

time (vertical black line). Note that unit modulation times follow the same pattern as BATs in figure A1.1E. **D-F.** Three additional units that show significant modulation during movement preparation following the same conventions as **A**-**C** except data are aligned to the instruction cue. In both epochs, we have attempted to show units that are modulated early (left column), intermediately (center), or late (right) in their respective task epochs. During execution, the spatio-temporal sequence of unit modulation is consistent with the spatio-temporal sequence of beta attenuation; however, this consistency is absent during motor preparation.

Of the 82, 113, and 51 neurons with firing rates greater than 2 Hz in monkeys Rs, Rj, and V (no unit spiking data were available from animal Rx), we found that 63 (77%), 71 (63%), and 40 (78%) of neurons, respectively, exhibited significant direction-specific modulation during movement execution (Fig. A1.3A-C show three exemplar neurons recorded on the red, green, and blue electrodes from Fig. A1.1D-E), consistent with the previous observation that approximately 75% of MI neurons are directionally tuned (Georgopoulos et al. 1982). Significant direction-specific modulation was also observed during motor preparation (i.e. during the instructed-delay epoch); however, fewer cells were modulated (46 (56%), 38 (34%), and 13 (25%) for Rs, Rj, and V, respectively; Fig. A1.3D-F show three different exemplar neurons whose spatial position is indicated in Fig. A1.1D. Neurons were generally less informative about movement direction during motor preparation(KS test on entropy values, $p < 0.00005$ for all animals).

As a preliminary analysis to relate unit engagement order with mesoscopic patterning of beta activity, we computed the scalar projection of each unit's spatial location onto the BAO and compared that unit's position along the BAO to its modulation time (Fig. A1.4A).

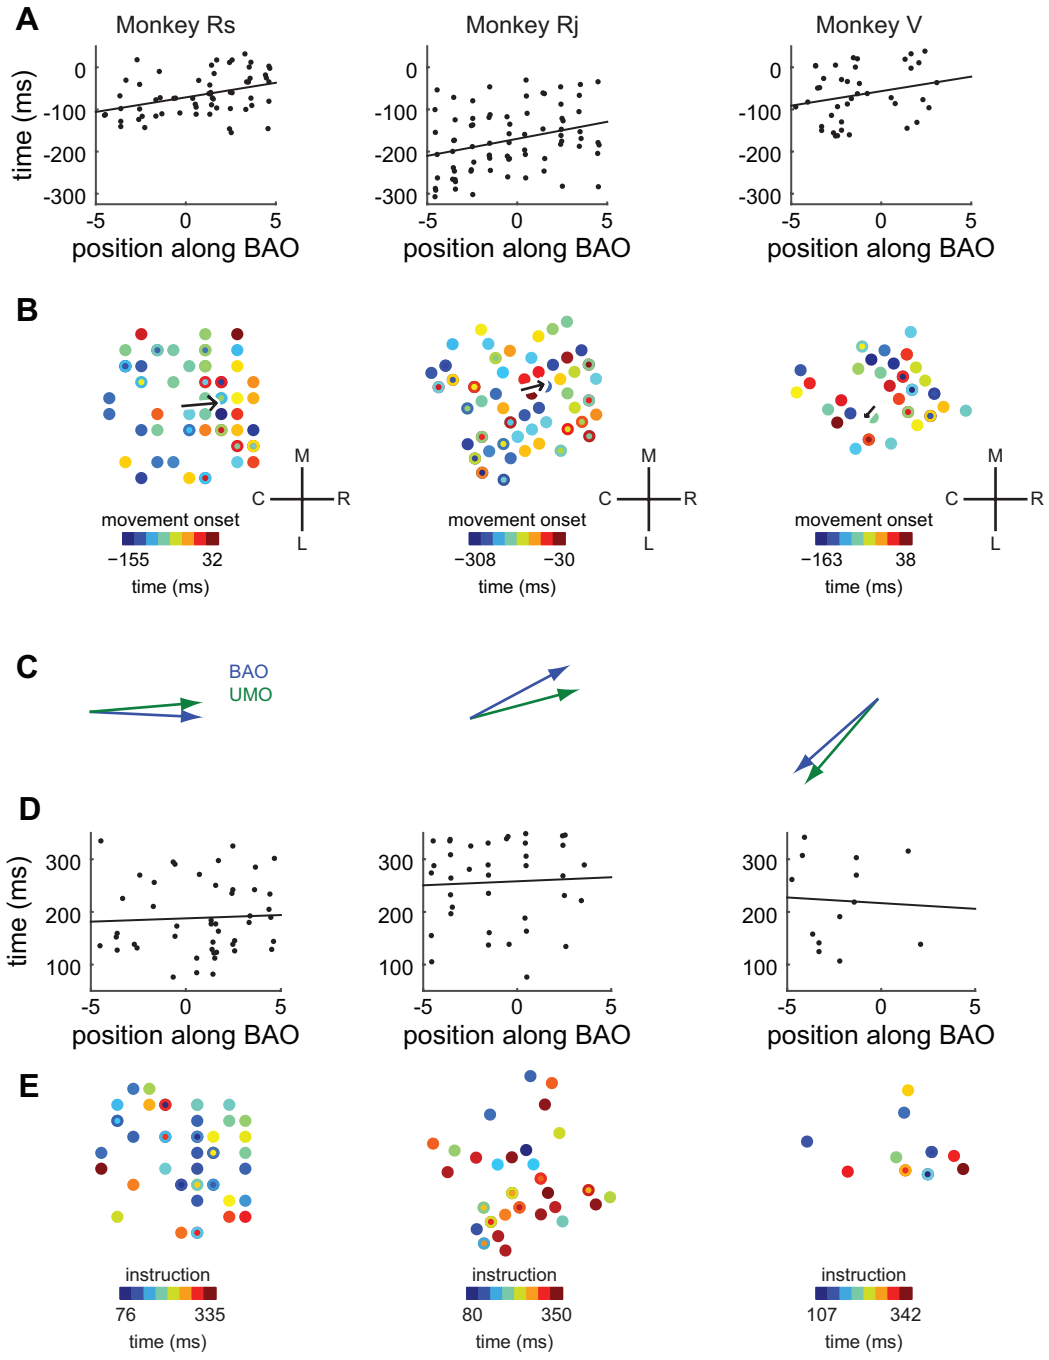

Figure A1.4 Relationship of beta attenuation to unit modulation. A. We performed an analysis to test whether units were sequentially engaged along the BAO. We compared the unit modulation time (UMT) of each cell to its position on the BAO and found a significant linear relationship across animals (see main text for statistics). Here, UMTs are relative to movement onset, i.e. movement onset occurs at $t = 0$. **B.** Raw heat maps of unit modulation times relative to movement onset as a function of their spatial location on the electrode array. The black arrow indicates the unit modulation orientation (UMO) for each dataset (length proportional to goodness of fit as in Fig. A1.2). Concentric circles indicate UMTs for multiple units on a given electrode.

Anatomical coordinates are the same as Fig. A1.2. **C.** We compared the similarity of the BAO and UMO in each animal and found that they were closely aligned during movement execution. **D-E.** Same as **A-B** except all unit activity is relative to the instruction cue. During preparation, there is no significant relationship between the BAO and unit modulation (**D**), and unit modulation times exhibit no significant spatial gradient (**E**).

Here, a unit's modulation time was operationally defined to be the moment when that cell was most informative about movement direction. We found that position along the BAO predicted unit modulation times during movement execution when data were pooled across animals (*F*-test on regression slope, $F_{1,175} = 18.6$, $p < 0.00003$). A similar result was found in two of the three animals individually, while the third showed a consistent, but insignificant trend (Rs: $F_{1,61} = 10.5$ $p < 0.002$; Rj: $F_{1,74} = 7.64$ $p < 0.007$; V: $F_{1,38} = 2.34$ $p < 0.13$). We also examined the relative timing of BATs and UMTs and found that UMTs significantly precede BATs (Fig. S3).

We then examined unit modulation times for direct evidence of spatial patterning and used the same regression procedure that was used to find the BAO to identify a unit modulation orientation (UMO). We found a spatial progression of unit modulation times across the cortical sheet along the rostro-caudal axis (Fig. A1.4B; UMO model statistics: Monkey Rs: $R^2 = 0.15$, $F_{2,60} = 5.27$, $p < 0.008$; Rj: $R^2 = 0.10$, $F_{2,73} = 3.94$, $p < 0.024$; V: $R^2 = 0.06$, $F_{2,37} = 1.16$, $p < 0.32$; additional controls see Fig. S4).

There appeared to be a close correspondence between the BAO and UMO during movement onset (Fig. A1.4C), so we developed a statistical test to quantitatively assess the similarity between these two vectors based on the angular difference between them. To apply a test to the group of animals as a whole, the angular differences between the BAO and UMO were added across animals to generate a test

statistic, *h*. We found strong evidence that the BAO and UMO were similarly oriented in the three animals ($h = 28.62$, $p < 0.0007$).

Directional modulation of single MI neurons is present not only during movement initiation but also during motor preparation, an epoch where overt movement does not occur. We compared these two task epochs to verify if the spatio-temporal patterning of unit modulation was a unique feature of movement initiation. During preparation, modulation timing was not linearly related to position along the BAO (Fig. A1.4D, $F_{1,94} = 0.13$, $p < 0.72$) and exhibited no significant spatial gradient (Fig. A1.4E, Rs: $F_{2,44} = 1.07$ $p < 0.35$; Rj: $F_{2,35} = 1.02$ $p < 0.37$; V: $F_{2,10} = 0.23$ $p < 0.80$), in contrast to the UMO that was identified at movement onset.

Discussion

Relationship of beta LFP to unit spiking activity

Beta attenuation has often been interpreted as a reflection of activated motor cortex (Pfurtscheller and Lopes da Silva 1999). Here, we have provided further support for this perspective by showing that the sequential engagement of unit spiking activity is aligned with the spatio-temporal progression of beta attenuation during movement onset. We determined the engagement order of cells by measuring when they were most informative about movement direction. Our findings do not, however, imply that directionally selective unit modulation is sufficient to cause beta attenuation. We found that cells in MI also exhibited significant directional modulation during motor preparation, and yet no beta attenuation occurred during this interval. In fact, beta oscillations are

highest in amplitude and propagate as travelling waves during motor planning (Rubino, Robbins, and Hatsopoulos 2006).

How, then, might directional modulation in cells give rise to beta attenuation during execution, but not during preparation? One possibility is that directional modulation is sufficiently gated during preparation to prevent beta attenuation, though recent studies have demonstrated this explanation is unlikely (Kaufman et al. 2014). An alternative explanation is that the behavior of MI neurons during preparation and execution is fundamentally different. In particular, it has been previously noted that a majority of MI neurons become entrained to high amplitude beta oscillations (Murthy and Fetz 1996b), i.e. during preparation. Yet, at movement onset (when beta attenuates), the temporal firing rate profiles (Churchland and Shenoy 2007) and tuning properties (Suminski et al. 2015) of these cells change, becoming heterogeneous and no longer phase locked to beta oscillations. Thus, beta attenuation represents the desynchronization of these cells around movement onset.

We observed that the temporal range of unit modulation times is substantially longer than the range of beta attenuation times. This finding is not necessarily inconsistent with our interpretation of beta attenuation. In previous studies, it has been shown that modulation times of MI neurons span several hundred milliseconds in reaching tasks (Murphy, Wong, and Kwan 1985; Lecas et al. 1986), and more generally, firing rate profiles are extremely heterogeneous (Churchland and Shenoy 2007). We have shown that the spatial location of a unit is one important factor for determining when that unit will become modulated, however, other factors are also involved. We

speculate that at least some of these other factors are intrinsic to the neuron and not shared amongst neighboring cells. The beta LFP, being an aggregate signal, then, will average away this intrinsic component. Thus, we might expect that if we could record from tens to hundreds of cells on an electrode the average of their UMTs would converge to the BAT for that electrode. In this view, the temporal range of UMTs is larger than the range of BATs because the UMTs are inherently more variable. Nevertheless, the relationship between BATs and UMTs remains tentative, and represents an opportunity for further inquiry. Specifically, a better understanding of the mechanistic link between beta attenuation and unit modulation may explain why the duration of UMTs is longer than BATs.

Several of the unit spiking results that were significant in animals Rj and Rs were weaker in animal V. However, all of animal V's results were, at minimum, qualitatively consistent with the other two animals. In this animal, we recorded from only 64 electrodes on the Utah array distributed over a smaller spatial area than the other two animals. Accordingly, we speculate that unit spiking results in animal V were weaker because we sampled fewer units over a smaller area than the other two animals.

Spatio-temporal patterning of motor cortical activity

Considering the motor cortex is a physical neural substrate with intrinsic connectivity, it is perhaps not surprising that the regular temporal dynamics within neural state space (Churchland et al. 2012) are manifested in specific spatio-temporal patterns on the cortical surface. And yet, any functional spatial structure within MI has remained elusive. Unlike most primary sensory cortical areas that exhibit clear topographic

organization, the motor cortex reveals highly distributed responses over a large area (M. H. Schieber and Hibbard 1993; Sanes et al. 1995; Rubino, Robbins, and Hatsopoulos 2006; Mollazadeh et al. 2011; Riehle et al. 2013; Peters, Chen, and Komiyama 2014). For example, whereas the somatosensory cortex possesses distinct representations of the segments of the upper limb, the somatotopic organization of upper limb MI is rather crude with highly overlapping representations, particularly between the elbow and shoulder (Kwan et al. 1978; Park, Belhaj-Saïf, and Cheney 2004). And yet, we have shown that a regular, spatio-temporal pattern of neural activity accompanies movement onset.

Previously, spatio-temporal patterning in MI was interpreted as a cortical correlate of the proximal-to-distal sequence of muscle activation during prehension (Murphy, Wong, and Kwan 1985; Riehle et al. 2013). Such an interpretation, however, is unlikely to explain the spatio-temporal phenomena we observed in this study. The planar reaching movements used in this study only involved arm movements of the shoulder and elbow which have been shown to evoke a diversity of muscle activation patterns including proximal-to-distal and distal-to-proximal patterns (Karst and Hasan 1991; Scott 1997). And yet, we observed that the BAO was relatively consistent across all movement directions. Moreover, there is no clear delineation between shoulder and elbow representations on the cortical surface, although the extent to which they overlap remains the subject of some debate (Marc H. Schieber 2001). This does not, however, imply that the spatio-temporal patterning we have documented here is independent of any putative somatotopic map. Indeed, we speculate that the precise orientation of the

BAO may be influenced by the local somatotopy underneath the Utah array. The BAO was oriented in the caudal direction for the arrays placed in a more medial position (monkeys V and Rx) whereas the BAO was in the rostral direction for the arrays placed more laterally (monkeys Rs and Rj). Several papers have suggested a horseshoe somatotopic organization where the proximal sites in MI form a horseshoe pattern surrounding a central distal core (Park et al. 2001; Park, Belhaj-Saïf, and Cheney 2004). We speculate that the arrays in V and Rx may be situated in the medial arm of the horseshoe whereas the arrays in Rs and Rj are located in the lateral arm, and this differential positioning with respect to the underlying somatotopic representation drives the different orientations of the BAO. Further research will be needed to explain why different beta attenuation orientations were observed across animals.

Rather than interpreting spatio-temporal activity patterns on the motor cortical surface as correlates of sequential muscle activation, we suggest that these patterns around movement onset reveal that MI performs a spatially distributed computation to initiate a movement. Inputs to MI from other cortical areas and subcortical structures, such as the cerebellum and basal ganglia, are presumably responsible for activating the motor cortex during movement initiation (Donoghue and Sanes 1994); yet, how these inputs interact with the intrinsic spatio-temporal activity patterns in MI is unknown (Khanna and Carmena 2015). We speculate that inputs to motor cortex are spatially distributed across large regions of MI and integrated via the dense network of horizontal connectivity (Gatter and Powell 1978; Huntley and Jones 1991). Using an information theoretic analysis, it has been previously shown that single units exhibit a

preponderance of directed, functional connections along the rostro-caudal axis after the onset of a visual cue to initiate movement (Quinn et al. 2011). Moreover, it has recently been demonstrated that the sequential spiking activity of pairs of functionally connected neurons distributed along this axis contains more information about movement direction (Takahashi et al. 2015). Thus, the spatio-temporal dynamics documented here may be supported, in part, by these functional connections aligned along the rostro-caudal axis.

The spatio-temporal patterning that we observed during movement initiation could serve to optimally drive downstream targets of MI, and specifically, muscles. Transcranial magnetic stimulation (TMS) experiments in humans have documented that motor evoked potentials (MEPs) in the muscles are differentially affected depending on the direction of current delivered across the motor cortex. In particular, the MEP amplitude is larger when the induced cortical currents occur along the anterior-posterior axis (Brasil-Neto et al. 1992). Moreover studies have observed differential MEP thresholds, latencies, and amplitudes when current flows in the anterior-posterior direction versus posterior-anterior direction (Brasil-Neto et al. 1992; Kammer et al. 2001; Sommer et al. 2006; Jung et al. 2012). Thus, the natural rostro-caudal patterns we observed in motor cortical activity may explain why particular TMS current orientations are more effective to drive the motor periphery.

In summary, we found that movement initiation is characterized by a unique spatio-temporal sequence of neural activity, and that this sequence is evident in both LFPs and unit spiking. From these findings, we introduced a new hypothesis about the nature of the motor cortical activity required to initiate movements. Future experiments

using either electrical or optogenetic stimulation to perturb both the spatial and temporal dynamics of MI activity will be required to further test this hypothesis.

Appendix 2 - Comparing offline decoding performance in physiologically defined neuronal classes

This scientific content of this chapter was accepted for publication at The Journal of Neural Engineering and was in press at the time this thesis was submitted.

Abstract

Objective: Recently, several studies have documented the presence of a bimodal distribution of spike waveform widths in primary motor cortex. Although narrow and wide spiking neurons, corresponding to the two modes of the distribution, exhibit different response properties, it remains unknown if these differences give rise to differential decoding performance between these two classes of cells.

Approach: We used a Gaussian mixture model to classify neurons into narrow and wide physiological classes. Using similar-size, random samples of neurons from these two physiological classes, we trained offline decoding models to predict a variety of movement features. We compared offline decoding performance between these two physiologically defined populations of cells.

Main results: We found that narrow spiking neural ensembles decode motor parameters better than wide spiking neural ensembles including kinematics, kinetics, and muscle activity.

Significance: These findings suggest that the utility of neural ensembles in brain machine interfaces may be predicted from their spike waveform widths.

Introduction

Neural interface systems for control have recently made a number of important advances in recording capabilities, decoding algorithms, and output devices (Vargas-

Irwin et al. 2010; Schwartz 2004). In particular, these systems have seen nearly a doubling of simultaneously recorded neurons every seven years using either high density electrode arrays, or more recently, optical calcium fluorescence imaging (Clancy et al. 2014). These advances have provided an ever growing set of rich, high-dimensional signals for control (Stevenson and Kording 2011). And yet, decoding ability has not increased correspondingly with the growth of input signals, but rather has plateaued. This disparity arises, in part, because high-dimensional input signals require larger models to relate neural activity to motor features. These models are harder to train for a given size of data, are more prone to overfitting, and thereby less generalizable (Chen 2013; Wu et al. 2006; Malik et al. 2014). As such, there is some debate about the optimal size of decoding models (Carmena et al. 2003).

One approach to reduce the dimensionality of neural feature space is to take advantage of the fact that the state space of neural activity patterns is much smaller than the full dimensionality of the neural features being analyzed (Churchland et al. 2012; Yu et al. 2009; Sadtler et al. 2014). That is, the responses of individual neurons are correlated, and the number of latent dimensions needed to explain the variability in the ensemble activity is less than the total number of recorded neurons. Indeed, recent reports have shown that the activity of moderately large neural ensembles (tens to hundreds of cells) can be described by a few orthogonal latent dimensions in neural state space (Yu et al. 2009). Within this framework, these latent dimensions may be used as control axes for a prosthetic device (Sadtler et al. 2014; Kao et al. 2013).

Another approach to constrain the dimensionality of neural feature space is to perform model selection using some statistical selection criterion (Malik et al. 2014). The guiding principle of this approach is to identify relevant (i.e. predictive) neural features by fitting a model on training data. Several different criteria have been proposed to rank feature relevancy including correlation coefficient (Sanchez et al. 2004; Bansal et al. 2012), mutual information (Paninski et al. 2004), and decoding accuracy (Singhal et al. 2010; Westwick et al. 2006), while additional criteria, such as Akaike information criterion (Kang et al. 2015), Bayesian information criterion (Brockmeier and Príncipe 2013; Wu et al. 2006), and automatic relevance determination (Chen and Takahashi 2013; Chen 2013) are used to determine the optimal number of features to include in a model.

Instead of approaching model selection as a dimensionality reduction or statistical problem, we wondered if physiological properties of motor cortical neurons could guide the choice of which neurons to use in a decoder. One physiological property that has received much attention recently is spike waveform width. In particular, primary motor cortex (MI) exhibits a bimodal distribution of extracellular, spike waveform widths (Kaufman, Churchland, and Shenoy 2013; Vigneswaran, Kraskov, and Lemon 2011; Merchant, Naselaris, and Georgopoulos 2008). Although extracellular recordings cannot directly determine cell type or morphology, these two modes exhibit different physiological and functional properties (Takahashi et al. 2015). That study observed that population responses among narrow spiking neurons exhibited more pronounced oscillatory power in the beta frequency band (15-35 Hz) as was observed in

the local field potential. Moreover, during peri-movement epochs, narrow, unlike wide spiking neurons, spatially coordinate their firing activity in a way that is consistent with both the wave propagation observed in the local field potential oscillations and the underlying anatomical horizontal connectivity (Gatter and Powell 1978; Takahashi et al. 2015). We reasoned that the narrow spiking neural population may be more closely related to movement because narrow spiking neurons form a network of functional connections that is aligned with the spatial pattern associated with movement onset. Thus, we hypothesized that narrow spiking neurons would lead to more accurate movement decoding as compared to wide spiking neurons.

Material and Methods

Neurophysiology

All surgical and experimental procedures were approved by either the University of Chicago, or Northwestern University Animal Care and Use Committees, and conformed to the principles outlined in the Guide for the Care and Use of Laboratory Animals (NIH publication no 86-23, revised 1985). Five rhesus macaques (*macaca mulatta*) were implanted with 96 channel Utah electrode arrays in the upper limb area of motor cortex contralateral to their working arm (for details about the exact placement of the electrode arrays, see (Suminski, Tkach, and Hatsopoulos 2009; Suminski et al. 2015)). Neural signals were collected from these arrays using a Cerebus neural data acquisition system (Blackrock Microsystems, Salt Lake City, UT). Unit spiking activity was sorted offline using semi-manual spike sorting software (Offline Sorter, Plexon Inc., Dallas, TX).

Behavioral tasks

This experiment consisted of three different tasks involving movement of the upper limb. In the first task, two rhesus macaques were trained to play an instructed-delay, center-out reaching task (for a description of the task see (Suminski et al. 2015)). Briefly, animals were trained to control the position of a cursor using a 2-link robotic exoskeleton (BKIN Technologies, Kingston, Ontario, CA). The position of the cursor was projected directly above the position of the animal's hand. A trial began when the animal moved the cursor to a center target and maintained it there for 500 ms. After that time, the animal was cued to move to one of 8 possible peripheral targets positioned radially around the center target, then had to wait 1000 ms until a go cue appeared. At this point, the animal was free to move from the center target to the peripheral one. Upon hitting the peripheral target, the animal had to hold the cursor at the peripheral target for 500 ms to complete the trial successfully. Fluid reinforcement was delivered on each successful trial.

Two additional rhesus macaques were trained to play a random target pursuit (RTP) task in multiple experimental conditions (for a detailed description of the task and experimental design see (Suminski, Tkach, and Hatsopoulos 2009)). Briefly, animals used the same robotic exoskeleton to make planar reaching movements to square targets randomly distributed within a 10 x 6 cm workspace. Every time the cursor hit the target, a new target appeared at a random location. In order to complete a trial successfully, an animal had to sequentially hit seven targets. Failure to hit a target

within 5 s of its appearance resulted in an aborted trial. Fluid reinforcement was delivered for every successful trial.

One rhesus macaque was trained to perform an isometric wrist flexion task. The upper arm was constrained largely to a para-sagittal plane with the elbow at a 90° angle and the forearm horizontal, in an orientation midway between supinated and pronated. The monkey's wrist was maintained in line with its forearm by securing its hand in a box, which was custom-fit with padding to minimize movement. A six-degree of freedom torque cell was mounted on the box, such that the axes of measurement aligned with those of the wrist. Cursor movement was proportional to the force along the flexion-extension and radial-ulnar deviation axes. The task required the monkey to move the cursor from a central target to one of eight peripheral targets separated by 45°. The force targets were set for each monkey to be submaximal (approximately 20-30% MVC) in order to reduce fatigue. To initiate a trial, the monkey held the cursor in the central target (requiring no force) for 0.5 seconds, after which a randomly selected outer target appeared. The monkey was required to move the cursor to the outer target within 5 seconds, and to maintain that force for 0.5 seconds in order to receive fluid reinforcement.

Classification of narrow and wide spiking neurons

We classified units into two discrete physiological classes based on their spike waveform width. To determine the width of each sorted unit, we measured the difference in time between the peak and trough of the average waveform. An additional quantity, waveform signal-to-noise ratio (SNR), was defined as the magnitude of the

peak minus the trough in the average waveform divided by the average standard deviation of the waveform across time (Suminski, Tkach, and Hatsopoulos 2009). Only units with SNRs greater than 3 were used in the subsequent analyses.

For each dataset, a Gaussian mixture model (Bishop 2006) was used to classify spike waveforms into narrow and wide categories. Mathematically, the Gaussian mixture model attempts to describe the distribution of spike waveform widths as a sum of *K* Gaussian distributions. Each Gaussian in the mixture model is referred to as a component (indexed with the variable, *k*), and is fit with a unique mean and standard deviation, μ_k, and σ_k, respectively. Each component also has an additional parameter, π_k, representing the proportion of data described by that component. Expressed as an equation, this model may be specified as

$$p(w) = \sum_{k=1}^{K} \pi_k N(w|\mu_k, \sigma_k),$$

where $p(w)$ is the probability of observing a spike waveform width, *w*, and $N(w|\mu_k, \sigma_k)$ indicates a Gaussian distribution with mean, μ_k, and standard deviation, σ_k. This model included an additional regularization parameter, λ, that was added to each σ_k to ensure that σ_k remained strictly positive for every component (see (Bishop 2006) for a more complete treatment on fitting Gaussian mixture models. Matlab function fitgmdist, The Mathworks, Natick, MA).

To confirm that the spike waveform width distributions were bimodal, we varied the number of components, *K*, in the mixture model, and computed the Akaike Information criterion (AIC), a goodness of fit statistic for each model (Akaike 1974). As we increased *K*, we also increased λ proportionately to ensure that each additional

component was non-degenerate. A chi-square test of homogeneity was used to compare the proportion of narrow and wide neurons across recording sessions in a given animal (Rice 1995).

Computing other response properties of cells

In addition to determining the waveform width of each cell, we also measured its average firing rate, and, for center-out datasets, the preferred direction and tuning strength. Average firing rate was determined by dividing the spike counts of each cell by the duration of the recording. Firing rate variance was computed using the following formula:

$$var := \frac{1}{n-1}\sum_{b=1}^{n}(y_b - \bar{y})^2$$

where n is the number of 50 ms bins, y_b is the spike count in bin b, and $\bar{y}$ is the average spike count over all bins. To determine preferred direction and tuning strength, we fit a cosine-tuning model of the form:

$$y_i = \alpha + \beta\cos(x_i - \phi) + \epsilon$$

where y_i indicates the number of spikes between the go cue and target hit on trial i, α is the overall firing rate of the cell, β is the gain of the cosine tuning model, x_i is the angular location of the peripheral target on trial i, ϕ is the preferred direction of the cell, and ϵ is a normally distributed error term. This model was fit using the Matlab function lsqcurvefit. The tuning strength of the cell was defined as the proportion of variance in spike counts explained by this tuning model.

Decoding analysis

Input features

Spiking activity from every neuron was binned into 50 ms bins. Only neurons with firing rates > 1 Hz and waveform SNR > 3 were used in subsequent analyses. The number of neurons that satisfied these criteria is listed in table 1. In general, the spike counts of each neuron in the preceding 20 time bins (i.e. 20 filter taps, 1 second of history) were used as input features to the decoding model, however, we varied the number of taps between 4 and 32 in one analysis to explore the effect of the number of taps on decoding performance (Fig. A1.4). In total, the input dimensionality to the decoding model was equal to the number of neurons multiplied by the number of taps (which was 20, unless otherwise noted).

Output features

Several different motor related quantities were decoded including kinematic and kinetic features as well as muscle activity. Output features were decoded in 50 ms bins. In center-out datasets, we decoded shoulder and elbow (joint) torque (computed as described in (Fagg et al. 2009)), joint angular velocities, Cartesian *x* and *y* velocities of the cursor, and wrist speed. In the isometric wrist dataset, j141203, we decoded the activity of 11 muscles of the forearm and hand including extensor digitorum communis (EDC), adductor pollicis longus (APL), flexor digitorum profundis (FDP), extensor carpi radialis (ECR), extensor digitorum communis 2 (EDC2), brachioradialis (Brad), pronator teres (PT), flexor carpi ulnaris (FCU), flexor digitorum superficialis (FDS), flexor carpi radialis (FCR), and flexor digitorum superficialis 2 (FDS2).

Decoding model

All computations were carried out offline in the Matlab programming environment. We employed a standard causal Wiener filter model to decode movement related quantities from neural activity (Carmena et al. 2003; Serruya et al. 2002; Warland, Reinagel, and Meister 1997; Sanchez et al. 2002; Brockmeier and Príncipe 2013). Mathematically, this model satisfies the following objective:

$$\underset{\beta}{\operatorname{argmin}} \sum_{t=1}^{T} \left\| y_t - \alpha - \sum_{u=0}^{19} x_{t-u}\beta_u \right\|_2^2$$

where y_t is a motor quantity at time bin t, x_{t-u} is a vector of spike counts corresponding to time bin $t - u$, α is an intercept term, β_u is a vector containing the coefficient of each filter tap at time lag, u, and $\|\cdot\|_2$ denotes the ℓ_2 norm. This model is arguably one of the simplest neural decoding models, yet it has been widely used and has been shown to achieve a high degree of decoding accuracy (Carmena et al. 2003; Serruya et al. 2002; Brockmeier and Príncipe 2013). Model goodness of fit was quantified using the coefficient of determination, R^2 given by the following formula:

$$R^2 := 1 - \frac{\sum_t (y_t - \hat{y}_t)^2}{\sum_t (y_t - \bar{y})^2}$$

where y_t, $\hat{y}_t$, and $\bar{y}$ denote the observed motor quantity, the fitted motor quantity, and the time averaged motor quantity, respectively.

Models were trained on 75% of available data and tested on the remaining 25%. We found that the proportion of data allocated to training and test sets did not have an appreciable effect on subsequent analyses.

Bootstrap analysis

In order to compare decoding performance between narrow and wide spiking neural populations, we drew random ensembles of *N* neurons from each population repeatedly (100 times, with replacement), trained linear decoding models, and measured the decoding performance on a separate set of test data. This process was applied to each dataset individually. The number of neurons in the ensemble, *N*, was varied systematically to quantify how decoding performance scaled in each population.

Matching procedure

In order to control for underlying differences in response properties between narrow and wide spiking neural populations, we developed a greedy matching algorithm to select neurons with similar response properties. We chose an ensemble of *N* wide spiking neurons completely at random. Then, for each selected wide spiking neuron, we found the narrow spiking neuron whose response property (e.g. firing rate) was closest to the wide spiking neuron, where closeness was defined by a distance metric (described below for each feature). If the narrow spiking neuron that was closest to the current wide spiking cell was already matched to another wide spiking cell, the next closest unmatched narrow spiking neuron was matched to that wide spiking neuron. We matched several underlying response properties including firing rate, waveform SNR, preferred direction, and tuning strength. For firing rate, waveform SNR, and tuning strength, we used the absolute value of the difference as our distance metric (i.e. the Euclidean distance). For a circular variable like preferred direction, we used the absolute value of the angle between preferred directions as our distance metric.

Results

We recorded spiking activity from single units in primary motor cortex (MI) while monkeys engaged in a variety of tasks involving the upper limb. We computed the spike width of each sorted unit (Fig. A2.1) and classified it as either narrow or wide.

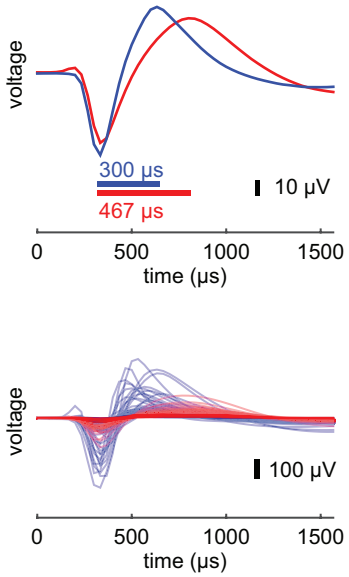

Figure A2.1 Quantifying spike waveform width. For each sorted unit, we computed its spike waveform width. Here, waveform width is defined as the difference in time between the peak and trough of the average waveform. Exemplary narrow (blue) and wide (red) waveforms (averaged over spikes) are shown as well as the time from trough to peak (top). The distribution of all recorded waveforms from dataset rs050225 (bottom). Color indicates either narrow or wide waveform width.

We then compared offline decoding performance of these two classes of cells across many different tasks. A preliminary version of these results was presented as a conference proceeding (Best, Takahashi, and Hatsopoulos 2015).

Narrow and wide spiking neural ensembles

We used a Gaussian mixture model to classify neurons as either narrow or wide spiking based on their spike waveform widths (Fig. A2.2).

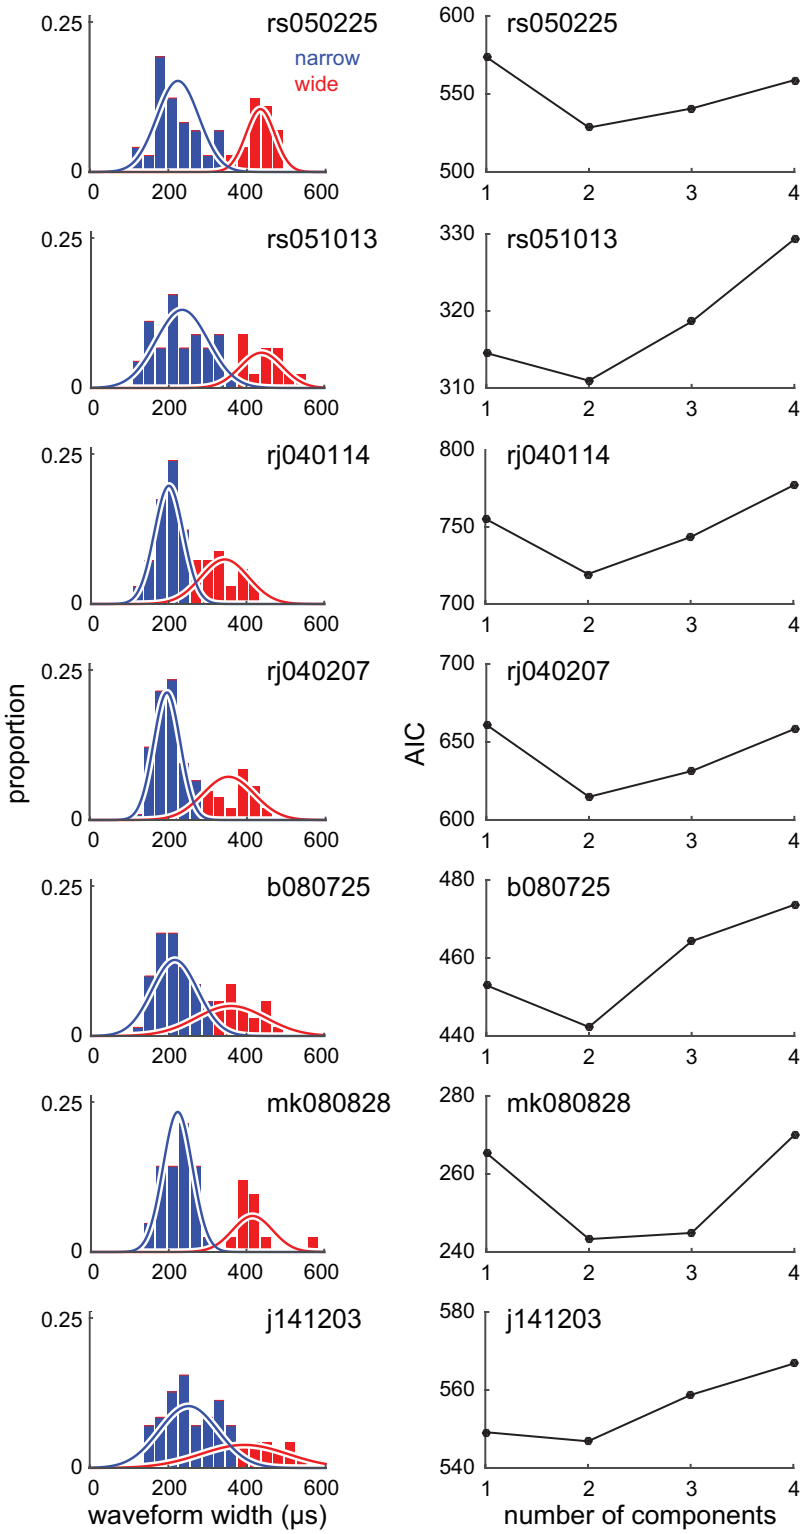

Figure A2.2 Bimodal distribution of spike waveform widths. A Gaussian mixture model was used to partition neurons from each dataset into narrow and wide spiking categories based on waveform width. To verify that each waveform distribution was indeed bimodal, we systematically varied the number of Gaussians in the mixture model and computed the Akaike information criterion (AIC) to perform model selection. For

each dataset, we found that a mixture model containing two components best described the data.

A separate model was trained on each recorded dataset (summary statistics of each model given in table 1).

Dataset	Task	# data points	# training points	# test points	# narrow	# neurons	μ_1	σ_1	μ_2	σ_2	π	λ
rs050225	center-out	4233	3175	1058	71	105	224	53	434	37	0.68	0
rs051013	center-out	4423	3317	1106	42	57	235	67	438	54	0.73	0
rj040114	center-out	2374	1781	593	98	163	200	38	344	62	0.62	0
rj040207	center-out	5095	3821	1274	88	130	201	38	367	57	0.65	0
b080725	RTP	4631	3473	1158	62	87	221	62	383	78	0.70	1
mk080828	RTP	6225	4668	1557	38	51	223	38	416	52	0.74	0
j141203	wrist	23982	17987	5995	58	75	253	76	401	109	0.67	1

Table 1 Summary of datasets Details regarding task, dataset size, number of neurons, and fit parameters for the Gaussian mixture model. All times are listed in microseconds.

To verify that each distribution was indeed bimodal, we fit additional Gaussian mixture models with varied numbers of Gaussian components. For all datasets, we found that a two-component model (i.e. a bimodal distribution) had optimal AIC values. The mean spatial locations of narrow and wide spiking neurons across the cortical sheet were not significantly different (Bonferroni corrected Hotelling's T^2 test) suggesting that subsequent decoding results are not due to differences in the location of the neurons on the cortical sheet.

We examined the consistency of the bimodal distribution across time. In monkeys Rs and Rj, we analyzed datasets that were collected 230 and 24 days apart, respectively. We performed a Chi-square test of homogeneity to assess whether the proportion of narrow spiking units was the same across datasets. We found no evidence of a significant difference in the proportion of narrow spiking units ($p < 0.42$ and $p < 0.18$ for animals Rs and Rj, respectively) across time.

The previous statistical test ensured that the relative proportion of narrow and wide spiking units was the same across time; however, we did not directly gauge whether the average waveform width of each population was similar across time. Accordingly, we performed a *t*-test on the average waveform width of each class across time. In both animals, we found no evidence to suggest that the average waveform width of the narrow spiking class was significantly different across time (Rs: $t_{49} = -0.82$, $p < 0.21$, Rj: $t_{182} = -0.18$, $p < 0.43$). With respect to wide spiking neurons, animal Rs showed no significant difference across time ($t_{20} = -0.26$, $p < 0.40$), although there was a significant difference in Rj ($t_{93} = -1.97$, $p < 0.026$).

Decoding kinetics and kinematics

We built simple linear decoding models to predict a variety of kinematic and kinetic motor features based on the activity of either narrow or wide spiking neural ensembles. For our initial analysis, we considered neural data that were collected while animals were performing an instructed-delay, center-out reaching task. We found that narrow spiking neural ensembles outperformed wide spiking neural ensembles at a

variety of different ensemble sizes (Fig. A2.3, see methods for details about model training and validation).

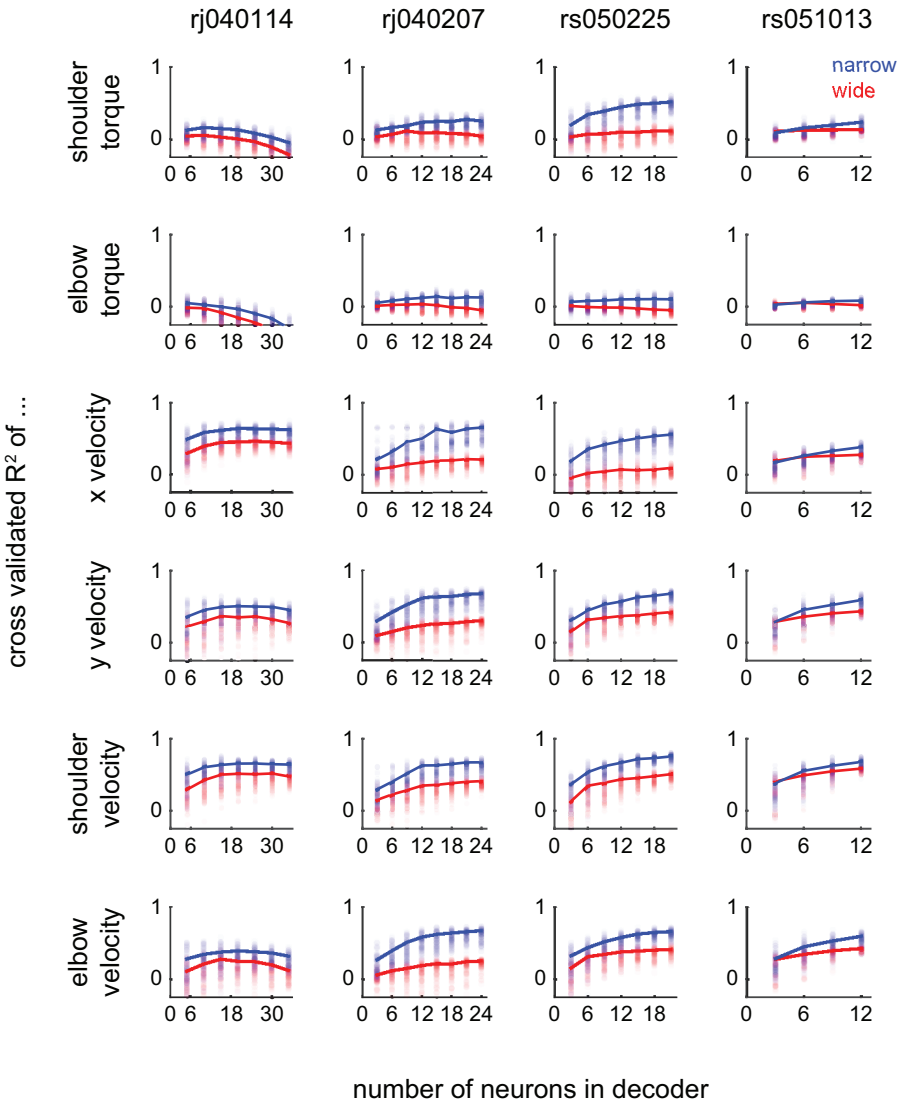

Figure A2.3 Decoding performance using narrow and wide spiking neural ensembles. We used a standard 20-tap causal Wiener filter to decode kinematic and kinetic quantities from neural data while two animals performed an instructed-delay, center-out task. We repeatedly (100 times) drew random samples of either narrow or wide spiking neurons, trained a decoding model, and then tested its performance on a separate set of data. We found that narrow spiking neural ensembles significantly outperformed wide spiking neural ensembles in a variety of coordinate frames (see text for summary statistics). Each column indicates a different dataset. Individual points correspond to each of the 100 random samples, while the solid lines indicate the upper 75th percentile of decoding performance.

We performed a two-way ANOVA using waveform class (i.e. narrow or wide) and ensemble size as factors. We observed a highly significant main effect of waveform class on decoding performance for each motor feature (median improvement in R^2 was 0.15 across datasets/motor features; $p < 1e\text{-}8$ for every dataset/motor feature combination, Bonferroni adjusted for multiple comparisons).

One potential explanation for the difference in decoding performance across waveform classes is that the optimal number of taps for each waveform class could differ. To test this possibility, we fixed the number of neurons in the decoder and systematically varied the number of filter taps from four (200 ms of history) up to 32 (1600 ms of history). Again we observed that narrow spiking neurons outperformed wide spiking neurons irrespective of the number of taps in the model (Fig. A2.4A, ANOVA, $F_{1,784} = 139.81$, $p < 1e\text{-}8$, $F_{1,784} = 407.77$, $p < 1e\text{-}8$, for x and y velocity, respectively), or model regularization (Fig. A2.4B, ANOVA, $F_{1,784} = 628.90$, $p < 1e\text{-}8$, $F_{1,784} = 270.81$, $p < 1e\text{-}8$, for x and y velocity, respectively).

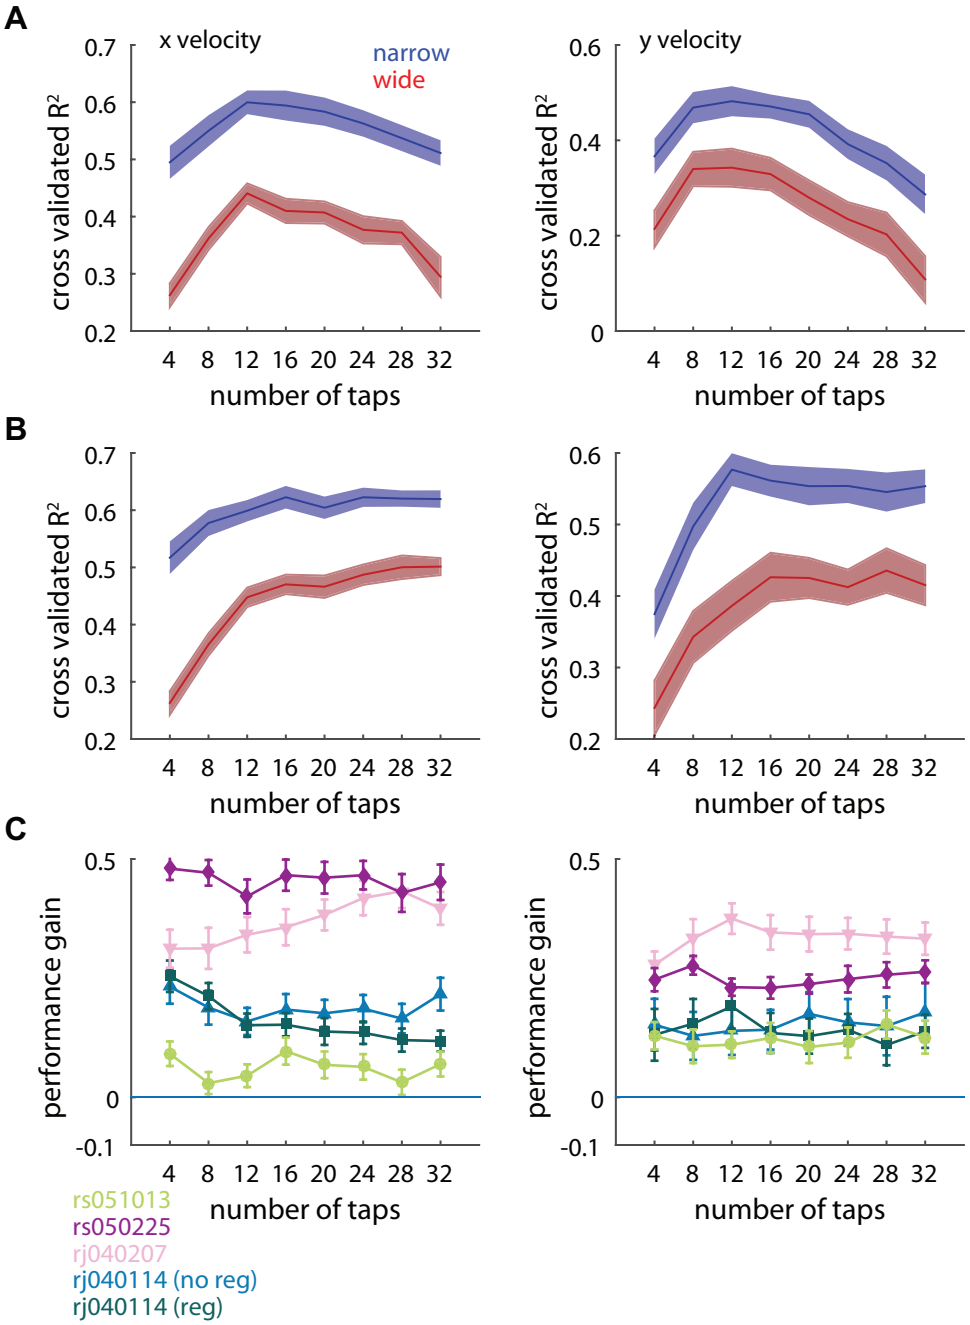

Figure A2.4 The number of taps does not explain the difference in decoding performance. A. We fit a linear decoding model containing 20 narrow or wide spiking neurons and systematically varied the number of filter taps. We observed that narrow spiking neurons could predict x and y velocities (left and right columns, respectively) better than wide spiking populations irrespective of the number of taps. Data shown are from one dataset, rj040114. Solid line indicates average performance across iterations of the bootstrap. Shaded area indicates +- 2 standard errors of the mean. Note that overfitting occurs when using many taps. **B.** To ensure that any performance gains were not due to overfitting, we repeated the previous analysis using ridge regression (Suminski, Tkach, and Hatsopoulos 2009). We observed that decoding performance no

longer declined with many taps suggesting that overfitting had been ameliorated by regularization, and that narrow spiking neurons still outperformed wide spiking neurons. **C.** We measured the performance gain, defined as the difference in R^2 values between narrow and wide for all datasets and found that the number of filter taps did not explain the disparity in decoding performance. Note that x velocity performance gains were slightly larger in the unregularized data from rj040114 suggesting that at least some of the improvement in performance at large numbers of taps may have been due to wide spiking neurons being more overfit than narrow spiking neurons. Nevertheless, in every case, narrow spiking neurons still significantly outperformed wide spiking neurons.

Although decoding performance varied with the number of taps in the model, the relative improvement from using narrow spiking ensembles was fairly constant across the range of taps with the narrow spiking populations always outperforming wide spiking populations (Fig. A2.4C).

Given that the number of taps in the decoder could not explain the difference in decoding ability, we next sought to control for several underlying response properties of these two populations. In general, narrow spiking neurons had higher firing rates, higher firing rate variance, and higher waveform SNRs (Fig. A2.5A-C), although in one dataset, mk080828, wide spiking neurons had higher firing rates and firing rate variance.

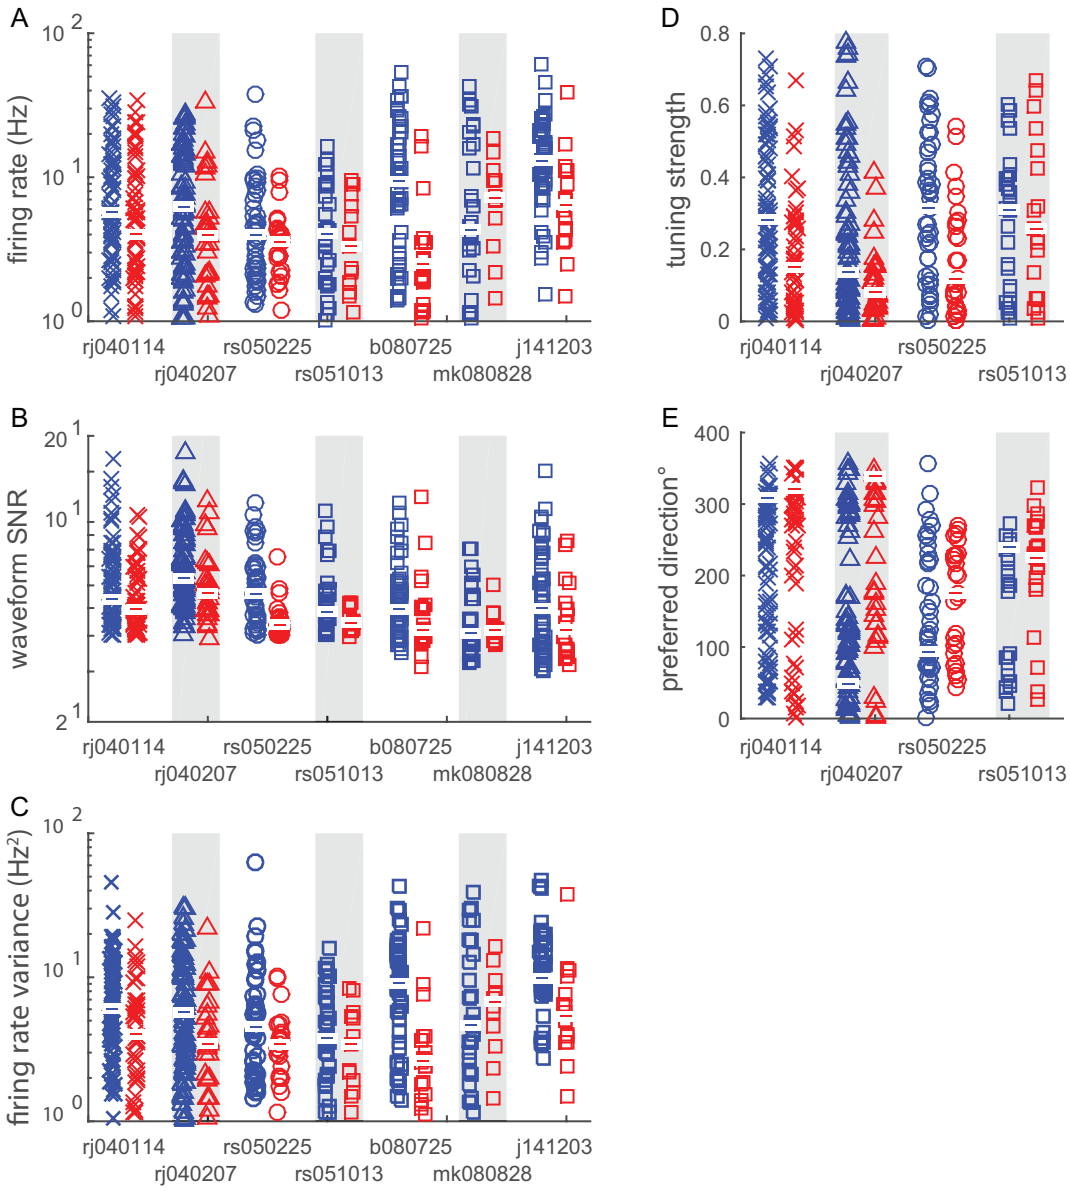

Figure A2.5 Response properties of narrow and wide spiking neural ensembles. The firing rate, firing variance, and waveform SNR for every neuron from each dataset was estimated (see Methods for details) and then compared based on waveform width category. Generally, narrow spiking units had significantly higher firing rates and waveform SNRs (all datasets except mk080828). Blue and red bars indicate median values for each dataset. For the center-out datasets, we estimated each neuron's tuning strength, and preferred direction. Again, tuning strength was significantly higher for narrow spiking neurons.

For two datasets, b080725 and j141203, narrow spiking neuron rates were significantly higher (KS test, $p < 0.0002$ and $p < 0.003$, respectively, Bonferroni adjusted for multiple comparisons), and, narrow spiking neuron rate variability was significantly

higher (KS test, $p < 0.00002$, and $p < 0.00006$, respectively). Narrow spiking neuron waveform SNRs were significantly greater than wide spiking SNRs in two datasets, rs050225, and rs051013 (KS test, $p < 0.0002$ for both datasets). Additionally, narrow spiking neurons showed stronger directional selectivity as revealed by their higher tuning strengths (Fig. A2.5D). This trend was significant in both datasets from animal R (KS test, $p < 0.007$ for both datasets). However, there was no significant difference in the distribution of preferred directions across waveform class in any dataset (circular medians test (Berens 2009), $p > 0.05$ for all datasets, Fig. A2.5E).

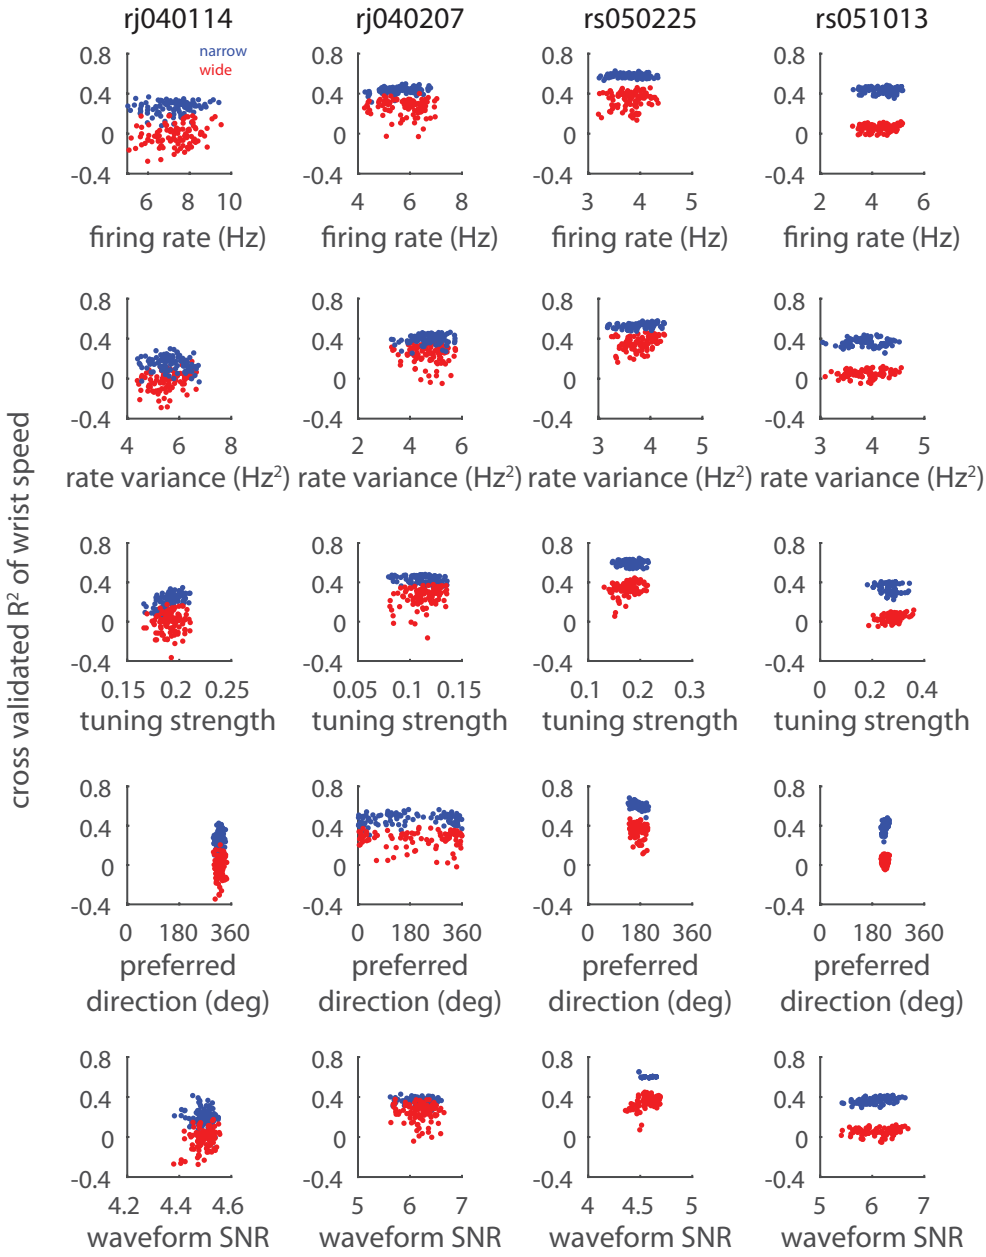

Figure A2.6 Narrow spiking neurons outperform wide spiking populations even after controlling for differences in response properties. We repeatedly drew random samples of narrow and wide spiking neurons (ensemble size of 30 for rj040114, 20 for rj040207 and rs050225, and 10 for rs051013) while controlling for either firing rate, firing variance, tuning strength, preferred direction, or waveform SNR using a matching procedure (see Methods for details). Narrow spiking neurons outperformed wide spiking neurons even after controlling for differences in response properties.

We developed a matching procedure to control for any putative differences between narrow and wide spiking neurons (see Methods for details). Random samples of wide spiking neurons were matched with narrow spiking units that exhibited the same firing rate, waveform SNR, preferred direction, or tuning strength each independently. This matching procedure yielded samples of narrow and wide spiking neurons that had statistically indistinguishable averages. Even after controlling for one underlying response property, narrow spiking units still almost always outperformed wide spiking units across a variety of motor features (Fig. A2.6 for wrist speed and Fig. A2.7 for *x* and *y*, velocities).

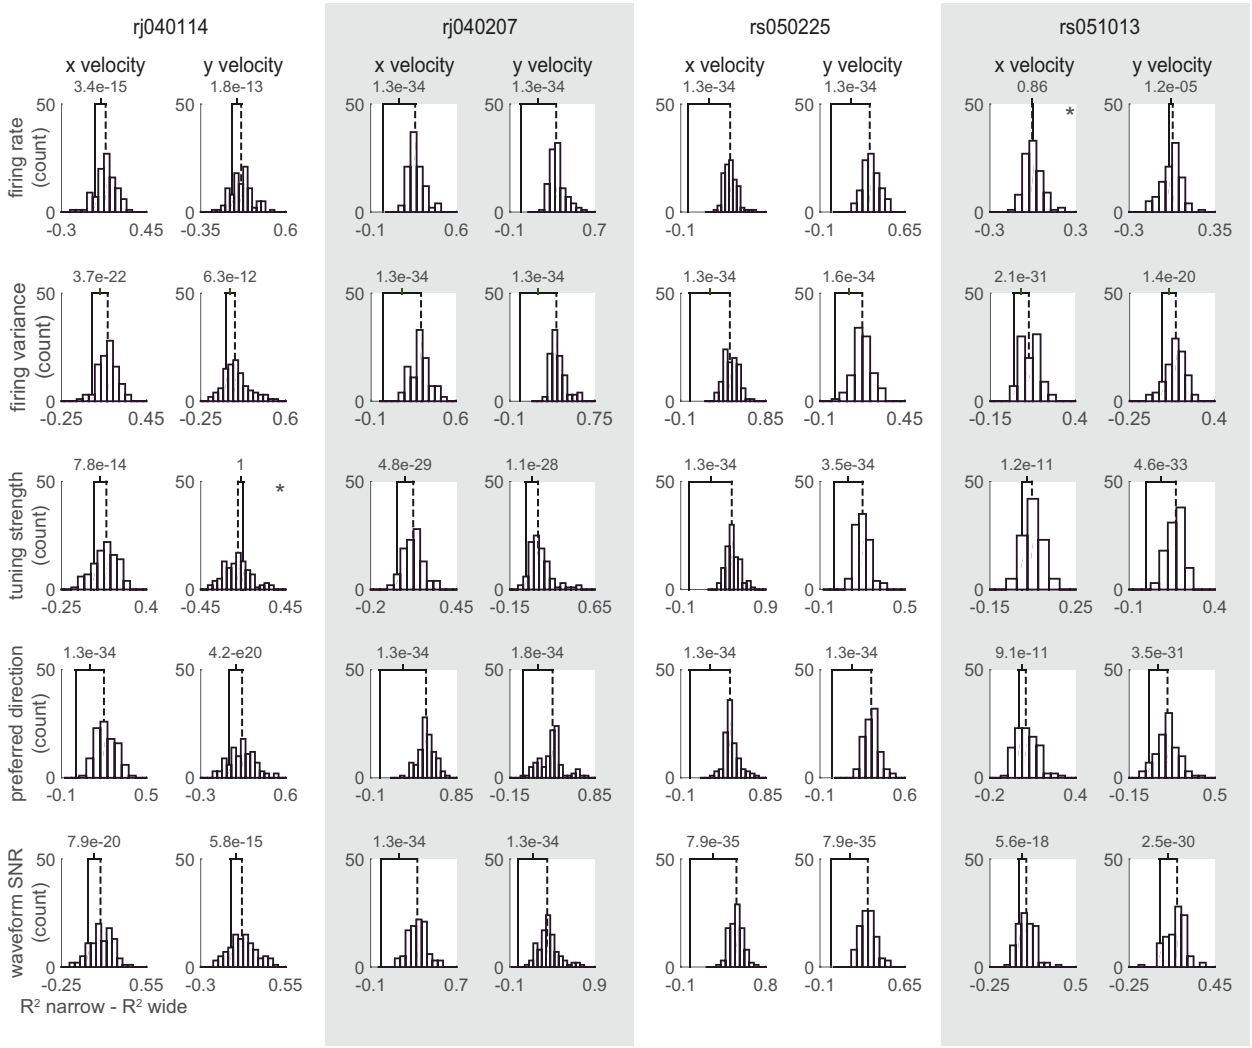

Figure A2.7 Underlying differences in response properties do not explain the difference in decoding performance. Here we show histograms of the difference in decoding performance (of *x* and *y* velocity) between narrow and wide spiking populations. The average difference is indicated by a vertical dashed line, while 0 is indicated by the solid vertical line. In almost every instance (except two indicated by stars), narrow spiking units outperformed wide spiking units even after controlling for one response property.

Decoding muscle activity

To further link narrow spiking neural activity with motor output, we examined data from an isometric center-out wrist task. Here, we attempted to predict the activity of 11 different upper limb muscles based on narrow and wide spiking neural ensembles (Fig.

A2.8). We found that each muscle's activity was also better predicted by narrow spiking ensembles (median improvement in R^2 was 0.06 across all motor features; ANOVA, EDC $F_{1,692}$ = 525; APL $F_{1,692}$ = 416; FDP $F_{1,692}$ = 213; ECR $F_{1,692}$ = 810; EDC2 $F_{1,692}$ = 453; Brad $F_{1,692}$ = 354; PT $F_{1,692}$ = 488; FCU $F_{1,692}$ = 219; FDS $F_{1,692}$ = 335; FCR $F_{1,692}$ = 90; FDS2 $F_{1,692}$ = 625; p < 1e-8 for all muscles).

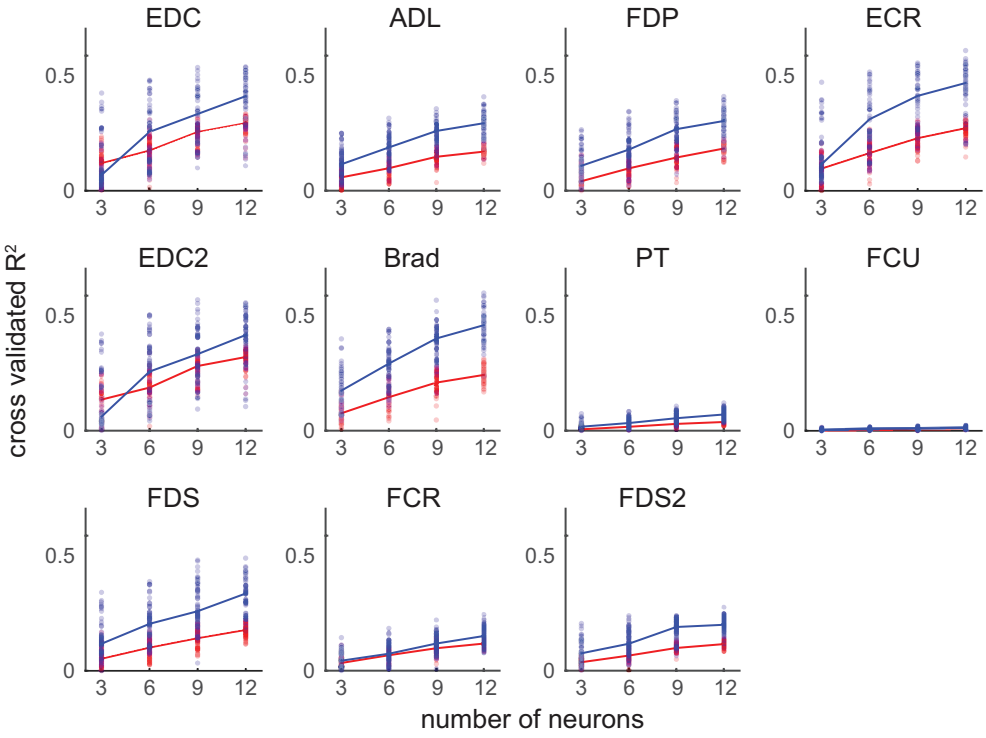

Figure A2.8 Decoding muscle activity using narrow and wide neural ensembles We used a standard 20-tap causal Wiener filter to decode muscle activity from neural data while an animal performed an isometric wrist flexion task. We repeatedly (100 times) drew random samples of either narrow or wide spiking neurons, trained a decoding model, and then tested its performance on a separate set of data. We found that narrow spiking neural ensembles outperformed wide spiking neural ensembles across all muscles.

Discussion

Interpretation of narrow and wide spiking neural ensembles

It is tempting to assume that narrow and wide spiking neurons correspond to inhibitory interneurons and pyramidal cells, respectively, because generally, inhibitory

interneurons exhibit narrow spike waveform widths while pyramidal cells have wider widths (Connors, Gutnick, and Prince 1982; McCormick et al. 1985; Contreras 2004; Taira and Georgopoulos 1993). However, recent evidence suggests that such a clear delineation is unlikely. One study found a relationship between Betz cells, projection cells in layer V of MI, and spike waveform width, such that the largest Betz cells had the narrowest waveform widths (Vigneswaran, Kraskov, and Lemon 2011). This population of cells is thought to comprise approximately 10-20% of neurons in layer V (Humphrey and Corrie 1978); however, due to their large size, they are oversampled, and may actually represent closer to 50% of recorded projection neurons (Humphrey and Corrie 1978; Towe and Harding 1970). Though the extent to which our data are subject to this sampling bias remains unknown, it is nevertheless likely that at least some of the narrow spiking neurons we recorded were indeed large projection neurons. Moreover, inhibitory interneurons exhibit a variety of spike widths including a small proportion with wide waveforms (Merchant, Naselaris, and Georgopoulos 2008). Thus, spike waveform width is not a reliable indicator of cell type.

In the present study, we found that narrow spiking neural ensembles substantially outperformed wide spiking ensembles in a variety of decoding contexts, and that this improvement in decoding performance was related to motor output. A fairly straightforward, albeit speculative explanation of this finding is that a substantial proportion of cells that we classified as narrow spiking neurons correspond to the largest Betz cells, and thus, the activity of the narrow spiking neural population contains more direct information about efferent motor activity.

Using Gaussian mixture models, we found that the distribution of spike waveform widths was best described by a mixture of two Gaussians based on AIC values; however, this finding does not imply that the true distribution of waveform widths is bimodal, nor are we arguing that it is. Indeed, the spike waveform width distribution was more clearly bimodal in some datasets than others. We used the Gaussian mixture model as a principled way of identifying the boundary between narrow and wide populations. In this way, it represents an improvement over previous methods based on specifying an arbitrary threshold (Kaufman, Churchland, and Shenoy 2013). The mixture model also provided a quantitative means of assessing the modality of the waveform width distribution rather than assuming bimodality.

Application to a clinically relevant BMI

Recent reports have shown that small ensembles of neurons are capable of achieving a high degree of decoding performance (Clancy et al. 2014; Law, Rivlis, and Schieber 2014). As BMIs scale to increasingly large degrees of freedom, these small ensembles may be used to control individual DoFs. In one study, individual control dimensions were allocated 10 neurons based on a statistical selection criterion (Balasubramanian et al. 2013). In our data, we observed that small ensembles of narrow spiking neurons could achieve performance comparable to, if not better than large ensembles of wide spiking neurons. This suggests that a dynamic allocation scheme could be devised based on the width of recorded neurons such that some DoFs would be controlled by small ensembles of narrow spiking neurons, while other DoFs would be controlled by larger populations of wide spiking neurons yet each DoF would

have the same expected level of performance despite being controlled by a different number of neurons.

An additional hurdle impeding the development of a clinically viable BMI is that few properties of the neural response are stable over long periods of time (Dickey et al. 2009). Here, we found that the bimodal distribution of spike waveform widths was similar across a timespan of several months. Additionally, in every dataset we analyzed, we observed a variety of spike waveform widths. Although there was some variability in the boundary between narrow and wide populations across animals, we observed both narrow and wide populations of cells in seven datasets recorded from five animals.

In summary, our approach has been to identify physiological properties of neurons that may reveal their utility in a neural decoder; this approach is not incompatible with other techniques aimed at improving decoding performance. Indeed, other approaches including linear dimensionality reduction and statistical model selection could be used in conjunction with waveform information to identify neurons within the narrow spiking neural population that are most relevant for decoding. More generally, we emphasize that neural decoding algorithms may be improved by using the underlying biological properties of neural signals to inform the design of these algorithms.

References

Aggarwal, Vikram, Mohsen Mollazadeh, Adam G. Davidson, Marc H. Schieber, and Nitish V. Thakor. 2013. “State-Based Decoding of Hand and Finger Kinematics Using Neuronal Ensemble and LFP Activity during Dexterous Reach-to-Grasp Movements.” *Journal of Neurophysiology* 109 (12): 3067–81. doi:10.1152/jn.01038.2011.

Akaike, H. 1974. “A New Look at the Statistical Model Identification.” *IEEE Transactions on Automatic Control* 19 (6): 716–23. doi:10.1109/TAC.1974.1100705.

Arce-McShane, Fritzie I., Nicholas G. Hatsopoulos, Jye-Chang Lee, Callum F. Ross, and Barry J. Sessle. 2014. “Modulation Dynamics in the Orofacial Sensorimotor Cortex during Motor Skill Acquisition.” *The Journal of Neuroscience* 34 (17): 5985–97. doi:10.1523/JNEUROSCI.4367-13.2014.

Baker, S. N., J. M. Kilner, E. M. Pinches, and R. N. Lemon. 1999. “The Role of Synchrony and Oscillations in the Motor Output.” *Experimental Brain Research* 128 (1-2): 109–17. doi:10.1007/s002210050825.

Balasubramanian, K., J. Southerland, M. Vaidya, Kai Qian, A. Eleryan, A.H. Fagg, M. Sluzky, K. Oweiss, and N. Hatsopoulos. 2013. “Operant Conditioning of a Multiple Degree-of-Freedom Brain-Machine Interface in a Primate Model of Amputation.” In *2013 35th Annual International Conference of the IEEE Engineering in Medicine and Biology Society (EMBC)*, 303–6. doi:10.1109/EMBC.2013.6609497.

Bansal, Arjun K., Wilson Truccolo, Carlos E. Vargas-Irwin, and John P. Donoghue. 2012. “Decoding 3D Reach and Grasp from Hybrid Signals in Motor and Premotor Cortices: Spikes, Multiunit Activity, and Local Field Potentials.” *Journal of Neurophysiology* 107 (5): 1337–55. doi:10.1152/jn.00781.2011.

Barbas, H., and D. N. Pandya. 1987. “Architecture and Frontal Cortical Connections of the Premotor Cortex (area 6) in the Rhesus Monkey.” *The Journal of Comparative Neurology* 256 (2): 211–28. doi:10.1002/cne.902560203.

Berens, Philipp. 2009. “CircStat: A MATLAB Toolbox for Circular Statistics.” *Journal of Statistical Software* 31 (10): 1–21.

Berger, Hans. 1933. “Über das Elektrenkephalogramm des Menschen.” *Archiv für Psychiatrie und Nervenkrankheiten* 99 (1): 555–74. doi:10.1007/BF01814320.

Berger, Professor Dr Hans. 1929. “Über das Elektrenkephalogramm des Menschen.” *Archiv für Psychiatrie und Nervenkrankheiten* 87 (1): 527–70. doi:10.1007/BF01797193.

Best, Matthew D., Aaron J. Suminski, Kazutaka Takahashi, and Nicholas G. Hatsopoulos. 2014. “Consideration of the Functional Relationship between Cortex and Motor Periphery Improves Offline Decoding Performance.” *Conference Proceedings: ... Annual International Conference of the IEEE Engineering in Medicine and Biology Society. IEEE Engineering in Medicine and Biology Society. Annual Conference* 2014: 4868–71. doi:10.1109/EMBC.2014.6944714.

Best, Matthew D., Kazutaka Takahashi, and Nicholas G. Hatsopoulos. 2015. “Comparing Decoding Performance between Functionally Defined Neural Populations.” In *2015 7th International IEEE/EMBS Conference on Neural Engineering (NER)*.

Best, M.D., K. Takahashi, Zhe Chen, N. Huh, K.A. Brown, and N.G. Hatsopoulos. 2013. “Integrating Neural Spiking and LFP Activity to Decode Kinematics of the Arm and Hand during Unconstrained Reach to Grasp Movements.” In *2013 6th International IEEE/EMBS Conference on Neural Engineering (NER)*, 1425–28. doi:10.1109/NER.2013.6696211.

Bishop, Christopher M. 2006. *Pattern Recognition and Machine Learning*. Information Science and Statistics. Springer. http://www.springer.com/computer/image+processing/book/978-0-387-31073-2.

Bonini, Luca, Monica Maranesi, Alessandro Livi, Leonardo Fogassi, and Giacomo Rizzolatti. 2014. “Space-Dependent Representation of Objects and Other’s Action in Monkey Ventral Premotor Grasping Neurons.” *The Journal of Neuroscience* 34 (11): 4108–19. doi:10.1523/JNEUROSCI.4187-13.2014.

Boudrias, Marie-Hélène, Rebecca L. McPherson, Shawn B. Frost, and Paul D. Cheney. 2010. “Output Properties and Organization of the Forelimb Representation of Motor Areas on the Lateral Aspect of the Hemisphere in Rhesus Macaques.” *Cerebral Cortex* 20 (1): 169–86. doi:10.1093/cercor/bhp084.

Brasil-Neto, J. P., L. G. Cohen, M. Panizza, J. Nilsson, B. J. Roth, and M. Hallett. 1992. “Optimal Focal Transcranial Magnetic Activation of the Human Motor Cortex: Effects of Coil Orientation, Shape of the Induced Current Pulse, and Stimulus Intensity.” *Journal of Clinical Neurophysiology: Official Publication of the American Electroencephalographic Society* 9 (1): 132–36.

Brendel, Wieland, Ranulfo Romo, and Christian K Machens. 2011. “Demixed Principal Component Analysis.” In *Advances in Neural Information Processing Systems 24*, edited by J. Shawe-Taylor, R. S. Zemel, P. L. Bartlett, F. Pereira, and K. Q. Weinberger, 2654–62. Curran Associates, Inc. http://papers.nips.cc/paper/4215-demixed-principal-component-analysis.pdf.

Brockmeier, Austin J., and José C. Príncipe. 2013. “Decoding Algorithms for Brain–Machine Interfaces.” In *Neural Engineering*, edited by Bin He, 223–57. Springer US. http://link.springer.com/chapter/10.1007/978-1-4614-5227-0_4.

Bucy, Paul C., and John F. Fulton. 1933. “Ipsilateral Representation in the Motor and Premotor Cortex Op Monkeys.” *Brain* 56 (3): 318–42. doi:10.1093/brain/56.3.318.

Buonomano, Dean V., and Rodrigo Laje. 2010. “Population Clocks: Motor Timing with Neural Dynamics.” *Trends in Cognitive Sciences*, Special Issue: Space, Time and Number, 14 (12): 520–27. doi:10.1016/j.tics.2010.09.002.

Carmena, Jose M, Mikhail A Lebedev, Roy E Crist, Joseph E O’Doherty, David M Santucci, Dragan F Dimitrov, Parag G Patil, Craig S Henriquez, and Miguel A. L Nicolelis. 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.

Cerri, G., H. Shimazu, M. A. Maier, and R. N. Lemon. 2003. “Facilitation From Ventral Premotor Cortex of Primary Motor Cortex Outputs to Macaque Hand Muscles.” *Journal of Neurophysiology* 90 (2): 832–42. doi:10.1152/jn.01026.2002.

Chen, Zhe. 2013. “An Overview of Bayesian Methods for Neural Spike Train Analysis.” *Computational Intelligence and Neuroscience* 2013 (November): e251905. doi:10.1155/2013/251905.

Chen, Zhe, and K. Takahashi. 2013. “Sparse Bayesian Inference Methods for Decoding 3D Reach and Grasp Kinematics and Joint Angles with Primary Motor Cortical Ensembles.” In *2013 35th Annual International Conference of the IEEE Engineering in Medicine and Biology Society (EMBC)*, 5930–33. doi:10.1109/EMBC.2013.6610902.

Chi, Zhiyi, Wei Wu, Zach Haga, Nicholas G. Hatsopoulos, and Daniel Margoliash. 2007. “Template-Based Spike Pattern Identification With Linear Convolution and Dynamic Time Warping.” *Journal of Neurophysiology* 97 (2): 1221–35. doi:10.1152/jn.00448.2006.

Churchland, Mark M., John P. Cunningham, Matthew T. Kaufman, Justin D. Foster, Paul Nuyujukian, Stephen I. Ryu, and Krishna V. Shenoy. 2012. “Neural Population Dynamics during Reaching.” *Nature* 487 (June): 51–56. doi:10.1038/nature11129.

Churchland, Mark M., John P. Cunningham, Matthew T. Kaufman, Stephen I. Ryu, and Krishna V. Shenoy. 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, Mark M., and Krishna V. Shenoy. 2007. “Temporal Complexity and Heterogeneity of Single-Neuron Activity in Premotor and Motor Cortex.” *Journal of Neurophysiology* 97 (6): 4235–57. doi:10.1152/jn.00095.2007.

Clancy, Kelly B., Aaron C. Koralek, Rui M. Costa, Daniel E. Feldman, and Jose M. Carmena. 2014. “Volitional Modulation of Optically Recorded Calcium Signals during Neuroprosthetic Learning.” *Nature Neuroscience* advance online publication (April). doi:10.1038/nn.3712.

Connors, B. W., M. J. Gutnick, and D. A. Prince. 1982. “Electrophysiological Properties of Neocortical Neurons in Vitro.” *Journal of Neurophysiology* 48 (6): 1302–20.

Contreras, Diego. 2004. “Electrophysiological Classes of Neocortical Neurons.” *Neural Networks*, Vision and Brain, 17 (5–6): 633–46. doi:10.1016/j.neunet.2004.04.003.

Conway, B A, D M Halliday, S F Farmer, U Shahani, P Maas, A I Weir, and J R Rosenberg. 1995. “Synchronization between Motor Cortex and Spinal Motoneuronal Pool during the Performance of a Maintained Motor Task in Man.” *The Journal of Physiology* 489 (3): 917–24. doi:10.1113/jphysiol.1995.sp021104.

Crawford, J. Douglas, Denise Y. P. Henriques, and W. Pieter Medendorp. 2011. “Three-Dimensional Transformations for Goal-Directed Action.” *Annual Review of Neuroscience* 34 (1): 309–31. doi:10.1146/annurev-neuro-061010-113749.

Crone, N E, D L Miglioretti, B Gordon, J M Sieracki, M T Wilson, S Uematsu, and R P Lesser. 1998. “Functional Mapping of Human Sensorimotor Cortex with Electrocorticographic Spectral Analysis. I. Alpha and Beta Event-Related Desynchronization.” *Brain: A Journal of Neurology* 121 (Pt 12) (December): 2271–99.

Dancause, Numa, Scott Barbay, Shawn B. Frost, Erik J. Plautz, Mihai Popescu, Philip M. Dixon, Ann M. Stowe, Kathleen M. Friel, and Randolph J. Nudo. 2006. “Topographically Divergent and Convergent Connectivity between Premotor and Primary Motor Cortex.” *Cerebral Cortex* 16 (8): 1057–68. doi:10.1093/cercor/bhj049.

Delp, S.L., F.C. Anderson, A.S. Arnold, P. Loan, A. Habib, C.T. John, E. Guendelman, and D.G. Thelen. 2007. “OpenSim: Open-Source Software to Create and Analyze Dynamic Simulations of Movement.” *IEEE Transactions on Biomedical Engineering* 54 (11): 1940–50. doi:10.1109/TBME.2007.901024.

Denker, Michael, Sébastien Roux, Marc Timme, Alexa Riehle, and Sonja Grün. 2007. “Phase Synchronization between LFP and Spiking Activity in Motor Cortex during Movement Preparation.” *Neurocomputing*, Computational Neuroscience: Trends

in Research 2007Computational Neuroscience 2006, 70 (10–12): 2096–2101. doi:10.1016/j.neucom.2006.10.088.

Dickey, Adam S., Aaron Suminski, Yali Amit, and Nicholas G. Hatsopoulos. 2009. "Single-Unit Stability Using Chronically Implanted Multielectrode Arrays." *Journal of Neurophysiology* 102 (2): 1331–39. doi:10.1152/jn.90920.2008.

Donoghue, John P., Jerome N. Sanes, Nicholas G. Hatsopoulos, and Gyöngyi Gaál. 1998. "Neural Discharge and Local Field Potential Oscillations in Primate Motor Cortex During Voluntary Movements." *Journal of Neurophysiology* 79 (1): 159–73.

Donoghue, J. P., and J. N. Sanes. 1994. "Motor Areas of the Cerebral Cortex." *Journal of Clinical Neurophysiology: Official Publication of the American Electroencephalographic Society* 11 (4): 382–96.

Dum, Richard P., and Peter L. Strick. 2005. "Frontal Lobe Inputs to the Digit Representations of the Motor Areas on the Lateral Surface of the Hemisphere." *The Journal of Neuroscience* 25 (6): 1375–86. doi:10.1523/JNEUROSCI.3902-04.2005.

Dum, R. P., and P. L. Strick. 1991. "The Origin of Corticospinal Projections from the Premotor Areas in the Frontal Lobe." *The Journal of Neuroscience* 11 (3): 667–89.

Engel, Andreas K, and Pascal Fries. 2010. "Beta-Band Oscillations — Signalling the Status Quo?" *Current Opinion in Neurobiology* 20 (2): 156–65. doi:10.1016/j.conb.2010.02.015.

Evarts, E. V., C. Fromm, J. Kroller, and V. A. Jennings. 1983. "Motor Cortex Control of Finely Graded Forces." *Journal of Neurophysiology* 49 (5): 1199–1215.

Fagg, A.H., G.W. Ojakangas, L.E. Miller, and N.G. Hatsopoulos. 2009. "Kinetic Trajectory Decoding Using Motor Cortical Ensembles." *IEEE Transactions on Neural Systems and Rehabilitation Engineering* 17 (5): 487–96. doi:10.1109/TNSRE.2009.2029313.

Fattori, Patrizia, Vassilis Raos, Rossella Breveglieri, Annalisa Bosco, Nicoletta Marzocchi, and Claudio Galletti. 2010. "The Dorsomedial Pathway Is Not Just for Reaching: Grasping Neurons in the Medial Parieto-Occipital Cortex of the Macaque Monkey." *The Journal of Neuroscience* 30 (1): 342–49. doi:10.1523/JNEUROSCI.3800-09.2010.

Fulton, J. F. 1935. "A Note on the Definition of the 'Motor' and 'Premotor' Areas." *Brain* 58 (2): 311–16. doi:10.1093/brain/58.2.311.

Gatter, K. c, and T. P. S. Powell. 1978. “The Intrinsic Connections of the Cortex of Area 4 of the Monkey.” *Brain* 101 (3): 513–41. doi:10.1093/brain/101.3.513.

Gentilucci, M., L. Fogassi, G. Luppino, M. Matelli, R. Camarda, and G. Rizzolatti. 1988. “Functional Organization of Inferior Area 6 in the Macaque Monkey.” *Experimental Brain Research* 71 (3): 475–90. doi:10.1007/BF00248741.

Georgopoulos, A. P., J. F. Kalaska, R. Caminiti, and J. T. Massey. 1982. “On the Relations between the Direction of Two-Dimensional Arm Movements and Cell Discharge in Primate Motor Cortex.” *The Journal of Neuroscience* 2 (11): 1527–37.

Gilbertson, Thomas, Elodie Lalo, Louise Doyle, Vincenzo Di Lazzaro, Beatrice Cioni, and Peter Brown. 2005. “Existing Motor State Is Favored at the Expense of New Movement during 13-35 Hz Oscillatory Synchrony in the Human Corticospinal System.” *The Journal of Neuroscience* 25 (34): 7771–79. doi:10.1523/JNEUROSCI.1762-05.2005.

Godschalk, M., R. N. Lemon, H. G. Kuypers, and H. K. Ronday. 1984. “Cortical Afferents and Efferents of Monkey Postarcuate Area: An Anatomical and Electrophysiological Study.” *Experimental Brain Research* 56 (3): 410–24.

Godschalk, M., R. N. Lemon, H. G. Kuypers, and J. van der Steen. 1985. “The Involvement of Monkey Premotor Cortex Neurones in Preparation of Visually Cued Arm Movements.” *Behavioural Brain Research* 18 (2): 143–57.

Godschalk, M., R. N. Lemon, H. G. Nijs, and H. G. Kuypers. 1981. “Behaviour of Neurons in Monkey Peri-Arcuate and Precentral Cortex before and during Visually Guided Arm and Hand Movements.” *Experimental Brain Research* 44 (1): 113–16.

Goel, A., and D. V. Buonomano. 2014. “Timing as an Intrinsic Property of Neural Networks: Evidence from in Vivo and in Vitro Experiments.” *Philosophical Transactions of the Royal Society B: Biological Sciences* 369 (1637): 20120460–20120460. doi:10.1098/rstb.2012.0460.

Grafton, Scott. 2013. “The Cognitive Neuroscience of Prehension: Recent Developments.” *Experimental Brain Research*. Accessed April 26. http://link.springer.com/article/10.1007/s00221-010-2315-2/fulltext.html.

Graziano, Michael S. A. 2016. “Ethological Action Maps: A Paradigm Shift for the Motor Cortex.” *Trends in Cognitive Sciences* 0 (0). Accessed January 8. doi:10.1016/j.tics.2015.10.008.

Graziano, M. S. A. 2005. “Arm Movements Evoked by Electrical Stimulation in the Motor Cortex of Monkeys.” *Journal of Neurophysiology* 94 (6): 4209–23. doi:10.1152/jn.01303.2004.

Griffin, Darcy M., Heather M. Hudson, Abderraouf Belhaj-Saïf, and Paul D. Cheney. 2014. “EMG Activation Patterns Associated with High Frequency, Long-Duration Intracortical Microstimulation of Primary Motor Cortex.” *The Journal of Neuroscience* 34 (5): 1647–56. doi:10.1523/JNEUROSCI.3643-13.2014.

Haggard, Patrick, and Alan Wing. 1995. “Coordinated Responses Following Mechanical Perturbation of the Arm during Prehension.” *Experimental Brain Research* 102 (3): 483–94. doi:10.1007/BF00230652.

Halsband, U., and R. E. Passingham. 1985. “Premotor Cortex and the Conditions for Movement in Monkeys (Macaca Fascicularis).” *Behavioural Brain Research* 18 (3): 269–77. doi:10.1016/0166-4328(85)90035-X.

Hatsopoulos, Nicholas G., Leo Olmedo, and Kazutaka Takahashi. 2010. “Proximal-to-Distal Sequencing Behavior and Motor Cortex.” In *Motor Control*. New York, NY, USA: Oxford University Press.

Hatsopoulos, Nicholas G., and Aaron J. Suminski. 2011. “Sensing with the Motor Cortex.” *Neuron* 72 (3): 477–87. doi:10.1016/j.neuron.2011.10.020.

Hatsopoulos, Nicholas G., Qingqing Xu, and Yali Amit. 2007. “Encoding of Movement Fragments in the Motor Cortex.” *The Journal of Neuroscience* 27 (19): 5105–14. doi:10.1523/JNEUROSCI.3570-06.2007.

Hendrix, Claudia M., Carolyn R. Mason, and Timothy J. Ebner. 2009. “Signaling of Grasp Dimension and Grasp Force in Dorsal Premotor Cortex and Primary Motor Cortex Neurons During Reach to Grasp in the Monkey.” *Journal of Neurophysiology* 102 (1): 132–45. doi:10.1152/jn.00016.2009.

He, S. Q., R. P. Dum, and P. L. Strick. 1993. “Topographic Organization of Corticospinal Projections from the Frontal Lobe: Motor Areas on the Lateral Surface of the Hemisphere.” *The Journal of Neuroscience* 13 (3): 952–80.

Hocherman, S., and S. P. Wise. 1991. “Effects of Hand Movement Path on Motor Cortical Activity in Awake, Behaving Rhesus Monkeys.” *Experimental Brain Research* 83 (2): 285–302. doi:10.1007/BF00231153.

Hoshi, Eiji, and Jun Tanji. 2007. “Distinctions between Dorsal and Ventral Premotor Areas: Anatomical Connectivity and Functional Properties.” *Current Opinion in Neurobiology* 17 (2): 234–42. doi:10.1016/j.conb.2007.02.003.

Humphrey, D. R., and W. S. Corrie. 1978. “Properties of Pyramidal Tract Neuron System within a Functionally Defined Subregion of Primate Motor Cortex.” *Journal of Neurophysiology* 41 (1): 216–43.

Huntley, G. W., and E. G. Jones. 1991. “Relationship of Intrinsic Connections to Forelimb Movement Representations in Monkey Motor Cortex: A Correlative Anatomic and Physiological Study.” *Journal of Neurophysiology* 66 (2): 390–413.

Jakobson, L. S., and M. A. Goodale. 1991. “Factors Affecting Higher-Order Movement Planning: A Kinematic Analysis of Human Prehension.” *Experimental Brain Research* 86 (1): 199–208.

Jasper, Herbert, and Wilder Penfield. 1949. “Electrocorticograms in Man: Effect of Voluntary Movement upon the Electrical Activity of the Precentral Gyrus.” *Archiv Für Psychiatrie Und Nervenkrankheiten* 183 (1-2): 163–74. doi:10.1007/BF01062488.

Jeannerod, Marc. 1981. “Intersegmental Coordination during Reaching at Natural Visual Objects.” In *Attention and Performance*, IX:153–68. Hillsdale, NJ: Lawrence Erlbaum Associates.

———. 1984. “The Timing of Natural Prehension Movements.” *Journal of Motor Behavior* 16: 235–54.

———. 1988. *The Neural and Behavioural Organization of Goal-Directed Movements*. Oxford Psychology Series 15. Oxford Science Publications.

Jeannerod, Marc, M.A. Arbib, G Rizzolatti, and H. Sakata. 1995. “Grasping Objects: The Cortical Mechanisms of Visuomotor Transformation.” *Trends in Neurosciences* 18 (7): 314–20. doi:10.1016/0166-2236(95)93921-J.

Jenkinson, Ned, and Peter Brown. 2011. “New Insights into the Relationship between Dopamine, Beta Oscillations and Motor Function.” *Trends in Neurosciences* 34 (12): 611–18. doi:10.1016/j.tins.2011.09.003.

Jung, Nikolai H., Igor Delvendahl, Astrid Pechmann, Bernhard Gleich, Norbert Gattinger, Hartwig R. Siebner, and Volker Mall. 2012. “Transcranial Magnetic Stimulation with a Half-Sine Wave Pulse Elicits Direction-Specific Effects in Human Motor Cortex.” *BMC Neuroscience* 13 (1): 139. doi:10.1186/1471-2202-13-139.

Kaas, Jon H., Omar A. Gharbawie, and Iwona Stepniewska. 2013. “Cortical Networks for Ethologically Relevant Behaviors in Primates: Cortical Networks and Behavior.” *American Journal of Primatology* 75 (5): 407–14. doi:10.1002/ajp.22065.

Kalaska, John F. 2009. “From Intention to Action: Motor Cortex and the Control of Reaching Movements.” In *Progress in Motor Control*, edited by Dagmar Sternad, 629:139–78. Boston, MA: Springer US. http://link.springer.com/10.1007/978-0-387-77064-2_8.

Kammer, Thomas, Sandra Beck, Axel Thielscher, Ulrike Laubis-Herrmann, and Helge Topka. 2001. “Motor Thresholds in Humans: A Transcranial Magnetic Stimulation Study Comparing Different Pulse Waveforms, Current Directions and Stimulator Types.” *Clinical Neurophysiology* 112 (2): 250–58. doi:10.1016/S1388-2457(00)00513-7.

Kandel, Eric, James Schwartz, and Thomas Jessell. 2013. *Principles of Neural Science, Fifth Edition*. McGraw Hill Professional.

Kang, X., s. sarma, S. Santaniello, M. Schieber, and N. Thakor. 2015. “Task-Independent Cognitive State Transition Detection From Cortical Neurons During 3D Reach-to-Grasp Movements.” *IEEE Transactions on Neural Systems and Rehabilitation Engineering* PP (99): 1–1. doi:10.1109/TNSRE.2015.2396495.

Kao, J.C., P. Nuyujukian, S. Stavisky, S.I. Ryu, S. Ganguli, and K.V. Shenoy. 2013. “Investigating the Role of Firing-Rate Normalization and Dimensionality Reduction in Brain-Machine Interface Robustness.” In *2013 35th Annual International Conference of the IEEE Engineering in Medicine and Biology Society (EMBC)*, 293–98. doi:10.1109/EMBC.2013.6609495.

Karst, G. M., and Z. Hasan. 1991. “Timing and Magnitude of Electromyographic Activity for Two-Joint Arm Movements in Different Directions.” *Journal of Neurophysiology* 66 (5): 1594–1604.

Kaufman, Matthew T., Mark M. Churchland, Stephen I. Ryu, and Krishna V. Shenoy. 2014. “Cortical Activity in the Null Space: Permitting Preparation without Movement.” *Nature Neuroscience* 17 (3): 440–48. doi:10.1038/nn.3643.

Kaufman, Matthew T., Mark M. Churchland, and Krishna V. Shenoy. 2013. “The Roles of Monkey M1 Neuron Classes in Movement Preparation and Execution.” *Journal of Neurophysiology* 110 (4): 817–25. doi:10.1152/jn.00892.2011.

Kennard, Margaret A., H. R. Viets, and J. F. Fulton. 1934. “The Syndrome of the Premotor Cortex in Man: Impairment of Skilled Movements, Forced Grasping, Spasticity, and Vasomotor Disturbance.” *Brain* 57 (1): 69–84. doi:10.1093/brain/57.1.69.

Khanna, Preeya, and Jose M Carmena. 2015. “Neural Oscillations: Beta Band Activity across Motor Networks.” *Current Opinion in Neurobiology* 32 (June): 60–67. doi:10.1016/j.conb.2014.11.010.

Kilavik, Bjørg Elisabeth, Manuel Zaepffel, Andrea Brovelli, William A. MacKay, and Alexa Riehle. 2013. “The Ups and Downs of Beta Oscillations in Sensorimotor Cortex.” *Experimental Neurology* 245 (July): 15–26. doi:10.1016/j.expneurol.2012.09.014.

Kobak, Dmitry, Wieland Brendel, Christos Constantinidis, Claudia E. Feierstein, Adam Kepecs, Zachary F. Mainen, Ranulfo Romo, Xue-Lian Qi, Naoshige Uchida, and Christian K. Machens. 2014. “Demixed Principal Component Analysis of Population Activity in Higher Cortical Areas Reveals Independent Representation of Task Parameters.” *arXiv:1410.6031 [q-Bio, Stat]*, October. http://arxiv.org/abs/1410.6031.

Kurata, K. 1991. “Corticocortical Inputs to the Dorsal and Ventral Aspects of the Premotor Cortex of Macaque Monkeys.” *Neuroscience Research* 12 (1): 263–80.

Kurata, K., and J. Tanji. 1986. “Premotor Cortex Neurons in Macaques: Activity before Distal and Proximal Forelimb Movements.” *The Journal of Neuroscience* 6 (2): 403–11.

Kwan, H. C., W. A. MacKay, J. T. Murphy, and Y. C. Wong. 1978. “Spatial Organization of Precentral Cortex in Awake Primates. II. Motor Outputs.” *Journal of Neurophysiology* 41 (5): 1120–31.

Laje, Rodrigo, and Dean V Buonomano. 2013. “Robust Timing and Motor Patterns by Taming Chaos in Recurrent Neural Networks.” *Nature Neuroscience* 16 (7): 925–33. doi:10.1038/nn.3405.

Law, Andrew J., Gil Rivlis, and Marc H. Schieber. 2014. “Rapid Acquisition of Novel Interface Control by Small Ensembles of Arbitrarily Selected Primary Motor Cortex Neurons.” *Journal of Neurophysiology* 112 (6): 1528–48. doi:10.1152/jn.00373.2013.

Lecas, J. C., J. Requin, C. Anger, and N. Vitton. 1986. “Changes in Neuronal Activity of the Monkey Precentral Cortex during Preparation for Movement.” *Journal of Neurophysiology* 56 (6): 1680–1702.

Lehmann, S. J., and H. Scherberger. 2013. “Reach and Gaze Representations in Macaque Parietal and Premotor Grasp Areas.” *Journal of Neuroscience* 33 (16): 7038–49. doi:10.1523/JNEUROSCI.5568-12.2013.

Li, Choh-Luh, and Herbert Jasper. 1953. “Microelectrode Studies of the Electrical Activity of the Cerebral Cortex in the Cat.” *The Journal of Physiology* 121 (1): 117–40.

Malik, W.Q., L.R. Hochberg, J.P. Donoghue, and E.N. Brown. 2014. “Modulation Depth Estimation and Variable Selection in State-Space Models for Neural Interfaces.” *IEEE Transactions on Biomedical Engineering* PP (99): 1–1. doi:10.1109/TBME.2014.2360393.

Mante, Valerio, David Sussillo, Krishna V. Shenoy, and William T. Newsome. 2013. “Context-Dependent Computation by Recurrent Dynamics in Prefrontal Cortex.” *Nature* 503 (7474): 78–84. doi:10.1038/nature12742.

Marzke, M. 1994. “Evolution.” In *Insights into the Reach to Grasp Movement*, 1st ed., 19–35. Amsterdam: Elsevier.

Mason, C. R., J. E. Gomez, and T. J. Ebner. 2001. “Hand Synergies during Reach-to-Grasp.” *Journal of Neurophysiology* 86 (6): 2896–2910.

Matelli, Massimo, Giuseppe Luppino, and Giacomo Rizzolatti. 1985. “Patterns of Cytochrome Oxidase Activity in the Frontal Agranular Cortex of the Macaque Monkey.” *Behavioural Brain Research* 18 (2): 125–36. doi:10.1016/0166-4328(85)90068-3.

Matsumura, Michikazu, and Kisou Kubota. 1979. “Cortical Projection to Hand-Arm Motor Area from Post-Arcuate Area in Macaque Monkeys: A Histological Study of Retrograde Transport of Horseradish Peroxidase.” *Neuroscience Letters* 11 (3): 241–46. doi:10.1016/0304-3940(79)90001-6.

McCormick, D. A., B. W. Connors, J. W. Lighthall, and D. A. Prince. 1985. “Comparative Electrophysiology of Pyramidal and Sparsely Spiny Stellate Neurons of the Neocortex.” *Journal of Neurophysiology* 54 (4): 782–806.

Mehring, Carsten, Jörn Rickert, Eilon Vaadia, Simone Cardoso de Oliveira, Ad Aertsen, and Stefan Rotter. 2003. “Inference of Hand Movements from Local Field Potentials in Monkey Motor Cortex.” *Nature Neuroscience* 6 (12): 1253–54. doi:10.1038/nn1158.

Merchant, Hugo, Thomas Naselaris, and Apostolos P. Georgopoulos. 2008. “Dynamic Sculpting of Directional Tuning in the Primate Motor Cortex during Three-Dimensional Reaching.” *The Journal of Neuroscience* 28 (37): 9164–72. doi:10.1523/JNEUROSCI.1898-08.2008.

Meyer, Gundela. 1987. “Forms and Spatial Arrangement of Neurons in the Primary Motor Cortex of Man.” *The Journal of Comparative Neurology* 262 (3): 402–28. doi:10.1002/cne.902620306.

Mollazadeh, Mohsen, Vikram Aggarwal, Adam G. Davidson, Andrew J. Law, Nitish V. Thakor, and Marc H. Schieber. 2011. “Spatiotemporal Variation of Multiple

Neurophysiological Signals in the Primary Motor Cortex during Dexterous Reach-to-Grasp Movements." *The Journal of Neuroscience* 31 (43): 15531–43. doi:10.1523/JNEUROSCI.2999-11.2011.

Mollazadeh, Mohsen, Vikram Aggarwal, Nitish V. Thakor, and Marc H. Schieber. 2014. "Principal Components of Hand Kinematics and Neurophysiological Signals in Motor Cortex during Reach to Grasp Movements." *Journal of Neurophysiology*, July, jn.00481.2013. doi:10.1152/jn.00481.2013.

Muakkassa, Kamel F., and Peter L. Strick. 1979. "Frontal Lobe Inputs to Primate Motor Cortex: Evidence for Four Somatotopically Organized 'premotor' Areas." *Brain Research* 177 (1): 176–82. doi:10.1016/0006-8993(79)90928-4.

Murphy, J. T., Y. C. Wong, and H. C. Kwan. 1985. "Sequential Activation of Neurons in Primate Motor Cortex during Unrestrained Forelimb Movement." *Journal of Neurophysiology* 53 (2): 435–45.

Murthy, V. N., and E. E. Fetz. 1992. "Coherent 25- to 35-Hz Oscillations in the Sensorimotor Cortex of Awake Behaving Monkeys." *Proceedings of the National Academy of Sciences* 89 (12): 5670–74. doi:10.1073/pnas.89.12.5670.

———. 1996a. "Oscillatory Activity in Sensorimotor Cortex of Awake Monkeys: Synchronization of Local Field Potentials and Relation to Behavior." *Journal of Neurophysiology* 76 (6): 3949–67.

———. 1996b. "Synchronization of Neurons during Local Field Potential Oscillations in Sensorimotor Cortex of Awake Monkeys." *Journal of Neurophysiology* 76 (6): 3968–82.

O'Doherty, J.E., M.A. Lebedev, Zheng Li, and M.A.L. Nicolelis. 2012. "Virtual Active Touch Using Randomly Patterned Intracortical Microstimulation." *IEEE Transactions on Neural Systems and Rehabilitation Engineering* 20 (1): 85–93. doi:10.1109/TNSRE.2011.2166807.

O'Leary, John G., and Nicholas G. Hatsopoulos. 2006. "Early Visuomotor Representations Revealed From Evoked Local Field Potentials in Motor and Premotor Cortical Areas." *Journal of Neurophysiology* 96 (3): 1492–1506. doi:10.1152/jn.00106.2006.

Paninski, Liam, Matthew R. Fellows, Nicholas G. Hatsopoulos, and John P. Donoghue. 2004. "Spatiotemporal Tuning of Motor Cortical Neurons for Hand Position and Velocity." *Journal of Neurophysiology* 91 (1): 515–32. doi:10.1152/jn.00587.2002.

Pani, Pierpaolo, Fabio Di Bello, Emiliano Brunamonti, Valeria D'Andrea, Odysseas Papazachariadis, and Stefano Ferraina. 2014. "Alpha- and Beta-Band

Oscillations Subserve Different Processes in Reactive Control of Limb Movements." *Frontiers in Behavioral Neuroscience* 8 (November). doi:10.3389/fnbeh.2014.00383.

Park, Michael C., Abderraouf Belhaj-Saïf, and Paul D. Cheney. 2004. "Properties of Primary Motor Cortex Output to Forelimb Muscles in Rhesus Macaques." *Journal of Neurophysiology* 92 (5): 2968–84. doi:10.1152/jn.00649.2003.

Park, Michael C., Abderraouf Belhaj-Saïf, Michael Gordon, and Paul D. Cheney. 2001. "Consistent Features in the Forelimb Representation of Primary Motor Cortex in Rhesus Macaques." *The Journal of Neuroscience* 21 (8): 2784–92.

Paulignan, Y., V. G. Frak, I. Toni, and M. Jeannerod. 1997. "Influence of Object Position and Size on Human Prehension Movements." *Experimental Brain Research* 114 (2): 226–34.

Paulignan, Y., M. Jeannerod, C. MacKenzie, and R. Marteniuk. 1991. "Selective Perturbation of Visual Input during Prehension Movements. 2. The Effects of Changing Object Size." *Experimental Brain Research* 87 (2): 407–20.

Paulignan, Y., C. MacKenzie, R. Marteniuk, and M. Jeannerod. 1991. "Selective Perturbation of Visual Input during Prehension Movements. 1. The Effects of Changing Object Position." *Experimental Brain Research* 83 (3): 502–12.

Pesaran, Bijan, Matthew J. Nelson, and Richard A. Andersen. 2006. "Dorsal Premotor Neurons Encode the Relative Position of the Hand, Eye, and Goal during Reach Planning." *Neuron* 51 (1): 125–34. doi:10.1016/j.neuron.2006.05.025.

Peters, Andrew J., Simon X. Chen, and Takaki Komiyama. 2014. "Emergence of Reproducible Spatiotemporal Activity during Motor Learning." *Nature* 510 (7504): 263–67. doi:10.1038/nature13235.

Petrides, M. 1986. "The Effect of Periarcuate Lesions in the Monkey on the Performance of Symmetrically and Asymmetrically Reinforced Visual and Auditory Go, No- Go Tasks." *The Journal of Neuroscience* 6 (7): 2054–63.

Pfurtscheller, Gert, and F.H. Lopes da Silva. 1999. "Functional Meaning of Event-Related Desynchronization (ERD) and Synchronization (ERS)." In *Handbook of Electroencephalography and Clinical Neurophysiology*, 6:51–56. Amsterdam: Elsevier.

Pfurtscheller, Gert, and Christa Neuper. 1997. "Motor Imagery Activates Primary Sensorimotor Area in Humans." *Neuroscience Letters* 239 (2–3): 65–68. doi:10.1016/S0304-3940(97)00889-6.

Pfurtscheller, G., B. Graimann, J.E. Huggins, S.P. Levine, and L.A. Schuh. 2003. "Spatiotemporal Patterns of Beta Desynchronization and Gamma Synchronization in Corticographic Data during Self-Paced Movement." *Clinical Neurophysiology* 114 (7): 1226–36. doi:10.1016/S1388-2457(03)00067-1.

Quinn, Christopher J., Todd P. Coleman, Negar Kiyavash, and Nicholas G. Hatsopoulos. 2011. "Estimating the Directed Information to Infer Causal Relationships in Ensemble Neural Spike Train Recordings." *Journal of Computational Neuroscience* 30 (1): 17–44. doi:10.1007/s10827-010-0247-2.

Raos, Vassilis, Gianfranco Franchi, Vittorio Gallese, and Leonardo Fogassi. 2003. "Somatotopic Organization of the Lateral Part of Area F2 (Dorsal Premotor Cortex) of the Macaque Monkey." *Journal of Neurophysiology* 89 (3): 1503–18. doi:10.1152/jn.00661.2002.

Raos, Vassilis, Maria-Alessandra Umiltá, Vittorio Gallese, and Leonardo Fogassi. 2004. "Functional Properties of Grasping-Related Neurons in the Dorsal Premotor Area F2 of the Macaque Monkey." *Journal of Neurophysiology* 92 (4): 1990–2002. doi:10.1152/jn.00154.2004.

Rathelot, Jean-Alban, and Peter L. Strick. 2006. "Muscle Representation in the Macaque Motor Cortex: An Anatomical Perspective." *Proceedings of the National Academy of Sciences* 103 (21): 8257–62. doi:10.1073/pnas.0602933103.

Reimer, Jacob, and Nicholas G. Hatsopoulos. 2009. "The Problem of Parametric Neural Coding in the Motor System." In *Progress in Motor Control*, edited by Dagmar Sternad, 243–59. Advances in Experimental Medicine and Biology 629. Springer US. http://link.springer.com/chapter/10.1007/978-0-387-77064-2_12.

Rice, John A. 1995. *Mathematical Statistics and Data Analysis*. Duxbury Press.

Rickert, Jörn, Simone Cardoso de Oliveira, Eilon Vaadia, Ad Aertsen, Stefan Rotter, and Carsten Mehring. 2005. "Encoding of Movement Direction in Different Frequency Ranges of Motor Cortical Local Field Potentials." *The Journal of Neuroscience* 25 (39): 8815–24. doi:10.1523/JNEUROSCI.0816-05.2005.

Riehle, Alexa. 2005. "Preparation for Action: One of the Key Functions of Motor Cortex." *Motor Cortex in Voluntary Movements: A Distributed System for Distributed Functions*.

Riehle, Alexa, Sarah Wirtssohn, Sonja Grün, and Thomas Brochier. 2013. "Mapping the Spatio-Temporal Structure of Motor Cortical LFP and Spiking Activities during Reach-to-Grasp Movements." *Frontiers in Neural Circuits* 7. doi:10.3389/fncir.2013.00048.

Riehle, A., and J. Requin. 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): 534–49.

Rigotti, Mattia, Omri Barak, Melissa R. Warden, Xiao-Jing Wang, Nathaniel D. Daw, Earl K. Miller, and Stefano Fusi. 2013. “The Importance of Mixed Selectivity in Complex Cognitive Tasks.” *Nature* 497 (7451): 585–90. doi:10.1038/nature12160.

Rizzolatti, G, R Camarda, L Fogassi, M Gentilucci, G Luppino, and M Matelli. 1988. “Functional Organization of Inferior Area 6 in the Macaque Monkey. II. Area F5 and the Control of Distal Movements.” *Experimental Brain Research. Experimentelle Hirnforschung. Expérimentation Cérébrale* 71 (3): 491–507.

Rizzolatti, Giacomo, Cristiana Scandolara, Massimo Matelli, and Maurizio Gentilucci. 1981. “Afferent Properties of Periarcuate Neurons in Macaque Monkeys. I. Somatosensory Responses.” *Behavioural Brain Research* 2 (2): 125–46. doi:10.1016/0166-4328(81)90052-8.

Rizzolatti, G., M. Matelli, and G. Pavesi. 1983. “Deficits in Attention and Movement Following the Removal of Postarcuate (area 6) and Prearcuate (area 8) Cortex in Macaque Monkeys.” *Brain* 106 (3): 655–73. doi:10.1093/brain/106.3.655.

Rizzolatti, G., C. Scandolara, M. Gentilucci, and R. Camarda. 1981. “Response Properties and Behavioral Modulation of ‘mouth’ Neurons of the Postarcuate Cortex (area 6) in Macaque Monkeys.” *Brain Research* 225 (2): 421–24. doi:10.1016/0006-8993(81)90847-7.

Roopun, Anita K., Steven J. Middleton, Mark O. Cunningham, Fiona E. N. LeBeau, Andrea Bibbig, Miles A. Whittington, and Roger D. Traub. 2006. “A beta2-Frequency (20–30 Hz) Oscillation in Nonsynaptic Networks of Somatosensory Cortex.” *Proceedings of the National Academy of Sciences* 103 (42): 15646–50. doi:10.1073/pnas.0607443103.

Rubino, Doug, Kay A. Robbins, and Nicholas G. Hatsopoulos. 2006. “Propagating Waves Mediate Information Transfer in the Motor Cortex.” *Nature Neuroscience* 9 (12): 1549–57. doi:10.1038/nn1802.

Sadtler, Patrick T., Kristin M. Quick, Matthew D. Golub, Steven M. Chase, Stephen I. Ryu, Elizabeth C. Tyler-Kabara, Byron M. Yu, and Aaron P. Batista. 2014. “Neural Constraints on Learning.” *Nature* 512 (7515): 423–26. doi:10.1038/nature13665.

Saleh, Maryam, Kazutaka Takahashi, Yali Amit, and Nicholas G. Hatsopoulos. 2010. “Encoding of Coordinated Grasp Trajectories in Primary Motor Cortex.” *The*

Journal of Neuroscience 30 (50): 17079–90. doi:10.1523/JNEUROSCI.2558-10.2010.

Saleh, Maryam, Kazutaka Takahashi, and Nicholas G. Hatsopoulos. 2012. “Encoding of Coordinated Reach and Grasp Trajectories in Primary Motor Cortex.” *The Journal of Neuroscience* 32 (4): 1220–32. doi:10.1523/JNEUROSCI.2438-11.2012.

Sanchez, J.C., J.M. Carmena, M.A. Lebedev, M.A.L. Nicolelis, J.G. Harris, and J.C. Principe. 2004. “Ascertaining the Importance of Neurons to Develop Better Brain-Machine Interfaces.” *IEEE Transactions on Biomedical Engineering* 51 (6): 943–53. doi:10.1109/TBME.2004.827061.

Sanchez, J.C., Sung-Phil Kim, D. Erdogmus, Yadunandana N. Rao, J.C. Principe, J. Wessberg, and M. Nicolelis. 2002. “Input-Output Mapping Performance of Linear and Nonlinear Models for Estimating Hand Trajectories from Cortical Neuronal Firing Patterns.” In *Proceedings of the 2002 12th IEEE Workshop on Neural Networks for Signal Processing, 2002*, 139–48. doi:10.1109/NNSP.2002.1030025.

Sanes, J N, and J P Donoghue. 1993. “Oscillations in Local Field Potentials of the Primate Motor Cortex during Voluntary Movement.” *Proceedings of the National Academy of Sciences of the United States of America* 90 (10): 4470–74.

Sanes, J. N., J. P. Donoghue, V. Thangaraj, R. R. Edelman, and S. Warach. 1995. “Shared Neural Substrates Controlling Hand Movements in Human Motor Cortex.” *Science (New York, N.Y.)* 268 (5218): 1775–77.

Sanides, F. 1964. “THE CYTO-MYELOARCHITECTURE OF THE HUMAN FRONTAL LOBE AND ITS RELATION TO PHYLOGENETIC DIFFERENTIATION OF THE CEREBRAL CORTEX.” *Journal Für Hirnforschung* 7: 269–82.

———. n.d. “Functional Architecture of Motor and Sensory Cortices in Primates in the Light of a New Concept of Neocortex Evolution.” In *The Primate Brain: Advances in Primatology*, 1:137–208. New York, NY, USA: Appleton-Century Crofts.

Santello, Marco, and John F. Soechting. 1998. “Gradual Molding of the Hand to Object Contours.” *Journal of Neurophysiology* 79 (3): 1307–20.

Schieber, Marc H. 2001. “Constraints on Somatotopic Organization in the Primary Motor Cortex.” *Journal of Neurophysiology* 86 (5): 2125–43.

Schieber, M. H., and L. S. Hibbard. 1993. “How Somatotopic Is the Motor Cortex Hand Area?” *Science (New York, N.Y.)* 261 (5120): 489–92.

Schwartz, Andrew B. 2004. “Cortical Neural Prosthetics.” *Annual Review of Neuroscience* 27 (1): 487–507. doi:10.1146/annurev.neuro.27.070203.144233.

Scott, Stephen H. 1997. “Comparison of Onset Time and Magnitude of Activity for Proximal Arm Muscles and Motor Cortical Cells Before Reaching Movements.” *Journal of Neurophysiology* 77 (2): 1016–22.

Serruya, Mijail D., Nicholas G. Hatsopoulos, Liam Paninski, Matthew R. Fellows, and John P. Donoghue. 2002. “Brain-Machine Interface: Instant Neural Control of a Movement Signal.” *Nature* 416 (6877): 141–42. doi:10.1038/416141a.

Shenoy, Krishna V., Matthew T. Kaufman, Maneesh Sahani, and Mark M. Churchland. 2011. “A Dynamical Systems View of Motor Preparation: Implications for Neural Prosthetic System Design.” *Progress in Brain Research* 192: 33–58. doi:10.1016/B978-0-444-53355-5.00003-8.

Shenoy, Krishna V., Maneesh Sahani, and Mark M. Churchland. 2013. “Cortical Control of Arm Movements: A Dynamical Systems Perspective.” *Annual Review of Neuroscience* 36 (1): 337–59. doi:10.1146/annurev-neuro-062111-150509.

Shimazu, H., M. A. Maier, G. Cerri, P. A. Kirkwood, and R. N. Lemon. 2004. “Macaque Ventral Premotor Cortex Exerts Powerful Facilitation of Motor Cortex Outputs to Upper Limb Motoneurons.” *The Journal of Neuroscience* 24 (5): 1200–1211. doi:10.1523/JNEUROSCI.4731-03.2004.

Singhal, Girish, Vikram Aggarwal, Soumyadipta Acharya, Jose Aguayo, Jiping He, and Nitish Thakor. 2010. “Ensemble Fractional Sensitivity: A Quantitative Approach to Neuron Selection for Decoding Motor Tasks.” *Computational Intelligence and Neuroscience* 2010 (February): e648202. doi:10.1155/2010/648202.

Sommer, Martin, Aránzazu Alfaro, Milena Rummel, Sascha Speck, Nicolas Lang, Tobias Tings, and Walter Paulus. 2006. “Half Sine, Monophasic and Biphasic Transcranial Magnetic Stimulation of the Human Motor Cortex.” *Clinical Neurophysiology* 117 (4): 838–44. doi:10.1016/j.clinph.2005.10.029.

Spinks, Rachel L., Alexander Kraskov, Thomas Brochier, M. Alessandra Umilta, and Roger N. Lemon. 2008. “Selectivity for Grasp in Local Field Potential and Single Neuron Activity Recorded Simultaneously from M1 and F5 in the Awake Macaque Monkey.” *The Journal of Neuroscience* 28 (43): 10961–71. doi:10.1523/JNEUROSCI.1956-08.2008.

Stancák, Andrej, and Gert Pfurtscheller. 1995. “Desynchronization and Recovery of β Rhythms during Brisk and Slow Self-Paced Finger Movements in Man.” *Neuroscience Letters* 196 (1-2): 21–24. doi:10.1016/0304-3940(95)11827-J.

Stancák Jr., Andrej, and Gert Pfurtscheller. 1996. “Event-Related Desynchronisation of Central Beta-Rhythms during Brisk and Slow Self-Paced Finger Movements of Dominant and Nondominant Hand.” *Cognitive Brain Research* 4 (3): 171–83. doi:10.1016/S0926-6410(96)00031-6.

Stevenson, Ian H., and Konrad P. Kording. 2011. “How Advances in Neural Recording Affect Data Analysis.” *Nature Neuroscience* 14 (2): 139–42. doi:10.1038/nn.2731.

Stoney, S. D., W. D. Thompson, and H. Asanuma. 1968. “Excitation of Pyramidal Tract Cells by Intracortical Microstimulation: Effective Extent of Stimulating Current.” *Journal of Neurophysiology* 31 (5): 659–69.

Suminski, Aaron J., Philip Mardoum, Timothy P. Lillicrap, and Nicholas G. Hatsopoulos. 2015. “Temporal Evolution of Both Premotor and Motor Cortical Tuning Properties Reflect Changes in Limb Biomechanics.” *Journal of Neurophysiology*, February, jn.00486.2014. doi:10.1152/jn.00486.2014.

Suminski, Aaron J., Dennis C. Tkach, and Nicholas G. Hatsopoulos. 2009. “Exploiting Multiple Sensory Modalities in Brain-Machine Interfaces.” *Neural Networks: The Official Journal of the International Neural Network Society* 22 (9): 1224–34. doi:10.1016/j.neunet.2009.05.006.

Sussillo, David, Mark M. Churchland, Matthew T. Kaufman, and Krishna V. Shenoy. 2015. “A Neural Network That Finds a Naturalistic Solution for the Production of Muscle Activity.” *Nature Neuroscience* 18 (7): 1025–33. doi:10.1038/nn.4042.

Taira, M., and A. P. Georgopoulos. 1993. “Cortical Cell Types from Spike Trains.” *Neuroscience Research* 17 (1): 39–45. doi:10.1016/0168-0102(93)90027-N.

Takahashi, Kazutaka, Sanggyun Kim, Todd P. Coleman, Kevin A. Brown, Aaron J. Suminski, Matthew D. Best, and Nicholas G. Hatsopoulos. 2015. “Large-Scale Spatiotemporal Spike Patterning Consistent with Wave Propagation in Motor Cortex.” *Nature Communications* 6 (May): 7169. doi:10.1038/ncomms8169.

Takahashi, Kazutaka, Maryam Saleh, Richard D. Penn, and Nicholas G. Hatsopoulos. 2011. “Propagating Waves in Human Motor Cortex.” *Frontiers in Human Neuroscience* 5: 40. doi:10.3389/fnhum.2011.00040.

Todorov, E., and Z. Ghahramani. 2004. “Analysis of the Synergies Underlying Complex Hand Manipulation.” In *26th Annual International Conference of the IEEE Engineering in Medicine and Biology Society, 2004. IEMBS ’04*, 2:4637–40. doi:10.1109/IEMBS.2004.1404285.

Tombini, Mario, Filippo Zappasodi, Loredana Zollo, Giovanni Pellegrino, Giuseppe Cavallo, Franca Tecchio, Eugenio Guglielmelli, and Paolo M. Rossini. 2009.

"Brain Activity Preceding a 2D Manual Catching Task." *NeuroImage* 47 (4): 1735–46. doi:10.1016/j.neuroimage.2009.04.046.

Towe, A. L., and G. W. Harding. 1970. "Extracellular Microelectrode Sampling Bias." *Experimental Neurology* 29 (2): 366–81. doi:10.1016/0014-4886(70)90065-8.

Truccolo, Wilson, Leigh R. Hochberg, and John P. Donoghue. 2010. "Collective Dynamics in Human and Monkey Sensorimotor Cortex: Predicting Single Neuron Spikes." *Nature Neuroscience* 13 (1): 105–11. doi:10.1038/nn.2455.

Vaidya, Mukta, Konrad Kording, Maryam Saleh, Kazutaka Takahashi, and Nicholas G. Hatsopoulos. 2015. "Neural Coordination during Reach-to-Grasp." *Journal of Neurophysiology* 114 (3): 1827–36. doi:10.1152/jn.00349.2015.

Vargas-Irwin, Carlos E., Gregory Shakhnarovich, Payman Yadollahpour, John M. K. Mislow, Michael J. Black, and John P. Donoghue. 2010. "Decoding Complete Reach and Grasp Actions from Local Primary Motor Cortex Populations." *The Journal of Neuroscience* 30 (29): 9659–69. doi:10.1523/JNEUROSCI.5443-09.2010.

Vigneswaran, Ganesh, Alexander Kraskov, and Roger N. Lemon. 2011. "Large Identified Pyramidal Cells in Macaque Motor and Premotor Cortex Exhibit 'Thin Spikes': Implications for Cell Type Classification." *The Journal of Neuroscience* 31 (40): 14235–42. doi:10.1523/JNEUROSCI.3142-11.2011.

Waldert, Stephan, Hubert Preissl, Evariste Demandt, Christoph Braun, Niels Birbaumer, Ad Aertsen, and Carsten Mehring. 2008. "Hand Movement Direction Decoded from MEG and EEG." *The Journal of Neuroscience* 28 (4): 1000–1008. doi:10.1523/JNEUROSCI.5171-07.2008.

Warland, David K., Pamela Reinagel, and Markus Meister. 1997. "Decoding Visual Information From a Population of Retinal Ganglion Cells." *Journal of Neurophysiology* 78 (5): 2336–50.

Weinrich, M, and S P Wise. 1982. "The Premotor Cortex of the Monkey." *The Journal of Neuroscience: The Official Journal of the Society for Neuroscience* 2 (9): 1329–45.

Weinrich, M., S. P. Wise, and K.-H. Mauritz. 1984. "A Neurophysiological Study of the Premotor Cortex in the Rhesus Monkey." *Brain* 107 (2): 385–414. doi:10.1093/brain/107.2.385.

Westwick, David T., Eric A. Pohlmeyer, Sara A. Solla, Lee E. Miller, and Eric J. Perreault. 2006. "Identification of Multiple-Input Systems with Highly Coupled

Inputs: Application to EMG Prediction from Multiple Intracortical Electrodes." *Neural Computation* 18 (2): 329–55. doi:10.1162/089976606775093855.

Wilkinson, G. N., and C. E. Rogers. 1973. "Symbolic Description of Factorial Models for Analysis of Variance." *Journal of the Royal Statistical Society. Series C (Applied Statistics)* 22 (3): 392–99. doi:10.2307/2346786.

Wise, S. P. 1985. "The Primate Premotor Cortex: Past, Present, and Preparatory." *Annual Review of Neuroscience* 8 (1): 1–19. doi:10.1146/annurev.ne.08.030185.000245.

Wise, Steven P., Driss Boussaoud, Paul B. Johnson, and Roberto Caminiti. 1997. "PREMOTOR AND PARIETAL CORTEX: Corticocortical Connectivity and Combinatorial Computations1." *Annual Review of Neuroscience* 20 (1): 25–42. doi:10.1146/annurev.neuro.20.1.25.

Woolsey, C N, P H SETTLAGE, D R MEYER, W SENCER, T PINTO HAMUY, and A M TRAVIS. 1952. "Patterns of Localization in Precentral and 'Supplementary' Motor Areas and Their Relation to the Concept of a Premotor Area." *Research Publications - Association for Research in Nervous and Mental Disease* 30: 238–64.

Wu, Wei, Yun Gao, Elie Bienenstock, John P. Donoghue, and Michael J. Black. 2006. "Bayesian Population Decoding of Motor Cortical Activity Using a Kalman Filter." *Neural Computation* 18 (1): 80–118. doi:10.1162/089976606774841585.

Wu, Wei, and Nicholas Hatsopoulos. 2006. "Evidence against a Single Coordinate System Representation in the Motor Cortex." *Experimental Brain Research* 175 (2): 197–210. doi:10.1007/s00221-006-0556-x.

Yu, Byron M., John P. Cunningham, Gopal Santhanam, Stephen I. Ryu, Krishna V. Shenoy, and Maneesh Sahani. 2009. "Gaussian-Process Factor Analysis for Low-Dimensional Single-Trial Analysis of Neural Population Activity." *Journal of Neurophysiology* 102 (1): 614–35. doi:10.1152/jn.90941.2008.