BRAIN COMPUTER INTERFACES: CHALLENGES TO CLINICAL VIABILITY ADDRESSED IN THE LABORATORY

by

Jessica L. Mischel

BS, Washington University in St. Louis, 2012

Submitted to the Graduate Faculty of

the Swanson School of Engineering in partial fulfillment

of the requirements for the degree of

Doctor of Philosophy

University of Pittsburgh

2016

UNIVERSITY OF PITTSBURGH

SWANSON SCHOOL OF ENGINEERING

This dissertation was presented

by

Jessica L. Mischel

It was defended on

October 14, 2016

and approved by

Aaron Batista, Ph.D., Associate Professor, Department of Bioengineering

Patrick Loughlin, Ph.D., Professor, Department of Bioengineering

Neeraj Gandhi, Ph.D., Associate Professor, Department of Bioengineering

Robert Gaunt, Ph.D., Assistant Professor, Department of Physical Medicine and Rehabilitation

Dissertation Director: Aaron Batista, Ph.D., Associate Professor, Department of Bioengineering

BRAIN COMPUTER INTERFACES: CHALLENGES TO CLINICAL VIABILITY ADDRESSED IN THE LABORATORY

Jessica L. Mischel, PhD

University of Pittsburgh, 2016

Paralysis following spinal cord injuries, amyotrophic lateral sclerosis, stroke, and other disorders can intervene with signal transduction from the brain to the motor periphery, and eliminate the ability to perform volitional movements. Brain computer interfaces (BCI) directly measure brain activity associated with the user's intent and translate the recorded brain activity into control signals for BCI applications, such as moving a computer cursor or a robot arm. While BCI technology has become an active and exciting field of research, much of the field's development and achievements to date have taken place in the laboratory. The translation of BCI technology to the clinical setting is still not a reality.

My thesis research has been dedicated to the objective of facilitating the translation of BCI systems from the primate lab to a clinical setting. That guiding objective has led me to work on several projects including: a technique that vastly improves the longevity of surgical implants in primate studies; a task that pushes the limits of sensorimotor performance – improving our knowledge of the function of primary motor cortex during realistic reaches and allowing us to quantify feedback effectiveness; characterizing the long-term tissue response to chronically implanted electrodes, and investigating how to optimally select parameters for neural information extraction. Each of these contributions will help bring BCI systems one step closer to clinical reality.

.

TABLE OF CONTENTS

LIST OF TABLES

LIST OF FIGURES

PREFACE

When I started graduate school, I had so many ideas in my head I could barely put a sentence together that made any kind of sense. My primary advisor, Dr. Aaron Batista, recognized that while my brain was all over the place, my ideas and intentions were good. He saw that I had the potential to be a great scientist, I just needed some direction. He gave me the tools to organize my thoughts and taught me to more clearly articulate my ideas. Aaron, you have helped me grow as a scientist and as a person, and I cannot thank you enough. Your genuine care for the people you work with promotes a comfortable and productive work environment. Your ability to dream big is what drew me to the lab, but your motivation and support are what kept me here. You believed in me, and taught me to believe in myself. Thank you for everything.

I would also like to acknowledge Dr. Patrick Loughlin for his contributions to the critical stability task (CST) project. Pat has been extremely important to the success of the CST. He supported the project in its early stages when others dismissed the idea and has recently co-authored an R01 with Aaron using the CST to explore sensorimotor integration. Pat has been a valuable resource throughout my graduate career, helping me understand the properties of a control system and providing insights based on his experience with human experiments. Thank you for believing in the CST and investing so much of yourself in the project. Working with you during your sabbatical was enlightening – not only did I become more comfortable with the

systems prospective of my work, but I also got to observe the work ethic of an academic badass. Thank you for your personal and academic support.

I would also like to thank the other two members of my thesis committee, Drs. Robert Gaunt and Neeraj Gandhi, for their direction and guidance. My thesis has changed considerably since I first proposed, and Rob and Raj have supported my work and my ideas in light of the changes. They have made themselves readily available for advice; I only wish I had taken advantage of this offer more frequently. Rob has been very responsive to the utility of the CST, implementing the task with a human subject. Watching the human subject perform a task that I had helped design using his neural signals brought tears to my eyes; I saw something that I helped develop work in a human. Thank you for giving me that opportunity. Raj has been a constant source of support throughout my graduate school career. Thank you for asking the hard questions (I will never forget that preliminary exam) and pushing me to do better science.

Importantly, I would like to thank Dr. Kristin Quick. You helped me find my way as an engineer in a neuroscience lab. I'll be honest; I wouldn't say the CST has answered most of the questions that brought me to graduate school, but it's opened a door so we can answer those questions the right way. Your composure, critical-thinking, and drive inspired me to work harder and more thoughtfully. I'd like to believe I'm a little less frenetic having worked under your guidance. Thank you for your patience, your advice, and your friendship.

I have to thank my lab spouse, Nick Pavlovsky, for continuously having my back. We've come a long way since summer 2012, and working so closely with you in classes and in lab has dramatically improved my graduate school experience. I'm honored to call you my science husband. I would also like to thank Dr. Berook Alemayehu who, aside from being an amazing roommate, taught me the importance of maintaining my hobbies and identity throughout the

hardest months of my graduate school career. He saw me through long nights building the new experimental rig, took my crazy dog out when I couldn't get home in time, and introduced me to the glory of frozen pizzas, which has proven critical while writing this thesis. Your presence in my house made coming home after a hard day so much more enjoyable.

I would like to thank the other Batista lab members who have helped me grow into the researcher I am today and not pull my hair out in the process. Dr. Jason Godlove taught me to find fun in the job and take life a little less seriously. Dr. Patrick Sadtler got me started in the lab, teaching me the ins and outs of our data acquisition software and providing essential support as I began performing my own analyses. Drs. Emily Oby and Alan Degenhart exemplified hard work, leading by example and providing valuable insight that improved my research. Jeffery Chiou – you have been an invaluable resource when it came to building the new rig computers, debugging analysis code, or simply trying to pass Systems Neuroscience and Neural Signal Processing. Nicole McClain – I can't think of a better individual to take over the CST project. Your coding prowess and experience with control systems will enable this project to thrive as it continues moving forward. I am excited to see the great science that will be achieved with the CST under your leadership. Erinn Grigsby – your work ethic is ridiculous. It's rare for a first year student to come in nights and weekends to train, stay up all night working on an RO1, and perform general lab maintenance all while taking classes. You've motivated me to push harder. Dylan Royston – while you may not be an official member of the Batista lab, your enthusiasm and passion are contagious and have recharged my excitement for my work on several occasions. You have given me valuable insight on my work and encouraged me to think out of the box. Carol Ehnerd, Jocelyn Stafura, and Melissa James – thank you for providing our animals with an unparalleled level of care and for all the training advice. The support fostered in the Batista lab is

truly special, and I am lucky to have worked with every one of you. For every computer fix, monkey training tip, PhaseSpace calibration session, presentation/research suggestion, and joke in the midst of a rough day, thank you from the bottom of my heart.

Additionally, I would like to thank my three college roommates and all my friends here in Pittsburgh. The people I have the honor of calling my friends are like family to me. Erica Lash, Sara Moskowitz, and Vaishnavi Hariprasad – you guys have stuck with me through everything. I would not have applied to graduate school if it wasn't for your support. We may be all over the country, but you guys have shaped me into the person I am today and continue to help me become a better person. Thank you. And to my Pittsburgh family, I am so fortunate to surround myself with such smart and dedicated individuals. Winers – you guys have been more than just a group of girls to drink wine with and whine about my problems. While I may question your taste in wine, I have never questioned your support. Festival friends – hanging out with you guys has helped me be comfortable being me. Even when the work has gotten rough, having confidence in my character and my science has gotten me through. Thank you for highlighting the importance of self-love and always being there for me.

I also need to thank my family, for hanging in there with me for so many years. Your love and support was essential for my academic survival. With each frustration I experienced, my family motivated me to keep my head up, stay true to myself and my research, and continue moving forward. First, I would like to thank my parents, Dr. Steven and Pamela Mischel. You guys have congratulated my successes and offered a shoulder to cry on when I've needed it. I'm lucky to have a family that's willing to fly to Pittsburgh to offer their support. I would also like to thank my sister. I'm so happy that we have gotten closer. You have been such a source of support for me over the last few years; you always knew exactly what to say to keep me pushing

forward and doing my best. I only hope that I can do the same as you move towards your degree. I love you all so much, and I hope you all are ready for another Dr. Mischel in the family.

Finally, I need to acknowledge my crazy dog, Ringo. Ringo has been my sounding board, and her inability to talk back has led to some of my most creative ideas yet. She's sat through each iteration of my various presentations, providing invaluable feedback with her gazes and yawns. I couldn't have done it without her.

1.0 GENERAL INTRODUCTION

Movement is essential for interacting with the environment. Starting at birth, we experience our surroundings using rudimentary limb movements, and over time our movements become more complex. Movements such as walking, speaking, or using a tool require that we tailor our actions to the environment in an ongoing manner, utilizing sensory information to determine the appropriate muscle forces and joint torques for the desired action. Imagine picking up a cup of water; we use visual information to locate the cup and identify its shape, use proprioceptive feedback to determine where our hand is in space, and use tactile feedback to determine the texture and weight of the glass. Sensory feedback is essential to accurately move our hand to the cup, shape our fingers to grasp the cup, and apply the correct amount of force to keep it safely in our grasp. Complex, sensory-guided movements such as this are common to everyday life, and enable the exploration and manipulation of our environment.

Paralysis following spinal cord injuries, amyotrophic lateral sclerosis, stroke, and other disorders can intervene with signal transduction from the brain to the motor periphery, and eliminate the ability to perform volitional movements. This forces paralyzed individuals to rely on care-takers to complete routine daily activities, such moving, eating, and bathing. Brain computer interfaces (BCI) offer a promising therapy for these individuals. BCIs directly measure brain activity associated with user intent and translate the recorded brain activity into control signals for BCI applications, such as a prosthetic limb, wheelchair control, and cursor movement.

This technology can help paralyzed individuals regain their independence and subsequently improve their quality of life.

Within the last decade, BCI technology has become an active and exciting field of research; however, much of the field's development and achievements to date have taken place in the laboratory. Work to date demonstrates the promising potential of BCIs to replace or restore movement to individuals with severe movement disorders, however control is limited to simple reaches and the longevity of recording modalities is limited. These challenges to clinical viability must be addressed before BCIs can be realized as a clinical therapy.

In this introduction, I will discuss natural sensorimotor control (Section 1.1), review the current state of BCI technologies (Section 1.2), and briefly describe some of the necessary advancements that must be made before BCIs can enter the clinical world (Section 1.3).

1.1 NATURAL SENSORIMOTOR CONTROL

Primary motor cortex (M1) plays an important role in voluntary motor behavior, and in combination with various forms of feedback, the sensorimotor loop enables the exploration and manipulation of our environment.

1.1.1 Motor control

Primary motor cortex (M1) is the final common pathway for voluntary motor commands, and thus plays an active part in natural sensorimotor control. These commands originate from corticospinal neurons in layer V of M1, which have axons that descend to the spinal cord. These

spinal motoneurons then make direct connections to the muscles (Dum and Strick, 2002; Rathelot and Strick, 2009; Porter and Lemon, 1993). Thus, neurons in M1 provide the most direct link between the brain and the muscles, and hence participate in the control movement. This makes M1 a prime location for studying motor control.

For decades, neuroscientists have been searching for the mechanism by which primary motor cortex controls voluntary movement. It remains controversial whether motor cortex encodes kinetic features of muscles and joints, such as the force and torque (Fetz and Cheney, 1980; Evarts, 1968; Kakei et al., 1999), or kinematic features such as position, velocity, and direction of the limb (Georgopoulos et al., 1982; 1986; Moran and Schwartz, 1999; Taylor et. al, 2002). Using different behavioral paradigms, experimental data have been interpreted to support both arguments. One possible explanation for this paradox is that M1 neurons code for context-specific, behaviorally relevant task features they directly influence. Another possible explanation is that the methods we are using to study M1 activity are not capturing the underlying mechanisms of *general* motor control.

Our current understanding of the neural mechanisms of arm movement is based on studies using single, relatively brief movements with a discrete start and end, such as a center-out reach (Georgeopoulos et al., 1982; Georgeopoulos et al., 1986; Schwartz et al., 1988; Georgeopoulos et al., 1988, Batista et al., 1999). These movements are highly practiced by the time neurophysiological data collection begins, and do not require ongoing sensory feedback for reasonably good performance (Gordon et al., 1995; Ghez et al., 1995). In contrast, real movements such as using a tool require that we tailor our actions to the environment in an ongoing manner. Thus, it is difficult to infer from these studies how M1 might function to control ongoing, realistic behaviors.

1.1.2 Sensory feedback

1.1.2.1 Visual feedback

Visual feedback is a primary mode for gathering information about our environment. As we move our eyes around a scene, light hits our retina and is processed to create the perception of our visual world. The flow of visual information begins at the retina and is then relayed to the lateral geniculate nucleus (LGN) of the thalamus before it arrives in the primary visual cortex (V1), the earliest visual processing area of the cerebral cortex (Callaway 1998). From V1, visual information is processed via connections with other cortical areas (Van Essen and Maunsell, 1983). One function of visual information is to guide and command subsequent eye movements, or saccades, during the process of scanning a scene. The superior colliculus plays a major role in the sensorimotor transformation of visual information to a saccade (Sparks and Mays, 1990; Hall and Moschovakis, 2004; Gandhi and Katnani, 2011). The superficial layers of the superior colliculus respond almost entirely to visual information, whereas the deep layers produce a burst of neural activity just prior to a saccade (Gandhi and Katnani, 2011). After a saccade is complete, a new set of visual information enters the brain through the retina so that the process may begin again.

Another function of visual information is to guide goal directed limb and body movements. Visual feedback is required for making complex movements, such as avoiding obstacles or adjusting reach trajectories. Even ballistic movements that do not require continuous feedback require knowledge of a reach goal. The successful completion of a reach relies on the ability of the visual-motor system to use visual stimuli to locate a reach goal and inform the correct motor command to guide the arm. Before visual information becomes a motor command, the reach goal must be transformed from visual coordinates into a coordinate frame that is

meaningful for reaches. An important region for this transformation is the parietal reach region (PRR) of the posterior parietal cortex. Here, the visual location of intended reach is represented in visual coordinates (Bremmer et al., 1998; Batista et al., 1999; Buneo et al., 2002). PRR modulates its activity based on the context of the task (Gail and Andersen, 2006), indicating that the neural activation is not purely a response visual stimuli. This provides evidence that PRR may be a node for sensorimotor transformation. The dorsal premotor cortex (PMd) is also a likely node of sensorimotor transformation. PMd is involved in planning and executing visually-guided reaches (Weinrich and Wise, 1982; He et al., 1993). Neuronal activation in PMd is robust while preparing to make a movement (Weinrich and Wise, 1982) and reflects the direction and amplitude of movement (Riehle and Requin, 1989; Caminiti et al., 1991; Fu et al., 1995; Crammond and Kalaska, 1994; 1996). It is not yet quite clear what role PMd plays in this sensorimotor transformation of reach goals from visual coordinates to reach coordinates, however it is well poised to be a key player. While it is still unclear exactly where in the brain this visual-motor transformation occurs, converting visual information into movement intention, such as a reach target or a corrective movement, is a crucial step in making accurate ongoing and realistic movements.

1.1.2.2 Non-visual feedback

Non-visual feedback is required when visual attention is diverted and in tasks that do not require visual feedback, such as grasping. During human motor control, tactile feedback and proprioceptive feedback are the dominant forms providing of non-visual information. Tactile sensation is the sense of touch produced by pressure receptors in the skin. We rely on tactile information to modulate grip force, understand textures, and in extreme cases, stop moving after running into an obstacle. Tactile signals such as discriminative touch, pressure, and vibration are

transmitted to higher brain centers, reaching consciousness, by way of the posterior column–medial lemniscus pathway. Damage to this pathway can render patients unable to hold objects, and causes patients to use inefficiently high grip forces when manipulating objects (Moberg 1962, Nowak et al. 2004).

While visual and tactile information are very good at conveying information about the external environment, proprioceptive feedback provides information on the movements of our own body (Paillard and Brouchon 1968, Weber et al. 2011). Patients who have lost proprioceptive sensation due to large-fiber neuropathy show clear motor deficits. They are still able to perform a reaching movement when visual feedback is removed, but make large errors related to extent and direction of the movement (Gordon et all 1995). These errors are significantly reduced when patients when visual information is available, however performance does not match that of control subjects.

1.2 CURRENT STATE OF BRAIN COMPUTER INTERFACES

Brain computer interfaces (BCIs) use neural activity to control external devices, enabling paralyzed individuals to interact with the environment. Instead of depending on peripheral nerves and muscles to achieve movement, a BCI directly measures brain activity associated with the user's intent and translates it into control signals for BCI applications, such as moving a computer cursor or a robot arm (Taylor et al., 2002; Carmena et al. 2003; Fabiani et al., 2004; Lebedev and Nicolelis, 2006; Hochberg et al., 2006; 2012; Velliste et al., 2008; O'Doherty et al., 2011; Schalk and Leuthardt, 2011; Gilja et. al., 2012; Collinger et al., 2013; Wang et al., 2013). Every BCI system must contain three components; it must record activity directly from the brain

(invasively or non-invasively), it must rely on intentional control – decoded into task parameters, and it must provide feedback to the user in real-time.

1.2.1 BCI signal acquisition

A critical component of a BCI system is the neural recording modality used to extract meaningful information from the brain. The primary modalities used for BCI research are intracortical electrode arrays, electrocorticography, and electroencephalography. These methods are common in their use of electrodes to record electrical activity from the cortex, but vary in electrode design and placement. These differences have implications for the characteristics of the recorded neural activity.

1.2.1.1 Intracortical Arrays

Penetrating intracortical electrode arrays offer the highest spatial and temporal resolution in neural recording. Intracortical electrodes can record from single neurons or small groups of neurons, demonstrating specificity and control unmatched by other recording modalities. This technology has been used successfully by non-human primates to control cursors in three dimensions using various decoding algorithms (Taylor et al., 2002; Carmena et al. 2003; Gilja et. al., 2012) and perform a reach-to-grasp task, enabling self-feeding (Velliste et al., 2008). Intracortical arrays have also been used in a handful of human studies, enabling paralyzed individuals to perform reach-to-grasp tasks, feed themselves, and control a virtual cursor making point-to-point reaches (Hochberg et al., 2006; 2012; Collinger et al., 2013).

While intracortical electrodes yield precise control, the implantation of these arrays disrupts brain tissue and vasculature, leading to a chronic inflammatory response hallmarked by

a glial scar around the probe and neuronal degeneration at the vicinity of the implanted electrodes (Polikov et al., 2005; McConnell et al., 2009; Biran et al., 2005). This ultimately causes signal deterioration, manifested as a reduction in the number of electrodes recording individual neurons and a decrease in signal amplitude (Barrese et al., 2013; Schwartz et al., 2006; Moran, 2010; Freire et al., 2011; Collinger et al., 2012; Schwartz, 2004; Simeral et al., 2011; Chestek et al., 2011). This results in decreased longevity of these electrodes and decreased signal stability across BCI sessions.

1.2.1.2 Electrocorticography (ECoG)

Electrocorticography (ECoG) is a recording modality where electrodes are placed either subdurally (below the dura) or epidurally (on the surface of the dura) to record electrical field potentials generated by aggregate cortical activity. Human and non-human primate subjects have demonstrated up to three-dimensional control of computer cursors or prosthetic limbs using ECoG (Leuthardt et al., 2011; Schalk and Leuthardt, 2011; Leuthardt et al., 2004; Schalk et al., 2008; Wilson et al., 2006; Yanagisawa et al., 2012; Wang et al., 2013; Hotson et al., 2016).

ECoG records the integrated activity of a much larger number of neurons that are in the proximity of the ECoG electrodes, and therefore does not match the spatial resolution of intracortical recordings. However, ECoG recordings are generally considered to have high longevity and stability since the placement of the electrodes does not compromise the integrity of cortical tissue. Additionally, the dependence on a large number of neurons is believed to prevent ECoG recordings from being sensitive to changes in the activity of individual neurons (Moran, 2010; Leuthardt et al., 2004). Despite the potential for ECoG in BCI applications, very few studies have evaluated the long-term tissue response to either epidural or subdural grids, and therefore the long term recording stability and safety are not well understood.

1.2.1.3 Electroencephalography (EEG)

Electroencephalography (EEG) is a neural recording technique in which electrodes are placed on the patient's scalp to record electrical field potentials from the brain. EEGs are minimally invasive, and therefore a good initial choice for brain computer interfaces. EEG has been used for two and three-dimensional cursor control (Wolpaw et al., 2000; 2002; Vaughan et al., 2003; Fabiani et al., 2004; Wolpaw and McFarland, 2004, McFarland et al., 2010, Foldes and Taylor, 2013), relating changes in the spectral power of electric field potentials in mu and beta frequency bands to movement parameters.

Since EEG electrodes reside on top of the skull, the distance of recording electrodes from the cortex results in low spatial resolution and low signal-to-noise ratio. This limits communication speed (bit/min) and therefore the performance of EEG-based BCI systems (Wolpaw et al., 2002). Still, the minimally invasive nature of EEGs means that electrodes are not susceptible to the tissue reaction believed to result in diminished recording quality over time for more invasive implanted microelectrodes, and thus may be considered a high-longevity option for BCI systems (Polikov et al., 2005).

1.2.2 BCI decoders

BCI systems interpret user intent by measuring and processing brain signals in real time. In order to control an external device, such as a cursor or robotic arm, recorded neural signals must somehow be translated into device commands. The origins of BCI decoders can be traced back to Fetz (1969), who showed for the first time that monkeys could be trained to volitionally modulate their neural activity using operant conditioning. This work opened the doors for the creation of BCI decoding algorithms, and in 1970, Humphrey et al. used neural activity recorded

on a handful of single electrodes and built the first BCI decoder. They recorded neural activity as animals moved a weighted joystick side-to-side by alternate flexion-extension of the wrist, and passed this data through their decoder offline. Comparing observed and predicted movements, they found that the accuracy of their decoder increased as they increased the number of simultaneously recorded neurons. This first BCI algorithm sparked the development of several mathematical models relating neural activity in M1 to movement parameters using several simultaneously implanted single electrodes (Georgopoulos et al., 1982; 1986; Chapin et al., 1999; Moran and Schwartz, 1999; Taylor et al., 2002), however the therapeutic potential of BCI was realized with the development of intracortical multielectrode arrays. Donoghue and colleagues were the first to implant Utah arrays in monkeys for BCI control (Serruya et al., 2002). Using a recording technique suitable for human use (Maynard et al., 1997; Kennedy and Bakey, 1998), their results provided support for the eventual use of BCIs to enable neurally based control of movement in paralyzed humans.

In the last few decades, intracortical multielectrode arrays have allowed us to record neural activity from 10s-100s of motor cortical neurons in real time. Many decoding algorithms have been proposed to translate neural population data into control signals for BCIs; exampes range from linear filters such as the population vector algorithm (PVA) (Georgopoulos et al., 1986, 1988) and optimal linear estimator (OLE) (Salinas and Abbott 1994) to various versions of Bayesian decoders such as the Kalman filter (Brockwell et al. 2004; Wu et al. 2006; Brockwell et al. 2007; Yu et al. 2007). The parameters for each of these decoders are calculated using neural activity from a calibration session. Calibration data is collected while the subject observes, imagines or preforms center-out movements (Carmena et al., 2003; Hochberg et al., 2006; 2012, Velliste et al., 2008; Collinger et al., 2012). This session allows experimenters to determine the

cosine tuning curve for each recorded neuron, and initialize decoding weights. Tuning curves are a function of firing rate based on reach direction. Since tuning curves are a construction of center-out reaches, current decoders often perform well for point-to-point movements, but have difficulty stopping and changing directions (Golub et al., 2014), limiting their usefulness in more complex task. This highlights the need for BCI decoding algorithms designed for and trained on realistic movements.

1.2.3 BCI feedback

Feedback is necessary for accurate movement control, be it by means of the natural sensorimotor system or a BCI. Today, all BCI systems use visual feedback for control; however, the majority of BCI's *only* use visual feedback. In contrast, able-bodied people gather information about their movements using both visual and non-visual forms of feedback. The absence of additional feedback modalities in BCI control forces users to maintain visual contact while performing tasks, which is especially problematic when visual attention is diverted and in tasks that do not require visual feedback, such as grasping.

In recent years, research groups have begun to incorporate additional feedback modalities to improve BCI control. Suminski et al. (2010) had monkey subjects receive congruent proprioceptive and visual feedback during BCI control. To achieve this, an exoskeletal robot moved the subject's arm in synchrony with the decoded cursor position. The congruent feedback significantly improved the BCI performance over visual feedback alone.

For patients with impaired sensory systems, vibrotactile stimulation and intracortical microstimulation (ICMS) are both candidate replacements for lost proprioceptive and cutaneous sensation. Vibrotactile feedback has been used to indicate sway for balance in patients with lack

of proprioception (Kadkade et al. 2003). For these patients as well as for prosthetics users with intact tactile sensations, tactile vibration can be used in place of proprioception during tasks that do not rely on vision, such as reaching for a glass out of view. For patients without tactile sensation, ICMS directly informs somatosensory cortex about arm/prosthetic sensations by bypassing damaged peripheral pathways. Studies using microstimulation on primary somatosensory cortex (S1) in monkeys (Mountcastle et al., 1990; Romo et al., 1998; 2000; Fitzsimmons et al., 2007; London et al., 2008; O'Doherty et al., 2011) and rats (Venkatraman and Carmena, 2011; Rebesco and Miller, 2011) have initiated research towards microstimulation-based feedback for neural prostheses. Romo et al. (2000) suggests that artificial and natural sensations can create perceptually similar sensations; monkeys could distinguish the higher of two applied frequencies, even when one was vibrotactile and the other ICMS. Conversely, Venkatraman and Carmena (2011) demonstrate that while whisker deflection and microstimulation of barrel cortex initially produce similar responses, they are not identical, and can be discriminated. Even if sensations are not perceptually identical, these studies suggest that ICMS can be used for similar applications in times when all tactile sensation is lost. For example, ICMS feedback can potentially be used in conjunction with BCI, enabling paralyzed individuals who have lost sensation or proprioception, to reach and grasp using a robotic arm.

A major challenge for BCIs that utilize ICMS is the interference caused by electrical stimulation artifacts on concurrent electrical recordings. Stimulation artifacts obscure ongoing neural activity, can be misinterpreted as neural activity, and can distort the shape of extracellular potentials, affecting action potential discrimination. Despite this challenge, scientists have demonstrated preliminary evidence of ICMS in monkeys as being useful for providing somatosensory feedback for upper limb neuroprosthetic devices. One method of artifact rejection

is multiplexing neuronal recordings and ICMS (O'Doherty et al. 2011). In O'Doherty et al. (2011), control signals derived from M1 were used to control a virtual-reality arm in a 2D environment while the artificial texture of different objects was conveyed to the animal via different ICMS patterns. These ICMS patterns instructed the animal to choose a specific object. In another study, Klaes et al. (2014) developed a bidirectional BCI by filtering the power spectrum of the recorded neural data to remove ICMS influence in real-time while monkeys used BCI control to complete a match-to-sample task. As ICMS becomes more commonly used for BCI feedback, choosing the optimal method of artifact rejection will be critical for BCI performance.

1.3 NECESSARY IMPROVEMENTS FOR CLINICAL BRAIN COMPUTER INTERFACES

Brain computer interfaces technology has the potential to improve the lives of individuals with severe disabilities, yet many of the striking achievements of the field have been made in the laboratory using able-bodied monkeys. Before BCI systems can be used as a clinical therapy, they must be able to reliably record neural signals for many years and enable users to make complex movements that interact with the environment.

1.3.1 BCI signal acquisition

All BCI systems depend on the sensors and associated hardware that record neural signals, yet recording longevity is limited, and the long-term effects of neural implants on underlying tissue is unclear.

Clinical BCI systems that utilize invasive electrodes should ensure several decades of usability. Extensive studies have been performed on Utah electrode microarrays in monkeys yielding viabilities ranging from a minimum of 3 months to greater than 1.5 years, where some electrodes lose signals early and others stay active for years (Williams et al., 1999; Nicolelis et al., 2003; Suner et al., 2005; Sharma et al., 2011). Signal loss has been associated with an immune response to the implant (Polikov et al., 2005), and several research groups are working towards combating this inflammation and scarring. Some groups have coated intracortical arrays with polymer coatings to improve recording capabilities (Kozai et al., 2015), whereas other have worked towards making arrays more flexible, as to mimic the dynamics of the brain and avoid scarring due to mechanical stress (Winslow et al., 2010). Some of the most promising systems utilize bioactive molecules, encouraging neural growth at the electrode (Kam et al., 2002; Cui et al., 2003; Singh et al., 2003; Szarowski et al., 2003). These methods yield increased neuron activity around the implant, yet the glial scars still form and limit the recorded signal. Many techniques have increased the longevity of neural recordings, but to date there is still not a solution to the impact of the immune response to neural implants. To increase recording longevity, more work must be done to characterize the tissue response using histology from non-human primate BCI studies.

ECoG based BCI systems are less invasive, and also have potential for bringing BCIs to the clinical realm, yet very little is known about long-term recording ability and the host tissue

response. A large portion of ECoG BCI studies are performed in epileptic patients while they are monitored after surgery. These studies are very informative, but are limited to 30 days due to FDA 501(K) regulations and do not provide insights into the long-term recording usability of this technology. Chronically implanted ECoG arrays in non-human primates circumvent this obstacle, yet very few studies have been conducted evaluating the long-term signal stability, recording ability, and tissue response of this technology.

1.3.2 BCI decoders

Clinically viable BCI systems must enable users to make complex, sensory guided movements that allow them to interact with their environment. To achieve this, we must understand M1 activation during these behaviors. As mentioned in Section 1.1.1, our current understanding of M1 is based on point-to-point reaches, which are much simpler than the majority of movements we make daily. To improve our understanding of M1, we must record multineuronal activity from M1 during ongoing, realistic behaviors. Many of these actions require sensory feedback, and therefore, we must also understand the neural mechanisms by which distinct sensory systems are integrated in the central nervous system and influence motor control. Recent studies show that sensory information activates M1 neurons, in many cases even more strongly than the motor command (Pruszynski et al., 2011; Hatsopoulos and Suminski, 2011; Pruszynski et al., 2014; Rao and Donoghue, 2014). However, in each of these experiments, only a single, brief sensory stimulus was administered. This leaves open the question of how M1 responds to the ongoing sensory information that guides real-world actions. A greater understanding of the underlying mechanisms of movement will enable us to design decoding algorithms that accurately represent neural function and thus, improve BCI control.

As we continue to build new and “better” decoders, we need a method to quantitatively assess decoder performance on realistic movements. With point-to-point reaches, various decoders show minimal differences in performance, with all decoders achieving high success rates. Current decoding algorithms, such as PVA and the Kalman filter, have difficulty stopping and changing directions (Golub et al., 2014). Point-to-point reaches only require an effector, such as a cursor or robotic limb, to move in one direction for a short period of time, and thus these decoders perform well on such tasks. For a BCI system to be clinically useful, patients must be able to perform everyday tasks that require interactions with the environment. These tasks require on-the-fly movement adjustments, which necessitates a BCI decoder’s ability to quickly change velocity. I believe that in the next few years, this ability will be what separates good BCI decoding algorithms from mediocre ones. Quantifying BCI performance using a task that requires a control signal to change directions quickly based on sensory information, such as an unstable system, will provide insight as to which decoders yield better control for realistic actions.

1.3.3 BCI feedback

BCI control currently relies solely on visual feedback. In contrast, natural muscle-based skills rely on diverse sensory inputs. Researchers have begun to explore additional feedback modalities to improve BCI control, but at this point in time, most artificial feedback methods are a proof-of-concept. We need to continue exploring additional methods of providing feedback, but also devise a way to quantitatively compare feedback modalities. As common as movement is to everyday life, there few tasks which can determine the effectiveness of someone's sensorimotor control. We need a task that pushes the sensorimotor loop to its limits. Such a task

would enable us to quantify someone's sensorimotor abilities, and also provide a means of comparing feedback modality effectiveness for BCI control.

In recent years, ICMS has gained attention as a means of providing sensory feedback to inform BCI control. Many paralyzed individuals suffer from spinal cord injuries, which cause damage to both efferent and afferent pathways. ICMS bypasses damaged peripheral pathways to directly informs somatosensory cortex about arm/prosthetic sensations. Most efforts to use ICMS to provide sensory information employ a biomimetic approach, which aims to recreate the patterns of neural activity that underlie natural somatosensation (Fagg et al., 2007; Choi et al., 2012; Daly et al., 2012; Weber et al., 2012). Truly biomimetic feedback is not currently attainable due to several technical and scientific considerations. These include the inability to access the full somatosensory neuronal population and an incomplete understanding of neuronal sensory encoding mechanisms (Daderlat et al., 2014). Several laboratories are recording from sensory cortices (Gardner and Costanzo, 1980; Warren et al., 1986; Johansson and Flanagan, 2009; Pei et al., 2010; 2011; Bensmaia and Miller, 2014; Saal et al., 2015) in hopes of gaining a deeper understanding of how natural sensations are represented and processed in the brain.

Biomimetic approaches to sensory substitution may be the eventual goal of BCI systems, however, alternative non-biomimetic have been shown to effectively provide sensory information by taking advantage of the inherent plasticity of the brain (Daderlat et al., 2014; Quick et al., 2015). Daderlat et al. (2014) demonstrate that a novel ICMS feedback signal can be used to continuously guide center-out reaches. In that study, movement vector direction was encoded on an array in S1 by the relative stimulation pulse rates across the electrodes. Within our own lab, we present proof-of-concept data for the usability of a non-biomimetic, non-visual artificial feedback to inform feedback driven reaches (Quick et al., 2015). In this study, monkeys

were instructed to stabilize a virtual, one-dimensional, unstable system using serial arm movements. During visual feedback, if the monkey saw the system drift right, the animal tried to stabilize the system by moving its hand to the opposite position on the left, and vice versa. During vibrotactile feedback, the animal did not see the system, but instead experienced the system via two vibrating motors on his non-reaching arm. Not only could monkeys perform the task with the non-biomimetic feedback, but they did so nearly as well as when they used visual feedback. Non-biomimetic artificial feedback is a first step towards incorporating non-visual feedback for BCI control, and has the potential to provide insight into how the brain adapts to a new source of feedback.

1.4 RESEARCH OBJECTIVES AND OUTLINE

The work in this thesis aims to facilitate the translation of BCI systems from the primate lab to a clinical setting. Much of the necessary research to bring BCIs into the clinical world will continue to rely primarily on animal studies before the initiation of human trials. Chapter 3 discusses a method for safe and durable head immobilization during animal experimentation. To address which BCIs are best for which purposes, we need a way of comparing BCIs and assessing performance. In Chapters 4, 5, and 8, I use a behavioral a task called the critical stability task (CST) that provides a way to quantitatively compare the effectiveness of different sensory feedback modalities and control methods. This task also enables researchers to probe the neurophysiological mechanisms of continuous, feedback-driven movements. Having a better understanding of the signals we're recording from will help us build better BCIs. Clinical BCI systems need hardware that is safe and that records neural signals reliably for many years. In

Chapter 6, we explore long term signal stability using electrocorticography (ECoG) in primary motor cortex and how long term implantation of this technology impacts the brain. In Chapter 7 we present data that suggests that the typical thresholds used in online decoding studies may discard information that could be useful for improving BCI performance. Finally, in Chapter 8, I discuss additional projects I worked on during my dissertation, including a project where I propose to understand how sensory feedback is represented in primary motor cortex. To begin, Chapter 2 describes the general methods used in Chapters 3 - 8.

2.0 GENERAL METHODS

In this chapter, I will describe the methods used throughout the experiments described in Chapters 4-8. In Section 2.1, I describe the electrophysiology data recorded from the non-human primates used in the following experiments. Sections 2.2 and 2.3 detail the motor control methods and feedback methods used in Chapters 4-8. Motor control methods include hand control and brain control, and feedback modalities include visual and vibrotactile feedback. In section 2.4, I discuss the critical stability task (CST). This task is used in Chapters 4 and 5, with non-human primates and humans. In all chapters utilizing non-human primates, experiments were approved by the University of Pittsburgh Institutional Animal Care and Use Committee.

2.1 ELECTROPHYSIOLOGY DATA

The data in Chapters 4 and in Chapters 6-8 were recorded from five male Rhesus macaques (Macaca mulatta), monkey I, monkey J, monkey K, monkey L, and monkey N. Behavioral data was collected from all monkeys, and analyzed to produce behavioral in addition to neural results. In all monkeys except monkey K, a 96-channel microelectrode array (Blackrock Microsystems) was implanted into cortex. Monkey I was equipped with an array in right hemisphere of dorsal premotor cortex (PMd), and neural data was used in Chapter 4 to expand on behavioral results. Monkey J had an array in the right hemisphere of primary motor cortex (M1), and contributed

behavioral results to Chapter 4, and neural results to Chapters 7 and 8. Monkey K had a custom 15-channel silicone micro-electrocorticography grid with 2mm platinum electrodes was placed over left hemisphere of premotor and motor cortex. Neural, histological and behavioral results for Monkey K are presented in Chapter 6. Monkey L had an array implanted in the right hemisphere of M1, and contributed to the data in chapter 7. Data from Monkey N was recorded from an array implanted in the left hemisphere of M1, and used as a control in Chapter 4.

2.2 MOTOR CONTROL METHODS

There were two methods of motor control used throughout the experiments of Chapters 4-8, hand control and brain control. During hand control, monkey subjects freely moved their hand in space in order to accomplish the task objectives. We attached a red LED to either the monkey's wrist or the monkey's index finger. The position of the LED marker was then recorded in three dimensions (3D) using a PhaseSpace, Inc. motion capture system.

During brain-computer interface control, termed `BCI control' or `brain control,' the monkeys controlled a cursor's movements using their neural activity. Brain control is used in Chapters 4 and 8. For brain control, neural activity was recorded via a Tucker-Davis Technologies recording system, which amplified and recorded the voltage from each of the 96 microelectrode channels. The spikes counts recorded from each channel were mapped to the brain-controlled cursor kinematics using a decoder. Monkeys I and J controlled velocity Kalman filter decoders that mapped spike counts to cursor velocity. The decoder used by Monkey I in Chapter 4 was trained on arm movements, and the decoder used by Monkey J in Chapter 8 was trained on observed movements.

2.3 FEEDBACK MODALITIES

There were two methods of feedback used throughout the experiments of Chapters 4-8. The primary method was visual feedback. For all monkey experiments, the animals were seated in a custom built chair that wheeled into and was secured fastened to the rig. In the rig, the monkeys were faced directly in front of a set of mirrors using a Wheatstone stereoscope configuration (Wheatstone 1838, Wheatstone 1852, Bacher 2009). When visual stimuli were presented on the mirrors using the Wheatstone configuration, to the monkey, the stimuli appeared to be in a 3D virtual reality (VR) environment. All visual stimuli were presented in the same fronto-parallel plane inside the 3D VR environment. When the monkeys performed hand control trials, the arm movements occurred behind the Wheatstone mirror configuration. The monkey's hand position was virtually rendered in the 3D VR environment as a round moving cursor.

The second method, used by monkeys in Chapter 4 and humans in Chapter 5, was vibrotactile feedback. For vibrotactile feedback, two coin tactors (Precision Micro-drives Limited) were attached to the monkey to communicate information in a non-visual manner. Tactors were placed on the monkey’s non-working forearm. One tactor would vibrate to indicate that the cursor was to the right of center, and the other would vibrate to indicate that it was to the left. The amount of tactor vibration was controlled via a supply voltage, and indicated how far the cursor was from the center of the workspace. As the supply voltage increased, the frequency and magnitude of vibration increased. An accelerometer was attached to the back of each tactor to measure the vibration.

2.4 CRITICAL STABILITY TASK

The central technical innovation of my research is the Critical Stability Task (CST), wherein subjects must utilize continuous sensory feedback (vision and/or tactile) in order to make corrective actions to maintain a drifting cursor at the center of a computer display for several seconds. Corrective actions to move the cursor are generated by the subject via hand movements, or via brain control. The CST was originally formulated (Jex et al. 1966) as a first-order system with unstable dynamics governed by the differential equation,

$$\frac{dx(t)}{dt} = \lambda\big(x(t) + u(t)\big),\ \lambda > 0,\ t \geq 0 \qquad (2.1)$$

where x(t) is the horizontal position of the cursor (with x > 0 representing the cursor to the right of center), u(t) is the horizontal position command generated by the subject, and t = 0 indicates the start of the trial. The system is characterized by the value of the parameter, λ. Because λ is positive, the cursor position will diverge over time unless appropriate compensating actions (control commands) are issued by the subject (Dorf 1967). As λ increases, the cursor diverges faster, and the CST is more challenging. The maximum λ (which has units of frequency) for which a subject can control the cursor is related to the bandwidth of the stabilized closed-loop system and is called the critical instability value (CIV) since, beyond that value, the closed-loop system becomes unstable. The greater the CIV, the more effective is the subject's sensorimotor control.

To appreciate how the CST allows the investigation of motor control that requires sensory feedback, consider the position command that is required to keep the cursor from drifting, i.e., what is u(t) such that the cursor velocity is zero? Plugging this condition $\left(\frac{dx(t)}{dt} = 0\right)$ into Eq. (2.1) shows that the subject must generate a position command that is exactly equal

and opposite to the current cursor position: $u(t) = -x(t)$. Thus, feedback is essential: the subject must know the state of the cursor. A subject who could issue this perfect motor command (instantaneous and 100% accurate) would be able to keep the cursor completely still.

In practice, this perfect control command is not possible because of sensorimotor noise and nervous system processing delays; hence, the cursor will continuously move, and maintaining control – keeping it from rapidly drifting off-screen – requires an ongoing sensory-guided motor action. With adequate sensory feedback, a subject can generate a command signal that keeps the cursor centered about an equilibrium point for λ values that are less than the subject's CIV. A subject's CIV captures important aspects of their sensorimotor skill; it is related to the speed and accuracy with which he or she can react to sensory input. The greater a subject's CIV, the more effective is their sensorimotor control. Thus the CIV is a single number that provides an objective, empirically-derived measure of performance. The CIV depends on several interesting factors, including the type of feedback (visual or tactile in our experiments), the quality of sensory information (which we can control), the nature of the effector (hand movements or BCI control in our experiments), and the subject's level of skill. A slower and less reliable sensory feedback modality or motor effector will yield a lower CIV than a faster and more reliable feedback modality or effector.

The CST is used in Chapters 4 and 5, with monkeys and humans respectively. Future directions for neuroscience research using this task are detailed further in Chapters 8 and 9. We are the first to use this task for non-human primate experiments and plan to record neural activity during the CST within the following year.

3.0 METHYL-FREE HALO-STYLE HEADPOSTS FOR STABILIZING BEHAVING MONKEYS

Figures and text in this chapter are from my first author manuscript that is ready to be submitted to the Journal of Neuroscience Methods. In this chapter, I describe our method of non-human primate head fixation for behavioral neuroscience experiments. I am the sole author of this manuscript, and in addition to documenting our techniques I have revised the surgical procedure for implanting the headposts and helped us move toward completely methyl-free surgeries.

I would like to acknowledge Dr. Hank Jedema from the Bradberry Lab, who helped develop the head stabilization technique we use based on the Thomas Recording Precision Positioning System. I would also like to acknowledge Andrew Holmes for machining our titanium headposts.

3.1 INTRODUCTION

Neural recordings in awake, behaving monkeys (Macaca mulatta) have led to numerous advances in understanding the neural basis of perception, cognition, and movement. During such experiments, it is often beneficial to brace the animal's head to ensure complete stability during recording or stimulation. Head movements made during a procedure can cause the brain to move relative to the electrode tip, destabilizing the recording, and potentially damaging to the cortical tissue. Eye tracking also necessitates head stabilization.

A variety and combination of materials, including stainless steel, titanium, and dental acrylic (methyl methacrylate) have been used to develop devices for head fixing nonhuman primates. The most widely used head immobilization method in the neuroscience community involves the construction of an acrylic headcap, housing the headpost and recording chambers. While this procedure provides a hard casing around the implanted hardware, dental acrylic is osteo-phobic and thus does not adhere to the skull sufficiently to be used alone. To circumvent this lack of bone-implant interface, metal screws are inserted into the skull surface, with a portion of each screw protruding to anchor the acrylic. Despite these anchor points, acrylic causes bone softening and necrosis, which can in turn cause the anchor screws to become loose within the bone. These complications cause an eventual failure of the entire implant. Spontaneous failure of such implants has scientific disadvantages, chiefly disruption of data collection. Foeller and Tychsen (2002) used these acrylic headcaps for head fixation in 20 animals and observed an average time to failure of 11 months (ranging from 5 months to 2 years) over the course of experimentation. Occasionally the bone cannot support another implant, which can end data collection in that animal sooner than might otherwise be possible. Another drawback to the use of an acrylic cap is that it requires a large area of the skull to be covered

with acrylic, leaving other brain areas less accessible for future experiments or recordings. The large cap also results in a “shoreline” between the scalp incision and the implant, requiring regular and thorough maintenance to abate infection. Since the acrylic bonds poorly to healthy bone, a capillary gap between the cap and skull surface can arise, providing an ideal environment for bacterial colonization. Once infected, the bone becomes necrotic and eventually expels remaining screws and causes failure of the implant. Even in the absence of infection, removing the cap after the completion of an experimental protocol involves a long, tedious surgery which involves drilling the original acrylic. Even when a well-maintained implant does last for years, softening of the bone due to the poor bone-implant interface may make it unsuitable for subsequent implants.

To overcome these problems, we and others have developed alternative methods of head fixation (Batelak et al., 2001; Isoda et al., 2005; Adams et al., 2007; Davis et al., 2007; Adams et al., 2011). Replacements for acrylic headcaps hold promise for shorter and easier implant surgeries, simpler explant surgeries, longer lifetimes of implants, and an improved prognosis for future implants. Additionally, a more biocompatible headpost device would improve the quality of life for monkeys used for neuroscience research. Most recently, techniques developed by the Horton lab using commercially pure titanium have paved the way to a minimally invasive, durable, and biocompatible head fixation device that is easy to implant and to use in daily experiments. These techniques circumvent the use of acrylic completely, thereby avoiding the bone softening and necrosis associated with its use. Our lab uses their findings as the basis of our design, however instead of using one, long, load bearing post, we use three shorter posts positioned around the top of the skull to be used with a halo. Each post has a titanium base with six to 12 holes. The bases are bent to conform to the shape of the skull, leaving little to no room

for infiltration of bacteria and tissue growth. The bone-implant interface is strengthened to the point of osseointegration, in which the bone envelops the implant, strengthening the stability further. By combining the halo head holder designs with a biocompatible titanium implant, we have achieved a simple, clean, streamlined, stable, and durable method of head immobilization. It is reversible under anesthesia if required, however removal may be more difficult once osseointegration has occurred.

3.2 METHODS

Our objective was to provide a robust and stable method of head fixation in macaques without the use of dental acrylic. The essence of our design is a three-point "halo" system, with pliable titanium posts attached to the animal's skull, an aluminum halo, which attaches to the posts, and then a head holder, which braces the head to the chair.

3.2.1 Headposts

We designed three short titanium headposts that protrude from the skull enabling head fixation while providing maximum unhindered skull real-estate for experimental purposes. These headpost designs are presented in Figure 3.1. Commercially pure titanium was used because it is readily osseointegrated into healthy bone tissue, and because it is pliable, enabling the surgeons to quickly customize to the skull curvature. (Adams et al. 2007; Brånemark et al. 1969; Linder et al. 1983). These headposts are machined from a single piece of titanium for maximal stability. We use either the “Y”-shaped (Figure 3.1a) or “T”-shaped (Figure 3.1b) footplates designed for

attachment to the skull with 12 bone screws. Each leg of the foot plate is 0.700 inches long, 0.178 inches wide, and 0.053 inches thick. Each leg contains three mounting holes that are countersunk to reduce the profile once the headposts are attached to the skull. The “Y” and “T” footplate options enable headpost placement individualization for various recording procedures. The ball at the tip of each post is used as the attachment point to the halo and together enable head fixation.

We use three radially spaced posts rather than a single post for several reasons. Using multiple points of attachment, we reduce the stress on the skull at any one attachment point. This lowers the likelihood of device failure due to tension on any individual post by distributing the force across several posts. The halo setup also enables the use of short posts, measuring 0.65 inches from skull to tip. This creates a low profile of hardware on each animal, thus not inhibiting normal movement in the cage and reducing the chance of damage in the cage (see results, Figure 3.6).

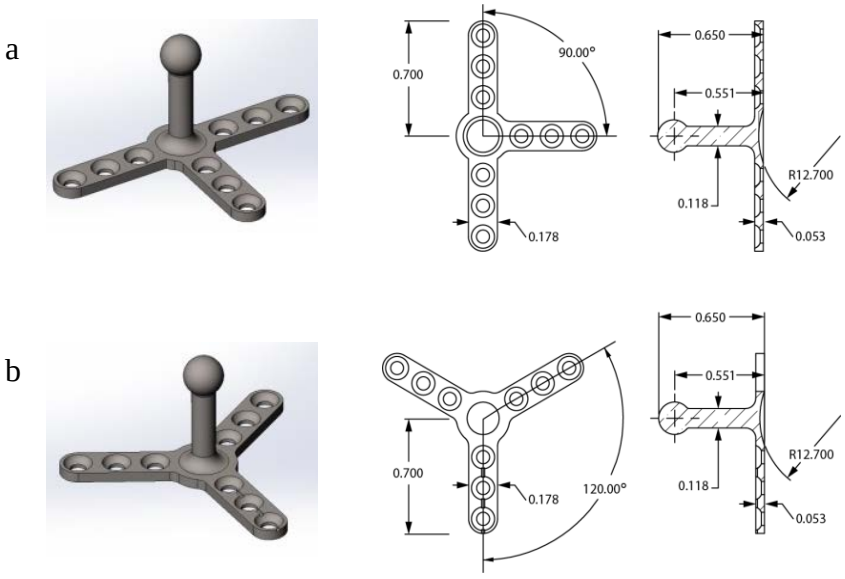

Figure 3.1: Headpost designs for headposts. Headpost with the a) “T”-shaped footplate and b) “Y”-shaped footplate. Dimensions remain consistent for both footplates and are in inches.

3.2.2 Halo

The halo attaches to the headposts, enabling head fixation. To build the halo, a three-dimensional model of the halo frame was created using Autodesk Inventor software and manufactured through eMachineShop. The halo has a thickness of 0.25 inches, inner diameter of 5 inches, and outer diameter of 6.75 inches (see Figure 3.2). The halo has one rigid offset post at the front that attaches to the brow ridge headpost, and two sliding offset posts in 3.846 in. channels to adjust to multiple headpost angle combinations (see Figure 3.3). Each offset post consists of a stainless steel block and an adjustable socket screw. These screws can be tightened or loosened to fit each headpost configuration, and to snugly attach to the three headposts. These screws are adjusted using a hex driver fit to the exterior end so that they can be manipulated by the experimenter at a safe distance from the animals head. Three holes with standard 1/4 in.-20 threads around the halo serve as attachment points to the halo holder on the primate chair, which is detailed in the next section.

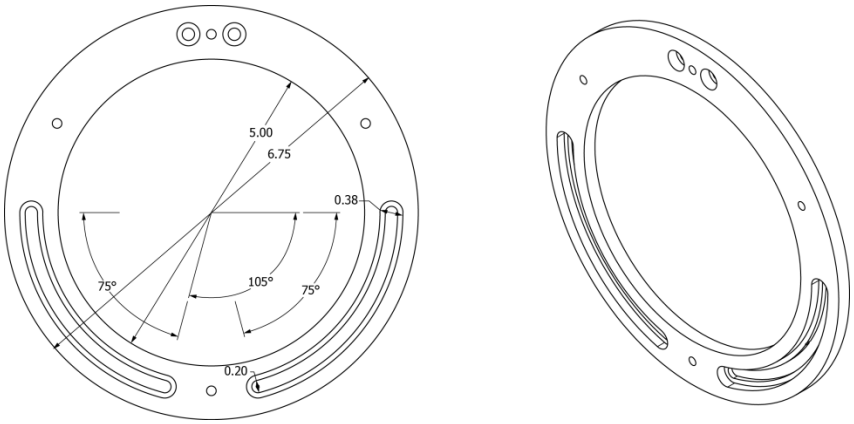

Figure 3.2: Machine drawing of the halo for head fixation. Dimensions are in inches.

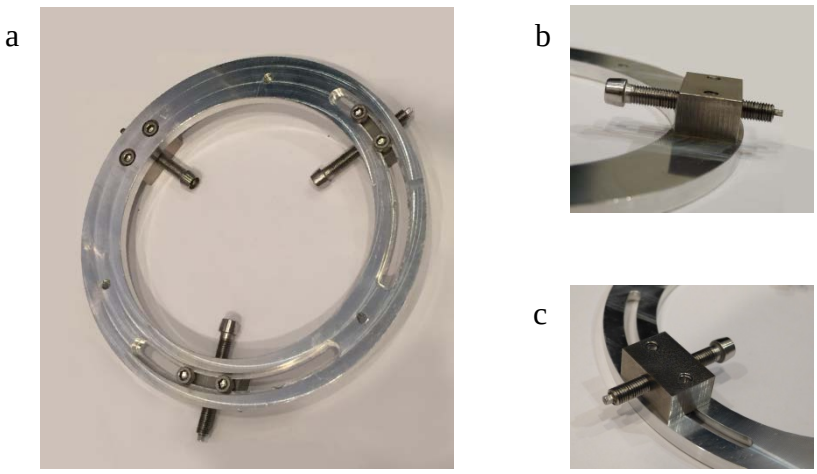

Figure 3.3: Halo and offset posts. a) Halo with offset posts and holder attachment holes. b) Fixed front offset post. c) Sliding side offset post.

3.2.3 Halo Holder

The halo holder is the device that attaches the halo to the primate chair, fixing the head. It consists of a rigid base that firmly attaches to the chair (see Figure 3.4). This base connects to the halo attachment ring via a sliding aluminum fitting that can adjust for the desired seated position of each animal. That is, the halo attachment ring can be fixed closer to the back of the chair for animals that sit far back, or can be extended away from the base for animals that prefer to sit forward. The halo attachment ring is attached to the base with three degrees of freedom by means of a ball and socket joint, so that it can be adjusted for the animal's comfort. The socket around the spherical bearing can be loosened or tightened to adjust the animals head position and fix it in place respectively.

Once the animal is in the chair, the experimenter secures the halo onto the headposts and the halo holder is lowered to the desired height and position for that animal. While the halo

holder is lowered, it is important to have the socket joint tightened so that the halo attachment ring does not hit the animal. Once the halo holder is in place, the experimenter loosens the socket using two screws, and attaches the halo attachment ring to the halo using three 1/4 in.-20 screws. While the socket is still loose, the experimenter can manipulate the animals head position by holding onto and maneuvering the halo. Once the animal is in a comfortable working position, the two screws are used to tighten the socket around the ball joint and fix the animals head in place.

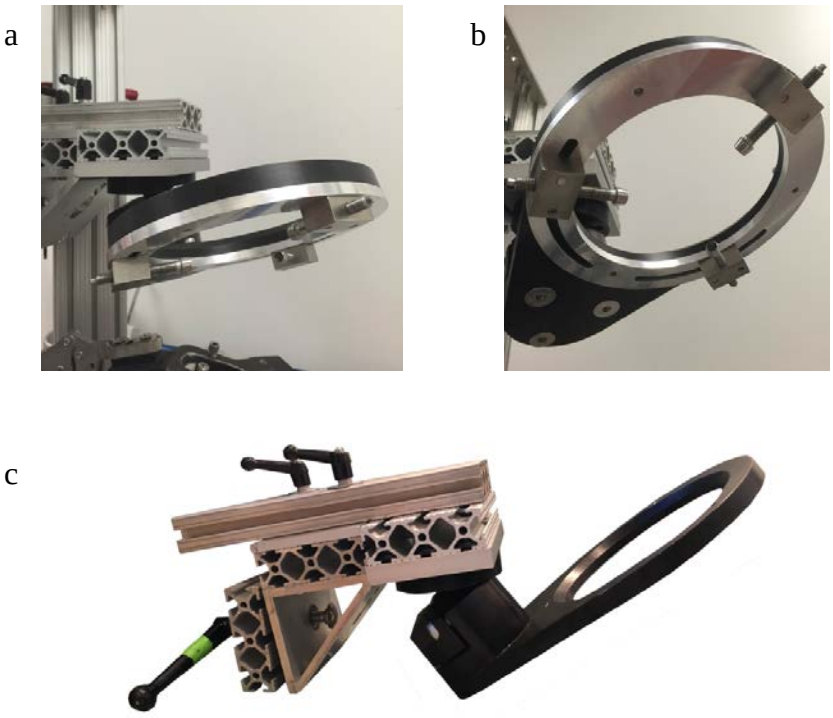

Figure 3.4: Halo holder attachment to primate chair. a) Halo holder attached to primate chair with the halo attached to it. b) Close up of how the halo attaches to the halo holder. c) Stand-alone halo holder device

3.2.4 Implantation Surgery

All surgical procedures have been approved by the University of Pittsburgh's Institutional Animal Care and Use Committee. The procedures meet the standards for humane animal care and use as set by the Animal Welfare Act and the NIH Guide for the Care and Use of Laboratory Animals. The animal facility is subject to voluntary inspections and is AALAC accredited.

During implantation, the animal is sedated with an injection of ketamine (15mg/kg) and transferred to a table for surgical preparation. There, he is intubated with an endotracheal tube of appropriate size and anesthesia is maintained with isoflurane (1-3%). The implantation site is prepared by shaving all hair from the head, including brow ridge and tops of pinnas, to the base of the skull. This provides the largest possible area for the surgeons to choose appropriate implant locations. The animal is then transferred to the facility dedicated surgery suite and placed into the stereotaxic frame. Betadine is used to disinfect the scalp and the monkey is draped sterilely.

Surgeons determine and mark locations of two medial incisions with a skin marker. The frontal incision is made coronally about 5 cm posterior to brow ridge, which is about 2 cm posterior to the location of the post. Connections between the fascia and skull are detached gently using periosteal elevators, and the tissue is secured out of the way. Care is taken not to damage the periosteum, the film-like substance that covers the surface of the skull. Many surgeons choose to remove this tissue by scraping it off of the skull surface as it interferes with adherence of traditional acrylic headcaps. It is not advisable to scrape the periosteum from the bone surface when using biocompatible implants, as it contains osteoprogenitor cells that aid in bone regeneration.

After the skin and fascia have been loosened and secured, the flat edge (for a T-shaped post) or distal end (for a Y-shaped post) of a headpost base is placed as close to the brow ridge as possible. With the headpost placed in the desired location and position, the post is under the skin, anterior to the incision. A pencil is used to mark this location on the skull, before headpost is removed for bending in order to match the shape of the skull. Subtle bends are made in the base straps, between the screw holes using various tools. Needlenose pliers are sufficient for this purpose, although specialized tools exist (Figure 3.5a). The headpost is replaced frequently on the skull within the pencil marks to ensure it was curved properly. It is not critical to bend the strap to exactly match the skull surface because it will continue to bend it is being screwed to the skull. We find that a tolerance of 1mm or slightly more is adequate.

When the base of the headpost conforms to the shape of the skull, the first drill hole is marked while the headpost is still in place. We begin with the screw-holes closest to the post. It is then removed and a hand drill is used to tap the first hole. Drill bits with a pre-set cuff (Synthes 310.141) prevent the drill from being inserted deeper than the skull. Bits range from 4mm to 12mm in length; however we find that a 5mm drill bit is most often used on our adult Rhesus monkeys. Holes are drilled at a slow and steady rate, since we believe heat can be generated by drilling too quickly, which can lead to bone degeneration. The surgeon may feel a slight resistance just before the drill bit reaches the inner table of the skull. For our purposes, the hole is drilled entirely through the skull, allowing the surgeon to visualize a pulse after flushing the newly drilled hole with saline (0.9% NaCl). The headpost is then replaced on the skull and a self-tapping titanium screw typically of 5 or 6mm length (Synthes 401.155 -156) is placed through the strap hole and into the skull (Figure 3.5b). It is important to insert the screw perpendicular to the skull, just as hole was originally drilled. Screws that are angled may not fit

into the designated slots and can cause a gap that is prone to tissue infiltration. We believe that in this scenario, the body may reject the screw and affect the integrity of the implant. We tighten the screw so that the head of it is touching the surface of the strap, but do not tighten it all the way at this time. Working our way around the post, we repeat the drilling and screw insertion with the headpost in place, starting with the most proximal strap holes and leaving the most distal holes for last (Figures 3.5c). When all nine screws are in place, they are tightened, working around the proximal perimeter then moving distally until all are tight. This area is then packed with gauze soaked in saline to keep the tissue moist until the skin is sutured at the end of the procedure.

The second incision is made coronally and is longer, so that it can incorporates both of the posterior posts. Skin and fascia are detached from bone and retracted from the skull, as performed in the previous incision. If temporalis muscle is present and preventing placement of the headpost on the skull, we blunt dissect it. The objective is to place the headposts out of the way of areas that will be utilized for neural recordings or locations of other implants. For our experiments, it is desirable to place the flat side (T), or distal ends of straps (Y) against the occipital ridge, leaving the maximum amount of space possible for future implants. We also use a halo template to make sure that the headpost placement is compatible with our restraint device. The halo template is the same dimensions of the halo, but machined from plastic.

The bending, drilling, and insertion of screws are repeated as above for both posterior headposts. For the anterior post, the skin anterior to the implant is stretched posteriorly over the headpost. A new, very small incision is made directly over the ball of the post. This incision should be just large enough for the post to be inserted through to the exterior of the skin (Figure 3.5d). Then, the initial anterior incision is closed. The fascia is sutured (4-0 Monocryl) before

suturing the skin (3-0 Monocryl). This procedure is repeated for both of the posterior headposts (using the skin posterior to the posts) and the posterior incision is then closed (Figure 3.5e). We may apply triple antibiotic ointment to the suture lines as well as the skin around the protruding posts. For animals that have a tendency to remove sutures prematurely, a layer of vet bond (3M Vetbond 1469SB) may be applied to provide an additional barrier, thought with the placement of the sutures away from the headposts, animals appear to be more interested in the posts than in the sutures.

The animal is then removed from anesthesia, extubated, and transported to his home cage for recovery. He is monitored until he is able to sit up and take food. At the conclusion of surgery, post-operative pain medication (buprenorphine 0.01-0.02mg/kg) is administered and he is maintained on a course of antibiotics (amoxicillin clavulate 20mg/kg) for a period of seven days.

3.3 RESULTS

We have used this head fixation procedure on 11 Macaque monkeys reported in Table 1. Previously trained behavioral tasks may be resumed a few days after surgery, however we allow 4-6 weeks of healing time before attaching the halo and restraining the head. Animals show no signs of discomfort or pain while in their cages or during training. Animals are only attached to the halo during training, and the headposts do not appear to be cumbersome to the animals or inhibit normal movement in their cages. The wounds at post site and the two main incisions heal with minimal scarring. Headposts have remained stable for the duration of experimentation for all eleven animals since the initial attachment procedure.

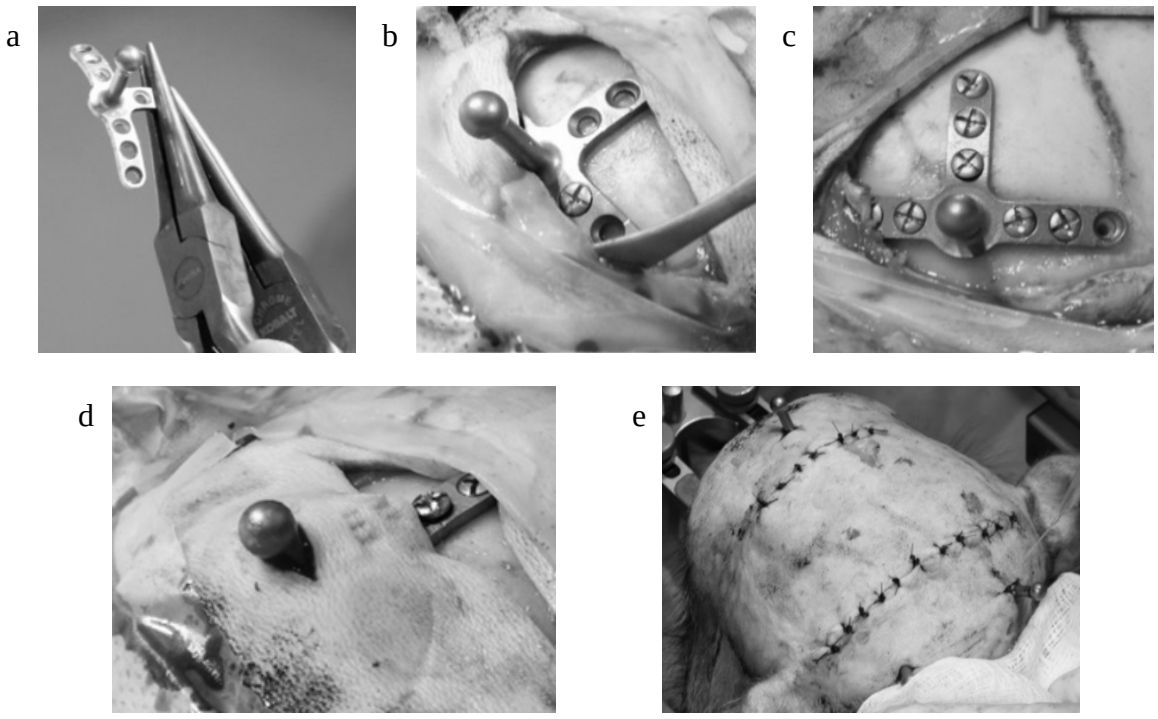

Figure 3.5: Headpost implantation surgery. a) Once a location for the headpost is chosen, the footplate is bent to match the curvature of the skull. b-d) The headpost is attached to the skull one screw at a time. Once the first screw is secure, we repeat the drilling and screw insertion working our way around the post from the most proximal strap holes to the most distal. e) Once the anterior post is secure, we pack the area and make the posterior incision. We use a template to ensure that our headposts will work with our restraint device. f) Before closing the main incisions the skin is pulled over the post. A very small incision is made directly over the ball of the post and the post is inserted through to the exterior of the skin. g) At the end of surgery both incisions are closed.

Table 3.1: Head fixation information for eleven animals with headposts

Monkey	Sx Date	Head fixation Date	Sx Weight	Current Weight*
D	8/5/2014	2/6/2015	7.6 kg	9.6 kg
E	6/17/2014	7/30/2014	7.6 kg	8.6 kg
F	11/6/2013	1/13/2014	7.0 kg	10.8 kg
G	11/27/2012	1/14/2013	7.0 kg	**
H	7/11/2011	9/6/2011	6.2 kg	9.1 kg
I	8/26/2010	11/5/2010	7.0 kg	9.5 kg
J	12/15/2009	1/26/2010	6.8 kg	13.1 kg
K	10/15/2008	1/14/2009	7.0 kg	13.3 kg
L	7/15/2009	9/9/2009	4.5 kg	7.2 kg
M	7/1/2009	8/20/2009	6.6 kg	15.2 kg
N	7/9/2013	8/14/2013	6.4 kg	10.1 kg

*Current weight or weight when deceased **Animal was transferred to another lab

a  b  c 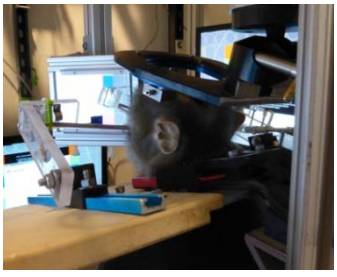

Figure 3.6: Daily appearance and use of the headpost-halo system. a) Minimal appearance of headposts during normal daily life. b) Halo attachment to the implanted headposts. c) Animal performing behavioral task with head fixed.

To demonstrate the stability of this head fixation device, we placed an animal in the primate chair for “head-free” and “head-fixed” conditions. The halo was placed on the animal while sitting in the primate chair, and an accelerometer (Freescale Semiconductor, MMA7361L) was attached to both the halo and the chair. The accelerometers were recorded from for three minutes while the monkey sat in the rig. During the “head-free” condition, he remained unattached to the halo holder and primate chair, and in the “head-fixed” condition he was fully attached. Figure 3.7 shows that with our device in place, the head remains stationary.

Not only do we circumvent bone softening and necrosis, we observe osseointegration during subsequent surgeries after periods of time with our headposts. This is shown for a few animals in Figure 3.8, and we have seen it in most animals implanted. The bone envelops the implant, strengthening the bone-implant interface and the overall stability of the device. We have seen osseointegration of the headposts both during additional experimental surgeries and post mortem. We do not know the exact timeline that bone growth occurs after the implant, however we have seen it during subsequent surgeries as early as seven months after implantation.

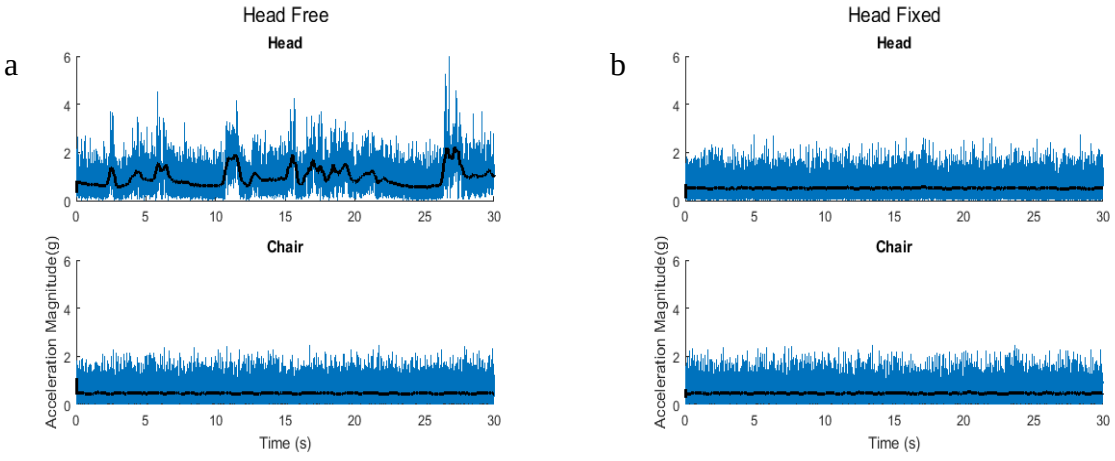

Figure 3.7: Accelerometer traces for a) head free and b) head fixed conditions. The black trace shows the moving average of the acceleration in 400ms bins.

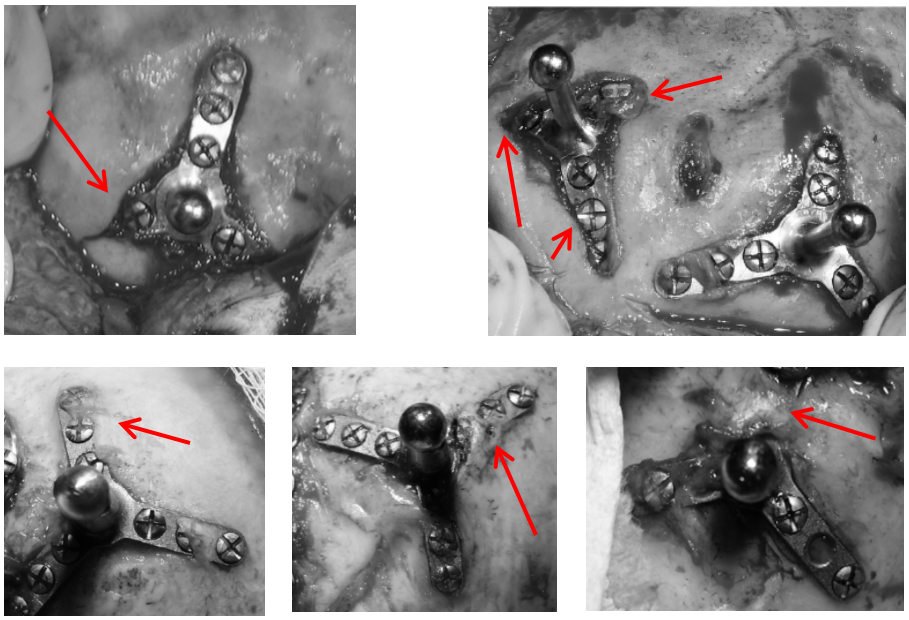

Figure 3.8: Examples of osseointegration between the skull and the headposts. The top row shows intraoperative pictures from Monkey M, three years after the headposts were implanted. The bottom row shows intraoperative pictures from Monkey K, four and a half years after implantation.

3.4 DISCUSSION

This method of head fixation for nonhuman primate research offers a stable, methyl free alternative to common “head-cap” approaches. This method has proven stable in at least eleven animals and avoids the loosening and bone necrosis issues that plague methyl alternatives. This device performed well for neural recordings and stimulation during behavioral tasks and in conditions where eye tracking was required.

A shift away from methyl headcaps has the potential to make studying the brain easier and safer. Traditional acrylic headposts are prone to infection and often require daily maintenance. In the event of a chronic infection, the chance of eventual failure of the headcap is increased. Our headposts require minimal maintenance, as that the skin heals around each post.

While the wound margin varies from animal to animal, we spend significantly less time performing margin maintenance than average for acrylic headcaps. By avoiding the use of acrylic, we have experienced zero headpost failures. When acrylic headcaps fail, additional surgery is required, halting data acquisition and putting stress on the animals. If the head becomes free during a recording session, not only is that session ended but damage to the electrodes and/or brain can lead to much larger consequences. By reducing the frequency of complications, we are ensuring that data collection remains consistent and that we get the most out of each research animal.

Methyl free head fixation is an emerging technique in neural physiology labs. As primate research advances, researchers are becoming aware of the disadvantages associated with using dental acrylic. Research in awake behaving monkeys not only requires head fixation, but also requires the implantation of a recording chamber that allows electrodes to be inserted into the brain or the implantation of electrode arrays and a connector mounted to the skull. Even if acrylic is still used for chamber or connector implantation, the size of the implant is reduced substantially when the headposts are fixed independently without acrylic. Our halo design provides ample real estate for hardware implantation independent of headpost fixation. With more labs moving towards methyl free techniques, new systems for safe and stable skull mounted recording techniques are sure to develop.

4.0 RHESUS MACAQUE PERFORMANCE ON CRITICAL STABILITY TASK

This chapter describes the work detailed in Quick et al. 2016, a manuscript that is ready for submission to the Journal of Neurophysiology for which I am second author on this manuscript. In this chapter, I show how the Critical Stability Task (CST) can be used as a tool in monkey subjects to measure natural and BCI sensorimotor control. Dr. Kristin Quick will be the first author for the manuscript describing the following work. The CST was hers brain-child, however my work expanding on her idea has helped the neuroscience community accept the CST as an alternative to the classic center-out task for studying motor control. The work presented below encompasses the bulk of my PhD thesis; I collected every data set used for analysis, trained both animals to perform the CST using visual and vibrotactile feedback, calibrated/troubleshot the hardware used for experiments, developed our method of measuring a "no feedback" condition with behaving monkeys, created the current version of figures for the manuscript, and reworked the manuscript text after Kristin graduated. My work has built a strong foundation for using the Critical Stability Task to test the effectiveness of different intracortical microstimulation feedback mappings and the effectiveness of different BCI decoders.

4.1 INTRODUCTION

Multisensory processing is vital for daily activities such as walking, reaching for and grasping objects, and interacting with the environment in general. Traditional approaches in behavioral neuroscience sought to isolate sensory modalities from one another (Andersen and Mountcastle 1983; Jay and Sparks, 1987), and to isolate sensory signals from motor commands (Funahashi et al., 1989). These approaches built an important knowledge base of sensorimotor action, yet much remains unknown about the mechanisms by which continuous multi-sensory feedback is integrated to influence ongoing motor control. Closing this knowledge gap could lead to advanced prostheses and aids for persons with sensorimotor impairments. For example, a brain-computer-interface (BCI) is a promising technology, still in development, the objective of which is to enable physically paralyzed individuals to interact with the environment in as dexterous and flexible a manner as do able-bodied individuals. However, unlike our inherently multisensory natural interactions with objects and the environment, a major limitation of current BCI technology is that sensory feedback is unimodal: in the vast majority of BCI devices, users rely solely on visual feedback to guide their actions.

The development of non-visual feedback for BCI, especially to provide touch and proprioception, is an active area of research (Weber et al. 2012). These various “sensory substitution" methods include intracortical microstimulation (O'Doherty et al., 2011; Dadarlat et al., 2015), spinal cord stimulation (Weber et al. 2011), peripheral nerve stimulation (Dhillon and Horch, 2005; Horch et al., 2011; Ledbetter et al., 2013), and vibrotactile (Cincotti et al., 2007; Chatterjee et al., 2008; Rombokas et al., 2013; Leeb et al., 2013; Lin et al., 2015) and electrotactilefeedback (Kaczmarek et al., 1991; Tyler et al., 2003; Ptito et al., 2005). In addition to the different feedback modalities themselves, sensory substitution methods also have a

number of design parameters to select. For example, with vibrotactile feedback, the vibrations can encode the position of an object (or the effector), or its velocity, or some combination of position and velocity (Goodworth et al., 2009; Kadkade et al., 2003; Wall and Kentala, 2010). In addition, there is the open question of how this information should be encoded. How do we determine which stimulation frequencies, amplitudes, and spatio-temporal patterns are the best for feedback? With such a large combination of possible feedback modalities, encoding schemes, and information to encode, the task of comparing the different possibilities is daunting. There is a need for flexible yet efficient and systematic ways to quantify and compare the effectiveness of various forms of sensory substitution options.

We introduce a new animal behavioral task that involves prolonged motor actions for which continuous sensory feedback is essential. We adapted the Critical Stability Task (CST) from the human performance literature (Burke et al., 1980; Jex et al., 1966; Kadkade et al., 2003; Kondraske et al., 1984; Potvin et al., 1977) to test the ability of Rhesus macaque monkeys to control a virtual object utilizing different sensory feedback modalities. In our implementation of the CST, monkeys were required to keep a horizontally moving cursor near the center of a computer monitor (within fixed (unseen) boundaries), by generating hand movements in opposition to the drifting cursor position. The dynamics of the cursor were inherently unstable (i.e., divergent), in that inadequate or erroneous hand position commands result in the cursor drifting out of bounds. The CST cannot be performed without continuous feedback about the cursor position. Hence, the CST allows one to directly quantitatively evaluate the effectiveness of different sensory feedback modalities.

We trained Rhesus macaques to perform the CST utilizing either continuous visual feedback or continuous vibrotactile feedback about the cursor state. The monkeys learned to

perform the task under each feedback modality. Their performance with their natural visual feedback was superior to that with the more abstract and unfamiliar vibrotactile feedback. Even so, their performance with vibrotactile feedback was 47-62% as effective as their performance with visual feedback, highlighting the potential for this sensory substitution modality, as well as the utility of the CST for systematically evaluating different feedback modalities in continuous motor action. Additionally, we demonstrate the ability of CST as a general tool to test the effectiveness of any `effector-feedback' pairing. With the eventual goal of testing BCI control with intracortical microstimulation feedback in mind, we determined a monkey's ability to use a BCI decoder with visual feedback. We found that using the BCI decoder as an effector was 22.7% as effective as the use of natural hand movements.

4.2 BACKGROUND: THE CRITICAL STABILITY TASK (CST)

The CST was originally developed to study sensorimotor performance limits in pilots (Jex et al., 1966) and it has been used since to study neurological health (Kondraske et al., 1984; Potvin et al., 1977), mental workload (Burke et al., 1980), drug impairment of executive function and motor control (Ramaekers et al., 2006), and to assess vibrotactile feedback for balance assistance (Kadkade et al., 2003). To gain intuition for this task and the sensorimotor actions required to perform it, think of manually balancing a ball on top of an inverted bowl held in your hands. The apex of the inverted bowl is an unstable equilibrium point: small random movements will cause the ball to roll down the side. Keeping the ball at the top requires careful adjustments, namely, well-timed movements of the bowl to counter the motion of the ball. Sensory feedback about the current state of the ball is essential to keep it from rolling off the top. Success at the task is

dependent on the sensorimotor abilities of the individual, such as the quality of their sensory feedback and their motor skill, among other factors (e.g., the curvature of the apex).

The CST simulates a motor control task akin to balancing a ball at the top of an inverted bowl, but for which the dynamics are simpler and under the experimenter's control. In the CST, the object that the user must balance is a virtual system consisting of a cursor on a computer monitor that will rapidly drift out-of-bounds without appropriate corrective sensory-guided action by the user. The velocity of the cursor, and hence the difficulty of the task, is set by a single parameter that enables one to quantify sensorimotor performance limits. The physics are simplified relative to balancing a ball at the top of an inverted bowl, in that the CST operates in one dimension in our implementation, and the dynamics are first-order, meaning that only the position of the cursor (rather than its position, velocity, acceleration, etc.) must be known in order to control the cursor. The virtual environment allows for such simplifications while maintaining aspects critical to sensory-guided motor control.

As originally proposed by Jex et al. 1966, the CST is a first-order system with unstable dynamics governed by the differential equation,

$$\frac{dx(t)}{dt} = \lambda\big(x(t) + u(t)\big),\ \lambda > 0,\ t \geq 0 \qquad (4.1)$$

where x(t) is the horizontal position of the cursor (with x > 0 representing the cursor to the right of center), u(t) is the horizontal position command generated by the subject, and t = 0 indicates the start of the trial. The system is characterized by the value of the parameter, λ. Because λ is positive, the cursor position will diverge over time unless appropriate compensating actions (control commands) are issued by the subject (Dorf 1967). Because these control commands must be functionally dependent on the cursor state in order to keep the cursor from diverging out-of-bounds, the CST can be envisioned as a closed-loop (feedback) system (Figure 4.1). As λ

increases, the cursor diverges faster, and the CST is more challenging. The maximum λ (which has units of frequency) for which a subject can control the cursor is related to the bandwidth of the stabilized closed-loop system and is called the critical instability value (CIV) since, beyond that value, the closed-loop system becomes unstable and the cursor diverges out of bounds. The greater the CIV, the more effective is the subject's sensorimotor control.

To appreciate how the CST allows the investigation of motor control that requires sensory feedback, consider the position command that is required to keep the cursor from drifting, i.e., what is u(t) such that the cursor velocity is zero? Plugging this condition $\left(\frac{dx(t)}{dt} = 0\right)$ into Eq. (4.1) shows that the subject must generate a position command that is exactly equal and opposite to the current cursor position: u(t) = −x(t). Thus, feedback is essential: the subject must know the state of the cursor. A subject who could issue this perfect motor command (instantaneous and 100% accurate) would be able to keep the cursor completely still.

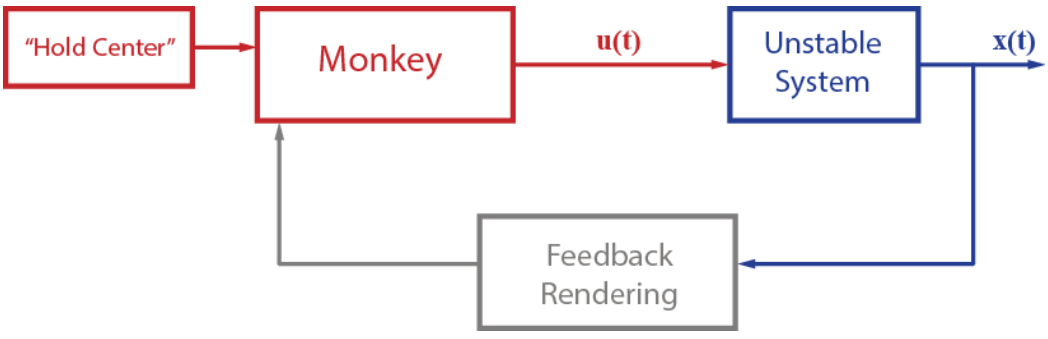

Figure 4.1: Closed-loop (feedback) control diagram depiction of the CST.

In practice, this perfect control command is not possible because of sensorimotor noise and nervous system processing delays; hence, the cursor will continuously move, and maintaining control – keeping it from rapidly drifting off-screen – requires an ongoing sensory-guided motor action. With adequate sensory feedback, a subject can generate a command signal that keeps the cursor centered about an equilibrium point for λ values that are less than the subject's CIV. The greater a subject's CIV, the more effective is their sensorimotor control. Thus the CIV is a single number that provides an objective, empirically-derived measure of performance. The CIV depends on several interesting factors, including the type of feedback (visual or tactile in our experiments), the quality of sensory information (which we can control), the nature of the effector (hand movements or BCI control in our experiments), and the subject's level of skill. Importantly, a slower or less reliable sensory feedback modality will yield a lower CIV than a faster and more reliable feedback modality.

4.3 METHODS

Three adult male Rhesus monkeys (macaca mulatta) were used in this study (monkeys I, J, and N – 9 kg, 12.6kg, 8.6kg respectively). All animal procedures were approved by the University of Pittsburgh's Institutional Animal Care and Use Committee, in accordance with the guidelines of the US Department of Agriculture, International Association for the Assessment and Accreditation of Laboratory Animal Care, and the National Institutes of Health. Monkey J was being administered diazepam and fluoxetine at a veterinarian's recommendation for reasons unrelated to the study's objectives.

4.3.1 Critical Stability Task implementation

To perform the CST (Jex et al., 1966), monkeys must maintain a drifting cursor at the center of a computer display for several seconds. The monkey must act to stabilize the cursor via hand movements, and if no corrective action is made, the cursor will drift rapidly off-screen. We implemented the CST on a digital computer via the discrete-time algorithm (Quick et al., 2014),

$$x(k+1) = ax(k) + (a-1)u(k),\ a > 1,\ k = 1,2,3,\ldots \quad (4.2)$$

where $x(k) = x(t)|_{t=kT_s}$ is the current horizontal cursor position at time $t = kT_s$ and $T_s = \frac{1}{120}$s is the update interval (i.e. sampling period); u(k) is the current horizontal hand position of the subject; x(k + 1) is the cursor position to be rendered at the next update; and $a = e^{\lambda T_s} > 1$ is the discrete system instability parameter as determined by the value of λ.

Two example trials of the CST are shown in Figure 4.2. The left top panel schematizes what the animal sees on the computer monitor at five time-points during a visual feedback trial, and the right top panel shows how the vibrotactile feedback is rendered during a vibrotactile feedback trial. The middle panel shows the cursor position (in blue) and the hand position (in red). The cursor is initialized to the center of the screen at the start of each trial, but due to inevitable motor noise, it soon begins to drift. The animal does his best to move his hand equal and opposite to the cursor position, but no correction is perfect, so the cursor continues to drift. Thus, ongoing movement is required to keep the cursor stabilized.

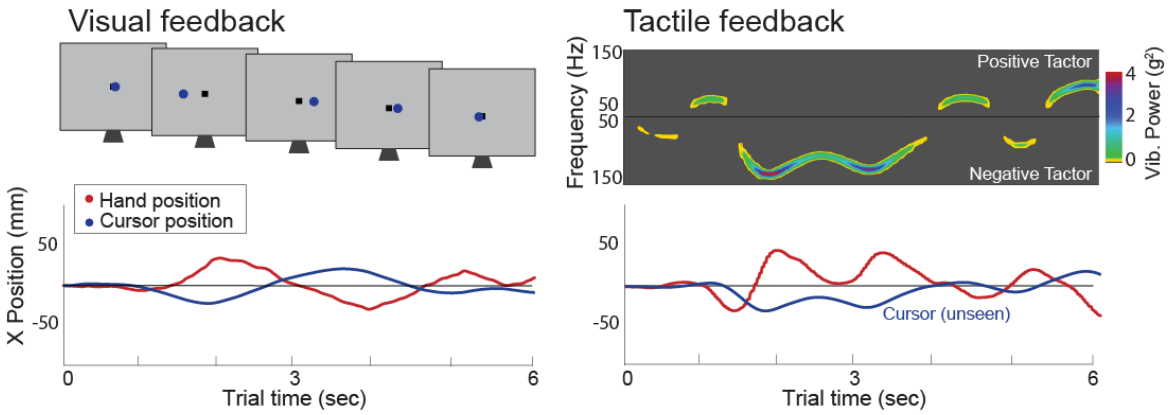

Figure 4.2: Two CST trials. 1st row: schematic of visual feedback (left) and spectrogram of vibrotactile feedback (right). 2nd row: cursor and hand position traces. Bottom left: raster of multichannel neural activity from M1.

4.3.2 Feedback modalities

Monkeys worked in a virtual environment rendered in a frontoparallel plane for all experiments. The monkeys did not see their hand because it traveled in the space behind the virtual 3D environment. Monkeys performed the CST using unconstrained hand movements and either visual or vibrotactile feedback of the system's position. Hand position was recorded using an active motion capture system (PhaseSpace Inc., San Leandro, CA). A powered LED marker was attached either to the monkey's finger (Monkeys N and J) or the monkey's wrist (Monkey I). Six motion capture cameras recorded the 3D position of the LED marker, which was then used for task control.

We examined performance in the CST during hand control with visual or vibrotactile feedback about the cursor state. The monkey's task was to keep a cursor at screen-center on the monitor. For the visual trials, feedback about the cursor state was provided by rendering the horizontal cursor position every 1/120 s (~8.33 ms) on the monitor. To initiate trials, the monkey moved his hand to move a red circle (5 mm diameter) displayed on the monitor so that it was on

top of a gray square (5 mm x 5 mm) at screen-center. After a 500 ms hold period, the red circle changed to blue, indicating the start of the trial, and now represented the cursor with unstable dynamics that the animal needed to keep centered at the gray square. The monkey generated cursor position corrections by moving his hand, which was unconstrained and free to move in 3D. However, only the 1D left/right hand movements affected the cursor position, which moved horizontally on the monitor. A trial lasted up to six seconds; if the monkey was able to keep the cursor from drifting too far from the center square (±50 mm) for six seconds, the trial was recorded as a success, the monkey received a water reward, and the next trial commenced. If the cursor deviated too far from the center square before six seconds, the trial was a failure, the screen flashed a large red square, no reward was given, and a 3 sec timeout occurred before the next trial commenced.

During vibrotactile trials, the monkey initiated trials as in the visual trials, by moving his hand to position the red circle on top of the gray square at screen-center. After a 500 ms hold period, the red circle disappeared indicating the start of the trial. The static center square remained visible to the monkey to indicate that the trial was ongoing. Feedback about the cursor position (not displayed visually) relative to the visible gray square at screen-center was delivered to the monkey via two vibrating coin tactors (model 312-101, Precision Microdrives Limited, UK). The tactors were attached to the forearm of the monkey's non-working arm using Velcro against a Spandex sleeve. One tactor was positioned at the wrist and one at the elbow. The wrist tactor encoded cursor position to the right of screen-center, and the elbow tactor encoded cursor position to the left of screen-center. Tactor vibrations were activated by a supply voltage signal proportional to the magnitude of the distance of the cursor from screen-center. An increase in supply voltage (corresponding to greater distance from screen-center) generated a higher

intensity and frequency of vibration. Tactor vibrations were measured using an accelerometer attached to the back of each tactor. The vibration frequency ranged from 50-175 Hz, and the vibrational power ranged from 0.5-20g^2, where 1g equals the force of gravity.

4.3.3 Training the CST

Monkeys were trained over several sessions to perform the CST with hand control under each feedback modality. Visual training preceded vibrotactile training, as it is more natural to the animal and allowed him to learn the essence of the task before introducing the abstract vibrotactile feedback. Initial training started with a very low instability parameter value (on the order of $\lambda \sim 0.1$ rad/s), for which the cursor moved very slowly. As the monkey learned to hold his hand such that the cursor stayed close enough to the center square for the six second trial, thereby earning the monkey a reward, the value of λ was slowly incremented over blocks of consecutive trials. At least ten trials, but sometimes dozens of trials, were run at each fixed λ value before increasing it for the next block of trials. It was important during training to keep the animal engaged and not increase λ too rapidly, lest the difficulty of the task vastly exceeded the animal's current abilities and cause him to give up for the session. With each training session, monkeys gradually developed the ability to control greater instability values.

After the monkeys became proficient with the task over several hundred training trials, we collected experimental data over five testing sessions. The data from these testing sessions were used to score the monkey's ability to perform the CST under that feedback modality. Each testing session began at an instability value for which the animal could easily succeed, and was increased by 0.1 rad/s every 10-20 trials. We stopped increasing λ when an animal's success rate at the current λ was ~30-40%. Once this upper level was reached, λ was reset to what it had been

at the beginning of the session and we repeated the incrementing process. For each experimental session, we attempted to get two sweeps through the λ values.

Before performing the CST using vibrotactile feedback, the monkeys were taught to understand vibrotactile information using a horizontal reaching task. In this task, monkeys moved their hand-controlled cursor to a start target, held for 250ms, and then made a horizontal reach to an illuminated peripheral target as soon as it appeared. If the subject held his hand within a 10mm acceptance window of the target for 250ms, the trial was successful and the subject was rewarded. This task was initially taught using visual feedback only, but once the animals understood the task vibrotactile feedback was added and the task was extended to include three possible targets to the right and left of the start target. For each trial, the vibrotactile workspace was re-centered around the target. For example, if the peripheral target appeared to the right of the start target, the cursor would initially be far left of the target, so the tactor near the elbow would vibrate at a high frequency and magnitude. The magnitude and frequency of vibration would decrease as the subject moved his hand-controlled cursor rightwards to the target. If the subject went too far right, such that the cursor became right of the target, the tactor near the wrist would vibrate at a magnitude and frequency proportional to his error.

The vibrotactile information first occurred simultaneously with the visual feedback. Then as subjects improved, the visual feedback was taken away. In the beginning of training, the monkeys abruptly stopped moving their hand when the visual feedback was removed. To remedy this, visual cursor feedback was removed as the subject approached the still visible target. As the monkeys became more accustomed to the lack of visual information, visual cursor feedback was removed earlier and earlier in the trial until it was only used to initiate the trials. Once the monkeys were able to reach without viewing the cursor, the visual peripheral target was removed

in the same manner. Eventually, monkeys performed the task using solely vibrotactile information to reach to and hold at an unseen target.

After successfully completing the horizontal reaching task, monkeys progressed onto learning how to perform the CST with vibrotactile feedback. Just as with visual feedback, the monkeys progressed to higher values of λ over several training sessions, and once they become proficient at performing CST under vibrotactile feedback, we collected data over five testing sessions, analogous to the visual feedback testing sessions.

4.3.4 "Hold Still" trials

In principle, the monkeys could have attempted to perform the CST by ignoring the sensory feedback that was provided to them. To ensure that they were using feedback to inform their movements, we sought to determine the CIV for the CST performed with hand control in the absence of feedback. We could not ask the animals to perform the CST without feedback, as that they became confused when they received no feedback and would immediately put their arm down as though there was no task occurring. The optimal strategy for succeeding at the CST without using feedback is holding perfectly still at the center square. If the monkey were able to hold perfectly at the center square, the monkey could attain an infinitely large λ_C score because the system would not have received any input to cause it to deviate from the center. However, motor noise and sensor noise make this impossible. However, for small values of λ, the cursor moves slowly, and it is possible for the cursor to remain within-bounds for six seconds without any corrective control commands being issued.

We collected data during trials in which the animal was trained to hold his hand still at the center target. In this "hold still" task, the subject would initiate trials by moving his hand to

move a red circle displayed on the monitor so that it was on top of a gray square at screen-center, just as in the CST. However, during this task, the animal was required to keep the red cursor at the center target for six seconds. Twenty-four such trials were collected from Monkey I, and 33 from Monkey N. These hand movement data were then used offline as input control commands to a CST simulation over various λ values. We deemed trials to be successful or not using the same criteria as during the CST task (6s trial, ±50mm boundary) and calculated the CIV as we would with online CST data. This allowed us to compare the CIV from holding still, with the CIV from visual and vibrotactile feedback to ensure that the feedback was being used.

Failure in the simulated CST task with the "hold still" hand position data is due predominately to motor noise. However, sensor noise alone would also cause the system to go unstable, if corrective movements were not made.

4.3.5 Data analysis

4.3.5.1 The Critical Instability Value (CIV)

For each feedback modality and the "hold still" condition, we calculated the value of λ below which the majority of trials were successful. We termed this value the critical instability value (CIV) since, above this value, the animal was unsuccessful at controlling the cursor (i.e., stabilizing the closed-loop system) for more than 50% of trials. We designate this value by λ_C, which was determined by a least-squares fit of a psychometric function (a Gaussian cumulative distribution function (GCDF)) to the behavioral data (percent of successful trials versus λ).

In order to obtain confidence intervals for the CIV for each feedback modality and the "hold still" condition, we used a bootstrap procedure to fit the distribution. Specifically, we generated 1000 surrogate data sets from the experimental data by resampling the experimental

data with replacement, such that each surrogate data set contained the same number of data samples as the original experimental data set. Then, for each surrogate data set n = 1…1000, we fit the surrogate data to the GCDF,

$$\hat{p}_n(\lambda, \mu_n, \sigma_n) = 50\left(1 - \mathrm{erf}\left(\frac{\lambda - \mu_n}{\sqrt{2}\sigma_n}\right)\right) \quad (4.3)$$

to obtain the critical lambda $\lambda_{C_n} = \mu_n$ for that set. We computed the mean and standard deviation of these 1000 values to obtain the CIV (λ_C) and 95% confidence intervals for each feedback modality.

4.3.5.2 Sensorimotor Lag

For trials in which the monkey was successful, the motor and cursor position signals were typically inverted with respect to one another, and the motor signal lagged the cursor signal (Figure 4.3). We calculated sensorimotor lag by cross-correlating the cursor position signal with the motor command signal,

$$r_{xu} = \int x(t)u(t+\tau)dt \quad (4.4)$$

and finding the time corresponding to the peak negative correlation value. A positive lag time indicates that the motor command lagged the cursor position. We accounted for lags in the rendering of the different feedback modalities, as well as the input/output CPU delay. Specifically, we determined the delay in visual rendering of the cursor position using a photodetector, and we determined the delay in tactor activation by measuring the delay between a step increase in the driving voltage and the time when a vibration was registered by the accelerometer mounted to the tactor. The visual rendering delay was 65ms, and the vibrotactile rendering delay was 33ms. These delays were subtracted from the respective cross-correlation lag for each feedback modality, as was the CPU delay (20ms).

4.3.5.3 Sensorimotor Gain

A sensorimotor gain was calculated as a function of λ for each feedback modality by computing the ratio of the root-mean-square (RMS) value of hand position to the RMS cursor position. The RMS cursor position value also serves as a measure of control "effectiveness," since the task was to keep the cursor at screen-center; hence, lower RMS cursor position indicates more effective control. In addition, the RMS value of hand velocity was calculated as a measure of control "effort," since the kinetic energy of hand motion is proportional to its squared-velocity; hence, lower RMS hand velocity indicates lower control effort.

4.4 RESULTS

4.4.1 Critical Instability Values (CIV)

Representative trials from monkey I are shown in Figure 4.3, which plots the cursor and hand position data during a trial. Also shown are examples (spectrograms) of the vibrotactile feedback signal. Figure 4.4 is a plot of the average success rate versus λ across all testing trials for monkey I, and the curve fits, for each feedback condition. The average CIV for each monkey and each feedback condition is shown in Figure 4.5. For visual feedback, Monkey I achieved a CIV of λ_C = 4.57 ± 0.07 rad/s (standard error of the mean, s.e.m.) over 1523 testing trials, and Monkey J achieved a CIV of λ_C = 3.03 ± 0.04 rad/s over testing 1863 trials. For vibrotactile feedback, Monkey I achieved a CIV of λ_C = 2.13 ± 0.04 rad/s over 1066 testing trials, and Monkey J achieved a CIV of λ_C = 1.89 ± 0.03 rad/s over 2090 testing trials. For each monkey, the CIV for visual feedback was significantly greater than the CIV for vibrotactile feedback ($p < 0.05$).

The "hold still" CIV was λ_C = 0.72 ± 0.06 rad/s for Monkey I. We were unable to perform the "hold still" task with Monkey J, so we collected data from another monkey for this condition (Monkey N, λ_C = 0.53 ± 0.04 rad/s). The "hold still" CIVs were significantly lower than all of the feedback CIVs.

4.4.2 Sensorimotor Lag

The sensorimotor lag tended to decrease as λ increased (Figure 4.6). The value of $1/\lambda$ is plotted as the dashed line in Figure 4.6, and represents the effective maximum delay that could yield stable control in the closed-loop system, assuming a linear time-invariant system with motor command $u(t) = -x(t - \tau)$ (Jex et al., 1966).

The sensorimotor lag was generally larger during vibrotactile feedback than during visual feedback. In particular, for successful trials when the monkeys performed at their CIV, we found that Monkey I's vibrotactile lag at his vibrotactile λ_C was 207.2 ± 10.3ms, while his visual lag at his visual λ_C was 136.6 ± 6.5ms, which was significantly faster than his vibrotactile lag (one-tailed t-test, p-value = 9.18e-9). The visual feedback data presented in Figure 4.6B (blue closed circles) was collected while the photodiodes in the lab were not functioning. The photodiodes enable us to determine the delay between when the RT computer executes a state transition and when that state transition is presented to the animal, which is used to calculate sensorimotor delays. Thus, we cannot make any definitive claims about Monkey J's visual lags. However, the data plotted in Figure 4.6B approximates the photo delay based on data collected after the photodiodes were fixed, and is likely a close approximation. The vibrotactile delay was accurate, and Monkey J's vibrotactile lag at his vibrotactile λ_C was 243.0 ± 11.0ms. We are unable to claim that this is significantly higher than his visual lag, while this is likely the case.

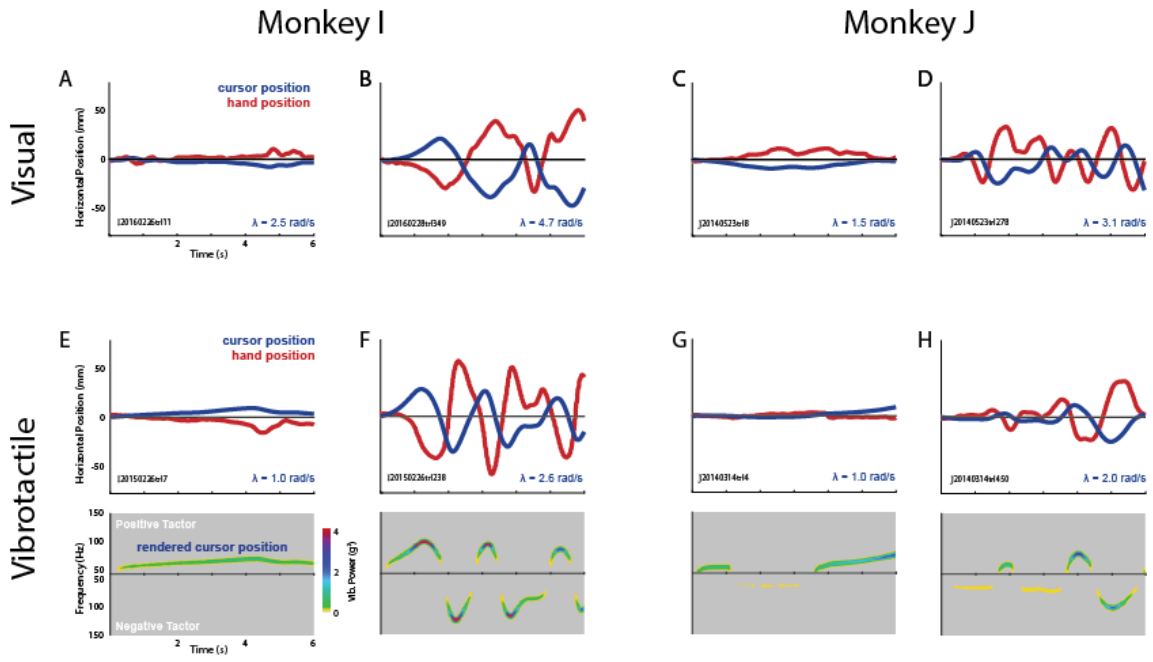

Figure 4.3: Example trials for each monkey under the different feedback conditions. (A-D): Visual feedback trials. (A) Monkey I, low instability level (λ = 1.5). (B) Monkey I, high instability level (λ = 3.1). (C) Monkey J, low instability level (λ = 1.5). (D) Monkey I, high instability level (λ = 3.1). (E-H): Vibrotactile feedback trials. Bottom subplot in each case shows spectrogram of the tactor feedback vibrations. (E) Monkey I, low instability level (λ = 1.0). (F) Monkey I, high instability level (λ = 2.8). (G) Monkey J, low instability level (λ = 1.0). (H) Monkey J, high instability level (λ = 2.0).

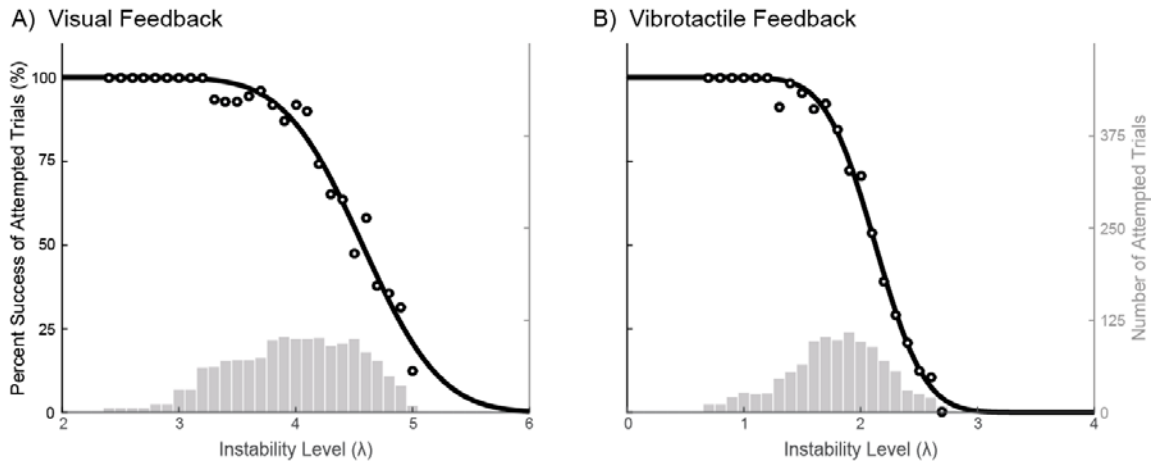

Figure 4.4: Example data for Monkey I. Success versus λ and curve fits for (a) visual feedback and (b) vibrotactile feedback.

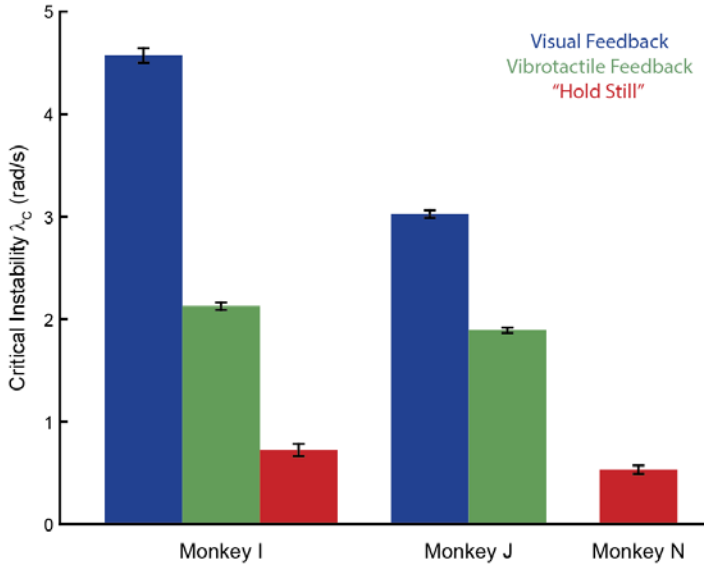

Figure 4.5: Average critical instability values (CIV) across feedback condtions. The visual feedback condition is blue, the vibrotactile feedback condition is in green, and the "hold still" condition is in red. Error bars show 95% confidence intervals.

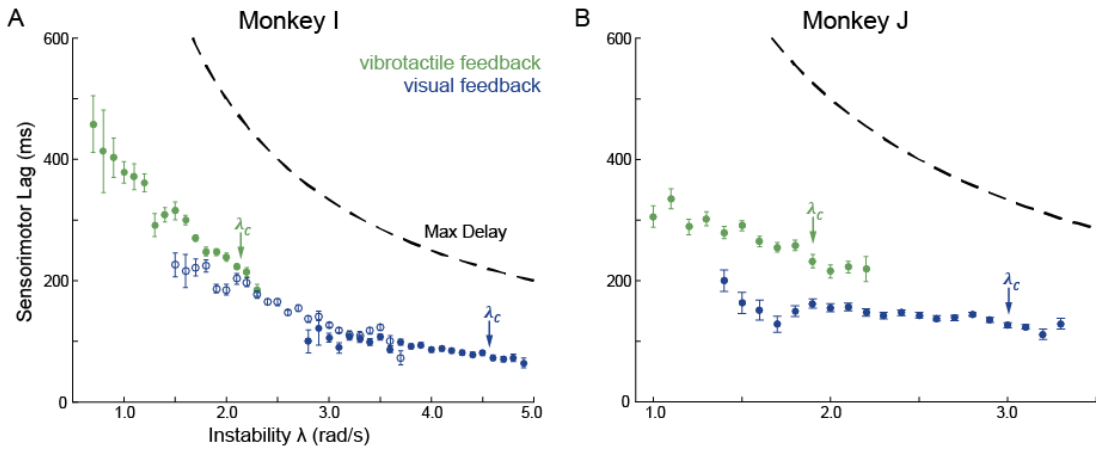

Figure 4.6: Sensorimotor lag versus λ. (A) Monkey I's vibrotactile (green circles) and visual (blue circles) sensorimotor lag (mean s.e.m.) v. λ. Each point represents a minimum of ten successful trials, obtained during testing sessions. The arrows labeled λ_c point to the CIV for each feedback modality. The dashed line depicts the theoretical maximum lag that could be tolerated. (B) Monkey J's vibrotactile and visual sensorimotor lag values. In (A), there was no overlap in λ values between vibrotactile and visual testing sessions (solid green and blue circles) for Monkey I, so we have also included results from sessions preceding the testing sessions (open blue circles), for comparison between vibrotactile and visual performance at similar values of λ.

4.4.3 Sensorimotor Gain

The RMS position of the cursor and hand, and the RMS hand velocity, for each monkey and feedback modality, are shown in Figures 4.7 and 4.8. In general, the RMS hand and cursor position, and the RMS hand velocity, tended to increase with λ for both monkeys and both feedback modalities. These effects were generally larger for vibrotactile feedback with Monkey I, whereas Monkey J's RMS results were similar between visual and vibrotactile feedback.

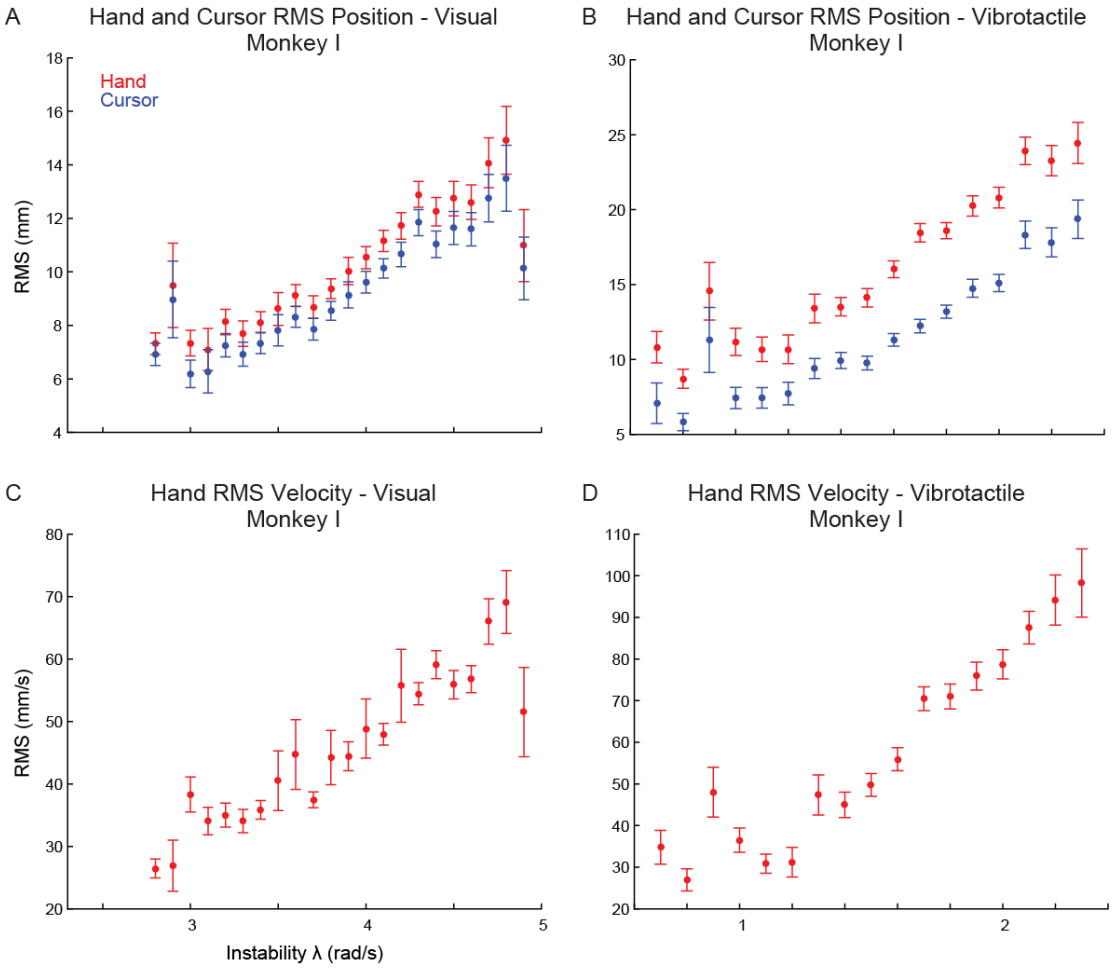

Figure 4.7: RMS position and velocity for Monkey I. Top Row: RMS position of the cursor (blue) and hand (red). Bottom Row: RMS hand velocity.

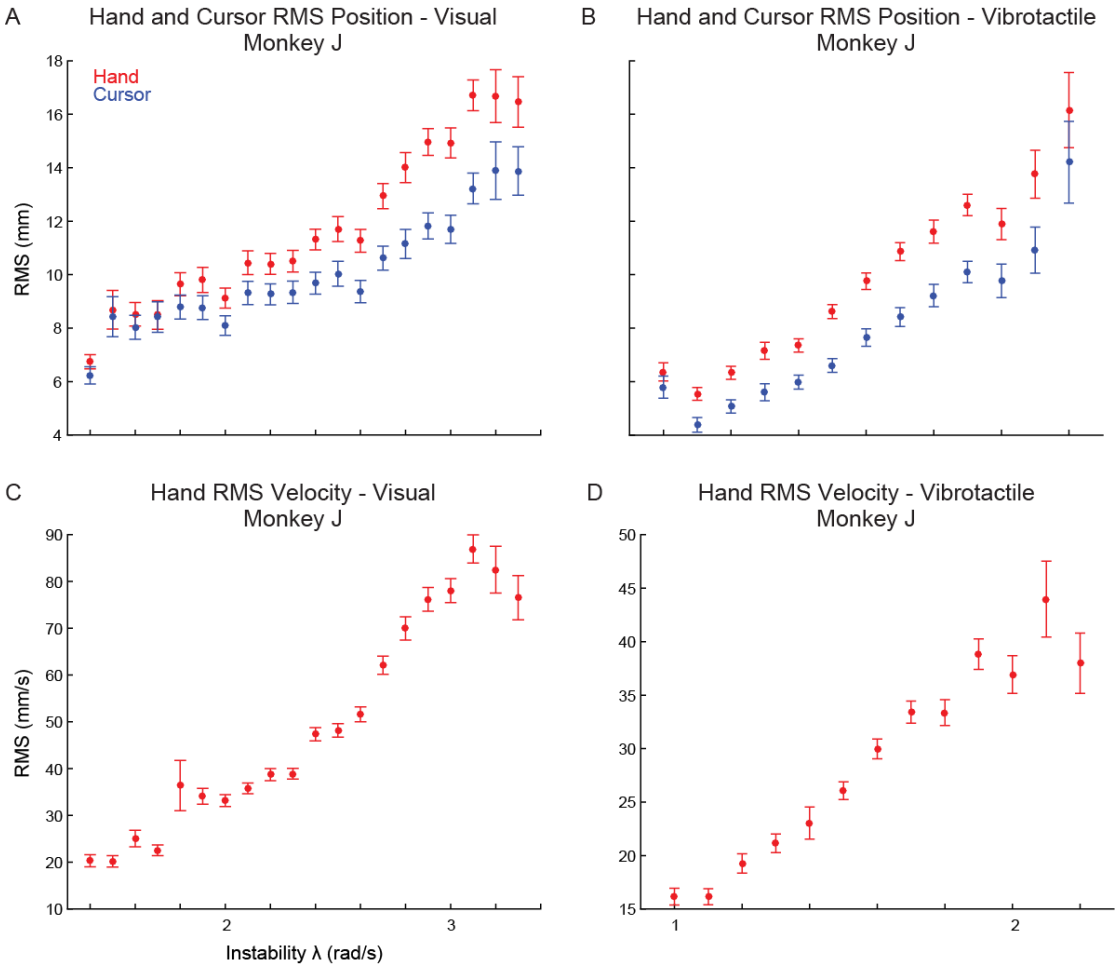

Figure 4.8: RMS position and velocity for Monkey J. Top Row: RMS position of the cursor (blue) and hand (red). Bottom Row: RMS hand velocity.

The sensorimotor gain, computed as the ratio of RMS hand position to RMS cursor position at each λ, is shown in two formats in Figures 4.9 and 4.10. Also shown (dotted lines) are the theoretical minimum and maximum gain for stable control under the assumption of an LTI system with position command of $u(t) = -Kx(t-\tau)$; under these assumptions, the minimum gain for stable control is K = 1, and the maximum gain is inversely related to λ (Jex et al., 1966).

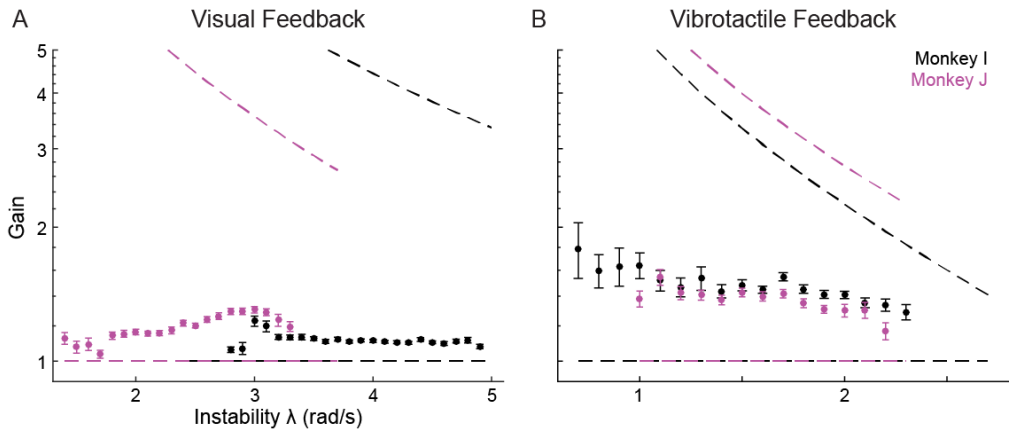

Figure 4.9: Sensorimotor gains – Visual feedback vs. Vibrotactile Feedback.

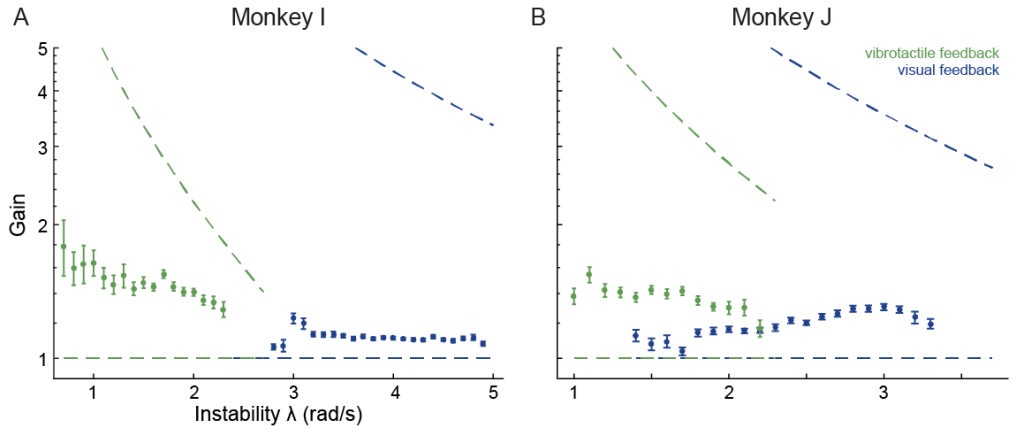

Figure 4.10: Sensorimotor gains – Monkey I vs. Monkey J.

4.5 DISCUSSION

The study of motor control in non-human primates has largely consisted of point-to-point reaches, which can be accomplished with intermittent sensory feedback (Ghez et al., 1995). In order to explore how continuous sensory feedback, both natural and artificial, shapes ongoing motor control, we trained monkeys to perform the Critical Stability Task (CST), in which monkeys made serial hand movements to try to keep a horizontally drifting cursor at the center of a monitor. The 1D dynamics of the cursor were such that the cursor would diverge from screen-center and drift off screen in the absence of hand movements, or for inadequate hand movements. The CST cannot be successfully accomplished without continuous sensory feedback about the cursor state, and therefore provides a salient method to evaluate the effectiveness of individual feedback modalities and a platform for studying the neural basis of sensorimotor integration.

4.5.1 Critical Instability Value (CIV)

Both the CIVs for both visual and vibrotactile feedback were higher than the CIVs for the “hold still” condition, demonstrating that both animals were using sensory feedback for control. Our main finding is that monkeys can learn to use the artificial sensory stimulus of vibrotactile feedback to shape a long-duration motor action that requires sensory feedback. Specifically, as quantified by the CIV, the monkeys' control of the cursor using vibrotactile feedback was 47-62% as good as their control using visual feedback (Figure 4.5). For comparison, human subjects performing a similar CST experiment achieved visual CIV of $\lambda_C = 2.83 \pm 0.55$ and vibrotactile CIV of $\lambda_C = 1.16 \pm 0.23$, demonstrated in Chapter 5 (Quick et al., 2014). This difference,

especially the lower CIV for vibrotactile feedback in humans, is likely due to the greater amount of training that the monkeys had with the CST (under both feedback modalities, but especially vibrotactile feedback).

4.5.2 Learning a novel feedback modality

We found that monkeys could learn an abstract form of vibrotactile feedback to shape complex, long duration movements. Recently, Dadarlat et al. also found that monkeys could learn to use abstract intracortical microstimulation feedback to guide reaches to unseen targets (2015). These findings bode well for teaching monkeys other forms of abstract feedback. It is interesting to note that even once the monkeys learned the meaning of the vibrotactile signal, they still had to learn how to use that signal to inform their movements within the CST. Monkeys first learned each tactor's direction, and how the vibration's frequency/magnitude corresponded to the distance their hand needed to travel. Even though the monkeys were proficient making simple, straight trajectory reaches using the vibrotactile mapping, they were not initially successful when they were transitioned to using this non-visual feedback to perform the CST. The monkeys needed many more weeks to learn how to quickly utilize these vibrotactile signals in order to perform complex, feedback driven movements. Their learning process demonstrates that simply knowing the meaning of a feedback signal does not directly correspond to effective use of the feedback signal for complex motor control. By asking the monkeys to control an unstable system, we pushed them to use the vibrotactile feedback as quickly and accurately as possible.

4.5.3 Sensorimotor Lag

Although monkeys and humans have been shown to have shorter reaction times to vibrotactile stimuli than to visual stimuli (Godlove et al., 2014), sensorimotor lag was generally larger during vibrotactile feedback than during visual feedback. This may be due to the fact that vibrotactile feedback is a learned feedback modality whereas visual feedback is native feedback modality. While the animals received ample practice with vibrotactile feedback, they have used visual feedback every day since birth. Animals may have *felt* the vibrotactile signal more quickly than they *saw* the visual signal, however they required additional time to cognitively process and interpret the vibrotactile feedback before generating the appropriate corrective hand movement.

While we cannot make any definitive claims about the comparison of Monkey J's vibrotactile lags to his visual lags, the trends seen in the data in Figure 4.6B are similar to those seen in Monkey I's data in Figure 4.6A, for which the visual lag times are accurate. We found that the sensorimotor lag generally decreased with increasing task difficulty for both feedback modalities. This finding is consistent with a similar (but smaller) negative slope in the effective sensorimotor lag with increasing task difficulty in human subjects performing a similar CST with visual feedback (Jex et al., 1966). This change in sensorimotor lag with task difficulty indicates that the animal's control command generation is more sophisticated than linear time-invariant (LTI) proportional control with fixed delay. If control command generation functioned as LTI proportional control with a fixed delay, we would expect that the sensorimotor lag would be constant across the instability levels, equaling their respective modality's reaction time, around 220ms for visual and 185ms for vibrotactile (Godlove et al. 2014).

It is possible that the sensorimotor lag is larger at lower instability values precisely because it can be: for λ substantially below the CIV ($\lambda << \lambda_C$), the cursor moves slowly and the

task is not very challenging, such that rapid cognitive processing and rapid hand adjustments are not required to keep the cursor from drifting out-of-bounds. Thus, one possibility is that the subject's sensorimotor delay is a function of task difficulty: $\tau = \tau(\lambda)$. It is interesting to note that at high instabilities, the visual lag for Monkey I (and Monkey J, using the approximated values) is lower the visual reaction time. This likely means that the monkeys were using a forward model to compensate for the visual feedback delay (Miall and Wolpert, 1996). Using a forward model, they could estimate the next system position and generate a movement in response to this estimate. With vibrotactile feedback however, even at high instability values, the vibrotactile lag at both monkeys was higher than the vibrotactile reaction time. This may mean that the monkeys were not able to form a salient, accurate perception or the vibrotactile feedback signal, and therefore were not able to generate a forward model utilizing the vibrotactile feedback information. Without a forward model, the monkeys could not achieve sensorimotor delays less than their vibrotactile reaction time, which may contribute to the fact that animals achieved lower CIVs using vibrotactile feedback.

4.5.4 Sensorimotor Gain

The sensorimotor gain reflects how accurately the subject's hand position counteracts the cursor position. Under the assumption of an LTI system with position command of $u(t) = -Kx(t - \tau)$, a sensorimotor gain where K = 1 suggests that $u(t) = -x(t)$. This would require the subject's hand position to be equal and opposite the cursor position with no delay. In practice, this unity gain is not possible because of sensorimotor noise and nervous system processing delays. This is supported by the fact that the animals' gains in Figures 4.9 and 4.10 are slightly above the theoretical minimum necessary gain. Gains were relatively constant across task difficulties due

to the fact that the task requires the subject to inversely track the cursor. The sensorimotor gain was generally higher for vibrotactile feedback than visual feedback. This may be due to the fact that the animals could more accurately perceive and counteract the visual feedback signal. The perception of where the cursor was using vibrotactile feedback may have been less accurate. In response, animals would make arm movements to counteract their perception of where the cursor is rather than where it actually is, causing the gain to increase.

4.5.5 Control Effort and Effectiveness

Control effort and effectiveness were assessed by RMS hand velocity and RMS cursor position, respectively. The results in Figures 4.7 and 4.8 show that, for both monkeys and both feedback modalities, control effort increased and control effectiveness decreased as the task become more difficult. This is expected. Control effort is calculated as the RMS value of hand velocity since the kinetic energy of hand motion is proportional to its squared-velocity. As the task becomes more difficult, λ increases; λ is the inverse of the effective time delay, τ, so as λ increases, the subject has less time to react to deviations of the cursor position from screen center and must move his hand more quickly. Control effectiveness is measured as the RMS cursor position since the task asks the subject to keep the cursor at screen-center. If control was perfectly effective, the subject would be able to counteract every cursor deviation from screen center quickly and accurately, causing the cursor to stay within one or two millimeters from the center. As the task becomes more difficult, the deviations of the cursor from screen center increase indicating that the subject's effectiveness decreases and he or she can no longer respond to the feedback as quickly and/or as accurately.

4.6 ADDITIONAL WORK: BRAIN CONTROL

While performing the behavioral analyses described above, Kristin and I also trained one monkey to perform the CST with BCI. Monkey I was implanted with a 96-channel multi-electrode array (Blackrock Microsystems) into dorsal premotor cortex (PMd), contralateral to the arm used during hand control CST trials. We recorded from 95 channels (Tucker-Davis Technologies). Threshold crossings at 4.0x RMS (measured at the beginning of the session) and were used to control the BCI.

A velocity Kalman filter (Wu et al. 2006) with a 30ms bin width was used for BCI control. The BCI was calibrated using hand movements during a standard 8-target center out task (Moran and Schwartz, 1999; Sadtler et al. 2014). The velocity Kalman filter used during BCI control is explained further Section 4.7.1. During the BCI CST trials, decoded velocities were integrated after each time bin to obtain the position signal that was input to the unstable system. The arm was unrestrained during BCI trials. At the beginning of each trial, the decoder was initialized 20mm above the center square. To start the CST, the monkey moved the BCI downward to acquire the center square (30mm acceptance window) and hold the decoder within the square for 100ms. This was to ensure that the monkey was engaged before the start of each BCI trial. When the monkey was not engaged, the decoder drifted upwards. As in hand control, the cursor changed colors once the CST began. Due to neural noise, the cursor began to drift immediately. The monkey generated neural signals to stabilize the system. The monkey was rewarded for keeping the system within the ±50mm boundary around the center square for six seconds.

4.6.1 BCI decoder training

Monkey I used a BCI decoder (velocity Kalman filter) for online control of the CST. The decoder was trained on neural activity collected while Monkey I performed a standard center-out task. We collected 80 successful hand control trials, 10 per target. The entire trial was used to train the velocity Kalman filter decoder, which included moving to the peripheral target and back.

We also performed an offline decode of Monkey I's the neural activity during "hold still" trials to determine the CIV for the CST performed using brain control in the absence of feedback. The offline decoded position was used as input into simulated CST trials. The decoder was trained and tested on the 24 "hold still" trials. These were the same trials that were used for the "hold still" condition for hand control, but here recorded neural activity was used as the input for offline simulations rather than the hand position. Initially, we tried to build the decoder offline using center-out trials, however this yielded disproportionately large decoded movements during "hold still" trials. So as to achieve the conservative, highest estimate of the critical instability using offline BCI movements, we decided to proceed with a decoder trained on the "holding still" portion of the trial.

A velocity Kalman filter (vKF) was used to decode horizontal and vertical cursor velocity:

$$\boldsymbol{x}_k|\boldsymbol{x}_{k-1} \sim N(A\boldsymbol{x}_{k-1} + \boldsymbol{b}, Q) \tag{4.5}$$

$$\boldsymbol{u}_k|\boldsymbol{x}_k \sim N(C\boldsymbol{x}_k + \boldsymbol{d}, R) \tag{4.6}$$

where $\boldsymbol{x}_k \in \mathbb{R}^{2x1}$ is a vector of horizontal and vertical decoder velocity at time step k, $\boldsymbol{u}_k \in \mathbb{R}^{qx1}$ 1 vector of z-scored spike counts (z-scoring performed separately for each neural units) and taken in non-overlapping 30ms bins (monkey I and N) across the q neural units. We fit the

parameters of A, b, Q, C, d, R using maximum likelihood (Wu et al. 2006) by relating the spike count vector to the monkey's recorded hand velocity from the training data. Since the spike count vectors were z-scored before decoding velocity, d=0. Because calibration kinematics were centered about the center of the workspace, b=0. The decoded velocity at time step k was $\hat{\boldsymbol{x}}_k = E[\boldsymbol{x}_k|\boldsymbol{u}_1, \ldots, \boldsymbol{u}_k]$. We can express $\hat{\boldsymbol{x}}_{k+1}$ in terms of the decoded velocity at the previous time step $\hat{\boldsymbol{x}}_k$and the current z-scored spike count vector $\boldsymbol{u}_k$, where K is the steady-state Kalman gain matrix:

$$\hat{\boldsymbol{x}}_k = M_1\hat{\boldsymbol{x}}_{k-1} + M_2\boldsymbol{u}_k \tag{4.7}$$

$$M_1 = A - KCA \tag{4.8}$$

$$M_2 = K \tag{4.9}$$

4.6.2 Results

Monkey I performed the CST using a closed-loop BCI decoder (velocity Kalman filter) while receiving visual feedback. He achieved a λ_C score of 1.036 ± 0.028 rad/s across 1278 trials over 5 sessions. Several example task trajectories are shown in Figure 4.11. This score is low, however the CIV for visual feedback with BCI is higher than that of the “hold still” task for an open loop, offline BCI decoder. With no feedback about the system, Monkey I’s neural activity produced a CIV of λ_C = 0.346 ± 0.03 rad/s. This score is significantly lower than when the monkey used BCI control with visual feedback about the systems position, assuring that the monkey is actively using the visual information to stabilize the CST using BCI control.

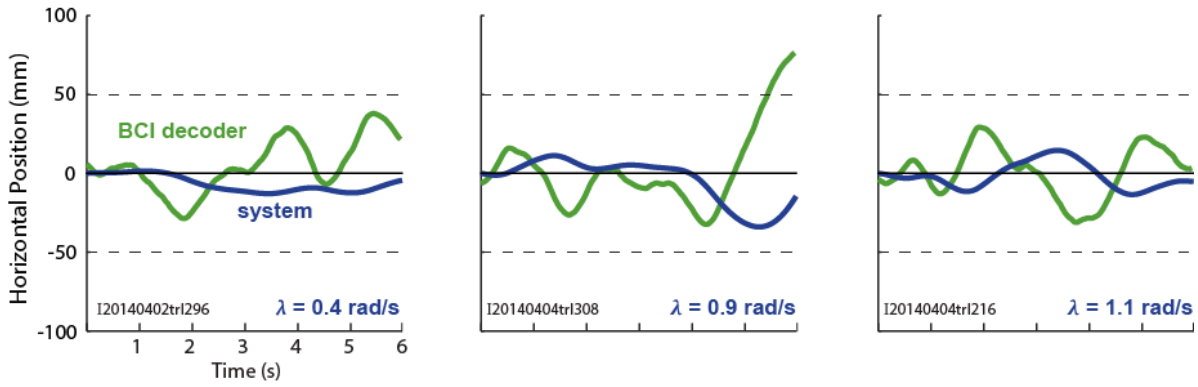

Figure 4.11: Example trials of monkey I using the BCI decoder to perform the CST. The green line indicates the position of the BCI decoder and the blue line indicates the systems position, which was rendered visually to the monkey.

4.6.3 Discussion

While Monkey I was able to use the BCI decoder to stabilize instabilities, the BCI decoder was not very effective. Monkey I's BCI control was only 22.7% as effective as when he used hand control. The large discrepancy between hand control and BCI control demonstrates that the CST provides a unique way to assess decoder effectiveness. Many other virtual BCI tasks, such as center-out, demonstrate that BCI control is only slightly lower than, or on-par with natural arm movements. It is difficult to design better decoders when the evaluation metrics only show slight improvements between decoders, or different metrics supply contradictory results (i.e. a BCI decoder which can move faster to the target, but will then have trouble holding on the target). With point-to-point reaches, various decoders show minimal difference in performance, with all decoders achieving high success rates. In these tasks, decoders only need to move an effector in one direction, and reach adjustments due to feedback occur over the duration of the reach. By contrast, in the CST subjects are required to change directions quickly as to keep the unstable

cursor in-bounds, yielding a low CIV even using one of most advanced BCI decoders currently available. Thus, this single metric demonstrates that there is still much room to design better decoders, and points to the usefulness of the CST to complement traditional BCI tasks.

It is interesting to note that for the CST, the BCI decoder makes similarly-sized movements across the three instability levels, as shown in Figure 4.11. This is in contrast with hand control (Figure 4.3), where animals stabilize low instabilities with small arm movements and large instabilities with larger arm movements. This may be due to the inability of today's BCI decoders to scale movement size. When I began working in the Batista lab, Kristin and I noticed that offline, the BCI decoder for CST was able to decode large movements, but had trouble decoding small movements. To date, the BCI community has been focused on gross movements, and thus current algorithms are trained to make large movements. Fine motor control is necessary for making corrective movements, and performing tasks such as hand writing, painting, or preparing a meal. In order for BCI interfaces to function seamlessly in everyday life, they will need to successfully make large and small movements. Dr. Quick and Nicolas Card, an undergraduate student she was mentoring, began looking into the neural representation of reach amplitude, and results are described in her thesis (Quick, 2016). Our choice to use a velocity Kalman filter decoder despite the recognition that today's algorithms are not equipped to decode small movements may be the reason for this observation.

5.0 HUMAN PERFORMANCE ON THE CRITICAL STABILITY TASK

Figures and text in this chapter have been reproduced with permission from Quick et al. 2014 (Quick KM, Card NS, Whaite SM, Mischel JL, Loughlin P, Batista AP (2014) Assessing Vibrotactile Feedback Strategies by Controlling a Cursor with Unstable Dynamics. EMBS 2014). In this chapter, I describe our research on human subjects performing the Critical Stability Task to assess sensorimotor control. I entered the graduate program as the critical stability task was in its foundational stage, and worked closely with Kristin as she spearheaded the work in this chapter. This work created a foundation for the monkey work in the previous chapter (for which I had a larger contribution), as well as highlighting the importance of both the rendered feedback and the dynamics of the control signal on a person's ability to control an unstable system. My contributions to this chapter are largely supportive; I did not participate in data collection or the writing of the manuscript. However, I participated in weekly meetings with the other authors where we discussed analyses, experimental design, and results. I have been a driving force for CST since joining the project, and subsumed the role of Principle Investigator for the IRB protocol after Dr. Kristin Quick graduated. I am passing this role on to Ms. Nicole McClain, so we have the opportunity to continue human studies.

5.1 ABSTRACT

Brain computer interface (BCI) control predominately uses visual feedback. Real arm movements, however, are controlled under a diversity of feedback mechanisms. The lack of additional BCI feedback modalities forces users to maintain visual contact while performing tasks. Such stringent requirements result in poor BCI control during tasks that inherently lack visual feedback, such as grasping, or when visual attention is diverted. Using a modified version of the Critical Tracking Task (Jex et al. 1966) which we call the Critical Stability Task (CST), we tested the ability of 9 human subjects to control an unstable system using either free arm movements or pinch force. The subjects were provided either visual feedback, 'proportional' vibrotactile feedback, or 'on-off' vibrotactile feedback about the state of the unstable system. We increased the difficulty of the control task by making the virtual system more unstable. We judged the effectiveness of a particular form of feedback as the maximal instability the system could reach before the subject lost control of it. We found three main results. First, subjects can use solely vibrotactile feedback to control an unstable system, although control was better using visual feedback. Second, 'proportional' vibrotactile feedback provided slightly better control than 'on-off' vibrotactile feedback. Third, there was large intra-subject variability in terms of the most effective input and feedback methods. This highlights the need to tailor the input and feedback methods to the subject when a high degree of control is desired. Our new task can provide a complement to traditional center-out paradigms to help boost the real-world relevance of BCI research in the lab.

5.2 INTRODUCTION

Brain computer interfaces (BCI) predominately rely on visual feedback. While new decoding algorithms have improved control and increased the number of simultaneously controlled degrees of freedom (Velliste et al. 2008, Gilja et al. 2012, Hochberg et al. 2012), subjects must continuously watch their effector. However, real arm movements do not rely solely on vision for control, and grasping movements use very little, if any, visual feedback. Many researchers have begun to study non-visual feedback, whether it be rendered through vibrating tactors or intracortical microstimulation. Of the studies using tactile feedback, the feedback signal has represented the grasping force on real (Chatterjee et al. 2008, Pylatiuk et al. 2004) or virtual objects (Cheng et al. 1996), when an object is slipping from grasp (Jiang et al. 2009, Damian et al. 2012), and dynamics during virtual object manipulation (Rombokas et al. 2013, Bark et al. 2008, Leeb et al. 2013). In these studies, subjects used a variety of actions to control their task, including natural arm movements, EMG signals to simulate myoelectric prosthesis use, actual myoelectric prosthesis movements, or EEG signals. Researchers are beginning to investigate how to combine intracortical BCI control with non-visual information (Fitzsimmons et al. 2007, London et al. 2008, Venkatraman and Carmena 2011, O'Doherty et al. 2011). Real-world feedback depends intimately on the type of movement the user is making. We sought a task paradigm that could capture the interaction between the subject and his or her environment.

To create a virtual environment in which the interaction between a user and an object can be studied in its essence, we designed the Critical Stability Task (CST), which is based on the Critical Tracking Task introduced by Jex et al. in 1966 (Jex et al. 1966). The Critical Tracking Task has been used to assess motor performance during drug use and teleoperation, and to design vibrotactile feedback displays for balance prostheses (Ramaekers et al. 2006a, Zhai and Milgram

1993, Kadkade et al. 2003). In the CST, subjects must stabilize a first order unstable linear system. A familiar example of an unstable first order system is compounding interest, whereby the debt grows exponentially over time, and the larger the interest rate, the faster the debt grows. In the absence of external factors (payments, in this example), the account balance can be modeled mathematically by $y(t) = y_o e^{\lambda t}$ where y_o is the initial loan, and $\lambda > 0$ is the interest rate. We implemented this model to map the one dimensional position of a cursor on a screen; without external control, the cursor will rapidly drift off the screen. Subjects were required to maintain the cursor near the center of the screen to the best of their ability. The system was made more difficult to control by gradually increasing the parameter λ over time. We determined the largest λ that subjects could control, λ_C, and examined how that parameter depended on different forms of feedback.

During our task, subjects used either unconstrained hand movements or pinch force to control the unstable system. The position of the unstable system was rendered using visual feedback, two forms of vibrotactile feedback, or not rendered at all. Within each control method, the quality of the feedback rendering was assessed using the CST. Feedback renderings that allowed subjects to control greater instability (larger λ) were better feedback renderings. We found that subjects can use solely vibrotactile feedback to control an unstable system, although control was better using visual feedback. We also found that 'proportional' vibrotactile feedback provided marginally better control than 'on-off' vibrotactile feedback. Our final observation is that there was large intra-subject variability in the effectiveness of input and feedback methods. This highlights the need to tailor the input and feedback methods to the subject when a high degree of control is desired (Loughlin et al. 2011).

5.3 METHODS

Subjects performed the Critical Stability Task wherein they had to compensate for their motor or perceptual errors in order to control an increasingly unstable system. Subjects alternately used two methods to control the system and were alternately provided with four methods of feedback on the system's position.

5.3.1 Subjects

We tested 9 healthy subjects between the ages of 18 – 40 years without any history of any motor disorders. The subjects were 4 males and 5 females. All subjects gave their informed consent before being tested using a protocol approved by the Institutional Review Board at the University of Pittsburgh.

5.3.2 Implementing CST

Subjects controlled an unstable system G(s), as seen in Figure 5.1, whose transfer function is:

$$G(s) = \frac{\lambda(\mathrm{k})}{s - \lambda(\mathrm{k})}$$

where instability $\lambda(\mathrm{k}) > 0$ increases at a constant rate of 0.10 rad/second until the subject loses control. After converting the transfer function into state-space representation using observable canonical form, we discretized the continuous system for implementation on a computer:

$$x(k+1) = e^{\lambda(\mathrm{k})\tau}x(k) + \left(1 - e^{\lambda(\mathrm{k})\tau}\right)u(k)$$

$$y(k) = x(k)$$

In this representation, u(k) is the input signal to the unstable system at time step k, x(k) is the current state of the unstable system, y(k) is the output of the unstable system, and T is the sampling period in seconds, which was 5ms for force control and 10ms for hand control.

In our experiments described below, subjects interacted with the system in a variety of ways. We recorded the critical instability λ_C when subjects lost control of the system. Loss of control occurred when the unstable system's position surpassed a predetermined threshold, in our case, +/-5cm from the center. As a note, λ_C is robust to different thresholds; once the system becomes unstable, it is only a short period of time until the system would cross any threshold.

In the CST, the subjects received feedback on the position of the unstable system, y(k). For example, when controlling the system with hand movements and the cursor was right of center, the subject had to move his or her hand to the opposite position left of center to stabilize the system.

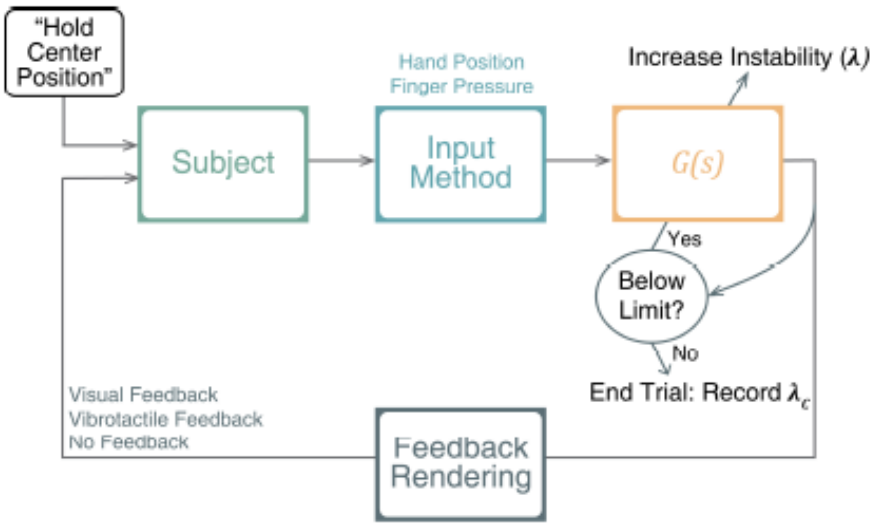

Figure 5.1: Diagram of Critical Stability Task. The possible input methods are shown in blue, the unstable system G(s) = λ(k)/(s- λ(k)) in yellow, and feedback rendering methods in gray.

5.3.3 Input methods to control system

Subjects were tested using two different control strategies with their dominant hand. Unconstrained hand movements were recorded using motion capture via an Improv system (PhaseSpace Inc., San Leandro, CA) and pinch force was measured using a force sensitive resistor (A201, FlexiForce, Tekscan, Boston, MA). The unconstrained hand movements involved mainly shoulder rotations that moved the hand through approximately 20cm of space. Only the horizontal component of their movement was used as input to the unstable system. A hand position of zero corresponded to a position directly in front of the subject. When subjects used pinch force, we subtracted an offset and multiplied by a gain such that - 10cm and +10cm were represented by 0N and 9.3N, respectively.

Subjects initiated each trial by positioning their input signal at zero, whether by moving their hand to this location or generating the pinch force that corresponds to a position of zero.

5.3.4 Feedback methods to render system

In order to control an unstable system, subjects must have feedback about the system. Four different types of feedback were rendered to the subjects. It is important to highlight that the feedback was about the state of the system, not the position of the hand or force on the transducer. The first feedback method was visual. The position of the unstable system was displayed on a monitor approximately 1m in front of the subject. When the subject used hand movements to control the system, the system's position was displayed along the horizontal axis. When the subject used pinch force to control the system, the system's position was displayed along the vertical axis, to make the visual feedback more congruent with the control scheme.

The second feedback method was ‘proportional’ vibrotactile feedback. The system’s position was rendered as the vibration intensity of two coin tactors (312-101, Precision Microdrives Limited, UK) on the non-dominate hand which the subject held motionless in their lap. During both hand control and pinch force control, one tactor attached to the thumb to indicate positive deviations and one tactor attached to the pinkie finger to indicate negative deviations.

Subjects oriented their hand with tactors so that the vibrotactile feedback was perceptually congruent to the control method. Thus, for hand control with the right hand, the left hand was palm-down with the thumb to the right of the pinkie. Then for pinch force control, the left hand was rotated vertically so the thumb was above the pinkie. The feedback was not meant to mimic proprioceptive or tactile feedback, but rather to provide information concerning the state of the system. We modulated each tactor’s “intensity” using a command voltage that was proportional to the system’s position. The tactor’s full voltage range was used to maximize amplitude and frequency modulation. An accelerometer attached to each tactor measured the feedback signal. The amplitude and frequency varied together, with amplitudes between 0.5- 12g and with frequencies between 50-170Hz.

The third feedback method was ‘on-off’ vibrotactile feedback. For any positive deviation of the system, the thumb tactor would vibrate at a fixed intensity near the middle of its operating range. Likewise, the pinky tactor would vibrate at a fixed intensity for all negative deviations. The fourth feedback method was no feedback. Subjects had to attempt to control the system without being given its current position.

5.3.5 Experimental design

Subjects completed one block of trials using pinch force control and one block of trials using hand control. Within each block, subjects completed five consecutive trials of each feedback method in random order, for a total of 40 trials per subject. Subjects had one practice trial before each block where they used visual and vibrotactile feedback. Trials lasted 5-45s depending on the subject's ability to control the system. Subjects had 2s of rest between hand movement trials and 15s between pinch force trials to reduce fatigue. Experimental sessions lasted approximately 35min. To determine the effectiveness of the eight different feedback-control methods, we compared the mean and standard deviation of λC using data from all subjects.

5.4 RESULTS

5.4.1 Visual feedback results

Figure 5.2 shows the mean λ_C scores for each of the eight different feedback-control methods. Subjects achieved a higher λ_C using pinch force control (3.19 ± 0.70) rather than hand control (2.83 ± 0.55) under visual feedback, as tested by Welch's t-test where the λ_C variance of each feedback – control method is not assumed to be equal ($p = 0.009$).

5.4.2 Vibrotactile feedback results

Figure 5.2 also shows how well subjects used vibrotactile feedback. For both control methods, the mean 'proportional' vibrotactile feedback λc scores were higher than 'on-off' vibrotactile feedback scores (Welch's t-test, hand movement, $p = 0.013$; pinch force, $p = 0.042$). Additionally, subjects used pinch force control better than hand movement control during both types of vibrotactile feedback (Welch's t-test, 'proportional', $p = 1.10e\text{-}5$; 'onoff', $p = 3.51e\text{-}7$).

These results might not be surprising, given that pinch force control afforded better control even under visual feedback. However, if we normalize the vibrotactile feedback results by the average λ_C achieved during visual feedback, both vibrotactile feedback methods still generate better control when using pinch force rather than using hand movements (Welch's t-test, 'proportional', $p = 0.0037$; 'on-off', $p = 9.68e\text{-}4$). These results can be seen in Table 5.1.

5.4.3 No feedback results

Finally, we show that all forms of feedback allow for better control than no feedback. "No feedback" means that we provided no feedback about the current state of the unstable system. However, subjects could still see their hand movements and feel their own pinch force. These results provide a baseline level of control when only the input to the system is known. It could be possible that subjects may take their input signal, run it through a mental simulation of the unstable system, and react appropriately. As we can see from these results, if such an internal model does occur, it does not work very well in this task.

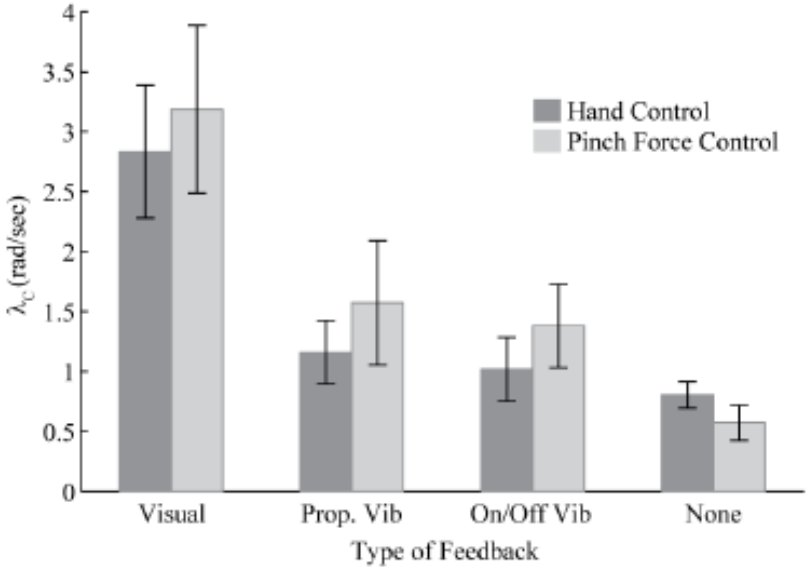

Figure 5.2: Mean and standard deviation of λ_C for each feedback-control method. These values are shown in Table 5.1, calculated using the five associated scores from each subject, for a total of 45 λ_C scores used in each bar.

Table 5.1: Vibrotactile Critical Instability Scores Normalized By Visual Feedback Scores

Control Method	**Feedback Method (Normalized λ_C)**	
	Proportional Vibrotactile	***On-Off Vibrotactile***
Hand Movement	0.41 ± 0.09	0.36 ± 0.09
Pinch Force	0.49 ± 0.16	0.43 ± 0.11

5.4.4 Intra-subject variability

Another finding is that there were striking differences between subjects in how well individuals could use a given type of control and feedback (Figure 5.3). Under visual feedback, pinch force was the best strategy on average, however, subjects K and P were able to use hand control better than pinch force control. Additionally, subjects G and H were able to use pinch force with vibrotactile feedback very well. These discrepancies in subjects' optimal control and feedback methods highlights the potential need to customize the feedback rendering strategy to the subject.

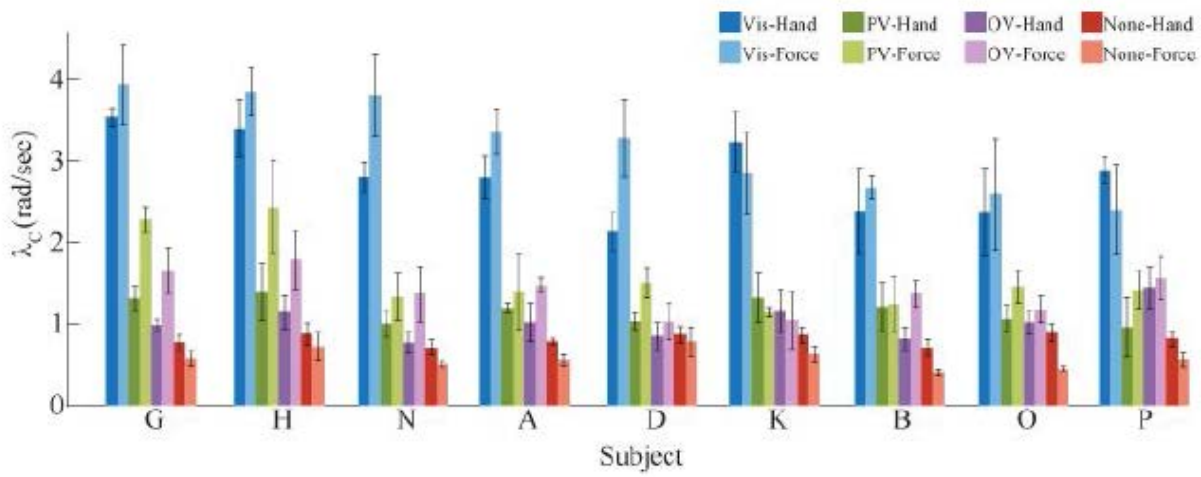

Figure 5.3: CST λ_C results ordered by the subject's performance on visual feedback – force control.

5.5 CONCLUSION

We interact with the objects around us through arm and hand movements. Only recently are BCIs being developed that interact with objects. As a first step towards a BCI that can interact with the environment through non-visual feedback, we sought to develop a suitable experimental paradigm. Our Critical Stability Task is a novel paradigm that bears some similarities to the control of grasp, in that the feedback depends on the interaction between the object and the user's actions, and that feedback is not necessarily visual. We investigated different vibrotactile feedback methods, and also different input methods. We found that the optimal input and feedback type differed somewhat between subjects. In this way, our approach provides a blueprint for the customization of tasks and feedback type for different applications, eventually including BCI control under non-visual feedback.

6.0 HISTOLOGICAL EVALUATION OF A CHRONICALLY-IMPLANTED ELECTROCORTICOGRAPHIC ELECTRODE GRID IN A NON-HUMAN PRIMATE

Figures and text in this chapter have been reproduced with permission from Degenhart et al. 2016 (Degenhart AD, Eles J, Dum R, Mischel JL, Smalianchuk I, Endler B, Ashmore RC, Tyler-Kabara EC, Hatsopoulos N, Wang W, Batista AP, Cui XT (2016) Histological Evaluation of a Chronically-implanted Electrocorticographic Electrode Grid in a Non-human Primate. J Neural Eng 13(4): 046019). In this chapter, I describe our research characterizing the effects of long-term implantation of electrocortigraphic (ECoG) electrodes on underlying cortical tissue. Throughout the duration of this project, I trained Monkey K to perform center out reaches while recording neural ECoG signals. The authors and I met weekly to discuss task advancements/parameters, training, and the direction of this work. Behavioral and neural analysis was spearheaded by Dr. Alan Degenhart, and histology and imaging analysis was performed by Mr. James Eles. For the manuscript, I worked on the introduction and performed the literature review.

6.1 ABSTRACT

Electrocorticography (ECoG), used as a neural recording modality for brain-machine interfaces (BMIs), potentially allows for field potentials to be recorded from the surface of the cerebral cortex for long durations without suffering the host-tissue reaction to the extent that it is common with intracortical microelectrodes. Though the stability of signals obtained from chronically-implanted ECoG electrodes has begun receiving attention, to date little work has characterized the effects of long-term implantation of ECoG electrodes on underlying cortical tissue. We implanted a high-density ECoG electrode grid subdurally over cortical motor areas of a Rhesus macaque for 666 days. Histological analysis revealed minimal damage to the cortex underneath the implant, though the grid itself was encapsulated in collagenous tissue. We observed macrophages and foreign body giant cells at the tissue-array interface, indicative of a stereotypical foreign body response. Despite this encapsulation, cortical modulation during reaching movements was observed more than 18 months post-implantation. These results suggest that ECoG may provide a means by which stable chronic cortical recordings can be obtained with comparatively little tissue damage, facilitating the development of clinically-viable brain-machine interface systems.

6.2 INTRODUCTION

Brain-machine interfaces (BMIs) utilize cortical signals to directly control external devices for restoration of motor function in individuals with disabilities (Moran, 2010; Schwartz et al., 2006; Wolpaw et al., 2002; Leuthardt et al., 2006). A critical component of a BMI system is the neural

recording modality used to extract meaningful information from the brain. The primary modalities used in clinical BMI systems are electroencephalography, electrocorticography (ECoG), local field potentials, and single/multi-unit activity. These methods can be characterized by factors such as performance, decoding stability, longevity, and invasiveness. The choice of a neural recording modality for a particular BMI application must appropriately weigh these factors.

Penetrating intracortical electrode arrays offer the highest spatial and temporal resolution in neural recording. However, the implantation of these arrays disrupts brain tissue and vasculature, and leads to a chronic inflammatory response hallmarked by an aggregation of astrocytes and microglia in a glial scar around the probe, as well as progressive neuronal degeneration at the vicinity of the implanted electrodes (Polikov et al., 2005; McConnell et al., 2009; Biran et al., 2005). This ultimately leads to recorded signal deterioration, manifested as a reduction in the number of electrodes recording individual neurons or a decrease in signal amplitude (Barrese et al., 2013; Schwartz et al., 2006; Moran, 2010; Freire et al., 2011; Collinger et al., 2012; Schwartz, 2004; Simeral et al., 2011; Chestek et al., 2011). Further, meningeal tissue proliferation and fibrous encapsulation have the potential to dislodge the implanted intracortical devices (Barrese et al., 2013).

ECoG is a recording modality where electrodes are placed either subdurally (below the dura) or epidurally (on the surface of the dura) to record electrical field potentials generated by aggregate cortical activity. Since ECoG arrays do not penetrate the cortex, they avoid blood-brain barrier disruption and mechanical strain between the brittle electrode and soft neural tissue, which are common for intracortical electrodes. This potentially mitigates some inflammatory burden on the brain. ECoG signals have been found to encode information about arm and hand

movements (Leuthardt et al., 2004; Schalk et al., 2007; Crone et al., 1998; Miller et al., 2007; Pistohl et al., 2008; Wang et al., 2009; Kub´anek et al., 2009; Ball et al., 2009; Miller et al., 2009; Chao et al., 2010; Acharya et al., 2010; Degenhart et al., 2011a; Shimoda et al., 2012; Chestek et al., 2013; Nakanishi et al., 2013), as well as auditory (Edwards et al., 2005; Trautner et al., 2006), visual (Lachaux et al., 2005), language (Crone et al., 2001; Mainy et al., 2007; Kellis et al., 2010; Wang et al.; 2011, Pei et al., 2011), and attentional processes (Tallon-Baudry et al., 2005; Jung et al., 2008; Ray et al., 2008). Encouraged by these findings, researchers have begun to investigate ECoG as a potential source of control signals for BMI devices. Human and non-human primate subjects have demonstrated up to three-dimensional control of computer cursors or prosthetic limbs using ECoG (Leuthardt et al., 2011; Schalk and Leuthardt, 2011; Leuthardt et al., 2004; Schalk et al., 2008; Wilson et al., 2006; Yanagisawa et al., 2012; Wang et al., 2013; Hotson et al., 2016).

Despite the promise of ECoG in BMI and neuroscience applications, very few studies have evaluated the long-term host-tissue response to either epidural or subdural grids. In humans and non-human primates, subdural and epidural arrays have been implanted for over one year with viable neural recording, however these studies did not report end-term tissue health (Shimoda et al., 2012; Morrell and RNS System in Epilepsy Study Group, 2011). Additionally, electrode grids are frequently implanted subdurally for up to 30 days in humans for purposes of epilepsy monitoring. While limited incidents of bleeding, infection, infarction, and functional deficits have been reported in association with these implants, no macroscopic fibrotic growth has been reported in otherwise uncomplicated surgeries (Fountas and Smith, 2007; Van Gompel et al., 2008; Wong et al., 2009). Some longer-term studies (> 1 year) of subdural spinal and cortical stimulators and probes reported encapsulation or dural thickening in the vicinity of the

implants (Nashold and Friedman, 1972; Pineda, 1978; Saitoh et al., 2000; Sillay et al., 2013). In rats, just one week after an epidural implantation, connective tissue overgrowth was observed (Schendel et al., 2014; Schendel et al., 2013). These studies followed tissue growth around epidural ECoG grids implanted in rats for up to 419 days, showing dural thickening under the arrays and tissue encapsulation over the top of the array as early as one month following implantation (Schendel et al., 2014). These findings are corroborated by findings of dural thickening at 6 months post subdural implant. Tissue proliferation was correlated with a rise in 1kHz electrical impedance as early as 1 week, with stabilization at 18 weeks post-implant (Henle et al., 2011). This would presumably reduce the quality of any neural signal recorded by the ECoG array, though no study has correlated long-term tissue reaction with neural signal quality. While Schendel et al. and Henle et al. investigated possible glial reaction to epidural and subdural implantation in the superficial layers of cortex, few studies explore the impact of subdural grids on deeper layers of the cortical tissue, particularly on neuronal health.

The current study explores the host-tissue response to a subdural ECoG grid implanted for 666 days, focusing on both cortical tissue health and fibrosis at the implant site, while also validating device performance by examining neural responses to overt reaching movements. We found that cortical thickness and neuronal density was unaffected by array implantation. Furthermore, while microglial density was increased in superficial cortical layers, they were in a resting stage morphology, and astrocyte activation was consistent with tissue not under the implant. Finally, though the grid itself was found to be encapsulated in a fibrous envelope upon explantation, robust modulation of ECoG signals was observed over 18 months post-implantation.

6.3 METHODS

All experimental procedures were approved by the Institutional Animal Care and Use Committee of the University of Pittsburgh and were in accordance with the National Institutes of Health's Guidelines for the Care and Use of Laboratory Animals.

6.3.1 ECoG grid implantation surgery

An adult male Rhesus monkey (macaca mulatta) was anesthetized, and a craniotomy was performed to expose the left motor and premotor cortex. The dura was retracted to expose an area approximately 2x2 cm between the arcuate and central sulci. A custom-built 15-channel, 1 mm thick silicone ECoG grid with 2mm diameter platinum electrode sites (Figure 6.1A, PMT Corp, Chanhassen, MN, USA), identical in construction to FDA-approved electrode grids commonly used for epilepsy monitoring, was placed directly on the exposed brain surface (Figure 1B-C). After positionin, the dura and the bone were reapproximated. Wires from the grid were connected to a Cereport pedestal connector (Blackrock Microsystems, Salt Lake City, UT, USA) affixed to the skull.

6.3.2 Neural recording and task control

Signals from the ECoG grid were recorded with a g.USBamp Biosignal Amplifier (g.tec Medical Engineering) and sampled at 1200 Hz. All recording, online processing, task control and presentation was performed using the Craniux Brain Computer Interface system (Degenhart et al., 2011b). Dura-facing electrodes 4 and 13 were used as reference and ground electrodes,

respectively, for all recordings (Figure 6.1A). Visual stimuli were presented on a 22-inch computer monitor placed approximately 0.8 meters from the monkey.

6.3.3 Hand control task

In order to validate device performance and ECoG signal modulation at long-term time points, we analyzed data recorded approximately 18 months post-implantation (day 542 to day 562 post-implant). During these experiments, the animal performed a standard 2D center-out task in a virtual environment, with the position of the hand controlling the location of a computer cursor in a two-dimensional frontoparallel plane. Hand position was tracked in real-time using an optical tracking system (Phasespace, San Leandro, CA) and rendered on a computer screen as a sphere in a virtual workspace. Trials began with the appearance of the cursor and central target; the animal was then required to move the cursor to the central target, and hold it there for 400–600 ms. One of eight peripheral targets would then appear, to which the animal was required to reach. A target hold time of 200 ms was enforced. The animal was provided a water reward immediately following successful completion of a trial.

Prior to offline analysis, hand control trials were visually examined for artifacts in both the time and time-frequency domains; all trials exhibiting artifacts during the central hold or target acquisition periods were excluded from further analysis. These artifacts were characterized by large-amplitude, broadband transients across the majority of recording electrodes, and are believed to be the result of jaw movements based on their consistent appearance during the reward period of each trial. Time domain data from the remaining trials were transformed into the time-frequency domain using the Burg autoregressive method (0 - 200Hz range, 2Hz frequency bands, 100th order, 100ms non-causal window, 33ms step size), log-transformed, then

normalized to pseudo Z-scores relative to the spectral power during the central target hold period. Trials were then manually aligned to movement onset using the cursor speed profile for each trial.

6.3.4 Explant

Electrodes remained implanted for a total of 666 days, after which the animal was sacrificed and the electrode grid removed. Surgical complications unrelated to the ECoG grid negated the possibility of perfusing the animal before removing the brain. After exposure of the skull, the original bone flap was removed to expose the dura. The skull proximal to the connector, the dura, and the encapsulated electrode grid were then removed as a single piece and the entire brain was extracted. The brain and the encapsulated array were then placed in a 10% formalin + 10% glycerin solution for 8 days followed by 10% formalin + 20% glycerin for 26 days to fixate the tissue. The brain was bisected along the midline and 3D renderings of each hemisphere were generated with a 3D scanner (Faro Platinum Arm, Faro, Warwickshire, UK). Surface topography of the implanted hemisphere was quantitatively compared to the mirror image of the non-implanted hemisphere using Geomagic Studio (Geomagic. Rock Hill, SC). Fixated tissue was then frozen and sectioned into 50µm sections for staining. Sections were cut perpendicular to the central sulcus. The electrode grid was carefully removed from the encapsulation "envelope," which was similarly fixed for 26 days and then stored in phosphate tris azide (PTA) solution until it was cut into 50µm sections for staining. Encapsulation tissue was cut perpendicular to the placement of the grid.

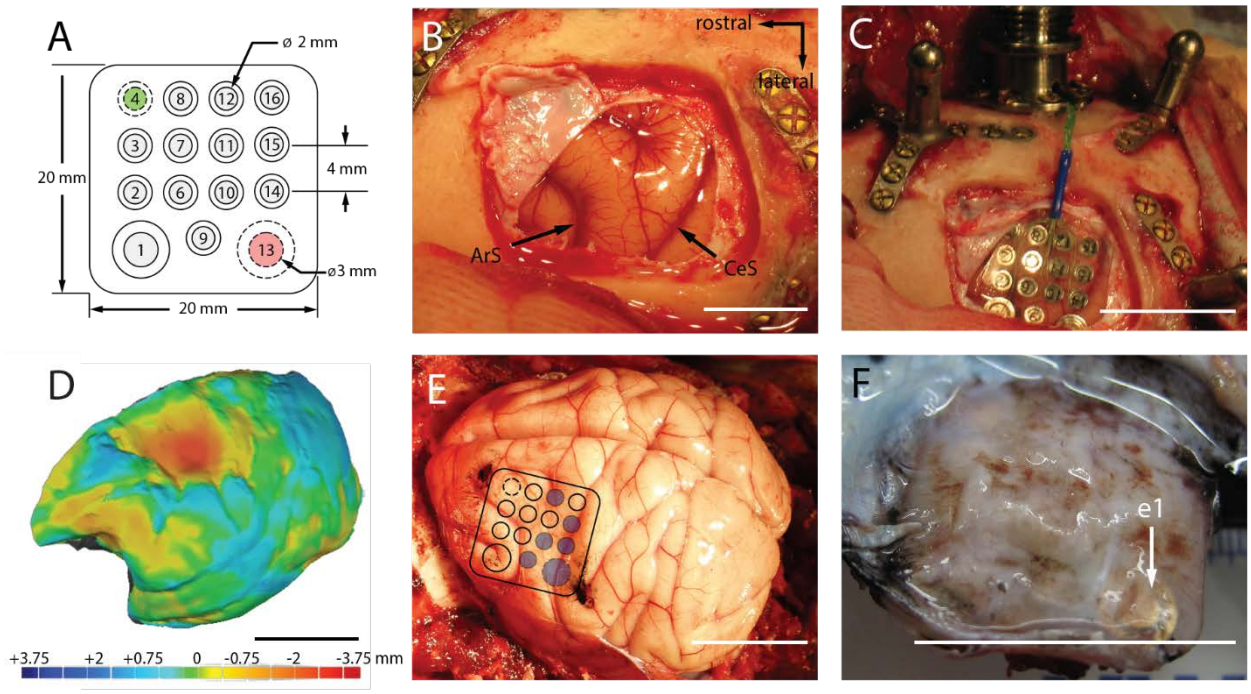

Figure 6.1: Electrode implantation and explantation. (A). Top view of the electrode grid. The neural recording electrodes (gray) face the cortical surface. The reference (#4, green) and ground (#13, red) electrodes face the dura. (B). Exposure of the left motor cortex prior to implantation (ArS: arcuate sulcus, CeS: central sulcus). (C). Placement of the electrode grid. (D). 3D rendering of the left-hemisphere superimposed on the mirror image of the right-hemisphere. Heat map denotes difference in surface topography between hemispheres in mm. (E). ECoG grid location superimposed on the postmortem brain. Blue circles indicate electrode sites targeted for histological analysis. Black ink marks the observed location of the rostral-medial and caudal-lateral corners of the grid. (F). Underside of the encapsulated grid following explantation. The location of electrode 1 (e1) is indicated by the white arrow. All scale bars are approximately 2cm unless otherwise indicated.

6.3.5 Immunohistochemistry

Cortical sections from implanted (left) and non-implanted (control, right) hemispheres were mounted on the same slide for comparison, and all slides were processed in the same session to minimize variability. A sample of dura mater distant from the edges of the tissue encapsulation (> 2cm) served as control dura mater for analysis of the encapsulation tissue.

Antibodies for cortical tissue were directed to neurons (NeuN, 1:200, Millipore), astrocytes (GFAP, 1:200, SeroTec), or microglia (Iba-1, 1:500, Fisher); antibodies for encapsulation/dura mater tissue were directed to macrophages (Iba-1, 1:500, Fisher) or fibroblasts/macrophages (Vimentin, 1:250, Millipore). Tissue was first blocked for 30 minutes in sodium citrate buffer (0.1M citric acid, 0.1M sodium citrate, pH 6.0) in room temperature followed by a peroxidase block (10% methanol, 3% hydrogen peroxide) for 20 minutes in room temperature on a shaker. Then, tissue was blocked in a serum containing blocking buffer (5% normal goat serum, Jackson Labs; 0.1% Triton X-100, Sigma) for one hour. Tissue was incubated in primary antibody for 12-24 hours. Following washes in phosphate buffer saline (PBS), tissue was incubated in 1:250 Alexa Fluor 488 and/or 633 (Invitrogen) for 2 hours at room temperature, followed by more PBS washes, 10 minute incubation in Hoescht 33342 (1uL/1mL; Invitrogen) stain, and more PBS washes. Coverslips were then mounted with Fluoromount-G (Southern-Biotech).

6.3.6 Confocal imaging

All confocal imaging was performed with an Olympus Fluoview 1000 confocal scanning microscope (Olympus Corporation, Tokyo, Japan). All images were taken with a 20X or 40X

objective to optimize cellular resolution and image frame size. Images were taken at multiple focal depths for each frame in order to image the full depth of a tissue slice. This ensured that image analysis was not biased by choice of a single image depth. Confocal laser power, photomultiplier tube voltage (the sensitivity of the image detector), and photomultiplier offset (background level of image detector) were selected to ensure that image pixels did not exceed upper or lower detection limits. Images (n = 5 tissue sections per stain) were collected from cortical regions directly under electrode sites in Brodmann Area 4 (specific sites denoted by blue circles in Figure 1E) on the ECoG array or from the thickest region of both the dorsal and ventral sides of the center of the tissue encapsulation. Images collected were only excluded from analysis on grounds of poor quality of signal, photobleaching, or severe tissue tears during processing. For cortical tissue imaging, images from the contralateral hemisphere were collected for comparison. Images were matched to the same sagittal slice depth and anterior-posterior region of interest as the ipsilateral hemisphere. Tissue encapsulation images were compared to images from random regions of interest of control dura mater retrieved from > 2cm from the tissue encapsulation.

To determine cortical layers and cortical thickness, disconnected images of cortex were stitched to create continuous high-resolution images of the entire cortical depth using Fiji, an Image-J (NIH) plug-in (Preibisch et al., 2009). Layers I/II-III were discerned from layer V by the location of layer V giant pyramidal cells (Matelli et al., 1991). Stitched images were used to measure cortical depth (n = 5) between conditions. Neuronal and microglial cell densities were determined for layers I/II-III and V with manual counting facilitated by Image J Cell Counter (n = 5). Layer I microglia morphology was assessed as previously described (Stence et al., 2001; Nimmerjahn et al., 2005; Kozai et al., 2012). Microglia were deemed to be ‘unresponsive’ if they

were ramified (resting) or activated but not extending processes to the cortical surface, and 'surface directed' if they had activated or amoeboid morphology, with processes extended to or along the cortical surface. Because GFAP labels extensive networks of astrocytic processes, discerning individual cell bodies for cell counting was not possible. Instead, the proportion of cortex occupied by reactive astrocytic signal (% GFAP signal) was determined by first setting a pixel intensity threshold to the mean pixel value of layer I/II-III, where the most intense signal was localized. Because the majority of pixels in a given image are not GFAP-signal, the pixels below the mean can be discounted as noise. Once thresholded, the GFAP signal was determined by automating a count of the non-zero pixels (n = 5). Implanted cortex and contralateral cortex were compared for all metrics by t-tests with significance level of $\alpha = 0.05$.

We identified cell-types in the encapsulation tissue based on morphology and antibody staining. Vimentin(+)/Iba-1(+) and vimentin(-)/Iba-1(+) cells were considered to be macrophages if found in the meninges or microglia if found in the parenchyma. Vimentin(+)/Iba-1(-) cells were considered to be fibroblasts. Multi-nucleated cells were considered cells that contained more than one Hoescht 33342 labeled nuclei in a single cell body. These cells are often found in pathological inflammatory conditions or in the presence of chronically implanted foreign bodies (Lynn et al., 2011; Brodbeck et al., 2002; Anderson et al., 2008).

6.3.7 Collagen-I imaging

Collagen I, a key component of tissue encapsulation, can be visualized using second- harmonic generation (SHG) imaging. SHG imaging takes advantage of a second-order nonlinear optical property of collagen type I to visualize an intrinsically-generated optical signal that can be used to locate and quantify collagen I in tissue. This is preferred to traditional histological staining

protocols, which have been shown to have less signal specificity and require chemical processing that may alter the tissue quality (Strupler et al., 2007; Cox et al., 2003).

SHG images of tissue encapsulation and dura mater were captured using a laser with a Nikon A1Plus multiphoton scanning confocal microscope and Nikon NIS- Elements Microscope Imaging Software. SHG was generated at an 830nm wavelength, and signal was collected via a bandpass filter that isolated tissue auto-fluorescence (435 − 700nm) and a low-pass filter that isolated SHG signal (<492nm). Images were taken with a 25X objective to maximize signal resolution and imaging frame; stitching software (EIS-Elements Microscope Imaging Software, Nikon) was used to consolidate disconnected images to make a seamless, high-resolution image of the encapsulation through the dorsal-ventral plane.

Encapsulation and dura mater thickness were determined by measuring average thickness of tissue extent as denoted by auto-fluorescence. Because SHG signal was confined within an uninterrupted, fibrous area, percent SHG-signal was measured by dividing the average thickness of SHG area by the total tissue thickness. Measures were generated for encapsulation tissue from the ventral and dorsal sites, as well as for control dura mater ($n = 5$ tissue sections for all groups). Encapsulation and dura mater thickness and percent SHG signal were compared between ventral encapsulation, dorsal encapsulation, and control dura mater groups by one-way ANOVA tests with Tukey's post-hoc tests. Significance for all comparisons was defined to be $\alpha = 0.05$.

6.4 RESULTS

6.4.1 Cortical architecture

Upon sacrifice and explantation, we found that the ECoG grid was fully encapsulated in fibrous tissue that was contiguous with the dura mater (Figure 6.1F). The brain underneath the encapsulated ECoG grid was mechanically depressed. We assessed the extent of this depression by generating a 3D rendering of the brain's surface topography. Then, the image of the implanted hemisphere was superimposed onto the non-implanted hemisphere. This allowed us to quantify the topographic differences between the two hemispheres (Figure 6.1D). We found that the brain region under the ECoG grid was mechanically depressed by as much as 3.63 mm relative to the same region of the non-implanted hemisphere.

To determine if this gross morphological change resulted in alterations of cortical cytoarchitecture, we evaluated neuronal and glial density as well as cortical thickness under the grid. We compared these metrics of cortical structure to those of the corresponding cortex in the opposite hemisphere (Figure 6.2). The density of neurons labeled with the NeuN antibody in layers I/II-III or layer V was not significantly higher (Student's t-test, $p = 0.5$ and 0.32, respectively) in the cortex under the grid versus the contralateral cortex (Figure 6.2A). Similarly, the signal intensity of GFAP antibody labeling of reactive astrocytes in layers I/II-III or layer V was not significantly different between the two hemispheres (Figure 6.2B; 0.18 and 0.73, respectively). Only the density of microglia labeled with the Iba-1 antibody and located in the superficial layers (I/II–III) exhibited a significant increase under the array (Figure 6.3A; t-test $p = 0.027$ for layers I/II-III; 0.24 for layer V). The microglia both under the array and in the contralateral cortex exhibited a qualitatively similar morphology, with only cell density along the

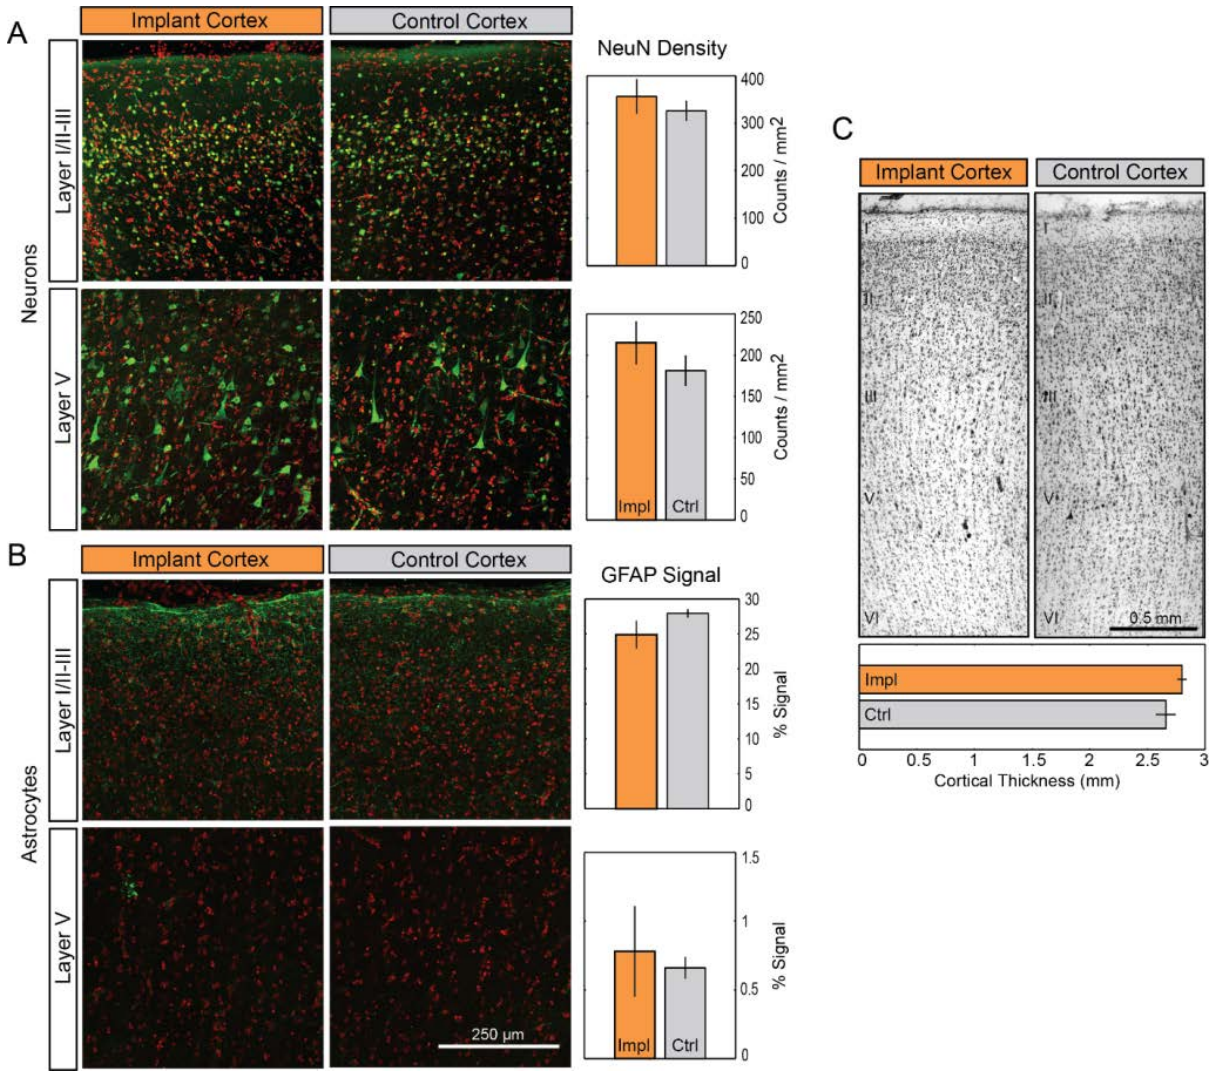

Figure 6.2: Long-term ECoG grid implantation causes minimal changes in cortical cytoarchitecture. (A-B). Neither the density of NeuN-labeled neurons (A; green) nor the signal intensity of GFAP-labeling in astrocytes (B; green) located in layers I/II-III or layer V were significantly affected by implantation. Cell nuclei (red) counterstained with Hoescht 33342. Data presented as mean ± SEM; * denotes significant difference from control ($p < 0.05$). (C). Comparison of Nissl-stained motor cortex between implanted and control hemispheres. Cortical layers are indicated by roman numerals I- VI. Impl: implanted cortex. Ctl: control cortex.

dorsal surface of the cortex of either hemisphere (7.6 ± 1.6 versus 6.6 ± 1.8 cells/mm for the implanted and control cortical surface, respectively) showing surface-directed morphology (defined in Methods) (Figure 6.3B). These cells typically had processes polarized parallel to the cortical surface. Finally, the thickness of the Nissl-stained cortical tissue under the grid (Figure 2C, top; 2.8 ± 0.04mm) was not statistically different from that of the contralateral cortex (Figure 2C, bottom; 2.7 ± 0.09mm; t-test: p = 0.34). These findings were qualitatively corroborated by observing the transition region of tissue at the edge of the ECoG array and tissue immediately adjacent to the implanted region, where limited changes were observed (Supplemental Figure 6.7). Taken together, these tests provide evidence of little to no cytological changes in the cortex underlying the ECoG array.

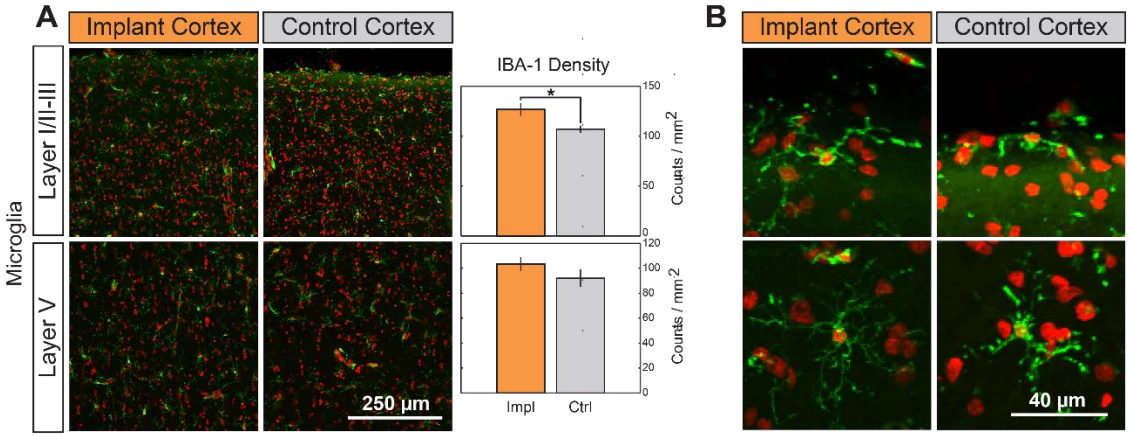

Figure 6.3: Chronic implantation yields higher microglial density with no change in cell morphology. (A). The density of microglia (green; nuclei in red) was significantly increased in layers I/II-III but not in layer V following implantation. Data presented as mean ± SEM; * denotes significant difference from control (P < 0.05). (B). Layer I microglia show no morphological indication of inflammatory response. A small population at the cortical surface of implant and control cortices are polarized along the curvature of the brain, all microglia beneath the surface are unresponsive.

6.4.2 Fibrous encapsulation

Chronic subdural ECoG implantation resulted in fibrous encapsulation of the grid. The grid was removed by making an incision along the anterior portion of the encapsulation and pulling the grid with forceps. Surprisingly, the grid offered little mechanical resistance to removal, indicating that adhesion between the grid and encapsulation tissue was minimal. Using second-harmonic generation (SHG) imaging, we detected collagen I in sections of both the tissue encapsulation and control dura mater (> 2cm from implantation site). Using filters to image second-harmonic signals and tissue autofluorescence simultaneously, we quantified both the thickness of encapsulation tissue and the percentage of encapsulation tissue that was collagen I-positive (Figure 6.4). Because the dorsal portion of the encapsulation emerged from the original, autografted dura mater, we analyzed it separately from the ventral portion of the encapsulation, which grew de novo following initial implantation. Both sides of the tissue encapsulation were compared to control dura mater taken more than 2cm from the implantation site. There were statistically significant differences in the thicknesses of the tissues (one-way ANOVA: $F(2, 14) = 136.13$, $p < 0.001$), with both dorsal encapsulation (0.82 ± 0.04mm) and ventral encapsulation (1.76 ± 0.09mm) being thicker than control dura mater (0.36 ± 0.03mm; Tukey's post-test: $p = 0.001$). The ventral encapsulation was also significantly thicker than dorsal encapsulation ($p < 0.001$). The encapsulation was presumably the major contributor to the visible depression of the cortex under the grid.

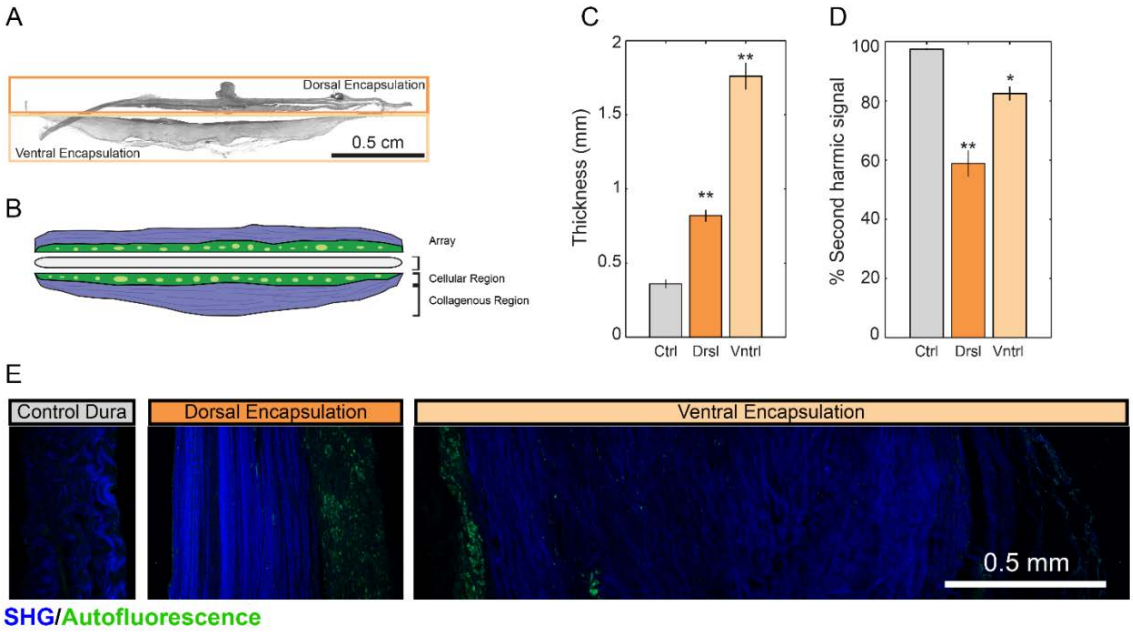

Figure 6.4: Second-harmonic imaging of fibrous encapsulation reveals fibrous, cell-sparse regions and cell-dense regions in both dorsal and ventral aspects of encapsulation. (A). Sample image of full tissue encapsulation slice. (B) Schematic representation of encapsulation components. (C). Comparison of thickness of dorsal and ventral aspects of encapsulation tissue to control dura. (D). Percentage of SHG(+) tissue was significantly reduced in encapsulation tissue. (E). Sample images of dorsal encapsulation, central encapsulation, and control dura with SHG signal shown in blue and tissue autofluorescence shown in green. Data presented as mean ± SEM. Asterisks * and ** denote significant differences from control at $p < 0.01$ and $p < 0.001$, respectively.

SHG imaging revealed encapsulation tissue to be comprised of a cellular region that did not express strong SHG signal and a collagenous region that was strongly SHG(+) (Figure 4B,E). Using the tissue thickness derived above, we were able to assess the relative proportions of cellular and collagenous regions by measuring the area of collagenous region (SHG(+) region) and dividing it by tissue thickness. This showed that the proportion of collagenous region was significantly different between the tissues (one-way ANOVA: $F(2, 14) = 44.33$; $p < 0.001$). Control dura mater had a significantly higher percentage of collagenous tissue (96.4 ± 0.33%) than either ventral encapsulation (82.5 ± 2.3%; Tukey's post-test: $p < 0.01$) or dorsal encapsulation (58.9 ± 4.5%; $p < 0.001$). The percentage of collagenous tissue in the ventral encapsulation was also significantly greater than that of the dorsal encapsulation ($p < 0.001$).

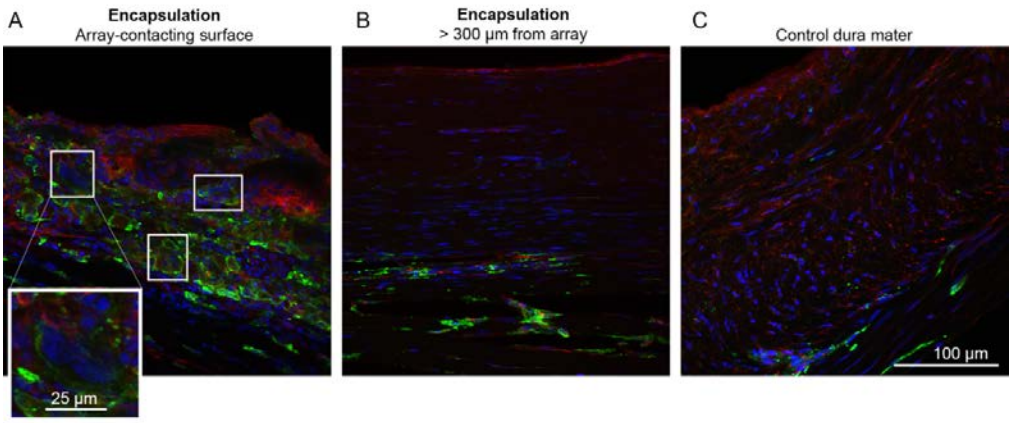

Figure 6.5: Immunohistochemical staining of encapsulation tissue. Tissue was stained for nuclei (blue; Hoescht 33342) and antibodies directed to macrophages (green; Iba-1) or macrophages/fibroblasts (red; vimentin). (A). Array-contacting aspects of the encapsulation were highly cell dense, populated with macrophages (vimentin(+ or -)/ Iba-1(+)) as well as fibroblasts (vimentin(+)/Iba-1(-)). Boxes indicate multi-nucleated giant cells. Inset: Magnification of a multinucleated giant cell. (B). Distal portions of encapsulation were hallmarked by elongated fibroblasts and macrophages (vimentin(-)/Iba-1(+)). (C). Control dura mater is largely composed of elongated fibroblasts with infrequent macrophages.

In order to determine the cellular composition of the encapsulation we used immunohistochemistry (described in section 6.2.5). We identified fibroblasts and macrophages in all tissue groups. Control dura mater was largely composed of fibroblasts, many of which exhibited elongated nuclei (Figure 6.5C), consistent with previous literature (Adeeb et al., 2012). Macrophages were sparsely distributed. This resembled the "collagenous" region (>300μm from the array; Figure 6.6B) of the tissue encapsulation, which also contained elongated fibroblasts and macrophages. The "cellular" region of encapsulation (< 300μm from the array; Figure 6.5A) was highly cell dense with round, mononuclear macrophages as well as multinuclear, foreign body giant cells (vimentin(+)/Iba-1(+)). We made the mononuclear/multinuclear distinction based on nuclei count (Figure 6.5A inset). The encapsulation was organized as a gradient, with "cellular" tissue closest to the array exhibiting more inflammatory cell activity, and "collagenous" tissue more distal to the array more closely resembling healthy dura mater.

6.4.3 Physiological recordings

In order to validate long-term signal modulation, we examined ECoG signals during center-out reaching task trials conducted between day 542 and 562 post-implantation. Signals exhibited clear modulation with target direction (Figure 6.6, Supplemental Figures 7.8, 7.9). Characteristic decreases in the mu and beta frequency bands (10-30 Hz), in conjunction with increases in the high-gamma band (> 60 Hz), were observed. High-gamma band modulation was found to be the strongest over the 70 – 100 Hz frequency range, and was tightly locked to movement onset. Thirteen of the fifteen electrodes exhibited reach-modulated signals. Of the two electrodes not exhibiting reach-related activity, one was a reference (dura-facing) electrode and the other an

electrically intact electrode that was not recording due to a failure in the cabling connecting the Cereport adapter and neural recording amplifier.

Prior to hand control experiments, a number of recording sessions devoted to BMI control experiments were conducted. However, during post-hoc analysis of these data we discovered that the animal had developed a strategy of using artifacts, possibly the result of jaw movement, to generate directionally-modulated broadband increases in spectral power. We now believe that our earlier report of stable multi-day BMI control was due in part to this strategy (Ashmore et al., 2012). The presence of these artifacts precludes further analysis of the brain control data, apart from a baseline confirmation of the stability and robustness of the ECoG recordings. We have demonstrated this with representative mean electrode root-mean speed (RMS) amplitude measurements of ECoG signals, which initially dropped, but stabilized by day 300 post-implant (Supplemental Figure 6.10C,D).

We also tracked mean 20Hz impedance of all functional electrode sites during some recording sessions from day 52-562 post-implant as well as after electrode grid explantation (Section 7.5.3: Supplemental Section C). Impedance was relatively stable through this time frame, though it fluctuated following surgical interventions (Supplemental Figure 6.10A). Depending on the impedance measurement method, mean impedance of the grid dropped 6-36 kOhms following removal from the tissue encapsulation (Supplemental Figure 6.10B).

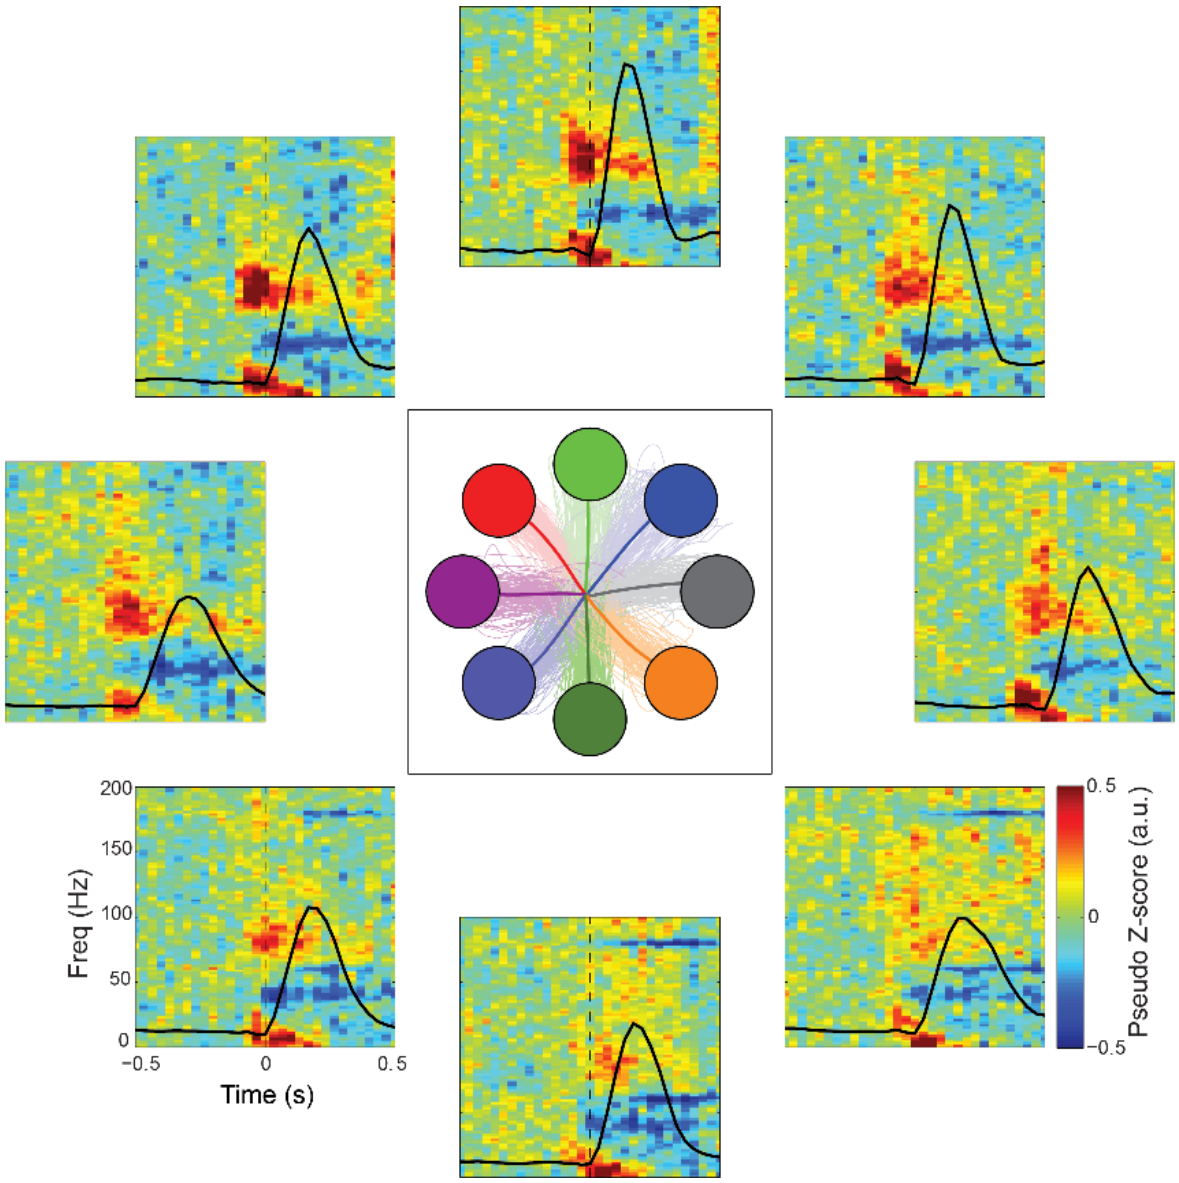

Figure 6.6: ECoG signal modulation during 8-target center-out reach tasks. Average time-frequency data are shown for a single electrode (e10) for all reach directions. Averaged (thick lines) and individual trajectories (thin lines) for each target are shown in the center panel. Time-frequency data were normalized with respect to the spectral data during a central hold period preceding each trial. Black lines show average speed profiles for each target.

6.5 DISCUSSION

6.5.1 Histological findings

We examined the foreign-body response to a subdural ECoG electrode grid nearly two years after implantation. There was fibrotic growth around the electrode grid, resulting in a shallow mechanical depression of the cortex under the array. Despite this, the cellular characteristics of the cerebral cortex underneath the ECoG grid were consistent with that of the tissue of the contralateral hemisphere, where no grid had been placed. Most importantly, cortical thickness and neuronal density of the tissue under the array were statistically indistinguishable from the contralateral tissue, with no morphological differences apparent at any spatial scale. After accounting for discrepancies in methods, our measurements of thickness and neuronal density for both hemispheres agreed with those of previous anatomical studies of primate frontal cortex (Matelli et al., 1991; Gittins and Harrison, 2004). Noting that cortical thickness and layering was unaffected by the gross mechanical deformation of the brain's surface, it is plausible that the displaced brain was merely pushed into the ventricles, as is observed in cerebral edema and subdural hematoma (Kim and Gean, 2011). Only the microglial density of the superficial cortical layers was significantly different between the implanted and non-implanted hemispheres. It is possible that the persistent, unactivated microglia population was part of the foreign body response to the implanted ECoG array, where increased macrophage density persists in the vicinity of the implant for its lifetime (Anderson, 2001; Sanders et al., 2000). It is also conceivable that this activation was not in response to a foreign body, but rather from pervasive mechanical stress caused by the fibrous encapsulation compressing the brain (Roth et al., 2014; Ding et al., 2008). With the exception of a population of cells at the cortical surface of both

hemispheres, layer I microglia were either ramified or polarized but without processes that extended to the cortical surface, indicating that these microglia were not actively responding to trauma or other noxious stimuli at the surface of the brain (Stence et al., 2001; Roth et al., 2014; Nimmerjahn et al., 2005; Kozai et al., 2012). The microglia at the surface of the brain were polarized with processes extending parallel to the cortical surface, similar to a cell type that has been described in healthy mouse cerebellum and rat prefrontal cortex, suggesting the observed cell type is not a result of pathology (Vela et al., 1995; Kongsui et al., 2014). The astrocytic GFAP expression between control and implanted hemispheres was not different. For both hemispheres, we observed low GFAP expression in the gray matter and relatively strong expression in the glia limitans. This expression pattern has been well documented in healthy macaque brain, with gray matter GFAP expression increasing only in response to trauma or chronic foreign body implantation (Eng et al., 2000; Griffith and Humphrey, 2006; Peters and Sethares, 2002). Qualitatively, our finding of low levels of gray-matter GFAP(+) cells under the array suggests that the array was not actively causing trauma to the cortex.

During the grid implantation surgery, we resected the dura mater, replaced it over the ECoG array, and sutured it in place. After the 666 days of implantation, dura mater/fibrous encapsulation tissue was found in a contiguous piece surrounding both the top and bottom of the ECoG. Since there were only leptomeninges separating the brain and array at the time of the implant, we assume that the ventral encapsulation grew de novo post-implantation. This is similar to recent findings by Schendel et al., who reported progressive fibrous overgrowth of epidural ECoG grids, with complete encapsulation as early as one month post-implantation (Schendel et al., 2014; Schendel et al., 2013). The cellular distribution in the dorsal and ventral tissue encapsulation was distinct from dura mater elsewhere in the brain, and implied that the

wound-healing response to implantation consisted of a stereotypic foreign body response, which involves aggregation of mononuclear macrophages and multinucleated foreign body giant cells to the implant site and encapsulation of the device in a collagenous envelope. Aggregated cells and tissue encapsulation generally persist through the lifetime of an implant, with pro-inflammatory cytokine expression diminishing within the first month as anti-inflammatory/pro-wound healing cytokines are expressed (Lynn et al., 2011; Brodbeck et al., 2002; Anderson et al., 2008). The fibrous encapsulation demarcates the final stage of wound healing in which the tissue disrupted by implantation is either regenerated from cells of the original cell type, or replaced with fibrous connective tissue. Given that dura mater is already largely fibrous connective tissue and mesenchymally derived fibroblasts, it was not clear to us the extent to which the tissue encapsulation was fibrous encapsulation or remodeled/regrown dura mater (Adeeb et al., 2012; Anderson, 2001; Anderson et al., 2008). We observed a gradient where tissue proximal to the implant more closely resembled fibrous encapsulation, and tissue distal to the implant more closely resembled control dura mater.

Both the dorsal encapsulation and ventral encapsulation were thicker than the control dura mater, which would be expected of a foreign body tissue encapsulation. In the case of autografted, dorsal encapsulation, dural thickening may have also been an inevitable consequence of craniotomy and/or durotomy that was simply exacerbated by the presence of a foreign body. This is seen in epidural ECoG implants, where encapsulation with ventral and/or dorsal dural thickening has been reported in long-term implants (Schendel et al., 2014). Merely performing a craniotomy triggered a 3.8 fold increase in dural thickness at 3 weeks, with a reduction to a 2.6 fold increase at 3 months in New Zealand white rabbits (Nunamaker and Kipke, 2010). Replacing dura with an alginate hydrogel resulted in a 2.8 fold increase of dural

thickness of regrown dura at 3 weeks and a 3.1 fold increase at 3 months. Dural thickening of 2mm 8 weeks after a 2cm dural incision has also been observed in coonhound dogs; following application of a poly(ethylene) glycol based dural sealant, the healed dura was found to have thickened as much as 4mm (Preul et al., 2003). Meningeal cells almost double collagen production following sub-arachnoid hemorrhage in rats (Sajanti et al., 1999); computational models of collagen I fibrosis following biomaterial implantation corroborate this (Su et al., 2011). In these experiment-validated models, increasing numbers of fibroblasts at the implant site results in significantly increased collagen deposition. Since the predominating cell type of dura mater is the fibroblast, it is plausible that we might expect pronounced collagen I production following implantation.

Despite the degree of encapsulation, the ECoG grid was extracted from the fibrous tissue with little effort, indicating relatively minor adhesions between the encapsulation and device. This is not surprising given the lack of porosity and surface features on the silicone grid. Previous studies have demonstrated that smooth, non-porous dural substitutes are less susceptible to fibrosis and adhesion formation (Barbolt et al., 2001; Sayama et al., 2014).

While we have shown that the foreign body response to chronic ECoG grid implantation can result in grid encapsulation after approximately 22 months, we were unable to determine its exact time course. It is unclear whether the encapsulation was stable, still growing, or perhaps shrinking, at the time of electrode explantation. Subdural ECoG electrodes implanted up to 30 days clinically for epilepsy monitoring do not exhibit such encapsulation (Fountas and Smith, 2007, Van Gompel et al., 2008, Wong et al., 2009, Wang et al., 2013), so it is likely that the subdural implant encaspulation response occurs on the order of months rather than weeks. Meningeal thickening without encapsulation has been observed in mirco-ECoG arrays in rats at 6

months post-implant, though no other time-points were assessed (Henle et al., 2011). In contrast, epidural implant encapsulation has been observed as soon as one month post-implant, with dramatically slower tissue encapsulation observed under an epidural array that had torn the dura mater during implantation (Schendel et al., 2014; Schendel et al., 2013). This suggests that there may be different foreign body response mechanisms for implants with different degrees of invasiveness. To our knowledge, there is no study directly comparing implantation depth to explore possible foreign body response mechanisms.

6.5.2 ECoG recording quality

ECoG electrodes provided recordings of physiological signals for nearly two years. Issues with the animal that were unrelated to the ECoG grid determined the termination date of the study, but we believe that signal quality may have persisted past two years. Of the fifteen electrodes on the array, only one lost recording capability during the course of our study. Post-explantation, we determined that failure was on account of a faulty wire connecting the Cereport adapter to the neural amplifier, and not due to the tissue response. We found that all functional electrodes showed signals that were temporally modulated and spatially tuned during a reaching task.

We conducted extensive BMI control tasks as part of this study. By the time we detected that the animal had developed a strategy of generating an electromyographic artifact (believed to be the result of jaw or face muscle contraction) to control the cursor, we were unable to train him to abandon this strategy. This made it difficult to study the functional properties of ECoG signals during the BMI control sessions. Despite the lack of longitudinal BCI performance data, both impedance and RMS amplitude measurements were relatively stable from day 56 (our earliest time-point measured) until day 562 post-implantation. The stable impedance and RMS amplitude

suggest that the encapsulation did not significantly compromise the device functionality. However, future studies using a stereotypical experimental paradigm, such as the center-out task, will likely be able to better characterize changes in ECoG signal properties throughout the entire lifespan of an ECoG implant.

6.5.3 Implications and future directions

We believe our results have implications for the viability of ECoG for long-term high-resolution brain recording. In addition to its use as a recording modality for brain-machine interfaces, ECoG has increasingly become a neuroimaging method of choice in a variety of neuroscience fields and non-BMI neural recording and neuromodulation applications. The potential for subdural ECoG grids to remain implanted for extended periods of time without damaging the cortex could facilitate the study of cognitive processes over long timescales. Lack of cortical damage combined with the ease of removal of the ECoG grid from encapsulation tissue may provide the possibility for re-implantation in case of device failure. This is not practical for intracortical electrodes, which typically damage neural tissue upon insertion (Barrese et al., 2013). While ECoG grid encapsulation presents as a potentially detrimental consequence of implantation, many strategies can and have been pursued to minimize the foreign body response. These include altering the shape of array substrate (Schendel et al., 2014; Schendel et al., 2013; Yamakwa et al., 2010), increasing array flexibility (Kim et al., 2010; Yeager et al., 2008; Rubehn et al., 2009), applying anti-fouling or biomimetic surface treatments (Collier et al., 2004; Kolarcik et al., 2012), and releasing anti-inflammatory drugs from the array substrate or electrodes (Norton et al., 2005; Weaver et al., 2014). Use of such strategies may help to further

increase the stability of long-term ECoG recordings by eliminating changes in recording quality resulting from the foreign body response to subdural ECoG grids.

As the presented work constitutes a case study of long-term grid implantation in a single animal, future studies are required to fully assess the impact of chronically-implanted ECoG electrodes. Nevertheless, our findings of meningeal thickening, encapsulation, and fibrosis echo studies on subdural and epidural implants in rats (Henle et al., 2011; Schendel et al., 2013; Schendel et al., 2014) as well as in long-term (> 1 year) subdural and epidural implants in humans (Nashold and Friedman, 1972; Pineda, 1978; Saitoh et al., 2000; Sillay et al., 2013). There are fewer studies on the health of neural tissue underlying these implants. Additionally, we believe that our results provide an analytical framework for further investigation into the effects of chronic implantation of ECoG electrodes on the health of cortical tissue.

This study is an important first step toward fully assessing the long-term use of chronically-implanted ECoG electrode grids. Minimal cortical tissue damage from chronic electrode implantation suggests that ECoG may provide the capability to record physiological signals from the cortex for extended periods of time. Ultimately, this highlights the utility of ECoG as a valuable tool for long-term BMI, clinical, and neuroscientific studies.

6.6 SUPPLEMENTAL MATERIAL

6.6.1 Supplemental Section A. Histology at the edge of the ECoG grid

Fluorescence microscopy was used to assess the transition area between the edge of ECoG grid and the tissue immediately adjacent to the device (Supplemental Figure 6.A1). We observed anatomical continuity in the transition area, with limited changes in neuron (red) and microglia (green) populations as well as the cortical layer structure.

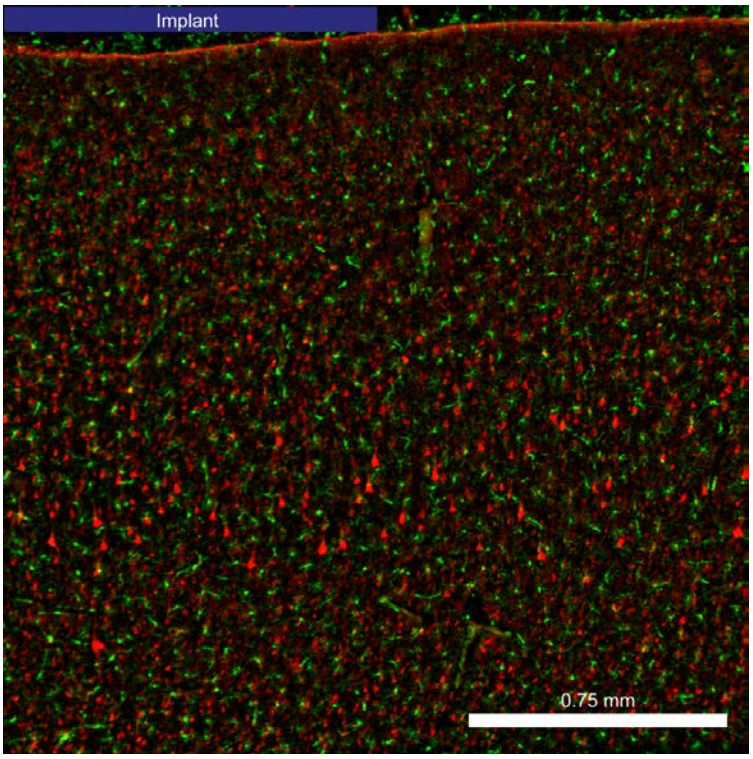

Figure 6.7: Fluorescence microscopy of neural tissue under the edge of the ECoG array and the adjacent tissue. No observable changes in cortical architecture as determined by neuronal morphology (red). Microglia are in green.

6.6.2 Supplemental Section B. ECoG modulation during arm movements

In order to assess the functional properties of ECoG signals recorded during the center-out reaching task, ECoG signals were fit to a standard cosine tuning model (Georgopoulos et al., 1982). Normalized spectral data was first averaged over the [70 - 110] Hz frequency band over the [-100ms, 100ms] interval relative to movement onset for each trial. Cosine tuning curves were fit to this averaged data using equation B.1:

$$f = b0 + bxx + byy + \varepsilon \qquad \text{(B.1)}$$

Supplemental Figure 6.B1 shows normalized spectral data for each target and electrode during center-out hand control trials. Despite differences in electrode position on the cortex, nearly all electrodes show preferential modulation for upper-left targets. This is confirmed by the results of the cosine tuning analysis shown in Supplemental Figure 6.B2. The preferred directions of high-gamma band activity from all electrodes were tightly clustered during the center-out task, as would be expected from the highly correlated time-frequency data shown by Figure 6 and Supplemental Figure 6.B1. It is unclear what role the fibrotic encapsulation observed upon explantation of the electrode grid may have played in the observed inter-electrode correlation. It is possible that the increased distance between the electrodes and the cortex resulting from the growth of the encapsulation tissue may have resulted in an increase in inter-electrode correlation. Additional studies are required in order to better understand how the functional specificity of ECoG signals recorded from motor cortex varies as electrodes are moved further from the cortical surface.

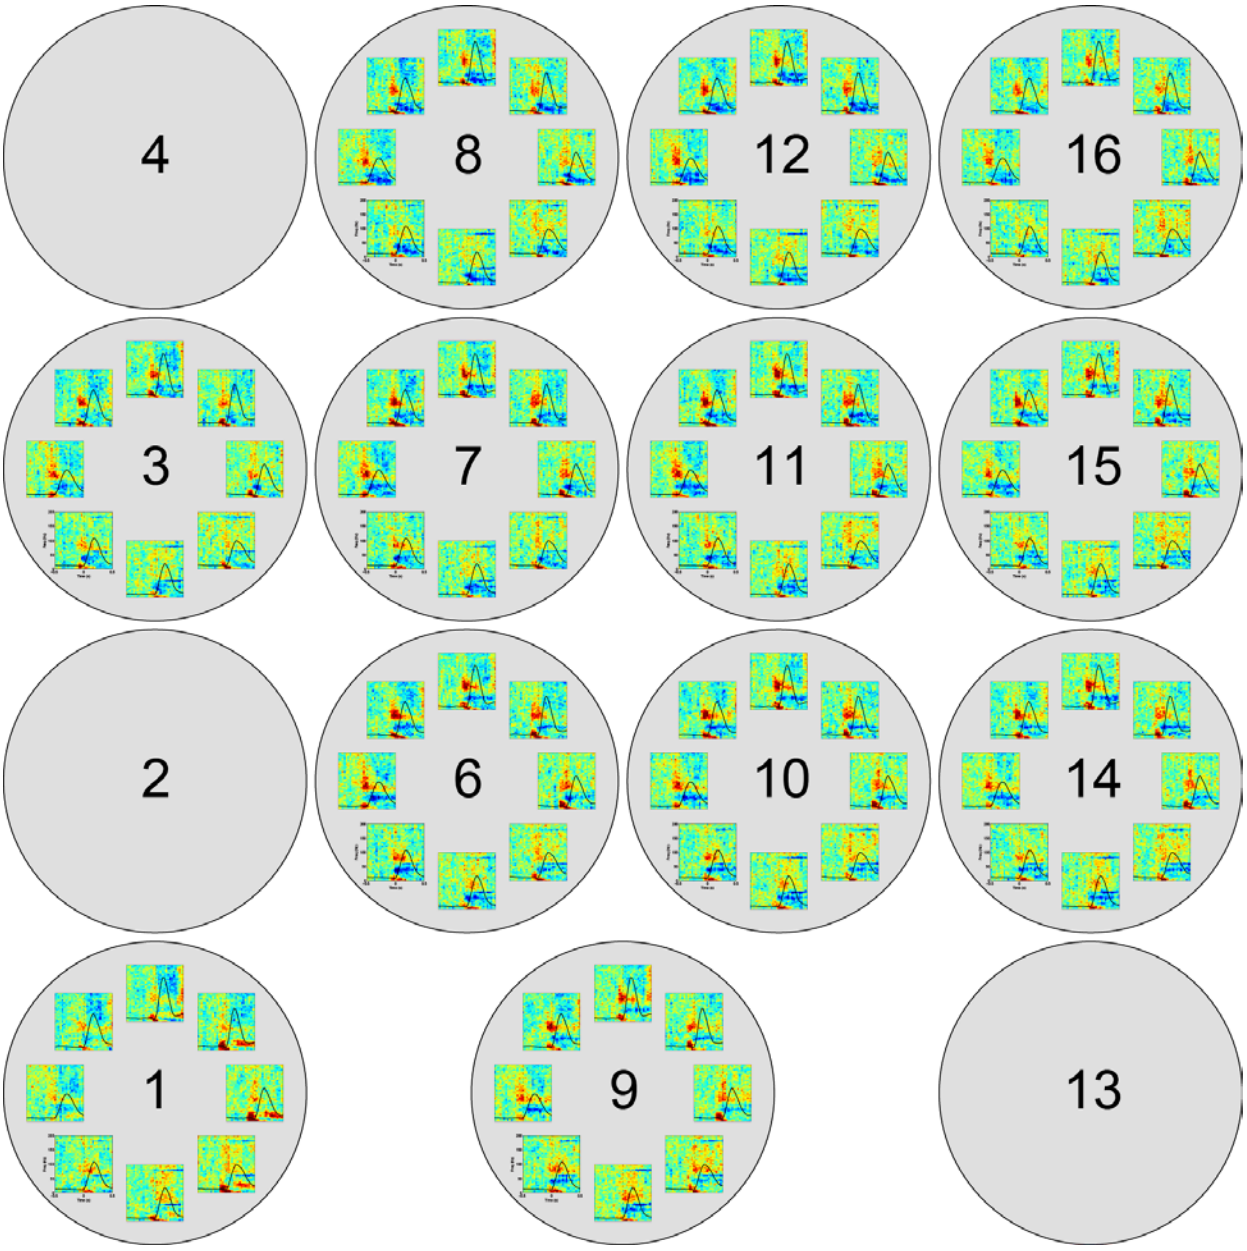

Figure 6.8: Directional modulation of ECoG signals during reaching. Data are shown for all successful reaching trials (N = 1,145). Movement onset is indicated by the dashed black line for each plot. Data for electrodes 4 and 13 are not available on account of these electrodes serving as ground and reference electrodes, respectively. Data for electrode 2 is not available on account of connectivity of this electrode being lost shortly after grid implantation.

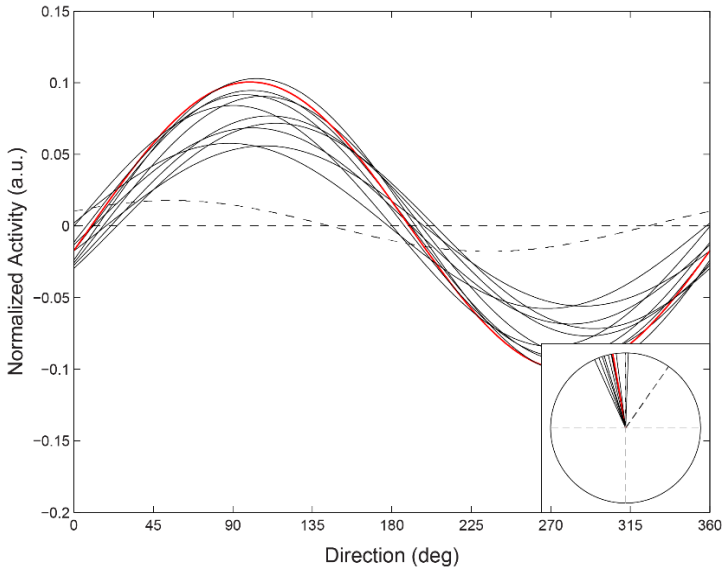

Figure 6.9: Cosine tuning of ECoG signals during reaching. The tuning curve for electrode e10 is shown in red; all other tuning curves are shown in black. Solid lines indicate electrodes exhibiting significant tuning to target direction ($p < 0.05$), while dashed lines indicate non-significant fits. Inset: Normalized preferred directions for each channel plotted on a unit circle. The mean of each curve (b0) has been removed prior to plotting.

6.6.3 Supplemental Section C. Longitudinal impedance and RMS results

6.6.3.1 Electrode impedance

Electrode impedances were measured both over the course of electrode implantation (*in vivo*) and after explantation (*ex vivo*). Impedance was measured at 20Hz using the g.USBamp biosignal amplifier (g.tec Medical Engineering). In vivo impedance measurements were made for each ECoG grid electrode using the recording ground and reference electrodes (e4 and e13, respectively) shorted together with one of the animal's headposts as the counter electrode. Ex vivo impedance was measured for each ECoG grid electrode in saline with two counter electrode configurations: one using the recording ground and reference electrodes (e4 and e13,

respectively) shorted together, and one using an external calomel electrode (reference) and platinum foil electrode (ground) shorted together. For both *ex vivo* configurations, impedance was measured with the grid in the encapsulating tissue envelope and after removal from the encapsulation, with the goal of assessing the impedance contribution of the encapsulating tissue. Analysis of impedance measurements were restricted to the mean of 2mm-diameter electrodes which retained electrical connectivity throughout the course of the experiment.

In vivo mean electrode impedance (Supplemental Figure 6.C1A) was found to be relatively stable over the course of the study. However, electrode impedance was observed to decrease following surgical intervention on day 271 post-implantation to repair an acrylic island used to cover the electrode grid wire bundle. The cause of this decrease in impedance is unclear, though it may be the case that this surgical intervention resulted in unintentional movement of the electrode grid relative to the cortex. *Ex vivo* impedance measurements (Supplemental Figure C1B) indicate that the presence of the encapsulation tissue likely increased electrode impedance, as extraction of the electrode grid from the encapsulating tissue drastically reduced measured impedance values. However, *ex vivo* measurements were performed on fixed encapsulation tissue; the actual contribution of the encapsulation to the *in vivo* impedance would most likely be less than the observed values.

6.6.3.2 RMS amplitude analysis

In order to assess changes in electrophysiological recordings over the course of the study, mean electrode root-mean-squared (RMS) amplitudes were calculated for selected datasets. For each dataset analyzed, 10-minute long segments at the beginning of the recording session were inspected for the presence of artifacts; these artifacts were excluded from subsequent analyses. ECoG signals were then re-referenced with respect to the common-mode average, notch-filtered

to remove power-line contamination, segmented into 1-second epochs, and band-pass filtered in 5Hz frequency bands over the 5Hz-195Hz range using 4th-order Butterworth filters. RMS values were computed for the resultant band-pass-filtered segments for each frequency band.

Mean electrode RMS amplitudes of recorded ECoG signals were found to decrease over the course of the study (Supplemental Figure 6.C1C,D), appearing to stabilize approximately 300 days post-implantation. The gradual decrease in signal amplitude observed is consistent with the development of encapsulation tissue over time. This development would increase the distance between the electrodes and cortex, which based on volume conductor models would be expected to reduce measured electrical potentials (Rattay 1988).

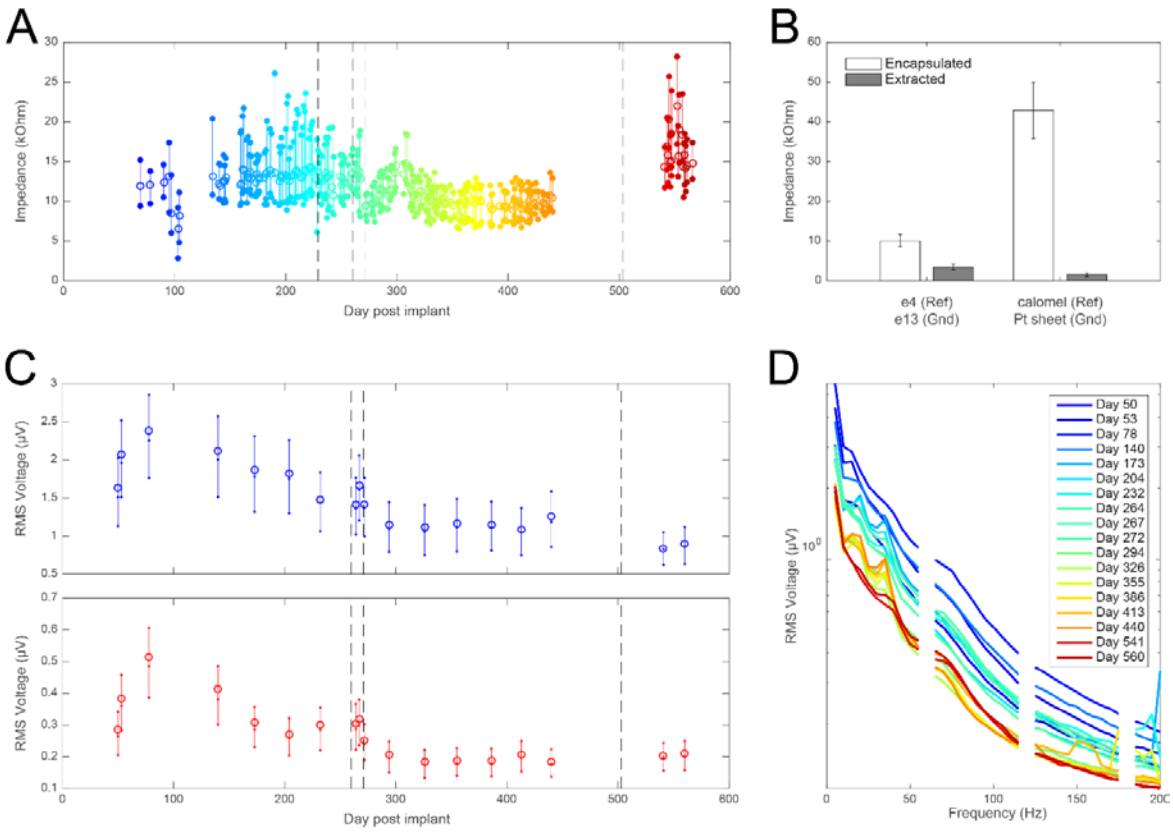

Figure 6.10: Electrode impedance and RMS amplitude of electrophysiological recordings. (A). In vivo electrode impedance values (measured at 20Hz) over the course of electrode grid implantation. For each day impedance was measured, the range of impedance values is indicated by colored vertical lines, while the mean impedance is indicated by open circles. Dashed black lines indicate days on which surgical interventions were performed. (B). Ex vivo electrode impedance values. Impedance data are shown for both 'e4 + e13' and 'calomel + Pt' measurement configurations (see supplemental methods) for the electrode grid during encapsulated (white) and extracted (gray) conditions. Error bars indicate standard deviations from the mean impedance values for each condition. (C). RMS amplitude over the course of electrode implantation. RMS amplitude at 20Hz (blue) and 100Hz (red) is shown for selected recording sessions. Solid vertical lines indicate the [25th – 75th] quantile range of RMS amplitudes, with filled circles indicating the 25th, 50th, and 75th quantile values. Open circles indicate the mean RMS amplitude. (D). RMS amplitude spectra for selected recording sessions. Mean RMS spectra are provided for the same recording days shown in (C). Individual spectra colors correspond to those of the in-vivo impedance data in (A).

7.0 EXTRACELLULAR VOLTAGE THRESHOLD SETTINGS CAN BE TUNED FOR OPTIMAL ENCODING OF MOVEMENT AND STIMULUS PARAMETERS

Figures and text in this chapter have been reproduced with permission from Oby et al. 2016 (Oby ER, Perel S, Sadtler PT, Ruff DA, Mischel JL, Montez DF, Cohen MR, Batista AP, Chase SM (2016) Extracellular Voltage Threshold Settings Can Be Tuned for Optimal Encoding of Movement and Stimulus Parameters. J of Neural Eng 13(3):036009). In this chapter, I describe a procedure to determine the best threshold for extracting information from extracellular recordings. Dr. Emily Oby was the main contributor to this work. My contribution was data collected in August 2012, in the first few months of my graduate schooling. While recording the data used in this manuscript, I learned how to set up brain computer interface experiments and became comfortable training the animals in the Batista lab.

7.1 ABSTRACT

A traditional goal of neural recording with extracellular electrodes is to isolate action potential waveforms of an individual neuron. Recently, in brain-computer interfaces (BCIs), it has been recognized that threshold crossing events of the voltage waveform also convey rich information. To date, the threshold for detecting threshold crossings has been selected to preserve single-neuron isolation. However, the optimal threshold for single-neuron identification is not

necessarily the optimal threshold for information extraction. Here we introduce a procedure to determine the best threshold for extracting information from extracellular recordings. We apply this procedure in two distinct contexts: the encoding of kinematic parameters from neural activity in primary motor cortical (M1), and visual stimulus parameters from neural activity in primary visual cortical (V1).

7.2 INTRODUCTION

Brain-Computer Interfaces (BCIs) extract information about motor intentions from recordings of neural signals to control an external device, with the goal of assisting patients with paralysis or other sensory-motor deficits. The recorded signals have taken the form of EEG, MEG, and intracortical signals (Schwartz et al., 2006). A promising class of BCIs extracts information directly from action potentials, or "spikes," identified from the voltage traces recorded from chronically implanted extracellular electrodes. To identify these spikes, the voltage trace is typically band-pass filtered, thresholded to identify transients in the voltage signal, and then sorted based on the shape of the transient waveform into clusters corresponding to individual neurons. This final pre-processing step, "spike sorting", has received considerable attention because it is time consuming, prone to inaccuracies, and difficult to perform in clinical settings (Lewicki, 1998; Rey, 2015). Fortunately, it appears that accurate spike sorting may not be necessary for good BCI performance (Ventura, 2008; Fraser et al., 2009; Chestek et al., 2011; Malik et al., 2014). Rather, a threshold can be set, and all voltage transients that exceed that threshold (that is, "threshold crossings") can be counted, regardless of the waveform shape. Evidence is accumulating that there is information in such non-spike signals recorded from

microelectrodes. In one example, Stark and Abeles (2007) used a multiunit activity signal, processed by computing the root mean square of the voltage signal in the 300 - 6000 Hz frequency band, to predict reach direction and grasp with better accuracy than either spike activity or local field potentials. With this knowledge, some researchers have investigated the possibility of moving away from using sorted units as inputs to BCI decoders and instead using threshold crossings (Fraser et al., 2009). Many studies agree that BCI performance is substantially degraded when the non-spike parts of the signal are discarded (Todorova et al., 2014; Kloosterman et al., 2014; Deng et al., 2015), raising the intriguing possibility that the threshold could be adjusted to maximize BCI performance.

Here we assess how the voltage detection threshold setting ("threshold") affects the encoding of movement parameters in primary motor cortex (M1). We then assess the generality of this approach by using it to examine the information present in recordings from primary visual cortex (V1). To interpret our observations, we reason that the choice of threshold impacts the effective sampling radius of the electrode. For example, choosing a more permissive threshold presumably enlarges the effective sampling radius of the electrode and, thus increases the number of neurons contributing to the threshold crossing signal (Martinez et al., 2009, Pedreira et al., 2012). At high detection thresholds, threshold crossings comprise the spikes from individual neurons close to the electrode. At low detection thresholds, threshold crossings comprise multi-unit activity from smaller neurons or neurons farther from the electrode. How the detection threshold is chosen impacts the neural contributions to the signal, and potentially, what information is contained in the signal.

Traditionally, thresholds have been chosen to maximize spike-sorting performance. However, the optimal threshold for single-neuron identification is not necessarily the optimal

threshold for information extraction. We hypothesized that the optimal threshold would depend on the parameter of interest. We assessed the impact of the detection threshold by systematically sweeping the detection threshold and evaluating the information content of threshold crossings about two different parameters of interest, velocity and speed, recorded from primary motor cortex (M1). We find that the type of information encoded by threshold crossings depends strongly on threshold, and the optimal threshold depends on the parameter of interest. In particular, we find that velocity, a directional parameter, is better represented at higher thresholds, whereas speed, a scalar quantity, is better represented at lower thresholds. Additionally, we show that optimal thresholds are surprisingly low, considerably below the thresholds commonly used in closed-loop BCI studies. This means that the optimal thresholds for extracting information are not typically the best thresholds for isolating single neurons.

We can understand these results in the context of the topographical representation of speed and velocity in M1. The scale of the topographic organization and the homogeneity of a parameter's representation across cortical tissue influence the optimal threshold. This observation could generalize to other areas of cortex, such that knowledge of the topographic representation of different parameters should predict the choice of threshold for maximizing the information available in neural recordings. We tested this hypothesis with recordings from V1, a cortical area with a distinctly different topographical representation of its relevant parameters. By applying our method of sweeping the threshold, we were able to predict the relative optimal thresholds for the parameters orientation and contrast of a visual stimulus. We conclude that the type and quality of information that can be extracted from extracellular signals depends on the threshold setting; there is more information present in extracellular voltage recordings than is typically extracted.

7.3 METHODS

All animal procedures complied with the National Institutes of Health Guide for Care and Use of Laboratory Animals, and were approved by the University of Pittsburgh's Animal Care and Use Committee. To assess the generality of our predictions, we analyzed data collected from two different cortical areas in two monkeys each, and in the context of two different behaviors.

7.3.1 M1 tasks and recordings

Two male monkeys (Macaca mulatta, 11.6 and 7.3 kg) were trained to perform an 8-target center-out reach task (Figure 7.1(a)). The position of an LED marker attached to the fingertip of the reaching hand was tracked at 120 Hz (<1 mm resolution; Phasespace Inc., San Leandro, CA). The position of the marker was visible to the monkey as a cursor on a frontoparallel screen. The hand was not visible to the monkey, because it moved in the space behind the screen. At the start of each trial, the monkey had to move the cursor to a central target and hold for 200-400 ms. Then, one of eight peripheral targets (arranged at 45° intervals and spaced 9 cm from the center) appeared, and he had ~800 ms to acquire it with the cursor. After holding the cursor on the peripheral target for 200-500 ms (randomized) he received a liquid reward. A failure at any point caused the trial to terminate without reward, and there was a 1.5 second timeout before the next trial began.

When the monkey was proficient at the task, we implanted a 96-electrode array (Blackrock Microsystems) in the arm area of M1 (as determined intraoperatively via cortical landmarks) contralateral to the reaching hand (Figure 7.1(a)). As the monkey performed the task we recorded neural data from M1 using a Tucker-Davis Technologies RZ2 system. During each

recording session, we streamed the filtered broadband signal (700-3000Hz band-pass, Kaiser window) from 4-10 different channels directly to disk at a 24 kHz sampling rate. In some cases we streamed an unfiltered broadband signal at a 24 kHz sampling rate and applied a 700-3000Hz bandpass filter offline. Because of system limitations, we could not record broadband signals from all 96 channels each day. In total, we recorded 20 unique channels over 5 experimental sessions from monkey J (26 months post-implant) and 53 unique channels over 9 experimental sessions from monkey L (2 weeks to 9 months post-implant).

In this data set, we analyzed the representation of two kinematic parameters -velocity and speed - which are known to correlate well with neural firing in M1 (Moran and Schwartz, 1999; Churchland and Shenoy, 2007; Golub et al., 2014).

7.3.2 V1 stimuli and recordings

Two different male monkeys (Macaca mulatta, 9.25 and 8.0 kg) were trained to fixate on a central spot while visual stimuli were presented peripherally (Figure 7.1(b)). The animals had been trained to perform an orientation change detection task over the course of several months and were able to stably maintain fixation for 3 - 5 seconds. Before electrophysiological recording, the animals were implanted with a custom titanium head post, and a 96-electrode array (Blackrock Microsystems) in V1 (as determined by cortical landmarks, Figure 7.1(b)). Eye position was monitored using an infrared optical recording system (Eyelink, SR Research) sampling at 1 kHz.

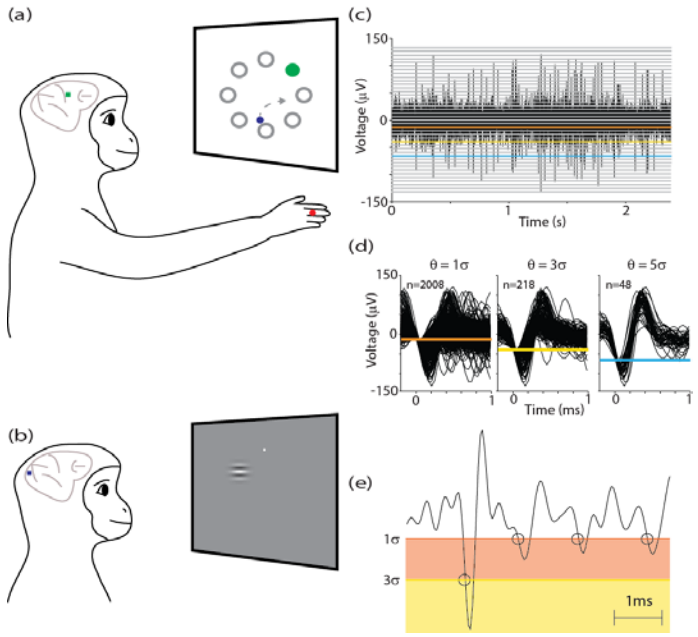

Figure 7.1: Schematic of tasks and neural recordings. (a) During the M1 recordings, a monkey performed an 8-target center-out reaching task. An LED marker (red) was attached to the monkey's finger tip to track his movements, which were displayed as a cursor on the screen (blue). The monkey made reaches from the center of the screen to one (green) of eight peripheral targets (gray). The array placement in M1 is shown by the green square. (b) During the V1 recordings, a monkey fixated on a central spot (white) while drifting Gabor patches were presented peripherally. The array placement in V1 is shown by the blue square. (c) Voltage trace from M1 during a single reach trial with detection threshold settings from $\theta = 10\sigma$ to -10σ. $\theta = 1\sigma$ (orange) is permissive, capturing low voltage transients. $\theta = 5\sigma$ (light blue) is more restrictive, capturing only high voltage transients which likely correspond to spikes from a single neuron. (d) Waveform snippets for threshold crossings of 1σ, 3σ, and 5σ in D. As the threshold becomes more permissive (1σ, orange) there are more threshold crossings. As the threshold becomes more selective (5σ, light blue) the waveform becomes more consistent. (e) Using the exclusive window categorization method, threshold crossings for the channel are identified when the voltage trace passes into and out of the window defined by a particular threshold without passing into higher-threshold windows. A 1σ window (orange) and a 3σ window (yellow) are shown in this example. If the voltage trace crosses the 1σ threshold but not the 3σ threshold, it is classified as a 1σ crossing. As indicated with the black circles, we can successfully select the larger voltage fluctuations with the exclusive $\theta = 3\sigma$ and we capture the smaller fluctuations with exclusive $\theta = 1\sigma$.

To begin each trial, the monkey would acquire fixation on a central spot. After 200-400ms of stable fixation within a 1 degree window, stimulus presentation began. A total of seven stimuli were flashed for 200 ms each with an interstimulus interval of 100ms. If the animal maintained fixation for the duration of the stimulus presentations, he was rewarded with a drop of juice. If the animal's eye position left the fixation window during stimulus presentation, the trial was aborted and no reward was given. Stimuli were presented on a mean gray luminance screen (1024 x 768; 27.9 pixels/degree; 120Hz refresh rate) placed 635mm in front of the animal. The stimuli were drifting oriented Gabor patches that varied in contrast (contrast values = 0.06, 0.12, 0.25, 0.5, 1) or orientation (orientation values ranged from 0 o to 330 o in 30o intervals). When orientation was varied, contrast = 1. When contrast was varied, orientation = 90o for monkey B and 180o for monkey G. The receptive fields of the V1 neurons recorded on the array were located approximately 3.5o eccentric from fixation, in the lower right visual field, and they spanned approximately 2 degrees of visual angle. Spatial dimensions of the stimuli were selected to envelop the receptive fields of all V1 neurons recorded by the array. In the initial frame of each stimulus, the grating had odd spatial symmetry. The phase velocity of the stimulus was selected so that upon presentation of the final frame, the stimulus had drifted one complete cycle.

Electrophysiological recordings were performed 2.5 months (monkey B) and 2 weeks (monkey G) post-implant. Data were collected with a Grapevine system (Ripple, Inc.). Broadband signals were recorded on all 96 channels on one day. Each channel was sampled at 30kHz and raw signals were bandpass filtered (Highpass filter: 0.3Hz; Lowpass filter: 7.5kHz, 3rd Order Butterworth) and streamed to disc. The saved signals were subsequently filtered offline in the same way that the M1 signals were, using a Kaiser window with a 700-3000Hz passband.

In this task, we analyzed the neural representation of the orientation and contrast of the drifting Gabor patch stimulus. Both of these parameters are known to drive neural firing in V1 (Hubel and Wiesel, 1959).

7.3.3 Threshold crossings

Our central analysis assesses the information content present in neural recordings at varying voltage thresholds. To do this, we systematically swept the level of the voltage detection threshold to extract threshold crossings (Figure 7.1(c)). At each threshold we evaluated the signal-to-noise ratio of the information about movement or stimulus parameters encoded by the corresponding threshold crossings. We defined threshold settings with respect to the standard deviation of the filtered signal (σ), computed as the average standard deviation of the recording over 100-200 trials. We considered threshold settings ranging from 0 (mean) to -10σ for the M1 data and -6σ for the V1 data, at intervals of 0.5σ. These negative threshold settings correspond with the depolarizing phase of the action potential. Results from positive-going thresholds were comparable, and thus we use only the negative thresholds in our analyses. We defined a threshold crossing as the time at which the recorded signal crossed the threshold voltage in a negative-going direction, with 100 μs resolution. For clarity, figures and the following text will refer to the absolute value of the multiplier of the threshold setting (e.g. 3σ).

7.3.4 Quantifying information content with Signal-to-Noise Ratio

We use the signal-to-noise ratio (SNR) to quantify the information content conveyed by the threshold crossings. Intuitively, SNR can be thought of as the ratio of useful information to

irrelevant information. Formally, “signal” is defined as the variance in the data that is explained by a parameter of interest (e.g. velocity, speed, orientation or contrast), and “noise” is defined as the residual, unexplained variance after accounting for that parameter. Here, our recorded data, Y_θ, is the number of threshold crossings recorded at a particular threshold θ. We can decompose the variance in our data, $\mathrm{Var}[Y_\theta]$, into a component explained by a stimulus X and a component remaining after accounting for X. This decomposition is exact, and is given by the Law of Total Variance:

$$\mathrm{Var}[Y_\theta] = \mathrm{Var}(\mathrm{E}[Y_\theta|X]) + \mathrm{E}(\mathrm{Var}[Y_\theta|X]), \tag{8.1}$$

where $\mathrm{Var}(\mathrm{E}[Y_\theta|X])$ is the variance of the expected value of Y_θ conditioned on X, and $\mathrm{E}(\mathrm{Var}[Y_\theta|X])$ is the expected value of the variance of Y_θ conditioned on X. The first term quantifies the variation in Y_θ that is explained by X (i.e., the signal variance); the second quantifies the residual variation in Y_θ that remains after accounting for X (i.e., the noise variance). The SNR is the ratio of these two quantities:

$$SNR_\theta \equiv \frac{Var(E[Y_\theta|X])}{E[Var(Y_\theta|X)]}. \tag{8.2}$$

For the M1 studies, the parameters of interest are velocity and speed, which both vary in a continuous fashion over the range of natural reaching movements. To compute the SNR in this case, we first fit linear tuning curves by regressing neural activity against kinematics, and then quantified how well these linear fits accounted for the variance of the threshold crossings with the SNR. We considered a separate encoding model for velocity (Eq. 8.3) and speed (Eq. 8.4), and fit an ordinary linear regression at each threshold setting:

$$y_\theta(t) = b_0 + b_x v_x(t) + b_y v_y(t) + \varepsilon_v(t) \tag{8.3}$$

$$y_\theta(t) = b_0 + b_s s(t) + \varepsilon_s(t) \tag{8.4}$$

where $y_\theta(t)$ is the number of threshold crossings for a given threshold in a 100 ms bin centered at time t, $v_x(t)$ and $v_y(t)$ are the x- and y-components of the velocity of the cursor averaged over a 100 ms bin, s(t) is the speed of the cursor averaged over a 100 ms bin, and $\varepsilon(t)$ is an error term that captures deviations from the model. These models can be fit at varying temporal offsets between the neural and kinematic data. We used a 100 ms offset (neural activity leading kinematics) because we have found this offset yields the best correlation with behavior for the data sets analyzed here (Perel et al., 2015). For each encoding model, we used the model estimates to compute the signal variance and the residuals of the regression to compute the noise variance. As an example, for speed the signal variance is the variance of $b_0+b_s s(t)$ over all recorded speeds, and the noise variance is the variance of $\varepsilon_s(t)$. Graphical depictions of these quantities are provided in Figure 7.2(c).

For the V1 studies, the parameters of interest are orientation and contrast. Each of these varied over discrete levels in our experiments, and firing rates were measured for multiple repetitions of each particular orientation or contrast. In this case $E[Y_\theta|X]$ and $Var[Y_\theta|X]$ can be measured directly from the data (as depicted in Figure 7.8(a)), without the need for linear regression.

Although SNR is not a common metric in either M1 or V1 studies, it provides a simple, intuitive metric of information content, it is relatively straightforward to compute, and it allows for relatively direct comparisons across brain areas even when the parameters of interest are quite different. A more common metric of goodness of fit in motor neurophysiology is the coefficient of determination (R^2) (e.g., Georgopoulos et al., 1982), which is a statistical measure of how well a model approximates the data. Qualitatively our M1 results are the same with either measure. However, since neurons in V1 do not respond in a linear fashion to stimuli of different

orientations, the R^2 would have been less appropriate for those data. Another possibility would have been to compute the mutual information between threshold crossings and parameters directly. However, comparisons of mutual information across different stimulus sets are difficult to interpret when those sets are not entropy-matched (Golub et al., 2014; Chase and Young, 2008). For these reasons, we favor the SNR metric for this study. Finally, we note that SNR values less than 1 are not uncommon in neural responses, especially when analyzed at fine temporal resolution, and low SNR values still signal the presence of meaningful information.

7.3.5 Exclusive windows approach to spike sorting

We reasoned that the small-amplitude fluctuations of the voltage trace might contain information that was distinct from the information contained in the high-amplitude fluctuations. To this end, we performed an "exclusive window" analysis. In this analysis, a threshold crossing was registered only if it crossed a defined threshold in the negative direction and re-crossed it in the positive direction before crossing another more-negative threshold (Figure 7.1(e)). With this definition, a given excursion of the voltage trace is exclusively categorized as crossing only one threshold. This is in contrast to our basic threshold analysis in which a threshold crossing that crossed a given threshold was counted at all smaller thresholds as well. To differentiate these choices in the text, we refer to exclusive threshold crossings as xTCs. The exclusive windows can act as a crude approach to spike sorting (Todorova et al., 2014), when large thresholds are selected. Here we examine two exclusive thresholds: a low threshold at 1σ to select the small voltage fluctuations, and a high threshold, which captures the large voltage fluctuations associated with spikes. We considered two possible high thresholds, 3σ or 4.5σ. Using these xTCs, we repeated the SNR analysis as described above.

7.4 RESULTS

Our central finding is that in extracellular recordings the detection threshold can be tuned to maximize information about parameters of interest, with different parameters exhibiting different optimal thresholds. Further, the threshold setting that maximizes information is usually not the setting that yields the best spike sorting. We show this in two cortical areas, with two parameters of interest for each area. Our main focus is on primary motor cortex (M1), where we consider the selection of optimal thresholds for the neural encoding of velocity or speed. To examine the generality of this approach, we also apply it to neural recordings from primary visual cortex (V1), where the parameters of interest are orientation and contrast. In both cases, recordings are collected with 96-electrode arrays. Broadband data are saved, and analyses are conducted offline. For each recorded channel, we swept the voltage detection threshold, and measured the number of threshold crossings at each threshold. At each threshold, we quantified the amount of information about the parameter of interest as a signal-to-noise ratio (SNR).

7.4.1 Information content depends on threshold

Figure 7.2(a) shows the density of threshold crossings at a low threshold ($\theta=1\sigma$) and a high threshold ($\theta=5\sigma$) during reaches to eight different target directions for an example M1 channel. It can be seen that different information is manifested in the tuning at different threshold settings. At $\theta = 5\sigma$ (cool color scale), this channel shows velocity tuning, with a preferred direction up and to the right. At $\theta = 1\sigma$ (warm color scale), this channel modulates similarly for all 8 reach directions. The velocity tuning is weak at this threshold, but, instead, the signal reflects speed (gray lines): it is active during the reach regardless of direction. This is an

exemplary channel which visually highlights our central finding: by adjusting the threshold setting, we differentially extract information about each parameter (not just different amounts of information about a given parameter) from the neural signal. Separately for each threshold setting, we modeled the relationship between threshold crossings and each kinematic parameter with linear regression. We quantified the information content with the SNR. As an example, figure 2(b) shows the linear regressions for speed at $\theta=1\sigma$ and $\theta=5\sigma$. As defined in Methods, the signal is the variance of the estimated threshold crossings and the noise is the variance of the residuals, as shown by the histograms in Figure 7.2(c). We calculated the signal and noise for each threshold setting from 0 to 10σ in 0.5σ increments (Figure 7.2(d) and (e)). Importantly, the SNR depends on the threshold setting (Figure 7.2(f)). Specifically, this channel has more speed information at low thresholds and more velocity information at high thresholds. Perhaps a more familiar metric of goodness of fit in motor neurophysiology is the coefficient of determination (R^2) (e.g., Georgopoulos et al., 1982). Figure 7.2(g) plots the dependence of R^2 on threshold. Qualitatively, we see the same dependence of information content on threshold regardless of which measure of goodness of fit we choose. This reassures us that quantification of information with SNR is an appropriate measure for neural recordings from M1, and it has the advantage that it can be applied more broadly to neural recordings from other brain areas.

Figure 7.3 shows the SNR dependence on threshold for three representative M1 channels. The curves for both speed and velocity show an inverted-U shape with respect to threshold. The lowest SNR values occur at $\theta = 0$, when there are so many threshold crossings that the signal does not provide clear information about the reach kinematics. Similarly, we see low values of SNR at high values of θ, when there are not enough threshold crossings to provide a clear

relationship between neural events and the velocity or speed of the reaches. The peak SNR is between these extremes.

From the SNR dependence on threshold we can extract the optimal threshold for velocity and speed information. We computed SNRs for the 73 M1 channels we recorded. We only included a channel in subsequent analyses if it exhibited a statistically significant regression (α=0.05) for at least one threshold setting for at least one of the kinematic parameters. This resulted in 0 discarded channels from monkey J and 6 from monkey L, leaving a total of 20 channels from monkey J and 47 channels from monkey L. The normalized average SNR relationship with threshold for those 67 channels is plotted in Figure 7.4(a). Each channel was normalized to its maximum SNR and then averaged. Normalization emphasizes the relative thresholds at which the peaks occur, regardless of differences in the absolute SNR values across channels. The deviation of the normalized peak from a value of one reflects the variability in the peak threshold across the population. The peak SNR varies for speed and velocity: speed is optimally encoded at a low threshold setting ($\theta = 2\sigma$), while velocity is optimally encoded at higher thresholds ($\theta = 2.5\sigma$). The optimal threshold depends on the information one wishes to extract, and is often lower than the threshold that is typically applied to isolate the activity of a single neuron.

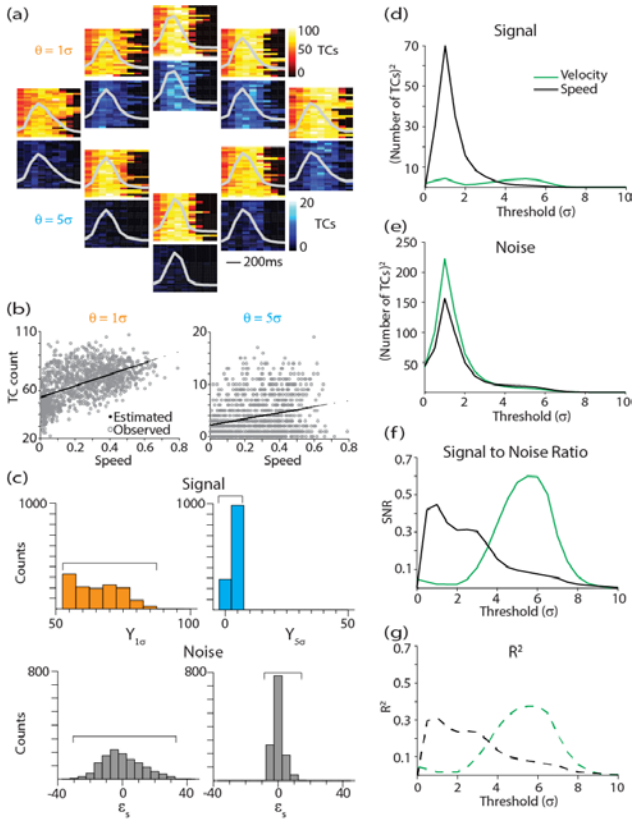

Figure 7.2: SNR quantifies information content in M1. (a) Single channel example of threshold crossing activity as a function of reach. Each plot shows the number of threshold crossings in 100ms bins for one of eight reach directions. The color indicates the number of threshold crossings, where the red scale is for $\theta = 1\sigma$ and the blue scale is for $\theta = 5\sigma$. Each row is a trial. The top plots are for a permissive threshold ($\theta = 1\sigma$) and the bottom plots are for a selective threshold ($\theta = 5\sigma$). Average speed profiles for each reach direction are plotted in gray for reference. Note the strong directional tuning for $\theta=5\sigma$ (with an upwards preferred direction), and the strong speed modulation for $\theta=1\sigma$. Panels (b)-(f) step through the SNR calculation which we use to quantify this. (b) The observed number of threshold crossings (gray) is plotted against the corresponding reach speed for a permissive (left) and selective (right) threshold. In black, we show the linear regression. (c) We take the variance of Y_θ (as described in Equation 8.4, black in (b)) to be the signal and the variance of the residuals (ε_s as described in Equation 8.4) to be the noise. The histograms show the distributions of these measurements from which the variance is calculated. The signal (d) and noise (e) arising from these variance calculations vary with threshold. (f) Combining signal and noise, velocity and speed SNRs show an inverted-U shaped relationship with threshold with peaks at different thresholds. (g) A common metric of tuning in M1 is R^2, plotted here for comparison.

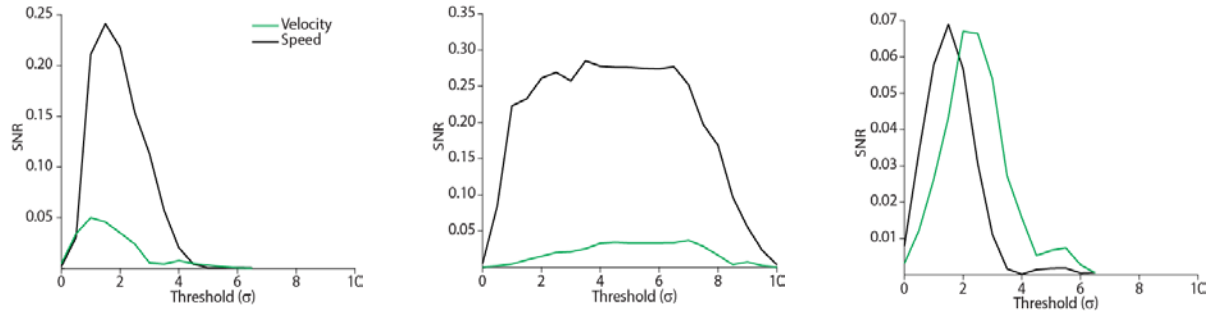

Figure 7.3: SNR in M1 depends on threshold. SNR dependence on threshold for three representative M1 channels. At each threshold, SNR is computed separately for velocity tuning (green) and for speed tuning (black).

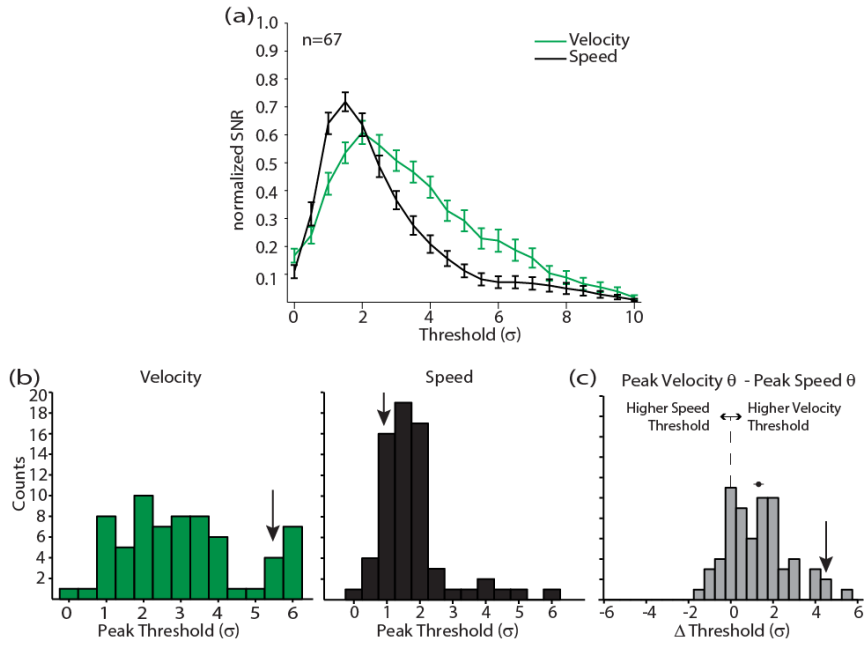

Figure 7.4: Optimal thresholds for a given parameter differ across channels. (a) Normalized mean ± SE of SNR dependence on threshold for all 67 M1 channels with significant tuning. Velocity (green), speed (black). (b) Optimal thresholds for velocity (green) and speed (black). The arrows point to the channel shown in Figure 7.2. (c) The per-channel difference between the optimal thresholds for velocity and speed. The mean ± SE is indicated by the dot with the line through it above the histogram. The mean is significantly different from zero (t-test, $p<10^{-7}$).

7.4.2 Optimal SNR thresholds are lower than typically used for recording

In multi-electrode systems where it is possible to adjust the threshold independently for each channel, even more information can be extracted. Figure 7.4(b) shows histograms for the optimal thresholds for velocity (left) and speed (middle). The distribution of optimal thresholds for speed is narrow with relatively low thresholds. The distribution of optimal thresholds for velocity is broader than is the distribution of optimal thresholds for speed and it includes channels with higher optimal thresholds. The distributions have statistically different means (t-test, p=10-7). The mean pairwise difference between the optimal velocity threshold and the optimal speed threshold is $1.28\sigma \pm 0.18\sigma$, and the distribution is shown in the histogram in Figure 7.4(c). The optimal velocity threshold is higher than the optimal speed threshold for 48 of 67 (72%) channels.

7.4.3 Using exclusive thresholds to highlight information content of low amplitude fluctuations

It is conceivable that large-voltage “spikes” are the sole source of information in an extracellularly recorded signal, and lower thresholds are just capturing these spikes with greater reliability. Alternatively, the lower amplitude fluctuations which are not readily attributable to the spiking of nearby neurons may contain information that is distinct from that carried by the high-amplitude events. We addressed this through an exclusive threshold analysis. We ask whether single-unit activity and the residual multi-unit hash contribute differently to the speed and velocity encoding models. Figure 7.5 shows how setting two exclusive thresholds can act as simple spike identifier, using the channel depicted in Figure 7.2 as an example. The black circles

identify exclusive threshold crossings (xTCs) for thresholds of 1σ and 3σ. Setting the threshold high has a similar effect as spike sorting, in that it captures single unit activity, whereas the low threshold captures non-single unit activity that might typically be discarded under a sorting paradigm, as evidenced by the waveform snippets shown in Figure 7.5(b). The SNR for the xTCs from this example channel at $\theta = 1\sigma$ shows that there is speed information contained in the non-single unit activity (Figure 7.5(c)). The single unit activity captured by $\theta = 3\sigma$ shows better velocity encoding. This supports the idea that low-voltage events contain information that is distinct from the information present in spiking activity.

The mean exclusive threshold SNR for all channels is shown in Figure 7.6. On average (Figure 7.6(a), left), the non-spike parts of the signal represented by the xTCs at $\theta = 1\sigma$ encode speed better than the xTCs at $\theta = 3\sigma$. Velocity is better encoded at more restrictive (higher) thresholds. To highlight the impact that the choice of threshold has on the information content of the threshold crossings, we repeated the exclusive threshold analysis for M1 at $\theta = 1\sigma$ and $\theta = 4.5\sigma$. Such a high threshold should isolate single units and is thought to obtain better encoding of kinematic information. However, we found that this threshold is quite restrictive and misses some of the available information (Figure 7.6(b)). The exclusive window analysis highlights that there is information contained in the low-amplitude fluctuations of the signal that is often discarded as noise.

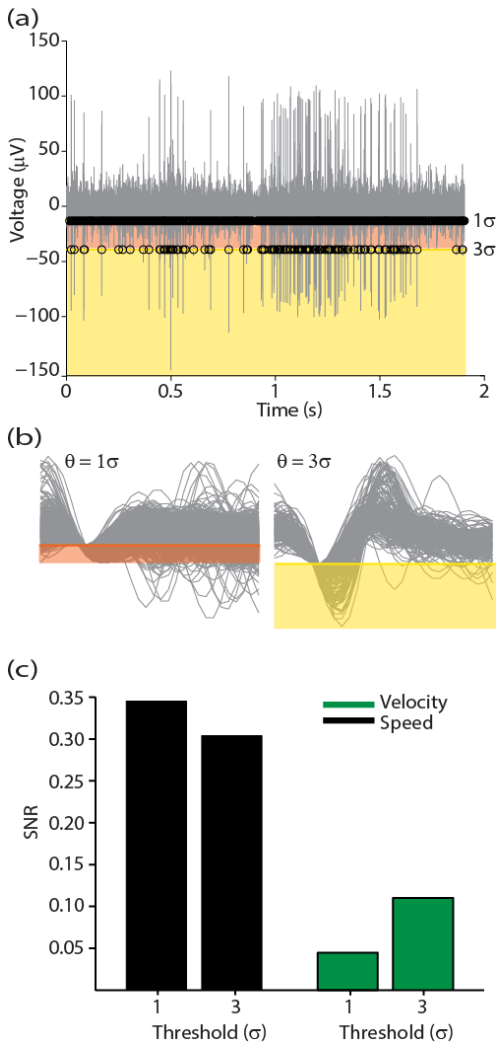

Figure 7.5: An exclusive window analysis reveals substantial information in small voltage fluctuations. (a) Using the exclusive window categorization method, threshold crossings for the channel depicted in Figure 7.2 are identified when the voltage trace passes into and out of the window defined by a particular threshold without passing into higher-threshold windows. A 1σ window and a 3σ window are shown in this example. If the voltage trace crosses the 1σ threshold but not the 3σ threshold, it is classified as a 1σ crossing. As indicated with the black circles, we can successfully select the larger voltage fluctuations with the exclusive θ = 3σ and we capture the smaller fluctuations with the exclusive θ= 1σ. (b) Waveform snippets corresponding to the xTCs for exclusive thresholds θ = 1σ (left) and θ = 3σ (right) for the channel shown in (a). (c) The SNR for velocity and speed at exclusive thresholds θ = 1σ and θ = 3σ for the channel shown in (a). Note that the speed SNR is higher at θ=1 than θ=3 even though those waveforms look like noise.

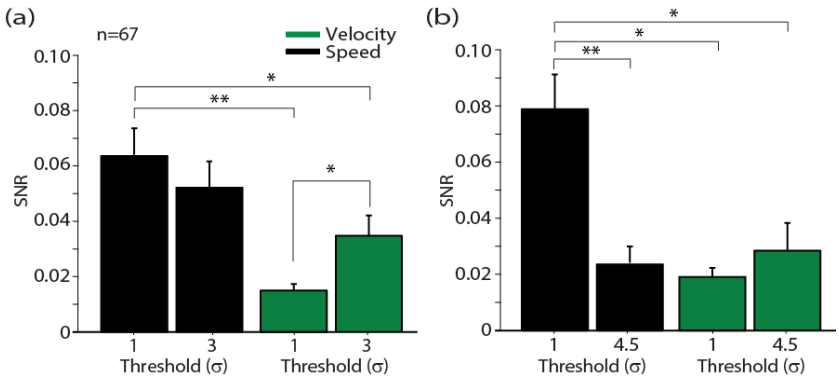

Figure 7.6: Distinct information is encoded by small and large voltage fluctuations. (a) SNR as a function of exclusive threshold for (mean ± SE). Different information is contained in putative spikes classified with $\theta = 3\sigma$ and in the low voltage fluctuations at $\theta = 1\sigma$. (b) The exclusive window SNR for exclusive windows of $\theta = 1\sigma$ and $\theta = 4.5\sigma$ (a threshold commonly chosen in BCI studies.) It is important to note that because of the exclusive nature of the thresholds, adjusting the high threshold also impacts the xTCs at the low threshold. Thus, the SNR at 1σ changes when the high threshold is different. This is not true for the inclusive thresholds used in the other analyses. (Data are from M1, n=67.) Significant differences are indicated with * for $p < 0.01$ and ** for $p < 10^{-4}$.

7.4.4 Information content in V1

To test the generality of our finding that the optimal threshold depends on the parameter of interest, we examined recordings from primary visual cortex (V1). We selected V1 for comparison in part because its function is markedly different from M1, and also because the topography of V1 is well-established. In V1, nearby neurons are tuned similarly to stimulus orientation, with orientation tuning changing in a systematic way across the cortical surface (Hubel, 1982). However, all V1 neurons are tuned similarly to contrast, showing increased firing rates with increasing stimulus contrast (Albrecht and Hamilton, 1982). The topographic organization of V1 led us to predict that the optimal threshold for contrast information would be lower than the optimal threshold for orientation information.

We recorded from two monkeys with multi-electrode arrays implanted in V1 while they viewed drifting gratings, and investigated how information about orientation and contrast depended on threshold. The channels were tuned to different orientations, but all channels showed a similar response to contrast, wherein the maximal response was for contrast = 1. Figure 7.7 plots tuning curves for contrast and orientation at three thresholds for an example channel. Each point in the tuning curve is the number of threshold crossings occurring during a single presentation of a stimulus with a particular orientation or contrast.

To test our prediction that orientation and contrast show different optimal thresholds, we calculated the SNR at each threshold to quantify the information content of the threshold crossings. We break down this calculation into its components in Figure 7.8. For orientation, signal is the variance of the mean number of threshold crossings over each orientation (Figure 7.8(a), orange). Noise is the mean of the variance in threshold crossings at each orientation (Figure 7.8(a), black). As shown for this example channel in Figure 7.8(c), orientation and contrast SNR depend on threshold, with both curves showing an inverted-U shape. For this channel, contrast shows a peak SNR at $\theta = 2\sigma$, and orientation shows a peak SNR at $\theta = 2.5\sigma$.

We calculated how SNR depends on threshold for a population of 49 channels (Figure 7.9(a)). Only channels which were well-tuned to orientation (SNR > 0.75) were included in this and subsequent analyses. Every channel that showed a response to the stimulus demonstrated an SNR greater than 0.75 for contrast, and thus we chose the significant channels conservatively, based on orientation tuning. Like the individual channel example, there is an inverted-U shaped curve with the peak occurring between the extremes. The SNR curves for both orientation and contrast depend on threshold similarly, with contrast optimally represented at $\theta = 2\sigma$ and orientation optimally represented at $\theta = 2.5\sigma$ on average for the population.

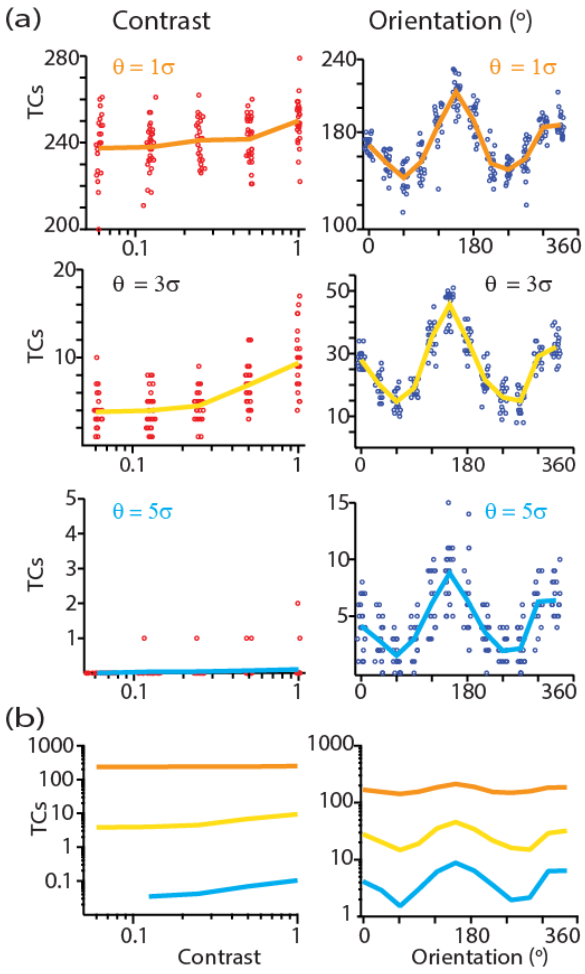

Figure 7.7: A single V1 channel example of how contrast and orientation tuning change with threshold. (a) Each data point represents the number of threshold crossings from one representative electrode recorded during a single trial. For visualization purposes, the data points are jittered with respect to orientation angle or contrast, respectively. To highlight the tuning, the mean threshold crossings to each orientation are connected and plotted using the color scheme in figure 2 (orange = 1σ, yellow = 3σ, light blue = 5σ). These curves are overlaid on the same plot in (b). Note the log scale on the TC axis.

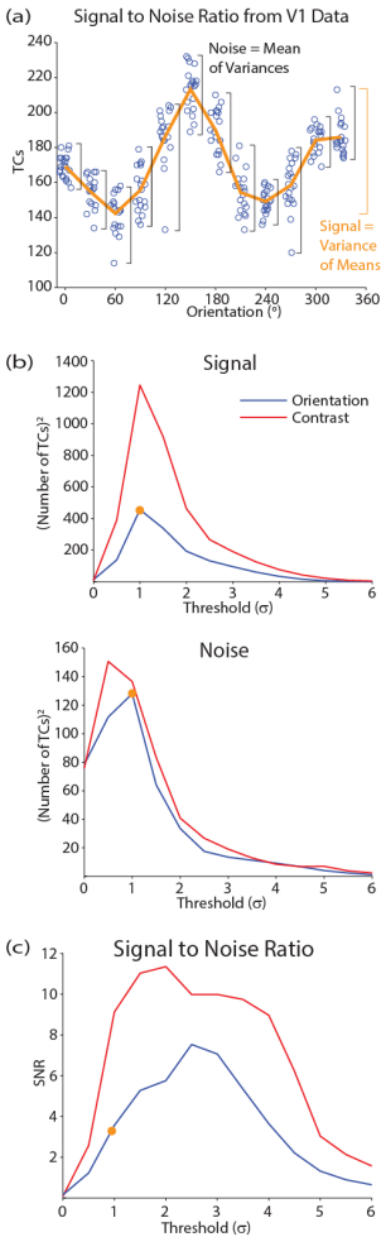

Figure 7.8: SNR quantifies information content in V1. (a) The total number of threshold crossings is plotted against orientation angle for a single channel at a threshold of θ=1. (For visualization purposes, the data points are jittered around the true orientation angle.) Signal is defined as the variance of the mean number of threshold crossings across each orientation (orange). Noise is defined as the mean of the variance of the number of threshold crossings across each orientation (black). (b) The calculations are performed at each threshold for orientation (blue) and contrast (red). Signal and noise both vary with threshold setting. The orange dot highlights the values that come from the tuning curve in A. (c) Combining the relationships in B shows that SNR exhibits an inverted-U shaped relationship with threshold.

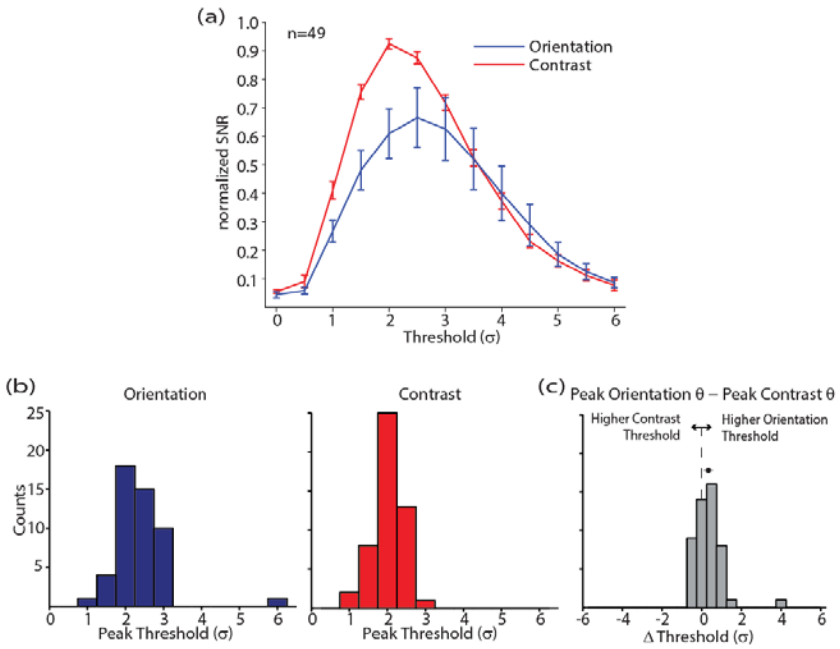

Figure 7.9: SNR in V1 depends on threshold. (a) SNR dependence on threshold for orientation (blue) and contrast (red) for all V1 electrodes with significant tuning, n=49 (normalized mean ± SE). (b) Optimal thresholds for encoding orientation and contrast. (c) The per-channel difference between the optimal thresholds for orientation and contrast. The mean ± SE is shown above the histogram. The mean is significantly different than 0 (t-test, p=0.004), and significance still holds when the outlier at $\Delta\theta=4\sigma$ is removed (p=0.0045).

The optimal threshold histograms for orientation and contrast are plotted in Figure 7.9(b). Both distributions of optimal contrast thresholds and optimal orientation thresholds are narrow with primarily low thresholds. However, the optimal threshold for orientation is higher than the optimal threshold for contrast for 26 of the 49 (53%) channels and the distributions have statistically different means (t-test, p=0.004). The mean pairwise difference between the optimal orientation threshold and optimal contrast threshold is $0.35\sigma \pm 0.11\sigma$ (Figure 7.9(c)).

7.5 DISCUSSION

We assessed the information content of extracellular recordings from M1 and V1 by systematically sweeping the voltage detection threshold, counting the number of threshold crossing events at that threshold setting, and evaluating how much information those threshold crossings provided about external parameters of interest. We found that optimal threshold depends on the parameter of interest. Specifically, directional parameters, like velocity and orientation, have higher optimal thresholds than scalar parameters, like speed and contrast. Regardless of the parameter of interest, the optimal thresholds for information were lower than the thresholds typically used in closed-loop BCI studies in which threshold crossings are used in lieu of spike sorting. We can make sense of these observations with a consideration of cortical topography. These results have implications for the optimal decoding of neural signals.

7.5.1 Cortical topography can explain optimal thresholds

How a stimulus parameter is represented in an extracellular voltage trace will depend in part on how the topographic scale of tuning to that parameter in the cortex relates to the effective sampling radius of the electrode, as determined by the detection threshold. At high detection thresholds, threshold crossings reflect the tuning of individual neurons. At low detection thresholds, threshold crossings comprise multi-unit activity and tuning likely reflects the homogeneity of the tuning of local neurons. Modeling studies have suggested that single unit activity arises from neurons within 50 μm of an extracellular electrode and multi-unit activity arises within 50-140 μm of the electrode (Martinez et al. 2009; Pedreira et al. 2012). Thus, it is

reasonable to expect that the topographic scale at which a stimulus parameter is represented impacts threshold crossing tuning, particularly at low thresholds.

Figure 7.10 schematizes a putative explanation for the effects of threshold selection that we observed. As the detection threshold of an electrode is moved toward 0, its effective sampling radius increases (Figure 7.10(a)). (Note that while the relationship between detection threshold and effective sampling radius is probably not linear, it is likely to be monotonic). As the threshold is lowered, the number of threshold crossings increases, as does the variability in the waveform shapes. A strict threshold, like $\theta = 5\sigma$ (blue), yields waveforms that likely originate from a single neuron. On the other hand, if we relax the threshold to $\theta = 1\sigma$ (orange), the waveforms are almost certainly not from a single unit.

A schematic example of parameters with different topographic scales relative to the "listening sphere" of an electrode is shown in Figure 7.10(b). Here the black arrows represent a directional parameter that has a small tuning scale relative to the listening sphere of the electrode, meaning the correlation in tuning among neurons falls off relatively quickly with distance. The light and dark gray regions represent a parameter that has a large tuning scale relative to the listening sphere of the electrode, meaning the correlation in tuning falls off relatively slowly with distance. The topographic scale of the stimulus parameter impacts the information present in the extracellular recording at different thresholds. At low detection thresholds, threshold crossings comprise multi-unit activity and tuning should better reflect those parameters that are homogeneously encoded among the population of neurons local to the electrode (Figure 7.10(c), gray). In contrast, we expect parameters that are more heterogeneously encoded to be better represented at high detection thresholds, where threshold

crossings reflect the tuning of individual neurons (Figure 7.10(c), black). Accordingly, in M1 we observed that velocity has a higher optimal threshold than speed.

7.5.2 The influence of topographic scale on information encoding by threshold crossings

We introduce a conceptual model based on the topographic scale of information encoding to explain our results (Figure 7.10). If we apply that model to V1's pinwheel organization of orientation preference, we should expect a change of <30 degrees for neurons within the putative sampling radius of our electrode at the lowest detection threshold. So, the topographic scale of orientation is on the order of the sampling radius. The topographic scale of contrast is larger than for orientation: nearly all V1 neurons increase their firing with increases in contrast. In accordance with this understanding of V1 topography, we found the optimal orientation threshold to be similar to but slightly larger than the optimal contrast threshold. Additionally, the optimal thresholds in V1 were relatively low, suggesting that including threshold crossings from more neurons provides more information than does a single neuron.

The heterogeneity of the local M1 preferred direction (PD) map is in stark contrast to the large-scale topography of V1 orientation columns (Schieber and Hibbard, 1993). However, in an effort to make sense of the structure of M1, a columnar organization similar to that observed in V1 has been proposed (Amirikian and Georgopoulos, 2003). This hypothesized structure of M1 consists of mini-columns of neurons with similar PDs 30 μm in width and repeating every 240 μm (Georgopoulos et al., 2007). Such structure would lead to a nearly complete set of PDs represented by neurons within the ~200 μm sampling radius of an electrode. This is a far less homogeneous local structure than that seen in V1. Consistent with this model, at high detection thresholds, velocity is encoded well by threshold crossings. This recapitulates the well-known

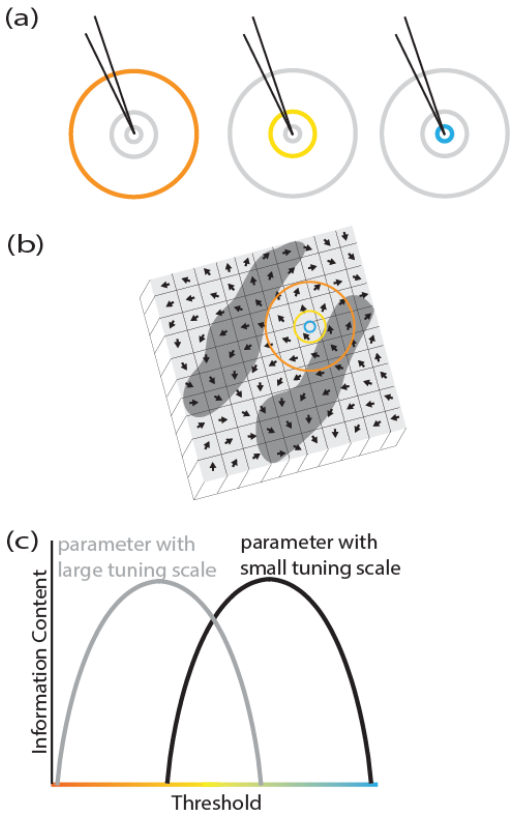

Figure 7.10: Information content depends on the voltage detection threshold and the topographic scale of the parameter of interest. (a) A change in the detection threshold might change the effective sampling radius of the electrode. As we decrease the detection threshold of an electrode (move from blue to yellow to orange), we increase its effective sampling radius. (b) The relationship between effective sampling radius and the topographic scale of an encoded parameter. The black arrows represent a parameter that is encoded on a small scale. The gray regions represent a parameter that is encoded on a larger scale. The color scheme of the sampling radii is the same as above. (c) The information content of a signal depends on the threshold setting and the local topography such that a parameter encoded on a large scale (gray) has a lower optimal threshold than a parameter encoded on a small scale (black).

preferred direction tuning of individual neurons in M1 (Schwartz et al., 1988). However, within the larger effective sampling radius specified by a low threshold, the diversity of PDs of the contributing neurons weakens the measured velocity tuning. On the other hand, most M1 neurons tune monotonically to speed (Moran and Schwartz, 1999). Thus, speed encoding is strongest at low thresholds, since many neurons contribute to the threshold crossings. This can explain our observation of higher optimal thresholds for velocity than for speed.

The information available at low voltage threshold settings is not just a watered-down version of the information available at higher thresholds. Although speed may be thought of as a less specific version of velocity (speed can be derived from velocity, but velocity is not uniquely specified by speed), speed and direction are independent quantities, and when we repeat our analyses using movement direction, we find direction and velocity have similar optimal thresholds (data not shown). Further, contrast cannot be derived from orientation, and we find the best threshold for orientation information is higher than for contrast. The primary characteristic influencing the optimal threshold of a parameter is whether it is represented homogeneously by the cortical population, or heterogeneously. The directional quantities, velocity and orientation, are heterogeneously represented, while the scalar parameters, speed and contrast, are homogenously represented.

Our results imply that even single electrodes might be useful for inferring the topography of tuning properties in brain areas where it is not known. By sweeping the event detection threshold and computing the SNR to various parameters of interest, some notion of the relative homogeneity of tuning to different parameters can be gained. Parameters that drive neurons in a heterogeneous, uncorrelated way over short spatial scales should be best represented at relatively high thresholds. In contrast, parameters whose tuning correlates over larger spatial extents should

be better represented at lower thresholds. This knowledge could be critical in designing more effective extracellular recording experiments to reveal the nature of the information present in a given brain area.

7.5.3 Implications for online decoding

In many successful BCIs to date, information is extracted directly from sorted spikes recorded from chronically implanted extracellular electrodes. Although BCIs based on sorted spikes have shown impressive performance both in the lab (Wessberg et al., 2000; Taylor et al., 2002; Velliste et al., 2008; Ethier et al., 2012; Gilja et al., 2012; Ifft et al., 2013) and in controlled clinical trials (Simeral et al., 2011; Collinger et al., 2013), spike-sorting is widely acknowledged to be time-consuming, and hard to automate (Lewicki, 1998). Because of these challenges, the spike-sorting step, once thought to be critical to BCI performance, may actually inhibit the translation of BCIs from the lab to the clinic. Here we add to that perspective by showing that better information extraction might be possible if thresholds are tuned for the parameter of interest, rather than set as if for spike sorting. The benefits should be especially salient for electrode channels where no identifiable single neuron is present.

The use of threshold crossings is becoming more prevalent in online decoding studies. This is not surprising given that in offline analyses multiunit activity and threshold crossings have yielded decoding performance and encoding fidelity that is comparable to or better than sorted spikes or local field potentials (Kloosterman et al., 2014; Stark and Abeles, 2007; Ventura, 2008; Malik et al., 2014; Todorova et al., 2014; Chestek et al., 2011; Christie et al., 2015; Perel et al., 2015). Recently, we and others have begun to recognize the need to investigate threshold setting in a principled way. Christie and colleagues (Christie et al., 2015) found optimal

thresholds for decoding performance to be between 3 - 4.5 times the rms voltage (V_{rms}). Importantly, they only considered threshold settings from 3 - 18 x V_{rms}; they did not consider threshold crossings at lower threshold settings. A separate study by our team included lower voltage fluctuations in their threshold crossings by using an approach similar to the exclusive windows analysis presented here, and found that threshold crossings at $\theta = 3\sigma$ actually improved decoding as compared to only well-sorted spikes (Todorova et al., 2014). This result corroborates our finding that the low voltage fluctuations are not noise, but rather, they do contain useful information: speed-related information that is distinct from the velocity information present at higher thresholds. For the thresholds and parameters we considered, the optimal thresholds were lower than typically used in online decoding studies (Gilja et al., 2012; Hochberg et al., 2012; Sadtler et al., 2014), suggesting that there is information available in extracellular recordings that is being discarded, and could be useful for improving BCI performance. Ideally, the detection threshold would be customized for each channel. In fact, each channel's signal could be duplicated and thresholded separately for each parameter used in a BCI.

Notably, the thresholds that we found to be optimal are lower than thresholds typically reported in published studies. In M1, particularly for online brain-computer interface experiments using threshold crossings, a fixed threshold of $\theta = 4.5\sigma$ is commonly chosen (Gilja et al., 2012; Hochberg et al., 2012) presumably because it approximates spike sorting. Some of the best online BCI control has been achieved with this commonly chosen threshold. However, we observed only 14 of 67 (20%) channels that showed optimal velocity thresholds greater than or equal to $\theta = 4.5\sigma$. This was even more apparent for speed, where only three channels had optimal thresholds at the level commonly chosen. Optimal thresholds are likely to depend on

many factors, including recording quality and the age of the implant. It stands to reason that even better BCI decoding may be possible if the threshold is chosen with information content in mind.

7.6 CONCLUSIONS

Historically, neurophysiologists have processed extracellular voltage recordings to extract action potentials from isolated single neurons, essentially ignoring small amplitude voltage fluctuations. This makes sense if the focus is on a careful characterization of the properties of single neurons. However, if the goal is to get as much information as possible from a recorded signal, processing can only reduce available information (Cover and Thomas, 1991). Accordingly, we have shown that non-spike parts of the recorded signal, in particular the low voltage fluctuations, include useful information about some parameters, and should not be discarded as noise. Our results suggest that signal preprocessing in neurophysiology experiments deserves careful consideration: one approach does not necessarily fit all applications. For recordings from a given brain area, it would be advantageous to sweep a range of thresholds to find the optimal choice for the desired information and planned experiment. For applications that do not require real-time processing, there is value in streaming the entire raw voltage signal to disk for offline analysis, and then considering the information content at different threshold settings. The practice of adjusting the detection threshold to the parameter of interest may improve our ability to determine how the brain is organized to encode sensory information, and it may improve our ability to accurately decode motor intentions.

8.0 ADDITIONAL PROJECTS AND CONTRIBUTIONS

8.1 ADDITIONAL PROJECTS

In addition to the larger projects described in Chapters 3 - 7, I was also involved with several smaller projects and science outreach to middle school and high school students. The additional projects as well as my outreach are described in this chapter.

8.1.1 Sensory Motor Index: How does sensory information influence M1 when sensory feedback is essential for behavior?

This work began as an NRSA fellowship grant application and has developed into the immediate future for the Critical Stability Task in the Batista lab. This project is being continued by a fellow graduate student, Nicole McClain, who will be leading the critical stability task (CST) project in Dr. Batista's lab after I graduate. The CST enables us to study ongoing behavior is a rich and broad new experimental approach, and I could not hope to explore every aspect of it during my PhD. As part of my graduate experience, I have trained Nicole to use CST as our lab delves into the multiple ways the CST can help us understand sensorimotor integration and improve BCI control. Her skills in control theory, signal processing, and programming will help solidify the impact of the CST on neuroscience. This project addresses one of the multiple paths CST will lead us down. I developed the idea for this project, performed preliminary analysis, trained

animals to perform the proposed tasks and created simulations to verify my analyses and ideas. Nicole will collect monkey data for this study and publish a manuscript as part of her dissertation, for which I will be second author.

The goal of this project was to determine how the brain controls continuous movements that require sensory feedback, such as walking or manipulating an object. Doing so must require a careful neural orchestration between motor commands and sensory information. Historically, researchers have focused on the role of primary motor cortex (M1) in generating motor commands (Evarts, 1968; Georgopoulos et al., 1982; Georgopoulos et al., 1986; Schwartz et al., 1988; Georgopoulos et al., 1988; Ashe and Georgopoulos, 1994), but more recent research has made it evident that there is a sensory component to M1's function (Ashe and Georgopoulos, 1994; Pruszynski et al., 2008; Pruszynski et al., 2011; Hatsopoulos and Suminski, 2011; Pruszynski et al., 2014; Rao and Donoghue, 2014). I proposed to extend this line of inquiry to examine how M1 uses sensory information to guide and shape ongoing movement, using multineuronal recordings during a novel behavioral task in which movements must be shaped continually through sensory feedback.

Sensory feedback is vital for activities of daily living. Behaviors such as using a tool require that we tailor our actions to the environment in an ongoing manner. Yet, the dominant paradigm for studying the neurophysiology of sensorimotor control is to use single, relatively brief movements with a discrete start and end, such as a center-out reach (Georgopoulos et al., 1982; Georgopoulos et al., 1986; Schwartz et al., 1988; Georgopoulos et al., 1988, Batista et al. 1999). These movements are highly practiced by the time neurophysiological data collection begins, and do not require ongoing sensory feedback for reasonably good performance (Gordon et al., 1995; Ghez et al., 1995). I believe that the knowledge we have gained about M1 from such

studies does not enable us to fully describe how the brain functions to perform continuous movements that require ongoing sensory feedback. To explore sensorimotor control at its essence, I have implemented a task that places Rhesus monkey subjects "in the loop" of a prolonged motor control task that cannot be performed without sensory feedback. In this task, the Critical Stability Task (CST), monkeys must maintain a drifting cursor at the center of the workspace for several seconds by making corrective hand movements, or by issuing neural commands via a brain-computer-interface (BCI). Recording multineuronal activity from M1 during the CST will enable us to probe how sensory feedback impacts motor commands in an ongoing manner. My central hypothesis is that sensory information exerts a powerful influence on M1 activity when sensory feedback is essential for behavior.

I hypothesize that primary motor cortex (M1) will exhibit a strong sensory response to visual feedback during a prolonged motor task where such feedback is critical for task success. Throughout the duration of this project (continuing beyond my time in the lab), we will record neural activity in M1 with a 96-electrode array while monkeys perform the CST using arm movements. The CST allows us to cleanly distinguish neural responses to visual feedback from motor commands (and proprioceptive feedback) because visual input and motor output are oppositely directed: success requires that the subject's corrective actions oppose the direction of the cursor motion; a rightward error must be corrected with a leftward movement, and vice versa. A strong prediction of my central hypothesis is that as the need for sensory feedback is increased, the sensory representation in M1 neurons grows stronger. We will test this by increasing the degree of difficulty of the CST, and looking for an increase in the amount of visual sensory information present in M1.

8.1.1.1 Significance

The central hypothesis of this research is that sensory information exerts a powerful influence on activity in primary motor cortex (M1) when sensory feedback is essential for behavior. M1 is the final common pathway for voluntary motor commands – 30-50% of the cortical projections to the spinal cord originate from M1 (Porter and Lemon, 1993). While much is known about the role of M1 in motor function, we know comparatively little about how sensory feedback shapes the motor commands issued by M1.

The canonical view of M1's function is perhaps best captured in the 1999 study by Moran and Schwartz (Moran and Schwartz, 1999), which showed that the activity of individual M1 neurons is tightly correlated with the velocity of the hand. The activity of a population of M1 neurons can predict the details of a hand movement (Georgeopoulos et al., 1982). Beyond this canonical view, evidence exists that points to a more complex role for M1 in motor control (Kakei et al., 1999; Evarts, 1972, Churchland and Shenoy, 2007; Schieber and Hibbard, 1993). However, all of these studies investigate the neural mechanisms of movements that have a discrete beginning and ending and that are highly stereotyped through extensive practice. It is difficult to infer from these studies how M1 might function to control ongoing behaviors. Many activities of daily living require delicate adjustments to behavior based on sensory feedback (e.g., consider carrying a cafeteria tray or steering a car). It stands to reason that M1 is doing far more than issuing velocity commands to control these behaviors.

Recent studies documenting sensory responses in M1 are expanding our understanding of its function (Pruszynski et al., 2011; Hatsopoulos and Suminski, 2011; Pruszynski et al., 2014; Rao and Donoghue, 2014). Together, these studies show that sensory information activates M1 neurons, in many cases even more strongly than the motor command. However, in each of these

experiments, only a single, brief sensory stimulus was administered. This leaves open the question of how M1 responds to the ongoing sensory information that guides real-world actions. My research is intended to understand the function of M1 during continuous behavior. In particular, I focus on the sensory aspects of motor control.

Perhaps the greatest benefit of a richer understanding of the function of M1 will be dramatic improvements in the capabilities of Brain-Computer Interfaces (BCI). BCI is an emerging clinical technology that holds promise for the restoration of motor function to paralyzed individuals (Hochberg et al., 2006). However, current BCI systems cannot achieve the speed, accuracy, and flexibility of natural movements. Nearly all current BCI systems are designed to provide point-to-point cursor or robotic movements and a small number of systems can provide grasp (Hochberg et al., 2012; Collinger et al., 2013). The limited function of clinical BCI systems may stem in part from the fact that our understanding of M1 is built almost entirely from point-to-point tasks. To enable BCI systems to restore the full repertoire of motor function, we must understand how M1 functions during continuous, feedback-driven behaviors.

8.1.1.2 Innovation

This work includes two chief innovations – a new scientific approach and a new experimental technique. The innovation in the scientific approach is to pursue the neural mechanisms of ongoing actions that are tightly orchestrated by sensory feedback. This is in contrast to traditional approaches to studying M1, which use discrete actions that are highly practiced and that require feedback only minimally (Gordon et al., 1995; Ghez et al., 1995). My innovative technique is a new behavioral task, the Critical Stability Task (CST), which enables me to probe the role of M1 in sensory-guided action. The CST was originally developed to study sensorimotor performance in pilots (Jex et al., 1966) and has been used since to study

neurological health (Potvin et al., 1977; Kondraske et al., 1984), to assess vibrotactile feedback for balance assistance (Kadkade et al., 2003), and to study mental workload (Burke et al., 1980). The CST captures a simple but challenging interaction with an object: the user must balance a virtual system.

I believe that the CST provides a powerful new framework for studying sensory-motor coordination. I have been helping to develop this technique since I began graduate school in 2012. I contributed to proof-of-concept human behavioral studies, presented in Chapter 5 (Quick et al., 2014), and I have been leading the effort to perform these experiments in Rhesus monkeys, presented in chapter 4. So far I have trained two animals to perform the CST (Quick et al. 2015), and I have trained one to perform all of the experiments proposed here. We were going to implant Monkey I with an array in primary motor cortex and perform these experiments before I graduated, but I will instead use this opportunity to train my predecessor, Nicole McClain, to train the CST and collect data.

8.1.1.3 Approach

We have adopted a behavioral paradigm that is new for animal experiments, the Critical Stability Task (CST), which requires continuous sensory-driven interaction with a simple but challenging virtual system. The CST allows us to observe and quantify neural responses in primary motor cortex (M1) when sensory information is essential for motor control.

The Critical Stability Task

To perform the CST, monkeys must maintain a drifting cursor at the center of a computer display for several seconds. The monkey must act to stabilize the cursor via hand movements or

via BCI control. If no corrective action is made, the cursor will drift rapidly off-screen. The CST is a first-order system with unstable dynamics governed by the differential equation,

$$\frac{dx(t)}{dt} = \lambda\big(x(t) + u(t)\big), \quad \lambda > 0 \tag{8.1}$$

where x(t) is the horizontal position of the cursor on the computer monitor, u(t) is the horizontal position command generated by the subject, t is time with t≥0 and t=0 indicates the start of the trial. The system is characterized by the value of the parameter λ (which has units of frequency). Because λ is positive, the cursor position will diverge over time unless appropriate corrective actions are made by the monkey (Dorf, 1967). If λ is large, the cursor will diverge rapidly, and the CST is more challenging. We increase λ until we determine the maximum λ for which a subject can control the cursor. This is called the subject's "critical instability value" (CIV) since beyond that value the closed-loop system becomes unstable. The greater their CIV, the better the subject is at the task. Balancing an unstable system to understand motor control has precedents (Peterka, 2002; Peterka and Loughlin, 2004) and it is an important new approach in neurophysiology because it pushes sensory-motor control to its limits.

To appreciate how the CST enables my studies, consider the position command that is required to keep the cursor from drifting, i.e., what is u(t) such that the cursor velocity is zero? Plugging the condition dx(t)/dt=0 into Eq. 1 shows that the subject must generate a position command that is exactly equal and opposite to the current cursor position: u(t)= -x(t).Thus, feedback is essential: the subject must know the current state of the cursor to keep the system in balance. A subject who could issue this perfect motor command would be able to keep the cursor completely still. In practice, this perfect control command is not possible because of sensorimotor noise and nervous system processing delays; hence, the cursor will continuously move, and maintaining control requires ongoing sensory-guided action.

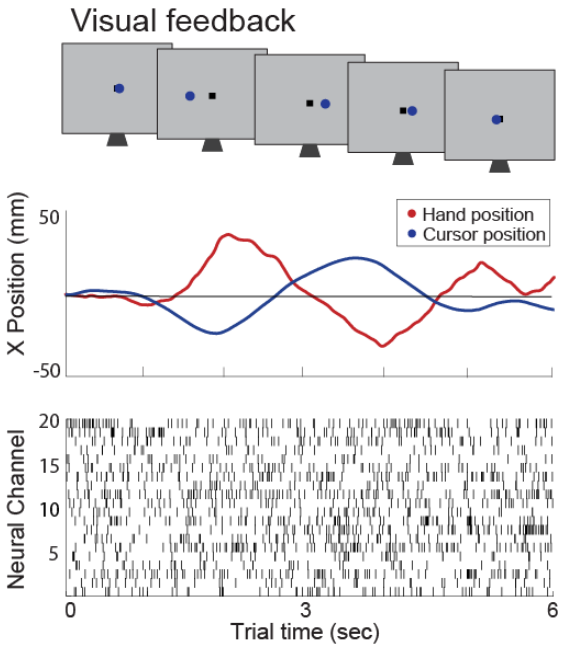

Figure 8.1: Example CST trial with hand control. Top row: schematic of visual feedback. Middle row: cursor and hand position traces. Bottom row: raster of multichannel neural activity from M1.

One trial of the CST is shown in figure 8.1. The top panel schematizes what the animal sees on the computer monitor at five time points during the trial. The middle panel shows the cursor position (in blue) and the hand position (in red). The cursor is initialized to the center of the screen at the start of each trial, but due to inevitable motor noise, it soon begins to drift. The animal does his best to move his hand equal and opposite to the cursor position, but no correction is perfect, so the cursor continues to drift. Thus, ongoing movement is required to keep the cursor stabilized. The third panel shows neural activity for this trial on 20 M1 electrodes.

I believe the CST offers a new, complementary avenue to study the neurophysiology of sensory-motor control in primates. It is illuminating to contrast it with a dominant approach in behavioral neurophysiology: the point-to-point reach task. That task was developed to provide a highly repeatable behavior, in part to facilitate data collection in single-unit neurophysiology,

where trial-averaging was a necessity. In contrast, the CST allows for study of neural signals in a setting where sensory and motor signals interact. Behavior is impacted by a changing environment, and so although the goal of the task is always the same, no two trials are identical. Thus, the CST captures many important aspects of environmental interactions such as manipulating an object. The CST provides an ideal initial testbed to explore the neurophysiology of flexible sensorimotor interactions.

I have employed the CST with nine human subjects (see Chapter 5) (Quick et al., 2014), and two monkey subjects using arm movements (see Chapter 4) (Quick et al., 2015). These studies provide important proof-of-concept demonstrations that the CST can be used to study behavior and neurophysiology in Rhesus monkeys.

Neural Recordings

Monkey I has been trained to perform all the proposed experiments and will be implanted with a 96-channel electrode array in the shoulder and forearm portion of primary motor cortex (M1) contralateral to the arm the animal uses to perform the task. Multielectrode arrays tend to yield 80-96 electrode channels with good-quality neural recordings. On a subset of channels, single-unit action potentials can be isolated. On the remaining channels, multiunit waveforms (probably consisting of action potentials from several neurons near the electrode tip) can be recorded by setting a voltage detection threshold at about 3.0x the root-mean squared signal amplitude (Oby et al., 2016). All analyses described below can be performed on single-unit isolations and multiunit recordings. For simplicity, henceforth I refer to both types of recordings as “cells” or “neurons”. Each analysis was performed using only a single day’s data - my hypothesis do not require us to assure that recordings are the same from one day to another, nor can I assert that they are independent.

Neural Data Analysis

In the CST, the details of the animals' behavior and sensory experience differ from trial to trial. This is a virtue of the CST, because real movements tailored to the environment are not stereotyped. The brain operates in a "single-trial" mode (Ganguly et al., 2011): motor areas do not have the luxury of waiting for an average neural response to accrue over many repetitions of the same stimulus. Traditional trial-averaged analyses are usually justified by the reasoning that cells with similar tuning properties are active at once, so single-neuron average responses may form a reasonable proxy to the population-wide single trial response (Britten et al., 1996). Multielectrode arrays enable single-trial views of neural activity (Kaufman et al., 2015; Kao et al., 2015). Here, I employ analysis techniques that can detect stable properties of neural tuning despite each trial of the CST being unique. I believe that studying neural activity during the CST will break new ground in learning how the brain controls ongoing, sensory-guided behavior. The particular research questions and analyses described here are just one approach among many that these data will serve.

The CST allows me to cleanly distinguish neural responses to visual feedback from motor commands (and proprioceptive feedback) because the two point in opposite directions: while performing the CST, a rightward error must be corrected with a leftward movement. In this work, I distinguish visual responses from motor commands (and proprioceptive responses). This allows me to ascribe neural activity to visual sensory responses and motor commands/proprioceptive responses.

Each experiment session begins with 80 trials of a conventional 8-target center-out reach task (Georgeopoulos et al., 1982; Churchland and Shenoy, 2007). This data is used to establish the response properties of the neurons on the array in a conventional paradigm. Particularly, I use

this data to extract the preferred direction (PD) of each neuron I record. For example, the cell in figure 8.2A responds best for rightward reaches. In the CST, does this cell respond best when the hand is moving to the right (and the cursor moves left), or when the cursor moves right (and the hand moves left)? The former indicates that the cell is more motor than sensory, and the latter indicates that the cell is more sensory than motor. Figure 8.2B shows the response of a neuron during the CST. To determine the proportion of a neuron's firing rate that is well-explained by the visual component of the task, and the proportion of the neuron's firing rate that is well-explained by the motor component of the task, I will compute a sensory motor index (SMI) for each cell. I will define the task-relevant sensory contribution (S) of a cell as that cell's average firing rate when the cursor velocity is in the cell's preferred direction. In the same fashion, I will define the motor contribution (M) of a cell as that cell's average firing rate when the hand velocity is in the cells preferred direction I will then compute SMI=(S-M)/(S+M). An SMI of 1 indicates a cell responds solely to visual feedback, while an SMI of -1 indicates a cell that drives arm movements, or that the cell responds to a combination of motor command and proprioceptive feedback. An SMI near zero indicates the presence of comparable visual and motor components to the neuron's response.

The data presented in Figure 8.2 is preliminary data collected in 2013. The animal's array was not in good condition, with only a third of the channels demonstrating modulations depths above 5Hz. To demonstrate the power of this analysis technique, we have simulated neurons with purely motor and purely sensory responses and calculated the SMI, yielding SMI values of -0.9997 and 0.9811 respectively.

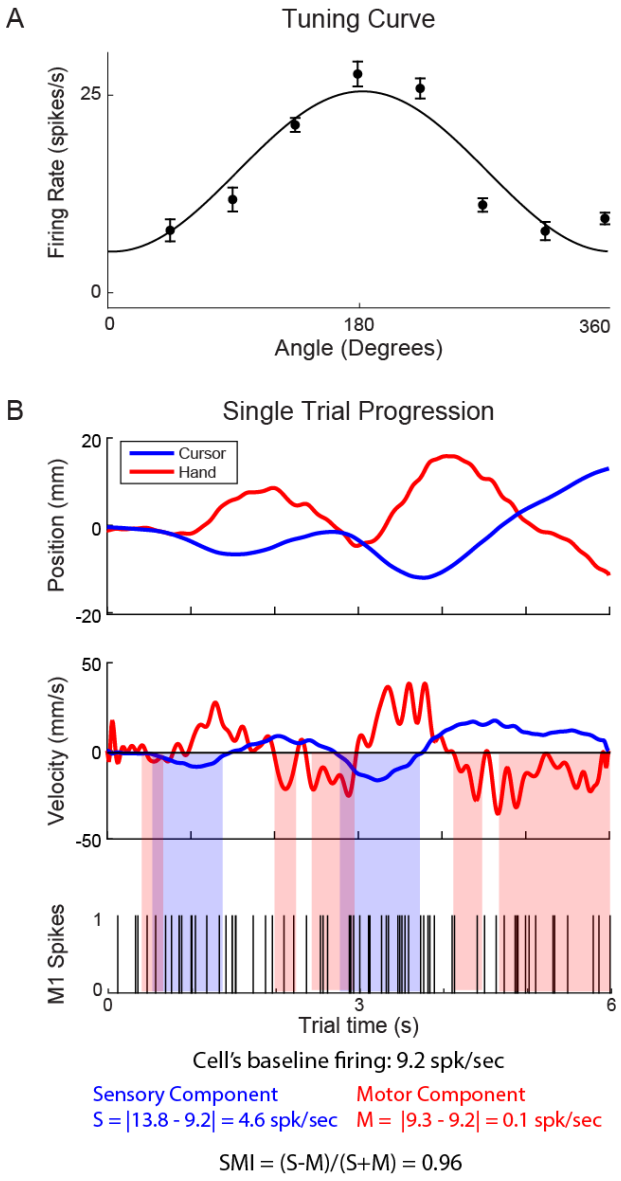

Figure 8.2: SMI calculation. A) tuning curve for M1 cell with PD to the left (180 degrees). B) Response during CST. Colored bands illustrate data used in the SMI calculation. When the cursor velocity is leftward (negative), the neural response contributes to the S measurement (blue bands.) When the hand velocity is leftward, the neural response counts toward the M measurement (red bands).

8.1.1.4 Discussion

Expected results

The overarching goal of this project is to understand how the brain controls ongoing actions guided by sensory information, and determine the extent to which sensory information is represented in M1. I hypothesize that M1 will exhibit a strong sensory response to visual feedback during a prolonged motor task where feedback is critical, and that as task difficulty increases the sensory representation in M1 neurons will become stronger due to the greater reliance on sensory feedback.

Calculating the sensory-motor index (SMI) as detailed above for each M1 neuron recorded from during the CST, I expect to see the distribution of SMIs shown in Figure 8.3. While most of the cells have a predominantly-motor response (SMI<0), some cells will exhibit stronger visual responses (SMI>0). This subpopulation of M1 neurons responds strongly to visual feedback during the CST.

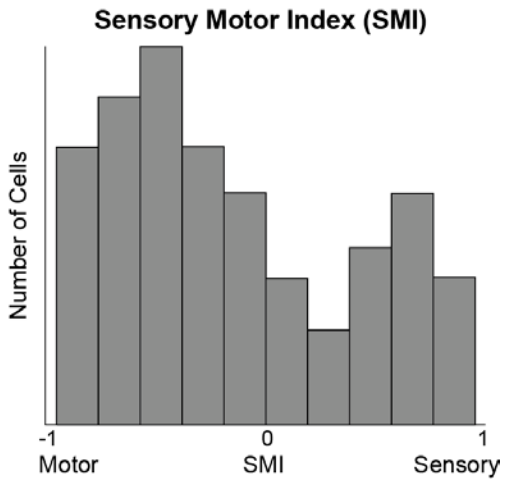

Figure 8.3: Expected SMI for M1 population.

If visual information is present in M1, it probably serves to help drive accurate movements. As the task becomes more demanding, sensory information becomes more essential for good performance; and therefore I expect SMI will increase among the cells that have a positive SMI – cells with a strong sensory component. To test this, animals will perform the CST at various λ values. As λ increases, the cursor will diverge more rapidly, requiring faster and more precise movements to maintain balance, and the strength of sensory information in M1 may increase accordingly. Calculating the SMI of each neuron at each tested value of λ, my hypothesis is supported if the SMI of visual cells to increases as a function of λ. I expect a larger number of cells to exhibit a visual response (SMI>0) at higher values of λ (Figure 8.4A) and the cells that are visual at lower values of λ will increase their SMI at higher λ values (Figure 8.4B).

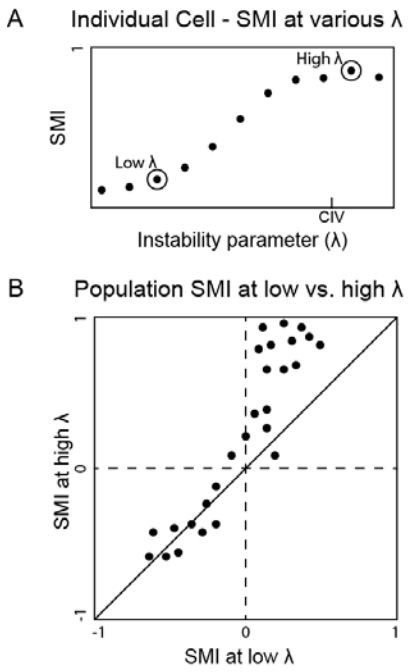

Figure 8.4: SMI changes as a function of λ. A) Hypothetical SMI for a visual cell at multiple difficulties. B) Hypothetical SMI of M1 population at low λ vs high λ.

Proprioceptive input to M1

When the CST is performed with hand control, visual information is directed opposite to motor commands, by the nature of the task. However, proprioceptive feedback about arm movement is aligned with the movement direction. Therefore the claims above can only determine if a cell is *visual* sensory vs motor/proprioceptive. Performing the CST using brain control will allow us to perform the same experiments and analysis, dissociating motor driving signals from proprioceptive responses. During BCI control, the hand is still, or makes small movements not consistently related to the movement direction (Sadtler et al., 2014). Thus, the amount of neural activity during BCI control can be compared to the amount of activity during hand control in putative motor neurons. I hypothesize that proprioceptive feedback is represented in M1, and thus the amount of activity in some putative motor neurons will decrease during BCI control. If the amount of activity in a given neuron decreases during BCI control, then one could conclude that a portion of its activity was due to proprioception. Performing the experimental paradigm and analysis described above while an animal performs the CST using hand control and brain control using should allow us to quantify the relative contributions of vision, proprioception, and motor command on each M1 neuron.

8.1.2 Visuomotor rotation BCI experiments

My first graduate school project explored the way in which the brain adapts to a visual-motor perturbation. Collaborating with Drs. Steven Chase and Byron Yu from Carnegie Mellon University, I developed an experimental paradigm that examines whether motor learning occurs in parallel throughout the brain, or is organized hierarchically. I began my experiments by training a monkey to adapt to a visual-motor rotation task using arm movements, dissociating his

arm movements from his visual perception of the reach. During this reaching task, I identified neurons in motor cortex (M1) that mediated motor learning and those that were static. I then repeated the adaptation task using brain-computer interface control, restricting control to either the adapting population, or the non-adapting population. We developed two hypotheses. First, we predicted that the neurons capable of adapting during actual reaches would contribute to greater motor learning in the BCI task. Second, we predicted that neurons not controlling the BCI cursor would only exhibit plasticity if they were involved in motor learning during the reaching task. These hypotheses were motivated by a classic hierarchical view of motor control: some neurons are involved in goals, and others in execution. To our surprise, neither hypothesis was borne out by the data. Both populations of neurons showed learning in the BCI task, and neurons not directly involved in the BCI control also showed plasticity. I believe this data supports an "equipotential" view of M1: neurons can "step to the plate" as they are needed to accomplish a goal. Understanding how the brain controls movement in a flexible manner will help us learn more about diseases that affect motor control and help us design restorative solutions. This work was abandoned shortly after it began, but the studies have become preliminary data for an RO1 grant application for Dr. Chase.

8.1.3 Hand control vs. observation vs. brain control

As I was working for Dr. Steve Chase exploring how the brain adapts to a visual-motor perturbation, I began to wonder what the brain does during brain control generally. Do neurons in motor cortex act similar in brain control tasks as they do during reaching tasks? Or do they act like they do during the observation task that's used to train the decoder? To explore this question, I collected data while Monkey J performed 80 center-out reaches, passively watched a

cursor make 80 center-out movements, and controlled a cursor using brain control to make 80 center-out movements. For each condition, I fit a cosine tuning curve to each channel yielding the modulation depth, the preferred direction, and the baseline of each channel. We found that preferred directions of channels are most similar between observation and brain control conditions. This is what I expected because we use observation to estimate initial decoder parameters. The preferred directions are not exactly the same between these two conditions, showing that the animals use visual feedback during brain control trajectories to learn and adjust neural weights and preferred directions. I also found that observation training has higher baseline firing rates than hand control and brain control. This may be due to the fact that the animal is not trying to control the cursor, but instead just watching the task and being stimulated by the visual stimulus. That is, his attention is high, and therefore his neural activity increases globally. Additionally, modulation depths appear most similar between observation and brain control. Even though baselines are different, modulation depths remain similar. This is expected; neurons learn to modulate their firing rates based on the observation condition, and therefore need to recreate those modulation depths for accurate control.

The results of this preliminary work were expected, however more work must be done to more fully characterize how neurons behave during control of a BCI decoder. Do neurons behave more similarly to their training condition when a decoder is trained on observation or when it is trained on arm movements? When the decoder is trained on arm movements, it is trained on a condition where the arm is moving and then tested on a condition when they arm is not moving. When the decoder is trained on observation, it is trained and tested on a condition where the arm is not moving. How does the arms condition (moving or stationary) effect decoder performance? Does neural activity change as a decoder is learned and performance increases? To

address these questions, I propose the following experiment. During each experimental session, animals will perform 80 center-out reaches, passively watch a cursor make 80 center-out movements, and then perform 400 center-out BCI movements, using a decoder trained on either the reaching data or the observation data. A cosine tuning curve will be fit to each channel to characterize the neural activity during the hand control trials (10 trials to each target), observation trials (10 trials to each target), and then every 64 trials (first 8 trials to each target) of the BCI trials. I hypothesize that neural behavior will be variable for early BCI trials, but will converge to more closely resemble the neural activity that the decoder was trained on as the subject becomes more familiar and proficient with the decoder. I also think that when the decoder is trained on observation data, neural activity during control will more closely resemble training activity since the arm position is consistent across the conditions.

8.2 GENERAL LAB WORK

8.2.1 Animal training

During my time in Dr. Aaron Batista's lab, I trained several animals of many ages on various tasks. Training the animals in the lab has been a highlight of my PhD. I love working with animals, and find working with difficult animals especially rewarding. I have worked closely with five monkeys while at the University of Pittsburgh, each with their own personalities and obstacles to overcome.

Monkey J

Monkey J was the first animal I trained when I joined the Batista lab. Dr. Patrick Sadtler taught me how to work with Monkey J, how to collect data and to run brain control experiments in August 2012. I worked with Monkey J again starting in July 2013 when his array began to lose viable neural signals. I trained him to perform the critical stability task (CST) using hand control and visual or vibrotactile feedback, collecting behavioral data until June 2014. This work is demonstrated in Chapter 4. During my time working with Monkey J, he demonstrated severe self-injurious behavior (SIB). He had shown some signs of SIB before I joined the lab, but we had not seen any signs of this behavior until mid-late 2013 when we began seeing intermittent puncture wounds. These incidents were not severe and did not require surgery. However by the fall of 2014, Monkey J had two severe SIB incidents requiring surgery and body casts. He did not seem unhappy, yet he was put diazepam and fluoxetine at a veterinarian's recommendation. Studies continued as I regularly met with the enrichment specialist and the vet to attend to his increased need for enrichment and to tend to his wounds. We transferred Monkey J into a mixed colony which seemed to help for a while, but after a third severe incident in May, the IACUC decided that Monkey J should be enrolled in a terminal experiment. I finished recording data for my immediate project, and while I would not get to use Monkey J for the next stage of my study, Monkey J's data helped confirm the usability of the critical stability task for neuroscience applications. My experiences with Monkey J were the most transformative I've had in graduate school; I saw and performed my first brain control experiments, I became an independent researcher as I finished classes, I collected proof of concept data for a new behavioral task in the neuroscience field, I learned the importance of the relationship between researcher and animal in order to get clean data, and I learned that sacrifice is crucial in science.

Monkey K

During late 2012 and early 2013, I was training Monkey K to perform center out with brain control and hand control using an electrocortigraphic (ECoG) grid. This work is detailed in chapter 6. This project started before I entered graduate school, and was initially intended to study the long-term stability of ECoG recordings and BCI performance. While analyzing the BCI data, we noticed that Monkey K was using jaw movements to move the cursor. These artifacts in the data lead us to move Monkey K to hand control and report on the modulation of the ECoG signals while performing arm movements. Monkey K was abnormally aggressive, and in March, I was bitten requiring surgery, and was hospitalized for a week. His behavior was monitored for a few weeks to a month afterward, and his interactions with other staff and lab members lead to the termination of our experiments. Monkey K was used in a terminal experiment, and we changed the focus of our paper to the histological impact of a two year subdural ECoG array on the brain.

Monkey G

In 2013, while I was working with Monkey K, I trained Monkey G to perform center out reaches in a virtual reality (VR) environment. He was slated to be our second ECoG animal for publishing our ECoG BCI and hand control results. Monkey G was young when I began working with him, and was difficult to train, as that he was not keen on wearing a marker for motion tracking. After he was trained on center-out, I began scheduling an ECoG array implant, but after Monkey K was sacrificed, Monkey G was transferred to Dr. Raj Gandhi's lab for saccade studies. After his death, autopsy demonstrated that Monkey G had a detached retina and which may have been the cause of his problems.

Monkey I

I began working with Monkey I in June 2014. Ike's PMd array no longer yielded signals, and I trained him as the second monkey for the CST visual and vibrotactile behavioral experiments. When I started working with Monkey I, his behavior was inconsistent, aggressive, and unmotivated. Working with him was frustrating, but I worked 92 days in a row without a missing a day, and his behavior turned around. I collected the CST data sets for the work presented in chapter 4, and then Monkey I was given a few months off while I trained a second monkey for collecting neural activity during CST. Monkey I would be the first animal, and while he wasn't working, we explanted his PMd array and all the hardware on his skull in order to allow the skull and the brain to heal before implanting an M1 array for my work. I began training Monkey I again in January 2016, and his great worth ethic did not dissolve while he wasn't working. He was able to jump right back into CST with visual or vibrotactile feedback and can switch between tasks without impacting the data. Since January, I have collected center-out data, saccade target data, center-out with eye fixation, and CST with eye fixation from Monkey I. I scheduled to implant an array into left primary motor cortex in July and begin collecting neural data during the CST, but this surgery has been postponed and will likely be completed by Nicole McClain, who is continuing the CST project after I graduate. I will transfer my knowledge to her as she learns to take over Ike's training.

Monkey D

I worked with Monkey D while in 2014 and 2015. He was going to serve as my second monkey for collecting neural data during the CST. Therefore, I needed to train him to perform center-out reaches, the CST using visual feedback, and the CST using vibrotactile feedback. I started with Monkey D when he first entered the lab, and spent the first few months with him

getting him used to human interaction and poling/chairing. Monkey D was not like any of the monkeys I had worked with in undergraduate and graduate school; he was very timid and slow. He was not food motivated or water motivated, and really needed to learn things at his own pace. More than ever before, I found myself having to step back and try to understand what Monkey D thought I was asking him to do. It took a long time to get him making center-out reaches with straight trajectories. When I started to train him on the CST, it became clear that we would need to make some changes to the software in order to get Monkey D to understand what we were asking him to do. During the CST only the unstable cursor is visible, and therefore Monkey D was not able to understand that he was still in control. Rather than making the necessary changes to the code, Monkey D was transferred to a project where he would only be required to do center-out reaches and I returned to working with Monkey I.

8.2.2 Headpost and halo evolution

Chapter 3 details the current head fixation technique used in the lab for behavioral experiments. I joined the Batista lab in 2012, and during my time in the lab, I have fine-tuned the surgery we perform to attach the halos to each animal's skull, and redesigned the halo that attaches the monkey's headposts to the chair.

The Batista lab has been using three-short posts positioned around the top of the skull to be used with a halo since its inception in 2007. The first headpost design contained a titanium base with a shaft that screwed into the base and a ball that screwed into the shaft. During implantation surgery, small incisions were made at the location of each headpost. Once each headpost was firmly attached to the skull, the incision was closed up with sutures meeting the headpost at each side. This occasionally led to skin recession at the base of each headpost, and to

the animals picking at sutures and introducing bacteria around the post. Additionally, the fact that each headpost was made out of three separate components lead to some animals screwing off the shaft and thus requiring headpost repairs.

Currently, headposts are created from one piece of commercially pure titanium to increase the strength and prevent animals from needing repairs. The biggest impact I have made on headpost implantation is a small change to our surgical procedure. Rather than creating three incisions, one for each headpost, and closing the skin around the headpost shaft, we now use two coronal incisions, an interior incision about two inches from the brow ridge and a posterior incision that spans from ear bar to ear bar. This is explained fully in Chapter 3. Before closing these incisions, headposts are poked through to the exterior of the skin using "the t-shirt method," where a small incision is made directly over the ball of the post, just large enough for the post to be inserted through to the exterior of the skin. This allows the skin to heal directly to the post and prevents the animal from putting its fingers into the area around the post. This minimizes the chance of infection and recession of the tissue around the post.

The halo has also changed since I joined the lab in 2012. When I joined the lab, we were using the halo shown in Figure 8.5, rather than the halo we currently use shown in Chapter 3, Figures 3.2 and 3.3. This halo had a thickness of 0.25 in., outer diameter of 6 in. and inner dimensions measuring 4.375/3.875 in. (major axis/minor axis). The halo has one rigid offset post at the front that attaches to the brow ridge headpost, and two sliding offset posts to adjust to multiple headpost angle combinations (see Figure 8.5b,c). These channels are approximately 2.5 in. long, but each of these original halos was made by hand, and therefore lengths were not standardized. Each offset post consists of a stainless steel block and an adjustable socket screw.

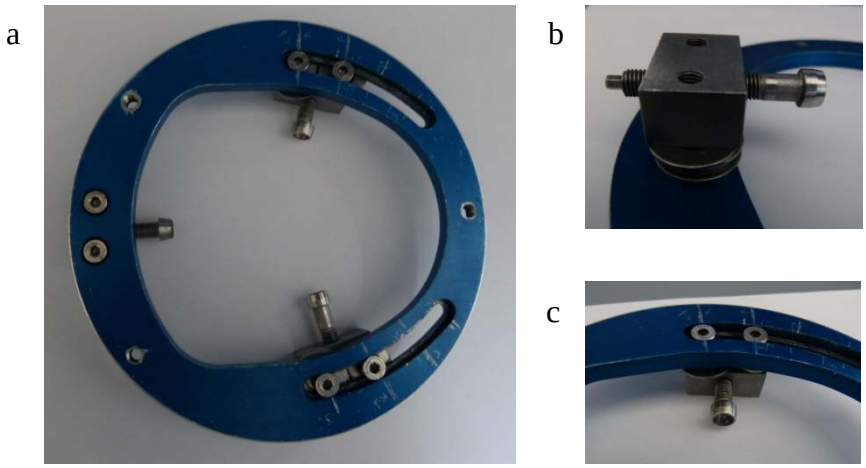

Figure 8.5: Old halo and offset posts. a) Halo with offset posts and holder attachment holes. b) Fixed front offset post. c) Sliding side offset post.

I performed my first headpost surgery on Monkey G, and since we didn't know where he would be getting a neural implant down the line, we placed his headposts in such a way that yielded the most real estate on the skull for later surgeries. I had to use a larger halo and halo holder stabilize Monkey G's head for behavioral experiments. This halo was larger than the one we were using at the time (7 in. outer diameter), and had a completely circular design providing more space on the inside of the halo to attach any neural recording hardware.

As new monkeys received headposts using the updated surgical procedure, more animals required the larger halo for head fixation. And even when animals could use the smaller halos, the graduate students and post docs in the lab noticed the additional room that the larger halo allotted for recording hardware. Thus, in January 2016 this larger halo became the only halo used for the animals in our recording rig. However, while this halo worked for all the animals in the recording rig at the time, I was planning to begin recording from Monkey I within the next few months. Monkey I had received his headposts before I entered the lab in 2012, and had a halo that was customized for him. The larger halo did not fit his posts, which lead me to design a halo

that could be adjusted for any headpost configuration. This halo is discussed in chapter 3, and then again below.

The machinist for the Systems Neuroscience Institute, Kevin Thiel, and I built a three-dimensional model of the halo frame was created using Autodesk Inventor software. I transferred this inventor file into Autodesk CAD and ordered one test halo from eMachineShop. The halo has a thickness of 0.25 inches, inner diameter of 5 inches, and outer diameter of 6.75 inches (see Chapter 3, Figures 3.2 and 3.3). Just like the old halo, this halo has one rigid offset post at the front that attaches to the brow ridge headpost, and two sliding offset posts to adjust to multiple headpost angle combinations. These channels are 3.846 in. long and are much longer than those in the old halos, and are a bit wider so that the offset posts slide more easily. Each offset post consists of a stainless steel block and an adjustable socket screw. For the new halo, we have two versions of the offset screws, one where the socket screw is perpendicular to the block, and one where it goes through the block at an angle. The angled offset blocks were made in case we place headposts on the frontal skull region, beyond the scope of the channels for the sliding offset posts. These changes will allow the new halos to be used for any headpost configuration, and we will be ordering more in the next few months so that each animal has a designated halo.

8.2.3 Rig rebuild

In order for a research lab to function successfully, data collection must be reliable. In the Batista lab, we have two experimental rigs; one rig is used primarily for behavioral experiments and animal training, while the other has a Tucker-Davis Technologies RZ2 system for collecting neural data and performing brain control experiments. Henceforth, these rigs will be referred to as the behavioral rig and the neural rig respectively. Each rig is comprised of a combination of

third party hardware and software for data acquisition, a custom LabVIEW-based (National Instruments Inc., Austin, TX) software package for system control, and a 3D virtual reality (VR) display and vibrotactile feedback device. Data acquisition sub-systems record and transmit neural ensemble data, limb position, eye position, and/or other general-purpose data to a central control unit consisting of real-time (RT) controller and Host PCs. The central control unit synchronizes and processes incoming data and generates the appropriate output commands to the VR and vibrotactile displays as dictated by the experimental paradigm design. The rigs were built in 2007 when Dr. Batista started his lab, and each component is described in Mr. Daniel Bacher's Master of Science dissertation (Bacher 2009).

By 2015, the rigs were in desperate need of both hardware and software updates. The original rigs used off-the-shelf PCs and had begun to slow down causing software lags. Additionally, the custom LabVIEW data acquisition software, called Monkey Host, was crashing regularly leading to unreliable data collection. Over the years, multiple lab members ranging from experienced programmers to those with nontechnical backgrounds, have designed and implemented experimental paradigms from scratch into Monkey Host without having an in depth knowledge of the underlying code that makes up the software. This led to messy code and difficulty when trying to find the source of the software crashes.

Nick Pavlovsky and I took charge in organizing a complete overhaul of the rigs, starting with the neural rig. His focus was on getting Monkey Host in working order while mine was on purchasing and building new computers for the rig and making sure all the hardware was wired properly for data acquisition. I decided to customize and build each PC, so that in the future, parts could be replaced rather than performing another complete overhaul. The parts required for each machine are documented on our lab's server.

To make the transition to a new rig as seamless as possible, we housed the new computers and hardware in a development rig that enabled software development while experiments continued in the existing neural rig. Nick worked closely with Nicole McClain and Alan Degenhart to make the major structural changes to the software, and once Monkey Host was working in the development rig, Nick and I moved the new system into the experimental rig. After physically moving the new computers and hardware into the rig, I updated the wiring and prepared the rig for testing. Wiring diagrams for both the neural and the behavioral rigs are on our lab's server.

After successfully updating the neural rig, I took the lead on updating the behavioral rig. At the time, my plan was to collect thesis data in the behavioral rig, and wanted a rig I could trust. For this update, we did not create a development rig, as that the new software code was tested in the neural rig and designed to be compatible for both the neural rig and the behavioral rig. Nick and Nicole incorporated "configuration files" that would initialize Monkey Host based on the hardware available. For example, Monkey Host would look for the neural acquisition PC in the neural rig but not in the behavioral rig. I built new computers using the same specifications that were used for the neural rig, rewired the connections, and tested the software.

Once both rigs were running the new software code running reliably on the new machines, I went through the software to create a Monkey Host user manual that will act as a reference for new lab members. This user manual is not intended to replace training with senior graduate students, but is meant to be a resource with answers to some of the common questions that arise and some troubleshooting solutions. I was the lead author of this manual, but received input from Nick, Nicole, Alan and Erinn Grigsby during its completion. This manual is included in my thesis in Appendix A.

8.3 MENTORING AND OUTREACH

A substantial part of my graduate student career was spent mentoring undergraduates, medical students and junior graduate students. Through these experiences, I went from being anxious and unsure of my teaching abilities to actually enjoying teaching and helping younger students find their place in science. During my time in the Batista lab, I mentored: Nathan Flemming (undergraduate), Matthew Ginsberg (medical student), Erinn Grigsby, and Nicole McClain (graduate students). Nicole McClain will be continuing my project, and will use the critical stability task to explore the neural basis of long-duration, feedback driven movements. It has been such a pleasure to work with such motivated, hard-working individuals, and to watch them work towards their goals.

I was also a volunteer member of the CNUP's Brain Program. The Brain Program is a group of volunteers who go into middle schools and high schools to give short interactive presentations about the brain. Topics ranged from the organization and function of the brain to what happens to the brain in disease states. I helped lead the neural prosthetics and motor control station. At the end of the program, we talked with the students about how we reached our current career path. Through this experience I have had the pleasure of teaching kids about the benefits of melding neuroscience with engineering, and inspiring them to continue learning.

I also volunteered to judge multiple science fairs and poster competitions throughout my time at the University of Pittsburgh. Investing Now is a program for high school kids during the summer where graduate students lead classes in bioengineering. At the end of their class, teams of sophomore students completed science projects where they tackled one of today’s medical problems using the bioengineering techniques they learned throughout the summer. They created and presented scientific posters detailing the symptoms and scope of their medical condition,

current solutions or therapies, their novel solution to the problem, and a conclusion. I judged these posters, asking questions that encouraged students to think deeply and creatively, and used this opportunity to talk to the students about college plans and my path to higher education in bioengineering. In addition to judging posters, I also volunteered to judge the Pittsburgh Regional Science and Engineering Fair (PRSEF). I interviewed students at the fair and evaluated their work. Many of the students in the Junior (6^{th} grade) or Intermediate (7^{th}-8^{th} grade) Divisions are "doing" science for the first time, and as a judge I got to encourage, teach and inspire these young people as they take early steps into science and technology. The Senior Division was composed of projects from high-school students. The work was high quality, and led to questions that not only expanded my own mind, but also may have helped develop a young scientist's idea develop into something larger. This experience taught me the importance of being encouraging and constructive. I wanted to motivate these students to advance in science, but also help them realize the shortcomings and possible alternatives of their work. Being able to look critically at my own work has stemmed from my experiences judging and has made me a better scientist.

9.0 CONCLUSIONS AND FUTURE DIRECTIONS

Brain-computer interface technology has given researchers and clinicians the hope of providing individuals with severe motor impairments the ability to interact with the environment. However, BCI systems are still far from being a usable clinical therapy. To bring BCIs to a clinical setting, we need to make improvements in a diversity of scientific areas. The work in this thesis begins to address the translation of BCI systems from the primate lab to a clinical setting.

Laboratory work must continue for BCIs to improve and move towards clinical usability. Chapter 3 discusses a method for safe and durable head immobilization during animal experimentation. To date, the most widely used head immobilization method in the neuroscience community involves the construction of an acrylic headcap, housing the headpost and recording chambers. Dental acrylic is osteophobic and can cause bone softening and necrosis, leading to an eventual failure of the entire implant. Infection of the skin-implant barrier can also lead to eventual failure. Spontaneous failure of such implants disrupts data collection and occasionally the bone cannot support a second implant and end data collection ceases entirely. Our method of head fixation circumvents these problems, yielding reliable and consistent data which in turn leads to BCI advancements.

BCI systems depend on sensors and associated hardware to acquire signals from the brain. Regardless of which recording modality is chosen, understanding the impact of such a method on the brain and how we can extract maximal information for control will be important

for bringing any BCI system to clinic. The work presented in Chapters 6 and 7 addresses BCI signal acquisition; in Chapter 6 we look at the long term recording abilities of an ECoG array and determine the impact on the brain histologically, and in Chapter 7 we use an intracortical electrode array to explore the optimal way to for extracting information from extracellular recordings.

The development of BCIs for people with disabilities requires clear validation of their real-life value in terms of efficacy, practicality and impact on quality of life. Two major limitations in the usability of BCIs are the focus of today's BCI research on making point-to-point reaches and almost exclusive reliance on visual feedback for control. Natural movements often require that we tailor our actions to the environment in an ongoing manner, utilizing feedback to make successful movements. The absence of non-visual feedback modalities forces users to maintain visual contact while performing tasks, and is especially problematic when visual attention is diverted and in tasks that do not require visual feedback, such as grasping. Chapters 4, 5, and 8.1.1 utilize a novel animal behavioral task that involves prolonged motor actions for which continuous sensory feedback is essential. Using this task, we evaluate the usefulness of different feedback modalities for continuous motor action and propose to probe the motor system during long-duration, feedback driven movements.

The CST complements classic point-to-point reaching paradigms in several ways. First, it allows one to probe the neurophysiological mechanisms of continuous, feedback-driven movements, by making neural recordings in animals performing the CST. Second, it provides a way to quantitatively compare the effectiveness of different sensory feedback modalities or control methods. Third, the implementation of the CST is fairly straightforward such that

different laboratories should be able to conduct experiments for which the results can be effectively compared.

The CST enables researchers to probe the neurophysiological mechanisms of continuous, feedback-driven movements. In Chapter 8.1.1, I propose to find evidence of sensory information in primary motor cortex (M1) when sensory feedback is essential for behavior. Even if sensory information is not found in M1, does what we know about M1 from point-to-point reaches fully describe how M1 neurons behave during continuous, long-duration, feedback driven movements? It has been shown that the activity of individual M1 neurons is tightly correlated with the velocity of the hand and that the activity of a population of M1 neurons can predict the details of a hand movement (Georgopoulos et al., 1982; Moran and Schwartz, 1999; Taylor et al., 2002). Can we use preferred directions to recreate continuous movements? With neural data collected during the CST, we will be able to answer this question. For the studies proposed in Chapter 8.1.1, animals will perform center-out reaches and CST in the same experimental sessions. Using the preferred directions calculated from a subset of the center-out reaches, we can use a PVA decoder offline to recreate trajectories for both center-out and CST trials. I hypothesize that the decoder will successfully recreate center-out reach trajectories, but will not be able to reconstruct the hand trajectories for CST trials. Neural preferred directions are calculated from a task where motor action and sensory feedback are congruent, and therefore even if cells are encoding both sensory and motor parameters, the PVA decoder will not fault. In the CST, visual feedback and motor action oppose each other, and therefore any response to feedback may hinder decoder performance. It is possible that a PVA decoder will accurately reconstruct CST hand trajectories. This is also an exciting result because it would demonstrate

the robust nature of preferred directions in M1 and show that our current understanding of how M1 controls voluntary movement extends beyond point-to-point reaches.

The CST provides a way to quantitatively compare the effectiveness of virtually any effector-feedback pairing. In our experiments, the monkeys received either visual or vibrotactile feedback and used either hand control or brain control. The feedback signals could also have been other vibrotactile feedback methods, intracortical microstimulation, or abstract visual signals and the effector could have also been pinch force, a different decoding algorithm, saccades or different neural recording technique. All of these sensorimotor pairings could be assessed using the same critical instability metric, the critical instability value (CIV). Along these lines, the CST can to assess decoder effectiveness as we work towards clinically viable BCI systems. With point-to-point reaches, various decoders show minimal differences in performance, with all decoders achieving high success rates. For a BCI system to be clinically useful, patients must be able to perform everyday tasks that require interactions with the environment and corrective actions. Current decoding algorithms, such as PVA and the Kalman filter, have difficulty stopping and changing directions (Golub et al., 2014). Point-to-point reaches only require an effector, such as a cursor or robotic limb, to move in one direction for a short period of time, and thus these decoders perform well on such tasks. By contrast, the CST requires the system input to change directions quickly in response to feedback in order to keep the cursor in-bounds. In Chapter 4, a monkey used a velocity Kalman filter for BCI control to perform the CST, which is the most widely used BCI decoding algorithm for reaching tasks, yet was only 22.7% as effective as when he used hand control for the CST. This example demonstrates that there is still room to design better decoders, and that the CST can complement traditional BCI tasks as we move towards clinical viability.

This thesis utilizes a one dimensional, first order implementation of the CST, however the CST can be adapted for increasing physical complexity to be added as desired (e.g., (Hogan and Sternad 2012; Kadkade et al. 2003; Nasseroleslami et al. 2014)). For example, it is straightforward to implement a 2D version of the CST. Likewise, one can increase the complexity of the system dynamics (and hence the sensorimotor demands) by implementing the differential equation,

$$\frac{d^n x(t)}{dt^n} = \lambda^n\big(x(t) + u(t)\big),\ \lambda > 0,\ t \geq 0 \qquad (9.1)$$

where n = 1, 2, … is the order of the system. For example, n = 2 results in a system with second-order dynamics analogous to the classic inverted pendulum problem (Kadkade et al. 2003). In another example, the CST can be readily implemented as a tracking task (the critical tracking task (CTT)) by changing the polarity on the motor command signal. Alternatively, noise could also be injected into the system instead of relying on the monkey's motor noise to generate instabilities. This flexibility of the CST coupled with a sparse parameter set to quantify behavioral performance and its amenability to closed-loop control systems analysis makes it a powerful paradigm for studying continuous sensory feedback in motor control.

Additionally, the linear equation used to model control during the CST was chosen for simplicity and assumes the subject only has information about the position of the cursor to inform his hand movements. This model could be expanded to include other kinematic parameters such as velocity and acceleration. Additionally, rather than assuming proportional control, the model can be expanded to include integral and derivative terms. The integral term accounts for past system positions and enables control based on previous knowledge of system dynamics, and the derivative term accounts for possible future trends in the system position. The animal may also utilize different control strategies for different methods of feedback or control.

By exploring more complex control models, we may be able to more accurately describe the animals control strategy given specific control and feedback modalities and discover trends that lead to improved sensorimotor control.

APPENDIX A

MONKEY HOST USER MANUAL

Monkey Host is the Batista lab's custom LabVIEW software for task control. Through a series of intuitive windows, the researcher is able to easily design behavioral control paradigms by constructing and modifying state tables. Other parameters are set through a series of pop-up windows. Stimulus timing, location, and appearance can all be dynamically adjusted on the fly during experiments through the graphical user interface (GUI). The GUI also provides real-time plots of behavioral data, including hand position, eye position, in addition to displaying state parameters, and other trial information. This user manual details the step-by-step procedure for using the Monkey Host software. It describes how to run already defined experiments as well as create new experimental paradigms and troubleshoot some of the more common errors. I wrote this document with input from Dr. Alan Degenhart, Nick Pavlovsky, Nicole McClain, and Erinn Grigsby.

A.1 GETTING STARTED

Before starting the Monkey Host application on the Host computer, there are a few things that you need to do.

1. Make sure LabVIEW is open on all computers, except Host.
2. Make sure that the RT computer is running.
3. Turn on the PhaseSpace system.
    a. Turn the KVM switch to the PhaseSpace computer.
    b. Click on the ‘Connect’ button in the upper left corner of the main screen.

A.2 STARTING MONKEY HOST

Now we can get started with running the Monkey Host Application.

1. Open the Monkey Host Application from the desktop of the Host computer.
2. A file window will pop up asking to select a Configuration file.
    a. Select the Configuration file that you need from this list.
3. Wait for all windows to pop up on every computer.
    a. IT IS VERY IMPORTANT THAT YOU DO NOT CLICK ANYTHING UNTIL ALL WINDOWS ARE OPENED ON ALL COMPUTERS.
    a. GPU, TDT, and Internet computers all have a window that should open.
4. Once all windows have opened, load a parameter file into Monkey Host.
    a. Select Load Parameter File from the File menu on Monkey Host Main.
    b. Choose the appropriate parameter file for your task.

A.3 STARTING TDT TO RECORD NEURAL DATA

If you are recording neurons, you will need to have the TDT hardware and software running.

1. Turn on the RZ2 base-station.
2. Open OpenProject on the TDT computer.
3. Load the project associated with your animal from the drop-down menu.
4. Create a new data tank to save your data in.
    a. Go to File -> Data Tank Management
    b. Click the button that looks like a box with a sunburst on the line that says TANK:.
    c. Name your Data Tank in the common convention, and click OK.
    d. Right click the newly created data tank, and then click Register Tank.
    e. Click OK to finish creating your new data tank.
5. Load your RCO circuit in the RPvdsEx tab.
6. Edit which broadband channels you are recording from in the RCO circuit.
7. Load your controller file in the Controller tab.
    a. Click the exclamation point (!) on the toolbar near the upper left corner.
    b. Select your controller file.
8. Once the controller file has loaded, save it and rename it.
    a. To save, click the button with a floppy disk and magenta arrow pointing at the disk.
9. Click the workbench tab, and you are now ready to begin recording from your animal.

A.4 RUNNING TRIALS

Now you are ready to begin recording.

1. On Monkey Host Main, fill out the data logging information.
    a. Check to make sure “Subject Name” is correct. This field is saved/loaded when you save/load a parameter file.
    b. Select a destination to save data. Click the folder icon next to the “Trial Logging Directory” and choose where you would like to save your data.
    c. Select the red indicator labeled “Log Data?” This indicator should turn green, signifying that data will be saved.
2. Reset trial number. From the System menu, select Reset Trial Number and specify the number for your first trial.
3. Reset the trial success counter. From the Trials menu, select Reset Trial Success Counter
4. Run Trials. From the System menu, select Run Trials.

Now your trials are running. To stop trials, select the Stop Trials from the System menu on Monkey Host Main.

A.5 MENU ITEMS

A.5.1 File

Load System Parameter File - loads a parameter (.prm) file into Monkey Host

Save Current System Parameters - saves current parameters to .prm file

Save Default System Parameters File - host loads default parameters during initialization

Exit - closes Monkey Host

A.5.2 Config

Load System Configuration File - loads a configuration (.cfg) file into Monkey Host. Configuration parameters include: IP addresses for each computer, important file paths, whether or not to use simulators, whether or not using TDT

Edit System Configuration Parameters - alter current configuration parameters

Save Current System Configuration - saves current configuration parameter to .cfg file

Set Default Configuration File - set default pathway for configuration file

A.5.3 Intervals

Edit Intertrial Interval - changes the length of the intertrial interval and specify how many possible ITIs

Edit Intervals - allows the user to specify/modify task variables used in the state table

A.5.4 Targets

Edit Tubes - loads tube files from matlab to be used in the state table and specifies the rotation

Edit Tube Windows - specifies the tolerance windows to achieve tube conditions

Edit Targets - specifies targets to be used in the state table. This VI does not impact target presentation, but specifies the target acceptance windows.

Edit Target Limits - specifies where targets can be rendered in VR.

Edit Target Window Tolerance - This value specifies the maximum number of consecutive points that must lie outside a target window for a failure to occur. If the marker is jumpy this should be increased.

A.5.5 Trials

Edit Trials - opens the list of trial files and specifies how often a certain trial occurs as well as other trial parameters. The “View/Edit” trials button opens the state table for the selected trial

Edit Maximum Trial Failures - specifies the maximum number of consecutive trial failures that will occur before trials are stopped from running

Edit Trial Randomness - switch either executes trials in order or randomizes trials by specified percentages in the “Edit Trials” menu

Edit State Table - opens the state table (.trl) for current trial

Display/Hide Trial Successes - displays or hides the success count for each trial in the status window

Reset Trial Success Counter - sets the trial success counter back to zero

A.5.6 Signals

Several fields in this menu remain from Raj’s software. Only relevant menu items are described.

Edit TTL Outputs - allows you to change the probability of each TTL output

Edit Motor Parameters - adjusts the voltage output waveform to vibration motors

A.5.7 Systems

Run/Stop Trials - starts/stops trials.

Reset Trial Number - resets the current trial number to the value specified.

Use 2-D Gaze Calibration - changes the default 3D eye calibration window to a 2D calibration window. This calibration allows eye tracker data to be viewed/used in Monkey Host.

Show/Hide Calibration Window - show or minimize window.

Show/Hide Current Parameters Window - show or minimize window.

Show/Hide Online Analysis Window - show or minimize window.

Edit Pushbutton Parameters - specifies the interval of the manual reward output and the TTL output to increment when the manual reward pushbutton is depressed .

Edit Auto-Monkey Speed Profiles - changes the speed profile for the cursor during the Auto-Monkey task. For example, the cursor velocity can be constant throughout the “reach” or have a bell-shaped velocity curve to replicate real reaches.

Edit Sounds - switches between a success tone and a failure tone.

Edit VR Parameters - specifies the settings for the VR environment. This menu item is not used much, as that it stays rather constant across tasks, however is useful for changing the placement of the photodiode. There is also a button that says “Set Bounds to GPU Resolution” that resizes the visual environment to fit the monitors in the rig room.

Edit Phasespace-VR Transformation Parameters - transforms Phasespace coordinates to VR coordinates for task control. This VI is used to calibrate the signal recorded from the animal’s arm position so that it can be used in the VR environment.

Edit Phasespace Parameters - sets the Phasespace parameters, such as IP address, number of markers, and frequency.

Edit Wii-Phasespace Calibration - transforms Wii-mote coordinates to Phasespace coordinates for task control

Edit Z-Lock Parameters - locks the z-position of the cursor so that the position of the animals hand does not cause the visual cursor to change planes. For example, if a target is cylindrical at Z=750 with depth 20. Setting the Z-Lock to 800 would ensure that the cursor was in front of the target, a Z-Lock of 750 would allow the cursor to go into the target, and a Z-Lock of 700 would mean the cursor would no longer be visible when it intersects with the target, as that it is behind it (good for making an occluder).

Edit Visual Lag - adds a visual lag to task events. When set to random, it will chose random visual delays between the specified minimum and maximum lags. When not random, it implements the specified visual lag.

Edit Markers - specifies the markers that are used in a task and what channel they correspond to. This is most commonly used to change the phasespace marker number that is being used for control.

Edit Pseudorandomness - allows user to specify how targets are chosen from the targets list.

Edit Marker Jitter - adjusts the cursor jitter. This is commonly used during observation training for BCI control so that the animal thinks he is controlling a cursor during observation and pays attention.

Load/Switch Decoder - allows the user to load a set of decoding parameters into memory. Also allows assist levels and visuomotor rotation to be specified.

Auto-Load Decoder - allows decoding parameters to be automatically loaded into MonkeyHost. This monitors a specified folder for new decoding parameters (saved as *.mat files), and

loads them into memory, incrementing the block number when the decoder is updated. Primarily used during gradual decoder calibration.

Decoder Update Scheduler - automatically buffers data and calls MATLAB scripts on a specified interval. Used for executing MATLAB code on regular intervals, particularly for purposes of automatically calibrating decoders.

Load/Switch Force Parameters - specifies the force parameters for when pinch force is used in task control, and specifies the Lambda Value for the CST task (error cursor parameters).

Edit TDT Parameters - specifies the number of channels, the brain area, and IP address for TDT.

A.5.8 Plotting

Edit Plot Colors - changes the display colors for task parameters in the XY plot and the 3D plot.

Edit Phasespace/Decoder Channel for Plots - specifies the phasespace or decoder to be displayed on XY plot and 3D plots. Currently there is only one option for each channel.

Edit Trace Plot Initial Step - specifies the first state table step to be displayed onto the trace plot

Edit Trace Plot Offsets - specifies the offsets applied to the signals on the trace plot (The specified values seem to be left from Raj's software)

Edit Trace Plot Axis Limits - specifies the maximum and minimum axis values for the trace plot

Edit XY Plot Display Data - enables the experimenter to turn on/off data channels to be displayed on the XY plot

Edit XY Plot Axis Limits - specifies the x-axis origin, y-axis origin, and axis length for the XY plot

Edit 3D Plot Limits - specifies the minimum and maximum axis values for the 3D plot

Edit Displayed TTL Output Count - displays the TTL output count on the trace and XY plots - note: also changes the reward sound

Edit PSTH/ISIH Plot - specifies parameters for the PSTH and Raster window. Enables the experimenter to change the initial PSTH step, the PSTH channel and sort numbers, the PSTH bin width. and the number of ISIH intervals. There is also the option to have the trials averaged.

A.5.9 Analysis

** This is a menu that remained from Raj's Lab when we initially created/modified this software for our experiments. To my knowledge, it has not been used by anyone in the lab and can be removed or changed for the next Monkey Host iteration*

A.6 SAVING PARAMETER FILES

Changes made during individual training sessions (such as interval lengths, plotting specs, target locations, etc) are not saved automatically. To save changes to the current parameter file, select Save Current System Parameters from the File menu, and save the new parameter file to the desired location. It is good practice to change the date and save a new parameter file when substantial changes are made.

A.7 EDITING PHASESPACE – VR TRANSFORMATION PARAMETERS

To enable task control using the PhaseSpace motion tracking system, we create a mapping between PhaseSpace coordinates and units in our virtual reality (VR) environment. We adjust this mapping to accommodate the arm span of individual animals so that their reaching arm is visible by the cameras at all locations in the workspace.

Select Edit Phasespace-VR Transformation Parameters from the Systems menu.

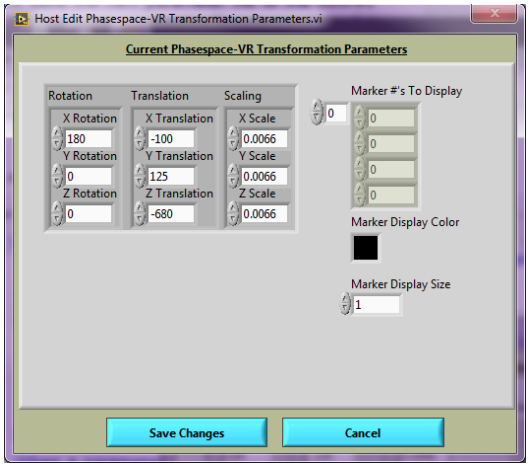

Order of operations:

1. Rotation about z axis
2. Rotation about y axis
3. Rotation about x axis
4. Translation
5. Scaling

Rotation: Angle of rotation about each of the x, y, and z axes. Tasks generally use a 180° rotation about the x axis to correct for the phasespace system's definition of positive y as downward and positive z as toward the wall. Alternatively, some tasks use a 180° rotation about the z axis and a negative scale factor for the x axis.

Translation: Corrects for the offset of the phasespace origin in the VR space.

Scaling: Scaling relates to how arm movements in space are shown on the screen. Tasks we commonly use have utilize a 0.0066-0.0076 scale. A negative scale flips the direction of the corresponding axis.

A.8 EDITING TRIALS IN THE STATE TABLE

When a parameter file is loaded, it loads a set of trial files that control the progression of the task. Loaded trial files are executed according to the set probability, and often only one trial file is used at a time. Changing trial parameters such as intervals (reward length, hold times, etc) can be done without accessing the state tables but changes to the task itself (order of states, visually rendered objects, target locations, etc) must occur within the state table and saved as a trial file.

1. Select Edit Trials from the Trials menu.
2. Select the trial you are interested in changing from the Current Trials list so that the file path appears in the Trial Filepath box in the top right corner of the Edit Trials window.

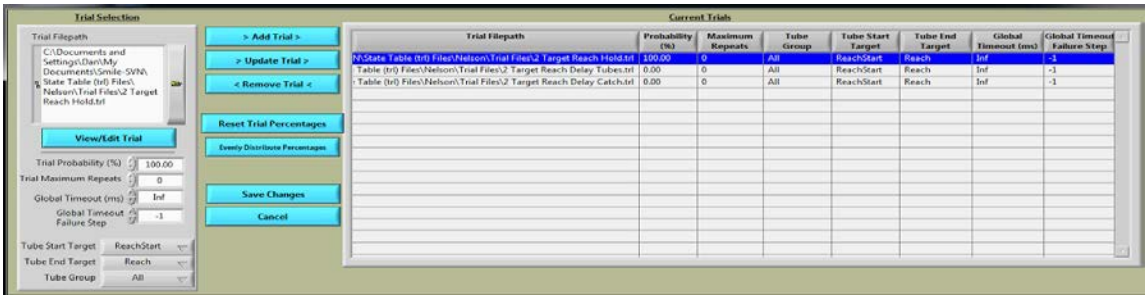

3. Hit the View/Edit Trials button underneath the Trial Filepath box to open the .trl file in the state table editor.

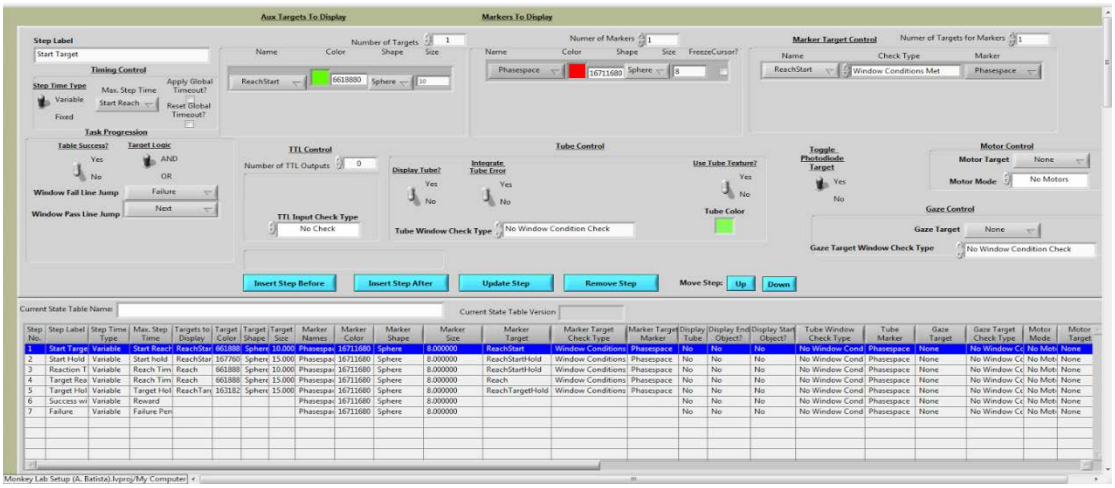

4. Click on a state to change the state parameters.

   Note: Changing the size of the targets in the state table does not change the size of the tolerance window. The tolerance window size must be changed in the Edit Targets menu (see Editing Targets section).

5. Click Update Step before selecting another state to save changed parameters for the edited state
6. After all state changes have been made, save the .trl file by clicking File > Save to save over loaded .trl file or File > Save As to save as new .trl file. Recommended practice is to save a new state table file with a descriptive name and the date of creation.
7. Exit window by clicking File > Exit

A.9 EDITING TARGETS

Target presentation is controlled in two places, and therefore both saved in the parameter file as well as in individual trial files. The boundaries and location for targets are specified by the parameter file, whereas target presentation is controlled by trial files.

A.9.1 Control of Target Location and Boundaries

Changing the boundaries and the location of targets for the task occurs in the parameter file.

1. Select Edit Targets from the Targets menu.

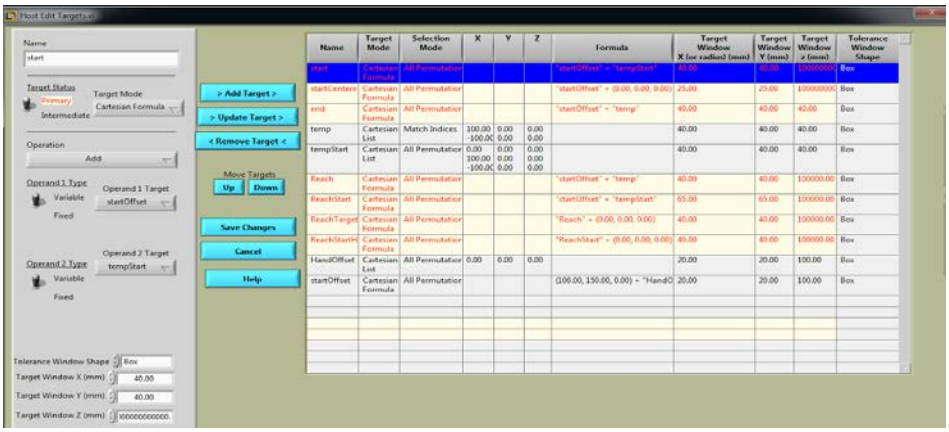

2. Target Status:
    a. Primary targets (red) are used within the state table for task control
    b. Intermediate targets (black) are used as variables for primary targets
3. Target Mode:
    a. Cartesian list - targets can be specified using x, y, and z coordinates
        i. Target Selection Mode allows the user specify if targets include all permutations of indices, or if indices should remain matched
    b. Cartesian formula - select the “help” button for information on how specify target locations based on other targets. “Add” is the most common operator, as that it enables you to specify targets based on a center location. This is equivalent to vector addition.
4. Target Tolerance Windows
    a. Tolerance window shape - Specify the shape of the tolerance window (NOT THE PRESENTED TARGET)
    b. Radius (or Target Window X and Target Window Y) - this sets the boundary value such that once the center of the cursor (phasespace or neural) crosses this boundary the target has been “acquired”

c. Target Window Z - this boundary is specified for 3D tasks. When we work in 2D, this value is made very large (100,000 or so).

5. ‘start’ and ‘end’ targets are present in most parameter files
    a. ‘start’ denotes the center of the workspace, and since it is used to define several targets, adjusting its position will effectively move all targets.
    b. ‘end’ is not used often in the state table, but is a cartesian formula that combines the ‘start’ target and some intermediate target
6. Reaching tasks commonly include primary targets ‘Start Target’ and ‘Reach Target’ and an intermediate target that defines the reach locations.
    a. The ‘Start Target’ is a cartesian formula adding the ‘start’ target (center of the work space) to a cartesian list. This can be (0,0,0) or be made to include an offset (such as intermediate targets that define handOffset or startOffset)
    b. The ‘Reach Target’ is a cartesian formula that adds the intermediate targets to the ‘Start Target’
    c. The intermediate target is comprised of a cartesian list that defines all possible reach locations. In a center out reaching task, this list contains a list of coordinates that define targets at a fixed radius from (0,0,0). When this is added to the ‘Start Target’, it defines reach targets to be equidistant from the center of the workspace.

A.9.2 Control of Target Presentation

Changing the target presentation for the task occurs in the trial file.

1. Select Edit Trials from the Trials menu.

2. Select the trial you are interested in changing from the Current Trials list so that the file path appears in the Trial Filepath box in the top right corner of the Edit Trials window

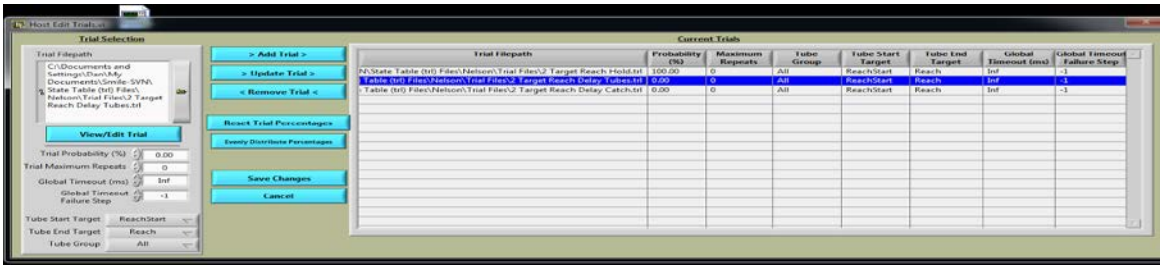

3. Hit the View/Edit Trials button underneath the Trial Filepath box to open the state table editor.

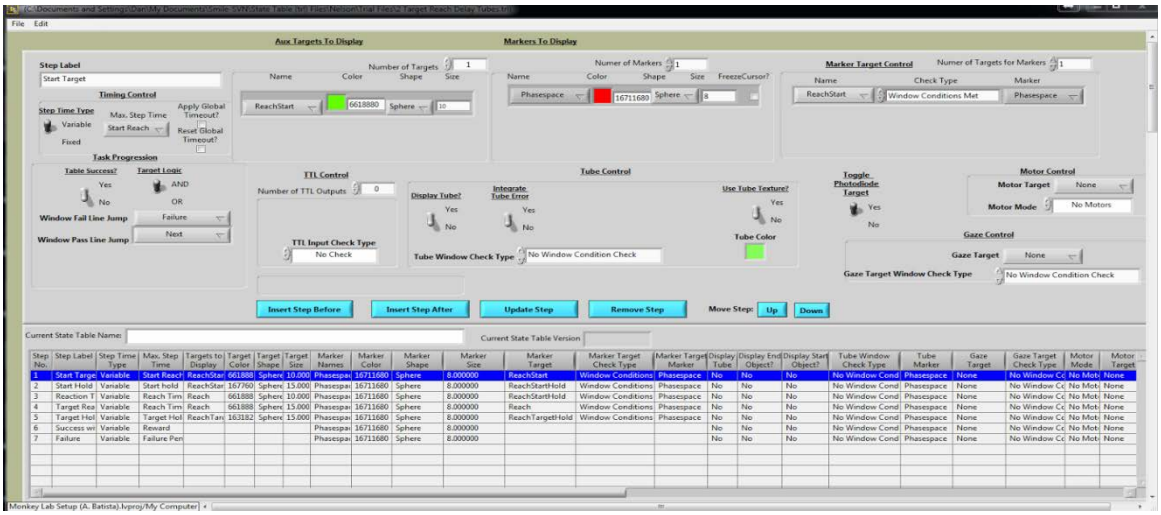

4. For each state, you can specify the number of targets that are displayed, the shape and color of each target, and the size of the targets. Specifying the target name from the drop-down menu specifies target location and boundaries by linking to the list of targets in the parameter file.

5. Once you have made changes to targets for a given state, click Update Step to save these changes before selecting the next state to edit.

A.10 USING THE PHOTODIODES

Photodiodes enable us to determine the delay between when the RT computer executes a state transition and when that state transition is presented to the animal. A photodiode is taped to the top-left corner of each monitor and is activated by a photodiode target, causing it to switch between 0V and 3.5-4V. This voltage is recorded as an analog signal, and processed offline to determine when the photodiode goes from the on-state to the off-state. The render delay is the difference between the RT state transition times and the photodiode transition times.

The photodiode target is a white square that activates the photodiode circuit. For the photodiodes to work properly, this target must be located in the top-left corner of the render window so that it lines up with the photodiode LED. To adjust the photodiode target location, select Edit VR Parameters from the System menu. The coordinates of the photodiode target are on the right. Changing these values and pressing the Save Changes button on the bottom of the window will move the photodiode target for the next trial. Once the photodiode target is in the correct location, save the parameter file so that the photodiodes are useful.

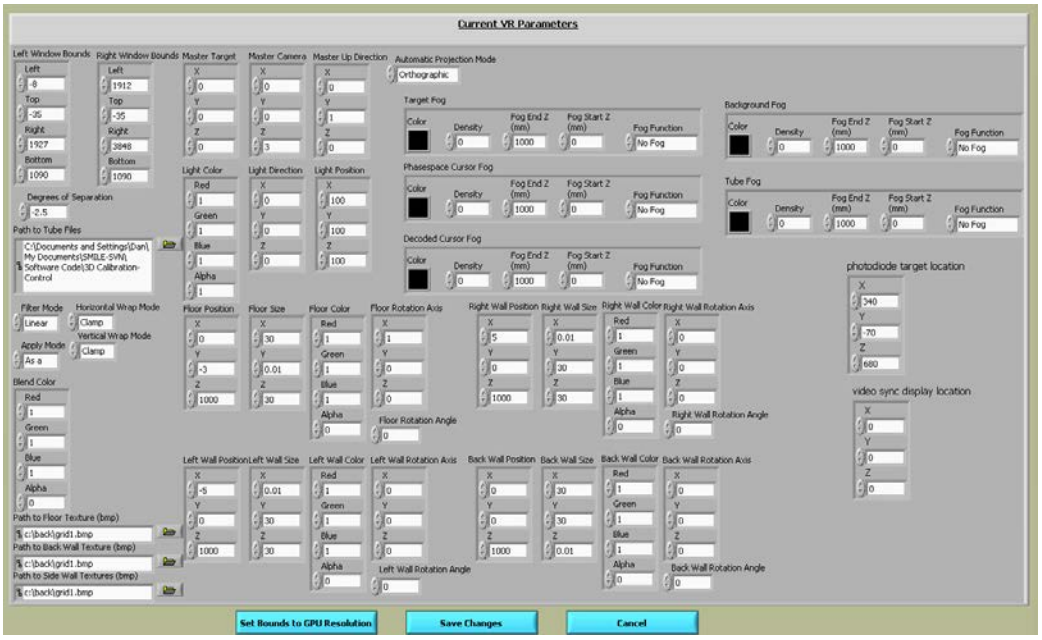

A.11 RESTARTING THE PHASESPACE COMPUTER

The Phasespace computer is not restarted often, but needs to be restarted after power outages and when the cameras are not being seen during calibration.

1. Press the “on/off” button on the PhaseSpace tower to turn off the Phasespace computer.
    a. Do not simply press “reset” as that this will get you stuck in the bios.
2. Wait a few seconds for the machine to shut down completely, then press the “on/off” button again to turn on the Phasespace computer.
3. Turn the KVM switch to the PhaseSpace computer.
4. In DOS mode, you will be prompted to input impulse login and password information. When you are typing in the password, no letters will appear on the screen.
    a. Login: demo
    b. Password: demo
5. Once you are logged into the impulse server, you will see “impulse:~$”
6. Type ‘startx’
7. A linux screen will open. Go to the Start menu and open a terminal
8. Type into terminal ‘cd phasespace’ and press enter.
9. Type ‘./master’. This will open the PhaseSpace master window.

A.12 CALIBRATING PHASESPACE

If the marker is dropping out or appears jumpy, recalibrating the PhaseSpace cameras can be a solution. Before recalibrating, check the marker location on the animal's hand to ensure that it is not being occluded during reaches.

A.12.1 Rig 1 Calibration

1. Turn off the large, green DAQ box by holding down of the on/off button for a couple seconds.
2. Disconnect the 2-pin white wire from the the small, square circuit board.
3. Place the DAQ onto the calibration wand. Plug the 2-pin connector into the small, square circuit board already on the calibration wand. The plug geometry on the wire should match up with the geometry on the circuit board.
4. Power on the DAQ back on by depressing the on/off button once.
5. Switch to the Phasespace computer on the KVM switch.
6. If the Master screen is on and 'Connected', disconnect the Master screen.
7. Open up the Calibration screen (open xterm; type 'cd /home/demo/phasespace'; type './calib'), click Advanced Calibration. Open the '/home/demo/phasespace/wand.rb' file. Hit the 'Connect' button.
    a. If the LEDs on the calibration wand come on, you're in business!
    b. If the LEDS do not come on, power off the DAQ and try flipping the orientations of the 2-pin connectors on the DAQ and/or small, square circuit board.

8. The six different camera views on the calibration screen should display the LEDS on the calibration wand. Hit 'Capture'.
    a. Note: If you hit 'Capture' and the six cameras do not show up, disconnect' and hit 'Connect' again.
9. In the monkey's workspace, hold the calibration wand vertical and move it around within the workspace so as to fill in the camera's view with green.
10. Once at least the center box of each camera's view is covered in with green, hit the 'Calibrate' button. If the calibration passes all four tests and the last test passes with a score less than 1.0, hit 'Save.'
11. Power off the green DAQ box. Unplug the 2-wire plug from the small, square circuit board attached to the wand. Remove the DAQ/wires from the calibration wand. Place DAQ back onto the rig. Restick the original small, square circuit board that was first attached to the velcro spot on the back of the DAQ. Plug the 2-wires from the DAQ into the appropriate end of the small square circuit board.
12. On the opposite end of the small, square circuit board, plug in the array of wires that contains wires with two blue, two pink, and two white flagged wires.
13. On the Phasespace computer, from the drop down menu, switch 'Calibration' to 'Alignment.'
14. Power DAQ back on. When on, the three LEDs attached to the rig should be on.
15. Under the 'Alignment Mode' drop down menu, choose first option 'Rigid Body.'
16. Upon selection, a menu will pop up. Select the file named '/home/demo/phasespace/Batista_Calibration_01_05_2010.rb' from the menu. If for some

reason you need to reselect the correct file, go to OWL<Change Trackers<Alignment Rigid to bring the menu back up again.

17. The three alignment makers on the rig should appear on the calibration screen. The images of the LEDs should be constantly on without much jitter. If there is a lot of jitter, then try small adjustments to the physical LEDs on the rig to see if they will stop the jitter in the Phasespace markers. You can try any and all trouble shooting methods. You can try recalibrating, sometimes that helps. The only requirement is that the backs of the LEDS should line up with black dots made on the rig.
18. Hit 'Align' when all the LED markers are stable. After hitting 'Align' the image should look upside down but comparable to what Rig 1 actually looks like (what you would normally see in the rig while sitting in the control room).
19. If the image looks good, then hit 'Save' and you are done with calibration. If the image does not look good, then depress the 'Align' button and start the calibration process again.
20. Once satisfied with calibration, disconnect the Calibration screen and reconnect the Master screen. If Master is not open: open xterm; type 'cd phasespace'; type './master'.
21. Power off the DAQ, unplug the array of wires that connected to the alignment markers. Instead, plug in the array of wires that allows the monkey gloves to be plugged in. Plug the ethernet power cord back into the DAQ box. Power DAQ box back on. You will know when you are successful if the LED marker on the glove turns on.
22. Test calibration with the Monkey Host program. If bad, recalibrate. Under 'Alignment', reselect the '/home/demo/phasespace/Batista_Calibration_01_05_2010.rb' file before pressing 'Align.'

A.12.2 Rig 2 Calibration

1. Turn off the large, green DAQ box by holding down of the on/off button for a couple seconds.
2. Disconnect the 2-pin white wire from the the small, square circuit board.
3. Place the DAQ onto the calibration wand. Plug the 2-pin connector into the small, square circuit board already on the calibration wand. The plug geometry on the wire should match up with the geometry on the circuit board.
4. Power on the DAQ back on by depressing the on/off button once.
5. Switch to the Phasespace computer on the KVM switch.
6. If the Master screen is on and 'Connected', disconnect the Master screen.
7. Open up the Calibration screen (open xterm; type 'cd /home/demo/phasespace'; type './calib'), click Advanced Calibration. Open the '/home/demo/phasespace/wand.rb' file. Hit the 'Connect' button.
    a. If the LEDs on the calibration wand come on, you're in business!
    b. If the LEDS do not come on, power off the DAQ and try flipping the orientations of the 2-pin connectors on the DAQ and/or small, square circuit board.
8. The six different camera views on the calibration screen should display the LEDS on the calibration wand. Hit 'Capture'.
    a. Note: If you hit 'Capture' and the six cameras do not show up, disconnect' and hit 'Connect' again.
9. In the monkey's workspace, hold the calibration wand vertical and move it around within the workspace so as to fill in the camera's view with green.

10. Once at least the center box of each camera's view is covered in with green, hit the 'Calibrate' button. If the calibration passes all four tests and the last test passes with a score less than 1.0, hit 'Save.'
11. Power off the green DAQ box. Unplug the 2-wire plug from the small, square circuit board attached to the wand. Remove the DAQ/wires from the calibration wand. Place DAQ back onto the rig. Restick the original small, square circuit board that was first attached to the velcro spot on the back of the DAQ. Plug the 2-wires from the DAQ into the appropriate end of the small square circuit board.
12. Attach the rigid body to the mirrors. This is the object that has what looks like 3 axes sticking out of it with LEDs on the ends of the axes. Place this on the left side of the mirrors in the small hole drilled into the frame. Make sure that it is stable and not moving around once it is on the mirrors.
13. Plug the wire of the rigid body into the DAQ.
14. Power on the DAQ. The LEDs on the object should light up. If they don't, try moving the connector to a different spot on the DAQ or flipping the connector around (it may be backwards).
15. On the Phasespace computer, from the drop down menu, switch 'Calibration' to 'Alignment.'
16. Under the 'Alignment Mode' drop down menu, choose first option 'Rigid Body.'
17. Upon selection, a menu should pop up. Select the file '/home/demo/phasespace/alignobjs/rig20140109.rb' from the menu. If for some reason you need to reselect the correct file, go to OWL<Change Trackers<Alignment Rigid to bring the menu back up again.

18. Hit 'Align' when all the LED markers are stable. After hitting 'Align' the image should look upside down but comparable to what Rig 1 actually looks like (what you would normally see in the rig while sitting in the control room).
19. If the image looks good, then hit 'Save' and you are done with calibration. If the image does not look good, then depress the 'Align' button and start the calibration process again.
20. Once satisfied with calibration, disconnect the Calibration screen and reconnect the Master screen. If Master is not open: open xterm; type 'cd phasespace'; type './master'.
21. Power off the DAQ, unplug the array of wires that connected to the alignment markers. Instead, plug in the array of wires that allows the monkey gloves to be plugged in. Plug the ethernet power cord back into the DAQ box. Power DAQ box back on. You will know when you are successful if the LED marker on the glove turns on.
22. Test calibration with the Monkey Host program. If bad for any reason, recalibrate. Under 'Alignment' menu, reselect the '/home/demo/phasespace/alignobjs/rig20140109.rb' file before pressing 'Align.'

A.13 TROUBLE-SHOOTING

Common errors:

1. Trials stop being recorded, the reward sound stops, the XY plot freezes, and the time between trials increases until trials stop being presented all together.
    a. Reason: This problem is yet to be diagnosed and is not regularly reproducible. Hopefully we will have something more complete to go here within the next few months.

b. What to do:

    i. Force close Monkey Host using the task manager

    ii. On the GPU computer, close the render windows and ‘GPU main’ so that the LabVIEW welcome window appears on the screen.

    iii. On the Auxiliary computer, close the online analysis window so that the LabVIEW welcome window appears on the screen.

    iv. Restart the RT computer.

    v. Restart Monkey Host and start trials again

2. Host main will present the following error every second: Error 1 occurred at TCP Read in Host Retrieve Idle Data From RT.vi->Host Retrieve Data From RT.vi->Host Main.vi

    a. Reason: This error occurs when the system is not initialized properly. This can happen when computers are not started in the proper order, or when windows such as the Online Analysis Window or GPU Main are open before the Monkey Host executable is opened.

    b. What to do:

        i. Exit Monkey Host

        ii. Restart RT

        iii. Once the RT monitor says “Welcome to LabVIEW 14.0,” you can reopen Monkey Host

3. Marker is not showing up on the render screen

    a. Reason: This can happen for various reasons, each being addressed below. The most common reason this occurs is that the marker ID specified in Monkey Host does not match the marker being seen by the PhaseSpace cameras. This problem

may also arise if the system is not initialized properly or if PhaseSpace is not properly recording marker positions.

b. What to do:

i. First, make sure that PhaseSpace is registering marker position

1. Turn the KVM switch to the PhaseSpace computer
2. Have someone move the physical marker in the rig (or place it in the field of view - this can be done by hanging the marker from the mirrors or placing it on one of the monkey chairs) and check to see that the marker is visible in the PhaseSpace master interface.
3. If the marker is not being seen by PhaseSpace, disconnect the marker by clicking the “connect” button on the upper left corner of the main window and then reconnect the marker by clicking the “connect” button again so the indicator is green.
4. If the marker is still not being registered, restart the PhaseSpace computer (see “Restarting the PhaseSpace computer” section)

ii. If the marker is being registered by PhaseSpace but is not appearing on the Monkey Host render screen

1. Determine which marker is being registered by PhaseSpace
    a. Place the physical marker in the camera’s field of view in the rig - this can be done by hanging the marker from the mirrors or placing it on one of the monkey chairs
    b. In the PhaseSpace master window, select the Markers tab on the right side of the

window. The marker that is seen by PhaseSpace will be highlighted in green.

2. In Monkey Host, select Edit Markers from the Systems menu

3. Make sure that the Channel Number associated with the phasespace marker is the same as the highlighted marker in PhaseSpace.

   a. If the Channel number does not match the active marker number in PhaseSpace, select the Phasespace marker and change the Channel Number and hit Update Marker

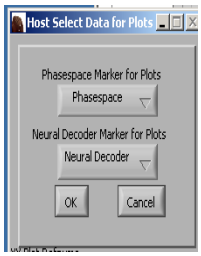

   b. Select Save Changes, causing the Host Select Data for Plots window to open.

   c. Make sure the Phasespace option is selected and hit OK.

4. If the Marker ID is consistent between Monkey Host and Phasespace, the Marker ID is not the cause of the problem and this error is likely due to how the system was initialized

   a. Exit Monkey Host

   b. On the phasespace computer, disconnect the marker by clicking the “connect” button on the upper left corner of the main window.

c. Reconnect the marker by clicking the “connect” button again so the indicator turns green. Verify that Phasespace is registering the marker by placing the physical marker in the rig and observing in the PhaseSpace master window.

d. Turn the KVM switch to the Internet computer

e. Restart Monkey Host

f. If this does not solve the problem, disconnect phasespace and restart RT. Once it says “Welcome to LabVIEW 14.0” reconnect phasespace, turn the KVM switch to the Internet computer and start Monkey Host

4. Error 5005: Real-time could not be maintained

    a. Reason: This error occurs due to system initialization. It happens every time Monkey Host is opened after the real-time computer is restarted.

    b. What to do: This error will occur on the first trial, and then the system will run normally.

5. Phasespace cannot find server. Server field in xterm shows <find servers>.

    a. Reason: Connection was disrupted due to disconnecting ethernet cords.

    b. What to do: Type the phasespace IP address (192.168.1.8) into the server field.

BIBLIOGRAPHY

My Citations

Degenhart AD, Eles J, Dum R, Mischel JL, Smalianchuk I, Endler B, Ashmore RC, Tyler-Kabara EC, Hatsopoulos N, Wang W, Batista AP, Cui XT (2016) Histological Evaluation of a Chronically-implanted Electrocorticographic Electrode Grid in a Non-human Primate. J Neural Eng 13(4): 046019.

Oby ER, Perel S, Sadtler PT, Ruff DA, Mischel JL, Montez DF, Cohen MR, Batista AP, Chase SM (2016) Extracellular Voltage Threshold Settings Can Be Tuned for Optimal Encoding of Movement and Stimulus Parameters. J of Neural Eng 13(3): 036009.

Quick KM, Card NS, Whaite SM, Mischel JL, Loughlin P, Batista AP (2014) Assessing Vibrotactile Feedback Strategies by Controlling a Cursor with Unstable Dynamics. Engineering in Medicine and Biology Society. EMBS 2014.

Quick KM, Mischel JL, Loughlin P, Batista AP. Monkeys can use continuous haptic feedback to stabilize an unstable cursor. Chicago IL: Society for Neuroscience Conference, October 2015, Poster.

Other References

Acharya S, Fifer, MS, Benz HL, Crone NE, Thakor NV (2010) Electrocorticographic amplitude predicts finger positions during slow grasping motions of the hand. J Neural Eng 7(4): 046002.

Adams DL, Economides JR, Jocson CM, Horton JC (2007) A biocompatible titanium headpost for stabilizing behaving monkeys. J Neurophysiol 98:993-1001.

Adams DL, Economides JR, Jocson CM, Parker JM, Horton JC (2011) A watertight acrylic-free titanium recording chamber for electrophysiology in behaving monkeys. J Neurophysiol 106:1581-1590.

Adeeb N, Mortazavi MM, Tubbs RS, Cohen-Gadol AA (2012) The cranial dura mater: a review of its history, embryology, and anatomy. Childs Nerv Syst 28(6): 827-37.

Albrecht DG, Hamilton DB (1982) Striate cortex of monkey and cat: contrast response function. J Neurophysiol 48: 217-37.

Amirikian B, Georgopoulos AP (2003) Modular organization of directionally tuned cells in the motor cortex: is there a short-range order? PNAS 100: 12474-9.

Andersen RA, Mountcastle VB (1983) The influence of the angle of gaze upon the excitability of the light-sensitive neurons of the posterior parietal cortex. J Neurosci 3(3):532{548.

Anderson JM (2001) Biological Response to Materials. Annu Rev Mater Res 31: 81-100.

Anderson JM, Rodriguez A, Chang DT (2008) Foreign Body Reaction to Biomaterials. Semin Immunol 20(2): 86-100.

Ashe J, Georgopoulos AP (1994) Movement parameters and neural activity in motor cortex and area 5. Cerebral Cortex 4:590-600.

Ashmore R, Endler B, Smalianchuk I, Degenhart A, Hatsopoulos N, Tyler-Kabara E, Batista AP, Wang W (2012) Stable Online Control of an Electrocorticographic Brain-Computer Interface Using a Static Decoder. Conf Proc IEEE Eng Med Biol Soc 2012:1740-4.

Bacher DH (2009) Real-time somatosensory feedback for neural prosthesis control: system development and experimental validation. Master's Thesis – University of Pittsburgh.

Ball T, Schulze-Bonhage A, Aertsen A, Mehring C (2009) Differential representation of arm movement direction in relation to cortical anatomy and function. J Neural Eng 6(1): 016006.

Barbolt TA, Odin M, Léger M, Kangas L, Holste J, Liu SH (2001) Biocompatibility evaluation of dura mater substitutes in an animal model. Neurol Res 23(8): 813-20.

Bark K, Wheeler JW, Premakumar S, Cutkosky MR (2008) Comparison of skin stretch and vibrotactile stimulation for feedback of proprioceptive information. Symp Haptic Interfaces Virtual Environ Teleoperator Syst 2008: 17-78.

Barrese JC, Rao N, Paroo K, Triebwasser C, Vargas-Irwin C, Franquemont L, Donoghue JP (2013) Failure mode analysis of silicon-based intracortical microelectrode arrays in non-human primates. J Neural Eng 10(6): 066014.

Batista AP, Buneo CA, Snyder LH, Andersen RA (1999) Reach plans in eye-centered coordinates. Science 285(5425):257–260.

Bensmaia SJ, Miller LE (2014) Restoring sensorimotor function through intracortical interfaces: progress and looming challenges. Nat Rev Neurosci 15: 313-325.

Ben-Shaul Y, Stark E, Asher I, Drori R, Nadasdy Z, Abeles M (2003) Dynamical organization of directional tuning in the primate premotor and primary motor cortex. J Neurophysiol 89: 1136-42.

Betelak KF, Margiotti EA, Wohlford ME, Suzuki DA (2001) The use of titanium implants and prosthodontic techniques in the preparation of non-human primates for long-term neuronal recording studies. J Neurosci Methods 112(1):9–20.

Biran R, Martin DC, Tresco PA (2005) Neuronal cell loss accompanies the brain tissue response to chronically implanted silicon microelectrode arrays. Exp Neurol 195(1): 115-26.

Biran R, Martin DC, Tresco PA (2007) The brain tissue response to implanted silicon microelectrode arrays is increased when the device is tethered to the skull. J Biomed Mater Res A 28(1): 169-78.

Bjornsson CS, Oh SJ, Al-Kofahi YA, Lim YJ, Smith KL, Turner JN, De S, Roysam B, Shain W, Kim SJ (2006) Effects of insertion conditions on tissue strain and vascular damage during neuroprosthetic device insertion. J Neural Eng 3(3): 196-207.

Bouchard KE, Mesgarani N, Johnson K, Chang EF (2013) Functional organization of human sensorimotor cortex for speech articulation. Nature 495: 327-332.

Brånemark PI, Adell R, Breine U, Hansson BO, Lindstrom J, Ohlsson A (1969) Intra-osseous anchorage of dental prostheses. I. Experimental studies. Scand J Plast Reconstr Surg 3(2):81–100.

Bremmer F, Pouget A, Hoffmann KP (1998) Eye position encoding in the macaque posterior parietal cortex. Eur J Neurosci 10(1): 153-160.

Britten KH, Newsome WT, Shadlen MN, Celebrini S, Movshon JA (1996) A relationship between behavioral choice and the visual responses of neurons in macaque MT. Vis Neurosci 13:87-100.

Brockwell AE, Rojas AL, Kass RE (2004). Recursive Bayesian decoding of motor cortical signals by particle filtering. J Neurophysiol 91: 1899–1907.

Brockwell AE, Schwartz AB, Kass RE (2007). Statistical signal processing and the motor cortex. Proceeding of the IEEE 95: 882–898.

Buneo CA, Jarvis MR, Batista AP, Andersen RA (2002) Direct visuomotor transformations for reaching. Nature 416(6881): 632-636.

Burke MW, Gilson RD, Jagacinski RJ (1980) Multi-modal information processing for visual workload relief. Ergonomics 23(10):961-975.

Callaway EM (1998) Local circuits in primary visual cortex of the macaque monkey. Annu Rev Neurosci 21:47-74.

Caminiti R, Johnson PB, Galli C, Ferraina S, Burnod Y (1991) Making arm movements within different parts of space: the premotor and motor cortical representation of a coordinate system for reaching to visual targets. J Neurosci 11(5): 1182-1197.

Cervenka MC, Boatman-Reich DF, Ward J, Franaszczuk, PJ, Crone NE (2011) Language mapping in multilingual patients: electrocorticography and cortical stimulation during naming. Front Hum Neurosci 22(5): 13.

Chaudhary U, Birbaumer N, Ramos-Murguialday A (2016) Brain–computer interfaces for communication and rehabilitation. Nat Rev Neurology 12: 513–525.

Chao ZC, Nagasaka Y, Fujii N (2010) Long-term asynchronous decoding of arm motion using electrocorticographic signals in monkeys. Front Neuroeng 3: 3.

Chase SM, Young E (2008) Cues for sound localization are encoded in multiple aspects of spike trains in the inferior colliculus. J Neurophysiol 99: 1672-82.

Chatterjee A, Chaubey P, Martin J, Thakor N (2008) Testing a prosthetic haptic feedback simulator with an interactive force matching task. J Prosthetics Orthot 20(2): 27–34.

Cheng L, Kazman R, Robinson J (1996) Vibrotactile feedback in delicate virtual reality operations. ACM Multi 1996: 243-251.

Chestek CA, Gilja V, Nuyujukian, P, Foster JD, Fan JM, Kaufman MT, Churchland MM, Rivera-Alvidrez Z, Cunningham JP, Ryu SI, Shenoy KV (2011) Long-term stability of neural prosthetic control signals from silicon cortical arrays in rhesus macaque motor cortex. J Neural Eng 8(4): 045005.

Chestek, CA, Gilja V, Blabe CH, Foster BL, Shenoy KV, Parvizi J, Henderson JM (2013) Hand posture classification using electrocorticography signals in the gamma band over human sensorimotor brain areas. J Neural Eng 10(2): 026002.

Choi JS, DiStasio MM, Brockmeier AJ, Francis JT (2012) An electric field model for prediction of somatosensory (S1) cortical field potentials induced by ventral posterior lateral (VPL) thalamic microstimulation. IEEE Trans Neural Syst Rehabil Eng 20: 161–169.

Christie BP, Tat DM, Irwin ZT, Gilja V, Nuyujukian P, Foster JD, Ryu SI, Shenoy KV, Thompson DE, Chestek CA (2015) Comparison of spike sorting and thresholding of voltage waveforms for intracortical brain-machine interface performance. J Neural Eng 12: 016009.

Churchland MM, Shenoy KV (2007) Temporal complexity and heterogeneity of single-neuron activity in premotor and motor cortex. J Neurophysiol 97: 4235-57.

Cincotti F, Kauhanen L, Aloise F, Palomaki T, Caporusso N, Jylanki P, Mattia D, Babiloni F, Vanacker G, Nuttin M, Marciani MG, Del R Millan J (2007) Vibrotactile feedback for brain-computer interface operation. Comp Intel Neurosci 2007:48937.

Collier TO, Anderson JM, Brodbeck WG, Barber T, Healy KE (2004) Inhibition of macrophage development and foreign body giant cell formation by hydrophilic interpenetrating polymer network. J Biomed Mater Res A 69(4): 644-50.

Collinger JL, Wodlinger B, Downey JE, Wang W, Tyler-Kabara EC, Weber DJ, McMorland AJ, Velliste M, Boninger M. L., and Schwartz AB (2012) High-performance neuroprosthetic control by an individual with tetraplegia. Lancet 381: 507-598.

Cover TM, Thomas JA (1991) Elements of Information Theory. New York: Wiley-Interscience.

Cox G, Kable E, Jones A, Fraser I, Manconi F, Gorrell MD (2003) 3-Dimensional imaging of collagen using second harmonic generation. J Struct Biol141(1): 53-62.

Crammond DJ, Kalaska JF (1994) Modulation of preparatory neuronal activity in dorsal premotor cortex due to stimulus-response compatibility. J Neurophysiol 71(3): 1281-1284.

Crammond DJ, Kalaska JF (1996). Differential relation of discharge in primary motor cortex and premotor cortex to movements versus actively maintained postures during a reaching task. Exp Brain Res 108(1): 45-61.

Crone NE, Hao L, Hart J, Boatman D, Lesser RP, Irizarry R, Gordon B (2001) Electrocorticographic gamma activity during word production in spoken and sign language. Neurology 57(11): 2045-53.

Cui XY, Wiler J, Dzaman M, Altschuler RA, Martin DC (2003) In vivo studies of polypyrrole/peptide coated neural probes. Biomaterials 24: 777–87.

Dadarlat MC, O'Doherty JE, Sabes PN (2015) A learning-based approach to artificial sensory feedback leads to optimal integration. Nat Neurosci 18(1):138-144.

Daly J, Liu J, Aghagolzadeh M, Oweiss K (2012) Optimal space–time precoding of artificial sensory feedback through mutichannel microstimulation in bi-directional brain–machine interfaces. J Neural Eng 9:065004.

Damian DD, Hernandez AA, Martinez H, Pfeifer R (2012) Slip speed feedback for grip force control. IEEE Trans Biomed Eng 59(8): 2200-2210.

Davis TS, Torab K, House P, Greger B (2009) A minimally invasive approach to long-term head fixation in behaving nonhuman primates. J Neurosci Methods 181:106-110.

Degenhart AD, Collinger JL, Vinjamuri R, Kelly J, Tyler-Kabara EC, Wang W (2011a) Classification of hand posture from electrocorticographic signals recorded during varying force conditions. Conf Proc IEEE Eng Med Biol Soc 2011: 5782-5.

Degenhart AD, Kelly JW, Ashmore RC, Collinger JL, Tyler- Kabara EC, Weber DJ, Wang W (2011b) Craniux: A LabVIEW-Based Modular Software Framework for Brain-Machine Interface Research. Comput Intell Neurosci 2011: 363565.

Deng X, Liu DF, Kay K, Frank LM, Eden UT (2015) Clusterless decoding of position from multiunit activity using a marked point process filter. Neural Comput. May 14: 1-23.

Dhillon GS, Horch KW (2005) Direct neural sensory feedback and control of a prosthetic arm. IEEE Trans Neural Syst Rehabil Eng 13(4):468-72.

Ding MC, Lo EH, Stanley GB (2008) Sustained focal cortical compression reduces electrically induced seizure threshold. Neuroscience 154(2): 551-555.

Dorf, R. C. (1967). Modern Control Systems. Addison-Wesley.

Dum RP, Strick PL (1991) The Origin of Corticospinal Projections from the Premotor Areas in the Frontal Lobe. J Neurosci 11(3):667-689.

Dum RP, Strick PL (2002) Motor areas in the frontal lobe of the primate. Physiology & Behavior 77(4-5): 677–682.

Edwards E, Soltani M, Deouell LY, Berger MS, Knight RT (2005) High gamma activity in response to deviant auditory stimuli recorded directly from human cortex. J Neurophysiol 94(6): 4269-80.

Eng LF, Ghirnikar RS, Lee YL (2000) Glial fibrillary acidic protein: GFAP-thirty-one years (1969-2000). Neurochem Res 25(9-10): 1439-1451.

Ethier C, Oby ER, Baumann MJ, Miller LE (2012) Restoration of grasp following paralysis through brain-controlled stimulation of muscles. Nature 485: 361-5.

Evarts EV (1966) Methods for recording activity of individual neurons in moving animals. Meth Med Res 11:241–50.

Evarts EV (1968) A technique for recording activity of subcortical neurons in moving animals. Electroenceph Clin Neurophysiol 24:83–6.

Evarts EV (1972) Activity of Motor Cortex Neurons in Association with Learned Movement. International J Neurosci 3(3):113-124.

Fabiani GE, McFarland DJ, Wolpaw JR, Pfurtscheller G (2004) Conversion of EEG activity into cursor movement by a brain-computer interface (BCI). IEEE Trans Neural Syst Rehabil Eng 12(3):331-8.

Fagg AH, Hatsopoulos NG, de Lafuente V, Moxon KA, Nemati S, Rebesco JM, Romo R, Solla SA, Reimer J, Tkach D, Pohlmeyer EA, Miller LE (2007) Biomimetic brain machine interfaces for the control of movement. J Neurosci 27: 11842–11846.

Fetz EE (1969) Operant conditioning of cortical unit activity. Science 163(3870):955-8.

Fetz EE, Cheney PD (1980) Postspike facilitation of forelimb muscle activity by primate corticomotoneuronal cells. J Neurophysiol 44(4): 751-72.

Fitzsimmons NA, Drake W, Hanson TL, Lebedev MA, Nicolelis MAL (2007). Primate reaching cued by multichannel spatiotemporal cortical microstimulation. J Neurosci 27(21): 5593-602.

Flint RD, Wright ZA, Scheid MR, Slutzky MW (2013) Long term, stable brain machine interface performance using local field potentials and multiunit spikes. J Neural Eng 10(5): 056005.

Fogassi L, Raos V, Franchi G, Gallese V, Luppino G, Matelli M (1999) Visual responses in the dorsal premotor area F2 of the macaque monkey. Exp Brain Res 128: 194-199.

Fountas KN, Smith JR (2007) Subdural Electrode-Associated Complications: A 20-Year Experience. Stereotact Funct Neurosurg 85(6): 264-72.

Fraser GW, Chase SM, Whitford AS, Schwartz AB (2009) Control of a brain-computer interface without spike sorting. J Neural Eng 6: 055004.

Freire MAM, Morya E, Faber J, Santos JR, Guimaraes JS, Lemos NA, Sameshima K, Pereira A, Ribeiro S, Nicolelis MA (2011) Comprehensive analysis of tissue preservation and recording quality from chronic multielectrode implants. PLoS ONE 6(11): e27554.

Friendlich AR (1973) Primate head restrainer using a nonsurgical technique. J Appl Physiol 35:934–5.

Fu QG, Flament D, Coltz JD, Ebner TJ (1995) Temporal encoding of movement kinematics in the discharge of primate primary motor and premotor neurons. J Neurophysiol 73(2): 836-854.

Funahashi S, Bruce CJ, Goldman-Rakic PS (1989) Mnemonic coding of visual space in the monkey's dorsolateral prefrontal cortex. J Neurophysiol 61(2):331-349.

Gail A, Andersen RA (2006) Neural dynamics in monkey parietal reach region reflect context-specific sensorimotor transformations. J Neurosci 26(37): 9376-84.

Gandhi N, Katnani H (2011) Motor Functions of the Superior Colliculus. Annu Rev Neurosci 34: 205-231.

Ganguly K, Carmena JM (2009) Emergence of a stable cortical map for neuroprosthetic control. PLoS Biol 7(7): e1000153.

Gardner EP, Costanzo RM (1980) Spatial integration of multiple-point stimuli in primary somatosensory cortical receptive fields of alert monkeys. J Neurophysiol 43(2):420-43.

Georgopoulos AP, Kalaska JF, Caminiti R, Massey JT (1982) On the relations between the direction of two-dimensional arm movements and cell discharge in primate motor cortex. J Neurosci 2: 1527-1537.

Georgopoulos AP, Schwartz AB, Kettner R (1986) Neuronal population coding of movement direction. Science 233:1416-1419.

Georgopoulos AP, Kettner RE, Schwartz AB (1988) Primate Motor Cortex and Free Arm Movements to Visual Targets in Three-Dimensional Space. II. Coding of the direction of Movement a Neuronal Population. J Neurosci 8:2928-37.

Georgopoulos AP, Merchant H, Naselaris T, Amirikian B (2007) Mapping of the preferred direction in the motor cortex. PNAS 104: 11068-72.

Ghez C, Gordon J, Ghilardi MF (1995) Impairments of reaching movements in patients without proprioception. II. Effects of visual information on accuracy. J Neurophysiol 73:361-72.

Gilja V, Nuyujukian P, Chestek CA, Cunningham JP, Yu BM, Fan JM, Churchland MM, Kaufman MT, Kao JC, Ryu SI, Shenoy KV (2012) A high-performance prosthesis enabled by control algorithm design. Nat Neurosci 15: 1752-7.

Gittins R, Harrison PJ (2004) Neuronal density, size and shape in the human anterior cingulate cortex: a comparison of Nissl and NeuN staining. Brain Res Bull 63(2): 155-60.

Godlove JM, Whaite EO, Batista AP (2014) Comparing temporal aspectsof visual, tactile, and microstimulation feedback for motor control. J Neural Eng 11(4):046025.

Golub MD, Yu BM, Schwartz AB, Chase SM (2014) Motor cortical control of movement speed with implications for brain-machine interface control. J Neurophysiol 112: 411-29.

Golub MD, Yu BM, Chase SM (2015) Internal models for interpreting neural population activity during sensorimotor control. eLife 4:e10015

Golub MD, Chase SM, Batista AP, Yu BM (2016) Brain–computer interfaces for dissecting cognitive processes underlying sensorimotor control. Curr Op Neurobiol 37:53-58.

Goodworth AD, Wall C, Peterka RJ (2009) Influence of feedback parameters on performance of a vibrotactile balance prosthesis. IEEE Trans Neural Syst Rehabil Eng 17(4):397-408.

Gordon J, Ghilardi MF, Ghez C (1995) Impairments of reaching movements in patients without proprioception I. Spatial errors. J Neurophysiol 73:347-360.

Griffith GW, Humphrey DR (2006) Long-term gliosis around chronically implanted platinum electrodes in the Rhesus macaque motor cortex. Neurosci Lett 406: 81-6.

Hall WC, Moschovakis AK (2004) The Superior Colliculus: New Approaches for Studying Sensorimotor Integration. CRC Press.

Hatsopoulos NG (2010) Columnar organization in the motor cortex. Cortex 46: 270-1.

Hatsopoulus NG, Suminski AJ (2011) Sensing with motor cortex. Neuron 72:77-87.

He SQ, Dum RP, Strick PL (1993) Topographic organization of corticospinal projections from the frontal lobe: motor areas on the lateral surface of the hemisphere. J Neurosci 13(3): 952-80.

Henle C, Raab M, Cordeiro JG, Doostkam S, Schulze-Bonhage A, Stieglitz T, Rickert J (2011) First long term in vivo study on subdurally implanted micro-ECoG electrodes, manufactured with a novel laser technology. Biomed microdevices 13(1):59-68.

Hochberg LR, Serruya MD, Friehs GM, Mukand JA, Saleh M, Caplan AH, Branner A, Chen D, Penn RD, Donoghue JP (2006) Neuronal ensemble control of prosthetic devices by a human with tetraplegia. Nature 442(7099):164-171.

Hochberg LR, Bacher D, Jarosiewicz B, Masse NY, Simeral JD, Vogel J, Haddadin S, Liu J, Cash SS, van der Smagt P, Donoghue JP (2012) Reach and grasp by people with tetraplegia using a neutrally controlled robotic arm. Nature 485: 372-5.

Hogan N, Sternad D (2012) Dynamic primitives of motor behavior. Biol Cybernetics 106(11-12):727-739.

Horch K, Meek S, Taylor TG, Hutchinson DT (2011) Object discrimination with an artificial hand using electrical stimulation of peripheral tactile and proprioceptive pathways with intrafascicular electrodes. IEEE Trans Neural Syst Rehabil Eng 19(5):483-489.

Hotson G, McMullen DP, Fifer MS, Johannes MS, Katyal KD, Para MP, Armiger R, Anderson WS, Thakor NV, Wester BA, Crone NE (2016) Individual finger control of a modular prosthetic limb using high-density electrocorticography in a human subject. J Neural Eng 13(2): 026017.

Hubel DH, Wiesel TN (1959) Receptive fields of single neurons in the cat's striate cortex. J Physiol 148: 574-91.

Hubel DH, Wiesel TN (1962) Receptive fields, binocular interaction and functional architecture in the cat's visual cortex. J Physiol 160(1): 106-154.

Hubel DH (1982) Exploration of the primary visual cortex, 1955-78. Nature 299: 515-524.

Hulliger M, Nordh E, Thelin AE, Vallbo AB (1979) The responses of afferent fibers from the glabrous skin of the hand during voluntary finger movements in man. J Physiol 291: 233-249.

Humphrey DR, Schmidt EM, Thompson WD (1970) Predicting Measures of Motor Performance from Multiple Cortical Spike Trains. Science 170(3959): 758-762.

Huston SJ, Jayaraman V (2011) Studying sensorimotor integration in insects. Current Opinion in Neurobiology 21(4): 527–534.

Ifft PJ, Shokur S, Lebedev MA, Nicolelis MA (2013) A brain-machine interface enables bimanual arm movements in monkeys. Sci Transl Med 5: 210ra154.

Ikezoe K, Mori Y, Kitamura K, Tamura H, Fujita I (2013) Relationship between the local structure of orientation map and the strength of orientation tuning of neurons in monkey V1: a 2-photon calcium imaging study. J Neurosci 33: 16818-27.

Isoda M, Tsutsui K, Katsuyama N, Naganuma T, Saito N, Furusawa Y, Mushiake H, Taira M, Tanji J (2005) Design of a head fixation device for experiments in behaving monkeys. J Neurosci Methods 141:277–282.

Jay MF, Sparks DL (1987) Sensorimotor integration in the primate superior colliculus. I. Motor convergence. J Neurophysiol 57(1):22-34.

Jex HR, McDonnell JD, Phatak AV (1966) A “Critical’' Tracking Task for Manual Control Research. IEEE Transactions on Human Factors in Electronics 7:138-145.

Jiang L, Cutkosky M, Ruutiainen J, Raisamo R (2009) Using haptic feedback to improve grasp force control in multiple sclerosis patients. IEEE Trans on Robot 25(3): 593-601.

Johansson RS (1978) Tactile sensibility in the human hand: receptive field characteristics of mechanoreceptive units in the glabrous skin area. J Physiol 281: 101-125.

Jung J, Mainy N, Kahane P, Minotti L, Hoffmann D, Bertrand O, Lachaux JP (2008). The neural bases of attentive reading. Hum Brain Mapp 29(10): 1193-206.

Kaczmarek KA, Webster JG, Bach-Y-Rita P, Tompkins WJ (1991) Electrotactile and vibrotactile displays for sensory substitution systems. IEEE Trans Biomed Eng 38(1):1-16.

Kadkade PP, Benda BJ, Schmidt PB, Wall C (2003) Vibrotactile display coding for a balance prosthesis. IEEE Trans Neural Syst Rehabil Eng 11(4): 392-9.

Kakei SI, Hoffman DS, Strick PL (1999) Muscle and movement representations in the primary motor cortex. Science 285(5436):2136-9.

Kam L, Shain W, Turner JN, Bizios R (2002) Selective adhesion of astrocytes to surfaces modified with immobilized peptides. Biomaterials 23: 511–5.

Kandel ER, Schwartz JH, Jessell TM (2000) The Bodily Senses. In The Principles of Neural Science, pages 430-471. McGraw-Hill, New York, fourth edition.

Kao JC, Nuyujukian P, Ryu SI, Churchland MM, Cunningham JP, Shenoy KV (2015) Single-trial dynamics of motor cortex and their applications to brain-machine interfaces. Nat Comm 6:1-12.

Kaufman MT, Churchland MM, Ryu SI, Shenoy KV (2014) Cortical activity in the null space: permitting preparation without movement. Nat Neurosci 17(3):440-448.

Kaufman MT, Churchland MM, Ryu SI, Shenoy KV (2015) Vacillation, indecision and hesitation in moment-by-moment decoding of monkey motor cortex. eLife

Kellis S, Miller K, Thomson K, Brown R, House P, Greger B (2010). Decoding spoken words using local field potentials recorded from the cortical surface. J Neural Eng 7: 056007.

Kelly JW, Siewiorek DP, Smailagic A, Wang W (2013) Automated filtering of common-mode artifacts in multichannel physiological recordings. IEEE Trans Biomed Eng 60(10): 2760-70.

Kelly R, Strick P (2003) Cerebellar loops with motor cortex and prefrontal cortex of a nonhuman primate. The J Neurosci 23(23): 8432–8444.

Kemere C, Shenoy KV, Meng TH (2004) Model-based neural decoding of reaching movements: A maximum likelihood approach. IEEE Trans Biomed Eng 51:925-32.

Kennedy PR, Bakay RA (1998) Restoration of neural output from a paralyzed patient by direct brain connection. NeuroReport 9:1707–1711.

Kim DH, Viventi J, Amsden JJ, Xiao J, Vigeland L, Kim YS, Blanco JA, Panilatis B, Frechette ES, Contreras D, Kaplan DL, Omenetto FG, Huang Y, Hwang KC, Zakin MR, Litt B, Rogers JA (2010) Dissolvable films of silk fibroin for ultrathin conformal bio-integrated electronics. Nat Mater 9: 511-517.

Kim JJ, Gean AD (2011) Imaging for the diagnosis and management of traumatic brain injury. Neurotherapeutics 8: 39-53.

Kim SP, Simeral JD, Hochberg LR, Donoghue JP, Black MJ (2008) Neural control of computer cursor velocity by decoding motor cortical spiking activity in humans with tetraplegia. Journal of Neural Engineering 5:455-76.

Kloosterman F, Layton SP, Chen Z, Wilson MA (2014) Bayesian decoding using unsorted spikes in the rat hippocampus. J Neurophysiol 111: 217-27.

Kolarcik CL, Bourbeau D, Azemi E, Rost E, Zhang L, Lagenaur CF, Weber DJ, Cui XT (2012) In vivo effects of L1 coating on inflammation and neuronal health at the electrode/tissue interface in rat dorsal root ganglion and spinal cord. Acta Biomater 8(10): 3561-75.

Kondraske GV, Potvin AR, Tourtellotte WW, Syndulko K (1984) A computer-based system for automated quantitation of neurologic function. IEEE Trans Biomed Eng 5:401-414.

Kongsui R, Beynon SB, Johnson SJ, Walker FR (2014) Quantitative assessment of microglial morphology and density reveals remarkable consistency in the distribution and morphology of cells within the healthy prefrontal cortex of the rat. J Neuroinflammation 11: 182.

Koyama S, Chase SM, Whitford AS, Velliste M, Schwartz AB, Kass RE (2010) Comparison of brain–computer interface decoding algorithms in open-loop and closed-loop control. J Comp Neurosci 29(1-2): 73-87.

Kozai TDY, Vasquez AL, Weaver CL, Kim SG, Cui XT (2012) In vivo two-photon microscopy reveals immediate microglial reaction to implantation of microelectrode through extension of processes. J Neural Eng 9: 066001.

Kozai TDY, Jaquins-Gerstl AS, Vazquez AL, Michael AC, Cui XT (2015) Brain Tissue Responses to Neural Implants Impact Signal Sensitivity and Intervention Strategies. ACS Chem Neurosci 6(1): 48–67.

Kubanek J, Miller KJ, Ojemann JG, Wolpaw JR, Schalk G (2009) Decoding flexion of individual fingers using electrocorticographic signals in humans. J Neural Eng 6(6): 066001.

Lachaux JP, George N, Tallon-Baudry C, Martinerie J, Hugueville L, Minotti L, Kahane P, Renault B (2005) The many faces of the gamma band response to complex visual stimuli. Neuroimage 25: 491-501.

Lebedev MA, Carmena JM, O'Doherty JE, Zacksenhouse M, Henriquez CS, Principe JC, Nicolelis MAL (2005) Cortical ensemble adaptation to represent velocity of an artificial actuator controlled by a brain-machine interface. J Neurosci 25:4681-93.

Lebedev MA, Nicolelis MA (2006) Brain-machine interfaces: past, present and future. Trends Neurosci 29(9):536-46.

Ledbetter NM, Ethier C, Oby ER, Hiatt SD, Wilder AM, Ko JH, Agnew SP, Miller LE, Clark GA (2013) Intrafascicular stimulation of monkey arm nerves evokes coordinated grasp and sensory responses. J Neurophysiol 109:580-590.

Leeb R, Gwak K, Kim DS, Millan JDR (2013) Freeing the visual channel by exploiting vibrotactile BCI feedback. 35th Annu Int Conf IEEE 2013: 3093-6.

Leuthardt EC, Gaona C, Sharma M, Szrama N, Roland J, Freudenberg Z, Solis J, Breshears J, Schalk G (2011) Using the electrocorticographic speech network to control a brain-computer interface in humans. J Neural Eng 8: 036004.

Leuthardt EC, Schalk G, Moran D, Ojemann JG (2006) The emerging world of motor neuroprosthetics: a neurosurgical perspective. Neurosurgery 59: 1-14.

Leuthardt EC, Schalk G, Wolpaw JR, Ojemann JG, Moran DW (2004) A brain-computer interface using electrocorticographic signals in humans. J Neural Eng 1(2): 63-71.

Lewicki MS (1998) A review of methods for spike sorting: the detection and classification of neural action potentials. Network 9: R53-78.

Li CR, Padoa-Schioppa C, Bizzi E (2001) Neuronal correlates of motor performance and motor learning in the primary motor cortex of monkeys adapting to an external force field. Neuron 30(2): 593-607.

Lin CC, Whitney SL, Loughlin P, Furman JM, Redfern MS, Sienko KH, Sparto PJ (2015) The effect of age on postural and cognitive task performance while using vibrotactile feedback. J Neurophysiol 113(7):2127-2136.

Linder L, Albrektsson T, Brånemark PI, Hansson HA, Ivarsson B, Jonsson U, Lundstrom I (1983) Electron microscopic analysis of the bonetitanium interface. Acta Orthop Scand 54:45–52.

London BM, Jordan LR, Jackson CR, Miller LE (2008) Electrical stimulation of the proprioceptive cortex (area 3a) used to instruct a behaving monkey. IEEE Trans Neural Syst Rehabil Eng16(1): 32-6.

Loughlin P, Mahboobin A, Furman J (2011) Designing vibrotactile balance feedback for desired body sway reductions. 33rd Annu Int Conf IEEE Eng Med Biol Soc 2011:1310-3.

Lynn AD, Blakney AK, Kyriakides TR, Bryant SJ (2010) Temporal progression of the host response to implanted poly(ethylene glycol)-based hydrogels. J Biomed Mater Res A 96(4): 621-631.

Mainy N, Kahane P, Minotti L, Hoffmann D, Bertrand O, Lachaux JP (2007) Neural correlates of consolidation in working memory. Hum Brain Mapp 28(3): 183-93.

Malik WQ, Donoghue JP, Hochberg LR (2014, August) Decoding Analog Multiunit Activity in a Brain-Machine Interface for Individuals with Tetraplegia. Paper presented at the IEEE Engineering in Medicine and Biology Society Conference, Chicago, IL.

Martinez J, Pedreira C, Ison MJ, Quian Quiroga R (2009) Realistic simulation of extracellular recordings. J Neurosci Methods 184: 285-93.

Matelli M, Luppino G, Rizzolatti G (1991) Architecture of superior and mesial area 6 and the adjacent cingulate cortex in the macaque monkey. J Comp Neurol 311(4): 445-462.

McConnell GC, Rees HD, Levey AI, Gutekunst CA, Gross RE, Bellamkonda RV (2009). Implanted neural electrodes cause chronic, local inflammation that is correlated with local neurodegeneration. J Neural Eng 6(5): 056003.

Miall RC, Wolpert DM (1996) Forward Models for Physiological Motor Control. Neural Networks 9(8):1265-1279.

Miller KJ, Zanos S, Fetz EE, Den Nijs M, Ojemann JG (2009) Decoupling the Cortical Power Spectrum Reveals Real-Time Representation of Individual Finger Movements in Humans. J Neurosci 29(1): 3132-7.

Moberg E (1962) Criticism and study of methods for examining sensibility in the hand. Neurology 12:8-19.

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

Moran D (2010) Evolution of brain-computer interface: action potentials, local field potentials and electrocorticograms. Curr Opin Neurobiol 20(6): 741-5.

Moreno-López Y, Olivares-Moreno R, Cordero-Erausquin M, Rojas-Piloni G (2016) Sensorimotor Integration by Corticospinal System. Front Neuroanat 10: 24.

Morrell MJ, RNS System in Epilepsy Study Group (2011) Responsive cortical stimulation for the treatment of medically intractable partial epilepsy. Neurology 77(13): 1295-304.

Mountcastle VB (1975) Posterior parietal association cortex of the monkey: command functions for operations within extrapersonal space. J Neurophysiol 38(4):871–908.

Mountcastle VB, Steinmetz MA, Romo R (1990) Frequency discrimination in the sense of flutter: psychophysical measurements correlated with postcentral events in behaving monkeys. J Neurosci 10(9):3032-3044.

Nakanishi Y, Yanagisawa T, Shin, D, Fukuma R, Chen C, Kambara H, Yoshimura N, Hirata M, Yoshimine T, Koike Y (2013) Prediction of Three-Dimensional Arm Trajectories Based on ECoG Signals Recorded from Human Sensorimotor Cortex. PLoS ONE 8(8): e72085.

Nashold BS, Friedman H (1972) Dorsal column stimulation for control of pain: Preliminary report on 30 patients. J Neurosurg36: 490-7.

Nasseroleslami B, Hasson CJ, Sternad D (2014) Rhythmic manipulation of objects with complex dynamics: predictability over chaos. PLoS Comput Biol 10(10):1-19.

Nauhaus I, Benucci A, Carandini M, Ringach DL (2008) Neuronal selectivity and local map structure in visual cortex. Neuron 57: 673-9.

Nicolelis MAL, Dimitrov D, Carmena JM, Crist R, Lehew G, Kralik JD, Wise SP (2003) Chronic, multisite, multielectrode recordings in macaque monkeys. Proc Natl Acad Sci USA 100: 11041–6.

Nimmerjahn A, Kirchhoff F, Helmchen F (2005) Resting microglial cells are highly dynamic surveillants of brain parenchyma in vivo. Science 308: 1314-18.

Norton LW, Tegnell E, Toporek SS, Reichert WM (2005) In vitro characterization of vascular endothelial growth factor and dexamethasone releasing hydrogels for implantable probe coatings. Biomaterials 26(16): 3285-97.

Nowak LG, Munk MH, Girard P, Bullier J (1995) Visual latencies in areas V1 and V2 of the macaque monkey. Vis Neurosci 12: 371-384.

Nunamaker EA, Kipke DR (2010) An alginate hydrogel dura mater replacement for use with intracortical electrodes. J Biomed Mater Res 95(2): 421-9.

Oby ER, Ethier C, Miller LE (2013) Movement representation in the primary motor cortex and its contribution to Generalizable EMG predictions. J Neurophysiol 109(3): 666–678.

O'Doherty JE, Lebedev MA, Ifft PJ, Zhuang KZ, Shokur S, Bleuler H, Nicolelis MAL (2011) Active tactile exploration using a brainmachinebrain interface. Nature 479: 228-31.

Paillard J, Brouchon M (1968) Active and passive movements in the calibration of position sense. In The Neuropsychology of Spatially Oriented Behavior, pages 37-55. Dorsey Press, Homewood, IL.

Pasley BN, David SV, Mesgarani N, Flinker A, Shamma SA, Crone NE, Knight RT, Chang EF (2012) Reconstructing speech from human auditory cortex. PLoS Biol 10(1): e1001251.

Pedreira C, Martinez J, Ison MJ, Quian Quiroga R(2012) How many neurons can we see with current spike sorting algorithms? J Neurosci Methods 211: 58-65.

Pei YC, Hsiao SS, Craig JC, Bensmaia SJ (2011) Neural mechanisms of tactile motion integration in somatosensory cortex. Neuron 69: 536-547

Pei YC, Hsiao SS, Craig JC, Bensmaia SJ (2010) Shape invariant coding of motion direction in somatosensory cortex. Public Library of Science Biology, 8: e1000305.

Pei X, Barbour DL, Leuthardt EC, Schalk G (2011) Decoding vowels and consonants in spoken and imagined words using electrocorticographic signals in humans. J Neural Eng 8(4): 046028.

Perel S, Sadtler PT, Oby ER, Ryu SI, Tyler-Kabara EC, Batista AP, Chase S (2015) Single-unit activity, threshold crossings, and local field potentials in motor cortex differentially encode reach kinematics. J Neurophysiol. Epub.

Peterka RJ (2002) Sensorimotor integration in human postural control. J Neurophysiol 88:1097-1118.

Peterka RJ, Loughlin PJ (2004) Dynamic regulation of sensorimotor integration in human postural control. J Neurophy 91:410-423.

Peters A, Sethares S (2002) The effects of age on the cells in layer 1 of primate cerebral cortex. Cereb Cortex 12: 27-36.

Pfingst BE, Albrektsson T, Tjellstrom A, Miller JM, Zappia J, Xue X, Weiser F (1989) Chronic skull-anchored percutaneous implants in non-human primates. J Neurosci Methods 29:207–16.

Pineda A (1978) Complications of dorsal column stimulation. J Neurosurg 48: 64-8.

Pistohl T, Ball T, Schulze-Bonhage A, Aertsen A, Mehring C (2008) Prediction of arm movement trajectories from ECoG-recordings in humans. J Neurosci Methods 167(1): 105:14.

Polikov VS, Tresco PA, Reichert WM (2005) Response of brain tissue to chronically implanted neural electrodes. J Neurosci Methods 148(1): 1-18.

Porter R, Lemon R (1993) Corticospinal Function and Voluntary Movement. Oxford: Clarendon Press.

Potvin AR, Doerr JA, Estes JT, Tourtellotte WW (1997) Portable clinical tracking-task instrument. Med Biol Eng Comp 15(4):391-397.

Preibisch S, Saalfeld S, Tomancak P (2009) Globally optimal stitching of tiled 3D microscopic image acquisitions. Bioinformatics 25(11): 1463-1465.

Preul MC, Bichard WD, Spetzler RF (2003) Toward optimal tissue sealants for neurosurgery: use of a novel hydrogel sealant in a canine durotomy repair model. Neurosurgery 53(5): 1189-98.

Pruszynski JA, Kurtzer I, Scott SH (2008) Rapid motor responses are appropriately tuned to the metrics of a visuospatial task. J Neurophysiol 100:224-238.

Pruszynski JA, Kurtzer I, Nashed JY, Omrani M, Brouwer B, Scott SH (2011) Primary motor cortex underlies multi-joint integration for fast feedback control. Nature 478:387-390.

Pruszynski JA, Omrani M, Scott SH (2014) Goal-dependent modulation of fast feedback responses in primary motor cortex. J Neurosci 34(13):4608-4617.

Ptito M, Moesgaard SM, Gjedde A, Kupers R (2005) Cross-modal plasticity revealed by electrotactile stimulation of the tongue in the congenitally blind. Brain 128(3):606-614.

Pylatiuk CI, Mounier S, Kargov A, Schulz S, Bretthauer G (2004) Progress in the development of a multifunctional hand prosthesis. Conf Proc IEEE Eng Med Biol Soc 6:4260-3.

Quick KM (2016) Investigation of methods for assessing sensorimotor performance in humans and monkeys. PhD Thesis – University of Pittsburgh.

Ramaekers JG, Kauert G, van Ruitenbeek P, Theunissen EL, Schneider E, Moeller MR (2006a) High-potency marijuana impairs executive function and inhibitory motor control. Neuropsychopharmacology, 31(10): 2296-303.

Ranck JB, Jr (1975) Which elements are excited in electrical stimulation of mammalian central nervous system: A review. Brain Research, 98:417-440.

Rao NG, Donoghue JP (2014) Cue to action processing in motor cortex populations. J Neurophysiol 111:441-453.

Rathelot JA, Strick PL (2009) Subdivisions of primary motor cortex based on cortico-motor neuronal cells. Proc Natl Acad Sci USA 106(3): 918–923.

Rattay F (1988) Modeling the excitation of fibers under surface electrodes. IEEE Trans Biomed Eng 35(3): 199–202

Ray S, Niebur E, Hsiao SS, Sinai A, Crone NE (2008) High-frequency gamma activity (80-150Hz) is increased in human cortex during selective attention. Clin Neurophysiol 119(1): 116-33.

Rebesco JM, Miller LE (2011) Enhanced detection threshold for in vivo cortical stimulation produced by Hebbian Conditioning. Journal of Neural Engineering 8:1-11.

Riehle A, Requin J (1989) Monkey primary motor and premotor cortex: single-cell activity related to prior information about direction and extent of an intended movement. J Neurophysiol 61(3): 534-549.

Reina GA, Moran DW, Schwartz AB, Snider J, Lee D, Harrington DL, Poizner H, Velliste M, Kennedy SD, Whitford AS, Sohn JW, McMorland AJC (2001) On the relationship between joint angular velocity and motor cortical discharge during reaching. J Neurophysiol 85: 2576–2589.

Rey HG, Pedreira C, Quian Quiroga R (2015) Past, present and future of spike sorting techniques. Brain Res Bull. Epub.

Roth TL, Nayak D, Atanasijevic T, Koretsky AP, Latour LL, McGavern DB (2014) Transcranial amelioration of inflammation and cell death after brain injury. Nature 505: 223-228.

Rombokas E, Stepp CE, Chang C, Malhotra M, Matsuoka Y (2013) Vibrotactile sensory substitution for electromyographic control of object manipulation. IEEE Trans Biomed Eng 60(8): 2226-32.

Romo R, Hernández A, Zainos A, Salinas E (1998) Somatosensory discrimination based on cortical microstimulation. Nature 392:387-90.

Romo R, Hernández A, Zainos A, Brody CD, Lemus L (2000) Sensing without touching: psychophysical performance based on cortical microstimulation. Neuron 26:273-8.

Rousche PJ, Normann RA (1998) Chronic recording capability of the Utah Intracortical Electrode Array in cat sensory cortex. J Neurosci Methods 82(1): 1-15.

Rubehn B, Bosman C, Oostenveld R, Fries P, Stieglitz T (2009) A MEMS-based flexible multichannel ECoG-electrode array. J Neural Eng 6(3): 036003.

Saal HP, Bensmaia SJ (2014) Touch is a team effort: interplay of submodalities in cutaneous sensibility. Cell 37(12): 689–697.

Saal HP, Harvey MA, Bensmaia SJ (2015) Rate and timing of cortical responses driven by separate sensory channels. eLife 4: e10450.

Sadtler PT, Quick KM, Golub MD, Chase SM, Ryu SI, Tyler-Kabara EC, Yu BM, Batista AP (2014) Neural constraints on learning. Nature 512(7515): 423-6.

Saitoh Y, Shibata M, Hirano SI, Hirata M, Mashimo T, Yoshimine T (2000) Motor cortex stimulation for central and peripheral deafferentation pain: Report of eight cases. J Neurosurg 92: 150-5.

Sajanti J, Bjorkstrand AS, Finnila S, Heikkinen E, Peltonen J, Majamaa K (1999) Increase of collagen synthesis and deposition in the arachnoid and the dura following subarachnoid hemorrhage in the rat. Biochim Biophys Acta 1454(3): 209-16.

Salinas E, Abbott LF (1994) Vector reconstruction from firing rates. J Comp Neurosci 1:89-107.

Sanders JE, Stiles CE, Hayes CL (2000) Tissue response to single-polymer fibers of varying diameters: evaluation of fibrous encapsulation and macrophage density. J Biomed Mater Res 52(1): 231-7.

Sayama CM, Sorour M, Schmidt RH (2014) Dural adhesion to porous cranioplastic implant: A potential safety concern. Surg Neurol Int 5: 19.

Saxena T, Karumbaiah L, Gaupp EA, Patkar R, Patil K, Betancur M, Stanley GB, Bellamkonda RV (2013) The impact of chronic blood-brain barrier breach on intracortical electrode function. Biomaterials 34(20): 4703-13.

Schalk, G, Miller KJ, Anderson, NR, Wilson JA, Smyth MD, Ojemann JG, Moran DW, Wolpaw JR, Leuthardt EC (2008) Two-dimensional movement control using electrocorticographic signals in humans. J Neural Eng 5(1): 75-84.

Schalk G, Lethardt EC (2011) Brain-computer interfaces using electrocorticographic signals. IEEE Rev Biomed Eng 4: 140-54.

Schendel AA, Nonte MW, Vokoun C, Richner TJ, Brodnick SK, Atry F, Frye S, Bostrom P, Pashale R, Thongpang S, Eliceirl KW, and Williams JC (2014) The effect of micro-ECoG substrate footprint on the meningeal tissue response. J Neural Eng 11(4): 046011.

Schendel AA, Thongpang S, Brodnick SK, Richner TJ, Lindevig BDB, Krugner-Higby L, Williams JC (2013) A cranial window imaging method for monitoring vascular growth around chronically implanted micro-ECoG devices. J Neurosci Methods 218(1): 121-30.

Schieber MH, Hibbard LS (1993) How somatotopic is the motor cortex hand area? Science 261: 489-92.

Schummers J, Marino J, Sur M (2002) Synaptic integration by V1 neurons depends on location within the orientation map. Neuron 36: 969-78.

Schwartz AB, Kettner RE, Georgopoulos AP (1988) Primate motor cortex and free arm movements to visual targets in three-dimensional space. J Neurosci 8: 2913-2927.

Schwartz, AB (2004) Cortical neural prosthetics. Annual Rev Neurosci 27: 487-507.

Schwartz AB, Cui XT, Weber DJ, Moran DW (2006) Brain-controlled interfaces: movement restoration with neural prosthetics. Neuron 52: 205-20.

Scott SH, Loeb GE (1994) The computation of position sense from spindles in mono- and multiarticular muscles. J Neurosci 14(12):7529-7540.

Scott SH (2004) Optimal feedback control and the neural basis of volitional motor control. Nat Rev Neurosci 5(7): 532-546.

Serruy MD, Hatsopoulos NG, Paninski L, Fellows MR, Donoghue JP (2002) Brain-machine interface: Instant neural control of a movement signal. Nature 416: 141-142.

Sharma A, Rieth L, Tathireddy P, Harrison R, Oppermann H, Klein M, Töpper M, Jung E, Normann R, Clark G, Solzbacher F (2011) Long term in vitro functional stability and recording longevity of fully integrated wireless neural interfaces based on the Utah Slant Electrode Array. J Neural Eng 8(4):045004.

Shih JJ, Krusienski DJ, Wolpaw JR (2012) Brain-computer interfaces in medicine. Mayo Clin Proc 87(3): 268–279.

Shimoda K, Nagasaka Y, Chao ZC, Fujii N (2012) Decoding continuous three-dimensional hand trajectories from epidural electrocorticographic signals in Japanese macaques. J Neural Eng 9(3): 036015.

Shoham S, Fellows MR, Normann RA (2003) Robust, automatic spike sorting using mixtures of multivariate t-distributions. J Neurosci Methods 127: 111-122.

Sillay KA, Rutecki P, Cicora K, Worrell G, Drazkowski J, Shih JJ, Sharan AD, Morrell MJ, Williams J, Wingeier B (2013) Long-Term Measurement of Impedance in Chornically Implanted Depth and Subdural Electrodes During Responsive Neurostimulation in Humans. Brain Stimul 6: 718-26.

Simeral JD, Kim SP, Black MJ, Donoghue JP, Hochberg LR (2011) Neural control of cursor trajectory and click by a human with tetraplegia 1000 days after implant of an intracortical microelectrode array. J Neural Eng 8: 025027.

Singh A, Ehteshami G, Massia S, He JP, Storer RG, Raupp G (2003) Glial cell and fibroblast cytotoxicity study on plasma-deposited diamond-like carbon coatings. Biomaterials 24: 5083–9.

Skoglund S (1973). Joint Receptors and Kinaesthesia. In Somatosensory System, pages 111-136. Springer.

Slutzky MW, Jordan LR, Krieg T, Chen M, Mogul DJ, Miller LE (2010) Optimal spacing of surface electrode arrays for brain-machine interface applications. J Neural Eng 7(2): 26004.

Smith MA, Kohn A (2008) Spatial and temporal scales of neuronal correlation in primary visual cortex. J Neurosci 28: 12591-12603.

Sommer MA, Wurtz RH (2000) Composition and topographic organization of signals sent from the frontal eye field to the superior colliculus. Journal of Neurophysiology 83:1979-2001.

Sparks DL, Mays LE (1990) Signal transformations required for the generation of saccadic eye movements. Annl Rev Neurosci 13: 309-336.

Stark E, Abeles M (2007) Predicting movement from multiunit activity. J Neurosci 27: 8387-94.

Stark E, Drori R, Abeles M (2009) Motor cortical activity related to movement kinematics exhibits local spatial organization. Cortex 45: 418-31.

Stence N, Waite M, Dailey ME (2001) Dynamics of microglial activation: A confocal time-lapse analysis in hippocampal slices. Glia 33(3): 256-66.

Strupler M, Pena AM, Hernest M, Tharaux PL, Martin JL, Beaurepaire E, Schanne-Klein MC (2007) Second harmonic imaging and scoring of collagen in fibrotic tissues. Opt Express 15(7): 4054-65.

Su J, Todorov M, Gonzales HP, Perkins L, Kojouharov H, Weng H, Tang L (2011) A Predictive Tool for Foreign Body Fibrotic Reactions Using 2-Dimensional Computational Model. Open Access Bioinformatics 3: 19-35.

Suminski AJ, Tkach DC, Fagg AH, Hatsopoulos NG (2010) Incorporating feedback from multiple sensory modalities enhances brain-machine interface control. J Neurosci 30:16777-16787.

Suner S, Fellows MR, Vargas-Irwin C, Nakata GK, Donoghue JP (2005) Reliability of signals from a chronically implanted, silicon-based electrode array in non-human primate primary motor cortex. IEEE Trans Neural Syst Rehabil Eng 13: 524–41.

Szarowski DH, Andersen MD, Retterer S, Spence AJ, Isaacson M, Craighead HG, Turner JN, Shain W (2003) Brain responses to micro-machined silicon devices. Brain Res 938(1-2): 23-35.

Tallon-Baudry C, Bertrand O, H´enaff MA, Isnard J, and Fischer C (2005) Attention modulates gamma-band oscillations differently in the human lateral occipital cortex and fusiform gyrus. Cereb Cortex 15(5): 654-62.

Taylor DM, Tillery SI, Schwartz AB (2002) Direct cortical control of 3D neuroprosthetic devices. Science 296: 1829-32.

Todorova S, Sadtler P, Batista A, Chase S, Ventura V (2014) To sort or not to sort: the impact of spike-sorting on neural decoding performance. J Neural Eng 11: 056005.

Trautner P, Rosburg T, Dietl T, Fell J, Korzyukov OA, Kurthen M, Schaller C, Elger CE, Boutros NN (2006) Sensory gating of auditory evoked and induced gamma band activity in intracranial recordings. Neuroimage 32(2): 790-798.

Turner JN, Shain W, Szarowski DH, Andersen M, Martins S, Isaacson M, and Craighead H (1999) Cerebral astrocyte response to micromachined silicon implants. Exp Neurol 156(1): 33-49.

Tyler M, Danilov Y, Bach-Y-Rita P (2003) Closing an open-loop control system: vestibular substitution through the tongue. Jof Integ Neurosci 2(2):159-164.

Tyler-Kabara, EC (2014) Brain Computer Interface (BCI) Controlled prosthetic arm movement is possible in the absence of visual input with proprioceptive feedback. 82nd Annual Meeting of the American Association of Neurological Surgeons. Tuesday, April 8, 2014.

Vallbo AB, Hagbarth K, Torebjork H, Wallin B (1979) Somatosensory, proprioceptive, and sympathetic activity in human peripheral nerves. Physiol Rev 59(4):919-957.

Van Essen DC, Maunsell JHR (1983) Hierarchical organization and functional steams in the visual cortex. Trends Neurosci 6(9):370-375.

Van Gompel JJ, Worrell GA, Bell ML, Patrick TA, Cascino GD, Raffel C, Marsh WR, Meyer FB (2008) Intracranial electroencephalography with subdural grid electrodes: techniques, complications, and outcomes. Neurosurgery 63(3): 498-505.

Vaughan TM, Heetderks WJ, Trejo LJ, Rymer WZ, Weinrich M, Moore MM, Kubler A, Dobkin BH, Birbaumer N, Donchin E, Wolpaw EW, Wolpaw JR (2003) Brain-computer interface technology: a review of the second international meeting. IEEE Trans Neural Syst and Rehabil Eng 11(2):94-109.

Vela JM, Dalmau I, Gonzalez B, Castellano B (1995) Morphology and distribution of microglial cells in the young and adult mouse cerebellum. J Comp Neuro 361: 602-16.

Velliste M, Perel S, Spalding MC, Whitford AS, Schwartz AB (2008) Cortical control of a prosthetic arm for self-feeding. Nature 453: 1098-101.

Venkatraman S, Carmena, JM (2011) Active sensing of target location encoded by cortical microstimulation. IEEE Trans Neural Syst Rehabil Eng 19(3): 317-24.

Ventura V (2008) Spike train decoding without spike sorting. Neural Comput 20: 923-63.

Wall C, Kentala E (2010) Effect of displacement, velocity, and combined vibrotactile tilt feedback on postural control of vestibulopathic subjects. J Vestib Res 20(1,2):61-69.

Wang W, Collinger JL, Degenhart AD, Tyler-Kabara EC, Schwartz AB, Moran DW, Weber DJ, Wodlinger B, Vinjamuri RK, Ashmore RC, Kelly JW, Boninger ML (2013) An electrocorticographic brain interface in an individual with tetraplegia. PLoS ONE 8(2): e55344.

Wang W, Degenhart AD, Collinger JL, Vinjamuri R, Sudre GP, Adelson PD, Holder DL, Leuthardt EC, Moran DW, Boninger ML, Schwartz AB, Crammond DJ, Tyler-Kabara EC, and Weber DJ (2009) Human motor cortical activity recorded with Micro-ECoG

electrodes, during individual finger movements. Conf Proc IEEE Eng Med Biol Soc 2009: 586-9.

Wang W, Degenhart AD, Sudre GP, Pomerleau D, Tyler-Kabara EC (2011) Decoding semantic information from human electrocorticographic (ECoG) signals. Conf Proc IEEE Eng Med Biol Soc 2011: 6294-8.

Wang Y, Papadimitrakopoulos F, Burgess DJ (2013) Polymeric "smart" coatings prevent foreign body response to implantable biosensors. J Control Release 169(3): 341-347.

Weaver CL, LaRosa JM, Luo X, Cui XT (2014) Electrically Controlled Drug Delivery from Graphene Oxide Nanocomposite Films. ACS Nano. 8(2): 1834-1843.

Weber DJ, London BM, Hokanson JA, Ayers CA, Gaunt RA, Torres RR, Zaaimi B, Miller LE (2011) Limb-state information encoded by peripheral and central somatosensory neurons: implications for an afferent interface. IEEE Trans Neural Syst Rehabil Eng 19(5):501-13.

Weber DJ, Friesen R, Miller LE (2012) Interfacing the somatosensory system to restore touch and proprioception: essential considerations. J Mot Behav 44(6):403-18.

Weinrich M, Wise SP (1982) The premotor cortex of the monkey. J Neurosci 2(9): 1329-1345.

Wessberg J, Stambaugh CR, Kralik JD, Beck PD, Laubach M, Chapin JK, Kim J, Biggs SJ (2000) Real-time prediction of hand trajectory by ensembles of cortical neurons in primates. Nature 408: 361-5.

Westling G, Johansson RS (1987) Responses in glabrous skin mechanoreceptors during precision grip in humans. Exp Brain Res 66(1): 128-140.

Williams JC, Rennaker RL, Kipke DR (1999) Long-term neural recording characteristics of wire microelectrode arrays implanted in cerebral cortex. Brain Res Protoc 4:303–13.

Williams JC, Rennaker RL, Kipke DR (1999) Stability of chronic multichannel neural recordings: implications for a long-term neural interface. Neurocomp 26–27:1069–76.

Wilson JA, Felton EA, Garell PC, Schalk G, Williams JC (2006) ECoG factors underlying multimodal control of a brain-computer interface. IEEE Trans Neural Syst Rehabil Eng 14(2): 246-50.

Winslow BD, Christensen MB, Yang WK, Solzbacher F, Tresco PA (2010) A comparison of the tissue response to chronically implanted Parylene-C-coated and uncoated planar silicon microelectrode arrays in rat cortex. Biomaterials 31: 9163−9172.

Wise SP, Boussaoud D, Johnson PB, Caminiti R (1997) Premotor and parietal cortex: cortico-cortical connectivity and combinatorial computations. Annu Rev Neurosci 20: 25-42.

Wolpaw JR, McFarland DJ, Vaughan TM (2000) Brain-computer interface research at the wadsworth center. IEEE Trans Rehabil Eng 8(2):222–226.

Wolpaw JR, Birbaumer N, McFarland DJ, Pfurtscheller G, Vaughan TM (2002) Brain-computer interfaces for communication and control. Clin Neurophysiol 113(6): 767-91.

Wong CH, Birkett J, Byth K, Dexter M, Somerville E, Gill D, Chaseling R, Fearnside M, Bleasel A (2009) Risk factors for complications during intracranial electrode recording in presurgical evaluation of drug resistant partial epilepsy. Acta Neurochir 151(1): 37-50.

Wood F, Black MJ, Vargas-Irwin C, Fellows M, Donoghue JP (2004) On variability of manual spike sorting. IEEE Trans Biomed Eng 51: 912-8.

Wu W, Gao Y, Bienenstock E, Donoghue JP, Black MJ (2006) Bayesian population decoding of motor cortical activity using a Kalman filter. Neural Comp 18:80-118.

Xu HT, Pan F, Yang G, Gan WB (2007) Choice of cranial window type for in vivo imaging affects dendritic spine turnover in the cortex. Nat Neurosci 10: 549-551.

Yamakawa T, Yamakwa T, Aou S, Ishizuka S, Suzuki M, and Fujii M (2010) Subdural electrode array manipulated by a shape memory allow guidewire for minimally-invasive electrocorticogram recording. World Automation Congress 2010: 1-6.

Yanagisawa T, Hirata M, Saitoh Y, Kishima H, Matsushita K, Goto T, Fukuma R, Yokoi H, Kamitani Y, Yoshimine T (2012) Electrocorticographic control of a prosthetic arm in paralyzed patients. Ann Neurol 71(3): 353-361.

Yeager JD, Phillips DJ, Rector DM, Bahr D (2008) Characterization of flexible ECoG electrode arrays for chronic recording in awake rats. J Neurosci Methods 173(2): 279-85.

Young FA, Spector M, Kresch CH (1979) Porous titanium endosseous dental implants in rhesus monkeys: microradiography and histological evaluation. J Biomed Mater Res 13:843.

Zhai S, Milgram P (1993) Human performance evaluation of isometric and elastic rate controllers in a 6 dof tracking task. Proceedings of SPIE Telemanip Tech 2057:1-12.