NORTHWESTERN UNIVERSITY

The Contribution of Whole-Limb Kinematics to Proprioceptive Representations in the Central Nervous System

A DISSERTATION

SUBMITTED TO THE GRADUATE SCHOOL
IN PARTIAL FULFILLMENT OF THE REQUIREMENTS

for the degree

DOCTOR OF PHILOSOPHY

Field of Biomedical Engineering

By
Raeed Hasan Chowdhury

EVANSTON, ILLINOIS

September 2019

Abstract

We have a remarkable ability to perform complex, coordinated movements without much conscious effort. In addition to the computations required to generate commands for muscles, a key aspect of coordinated motor control is incorporating sensory feedback about the movement. One of the most important feedback routes is through proprioception, the sense of body position, movement, and related forces. This sense is so critical that loss of proprioception typically leaves a person wheelchair-bound, despite the retained ability to activate muscles.

The sensory organs for proprioception lie in the muscles, sensing their length and the forces on them. However, classic studies of proprioceptive neurons in primary somatosensory cortex (S1) suggest that activity relates to hand movement, implying that on the way to S1, proprioceptive signals are transformed from representing kinematics in terms of muscle to one in terms of the hand. Surprisingly, another classic study in the dorsal spinocerebellar tract (DSCT) suggested that this transformation may take place as early as the spinal cord. However, these classic studies did not consider how musculoskeletal geometry might contribute to the proprioceptive activity in DSCT and S1.

This dissertation outlines my work in examining how proprioceptive information is processed as it travels from muscle receptors to S1. Chapter 2 details a simulation study of DSCT, in which I found that an apparent limb-endpoint-based representation can arise simply from a convergence of muscle inputs. Chapter 3 extends these results to an electrophysiological study in monkeys, where I found that neurons in S1 represent not just the kinematics of the hand, but the whole arm. Taken together, these studies suggest that if proprioceptive signals transform to a hand-based representation in the brain, this transformation likely occurs beyond S1.

Acknowledgements

Throughout my time at Northwestern and in the Miller lab, one thing has stood out to me above all else: this is a good environment for science. Like most graduate students, I've had my share of setbacks and disappointments with my research, but there has always been at least a handful of people around to pick me up and push me forward. For this I must thank current and former members of the Miller lab, many of whom have become close friends, and my dissertation committee, which includes Dr. Eric Perreault, Dr. Wendy Murray, Dr. Sliman Bensmaia, and (briefly) Dr. Konrad Kording.

Special thanks go to Dr. Lee Miller, my PhD advisor. I feel strongly that the supportive environment I encountered when I joined his lab follows directly from the example he sets. It seems to me that very few advisors care about mentoring their trainees as much as Lee, and as I embark on a career in academia, it is this mentorship and support that I aspire to provide my future trainees, should I be lucky enough to have any.

Of course, none of this work would have been possible without the loving support from my family and friends, who encouraged me both before and during my PhD. This goes double for my wife, Nameera, who has been in my corner through this whole process. She encouraged me through both the ups and downs of this research, with a compassionate understanding that I can only hope to reciprocate as we move forward with our life together.

Table of Contents

Table of Figures

Chapter 1 - Introduction

Aristotle believed there were five senses: sight, hearing, smell, taste, and touch. However, absent from the list is the sense of proprioception, or the sense of body state. Despite our being largely unaware of this sense in our normal lives, proprioception is critical for making coordinated movements. This importance is easiest to see in cases where proprioception is lost. As depicted in the BBC documentary "The Man Who Lost His Body", patients who have lost their sense of proprioception show dramatic loss of body control, and in the great majority of cases become wheelchair-bound, even when they retain the ability to activate their muscles.

While the field has made great strides in identifying where proprioception comes from, we are still relatively ignorant of how proprioceptive signals are processed and represented by the central nervous system (CNS), especially when compared with other senses like vision and touch. In my doctoral work, I examined how neurons in the proprioceptive system represent limb movements, starting with a computational study of a proprioceptive spinal cord tract called the dorsal spinocerebellar tract (DSCT), described in Chapter 2. I then used insights gained from this computational study in Chapter 3, which details an electrophysiological study of primary somatosensory cortex (S1), one of the main brain regions involved with proprioception. This introduction will review background information relevant to my doctoral work, first with an overview of the anatomy and physiology of limb proprioception in the nervous system, followed by an overview of how neural representations of movement have been addressed in the past. Finally, I will summarize ideas suggesting how proprioceptive feedback could be used to generate movements and how these ideas may shed light on proprioceptive neural activity.

Anatomy and physiology of proprioception in the nervous system

Physiology of sensors

Unlike with vision or hearing, the origins of the neural signals that make up the sense of proprioception are not immediately obvious. In fact, there was some debate in the 19th century about whether the sense of proprioception arose mainly from motor commands, originating in the brain or from peripheral receptors. Since the time of this debate, however, many different sensors contributing to proprioception have been discovered in vertebrates, including receptors embedded in the skin, joints, tendons and muscle bellies. While motor commands certainly play a role in proprioception, we now know that, at least for limbs, the sense of proprioception originates mainly from these receptors, with those embedded in the muscles and tendons playing the most important role (see (Proske & Gandevia, 2012) for review of this historical debate).

One group of muscle receptors, called muscle spindles, lie parallel to the fibers of the muscle belly, and are classically thought to sense the length of the muscle and how it changes (Houk, et al., 1981). Each muscle spindle consists of several types of fibers, each of which has a different response to muscle stretch (Proske, 1997). Afferent nerves innervating these fibers transduce movements into action potentials, which propagate to the spinal cord. This transduced neural activity responds strongly to vibratory stimuli applied to the muscle (Brown, et al., 1967), indicating that muscle spindles are highly sensitive to rapid changes in muscle fiber length. Incidentally, such vibration also creates a perceptual illusion of movement (Goodwin, et al., 1972; Cordo, et al., 1995), adding to the evidence that muscle sensors are the primary drivers of proprioceptive sensation.

In the studies presented in Chapters 2 and 3, I used rather simple models of muscle spindle activity, in which muscle length and stretch velocity were represented linearly in discharge rate. However, muscle spindle activity has been shown to have a nonlinear relationship to muscle length and stretch velocity (Houk, et al., 1981), and it even appears to be dependent on the history of muscle stretching (Proske & Stuart, 1985; Haftel, et al., 2004). A further complication to the relationship between spindle activity and muscle state is the fact that, like the muscles in which they reside, muscle spindles are also innervated by motor neurons, called gamma motor neurons. Gamma motor drive causes spindle fibers to contract, thereby increasing their sensitivity to movement (Kuffler, et al., 1951; Prochazka, 1981; Hulliger, et al., 1989; Hulliger, 1993; Macefield & Knellwolf, 2018). One hypothesized function for gamma drive is to keep the spindle in its operating range as the muscle contracts or relaxes. This might be achieved through a linkage between gamma and alpha motor neuronal activity, the latter providing the primary drive for the spindle's parent muscle (Hagbarth & Vallbo, 1968; Burke, et al., 1978). However, many experiments have shown that alpha and gamma drive can be dissociated during voluntary movement, revealing that the role of gamma drive is more complex than this hypothesis would suggest (Prochazka, et al., 1976; Prochazka, 1981; Hulliger, et al., 1989).

A second class of muscle receptors, known as Golgi tendon organs (GTOs), also appears to play a large role in proprioception. Because they reside in the tendons, and are thus in series with the muscle fibers, GTOs are thought to signal the force exerted by the muscles, and likely contribute to a sense of effort (Houk & Simon, 1967; Crago, et al., 1982). These sensors also respond to vibration, particularly when the connected muscle is active (Brown, et al., 1967; Fallon & Macefield, 2007), making vibration a potentially poor stimulus to determine

experimentally whether an axon carries information from a muscle spindle or from a GTO. One stimulus that that does distinguish these two sensors, however, is a brief electrical stimulation of the muscle (Hunt & Kuffler, 1951; Edin & Vallbo, 1987). This stimulation causes the muscle to twitch, simultaneously shortening and silencing the muscle spindles, and applying tension to the tendon, thereby activating the GTOs. This twitch test has been called the "gold standard" for distinguishing muscle spindle axons from those of GTOs (Fallon & Macefield, 2007).

Overview of pathways to cortex

As noted above, proprioception, like tactile sensation, originates from sensors all over the body. Signals from muscle and cutaneous receptors arrive at the spinal cord via structures known as dorsal root ganglia (DRG). These ganglia include bundles of neurons that innervate muscle spindles and GTOs, as well as sensors in the skin and joints. From there, DRG signals enter the dorsal horn of the spinal cord and split into two paths, one intended for fast, spinal reflex loops and the other to carry signals to the brain. Signals used for reflexes are relayed to spinal interneurons and motor neurons, while signals intended for the brain mostly travel up the dorsal column tract. For the arms or forelimbs, these DRG axons synapse onto neurons in the cuneate nucleus of the medulla (Rosén, 1969; Rosén, 1969). From there, the cuneate nucleus sends signals through the medial lemniscus to the thalamus, which routes the signals to primary somatosensory cortex (S1) (Rosén, 1969; Oscarsson & Rosén, 1963; Padberg, et al., 2009). Cutaneous signals from the legs or hindlimbs follow a similar pathway through the gracile nucleus rather than the cuneate (Gordon & Paine, 1960), but muscle signals from these lower limbs take a detour. Instead of traveling up the dorsal tract directly to a brainstem nucleus, the hindlimb muscle DRG neurons synapse onto neurons in Clarke's column, which then sends

signals through the dorsal spinocerebellar tract (DSCT) (Lloyd & McIntyre, 1950) (see (Bosco & Poppele, 2001) for review of DSCT). As the name suggests, this tract sends proprioceptive signals to the cerebellum, but it also sends signals to another brainstem nucleus called Nucleus Z, which in turn sends signals through thalamus to S1 (Landgren & Silfvenius, 1971).

Physiology of proprioception in cortex

S1 itself resides on the postcentral gyrus, just posterior to primary motor cortex (M1). It can be anatomically divided into four areas, named for their initial classification by Brodmann (Brodmann, 1909). In order from anterior to posterior, they are designated areas 3a, 3b, 1, and 2. Of these, areas 3b and 1 contain mostly tactile information, while area 3a and area 2 are thought to be more proprioceptive areas (Jennings, et al., 1983; Kaas, et al., 1979; London & Miller, 2013). Each of these four areas contains a rough somatotopic map of the body, often called the homunculus, arranged medial to lateral along the gyrus (Penfield & Boldrey, 1937; Kaas, et al., 1979; Pons, et al., 1985).

Area 3a is typically centered near the bottom of the central sulcus for old world primates, including humans. Early electrophysiology studies of area 3a showed that the area received input mostly from muscle receptors (Phillips, et al., 1971; Yumiya, et al., 1974; Tanji, 1975; Lucier, et al., 1975), though later studies showed that areas 3b and 1 do appear to project some cutaneous information into area 3a (Heath, et al., 1976; Yamada, et al., 2016). In contrast, area 2 receives strong input from areas 3a, 3b, and 1, and as a result, contains a combination of cutaneous and muscle information (Padberg, et al., 2018; Pons, et al., 1985; Pons & Kaas, 1986; Hyvärinen & Poranen, 1978). Similarly, while neurons in anterior S1 tend to have simple receptive fields, with response properties that match peripheral receptors, neurons in area 2 tend to have larger and

more complex receptive fields (Hyvärinen & Poranen, 1978) and together appear to have a more complex somatotopy than found in more anterior areas of S1 (Pons, et al., 1985). Altogether, this suggests that information in area 2 may be more processed than in anterior areas of S1.

While several studies have examined how S1 relates to proprioception of the hand or wrist (e.g. (Iwamura & Tanaka, 1978; Costanzo & Gardner, 1981; Gardner & Costanzo, 1981; Jennings, et al., 1983; Nelson, 1987; Gardner, et al., 1999; Ro, et al., 2000; Debowy, et al., 2001; Goodman, et al., 2019), relatively few have examined how S1 relates to proprioception of the proximal arm. The majority of these studies show that neural activity in S1 modulates strongly with the position of the hand (Tillery, et al., 1996), the direction of hand movement (Prud'homme & Kalaska, 1994; Weber, et al., 2011; London & Miller, 2013), or the direction of a load applied to the hand (Prud'homme & Kalaska, 1994; London, et al., 2011). In short, these studies implicitly assume that neural activity in S1 represent the state of the hand, despite proprioception arising from muscle sensors. This observation matches with psychophysical results suggesting that people are better at estimating the location of their hands than estimating their joint angles (Fuentes & Bastian, 2010), and surprisingly, this limb endpoint-based representation also seems to explain neural activity in the DSCT, one of the earliest stages of processing for proprioception of the leg (Bosco, et al., 1996; Bosco, et al., 2000).

Another cortical area involved in proprioception is area 5, lying just posterior to S1, and part of the posterior parietal cortex (PPC). Like area 2, area 5 responds to passive joint manipulation and appears to be tuned to the direction of hand movement during reaching (Mountcastle, et al., 1975; Kalaska, et al., 1983; Seelke, et al., 2011). However, area 5 appears to have much broader cortical connectivity than area 2, which mostly receives cortical input from

other areas of S1, primary motor cortex (M1), and the secondary somatosensory cortex. Area 5 receives input from these areas as well, but it also receives strong input from other areas in PPC, along with premotor and supplementary motor cortex (Padberg, et al., 2018). As might be expected from this broad range of inputs, neurons in area 5 have complex response properties, with subpopulations dependent on task engagement (Omrani, et al., 2016) and whether the movement was actively generated or a passive perturbation (Mountcastle, et al., 1975; Chapman, et al., 1984). This broad connectivity, coupled with the fact that parts of area 5 also project to neurons in the spinal cord (Rathelot, et al., 2017), suggests that this area plays a direct role in using proprioceptive information to guide movement.

History of neural representations of movement

The hand-based representational framework of S1 and DSCT draws directly from that proposed by Georgopoulos et al. for explaining M1 activity (Georgopoulos, et al., 1982). In that study, Georgopoulos et al. used the now classic center-out reaching paradigm, in which the monkey reached from a target in the center of a planar workspace to a cued target at the edge of the workspace in one of eight directions. Using this experiment, they found that the activity of many neurons could be explained by a cosine-shaped tuning curve based on the target direction, with the direction of maximum activity defined as the neuron's preferred direction (PD). The tuning curve could be parameterized by the neuron's PD, as well as a baseline firing rate and a modulation depth. (Georgopoulos, et al., 1986) expanded this framework to three-dimensional reaches, showing that neural activity could often be explained by how much the reaching direction differed from the three-dimensional neural PD.

However, other experiments cast doubt on this idea of a hand-based representation of motor control. In old-world primates, M1 contains a large population of cells known as corticomotorneurons, which project directly to spinal motor neurons (Rathelot & Strick, 2009), suggesting that a muscle-based representation might better characterize M1. Because muscle activity also shows cosine-shaped tuning curves to movement direction, it was unclear whether primary motor cortex truly represented movements in terms of the hand or whether it just appeared to do so because signals in hand-based and muscle-based coordinate frames were highly correlated in the original experiments (Mussa-Ivaldi, 1988). In fact, the very first single-neuron recordings in motor areas of behaving animals were those of Ed Evarts. His experiments suggested that identified corticospinal tract neurons in M1 were better related to force than movement (Evarts, 1968). Later reaching experiments showed that changing the location of the workspace in relation to the monkey's body changed neural PDs (Caminiti, et al., 1990; Morrow, et al., 2007), an observation inconsistent with the idea that M1 controls the movement of the hand in a hand-centered, Cartesian coordinate frame. Other experiments showed that the neural PD appeared to change even during the course of the reach (Sergio, et al., 2005; Churchland, et al., 2012) or as the result of forces imposed on the monkey's hand (Kalaska & Hyde, 1985; Kalaska, et al., 1989; Rokni, et al., 2007; Cherian, et al., 2011; Perich & Miller, 2017). These results suggest that while PD analysis might be a useful tool to characterize neural activity as it relates to arm kinematics in particular dynamic conditions, the neurons in motor cortex clearly do not just represent where the hand was going.

Recent studies have begun to examine motor areas in terms of how they process information to generate movements, instead of using the neural PD models of movement

representation (Churchland, et al., 2012; Kaufman, et al., 2014; Sussillo, et al., 2015; Perich, et al., 2018; Dekleva, et al., 2018). However, the limb endpoint-based model remains the dominant model for describing proprioceptive neural activity related to the proximal limb. Studies that assume the endpoint-based model generally neglect the fact that these signals originate from the muscles of the limb, implicitly assuming the system has integrated the muscle afferents to compute the location of the limb's endpoint. However, it is unclear whether such a computation actually occurs before the signals reach S1, let alone DSCT. Thus, the central question of this dissertation is whether there is evidence of such a computation in the proprioceptive system.

The role of proprioception in generating movement

The idea the early proprioceptive signals may be integrated to form representations of movements frames proprioceptive areas in terms of a hierarchy of processing, like that described for visual cortices (Hubel & Wiesel, 1959; Hubel & Wiesel, 1962; Felleman & Essen, 1991). However, another way to frame proprioceptive areas is in terms of how they contribute to movement generation. Proprioception is critical for generating coordinated movement, as evidenced by studies of patients who have lost it (Sainburg, et al., 1993; Sainburg, et al., 1995; Ghez & Sainburg, 1995; Gordon, et al., 1995). Thus, by considering how proprioceptive areas communicate with motor areas, we may gain insight into how these areas encode feedback about the limb.

Spinal reflexes were among the first studied neural circuits for motor control. In particular, proprioceptive feedback from muscle spindles underlies a reflex so well-known that it has its own cliché: the "knee-jerk response", otherwise known as the monosynaptic stretch reflex (Liddell & Sherrington, 1924). In this reflex, spindle activity directly recruits the spinal motor

neurons associated with the muscle in which the spindle resides, causing that muscle to contract when it is stretched. This can be thought of as a negative feedback loop, encouraging muscles to stay at a set length, in a spring-like fashion (Nichols & Houk, 1973). This stretch reflex, along with other spinal reflexes, can also be modified by motor commands, for example through gamma motor drive to spindles (Kuffler, et al., 1951) or through changing spinal motor neuron excitability (Hounsgaard, et al., 1988). Surprisingly, one recent study also demonstrated that in some cases, these spinal reflexes can respond to perturbations in a task-dependent manner, a function usually thought to require a feedback loop through cortical areas (Weiler, et al., 2019).

The spring-like nature of the muscles, brought about by the stretch reflex, became one basis for a theory of motor control called the equilibrium point hypothesis (Asatryan & Feldman, 1965; Feldman, 1966; Feldman, 1986; Flanagan, et al., 1993). The idea behind this hypothesis was that upstream motor commands could set an equilibrium point for each of the muscles in the arm, resulting in a restoring force to carry the hand to the planned set point. This idea was attractive, given that it offloaded the necessary computations for generating movements to the feedback system – the motor system only had to set an end goal to generate the movement. However, this idea became controversial, with some studies supporting the hypothesis (Bizzi, et al., 1984; Gribble, et al., 1998), and others exposing flaws in it (Lackner & Dizio, 1994; Gomi & Kawato, 1996; Hinder & Milner, 2003). As with the controversy concerning the coordinate frame of movement representation in motor cortex, this equilibrium point hypothesis controversy has not been completely resolved; instead, the field has turned towards other views of how feedback is integrated into motor control.

Arguably the most influential of these views is one that suggests the motor control system uses internal models of the limb to generate coordinated movements and integrate feedback (Shadmehr & Mussa-Ivaldi, 1994; Wolpert, et al., 1995; Ghez & Sainburg, 1995; Wolpert, et al., 1998; Thoroughman & Shadmehr, 1999; Hinder & Milner, 2003; Hwang, et al., 2003; Kurtzer, et al., 2008; Therrien & Bastian, 2015; Maeda, et al., 2018). Central to this internal model framework is the idea that pure feedback control of the limb is difficult, due to delays in both the sensory and motor systems. To overcome these delays, the sensorimotor system may use a set of internal models of the limb. In the simplest realization, this internal model framework boils down to a few key features. First the sensorimotor system chooses a desired trajectory of movement (e.g. "move the right hand to that cup of tea"). An inverse dynamics model of the limb converts this desired trajectory into low-level motor commands, to be sent to the muscles (Shadmehr & Mussa-Ivaldi, 1994). Muscles contract to move the arm, generating proprioceptive and visual feedback. At the same time, a forward model uses the motor commands to predict the sensory consequences of the generated movement, potentially to be used as a surrogate for rapid feedback about the movement (Wolpert, et al., 1995; Wolpert, et al., 1998). Finally, the actual sensory feedback and predicted feedback are compared against each other to estimate the sensory error, which is used to update the internal models to reflect a change in the limb dynamics (Shadmehr & Mussa-Ivaldi, 1994; Ghez & Sainburg, 1995; Mathis, et al., 2017; Maeda, et al., 2018).

How and where might these internal models be implemented in the brain? One important brain area for this internal model framework is the cerebellum, which is potentially important for instantiating both the inverse and forward models and for updating them using sensory feedback

(Wolpert, et al., 1998; Izawa, et al., 2012). On the motor command side specifically, the dorsal premotor cortex (PMd) and M1 are important for planning and generating motor commands. As mentioned in the previous section, there has been some debate about how exactly these areas represent movements, but the most recent studies suggest that they are most accurately modeled as complex dynamical systems that generate commands for the muscles (Churchland, et al., 2010; Churchland, et al., 2012; Sussillo, et al., 2015; Kao, et al., 2015; Kaufman, et al., 2016; Russo, et al., 2018; Pandarinath, et al., 2018). One recent study also shows that changing activity in PMd underlies adaptation to altered dynamics of movement, suggesting that PMd may employ an inverse model, updated by sensory prediction errors, to inform how it recruits M1 in different dynamic conditions (Perich, et al., 2018).

On the sensory side, S1 plays a crucial role in both learning new motor skills (Pavlides, et al., 1993) and adapting to different movement dynamics (Mathis, et al., 2017), pointing to a role for S1 in comparing proprioceptive feedback with the predictions from the forward model. Furthermore, S1 activity often changes just before active movements (Nelson, 1987; London & Miller, 2013), suggesting that it receives information about intended movements from motor areas, as a forward model would. Given that M1 contains at least some information about muscles, we might also expect neurons in S1 to contain a muscle-like representation of feedback, to aid in its communication with M1.

Furthermore, planning a reaching movement requires integration of proprioceptive feedback about where the hand is with information about targets for reaching, the latter typically acquired by vision (Flash & Hogan, 1985; Sainburg, et al., 2003). Supporting this is the fact that patients who have lost proprioception tend to make movement errors seemingly related to an

inability to accurately plan movements (Gordon, et al., 1995). The most likely area for this integration to occur is in PPC, where neurons with complex responses to both proprioception and vision are found (Mountcastle, et al., 1975). In addition to the broad range of inputs to PPC described earlier, this area appears to play an important role in selecting actions and in guiding reaches (Snyder, et al., 1997; Batista & Andersen, 2001; Buneo & Andersen, 2006; Cui & Andersen, 2007), suggesting that PPC may be at least partially responsible for sending high-level commands to an inverse model. As such, we might expect proprioceptive information in PPC to be related more to hand kinematics than forces or muscle lengths, an idea supported by electrophysiological studies (Kalaska & Hyde, 1985; Buneo & Andersen, 2006).

Summary

This chapter has given a brief introduction to how proprioceptive information reaches the brain, a brief summary of the study of neural representations of movement, and finally outlined a few ideas on how proprioceptive feedback plays into the generation of movement. The following chapters will detail my doctoral work examining neural representations in two different proprioceptive areas of the CNS: the DSCT and area 2 of S1.

Chapter 2 details a computational study focused on the DSCT. Earlier studies of this area showed that the DSCT integrates information from the whole hindlimb and seemingly computes a representation of the limb's endpoint, only one or two synapses away from the muscle receptors themselves. However, the simulations described in Chapter 2 show that this apparent endpoint representation might simply arise from the biomechanics of the limb and not due to any specific neural computation.

Chapter 3 details an electrophysiological study of area 2 in S1. In this study, I explored the question of whether S1 represents more than just the movement of the hand. I found that whole-arm-based models could predict features of neural activity that the classic hand-based models could not, suggesting that S1 does indeed represent the movement of the whole arm. Thus, if there truly is a transformation into hand-based coordinates, as conscious experience and psychophysics experiments suggest, it likely has not occurred before the signals get to S1.

Finally, Chapter 4 examines two different perspectives from which to examine the results presented in Chapters 2 and 3. The first concerns how proprioceptive information is processed hierarchically by the proprioceptive system, while the second concerns how proprioception contributes to motor control. In light of these perspectives, this final chapter will discuss the implications of these results, along with future directions in which to extend this work.

Chapter 2 - Musculoskeletal geometry accounts for apparent extrinsic representation of paw position in dorsal spinocerebellar tract

Raeed H. Chowdhury, Matthew C. Tresch, and Lee E. Miller

Foreword

The following chapter has been adapted from a manuscript published in the *Journal of Neurophysiology* in April 2017. The purpose of this project was to understand the results from a series of studies of cat dorsal spinocerebellar tract (DSCT) conducted by Bosco and Poppele. One of these experiments found that neural activity related to the position of the hindlimb paw, even when a constraint was added to change how the joint angles related to the paw position. This result was surprising -- do neurons in the spinal cord, only one synapse removed from the muscle receptors, truly compute the position of the hindlimb paw, or is this apparent representation a consequence of musculoskeletal geometry? In this project, I replicated this experiment in simulation to address this question. By simulating artificial DSCT neurons as weighted linear combinations of muscle lengths, I explored how musculoskeletal geometry alone might contribute to this surprising empirical result. These simulated neurons exhibited many tuning properties similar to those found in the actual DSCT, suggesting that a transformation to a limb endpoint representation is unlikely to have occurred at this stage in proprioceptive processing.

Abstract

Proprioception, the sense of limb position and motion, arises from individual muscle receptors. An important question is how and where in the neuroaxis our high level “extrinsic” sense of limb movement originates. In the 1990s, a series of papers detailed the properties of neurons in the dorsal spinocerebellar tract (DSCT) of the cat. Despite their direct projections from sensory receptors, it appeared that half of these neurons had consistent, high-level tuning to paw position rather than to joint angles (or muscle lengths). These results suggested that many DSCT neurons compute paw position from lower level sensory information. We examined the contribution of musculoskeletal geometry to this apparent extrinsic representation by simulating a three-joint hindlimb with mono- and biarticular muscles, each providing a muscle spindle-like signal, modulated by the muscle length. We simulated neurons driven by randomly weighted combinations of these signals and moved the paw to different positions under two joint-covariance conditions similar to the original experiments. Our results paralleled those experiments in a number of respects: 1) Many neurons were tuned to paw position relative to the hip under both conditions. 2) The distribution of tuning was strongly bimodal, with most neurons driven by whole-leg flexion or extension. 3) The change in tuning between conditions clustered around zero (median absolute change ~20°). These results indicate that, at least for these constraint conditions, extrinsic-like representation can be achieved simply through musculoskeletal geometry and convergent muscle length inputs. Consequently, they suggest a reinterpretation of the earlier results may be required.

Introduction

There is an intimate interplay between somatosensation and the control of movement that is evident in the movement deficits that result when somatosensation is lost. Patients make large errors of extent while making reaching movements in different directions (Gordon et al. 1995) and are largely unable to control the dynamics of the limb (Sainburg et al. 1995). Part of the complexity of limb movement control arises from the redundancy that allows us to place our hand or foot in a particular location using an infinite variety of limb configurations. In order to complete a simple reaching movement, a high-level command, e.g. “reach for that book”, must be transformed into commands for dozens of individual muscles. The cerebellum monitors and processes the movement commands themselves, as well as somatosensory feedback from the limb in a process that makes the movement execution fluid and effortless. An important question in motor neurophysiology has been how these sensorimotor signals are represented by populations of neurons, and how these representations are transformed through the central nervous system.

This question has received far more attention in the motor system than it has in the somatosensory system. In a highly cited classic study, Georgopoulos and his colleagues found that neurons in primary motor cortex (M1) appear to represent arm kinematics in an extrinsic, hand-based coordinate frame, where each neuron has broad tuning to a particular direction of hand movement (Georgopoulos et al. 1982). Such a system would require subsequent transformation of this high level control signal into signals appropriate to activate muscles.

Subsequent experiments examined the question of representation more closely, taking advantage of the limb’s redundancy to alter limb posture and change the relation between hand movement and joint rotation (Caminiti et al. 1990; Kakei et al. 1999; Morrow et al. 2007; Scott

and Kalaska 1997). Most of these studies found many M1 neurons with tuning that was more closely related to the "intrinsic" coordinates of joint rotations or muscle lengths. Using similar methods, there is now evidence that much of the transformation from extrinsic to intrinsic coordinates occurs between the premotor and motor cortices (Kakei et al. 2001; Shen and Alexander 1997) and is completed between motor cortex and the spinal cord (Yanai et al. 2008).

Compared to the motor system, the representation of proprioception, the sense of body position, movement, and related forces, has received far less experimental attention. Proprioception arises from muscle spindles, situated parallel to muscle fibers and sensitive to muscle length and length change, and Golgi tendon organs (GTOs), lying in series with muscles and sensitive to active muscle force, and even cutaneous receptors that respond to skin deformation and stretch (Kandel et al. 2012). However, psychophysical experiments in humans (and common experience) suggest that perception of the arm is focused on hand position or overall limb orientation rather than individual joint angles or muscles lengths (Fuentes and Bastian 2010), thus matching vision and the other exteroceptive senses. Consistent with these studies, proprioceptive neurons in primary somatosensory cortex (S1) have been thought to exhibit extrinsic tuning (Prud'homme and Kalaska 1994) quite similar to that suggested for M1 neurons by the experiments of Georgopoulos and colleagues (Georgopoulos et al. 1982). Thus, the system mediating proprioception has the opposite requirement of the motor system: low level sensors in muscles must ultimately generate high level perception in extrinsic coordinates.

Remarkably, experiments involving recordings from neurons in the dorsal spinocerebellar tract (DSCT) of the cat, a pathway carrying proprioceptive information from the cat's hindlimb to the cerebellum, suggested that such a transformation had already occurred at that level (Bosco and

Poppele 1997; Bosco et al. 2000; Bosco et al. 1996). As the hind paw of an anesthetized cat was moved throughout a range of positions in the sagittal plane, DSCT neurons tended to represent the limb position not in terms of individual joint angles or muscle lengths, but in terms of a vector drawn from the hip to the paw. For approximately half of DSCT neurons, this extrinsic representation persisted even in the presence of a constraint that fixed the knee angle and changed how the joints covaried during hindlimb movement.

This result was surprising for a number of reasons. For one, the significant amount of computation implied by this transformation must occur very early in sensory systems, at DSCT neurons, which receive inputs directly from sensory afferents. Furthermore, the posited cerebellar role in the coordination of intersegmental dynamics during movement (Cooper et al. 2000; Sainburg et al. 1995), might well profit from muscle-related feedback, rather than a simple representation of paw position. Finally, arguments for the need for this transformation based on characteristics of our conscious perception are much less convincing in the context of cerebellar processing than they are for signals in the cerebral cortex.

The idea that neurons in the DSCT encode the hindlimb state in terms of the paw position has been well accepted into proprioceptive literature (Daley and Biewener 2006; Fuentes and Bastian 2010; Kim et al. 2010; Morton and Bastian 2006; Ting and Macpherson 2005; Weber et al. 2007). However, the original experiment did not explore the contribution of musculoskeletal geometry to the tuning of neural activity. While the knee-fixed constraint changed the hindlimb joint covariance, hindlimb muscle length changes do not have a simple linear relationship with joint rotation. As such, the effect of this constraint on convergent muscle-length inputs remains unclear. Consequently, we replicated the experiment in simulation, with neurons that drew their

activity directly from randomly weighted combinations of spindle-like muscle-length inputs. Surprisingly, the results of this simple neural model closely paralleled the earlier experimental results. This outcome suggests a very different interpretation of the representation of limb state by DSCT neurons. These results suggest that there is little evidence that the DSCT neural signals have been transformed to represent limb state in anything other than simple muscle coordinates. This outcome also highlights the importance of considering the properties of the musculoskeletal system when interpreting signals recorded in the central nervous system.

Methods

Musculoskeletal Model

To simulate the cat hindlimb, we used a four-segment, three-joint musculoskeletal model, with segment lengths and muscle insertion points adapted from the musculoskeletal model used by (Bunderson et al. 2010), based on an anatomical study performed by (Burkholder and Nichols 2004). We used eight muscles (five monoarticular and three biarticular), and constrained the limb to the sagittal plane. Figure 2.1 shows a schematic representation of the model.

Our simulation experiment included two constraint conditions meant to simulate those of the original experiment: an elastic constraint that approximated normal limb mechanics and a knee-fixed constraint. We implemented the elastic constraint by placing identical springs at each joint to mimic the passive elastic effects of muscles and fascia, thus defining a minimum energy state for any given paw position. This constraint fully determined the hindlimb posture for any given paw position. As in the baseline condition with the anesthetized cat (Bosco and Poppele 1997; Bosco et al. 2000), the joints covaried essentially linearly, with a plane accounting for 94% of the joint covariance. In our second constraint condition we fixed the knee angle, forcing a linear joint covariance different from that of the elastic condition. The consistency of neural

tuning to paw position between these two conditions can be considered a measure of how nearly a given neuron is tuned to extrinsic coordinates.

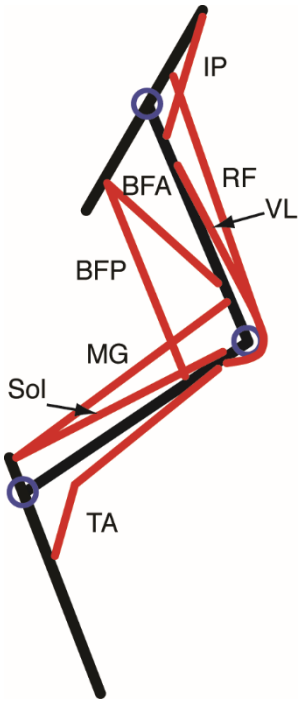

Figure 2.1 - Schematic of simulated leg, showing segments(black), joints(blue), and muscles(red). Muscles are: BFP-Biceps Femoris Posterior, BFA-Biceps Femoris Anterior, IP-Illiopsoas, RF-Rectus Femoris, VL-Vastus Lateralis, MG-Medial Gastrocnemius, Sol-Soleus, TA-Tibialis Anterior. We conducted the simulation by moving the paw to different positions in the sagittal plane in two conditions: elastic joint and knee-fixed.

Generation of neural activity

In our simulation experiments we moved the paw in the sagittal plane through 100 equally spaced locations in a 10x10 polar grid centered on the hip, and spanning angles from 245° to 300° with radii from 18 cm to 22 cm. This matched the grid pattern used in the original experiments (Bosco and Poppele 1997; Bosco et al. 2000; Bosco et al. 1996), and resulted in joint angle excursions of 40°, 49°, and 50° for hip, knee, and ankle, respectively. These excursions are similar to those reported in the original experiments (25-50°, 45-60°, 60-80° excursions for hip, knee, and ankle joints across several experiments). Varying the stiffness of each joint in the model (±25%) altered the simulated joint excursions but did not affect the overall conclusions reported here (see Results).

For each paw position, we found the corresponding hindlimb configuration and calculated the length of each muscle from origin to insertion from the musculoskeletal model. The maximum change in muscle length for a given paw displacement corresponds to the muscle's pulling direction. Muscle lengths were normalized to lie between 0 and 1, corresponding to the minimum and maximum lengths achieved throughout the full workspace. We transformed the resultant muscle lengths into simulated neural discharge by taking a weighted sum of the normalized lengths to simulate the convergence of length-sensitive afferents onto a given DSCT neuron (Figure 2.2). Each neuron was thus characterized fully by its eight muscle weights, which we drew randomly from a standard normal distribution ($\mu=0$, $\Sigma=I$). The zero mean allowed negative (inhibitory) weights. This raw neural activation was passed through a sigmoidal function to produce an average firing rate for each neuron that ranged between 0 and 60 spikes/second. Lastly, we used this average firing rate as the intensity parameter of a Poisson

process to simulate spike trains for 10,000 neurons as the paw was held for two seconds in each of the 100 positions of the 10x10 workspace.

Given the muscle model used to generate the activity of these neurons, the comparison across conditions reveals how consistent paw-position tuning might be if DSCT neural responses were driven directly by muscle sensor outputs, where the only neural computation is the summation of randomly scaled inputs. While it is unlikely that real neurons in the DSCT have fully random connectivity from muscle lengths, the use of random weights allowed us to probe the limits of the potential role of musculoskeletal geometry, as opposed to more complex learned weights, on the neural representation of limb state in DSCT.

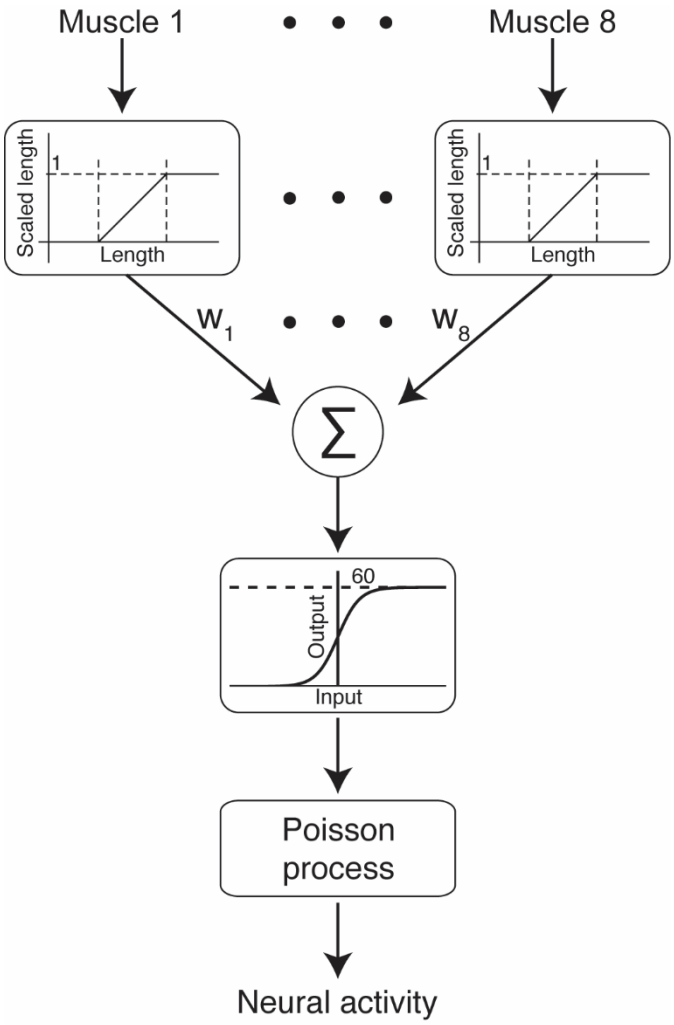

Figure 2.2 - Block diagram for generating neural activity. Muscle lengths are first normalized to values between 0 and 1, then linearly combined with random weights drawn from a standard normal distribution. Output is then passed through a sigmoid function and a Poisson process to simulate neural activity.

Analysis of neural representation

Unless noted otherwise, we analyzed the simulated firing rates using the methods from the original experimental paper (Bosco et al. 2000). To find tuning in Cartesian coordinates, we fit the simulated firing rates to the x and y coordinates of paw position. These planar fits yielded a gradient direction (θ_G), which denotes the direction of maximal change in firing rate, and a gradient magnitude (ρ_G), which denotes the overall sensitivity of the firing rate to displacement from the center, shown in Equation (2.1):

$$f = \alpha_0 + \rho_G \cdot d \cdot \cos(\theta - \theta_G) \tag{2.1}$$

Where:

$$d = \sqrt{x^2 + y^2}\,, \theta = \tan^{-1} {}^{y}/_{x} \tag{2.2}$$

In Eq. 2.1, d represents the Euclidean distance from the center of the grid, while θ represents the direction of displacement from the center, defined in Eq. 2.2. If a neuron changes its Cartesian representation between experimental conditions, one or both of these parameters would change in some way. A change in θ_G might be considered more interesting than a change in ρ_G, as altered sensitivity could be caused simply by global effects on the neurons in the DSCT that result in an overall increase or decrease in activity. On the other hand, a change in θ_G represents a differential change in sensitivity to the two cardinal axes of movement. Such a change resulting from different biomechanical conditions (as in the two conditions in this study) would contradict the hypothesis that the neuron encodes paw position. For this reason, our analysis focuses on changes in θ_G between the two experimental conditions.

We only considered neurons that were well tuned in both the elastic and knee-fixed conditions, as the θ_G is otherwise undefined. As in the original experiment, we considered a neuron to be tuned if it satisfied two conditions: first, using an F-test, the fit of Eq. 2.1 needed to be significant ($p < 0.05$) compared with a constant firing rate model; second, the R^2 of the model fit by Eq. 2.1 needed to be greater than 0.4 (Bosco et al. 2000).

Results

Tuning of intrinsic neurons to paw position

Within either joint covariance condition, the relationship between the intrinsic and extrinsic coordinate systems is locally linear (Bosco et al. 1996; Mussa-Ivaldi 1988). In motor systems, it has been common to impose a postural perturbation as a means of disambiguating these two possible representations (Caminiti et al. 1990; Morrow et al. 2007; Oby et al. 2013; Scott and Kalaska 1997). While less common in the somatosensory system, this approach has been used to suggest that at just one synapse removed from the afferent receptors at the level of the dorsal spinal cerebellar tract, neurons encode limb state in extrinsic coordinates (Bosco and Poppele 1996; Bosco et al. 2000). We sought to better understand the origin of this remarkable result. To this end, we conducted a simulation to determine the conditions under which muscle-length sensitive neurons might exhibit apparent tuning to paw position in extrinsic coordinates.

Figure 2.3A shows the activity of a well-tuned neuron as a function of paw position in the elastic condition. When fit with Eq. 2.1, R^2 was 0.64 ($p<0.001$). θ_G for this neuron (indicated by the arrow) was 80°. Sixty-one percent of all simulated neurons were well tuned to paw position ($R^2>0.4$, $p<0.001$) in the elastic condition.

Among this group of well-tuned neurons, the distribution of θ_G was strongly bimodal. Figure 2.4A shows this distribution in the elastic constraint condition, with a significant mean axis running from 82° to 262° (circular mean test $p<0.0001$), which approximately matched the axis running between the hip and paw (which we call the "limb axis"), indicating that neurons were much more likely to be tuned to whole leg extension and withdrawal, than the orthogonal direction of movement (the "orientation axis"). This distribution was remarkably similar to that found in the studies of Bosco and colleagues (Figure 2.4B). Because the tuning of our simulated neurons resulted directly from that of the muscles, this nonuniformity in the distribution of θ_G suggests that the hindlimb muscles also have their greatest length change for paw movement along this axis. Intuitively this is quite reasonable; while the workspace spans a greater range along the orientation axis, only the hip muscles are strongly affected by orientation changes. In contrast, paw movement along the limb axis affects all of the muscles, leading to greater overall muscle length change and a greater sensitivity of the neurons to these movements.

Tuned and untuned neurons received essentially the same magnitude weight of inputs from all muscles (medians 2.75 and 2.65, respectively). While we did not find any clear biases in tuned neurons for higher weights on certain muscles than others, we do note that some combinations of inputs (for example, equal inputs from antagonistic muscles with opposing pulling directions) contribute no net neural modulation, while others are much more effective. These combinations of inputs were quite different for tuned and untuned neurons. We computed a net input vector for each neuron, calculated as the vector sum of all muscle pulling directions weighted by each muscle's input strength for that neuron. The length of this vector, normalized by the scalar sum of the input weights, provides a measure of the "effectiveness" of the inputs.

The input to well-tuned neurons was much more effective than that for untuned neurons (a 36% increase for the elastic condition and a 48% increase for knee-fixed).

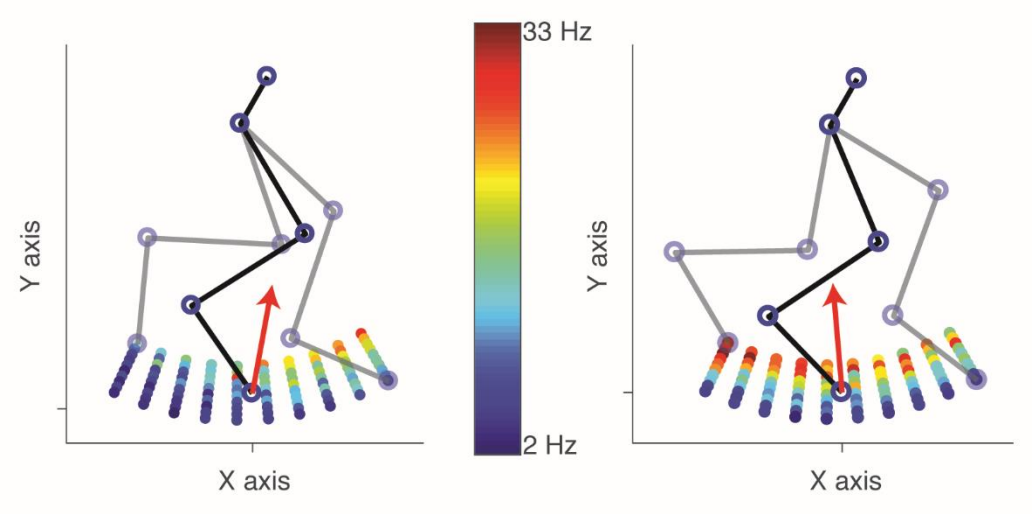

Figure 2.3 - Heat map diagram showing activity of an example simulated neuron at different foot positions during elastic joint (A) and knee-fixed (B) conditions. Red indicates high activity, while blue indicates low activity. Red arrows show the "gradient direction" (θ_G) of this neuron in both conditions. For this neuron, θ_G changed by 15°.

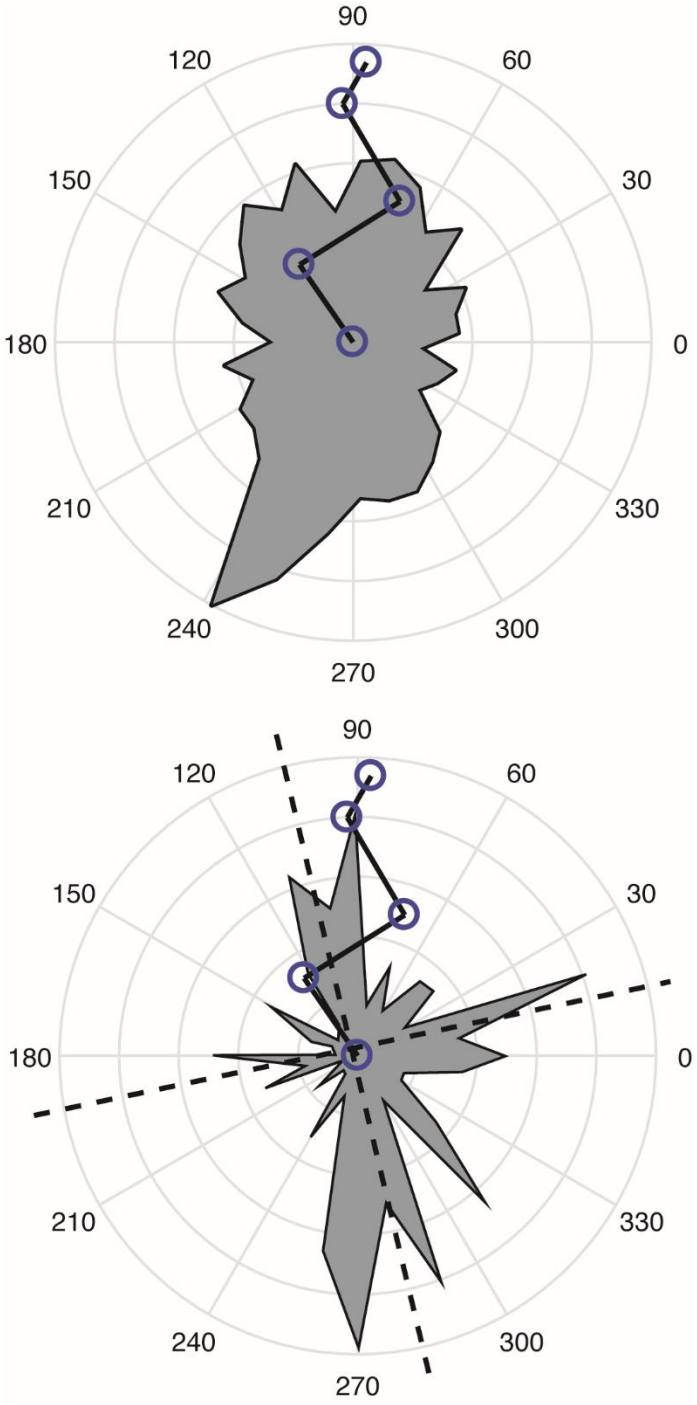

Figure 2.4 - Distribution of gradient directions for tuned neurons in simulation (A) and in empirical recordings of previous studies (B). In both cases, the distribution was mostly bimodal along the axis of the leg (the "limb axis"). Panel B, redrawn from (Bosco and Poppele 2001). Dashed lines represent the primary axes of the distribution's lobes.

Stability of cross-condition gradient directions

Among our simulated neurons, 61% had statistically significant tuning to paw position in the elastic condition, and 69% were tuned in the knee-fixed condition. Fifty-two percent were tuned in both conditions. For these neurons to be considered endpoint-tuned, θ_G must be invariant across these two conditions, for which the mappings between paw position and joint angles differ. This is indeed what Bosco and colleagues found: θ_G for many DSCT neurons remained largely unchanged across the elastic and knee-fixed conditions (Bosco et al. 2000). We made a similar comparison here, examining θ_G among the 52% of simulated neurons that were well tuned to paw position under both conditions.

The θ_G of the example neuron of Figure 2.3 rotated from 80° in the elastic condition to 95° in the knee-fixed condition. (Figure 2.3B). Unexpectedly, among neurons tuned to both conditions, the majority had similarly small changes in θ_G between conditions. Figure 2.5A shows a histogram of the cosines of these $\Delta\theta_G$, a typical measure of the alignment of two vectors. Values close to 1 represent neurons that had very small absolute $\Delta\theta_G$, while numbers close to 0 represent neurons with nearly orthogonal θ_G. In our data, the median of the cosine of $\Delta\theta_G$ was 0.94, corresponding to $|\Delta\theta_G|$ = 20°. For comparison, Figure 2.5B shows the corresponding original experimental result: a median change in the cosine of $\Delta\theta_G$ of 0.91 (23°) for all tuned neurons (Bosco et al. 2000).

Combined with the observation of a bimodal distribution of θ_G along the limb axis, the stability of θ_G between experimental conditions constitutes a surprisingly close match between actual and simulated DSCT neurons. These results beg the question of what mechanism generates the consistent endpoint tuning across the joint coupling conditions. Because the neural

activity in our simulation was generated from randomly weighted combinations of muscle lengths, it weakens the argument that the apparent endpoint tuning arises from a specific weighting of inputs.

One potentially simple explanation for this stability of extrinsic tuning is the effect of multiarticular muscles. Because they span multiple joints, their length changes might be more directly related to paw position than individual joint angles. However, when we removed the biarticular muscles from our musculoskeletal model, the stability of endpoint representation largely persisted: 46% of neurons were well tuned to paw position in both constraint conditions, and the median $|\Delta\theta_G|$ was 15°. Thus, while there were fewer well-tuned neurons without the biarticular muscles, the tuning direction change was actually smaller. Evidently, there is a different explanation for the stable extrinsic tuning that is related to how muscle lengths are affected by paw position.

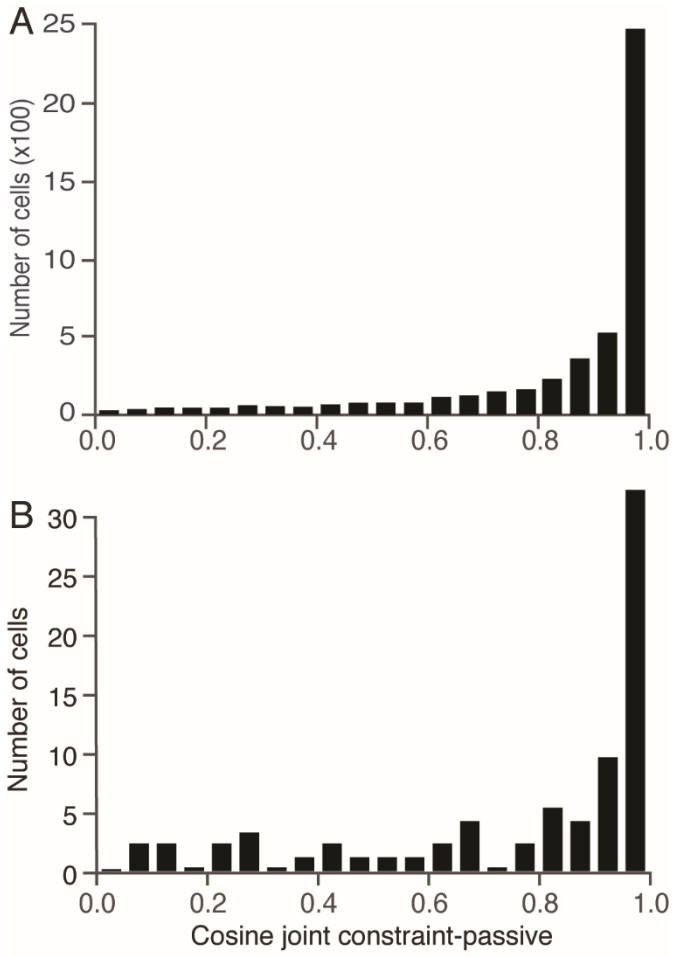

Figure 2.5 - Histogram of cosine of change in θ_G between experimental conditions in simulation (A) and in previous empirical recordings (B). Both distributions were heavily biased towards 1, indicating that most neurons did not change the gradient direction substantially. For simulated neurons (A), the median value was 0.94, corresponding to a change of ±20°, and for empirically recorded neurons (B), the median value was 0.91, corresponding to a change of ±23°. Panel B redrawn from (Bosco et al. 2000).

The effect of muscle tuning across constraint conditions

To investigate the role of musculoskeletal geometry in the bimodal tuning distribution and stability of θ_G across constraint conditions, we examined muscle length tuning in both conditions. Figure 2.6A shows the tuning to paw position of all muscles in the elastic condition. The length of each vector denotes the sensitivity (ρ_G) of the muscle length to paw movement, and its direction denotes the θ_G. As shown graphically by the ellipse, the principal axis of these vectors, calculated using PCA, was along an axis between 92° to 272°, only 10° away from the axis of Figure 2.4A. This principal axis, accounting for 80% of the total variance of these tuning vectors, forms the major axis of the ellipse. Likewise, the minor axis corresponds to the magnitude of the second principal component. Thus, in the elastic condition, the hindlimb muscles changed length mostly for paw movements along the limb axis.

Figure 2.6B shows the corresponding muscle tuning for the knee-fixed condition, which altered the tuning vectors of individual muscles. Vastus lateralis (VL), a monoarticular knee extensor, could not change length, and thus lost all relation to paw movement. In compensation for the lack of knee movement, the ankle needed to move through a much greater range in order for the paw to reach all the positions in the workspace. Consequently, the tuning vectors of soleus (Sol) and tibialis anterior (TA), monoarticular ankle extensors and flexors, respectively, increased considerably in length, as did that of the medial gastrocnemius (MG), a knee-ankle biarticular muscle. Additionally, the monoarticular hip muscles, iliopsoas (IP) and biceps femoris anterior (BFA) changed in tuning direction to reflect the change in hip movement under the knee-fixed constraint. All three ankle-related muscles (MG, TA, Sol) took on essentially identical tuning direction (disregarding sign). Likewise, all hip-related muscles, including the hip-knee bi-articular muscle rectus femoris (RF), had the same tuning directions. This grouping

of muscle tuning directions in the knee-fixed condition reflects the fact that all muscles were essentially reduced to monoarticular flexors or extensors; the two groups of tuning vectors represent the influence of the hip and the ankle rotation on paw position.

Despite these changes, the overall biased distribution of muscle tuning vectors remained largely unchanged across constraint conditions. The orientation of the covariance ellipse with the knee constraint (Figure 2.6B) was very similar to the orientation under the elastic condition (Figure 2.6A); their principle axes differed by only 5 degrees. The similar distribution of preferred neural tuning vectors in the two constraint conditions (Figure 2.4) and the small change in tuning direction for individual neurons (Figure 2.5) are both likely a direct consequence of the preserved biased distribution of muscle tuning vectors. Neurons with significant tuning in both conditions were very likely to have had a tuning vector along the principle axes shown in Figure 2.6; because the direction of this principle axis didn't change substantially across the two conditions, the preferred directions of most neurons didn't change either.

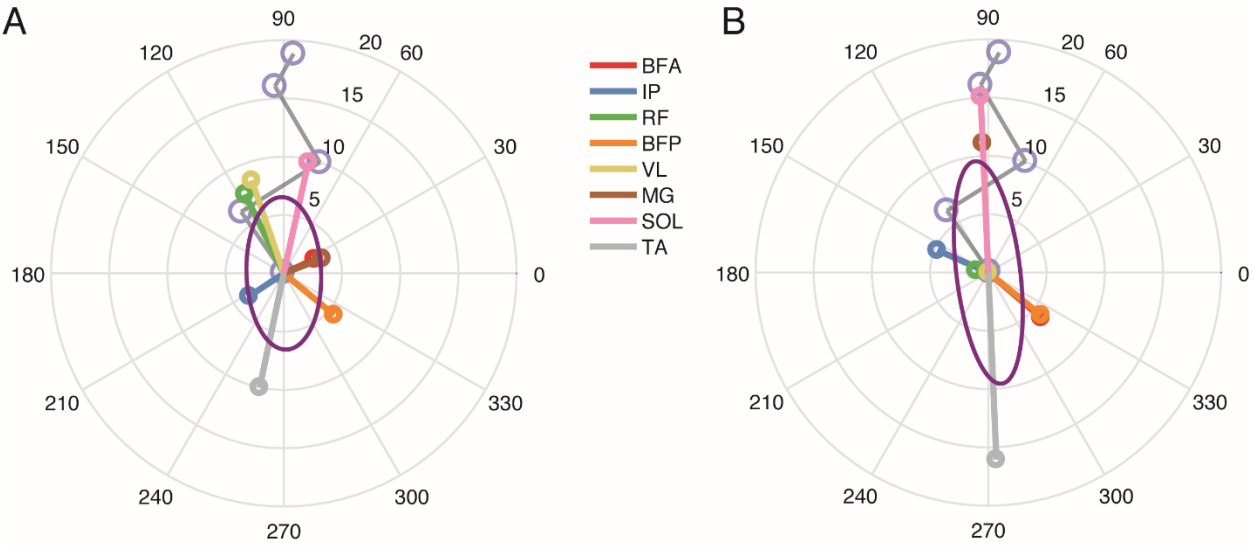

Figure 2.6 - Muscle tuning vectors in elastic joint (A) and knee-fixed (B) conditions. Ellipses in both panels illustrate the variance of the muscle tuning vectors along the principal axes. Tuning vectors were strongly bimodal along the limb axis, in essentially the manner of the neurons. In the knee-fixed condition, VL, a monoarticular knee muscle, became unmodulated, and the ankle muscles (TA, Sol, MG) changed modulation and hip muscles (IP, BFA) changed tuning direction to compensate. Despite this, the overall distribution of muscle tuning did not change, leading to only small changes in neural tuning direction.

Sensitivity to simulation parameters

Muscle tuning vectors reflect the changes in muscle length resulting from displacements of the paw. The relationship between endpoint displacement and joint rotation is given by the limb Jacobian: $dx = Jdq$. Given a set of joint angle changes we can find the resulting endpoint displacement. However, because the limb is redundant, the opposite relationship is ill-posed and only becomes well-defined by the addition of constraints, such as the joint stiffnesses included in the elastic condition, reflecting contributions from passive musculoskeletal properties. The further effect of the knee constraint therefore depended on the particular joint stiffnesses used in our simulation. As a trivial example, if the knee joint stiffness were much higher than the stiffness of the other joints, then fixing the knee would have minimal effect on neural tuning vectors. We examined the extent to which our results depended on the particular values of stiffness by repeating the analyses after varying the stiffness of each joint by ±25%. The number of neurons tuned to both conditions varied only slightly, between 52% and 56%, and median $|\Delta\theta_G|$ for the tuned neurons varied only between 20° and 25°.

In order to evaluate whether a neuron changed its tuning across constraint conditions it had to have statistically significant modulation of its tuning across paw positions. These statistical tests are dependent on the noise inherent in the neural discharge and the confidence with which we could measure tuning. We therefore examined whether the main results depended on the amount of uncertainty in neural tuning. To change the measurement uncertainty, we varied the measurement time from two seconds, the length time used to compute firing rates in most of the original analyses, to a range between 0.1 and 5 seconds. Longer measurement times meant less noise. As we increased the measurement time, the number of tuned neurons increased tremendously from 7.7% at 0.1 seconds, to 69% at 5 seconds. However, despite this wide range

in number of tuned neurons, the median $|\Delta\theta_G|$ remained largely unchanged, ranging from 16° to 21°. For simulations in which more than 20% of neurons were tuned (measurement times >0.5 seconds), the median $|\Delta\theta_G|$ ranged between 19° and 21°. This indicates that Poisson noise, apart from determining how many neurons were well tuned to paw position, did not change the stability of their tuning or the major conclusions of the study.

Discussion

Summary

Because of the nearly linear relation between joint angles and paw position in any restricted workspace, determining the coordinates in which proprioceptive neurons are tuned requires an experiment that alters this relation. Bosco et al. altered this relation by fixing the knee and thereby altering the joint covariance. They found relatively little change in neural tuning with respect to paw position across the naturally constrained limb and the knee-fixed limb and reasoned that neurons with this property should be considered extrinsically tuned.

Unexpectedly, we found that their analyses applied to our neurons, which received only random combinations of muscle length inputs, closely matched many tuning characteristics of the actual DSCT neurons. We focused on two specific analyses used in the original experiments. The first determined the distribution of neural tuning direction, which was heavily biased to be along the limb axis direction. The second determined the overall change in this tuning direction between two different joint covariance conditions, which was generally small for most neurons. The results of both analyses closely matched the empirical results.

In addition to these analyses, Bosco et al. performed a set of regression analyses, which they used to classify neurons into groups that responded consistently to paw position, joint

angles, or neither. We chose not to replicate these analyses, which would have been sensitive to a number of parameters (e.g., measurement errors using skin markers, noise characteristics of DSCT firing rates, etc.) that we had no means to estimate accurately. Since we did not want simply to tweak the simulation parameters to achieve a match with their results, we sought to assess the overall tuning characteristics of neurons without actually classifying them.

Our results serve a dual purpose. First, they provide necessary context for the previous experiments in the cat DSCT. The substantial signal transformation implied by those results was quite surprising, given the close proximity of DSCT to the periphery. Apparently, despite its seeming severity, the knee-fixed constraint did not cause a sufficient change in musculoskeletal mechanics to determine whether DSCT neurons were extrinsically tuned. Consequently, a simple linear convergence of muscle spindle-like inputs, given the musculoskeletal geometry, provides an adequate explanation for the apparent tuning to paw position.

Second, these results show the importance of simulation to inform the design of similar experiments, providing the basis for an appropriate null hypothesis. In this case, the gradient direction analysis shows very little change between conditions; thus, in order for DSCT neurons to be considered extrinsic, they must demonstrate even less change than this null hypothesis. An appropriate use of this simulation for future experiments could be to choose a more effective constraint that maximized the change in endpoint tuning given muscle-based neurons.

Contributions of musculoskeletal properties to apparent neural coding

A key question of our study was to determine why the neural tuning was stable across two seemingly very different joint constraints. We considered whether the stable endpoint tuning of length-sensitive DSCT neurons might have arisen because of the geometry of bi-articular

muscles, the length of which is not a simple function of individual joint angles. However, when we removed the bi-articular muscles from our model, the results remained qualitatively unchanged.

Instead, the invariance appears to be due to the strongly bimodal distribution of the muscle tuning vectors along the limb axis for both constraint conditions (Figure 2.6). In the normal, elastic condition, this bimodality is understandable biomechanically, as all three joints rotate during whole-leg flexion and extension, but for anterior-posterior paw movement, muscles other than those at the hip, changed very little. This strong bimodality was not substantially affected by the knee constraint. Several muscle tuning vectors changed, as hip and ankle rotation compensated for the lost knee rotation. However, the greatest effect on muscle tuning vectors was on their length, most obviously for the ankle muscles tibialis anterior, medial gastrocnemius, and soleus (gray, pink, and brown vectors, respectively), which all lengthened dramatically. The overall bimodality along the limb axis was preserved.

As context for the empirical results, this simulation makes the discovery of constraint-invariant neurons much less surprising. The musculoskeletal geometry of the hindlimb can explain both the non-uniform distribution of tuning to paw positions and the stability of tuning between the knee-fixed condition and a naturally-moving condition. As such, the most straightforward explanation for the behavior of DSCT neurons is that it is the result of a convergence of muscle inputs.

It is important to note once again that our analysis, as that of the original study, considered only neurons that were significantly tuned to paw position. Only 61% of the simulated neurons were tuned in the elastic condition, a number that differed greatly from the

95% in the original study. This large difference is at least in part, a product of the random connectivity used in our simulations. In the simulation, a neuron could receive equal inputs from antagonistic muscles, essentially nulling the effect of paw movement on the neuron's firing rate. Indeed, the input to tuned neurons was 35-50% more effective at modulating neural activity as input to untuned neurons, despite there being no difference in the magnitude of the weights themselves. A significant amount of weight cancellation is unlikely to occur in the actual DSCT; such ineffective combinations of inputs are likely to be pruned during development, leaving a bias towards tuned neurons.

One caveat to this study, as well as in the original experimental study, is that the relationship between limb configuration and neural activity during normal behaviors can also be altered by fusimotor drive. In the anesthetized conditions of the origin experiments, this drive would be significantly reduced (Prochazka et al. 1977), and our simulations matched these experimental conditions by deriving neural activity directly from muscle length. Nonetheless, an important question is how variations in this drive might have changed the apparent representation in DSCT. Normally, both dynamic and static gamma drive appear to be modulated during the step cycle. Although the time course of this modulation remains uncertain, gamma static drive (most relevant for the postures represented in this study) increases abruptly at movement onset, with a sustained, dramatic increase during muscle shortening (Taylor et al. 2000). A further complication is that the particular pattern of gamma activation appears to depend on muscle type, be it flexor, extensor, or bi-articular (Loeb 1981). More study is required to understand how this changing fusimotor drive would affect the neural representation in DSCT.

Simulation for experimental design

In addition to shedding light on the previous empirical results, this simulation can help in designing experiments to test the coordinate frame of DSCT or other neurons. Simulation can be used for this in three separate ways: choosing the best analyses for assessing tuning stability, generating a null hypothesis to test results against, and optimizing the constraint to cause maximum tuning change.

With respect to analytical approaches, the original experiment (and our simulation) examined changes only in the *direction*, θ_G of the neural response gradient, not its magnitude (ρ_G). As we suggested above, it is not unreasonable to discount changes in ρ_G as they represent simply an overall change in sensitivity to movement, rather than a differential change in how neurons respond to movement in different directions. However, because the changes imposed by the knee-fixed constraint were focused mostly along the limb axis, neurons with gradient vectors in that direction would be predicted to change more in length than direction. In our simulation, we found a median change in ρ_G of 38% from the elastic to the knee-fixed condition. Hence, for this particular constraint, considering changes in ρ_G as well as θ_G might have clarified the interpretation of whether or not DSCT neurons are extrinsically tuned.

Secondly, such a simulation would be useful in generating a null hypothesis based on intrinsic tuning to which empirical results could be compared. For example, in addition to an analysis of Cartesian tuning direction, the original experiments included an analysis in polar coordinates, in which each neuron's firing rate was fit to the length and orientation of a vector that pointed from the hip to the paw. Polar tuning of approximately half of the DSCT neurons did not change significantly across conditions, which led the authors to conclude that they were

extrinsically tuned. However, the t-test used for that analysis should be used only to show that two distributions *are* significantly different; in this case, that a particular neuron is *not* extrinsically tuned. The fact that a neuron fails the t-test is not sufficient to conclude that it has the same tuning in both conditions. A more appropriate test would be an equivalence test, which, when passed, shows that two quantities are within some given bound of each other. This bound represents a difference that is scientifically (as opposed to statistically) significant. In the θ_G analysis, the equivalence test bound on the median $\Delta\theta_G$ might be 20°, as that value results from randomly connected DSCT neurons. If an equivalence test shows a median change significantly less than this bound, it provides good evidence that the neurons are extrinsically tuned.

Finally, such a simulation could be used for basic experimental design. Our simulation showed that the knee-fixed constraint did not significantly change how the muscles respond to paw movement compared to the natural elastic constraint condition. Invariance in neural tuning despite a constraint that caused significant changes in muscle tuning along the orientation axis would be much stronger evidence of extrinsic tuning in the DSCT.

Broader implications of these modeling results

An important question is the extent to which our conclusions apply more broadly than to the particular sagittal plane motions and knee-fixed constraints of this experiment in the somatosensory system. The classic studies of Georgopoulos revealed hand-centered tuning of M1 neurons during reaching (Georgopoulos et al. 1982) that was analogous to the apparent paw-centered tuning of DSCT. However, numerous subsequent studies found clear evidence that the tuning of most M1 neurons is not consistent with a simple endpoint model across different workspaces (Caminiti et al. 1990; Morrow et al. 2007) or limb postures (Cherian et al. 2011;

Scott and Kalaska 1997). A recent simulation study demonstrated the importance of a variety of musculoskeletal properties on the distribution of M1 preferred directions (Lillicrap and Scott 2013). Interestingly, they found that limb geometry was more important than the presence or absence of biarticular muscles, an observation that is consistent with ours.

The Brain-Machine Interface (BMI) field currently faces similar difficulties. Remarkably accurate predictions of hand trajectory can be made using simple linear combinations of M1 firing rates (Carmena et al. 2003; Serruya et al. 2002; Taylor et al. 2002; Wessberg et al. 2000). It is likely that BMIs based on such extrinsic coding models will perform poorly when used to predict movements following postural perturbations that dissociate extrinsic from intrinsic variables, as suggested by previous studies (Cherian et al. 2011; Oby et al. 2013). Likewise, recordings from small ensembles of afferent neurons in the dorsal root ganglia in cats have been used to predict the kinematics of the paw in extrinsic coordinates during passive limb displacement under anesthesia (Stein et al. 2004) and walking on a treadmill (Weber et al. 2006), but the success of these methods may be due at least partially to the constrained biomechanics of the tested movements. The results of the present study suggest that properties of the musculoskeletal system being studied should be carefully considered when evaluating the ability of BMI interfaces to predict limb movements.

Conclusion

We have shown that under the conditions of the original experiments (sagittal plane motion, unchanging gamma drive), limb biomechanics and a simple linear convergence of muscle-length inputs can replicate the apparent extrinsic representation of DSCT. Given that our muscle-based simulation replicated key features of neurons in the DSCT, there is no longer justification to

conclude that the signals carried by these neurons have undergone a transformation to produce a high-level representation of limb state.

Chapter 3 - Area 2 of primary somatosensory cortex encodes kinematics of the whole arm

Raeed H. Chowdhury, Joshua I. Glaser, and Lee E. Miller

Foreword

This chapter is an adaptation of a manuscript submitted to *bioRxiv* as a preprint and to *eLife* for publication as a research article. The manuscript details two projects that were developed in parallel, but in the end told similar stories. One of these projects was a follow-up of the simulation study presented in Chapter 2, applied to neural activity in area 2 of primary somatosensory cortex (S1). Like the classic DSCT literature, previous literature studying this area related neural activity to the movement of and forces on the hand, rather than the muscles. I developed an experiment in which a monkey reached to targets in two separate workspaces, finding that models incorporating the kinematics of the whole arm could predict area 2 neural activity better than a model based on just hand kinematics.

In the second project, I studied the representation of active and passive movements in area 2 neural activity. Previous studies suggested that these two movements are represented similarly by individual neurons in area 2. However, Dr. Brian London, a previous student in Dr. Lee Miller's lab, found that area 2 neurons represent these movements differently at the population level, suggesting a component of area 2 activity related to volition. While trying to characterize this volitional component, I found that my whole-arm kinematic models from the two-workspace project could explain this representational difference, suggesting that it stemmed not from volition, but from a difference in how the whole arm moved in active and passive movements. As the two projects told similar stories, we combined them into this one manuscript.

Abstract

Proprioception, or the sense of body position, movement, and associated forces, remains poorly understood, despite its clear importance to making coordinated movements. In the classic model of neurons in area 2, a proprioceptive area of primary somatosensory cortex, neural activity is related simply to the movement of the hand and interaction forces it encounters during movement, similar to our conscious experience of proprioception. However, in two separate experiments, examining 1) active and passive movements and 2) reaches to targets in two different workspaces, we found that a model of area 2 activity built on kinematics of the whole arm successfully predicted how features of neural activity changed across movement conditions, while the classic hand-based model was unable to. This suggests that, unlike our conscious experience and the classic model of proprioception in cortex, neural activity in area 2 represents movement of the whole arm.

Introduction

Moving around in an uncontrolled environment is a remarkably complex feat. In addition to the necessary computations on the efferent side to generate movement, an important aspect of motor control is the afferent information we receive from our limbs, essential both for movement planning and for the feedback it provides during movement. Of the sensory modalities we receive, proprioception, or the sense of body position, movement and forces, is arguably the most critical for making coordinated movements (Ghez and Sainburg 1995; Gordon et al. 1995; Sainburg et al. 1995; Sainburg et al. 1993; Sanes et al. 1984). However, despite its importance, few studies have explicitly addressed how proprioception is represented in the brain during natural movement. In comparison, touch, vision, and the motor areas of the brain have received far more attention.

Our conscious experience of proprioception typically focuses on where our hands are going, rather than the rotations of our joints. This also matches with psychophysical data showing that humans are better at estimating the location of the hand than estimating joint angles (Fuentes and Bastian 2010). Perhaps consequently, of the few studies that examine central nervous system (CNS) activity underlying limb proprioception, most assume that individual neurons represent the movement of the limb's endpoint, rather than joint angles or muscle lengths (Bosco et al. 2000; Bosco et al. 1996; London and Miller 2013; London et al. 2011; Prud'homme and Kalaska 1994; Weber et al. 2011). However, to achieve such a hand-centric view, neurons would need to integrate the proprioceptive signals from sensors in the muscles, joints, and skin of the whole limb (Goodwin et al. 1972; Proske and Gandevia 2012).

Area 2 of primary somatosensory cortex (S1) receives a combination of muscle and cutaneous information (Hyvärinen & Poranen, 1978; Padberg, et al., 2018; Pons, et al., 1985) and is thought to be important for proprioception (Jennings et al. 1983; Kaas et al. 1979; London and Miller 2013). Interestingly, recent computational studies have shown that while neural activity may appear to be tuned to the state of the limb's endpoint, features of this tuning might be a direct consequence of the biomechanics of the limb (Chowdhury et al. 2017; Lillicrap and Scott 2013). Consistent with those results, we have recently observed, using artificial neural networks, that that muscle lengths were better predictors of S1 activity than were hand kinematics (Lucas et al. 2019). Thus, in this study, we set out to understand what information is represented in area 2 of S1, with two experiments that altered the relationship between hand and whole limb kinematics. Using both classic single neuron analysis techniques like tuning curves and preferred movement directions (Georgopoulos et al. 1982; Prud'homme and Kalaska 1994; Sergio and Kalaska 2003), as well as more recently developed neural population analyses (Churchland et al. 2012; Cunningham and Byron 2014), we discovered several features of neural activity that could not be explained by the classic hand-based model. However, using biomechanical modeling, we show that models relating neural activity to kinematics of the whole arm can explain these features. Our results indicate that if there is a transformation towards representing reaching in terms of only the hand, it likely occurs beyond S1.

Results

For the experiments detailed in this paper, we recorded neural signals from three monkeys (H, C, and L) using Utah multi-electrode arrays (Blackrock Microsystems) implanted in the arm representation within Brodmann's area 2 of S1 (Figure 3.1A). We trained each of these monkeys to grasp a two-link planar manipulandum and make reaching movements to targets presented on

a screen in front of them (Figure 3.1B). During these behavioral sessions, we also tracked the locations of markers on the monkey's arm using a custom motion tracking system and registered these markers to a musculoskeletal model in OpenSim (SimTK) to estimate joint angles and muscle lengths. Our experiments included two components: one comparing active and passive movements and one comparing reaching movements in two different workspaces.

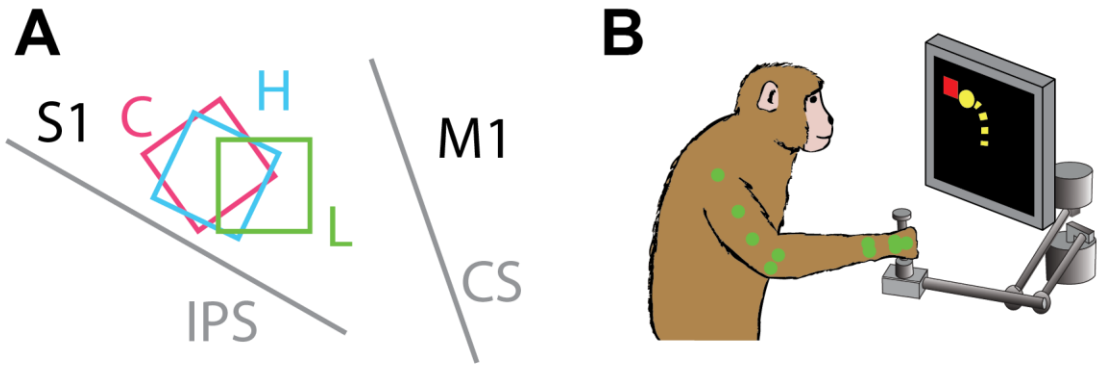

Figure 3.1 - General experimental setup. A –locations of Utah arrays implanted in primary somatosensory cortex (S1) of Monkeys C, H, and L. M1, primary motor cortex; IPS, intraparietal cortex; CS central sulcus. B – Behavioral task. Monkey controls a cursor on screen (yellow) with a two link manipulandum to reach to visually presented targets (red). We track the locations of markers on the monkey's arm (green) during the task.

Whole arm, rather than hand-based models, explain S1 representation of active and passive movements

We reported previously that the direction of maximal response (the "preferred direction"; PD) of single neurons in area 2 is typically similar during active reaching movements and passive perturbations of the hand (London and Miller 2013). However, this section shows that despite the apparent similarity in directional tuning of individual neurons, information from a population of neurons can be used to clearly distinguish active and passive movements, which is unexpected given our previous results. We set out to explore whether this separation could be explained by modeling the neural activity in area 2 purely in terms of behavioral variables.

In this experiment, the monkey performed a center-out reaching task to four targets. On half of these reaching trials, the monkey's hand was bumped by the manipulandum during the center-hold period in one of the four target directions (Figure 3.2). See Methods section or (London and Miller 2013) for task details. This experiment involved two sessions with each of Monkeys C and H.

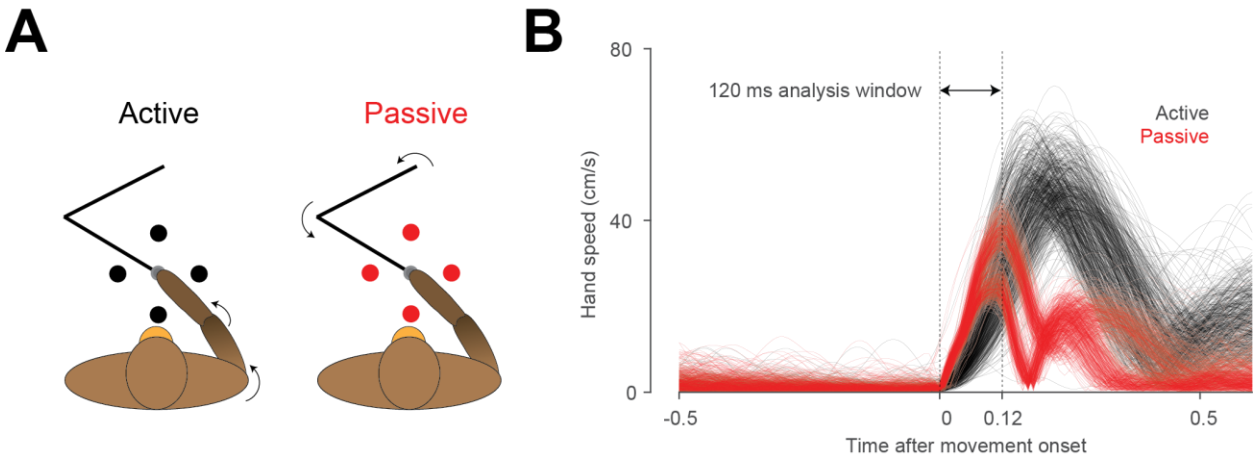

Figure 3.2 - Active vs. passive behavior. A – Schematic of task. On active trials (black), monkey reaches from center target to a visually presented target in one of four directions. On passive trials, manipulandum bumps monkey's hand in one of the four target directions. B – Speed of hand during active (black) and passive (red) trials, plotted against time, for one session. Starting around 120 ms after movement onset, a bimodal distribution in passive movement speed emerges. This bimodality reflects differences in the stiffness of the arm for different directions of movement. Perturbations towards and away from the body tended to result in a shorter overall movement than those to the left or right. However, average movement speed was similar between active and passive trials in the 120 ms after movement onset, the window used for neural analysis.

Of necessity, the similar active and passive tuning described by London and Miller included only neurons that had significant sinusoidal tuning to both active and passive movement. Consequently, it was based on fewer than 30% of the recorded cells. Similarly, in this study, only 29 of 57 recorded neurons were significantly tuned to both types of movements for Monkey C and only 12 of 83 for Monkey H (Figure 3.3). The low percentage of tuned neurons is at least partially due to measurement error (Stevenson, et al., 2011), exacerbated by the small amount of data available in the first 120 ms after movement onset when the kinematics of the two types of movements were similar. Thus, there are clear limitations inherent to single neuron tuning curve-based analyses.

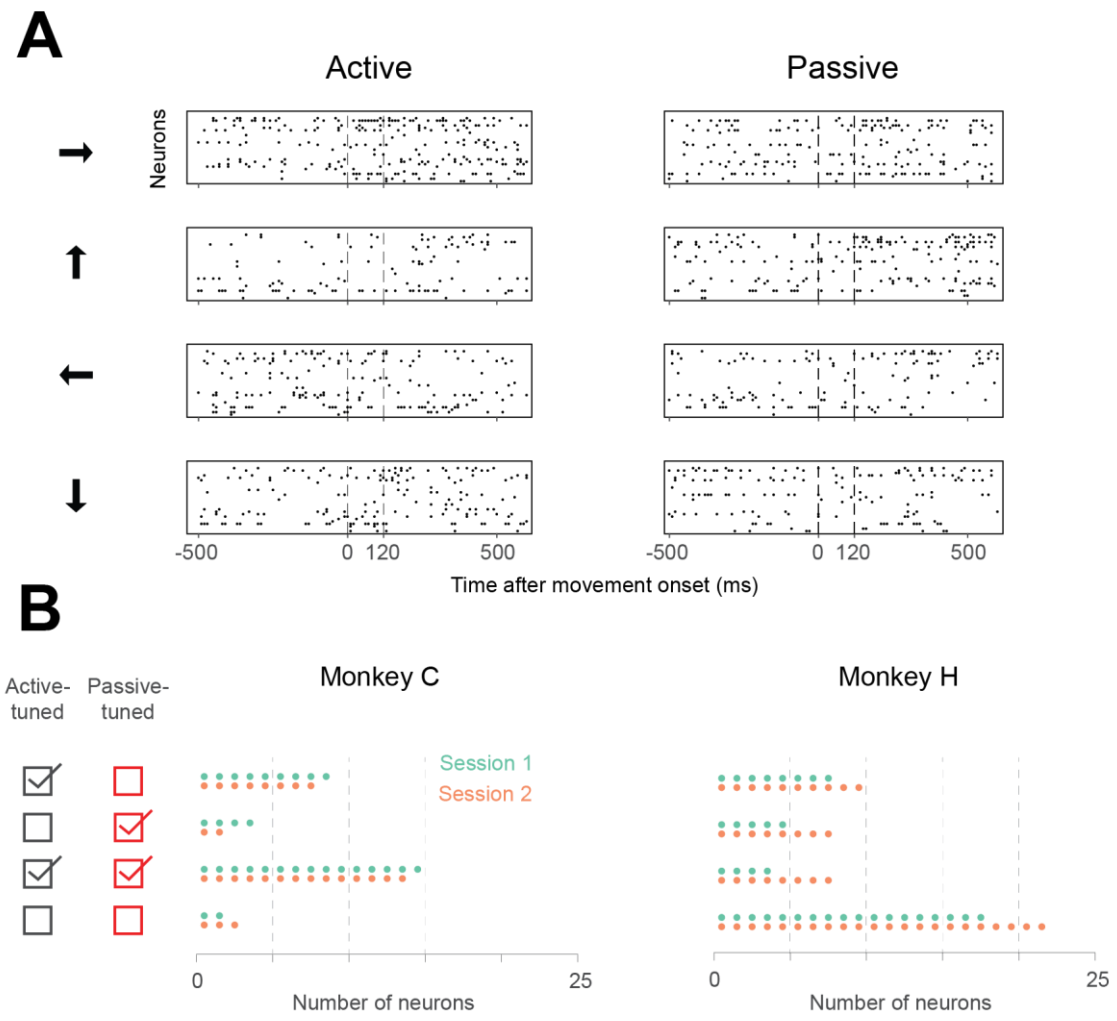

Figure 3.3 - Neural Activity during active vs. passive task. A – Neural raster plots for example active and passive trials in each movement direction (indicated by arrows to the left of each raster plot). In each plot, rows indicate spikes recorded from different neurons, plotted against time. Vertical dashed lines indicate boundaries of analysis window. B – Tuning properties for neurons. Dot histograms indicate counts of neurons tuned only to active movements, those tuned only to passive movements, those tuned to both, and those tuned to neither. Note that tuning estimates from Monkey H have more error due to the fewer trials that were typically available in a given session, leading to a higher proportion of apparently untuned neurons.

Because we had simultaneous recordings from many neurons, we could circumvent this limitation by analyzing the *population* responses to active and passive movements. Moreover, by aggregating information across neurons instead of averaging across trials, the neural population analysis allows the analysis of single trials. Recent studies have uncovered a host of neural coding properties that are difficult to infer from single neuron analyses, yet are revealed with single-trial, population-level analyses (Cunningham & Yu, 2014; Dekleva, et al., 2018; Perich, et al., 2018). In this case, our population analysis treats each active or passive trial as a point in a high-dimensional neural state space, where the value of each dimension corresponds to the activity of a single neuron. Using Principal Components Analysis (PCA), which extracts the sequential, orthogonal dimensions of highest covariance from the neural population activity, we uncovered an unexpected linear boundary between active and passive movements within the first three PCs (Figure 3.4). To quantify this separation, we trained a linear discriminant analysis (LDA) model to classify the movement type in each of the four sessions. On average, for data not used to train the LDA model, 84% of movements were correctly classified as active or passive, indicating that despite the similarity of hand-movement coding by single neurons in area 2, the neural population distinguishes the different movement types.

To explore potential causes of this separation in S1 neural state space, we attempted to replicate it using synthetic state spaces constructed from simplified neuron activity models that assume different representations in S1 (see Methods for modeling details). As a baseline, we fit a purely hand kinematics model, in which neural activity was driven by the position and velocity of hand movement. Classification rate of this simple model was slightly above chance but well under the actual data, with only 65% correct movement classification (Figure 3.4). Thinking that

the difference between active and passive movements may have been related to the differing forces on the hand in the two conditions, we also fit a hand kinematics+force model, like the one described in (Prud'homme & Kalaska, 1994), where neural activity was driven by both the velocity of the hand and the forces applied to it. Reaching only 64% correct classification, the addition of force made no difference, indicating that S1 contains some additional information that allows it to distinguish between active and passive movements.

Finally, we fit a model that represented the Cartesian positions and velocities of the elbow as well as the hand. Unlike the others, this model resulted in an average movement classification accuracy of 82% (Figure 3.4). This result provides a possible explanation of how S1 separates the two types of movement: while hand movement is similar, the movement of the whole arm differs between active and passive conditions. While there might be additional explanations for this separability, like an efference copy signal from motor areas specifically during active movements (Bell 1981; London and Miller 2013; Nelson 1987), this finding suggests that the parsimonious hand/elbow model is sufficient to explain it, unlike the hand-based models.

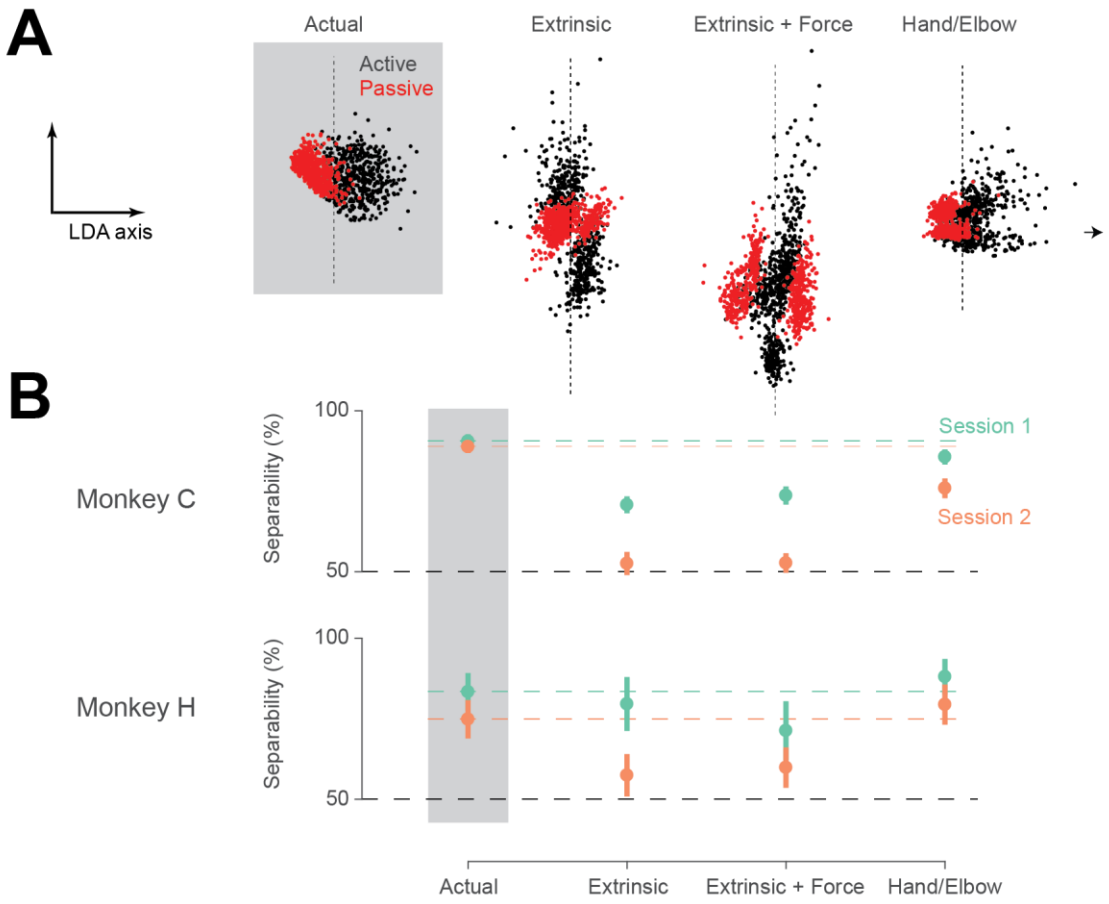

Figure 3.4 - S1 population separability. A – Population activity plots for an example session with Monkey C. Each dot in the scatter plots represents population neural activity during one trial, projected onto the LDA axis of highest separation between movement types (within the first three PCs of neural population activity). Vertical axis is an arbitrary axis orthogonal to the LDA axis in the three-dimensional PC space. Black dots indicate active trials, and red dots indicate passive trials. Leftmost panel shows actual population activity, while other three panels show population activity predicted by extrinsic, extrinsic + force, and hand/elbow models, from left to right. Black arrow in hand/elbow plot indicates vertical axis location of one active trial projected far to the right of the main cluster; this point was excluded from the figure for visual clarity of the main cluster. Note that actual population activity and activity predicted by the hand/elbow model separate active and passive movements, while activity predicted by the extrinsic and extrinsic + force models do not. B – Percent separability of active and passive movements for actual neural population activity and activity predicted by the three models. Error bars indicate 95% confidence intervals (derived from cross-validation – see Methods), and colored dashed lines indicates actual population separability for a given session. Black dashed line indicates chance separability. Extrinsic and extrinsic + force models generally yielded lower separability than the hand/elbow model, which predicts similar separability to that found in actual population activity.

Whole arm, rather than hand-based models, explain S1 representation of movements in two separate workspaces

Given the importance of whole-arm kinematics for explaining neural population activity revealed by the active vs. passive experiment, we conducted a second experiment to further examine how single neurons represent reaching movements that are constrained to two disjoint areas in front of the monkey, henceforth called workspaces (Figure 3.5A). We found that whole-arm models and hand-based models made different predictions of neural activity across these two conditions, allowing us to further characterize the extent to which S1 represented whole-arm movements rather than just the movements of the hand.

In this experiment, we tested four different models of S1 encoding, which we titled "extrinsic", "egocentric", "muscle", and "hand-elbow", schematized in Figure 3.5D. The extrinsic and egocentric models stem from classic, endpoint models of limb movement-related neural activity (Bosco and Poppele 2001; Bosco et al. 1996; Caminiti et al. 1990; Georgopoulos et al. 1982; Prud'homme and Kalaska 1994). Both represent the kinematics of the hand, but in different coordinate frames: the extrinsic model assumes neurons relate to Cartesian coordinates of hand position and velocity, while the egocentric model assumes neurons relate to spherical coordinates with respect to the shoulder. On the other hand, the muscle and hand-elbow models account more fully for how the whole arm moves. In the muscle model, neural activity was assumed to be driven by the length and stretch velocity of the muscles (see Methods for details). In the hand-elbow model, the activity was related to the Cartesian kinematics (position and velocity) of both the hand and the elbow markers. As in the active vs. passive experiment, we aimed to test how well these models predicted features of neural activity during reaching.

However, a challenge in comparing these models of neural representation is that for the typical, center-out reaching task in a small workspace, kinematic signals in the various coordinate frames are highly correlated. Because a high correlation means that a linear transform can accurately convert one coordinate frame into another, all four models make very similar predictions of neural activity. To deal with this problem, we trained the monkeys to reach to randomly generated targets presented in two different workspaces: one close to the body and contralateral to the reaching hand, and one distant and ipsilateral. This had two important effects. First, the random locations of the targets lessened the stereotopy of the movements, allowing for the collection of more varied movement data than from the center-out paradigm. Second, the average posture in each workspace was significantly different, such that while signals in the different model coordinate frames were still correlated within each workspace, this correlation (and the mapping between coordinate frames) changed significantly between workspaces. This forced the models to make different predictions of neural activity across the two workspaces. Figure 3.5A schematizes the two-workspace random target behavior, and Figure 3.5B shows an example raster plot of neurons while a monkey performed the task. We recorded three sessions with each of Monkeys C and H and two sessions with Monkey L.

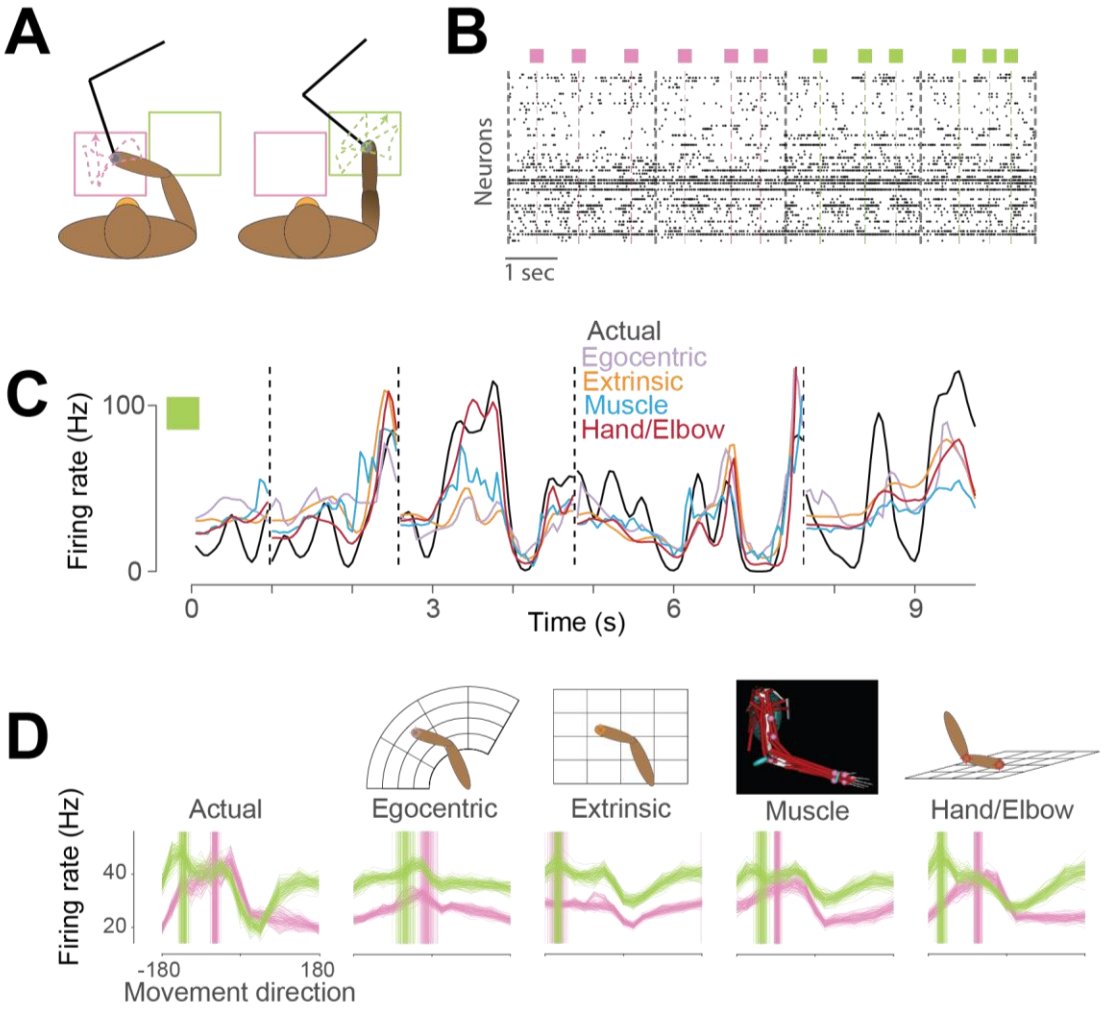

Figure 3.5 - Behavior and example neural activity for two-workspace task. A – Two-workspace behavior. On each trial, monkey makes reaches to random targets in one of two workspaces: one close to the body and contralateral to the reaching hand (pink) and the other distant and ipsilateral (green). B – Example neural raster plot from one session for two randomly drawn trials in each workspace. Dots in each row represent activity for one of the simultaneously recorded neurons. Black dashed lines indicate starts and ends of trials, and colored dashed lines and boxes indicate times of target presentation, with color indicating the workspace for the trial. C – Firing rate plot for an example neuron during five randomly drawn trials from the distal (green) workspace. Black trace represents neural firing rate, smoothed with a 50 ms gaussian kernel. Colored traces represent firing rates predicted by egocentric (purple), extrinsic (orange), muscle (blue), and hand/elbow (red) models (not smoothed). D – Actual and predicted tuning curves and preferred directions (PDs) for an example neuron in the two workspaces. Leftmost plot shows tuning curves and PDs for actual neural firing rate, while other plots show curves and PDs for activity predicted by each of the models. Each panel plots firing rate against direction of hand movement for each workspace. Vertical bars indicate preferred directions calculated in each workspace. Individual traces and vertical bars are tuning curves and preferred directions calculated during one cross-validation run. Note that both the actual tuning curve and PD change between workspaces. Neither the egocentric nor extrinsic models predict these changes, but the muscle and hand/elbow models can.

The simplest evaluation of these models is to compare how well they predicted actual neural firing rates (Figure 3.5C). To assess this, we used repeated k-fold cross-validation of a goodness-of-fit metric (see Methods for more details). Here, normal goodness-of-fit metrics like R^2 or variance-accounted-for (VAF) are ill-suited to the Poisson-like statistics of neural activity; instead, we used the likelihood-based pseudo-R^2 (Cameron and Windmeijer 1997; 1996; McFadden 1977). Like VAF, psuedo-R^2 has a maximum value of 1, but can also be negative for models that do worse than predicting the mean firing rate during cross-validation. In general, the values corresponding to what constitutes a good fit are lower for pR^2 than for either R^2 or VAF. Of our four models, the whole-arm models out-performed the hand models (Figure 3.6). Generally, the hand-elbow model was the most predictive of actual neural firing rate across the two workspaces, winning the great majority of the pairwise comparisons with the other models (Figure 3.6B, Figure 3.7). Of the 288 neurons recorded across the 8 sessions, in only 47 did either of the hand-based models win a pairwise comparison against either of the whole-arm models (using $\alpha = 0.05$ and a Bonferroni correction to account for six total pairwise comparisons for each neuron). For the other 241 cells, neither hand-based model could out-predict either whole-arm model.

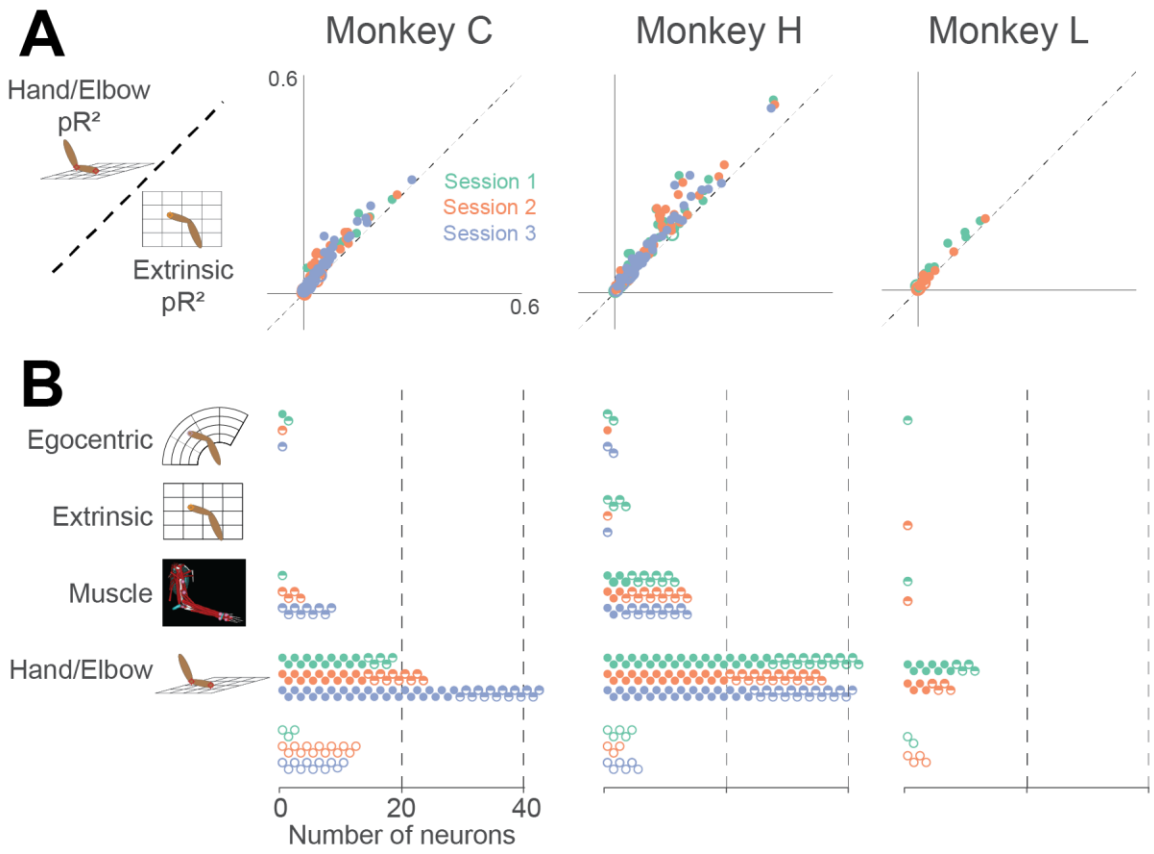

Figure 3.6 - Goodness-of-fit comparison analysis, with one column for each monkey. A – Scatter plots comparing the pseudo-R2 of the hand/elbow model to that of the extrinsic model for each monkey. Each point in the scatter plot represents the pseudo-R2 values of one neuron, with hand/elbow pseudo-R2 on the vertical axis and extrinsic pseudo-R2 on the horizontal. Different colors represent neurons recorded during different sessions. Filled circles represent neurons for which one model's pseudo-R2 was significantly higher than that of the other model. Conversely, open circles represent neurons for which the model pseudo-R2 values were not significantly different. In this comparison, most filled circles lie above the dashed unity line, indicating that the hand/elbow model performed better than the extrinsic model. All six pairwise comparisons for pseudo-R2 are shown in Figure 3.7. B – Dot plot of pairwise comparison winners. Each filled circle in this plot represents a neuron for which a model won all pairwise comparisons with the other three models. Each half-filled circle represents a neuron for which a model won all but one comparison. Empty circles, in the bottom group, correspond to neurons for which there was no model that won at least two of the comparisons. Color indicates the session in which a neuron was recorded. This figure shows that the hand/elbow model won at least two out of three pairwise comparisons for most neurons.

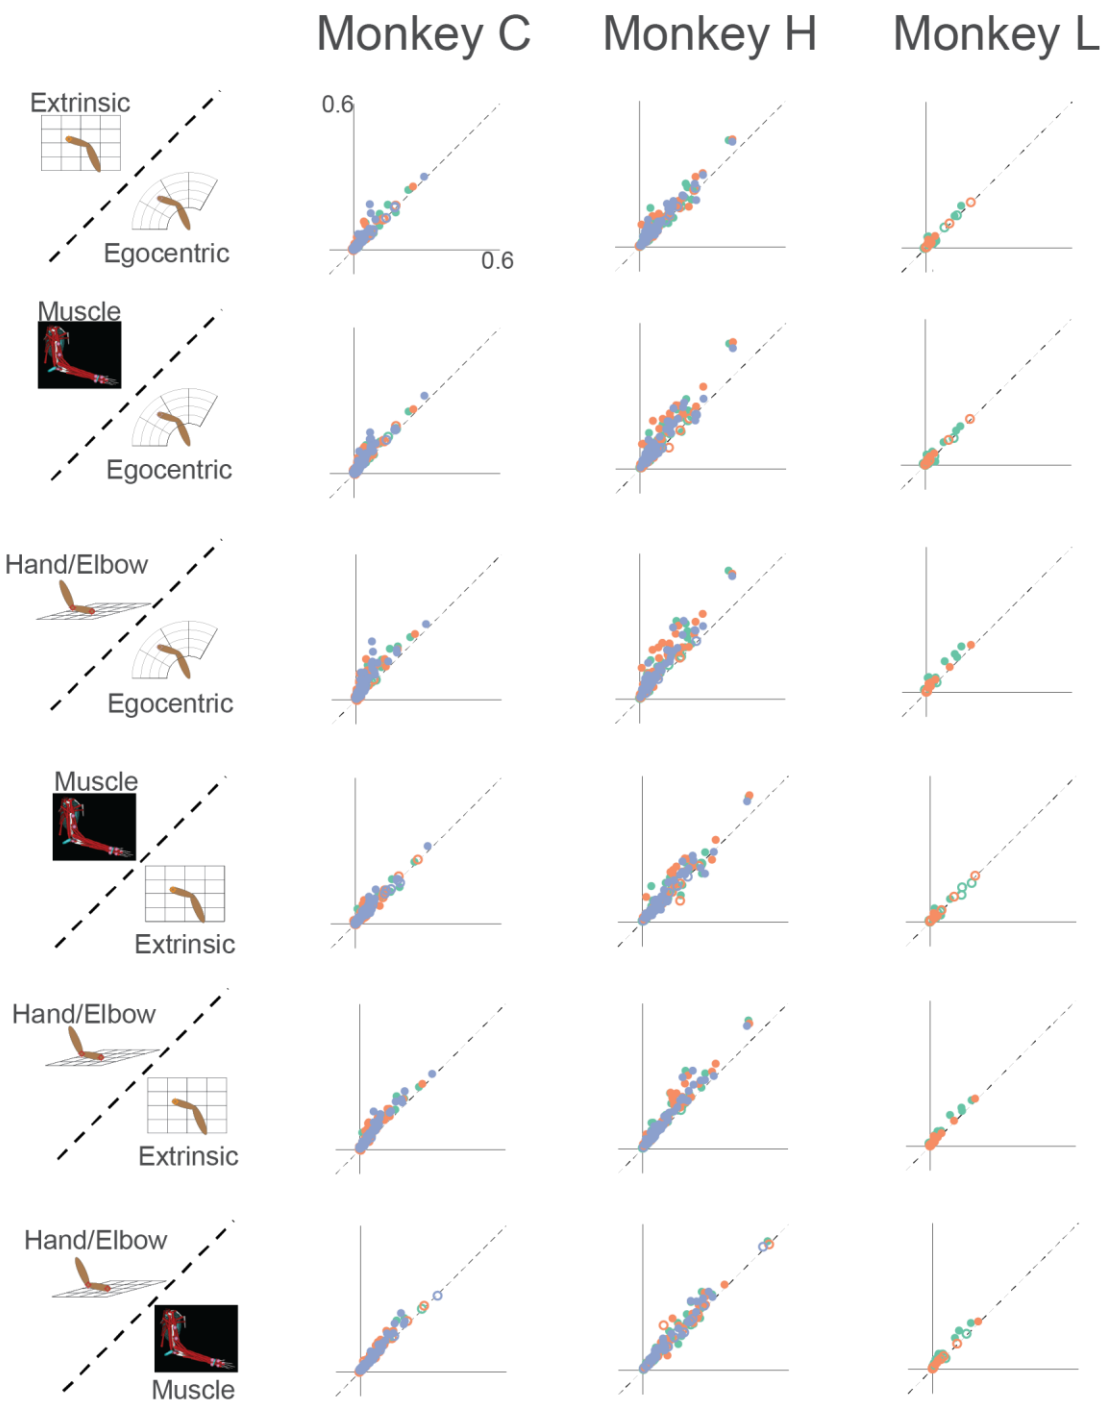

Figure 3.7 - All pairwise comparisons of model pseudo-R^2 values. Each column corresponds to one of the monkeys, and each row corresponds to the comparison between two models, indicated by the comparison label on the right. Dots in the scatter plots correspond to neurons, with pseudo-R2 values for one model plotted against those for the other model. Different colors correspond to different sessions, and filled circles correspond to neurons for which one of the two models is significantly better than the other one (i.e. significantly different from the dashed line of unity). Open circles correspond to neurons for which neither model is significantly better than the other one.

We further tested our models on how well features (e.g., tuning curves and PDs) computed from the model-predicted firing rates matched those of the actual data. Figure 3.5D shows the directional tuning curves for an example neuron, along with the tuning curves predicted by each model. We calculated the correlation between the predicted and the actual tuning curves in different workspaces as a measure of their similarity. With this measure, the hand-elbow model resulted in the best reconstruction of tuning curves, once again winning most of the pairwise comparisons with the other models (Figure 3.8, Figure 3.9). In this case, for only 37 neurons did either of the hand-based models win any pairwise comparison against a whole-arm model (using $\alpha = 0.05$ and a Bonferroni correction to account for six total pairwise comparisons for each neuron) For the other 251 cells, neither hand-based model could out-predict either whole-arm model. Therefore, this experiment further suggests that area 2 of S1 encodes whole-arm rather than just hand kinematics.

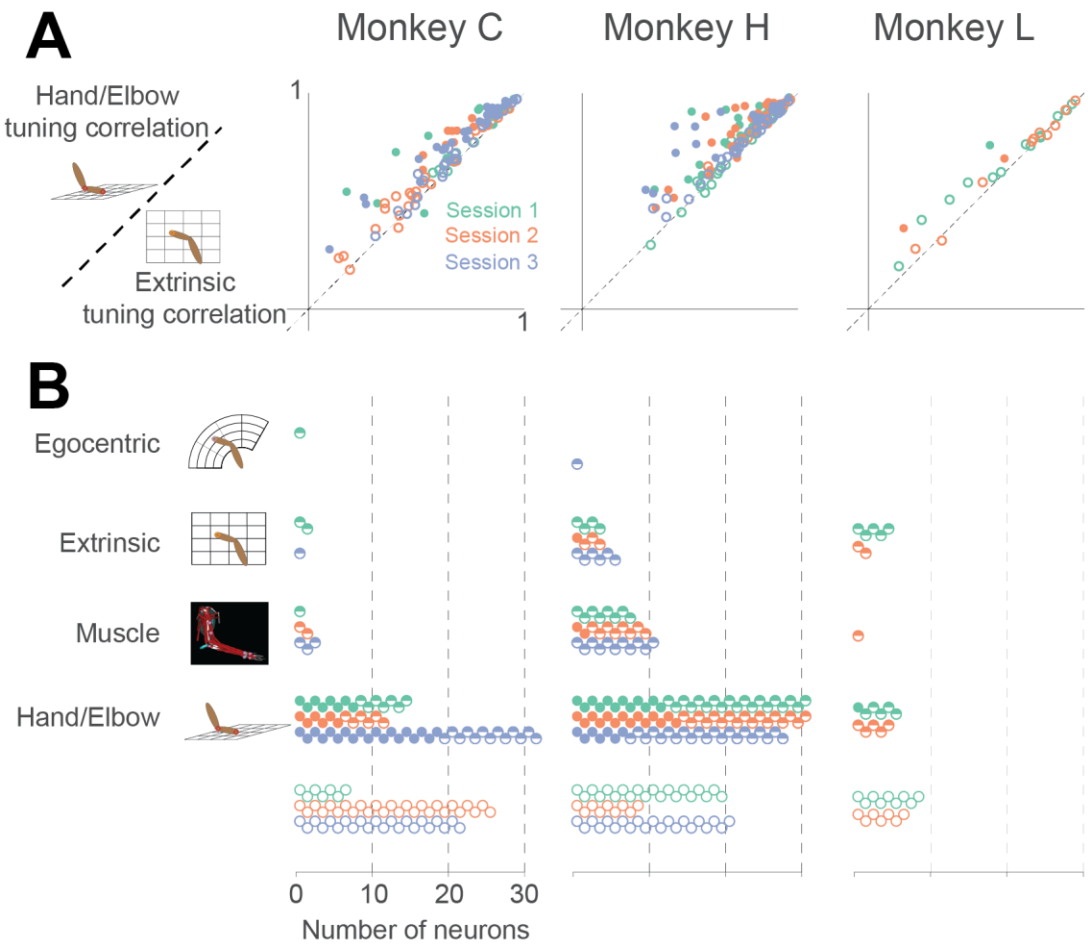

Figure 3.8 - Tuning curve shape correlation analysis, with one column for each monkey, in the same layout as Figure 3.6. A – Scatter plot for one of the six pairwise comparisons of tuning curve shape correlation. Once again, most filled circles lie above the dashed line of unity, indicating that the hand/elbow model was better at predicting tuning curve shape than the extrinsic model. All six pairwise comparisons for tuning curve correlation are shown in Figure 3.9. B – Dot plot of pairwise comparison winners. Compared to the pseudo-R^2 winner plot (Figure 3.6B), many more neurons fell into the 'no winner' category for this analysis, where there was no model that won two or more pairwise comparisons. Still, for most of the remaining neurons, the hand/elbow model appeared to predict tuning curve shape better than the other models.

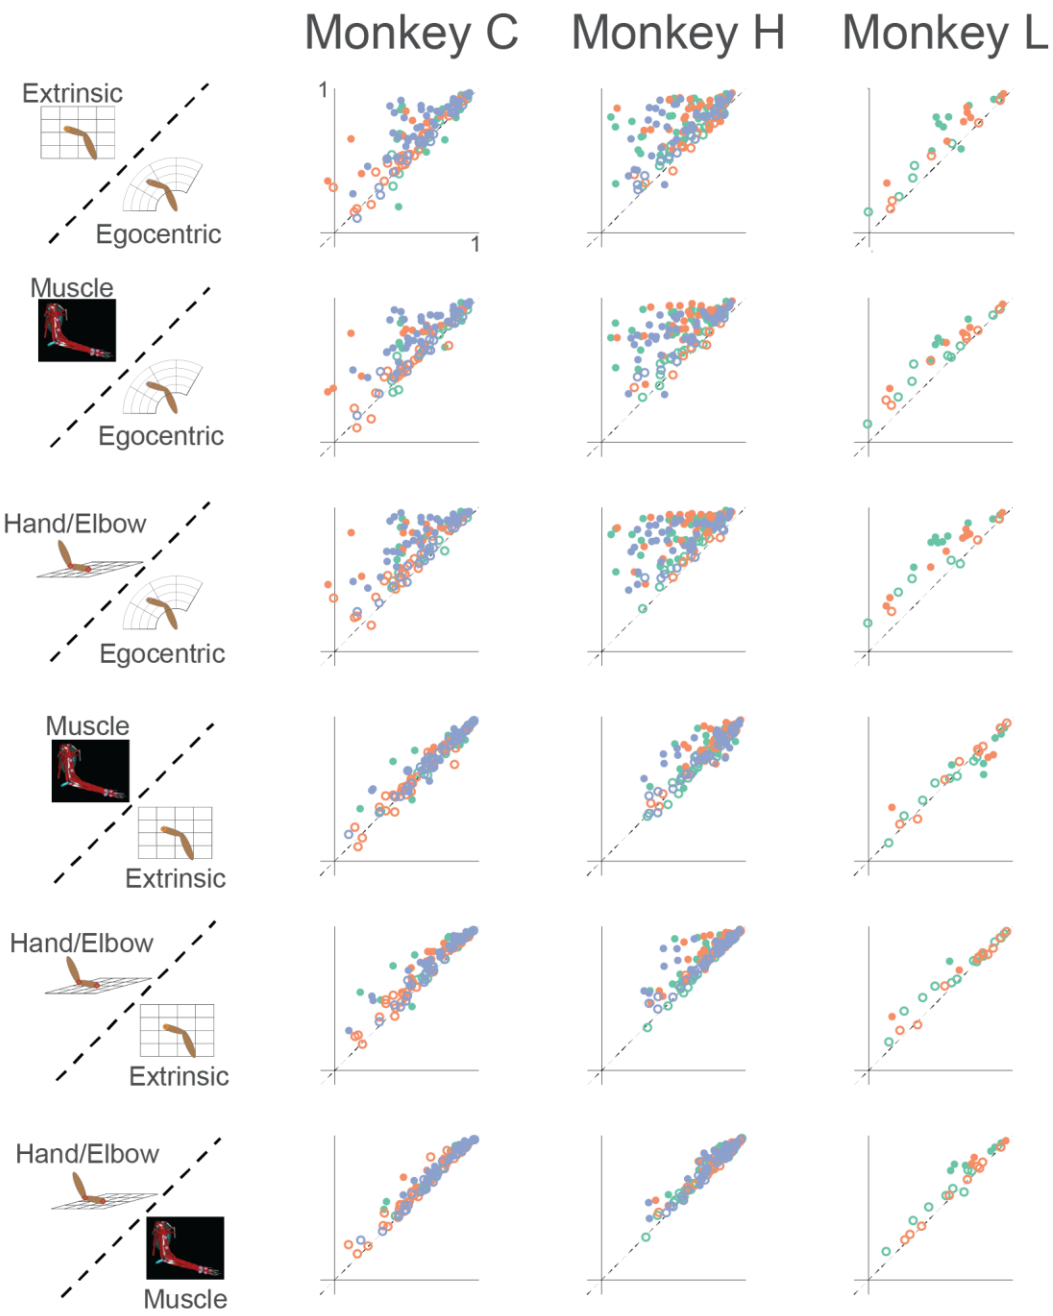

Figure 3.9 - All pairwise comparisons between models of tuning curve shape correlation. Same format as in Figure 3.7.

Of the 288 recorded neurons, 260 were significantly tuned to movement direction in both workspaces. This fraction of tuned neurons is much higher than the fraction we found in the active vs. passive task, likely for two reasons. First, unlike with the active vs. passive task, we used the entirety of each trial in the two-workspace experiment to calculate PDs. Second, monkeys tended to move faster during this random target task than during the center-out task, which likely led to a higher modulation depths and less measurement error, as predicted by (Stevenson, et al., 2011). Thus, in addition to the goodness-of-fit and tuning curve correlation analyses, we were also able to examine the properties of the neural PDs in the two workspaces.

An interesting feature of this task is that for many neurons, the PD of movement changed significantly between workspaces, exemplified by the difference between the vertical bars in the leftmost panel of Figure 3.5D. Figure 3.10A shows the actual PD shifts for these neurons plotted against the PD shifts predicted by each model. The large changes in PD, shown on the vertical axes of the scatter plots are a clue that the extrinsic model does not explain neural activity correctly; if it did, the preferred direction changes should have been insignificant (in principle, zero), as shown by the generally small extrinsic model-predicted changes (second column of Figure 3.10A). Additionally, and perhaps counterintuitively, the actual changes included both clockwise and counter-clockwise rotations, so it is also unlikely that they arose from a rotation of the extrinsic coordinate frame about the shoulder or the egocentric model. However, we found that the whole-arm models did predict both clockwise and counter-clockwise PD changes. Based on the circular VAF (cVAF) of the PD change prediction, Figure 3.10B shows that the hand-elbow model once again out-predicted the other models, with hand/elbow having the highest average cVAF (0.73), followed in order by muscle (0.63), extrinsic (0.55), and egocentric (0.38).

We made pairwise comparisons between model cVAF for each session. In every session but one, the hand-based models lost pairwise comparisons to at least one of the whole-arm models, and in the outlier session, no pairwise comparisons showed significance (again using $\alpha = 0.05$ and a Bonferroni correction to account for six comparisons for each session).

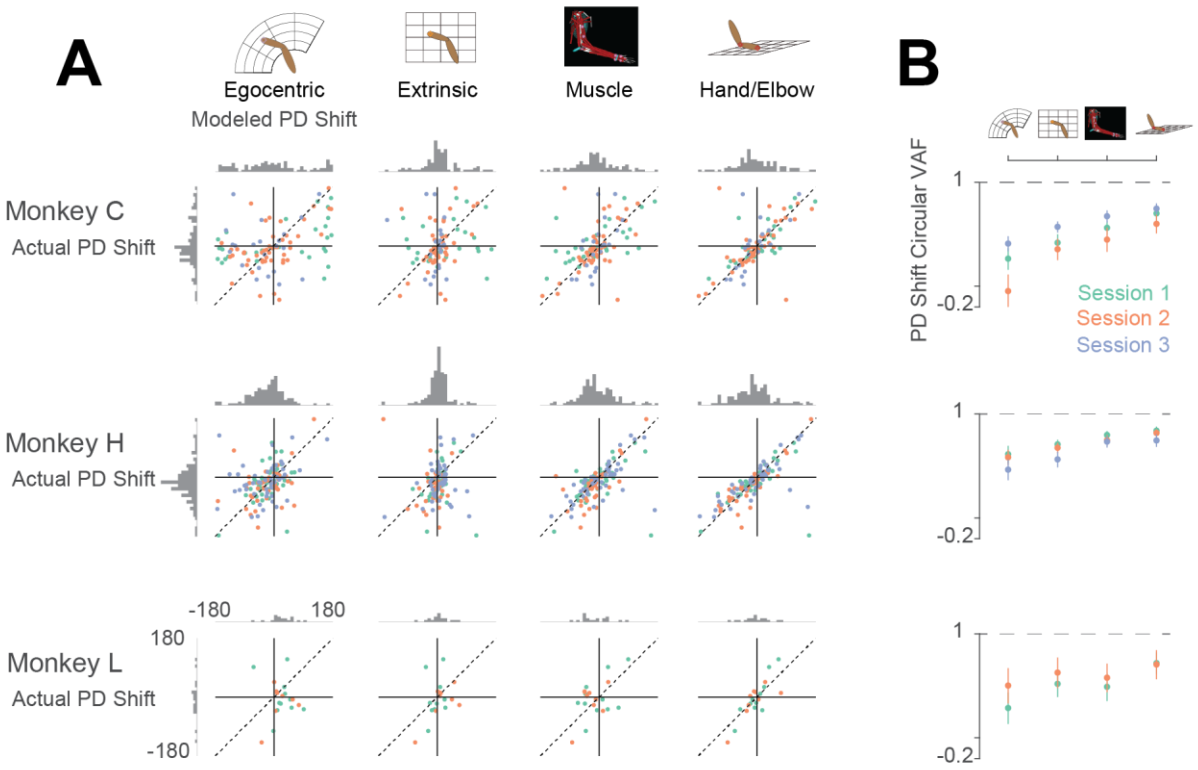

Figure 3.10 - Model predictions of PD shift. A – Scatter plots of actual PD shifts plotted against model-predicted shifts. Each dot represents the actual and modeled PD shifts of a single neuron, where different colors correspond to neurons recorded during different sessions. Dashed diagonal line shows ideal prediction. Vertical histograms indicate distributions of actual PD shifts for each monkey. Horizontal histograms indicate distributions of modeled shifts. Note that both horizontal and vertical axes are circular, meaning that opposing edges of the plot (top/bottom, left/right) are the same. Vertical histograms show that the distribution of actual PD shifts spanned both clockwise and counter-clockwise shifts. Clustering of scatter plot points on the diagonal line for the hand/elbow model indicates that it was most predictive of PD shift. B – plot showing circular VAF (cVAF) of scatter plots in A, an indicator of how clustered points are around the diagonal line (see Methods for details). Each point corresponds to the average cVAF for a model in a given session (indicated by color), and the horizontal dashed lines indicate the cVAF for perfect prediction. Error bars show 95% confidence intervals (derived from cross-validation – see Methods). Pairwise comparisons between model cVAFs showed that in no sessions did either of the hand-based models significantly out-predict either of the whole-arm models.

As a control for errors introduced into the muscle model by processing marker data with OpenSim, we also performed this cVAF analysis on a hand/elbow model where hand and elbow kinematics were derived from joint angles of the musculoskeletal model, rather than directly from the marker locations captured by the motion tracking system. We re-ran the model prediction analysis for only the muscle model, marker-derived hand/elbow model, and OpenSim-based hand/elbow model. Unsurprisingly, we found average cVAFs similar to those from the main analysis for the marker-derived hand/elbow model (0.75) and the muscle model (0.67). However, the average cVAF for the OpenSim-based hand/elbow model (0.67) dropped to that for the muscle model. This suggests that the difference in predictive capability between the muscle and hand/elbow models stems at least in part from errors introduced in OpenSim modeling, rather than from the hand/elbow model being the better model for S1 neural activity. Even so, as shown through the main goodness-of-fit (Figure 3.6, Figure 3.7), tuning curve correlation (Figure 3.8, Figure 3.9), and cVAF (Figure 3.10) analyses, both of these whole-arm models are better than the hand-based models, supporting the fact that S1 encodes whole arm kinematics.

Discussion

Summary

In this study, we explored how S1 represents reaching movements using two separate experiments. In the first experiment, we found that although single neuron directional tuning is largely preserved between active and passive movements, the two types of movements can be differentiated in the S1 neural state space. While the classic hand-centric models cannot explain this separation, we found that a model of S1 activity built from the kinematics of both the hand and elbow did. We further explored S1 encoding of whole-arm kinematics in a second

experiment, in which the monkey reached to targets in a pair of workspaces that changed the relationship between hand-based and whole-arm models. Across these two conditions, features of neural activity, including the dynamics of neural discharge, tuning curve shape and preferred direction, were better explained by a model incorporating the kinematics of the whole arm than by the classic hand-based models. Altogether, these results suggest that S1 represents movements in terms of not just the hand, but the whole arm.

Model complexity

A significant difference between the hand and whole-arm models is their number of parameters, which make the whole-arm models more complex and expressible. There are two concerns with testing models of differing complexity, the first dealing with model training and evaluation, and the second with interpretation of the results.

In training and evaluating our models, we had to make sure that the complex models did not overfit the data, resulting in artificially high performance on the training dataset but low generalizability to new data. However, because we found through cross-validation that the more complex models generalized to test data better than the simpler models, they were not overfitting. Consequently, the hand-based models are clearly impoverished compared to the whole-arm models.

The second concern is in interpreting what it means when the more complex models performed better. One interpretation is that this is an obvious result; if the added degrees of freedom have anything at all to do with S1 neural activity, then the more complex models should perform better. However, the choice of the two less complex hand-based models was not arbitrary. Both are classic models within the canon of the cortical representation of limb state

(Bosco et al. 2000; Prud'homme and Kalaska 1994). Additionally, as shown in the active/passive study, if the classic extrinsic model is to be believed, then the separation between neural representations of active and passive movement must have an explanation not rooted in the behavioral kinematics. Thus, the point of this study was to determine whether either of these now classic models adequately encapsulates what S1 represents. Because the whole-arm models outperformed the classic models, we conclude that the classic models are incomplete, and in the worst cases, potentially incorrect or misleading.

Influence of modeling errors

One surprising result from the two-workspace experiment was that the hand/elbow model, based on the addition of one arbitrary point on the proximal limb to the classic hand model, performed as well or better than the muscle-based model. As proprioceptive signals originate in the muscles, arising from muscle spindles and Golgi tendon organs, we expected to find that the muscle model would outperform the other models. However, there are several potential reasons why this was not so. The most important ones can be divided into two categories loosely tied to 1) errors in estimating the musclulotendon lengths, through motion tracking and musculoskeletal modeling, and 2) the fidelity of the muscle model to the actual signals sent by the proprioceptors.

In the first category, the main issue is that of error propagation. The extra stages of analysis required to compute musculotendon lengths (registering markers to a musculoskeletal model, performing inverse kinematics to find joint angles, and using modeled moment arms to estimate musculotendon lengths) introduce errors not present when simply using the positions of markers on the arm. As a control, we ran the hand/elbow model through two of these extra steps by computing the hand and elbow positions from the joint angles of the scaled model, estimated from inverse kinematics. The results of this analysis showed that the performance of the

hand/elbow model with added noise dropped to that of the muscle model, indicating that there are, in fact, errors introduced in even this portion of the processing chain.

The other potential source of error in this processing chain stems from the modeled moment arms, which might not accurately reflect those of the actual muscles. In developing their musculoskeletal model, Chan and Moran collected muscle origin and insertion point measurements from both cadaveric studies and existing literature (Chan & Moran, 2006). However, due to the complexity of some joints, along with ambiguity of how the muscle wraps around bones and other surfaces, determining moment arms purely by bone and muscle geometry is a difficult problem (An, et al., 1984). Because moment arms are irrelevant for determining hand and elbow kinematics, we could not subject the hand/elbow model to the error introduced by this step.

In addition to the questions of error propagation and musculoskeletal model accuracy is the question of whether our muscle model was truly representative of the signals sensed by the proprioceptors. The central complication is that spindles sense the state of the intrafusal fibers in which they reside, and have a complex, nonlinear relation to the musculotendon length that we used in our muscle model. Factors like load-dependent fiber pennation angle (Azizi et al. 2008), or tendon elasticity (Rack and Westbury 1984) can decouple muscle fiber length from musculotendon length. Additionally, intrafusal fibers receive motor drive from gamma motor neurons, which continuously alters muscle spindle sensitivity (Loeb et al. 1985; Prochazka and Wand 1981; Prochazka et al. 1976), and spindle activity also depends on the history of strain on the fibers (Haftel et al. 2004; Proske and Stuart 1985). Altogether, this means that while the musculotendon lengths we computed provide a reasonably good approximation of what the arm

is doing, they may not be a good representation of the spindle responses themselves. Spindle activity might be more accurately modeled when given enough information about the musculotendon physiology. However, to model the effects of gamma drive, we would either have to record directly from gamma motor neurons or make assumptions of how gamma drive changes over the course of reaching. In developing models of neural activity, one must carefully consider the tradeoff between increased model complexity and the extra error introduced by propagating through the additional requisite measurement and analysis steps. Given our data obtained by measuring the kinematics of the arm with motion tracking, it seems that the coordinate frame with which to best explain S1 neural activity is simply the one with the most information about the arm kinematics and the fewest steps in processing. However, this does not rule out the idea that S1 more nearly represents a different whole-arm model that may be less abstracted from physiology, like musculotendon length or muscle spindle activity.

Coordinate frame vs. informational content

Because of their differing dimensionality, the signals from hand-based models and those from whole-arm models do not have a one-to-one relationship: there are many different arm configurations that result in a given hand position. Thus, a comparison between the hand-based and whole-arm models is mainly a question of information content (do S1 neurons have information about more than just the hand?). In contrast, signals of the two whole-arm models do have a one-to-one (albeit nonlinear) relationship to each other. Knowledge of the hand and elbow position should completely determine the estimated musculotendon lengths, indicating that the two models have the same informational content. As such, a comparison between the muscle model and the hand/elbow is purely one of coordinate frame. While the interpretation for a comparison of information content is straightforward, interpreting the results of a comparison

between coordinate frames is not. One major issue is that these comparisons only make sense when using linear models to relate neural activity to behavior. Once nonlinear models are considered, as in our study with artificial neural networks (Lucas et al. 2019), coordinate frames with one-to-one correspondence become nearly equivalent, and much more difficult to compare meaningfully.

Clear parallels exist between this and earlier studies seeking to find a unique representation of movement in motor areas. Over the last few decades, a controversy involving the exact nature of the neural representation of movement has played itself out in the literature surrounding motor cortex, with some advocating a hand-based representation of motor control (Georgopoulos et al. 1982; Georgopoulos et al. 1986; Moran and Schwartz 1999) and others a muscle-based representation (Evarts 1968; Fetz et al. 1989; Morrow et al. 2007; Oby et al. 2012). Recently, the motor control field started turning away from questions of coordinate frame and towards questions of neural population dynamics and information processing (Churchland et al. 2010; Elsayed et al. 2016; Gallego et al. 2017; Kaufman et al. 2014; Perich et al. 2018; Russo et al. 2018; Sussillo et al. 2015). Part of the motivation for this pivot in viewpoint is that it became increasingly clear that a "pure" coordinate frame of movement representation is unlikely to exist (Fetz 1992; Kakei et al. 1999). Further, studies tended to use correlation between neural activity and behavioral variables as evidence that the neurons represent movements in a particular coordinate frame. However, as noted above, these correlations could often be explained by multiple coordinate frames, casting doubt on the conclusiveness of the exact coordinate frame of representation (Mussa-Ivaldi, 1988). Consequently, in our study, we put aside the question of the coordinate frame of S1, focusing instead on the fact that S1 contains information not included

in the classic model for proprioceptive neural activity during reaching; it relates to movements of the whole arm, rather than just the hand.

A major question this study leaves open is that of how information about reaching is processed by different areas of the proprioceptive neuraxis. While we might expect a muscle spindle-like representation at the level of the dorsal root ganglia (DRG) or the cuneate nucleus, removed from the receptors by one and two synapses, respectively, this representation likely changes as the signals propagate through thalamus and into S1. Even different areas of S1 may have different representations. Area 3a, which receives input mostly from muscle afferents (Heath et al. 1976; Kaas et al. 1979; Phillips et al. 1971; Yamada et al. 2016), seems more likely to retain a muscle-like representation than is area 2, which integrates muscle afferent input with that from cutaneous receptors (Hyvärinen & Poranen, 1978; Padberg, et al., 2018; Pons, et al., 1985). Likewise, area 5 may have an even higher-level representation, as it receives input from both somatosensory (Mountcastle et al. 1975) and motor cortices (Padberg et al. 2018), and appears to depend on attention (Chapman et al. 1984; Omrani et al. 2016). As it becomes increasingly feasible to record from several of these areas simultaneously (Richardson et al. 2016; Suresh et al. 2017; Weber et al. 2006), future experiments could examine how these areas project information to each other, as has been explored in motor and premotor cortices (Churchland et al. 2010; Elsayed et al. 2016; Kaufman et al. 2014; Perich et al. 2018), without modeling the more complex cortical areas explicitly in terms of particular behavioral variables "encoded" by single neurons.

Relevance for BCI

One motivation for this work is its potential to augment brain-computer interfaces (BCI) for restoring movement to persons with spinal cord injury or limb amputation. As BCI for motor

control gets more advanced (Collinger et al. 2013; Ethier et al. 2012; Kao et al. 2015; Young et al. 2018), it will become more necessary to develop a method to provide feedback about movements to the brain, potentially using intracortical microstimulation (ICMS) to activate somatosensory areas. While ICMS in S1 has seen some success in providing feedback about touch (Flesher et al. 2016; Romo et al. 1998; Salas et al. 2018; Tabot et al. 2013), the path towards providing proprioceptive feedback remains relatively unexplored. At least one study did use electrical stimulation in S1 for feedback during movement, using the stimulation to specify target direction with respect to the evolving hand position (Dadarlat et al. 2015). In that study, monkeys used the ICMS to reach to targets, even in the absence of visual feedback. However, target-location information is very different from the information normally encoded by S1, and the monkeys required several months to learn to use it. To our knowledge, no study has yet shown a way to use ICMS to provide more biomimetic proprioceptive feedback during reaching. Previously, our lab attempted to address this gap by stimulating a small number of electrodes in S1 based on neural activity recorded from them during normal reaching movements. In that experiment, the monkey reported the direction of a mechanical bump to his arm that occurred simultaneously with the ICMS. The ICMS biased one monkey's reports of the mechanical bump direction toward the PDs of the stimulated electrodes. Key to this finding was the fact that any bias in reporting actually decreased the reward rate, suggesting that the ICMS was indistinguishable from the perception of the bump itself (Tomlinson and Miller 2016). Unfortunately, the result could not be replicated in other monkeys; while the ICMS often biased their reports, the direction of the bias could not be explained by the PDs of the stimulated electrodes. One potential reason may be that the stimulation paradigm in those experiments was derived from the classic, hand-based model and the assumption that S1 represents active and

passive movements similarly. As this paper has shown, both of these assumptions have important caveats. It is possible that a stimulation paradigm based on a whole-arm model may be more successful, due to its greater accuracy at predicting neural PDs (Figure 3.10). It is also possible that the stimulus model would need to include information about forces in addition to kinematics. Regardless of the exact model, prospects for stimulating S1 to create natural proprioceptive sensations would likely improve given a more accurate generative model of S1 activity.

In addition to developing better models for S1 activity, it will be important to consider the implications of the difference between sensation for perception versus action. These two broad purposes for sensation are thought to involve distinct pathways in both vision and touch (Dijkerman and De Haan 2007; Mishkin and Ungerleider 1982; Sedda and Scarpina 2012). It is quite plausible that this distinction exists for proprioception as well (Dijkerman and De Haan 2007). However, studies of the effects of ICMS in S1 tend to use perceptual reporting to test the effect of stimulation (Salas et al. 2018; Tomlinson and Miller 2016; Zaaimi et al. 2013), thereby not directly addressing how effectively ICMS can be used as feedback for action. Even in the study conducted by Dadarlat et al., movements guided by ICMS were slower and contained more submovements that those guided by even a noisy visual signal, suggesting that monkeys used the ICMS as a learned sensory substitute, rather than as a biomimetic replacement for proprioception. As such, that study was also likely a cognitive one, engaging the perceptual stream rather than the action stream of proprioception (see (Deroy and Auvray 2012; Elli et al. 2014) for discussion of the limits of sensory substitution). As we better characterize how S1 represents movements, we hope to develop a stimulation paradigm in which we can engage both

streams, to enable users of a BCI both to perceive their limb, and to respond rapidly to movement perturbations.

Conclusion

This study began with an observation: the classic, hand-based cortical model of proprioception could not explain the separability of active and passive movements we observed in S1 neural state space. We found, however, that this feature could be explained by extending the classic model to include the kinematics of the whole arm. In a second experiment, we found that predictions of S1 neural activity from such whole-arm models generalized to different behavioral conditions better than those of the classic model. This suggests that even though our perception of our arm is typically centered on the hand, this area of S1 still appears to represent movement of the whole arm.

Methods and Materials

Behavior

We recorded data from a monkey while it used a manipulandum to reach for targets presented on a screen within a 20 cm x 20 cm workspace. After each successful reaching trial, the monkey received a pulse of juice or water as a reward. We recorded the position of the handle using encoders on the manipulandum joints. We also recorded the interaction forces between the monkey's hand and the handle using a six-axis load cell mounted underneath the handle.

For the active vs. passive experiment, we had the monkey perform a classic center-out (CO) reaching task, as described in (London and Miller 2013). Briefly, the monkey held in a target at the center of the full workspace for a random amount of time, after which one of four outer targets was presented. The trial ended in success once the monkey reached to the outer target. On 50% of the trials (deemed "passive" trials), during the center hold period, we used

motors on the manipulandum to deliver a 2 N perturbation to the monkey's hand in one of the four target directions. After the bump, the monkey returned to the center target, after which the trial proceeded like an active trial. From only the successful passive and active trials, we analyzed the first 120 ms after movement onset. Movement onset was determined by looking for the peak in handle acceleration either after the motor pulse (in the passive condition) or after 200 ms post-go cue (in the active condition) and sweeping backwards in time until the acceleration was less than 10% of the peak.

For the two-workspace experiment, we partitioned the full workspace into four 10cm x 10cm quadrants. Of these four quadrants, we chose the far ipsilateral one and the near contralateral one in which to compare neural representations of movement. Before each trial, we chose one of the two workspaces randomly, within which the monkey reached to a short sequence of targets randomly positioned in the workspace. For this experiment, we only analyzed the portion of data from the end of the center-hold period to the end of the trial.

Motion tracking

Before each reaching experiment, we painted 10 markers on the outside of the monkey's arm, marking bony landmarks and a few points in between, a la Chan and Moran 2008. Using a custom motion tracking system built from a Microsoft Kinect, we recorded the 3D locations of these markers with respect to the camera, synced in time to the other behavioral recordings. We then aligned the Kinect-measured marker locations to the lab frame by aligning location of the Kinect hand marker to the location of the handle in the manipulandum coordinate frame. Code for motion tracking can be found at https://github.com/limblab/KinectTracking.git.

Musculoskeletal modeling

We registered the Kinect marker locations to a monkey arm musculoskeletal model in OpenSim (SimTK), based on a model published by (Chan & Moran, 2006). After scaling the limb segments of the model to match those of each monkey, we used the inverse kinematics analysis tool provided by OpenSim to estimate the joint angles (and corresponding muscle lengths) required to match the model's virtual marker positions to the positions of the actual recorded markers. The OpenSim model we used can be found at https://github.com/limblab/monkeyArmModel.git.

Neural recordings

We implanted 100-electrode arrays (Blackrock Microsystems) into the arm representation of area 2 of S1 in these monkeys. For more details on surgical techniques, see (Weber et al. 2011). In surgery, we roughly mapped S1 by recording from the cortical surface while manipulating the arm and hand to localize their representations. To record neural data for our experiments, we used a Cerebus recording system (Blackrock). This recording system sampled signals from each of the 96 electrodes at 30 kHz. To conserve data storage space, the system detected spikes online using a threshold set at -5 x signal RMS, and only wrote to disk a time stamp and the 1.6 ms snippet of signal surrounding the threshold crossing. After data collection, we used Plexon Offline Sorter to manually sort these snippets into putative single units, using features like waveform shape and inter-spike interval. In addition to these recording sessions, we also occasionally performed sensory mapping sessions to identify the neural receptive fields by manipulating the monkey's arm while listening to neural activity. In all monkeys, we found a roughly equal mix of cutaneous and deep (muscle) receptive fields, suggesting that we were recording primarily from area 2 (Hyvärinen & Poranen, 1978; Padberg, et al., 2018; Pons, et al.,

1985; Seelke, et al., 2011). It is possible that the more posterior regions of the array were at least in transition regions of area 5.

Neural analysis

Code for the following neural analyses can be found at https://github.com/raeedcho/s1-kinematics.git.

Preferred directions

We used a simple bootstrapping procedure to calculate PDs for each neuron. On each bootstrap iteration, we randomly drew timepoints from the reaching data, making sure that the distribution of movement directions was uniform to mitigate the effects of any potential bias. Then, as in (Georgopoulos et al. 1982), we fit a cosine tuning function to the neural activity with respect to the movement direction, using equations 3.1a-b.

$$f_i(\tau) = b_0 + b_1 * \sin\big(\theta_m(\tau)\big) + b_2 * \cos(\theta_m(\tau)) \qquad (3.1a)$$

$$= b_0 + r_i * \cos(\theta_m(\tau) - PD_i) \qquad (3.1b)$$

where

$$PD_i = atan2(b_1, b_2) \text{ and } r_i = sqrt(b_1^2 + b_2^2)$$

Here, $f_i(\tau)$ is the average firing rate of neuron i for a given time point τ, and $\theta_m(\tau)$ is the corresponding movement direction, which for the active/passive task was the target or bump direction, and for the two-workspace experiment was the average movement direction over a time bin. We took the circular mean of PD_i and mean of r_i over all bootstrap iterations to determine the preferred direction and the modulation depth respectively, for each neuron.

As the PD analysis is meaningless for neurons that don't have a preferred direction of movement, we only analyzed the PDs of neurons that were significantly tuned. We assessed tuning through a separate bootstrapping procedure, described in (Dekleva et al. 2018). Briefly, we randomly sampled the timepoints from reaching data, again ensuring a uniform distribution of movement directions, but this time also randomly shuffled the corresponding neural activity. We calculated the r_i for this shuffled data on each bootstrap iteration, thereby creating a null distribution of modulation depths. We considered a neuron to be tuned if the true r_i was greater than the 95th percentile of the null distribution.

Models of neural activity

For the active/passive analyses, we averaged behavioral variables and neural firing rates over the 120 ms period following movement onset in each trial. For the two-workspace analyses, both behavioral variables and neural firing rate were averaged over 50 ms bins. We modeled neural activity with respect to the behavior using Poisson generalized linear models (outline in (Truccolo et al. 2005)) shown in equation 3.2a, below.

$$f \sim Poisson(\lambda), \lambda = \exp(X\beta) \qquad (3.2a)$$

In this equation, f is a T (number of time points) x N (number of neurons) matrix of average firing rates, X is a T x P (number of behavioral covariates, explained below) matrix of behavioral correlates, and β is a P x N matrix of model parameters. We fit these GLMs by finding maximum likelihood estimation of the parameters, $\hat{\beta}$. With these fitted models, we predicted firing rates ($\hat{f}$) on data not used for training, shown in equation 3.2b, below.

$$\hat{f} = \exp(X\hat{\beta}) \qquad (3.2b)$$

We tested five firing rate encoding models, detailed below. Note that each model also includes an offset term, increasing the number of parameters, P, by one.

- Extrinsic kinematics: behavioral covariates were position and velocity of the hand, estimated by using the location of one of the hand markers, in three-dimensional Cartesian space, with origin at the shoulder ($P = 7$).
- Extrinsic kinematics+force: behavioral covariates were position and velocity of the hand, as well as forces on the hand, in three-dimensional Cartesian space (P =10).
- Egocentric kinematics: behavior covariates were position and velocity of the hand marker in spherical coordinates (θ, ϕ, and ρ), with origin at the shoulder ($P = 7$).
- Hand/elbow kinematics: behavior covariates were position and velocity of both the hand and elbow markers in three-dimensional Cartesian space, with origin at the shoulder. This is the simplest extension of the extrinsic model that incorporates information about the configuration of the whole arm ($P = 13$)
- Muscle kinematics: behavioral covariates were derived from the length of the 39 modeled muscles (Chan and Moran 2008) and their time derivatives. However, because this would result in almost 78 (highly correlated) covariates, we used PCA to extract 5-dimensional orthogonal basis sets for both the lengths and their derivatives. On average, five components explained 99 and 96 percent of the total variance of lengths and length derivatives, respectively. Behavioral covariates of this model were the projections of the muscle variables into these spaces during behavior ($P = 11$).

We used repeated 5-fold cross-validation to evaluate our models of neural activity, given that the models had different numbers of parameters, P. On each repeat, we randomly split trials into five

groups (folds) and trained the models on four of them. We used these trained models to predict neural firing rates ($\widehat{f_i}$) in the fifth fold. We then compared the predicted firing rates from each model to the actual firing rates in that test fold, using analyses described in the following sections. This process (including random splitting) was repeated 20 times, resulting in n=100 sample size for each analysis result. Thus, if a more expressive model with more parameters performs better than a simpler model, it would suggest that the extra parameters do provide relevant information about the neural activity not accounted for by the simpler models.

Statistical tests and confidence intervals

To perform statistical tests on the output of repeated 5-fold cross-validation, we used a corrected resampled t-test, outlined in (Ernst 2017) and (Nadeau and Bengio 2003). Here, sample mean and variance are calculated as in a normal t-test, but a correction factor needs to be applied to the standard error, depending on the nature of the cross-validation. Equation 3.3a-c shows a general case of this correction for R repeats of K-fold cross-validation of some analysis result d_{kr}.

$$\hat{\mu}_d = \frac{1}{K \times R} \sum_{k=1}^{K} \sum_{r=1}^{R} d_{kr} \qquad (3.3a)$$

$$\hat{\sigma}_d^2 = \frac{1}{(K \times R) - 1} \sum_{k=1}^{K} \sum_{r=1}^{R} (d_{kr} - \widehat{\mu_d})^2 \qquad (3.3b)$$

$$t_{stat} = \frac{\hat{\mu}_d}{\sqrt{(\frac{1}{K \times R} + \frac{1/K}{1 - 1/K})\hat{\sigma}_d^2}} \qquad (3.3c)$$

We then compare the t-statistic here (t_{stat}) to a t-distribution with $K \times R - 1$ degrees of freedom. Note that the correction applied is an extra term (i.e., $\frac{1/K}{1-1/K}$) under the square root, compared to the typical standard error calculation.

Active/passive analyses

Estimating neural state-space separability

The neural response to movement, whether active or passive, can be represented as a single datapoint in a neural population space defined by the activity of all neurons in the relevant 120 ms period following movement onset. As described above, we estimated the separability of the active and passive movements in neural space using 20x repeated, 5-fold cross-validation. We did this in three steps. First, for each training set, we characterized the population response to each trial by finding the first three modes of neural activity using principal component analysis (PCA). We then projected the neural activity onto these three principal components and trained a linear discriminant analysis (LDA) model to find the axis of maximal separation between active and passive trials. Finally, we sequentially projected each test fold's neural data into the PC space and then onto the LDA axis. This resulted in a scalar value for each trial, with the sign indicating whether LDA classified the trial as active or passive. We took the average classification accuracy of the test fold's data as the percent separability for the fold, giving us 100 total samples from the 20x5-fold cross-validation. By averaging these samples, we estimated the overall neural separability of active and passive movements in a given session.

Estimating model-predicted separability

We also trained encoding models to predict neural firing from behavior (see equation 2a for procedure) using three different models of neural activity: the extrinsic kinematics, extrinsic kinematics+force, and hand/elbow kinematics models. See "Models of neural activity" section

for more details on the specific models. After training each model over the four training folds, we estimated firing rates in both the training and the test folds. Subsequent analysis mirrored that of the actual neural data: we found the three leading PCs and LDA axis of highest separability in the training folds and then sequentially projected the test-fold data through the PC space and onto the LDA axis. This resulted in 100 samples with which to estimate the model-predicted separability of active and passive movements.

Neural space dimensionality reduction

To visualize the population neural activity for figures, we used a combination of LDA and PCA. For the horizontal axis, we used LDA to find an axis in the three PCs of neural population space along which active and passive trials were most separated. For the vertical axis, we projected all activity onto the hyperplane orthogonal to the LDA axis and used PCA again to find the remaining axis of highest variance.

Two-workspace analyses

These analyses examined how well models of neural activity could predict neural activity as the monkey reached to targets in different workspaces. As such, we analyzed firing rate goodness-of-fit, along with how well the models could replicate the tuning curves and preferred directions (PDs) of neurons.

Goodness-of-fit

We evaluated goodness-of-fit of these models for each neuron by using a pseudo-R^2 (pR^2) metric. We used a formulation of pseudo-R^2 based on a comparison between the deviance of the full model and the deviance of a "null" model, i.e., a model that only predicts the overall mean firing rate (Heinzl and Mittlböck 2003; Perich et al. 2018) (Cameron and Windmeijer 1997; 1996).

$$pR^2 = 1 - \frac{D\left(f_i; \widehat{f_i}\right)}{D\left(f_i; \overline{f_i}\right)} \qquad (3.4a)$$

$$= 1 - \frac{logL(f_i) - logL\left(\widehat{f_i}\right)}{logL(f_i) - logL\left(\overline{f_i}\right)} \qquad (3.4b)$$

When computing the likelihood of a Poisson statistic, this is:

$$= 1 - \frac{\sum_{\tau=1}^{T} f_i\,(\tau) \log\left(\frac{f_i(\tau)}{\widehat{f_i}(\tau)}\right) - \left(f_i(\tau) - \widehat{f_i}(\tau)\right)}{\sum_{\tau=1}^{T} f_i\,(\tau) \log\left(\frac{f_i(\tau)}{\overline{f_i}}\right) - \left(f_i(\tau) - \overline{f_i}\right)} \qquad (3.4c)$$

This pR^2 metric ranges from $-\infty$ to 1, with a value of 1 corresponding to a perfectly fit model and a value of 0 corresponding to a model that only fits as well as the "null" model. In contrast with the general intuition for regular R^2, a pR^2 of ~0.2 is considered a "good" fit (McFadden 1977).

Tuning curves

We binned the trajectory into 16 bins, each 22.5 degrees wide, based on the mean direction across 50 ms of hand motion. For each directional bin, we calculated the sample mean and 95% confidence interval of the mean. In figures, we plotted this mean firing rate against the center-point of the bin.

Preferred direction shift

We calculated PDs for each neuron in each workspace and found the predicted change in PD from the contralateral workspace to the ipsilateral workspace, given each model. We compared these changes to those observed for each neuron. The values of these PD shifts are shown in

Figure 3.10 for all neurons tuned to movements in both workspaces, averaged over all 100 test folds.

We computed a variance-accounted-for (VAF) metric, here called the "circular VAF" (cVAF) for each neuron (i) in each fold as:

$$cVAF_i = cos(\Delta\theta_{PD,i} - \Delta\hat{\theta}_{PD,i}) \qquad (3.5)$$

As the cVAF metric is essentially the inner product of unit vectors with direction $\Delta\theta_{PD,i}$ and $\Delta\hat{\theta}_{PD,i}$, it accounts for the circular domain of the PD shifts. Like regular VAF, the cVAF has a maximum value of 1 when $\Delta\theta_{PD,i}$ and $\Delta\hat{\theta}_{PD,i}$ are the same, and decreases in proportion to the squared difference between $\Delta\theta_{PD,i}$ and $\Delta\hat{\theta}_{PD,i}$. We took the average cVAF over all neurons as the cVAF for the fold. In total, given the 20 repeats of 5-fold cross-validation, this gave us 100-samples of the cVAF for each model in a given session.

Chapter 4 - Discussion

Chapters 2 and 3 described my work towards examining the representations of movement in the dorsal spinocerebellar tract (DSCT) and in primary somatosensory cortex (S1), respectively. Classic models suggested that neurons in both of these areas represent the state of the limb's endpoint, implying a neural computation that integrates muscle information to form this endpoint representation. However, in Chapter 2, I showed in a simulation study that what looks like an endpoint representation can arise simply from biomechanics and convergence of information from muscles, suggesting that the evidence for neural computation towards a true endpoint representation at the level of DSCT is inconclusive. Chapter 3 built on this observation to examine the validity of the hand-based model of representation in S1. In that chapter, I explored the limits of the hand-based model with two separate experiments. In the first, I found that the hand-based model could not explain a distinct difference in population neural activity between active and passive movements. In the second experiment, I also found that the hand-based model could not explain differences in neural activity between reaches in two different workspaces. In both of these cases, however, a model of neural representation that incorporates kinematics of the whole arm did explain these differences, suggesting that S1 represents more than just the state of the hand. Taken together, these two chapters reveal that if a transformation to a pure endpoint representation occurs somewhere in the proprioceptive information processing chain, then it must take place after S1.

The questions addressed in this dissertation bear a resemblance to a controversy in the motor cortex literature over the last few decades. As discussed in Chapter 1, a debate emerged,

starting in the 1980s, about which coordinate frame primary motor cortex (M1) represented movement in, with some studies on the side of a hand-based representation (Georgopoulos, et al., 1982; Georgopoulos, et al., 1986; Moran & Schwartz, 1999) and others on the side of a muscle-based representation (Evarts, 1968; Fetz, et al., 1989; Morrow, et al., 2007; Oby, et al., 2012). This decades-long controversy has still not been satisfactorily resolved, though it now seems unlikely that motor cortex represents movement in any one "pure" coordinate frame (Fetz, 1992; Kakei, et al., 1999). Consequently, recent studies have turned away from trying to find that coordinate frame, examining instead how neurons communicate with each other and with muscles to generate movements. One recent example of this shift in approach is in studies of how activity in motor cortex changes during motor adaptation to a novel force field. One study, examining M1 using a hand-based model of movement representation, suggested that force field adaptation involved changes in single-neural tuning (Rokni, et al., 2007). However, more recent studies have examined, not neural movement tuning, but communication between dorsal premotor cortex (PMd) and M1. In doing so, they showed that the observed tuning changes were likely the indirect result of changes in how PMd neurons recruited M1 neurons (Perich & Miller, 2017; Perich, et al., 2018).

This change in approach within the motor literature provides an important lesson for studying proprioception. Specifically, trying to determine the exact coordinate frame of movement representation for single somatosensory neurons is unlikely to be fruitful. Instead, this chapter will examine what the results presented in this dissertation imply about how proprioceptive areas of the nervous system process information and contribute to the control of movement. In doing so, this chapter will also lay out ideas on future work to expand on these

results, considering especially the rapidly advancing technology allowing researchers to record from hundreds to thousands of neurons in multiple brain areas simultaneously (Stevenson, et al., 2011; Jun, et al., 2017).

In the context of proprioceptive processing, there are two principle perspectives from which to examine and build on the results presented in this dissertation. Most applicable is the hierarchical sensory processing perspective, which is concerned with how proprioceptive information is processed as it moves from the periphery to S1. The second perspective is that of motor control, which is concerned with how S1 communicates with visual and motor areas to plan and generate movements.

Perspective 1: How are signals transformed on the way to S1?

Chapter 1 described broadly the path taken by proprioceptive information as it travels from peripheral receptors to S1. In short, muscle spindles and Golgi tendon organs transduce muscle length and force respectively, into action potentials. Neurons in the dorsal root ganglia (DRG) transmit this muscle receptor activity to brainstem nuclei (Rosén, 1969; Rosén, 1969). These then project information to the thalamus, which projects to S1 (Rosén, 1969; Oscarsson & Rosén, 1963; Padberg, et al., 2009). S1 itself has two major proprioceptive areas, named areas 3a and 2 (Jennings, et al., 1983; Kaas, et al., 1979; London & Miller, 2013). The main goals of my doctoral work were to explore how proprioceptive information may be processed by the time it gets to S1, and to understand what form the neural representation of movement takes.

An important aspect of exploring the proprioceptive representation of movement is to understand the contributions of peripheral mechanics, including both biomechanics and muscle receptor physiology, to proprioceptive neural activity. This is made particularly clear by the

results of the studies presented in Chapters 2 and 3, which specifically examined the biomechanics of the limb. Both the empirical finding that DSCT neurons appear tuned to endpoint and the finding that active and passive movements have distinct representations in S1 neural activity at first suggested explicit neural computations. In the first case, neurons in the DSCT appeared to compute the location of the hindlimb paw, despite the addition of a constraint that changed how joint angles related to paw position (Bosco, et al., 2000). In the second case, because the classic hand-based model could not predict the separation between active and passive movement representations, we initially hypothesized that S1 integrated efference copy information from motor areas to distinguish the movements. However, in both cases, we found that models of neural activity built on kinematics from the whole limb, rather than just the hand, were sufficient to explain these features, suggesting that they were direct consequences of biomechanics and simple receptor properties, rather than explicit neural computation.

Moreover, the preceding chapters show that even after proprioceptive signals reach area 2, neural activity can still be predicted well by a linear convergence of muscle-like signals. Why is this the case, if the signals have been processed by several sensory areas along the way to area 2? One potential answer is that at each level of processing, neurons in the current stage simply integrate information from many neurons of the previous stage, effectively creating more complex response properties as the signals move up a simple hierarchical system. This idea of hierarchical processing was first described for the visual system, to explain how features like edge detection and orientation tuning might develop from a spatial integration of photoreceptor responses (Hubel & Wiesel, 1959; Hubel & Wiesel, 1962; Felleman & Essen, 1991). This became the inspiration for the design of deep convolutional artificial neural networks, which

have since become the state of the art in machine learning for image classification (Krizhevsky, et al., 2012). Unlike previous machine learning methods for image recognition, these feedforward neural networks are not designed to extract specific, human-defined features of images. Instead, they rely on intermediate layers to integrate information from earlier layers in particular spatial combinations, building up a library of neurons that respond to different features in the input. From a different perspective, the library of feature detectors in one intermediate layer can be read out by the next layer in various ways to build up more complex feature detectors. Eventually, in the final layers, neurons become object detectors. In the proprioceptive system, such integration, without explicit transformation to some intermediate movement representation, might allow neurons in S1 to serve as a general-purpose library of proprioceptive feedback features, whose activity is read out in different ways for either perception or use in motor control.

Surprisingly, the studies presented in the preceding chapters also revealed that even though muscle-based models could explain neural activity well, joint-based models could not. The study in Chapter 2 showed that despite an experimental manipulation that altered how joint angles relate to endpoint position, a single model of neural activity based on musculotendon lengths still produced neurons with an apparent endpoint representation. To further compare muscle and joint models, I trained a joint-based model of area 2 neural activity during the two-workspace task, in an analysis not included in Chapter 3. Aligned with the result from Chapter 2, I found that this joint-based model performed worse than the muscle and hand/elbow models at predicting the shifts in neural preferred direction between the two workspaces. Together, these analyses suggest that even though musculotendon lengths and joint angles are closely related,

joint angles are a poorer model for proprioceptive neural activity. This may be unsurprising, given that proprioception arises mostly from muscle sensors, and that the relationship between musculotendon lengths and joint angles, while one-to-one, is fairly nonlinear.

More surprising is the fact that in both of these studies, models involving the locations of points on the limb (e.g., the hand/elbow model) also outperform models involving joints. This suggests that the relationship between musculotendon length and the kinematics of points on the limb is more nearly linear over the range of normal reaching movements than the one between musculotendon lengths and joint angles. This may be because calculating musculotendon lengths from joint angles is in many cases equivalent to calculating the distance between two points on the limb. Interestingly, this suggests a potential reason behind the psychophysical results showing that people are worse at estimating joint angles than hand location (Fuentes & Bastian, 2010) or the angles of limb segments with respect to the direction of gravity (Soechting & Ross, 1984). Given the conclusions of the studies in this dissertation, these psychophysical results may arise simply from the fact that joint angle is not well represented at any level of the proprioceptive system.

Future experiments

These findings specifically highlight the importance of considering the biomechanics of movement when studying proprioception. However, in examining the question of how muscle geometry contributes to the neural representation of movement, I abstracted away many intermediate structures, simply modeling neural activity in the DSCT and S1 in terms of musculotendon lengths and their time derivatives. This presents three future directions in which to explore the hierarchical processing of proprioceptive information. The first and second directions involve considering how muscle forces and muscle receptor physiology, respectively,

might augment the findings presented in this dissertation. The third future direction involves using new experimental and analytical techniques to explore how different areas of the proprioceptive system communicate and process information on the way to cortex.

Modeling muscle force contributions to proprioceptive neural activity

One aspect of proprioception left unexplored by the studies presented in Chapters 2 and 3 is the representation of forces on the limb, which have been shown to play an important role in S1 neural activity (Prud'homme & Kalaska, 1994; London, et al., 2011). As with kinematics, the studies investigating the effects of limb forces on S1 activity assume a hand-based model. In this model, neural activity is cosine-tuned to the direction of the load on the hand (Prud'homme & Kalaska, 1994). Given that I found muscle kinematics to be an important component of the limb-state representation in proprioceptive areas, an obvious follow-up study would be to examine how muscle forces contribute to neural activity in these areas.

Unfortunately, measuring muscle force during behavior is non-trivial. Unlike musculotendon lengths, which are unique for a given posture or movement of the arm, there are many different sets of muscle forces that result in the same joint torques and forces on the hand. This issue stems from the fact that the musculature of the limb is highly redundant, with many muscles having similar pulling directions. As such, muscle forces are impossible to model without making assumptions to resolve the redundancy.

One of the few computationally tractable approaches to this problem is called Computed Muscle Control (CMC), which attempts to activate the muscles of a musculoskeletal model in a forward dynamics simulation to match a given kinematic trajectory (Thelen, et al., 2003). While modeling these muscle forces, CMC constrains the redundancy problem by optimizing for

minimum total muscle activation or minimum fatigue. As it is unclear whether these minimum energy assumptions are completely valid (Kistemaker, et al., 2010), the CMC algorithm can use electromyograms (EMG) to further constrain the solution space of possible muscle forces.

Another option to constrain the problem of muscle force redundancy is to measure the forces from some of the muscles. Over the last few decades, implantable transducers have been developed to measure the strain on tendons and ligaments in vivo (see (Ravary, et al., 2004) for review). Our lab has successfully used implanted strain transducers to measure leg muscle force in rabbits, but we have yet to use them successfully in monkeys. Once we develop the technology further, we can use direct force measurement of a few large muscles, in combination with EMG recordings and CMC, to more accurately estimate arm muscle forces during reaching.

Once muscle forces can be estimated, we could explore their contributions to proprioceptive neural activity. An example experiment could take the same form of those presented in Chapter 3, examining the generalizability of different models of proprioceptive neural activity. However, instead of using different kinematic conditions, as in Chapter 3, this hypothetical experiment could use multiple dynamic conditions, like reaches through force fields of various strengths, or comparing fast and slow movements of varying curvature.

Exploring how muscle receptor physiology contributes to proprioceptive neural activity

The studies presented in this dissertation modeled muscle receptor activity simply as the length of the musculotendon unit, along with its time derivative. However, as described in Chapter 1, the relationship between spindle activity and musculotendon length is highly nonlinear, depending on several factors, like muscle pennation (Azizi, et al., 2008), tendon elasticity (Rack & Westbury, 1984), gamma drive (Loeb, et al., 1985; Prochazka, 1981; Prochazka, et al., 1976),

and even the history of muscle stretch (Proske & Stuart, 1985; Haftel, et al., 2004). As such, one potential extension to the work presented in this dissertation is to estimate muscle spindle activity, using one of the many published muscle spindle models (Hasan, 1983; Lin & Crago, 2002; Mileusnic, et al., 2006). While these models would add several factors that may be difficult to infer, experiments could examine how important each of these spindle model nonlinearities are to the neural activity in various proprioceptive areas. Given that a muscle length model performed well at predicting S1 activity, it may even be that higher proprioceptive areas compensate for these nonlinearities in order to provide more veridical feedback of where the limb is.

Recording from different stages of the proprioceptive system

One limitation of the muscle force and receptor activity estimation approaches is their reliance on modeling peripheral signals. A different approach to the question of how proprioceptive signals are processed is to circumvent the peripheral modeling and record directly from early areas of the proprioceptive system, like DRG or cuneate nucleus. With simultaneously implanted arrays in combinations of DRG, cuneate, and S1, it would be possible to build models of neural activity at one stage of processing from the neural activity in a different stage. The specific structure of these resulting models would likely shed light on how information is communicated between the two areas.

A good example of this approach comes from a study of how neural activity associated with preparing a movement causes no muscle activation (Kaufman, et al., 2014). By modeling recorded muscle activity as a linear combination of M1 and PMd neural activity, this study found that the predictions of muscle activity during the preparatory period canceled out, even though neural activity in M1 and PMd was non-zero. This effectively showed that the neural activity of

the population could be decomposed into “output-potent” dimensions, in which neural activity would result in muscle activity, and “output-null” dimensions, along which neural activity did not affect muscle activity (Kaufman, et al., 2014). Similar approaches have been used to examine how PMd communicates with M1 during motor adaptation (Perich, et al., 2018) and to investigate how primary visual cortex projects information to secondary visual cortex (Semedo, et al., 2019).

One complication for this approach is the difficulty in obtaining chronic recordings from early proprioceptive areas, compared to recording from area 2 of S1. However, this is beginning to change; in a collaboration with Dr. Sliman Bensmaia’s lab at the University of Chicago, our lab has recently developed a method to record chronically from the cuneate nucleus of monkeys (Semedo, et al., 2019). With this technique, we have successfully implanted four monkeys with chronic cuneate arrays, allowing us to record from as many as 80 neurons simultaneously during reaching behavior. We have also started a collaboration with Dr. Doug Weber at the University of Pittsburgh to develop methods for implanting chronic electrode arrays in monkey cervical DRG. Recording from combinations of these proprioceptive areas simultaneously will allow us to directly examine the hierarchy of proprioceptive processing, from muscle receptors to cortex.

Perspective 2: Role of proprioception in motor control

Generating movements requires a coordination between sensory and motor areas. One way of examining this coordination is through the internal model framework described in Chapter 1. In this framework, during movement, internal models integrate sensory feedback with a copy of the motor command (called an efference copy signal) both to estimate the state of the limb, and to update the control policy in the case of a new dynamic environment (Wolpert, et al., 1995;

Wolpert, et al., 1998). Thus, one potential reason for keeping a more muscle-like representation of proprioceptive feedback in S1 is for easy communication of sensory information with motor areas, like M1. As described in the introduction, despite the "motor encoding" controversy in the literature, it is clear that a large population of M1 neurons project directly to spinal motor neurons (Rathelot & Strick, 2009), and therefore relate directly to muscle activity. Furthermore, several other studies including more than just these corticomotor neurons suggest that information about muscle activity is present in M1 (Fetz, et al., 1989; Kakei, et al., 1999; Morrow, et al., 2007; Cherian, et al., 2011; Oby, et al., 2012). As such, if S1 is to communicate feedback about movements to and receive efference copy information from M1, retaining a representation that contains muscle-like information would be beneficial.

In addition to this sensorimotor integration for controlling movement, another important aspect of motor control is movement planning, in which proprioceptive feedback appears to play an important role (Gordon, et al., 1995). In the case of visually-guided reaching, this planning requires an integration between the current state of the limb and a target location. Psychophysical studies suggest that this planning optimizes for dynamically smooth, straight hand trajectories (Flash & Hogan, 1985; Wolpert, et al., 1995). Furthermore, reaches appear to be planned with respect to the displacement between the hand and the intended target (Sainburg, et al., 2003). Together, these studies imply that reach planning takes place with respect to the motion of the hand. Given this, along with the psychophysical observation that we have better perceptual access to the position of our hands than to joint angles of the arm, might there be an area in the brain that represents proprioceptive information explicitly in terms of the hand? If so, one candidate location is the posterior parietal cortex (PPC), which integrates both proprioceptive

and visual information (Mountcastle, et al., 1975). Studies of arm-movement-related neurons in PPC suggest that PPC represents the location of both the reach target and the hand in an apparently visual coordinate frame (Buneo & Andersen, 2006) and is at least partially responsible for planning movements (Snyder, et al., 1997; Batista & Andersen, 2001; Cui & Andersen, 2007).

This suggests a dual purpose for the neural activity in S1: it may contain a feature library of proprioceptive feedback, read out in different ways, for different purposes. While M1 may read out S1 activity in terms of feedback about the limb for movement, PPC appears to integrate a readout of S1 activity with vision for movement planning. This potential flexibility in the interpretation of S1 activity is at least partially evidenced by the results presented in Chapter 3. There, we found that linear models of neural activity with inputs of either musculotendon length or the Cartesian kinematics of the hand and elbow predict S1 neural activity well. Thus, with a simple linear filter, M1 could use S1 activity for muscle-related feedback, while PPC uses S1 activity to infer the position of the hand in Cartesian coordinates.

Future experiments

Recording simultaneously from sensory and motor areas

Just as in the hierarchical processing perspective, one of the potential future directions for the motor control perspective is to record simultaneously from multiple areas to examine how proprioceptive feedback is processed. For example, one experiment could test the hypothesis that M1 and PPC have different read-outs from S1 activity, using simultaneous recordings from S1, M1, and PPC, and by using potent/null space analyses like those described above.

We can further interrogate the interplay between S1 and M1 activity by using tasks that require responses to perturbation. An early study of M1 showed that output-layer neurons

responded rapidly to a perturbation in two phases: the first dependent on the perturbation only, and the second slightly delayed and dependent on the task instructions (Evarts & Tanji, 1976). Since this seminal experiment, many psychophysics experiments have extended this finding to show relatively fast perturbation responses dependent on task parameters like target shape, the presence of obstacles, and how a cursor was controlled (Nashed, et al., 2012; Omrani, et al., 2013; Nashed, et al., 2014; Weiler, et al., 2018). These studies suggest a transcortical feedback loop, likely through S1 and M1, as hinted at by a single unit recording study of the two areas' responses to perturbation (Wolpaw, 1980). An experiment simultaneously examining M1 and S1 activity in response to perturbations during different tasks would shed light on the nature of M1-S1 communication during online feedback control.

Exploring changes in sensory areas during motor adaptation

Another question in this vein is how motor learning might change the activity of sensory areas, as well as the communication between sensory and motor areas. As described in Chapter 1 and alluded to above, internal models for movement generation are thought to be updated when there is a mismatch between predicted sensory feedback and actual feedback. This is evident in experiments where the dynamics of movement are suddenly altered by, for example, turning on a velocity-dependent force field. While this initially causes errors in reaching, eventually, the subjects adapt to the new dynamics, presumably by updating their internal models (Shadmehr & Mussa-Ivaldi, 1994).

Most studies examining the neurophysiology of this adaptation tend to focus on how motor cortical activity changes (Rokni, et al., 2007; Perich & Miller, 2017; Perich, et al., 2018), leaving proprioceptive feedback during motor adaptation relatively unexplored. However, several studies indicate that somatosensory cortex plays an important role in motor adaptation.

One recent study found that mice, like humans and monkeys, could adapt reaches to novel dynamics, but not when S1 was optogenetically inactivated. This deficit in adaptation was remarkably consistent with an inability to update an inverse model of the limb dynamics (Mathis, et al., 2017). Furthermore, a set of psychophysical experiments conducted by Dr. David Ostry's lab suggest that not only is sensory feedback important for motor adaptation, but the processing of the feedback appears to change as well. One of these studies found that adaptation to a force field caused subjects' perceptions of their hand position to shift in a direction dependent on the direction of the learned force field (Ostry, et al., 2010). A follow-up study showed that somatosensory evoked potentials changed in proportion to the degree of adaptation to the field (Nasir, et al., 2013).

One clear follow-up to these studies would be to replicate the psychophysical results in monkeys by having them report the direction of a perturbation to the hand before, during, and after adaptation to a force field. Subsequently, the study could look for changes in how S1 represents movement or in how S1 projects information other areas, like M1 and PPC.

Using normative models to characterize proprioceptive processing

There are many ways to study how a sensory system processes information, but one of the most common is to model neural activity as simple linear functions of sensory inputs. This approach, pioneered by Hubel and Wiesel in their study of early visual cortex (Hubel & Wiesel, 1959), was the basis of the analyses presented in Chapters 2 and 3 to examine how DSCT and S1 represent limb state. However, a downside of such a bottom-up approach is that in simply trying to get the best model fits, a study may become merely descriptive of neural activity and fail to provide an understanding of how it contributes to a high-level computational task, like image recognition or movement coordination (Marr & Poggio, 1976; Krakauer, et al., 2017). One recently popularized

alternative approach is the use of goal-based models to explain neural activity (Yamins & DiCarlo, 2016). In this top-down approach, researchers use modern machine learning techniques, like artificial neural networks, to build a system to complete a given computational task, like image recognition. In contrast to the bottom-up methods, this goal-based method essentially provides a computational model of how a high-level task might be completed. This gives experimenters the ability to examine the inner workings of the computational model, comparing the model's implementation to that of the actual neural system.

A recent example of this goal-driven modeling approach is a study examining deep neural networks trained to recognize images. Early, middle, and late layers of these networks were highly reminiscent of actual neural activity in different layers of the ventral visual stream (Yamins, et al., 2014). On the motor control side, other goal-driven modeling studies found that artificial neurons in recurrent neural networks trained to control arm muscles had properties similar to neurons in motor cortex (Lillicrap & Scott, 2013; Sussillo, et al., 2015).

Thus, one future direction of the work in this dissertation might be to use a goal-driven neural network to model the proprioceptive system. There may be many "goals" for the proprioceptive system, but in the context of motor control, three potential uses of proprioceptive feedback are for planning movements, for online motor control, and for motor adaptation. These reflect the interactions of S1 with PPC and M1 and thus, a goal-driven model approach might shed light on how these regions communicate. Specifically, one experiment might model the proprioceptive system, from DRG to PPC, as a feed-forward hierarchical neural network, computing the displacement between the hand and a reaching target. Likewise, a different experiment might use a similar hierarchical model of proprioceptive processing as a feedback

input for a recurrent neural network controlling muscle activity, like those described by Lillicrap et al. and Sussillo et al. To address motor adaptation, another experiment might examine how the activity of such an artificial sensorimotor control network would adapt to changes in environmental dynamics. These goal-driven models frame proprioceptive activity in terms of how it contributes to the computational goal of controlling movement, and thus may uncover features of proprioceptive activity that might not be found by classic modeling techniques.

Stimulating S1 to restore sensation for a bi-directional brain-machine interface

One direct clinical application of this research is in the development of brain-machine interfaces (BMI) to restore movement after spinal cord injury. The BMI field began with simple studies showing that a monkey could be trained to modulate the activity of individually recorded motor cortical neurons (Fetz, 1969), and that the activity of a population of motor cortical neurons could be used to predict movements and forces at the wrist (Humphrey, et al., 1970). Since these early experiments, the field of BMIs has advanced considerably, having developed techniques to decode motor cortical signals into commands for an on-screen cursor (Serruya, et al., 2002), a robotic arm (Collinger, et al., 2013), and even the paralyzed muscles of the subject's own arm (Ethier, et al., 2012). Despite these advances in motor BMIs, however, the ability to provide artificial proprioceptive feedback to the brain for motor control remains elusive.

One promising method to provide this feedback to the brain is with intracortical microstimulation (ICMS) in S1. Subjects can both detect and discriminate between different patterns of ICMS in S1, making this a potential method for providing artificial tactile feedback (Romo, et al., 1998; O'Doherty, et al., 2011; Zaaimi, et al., 2013; Flesher, et al., 2016; Salas, et al., 2018). One study even showed that monkeys could learn how to use ICMS in S1 to reach to a target without visual feedback (Dadarlat, et al., 2014). However, in that study, while monkeys

could eventually reach targets guided by ICMS, they required a prolonged training period during which the learned how to use the stimulation. Furthermore, the movements were far slower and more segmented than normal reaching movements, suggesting that the stimulation required cognitive interpretation, rather than taking advantage of the close ties between proprioceptive feedback and motor control.

In contrast with this cognitive learning approach, a biomimetic approach to S1 stimulation may be more likely to elicit sensations of movement, rather than simply a discriminable stimulus. In this approach, electrodes are stimulated with the goal of eliciting neural activity that mimics the normal activity of S1 during movement. One study has attempted this biomimetic approach to alter a monkey's perception of a mechanical perturbation to the hand, with limited success (Tomlinson & Miller, 2016), but no published studies so far have examined whether the biomimetic approach can be used for online feedback during movements. However, as we better characterize how S1 responds to movement, as well as how S1 communicates with motor areas, we may be able to design ICMS paradigms that more closely mimic the area's natural activity, allowing us to provide proprioceptive feedback that is more readily useful for online motor control.

Final conclusions

Classic models of proprioceptive neural activity in S1 assume that neurons represent simply the direction of hand movement, rather than intrinsic variables, like joint angles or muscle lengths. Surprisingly, this limb endpoint model outperformed a joint-based model for explaining neural activity in the DSCT, one of the earliest stages of proprioceptive processing. In this work, I used musculoskeletal modeling to examine how much biomechanics and muscle geometry contribute

to the neural activity at both the DSCT (Chapter 2) and S1 (Chapter 3). In both studies, I found that features of neural activity, like changes in preferred direction, could be explained better by a model based on the whole limb than by one based only on the limb endpoint. In this final chapter, I discussed the potential implications of these results, drawing on wisdom from the history of similar experiments in the motor cortical literature. Given this history, it seems unlikely that more precisely determining neural representations of movement will yield useful insights. Instead, it is likely to be more fruitful to frame proprioception in terms of how feedback is processed as it travels from muscles to cortex and how this progressive transformation contributes to the control of movement.

References

An, K. N., Takahashi, K., Harrigan, T. P. & Chao, E. Y., 1984. Determination of Muscle Orientations and Moment Arms. *Journal of Biomechanical Engineering,* Volume 106, p. 280.

Asatryan, D. G. & Feldman, A. G., 1965. Functional tuning of the nervous system with control of movement or maintenance of a steady posture. 1. Mechanographic analysis of the work of the joint on execution of a postural task. *Biophysics USSR,* Volume 10, pp. 925-935.

Azizi, E., Brainerd, E. L. & Roberts, T. J., 2008. Variable gearing in pennate muscles. *Proceedings of the National Academy of Sciences,* 1, Volume 105, pp. 1745-1750.

Batista, A. P. & Andersen, R. A., 2001. The Parietal Reach Region Codes the Next Planned Movement in a Sequential Reach Task. *Journal of Neurophysiology,* 2, Volume 85, pp. 539-544.

Bell, C. C., 1981. An Efference Copy Which is Modified by Reafferent Input. *Science,* Volume 214, pp. 450-453.

Bizzi, E., Accornero, N., Chapple, W. & Hogan, N., 1984. Posture control and trajectory formation during arm movement. *The Journal of Neuroscience,* 11, Volume 4, pp. 2738-2744.

Bosco, G. & Poppele, R. E., 1996. Temporal features of directional tuning by spinocerebellar neurons: relation to limb geometry. *Journal of Neurophysiology,* 4, Volume 75, pp. 1647-1658.

Bosco, G. & Poppele, R. E., 1997. Representation of Multiple Kinematic Parameters of the Cat Hindlimb in Spinocerebellar Activity. *Journal of Neurophysiology,* 9, Volume 78, pp. 1421-1432.

Bosco, G. & Poppele, R. E., 2000. Reference Frames for Spinal Proprioception: Kinematics Based or Kinetics Based?. *Journal of Neurophysiology,* 5, Volume 83, pp. 2946-2955.

Bosco, G. & Poppele, R. E., 2001. Proprioception From a Spinocerebellar Perspective. *Physiological Reviews,* 1 4, Volume 81, pp. 539-568.

Bosco, G., Poppele, R. E. & Eian, J., 2000. Reference Frames for Spinal Proprioception: Limb Endpoint Based or Joint-Level Based?. *Journal of Neurophysiology,* 5, Volume 83, pp. 2931-2945.

Bosco, G., Rankin, A. & Poppele, R., 1996. Representation of passive hindlimb postures in cat spinocerebellar activity. *Journal of Neurophysiology,* 8, Volume 76, pp. 715-726.

Brodmann, K., 1909. Vergleichende Lokalisationslehre der Grosshirnrinde in ihren Prinzipien dargestellt auf Grund des Zellenbaues. *Barth.*

Brown, M. C., Engberg, I. & Matthews, P. B. C., 1967. The relative sensitivity to vibration of muscle receptors of the cat. *The Journal of Physiology,* 10, Volume 192, pp. 773-800.

Bunderson, N. E., McKay, J. L., Ting, L. H. & Burkholder, T. J., 2010. Directional constraint of endpoint force emerges from hindlimb anatomy. *Journal of Experimental Biology,* 5, Volume 213, pp. 2131-2141.

Buneo, C. A. & Andersen, R. A., 2006. The posterior parietal cortex: Sensorimotor interface for the planning and online control of visually guided movements. *Neuropsychologia,* 1, Volume 44, pp. 2594-2606.

Burke, D., Hagbarth, K. E. & Löfstedt, L., 1978. Muscle spindle activity in man during shortening and lengthening contractions.. *The Journal of Physiology,* 4, Volume 277, pp. 131-142.

Burkholder, T. J. & Nichols, T. R., 2004. Three-dimensional model of the feline hindlimb. *Journal of Morphology,* Volume 261, pp. 118-129.

Cameron, A. C. & Windmeijer, F. A. G., 1996. R-Squared Measures for Count Data Regression Models With Applications to Health-Care Utilization. *Journal of Business & Economic Statistics,* 4, Volume 14, pp. 209-220.

Cameron, A. C. & Windmeijer, F. A. G., 1997. An R-squared measure of goodness of fit for some common nonlinear regression models. *Journal of Econometrics,* 4, Volume 77, pp. 329-342.

Caminiti, R., Johnson, P. B. & Urbano, A., 1990. Making arm movements within different parts of space: dynamic aspects in the primate motor cortex. *Journal of Neuroscience,* 1 7, Volume 10, pp. 2039-2058.

Carmena, J. M. et al., 2003. Learning to Control a Brain-Machine Interface for Reaching and Grasping by Primates. *PLoS Biology,* 10, Volume 1, p. e42.

Chan, S. S. & Moran, D. W., 2006. Computational model of a primate arm: from hand position to joint angles, joint torques and muscle forces. *Journal of Neural Engineering,* 11, Volume 3, pp. 327-337.

Chapman, C. E., Spidalieri, G. & Lamarre, Y., 1984. Discharge properties of area 5 neurones during arm movements triggered by sensory stimuli in the monkey. *Brain Research,* 8, Volume 309, pp. 63-77.

Cherian, A., Krucoff, M. O. & Miller, L. E., 2011. Motor cortical prediction of EMG: evidence that a kinetic brain-machine interface may be robust across altered movement dynamics. *Journal of Neurophysiology,* 8, Volume 106, pp. 564-575.

Chowdhury, R. H., Tresch, M. C. & Miller, L. E., 2017. Musculoskeletal geometry accounts for apparent extrinsic representation of paw position in dorsal spinocerebellar tract. *Journal of Neurophysiology,* 7, Volume 118, pp. 234-242.

Churchland, M. M. et al., 2012. Neural population dynamics during reaching. *Nature,* 6, Volume 487, pp. 51-56.

Churchland, M. M. et al., 2010. Cortical Preparatory Activity: Representation of Movement or First Cog in a Dynamical Machine?. *Neuron,* 11, Volume 68, pp. 387-400.

Collinger, J. L. et al., 2013. High-performance neuroprosthetic control by an individual with tetraplegia. *The Lancet,* 2, Volume 381, pp. 557-564.

Cooper, S. E., Martin, J. H. & Ghez, C., 2000. Effects of Inactivation of the Anterior Interpositus Nucleus on the Kinematic and Dynamic Control of Multijoint Movement. *Journal of Neurophysiology,* 10, Volume 84, pp. 1988-2000.

Cordo, P., Gurfinkel, V. S., Bevan, L. & Kerr, G. K., 1995. Proprioceptive consequences of tendon vibration during movement. *Journal of Neurophysiology,* 10, Volume 74, pp. 1675-1688.

Costanzo, R. M. & Gardner, E. P., 1981. Multiple-joint neurons in somatosensory cortex of awake monkeys. *Brain Research,* 6, Volume 214, pp. 321-333.

Crago, P. E., Houk, J. C. & Rymer, W. Z., 1982. Sampling of total muscle force by tendon organs.. *Journal of Neurophysiology,* 6, Volume 47, pp. 1069-1083.

Cui, H. & Andersen, R. A., 2007. Posterior Parietal Cortex Encodes Autonomously Selected Motor Plans. *Neuron,* 11, Volume 56, pp. 552-559.

Cunningham, J. P. & Yu, B. M., 2014. Dimensionality reduction for large-scale neural recordings. *Nature Neuroscience,* 8, Volume 17, pp. 1500-1509.

Dadarlat, M. C., O'Doherty, J. E. & Sabes, P. N., 2014. A learning-based approach to artificial sensory feedback leads to optimal integration. *Nature Neuroscience,* 11, Volume 18, pp. 138-144.

Daley, M. A. & Biewener, A. A., 2006. Running over rough terrain reveals limb control for intrinsic stability. *Proceedings of the National Academy of Sciences,* 10, Volume 103, pp. 15681-15686.

Debowy, D. J., Ghosh, S., Gardner, E. P. & Ro, J. Y., 2001. Comparison of neuronal firing rates in somatosensory and posterior parietal cortex during prehension. *Experimental Brain Research,* 4, Volume 137, pp. 269-291.

Dekleva, B. M., Kording, K. P. & Miller, L. E., 2018. Single reach plans in dorsal premotor cortex during a two-target task. *Nature Communications,* 9.Volume 9.

Deroy, O. & Auvray, M., 2012. Reading the World through the Skin and Ears: A New Perspective on Sensory Substitution. *Frontiers in Psychology,* Volume 3.

Dijkerman, H. C. & Haan, E. H. F., 2007. Somatosensory processing subserving perception and action: Dissociations, interactions, and integration. *Behavioral and Brain Sciences,* 4, Volume 30, p. 224.

Edin, B. B. & Vallbo, A. B., 1987. Twitch contraction for identification of human muscle afferents. *Acta Physiologica Scandinavica,* 9, Volume 131, pp. 129-138.

Elli, G. V., Benetti, S. & Collignon, O., 2014. Is There a Future for Sensory Substitution Outside Academic Laboratories?. *Multisensory Research,* Volume 27, pp. 271-291.

Elsayed, G. F. et al., 2016. Reorganization between preparatory and movement population responses in motor cortex. *Nature Communications,* 10.p. 13239.

Ernst, D., 2017. *Get CI and p-values for cross validated performance measures (AUC, rho).* s.l.:s.n.

Ethier, C., Oby, E. R., Bauman, M. J. & Miller, L. E., 2012. Restoration of grasp following paralysis through brain-controlled stimulation of muscles. *Nature,* 4, Volume 485, pp. 368-371.

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

Evarts, E. V. & Tanji, J., 1976. Reflex and intended responses in motor cortex pyramidal tract neurons of monkey. *Journal of Neurophysiology,* 9, Volume 39, pp. 1069-1080.

Fallon, J. B. & Macefield, V. G., 2007. Vibration sensitivity of human muscle spindles and golgi tendon organs. *Muscle & Nerve,* Volume 36, pp. 21-29.

Feldman, A. G., 1966. Functional tuning of the nervous system with control of movement or maintenance of a steady posture-II. Controllable parameters of the muscle. *Biophysics USSR,* Volume 11, pp. 565-578.

Feldman, A. G., 1986. Once More on the Equilibrium-Point Hypothesis (Model) for Motor Control. *Journal of Motor Behavior,* 3, Volume 18, pp. 17-54.

Felleman, D. J. & Essen, D. C. V., 1991. Distributed Hierarchical Processing in the Primate Cerebral Cortex. *Cerebral Cortex,* 1, Volume 1, pp. 1-47.

Fetz, E. E., 1969. Operant Conditioning of Cortical Unit Activity. *Science,* 28 2, Volume 163, pp. 955-958.

Fetz, E. E., 1992. Are movement parameters recognizably coded in the activity of single neurons?. *Behavioral and Brain Sciences,* Volume 15, pp. 679-690.

Fetz, E. E., Cheney, P. D., Mewes, K. & Palmer, S., 1989. Chapter 36 Control of forelimb muscle activity by populations of corticomotoneuronal and rubromotoneuronal cells. In: *Progress in Brain Research.* s.l.:Elsevier, pp. 437-449.

Flanagan, J. R., Ostry, D. J. & Feldman, A. G., 1993. Control of Trajectory Modifications in Target-Directed Reaching. *Journal of Motor Behavior,* 9, Volume 25, pp. 140-152.

Flash, T. & Hogan, N., 1985. The coordination of arm movements: an experimentally confirmed mathematical model. *The Journal of Neuroscience,* 7, Volume 5, pp. 1688-1703.

Flesher, S. N. et al., 2016. Intracortical microstimulation of human somatosensory cortex. *Science Translational Medicine,* 10, Volume 8, pp. 361ra141--361ra141.

Fuentes, C. T. & Bastian, A. J., 2010. Where Is Your Arm? Variations in Proprioception Across Space and Tasks. *Journal of Neurophysiology,* 1, Volume 103, pp. 164-171.

Gallego, J. A., Perich, M. G., Miller, L. E. & Solla, S. A., 2017. Neural Manifolds for the Control of Movement. *Neuron,* 6, Volume 94, pp. 978-984.

Gardner, E. P. & Costanzo, R. M., 1981. Properties of kinesthetic neurons in somatosensory cortex of awake monkeys. *Brain Research,* 6, Volume 214, pp. 301-319.

Gardner, E. P., Ro, J. Y., Debowy, D. & Ghosh, S., 1999. Facilitation of neuronal activity in somatosensory and posterior parietal cortex during prehension. *Experimental Brain Research,* 8, Volume 127, pp. 329-354.

Georgopoulos, A. P., Kalaska, J. F., Caminiti, R. & Massey, J. T., 1982. On the relations between the direction of two-dimensional arm movements and cell discharge in primate motor cortex. *Journal of Neuroscience,* 1 11, Volume 2, pp. 1527-1537.

Georgopoulos, A., Schwartz, A. & Kettner, R., 1986. Neuronal population coding of movement direction. *Science,* 9, Volume 233, pp. 1416-1419.

Ghez, C. & Sainburg, R., 1995. Proprioceptive control of interjoint coordination. *Canadian Journal of Physiology and Pharmacology,* 2, Volume 73, pp. 273-284.

Gomi, H. & Kawato, M., 1996. Equilibrium-Point Control Hypothesis Examined by Measured Arm Stiffness During Multijoint Movement. *Science,* 5 4, Volume 272, pp. 117-120.

Goodman, J. M. et al., 2019. Postural Representations of the Hand in Primate Sensorimotor Cortex. *bioRxiv,* 3.

Goodwin, G. M., McCloskey, D. I. & Matthews, P. B. C., 1972. Proprioceptive Illusions Induced by Muscle Vibration: Contribution by Muscle Spindles to Perception?. *Science,* 3, Volume 175, pp. 1382-1384.

Gordon, G. & Paine, C. H., 1960. Functional organization in nucleus gracilis of the cat. *The Journal of Physiology,* 9, Volume 153, pp. 331-349.

Gordon, J., Ghilardi, M. F. & Ghez, C., 1995. Impairments of reaching movements in patients without proprioception. I. Spatial errors. *Journal of Neurophysiology,* 1, Volume 73, pp. 347-360.

Gribble, P. L., Ostry, D. J., Sanguineti, V. & Laboissière, R., 1998. Are Complex Control Signals Required for Human Arm Movement?. *Journal of Neurophysiology,* 1 3, Volume 79, pp. 1409-1424.

Haftel, V. K. et al., 2004. Movement Reduces the Dynamic Response of Muscle Spindle Afferents and Motoneuron Synaptic Potentials in Rat. *Journal of Neurophysiology,* 5, Volume 91, pp. 2164-2171.

Hagbarth, K.-E. & Vallbo, A. B., 1968. Discharge characteristics of human muscle afferents during muscle stretch and contraction. *Experimental Neurology,* 12, Volume 22, pp. 674-694.

Hasan, Z., 1983. A model of spindle afferent response to muscle stretch. *Journal of Neurophysiology,* 4, Volume 49, pp. 989-1006.

Heath, C. J., Hore, J. & Phillips, C. G., 1976. Inputs from low threshold muscle and cutaneous afferents of hand and forearm to areas 3a and 3b of baboon's cerebral cortex.. *The Journal of Physiology,* 5, Volume 257, pp. 199-227.

Heinzl, H. & Mittlböck, M., 2003. Pseudo R-squared measures for Poisson regression models with over- or underdispersion. *Computational Statistics & Data Analysis,* 10, Volume 44, pp. 253-271.

Hinder, M. R. & Milner, T. E., 2003. The Case for an Internal Dynamics ModelversusEquilibrium Point Control in Human Movement. *The Journal of Physiology,* 6, Volume 549, pp. 953-963.

Houk, J. C., Rymer, W. Z. & Crago, P. E., 1981. Dependence of dynamic response of spindle receptors on muscle length and velocity.. *Journal of Neurophysiology,* 7, Volume 46, pp. 143-166.

Houk, J. & Simon, W., 1967. Responses of Golgi tendon organs to forces applied to muscle tendon.. *Journal of Neurophysiology,* 11, Volume 30, pp. 1466-1481.

Hounsgaard, J., Hultborn, H., Jespersen, B. & Kiehn, O., 1988. Bistability of alpha-motoneurones in the decerebrate cat and in the acute spinal cat after intravenous 5-hydroxytryptophan.. *The Journal of Physiology,* 11, Volume 405, pp. 345-367.

Hubel, D. H. & Wiesel, T. N., 1959. Receptive fields of single neurones in the cat's striate cortex. *The Journal of Physiology,* 10, Volume 148, pp. 574-591.

Hubel, D. H. & Wiesel, T. N., 1962. Receptive fields, binocular interaction and functional architecture in the cat's visual cortex. *The Journal of Physiology,* 1 1, Volume 160, pp. 106-154.

Hulliger, M., 1993. Chapter 16 Fusimotor control of proprioceptive feedback during locomotion and balancing: can simple lessons be learned for artificial control of gait?. In: *Natural and Artificial Control of Hearing and Balance.* s.l.:Elsevier, pp. 173-180.

Hulliger, M., Dürmüller, N., Prochazka, A. & Trend, P., 1989. Chapter 8 Flexible fusimotor control of muscle spindle feedback during a variety of natural movements. In: H. J ., et al. eds. *Progress in Brain Research.* s.l.:Elsevier, pp. 87-101.

Humphrey, D. R., Schmidt, E. M. & Thompson, W. D., 1970. Predicting Measures of Motor Performance from Multiple Cortical Spike Trains. *Science,* 13 11, Volume 170, pp. 758-762.

Hunt, C. C. & Kuffler, S. W., 1951. Stretch receptor discharges during muscle contraction. *The Journal of Physiology,* 4, Volume 113, pp. 298-315.

Hwang, E. J., Donchin, O., Smith, M. A. & Shadmehr, R., 2003. A Gain-Field Encoding of Limb Position and Velocity in the Internal Model of Arm Dynamics. *PLoS Biology,* Volume 1, p. e5.

Hyvärinen, J. & Poranen, A., 1978. Receptive field integration and submodality convergence in the hand area of the post-central gyrus of the alert monkey.. *The Journal of Physiology,* 10, Volume 283, pp. 539-556.

Iwamura, Y. & Tanaka, M., 1978. Postcentral neurons in hand region of area 2: their possible role in the form discrimination of tactile objects. *Brain Research,* 7, Volume 150, pp. 662-666.

Izawa, J., Criscimagna-Hemminger, S. E. & Shadmehr, R., 2012. Cerebellar Contributions to Reach Adaptation and Learning Sensory Consequences of Action. *Journal of Neuroscience,* 3, Volume 32, pp. 4230-4239.

Jennings, V. A., Lamour, Y., Solis, H. & Fromm, C., 1983. Somatosensory cortex activity related to position and force. *Journal of Neurophysiology,* 5, Volume 49, pp. 1216-1229.

Jun, J. J. et al., 2017. Fully integrated silicon probes for high-density recording of neural activity. *Nature,* 11, Volume 551, pp. 232-236.

Kaas, J. et al., 1979. Multiple representations of the body within the primary somatosensory cortex of primates. *Science,* 5, Volume 204, pp. 521-523.

Kakei, S., Hoffman, D. S. & Strick, P. L., 1999. Muscle and Movement Representations in the Primary Motor Cortex. *Science,* 9, Volume 285, pp. 2136-2139.

Kakei, S., Hoffman, D. S. & Strick, P. L., 2001. Direction of action is represented in the ventral premotor cortex. *Nature Neuroscience,* 9, Volume 4, pp. 1020-1025.

Kalaska, J. F., Caminiti, R. & Georgopoulos, A. P., 1983. Cortical mechanisms related to the direction of two-dimensional arm movements: relations in parietal area 5 and comparison with motor cortex. *Experimental Brain Research,* 7.Volume 51.

Kalaska, J. F., Cohen, D. A., Hyde, M. L. & Prud'homme, M., 1989. A comparison of movement direction-related versus load direction- related activity in primate motor cortex, using a two-dimensional reaching task. *The Journal of Neuroscience,* 6, Volume 9, pp. 2080-2102.

Kalaska, J. F. & Hyde, M. L., 1985. Area 4 and area 5: differences between the load direction-dependent discharge variability of cells during active postural fixation. *Experimental Brain Research,* 6.Volume 59.

Kandel, E. R. et al., 2012. *Principles of Neural Science, Fifth Edition.* s.l.:McGraw-Hill Education,.

Kao, J. C. et al., 2015. Single-trial dynamics of motor cortex and their applications to brain-machine interfaces. *Nature Communications,* 7.Volume 6.

Kaufman, M. T., Churchland, M. M., Ryu, S. I. & Shenoy, K. V., 2014. Cortical activity in the null space: permitting preparation without movement. *Nature Neuroscience,* 2, Volume 17, pp. 440-448.

Kaufman, M. T. et al., 2016. The Largest Response Component in the Motor Cortex Reflects Movement Timing but Not Movement Type. *eNeuro,* 8.Volume 3.

Kim, S. H. et al., 2010. Robot-assisted modifications of gait in healthy individuals. *Experimental Brain Research,* 2, Volume 202, pp. 809-824.

Kistemaker, D. A., Wong, J. D. & Gribble, P. L., 2010. The Central Nervous System Does Not Minimize Energy Cost in Arm Movements. *Journal of Neurophysiology,* 12, Volume 104, pp. 2985-2994.

Krakauer, J. W. et al., 2017. Neuroscience Needs Behavior: Correcting a Reductionist Bias. *Neuron,* 2, Volume 93, pp. 480-490.

Krizhevsky, A., Sutskever, I. & Hinton, G. E., 2012. *Imagenet classification with deep convolutional neural networks.* s.l., s.n., pp. 1097-1105.

Kuffler, S. W., Hunt, C. C. & Quilliam, J. P., 1951. FUNCTION OF MEDULLATED SMALL-NERVE FIBERS IN MAMMALIAN VENTRAL ROOTS: EFFERENT MUSCLE SPINDLE INNERVATION. *Journal of Neurophysiology,* 1, Volume 14, pp. 29-54.

Kurtzer, I. L., Pruszynski, J. A. & Scott, S. H., 2008. Long-Latency Reflexes of the Human Arm Reflect an Internal Model of Limb Dynamics. *Current Biology,* 3, Volume 18, pp. 449-453.

Lackner, J. R. & Dizio, P., 1994. Rapid adaptation to Coriolis force perturbations of arm trajectory. *Journal of Neurophysiology,* 7, Volume 72, pp. 299-313.

Landgren, S. & Silfvenius, H., 1971. Nucleus Z, the medullary relay in the projection path to the cerebral cortex of group I muscle afferents from the cat's hind limb. *The Journal of Physiology,* 11, Volume 218, pp. 551-571.

Liddell, E. G. T. & Sherrington, C., 1924. Reflexes in Response to Stretch (Myotatic Reflexes). *Proceedings of the Royal Society B: Biological Sciences,* 5, Volume 96, pp. 212-242.

Lillicrap, T. P. & Scott, S. H., 2013. Preference Distributions of Primary Motor Cortex Neurons Reflect Control Solutions Optimized for Limb Biomechanics. *Neuron,* 1, Volume 77, pp. 168-179.

Lin, C.-C. K. & Crago, P. E., 2002. Structural Model of the Muscle Spindle. *Annals of Biomedical Engineering,* 1, Volume 30, pp. 68-83.

Lloyd, D. P. C. & McIntyre, A. K., 1950. Dorsal column conduction of group I muscle afferent impulses and their relay through Clarke's Column. *Journal of Neurophysiology,* 1, Volume 13, pp. 39-54.

Loeb, G. E., 1981. Somatosensory unit input to the spinal cord during normal walking. *Canadian Journal of Physiology and Pharmacology,* 7, Volume 59, pp. 627-635.

Loeb, G. E., Hoffer, J. A. & Pratt, C. A., 1985. Activity of spindle afferents from cat anterior thigh muscles. I. Identification and patterns during normal locomotion. *Journal of Neurophysiology,* 9, Volume 54, pp. 549-564.

London, B. M. & Miller, L. E., 2013. Responses of somatosensory area 2 neurons to actively and passively generated limb movements. *Journal of Neurophysiology,* 3, Volume 109, pp. 1505-1513.

London, B. M., Torres, R. R., Slutzky, M. W. & Miller, L. E., 2011. *Designing stimulation patterns for an afferent BMI: Representation of kinetics in somatosensory cortex.* s.l., IEEE.

Lucas, A. et al., 2019. Neural Networks for Modeling Neural Spiking in S1 Cortex. *Frontiers in Systems Neuroscience,* 3.Volume 13.

Lucier, G. E., Rüegg, D. C. & Wiesendanger, M., 1975. Responses of neurones in motor cortex and in area 3A to controlled stretches of forelimb muscles in cebus monkeys.. *The Journal of Physiology,* 10, Volume 251, pp. 833-853.

Macefield, V. G. & Knellwolf, T. P., 2018. Functional properties of human muscle spindles. *Journal of Neurophysiology,* 8, Volume 120, pp. 452-467.

Maeda, R. S., Cluff, T., Gribble, P. L. & Pruszynski, J. A., 2018. Feedforward and feedback control share an internal model of the arm's dynamics. *The Journal of Neuroscience,* 10.pp. 1709-18.

Marr, D. & Poggio, T., 1976. *From understanding computation to understanding neural circuitry.* s.l.:s.n.

Mathis, M. W., Mathis, A. & Uchida, N., 2017. Somatosensory Cortex Plays an Essential Role in Forelimb Motor Adaptation in Mice. *Neuron,* 3, Volume 93, pp. 1493-1503.

McFadden, D., 1977. *Quantitative methods for analyzing travel behavior of individuals: some recent developments.* s.l.:Institute of Transportation Studies, University of California.

Mileusnic, M. P., Brown, I. E., Lan, N. & Loeb, G. E., 2006. Mathematical Models of Proprioceptors. I. Control and Transduction in the Muscle Spindle. *Journal of Neurophysiology,* 10, Volume 96, pp. 1772-1788.

Mishkin, M. & Ungerleider, L. G., 1982. Contribution of striate inputs to the visuospatial functions of parieto-preoccipital cortex in monkeys. *Behavioural Brain Research,* 9, Volume 6, pp. 57-77.

Moran, D. W. & Schwartz, A. B., 1999. Motor Cortical Representation of Speed and Direction During Reaching. *Journal of Neurophysiology,* 11, Volume 82, pp. 2676-2692.

Morrow, M. M., Jordan, L. R. & Miller, L. E., 2007. Direct Comparison of the Task-Dependent Discharge of M1 in Hand Space and Muscle Space. *Journal of Neurophysiology,* 2, Volume 97, pp. 1786-1798.

Morton, S. M., 2006. Cerebellar Contributions to Locomotor Adaptations during Splitbelt Treadmill Walking. *Journal of Neuroscience,* 9, Volume 26, pp. 9107-9116.

Mountcastle, V. B. et al., 1975. Posterior parietal association cortex of the monkey: command functions for operations within extrapersonal space. *Journal of Neurophysiology,* 7, Volume 38, pp. 871-908.

Mussa-Ivaldi, F. A., 1988. Do neurons in the motor cortex encode movement direction? An alternative hypothesis. *Neuroscience Letters,* 8, Volume 91, pp. 106-111.

Nadeau, C. & Bengio, Y., 2003. Inference for the Generalization Error. *Machine Learning,* Volume 52, pp. 239-281.

Nashed, J. Y., Crevecoeur, F. & Scott, S. H., 2012. Influence of the behavioral goal and environmental obstacles on rapid feedback responses. *Journal of Neurophysiology,* 8, Volume 108, pp. 999-1009.

Nashed, J. Y., Crevecoeur, F. & Scott, S. H., 2014. Rapid Online Selection between Multiple Motor Plans. *Journal of Neuroscience,* 1, Volume 34, pp. 1769-1780.

Nasir, S. M., Darainy, M. & Ostry, D. J., 2013. Sensorimotor adaptation changes the neural coding of somatosensory stimuli. *Journal of Neurophysiology,* 4, Volume 109, pp. 2077-2085.

Nelson, R. J., 1987. Activity of monkey primary somatosensory cortical neurons changes prior to active movement. *Brain Research,* 3, Volume 406, pp. 402-407.

Nichols, T. R. & Houk, J. C., 1973. Reflex Compensation for Variations in the Mechanical Properties of a Muscle. *Science,* 13 7, Volume 181, pp. 182-184.

Oby, E. R., Ethier, C. & Miller, L. E., 2012. Movement representation in the primary motor cortex and its contribution to generalizable EMG predictions. *Journal of Neurophysiology,* 2, Volume 109, pp. 666-678.

O'Doherty, J. E. et al., 2011. Active tactile exploration using a brain-machine-brain interface. *Nature,* 10, Volume 479, pp. 228-231.

Omrani, M., Diedrichsen, J. & Scott, S. H., 2013. Rapid feedback corrections during a bimanual postural task. *Journal of Neurophysiology,* 1, Volume 109, pp. 147-161.

Omrani, M., Murnaghan, C. D., Pruszynski, J. A. & Scott, S. H., 2016. Distributed task-specific processing of somatosensory feedback for voluntary motor control. *eLife,* 4.Volume 5.

Oscarsson, O. & Rosén, I., 1963. Projection to cerebral cortex of large muscle-spindle afferents in forelimb nerves of the cat. *The Journal of Physiology,* 12, Volume 169, pp. 924-945.

Ostry, D. J. et al., 2010. Somatosensory Plasticity and Motor Learning. *Journal of Neuroscience,* 4, Volume 30, pp. 5384-5393.

Padberg, J. et al., 2009. Thalamocortical Connections of Parietal Somatosensory Cortical Fields in Macaque Monkeys are Highly Divergent and Convergent. *Cerebral Cortex,* 2, Volume 19, pp. 2038-2064.

Padberg, J. et al., 2018. Cortical connections of area 2 and posterior parietal area 5 in macaque monkeys. *Journal of Comparative Neurology,* 5.

Pandarinath, C. et al., 2018. Inferring single-trial neural population dynamics using sequential auto-encoders. *Nature Methods,* 9.

Pavlides, C., Miyashita, E. & Asanuma, H., 1993. Projection from the sensory to the motor cortex is important in learning motor skills in the monkey. *Journal of Neurophysiology,* 8, Volume 70, pp. 733-741.

Penfield, W. & Boldrey, E., 1937. Somatic motor and sensory representation in the cerebral cortex of man as studied by electrical stimulation. *Brain,* Volume 60, pp. 389-443.

Perich, M. G., Gallego, J. A. & Miller, L. E., 2018. A Neural Population Mechanism for Rapid Learning. *Neuron,* 11, Volume 100, pp. 964--976.e7.

Perich, M. G. & Miller, L. E., 2017. Altered tuning in primary motor cortex does not account for behavioral adaptation during force field learning. *Experimental Brain Research,* 6, Volume 235, pp. 2689-2704.

Phillips, C. G., Powell, T. P. S. & Wiesendanger, M., 1971. Projection from low-threshold muscle afferents of hand and forearm to area 3a of baboon's cortex. *The Journal of Physiology,* 9, Volume 217, pp. 419-446.

Pons, T. P., Garraghty, P. E., Cusick, C. G. & Kaas, J. H., 1985. The somatotopic organization of area 2 in macaque monkeys. *The Journal of Comparative Neurology,* 11, Volume 241, pp. 445-466.

Pons, T. P. & Kaas, J. H., 1986. Corticocortical connections of area 2 of somatosensory cortex in macaque monkeys: A correlative anatomical and electrophysiological study. *The Journal of Comparative Neurology*, 6, Volume 248, pp. 313-335.

Prochazka, A., 1981. Muscle spindle function during normal movement. *International Review of Physiology*, Volume 25, pp. 47-90.

Prochazka, A. & Wand, P., 1981. Independence of fusimotor and skeletomotor systems during voluntary movement. In: *Muscle Receptors and Movement.* s.l.:Palgrave Macmillan UK, pp. 229-243.

Prochazka, A., Westerman, R. A. & Ziccone, S. P., 1976. Discharges of single hindlimb afferents in the freely moving cat. *Journal of Neurophysiology,* 9, Volume 39, pp. 1090-1104.

Prochazka, A., Westerman, R. A. & Ziccone, S. P., 1977. Ia afferent activity during a variety of voluntary movements in the cat. *The Journal of Physiology,* 6, Volume 268, pp. 423-448.

Proske, U., 1997. The Mammalian Muscle Spindle. *Physiology,* 2, Volume 12, pp. 37-42.

Proske, U. & Gandevia, S. C., 2012. The Proprioceptive Senses: Their Roles in Signaling Body Shape, Body Position and Movement, and Muscle Force. *Physiological Reviews,* 10, Volume 92, pp. 1651-1697.

Proske, U. & Stuart, G. J., 1985. The initial burst of impulses in responses of toad muscle spindles during stretch.. *The Journal of Physiology,* 11, Volume 368, pp. 1-17.

Prud'homme, M. J. & Kalaska, J. F., 1994. Proprioceptive activity in primate primary somatosensory cortex during active arm reaching movements. *Journal of Neurophysiology,* 11, Volume 72, pp. 2280-2301.

Rack, P. M. & Westbury, D. R., 1984. Elastic properties of the cat soleus tendon and their functional importance.. *The Journal of Physiology,* 2, Volume 347, pp. 479-495.

Rathelot, J.-A., Dum, R. P. & Strick, P. L., 2017. Posterior parietal cortex contains a command apparatus for hand movements. *Proceedings of the National Academy of Sciences,* 4, Volume 114, pp. 4255-4260.

Rathelot, J.-A. & Strick, P. L., 2009. Subdivisions of primary motor cortex based on cortico-motoneuronal cells. *Proceedings of the National Academy of Sciences,* 1, Volume 106, pp. 918-923.

Ravary, B. et al., 2004. Strain and force transducers used in human and veterinary tendon and ligament biomechanical studies. *Clinical Biomechanics,* 6, Volume 19, pp. 433-447.

Richardson, A. G., Weigand, P. K., Sritharan, S. Y. & Lucas, T. H., 2016. A chronic neural interface to the macaque dorsal column nuclei. *Journal of Neurophysiology,* 5, Volume 115, pp. 2255-2264.

Ro, J. Y., Debowy, D., Ghosh, S. & Gardner, E. P., 2000. Depression of neuronal firing rates in somatosensory and posterior parietal cortex during object acquisition in a prehension task. *Experimental Brain Research,* 10, Volume 135, pp. 1-11.

Rokni, U., Richardson, A. G., Bizzi, E. & Seung, H. S., 2007. Motor Learning with Unstable Neural Representations. *Neuron,* 5, Volume 54, pp. 653-666.

Romo, R., Hernández, A., Zainos, A. & Salinas, E., 1998. Somatosensory discrimination based on cortical microstimulation. *Nature,* 3, Volume 392, pp. 387-390.

Rosén, I., 1969. Afferent connexions to Group I activated cells in the main cuneate nucleus of the cat. *The Journal of Physiology,* 11, Volume 205, pp. 209-236.

Rosén, I., 1969. Excitation of Group I activated thalamocortical relay neurones in the cat. *The Journal of Physiology,* 11, Volume 205, pp. 237-255.

Rosén, I., 1969. Localization in caudal brain stem and cervical spinal cord of neurones activated from forelimb group I afferents in the cat. *Brain Research,* 11, Volume 16, pp. 55-71.

Russo, A. A. et al., 2018. Motor Cortex Embeds Muscle-like Commands in an Untangled Population Response. *Neuron,* 2, Volume 97, pp. 953--966.e8.

Sainburg, R. L., Ghilardi, M. F., Poizner, H. & Ghez, C., 1995. Control of limb dynamics in normal subjects and patients without proprioception. *Journal of Neurophysiology,* 2, Volume 73, pp. 820-835.

Sainburg, R. L., Lateiner, J. E., Latash, M. L. & Bagesteiro, L. B., 2003. Effects of Altering Initial Position on Movement Direction and Extent. *Journal of Neurophysiology,* 1, Volume 89, pp. 401-415.

Sainburg, R. L., Poizner, H. & Ghez, C., 1993. Loss of proprioception produces deficits in interjoint coordination. *Journal of Neurophysiology,* 11, Volume 70, pp. 2136-2147.

Salas, M. A. et al., 2018. Proprioceptive and cutaneous sensations in humans elicited by intracortical microstimulation. *eLife,* 4.Volume 7.

Sanes, J. N. et al., 1984. Motor deficits in patients with large-fiber sensory neuropathy.. *Proceedings of the National Academy of Sciences,* 2, Volume 81, pp. 979-982.

Scott, S. H. & Kalaska, J. F., 1997. Reaching Movements With Similar Hand Paths But Different Arm Orientations. I. Activity of Individual Cells in Motor Cortex. *Journal of Neurophysiology,* 1 2, Volume 77, pp. 826-852.

Sedda, A. & Scarpina, F., 2012. Dorsal and ventral streams across sensory modalities. *Neuroscience Bulletin,* 6, Volume 28, pp. 291-300.

Seelke, A. M. H. et al., 2011. Topographic Maps within Brodmann's Area 5 of Macaque Monkeys. *Cerebral Cortex,* 9, Volume 22, pp. 1834-1850.

Semedo, J. D. et al., 2019. Cortical Areas Interact through a Communication Subspace. *Neuron,* 2.

Sergio, L. E., Hamel-Pâquet, C. & Kalaska, J. F., 2005. Motor Cortex Neural Correlates of Output Kinematics and Kinetics During Isometric-Force and Arm-Reaching Tasks. *Journal of Neurophysiology,* 10, Volume 94, pp. 2353-2378.

Sergio, L. E. & Kalaska, J. F., 2003. Systematic Changes in Motor Cortex Cell Activity With Arm Posture During Directional Isometric Force Generation. *Journal of Neurophysiology,* 1, Volume 89, pp. 212-228.

Serruya, M. D. et al., 2002. Instant neural control of a movement signal. *Nature,* 3, Volume 416, pp. 141-142.

Shadmehr, R. & Mussa-Ivaldi, F. A., 1994. Adaptive representation of dynamics during learning of a motor task. *Journal of Neuroscience,* 1 5, Volume 14, pp. 3208-3224.

Shen, L. & Alexander, G. E., 1997. Neural Correlates of a Spatial Sensory-To-Motor Transformation in Primary Motor Cortex. *Journal of Neurophysiology,* 3, Volume 77, pp. 1171-1194.

Snyder, L. H., Batista, A. P. & Andersen, R. A., 1997. Coding of intention in the posterior parietal cortex. *Nature,* 3, Volume 386, pp. 167-170.

Soechting, J. F. & Ross, B., 1984. Psychophysical determination of coordinate representation of human arm orientation. *Neuroscience,* 10, Volume 13, pp. 595-604.

Stein, R. B. et al., 2004. Coding of position by simultaneously recorded sensory neurones in the cat dorsal root ganglion. *The Journal of Physiology,* 10, Volume 560, pp. 883-896.

Stevenson, I. H. et al., 2011. Statistical assessment of the stability of neural movement representations. *Journal of Neurophysiology,* 8, Volume 106, pp. 764-774.

Suresh, A. K. et al., 2017. Methodological considerations for a chronic neural interface with the cuneate nucleus of macaques. *Journal of Neurophysiology,* 12, Volume 118, pp. 3271-3281.

Sussillo, D., Churchland, M. M., Kaufman, M. T. & Shenoy, K. V., 2015. A neural network that finds a naturalistic solution for the production of muscle activity. *Nature Neuroscience,* 6, Volume 18, pp. 1025-1033.

Tabot, G. A. et al., 2013. Restoring the sense of touch with a prosthetic hand through a brain interface. *Proceedings of the National Academy of Sciences,* 10, Volume 110, pp. 18279-18284.

Tanji, J., 1975. Activity of neurons in cortical area 3a during maintenance of steady postures by the monkey. *Brain Research,* 5, Volume 88, pp. 549-553.

Taylor, A., Durbaba, R., Ellaway, P. H. & Rawlinson, S., 2000. Patterns of fusimotor activity during locomotion in the decerebrate cat deduced from recordings from hindlimb muscle spindles. *The Journal of Physiology,* 2, Volume 522, pp. 515-532.

Taylor, D. M., Tillery, S. I. H. & Schwartz, A. B., 2002. Direct Cortical Control of 3D Neuroprosthetic Devices. *Science,* 6, Volume 296, pp. 1829-1832.

Thelen, D. G., Anderson, F. C. & Delp, S. L., 2003. Generating dynamic simulations of movement using computed muscle control. *Journal of Biomechanics,* 3, Volume 36, pp. 321-328.

Therrien, A. S. & Bastian, A. J., 2015. Cerebellar damage impairs internal predictions for sensory and motor function. *Current Opinion in Neurobiology,* 8, Volume 33, pp. 127-133.

Thoroughman, K. A. & Shadmehr, R., 1999. Electromyographic Correlates of Learning an Internal Model of Reaching Movements. *Journal of Neuroscience,* 1 10, Volume 19, pp. 8573-8588.

Tillery, S. I., Soechting, J. F. & Ebner, T. J., 1996. Somatosensory cortical activity in relation to arm posture: nonuniform spatial tuning. *Journal of Neurophysiology,* 10, Volume 76, pp. 2423-2438.

Ting, L. H. & Macpherson, J. M., 2005. A Limited Set of Muscle Synergies for Force Control During a Postural Task. *Journal of Neurophysiology,* 1, Volume 93, pp. 609-613.

Tomlinson, T. & Miller, L. E., 2016. Toward a Proprioceptive Neural Interface that Mimics Natural Cortical Activity. In: *Advances in Experimental Medicine and Biology.* s.l.:Springer International Publishing, pp. 367-388.

Truccolo, W. et al., 2005. A Point Process Framework for Relating Neural Spiking Activity to Spiking History, Neural Ensemble, and Extrinsic Covariate Effects. *Journal of Neurophysiology,* 2, Volume 93, pp. 1074-1089.

Weber, D. J. et al., 2011. Limb-State Information Encoded by Peripheral and Central Somatosensory Neurons: Implications for an Afferent Interface. *IEEE Transactions on Neural Systems and Rehabilitation Engineering,* 10, Volume 19, pp. 501-513.

Weber, D. J., Stein, R. B., Everaert, D. G. & Prochazka, A., 2006. Decoding Sensory Feedback From Firing Rates of Afferent Ensembles Recorded in Cat Dorsal Root Ganglia in Normal Locomotion. *IEEE Transactions on Neural Systems and Rehabilitation Engineering,* 6, Volume 14, pp. 240-243.

Weber, D. J., Stein, R. B., Everaert, D. G. & Prochazka, A., 2007. Limb-state feedback from ensembles of simultaneously recorded dorsal root ganglion neurons. *Journal of Neural Engineering,* 8, Volume 4, pp. S168--S180.

Weiler, J., Gribble, P. L. & Pruszynski, J. A., 2018. Rapid feedback responses are flexibly coordinated across arm muscles to support goal-directed reaching. *Journal of Neurophysiology,* 2, Volume 119, pp. 537-547.

Weiler, J., Gribble, P. L. & Pruszynski, J. A., 2019. Spinal stretch reflexes support efficient hand control. *Nature Neuroscience,* 2.

Wessberg, J. et al., 2000. Real-time prediction of hand trajectory by ensembles of cortical neurons in primates. *Nature,* 11, Volume 408, pp. 361-365.

Wolpaw, J. R., 1980. Correlations between task-related activity and responses to perturbation in primate sensorimotor cortex. *Journal of Neurophysiology,* 12, Volume 44, pp. 1122-1138.

Wolpert, D., Ghahramani, Z. & Jordan, M., 1995. Are arm trajectories planned in kinematic or dynamic coordinates? An adaptation study. *Experimental Brain Research,* 11.Volume 103.

Wolpert, D. M., Ghahramani, Z. & Jordan, M. I., 1995. An Internal Model for Sensorimotor Integration. *Science,* 29 9, Volume 269, pp. 1880-1882.

Wolpert, D. M., Miall, R. C. & Kawato, M., 1998. Internal models in the cerebellum. *Trends in Cognitive Sciences,* 9, Volume 2, pp. 338-347.

Yamada, H. et al., 2016. Representation of Afferent Signals from Forearm Muscle and Cutaneous Nerves in the Primary Somatosensory Cortex of the Macaque Monkey. *PLOS ONE,* 10, Volume 11, p. e0163948.

Yamins, D. L. K. & DiCarlo, J. J., 2016. Using goal-driven deep learning models to understand sensory cortex. *Nature Neuroscience,* 3, Volume 19, pp. 356-365.

Yamins, D. L. K. et al., 2014. Performance-optimized hierarchical models predict neural responses in higher visual cortex. *Proceedings of the National Academy of Sciences,* 5, Volume 111, pp. 8619-8624.

Yanai, Y. et al., 2008. Coordinate Transformation is First Completed Downstream of Primary Motor Cortex. *Journal of Neuroscience,* 2, Volume 28, pp. 1728-1732.

Young, D. et al., 2018. Closed-loop cortical control of virtual reach and posture using cartesian and joint velocity commands. *Journal of Neural Engineering,* 12.

Yumiya, H., Kubota, K. & Asanuma, H., 1974. Activities of neurons in area 3a of the cerebral cortex during voluntary movements in the monkey. *Brain Research,* 9, Volume 78, pp. 169-177.

Zaaimi, B., Ruiz-Torres, R., Solla, S. A. & Miller, L. E., 2013. Multi-electrode stimulation in somatosensory cortex increases probability of detection. *Journal of Neural Engineering,* 8, Volume 10, p. 056013.