THE UNIVERSITY OF CHICAGO

POPULATION DYNAMICS IN PRIMARY MOTOR CORTEX DURING COORDINATED REACH TO GRASP

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

MUKTA PRAVIN VAIDYA

CHICAGO, ILLINOIS

MARCH 2016

DEDICATION

To my mother, Swati Vaidya, my father, Pravin Vaidya, my brother, Bhaskar Vaidya, my aunt, Sangeeta Mehendale, my uncle, Ramkrishna Mehendale, & the entirety of my extended family who have been a constant source of love, encouragement, and support throughout my life.

If I have seen further, it is by standing on the shoulders of giants.

– Isaac Newton

To Bhaskar Vaidya, Thaddeus Cybulski, Kalpita Ainapure, and all of my friends, who have added vibrancy and color to my life, and have kept me curious and full of wonder.

There is a single light of science, and to brighten it anywhere is to brighten it everywhere.

– Isaac Asimov

To Nicholas Hatsopoulos, Jason MacLean, & Leslie Osborne, who made me a scientist.

Orolo was onto something; when I saw any of those kinds of beauty I knew I was alive, and not just in the sense that when I hit my thumb with a hammer I knew I was alive, but rather in the sense that I was partaking of something--something was passing through me that it was in my nature to be a part of.

– Neal Stephenson, **Anathem**

TABLE OF CONTENTS

LIST OF FIGURES

LIST OF TABLES

LIST OF EQUATIONS

ACKNOWLEDGEMENTS

Advisor

Nicho Hatsopoulous

Thesis Committee

Callum Ross

Jason MacLean

Leslie Osborne

Konrad Kording

Collaborators

Josh Southerland

Fagg Lab

Oweiss Lab

MacLean Lab

Kording Lab

Ross Lab

Bensmaia Lab

Hatlab

Maryam Saleh, Adam Dickey

Matt Best, Jeff Walker, Alex Rajan

Kelsey Shattuck, Suchin Gururangan

Aaron Suminski, Kazutaka Takahashi

Fritze Arce, Karthikeyan Balasubramaniam

Kevin Brown, Frank Willet, Phil Mardoum

Graham Smith, Kai Qian

Carrie Balcer, Josh Coles

Faculty

Stephanie Palmer

Wim VanDrongelen

Computational Neuroscience Community

Neuroscience Administration

ARC Staff

My Family & Friends

ABSTRACT

When reaching to grasp, we coordinate how we preshape the hand with how we move it. The temporal and spatial coordination of transport and grasp components during reach-to-grasp behavior has been studied extensively in psychophysics experiments (M. Jeannerod 1984; Haggard and Wing 1995). Little is known, however, on the role of neocortex in generating, driving, or modulating the coordination of such behavior. The results discussed in this thesis attempt to answer the question: how do neural ensembles in the primary motor cortex functionally interact to produce coordinated behavior? In the first study, we examined the interactions between reach- and grasp-related neuronal ensembles while monkeys reached-to-grasp a variety of different objects in different locations. By describing the dynamics of these two ensembles as trajectories in a low-dimensional state space, we examined their coupling in time. The development of this coordination has been studied in infants and children (Claes von Hofsten 1984; Kuhtz-Buschbeck, Stolze, Jöhnk, et al. 1998; Wimmers et al. 1998). As children age, their inter-joint coordination becomes more stereotyped (Claes von Hofsten 1984; Kuhtz-Buschbeck, Stolze, Jöhnk, et al. 1998). Less is known about the role of motor cortex in developing coordinated reach-to-grasp. Additionally, studies have shown that after amputation, the cortical area previously involved in the control of the lost limb undergoes reorganization (J. N. Sanes et al. 1988; J. N. Sanes, Suner, and Donoghue 1990; Schieber and Deuel 1997; Wu and Kaas 1999; Qi, Stepniewska, and Kaas 2000), but limited work has gone towards developing BMIs that use neurons that are not tuned for kinetic or kinematic features. In the second study, we probed the emergence of coordinated reach-to-grasp in macaques that were being taught to cortically control a robotic arm through operant conditioning, and probed the neural correlates of this emergence.

I. INTRODUCTION

Overview

The vast extent to which we can interact with the world around us has shaped our evolution as a species. Whether through language or through movement, this interaction occurs by way of our muscles, joints, and limbs. Human beings have an incredibly rich repertoire of motor behaviors that we take advantage of constantly, in addition to the ability to learn new ones through the plasticity afforded to us by our central nervous system. Nor do we limit ourselves to individual movements in isolation—how often do you find yourself moving just your elbow when making any sort of movement of the arm? How many tennis players can serve perfectly without keeping their eyes on the tennis ball? How many dancers are told to forget about controlling their facial expressions while doing pirouettes?

All of these movements require a remarkable degree of spatial and temporal coordination. Even primitive tool use in our ancestral history required coordination of many of the joints of the upper limb (Marzke 1997). The coordination involved in making complex and ethologically relevant movements has been an area of considerable interest and research; in particular, the coordination of the transport and grasp components of reach-to-grasp behavior has been studied extensively psychophysically. Little is known, however, on the role of neocortex in generating, driving, or modulating the coordination of such behavior. Given how much of a role the learning and execution of coordinated reach-to-grasp plays in both human history and in our daily lives, what could be more fundamental than understanding the role that neocortex, specifically primary motor cortex (M1), plays in this coordination?

Reach-to-Grasp: Psychophysics

There is an extensive body of psychophysics literature on the coordination of reach-to-grasp in adults. Coordination in the context of reach-to-grasp behavior refers to the spatiotemporal coupling of the proximal arm and the distal hand such that when the arm reaches the object, the hand shape matches the required configuration to grasp the object. One seminal study by Jeannerod observed such temporal coordination between the transport (reaching) and manipulation (grasping) components of prehension in the reach-to-grasp movements of seven right handed adults (M. Jeannerod 1984). The transport component was characterized by a fast velocity phase for the majority of the movement's time course, followed by a slow-velocity phase during the last quarter of the time course; this 3:1 ratio was maintained across distance. Preshaping of the hand occurred during transport, with the fingers initially stretched then closing in the appropriate configuration. The timing of finger closing correlated with the onset of the slow-velocity phase.

Another group found that the coordination observed between prehension and transport varied based on the type of prehension that subjects used. Gentilucci and colleagues conducted a study with eight subjects, looking at reach-to-grasp movements with different objects placed at different distances (Gentilucci et al. 1991). Objects were meant to require either prehension with the whole hand or prehension with the index finger and the thumb (precision grip). They found that the transport component was invariant to the type of grasp, whereas the temporal trajectory of prehension, as well as its temporal relationship with the transport component, changed depending on the type of grasp. When the object required precision grip, subjects reached the

maximum aperture earlier. In addition, they found that the temporal coupling between transport and prehension trajectories was stronger when the subjects were reaching to grasp objects not requiring precision grip.

Perturbation studies have been used to investigate important features of reach-to-grasp coordination and their robustness. In one study, Haggard and Wing examined the coordination between the transport component and aperture by observing the effects of aperture perturbations (Haggard and Wing 1995). An electric actuator would push or pull the subject's relevant arm during the course of movement in a quarter of the trials. They found that subjects compensated for this perturbation by adjusting their aperture accordingly to preserve the stereotypical relationship between the spatial trajectories of the transport component and aperture.

These three seminal studies demonstrated that there exists spatiotemporal coordination between the reach (transport) and the grasp (manipulation, aperture) components during reach-to-grasp movements, established that this coordination varies based on the type of prehension, and indicated that such coordination is robust to perturbations. This type of coordination would suggest that, at some level, control of aperture requires information about the state of transport and vice versa.

Reach-to-Grasp: Cortex

Traditionally, reaching and grasping are thought to be mediated by separate parieto-premotor pathways due to the fact that reaching and grasping involve the sensorimotor transformation of different kinds of environmental information (M Jeannerod 1999; Marc Jeannerod 1988;

Krakauer and Ghez 1991). Reaching requires visual information about the target location, whereas the kinematics of grasping largely depend on intrinsic properties of the object information and not the location of the object. As discussed earlier, the kinematic trajectories of reaching and grasping are highly coupled, but this coupling has been posited to occur between separated informational streams prior to primary motor cortex. This paradigm is posited to be analogous to the division between dorsal and ventral informational streams in visual processing, where the ventral stream is associated with object recognition and form representation, while the dorsal stream is associated with the processing of spatial information. Analogous to the cortical areas higher up in the visual hierarchy, primary motor cortex would have to coordinate these two allegedly independent streams of information for the production of coordinated reach-to-grasp behavior.

In motor neuroscience, the "dorsal pathway" is associated with reaching, and is involved in transforming visual information about target location in extrapersonal space into the direction of a reaching movement (Krakauer and Ghez 1991). Parameters for reaching movement depend on locations of the target relative to body, shoulder, and hand. Reaching requires that visual information about target location and the position of the upper limb be used to specify critical features of the upcoming arm movement. Research suggests that the memory of extrapersonal space is stored within an eye-centered frame of reference (Galletti et al. 2003; Batista et al. 1999; Cohen and Andersen 2002; Cohen and Andersen 2000). In particular, the medial intraparietal area (MIP) is thought be involved with encoding the location of objects with respect to this frame. MIP, together with the parieto-occipital cortex (PO), which contains both a retinotopically organized visual area and a visual/somatosensory area that is involved in form, motion, and

space processing (Galletti et al. 2003), is called the Parietal Reach Region (PRR) ((Cohen and Andersen 2002; Batista et al. 1999; Cohen and Andersen 2000). Posterior parietal cortex projects to premotor and motor cortices. There are also direct projections from parietal areas to dorsal premotor area (PMd) (Caminiti et al. 1999; Shipp, Blanton, and Zeki 1998; Johnson and Ferraina 1996; Steven P. Wise et al. 1997). During reaching, neurons in PRR code for direction of movement but discharge later than dorsal premotor neurons to which they are connected. These neurons could monitor ongoing movements and improve the planning and execution of subsequent reaches by premotor areas. PMd contains neurons that are active during both reaching movements and the preparation for reaching movements and exhibits tuning for the direction of reaching (Cisek and Kalaska 2005; Messier and Kalaska 2000; Kurata 1993; Gentilucci et al. 1988; Fu et al. 1995; Crammond and Kalaska 1989; S. P. Wise and Mauritz 1985; Weinrich and Wise 1982).

The "ventral pathway" is associated with transforming visual object-related information, like shape or size, into commands for grasping (Krakauer and Ghez 1991). Projections connect the anterior intraparietal areas (AIP) with the ventral premotor area (PMv) (Luppino et al. 1999; Matelli et al. 1986). AIP has neurons that are selective for object properties such as shape, size, and orientation, as well as neurons that are responsive to manipulation of the hand (Akira Murata et al. 2000), and has been posited to be involved in transforming visual object information into motor-related information. Neurons in PMv are selective for different hand manipulations, such as precision grip or power grip, and tend to be active throughout reach-to-grasp movements (A. Murata et al. 1997; Akira Murata et al. 2000; Rizzolatti et al. 1988).

More recently, this paradigm for separate parieto-premotor pathways for reaching and grasping has been called into question. Area PO, part of PRR, also contains information about object form and space, and is thought to play a role in grasping as well (Galletti et al. 2003). In general the dorsal visual stream connects to both pathways. Each premotor area receives input from posterior parietal cortex; in addition, there are dense connections between premotor areas (Dum and Strick 2005). To this effect, one study recorded from neurons in both primary motor cortex (M1) and PMv while monkeys made naturalistic reach to grasp movements. They assessed the ability of spiking activity and local field potentials in both areas in decoding 3D arm end point and grip aperture kinematics; they found that neither MI nor PMv showed a bias towards a reach or grasp decode (Bansal et al. 2012).

Regardless of whether separate, independent pathways exist in the processing of reach-related and grasp-related information, it is clear that primary motor cortex must integrate and coordinate reach-related and grasp-related information from multiple streams and cortical areas for the production of reach-to-grasp behavior. There are projections from both premotor areas to primary motor cortex, which is responsible for elaborating and executing both reach and grasp movements (Tokuno and Tanji 1993; Krakauer and Ghez 1991). In addition, premotor and motor cortex also receive input from the basal ganglia and cerebellum via nuclei in the ventrolateral thalamus (many reciprocal connections called cortico-subcortical loops) (Middleton and Strick 2000).

Within primary motor cortex there is a broad spectrum of functional representations. Ensembles of neurons involved in the processing of different types of information need to be coordinated for

the production of coordinated motor behaviors such as reach-to-grasp. Primary motor cortex (M1) is composed of the forelimb, hindlimb, trunk, and orofacial regions, the somatotopy of which is described in detail in the section **Somatotopic Organization of Primary Motor Cortex**. MI contains both kinematically and kinetically selective neurons, including for example, neurons that are tuned for direction of movement, muscle force, muscle synergies, kinetic synergies, joint velocities, or joint positions (Krakauer and Ghez 1991; Kakei, Hoffman, and Strick 1999; Holdefer and Miller 2002; Scott 2003; Georgopoulos et al. 1982; Georgopoulos, Schwartz, and Kettner 1986; Werner, Bauswein, and Fromm 1991; Paninski et al. 2004). Most neurons in primary motor cortex become active only shortly before and during movement (Krakauer and Ghez 1991).

Recent work has shown that neurons in primary motor cortex encode both reaching and grasping related kinematic trajectories, as opposed to single point-in-time kinematic variables. In a study by Saleh, Takahashi, and Hatsopoulos, two female rhesus macaques were trained on a reach to grasp task (Saleh, Takahashi, and Hatsopoulos 2012). They simultaneously recorded from multiple single units in primary motor cortex (MI) while the monkeys performed reach-to-grasp movements to four different objects presented in seven different locations by a robot. A 6-camera, Vicon motion capture system tracked the kinematics of a set of infrared reflective markers placed on the arm, wrist and fingers from which the kinematics of 17 joints and wrist position were reconstructed. Saleh and colleagues created generalized linear models using kinematic features, including joint angles and velocities, along with wrist position and velocity, as well as a neuron's own spike history, to predict the firing activity of single neurons. By using trajectories of kinematic features, instead of values of features at a single instance in time, the

models were able to better predict firing activity for most neurons. The same held true for using multiple features in a model instead of single features, as well as using velocity features over position features. In addition, this study found that although some neurons primarily encode either reach-related or grasp-related kinematic features, a plurality tended to encode both types of features.

In summary, primary motor cortex must coordinate reach-related and grasp-related information from multiple streams and cortical areas for the production of reach-to-grasp behavior. Given that neuronal ensembles within primary motor cortex encode both reaching and grasping related kinematic trajectories, studying how ensembles of neurons encoding reaching-related and grasping-related kinematics are coordinated can help shed light on the role of primary motor cortex in producing coordinated reach-to-grasp movements.

Cortical Coordination

Prior work has examined the cortical basis of coordinating reach-to-grasp, but has focused on single neuron responses or pair-wise correlations. It has been shown that M1 neurons modulate sequentially during a reach-to-touch task such that shoulder- and elbow-related cells (assessed via intracortical microstimulation effects on the same electrode sites from which the cells were recorded) begin firing on average about 60 ms before wrist and finger-related neurons consistent with the proximal-to-distal sequencing of muscle activation evident in this task as well as in many other multi-joint motor behaviors (Murphy, Wong, and Kwan 1985). In premotor cortex, pair-wise correlations between small groups of neurons have been shown to carry information about particular reaching and grasping combinations even though the constituent neurons' firing

rates only carried information about either reaching or grasping (Stark et al. 2008). This study speculates that correlations in primary motor cortex during movement could be even more informative in motor cortex, even though this was not experimentally verified. Neither of these studies has examined the dynamics of large neural ensembles during coordinated reach-to-grasp behavior.

The field of cortical motor neuroscience has recently been shifting towards the study of neural dynamics (Todorov and Jordan 2002; Scott 2004; Todorov 2004; Cheng and Sabes 2006; M. Churchland et al. 2012; Shenoy, Sahani, and Churchland 2013). A standard approach in analyzing the dynamics of neural ensembles uses state space-methods, where a high dimensional neural signal is described by a trajectory through a low-dimensional space. This method, if extended to multiple neural ensembles, could allow us to ask how the trajectories of multiple neural populations, such as reach- and grasp- related populations, coordinate with one another, broadening the literature beyond single-neuron or pair-wise correlation analyses.

This type of approach could also speak to work examining phase-locking and coherence in local field potential oscillations during cognition and coordinated behavior (Sternad, Turvey, and Schmidt 1992; Kelso 1994; Bressler and Kelso 2001). This school of inquiry focuses on the question of how non-linear coupling among some group of components can give rise to a variety of complex behaviors. This involves both the identification of key variables, "defined as a functional ordering among interacting components", and the dynamics of these variables, or "the rules that govern the stability and change of coordination patterns and the non-linear components

that give rise to them" (Bressler and Kelso 2001). The substrate for this coordination posited by researchers has been primarily local field potentials (Bressler and Kelso 2001).

Thus, the field is primed for new approaches to studying the role of primary motor cortex in producing coordinated reach-to-grasp movements—in particular, investigating the dynamics between ensembles of reach-related and grasp-related neurons during these movements (see **Specific Aim 1**). Similarly, little is known about the role of primary motor cortex in developing coordinated reach-to-grasp behavior; to date, there have been no cortical studies in human subjects or in any viable animal models investigating this development. However, as with fully-developed reach-to-grasp behavior, there is a psychophysics literature examining the development of coordinated reach-to-grasp behavior in infants and in children.

Reach-to-Grasp: Development of Coordination in Infants

Humans develop the ability to reach-to-grasp in infancy, and refine these movements through early adolescence (Claes von Hofsten 1984; Kuhtz-Buschbeck, Stolze, Jöhnk, et al. 1998). Several studies have examined when and how reach-to-grasp behavior emerges in infants. Neonates tend to open their hand while extending their arm forwards (Claes von Hofsten 1984). In fact, at this point, infants only do the opposite, extend their arm with a closed fist, about 8 percent of the time (DiFranco, Muir, and Dodwell 1978). This coupling only becomes less prevalent around the age of 2 months; at this point, infants tend to fist while extending their arms (Claes von Hofsten 1984).

A common framework that has been used over the past century to analyze human motor development is to study specific "milestones" (Wimmers et al. 1998). For example, infants typically exhibit reaching behavior at about 3 months of age. By this point, they can even make predictive movements for fast moving objects (Claes von Hofsten 1981). Infants typically exhibit reaching-to-grasp behavior around about 5 months of age (Wimmers et al. 1998). The occurrence of such milestones is thought to reflect some maturation point of the motor system. Less well understood is how infants transition between these milestones or what neural processes drive the emergence of new motor behaviors.

Researchers have proposed that the stability of a motor behavior is conducive to the emergence of new motor behaviors. In one longitudinal study, infants were examined every week from the age of 8 weeks to the age of 24 weeks (Wimmers et al. 1998). In each trial, a stick was presented to the infant; the stick was at a distance of about 2/3 of the length of the infant's arm. Infants were filmed as they were presented with and moved their hands towards the stick; they could either reach without grasping, or reach-to-grasp. Trials were then divided into the reaching without grasping and reaching with grasping categories; the reaching without grasping category included both trials in which the infant came within 5 cm of the object and trials in which the infant reached out and touched the object without grasping. The behavior underwent a distinct transition from when the former was predominant to when the latter was predominant. The mean first passage time, defined as the average amount of time following the onset of trial that it took the infant to first touch the object or come within 5 cm of it, tended to decrease until right after the transition to a higher percentage of reach-to-grasp trials. After this transition, mean first passage time tended to remain constant. Thus, these results lend support to the idea that the

stability of a certain motor activity is conducive to the emergence of new motor activities (Wimmers et al. 1998).

Other work has focused on studying the nature of prehension once infants begin to exhibit reach-to-grasp movements. One study found that infants have the capacity to precisely time and coordinate catching a moving object (Claes von Hofsten 1983). Another longitudinal study found that infants concurrently master the ability to reach for stationary and moving objects (C. von Hofsten and Lindhagen 1979). For this study, infants were scored on one of four behaviors while reaching: grasp and hold the object, grasp and let go of the object within a second, touch the object, and miss the object. Infants were seen at regular intervals from between 12 and 24 weeks of age to 30 weeks of age. The number of occurrences of the first two categories when the infant actually grasped increased with age, while the frequency of just touching remained roughly steady, and the number of misses decreased. In addition, infants were more likely to grasp and hold for objects that were closer, and more likely to grasp and let go of objects moving at slower speeds. By the time they were 18 weeks old, most of the reaches involved a grasp. In fact, several of the infants were able to catch an object moving at 30 centimeter per second at this age. Most infants moved their hand directly to the expected meeting point with the object.

Although infants are able to reach-to-grasp around 18 months of age, they fine-tune their prehension as they get older. Another longitudinal study by von Hofsten and Fazal-Zandy observed 15 infants every four weeks from when they were 18 weeks of age to 24 weeks of age (Claes von Hofsten and Fazel-Zandy 1984); these infants were presented with both horizontal and vertical rods. They found signs of infants trying to adjust their hand to the appropriate

orientation even at 18 weeks of age, but without success. The infants showed rapid improvements over the course of the study, indicating that although the development of prehension is accessible at an early age, these movements become more fine-tuned with age. In fact, the refinement of reach-to-grasp continues through the course of childhood.

Reach-to-Grasp: Refinement in Children

Although reaching-to-grasp is a behavior that is learned in infancy, studies have shown that the coordination of reaching-to-grasp continues to develop through childhood. One study found that the kinematic profiles of reach-to-grasp behavior becomes more stereotyped with age (Kuhtz-Buschbeck, Stolze, Jöhnk, et al. 1998). In this study, 54 children between the ages of 4 and 12 had to reach-to-grasp a cylindrical object with precision grip with their dominant hand. Both the trajectory of the reaching hand and the trajectory of the finger aperture were measured for each trial. Younger children were found to open their hands earlier in order to grasp objects. Figure 1.1 illustrates kinematics profiles of the hand velocity and grip aperture for each of the three age groups; these profiles become more stereotyped with age. A similar study of children who were 4 or 5 years old found that the reach trajectories of children reaching-to-grasp became straighter with age (Kuhtz-Buschbeck et al. 1999).

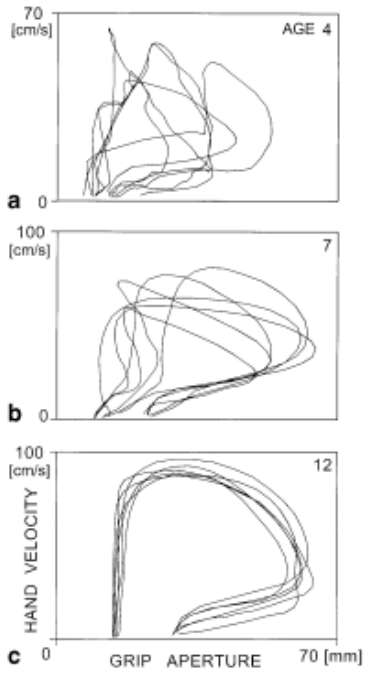

Figure 1.1 Prehension in Children Kinematic profiles of prehension at the age of 4, 7, and 12 years. The hand velocity is plotted against grip aperture in three children of difference ages. Six trials are superimposed for each subject." (SOURCE: Kuhtz-Buschbeck, Stolze, Jöhnk, et al. 1998)

Similarly, interjoint coordination during reach-to-grasp becomes more fine-tuned with age (Schneiberg et al. 2002). One study examined the development of stable, highly stereotyped kinematic trajectories that were characteristic of adults while reaching-to-grasp. This study was conducted on thirty-eight children, between the ages of 4 and 11, as well as nine adults. Subjects had to reach for cones that were placed in front of them at three different distances. Similar to the previous study, they found increased variability in younger children and that kinematic trajectories got more stereotyped with age. In addition, they found that interjoint coordination became more stereotyped, for example, between shoulder "horizontal adduction" and elbow extension.

Not only is reach-to-grasp behavior not fully developed by the age of 6 or 7—there are marked differences in the strategies that children and adults use to reach-to-grasp. One study found that children opened their grip wider than adults (Kuhtz-Buschbeck et al. 1999). Another study varied three parameters, the distance of the object, the size of the object, and the presence of visual feedback, and examined how each parameter affected reach-to-grasp movements in both children and adults (Kuhtz-Buschbeck, Stolze, Boczek-Funcke, et al. 1998). The experimental setup was scaled to reflect the differing body proportions of a child verses an adult. They found that the peak transport velocities scaled with object distance similarly in both children and adults. When it came to prehension, there were a number of differences between the two groups. The deceleration phase was shorter in children, who tended to open their hands with larger apertures than adults, consistent with results from other studies. Children were unable to scale their aperture successfully with object size without the presence of visual feedback, nor could they scale aperture with object distance regardless of whether visual feedback was present. Similar to with other studies, Kuhtz-Buschbeck and colleagues also found that reach-to-grasp trajectories were more variable in children than in adults. Another study also compared the importance of visual feedback in reach-to-grasp behavior in both adults and children (Zoia et al. 2006). This study found that children had a larger maximal aperture for the grasp, a longer duration for the transport component, and a longer deceleration time, the last of which stands in direct contrast to the Kuhtz Buschbeck study (Kuhtz-Buschbeck, Stolze, Boczek-Funcke, et al. 1998).

In contrast to this body of work on the development of coordination of reach-to-grasp behavior in infants and children in the psychophysics realm, little is known on the role of neocortex in producing, elaborating, or modulating the development of such coordination. Invasive studies

have not been conducted on human or non-human primates during the development of reach-to-grasp. As such, studying the learning of coordinated behavior in the context of a brain-machine interface could provide a unique window into the development of reach-to-grasp.

Learning in the Context of a Brain Machine Interface

Traditionally in the field of brain-machine interfaces, researchers have used kinematic or kinetic decoders to use neurons that are already tuned for particular kinematic or kinetic features, such as direction of motion to control those features. However, since somatotopic reorganization takes place after peripheral nerve damage or injury (see section **Somatotopic Reorganization of Primary Motor Cortex due to Injury**), the reliance on finding neurons already tuned for behavioral variables that we wish to control becomes a hindrance. Single and multi-unit activity recorded from arrays implanted into reorganized motor cortex may not be expected to be reliably "pre-tuned" for kinematic or kinetic features to nearly the same degree. As such, more recent work has started to focus on learning in the context of a brain-machine interface.

Studies have shown that subjects can be taught to modulate single neurons to control a brain machine interface through the use of operant conditioning. Seminal work on this subject has come out of the lab of Eberhard Fetz (E. E. Fetz 1969; E. E. Fetz and Baker 1973; Eberhard E. Fetz 2007). Monkeys were operantly conditioned to control the firing rates of individual motor cortical neurons with the use of biofeedback. Every time the monkey was able to drive a meter arm that reflected the firing activity of a single neuron, it was rewarded with food pellets. More recently, one study showed that subjects can learn to modulate groups of neurons as well. Ganguly and Carmena trained an algorithm to decode the direction of motion from neurons being

recorded from the primary motor cortex of a rhesus macaque performing a center-out task (Ganguly and Carmena 2009). After the monkey had learnt how to do this task successfully, they shuffled the decoder coefficients. The monkey was able to learn how to use the decoder with the new decoder coefficients after several days of practice. In other words, given enough practice, the monkey was able to use a new, perturbed mapping between neurons and direction of motion.

One common confound that makes this learning paradigm difficult to use over the course of multiple days is neural instability. That is, signals from neurons that are used for decoding on one day might have shifted or not be present on the next day, thus impairing both learning and performance. Several studies have attempted to address this issue through the use of adaptive decoding-- updating or recalibrating the decoder at least once on each day of training (Hochberg et al. 2012; Taylor, Tillery, and Schwartz 2002; Li et al. 2011). By doing so, these studies have been able to improve performance (Hochberg et al. 2012; Taylor, Tillery, and Schwartz 2002), improve the tuning of individual neurons (Taylor, Tillery, and Schwartz 2002), or prevent performance degradation (Li et al. 2011). However, since adaptive decoding involves shifting the neural mapping that subjects have to learn in order to control the brain machine interface every time that the decoder is recalibrated, it is not a useful paradigm for studying learning over the course of multiple days.

Recent work has attempted to move away from the adaptive decoding paradigm by examining the learning of decoder that is held fixed across days (Ganguly et al. 2011). On each day of experiment, neurons were checked and ensured to be stable; decoders were held stable for between 3 to 6 days. This study examined the effects of learning a fixed decoder for a BMI task

on the population activity of neurons directly involved in the decoder as well as other stable neurons that were not involved in the decoder, termed indirect neurons (Ganguly et al. 2011). They found changes in the preferred direction of both direct and indirect neurons; however, they found an overall relative decrease in modulation of indirect neurons as compared with direct neurons. In addition, they found that these changes in modulation remained stable as the macaques learned the decoder.

The 2011 Carmena study was only able to examine learning of the brain machine interface over the course of 3 to 6 days, and focused on changes in the properties of single neurons (firing rate, preferred direction of motion). Our work aims to address these particular shortcomings by examining the learning of a fixed decoder over the course of multiple weeks instead of multiple days, and focuses on how neurons coordinate with each other to drive the emergence of coordinated behavior in a brain machine interface task (see **Aims 2 & 3**). This paradigm provides a unique window into the study of functional plasticity at the level of motor cortex. By establishing a direct, causal connection between primary motor cortex and the device that subjects learn to control, bypassing spinal cord and neuromuscular dynamics, we can infer that any learning at the behavioral level is directly due to changes in primary motor cortical activity. Given that somatotopic reorganization takes place after peripheral nerve damage or injury (see section **Somatotopic Reorganization of Primary Motor Cortex due to Injury**), this paradigm could be particularly beneficial for patients that have been the recipients of amputation.

In Context: An Amputee Model

The ability to train subjects to control individual or groups of neurons for the purpose of executing a novel behavior through the use of operant conditioning presents a powerful tool in multiple regards. As mentioned before, this paradigm could be used to study the development of coordination in a completely novel behavior—something which the field has insofar not been able to do—concurrently at the level of neurons and the behavior itself. There is another domain in which this approach is imminently useful. This paradigm for training subjects to use a "new" set of neurons to reach-to-grasp a robotic arm could be beneficial for patients who have been the recipients of amputation. Just to put this into perspective, about 185,000 amputations are conducted in the United States each year-- there are almost 2 million people living in the United States who have been the recipients of amputations (Ziegler-Graham et al. 2008). Over 5,000 service members in the U.S. Armed Forces had traumatic amputations in the timespan of 2000 to 2011 (Armed Forces Health Surveillance Center (AFHSC) 2012). Much work has gone into showing that after amputation or peripheral nerve damage, there is somatotopic remapping in the cortical area previously involved in the control of the lost limb (J. N. Sanes et al. 1988; J. N. Sanes, Suner, and Donoghue 1990; Schieber and Deuel 1997; Wu and Kaas 1999; Qi, Stepniewska, and Kaas 2000).

Somatotopic Organization of Primary Motor Cortex

Since the mid-twentieth century, a great deal of work has gone into the elucidating the somatotopic organization of primary motor cortex. In 1952, Woolsey and colleagues published on an overall medial to lateral topography of the leg/hindlimb, arm/forelimb, and head and face (Woolsey et al. 1952). More recent studies from the 1970s, 1980s, and 1990s have upheld the broader somatotopic map of the placement of the forelimb, hindlimb, and orofacial regions, but

only upholding the idea of a homunculus in this very rough sense (Jerome N. Sanes and Donoghue 2000); most are stimulation experiments. One type of stimulation commonly used is intracortical microstimulation (ICMS), a technique first developed by Asanuma and colleagues in 1968, when they measured thresholds for stimulation and the extent of excitation of pyramidal tract neurons in cats (Stoney, Thompson, and Asanuma 1968). Researchers have used this method to map out causal relationships between focal electrical stimulation via microelectrodes in primary motor cortex and movement or muscles. The focal site of stimulation is stepped in a principled fashion across the surface of primary motor cortex in order to uncover underlying functional patterns of organization (Jerome N. Sanes and Donoghue 2000).

Others have used ICMS paired with electromyography (EMG) to measure similar causal relationships. One study used electromyography (EMG) to record evoked muscle activity from ICMS stimulation at 433 sites in layer 5 of the forelimb area in primary motor cortex in anesthetized squirrel monkeys (Sessle and Wiesendanger 1982). This study found that all of the muscles recorded from in the hand, forearm, and arm could be activated by ICMS, and that most often, multiple, separate muscles were activated from any given site. Another interesting result was that increasing the intensity of stimulation yielded additional areas of activation. In addition, they did not find any sort of orderly topographic organization of forelimb joints or muscles, except for a rough posterior to anterior trend for generally distal to proximal muscles, although the whole spectrum of muscles could be activated from "broadly distributed and overlapping areas" (Sessle and Wiesendanger 1982). A decade later, Nudo and colleagues published a study examining primary motor cortex topography in six adult squirrel monkeys using ICMS to evoke distal forelimb movement, comparing inter-animal variability with hand preference during a

motor task requiring the use of only one digit (Nudo et al. 1992). As with the previous study, the topography of primary motor cortex was mosaic like, and did not reflect an orderly point-to-point or continuous organization.

Other studies have yielded similar results supporting a distributed network representation in primary motor cortex. One study used ICMS in the motor cortex of owl monkeys to evoke a wide range of body part movements, spanning the spectrum of tail to mouth (Gould et al. 1986). This study found a rough somatotopic mediolateral organization of tail to mouth movements; however, the same movement could be elicited at multiple sites. Unlike any nested or strict orderly homunculus-like configuration, they stated instead that "M-I is more adequately described as a mosaic of regions, each representing movements of a restricted part of the body, with multiple representations of movements that tend to be somatotopically related" (Gould et al. 1986). Another study used ICMS and EMG to measure evoked responses in primary face motor cortex in crab-eating macaques (Huang, Hiraba, and Sessle 1989). This study also found that multiple, separated cortical sites could evoke the same type of movement.

Imaging studies using Positron Emission Tomography (PET) and Magnetic Resonance Imaging (MRI) have also shown overlapping activation patterns both between different types of proximal and distal movement as well as between proximal and distal movements. However, studies have also shown that after amputation or peripheral nerve damage, there is somatotopic remapping in the cortical area previously involved in the control of the lost limb (J. N. Sanes et al. 1988; J. N. Sanes, Suner, and Donoghue 1990; Schieber and Deuel 1997; Wu and Kaas 1999; Qi,

Stepniewska, and Kaas 2000). Understanding this reorganization is key for the development of brain machines for patients who have been the recipients of amputation.

Somatotopic Reorganization of Primary Motor Cortex due to Injury

Several studies have shown that chronic deafferentation induces reorganization in primary motor cortex, and that the stimulation thresholds for inducing a movement response are similar between normal and reorganized motor cortex. One group published a study on the reorganization of representations in motor cortex following peripheral nerve injury in rats, specifically facial nerve transection (J. N. Sanes et al. 1988). The branches that were lesioned only carried motor axons for innervating muscles controlling the vibrissae. By comparing primary motor cortex mapping in normal rats with the ones with facial nerve lesions, they were able to determine that the forelimb, the eye, and eyelid representations expanded into the area that typically mapped to the vibrissae within hours of the transection. A couple of years later, this group published another study on the reorganization of primary motor cortex in adult rats, looking at the effects of motor or mixed peripheral nerve lesions on long-term reorganization (J. N. Sanes, Suner, and Donoghue 1990). The rats received either forelimb amputation or a facial nerve transection from 1 week to 4 months before their reorganized motor cortex was mapped using ICMS, and compared with those of normal rats. In the rats with forelimb amputation, the general area that elicited shoulder movements expanded into the forelimb region. In the rats with facial nerve transection, the forelimb, eye, and eyelid area expanded into the vibrissae area, as in the previous study. This change in representation was apparent even after one week, and the stimulation thresholds for generating a movement response were comparable to normal thresholds.

Similar studies have been conducted in primates analyzing changes in representation after amputation. Schieber and Deuel published a study in macaque (Schieber and Deuel 1997); the macaque used in this study was a 15 years old whose right upper extremity was amputated at the shoulder joint two years prior. The size of the primary motor cortex upper extremity region for reorganized cortex contralateral to the amputated arm was comparable to its counterpart contralateral to the normal arm. However, throughout the reorganized primary motor cortex upper extremity region, the movements elicited were for the shoulder girdle and the stump, and required longer stimulus trains, higher amplitude current thresholds, or both. Qi and colleagues published a similar study in macaques that had received therapeutic amputations at different points in their development: two had been injured as adults, one as a juvenile, and one as in infant (Qi, Stepniewska, and Kaas 2000). The reorganized maps of in all four of the macaque showed the characteristic expansion of the remaining proximal body parts or joints and neighboring body regions, despite variation in when they received amputation.

Wu and Kaas published a similar study taking a novel approach to comparing somatotopic maps before and after amputation (Wu and Kaas 1999). This study also used ICMS, examining the reorganization of primary motor cortex in three squirrel monkeys and two galagos that had received prior therapeutic amputations, either forelimb or hindlimb, due to injury. Tracers injected into the lower cervical segments of the spinal cord helped identify corticospinal neurons. The distributions of the corticospinal neurons were used to help identify the location of the forelimb region prior to amputation; this rough somatotopic map was compared with the somatotopic map of reorganized primary motor cortex generated using ICMS. In each of the animals, primary motor cortex ipsilateral to the amputated limb was not any different from

normal; however, contralateral to the missing limb, stimulation elicited movements either from remaining proximal muscles or neighboring body regions (hindlimb, forelimb, or orofacial). Stimulation at every site resulted in some movement response, given enough current; on average, higher amplitude current pulses were required to elicit movement responses.

In summary, reorganization of motor maps in primary motor cortex is induced following peripheral nerve damage or limb amputation. This process involves an expansion of areas encoding for remaining proximal movements ipsilateral to the injury as well as neighboring gross body regions along the rough somatotopy into the area that previously mapped to the injured region. Thus, for patients that have been the recipients of amputations, traditional methods for training brain machine interfaces (BMIs) that necessitate finding neurons "pre-tuned" for a missing limb may fall quite short given that we observe a shift in cortical representation away from representation of the missing limb. Thus, the ability to train subjects to control individual or groups of neurons for the purpose of controlling a brain machine interface presents a powerful tool. This paradigm for training subjects to use a "new" set of neurons to reach-to-grasp a robotic arm could be beneficial for patients who have been the recipients of amputation (see **Specific Aims 2 & 3**).

Specific Aim 1

The first aim of my proposed work is to determine how populations of reach- and grasp- related neurons in primary motor cortex interact with each other during reach-to-grasp movements. The coordination of the transport and grasp components of reach-to-grasp behavior has been studied extensively psychophysically (M. Jeannerod 1984; Haggard and Wing 1995; Gentilucci et al.

1991). As stated earlier, psychophysical work has shown that the reaching of the proximal arm is temporally and spatially coordinated with the distal hand preshaping as the object is approached. However, it is not clear how populations of motor cortical neurons representing the movement of proximal and distal joints of the upper limb participate in this coordination.

A recently favored approach analyzes the dynamics of neural ensembles using state space-methods, where a high dimensional neural signal is described by a trajectory through a low-dimensional space (M. Churchland et al. 2012; Shenoy, Sahani, and Churchland 2013). This method, if extended to multiple neural ensembles, allows us to ask how the trajectories of multiple neural populations, such as reach- and grasp- related populations coordinate with one another.

Toward this aim, we simultaneously recorded from neurons in the primary motor cortices of two rhesus macaques that performed a reach-to-grasp task to multiple objects in different locations. By measuring temporal fluctuations in single-trial neural trajectories in reach- and grasp-related ensembles, we found evidence of a corrective mechanism which reduced the asychronies between the two ensembles' fluctuations particularly during the middle of the movement. Using Granger causality analysis on these fluctuations, we also found this coordination to be bi-directional.

Specific Aims 2 & 3

The second and third aims of my work examine how "mechanisms" for the coordination of reach-to-grasp movements at the level motor cortex are developed by looking at how they

emerge in a novel and artificial context through the use of brain-machine interfaces (BMIs). Traditionally in the BMI field, researchers have used kinetic or kinematic decoders that are tuned for particular kinetic or kinematic features in intact animals that have been trained for specific types of tasks. In this regard, the ability to train subjects to control individual or groups of neurons for the purpose of executing a novel behavior through the use of operant conditioning presents a powerful new paradigm.

The focus of my work then shifts to what it means to learn in the context of a brain-machine interface, and use this as a model to study learning to coordinate a novel motor behavior. In the 2009 study by Ganguly, a decoder was trained for direction of motion using neurons being recorded from primary motor cortex while rhesus macaques performed a center-out task (Ganguly and Carmena 2009). After the monkey had learnt how to do this task successfully, the decoder coefficients were shuffled. The monkey was able to learn how to use the decoder with the new decoder coefficients after just three days of training. However, it is important to note that the macaque had already been trained for the task at hand using the same neurons in the decoder, but with different coefficients.

The work of Eberhard Fetz motivated a slightly different approach that we used in this study. In his work, monkeys were operantly conditioned to control the firing rates of individual motor cortical neurons with the use of biofeedback (E. E. Fetz 1969; E. E. Fetz and Baker 1973; Eberhard E. Fetz 2007). This work presented a powerful possibility: the ability to train monkeys to control groups of neurons for the purpose of executing a particular task through the use of operant conditioning.

The motivations for this study were two-fold. First, this paradigm gave us a unique model for studying the development of coordination of a novel motor behavior through functional plasticity at the level of cortex. By creating a direct, causal connection between the primary motor cortex and the robotic plant, bypassing spinal cord and muscular dynamics, we ensured that any learning at the behavioral level was directly linked to changes in primary motor cortical activity. Second, if it is possible to train subjects to perform a naturalistic motor task with neurons using an essentially artificial mapping, this could have tremendous implications for patients who have suffered from peripheral nerve damage, or been the recipients of amputations. Much work has gone into showing that after amputation or peripheral nerve damage, there is somatotopic remapping in the cortical area previously involved in the control of the lost limb (J. N. Sanes et al. 1988; J. N. Sanes, Suner, and Donoghue 1990; Schieber and Deuel 1997; Wu and Kaas 1999; Qi, Stepniewska, and Kaas 2000). As such, single and multi-unit activity recorded from arrays implanted into reorganized motor cortex may not be reliably "pre-tuned" for kinematic or kinetic features to nearly the same degree. This paradigm could be used as a tool to help the almost 2 million people living in the United States who have amputations (Ziegler-Graham et al. 2008).

To address the second and third aims, this portion of my work will examine the emergent coordination underlying the learning of a novel reach-to-grasp movement in a brain machine interface context. Through the use of operant conditioning, we taught rhesus macaques to control a robotic arm to reach to, grasp, pull, and drop an object using neurons from their primary motor cortices. First, we demonstrated proof-of concept, testing whether monkeys could learn to coordinate independent reaching and grasping control dimensions in the BMI task. We found

that stereotypical patterns emerged and stabilized in the cross-covariance between the reaching and grasping velocity profiles, between pairs of neurons involved in decoding reach and grasp, as well as between other stable neurons in the network. The possibility of training subjects to perform a naturalistic motor task by modulating cortical neurons using an artificial mapping and of potentially influencing network activity in a lasting manner could have clinical implications for patients with amputations or strokes.

II. NEURAL COORDINATION DURING REACH-TO-GRASP

Abstract

When reaching to grasp, we coordinate how we preshape the hand with how we move it. To ask how motor cortical neurons participate in this coordination, we examined the interactions between reach- and grasp-related neuronal ensembles while monkeys reached-to-grasp a variety of different objects in different locations. By describing the dynamics of these two ensembles as trajectories in a low-dimensional state space, we examined their coupling in time. We found evidence for temporal compensation across many different reach-to-grasp conditions such that if one neural trajectory led in time, the other tended to catch up, reducing the asynchrony between the trajectories. Granger causality revealed bi-directional interactions between reach and grasp neural trajectories beyond that which could be attributed to the joint kinematics that were consistently stronger in the grasp to reach direction. Characterizing cortical coordination dynamics provides a new framework for understanding the functional interactions between neural populations.

Introduction

The focus of cortical motor neuroscience has been shifting from the question of representation to the question of neural dynamics (M. Churchland et al. 2012; Shenoy, Sahani, and Churchland 2013; Todorov and Jordan 2002; Todorov 2004; Scott 2004; Cheng and Sabes 2006). To analyze such neural dynamics, a standard approach is the use of state space methods, where a high-dimensional neural signal is described by a trajectory through a low-dimensional space. These methods can also allow us to ask how the trajectories of two neural populations, such as reach- and grasp-related populations, coordinate with one another. The temporal and spatial

coordination of transport and grasp components during reach-to-grasp behavior has been studied extensively in psychophysics experiments (M. Jeannerod 1984; Haggard and Wing 1995), and the development of this coordination has been studied in infants and children (Claes von Hofsten 1984; Kuhtz-Buschbeck, Stolze, Jöhnk, et al. 1998; Wimmers et al. 1998). The hand preshapes itself during transport such that when the arm reaches the object, the grasp aperture and orientation of the hand match the required configuration to grasp the object. Other work has posited that during reach-to-grasp, the nervous system is concerned with finding suitable positions for the thumb and fingers on the object's surface, and then moving them there (Smeets and Brenner 1999). In both scenarios, populations of motor cortical neurons may need to control the hand and arm in a tightly coordinated fashion.

Previous studies have touched on the cortical basis of coordinating reach-to-grasp movements but have focused on single neuron responses or pair-wise correlations. It has been shown that neurons in primary motor cortex (MI) modulate sequentially during a reach-to-touch task such that shoulder- and elbow-related cells (assessed via intracortical microstimulation effects on the same electrode sites from which the cells were recorded) begin firing on average about 60 ms before wrist and finger-related neurons consistent with the proximal-to-distal sequencing of muscle activation evident in this task as well as in many other multi-joint motor behaviors (Murphy, Wong, and Kwan 1985). In premotor cortex, pair-wise correlations between small groups of neurons have been shown to carry information about particular reaching and grasping combinations even though the constituent neurons' firing rates only carried information about either reaching or grasping (Stark et al. 2008). However, these studies have not examined how the dynamics of large neural ensembles coordinate during this reach-to-grasp behavior.

Here we simultaneously recorded from multiple MI neurons of two rhesus macaques who performed a reach-to-grasp task to multiple objects in different locations. By measuring temporal fluctuations in single-trial neural trajectories in reach- and grasp-related ensembles, we found evidence of a corrective mechanism which reduced the relative timing between the two ensembles' fluctuations particularly during the middle of the movement. Using Granger causality analysis on these fluctuations, we also found this coordination to be bi-directional.

Materials & Methods

The materials and methods used to collect this data were previously described in Saleh et al.; the same data were used for both studies (Saleh, Takahashi, and Hatsopoulos 2012). All surgical and behavioral procedures involved in this study were approved by the University of Chicago Institutional Animal Care and Use Committee and conform to the principles outlined in the Guide for the Care and Use of Laboratory Animals.

Behavioral Task

Two female rhesus macaques (Monkey O, 6.6 kg, age 7; Monkey A, 6.5 kg, age 7) were trained to reach for, grasp, hold, and release objects presented to them by a robot. Each trial was comprised of four periods: a pre-movement period, a movement period when reaching to grasp the object occurred, a hold period, and a release period, where the robot retracted the object (See Fig. 1a). The animal's vision was blocked during the pre-movement period, during which the macaque rested her hand on a button. The vision block was then retracted, and the animal proceeded to reach and grasp the object. Once the object was grasped, the animal was required to

hold the object until it was retracted by the robot at which point the object was released. Five different objects/orientations were used for this study, each of which was presented repeatedly in blocks of trials. The first was a key-like object that required a key grip, where the monkey held the object between the thumb and the side of the middle phalanx of the index finger. The second object, the D-ring object, was presented in a vertical and horizontal configuration that elicited a whole-hand power grip in two different wrist orientations (neutral, and pronated). The two configurations were treated as separate objects for the purposes of this study. The fourth object was a smaller D-ring, which elicited a precision grip, where the tips of the index and thumb were opposed. The fifth object, a sphere, required a whole-hand grip with a fanning of the fingers. Objects were presented at seven different locations (at varying azimuthal and elevation angles and depths) in pseudo-random order such that the macaque could not predict the object location on any given trial.

Neural Data Acquisition

Each macaque was chronically implanted with a Utah 100-electrode (400 mm inter-electrode distance, 1.5 mm electrode length; Blackrock Microsystems) microelectrode array in the upper limb area of M1 in the left hemisphere. During recording sessions, signals were amplified (gain 5000x), bandpass filtered between 0.3 Hz and 7.5 kHz, and digitized at 30kHz using the Cerebus Neural Data Acquisition system (Blackrock Microsystems). A high pass filter of 100Hz was applied for the spikes. The units were spike sorted offline using a semiautomated, Matlab routine (Mathworks) developed in our laboratory (courtesy of Jacob Reimer), incorporating some elements of a previously published algorithm (C. Vargas-Irwin and Donoghue 2007). We used 33

and 42 units for Monkey O's first and second datasets respectively, and 37 and 50 units for Monkey A's first and second datasets respectively.

Motion Tracking

The 3-D positions of the macaque's right arm, wrist, and fingers were recorded using a video-based motion analysis system (Vicon Motion Tracking System, Workstation 460); six M2 cameras were used at 1.2 megapixel resolution. Spherical retro-reflective markers (3 mm diameter) were glued to the macaque's fingers, hand, and forearm. A scaled version of a skeletal model of the arm (Holzbaur, Murray, and Delp 2005) was developed using the OpenSim platform (https://simtk.org/home/opensim) and used to compute joint kinematics of the shoulder, elbow, wrist, and fingers. The 17 joint angles are described in detail in Saleh et al (Saleh, Takahashi, and Hatsopoulos 2012). The joint angles computed for the shoulder were rotation about the coronal plane (abduction/adduction), rotation about the sagittal plane (flexion/extension), and rotation about the upper arm. We computed the angle of flexion/extension for the elbow. For the wrist, we computed the angles of flexion/extension, ulnar/radial deviation, and pronation/supination. We computed flexion/extension at the carpal-metacarpal joint of the thumb, at the metacarpophalangeal joints of the thumb, index, and ring fingers, and at the proximal interphalangeal joints of the thumb, index, and ring fingers. In addition, we computed the abduction/adduction of the thumb, index, and ring fingers (Saleh, Takahashi, and Hatsopoulos 2012). A synchronizing pulse was sent to both recording systems every 150 ms using a custom software package (TheGame, version 2) in order to align the kinematics with the neural data. The kinematic data were sampled at 250 Hz and bidirectionally filtered with a fourth-order Butterworth low-pass filter with a 6 Hz cutoff.

Analytical Methods: GLMs

We classified neurons as primarily related to either the reach component or the grasp component of movement by fitting generalized linear models with a log link function and Poisson error distributions relating either the set of proximal (shoulder and elbow) or distal (wrist and finger) joint velocities to the firing rates of individual neurons. Neurons whose activity could be better predicted through a model using only proximal joint velocities were labeled as reach-related (See Figure 2.1c), and those through a model using only distal joint velocities, grasp-related (See Figure 2.1d). For this part of the analysis, the neural data and the kinematics were downsampled to 50 Hz. For the remainder of the analysis, the neural and kinematic data were both analyzed at 250 Hz or a sampling interval, T, of 4 ms. To allow neural responses to have different latencies with respect to kinematics, we used a stack of kinematics with different delays as regressors. In line with a previous study (Saleh, Takahashi, and Hatsopoulos 2012), we used 11 different delays ranging from a lead time of 40 ms (neuron fires after movement) to a delay time of 160 ms (neuron fires before movement). The mean pseudo R-squared (Heinzl and Mittlböck 2003) across the runs of a ten-fold cross validation was used to select the neurons; those with a larger pseudo R-squared for the reach model were classified as reach-related and vice versa.

Analytical Methods: PCA

Instead of looking at the firing activity of each of the neurons in a population separately throughout the task, we used principal component analysis (PCA) to reduce the dimensionality of a population's firing activity in order to interpret the interaction between the two populations of neurons. First, each dataset was subdivided by condition, defined by the object type and location,

for a total of 35 (5 objects x 7 locations) conditions. The firing activities of all of the neurons were smoothed with a Gaussian kernel 100ms in width, and then all of the trials in the condition were concatenated, resulting in a matrix whose dimensions were number of neurons by the number of timepoints in all of the trials for that condition. We applied PCA to reduce the neuronal dimension of this matrix for every condition in each of the datasets. For each dataset, we chose between 3 and 8 components to describe the firing activity of the neural population in this reduced subspace, so as to account for at least 80% of the variance in the neural data throughout the course of the trial (See Figure 2.2a).

Results

We recorded simultaneously from between 33 to 50 MI neurons while macaques made reach-to-grasp movements to five different object types presented at seven different locations. Neurons were classified as primarily related to either the reach or grasp component of movement by fitting generalized linear models. Classifying neurons in M1 as either reach- or grasp-related based on limb kinematics is challenging because the kinematics of the proximal and distal joints of the upper limb are highly correlated during natural reach-to-grasp movements. In fact, we found that neurons whose activity could be well-predicted with a model using only the proximal (i.e. shoulder and elbow, see Methods) joint velocities tended to be well-predicted by a model using only the distal (i.e. wrist and fingers, see Methods) joint velocities (Figure 2.1b). Neurons whose activity could be better predicted with a model using only proximal joint velocities were labeled as reach-related (Figure 2.1c), and those better predicted with a model using only distal joint velocities as grasp-related (Figure 2.1d). For monkey O, 45% of the neurons were classified as reach-related and 55% of the neurons were classified as grasp-related for the first dataset, and 35% reach-related and 64% grasp-related for the second dataset. For monkey A , 32% of the

neurons were classified as reach-related and 66% were classified as grasp-related in the first dataset, and 34% reach-related and 66% grasp-related in the second dataset.

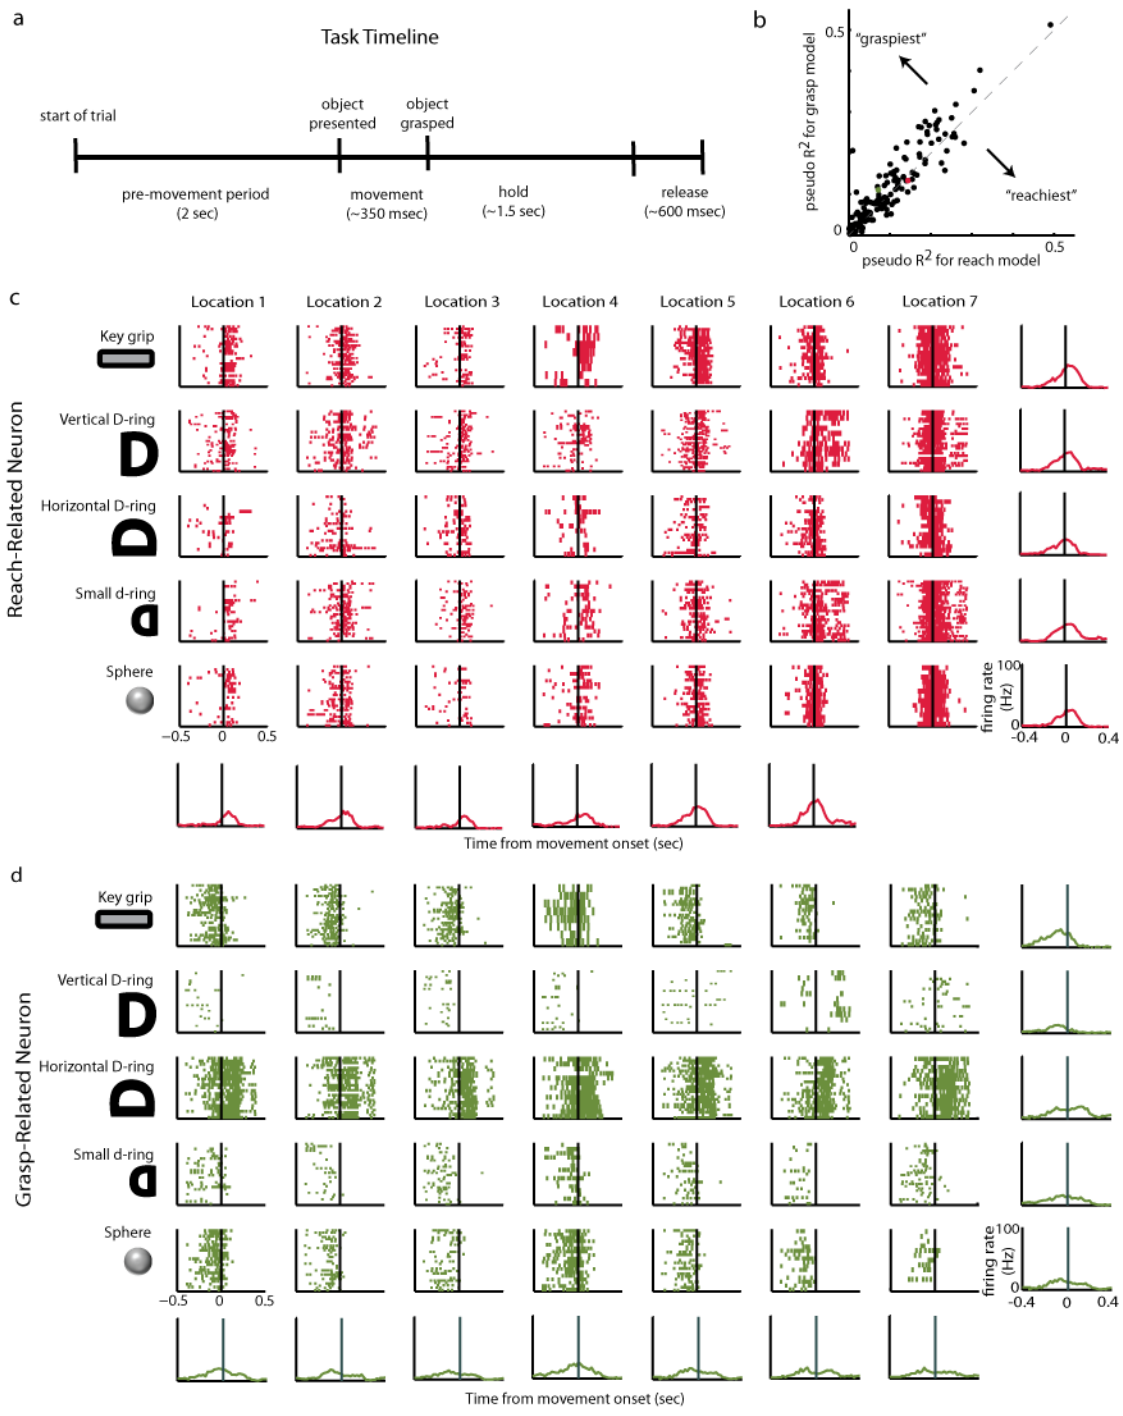

Figure 2.1 Reach- and Grasp-related Neurons (a) Each trial was comprised of four periods: a pre-movement period, reaching to grasp the object, a hold period, and a release period. (b) Pseudo R^2 values for reach encoding models versus the grasp encoding models for each neuron in all four datasets (c/d) Raster plots for one exemplary reach-related (c) and one exemplary grasp-related (d) neuron for monkey O reaching to grasp each of the different objects at each of the different locations. Peri-event time histograms shown for each object at each location.

In order to study the coordination between populations of M1 neurons during reach-to-grasp, we first verified that the macaques were temporally coordinating arm movement and hand prehension (Figure 2.2b). We found that the time of peak deceleration of the monkey's arm was significantly correlated with the time of maximum aperture of their hand, as is seen in human psychophysics studies (M. Jeannerod 1984). In particular, we found that the time of maximal flexion of the carpal-metacarpal joint of the monkey's thumb, a proxy for maximum aperture, was correlated with time of peak deceleration of the monkey's wrist, with a Pearson linear correlation coefficient of 0.298 (Monkey O, $p<.001$, t =12.245, DOF=1534) and 0.260 (Monkey A, $p<.001$, t =11.788, DOF=1918) for each monkey. Observation of this hallmark of coordinated reach-to-grasp allowed us to ask how populations of motor cortical neurons interacted with each other to achieve this coordination.

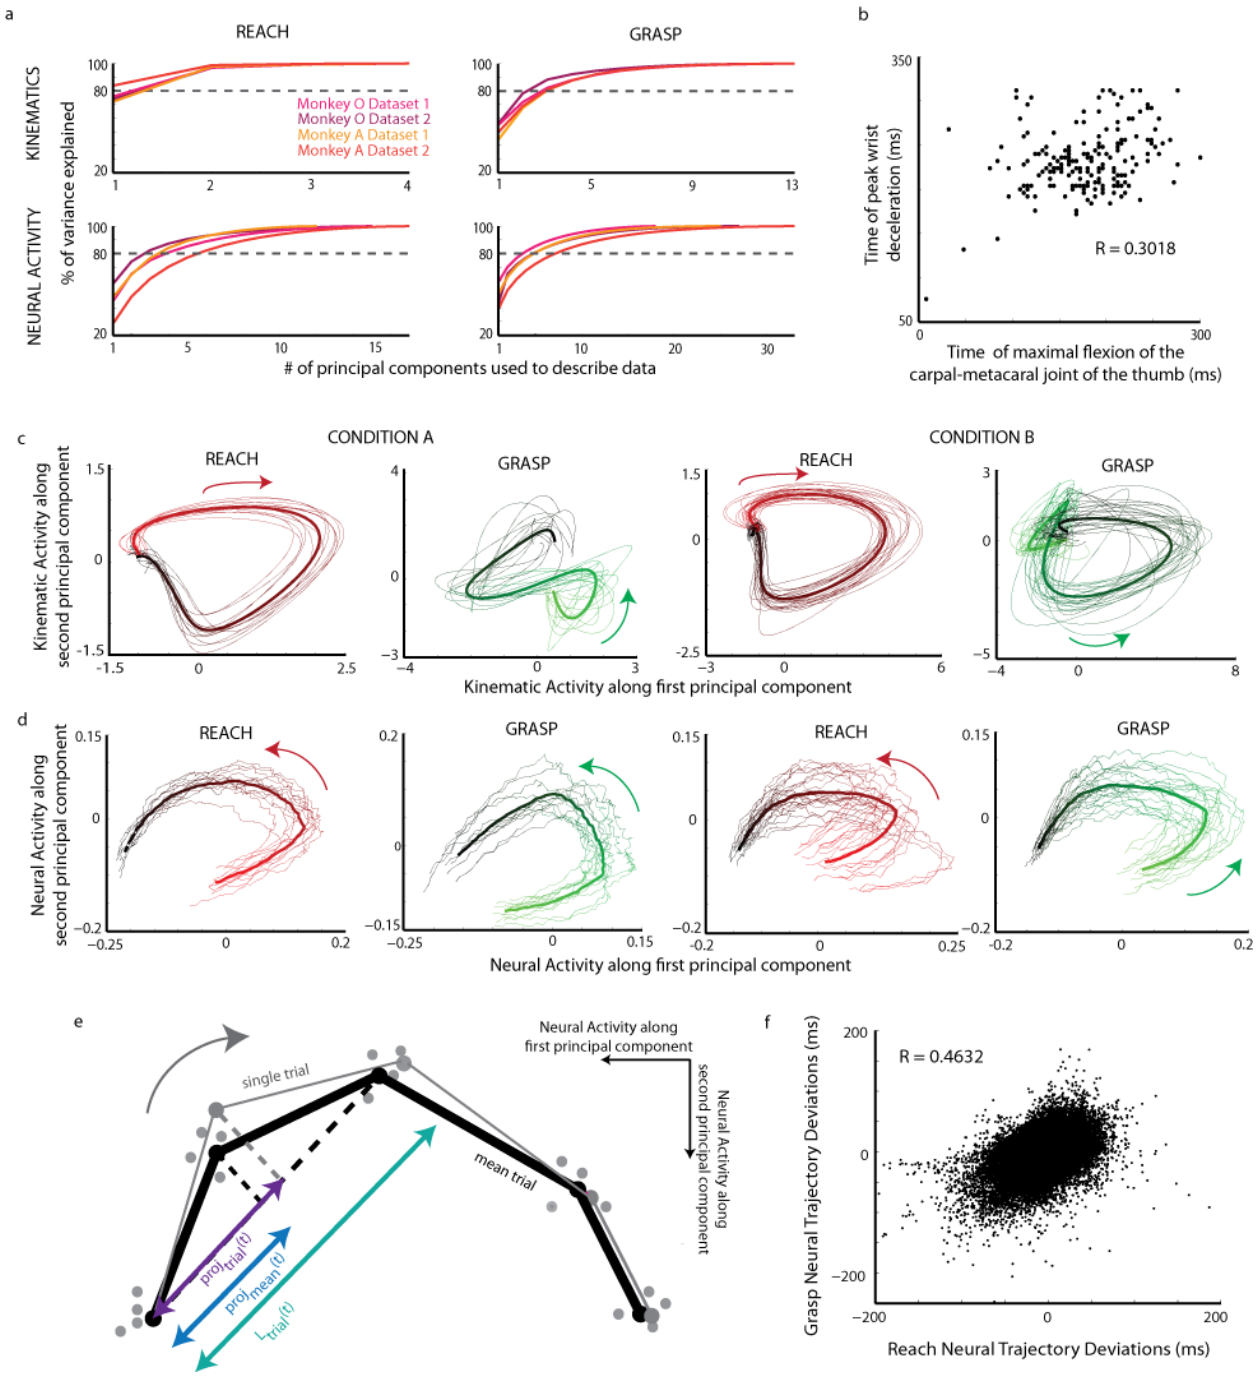

Figure 2.2 Trajectory Analysis (a) Percent of variance explained by reach-related and grasp-related neural and kinematic data for Monkey O (pink, purple) and Monkey A (orange, red) as a function of the number of principal components. (b) Time of maximal flexion of the carpal-metacarpal joint of the thumb, a proxy for maximum aperture, plotted against time of peak wrist deceleration for all of the trials in one dataset where Monkey O was reaching to grasp the vertical D-ring. (c) Kinematic trajectories of proximal (reach-related) and distal (grasp-related) joints in a reduced dimensional space. Trials shown are for condition A where one monkey reached to grasp a sphere at one location and condition B where the same monkey… (continued)

(continued) …reached to the vertical D-ring at a different location. Mean trajectories for the condition are shown as thick lines. (d) Neural population trajectories of reach-related and grasp-related activity in reduced dimensional space for the same two conditions. Mean trajectories for the condition are shown as thick lines. Color gradients in (c) and (d) reflect time from the beginning (red/green) to end (black) of the trial.(e) Cartoon illustration of a single trial (gray) reach-related trajectory and mean trajectory (black thick line) projected onto a local approximation of the mean trajectory (black dotted line) in two dimensional principal component space. The projections are in the purple and blue, respectively. (Samples from other trials are also shown as gray dots.) The ratio of the difference between the two projections (i.e. the single-trial onto the local approximation to the mean minus the mean onto the local approximation to the mean) and the length of the approximation of the mean trajectory (turquoise) is used to convert the deviations into units of time. (f) Deviations for reach and grasp neural trajectories plotted against each other for all of the trials in one dataset (Monkey O).

We examined single-trial neural trajectories of the reach and grasp populations in a reduced subspace using principal components analysis (see Methods). For each dataset, we chose between 3 and 8 components to describe the population activity so as to account for at least 80% of the variance in the data (Figure 2.2a). Each dataset was subdivided by condition, defined by the object type and location for a total of 35 (5 objects x 7 locations) conditions. Given that the trajectories of the proximal and distal joints were fairly stereotyped within a particular condition—i.e. the monkey made very similar movements to reach and grasp the same object presented at the same location (Figure 2.2c), the neural trajectories of the reach and grasp populations (Figure 2.2d) were also fairly stereotyped. We took advantage of this to compute mean reach and grasp neural trajectories for each of the conditions in the reduced dimensional space. The trial-averaged neural trajectory for a particular condition was thus used to represent the ideal, stereotyped dynamics of the neural ensemble when the monkey was reaching to grasp that particular object at that particular location.

Assuming that each of the neural populations would tend towards this stereotyped behavior, we compared the neural population trajectories for each trial to the trial-averaged trajectories by

measuring the temporal deviations of each trial's neural trajectory from the mean neural trajectory as a function of time. At each time point during the trial, we projected the single-trial trajectories and mean trajectory onto a local approximation to mean trajectory (Figure 2.2e). Thus, for trial *i*, every time point had an associated temporal deviation, $\delta_i(t)$, from the mean, defined as the amount of time the reach or grasp neural trajectory was leading (or lagging) the mean trajectory:

$$\delta_i(t) = 2T \cdot \frac{[proj_i(t) - proj_{mean}(t)]}{L_i(t)}$$

Equation 2.1

where *T* is the sampling interval (i.e. 4 ms), $proj_i(t)$ is the projection of trial i at time point *t* onto a local approximation to the mean trajectory, $proj_{mean}(t)$ is the projection of the mean trajectory at same time point onto the local approximation to the mean trajectory, and $L_i(t)$ is the length of the local approximation to the mean trajectory from time point *t-1* to *t+1*.

Measuring how far both neural populations were along their own stereotyped behavior gives us a window into how synchronized their behavior is. Overall, the reach-related and grasp-related populations' deviations from their respective means were significantly correlated, with correlation coefficients of 0.3856 (p=0, t =108.676, DOF=67611) and 0.4632 (p=0, t =140.585, DOF=72343) for monkey O's first (Figure 2.2f) and second datasets, and 0.107 (p=0, t =30.595, DOF=80806) and 0.0818 (p=0, t =25.167, DOF=94092) for monkey A's first and second datasets, indicating that the degree to which each population tends towards stereotyped behavior is linked.

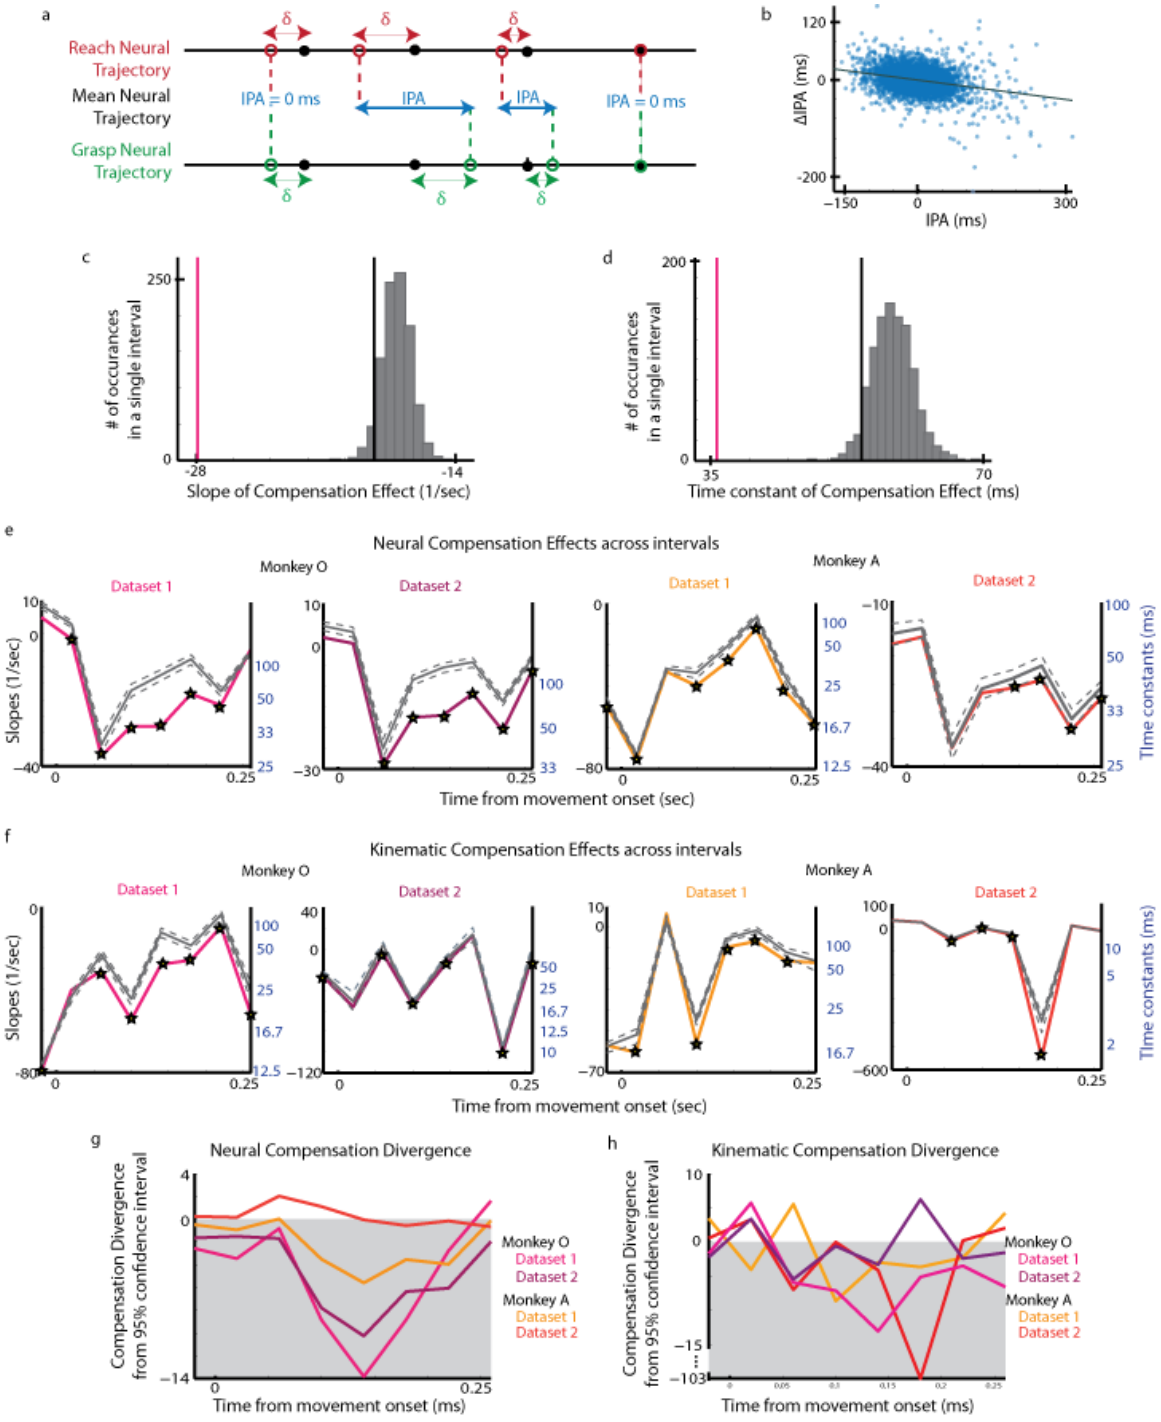

Figure 2.3 Inter-Population Asynchrony (IPA) and Compensation (a) Illustration of compensation effect when reach-related trajectory lags grasp-related neural trajectory. (b) Inter-population asynchrony (IPA) plotted against change in IPA for one interval in one dataset (Monkey O) (c) Slope of compensation (vertical pink line), distribution of shuffled slopes of compensation (gray), and 95% confidence level (vertical black line) for one interval in one dataset (Monkey O) (d) Time constants of compensation (pink), distribution of shuffled time constants of compensation (gray), and 95% confidence level (black) for one … (continued)

(continued)…interval in one dataset (Monkey O) (e) Slopes of compensation (colored) and mean shuffled slopes of compensation (gray, with 5% and 95% confidence levels after Bonferroni correction) for the neural data plotted throughout the course of movement. Intervals with significant compensation effects are starred; time constants are shown in blue. (f) Slopes of compensation and shuffled slopes of compensation (gray, with 5% and 95% confidence levels after Bonferroni correction) for the kinematic data plotted throughout the course of movement. Intervals with significant compensation effects are starred; time constants are shown in blue. (g) Divergence of neural compensation slope from 95% confidence level (after Bonferroni correction) of shuffled slopes for both monkeys. (h) Divergence of kinematic compensation slope from 95% confidence level (after Bonferroni correction) of shuffled slopes for both monkeys.

We defined a quantity called the Inter-Population Asynchrony (IPA) as the difference between the two populations' temporal deviations from their respective means in order to quantify this degree of coupling (Equation 2). Analogous to previous psychophysical work that has used perturbations to study the characteristic coupling of the transport and aperture components of reach-to-grasp (Haggard and Wing 1995), we examined fluctuations in $IPA_i(t)$ to examine the temporal coordination between the reach- and grasp-related neural trajectories (Figure 2.3a).

$$IPA_i(t) = \delta_i^{grasp}(t) - \delta_i^{reach}(t)$$

Equation 2.2

The primary question we addressed in this study was whether a compensatory mechanism existed to correct for transient increases in the asynchrony between the two populations' neural trajectories. That is, if one population's neural trajectory started to lead or lag the other, would either or both populations' neural trajectories compensate such that they would once again become synchronized (i.e. IPA=0). To do this, we analyzed how asynchronies tended to diminish in time. We found that an inverse relationship existed between the IPA and its instantaneous

temporal derivative, such that larger positive asynchrony was associated with a larger negative derivative in asynchrony and vice-versa (Figure 2.3b, Equation 2.3). That is to say, the larger the difference in the degree to which both neural populations were following their respective stereotyped behavior, the more quickly that both would work to correct for this difference. This compensation effect reflected a proportional control law, indicating an exponential decay in asynchrony (Equation 4) with a time constant of $1/\alpha$ inversely proportional to the relationship between the IPA and its instantaneous temporal derivative. In order to assess how the time constant associated with this decay in IPA varied over the course of movement, we divided trials into intervals and examined the relationship between the IPA and its instantaneous derivative in each interval. Each interval was 40 ms long; there was one interval prior to movement onset, and eight intervals after. Although trials were approximately the same length, there was variability in when individual trials ended during the last interval, making the data less interpretable. Thus, only the first eight intervals were considered for the remainder of the analysis.

$$\frac{\Delta IPA(t)}{dt} = -\alpha \cdot IPA(t)$$

Equation 2.3

$$IPA(t) \propto e^{-\alpha t}$$

Equation 2.4

Since we collected data from both populations of neurons simultaneously, this patterning could be a result of both populations independently and simultaneously regressing to their respective

means. To address this issue, we shuffled the trial order of the reach- and grasp-related neural population trajectories within each reach-to-grasp condition. For example, the reach-related neural population trajectory on a trial where the monkey reached to grasp a sphere at a specific location would be paired with a grasp-related neural population trajectory on a different trial in the same condition. After shuffling, any trial-specific, temporal coordination between the two populations' trajectories would be destroyed. We shuffled the trial order and performed the same analysis 1000 times, looking at the temporal evolution of the IPA and its temporal derivative over the course of the movement.

The distribution of the slopes generated by the analysis of the shuffled data was compared to the slope generated by the original data for each interval (Figure 2.3c). The same comparison was made between the time constants generated by the shuffled data and the original data for each interval (Figure 2.3d). If the original slope was less than 95% of the shuffled slopes for a particular interval after Bonferroni correction, it was considered to be significant. We found that the presence of a significant compensation effect towards the middle of the movement (between 60 ms to 220 ms after movement onset; average movement duration was 324 ms); this effect was present in at least one interval towards the middle of the movement in 88.57% and 80% of the 35 reach-to-grasp conditions for monkey O datasets, and 42.86% and 16% for monkey A's datasets ($p<0.05$, one-tailed, non-parametric shuffling test) respectively. In all of the datasets, all of the significant time constants ranged were100 ms or shorter (Figure 2.3e).

In order to understand the importance of such inter-neuronal dynamics during coordinated reach-to-grasp, we analyzed and compared their relationships with kinematics. We performed the same

analysis on the kinematic trajectories of proximal and distal joints by examining them in a reduced subspace using principal components analysis; we computed trial-averaged trajectories of the proximal and the distal joint velocities for each condition. As with the neural data analysis, we measured deviations from the mean by projecting single-trial kinematic trajectories onto the trial-averaged kinematic trajectories. Thus, much like the neural trajectories, every point on the kinematic trajectories of either the proximal or distal joints had an associated temporal deviation. By comparing the asynchrony between the proximal and distal joint deviations to its instantaneous temporal derivative, we found patterns of compensation with significant time constants in the kinematic data (Figure 2.3f).

In order to better assess the relationship between the kinematic compensation and the neural compensation, we examined how the slopes of the compensation effect diverged from the 95% confidence level of the slopes generated from the shuffled data by subtracting the 95% confidence level (after Bonferroni correction) from the unshuffled slopes. In both of Monkey O's datasets and Monkey A's 1st dataset, we found that this divergence increased during the movement, indicating a quicker compensation effect that cannot be explained by both trajectories independently regressing to their own means. This patterning was present in both the neural (Figure 2.3g) and the kinematic (Figure 2.3h) trajectory dynamics in all four datasets.

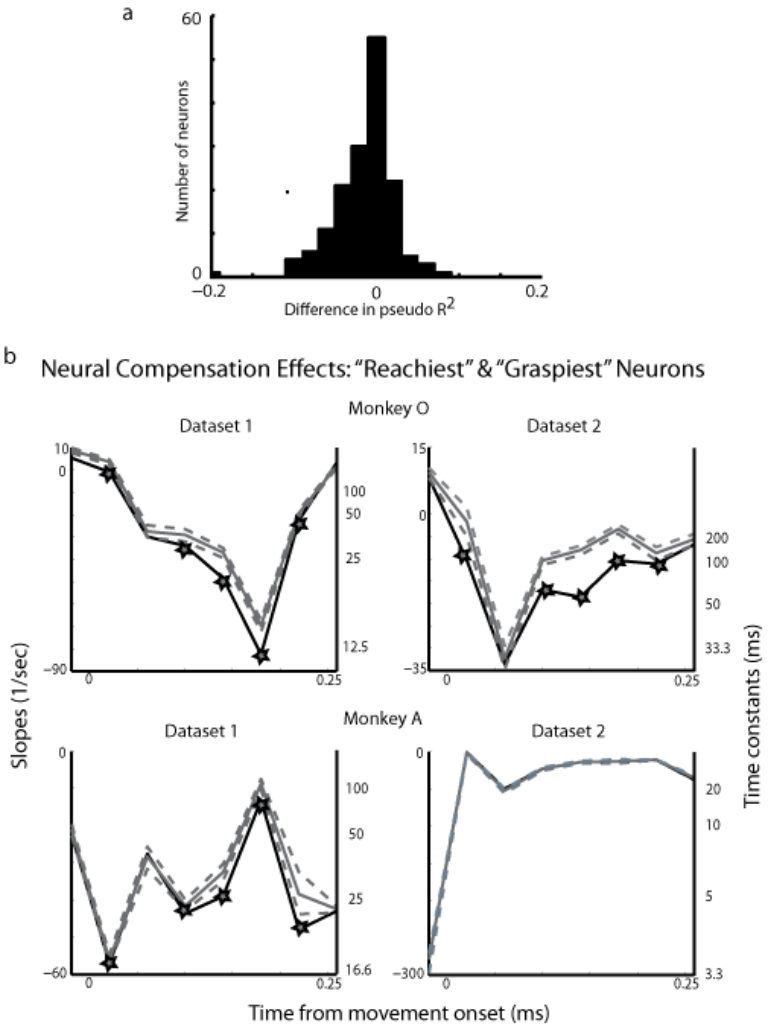

Figure 2.4 "Reachiest" & "Graspiest" Neurons (a) Histogram of differences between the pseudo R-squared values for the models based on reach-related and grasp-related kinematics for all of the datasets (b) The subpopulations were comprised of 10 "reachiest" and 10 "graspiests" neurons. Slopes of compensation (black) and mean shuffled slopes of compensation (gray, with 5% and 95% confidence levels after Bonferroni correction) are plotted throughout the course of movement. Intervals with significant compensation effects are starred; time constants are shown in black.

As stated earlier, classifying neurons in M1 as either reach- or grasp-related based on limb kinematics is challenging because the kinematics of the proximal and distal joints of the upper limb are highly correlated during natural reach-to-grasp movements. Given that this classification is actually a continuum (see Figure 2.1b), we reran this analysis on subpopulations

of the 10 “reachiest” and 10 “graspiest neurons; we found significant compensation effects in three of the four datasets (Figure 2.4). In addition, to measure how sensitive the compensation dynamics were to this partitioning into reach vs grasp subpopulations, we randomly partitioned the neurons into two subpopulations one hundred times, preserving the respective sizes of the subpopulations: one population was the same size as the original reach-related population and the other, the same size as the original grasp-related population. We repeated all of the previous analysis for each of the population partitions to generate the slopes and time constants associated with the compensation for each partition; examining the spread of these slopes and time constants gave us an idea of how sensitive our results were to the original reach or grasp population partition. Judging from the mean slopes for all of these partitions it was evident that the compensation dynamics were still present for random partitions of the population, but not to the same extent as the original reach- and grasp-related partition (Figure 2.5).

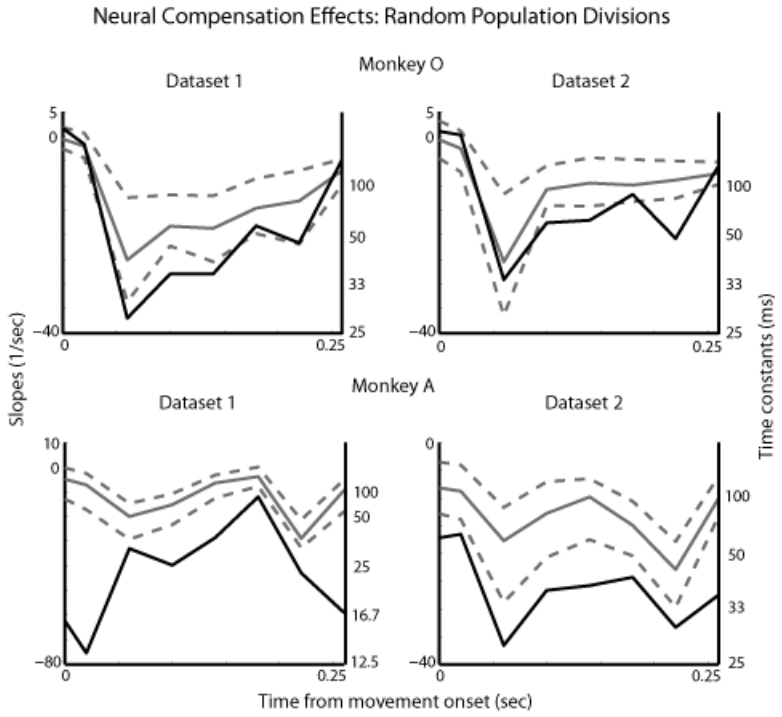

Figure 2.5 Random Population Divisions Neurons were randomly partitioned into the two subpopulations 100 different times. The original compensation slopes (and time constants) for the reach/grasp partitions are shown in black. The compensation slopes averaged over all random partitions are shown in gray, along with the 5% and 95% spread (after Bonferroni correction).

We next used Granger causality to assess directed causal influences between all combinations of the reach- and grasp-related neural deviations and the reach- and grasp-related kinematic deviations. We quantified the ability of the past deviations of one type to predict the present deviations of another type. We compared a full model that predicted one deviation using regressors that included the past of all four deviation trajectories with a model where the regressor of interest was removed (Seth 2010). We set the maximal model order to 10 lags (40 ms), and used the Bayesian Information Criterion to select the model order which resulted in a model order of 10 lags for each dataset. We found the significance of the Granger causal relationships using an F-test, and then corrected for multiple comparisons using Bonferroni

correction. This analysis was done using the Granger Causal Connectivity Toolbox and methods developed by Anil Seth (Seth 2010).

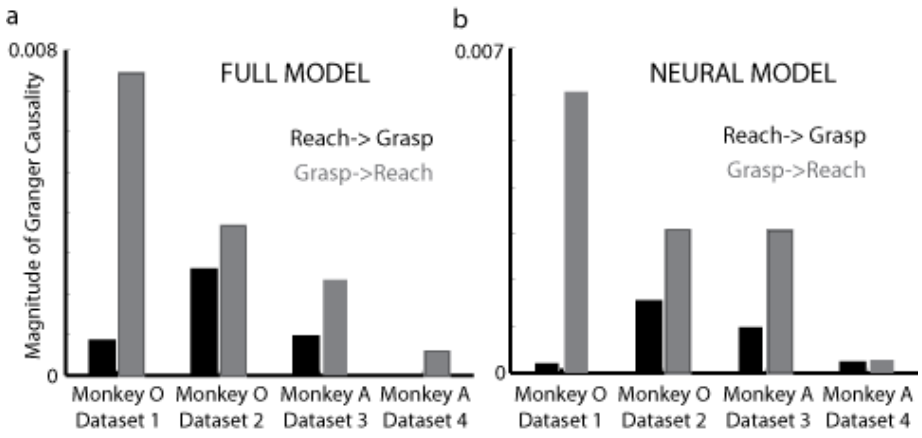

Figure 2.6 Granger Causality Relationships (a/b) Magnitude of significant Granger causality relationships computed using the full model (a) and neural model (b). Black bars indicate magnitude of Granger causality directed from Reach to Grasp neural deviations whereas gray bars indicate magnitude of Granger causality directed from Grasp to Reach neural deviations

For three of the datasets, we found significant bi-directional Granger causal relationships between the deviations of the reach- and grasp-related neural trajectories (Monkey O dataset 1 reach to grasp: $p<.001$, $F=7.575$, DOF=(1,66866); Monkey O dataset 1 grasp to reach: $p<.001$, $F=47.625$, DOF=(1,66866); Monkey O dataset 2 reach to grasp: $p<.001$, $F=18.949$, DOF=(1,71546); Monkey O dataset 2 grasp to reach: $p<.001$, $F=25.573$, DOF=(1,71546) ; Monkey A dataset 1 reach to grasp: $p<.001$, $F=8.731$, DOF=(1,79916); Monkey A dataset 1 grasp to reach: $p<.001$, $F=20.866$, DOF=(1,79916)). For the fourth dataset (i.e. 2nd dataset from monkey A), we only found a significant Granger causal relationship from the deviations of the grasp-related neural trajectories to the reach-related neural trajectories (Monkey A dataset 2 grasp to reach: $p<0.002$, $F=5.172$, DOF=(1,93056)). In all of the datasets, there was stronger "causal" influence from the grasp-related to the reach-related neural deviations than from the

reach-related to the grasp-related neural deviations (Figure 2.6a). In addition, using a model that incorporated only the neural deviation trajectories and setting the model order to 2 lags (8ms), we also found significant bi-directional Granger causal relationships between the deviations of the reach- and grasp-related neural trajectories that were stronger for the first three datasets (Monkey O dataset 1 reach to grasp: $p<.001$, $F=7.459$, DOF=(1,66868); Monkey O dataset 1 grasp to reach: $p<.001$, $F=198.992$, DOF=(1,66868); Monkey O dataset 2 reach to grasp: $p<.001$, $F=54.9$, DOF=(1,71548); Monkey O dataset 2 grasp to reach: $p<.001$, $F=110.607$, DOF=(1,71548); Monkey A dataset 1 reach to grasp: $p<.001$, $F=39.899$, DOF=(1,79918); Monkey A dataset 2 grasp to reach: $p<.001$, $F=124.684$, DOF=(1,79918)) than for the fourth dataset (Monkey A dataset 2 reach to grasp: $p<0.004$, F=9.899, DOF=(1,93058); Monkey A dataset 2 grasp to reach: $p<0.004$, F=12.606, DOF=(1,93058)). These "causal" relationships were also stronger from grasp- to reach-related neural deviations than vice-versa (Figure 2.6b).

Discussion

We have introduced a novel approach to quantifying the coordination dynamics of neural ensembles in the context of reach-to-grasp. We found that the trajectories of the reach-related and grasp-related neural populations were temporally coupled over the course of movement such that any transient asynchrony in time progression between the two trajectories was corrected for by a proportional control mechanism. These coordination dynamics between the two neural populations were observed in a large variety of different types of reach-to-grasp behaviors.

We found that the compensation dynamics were still present for random partitions of the population, but were slower as compared to the original reach- and grasp-related partition (see

Figure 2.5). The presence of compensation dynamics even for random partitions is perhaps not surprising given that the upper limb area of MI forms a highly interconnected network of neurons with highly overlapping movement representations. ICMS studies have shown that nearly half of stimulation sites in MI evoke twitches in both proximal and distal muscles particularly on the precentral gyrus where we implanted our electrode arrays (Park et al. 2001; Park, Belhaj-Saïf, and Cheney 2004). Several reach-to-grasp studies have shown that the responses of a large percentage of individual MI neurons were more accurately predicted by considering both proximal and distal kinematics, thereby accounting for the inherent correlations in proximal and distal kinematics in reach-to-grasp (C. E. Vargas-Irwin et al. 2010; Saleh, Takahashi, and Hatsopoulos 2012). According to this perspective, the neural coordination we observed is a broader property of the entire upper limb area of motor cortex.

The time constant of this compensation effect generally ranged from around 10 ms to 90 ms. The longer durations in this range were greater than the shortest possible transmission time from motor cortex down to the periphery and back up to cortex as measured by the long-latency loop time (Witham et al. 2011; Cheney and Fetz 1984). However, it is unlikely that the temporal coupling we observed between reach-related and grasp-related populations in MI is solely mediated by signaling though the periphery. The Granger causality analysis indicated that current temporal deviations in one population's neural trajectory could not be explained solely from the past deviations of limb kinematics but required knowledge of the past deviations of the second population's neural trajectory. This suggests that the temporal coupling between the two populations was not only due to interactions via the periphery. In addition, the Granger causality

analysis indicated that causal interactions between the two neural populations' deviations occurred at very fast timescales, on the order of 8 ms.

We found that these Granger causal relationships were also stronger from grasp- to reach-related neural deviations than vice-versa, although bidirectional relationships were present in three of the four datasets. The arm and hand constitute a biomechanically linked chain of limb segments, and as such motion about one joint induces interaction torques on the other. However, given the larger inertial properties of the arm as compared to the hand, interaction torques on the distal limb segments induced by motion of the shoulder and elbow are larger than interaction torques on the proximal limb segments due to motion of the wrist and fingers (Galloway and Koshland 2002; Ambike and Schmiedeler 2013). Therefore, when temporal deviations in the neural state associated with the distal joints occur, it may be particularly important for these deviations to inform the deviations of the proximal joints than vice-versa. In other words, the asymmetry in interaction torques from proximal to distal joints may be compensated for by an asymmetry in neural communication from distal to proximal joint representations.

The compensation effects were more evident for Monkey O's two datasets than for Monkey A's two datasets, particularly in the case of Monkey A's second dataset. One key difference in this dataset is that we had more neurons (50 as compared with 33, 37, and 42), and needed more principal components to describe this dataset, which may have contributed to differences in the results. It is also true that for datasets in which the reach- and grasp-related neural deviations are more strongly correlated, we see a greater degree of neural compensation divergence (Figure 2.3g). We hypothesize that a greater correlation between the deviation types could increase the

accuracy of our estimation of compensation effects, since we have more data points at which the trajectories are coupled. This effect is evident when we compare the results for data aggregated over all of the conditions to those calculated for each condition. In the aggregate case (Figure 2.3e), all of the datasets have between 4 and 6 intervals with significant compensation. The difference becomes stark when examining the compensation effects within each condition: Monkey A's datasets only have significant compensation present in 42% and 16% of the conditions, while Monkey O's datasets have effects presents in 80% and 88% of the conditions. We propose that the globally weaker results in the case of Monkey A do not necessarily indicate a lesser degree of compensation, but a decreased ability to accurately estimate compensation.

Our approach is limited in revealing how these dynamics causally affect coordination. In psychophysics studies, behavioral perturbations are used to establish causality (Haggard and Wing 1995). Similarly, electrical and optogenetic stimulation could be used to establish causality (Diester et al. 2011; Liu et al. 2012). In this study, as a first step towards establishing causality, we incorporated limb kinematics as a proxy for common input and found significant interactions between the two neural populations above and beyond the effects of the common kinematics.

This work extends previous research that has focused on quantifying the dynamics of individual populations (M. M. Churchland et al. 2010; M. Churchland et al. 2012; Mante et al. 2013; Shenoy, Sahani, and Churchland 2013; Kaufman et al. 2014). These approaches characterize the dynamics of neural activity through their trajectories in a low dimensional projection. Our approach generalizes this approach, by dividing the population into functionally defined subsets and analyzing the relations between trajectories during behavior. This could be adapted to

analyze coordination within any brain region and, more importantly, between brain areas. Being able to monitor cortical coordination dynamics is a first step towards understanding how interactions across populations of neurons give rise to coordinated behavior.

Acknowledgements

Maryam Saleh is currently affiliated with MATTER in Chicago, IL.

Grants

This work was supported by the National Institute of Neurological Disorders and Stroke Grant R01 NS045853, IGERT: Integrative Research in Motor Control & Movement NSF Grant#DGE-0903637, National Institutes of Health 5R01NS063399, and the Pritzker Fellowship in Neuroscience.

III. EMERGENT COORDINATION UNDERLYING LEARNING TO REACH-TO-GRASP WITH A BMI

Abstract

As children age, their inter-joint coordination becomes more stereotyped (Claes von Hofsten 1984; Kuhtz-Buschbeck, Stolze, Jöhnk, et al. 1998). Less is known about the role of motor cortex in developing coordinated reach-to-grasp. Similarly, studies have shown that after amputation, the cortical area previously involved in the control of the lost limb undergoes reorganization (J. N. Sanes et al. 1988; J. N. Sanes, Suner, and Donoghue 1990; Schieber and Deuel 1997; Wu and Kaas 1999; Qi, Stepniewska, and Kaas 2000), but limited work has focused on developing BMIs that use neurons that are not tuned for kinetic or kinematic features. Here we studied the emergence of coordinated reach-to-grasp behavior in rhesus macaques that had been the recipients of therapeutic amputations, and were taught to cortically control a robotic arm through operant conditioning using neurons that were not explicitly reach- or grasp-related. Over the course of training, stereotypical patterns emerged and stabilized in the cross-covariance between the reaching and grasping velocity profiles, between pairs of neurons involved in decoding reach and grasp, and to a lesser extent between other stable neurons in the network. In addition, we found that the degree to which a pair of neurons covary their modulation, either positively or negatively, has a direct effect on the extent to which the behavior is coordinated.

Introduction

Around 185,000 amputations are conducted in the United States each year-- there are almost 2 million people living in the United States who have been the recipients of amputations (Ziegler-

Graham et al. 2008). Over 5,000 service members in the U.S. Armed Forces had traumatic amputations in the timespan of 2000 to 2011 (Armed Forces Health Surveillance Center (AFHSC) 2012). A number of studies have shown that after amputation or peripheral nerve damage, the cortical area previously involved in the control of the lost limb gets remapped to areas that are somatotopically close by (J. N. Sanes et al. 1988; J. N. Sanes, Suner, and Donoghue 1990; Schieber and Deuel 1997; Wu and Kaas 1999; Qi, Stepniewska, and Kaas 2000). It follows that these patients might not have the same degree of representation of the lost limb in primary motor cortex. Moreover, because of the lack of an intact limb, a decoding approach that maps neural modulation to intact limb movements is not possible. Past cortically-controlled brain-machine interfaces (BMIs) have largely focused either on intact non-human primates or short-term nerve blocks of the arm to model severe motor disabilities. In such cases, neural decoders for limb/cursor control are initialized using a biomimetic technique where paired measurements of neural signals and intact limb kinematics are available (Li et al. 2011; Ganguly and Carmena 2009; Taylor, Tillery, and Schwartz 2002). Thus, we could not rely on finding "arm" neurons or "hand" neurons to use to control a prosthetic arm for patients that have undergone upper limb amputations.

Recent work has been shifting away from the traditional paradigm through the use of adaptive decoding or by studying what happens when tuned decoders are perturbed during training (Ganguly et al. 2011; Ganguly and Carmena 2009; Taylor, Tillery, and Schwartz 2002; Li et al. 2011; Mussa-Ivaldi and Miller 2003; Willett et al. 2013; Orsborn et al. 2012). In this study, we investigated whether we could teach animals to reach to grasp in a novel brain machine interface

task using neurons that were not originally identified as “reach-related” or “grasp-related” neurons.

This approach provides a unique window into plasticity at the level of motor cortex. By establishing a direct, causal connection between the primary motor cortex and the robotic plant that bypasses both spinal cord and neuromuscular dynamics, we can infer that any learning at the behavioral level is directly due to changes in primary motor cortical activity. To this effect, we examined the plasticity between pairs of neurons that are directly involved in BMI control, as well as the plasticity associated with those that are not directly involved to further elucidate their role in learning through our study.

In a similar vein, this paradigm also provided a way to study the development of a novel, coordinated motor behavior at the level of cortex. Extensive psychophysical work has shown that during reach-to-grasp, reaching of the proximal arm (reaching or transport component) is temporally and spatially coordinated with the preshaping of the hand (grasp component) (M. Jeannerod 1984; Haggard and Wing 1995). The development of this coordination has been studied in infants and children (Claes von Hofsten 1984; Kuhtz-Buschbeck, Stolze, Jöhnk, et al. 1998); they tend to have larger apertures, start prehension earlier, and rely more on visual feedback than adults. As children age, their kinematic profiles and inter-joint coordination become more stereotyped. Less well understood is the role of cortex in coordinating and learning to coordinate reaching to grasp.

Here we studied the emergence of coordinated reach-to-grasp behavior in rhesus macaques that had been the recipients of therapeutic amputations, and were being taught to control a robotic arm and hand via a cortically-controlled brain machine interface through the use of operant conditioning. We examined how the coordination of this behavior was exhibited, both at the behavioral and neural level. We focused our efforts on three primary questions: Can macaques learn to coordinate independent reach and grasp control dimensions in a brain machine interface task using neurons that are not explicitly "reach-related" or "grasp-related" to begin with? What are the neural correlates of learning this coordination among the neural ensembles controlling reaching and grasping? Finally, do we see similar correlates arise among neurons that are not directly involved in BMI control, and how do these compare?

Methods

All surgical and behavioral procedures involved in this study were approved by the University of Chicago Institutional Animal Care and Use Committee and conform to the principles outlines in the guide for the Care and Use of Laboratory Animals.

Neural Recordings

Data used for this analysis was collected from two female rhesus macaques (Macaca mulatta) monkeys that had been the recipients of therapeutic amputations. The Utah 100-microelectrode arrays (Blackrock Microsystems, Salt Lake City, UT) used for this analysis were placed in the forelimb region of primary motor cortex (M1) contralateral to the intact limb in Monkey K, and in M1 ipsilateral to the intact limb in Monkey Z. Informed by prior work on the reorganization of motor cortical somatotopy following amputation (J. N. Sanes et al. 1988; J. N. Sanes, Suner, and

Donoghue 1990; Schieber and Deuel 1997; Wu and Kaas 1999; Qi, Stepniewska, and Kaas 2000), we stimulated the surface of motor cortex contralateral to the amputation and implanted in a region that elicited movement of the shoulder and stump, bearing in mind the idea that this region could have previously been part of the forelimb region prior to amputation. The electrodes on the array were 1 mm in length. During a recording session, spike waveforms from up to 96 electrodes were amplified, filtered, and recorded digitally (14-bit) at 30kHz per channel using a Cerebus acquisition system (Blackrock Microsystems). Units were sorted online with a hoop-sorting algorithm (Santhanam et al. 2004); potential spikes were first identified when the filtered voltage dropped below a user-defined threshold. These spikes were sorted by placing lower and upper voltage thresholds (the "hoops") at specific times relative to the initial threshold crossing.

Experimental Setup

The macaques were trained on a brain-machine interface task using an operant conditioning paradigm. In this task, the macaque had to learn how to navigate two control dimensions of a robotic arm in order to perform a reach-to-grasp task. The robot was composed of a 7 DOF WAM arm attached to a 4 DOF Barrett Hand (Barret Technology, Inc.). The macaque simultaneously learned to control the reaching motion, towards and away from the base of the robot, and the grasping motion, opening or closing all three digits of the hand concurrently. A successful trial would involve reaching-to-grasp a sphere placed on a board in front of the robot, pulling it back, and finally dropping it. Using neurons identified using online spike sorting, distinct clusters of functionally connected groups of neurons were created for the purpose of controlling each control dimension (Eldawlatly, Jin, and Oweiss 2008). For Monkey K, each of the clusters was 15 neurons, for Monkey Z, 10 neurons. Twenty-tap Wiener filter decoders were

initialized in an unsupervised manner using spontaneous data from each ensemble of neurons (Badreldin et al. 2013), binned at 50 milliseconds, these decoders remained static over the course of the study. The macaques learned this mapping such that they could coordinate both control dimensions in order to perform successful trials. Instruction was provided through operant conditioning. In order to shape the behavior early on in the study, the macaques were rewarded for making portions of the complete reach-grasp-pull-drop movement. As the animals started to pick up the task, they were rewarded only for completed trials.

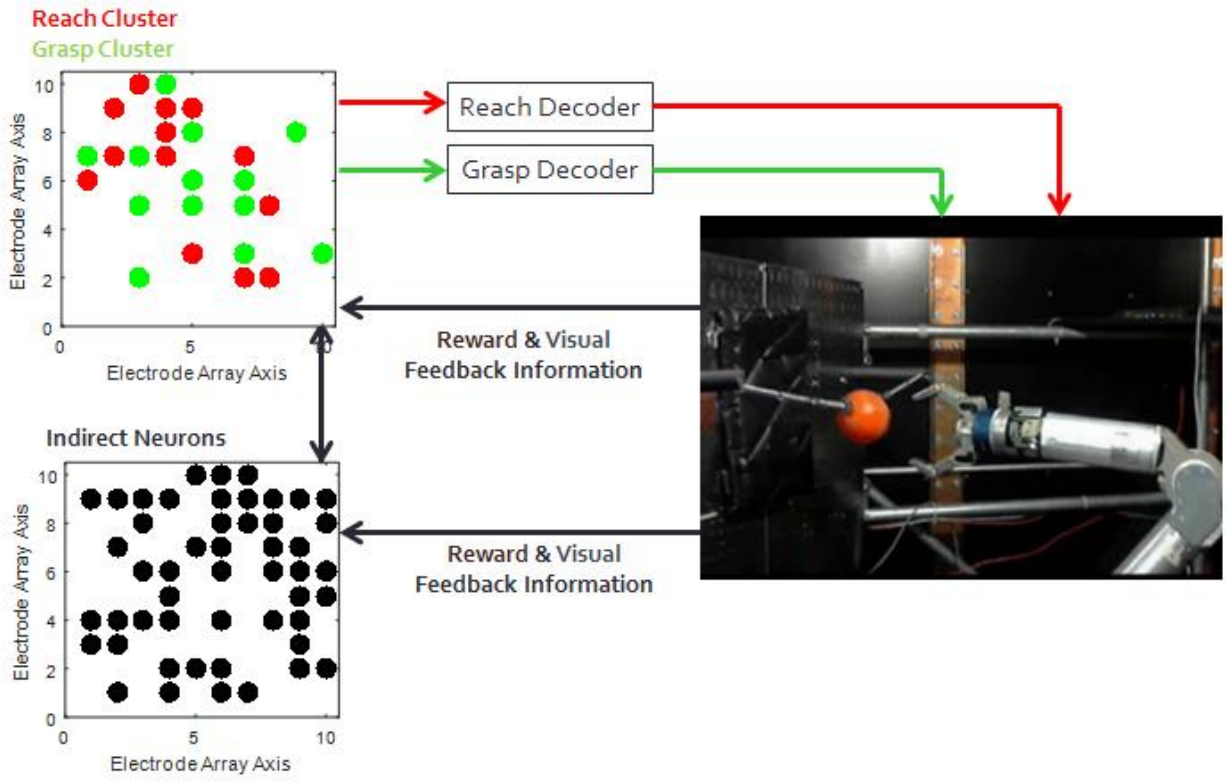

Figure 3.1 Experimental Setup The macaque was taught to use neural cluster to reach (red) and grasp (green) with a robotic arm. Twenty-tap Wiener filter decoders were initialized using spontaneous data from each ensemble and remained static over the course of the study. The macaques learned this mapping such that they could coordinate both control dimensions in order to perform successful trials. Instruction was provided through operant conditioning (reward feedback).

Results

Reach-to-Grasp Behavior

The macaques were operantly conditioned to make complete reach-grasp-pull-drop movements. For this study, we focused only on successful reach-to-grasp portions of each full trial-- that is from the beginning of every trial to the time of first successful grasp. The data were subdivided into shorter training periods based on frequency and length of daily training sessions: there were no more than three contiguous days within each training period where the macaques had been trained for less than 30 minutes, and we required at least two days of training in a period. There were two training periods for Monkey K of lengths 679 and 777 trials respectively and three training periods for Monkey Z of lengths 193, 238, 298 trials respectively. Towards the end of training, the animals were able to complete successful trials at faster rates than at the beginning of training (Figure 3.2), although the daily trial rates tended to also fluctuate based on the motivational state of the animal.

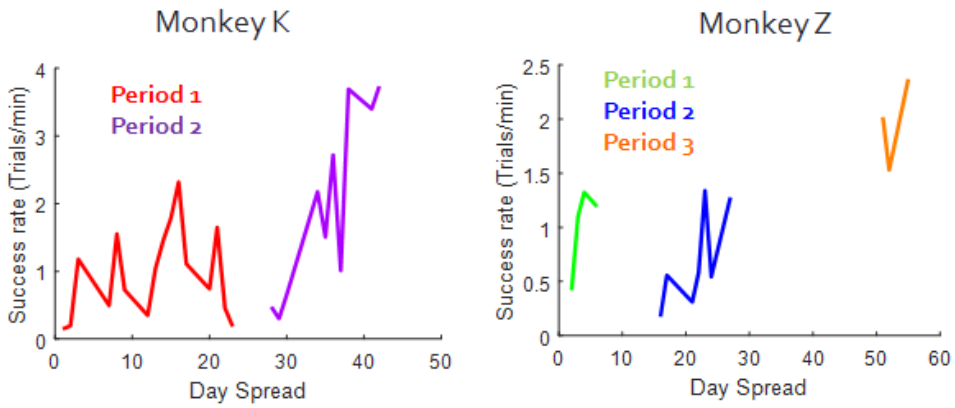

Figure 3.2 Success Rates Number of trials normalized by training time (trials/minute) over the course of the experiment. The data were subdivided into shorter training periods based on frequency and length of daily training sessions: there were no more than three contiguous days within each training period where the macaques had been trained for less than 30 minutes, and we required at least two days of training in a period.

Since no constraints were placed on the behavior during the trial, there were substantial inter-trial differences in the reach-to-grasp trajectories, both spatial and temporal in nature. This was especially evident when comparing the very first trial that both animals completed successfully with trials completed on the last day of training (Figure 3.3). Another potential confound in the data is that the macaques were not necessarily uniformly engaged throughout the course of a single training session. For example, some of the trials were on the order of minutes, while others were measured in seconds. Since each trial had to end with a successful grasp, followed by a pull and a drop, we focused on the very end of the reach-to-grasp movement: for the remainder of this analysis, we defined a trial as the last 500 milliseconds of the reach-to-grasp behavior before the time of first successful grasp. Similarly, since both static neural decoders for the reaching and grasping dimensions had twenty 50 millisecond taps, or 1 second of history, we examined the last 1500 ms of neural activity of both ensembles for each trial. In order to examine any learning effects within a single training period, each period was subdivided into three intervals composed of equal numbers of trials: beginning (1), middle (2), and end (3).

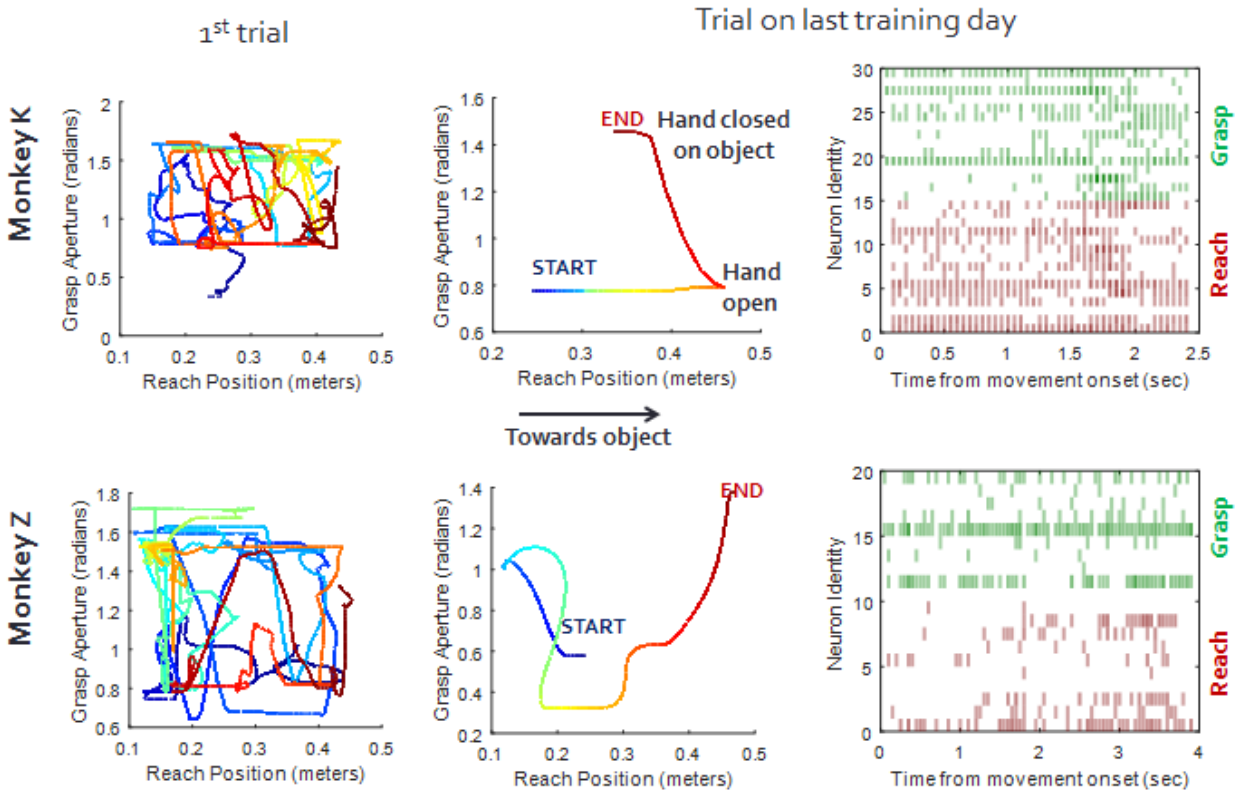

Figure 3.3 Example Trials Column 1: The robotic trajectories of the very successful first trials completed by each animal in this study plotted as a function of time (displayed as color gradient). Column 2: The robotic trajectories of the trials completed on the last day of training. Column 3: Raster plots of reach (red) and grasp (green) neurons corresponding the trials shown in Column 2.

For each interval, we used principal components analysis to reduce the dimensionality of the reaching and grasping velocity profiles concurrently. For all of the datasets, we found that one, two, or three principal components explained more of the variance in the last interval of each period as compared with the first interval of each period (Figure 3.4). This would indicate that the reaching and grasping trajectories that the macaques made with the robot became more structured over the course of each interval.

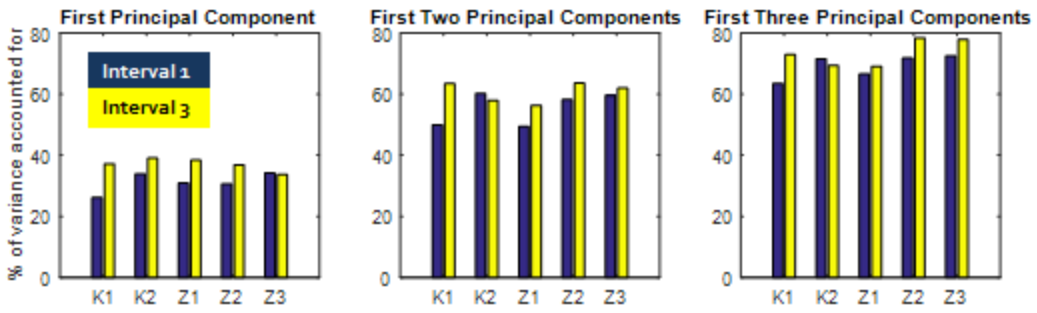

Figure 3.4 Principal Components Analysis Percentage of variance accounted by the first, first two, and first three principal components in the first (blue) and third (yellow) interval of each training period. (Xn = Monkey K: Period n)

We found stereotypical patterns in the how the macaques coordinated the reaching and grasping trajectories of the robot. The cross covariance patterns between the reaching and grasping velocity profiles tended to stereotyped for each animal. In most of the intervals, we found that Monkey K tended to open up the robotic hand while simultaneously moving the hand towards the object or, similarly, close the hand while moving the hand in the opposite direction (Figure 3.5). We also found that her closing movements of the hand tended to positively covary with her moving the arm forward or her opening of the hand tended to positively covary with the arm moving back, but with a of delay of 200 to 300 milliseconds (Figure 3.5). This patterning is consistent with a number of possible spatial trajectories. One strategy could be moving the arm forward while opening up the hand, and then closing around the object later on as the arm is still moving forward. Another strategy was made possible by the fact that the object was held in place by a spring, making it possible for the macaque to move the arm forward while opening up the hand, overshoot, and then close around the object while moving slightly back.

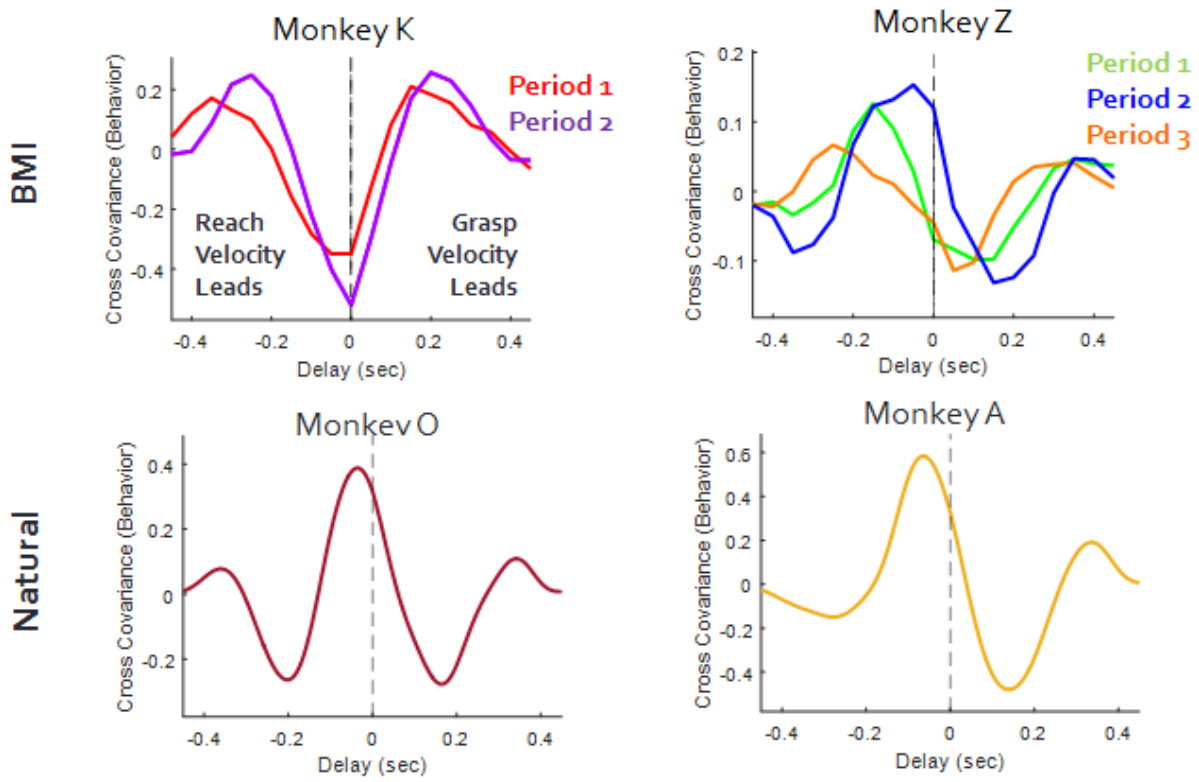

Figure 3.5 Cross Covariance Patterns: Velocity Profiles (top row) Cross covariance patterns between the reaching and grasping velocity profiles towards the end of each learning period for Monkey K and Monkey Z. (bottom row) Cross covariance patterns between the reaching and grasping velocity profiles of two animals, Monkey O and Monkey A, making reach-to-grasp movements with intact limbs.

In the case of Monkey Z, moving the hand forward tended to positively covary with opening of the hand and likewise moving the hand away from the robot tended to covary with closing the hand, but sometime with a delay of 50 or 100 milliseconds between the opening of the hand and reaching forward, or conversely, between the closing the hand and moving back (Figure 3.5). We also found the closing movements of the hands tended to positively covary with Monkey Z moving the arm forward, and conversely, opening movements positively covarying with the arm moving back, such that the grasping movement would lag the reaching movement by 50 to 250 milliseconds (Figure 3.5). As such, the cross covariance patterns between Monkey Z's reaching

and grasping velocity profiles tended to more closely reflect the cross covariance patterns of two monkeys, Monkey O and Monkey A, making natural reach-to-grasp movements with their intact limbs (Figure 3.5); Monkey O and Monkey A both tended to positively covary closing of their hands with reaching forward, but at a delay of around 50 milliseconds.

In order to examine how this patterning in cross covariance varied over the course of training, we calculated the Euclidean distance between the pattern of cross covariance on every day and the average pattern across up to three prior days of training (Figure 3.6). Over both the entire course of training and within each training period, the Euclidean distance between patterns tends to decrease edging towards zero, indicating that the cross covariance patterns were getting more stable.

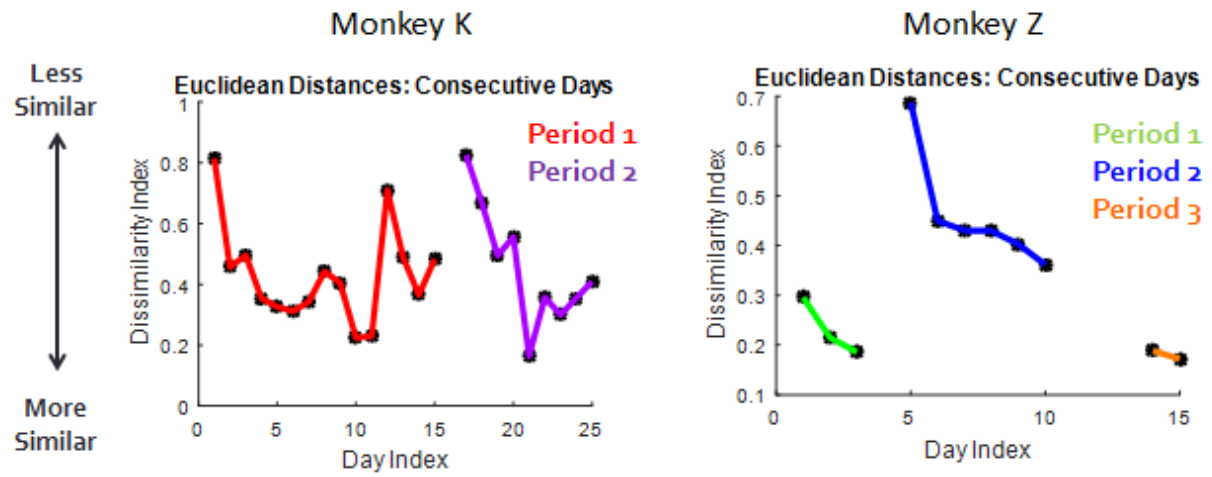

Figure 3.6 Stabilization: Behavior Euclidean distances between the patterns of cross covariance on every day and the average pattern across up to three prior days of training.

Reaching & Grasping Clusters

We probed the dynamics between pairs of neurons in the reaching and grasping clusters to investigate the neural correlates of this emerging coordinated behavior. For each pair, composed

of one reach neuron and one grasp neuron, we calculated the cross covariance between their firing activities in the last 1500 ms of every trial, and then found the average cross covariance across all of the trials for each day of training. Stereotyped patterns of cross covariance were observed to arise between pairs of neurons (Figure 3.7). For example, Figure 3.7 shows one pair of neurons that tend to fire in synchrony (*Monkey Z*) and one pair of neurons that tend to fire in close synchrony, but with the reach neuron leading the grasp neuron by about 50 milliseconds (*Monkey K*).

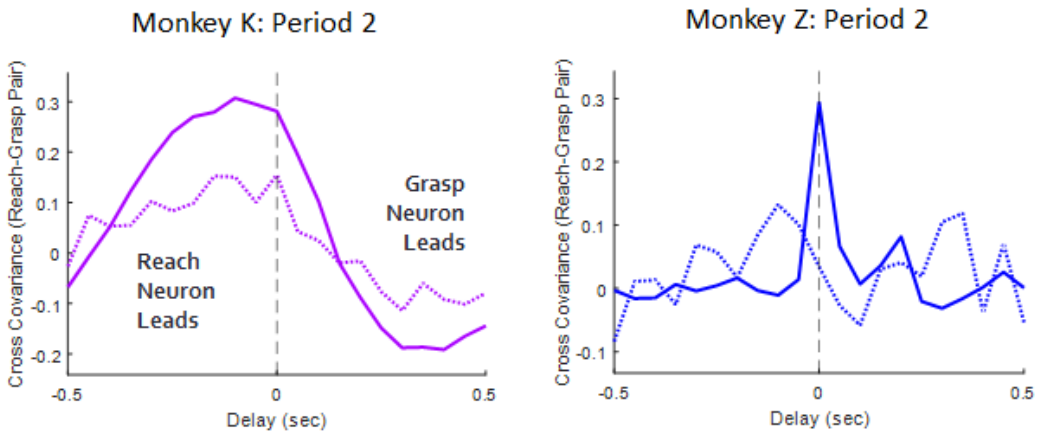

Figure 3.7 Cross Covariance Patterns: Example Reach Grasp Neural Pairs Patterns of cross covariance were observed to arise between pairs of reach and grasp neurons. Cross covariance on first day of training period (dotted line) vs cross covariance on last day of training period (solid line).

In order to examine how pair-wise cross covariance patterns evolved within training periods, we recalculated the cross variance patterns, this time grouping together trials during the beginning (1), middle (2), and end (3) intervals in each training period. In addition, being primarily interested in the coordination between reach and grasp neurons as it related to task-related behavior, we controlled for non-task-related correlated firing activity by shuffling the trial order 50 times for each pair of neurons, and then conducting the same analysis on the shuffled data for each interval.

For each interval, the cross covariance pattern between a reach-grasp neuron pair was considered significant if its peak value was greater than the peak values of 95% of the cross covariance patterns generated from the shuffled data (Figure 3.8), or if its trough value was less than the trough values of 95% of the cross covariance patterns generated from the shuffled data. Significant pairs would indicate that macaques tended to covary the modulation of that particular reach-grasp neuron pair to a greater extent, either positively or negatively, for the purpose of coordinating the reach-to-grasp behavior on a trial-by-trial basis. We probed the last interval in each of the training periods to find significant characteristic patterns that had emerged for each of the macaques. In the last interval of Monkey K's first and second training periods, we found that 61.33% and 64.44% of all potential reach-grasp pairs showed significant cross-covariance. In the last interval of Monkey Z's first and second training periods, we found that 45%, 58%, and 53% of all possible reach-grasp pairs showed significant cross-covariance.

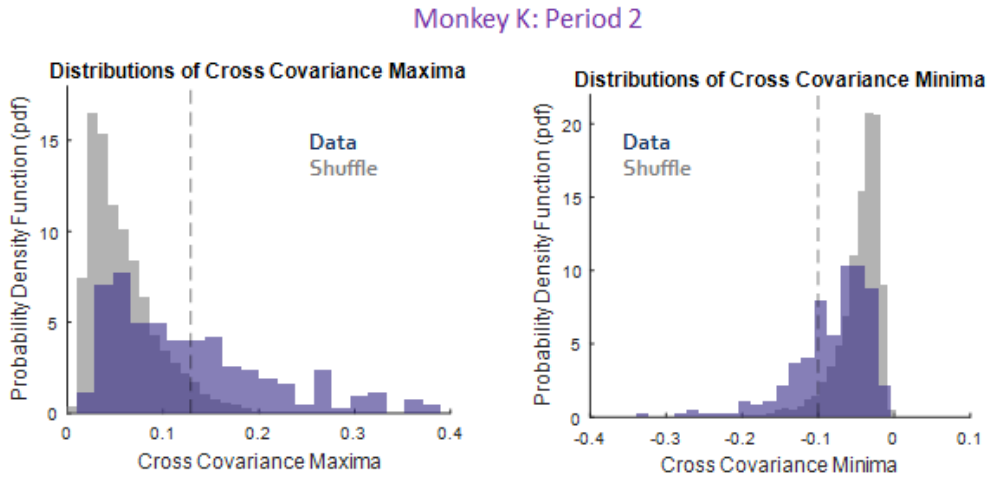

Figure 3.8 Shuffle Analysis We controlled for non-task-related correlated firing activity by shuffling the trial order 50 times for each pair of neurons. Distributions of cross covariance maxima and minima are shown for the original data and the shuffled data.

In order to examine the evolution of cross covariance patterns over the course of a training period, we investigated any shift in positive or negative extrema between the first and third intervals. First, we examined this shift in maxima and minima for all reach-grasp neural pairs. We found that shifts in the distribution such that the maxima would become more positive in Monkey K's 2nd period and Monkey Z's 2nd period (Wilcoxon rank sum test, $p<.05$). These shifts in maxima were evident when examining the average cross covariance across pairs in the first and last intervals (Figure 3.9). We also found shifts in the distributions such that the minima would become more negative in all of the training periods (Wilcoxon rank sum test, $p<.05$). Next, we examined shifts in the maxima and minima of reach-grasp neural pairs that showed significant cross-covariance in the third interval. We found shifts in the distributions such that the maxima would become more positive in both of Monkey K's periods and Monkey Z's 2nd and 3rd training periods (Wilcoxon rank sum test, $p<.05$) (Figure 3.10). We also found shifts in the distribution such that the minima would become more negative in both of Monkey K's periods and Monkey Z's 1st and 2nd training periods (Wilcoxon rank sum test, $p<.05$).

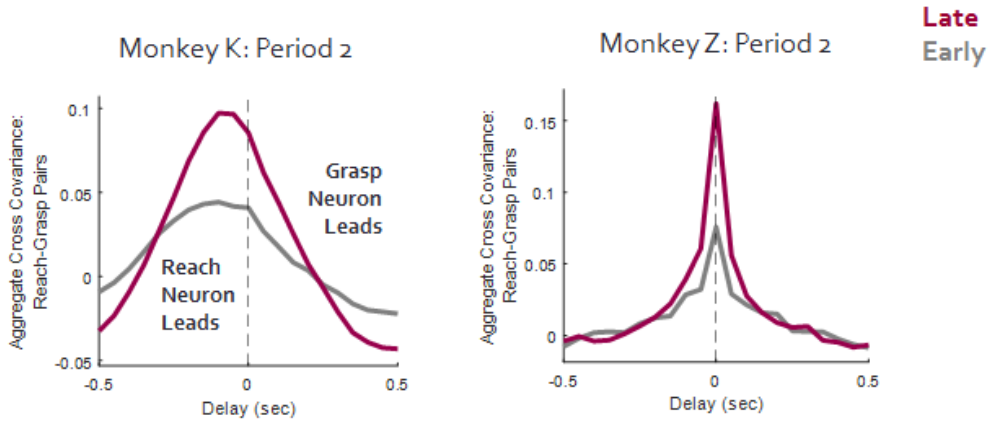

Figure 3.9 Aggregate Cross Covariance Patterns: Reach Grasp Pairs Average cross covariance of reach-grasp pairs in the first (gray) and third (magenta) intervals of Period 2 in each monkey.

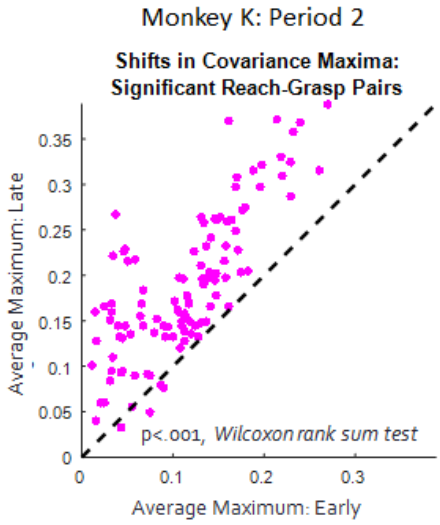

Figure 3.10 Shifts in Extrema: Reach Grasp Pairs Shifts in positive extrema between the first and third intervals, plotted for pairs of reaching and grasping neurons.

In order to examine the stability of cross covariance between pairs of reach neurons and grasp neurons in the same manner as with the reaching and grasping movements of the robot, for each reach-grasp neural pair, we calculated the Euclidean distance between the pattern of cross covariance on every day and the average pattern across up to three prior days of training (Figure 3.11). We then averaged these distances across all of reach grasp neural pairs. Over both the entire course of training and within each training period, the average Euclidean distance between patterns of cross covariance tends to decrease, indicating that the cross covariance patterns between pairs of reach and grasp neurons were becoming more stable.

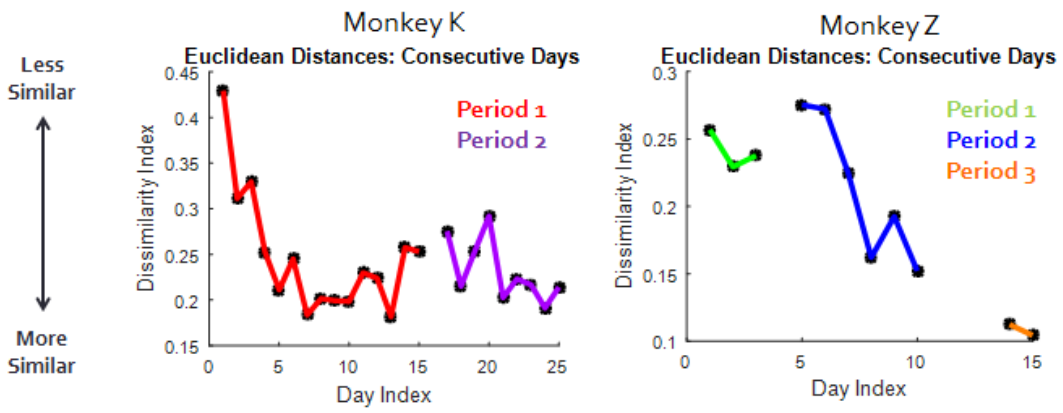

Figure 3.11 Stabilization: Reach Grasp Pairs Euclidean distances between the patterns of cross covariance on every day and the average pattern across up to three prior days of training.

Indirect Neurons

Given that the we saw significant patterning develop between reach-grasp neuron pairs, the next step was to investigate the emergence of similar patterning between neurons that were not included in either decoder, termed by others in the field as indirect neurons (Ganguly et al. 2011). We used all of the neurons that were not used for either decoder and were stable over the course of training for this analysis (Dickey et al. 2009b; Eleryan et al. 2014): 72 neurons for Monkey K and 70 neurons for Monkey Z.

We examined the dynamics between pairs of indirect neurons during this task to probe any larger network effects of the emergence of this coordinated reach-to-grasp behavior. As with the reach grasp neural pairs, for each pair of indirect neurons we calculated the cross covariance between their firing activities in the last 1500 milliseconds of every trial, and then found the average cross covariance across all of the trials for each day of training. Stereotyped patterns of cross covariance were observed to arise between pairs of indirect neurons (Figure 3.12). For example, Figure 3.12 shows two pairs of neurons that tend to fire in synchrony.

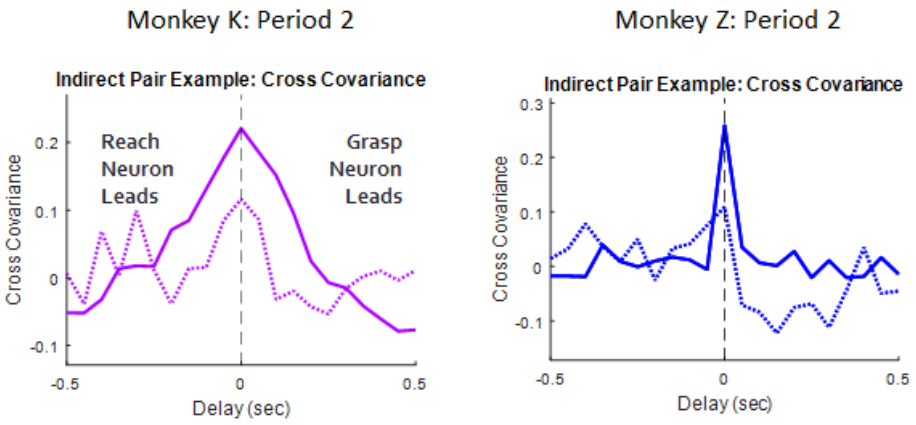

Figure 3.12 Cross Covariance Patterns: Example Indirect Neural Pairs Patterns of cross covariance were observed to arise between pairs of indirect neurons. Cross covariance on first day of training period (dotted line) vs cross covariance on last day of training period (solid line).

As with the analysis of reach-grasp neural pairs, we recalculated the cross variance patterns, grouping together trials during the beginning (1), middle (2), and end (3) intervals in each training period. We repeated the same shuffle analysis, controlling for non-task related correlated firing activity between indirect pairs, by shuffling the trial order 50 times for each pair of neurons, and then conducting the same analysis on the shuffled data for each interval.

The same rule was used to determine significance: the cross covariance pattern between an indirect pair was considered significant if its peak value was greater than the peak values of 95% of the cross covariance patterns generated from the shuffled data (Figure 3.8), or if its trough value was less than the trough values of 95% of the cross covariance patterns generated from the shuffled data. The existence of significant indirect pairs would indicate that the macaques tended to covary the modulation of neurons not directly involved in either decoder to a greater extent, either positively or negatively, for the specific purpose of coordinating the reach-to-grasp robotic task on a trial-by-trial basis. We probed the last interval in each of the training periods to

find significant characteristic patterns between indirect neurons that had emerged for each of the macaques. In the last interval of Monkey K's first and second training periods, we found that 57.2% and 64.44% of all potential indirect neuron pairs showed significant cross-covariance. In the last interval of Monkey Z's first and second training periods, we found that 32.34%, 64.93%, and 42.82% of all possible indirect neuron pairs showed significant cross-covariance.

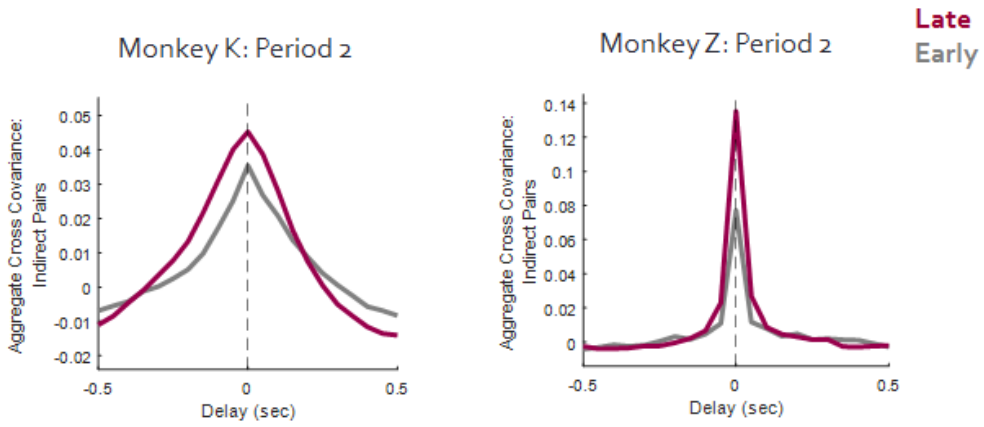

Figure 3.13 Aggregate Cross Covariance Patterns: Indirect Pairs Average cross covariance of indirect neuron pairs in the first (gray) and third (magenta) intervals of Period 2 in each monkey.

We examined shifts in the extrema of the cross covariance patterns between the first and third intervals, in order to examine the evolution of indirect neuron relationships. First, we examined this shift in maxima and minima for all indirect neuron pairs (Figure 3.13). We found shifts in the distribution such that the maxima would become more positive in both of Monkey K's periods and Monkey Z's 1st and 2nd training periods (Wilcoxon rank sum test, $p<.05$). These shifts in maxima were evident when examining the average cross covariance across pairs in the first and last intervals (Figure 3.13). We also found shifts in the distributions such that the minima would become more positive in both of Monkey K's periods and Monkey Z's 1st and 2nd training period (Wilcoxon rank sum test, $p<.05$). Next, we examined shifts in the maxima and

minima of indirect neuron pairs that showed significant cross-covariance in the third interval. We found shifts in the distributions such that the maxima would become more positive in all of the training periods (Wilcoxon rank sum test, $p<.05$) (Figure 3.14). We also found shifts in the distribution such that the minima would become more negative in all of the training periods (Wilcoxon rank sum test, $p<.05$).

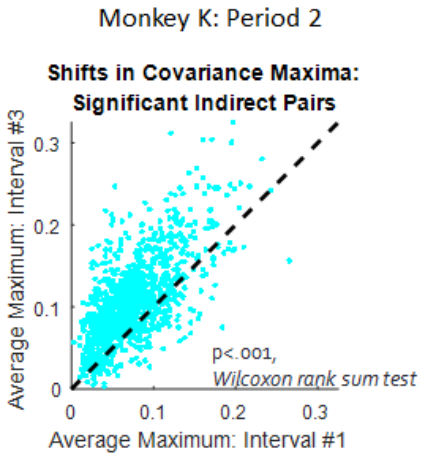

Figure 3.14 Shifts in Extrema: Indirect Pairs Shifts in positive extrema between the first and third intervals, plotted for pairs of indirect neurons

We examined the stability of cross covariance patterning between pairs of indirect neurons in the same manner as with the reach-grasp neural pairs; for each indirect neuron pair, we calculated the Euclidean distance between the pattern of cross covariance on every day and the average pattern across up to three prior days of training. We then averaged these distances across all of indirect neuron pairs. Both over the entire course of training and within each training period, the average Euclidean distance between patterns of cross covariance tends to decrease, indicating that the cross covariance patterns between pairs of indirect neurons were getting more stable (Figure 3.15). The sole training period in which this decrease was not observed, Monkey Z's 3rd

period, was the same period in which there was not a clear shift in the maxima or minima of the indirect neuron pair cross covariance patterns between first and third intervals.

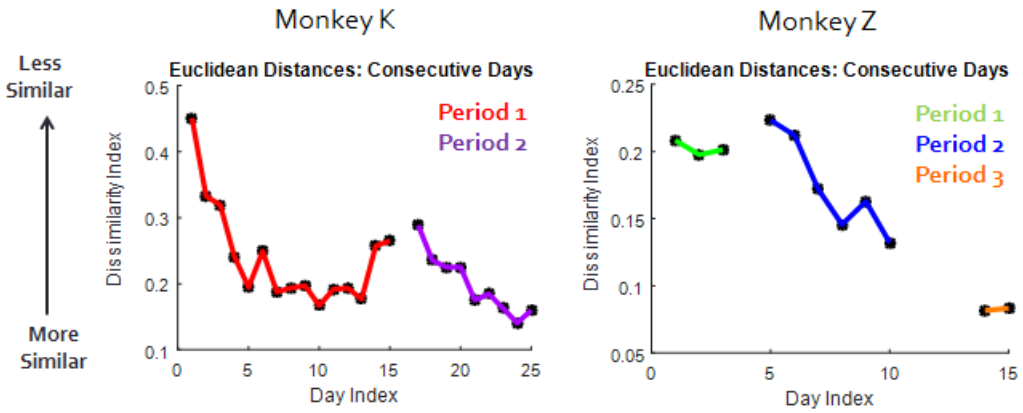

Figure 3.15 Stabilization: Indirect Pairs Euclidean distances between the patterns of cross covariance on every day and the average pattern across up to three prior days of training.

Network Comparison

Given that we observed similar trends and stabilization in the cross covariance patterns between pairs of reach and grasp neurons and pairs of indirect neurons, we directed our efforts to comparing patterning between pair types and with behavior. In the case of Monkey Z, we found that for both pair-types, neurons tended to modulate synchronously (Figure 3.15). In the case of Monkey K, we found that for both pair-types, pairs of neurons tend to modulate not only synchronously but also with one neuron leading or lagging the other over a range of delays (Figure 3.15).

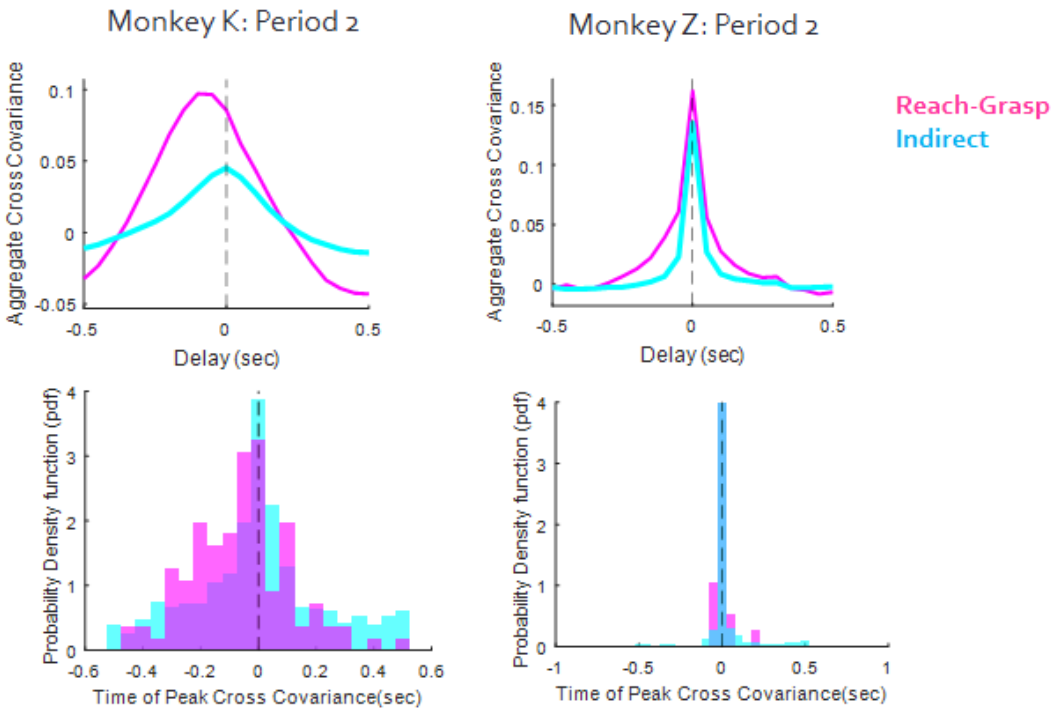

Figure 3.15 Cross Covariance Pattern Comparison (Row 1) Average cross covariance of reach-grasp neural pairs (pink) and indirect neuron pairs (cyan) in the third interval of Period 2 in each monkey. (Row 2) Distributions of peak cross covariance times between reach-grasp (pink) and indirect (cyan) pairs.

In order to compare the degree to which both the reaching and grasping and the indirect neural ensembles showed significant patterning, we repeated the shuffle analysis for every day of training. The cross covariance pattern between a pair on a given day was considered significant if its peak value was greater than the peak values of 95% of the cross covariance patterns generated from the shuffled data for that day (Figure 3.8), or if its trough value was less than the trough values of 95% of the cross covariance patterns generated from the shuffled data from that day. In the case of Monkey K, the percentage of all potential reach-grasp neural pairs that showed significant extrema tended to be larger than the percentage of all potential indirect neuron pairs that showed significant extrema (sign test, $p<.001$) (Figure 3.16). In the case of Monkey Z, this trend was not significant but close (sign test, $p<.09$); however, on the last days of all three

training periods, the percentage of all potential reach-grasp neural pairs that showed significant extrema was larger than the percentage of all potential indirect neuron pairs that showed significant extrema (Figure 3.16). In addition, in the case of Monkey K, we found that the percentage of significant reach-grasp pairs and the percentage of significant indirect neuron pairs on a given day were both significantly correlated with the performance of the animal on that day, specifically with the successful trial rate (reach-grasp: correlation coefficient= .649, p<.001; indirect: correlation coefficient= .706, p<.001) (Figure 3.17).

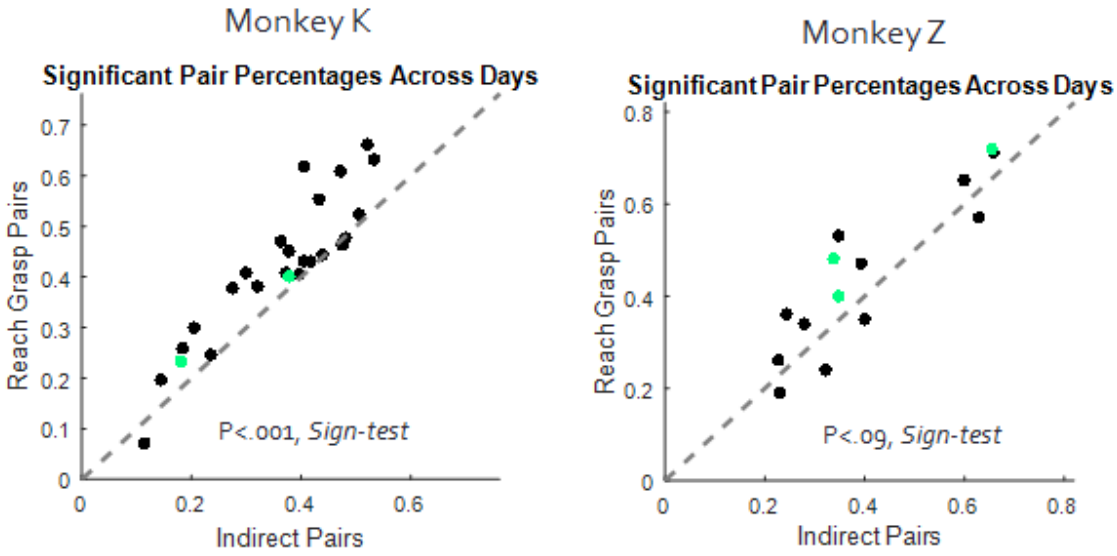

Figure 3.16 Significant Pairs Across Days The percentage of all potential reach-grasp neural pairs that showed significant extrema plotted against the percentage of all potential indirect neuron pairs that showed significant extrema for each day of training. Green points refer to the last training days in each training period.

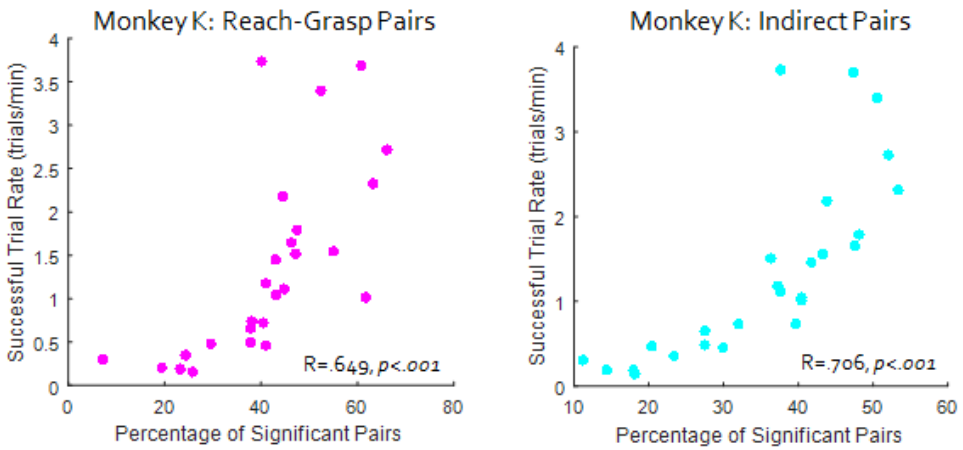

Figure 3.17 Significant Pairs & Performance For each day of training, the percentage of significant reach-grasp pairs (left) and the percentage of significant indirect neuron pairs (right) plotted against the inverse of the average time that it took the animal to reach-to-grasp the object on that day.

The extrema of the cross covariance patterns were compared between the reach-grasp and indirect neuron pairs on each training day. In the case of Monkey K, the maxima of the covariance patterns between the reach-grasp neural pairs were greater than the maxima of the covariance patterns between the indirect neuron pairs (sign test, $p<.001$) (Figure 3.18), and the minima of the covariance patterns between the reach-grasp neural pairs were more negative than the minima of the covariance patterns between the indirect neuron pairs (sign test, $p<.001$) (Figure 3.5d). In the case of Monkey Z, this trend was present for the maxima but not significant (sign test, $p<.09$) (Figure 3.18); in addition, on the last days of two of the three training period, the maxima of the covariance patterns between the reach-grasp neural pairs was greater than the maxima of the covariance patterns between the indirect neuron pairs (Figure 3.18).

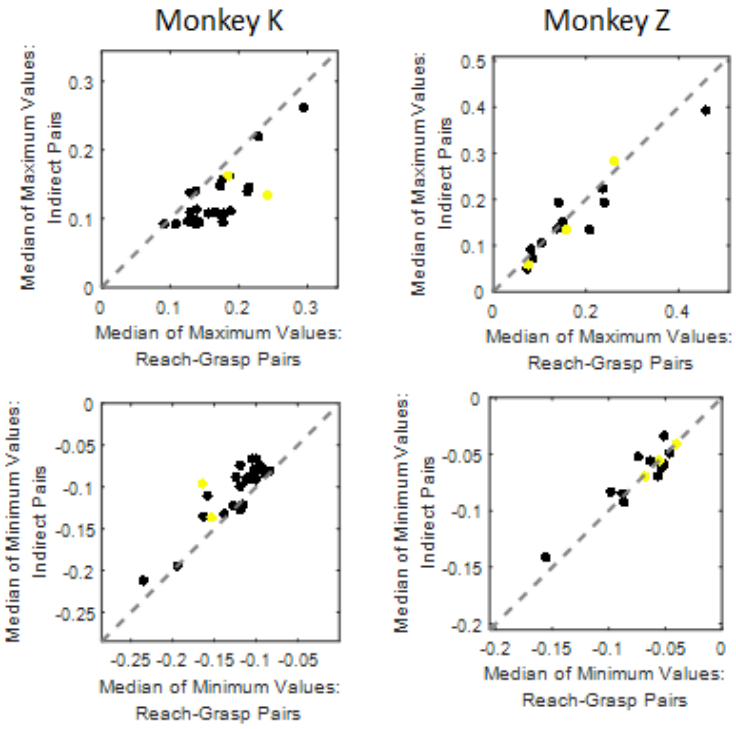

Figure 3.18 Extrema Comparison: Neural Pairs The median extremum of the cross covariance patterns between the reach-grasp neurons plotted against the median extremum of the cross covariance patterns between indirect neuron pairs for each day of training. Yellow points refer to the last training days in each training period.

Having compared the extrema between the reach-grasp neural pairs and the indirect neuron pairs, it is important to examine them both in the context of behavior. To do so, we examined the relationship of the cross covariance between the reaching and grasping robot trajectories and the cross covariance between pairs of neurons on each training day. In the case of Monkey K, we found that the maxima of the cross covariance patterns between both reach-grasp neuron pairs and indirect neuron pairs were both highly correlated with the maxima of the cross covariance patterns between the reaching and grasping robot trajectories (reach-grasp: correlation coefficient=0.69, $p<.001$; indirect: correlation coefficient=0.414, $p<.05$) (Figure 3.19). We also found that minima of the cross covariance between reach-grasp pairs was negatively correlated

with the maxima of the cross covariance between the reaching and grasping trajectories (correlation coefficient= -0.557, p<.005), and conversely, the maxima of the cross covariance between reach-grasp pairs were negatively correlated with the minima of the cross covariance between the reaching and grasping trajectories (correlation coefficient = -.445, p<.05). In the case of Monkey Z, we found that the minima of the cross covariance patterns between both reach-grasp pairs and indirect pairs were highly correlated with the minima of the cross covariance between the reaching and grasping robot trajectories (reach-grasp: correlation coefficient=.811, p<.001; indirect: correlation coefficient= .769, p<.005). We also found that the minima of the cross covariance patterns between both reach-grasp pairs and indirect pairs were negatively correlated with the maxima of the cross covariance between the reaching and grasping trajectories (reach-grasp: correlation coefficient= -.595, p<.05; indirect: correlation coefficient= -.695, p<.001).

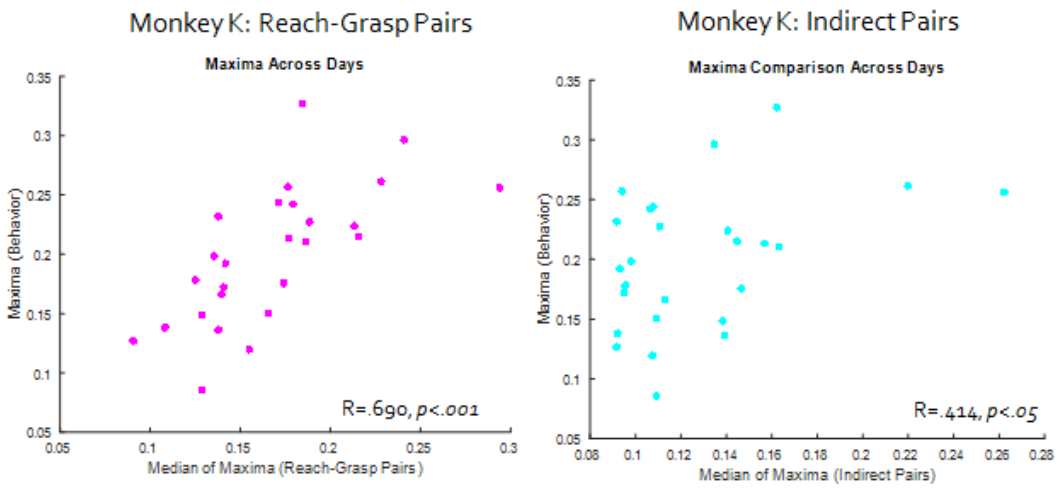

Figure 3.19 Extrema Comparison: Neural Pairs & Behavior Maxima cross covariance patterns between both reach-grasp neuron pairs (right, pink) and indirect neuron pairs (left, cyan) were both significantly correlated with the maxima of the cross covariance patterns between the reaching and grasping robot trajectories.

Discussion

We have introduced a paradigm for studying neural adaptation associated with the learning of a novel behavior at the level of neocortex: studying the emergence of coordinated reach-to-grasp behavior in rhesus macaques that have been the recipients of therapeutic amputations, and are being taught to control a robotic arm via a cortically controlled BMI through the use of operant conditioning. Our paradigm posits a new framework to study the learning of a novel coordinated motor behavior not only at the behavioral level, but also at the level of primary motor cortex.

Towards the end of training the animals were able to complete successful trials at faster rate than at the beginning of training (Figure 3.2), although the daily trial rates tended to also fluctuate based on the motivational state of the animal. We found that the reaching and grasping robot trajectories became more structured over the course of each training period, as evidenced by the increase in variance accounted for by one, two, or three principal components (Figure 3.4). We found stereotypical patterns in the cross covariance between the reaching and grasping velocity profiles. Both animals tended to open the hand while moving the arm towards the object, either concurrently, or with a slight day between the opening movement and reaching forwards (Figure 3.5). In both animals, the closing of the hand tended to occur concurrently with the arm moving back, or at a delay with respect to the arm reaching forward, strategy which more closely mirrors natural reach-to-grasp with intact limbs (Figure 3.5). We found that cross covariance patterns were stabilizing, both within each training period and over the entire course of training, consistent with human psychophysics studies in children (Kuhtz-Buschbeck, Stolze, Jöhnk, et al. 1998).

The macaques learned to coordinate reaching-to-grasp using neurons that were not initially marked as reach-related or grasp-related Several studies have shown that after amputation, the cortical area previously involved in the control of the lost limb undergoes reorganization (J. N. Sanes et al. 1988; J. N. Sanes, Suner, and Donoghue 1990; Schieber and Deuel 1997; Wu and Kaas 1999; Qi, Stepniewska, and Kaas 2000), indicating that patients may not have the same degree of representation of the lost limb in primary motor cortex. Instead of being reliant on finding "hand-neurons" or "arm-neurons" in a patient, our approach uses operant conditioning to train subjects to use groups of neurons to control a robotic arm. This kind of paradigm thus has tremendous implications for developing prosthetics for this community, by taking advantage of the plasticity afforded to us in motor cortex. It also provides us with a unique window into studying this plasticity. Since we established a direct, causal relationship between the neurons in the reaching and grasping clusters and the robotic arm, we can infer that any learning at the behavioral level is directly due to plasticity among those neurons. Further, we can examine any learning effects among neurons not involved in BMI control.

We found stereotyped patterns of cross covariance between pairs of reaching and grasping neurons and between pairs of indirect neurons. In one macaque, we found that for both pair-types, neurons tended to modulate synchronously, in the other, we found that for both pair-types, pairs of neurons tend to modulate synchronously or with one neuron leading or lagging the other over a range of delays (Figure 3.15). In all of the training periods, at least 45% of all potential reach-grasp pairs and 32% of all possible indirect neuron pairs had significant extrema in the last interval. We found that the percentage of all potential reach-grasp pairs that showed significance was larger than the percentage of all potential indirect neuron pairs that showed significance on

the last day of each training period. In the case of reach-grasp neural pairs, we found that the maxima of the cross covariance patterns became more positive and the minima became more negative from the first interval to the third interval in four of the five training periods. In the case of indirect neuron pairs, this held for all five training periods. The maxima for the reach-grasp pairs were generally larger than the maxima for the indirect neuron pairs, especially on the last day of each training period.

Findings from a previous study in this field are partially consistent with our results. In this study, animals that had extensive experience with brain control had to learn a fixed, novel decoder for a center-out reaching task (Ganguly et al. 2011). Learning the decoder involved modifications in the modulation of both neurons involved in the decoder and indirect neurons; a similar fraction of both experienced a change in preferred direction. This is consistent with our finding that comparable percentages of reach-grasp pairs and indirect pairs show significant cross covariance patterns in the last interval of each training period. However, this study found that the neurons directly involved in the decoder tended to modulate to a greater degree relative to indirect neurons (Ganguly et al. 2011). Although we found that cross covariance extrema generally tend to be more extreme between reach-grasp pairs than indirect pairs, this was not a strong trend in Monkey Z, and in Monkey K, the extrema were of comparable magnitudes. Of course, in our study, we are examining the relationships between neurons as opposed to the modulation of individual neurons, so the two findings may not be mutually exclusive.

Finally, we found that cross covariance patterns between pairs of neurons, both reach-grasp and indirect, were stabilizing both within training periods, and over the entire course of training,

much like the cross variance patterns in behavior. In fact, we found a high degree of correlation between the extrema of the neural cross covariance patterns and the extrema of the behavioral cross covariance patterns. Across all of the pairings of neural extrema and behavioral extrema where we found a high degree of correlation, there was one consistent trend: maxima tended to be positively correlated with maxima, minima tended to be positively correlated with minima, and maxima and minima were negatively correlated with one another. In essence, this would indicate two key findings: first, the degree to which a pair of neurons covary their modulation, either positively or negatively, has a direct effect on the extent to which the behavior is coordinated; second, this general relationship is consistent across animals. In Monkey K, we found another direct relationship with behavior: the percentage of significant reach-grasp pairs and the percentage of significant indirect neuron pairs on a given day were both significantly correlated with the Monkey K's performance on that day. In short, how well Monkey K performed on any given day directly corresponded with the degree to which neurons in the network significantly covarying their firing activities to engage in the task.

This paradigm gave us a unique model for both studying the development of novel, coordinated reach-to-grasp at the behavioral level and cortical level, and for examining the plasticity afforded to us by primary motor cortex through the lens of learning to use a BMI. In addition, the possibility of training subjects to perform a naturalistic motor task by modulating cortical neurons using an artificial mapping and of potentially influencing network activity in a lasting manner could have clinical implications for patients with amputations or strokes.

Acknowledgements

The authors would like to thank J. Coles for assistance with the training and care of laboratory

animals.

Grants

Research supported by DARPA Grant No. N66 001-12-1-4023, IGERT: Integrative Research in Motor Control and Movement. National Science Foundation Grant # DGE-0903637

IV. DISCUSSION & CONCLUDING REMARKS

Reaching to grasp requires a remarkable degree of spatial and temporal coordination. The psychophysics literature on coordinating reach-to-grasp has been well-flushed out: the temporal and spatial coordination of transport and grasp components during reach-to-grasp behavior has been studied extensively (M. Jeannerod 1984; Haggard and Wing 1995), and the development of this coordination has been studied in infants and children (Claes von Hofsten 1984; Kuhtz-Buschbeck, Stolze, Jöhnk, et al. 1998; Wimmers et al. 1998). The hand preshapes itself during transport such that when the arm reaches the object, the grasp aperture and orientation of the hand match the required configuration to grasp the object. However, our knowledge on the role of neocortex in generating, driving, or modulating reach-to-grasp is quite limited. In this work, we have both examined the coupled dynamics between ensembles of reaching-related and grasping-related neurons in primary motor cortex during a 3-dimensional reach-to-grasp task with an intact upper limb, and investigated the emergence of this coupling by studying the learning of a reach-to-grasp behavior in a brain-machine interface context in chronically amputated animals.

Neural Coordination During Reach to Grasp

In the first study, we introduced a novel approach to quantifying the coordination dynamics of neural ensembles. We used this approach to study the functional interaction between ensembles of reaching-related and grasping-related neurons in the primary motor cortices of macaques that reached-to-grasp a variety of different objects at different locations. By describing the dynamics of these two ensembles as trajectories in a low-dimensional state space, we examined their

coupling in time. We found that these trajectories were temporally coupled over the course of movement such that any transient asynchrony in time progression between the two trajectories was corrected for by a proportional control mechanism. For example, if the trajectory of the reaching-related neural ensemble led the grasping-related neural ensemble, the grasping-related neural ensemble tended to catch up—and at a speed proportional to the lead time.

One interesting finding that was echoed in the results of the second experiment is that these compensation dynamics were still present for random partitions of the population, but were slower as compared to the original reach-related vs grasp-related partition (Figure 2.5). Given that primary motor cortex forms a highly interconnected network of neurons with highly overlapping movement representations (Park et al. 2001; Park, Belhaj-Saïf, and Cheney 2004; C. E. Vargas-Irwin et al. 2010; Saleh, Takahashi, and Hatsopoulos 2012), this result is unsurprising, and points to these dynamics as a broader property of the upper limb area of motor cortex. What is telling is that the time constants for the compensation effect between reaching- and grasping-related ensembles was smaller than for the random partitions, and ranged from around 10ms to 90ms. Taken in conjunction with the Granger causality analysis, which demonstrated that current temporal deviations in the neural trajectory of one ensemble could not be explained solely from the past deviations of the limb kinematics but required knowledge of the past deviations of the second ensemble's trajectory, and that the causal interactions between the two ensembles deviations were present at very fast timescales, on the order of 8ms, these results suggest that these cortico-cortical interactions are direct and not mediated solely through indirect interactions through the periphery .

Another curious finding of this study is that in all of the datasets, the Granger causal relationships were stronger from grasp-related to reach-related neural deviations than vice versa, although bidirectional relationships were present in three of the four datasets. Proximal-to-distal sequencing has been well-explored in the psychophysics literature for movements such as throwing, jumping, and hitting (Danion and Latash 2010); proximal muscles activate before more distal muscles, and proximal joints begin moving before more distal joints. However, reaching-to-grasp with a precision grip, defined as when the thumb is "abducted in both the metacarpophalangeal and the carpometacarpal joints," is a complex and more evolutionarily recent type of movement (Landsmeer 1962; Goodale 1990). A potential hypothesis for the asymmetric granger causal relationships between the reach and the grasp neural deviations is that neocortex, in this case primary motor cortex, plays a larger role in the learning and control of more complex and evolutionary recent movements such as forming a precision grip, as opposed to reaching. It then follows that while subjects are reaching-to-grasp, a greater fraction of neurons in primary motor cortex would be reliant on grasp-related information or feedback about the current state of grasping-related neural ensembles as opposed to the current state of reaching ensembles.

In a similar vein, corticomotoneuronal (CM) cells have been thought of as the infrastructure that allows us the flexibility to sculpt novel, highly skilled movements since they directly covey cortical motor information to motoneurons and "bypass" spinal cord interneuron circuitry entirely (Rathelot and Strick 2009). The existence of such CM neurons can be correlated with a higher degree of skilled movement, as evidenced by comparing the cebus and the squirrel

monkey. Biomechanically, they have similar hands and ecologically, they live in similar environments. However, cebus monkeys can use their digits relatively independently; they can grasp small objects and manipulate tools. In contrast, squirrel monkeys "can only pick up small objects by using a sweeping motion of the hand that involves all of the fingers acting in concert" (Rathelot, 2009). Cebus monkeys have prominent CM neurons, while squirrel neurons have weak direct connections at best; cebus monkeys have dense corticospinal terminations in the ventral horn, while squirrel monkeys have sparse ones at best (Bortoff, 1993). Direct CM neuron connections are not present at birth, but start developing in the first few months and continue developing for the first two years, which parallels development of motor skills, including the ability to move digits relatively independently (Rathelot, 2009).

Corticomotoneuronal cells have been observed to make more potent terminations in the motoneuron pools associated with distal muscles (McKiernan et al. 1998). In addition, this study examined to what extent corticomotoneuronal cells produced post-spike effects in the proximal and distal muscles of the upper limb while monkeys performed a reach-to-grasp task. They found post-spike effects produced by 112 or the 174 neurons they examined; of these, less than 10% produced post-spike effects only in the proximal muscles of the upper-limb, around 45% produced post-spike effects in only the distal muscles, and the other 45% produced post-spike effects in both groups.

Further lending support to the idea that neocortex plays a larger role in the control of grasp are a set of experiments by Kuypers in the 1960s: in these studies, monkeys received bilateral lesions to their pyramidal tracts (Donald G. Lawrence and Kuypers 1968; D. G. Lawrence and Kuypers

1968). They found that these monkeys were able to generally recover independent upper limb movement; however, they never regained the ability to individually move their fingers, even after recovery periods of up to eleven months. This work, coupled with the frequency and potency of the corticomotoneuronal projections to the motoneuron pools of distal muscles, would support the idea that the motor cortex plays a more significant role in grasping as opposed to reaching.

It would be interesting and useful to be able to test such a hypothesis. Our approach is limited in revealing how neural dynamics between reaching-related and grasping-related neural ensembles causally affect coordination. In psychophysics studies, behavioral perturbations are used to establish causality (Haggard and Wing 1995). Similarly, as these techniques become more sophisticated in the future, electrical or optogenetic stimulation could be used to establish causality (Diester et al. 2011; Liu et al. 2012). These techniques could eventually provide the ability to more accurately measure the behavioral and neural responses to specific deviations in neural trajectories, as well as the information content of this compensation effect that we observed.

Our work extended previous research that has focused on quantifying the dynamics of individual populations by examining their trajectories in the low dimensional projection (M. M. Churchland et al. 2010; M. Churchland et al. 2012; Mante et al. 2013; Shenoy, Sahani, and Churchland 2013; Kaufman et al. 2014). Our approach generalizes this approach, by dividing the population into functionally defined subsets and analyzing the relations between trajectories during behavior. This could be adapted to analyze coordination within any brain region and, more importantly,

between brain areas. Future work (my own) will include looking at the dynamics between ensembles of neurons in premotor cortex, primary motor cortex, and prefrontal cortex, as well as looking at these dynamics between spikes and local field potentials.

In this study, we uncovered a proportional compensation mechanism for recovering from asynchronies between the trajectories of the reaching and grasping neural ensembles. Of course, this is just scratching the surface of the massive, highly connected, and anatomically diverse structure that is primary motor cortex. The field of dynamics and control has already provided us with insight and analytical tools for motor neuroscience (Todorov and Jordan 2002; Todorov 2004; Scott 2004; Cheng and Sabes 2006). Asking more intricate questions about the dynamics of microcircuits within primary motor cortex will involve using more recent techniques to record at a much finer spatial resolution, and from a far larger sample of neurons (Ahrens et al. 2013; Sadovsky et al. 2011).

Emergence Dynamics Underlying Learning to Reach-to-Grasp with a BMI

In the second study, we introduced a paradigm for studying neural adaptation associated with the learning of a novel behavior at the level of neocortex: studying the emergence of coordinated reach-to-grasp behavior in rhesus macaques that have been the recipients of therapeutic amputations, and were taught to control a robotic arm via a cortically controlled BMI through the use of operant conditioning.

We found that both macaques were able to complete more successful trials with training (Figure 3.2), and that the reaching and grasping robot trajectories became more structured over the

course of each training period, as evidenced by the increase in variance accounted for by one, two, or three principal components (Figure 3.4). Further analysis, along with more sophisticated dimensionality reduction techniques, could uncover specific structure that emerges in the reach-to-grasp behavior that is relevant to the development of coordination. We found that behavioral cross covariance patterns were stabilizing over the course of training, consistent with human psychophysics studies in children (Kuhtz-Buschbeck, Stolze, Jöhnk, et al. 1998). Future work (my own) will compare and contrast the stereotypical patterns observed in the cross covariance between the reaching and grasping velocity profiles in this learning task with the cross covariance patterns between the reaching and grasping kinematic trajectories of over-trained macaques.

The macaques learned to coordinate reaching-to-grasp using neurons that were not "pre-tuned" for reaching or grasping movements. Studies have shown that after amputation, the cortical area previously involved in the control of the lost limb undergoes reorganization (J. N. Sanes et al. 1988; J. N. Sanes, Suner, and Donoghue 1990; Schieber and Deuel 1997; Wu and Kaas 1999; Qi, Stepniewska, and Kaas 2000), indicating that patients may not have the same degree of representation of the lost limb in primary motor cortex. Instead of being reliant on finding "hand-neurons" or "arm-neurons" in a patient, our approach uses operant conditioning to train subjects to use un-marked groups of stable neurons to control a robotic arm. This kind of paradigm thus has tremendous implications for developing prosthetics for this community, by taking advantage of the adaptation afforded to us in motor cortex. In order to better cater to this community, further work must be conducted on any differential learning effects between using reorganized

cortex and using the cortex ipsilateral to the amputation, along with careful examination of any change in the behavior of the intact limb while this training is taking place.

Given that we established a direct, causal relationship between the neurons in both clusters and the robotic arm, we can infer that any learning at the behavioral level is directly due to adaptation among those neurons. This entire experiment only made use of two channels of feedback in the control of the robotic arm—visual feedback about the state of the arm, and reward feedback for the sake of shaping the behavior and upon completion of successful trials. Given that this paradigm affords us a unique window into studying plasticity in primary motor cortex during learning, by removing channels or adding feedback channels to provide auditory or somatosensory feedback, we could investigate the role of different types of sensory feedback in driving plasticity in motor cortex. In addition, this work could have potential bearing on how best to augment prosthetics with sensory feedback.

In this study, we found stereotyped patterns of cross covariance between significant portions of reach-grasp neural pairs and significant portions of indirect neuron pairs. In a prior study, animals that had extensive experience with brain control had to learn a fixed, novel decoder for a center-out reaching task (Ganguly et al. 2011). Learning the decoder involved modifications in the modulation of both neurons involved in the decoder and indirect neurons; a similar fraction of both experienced a change in preferred direction. This was consistent with our finding that comparable percentages of reach-grasp pairs and indirect pairs showed significant cross covariance extrema in the last interval of each training period.

On the other hand, this study also found that the neurons directly involved in the decoder tended to modulate to a greater degree relative to indirect neurons (Ganguly et al. 2011). Although we found that cross covariance extrema generally tend to be more extreme between reach-grasp pairs than indirect pairs, this was not a strong trend in Monkey Z, and in Monkey K, the extrema were of comparable magnitudes. However, in our study, we are examining the relationships between neurons as opposed to the modulation of individual neurons, so the two findings may not be mutually exclusive. Future work ought to attempt to decouple these two phenomena to see if these results can hold in parallel.

Recent BMI work has already been shifting away from the traditional paradigm of pre-tuned decoders through the use of adaptive decoding or by studying what happens when tuned decoders are perturbed during training (Ganguly et al. 2011; Ganguly and Carmena 2009; Taylor, Tillery, and Schwartz 2002; Li et al. 2011; Mussa-Ivaldi and Miller 2003; Willett et al. 2013; Orsborn et al. 2012). Our training paradigm relied on the animal learning a static, fixed decoder on its own; in reality, a hybrid approach that makes use of both the neural adaptation associated with subjects learning to use a fixed decoder and our capacity to adapt decoders to compensate for neural instabilities and to assist subjects in the process of learning, is highly preferable in a clinical setting.

Even with this goal in mind, there are benefits to studying neural adaptation on its own, and not in the context of a hybrid system. Future work could shed light on whether some neurons are better suited for decoder use, and on the characteristics of this subset of neurons if it exists. After all, in our study, we had up to 15 neurons controlling one dimension on the robot—there are still

a myriad of potential solutions to producing a given velocity even after taking the coefficient assignments at the beginning of training into account. Yet, the system converged to smaller subset of these solutions; future work could imagine what factors dictate convergence or influence the solution in greater depth.

Stability

Our second study, all of the future brain-machine interface work that I have proposed, as well as every cortically-controlled brain machine that has any semblance of a clinical future relies on our ability to implant robust, chronic, long-lasting electrode arrays that provide us with reliable and stable signals. The development of such instrumentation is the true bottleneck of the BMI field. Until that happens, a greater body of work needs to be built around characterizing the stability of signals on electrode arrays that are in use today; this includes factors such as neural lifespans on the electrode array (see **Appendix A**), implantation techniques, electrode impedances, sensitivity to spike sorting, immune responses of the subjects after surgery, amongst other actors.

Concluding Remarks

The results discussed in this thesis begin to answer the question: how do neural ensembles in the primary motor cortex functionally interact to produce coordinated behavior? In the first study, we looked at how motor cortical neurons participated in coordination by examining the interactions between populations of reach- and grasp-related neurons during reach-to-grasp behaviors in an intact primate model. We found that the characteristic dynamics between these two ensembles were also present, but to a lesser extent, between ensembles generated from

random partitioning of the population. We found evidence for temporal compensation across such that if one neural trajectory led in time, the other tended to catch up, reducing the asynchrony between the trajectories. In the second study, we probed the emergence of coordinated reach-to-grasp in chronically amputated macaques that were being taught to cortically control a robotic arm through operant conditioning. We found that stereotypical patterns emerged and stabilized in the cross-covariance between the reaching and grasping velocity profiles, between pairs of neurons involved in decoding reach and grasp, as well as between other stable neurons in the network. In addition, we found that the degree to which a pair of neurons covary their modulation, either positively or negatively, has a direct effect on the extent to which the behavior is coordinated. Taken together, these sets of results provide a framework for characterizing population dynamics in primary motor cortex during coordinated reach-to-grasp, and for probing their emergence during the development of coordinated reach-to-grasp.

REFERENCES

Ahrens, Misha B., Michael B. Orger, Drew N. Robson, Jennifer M. Li, and Philipp J. Keller. 2013. "Whole-Brain Functional Imaging at Cellular Resolution Using Light-Sheet Microscopy." *Nature Methods* 10 (5): 413–20. doi:10.1038/nmeth.2434.

Ambike, Satyajit, and James P. Schmiedeler. 2013. "The Leading Joint Hypothesis for Spatial Reaching Arm Motions." *Experimental Brain Research* 224 (4): 591–603. doi:10.1007/s00221-012-3335-x.

Armed Forces Health Surveillance Center (AFHSC). 2012. "Amputations of Upper and Lower Extremities, Active and Reserve Components, U.S. Armed Forces, 2000-2011." *MSMR* 19 (6): 2–6.

Badreldin, I., J. Southerland, M. Vaidya, A. Eleryan, K. Balasubramanian, A. Fagg, N. Hatsopoulos, and K. Oweiss. 2013. "Unsupervised Decoder Initialization for Brain-Machine Interfaces Using Neural State Space Dynamics." In *2013 6th International IEEE/EMBS Conference on Neural Engineering (NER)*, 997–1000. doi:10.1109/NER.2013.6696104.

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.

Batista, Aaron P., Christopher A. Buneo, Lawrence H. Snyder, and Richard A. Andersen. 1999. "Reach Plans in Eye-Centered Coordinates." *Science* 285 (5425): 257–60. doi:10.1126/science.285.5425.257.

Bressler, S. L., and J. a. S. Kelso. 2001. "Cortical Coordination Dynamics and Cognition." *Trends in Cognitive Sciences* 5 (1): 26–36.

Caminiti, R., A. Genovesio, B. Marconi, A. B. Mayer, P. Onorati, S. Ferraina, T. Mitsuda, et al. 1999. "Early Coding of Reaching: Frontal and Parietal Association Connections of Parieto-Occipital Cortex." *The European Journal of Neuroscience* 11 (9): 3339–45.

Chen, Daofen, and Eberhard E. Fetz. 2005. "Characteristic Membrane Potential Trajectories in Primate Sensorimotor Cortex Neurons Recorded In Vivo." *Journal of Neurophysiology* 94 (4): 2713–25. doi:10.1152/jn.00024.2005.

Cheney, P D, and E E Fetz. 1984. "Corticomotoneuronal Cells Contribute to Long-Latency Stretch Reflexes in the Rhesus Monkey." *The Journal of Physiology* 349 (1): 249–72. doi:10.1113/jphysiol.1984.sp015155.

Cheng, Sen, and Philip N. Sabes. 2006. "Modeling Sensorimotor Learning with Linear Dynamical Systems." *Neural Computation* 18 (4): 760–93. doi:10.1162/089976606775774651.

Churchland, Mark M., Byron M. Yu, John P. Cunningham, Leo P. Sugrue, Marlene R. Cohen, Greg S. Corrado, William T. Newsome, et al. 2010. “Stimulus Onset Quenches Neural Variability: A Widespread Cortical Phenomenon.” *Nature Neuroscience* 13 (3): 369–78. doi:10.1038/nn.2501.

Churchland, MM, JP Cunningham, MT Kaufman, JD Foster, P Nuyujukian, SI Ryu, and KV Shenoy. 2012. “Neural Population Dynamics during Reaching.” *Nature* 487 (7405): 51–56. doi:10.1038/nature11129.

Cisek, Paul, and John F. Kalaska. 2005. “Neural Correlates of Reaching Decisions in Dorsal Premotor Cortex: Specification of Multiple Direction Choices and Final Selection of Action.” *Neuron* 45 (5): 801–14. doi:10.1016/j.neuron.2005.01.027.

Cohen, Yale E., and Richard A. Andersen. 2000. “Reaches to Sounds Encoded in an Eye-Centered Reference Frame.” *Neuron* 27 (3): 647–52. doi:10.1016/S0896-6273(00)00073-8.

———. 2002. “A Common Reference Frame for Movement Plans in the Posterior Parietal Cortex.” *Nature Reviews Neuroscience* 3 (7): 553–62. doi:10.1038/nrn873.

Crammond, D. J., and J. F. Kalaska. 1989. “Neuronal Activity in Primate Parietal Cortex Area 5 Varies with Intended Movement Direction during an Instructed-Delay Period.” *Experimental Brain Research* 76 (2): 458–62.

Danion, Frederic, and Mark Latash. 2010. *Motor Control: Theories, Experiments, and Applications*. Oxford University Press.

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

Dickey, Adam S., Aaron J. Suminski, Yali Amit, and Nicholas G. Hatsopoulos. 2009b. “Single-Unit Stability Using Chronically Implanted Multi-Electrode Arrays.” *Journal of Neurophysiology*, June. doi:10.1152/jn.90920.2008.

Diester, Ilka, Matthew T. Kaufman, Murtaza Mogri, Ramin Pashaie, Werapong Goo, Ofer Yizhar, Charu Ramakrishnan, Karl Deisseroth, and Krishna V. Shenoy. 2011. “An Optogenetic Toolbox Designed for Primates.” *Nature Neuroscience* 14 (3): 387–97. doi:10.1038/nn.2749.

DiFranco, D., D. W. Muir, and P. C. Dodwell. 1978. “Reaching in Very Young Infants.” *Perception* 7 (4): 385–92.

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.

Eldawlatly, Seif, Rong Jin, and Karim G. Oweiss. 2008. “Identifying Functional Connectivity in Large-Scale Neural Ensemble Recordings: A Multiscale Data Mining Approach.” *Neural Computation* 21 (2): 450–77. doi:10.1162/neco.2008.09-07-606.

Eleryan, Ahmed, Mukta Vaidya, Joshua Southerland, Islam S. Badreldin, Karthikeyan Balasubramanian, Andrew H. Fagg, Nicholas Hatsopoulos, and Karim Oweiss. 2014. “Tracking Single Units in Chronic, Large Scale, Neural Recordings for Brain Machine Interface Applications.” *Frontiers in Neuroengineering* 7 (July). doi:10.3389/fneng.2014.00023.

Fetz, Eberhard E. 2007. “Volitional Control of Neural Activity: Implications for Brain-Computer Interfaces.” *The Journal of Physiology* 579 (Pt 3): 571–79. doi:10.1113/jphysiol.2006.127142.

Fetz, E. E. 1969. “Operant Conditioning of Cortical Unit Activity.” *Science (New York, N.Y.)* 163 (3870): 955–58.

Fetz, E. E., and M. A. Baker. 1973. “Operantly Conditioned Patterns on Precentral Unit Activity and Correlated Responses in Adjacent Cells and Contralateral Muscles.” *Journal of Neurophysiology* 36 (2): 179–204.

Fraser, George W., and Andrew B. Schwartz. 2012. “Recording from the Same Neurons Chronically in Motor Cortex.” *Journal of Neurophysiology* 107 (7): 1970–78. doi:10.1152/jn.01012.2010.

Fu, Q. G., D. Flament, J. D. Coltz, and T. J. Ebner. 1995. “Temporal Encoding of Movement Kinematics in the Discharge of Primate Primary Motor and Premotor Neurons.” *Journal of Neurophysiology* 73 (2): 836–54.

Galletti, Claudio, Dieter F. Kutz, Michela Gamberini, Rossella Breveglieri, and Patrizia Fattori. 2003. “Role of the Medial Parieto-Occipital Cortex in the Control of Reaching and Grasping Movements.” *Experimental Brain Research* 153 (2): 158–70. doi:10.1007/s00221-003-1589-z.

Galloway, James C., and Gail F. Koshland. 2002. “General Coordination of Shoulder, Elbow and Wrist Dynamics during Multijoint Arm Movements.” *Experimental Brain Research* 142 (2): 163–80. doi:10.1007/s002210100882.

Ganguly, Karunesh, and Jose M. Carmena. 2009. “Emergence of a Stable Cortical Map for Neuroprosthetic Control.” *PLoS Biol* 7 (7): e1000153. doi:10.1371/journal.pbio.1000153.

Ganguly, Karunesh, Dragan F. Dimitrov, Jonathan D. Wallis, and Jose M. Carmena. 2011. “Reversible Large-Scale Modification of Cortical Networks during Neuroprosthetic Control.” *Nature Neuroscience* 14 (5): 662–67. doi:10.1038/nn.2797.

Gentilucci, M., U. Castiello, M. L. Corradini, M. Scarpa, C. Umiltà, and G. Rizzolatti. 1991. “Influence of Different Types of Grasping on the Transport Component of Prehension Movements.” *Neuropsychologia* 29 (5): 361–78.

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.

Georgopoulos, A. P., A. B. Schwartz, and R. E. Kettner. 1986. "Neuronal Population Coding of Movement Direction." *Science* 233 (4771): 1416–19. doi:10.1126/science.3749885.

Goodale, Melvyn A. 1990. *Vision and Action: The Control of Grasping*. Intellect Books.

Gould, H. J., C. G. Cusick, T. P. Pons, and J. H. Kaas. 1986. "The Relationship of Corpus Callosum Connections to Electrical Stimulation Maps of Motor, Supplementary Motor, and the Frontal Eye Fields in Owl Monkeys." *The Journal of Comparative Neurology* 247 (3): 297–325. doi:10.1002/cne.902470303.

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.

Heinzl, Harald, and Martina Mittlböck. 2003. "Pseudo R-Squared Measures for Poisson Regression Models with over- or Underdispersion." *Computational Statistics & Data Analysis*, Special Issue in Honour of Stan Azen: a Birthday Celebration, 44 (1–2): 253–71. doi:10.1016/S0167-9473(03)00062-8.

Hochberg, Leigh R., Daniel Bacher, Beata Jarosiewicz, Nicolas Y. Masse, John D. Simeral, Joern Vogel, Sami Haddadin, et al. 2012. "Reach and Grasp by People with Tetraplegia Using a Neurally Controlled Robotic Arm." *Nature* 485 (7398): 372–75. doi:10.1038/nature11076.

Hofsten, Claes von. 1981. "Predictive Reaching to Moving Objects by Human Infants." *Journal of Experimental Child Psychology* 30 (3): 369–82. doi:10.1016/0022-0965(80)90043-0.

Holdefer, R. N., and L. E. Miller. 2002. "Primary Motor Cortical Neurons Encode Functional Muscle Synergies." *Experimental Brain Research* 146 (2): 233–43. doi:10.1007/s00221-002-1166-x.

Holzbaur, Katherine R. S., Wendy M. Murray, and Scott L. Delp. 2005. "A Model of the Upper Extremity for Simulating Musculoskeletal Surgery and Analyzing Neuromuscular Control." *Annals of Biomedical Engineering* 33 (6): 829–40. doi:10.1007/s10439-005-3320-7.

Huang, C. S., H. Hiraba, and B. J. Sessle. 1989. "Input-Output Relationships of the Primary Face Motor Cortex in the Monkey (Macaca Fascicularis)." *Journal of Neurophysiology* 61 (2): 350–62.

Jackson, Andrew, and Eberhard E. Fetz. 2007. “Compact Movable Microwire Array for Long-Term Chronic Unit Recording in Cerebral Cortex of Primates.” *Journal of Neurophysiology* 98 (5): 3109–18. doi:10.1152/jn.00569.2007.

Jeannerod, M. 1984. “The Timing of Natural Prehension Movements.” *Journal of Motor Behavior* 16 (3): 235–54. doi:10.1080/00222895.1984.10735319.

Jeannerod, M. 1999. “Visuomotor Channels: Their Integration in Goal-Directed Prehension.” *Human Movement Science* 18 (2–3): 201–18. doi:10.1016/S0167-9457(99)00008-1.

Jeannerod, Marc. 1988. *The Neural and Behavioural Organization of Goal-Directed Movements*. Vol. xii. Oxford Psychology Series, No. 15. New York, NY, US: Clarendon Press/Oxford University Press.

Johnson, P. B., and S. Ferraina. 1996. “Cortical Networks for Visual Reaching: Intrinsic Frontal Lobe Connectivity.” *The European Journal of Neuroscience* 8 (7): 1358–62.

Kakei, Shinji, Donna S. Hoffman, and Peter L. Strick. 1999. “Muscle and Movement Representations in the Primary Motor Cortex.” *Science* 285 (5436): 2136–39. doi:10.1126/science.285.5436.2136.

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.

Kelso, J. A. S. 1994. “The Informational Character of Self-Organized Coordination Dynamics.” *Human Movement Science* 13 (3): 393–413. doi:10.1016/0167-9457(94)90047-7.

Krakauer, John, and Claude Ghez. 1991. “Voluntary Movement.” In *Principles of Neural Science*, 4th ed.

Kuhtz-Buschbeck, J. P., A. Boczek-Funcke, M. Illert, K. Joehnk, and H. Stolze. 1999. “Prehension Movements and Motor Development in Children.” *Experimental Brain Research* 128 (1-2): 65–68. doi:10.1007/s002210050818.

Kuhtz-Buschbeck, J. P, H Stolze, A Boczek-Funcke, K Jöhnk, H Heinrichs, and M Illert. 1998. “Kinematic Analysis of Prehension Movements in Children.” *Behavioural Brain Research* 93 (1–2): 131–41. doi:10.1016/S0166-4328(97)00147-2.

Kuhtz-Buschbeck, J. P., H. Stolze, K. Jöhnk, A. Boczek-Funcke, and M. Illert. 1998. “Development of Prehension Movements in Children: A Kinematic Study.” *Experimental Brain Research* 122 (4): 424–32.

Kurata, K. 1993. “Premotor Cortex of Monkeys: Set- and Movement-Related Activity Reflecting Amplitude and Direction of Wrist Movements.” *Journal of Neurophysiology* 69 (1): 187–200.

Landsmeer, J. M. F. 1962. “Power Grip and Precision Handling.” *Annals of the Rheumatic Diseases* 21 (2): 164–70.

Lawrence, D. G., and H. G. Kuypers. 1968. “The Functional Organization of the Motor System in the Monkey. II. The Effects of Lesions of the Descending Brain-Stem Pathways.” *Brain: A Journal of Neurology* 91 (1): 15–36.

Lawrence, Donald G., and Henricus G. J. M. Kuypers. 1968. “The Functional Organization of the Motor System in the Monkey.” *Brain* 91 (1): 1–14. doi:10.1093/brain/91.1.1.

Liu, Xu, Steve Ramirez, Petti T. Pang, Corey B. Puryear, Arvind Govindarajan, Karl Deisseroth, and Susumu Tonegawa. 2012. “Optogenetic Stimulation of a Hippocampal Engram Activates Fear Memory Recall.” *Nature* 484 (7394): 381–85. doi:10.1038/nature11028.

Li, Zheng, Joseph E. O’Doherty, Mikhail A. Lebedev, and Miguel A. L. Nicolelis. 2011. “Adaptive Decoding for Brain-Machine Interfaces Through Bayesian Parameter Updates.” *Neural Computation* 23 (12): 3162–3204. doi:10.1162/NECO_a_00207.

Luppino, G., A. Murata, P. Govoni, and M. Matelli. 1999. “Largely Segregated Parietofrontal Connections Linking Rostral Intraparietal Cortex (areas AIP and VIP) and the Ventral Premotor Cortex (areas F5 and F4).” *Experimental Brain Research* 128 (1-2): 181–87.

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, Mary W. 1997. “Precision Grips, Hand Morphology, and Tools.” *American Journal of Physical Anthropology* 102 (1): 91–110. doi:10.1002/(SICI)1096-8644(199701)102:1<91::AID-AJPA8>3.0.CO;2-G.

Matelli, M., R. Camarda, M. Glickstein, and G. Rizzolatti. 1986. “Afferent and Efferent Projections of the Inferior Area 6 in the Macaque Monkey.” *The Journal of Comparative Neurology* 251 (3): 281–98. doi:10.1002/cne.902510302.

McKiernan, Brian J., Joanne K. Marcario, Jennifer Hill Karrer, and Paul D. Cheney. 1998. “Corticomotoneuronal Postspike Effects in Shoulder, Elbow, Wrist, Digit, and Intrinsic Hand Muscles During a Reach and Prehension Task.” *Journal of Neurophysiology* 80 (4): 1961–80.

Messier, J., and J. F. Kalaska. 2000. “Covariation of Primate Dorsal Premotor Cell Activity with Direction and Amplitude during a Memorized-Delay Reaching Task.” *Journal of Neurophysiology* 84 (1): 152–65.

Middleton, Frank A, and Peter L Strick. 2000. “Basal Ganglia and Cerebellar Loops: Motor and Cognitive Circuits.” *Brain Research Reviews* 31 (2–3): 236–50. doi:10.1016/S0165-0173(99)00040-5.

Murata, A., L. Fadiga, L. Fogassi, V. Gallese, V. Raos, and G. Rizzolatti. 1997. “Object Representation in the Ventral Premotor Cortex (area F5) of the Monkey.” *Journal of Neurophysiology* 78 (4): 2226–30.

Murata, Akira, Vittorio Gallese, Giuseppe Luppino, Masakazu Kaseda, and Hideo Sakata. 2000. “Selectivity for the Shape, Size, and Orientation of Objects for Grasping in Neurons of Monkey Parietal Area AIP.” *Journal of Neurophysiology* 83 (5): 2580–2601.

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.

Mussa-Ivaldi, Ferdinando A., and Lee E. Miller. 2003. “Brain–machine Interfaces: Computational Demands and Clinical Needs Meet Basic Neuroscience.” *Trends in Neurosciences* 26 (6): 329–34. doi:10.1016/S0166-2236(03)00121-8.

Nicolelis, Miguel A. L., Dragan Dimitrov, Jose M. Carmena, Roy Crist, Gary Lehew, Jerald D. Kralik, and Steven P. Wise. 2003. “Chronic, Multisite, Multielectrode Recordings in Macaque Monkeys.” *Proceedings of the National Academy of Sciences* 100 (19): 11041–46. doi:10.1073/pnas.1934665100.

Nudo, R. J., W. M. Jenkins, M. M. Merzenich, T. Prejean, and R. Grenda. 1992. “Neurophysiological Correlates of Hand Preference in Primary Motor Cortex of Adult Squirrel Monkeys.” *The Journal of Neuroscience: The Official Journal of the Society for Neuroscience* 12 (8): 2918–47.

Orsborn, A.L., S. Dangi, H.G. Moorman, and J.M. Carmena. 2012. “Closed-Loop Decoder Adaptation on Intermediate Time-Scales Facilitates Rapid BMI Performance Improvements Independent of Decoder Initialization Conditions.” *IEEE Transactions on Neural Systems and Rehabilitation Engineering* 20 (4): 468–77. doi:10.1109/TNSRE.2012.2185066.

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.

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.

Qi, H. X., I. Stepniewska, and J. H. Kaas. 2000. “Reorganization of Primary Motor Cortex in Adult Macaque Monkeys with Long-Standing Amputations.” *Journal of Neurophysiology* 84 (4): 2133–47.

Rathelot, Jean-Alban, and Peter L. Strick. 2009. “Subdivisions of Primary Motor Cortex Based on Cortico-Motoneuronal Cells.” *Proceedings of the National Academy of Sciences* 106 (3): 918–23. doi:10.1073/pnas.0808362106.

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* 71 (3): 491–507.

Rousche, Patrick J, and Richard A Normann. 1998. “Chronic Recording Capability of the Utah Intracortical Electrode Array in Cat Sensory Cortex.” *Journal of Neuroscience Methods* 82 (1): 1–15. doi:10.1016/S0165-0270(98)00031-4.

Sadovsky, Alexander J., Peter B. Kruskal, Joseph M. Kimmel, Jared Ostmeyer, Florian B. Neubauer, and Jason N. MacLean. 2011. “Heuristically Optimal Path Scanning for High-Speed Multiphoton Circuit Imaging.” *Journal of Neurophysiology* 106 (3): 1591–98. doi:10.1152/jn.00334.2011.

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.

Sanes, Jerome N., and John P. Donoghue. 2000. “Plasticity and Primary Motor Cortex.” *Annual Review of Neuroscience* 23 (1): 393–415. doi:10.1146/annurev.neuro.23.1.393.

Sanes, J. N., S. Suner, and J. P. Donoghue. 1990. “Dynamic Organization of Primary Motor Cortex Output to Target Muscles in Adult Rats. I. Long-Term Patterns of Reorganization Following Motor or Mixed Peripheral Nerve Lesions.” *Experimental Brain Research* 79 (3): 479–91.

Sanes, J. N., S. Suner, J. F. Lando, and J. P. Donoghue. 1988. “Rapid Reorganization of Adult Rat Motor Cortex Somatic Representation Patterns after Motor Nerve Injury.” *Proceedings of the National Academy of Sciences of the United States of America* 85 (6): 2003–7.

Santhanam, G., M. Sahani, S.I. Ryu, and K.V. Shenoy. 2004. “An Extensible Infrastructure for Fully Automated Spike Sorting during Online Experiments.” In *26th Annual International Conference of the IEEE Engineering in Medicine and Biology Society, 2004. IEMBS ’04*, 2:4380–84. doi:10.1109/IEMBS.2004.1404219.

Schieber, M. H., and R. K. Deuel. 1997. “Primary Motor Cortex Reorganization in a Long-Term Monkey Amputee.” *Somatosensory & Motor Research* 14 (3): 157–67.

Schneiberg, Sheila, Heidi Sveistrup, Bradford McFadyen, Patricia McKinley, and Mindy F. Levin. 2002. “The Development of Coordination for Reach-to-Grasp Movements in Children.” *Experimental Brain Research* 146 (2): 142–54. doi:10.1007/s00221-002-1156-z.

Scott, Stephen H. 2003. “The Role of Primary Motor Cortex in Goal-Directed Movements: Insights from Neurophysiological Studies on Non-Human Primates.” *Current Opinion in Neurobiology* 13 (6): 671–77. doi:10.1016/j.conb.2003.10.012.

Scott, Stephen H. 2004. “Optimal Feedback Control and the Neural Basis of Volitional Motor Control.” *Nature Reviews Neuroscience* 5 (7): 532–46. doi:10.1038/nrn1427.

Sessle, B. J., and M. Wiesendanger. 1982. “Structural and Functional Definition of the Motor Cortex in the Monkey (Macaca Fascicularis).” *The Journal of Physiology* 323 (February): 245–65.

Seth, Anil K. 2010. “A MATLAB Toolbox for Granger Causal Connectivity Analysis.” *Journal of Neuroscience Methods* 186 (2): 262–73. doi:10.1016/j.jneumeth.2009.11.020.

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.

Shipp, S., M. Blanton, and S. Zeki. 1998. “A Visuo-Somatomotor Pathway through Superior Parietal Cortex in the Macaque Monkey: Cortical Connections of Areas V6 and V6A.” *The European Journal of Neuroscience* 10 (10): 3171–93.

Smeets, J. B., and E. Brenner. 1999. “A New View on Grasping.” *Motor Control* 3 (3): 237–71.

Stark, Eran, Amir Globerson, Itay Asher, and Moshe Abeles. 2008. “Correlations between Groups of Premotor Neurons Carry Information about Prehension.” *The Journal of Neuroscience* 28 (42): 10618–30. doi:10.1523/JNEUROSCI.3418-08.2008.

Sternad, D., M. T. Turvey, and R. C. Schmidt. 1992. “Average Phase Difference Theory and 1:1 Phase Entrainment in Interlimb Coordination.” *Biological Cybernetics* 67 (3): 223–31.

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.

Taylor, Dawn M., Stephen I. Helms Tillery, and Andrew B. Schwartz. 2002. “Direct Cortical Control of 3D Neuroprosthetic Devices.” *Science* 296 (5574): 1829–32. doi:10.1126/science.1070291.

Todorov, Emanuel. 2004. “Optimality Principles in Sensorimotor Control.” *Nature Neuroscience* 7 (9): 907–15. doi:10.1038/nn1309.

Todorov, Emanuel, and Michael I. Jordan. 2002. “Optimal Feedback Control as a Theory of Motor Coordination.” *Nature Neuroscience* 5 (11): 1226–35. doi:10.1038/nn963.

Tokuno, Hironobu, and Jun Tanji. 1993. “Input Organization of Distal and Proximal Forelimb Areas in the Monkey Primary Motor Cortex: A Retrograde Double Labeling Study.” *The Journal of Comparative Neurology* 333 (2): 199–209. doi:10.1002/cne.903330206.

Tolias, Andreas S., Alexander S. Ecker, Athanassios G. Siapas, Andreas Hoenselaar, Georgios A. Keliris, and Nikos K. Logothetis. 2007. “Recording Chronically From the Same Neurons in Awake, Behaving Primates.” *Journal of Neurophysiology* 98 (6): 3780–90. doi:10.1152/jn.00260.2007.

Vargas-Irwin, Carlos, and John P. Donoghue. 2007. “Automated Spike Sorting Using Density Grid Contour Clustering and Subtractive Waveform Decomposition.” *Journal of Neuroscience Methods* 164 (1): 1–18. doi:10.1016/j.jneumeth.2007.03.025.

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.

Von Hofsten, Claes. 1983. “Catching Skills in Infancy.” *Journal of Experimental Psychology: Human Perception and Performance* 9 (1): 75–85. doi:10.1037/0096-1523.9.1.75.

———. 1984. “Developmental Changes in the Organization of Prereaching Movements.” *Developmental Psychology* 20 (3): 378–88.

Von Hofsten, Claes, and Shirin Fazel-Zandy. 1984. “Development of Visually Guided Hand Orientation in Reaching.” *Journal of Experimental Child Psychology* 38 (2): 208–19. doi:10.1016/0022-0965(84)90122-X.

Von Hofsten, C., and K. Lindhagen. 1979. “Observations on the Development of Reaching for Moving Objects.” *Journal of Experimental Child Psychology* 28 (1): 158–73.

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

Werner, W., E. Bauswein, and C. Fromm. 1991. “Static Firing Rates of Premotor and Primary Motor Cortical Neurons Associated with Torque and Joint Position.” *Experimental Brain Research* 86 (2): 293–302. doi:10.1007/BF00228952.

Willett, Francis R., Aaron J. Suminski, Andrew H. Fagg, and Nicholas G. Hatsopoulos. 2013. “Improving Brain–machine Interface Performance by Decoding Intended Future Movements.” *Journal of Neural Engineering* 10 (2): 026011. doi:10.1088/1741-2560/10/2/026011.

Wimmers, R. H., G. J. Savelsbergh, P. J. Beek, and B. Hopkins. 1998. “Evidence for a Phase Transition in the Early Development of Prehension.” *Developmental Psychobiology* 32 (3): 235–48.

Wise, S. P., and K.-H. Mauritz. 1985. “Set-Related Neuronal Activity in the Premotor Cortex of Rhesus Monkeys: Effects of Changes in Motor Set.” *Proceedings of the Royal Society of London. Series B, Biological Sciences* 223 (1232): 331–54.

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

Witham, Claire L., C. Nicholas Riddle, Mark R. Baker, and Stuart N. Baker. 2011. "Contributions of Descending and Ascending Pathways to Corticomuscular Coherence in Humans." *The Journal of Physiology* 589 (15): 3789–3800. doi:10.1113/jphysiol.2011.211045.

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, C. W., and J. H. Kaas. 1999. "Reorganization in Primary Motor Cortex of Primates with Long-Standing Therapeutic Amputations." *The Journal of Neuroscience: The Official Journal of the Society for Neuroscience* 19 (17): 7679–97.

Ziegler-Graham, Kathryn, Ellen J. MacKenzie, Patti L. Ephraim, Thomas G. Travison, and Ron Brookmeyer. 2008. "Estimating the Prevalence of Limb Loss in the United States: 2005 to 2050." *Archives of Physical Medicine and Rehabilitation* 89 (3): 422–29. doi:10.1016/j.apmr.2007.11.005.

Zoia, Stefania, Eva Pezzetta, Laura Blason, Aldo Scabar, Marco Carrozzi, Maria Bulgheroni, and Umberto Castiello. 2006. "A Comparison of the Reach-to-Grasp Movement between Children and Adults: A Kinematic Study." *Developmental Neuropsychology* 30 (2): 719–38. doi:10.1207/s15326942dn3002_4.

APPENDIX A: ULTRA-LONG TERM STABILITY OF SINGLE UNITS USING CHRONICALLY IMPLANTED MULTIELECTRODE ARRAYS

Abstract

Recordings from chronically implanted multielectrode arrays have become prevalent in both neuroscience and neural engineering experiments. To date, however, the extent to which populations of single-units remain stable over long periods of time has not been well characterized. In this study, neural activity was recorded from a Utah multielectrode array implanted in the primary motor cortex of a rhesus macaque during 18 recording sessions spanning nine months. We found that 67% of the units were stable through the first 15 days, 31% of units were stable through 47 days, 21% of units were stable through 106 days, and 8% of units were stable over 9 months. Thus not only were units stable over a timescale of several months, but units stable over 2 months were more likely to remain stable in the next 2 months.

Introduction

Chronically implanted multielectrode arrays are commonly used in neurophysiological experiments because of their ability to record from a large number of units over a long period of time. However, it is useful to know whether a unit recorded on the same electrode on different days might actually be originating from the same neuron. Initial attempts to track neurons involved visually inspecting the waveform shape across days (Rousche and Normann 1998), by computing the correlation between average waveforms (Jackson and Fetz 2007), or by comparing the clustering in principal component space (Nicolelis et al. 2003). However, relying

purely on the waveform shape along might lead to false positives, since different neurons may have a similar average waveform. Subsequent attempts at tracking neurons used addition information besides the waveform, such as the 3-D location relative to a tetrode (Tolias et al. 2007), the inter-spike interval histogram (Dickey et al. 2009a) or the correlation of neuronal firing with other stable neurons (Fraser and Schwartz 2012). Fraser and Schwartz reported tracking stable units in rhesus macaque motor cortex for over 100 days (Fraser and Schwartz 2012).

Here, we use the method described in Dickey et al. to track neurons over a series of datasets recorded over a longer time scale of 265 days. The recordings occurred during an experiment in which a naïve monkey was introduced to brain-machine interface (BMI) control of a robotic arm and hand. Thus the identification of stable units over these datasets would allow the examination of learning effects during the initial exposure of BMI real-time control.

Materials & Methods

Neural Recordings

All surgical and behavioral procedures involved in this study were approved by the University of Chicago Institutional Animal Care and Use Committee and conform to the principles outlines in the Guide for the Care and Use of Laboratory Animals.

Data used for this analysis were collected from a female rhesus macaque (*Macaca mulatta*) monkey that was implanted with a Utah 100-microelectrode array (Blackrock Microsystems, Salt Lake City, UT) in primary motor (MI) cortex. The macaque had been the recipient of a

therapeutic amputation 5 years prior due to injury. The array used for this analysis was placed contralateral to the intact limb. The electrodes on the array were 1 mm in length. During a recording session, spike waveforms from up to 96 electrodes were amplified (gain of 5,000), filtered between 0.3Hz and 7.5 kHz, and recorded digitally (14-bit) at 30kHz per channel using a Cerebus acquisition system (Blackrock Microsystems). On the first experimental session, units were sorted online with a hoop-sorting algorithm, described in Santhanam et al. (Santhanam et al. 2004). Potential spikes were first identified when the filtered voltage dropped below a user-defined threshold. These spikes were sorted by placing a lower and upper voltage threshold (the "hoop") at specific times relative to the initial threshold crossing. The same sorting rules were applied to the remainder of recording sessions, so that the number of sorted units remained constant over time. The recordings for this analysis were collected on 18 separate daily sessions over the course of nine months. The first session was recorded 11 months after implantation of the array.

Behavioral Task

The macaque performed the same behavioral task on all of the recording sessions. In this task, the macaque had to learn how to navigate two control dimensions of a robotic arm in order to perform a reach-to grasp task. The robot was composed of a 7 DOF WAM arm attached to a 4 DOF BarrettHand (Barret Technology, Inc.). Through the use of operant conditioning, the macaque learned to control the reaching motion towards and away from the base of the robot, as well as the grasping motion, opening or closing all three digits of the hand concurrently. A successful trial would involve reaching-to-grasp a sphere placed on a board in front of the robot, pulling it back, and finally dropping it.

Using neurons identified using online spike sorting, distinct clusters of functionally connected groups of neurons were created for the purpose of controlling each control dimension. Decoders using the neural population activity of either group, binned at 50 milliseconds, were initialized in an unsupervised manner using spontaneous data (Badreldin et al. 2013). Over the course of the study, the macaque learned this mapping such that she could coordinate movement along both dimensions in order to perform successful trials at an increasing rate.

Stability Analysis

The methodology outlined by Dickey et al. (Dickey et al. 2009a) was used to analyze the stability of single-units on the array over the course of nine months. For each sorted unit, we computed the average waveform and interspike interval histogram (ISIH). . Waveforms on two different datasets were compared by computing the Pearson's correlation coefficient to create a waveform score. The ISIHs were fit with a mixture of the three log-normal distributions, and the parameters of that fit were compared to create an ISIH score. The waveform and ISIH scores were then combined into a single score. A unit was classified as stable on a given day if this combined score fell below a fixed threshold.

Results

Using stability criteria outlined earlier, units were classified as either stable or unstable with respect to the first recording session. For a unit to be considered stable over the course of the entire study, it was necessary for the unit to be considered stable in every one of the recording sessions. Similarly, for a unit to be considered stable through a particular day, it must have been stable on all of the previous recording sessions prior up to that day. The survival curve in Figure

A.1 shows the percentage of stable units through a particular day for all of the recording sessions.

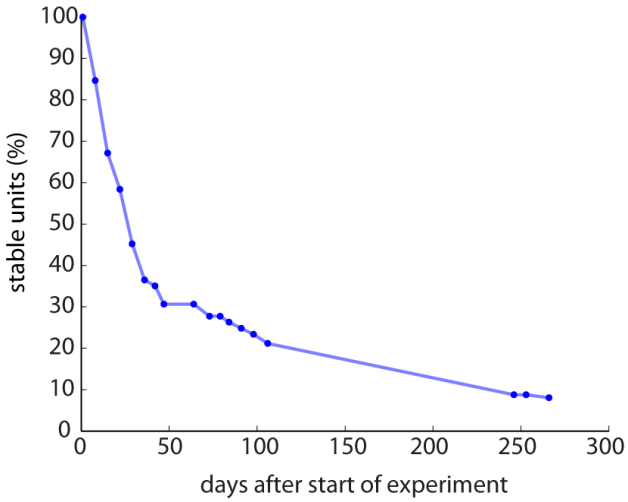

Figure A.1 Survival curve of tracked units The Y axis shows the percentage of stable units, relative to the 137 neurons sorted on day 0.

Figure A.1 illustrates that around 8% of neurons are stable over the entire nine months of the study. After 106 days have elapsed, over a fifth of the units are still stable. An inflection point can be observed around 50 days. Around 31% of the units are stable through day 47, while around 23% of the units are stable through day 98. Thus 75% of the units that are stable on day 47 survive the next 51 days. In contrast if we look at earlier recording sessions, about 58% of units are stable through day 22, while about 31% of units are stable through day 47. So only 53% of the units that are stable on day 22 survive the next 25 days.

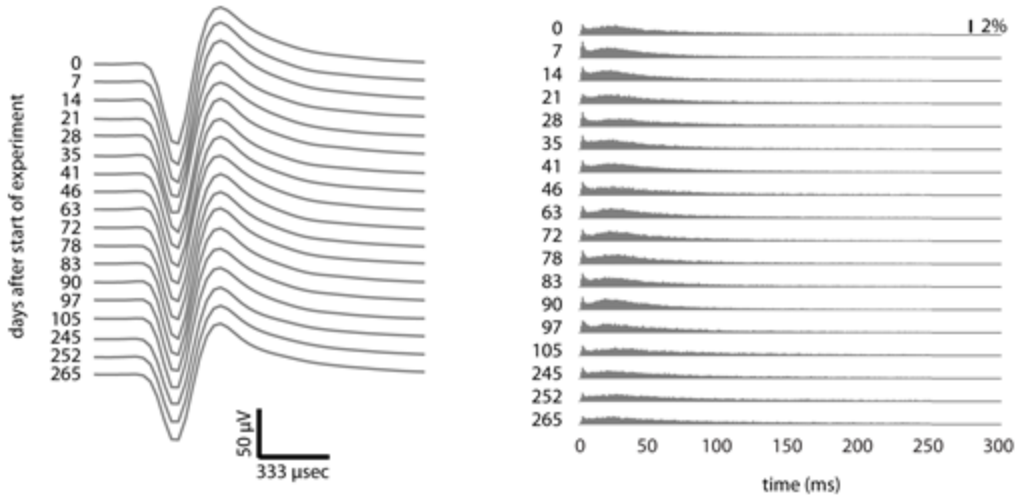

Figure A.2 Example Stable Unit (Left) Waveforms and (Right) Interspike interval histograms for a stable unit over 265 days

An example of a stable unit is provided in Figure A.2. The similarity in waveforms across all eighteen of the recording sessions over nine of months of recording is quite apparent. The same is true for the interspike interval histograms; the characteristic bimodal distribution (a Type II neuron as described in Chen & Fetz (Chen and Fetz 2005)) can be seen in each session. In contrast, Figure A.3 provides a clear example of a unit that changes the properties of both the waveform and the ISIH between the 15th and 16th dataset. This is when there was a 5 month gap in recording, and it appears that the first neuron was replaced with a completely different neuron. This unit was classified as stable through the 15th dataset on day 106, but not thereafter. Other unstable units change properties in both the waveform and the ISIH multiple times; in one example, the first change occurs between the first data of recording and seven days later, the second is, again, between the 15th and 16th dataset. Table A.1 provides the scores for the stability of each unit shown (Figure A.2, Figure A.3) when compared with the first day of recording. The

threshold for stability was below a score of 11.67 as defined by Dickey et al. (Dickey et al. 2009a).

Table A.1 Unit Stability Scores

# Days	Stable Unit (Fig.2)	Unstable Unit (Fig.3L)
7	-6.57	-6.03
14	-10.88	-12.49
21	-21.19	-3.64
28	-4.42	-11.05
35	-3.26	7.56
41	-6.20	-0.59
46	-6.60	-11.45
63	-0.97	0.43
72	-2.96	1.14
78	-3.35	2.79
83	-3.13	0.37
90	-2.75	1.33
97	-2.19	-13.16
105	-2.15	-6.86
245	-5.07	17.10
252	-3.21	21.80
265	-4.23	19.38

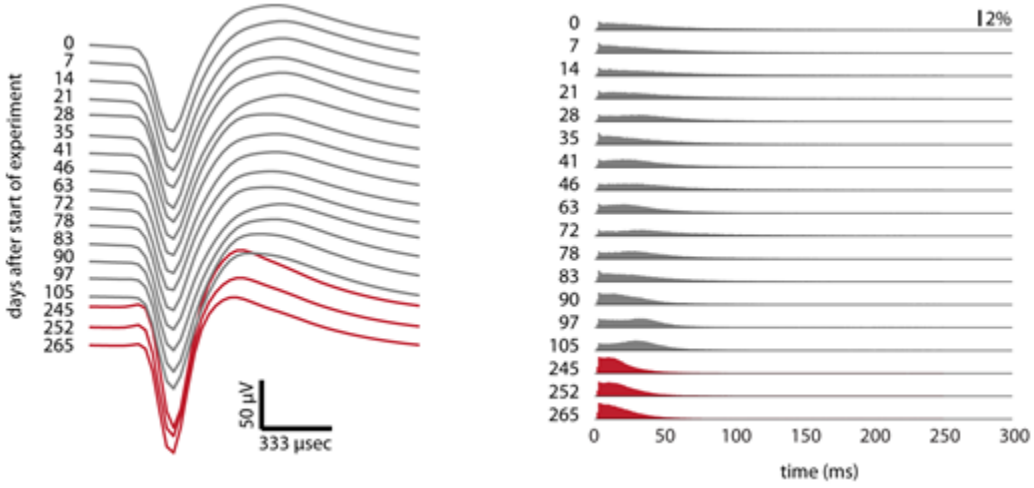

Figure A.3 Example Unstable Unit (Left) Waveforms and (Right) Interspike interval histograms for an unstable unit over the course of the study. Note the change in both waveform and ISIH between day 104 and 245

Since the stability criterion relies on scores from based on the similarity of both the waveforms and the interspike interval histograms, units can fail to be classified as stable if they exhibit similarity in one feature but not the other. For example, a can show stability in its interspike interval histograms throughout the recording sessions, but can fail to maintain a consistent waveform.

Discussion

The most telling result of this study is that we were able to identify a subset of neurons that were stable over the course of nine months. The stability criterion that we adopted is especially conservative, so it is likely that we are systematically underestimating unit stability. It could be that 8% is a more conservative estimate due to the fact that we saw units drop out for a couple of recording sessions, only to return for the remainder of the experiment. Anecdotally, it seems that an electrode can become noisy for a single recording session, perhaps because of mechanical disruption of the contact between the Utah array and the connector, or between the connector and the headstage. This would lead to a unit as being classified as unstable, even if the noise were to disappear the next day and the unit to return. One way to address this in future work is to allow a unit to drop out for a handful of sessions, as long as the unit reappears and is still classified as stable afterwards. This will require careful consideration of the false positive rate, given the multiple comparisons across units and days.

Recording every day instead of every week would give us a survival curve to a greater degree of precision. This presents the similar issue of determining how many datasets a neuron can “miss”

and still be classified as stable. Additionally, it was not possible to have the same elapsed time between datasets. We might expect of the probability of a stable unit disappearing during an interval to increase as a function of recording sessions in the interval—i.e. the true negative rate within an interval will depend on how often you are sampling within that interval. In this collection of datasets, five months elapsed between the 15^{th} and 16^{th} recording sessions as opposed to the otherwise weekly or biweekly recording sessions. Around 21% of units were stable through the 15^{th} recording session, whereas only about 8% were stable through the 16^{th} recording session.

This analysis was also completely dependent on the online spike sorting from the first recording session In the future, we intend to use the online spike sorting from an earlier or later dataset so that we can measure how sensitive the survival curve is to the online sorting on a particular day. In addition, neural data from the first recording session can be sorted offline; this sort could then to be applied to all of the datasets in lieu of the online sort to measure how sensitive the survival curve is to online vs offline spike sorting.

On a similar note, instead of measuring how sensitive the survival curve is to spike sorting, we could look at sensitivity to the stability threshold that was fixed as a value based on a previous stability study (Dickey et al. 2009a). In fact, rather than discretely classifying the units as a "stable" or "not stable" for each day, we could use the continuous stability score to determine the exact day a unit switched from stable to unstable.

Our motivation for applying the stability analysis to this particular collection of datasets was to allow us to investigate the effects of learning to control a BMI on stable neurons. In the learning study described here, the animal was able to not only learn, but improve her ability to coordinate and control different dimensions of a robotic arm in a reach-to-grasp task despite what would appear to be instabilities in her neuronal units. She was able to modify and maintain function of these clusters of neurons, despite changing properties of individual neurons. These experiments can speak to how populations of neurons in primary motor cortex can impact overall behavioral function in a robust fashion, in the presence of dynamic changes at the single-unit level – the type of plasticity that has long been implicated in motor learning experiments (Jerome N. Sanes and Donoghue 2000).

In addition, this kind of tracking could be of broader use to the science and engineering communities at large. The ability to pool neurons across different recording sessions for data analysis purposes, or to identify stable units that can be consistently used for brain machine interfaces without having to retrain decoders could be beneficial. Consequently, development of an algorithm that could analyze the stability of units online, taking into account their known properties and history, could prove very useful in a BMI-context (Badreldin et al. 2013).

Acknowledgements
The authors would like to thank J. Coles for assistance with the training and care of laboratory animals.

Grants
Research supported by DARPA Grant No. N66 001-12-1-4023, IGERT: Integrative Research in Motor Control and Movement. National Science Foundation Grant # DGE-0903637