INVESTIGATION OF METHODS FOR ASSESSING SENSORIMOTOR PERFORMANCE IN HUMANS AND MONKEYS

by

Kristin M. Quick

BS, Rose–Hulman Institute of Technology, 2009

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

2015

UNIVERSITY OF PITTSBURGH

SWANSON SCHOOL OF ENGINEERING

This dissertation was presented

by

Kristin M. Quick

It was defended on

July 7, 2015

and approved by

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

Patrick Loughlin, Ph.D, William Kepler Whiteford Professor of Bioengineering, Department of Bioengineering

Byron Yu, Ph.D, Departments of Electrical and Computer Engineering and Biomedical Engineering Carnegie Mellon University

Rakie Cham, Ph.D, Associate Professor, Department of Bioengineering

Douglas Weber, Ph.D, Associate Professor, Department of Bioengineering

Stephen Scott, Ph.D, Professor, Department of Biomedical and Molecular Sciences, Queen's University

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

INVESTIGATION OF METHODS FOR ASSESSING SENSORIMOTOR PERFORMANCE IN HUMANS AND MONKEYS

Kristin M. Quick, PhD

University of Pittsburgh, 2015

The sensorimotor system's beauty becomes evident when the system is put to the test. In order to understand the neural and physiological mechanisms of motor control, researchers need tasks where the task difficulty can be adjusted. By studying the interplay between difficulty and the motor parameter of interest, we can begin to tease apart the inner-workings of how the sensorimotor system is able to act in the face of sensorimotor delay, perform macro- and micro-scale movements, and learn to perform new tasks.

We tested the ability of human and monkey subjects to perform a new task, called the Critical Stability Task. In this task, we can pressure the sensorimotor loop to act as quickly as possible, in order to look for deficiencies in the motor or sensory components of control. Additionally, the difficulty of this task is decided by a single number, which enables the subject's sensorimotor performance to be evaluated by a single number. This task will be valuable for testing and evaluating new forms of sensory substitution, new non-visual feedback methods for brain-computer interface control, as well as serve as a tool for neurological assessment.

While using this task in monkey subjects, it became apparent how little we know about the control of very small movements. In the Critical Stability Task, the majority of the movements are under +/-10mm, and only exceed this window once control becomes difficult. As such, we began an investigation of the neural representation of small amplitude movements.

Finally, we tested the neural constraints of learning a new task. We setup two different types of perturbed brain-computer interface (BCI) decoders. The first type maintained the prior patterns of neural activity, the second type required the monkey to generate new patterns of neural activity. We found that within a day, it was difficult for monkey subjects to learn to generate new patterns of neural activity. Taken together, task difficulty is a powerful tool for assessing sensorimotor performance. The only way to *improve* poor motor performance, be it natural control or BCI control, is to *detect* poor motor performance with the right assessment tool.

TABLE OF CONTENTS

LIST OF FIGURES

PREFACE

Much of the work described here is the result of trying to scratch an itch. I believe that the best thinking comes about when trying to solve a problem, or questions, that are truly just bothering you. Of primary importance, I must thank my advisor, Dr. Aaron Batista, who let me scratch my itches. He gave me the freedom, support, and confidence and the means to try to answer the questions that were bothering me. He thankfully provided me with the level-headed guidance, ability to more clearly articulate my thoughts into a comprehensible statements, and informed me of which itches to scratch, and which were better left alone for another day. I cannot thank him enough for his genuine good-nature and care for everyone in his lab, and his always asking the big questions and dreaming big dreams. His big dreams are what pulled me to the lab in the first place (particularly a statement about 'over-clocking the brain'), and my experiences lived up to, and surpassed, any hopes that I had. His desire to innovate in the field enables him to harbor a one-of-a-kind lab which harbors open-thinking and creates a support system where students can become confident scientists. I wish him my best, and I hope that he'll keep letting me visit and see what new ideas he has brewing.

I would also like to acknowledge Dr. Patrick Loughlin for his help on control theory and helping me understanding the understand the properties of a control system. In addition to illuminating control theory in a clear manner, he also gave me a better framework for understanding how human and monkey subjects may use vibrotactile feedback. Additionally, he was extremely important to me and to the success of the Critical Stability Task because he took me under his wing, and supported the project in it's early stages when most others quickly dismissed the idea. I cannot convey how important this was to me and to the success of the project.

I would also like to thank Dr. Byron Yu for his unfailing patience with me and always asking the hard questions. Byron is the best teacher any student could ask for. His logical thinking enabled me to more easily step from simple concepts to more complex theories and practices. His questions always served as a way to delve deeper into understanding the root of any problem. He would continue to ask me questions when I was being unclear, long after others have thrown-in the towel. Because of his persistence and caring nature, I was able to finally hone-in on a set of explanations for the Critical Stability Task that I could take to others to bring them up to speed on the objectives of the task.

Next, I would like to thank Dr. Stephen Scott, first for putting up with all of the travel disruptions he experienced in order to serve on my committee. This was no small feat; there was always a hidden disaster. Second, his work on the motor system's use of feedback gave me courage and inspiration. His belief in the importance of feedback helped me to believe that I was not crazy for also thinking feedback was extremely important for control. He has led me to think in new and different ways. It has been an extreme privilege to have received his viewpoints and ideas on my work. His contributions will be utilized and felt for years ahead.

Importantly, I would also like to thank my lab-mates from Aaron's lab. Dr. Patrick Sadtler was indispensable for helping me complete every aspect of my projects in the lab. Whether it was a technical glitch that left me scratching my head or an equation I couldn't solve, or letting me take the last piece (ok, pieces) of chocolate, nothing would have been possible without him. I would also like to give a hearty thank-you to Jessi Mischel. Her enthusiasm, clear-thinking, and hard work helped me to tackle the Critical Stability Task project. In her, I had one of my greatest allies. Jessi, you're awesome. She spent many hard months, with few days off, training monkeys to do the CST. I could not have done any of this work without Jessi's help. I am so happy to be handing-off the project to such a capable and talented person.

I would also like to thank those who helped me in all of my crazy projects while they were undergraduate students in the lab, including Nicholas Card, Ivan Smalianchuk, Stephen Whaite, and Frank Petraglia. They all helped me through some of the toughest spots, including endless analyses with data, ill-tempered monkeys, parameter sweeps, and generally,

just putting up with me. You helped me troubleshoot so many problems (even though neither of us wanted to), I could not have collected or analyzed the amount of data that we did without your help. Thank you so much.

Additionally, other Batista lab members helped me to better hone my thoughts, fix problems in the rig or with analysis code, and deal with the problems inherent to life in graduate school. I would like to thank Dr. Jason Godlove, Berook Alemayehu, Dr. Emily Oby, Dr. Alan Degenhart, Nick Pavlovsky, Jeffery Chiou, Nicole McClain, Jocelyn Stafura, and Melissa James. Every computer fix, monkey training tip, PhaseSpace calibration session, and presentation and research suggestion was taken to heart and greatly appreciated. I couldn't have asked to be a part of a better team of people.

Finally, I need to thank my family, for hanging in there with me for so many years. Your love and support was essential for me to 'stick with it.' With each frustration I experienced, my family motivated me to find a way to solve the problem, keep my head up, and continue moving forward. First, I would like to thank my husband, Chris Quick, for dealing with me through thick and thin. It was great to be able to come home and continue to bounce ideas off of someone, and you always brought me a new perspective. Thank you for keeping my spirits high, I could not have done this work without your support. I would also like to thank my parents, Pam and Chris Wilson. They had been telling me to go to graduate school since I was in high school. Thankfully, I listened to them my senior year in college. When the going got tough, they lent a supporting hand, ear, eye, or brain, to whatever my concern was. This dissertation is for the best parents in the whole world by a wide margin.

I would also like to thank my sisters, Sarah, Valerie, Melissa, and Samantha, for being understanding of my sometimes strange schedule and times when I was frustrating to deal with. I would also like to thank my second family, Dee and Marty Quick, and Jon and Courtney. I know I bombarded you with plenty of my frustrations. Thank you for your faith in me and your loving support.

Finally, I would like to thank Dr. Allen and Cheryl White for their belief that I could do this, and encouraging me to persevere. It was very valuable for me to be able to talk with someone who had been in my shoes, but was in an outside field with a different perspective. From the bottom of my heart, thank you all very much.

1.0 GENERAL INTRODUCTION

All living organisms use movement to interact with the world, motion is necessary for nearly every aspect of life. Starting from birth, we are constantly testing and discovering our environment and ourselves through motion. It is through motion that we acquire sensations of physical properties of objects, such as texture and shape. Most importantly, we use motion to manipulate objects. At it's most basic level, motion is made up of the application of a force, and the result of that force on an object. Humans complete the cycle by sensing the object's reaction to that force.

While sensory processing is most commonly thought of as a passive process, it is actually tightly linked to motion. One primary purpose of motion is to acquire information through active sensing (Schroeder et al. 2010). For example, the visual system does not sit idle and wait for a meaningful stimulus. Instead, we seek out visual feedback by moving the eyes to specific locations to collect the most meaningful information (Rayner and McConkie 1976, McAuley et al. 1999, Andrews and Coppola 1999, Mennie et al. 2007). Furthermore, we do not passively wait for objects to interact with our fingers. Instead, we seek out tactile feedback by choosing the exploratory motion to identify an object's shape and texture (Johansson and Westling 1987, Birznieks et al. 2001, Jenmalm et al. 2003, Johansson and Flanagan 2009).

The second primary purpose of motion is to manipulate objects for an intended purpose. Whether the object is a tool to use or food to eat, we interact with the object in a specific manner. Behind the scenes when we manipulate an object, we understand both the object's position and orientation and our own. This understanding of position and orientation is called proprioception. Proprioceptive feedback is integral to successfully move to and interact with an object (Johansson and Cole 1992, Gordon et al. 1995, Sober and Sabes 2003, Peterka 2002, London and Miller 2013).

In summary, visual, tactile, and proprioceptive feedback are at the heart of motor control. Together, they complete a sensory-motor loop which enables the exploration and manipulation of our environment. In this work, I address methods to study and assess sensorimotor control. As common as movement is to everyday life, there few tasks which can determine the effectiveness someone's sensorimotor control. People have a fine appreciation for excellent sensorimotor control, evident by our fascination with professional athletes, artists, and musicians. However, there are few scientific metrics for measuring the sensorimotor control required to perform these tasks.

These metrics are needed for natural sensorimotor control as well as brain-computer interface control. Ideally, the same assessment metrics could be used. BCI control is a from of impaired motor control. In Chapter 3, I show how my task, the Critical Stability Task, can be used to assess sensorimotor performance in human subjects across two different motor control methods and two different feedback methods. In Chapter 4, I show how the same task, the CST, can be modified for use with monkey subjects to assess natural motor control and BCI control, as well as demonstrating monkeys can learn to use an abstract, non-visual form of feedback. In Chapters 5, 6, and 7, I show neural constraints on using a brain-computer interface decoder, constraints that should be considered in the future when trying to design better BCI decoders.

In the following sections, I will discuss three types of feedback which are integral to successful actions (Section 1.1): visual, proprioceptive, and tactile feedback. Next, I will review the current state of motor prosthetics and sensory prosthetics (Section 1.2).

1.1 NATURAL SENSORIMOTOR CONTROL

1.1.1 Visual feedback

For humans, visual feedback is a primary mode for gathering information. The flow of visual information begins in the retina, travels through the lateral geniculate nucleus (LGN) of the thalamus before entering primary visual cortex (V1) (Callaway 1998). Once entering cortex, it undergoes a tremendous processing by its many feedforward and feedback connections with other cortical regions (Van Essen and Maunsell 1983). One function of processing is to guide and command subsequent eye movements, or saccades. The superior colliculus is a major node in the sensorimotor transformation from visual information to saccades (Sparks and Mays 1990, Hall and Moschovakis 2004). The superior colliculus lies on the roof of the brain stem. The superficial layers respond almost entirely to visual information. By contrast, the deep layers produce a burst of neural activity just prior to a contraversive saccade (Gandhi and Katnani 2011). After the saccade is complete, a new set of visual information is read into the retina so that the process may begin again.

Visual information is also used to control reaches. Commonly, people will look to the object with which they intend to interact. After fixating on the object, the object is now represented in a visual coordinate frame. In order for this information to be useful to the motor system, the visual information must be transformed into a coordinate frame which is meaningful for reaches. In important region for this transformation is the parietal reach region (PRR) of the posterior parietal cortex. Here, the visual location of intended movement is represented in visual coordinates (Batista et al. 1999, Buneo et al. 2002). Evidence of the sensorimotor transformation is seen when PRR changes it's motor plan based on context of the task (Gail and Andersen 2006) which indicates that the neural activation is not purely visual information. In addition to PRR, visual information is also used to a certain extent in dorsal premotor cortex (PMd) (Pesaran et al. 2010). It is not yet quite clear what kind of a role PMd plays in this sensorimotor transformation, however, it is well poised to do so.

PMd responds to visual stimuli (Wise et al. 1997, Fogassi et al. 1999), receives connections from parietal cortex and has direct connection to the primary motor cortex (M1) (Wise et al. 1997, Dum and Strick 2002).

An important component of visually-guided reaching is amount of time it takes to sense, process, and act upon the information. The reaction time to visually guided movements can be variable, but ranges between 163-222ms (Diederich and Colonius 2004, Godlove et al. 2014). Part of this time includes the response execution, which may explain some of the variability. Another component of the reaction time is the delay between a visual stimulus and area V1. It was found that this delays is approximately 77ms (Nowak et al. 1995). It is likely that the remainder of the reaction time delay is processing time to prepare the motor plan. Having a single planned motor action in response to the stimulus helps to reduce the reaction time (Surwillo 1973, Miller and Low 2001).

1.1.2 Tactile feedback

We rely on tactile information to interact with and manipulate objects. Patients with impaired sensory capacity of the median nerve experience frustration at their motor deficiencies and inability to hold objects (Moberg 1962). While there are many types of receptors cell in the skin, for this short discussion we will discuss three, the Meissner's corpuscle, Merkel disk receptors, and Pacinian corpuscles. The Meissner's corpuscle and the Merkel disk receptor are located in the superficial layers of the skin. The Meissner's corpuscle is a rapidly adapting mechanoreceptor which conveys fine detail sensations, such as flutter. The Merkel disk receptor is a slowly adapting receptor better suited to detecting large, lower temporal frequency details. Because they are superficial, however, the have small receptive fields. This enables them to convey details with good spatial resolution. Finally, the Pacinian corpuscle is in the deep tissue layer. Like the Meissner's corpuscle, the Pacinian corpuscle quickly responds to a skin indentation. However, these deeper, larger receptors respond to a larger receptive field. As a result, they are better suited to detecting macro-scale vibrations (Johansson 1978, Hulliger et al. 1979, Westling and Johansson 1987).

The tactile signals are conveyed to the central nervous system via the dorsal root ganglion neurons, the spinal cord, thalamus, and finally the cerebral cortex. As with visually-guided movements, the reaction time is an important indicator of processing delays. Tactile reaction times can vary between 177-184ms (Diederich and Colonius 2004, Godlove et al. 2014). Thus, they are on the same scale as visual reaction times, but they are commonly thought of as being a slightly faster way to transmit information.

1.1.3 Proprioceptive feedback

Proprioceptive feedback provides us with a sense of the position and velocity of our limbs. 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). Muscle spindle fibers in the muscles sense the length of the muscle and the speed at which it is expanding or contracting (Kandel et al. 2000, Scott and Loeb 1994). Golgi tendon organs, in series with the muscle, provides information on the muscle tension (Vallbo et al. 1979). Joint capsule mechanorecpetors provide kinematic information on the joint angles (Skoglund 1973). When these proprioceptive sensors fail, motor deficits are clearly seen. Patients who have lost proprioceptive sensation due to large-fiber neuropathy are still able to perform movements with visual feedback (Ghez et al. 1995). However, patients make large movement errors and reach with a curvature.

1.2 SENSORIMOTOR NEURAL PROSTHETICS

Every year, people suffer from a disease or injury which causes paralysis. To combat paralysis, much research has gone into brain-computer interfaces (BCI) which translate a person's neural activity into the movement of assistive devices, such as a computer cursor or a robot arm (Taylor et al. 2002, Carmena et al. 2003, Velliste et al. 2008, Cherian et al. 2011, Collinger et al. 2013). A BCI user gathers information on the device's movements via visual feedback

only. In contrast, able-bodied people gather information about their movements using both visual and non-visual forms of feedback. To test whether the re-addition of proprioceptive feedback would be helpful for BCI users, researchers had monkey subjects receive congruent proprioceptive and visual feedback during BCI control. The congruent feedback significantly improved the BCI performance over visual feedback alone (Suminski et al. 2010).

As a way to provide an alternate form of feedback, intracortical microstimulation could be used to supply abstract tactile or proprioceptive signals. Studies on primary somatosensory cortex (S1) microstimulation in monkeys (Mountcastle et al. 1990, Romo et al. 1998, Fitzsimmons et al. 2007, London et al. 2008, O'Doherty et al. 2011, Dadarlat et al. 2015) have initiated and guided research towards microstimulation-based feedback for BCI use. To this point in time, most artificial feedback methods are a proof-of-concept. What is needed is a method to push a subject's performance to the limits while using a particular form of feedback. By increasing the difficulty of the tasks, we can truly begin to test out which methods are effective.

2.0 GENERAL METHODS

In this chapter, I will describe the methods used throughout the experiments in Chapters 3-7. In Section 2.1, I describe the electrophysiology data recorded from all the non-human primates used in the following experiments. In Section 2.2 and 2.3, I describe the motor control methods and feedback methods used by human and non-human primate subjects. In Chapter 3, data was collected from able-bodied human subjects in an experiment approved by the University of Pittsburgh Institutional Review Board (IRB). In Chapters 4-7, data was collected from non-human primates in experiments approved by the University of Pittsburgh Institutional Animal Care and Use Committee (IACUC).

2.1 ELECTROPHYSIOLOGY DATA

The data in Chapters 4, 5, 6, and 7 were recorded from five male Rhesus macaques (*Macaca mulatta*), monkey I, monkey J, monkey K, monkey L, and monkey N. Monkey I was studied in Chapters 4 and 5, monkey J was studied in Chapters 4 and 6, monkey K was studied in Chapter 7, monkey L was studied in Chapters 5 and 6, and monkey N was studied in Chapters 4 and 5. In all monkeys except monkey K, a 96-channel microelectrode array (Blackrock Microsystems) was implanted into cortex. For monkey I, the array was implanted into the right hemisphere of dorsal premotor cortex (PMd), monkey J was implanted into the right hemisphere of primary motor cortex (M1), and monkeys L and N were implanted into the left hemisphere of M1. For monkey K, a custom 15-channel silicone micro-electrocorticography grid with 2mm platinum electrodes (approximately 2cm by 2cm) was placed over left hemisphere of premotor and motor cortex.

2.2 MOTOR CONTROL METHODS

There were two methods of motor control used throughout the experiments of Chapters 3-6. One method, used by humans and monkeys, was hand control. During hand control, subjects freely moved their hand in space in order to accomplish the task objectives. For human subjects, we attached a red LED to the index finger of their dominant hand. For monkey subjects, we attached a red LED to either the monkey's wrist (monkey I) or the monkey's index finder (monkeys J, L, N). The position of the LED marker was then recorded in three dimensions (3D) using a PhaseSpace, Inc. motion capture system.

The second method, used by monkeys I, J, and L, was brain-computer interface control, termed 'BCI control' or 'brain control.' During BCI control, the monkeys controlled a cursor's movements using their neural activity. 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. For monkeys I (Chapter 4) and L (Chapter 5), a velocity Kalman filter decoder was used to map the spike counts to cursor velocity. For monkeys J and L (Chapter 6), a velocity Kalman filter using factor analysis was used. For both decoders, the decoded velocity signal was integrated (over the bin length for which the spike were counted) into a position signal and displayed as the brain-controlled cursor. Details on the decoding algorithms can be found in Chapters 5 and 6.

2.3 MOTOR FEEDBACK METHODS

There were two methods of feedback used throughout the experiments of Chapters 3-6. The primary method was visual feedback. Human subjects, while seated comfortably in a chair, received visual information from a computer monitor. During many experiments, the subject's hand position was visually rendered on the monitor as a round moving cursor. 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 both human subjects in Chapters 3 and monkeys in Chapter 4 was vibrotactile feedback. For vibrotactile feedback, two coin tactors (Precision Microdrives Limited) were attached to the human or monkey to communicate information in a non-visual manner. For human subjects, the tactors were attached to the non-working hand. One tactor was attached to the subject's thumb and one was attached to the subject's pinkie finger. During the experiment, the subject would hold their hand with their palm facing downward. In this fashion, one tactor would vibrate to convey movements to the right and the other would convey movements to the left. For monkey subjects, the tactors were placed on the non-working forearm. The tactors were too large to place on the monkeys fingers. In a similar manner, one tactor would vibrate to indicate movements to the right, and the other would vibrate to indicate movements to the left. The amount of tactor vibration was controlled via a supply voltage. As the supply voltage increased, the frequency and magnitude of vibration increased. An accelerometer was fixed to the back of each tactor to measure the vibration.

3.0 HUMAN PERFORMANCE ON CRITICAL STABILITY TASK

Figures and text in this chapter are from Quick et al. 2014. In this chapter, I describe my research on human subjects performing the Critical Stability Task using two different control methods and four different feedback methods. This work created a foundation for the monkey work in later chapters, 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. This was the foundational work to test the Critical Stability Task as a tool for assessing natural sensorimotor control. In order to test human subjects, I developed our lab's first protocol for the Institutional Review Board (IRB) of the University of Pittsburgh, for which I served as the Principle Investigator.

3.1 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 where the user can interact with an object, we designed the Critical Stability Task (CST), which is based on the Critical Tracking Task (Jex et al. 1966). The Critical Tracking Task has been used to assess motor performance during drug use and tele-operation, 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 stabilize a first order unstable linear system. A familiar example of an unstable first order system is compounding interest; the debt grows exponentially over time, and the larger the interest rate, the faster the debt grows. In the absence of external factors (i.e. payments), the account balance can be modeled mathematically by $y(t) = y_0 e^{\lambda t}$ where y_0 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.

$$G(s) = \frac{\lambda_k}{s - \lambda_k} \tag{3.1}$$

$$x(k+1) = (e^{\lambda_k T})x(k) + (e^{\lambda_k T} - 1)u(k) \tag{3.2}$$

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

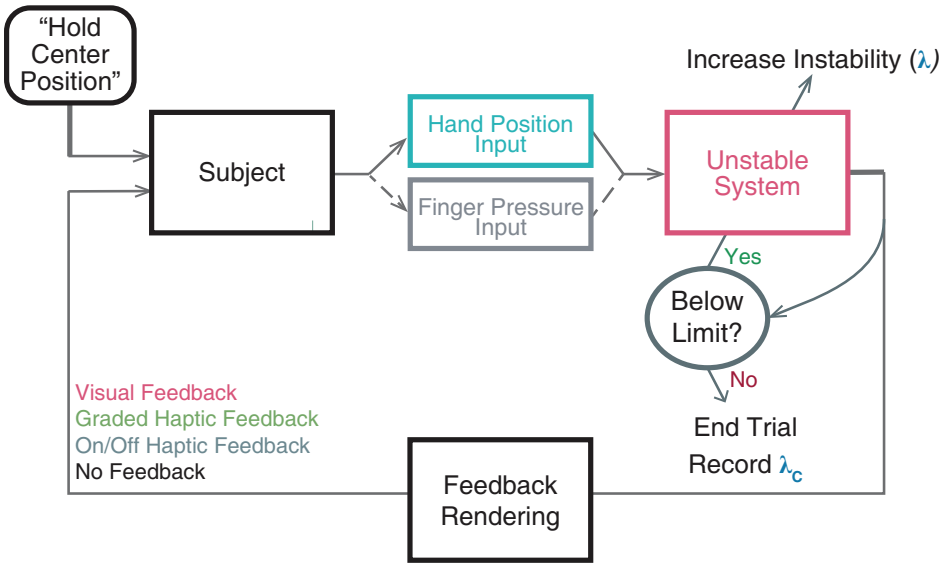

Figure 3.1: Diagram of the Critical Stability Task as used by human subjects. Subjects are told to attempt to keep their control signal, either their hand position or their finger pressure represented by position, on the center target. Their input signal serves as the input to the unstable system, which has it's own internal dynamics. After each loop iteration, the instability of the unstable system's dynamics increases, which makes the system more difficult to control over time. The updated unstable system's position is compared with a position threshold. If the unstable system's position crosses either the positive or negative threshold, the trial is terminated and the last instability is recorded as the critical instability. If the unstable system's position is below the threshold, the then position is rendered to the subject through either visual or vibrotactile feedback.

method, the quality of the feedback rendering was assessed using the CST. Feedback renderings that allowed subjects to control greater instability 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 graded 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).

3.2 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 systems position, Figure 3.1.

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.

3.2.1 Critical Stability Task implementation

Subjects controlled an unstable system $G(s)$, as seen in Equation 3.1, where instability $\lambda(k) > 0$ increases at a constant rate of 0.10 rad/s 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, Equation 3.2. 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 systems position surpassed a predetermined threshold, in our case, $\pm$ 50mm from the center. As a note, λ_C is robust to different thresholds [1]; 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.

3.2.2 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, as seen in Figure 3.2.

3.2.3 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

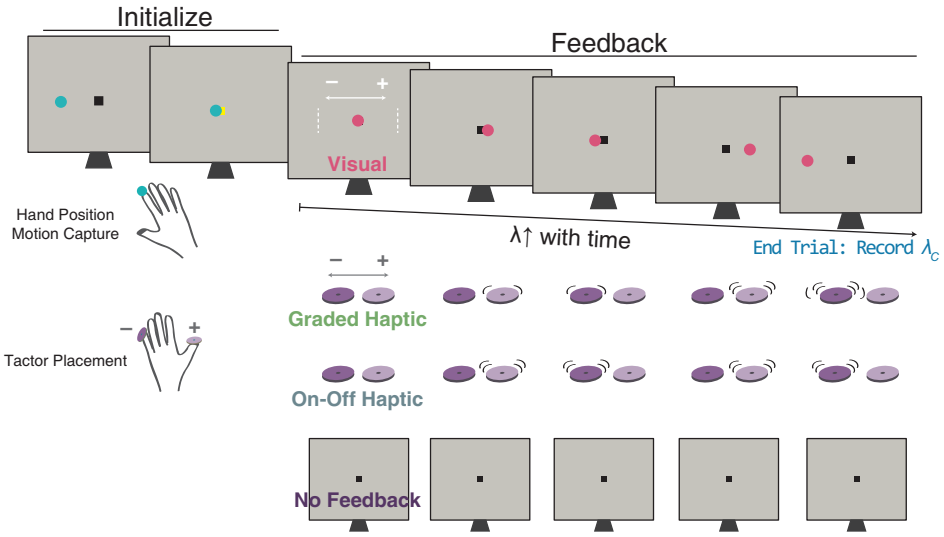

Figure 3.2: A depiction of the four feedback methods used by humans. Subject initiated each trial using visual feedback of their control signal. After a hold period, the CST began, which was indicated by a change in the cursor's color during visual feedback. During hand control, subject's hand an LED marker placed on their dominant hand's index finger. On the non dominant hand, the positive tactor was attached to the digit on the right, when the palm was facing down. The left tactor was placed on the furthest left digit. During graded vibrotactile feedback, the tactors would vibrate in proportion to the system's distance from the center. During on/off vibrotactile feedback, the appropriate tactor would vibrate at the highest intensity, regardless of the system's distance from the center.

was displayed on a monitor approximately 1m in front of the subject. When the subject used hand movements to control the system, the systems position was displayed along the horizontal axis. When the subject used pinch force to control the system, the systems position was displayed along the vertical axis, to make the visual feedback more congruent with the control scheme.

The second feedback method was graded vibrotactile feedback. The systems 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. In Figure 3.2, the graded nature of the vibrotactile feedback is indicated by the number of vibration lines. More vibration lines indicate that the tactor is vibrating at a higher frequency and amplitude.

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 tactors intensity using a command voltage that was proportional to the systems position. The tactors 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 50-175Hz.

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 pinkie tactor would vibrate at a fixed intensity for all negative deviations. In Figure 3.2, the on/off vibrotactile feedback is indicated by the constant number of vibration lines when the tactor is on. The fourth feedback method was no feedback. Subjects had to attempt to control the system without being given its current position. For

these trials, they remained seated comfortably in their chairs and were free to look about the room. However, a black square target was on the screen to indicate that a trial was occurring.

3.2.4 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 subjects 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 λ_C using data from all subjects.

3.3 RESULTS

3.3.1 Visual feedback results

Figure 3.3 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 rad/s) rather than hand control (2.83 ± 0.55 rad/s) 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 ($p = 0.009$). The instability starts off at zero, and the subject makes very small compensatory movements. As the instability increases during the trial, the movements get larger, as seen in Figure 3.4. For hand control, subjects had a tendency to stay at a particular position for a period of time, rather than move continuously. At the end of the trial, the subject makes a large rightward correction, anticipating the the system will soon cross the center target. However, this motor decision was made too soon, because the system had not yet cross the center. For

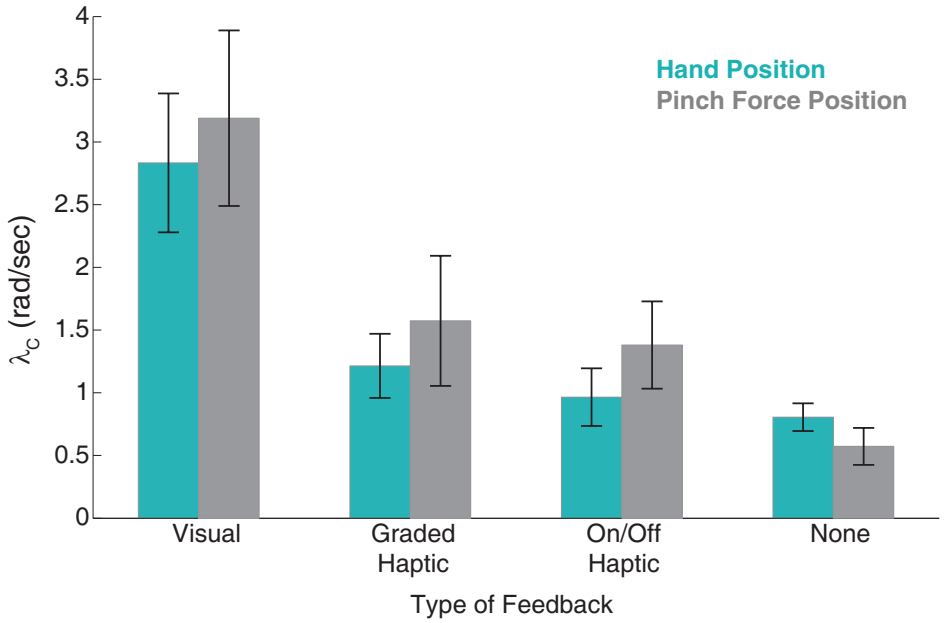

Figure 3.3: Human performance scores on the CST combined across subjects. The critical instability for each subject, for each of the eight feedback-control methods, was averaged together to determine the critical instability for each feedback-control method. Error bars indicate the standard deviation across subjects. The results from when the subjects were using their hand position are in blue, and the results from using pinch force are in gray. Vibrotactile control outperformed hand position control in all methods except the No Feedback condition.

this reason, he amplified the system's current position, and it quickly crossed the positive threshold. At the beginning of the trial, the finger pressure movements are slightly larger than subject P's hand position movements. After about fifteen seconds into the trial, subject D begins making much larger movements. These movements do not stay in the same place long, as seen by the trajectories that have more peaks in them than flat portions. Subject D was able to stabilize higher instabilities than subject P using hand position control, before failing at an instability of 3.25 rad/s.

3.3.2 Vibrotactile feedback results

Figure 3.3 also shows how well subjects used vibrotactile feedback. For both control methods, the mean graded vibrotactile feedback λ_C scores were higher than on/off vibrotactile feedback scores (Welch's t-test, hand movement, p = 0.013; pinch force, p = 0.042). As seen in Figure 3.5, subjects made very large movements during vibrotactile feedback during both hand control and pinch force control. A major difference between the two control methods, however, was that hand control took longer to change direction. The arm's inertia prevented the subject from making fast corrections. Conversely, with the pinch force control, the control signal could be change very rapidly.

Additionally, subjects used pinch force control better than hand movement control during both types of vibrotactile feedback (Welch's t-test, graded, p = 1.10e-5; on/off, p = 3.51e-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 (Figure 3.6) 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, graded, p = 0.0037; on/off, p = 9.68e-4).

3.3.3 '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

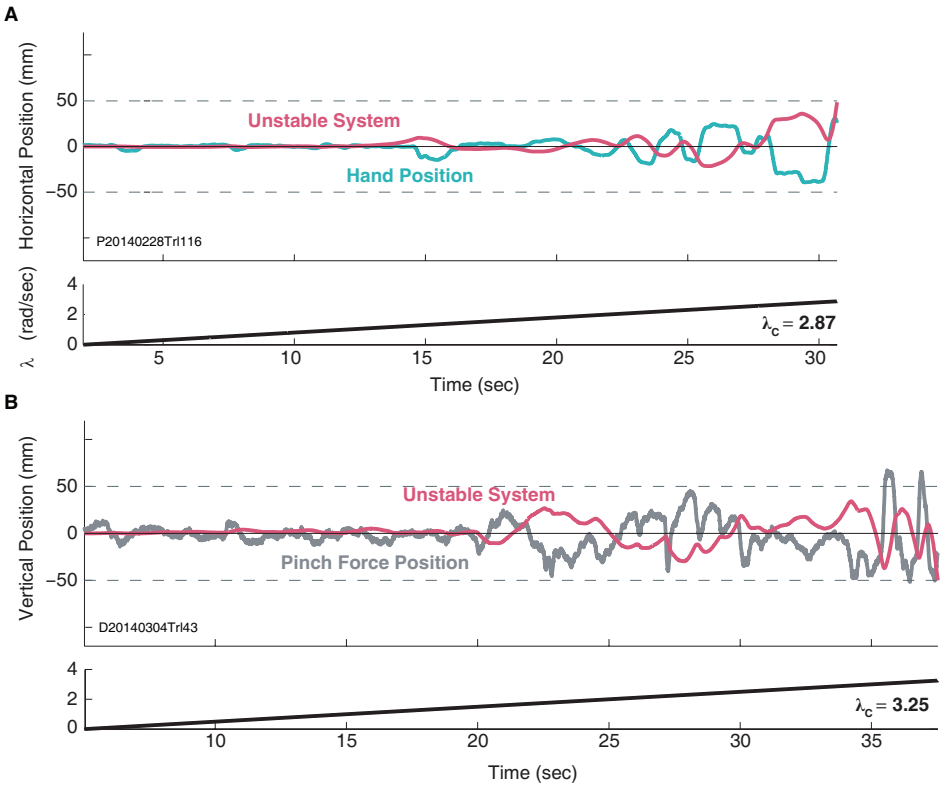

Figure 3.4: Sample trials from humans using visual feedback to control CST. (A) Subject P using hand position (blue trace)to control the unstable system (pink trace). The unstable system crosses the position threshold (dotted lines) and fails the trial when the instability is 2.87 rad/s. The subplot underneath the position trace shows the continuous increase of the system's instability. (B) Subject D is using finger pressure (gray) to control the unstable system (pink). The subject loses control of the system when the system reaches an instability of 3.25 rad/s.

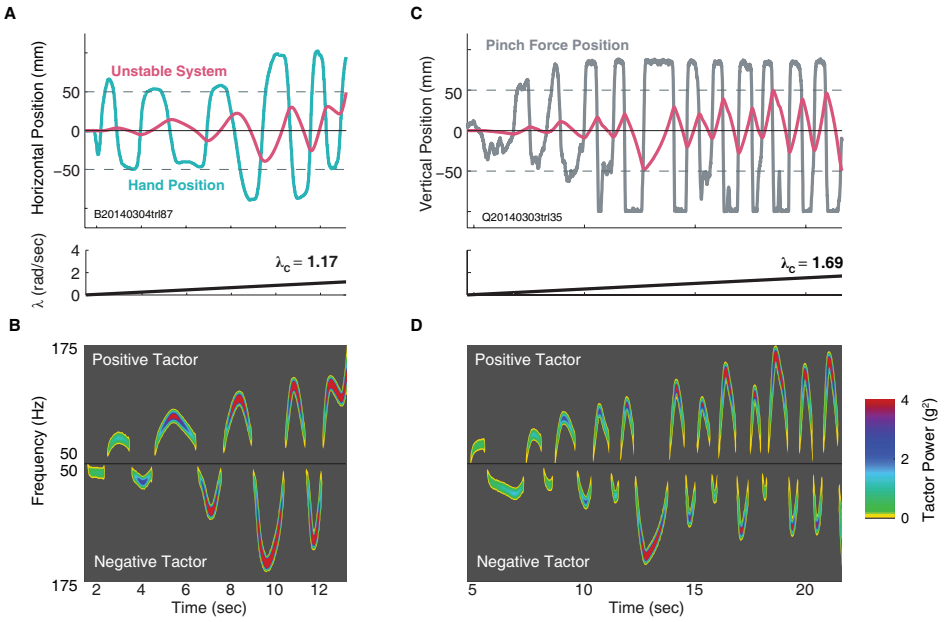

Figure 3.5: Sample trials from humans using graded vibrotactile feedback to control the CST. (A) Subject B using hand position (blue trace) to control the unstable system (pink trace). The subject's movements start off large, even though the system's actual position is small, and near the center target. The subject fails the trial at an instability of 1.17 rad/s. (B) The spectrograms from the positive and negative tactor during the trial in (A). The tactors start vibrating with a low frequency and low power, to indicate the system's small deviations away from the center. As the systems starts moving further away, both tactors increase their respective frequency and power. (C) Subject Q using pinch force (gray trace) to control the unstable system (pink trace). The control movements do not start out as large, but soon the subject hits the upper and lower limits of the force sensor. This trial is interesting because the subject narrowly avoids failing three times, around 12s, 17s, and 21s into the trial. However, the subject's movements are not accurate enough to prevent the trial from failing on the fourth time. The subject is able to attain a critical instability score of 1.69 rad/s. (D) The spectrograms from the positive and negative tactors during the trial in (C).

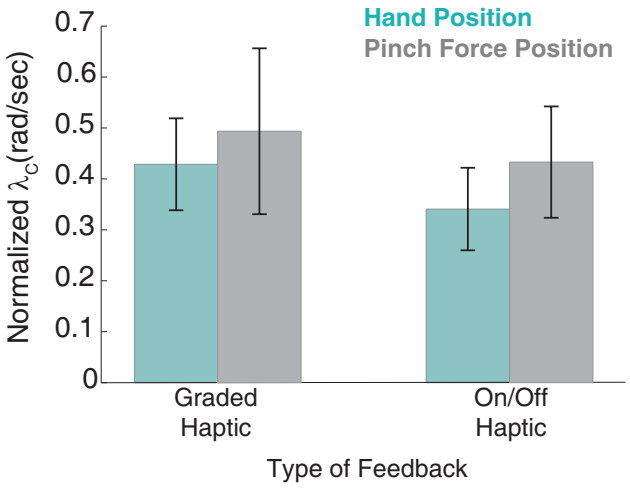

Figure 3.6: Human CST vibrotactile feedback scores normalized by visual feedback scores. The blue bars indicate the graded and on/off vibrotactile feedback scores as normalized by the visual feedback score during hand control. The gray bars indicate the vibrotactile and on/off feedback scores as normalized by the visual feedback scores during pinch force control. Error bars indicate standard deviation of the normalized scores. Even after normalizing, the pinch force control continuous to yield better control during both forms of vibrotactile feedback.

results provide a baseline level of control when only the input to the system is known. While running the 'no feedback' trials (Figure 3.7), subject attempted different control signal strategies. Most subjects tried to keep their control signal near zero. However, some subject seemed to have felt that they were doing a poor job of keeping their control signal at zero. When this occurred, they corrected their control signal by moving in the opposite direciton. 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. In general, the subjects performed better at keeping their hand control signal near zero than keeping their pinch force signal near zero.

3.3.4 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 3.8). 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 (Figure 3.9) 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 for to the subject.

3.4 DISCUSSION

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 users actions, and that feedback is not necessarily visual. We investigated

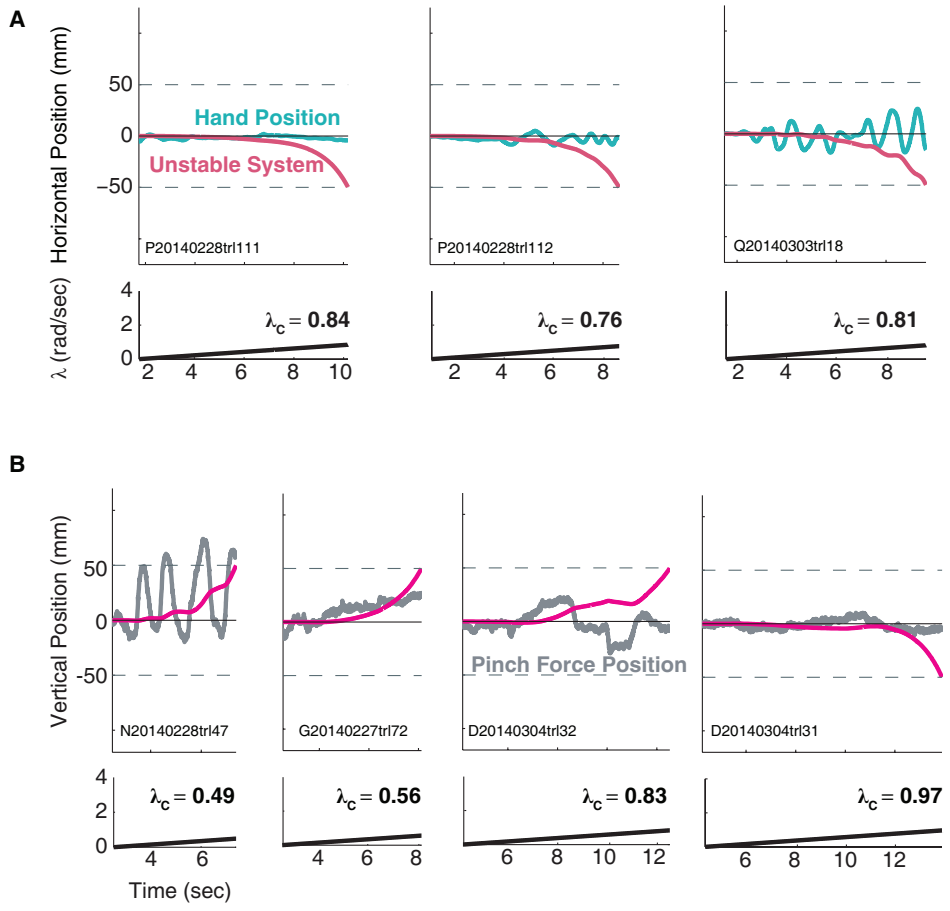

Figure 3.7: Human CST trials during the absence of feedback. (A) Three trials from human subjects using hand position to control the CST with no feedback. Three different strategies were used, all of which yielded similar CST scores. (B) Four trials from human subjects using pinch force to control the CST with no feedback. As compared to *(A)*, subjects do not keep their pinch force position as close to the center. This yielded smaller and more variable scores. The fourth trial, however, demonstrates that it is physically possible to keep pinch force position close to the center, and in doing so, yield a similar score as hand control.

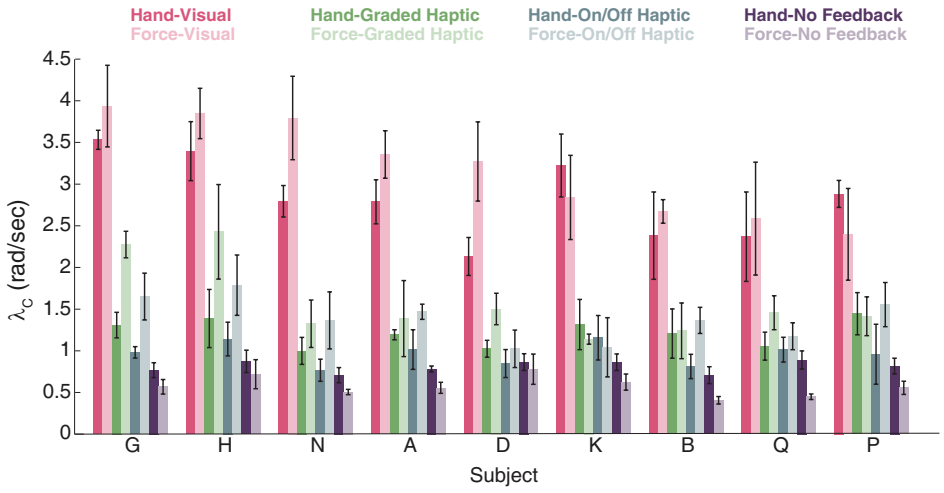

Figure 3.8: Individual human CST scores for each subject. Subjects are ordered by their ability to perform the CST during pinch force control and visual feedback. Subjects had a wide range of ability in terms of performing the CST with visual feedback. Except for subjects G and H, subjects had less variability in terms of using the vibrotactile feedback.

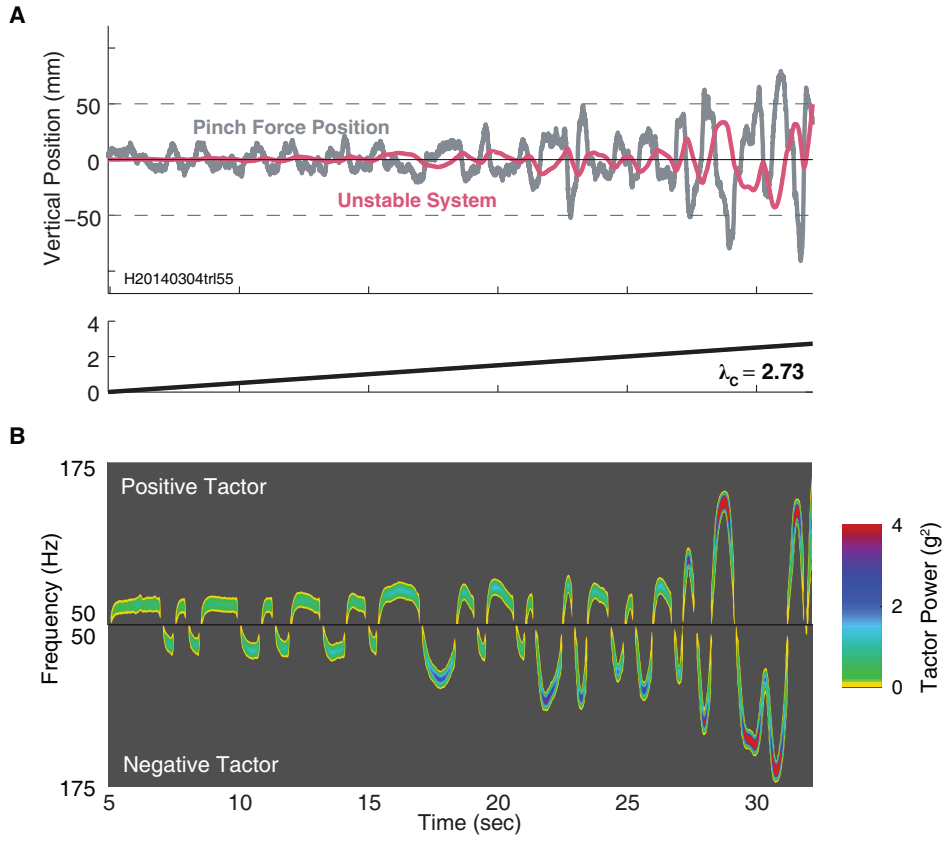

Figure 3.9: Subject H using graded vibrotactile feedback to control the CST. (A) Subject H was exceptionally good at using pinch force (gray trace)and graded vibrotactile feedback to control the system (pink trace). Compared to the previous figure in C, subject H is better able to match and compensate for the system's deviations. During the same instabilities as witnessed by subject Q, the subject produces small control signals. This indicates a better understanding and usage of the graded feedback. Subject H does not hit the sensor's limits, but rather, grades the control signal to the feedback being perceived. (B) Positive and negative tactor spectrograms rendered to subject H during the trial in (A). The vibrations remain at a lower frequency and power until larger instabilities are reached by the subject.

different vibrotactile feedback methods, and also different input methods. We found that the optimal input and feedback type differed somewhat between subjects. The CST is a fast, efficient method to calculate a subject's ability to control a system using a particular control method and feedback method – five trials. 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.

4.0 RHESUS MACAQUE PERFORMANCE ON CRITICAL STABILITY TASK

Figures and text in this chapter have been reproduced with permission from Quick et al. 2015. This chapter expands on the human work in Chapter 3 by amending and translating the task for use in monkey subjects. We made changes in the task design and how we determined of the final critical instability score. In addition to stabilizing instability levels using hand control, we also tested a monkey's ability to stabilize instability levels using a BCI decoder. This this chapter, I show how the Critical Stability Task can also be used as a tool in monkey subjects to measure natural and BCI sensorimotor control. This 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

A perfect brain-computer interface (BCI) would allow users to interact with objects as flexibly and dexterously as they did prior to injury. Able-bodied people use proprioceptive and tactile feedback to perform such interactions. BCI users lack non-visual feedback and must rely on vision to gather information about their prosthesis and how their prosthesis affects objects. Methods to provide non-visual feedback is an active area of research. Current methods being tested include mechanical or electrical stimulation (Weber et al. 2012), such as including 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 electrotactile feedback (Kaczmarek et al. 1991, Tyler et al. 2003, Ptito et al. 2005). In addition to the feedback methods themselves, each method has a vast number of possible parameters from which to choose. How do we determine which stimulation frequencies, amplitudes, and spatio-temporal patterns are the best for feedback? Although a biomimetic approach might be ideal in the long run, in the near term, we will need a way to quantify the effectiveness of various forms of artificial feedback. Here we report a task to assess a monkeys ability to use a particular form of sensory feedback.

Our task evaluates feedback at its essence: it tests the feedbacks ability to be used by the user to rapidly correct their movement. If the user takes too long to interpret the feedback signal before making a motor decision, the feedback is not effective. If the user misinterprets the information conveyed by the feedback signal, the feedback is not effective. A task to evaluate these qualities in a feedback signal must meet two design criteria: First, the task must be performed with as few sources of feedback as possible, to help us unambiguously ascribe performance to a particular source of feedback. Second, the task cannot be successfully performed using memorized movements, but must intrinsically rely on the feedback signal. Third, task performance should be characterized by a single metric, so that experimenters can easily interpret the results.

To achieve these two objectives, we developed a task in which our Rhesus monkey subjects performed a long-duration movement to stabilize an unstable system. Long-duration stabilizing movements are routine in human behavior, as salient examples: maintaining balance while standing, balancing a tray of food, steering a car, or pushing a stroller on a path. A classic example of stabilizing an unstable system is balancing an inverted pendulum.

In our new task is called the Critical Stability Task (CST) (Quick et al. 2014, Jex et al. 1966). In the CST, monkeys stabilized a one-dimensional virtual unstable system by making an appropriate series of hand movements. The objective of stabilizing an unstable system was depicted as trying to keep a cursor on a target which represented an unstable equilibrium point. If the cursor drifted to the right of the equilibrium point, the monkey stabilized the system by moving his hand to the equal and opposite position on the left, and vice versa.

The CST was made more difficult to stabilize by increasing the systems instability level. At higher instability levels, if the cursor was not exactly centered on the target, the cursor would move away from the target at an increasing speed. In order to bring the cursor back to the center, the subject must respond more quickly to the feedback.

This task is ideally suited for studying non-visual feedback and it achieves our two objectives. Subjects can perform the using only non-visual feedback and do not need any form of visual feedback. Additionally, by making the dynamics of the cursor unstable, the monkey must rely on the non-visual feedback and cannot perform memorized actionseach trial is unique.

The contributions of the present work are found in both the methods Section 4.2 and the results Section 4.3. We demonstrate the intuition behind using the CST to test the effectiveness of a feedback signal. Importantly, we show how the effectiveness of a feedback signal can be captured using a single metric. Next, we take a first pass at demonstrating this metric by measuring the effectiveness of an intuitive visual feedback signal and an abstract vibrotactile feedback signal. We hypothesized that monkey could learn to use an abstract non-visual feedback signal to continuously control a movement. We also hypothesized that the CST would be able to score the effectiveness of that vibrotactile feedback signal in such a way that it could be compared with natural visual feedback. We found that with learning, the monkeys use of the vibrotactile feedback signal was approximately 61–75% as effective as their use of visual feedback. Additionally, we report that the CST is 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 monkeys ability to use a BCI decoder with visual feedback. We found that using the BCI decoder as an effector was 37% as effective as the use of natural hand movements.

4.2 METHODS

Three adult male Rhesus monkeys (macaca mulatta) were used in this study (monkeys I, J, and N (8.8 kg, 12.6kg, 8.6kg). All animal procedures were approved by the University of

Pittsburghs 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 veterinarians recommendation for reasons unrelated to the studys objectives.

4.2.1 Critical Stability Task implementation

The monkeys used two methods for controlling the Critical Stability Task (CST): hand movements and a BCI decoder. During both control schemes, monkeys received feedback on the unstable systems position through either visual feedback or vibrotactile feedback.

The objective of performing the CST is to stabilize a one dimensional, first-order unstable system (Figure 4.1A). Monkeys were asked to keep the unstable system, as represented by a cursor, positioned on a center target. The system was successfully stabilized when the systems position remained within ±50mm of the center target for six seconds. To accomplish this goal, monkeys counteracted the unstable dynamics by moving their hand opposite to the unstable systems position (Figure 4.1C). In theory, the monkeys could instantly and perfectly stabilize the system if they could instantaneously move to the exact equal and opposite position of the systems current position. In practice, of course, each corrective movement intended to stabilize the system yielded a residual error, requiring another corrective movement. The unstable system transfer function ($G(s) = \frac{\lambda}{s-\lambda}$) (Jex et al. 1966, Kadkade et al. 2003, Quick et al. 2014) was converted into a time-domain equation and discretized, Equation 4.1.

$$x(k+1) = e^{\lambda T_s} x(k) + (e^{\lambda T_s} - 1)u(k) \tag{4.1}$$

In Equation 4.1 and Figure 4.1, $u(k)$ represents the position of the input signal at time step k, which was either the monkeys hand position or BCI decoders position, $x(k)$ represents the systems position at time step k, and T_s is sampling time which was 10ms for hand control and 30ms for BCI decoder control. We trained the monkeys to try to keep the cursor at

the center target, which is equivalently stated as trying to make $x(k+1) = x(k) = 0$. The systems instability level was represented by $\lambda > 0$ rad/s, where systems with larger λ values were more unstable and harder to control (Figure 4.2).

As the system's instability level increases, the more quickly the system responds to the input. In Figure 4.2A, a one hertz sine wave with an amplitude of 5mm is passed into three different systems with instability levels at 1 rad/s, 2 rad/s and 3 rad/s. The system with the 3 rad/s instability level deviates the most quickly from zero. We can close-the-loop by having a monkey supply the input signal rather than passing in the sine wave. The monkey will receive feedback on the systems position and respond with a corrective movement, Figure 4.2B. In order to stabilize an instability level of 1 rad/s, the monkey can respond and move at a low speed. In order to stabilize an instability of 3 rad/s, the monkey must respond quickly, before the systems position moves too far away from zero.

Because a systems instability level sets the motor response requirements, we will use the instability level to measure motor performance. Our fundamental performance metric is the maximal instability level at which the monkey can maintain control of the system. We deemed this the critical instability λ_c. Pragmatically, we defined λ_c as the maximal instability for which the monkey could maintain stability for six seconds on half of the trials. Six seconds was chosen as being a sufficiently long period of time to determine successful stabilization. The λ_c could be determined for almost any 'effector–feedback' pairing. Effector–feedback pairings which yield a high λ_c indicate effective motor responses. By using the same effector and changing the feedback signal, we can isolate the effectiveness of different feedback signals. By using the same feedback signal and changing the effector, we can isolate the effectiveness of different effectors.

For each effector–feedback pairing, monkeys had many practice sessions where they learned and progressed through the different instability levels. After their learning had plateaued, they performed five sessions which were analyzed to find the λ_c score for that effector–feedback pairing (Figure 4.8). Each experiment session began with systems with small instability levels. Every 10-20 trials, the instability level increased by 0.1 rad/s. We stopped incrementing the instability level when the monkey repeatedly failed or stopped attempting trials. Once this upper instability level was reached, we reset the instability level

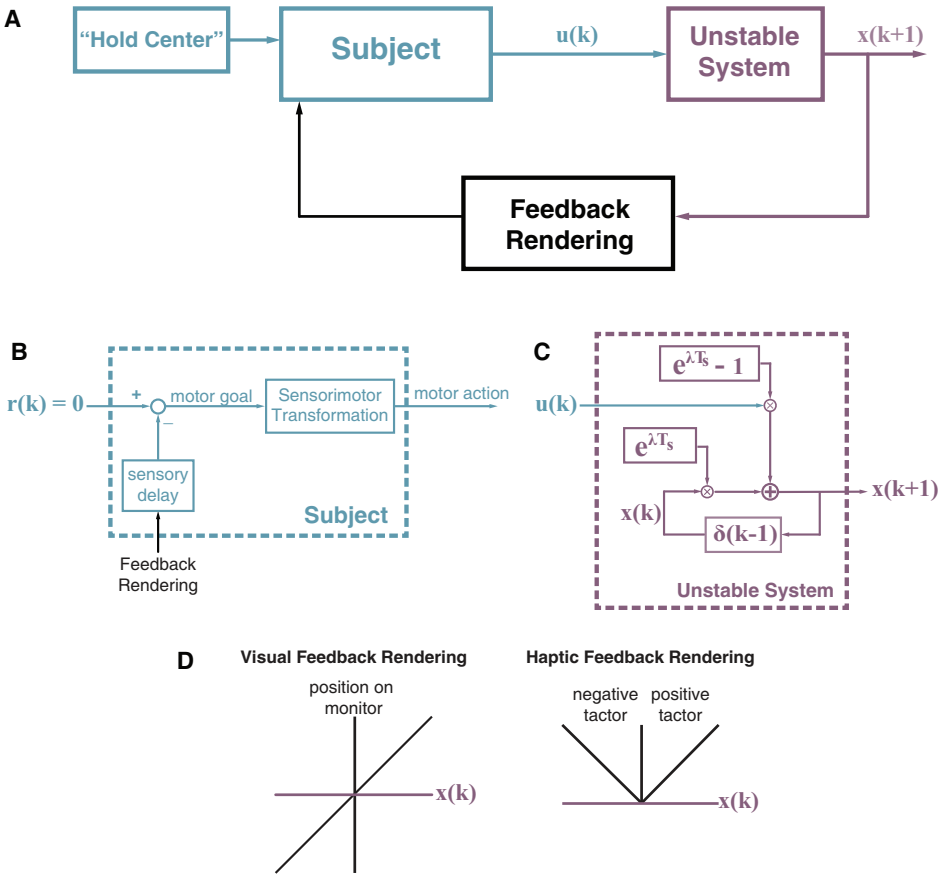

Figure 4.1: The feedback - control loop of the CST. (A) The four components of the CST. (B) The subject receives delayed feedback which is compared with the goal of keeping the systems position at the center. The distance from the current system position to the center becomes the motor goal and is transformed into the next motor command. (C) The current system position $x(k)$ is multiplied by a gain greater than 1, with a larger gain corresponding with larger instability levels λ. The motor command $u(k)$ is multiplied by the portion of the gain that exceeds 1. The dynamics are setup such that if the subject can match $u(k)$ with $-x(k)$, then the system will not change positions on the next time step $x(k+1)$. (D) There are two potential ways for the feedback to be rendered, through visual feedback or vibrotactile feedback. For visual feedback, there is a linear mapping from the system's position to coordinates on the computer monitor. For vibrotactile feedback, two tactors encode the systems distance from the center target. Each tactor increases its frequency and magnitude of vibration with the systems distance from the center.

to what it had been at the beginning of the session and repeated the incrementing process. For each experiment sessions, we attempted to get two full sweeps through the instability levels.

4.2.2 Effector–feedback pairings

4.2.2.1 Hand control with visual feedback 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 I and J performed the CST using unconstrained hand movements and visual feedback of the system's position. Hand position was recorded using an active motion capture system (PhaseSpace Inc., San Leandro, CA), Figure 4.3A. A powered LED marker was attached either to the monkey's finger (monkey J) or the monkey's wrist (monkey I). Six motion capture cameras recorded the 3D position of the LED marker.

The hand position was rendered in real-time to the monkey as a cursor in a virtual 3D environment. Each trial began with the appearance of a square (acceptance window 4mm monkey I, 10mm monkey J) in the center of the screen. The monkey moved his hand, as represented by the red cursor, to the center square. After a hold period of 1000ms (during which the acceptance window enlarged to 10mm monkey I, 14mm monkey J), the cursor changed colors to indicate that it now represented the position of the system. The monkey initialized and held his hand position at the center square so that the initial input (which acts like a disturbance) to the CST would be small. Once the CST began, only the monkey's horizontal hand movements controlled the unstable system, Figure 4.3. We wish to emphasize that the cursor's position during the CST represented the position of the system, and not the position of the hand. Hand movements altered the system's state according to Equation 4.1, but the hand position was not directly rendered to the monkey. For six seconds, the

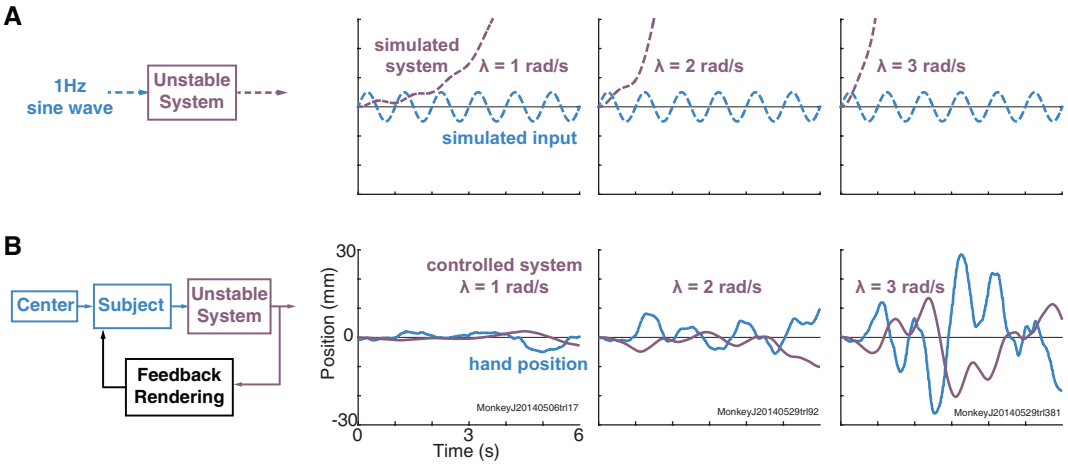

Figure 4.2: The effects of increasing system instability. (A) In simulation, a 1Hz sine wave with an amplitude of 5mm is passed through three different unstable systems, with increasing instability: 1 rad/s, 2 rad/s and 3 rad/s. With increasing instability, the systems output diverges from the center more rapidly. (B) Individual trials of monkey J controlling these systems. At an instability of 1 rad/s, the monkey can easily stabilize the system with visual feedback by making small, slow movements. At an instability of 2 rad/s, faster movements are required and the subject also makes larger movements. At a high instability of 3 rad/s, monkey J makes fast and large movements to stabilize the system.

monkey was required to keep the unstable system within a boundary of ±20mm (monkey I, ±50mm monkey J) around the center square. If the monkey succeeded in maintaining stability of the system for six seconds, a water reward was delivered. If the system's position drifted past the boundary (dashed lines in Figure 4.3) before the six seconds elapsed, the trial was deemed unsuccessful and a 2s failure penalty ensued.

4.2.2.2 Hand control with vibrotactile feedback Vibrotactile feedback was rendered via two coin tactors (312-101, Precision Microdrives Limited, UK) which vibrated when supplied a voltage. When a tactor's supply voltage was increased, the tactor vibrated with both increased frequency and magnitude. The tactors were attached to the forearm of the monkey's non-working arm using Velcro against a Spandex sleeve. The non-working arm was comfortably restrained on an armrest. When the systems position was right of center, the tactor near the wrist vibrated and when the unstable system was left of center, the tactor near the elbow vibrated. Each tactor's supply voltage was proportional to the system's position. Thus, when the system's position was far right of the center, the tactor near the wrist would provide a large- magnitude high-frequency signal. In principle, if the system were ever perfectly balanced at 0mm, both tactors would be off. In practice, one of the two tactors was always vibrating. We measured the tactors' activity level via an accelerometer attached to the back of each tactor. At steady-state, the vibration frequency ranged from 50 – 175 Hz, and the vibrational power ranged from 0.5 to 20g^2 (where 1g equals the force of gravity).

Monkeys initiated vibrotactile feedback trials by moving their hand to the center square under visual guidance. Once the CST began, the cursor vanished and the monkeys received vibrotactile feedback of the system's position. The center square remained visible to the monkey to indicate that the trial was ongoing. The success criteria were the same as the criteria for visual feedback trials (monkey J). Both monkeys had to maintain stability of the system within the boundaries of ±50mm around the center position for six seconds to earn a water reward.

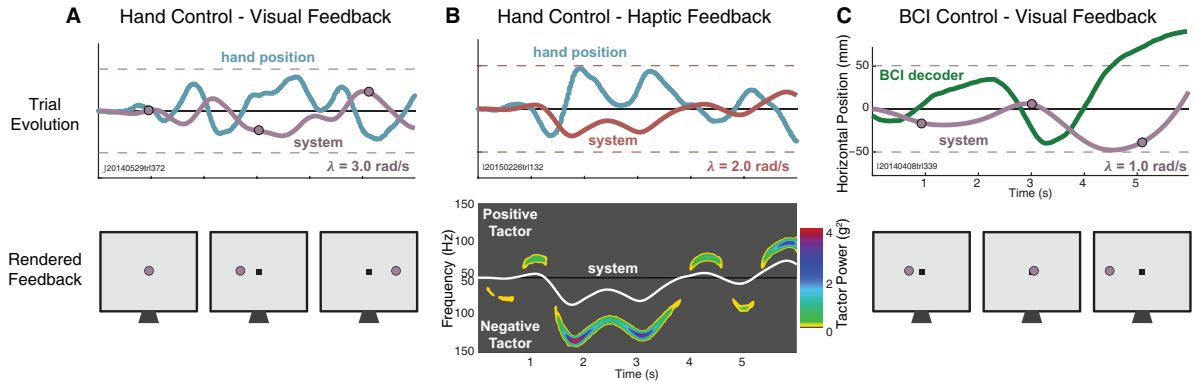

Figure 4.3: Three representative trials the Critical Stability Task under different effector – feedback pairings.(A) Hand control – visual feedback. For this trial, $\lambda = 3.0$ rad/s. Purple: system state (which is seen by the monkey) blue: hand position (which is not seen by the animal). Lower row: schematic of visual feedback rendered to the monkey at three time points (circles). Data from monkey J. (B) Hand control – vibrotactile feedback. For this trial, $\lambda = 2.0$ rad/s. Lower row: activity level of the two tactors rendering the vibrotactile feedback. The y-axis is the frequency of vibration and the color is the power of the vibration in g-force. One tactor encodes the positive deviations of the system from equilibrium and the other encodes negative deviations. The white trace is the systems position (same as the purple trace in the top row). Data from monkey I. (C) BCI control - visual feedback. For this trial, $\lambda = 1.0$ rad/s. Data from monkey I.

4.2.2.3 Visualizing the vibrotactile feedback signal

To plot the vibrotactile feedback signal, as seen in Figure 1, accelerometers were attached to the tactors and sampled at 1kHz. To analyze the accelerometer signal, the voltage signal had the baseline voltage during no vibrotactile feedback subtracted, and the resulting trace was multiplied by 4.854g/V as specified in the data sheet. The g-force trace was then prepared for the spectrogram by subtracting off the mean g-force for that trial. We then found the power in the g-force signal (g^2) using a short-time Fourier transform on 200ms segments of data, where each segment overlapped with 199ms of the previous segment. The vibrotactile power was then plotted as a function of frequency and time.

4.2.2.4 BCI control with visual feedback

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 4.0×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 4.2.2.5. 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.

4.2.2.5 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 the hold still trials. The offline decoded position was used as input into simulated CST trials. The decoder was trained and tested on the 24 hold still trials. Initially, we had tried training the decoder on center-out trials, however, this yielded disproportionately large decoded movements during 'hold still' trials. So as to achieve the 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:

$$\boldsymbol{x}_k|\boldsymbol{x}_{k-1} \sim \mathrm{N}(A\boldsymbol{x}_{k-1} + \boldsymbol{b}, Q) \tag{4.2}$$

$$\boldsymbol{u}_k|\boldsymbol{x}_k \sim \mathrm{N}(C\boldsymbol{x}_k + \boldsymbol{d}, R) \tag{4.3}$$

where $\boldsymbol{x}_k \in \mathbb{R}^{2\times 1}$ is a vector of horizontal and vertical decoder velocity at time step k, $\boldsymbol{u}_k \in \mathbb{R}^{q\times 1}$ 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 either 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 $\hat{\boldsymbol{x}}_k = \mathrm{E}[\boldsymbol{x}_k|\boldsymbol{u}_1, \ldots, \boldsymbol{u}_k]$. 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:

$$\hat{\boldsymbol{x}}_k = M_1\hat{\boldsymbol{x}}_{k-1} + M_2\boldsymbol{u}_k \tag{4.4}$$

$$M_1 = A - KCA \tag{4.5}$$

$$M_2 = K \tag{4.6}$$

4.2.2.6 'Holding still' trials In principle, the monkeys could have attempted to perform the CST by ignoring the sensory feedback that was provided to them. This strategy would have manifested in one of several behaviors such as flailing in a manner that is unconnected with the feedback, moving their hand left and right in a self-generated oscillatory movement, or holding perfectly still. We saw no evidence for uncoordinated flailing. Self-generated oscillatory movements would have been a poor strategy because they cause the system to quickly cross the failure boundaries (Figure 4.4A). That leaves the remaining strategy, holding perfectly still at the center square, as the monkeys best opportunity to attain a high λ_c score without using feedback about the systems position. 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.

To measure the λ_c score the monkeys could have achieved simply by holding as still as possible, we instructed the monkeys I and N to hold their hand at the center target for six seconds. In these trials, they were shown their hand position rather than the systems position. If the hand remained within the center squares acceptance window (7 mm monkey I, 20 mm for monkey J) for six seconds, a water reward was given. Twenty four successful trials were collected from monkey I, and 33 from monkey N. We also recorded neural activity from monkey I while he performed the 'hold still' trials.

Offline, we determined the λ_c score that the monkey would have earned by holding at the center, given the hand position data measured on each trial (Figure 4.4B). The recorded hand position signals were input into unstable system simulations with different instability levels (Figure 4.4A). Each trial was used as input to 41 system simulations with instabilities

ranging between $\lambda = 0.0$ - 2.0 rad/sec in 0.05 rad/sec increments. Thus 984 simulated CST trials were generated for monkey I and 1353 simulated trials were generated for monkey N. We deemed trials to be successful or not using the same criteria as during the CST task (6s trial, $\pm$50mm boundary). This allowed us to compare the λ_c from holding still, with the λ_c score from visual and vibrotactile feedback to ensure that the feedback was being used. For monkey I, we also applied an offline BCI decoder to the neural data recorded during these hold still trials to estimate the BCI performance if the monkeys strategy was to hold still.

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.

4.2.3 Training the CST

4.2.3.1 Visual feedback training regime Both monkey's first learned how to perform the CST using visual feedback. The monkey's first learned to hold their hand still on the center target. This objective was instilled in the monkeys by allowing them a small, but continuous, stream of water reward for holding on the center target. If they moved outside of the acceptance window, the water immediately stopped flowing. If they moved back into the window, the water began flowing again.

Once the monkey understood how to hold still at the target, they progressed to trials in which they were rewarded only after holding at the center for six seconds. Next, we switched to CST trials with low instability levels. At an instability level of around 0.4 rad/s, the trials appear similar to the original 'holding still' trials. Then, gradually over several sessions, we increased the instability level. In this manner, the monkeys' objective was consistent all through the course of training keep the cursor in the center.

4.2.3.2 Vibrotactile feedback training regime Before using vibrotactile feedback in the CST, the monkeys were taught to understand vibrotactile information using a horizontal reaching task. First, monkeys were trained to perform a standard two-target reaching task under visual feedback. In this task, monkeys moved their hand-controlled cursor to a start

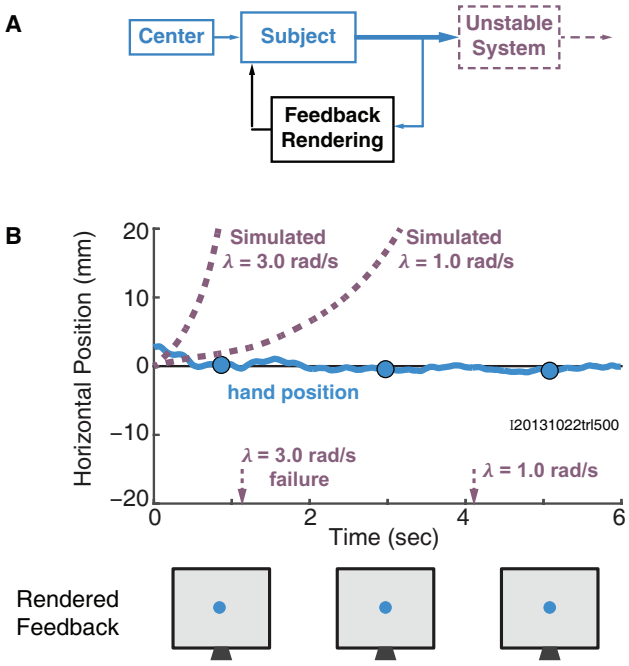

Figure 4.4: Schematic holding hand position at the center target. (A) In contrast to the CST, the monkey only receives feedback about his hand position. The unstable system is applied offline. This allows many different unstable system simulations to be tested. (B) An example trial of holding at the center (monkey I). The dashed lines indicate simulated system performance for systems with $\lambda = 3$ rad/s and $\lambda = 1$ rad/s. In both cases, the monkey would have failed the trial (failure times indicated by the arrows on x-axis).

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 15mm acceptance window for 250ms, the trial was successful and the subject was rewarded. Once proficient, 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 the subject 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.

4.2.4 Data analysis

4.2.4.1 Measuring the critical instability of an effector – feedback pairing Our single task performance metric is the critical instability score, λ_c. This is the system parameter λ for which the monkey can maintain stability for the full trial duration for half the trials. The critical instability score λ_c depends on the feedback signal as well as the effector

used to control the system. Thus, the λ_c score allows a direct comparison of the effectiveness of different feedback signals (when the input modality is held constant) and the effectiveness of different effectors (when the feedback signal is held constant.)

We estimated λ_c for each effector and feedback combination separately. The estimate was made using all trials from the last five sessions (Figure 4.8) for that effector – feedback pairing. The λ_c score for each effector – feedback pairing was computed using a bootstrap analysis (Figure 4.5). In the bootstrap, we used the instability level for each trial and whether the trial was a success. We then resampled the trials with replacement until the number of resampled trials equaled the original number of trials. Next, the monkeys percent success at each instability was calculated from the resampled trials. Using nonlinear regression, we fit a cumulative Gaussian distribution to the percent success $p_N(\lambda)$, using Equation 4.7, where N was the bootstrap iteration.

$$\hat{p}_N(\lambda, \mu_N, \sigma_N) = 0.5\left(100 - 100erf\left(\frac{\lambda - \mu_N}{\sqrt{2}\sigma_N}\right)\right) \tag{4.7}$$

The nonlinear regression estimated percent success $\hat{p}_N$ by solving for the distribution's mean μ_N and standard deviation σ_N. We then used the distribution's mean μ_N as the Nth bootstrap estimate of the critical instability level λ_{CN}, the instability level at which the monkey succeeded on half the trials. The bootstrapping procedure was repeated for 1000 iterations, to obtain 1000 estimates of λ_{CN}. We then calculated the mean and standard deviation across all λ_{CN} estimates, which we took as the mean and standard error of the mean of the true λ_c for that effector–feedback pairing. The λ_C score was found for each effector – feedback pairing, as seen in the results in Figure 4.6.

This bootstrap procedure was used in order to obtain the standard error of the mean of λ_c. An alternative method would have been to form a single Gaussian distribution of CST success as a function of the systems instability level. This would have yielded a λ_c and a standard deviation describing the spread of the distribution, however, we would not have had a confidence interval about λ_C. Without a confidence interval, we could not test if a particular effector – feedback pairing was significantly more effective than another effector – feedback pairing.

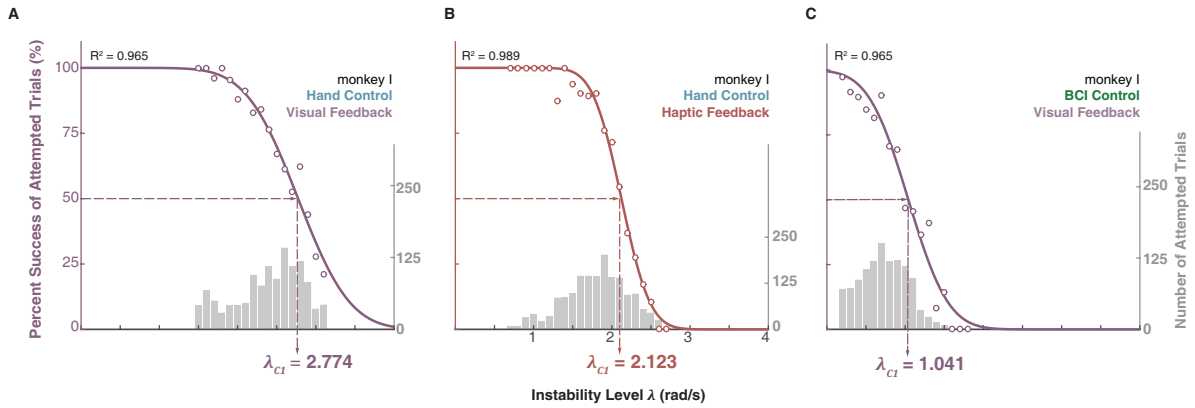

Figure 4.5: Illustration of bootstrap procedure. Three effector – feedback pairings are shown. Data from monkey I. Each circle represents the percentage of successful trials at a specific instability level, calculated from the bootstrap procedure. The curve is the psychometric function fit to those data. The critical instability for one iteration of the bootstrap λ_{C1} was determined as the x-axis value corresponding to the y-axis 50% point on the psychometric curve. The final critical instability is the mean of 1,000 iterations. The inset histograms show the number of trials monkey I attempted for each instability level. (A) Hand control – visual feedback. (B) Hand control – vibrotactile feedback. (C) BCI control – visual feedback.

4.2.4.2 Determining the sensorimotor delay

For unstable systems with high instability levels, monkeys must respond to the rendered feedback quickly. We reasoned that we might see behavioral evidence for more exacting control as the behavioral requirements exerted by the system become more stringent. For each effector–feedback pairing, we determined the sensorimotor delay between changes in the systems state and the monkeys corresponding corrective action. To determine the sensorimotor delay on each trial, we performed a cross-correlation between the rendered systems position and the monkeys effectors position, on successful trials only. To find the time at which the systems position was actually rendered, we measured the rendering delay of our rig. For visual feedback trials, we measured the graphics rendering delay using a photodetector. For vibrotactile feedback, we determined the average tactor onset time using accelerometers. We then incremented the systems position forward in time by these delays.

We examined all time lags (10ms spacing) in the range 500 to 2000ms, where a positive lag indicates the hand is lagging the system position. Since the corrective movement was made in the opposite direction from the systems position, we recorded the lag associated with the peak in negative correlation between the systems position and the effectors position. We then found the mean sensorimotor delay for each effector–feedback pairing for each instability level. This sensorimotor delay is similar to the reaction time measured in a tradition static-target task, such as center out, but the CST affords the ability to measure the sensorimotor delay in an ongoing basis through a long trial.

4.3 RESULTS

The Critical Stability Task allowed us to quantify sensorimotor performance across a variety of effector modalities (hand control and brain control) and feedback modalities (visual and vibrotactile.) To succeed at a CST trial, the monkeys had to maintain stability of an unstable system for six seconds. The task was made more difficult by increasing the instability level (λ) of the unstable system. The CST allows us to quantify motor control with a single number, λ_c, which is the maximal instability level that the monkey is able to control.

4.3.1 CST under hand control with visual feedback

When performing the CST under hand control with visual feedback, monkey I achieved a λ_c score of 2.784 ± 0.030 rad/s (mean ± s.e.m.) over 1274 trials and monkey J achieved a λ_c score of 3.026 ± 0.038 rad/sec over 1370 trials, Figure 4.6.

4.3.2 CST under hand control with vibrotactile feedback

After becoming proficient at the CST with visual feedback, both monkeys were taught to perform the task using vibrotactile feedback. The training procedure is described in Section 4.2.3.2. Vibrotactile feedback was delivered via two tactors placed on the non-working arm. Vibration of the tactor near the wrist signaled that the systems state was to the right of the center square, and vibration of the tactor near the elbow signaled that the systems state was to the left of the center square. The frequency and magnitude of the tactor signaled the distance of the system from the center square.

With vibrotactile feedback, monkey I achieved a λ_c score of 2.130 ± 0.019 rad/s, Figure 4.6, which was 76.5% as effective as his control using visual feedback. Monkey J was able to achieve a λ_c score of 1.858 ± 0.013 rad/s, which was 61.4% as effective as his control using visual feedback.

4.3.2.1 CST under BCI control with visual feedback 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 (5 sessions), Figure 4.6).

Several example stabilization trajectories are shown in Figure 8. The three subplots show how the monkey successfully controlled systems with instability levels at 0.4 rad/s, 0.9 rad/s, and 1.1 rad/s. The first two levels are below his λ_c score of 1.036 rad/s. The third level is higher than his critical instability level. This third plots illustrates some of his best control because he can successfully stabilize 1.1 rad/s less than half of the time. The second

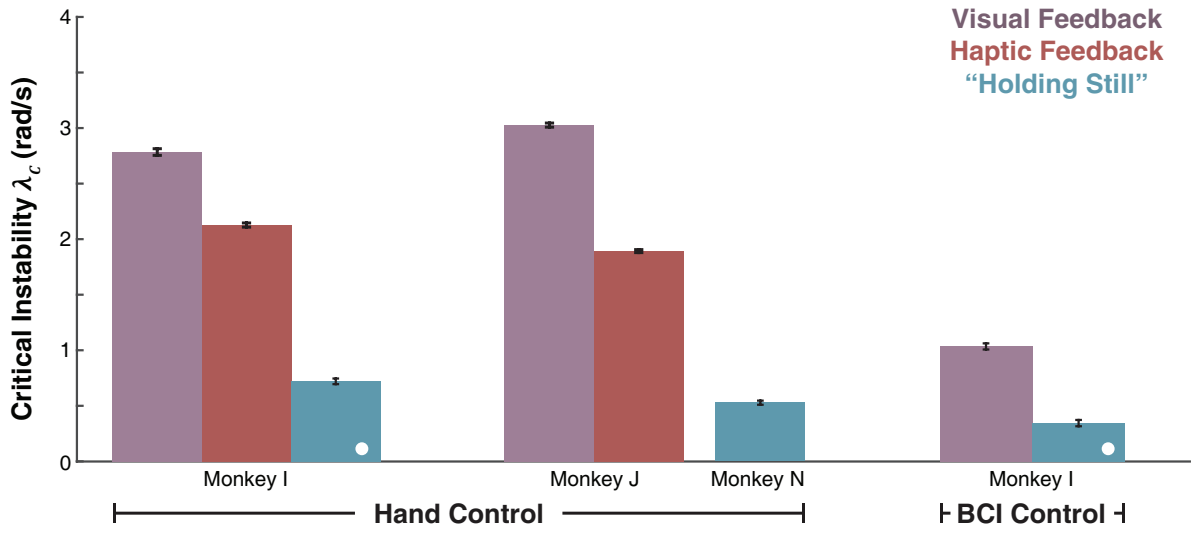

Figure 4.6: Critical instability scores using for all effector – feedback pairings. Feedback condition is represented by bar color, the effector condition is represented on the horizontal axis. The blue bars represent trials from the 'holding still' task where the monkey received visual feedback about their hand position. White dots indicate use of the same trials. Black error bars showing the mean ± s.e.m. of the critical instability.

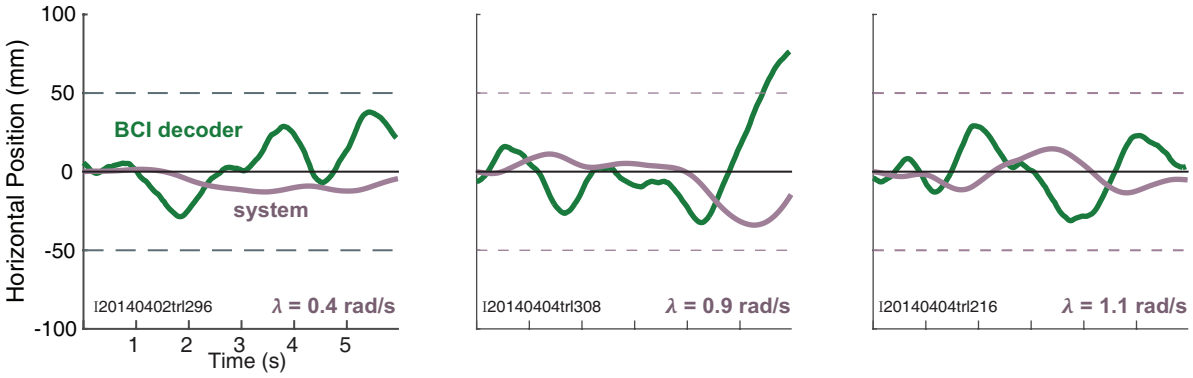

Figure 4.7: Example trials of monkey I using the BCI decoder to successfully stabilize systems with instability levels of 0.4 rad/s, 0.9 rad/s, and 1.1 rad/s. The green line indicates the position of the BCI decoder. The purple line indicates the systems position, which is rendered visually to the monkey.

subplot shows monkey I making a large corrective movement in order to prevent trial failure. A final point of interest is that the BCI decoder makes similarly-sized movements across the three instability levels. This is in contrast with hand control (Figure 4.2B), where a monkey stabilizes low-instability levels with small hand movements.

Even though monkey I was able to use the BCI decoder to stabilize instabilities, the BCI decoder was not very effective. Monkey Is BCI control was only 37.2% 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 show that BCI control that is slightly lower than, 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). By contrast, the single λ_c score demonstrates that there is still much room to design better decoders, such that the CST complements traditional BCI tasks.

4.3.2.2 CST performance under 'holding still' trials An important strength of the CST is that the rendered feedback signal (visual or vibrotactile) is dissociated from the monkeys proprioceptive feedback signals (Figures 4.2, 4.3, and 4.7). The monkeys' proprioceptive signals inform the monkeys only about their input signal, $u(k)$. To assess how well monkeys could perform the CST using only knowledge of their input signal, monkeys I and N performed the 'hold still' task where they held their hand still for six seconds using visual feedback of their hand position. We wished to verify that the monkeys were using the vibrotactile feedback about $x(t)$, and not using the next-best strategy - holding their hand still at the center target. If the monkeys ignored the vibrotactile feedback, and kept their arm still, then the 'hold still' λ_c score would equal the vibrotactile feedback λ_c score. If the monkeys successfully used the vibrotactile feedback, the vibrotactile feedback λ_c score would be higher than the 'hold still' λ_c score. Finally, it is also possible that the monkeys could use the vibrotactile feedback incorrectly, meaning, they could make inappropriate movements in response to the vibrotactile feedback. If they used the vibrotactile feedback incorrectly, the vibrotactile feedback critical instability would be less than the hold still critical instability. The notion that the vibrotactile feedback could perform worse than the 'hold still' λ_c score can be seen early in learning, where the monkeys attempt to stabilize instability levels around 0.5 rad/s (Figure 4.8D-E).

Monkey I's 'hold still' hand position was used as inputs to multiple simulations of unstable systems with different instability levels. Using these simulated trials, monkey I's 'hold still' λ_c score was 0.722 $\pm$ 0.031 rad/s (mean $\pm$ s.e.m.). Monkey N's 'hold still' λ_c score was 0.531 $\pm$ 0.021 rad/s. The fact that monkey I obtained a higher 'hold still' critical instability score than monkey N is due to the fact that monkey Is hand position remained closer to the center square. Both monkeys attempted to hold on the center square, but monkey Is hand position was more accurate and steady (-0.23mm $\pm$ 1.46mm, mean $\pm$ st. dev.) than monkey N's (-1.55mm $\pm$ 6.11mm). The kinematic differences in their 'hold still' abilities further highlights that holding steady is a better open-loop control strategy than oscillatory-like movements about the center. Thus, the oscillatory movements seen during the other effector – feedback pairings are made in response to the feedback signal and are not internally generated.

Both monkey I and monkey J achieved higher vibrotactile feedback λ_c score than the 'hold still' λ_c scores. This validates our hypothesis that monkeys can successfully use abstract vibrotactile feedback signals, rather than ignoring the vibrotactile feedback or using the vibrotactile feedback incorrectly. Additionally, this validates that our method can successfully differentiate the effectiveness of three different feedback conditions.

We also found the 'hold still' λ_c score for an open-loop, offline BCI decoder. With no feedback about the system, monkey Is neural activity produced a λ_c score of 0.346 ± 0.028 rad/s. This score is significantly lower than when the monkey used BCI control with visual feedback about the systems position. This assures us that the monkey is actively using the visual information to stabilize the CST using BCI control.

4.3.2.3 Learning to stabilize the CST

Monkeys I and J learned how to control the CST over a period of weeks. Both monkeys first learned to perform the CST using hand control with visual feedback, Figure 4.8A-B. After the initial training procedure for visual feedback (Section 4.2.3.1), monkey I learned to control the CST over 24 sessions [6.57 weeks]. Monkey J learned to control the CST over 39 sessions [57.7 weeks]. There was a large gap in monkey Js training, between sessions 19 and 20, which spanned 44.6 weeks. This was because monkey J was the first monkey we had trained to perform the CST, and we were not sure what λ_c score a monkey would be able to attain. After monkey I was trained, and had attained a λ_c score than monkey J, we went back and continued training monkey J. With the additional training, monkey J was able to improve his CST performance and stabilize higher instability levels.

After monkey I learned how to stabilize the CST using hand control with visual feedback, we moved him on to BCI control with visual feedback, Figure 4.8C. Again, we started with low instability levels and increased the level as his performance improved. He learned how to use BCI to control the CST over 15 sessions [4 weeks].

Finally, monkeys I and J learned how to perform the CST using hand control with vibrotactile feedback. Monkeys first went through a tactor familiarization process to introduce them to what the tactor vibrations signified. After this process, we switched them to the CST at low instability levels. Over 86 sessions [21.42 weeks], monkey I learned how to sta-

bilize the CST with vibrotactile feedback. Monkey J learned how to stabilize the CST in 27 sessions [12.14 weeks], Figure 4.8D-E. Both monkeys displayed an ability to learn this abstract form of feedback. It is possible that with further prolonged practice, the monkeys may have been able to use the vibrotactile feedback even more effectively.

4.3.2.4 Sensorimotor delays are task dependent Both monkeys learned how to use vibrotactile feedback. However, they attained a lower vibrotactile λ_c score than visual λ_c score. We were interested in learning why the monkeys were unable to use vibrotactile feedback as effectively as visual feedback. In the CST, one of the primary factors in trial success is the monkeys sensorimotor delay. The theoretical maximum sensorimotor delay for stabilizing a particular instability level is shown by the dotted lines in Figure 4.9. The theoretical maximum delay is the systems time constant (Jex et al. 1966), which is the inverse of the system's instability level. To stabilize an instability level of λ rad/s, the monkey's sensorimotor delay must be less than $1/\lambda$ seconds. If the monkey's sensorimotor delay was near the maximum of $1/\lambda$ seconds, the monkey must produce accurate movements in order to inversely track the system's position with a gain of 1: $u(t) = -x(t - \frac{1}{\lambda})$. When the monkeys sensorimotor delay is less than $1/\lambda$ seconds, the monkey's gain is less constrained and can become larger. The scope of this work does not include analysis of the gains used by the monkeys.

Using the cross-correlation analysis, we found that the monkey's sensorimotor gains were consistently lower than the theoretical maximum delay. On each trial, we found the sensorimotor delay corresponding to the peak in negative correlation. The mean sensorimotor delay for each effector–feedback pairing, at each instability level, is plotted in Figure 4.9. For hand control, the longest sensorimotor delays were measured when the monkeys used vibrotactile feedback to stabilize low instability levels. The shortest sensorimotor delays were measured when the monkeys used visual feedback to stabilize high instability levels. The overall shortest sensorimotor delay was found in Monkey I when stabilizing an instability level of 3.1 rad/s with visual feedback. At the same instability level, monkey J's visual-hand delay was longer. This may in part be because monkey I's success boundary thresholds were ±20mm and monkey J's were ±50mm.

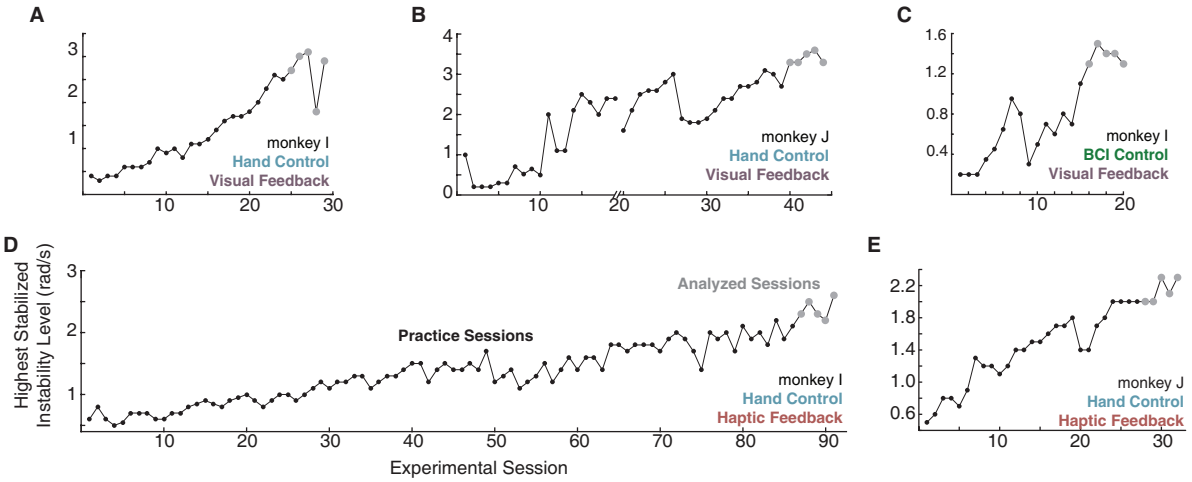

Figure 4.8: Learning to control the CST with five different effector–feedback pairings. (A) Monkey I learning to use hand control - visual feedback to control the CST. Each point indicates the highest instability level he was able to stabilize that day. The black points are sessions where he was practicing and learning to use hand-control - visual feedback. The gray points are sessions that were analyzed throughout the results. (B) Monkey J learning to use hand control - visual feedback. (C) Monkey I learning to use BCI control with visual feedback. (D) Monkey I learning to use hand control with vibrotactile feedback. (D) Monkey J learning to use hand control with vibrotactile feedback.

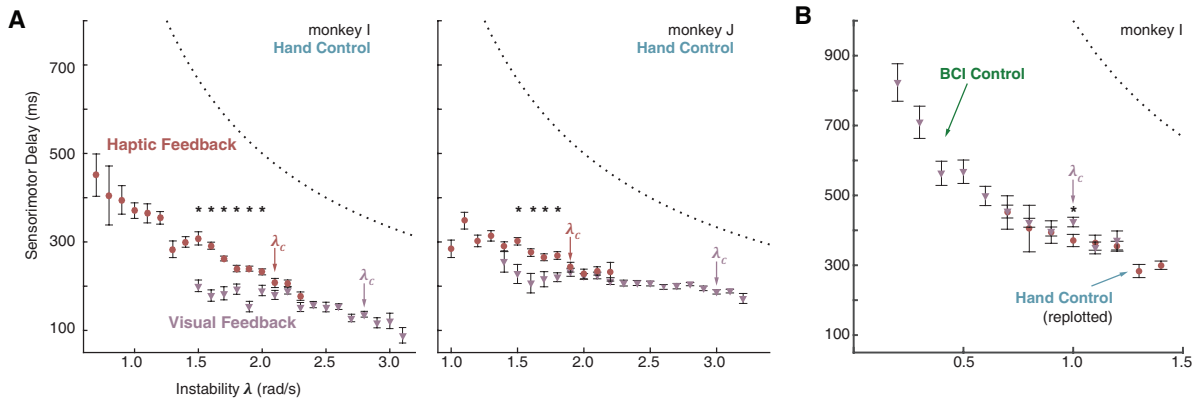

Figure 4.9: Sensorimotor delay during hand control and BCI control. (A) Monkey I and J's visual-hand delays (pink) and vibrotactile-hand delays (orange) as calculated by the cross-correlation analysis. Each point represents the delay (mean ± s.e.m.) at that instability for a minimum of 10 successful trials. The λ_c arrows point to that effector – feedback pairing's λ_c score. The asterisks indicate which vibrotactile-hand and visual-hand delays are significantly different (t-test, $p < 0.05$). (B) Monkey I's visual-BCI delays (pink) and his vibrotactile-hand delays (orange) (replotted from (A)). Both sensorimotor delays follow the same trend.

Next, we looked at the instability levels with overlap between the vibrotactile and visual feedback. At the lower instability levels, we saw that their vibrotactile-hand delays were significantly longer than their visual-hand delays. At the higher instability levels, we saw that their vibrotactile-hand delays were not significantly different than their visual-hand delay. Even with the same sensorimotor delay, the monkeys' performance using visual and vibrotactile feedback differed. At an instability level of 2.2 rad/s, monkeys used visual feedback to successfully complete more than 50% of trials, Figure 4.9. At the same instability level, monkeys stabilized less than 50% of trials using vibrotactile feedback. The difference in performance likely stemmed from a mixture of less accurate motor responses and difficulty in meeting the necessary sensorimotor delay. In order for the monkeys to stabilize higher instabilities with vibrotactile feedback, it is likely their vibrotactile delay would have needed to continue to track their visual delay trend. Unfortunately, the monkeys were not capable of accurately responding more quickly. When trying to make fast sensory judgments, it reasonable that a monkey may misinterpret the vibrotactile feedback and generate an inappropriate motor response.

Before performing the the cross-correlation analysis on the hand control trials, we hypothesized that the vibrotactile delay and the visual delay would be constant across the instability levels. We believed that each delay would equal their respective modality's reaction time, around 220ms for visual and 185ms for vibrotactile (Godlove et al. 2014). After the cross-correlations were performed, and we saw that the sensorimotor delays were not constant, but depended on the instability level. More specifically, we then found each monkey's visual λ_c level and compared his visual delay at that instability level to the nominal visual reaction time of 220ms. Monkey I's visual delay at 2.8 rad/s was 137ms, monkey Js visual-hand delay at 3.0 rad/s was 187ms. The visual-hand delay of both monkeys 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). By using a forward model, they could estimate the system's position and generate a movement in response to the estimated position. Next, we compared each monkeys vibrotactile delay at his vibrotactile λ_c level. Monkey I's vibrotactile delay at 2.1 rad/s was 207ms, monkey J's vibrotactile delay at 1.9 rad/s was 243ms. The vibrotactile delay of both monkeys was higher than

the vibrotactile reaction time. Our interpretation is the monkeys were not able to form a salient, accurate perception of the vibrotactile feedback. Without an salient feedback signal, the monkeys 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.

Finally, we looked at monkey I's visual-BCI delay. Because BCI control yielded the lowest λ_c score, we wondered how visual-motor delays brought about this performance. When looking at the visual-BCI delay, Figure 4.9B, the delays at first appear astoundingly large. However, if one compares the visual-BCI delay to the vibrotactile-hand delay (replotted in figure9B), the visual-BCI delay lies on top of the vibrotactile-hand delay. It is very surprising to find that two effector–feedback pairings, which share no overlap between effector or feedback modality, demonstrate the same sensorimotor delay trends. We believe this is because both pairings are abstract, rather than natural, and as such, the monkey responds to the constraints of the CST with the same sensorimotor delay. However, monkey I was not able to follow the sensorimotor delay trend; he attained a BCI control-visual feedback λ_c score of 1.0 rad/s. Because monkey I was able to attain a high λ_c score with hand control as the effector, we believe the BCI control-visual feedback suffered because of the BCI decoder. The BCI decoder's response time was too slow to generate the corrective movements necessary to stabilize high instability levels. This may in part be due to the integration step of the velocity Kalman filter. The mathematical process of integration acts like a low-pass filter, which smoothes and delays rapid changes. We may be able to improve BCI decoders by increasing their agility by decreasing the time they take to change directions.

4.4 DISCUSSION

Our goal was to create a single metric to simply quantify the effectiveness of a particular form of non-visual feedback. We accomplished our goal by developing the Critical Stability Task (CST). In the CST, monkey subjects stabilized unstable systems at different instability levels. By varying the instability level, we developed a continuous gradation of task difficulty.

More effective feedback enabled the monkeys to successfully complete more difficult motor tasks. For a particular form of feedback, we could find the maximal instability that could be stabilized 50% of the time, which was termed the critical instability level. As a proof of concept, we tested the effectiveness of visual feedback and vibrotactile feedback during hand control. We then used the visual feedback λ_c level to establish a standard for effective control. We found that monkeys were able to learn to use vibrotactile feedback 61-75% as effectively as visual feedback. Using the CST, we were able to successfully determine small differences in the effectiveness of a particular form of sensory-motor control. We were also able to compare hand control - visual feedback with BCI control - visual feedback. We found that with our BCI decoder, the monkey was only able to use the BCI decoder 37% as effectively as hand control. The low effectiveness of the BCI decoder as compared to hand control shows there is still much room for BCI decoders to improve before they approach natural motor control.

Additionally, 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. However, our experience with the vibrotactile learning process also provides a word of caution. The monkeys first became familiarized with the tactors. They 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 with the vibrotactile mappings, they were not initially successful at the CST. The monkeys needed many more weeks to learn how to quickly utilize these vibrotactile signals in order to perform complex movements. In a simple, straight trajectory task used to familiarize the monkeys with the tactors, the monkeys were less pressured to quickly use the feedback. Alternatively, with the CST, the monkeys were continuously pressured to react to the feedback within specific period of time. 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.

Unstable systems are a beautiful way to study and assess the sensorimotor system. Several psychomotor studies have looked at how humans control unstable systems. In 1966, Jex et al.developed the original Critical Tracking Task as a way to characterize a pilots ability to control an unstable aircraft. Researchers found the maximal instability a pilot could control using different information display systems, or in the face of a distracting task. Since then, Kadkade et al. used the Critical Tracking Task to test vibrotactile feedback displays for balance prosthesis research (2003). Modified versions of the Critical Tracking Task have been used in human subjects to study motor impairment during alcohol use and high potency marijuana use (Ramaekers et al. 2006a, Ramaekers et al. 2006b), diphenhydramine and alcohol use (Burns and Moskowitz 1980), assessment of vehicle driver distraction (Petzoldt et al. 2013), and assessing a patient's neurological function (Potvin et al. 1977).

A unique aspect of the CST is that the task is generalizable and customizable. The CST can assess virtually any effector–feedback pairing. In our experiments, the monkeys received either visual or vibrotactile feedback. The feedback signals could also have been other vibrotactile feedback methods, intracortical microstimulation, or an abstract visual signals. In our experiments, the effector was hand control and a BCI decoder. The effector could have also been pinch force, a different decoding algorithm, or different neural recording technique. Most importantly, all of these sensorimotor pairings could be assessed using the same critical instability metric. In addition to assessing various effector–feedback pairings, the CST can also be customized to study aspects of sensorimotor control. Instead of relying on the monkey's motor noise to generate instabilities, noise can also be injected into the system. Different types of noise could be injected, ranging from white noise, or noise within a certain frequency band. Jex et al. injected a sum of sinusoids of different frequencies into the system. He then looked at the system's position and the resulting stabilizing hand movements, and calculated the frequency response of the closed-loop system. Additional changes could include increasing the duration of the CST to capture longer sensorimotor relationships, or changing the system into a two-dimensional system, or a second-order system. Even with all of these possibilities, the consistent, underlying theme of the CST is to test a subject's ability to quickly and accurately respond to the feedback signal.

Because feedback controllers are so sensitive to delays, it was important to look at the sensorimotor delays of our tested effector–feedback pairings. At high instability levels (>2 rad/s) using hand control, the monkeys' vibrotactile delays were equivalent to their visual delays. However, the monkeys were less than 50% successful when using the vibrotactile feedback and over 50% successful when using the visual feedback. As the difficulty increased, he vibrotactile delays began following the same sensorimotor delay trend as the visual delays. However, the monkeys were not as successful at using vibrotactile feedback to generate accurate movements at these sensorimotor delays. After the vibrotactile λ_c level, the monkeys became less and less capable of meeting the delay and accuracy requirements using vibrotactile feedback. This may indicate that the monkeys' vibrotactile-motor delays were intrinsically higher, and they had to exert effort in order to reduce them. While the monkeys were able to reduce their visual-motor delays below the nominal visual reaction time, the monkeys were not able to make their vibrotactile delay less than the nominal vibrotactile reaction time. Forward models are a possible method for successfully mitigating sensory feedback delays (Miall and Wolpert 1996, Mehta and Schaal 2002).

There are two theorized components of forward models, a motor prediction and a sensory prediction. In our experiments, we varied the feedback signal (visual feedback or vibrotactile feedback) while holding the feedforward command constant (hand control). Because the motor component was the same across both feedback conditions, we could make comparisons between how the two forms of feedback were utilized to make motor decisions. Under the constraints of the CST, the monkeys could achieve the same sensory-motor latency for visual and vibrotactile feedback. However, the monkeys achieved different performance levels under visual and vibrotactile feedback. This first piece of evidence suggests that it was the feedback modality which hindered performance, not feedforward motor prediction errors. At this point, it is pertinent to reiterate that the motor prediction is the prediction of the unstable system's position, not the hand position. The second piece of evidence is that monkey I achieved a visual-motor delay of around 100ms, below what would be expected by visual reaction time. However, neither monkey was able to success use vibrotactile-motor delays lower than expected by vibrotactile reaction time. All vibrotactile-motor delays were above the expected vibrotactile reaction time. These two pieces of evidence suggest that an

internal sensory prediction is a component of the forward model (Mehta and Schaal 2002, Crevecoeur and Scott 2013). Furthermore, the monkeys built an internal sensory prediction of the system's state when provided visual feedback, and used this prediction to lower their visual-motor delay. Conversely, the monkeys were unable to build an internal sensory prediction of the system's state when the system's state was conveyed through vibrotactile feedback. When provided vibrotactile feedback, they were unable to lower their vibrotactile-motor delay while still achieving good motor performance. The differences in visual- and vibrotactile-motor delays suggests that in order to 'intuitively' use a new form of sensory feedback, a monkey needs to learn to how to incorporate that feedback signal into a forward model. In order to design adequate sensory substitution systems, it will be important to understand how internal sensory prediction models are used and learned (Weber et al. 2012).

Learning will be extremely important for the next sensory feedback modality we wish to test, intracortical microstimulation. Much research has been going into the area of intracortical microstimulation feedback (Mountcastle et al. 1990, Romo et al. 1998, Fitzsimmons et al. 2007, London et al. 2008, O'Doherty et al. 2011, Dadarlat et al. 2015). Many tests look at the discriminability of between two or three discrete feedback signals. Recently, Dadarlat et al. 2015 performed the first continuous ICMS feedback to guide reaching movements. Between these research projects, various stimulation frequencies, amplitudes, numbers of electrode, stimulation patterns, and cortical areas were used. In order to make objective decisions about which stimulation parameters actually produce meaningful feedback signals, the community needs to adopt a standardized method for testing feedback effectiveness. Standard point-to-point reaches have been the gold standard task for testing brain-computer interface control with visual feedback. However, point-to-point reaches are not well-suited for studying the effectiveness of feedback. The hallmark of good feedback is enabling the monkey to quickly and accurately correct their movements. Point-to-point reaching contains a high degree of ballistic, pre-planned motion. With such straight paths, it is difficult to discern the point in time at which the monkey updates his trajectory based on new information. Alternatively

with the CST, monkeys are stabilizing an unstable system, which requires feedback. Because the system is unstable, any slight motor error will be exponentially amplified until the corresponding corrective movement is made. As a result, the monkey's movements are always changing speed, and in a direction. The changes in speed and direction make it easier to correlate how their movements were a response to the feedback.

5.0 NEURAL REPRESENTATION OF REACH AMPLITUDE

Figures and text in this chapter serve as the beginning of a journal article concerning the neural representation of large and small arm movements. Much of the analyses were performed under my mentorship of an undergraduate student, Nicholas Card.

While I was using the Critical Stability Task during BCI control in Chapter 4, I noticed that the BCI decoder was able to decode large movements, but had trouble difficulty decoding small movements. This was concerning to me because small movements, are crucial for stabilizing small instability levels. Nicholas Card and I then endeavored to try to determine why the BCI decoder was not working well for small movements. The neural representation of large and small movements became a topic of interest when I realized that the BCI community has been focused on gross movements at the expense of studying fine motor control. 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.

5.1 INTRODUCTION

While researching ways to train decoders for the monkeys during the CST task, I debated between attempting observation training and hand control training. For the first attempt with observation training, the monkey did not have very good control over the CST cursor. For the second attempt with monkey I, I used hand control while performing center-out-center trials to train the decoder. After training the decoder, the monkey was allowed to

keep his arm free and began using the velocity Kalman filter decoder for the CST. I was seeing his arm make very small movements and the decoder making very large movements. I thought this was very strange. I had noticed this observation before with monkey J, however, I had assumed that it was because he did not understand how to do the task at that point in time, much less BCI control. Additionally, even though the BCI cursor made very large amplitude movements, his hand movements were small. He learned over the course of the day and following days how to reduce the size of the BCI decoder movement. This experience caused me to ponder how the brain actually encodes large verses small movements.

The large versus small movement question is particularly pertinent to BCI control. During BCI control, if the user makes a straight accurate reach towards the target, he is able to achieve success. Problems, however, come into play when the BCI cursor approaches the target but there is a small correction that must be made. BCI's typically struggle to make the small corrective movement. Instead, the cursor may 'spiral in' or make a series of small, but not quite small enough corrections until one of them lands the cursor in the target for the desired about of time. The action is similar to a game of putt-putt, where you make one pretty accurate put to reach the hole, and then you could spend the next four putts just slightly missing the hole. Golub et al. 2014 noted this BCI difficulty by the hold time requirement on the decoder. As the hold time increased from 50 to 600ms, successful target acquisition dropped from near 100% to near 30%.

So, how does the motor cortex actually handle movements of drastically different sizes? How will we create decoders that enable users to throw a baseball and write their name? Does the neural activity have a fixed firing rate to velocity mapping, as linear BCI's presuppose? Or is the mapping between velocity and neural activity scale depending on the movement that is about to be made? To investigate these questions,we studied two monkeys performing large and small magnitude reaches. We also looked at one monkey making large and small magnitude BCI reaches. We looked to see if the velocity to firing rate mapping held between the two different contexts of reaching. We found that the mappings were not constant, but the dynamic firing range of the neurons rescaled to cover a different set of velocities. This was found to be true for hand control and BCI control. While we did find this rescaling to be taking place, it did not fully 100% cover the previously large dynamic range, but it

covered considerably more than to be expected by a linear relationship between firing rates and velocity. This demonstrates the need to investigate non-linear BCI decoders that would better enable the switch between large and small BCI movements.

5.2 METHODS

There were three monkeys involved in total. Two monkeys, Monkey I and Monkey N performed natural arm reaches. Monkey I had his 96-channel Blackrock array implanted in PMd. Monkey N has his array implanted in M1. A third monkey, Monkey L, performed BCI reaches. His array was also implanted in M1. All three monkeys sat in the same virtual reality environment. Monkey I and Monkey N had their reaching arm free. An powered LED marker was attached to Monkey I's wrist and to Monkey N's finger. Six motion capture cameras recorded the three dimensional movements of the marker. The real-time position of the marker was rendered back to them as a cursor in a virtual 3D environment. All of the visual stimuli were displayed in a fronto-parallel plane at a fixed depth. The two monkeys did not see their hand because it traveled in the space behind the virtual 3D environment.

Monkey I had his arm restrained for the duration of the experiment. He still wore a powered LED marker on the back of his hand to record any hand movement, but his arm was gently restrained in an arm rest. He saw his decoded BCI position as a cursor moving in the same virtual 3D environment as Monkey's I and N.

5.2.1 Examining neural parameters during natural reaching

Monkey I performed a version of the standard center-out reaching task in which the targets were spaced around a circle with one of six radii, which included 7mm, 10mm, 13mm, 35mm, 50mm, and 65mm. The first three radii were variations on a small magnitude reach, each of which had a circular acquire window with a radius of 3mm. The last three radii were all variations on the large magnitude, and each had a circular acquire window of 5mm. The monkey first performed a block of trials with large magnitude reaches followed by a block of

trials with small magnitude reaches. Monkey I's neural data collected with a threshold of 4.0xRMS and sorted offline into 37 multi-units. Next, we combined the data together from the three small radii and made tuning curves for each multi-unit. We also combined the data from the three large radii and made tuning curves for each multi-unit. For all tuning curves, we used the neural activity between target onset and 200ms after target onset. Within this same time span, we also calculated the speed of the monkey's hand.

Monkey N also performed a center-out reaching task to three different radii, which included 15mm, 60mm, and 120mm, with circular acceptance windows of 5mm, 7mm, 9mm. Monkey N's neural activity was collected with a threshold of 3.0×RMS. We removed 16 bad channels to remain with 80 neural units composed of threshold crossings. Similarly, we calculated tuning curves for each neural unit at each of the three radii using the neural activity between target onset and 200ms after target onset. Within this same time span, we also calculated the speed of the monkey's hand.

We then wished to analyze the changes in the tuning curve parameters for the different size reaches. We compared the preferred direction, modulation depth, and baseline of tuning curves while the monkeys were reaching to targets situated around the small radius circle, or targets around the large radius circle. For monkey N, we also compared the neural activity while reaching to the small radius circle to the neural activity while reaching to the medium radius circle and to the large radius circle. For each neural unit, we subtracted the small radius' tuning curve parameters from the large radius' tuning curve parameters. We then used a one-sample t-test with a 95% confidence interval to test if there was a significant difference between the tuning curve parameters. A two-sample t-test with a 95% confidence interval was used to test for a difference in reach speeds for the different reach distances.

5.2.2 Examining online BCI control during large and small movements

We saw large and non-linear transformations between the neural activity during short distance movements and neural activity for long distance movements. Most decoders in use today assume a linear relationship between movement velocity and firing rate. Usually, the baseline firing rate of the neural decoders is fixed, the preferred direction is fixed, and the

modulation rate changes for different movement velocities. Even though this relationship did not hold true for natural reaches, we wanted to see if the relationship would hold true for BCI reaches. If the relationship did hold for BCI reaches, then BCI's are on solid footing for making long and short distance reaches. However, if the relationship also does not hold true for BCI reaches, then serious BCI redesigns will be necessary. BCI users must be able to use a BCI to make long distance movements, as well as short distance or corrective movements.

For this experiment, monkey L was implanted with an 96-channel array into primary motor cortex. He had undergone extensive training in using a velocity Kalman filter decoder use prior to the collection of the data for this experiment. We desired to have monkey L use a BCI decoder to make short and long distance movements, under two different contexts. First we wanted to see how well he could use a decoder trained on short distance movements to move short distances. We also wanted to see how well he could use a decoder trained on long distances to move long distances. Second, we wanted to see how a decoder simultaneously trained on short *and* long distance movements could move to short *and* long distances. Our idea was that if neural activity does show a linear relationship to movement velocity, then training a decoder on short and long movements would most adequately capture the span of this relationship. If the relationship between neural activity and movement velocity is non-linear, then the decoder trained on both short and long movements would do worse than individually training decoders on short or long movements.

In the four experimental sessions being analyzed from monkey L, we recorded threshold crossings from 88-89 channels at 3.09×RMS. The decoders trained solely on short distance movement or solely long distance movements were trained in similar fashion. The monkey first observed the cursor moving to each of the 8 targets two times. Using this observation data, we trained the first iteration of a velocity Kalman filter, Section 4.2.2.5. Next, the monkey used this velocity Kalman to reach to each of the 8 targets two times. The decoded cursor movement's velocity was restricted to 0m/s in the direction perpendicular to the straight-line movement towards the target. The second iteration decoder was trained off of this data. Next, the second iteration decoder was used to reach to each of the eight targets one time, without restrictions. The third iteration decoder was trained off of the second and

third iteration data. Three more times we repeated the process of having the monkey make BCI reaches to each of the 8 targets and retraining a decoder using the new and old data. The final decoder was a result of 48 trials, Section 6.2.4.

When the monkey was making BCI reaches to these targets, the short distance BCI reaches were 25mm away from the center (circular acceptance window, 5mm diameter). The long distance reaches were 200mm away from the center (circular acceptance window, 20mm diameter). Upon reaching the target, the monkey had to hold inside of the acceptance window for 50ms. During the first training round, When the monkey was observing the cursor move to these targets, the cursor always took 1s to reach the target. This made it so that the cursor moved with a smaller velocity when moving to the short distance targets and a higher velocity when moving to the long distance targets.

In addition to training a decoder solely on long or short distance movements, we also trained a decoder simultaneously on short and long distance movements. In order to do so, we went through a similar three-step gradual training process. First, the monkey observed the cursor moving to each of the 8 short distance targets twice, and each of the 8 long distance targets twice, for a total of 32 observation trials. We trained the first iteration decoder on the observation trials. Next, the monkey used these observation trials to move to each of the 16 targets four times, for a total of 64 trials. The decoded cursor movement's velocity was restricted to 0m/s in the direction perpendicular to the straight-line movement towards the target. The second iteration decoder was trained off of these BCI movements. Next, the second iteration decoder was used to reach to each of the 16 targets one time, without restrictions. The third iteration decoder was trained off of the second and third iteration data. Four more times we repeated the process of having the monkey make BCI reaches to each of the 16 targets and retraining a decoder using the new and old data. The final decoder was a result of 144 trials.

When the monkey was making BCI reaches to these targets, the short distance BCI reaches were 25mm away from the center (with an circular acceptance window with a diameter of 5mm). The long distance reaches were 200mm away from the center (with a circular acceptance window with a diameter of 20mm). Upon reaching the long distance targets, the monkey had to hold inside of the acceptance window for 50ms. Upon reaching the short

distance targets, the monkey had to hold inside of the acceptance window for 10ms. We made this change to the hold time for the small targets because the monkey was having a very difficult time meeting the hold requirements. During the first training round, When the monkey was observing the cursor move to these targets, the cursor always took 1s to reach the target. This made it so that the cursor moved with a smaller velocity when moving to the short distance targets and a higher velocity when moving to the long distance targets.

After these three types of decoder were trained, we tested each decoder type on the same long and short distance movements used during training. The hold time was 50ms for each of the targets. We calculated the percentage of trials in which the monkey successfully acquired the target.

5.3 RESULTS

5.3.1 Tuning curve changes during natural reaching

We measured the changes in the tuning curves between reaching a small distance and reaching a large distance. The first tuning curve parameter that we looked at was the baseline of the tuning curve. In Figure 5.1A, we can see the distribution of monkey I's baseline firing rates. The distributions are overlapping, but when we look at how the baseline changes for each neural unit, we find that there is a significant trend. Overall, the baseline firing rate for the large reaches was 4Hz greater than the baseline for the small magnitude reaches ($p = 3.32e\text{-}6$). This trend carries through for monkey N as well. For monkey N, we first compared the 15mm reaches to the 60mm reaches, Figure 5.1B. Overall, monkey N's baseline firing rates were more consistent from neural unit to neural unit. When we looked at the difference between the 15mm reach and the 60mm reach, we found a narrow difference distribution. Monkey N's baseline firing rates were on average 3.15Hz larger for the 60mm reaches than for the 15mm reaches ($p = 2.54e\text{-}17$). Additionally, we looked at the comparison between monkey N's 15mm reaches and 120mm reaches, Figure 5.1C. At the 120mm reach magnitude, the baseline firing rates was 3.1Hz larger than the 15mm reach magnitude's baseline firing

rates (p = 9.63e-15). The similar changes in baseline firing rates between monkey I and monkey N point to a neural change between making medium and large reaches, and make small reaches.

Next, we looked into how the modulation depth changes as a function of reaching distance. In Figure 5.2A on the left, we can see the distribution of the population's modulation depths from monkey I. Next, we take the difference between each neural unit's modulation depth for a long distance reach and a short distance reach. We see that this distribution of differences in modulation depth is not significantly different from zero. In Figure 5.2B, we look that the population's modulation depths during 15mm reaches and 60mm reaches in monkey N. When we look at the difference between these two distributions, we do find a significant difference. The modulation depth is 2.52Hz larger for the 60mm reaches than for the 15mm reaches (p = 2.31e-6). While the peak of this difference distribution is slightly greater than zero, there are several neural units that greatly increase their modulation depth with the longer distance reach. Finally, we look that the population's modulation depths during 15mm reaches and 120mm reaches in monkey N, Figure 5.2C. When we look at the difference between these two distributions, we do find a significant difference. The modulation depth is 3.59Hz larger for the 120mm reaches than for the 15mm reaches (p = 2.31e-6).

We also investigated the third parameter of the tuning curve, the preferred reaching direction. We looked at how the preferred directions changed from one reaching distance to another, but our findings were inconclusive. Monkey I's preferred directions were spaced fairly evenly around the circle. We also saw changes in preferred direction, sometimes large changes. Because of the size of these changes, and relatively few number of neurons, we cannot assess whether there were preferred direction changes. Monkey N's preferred directions were also distributed around the circle, but there was a clustering around 160°. The changes in monkey N's preferred direction centered around 0°, but again, we also saw many large changes in preferred direction. At this point in time, we cannot draw any solid conclusions about preferred direction changes as a function of reach distance.

The final portion of this analysis was to tie the neural data back to what the arm was doing during the same period of time. For both monkeys, we looked for the differences in reach speed between short distance reaches and long distance reaches. Starting with

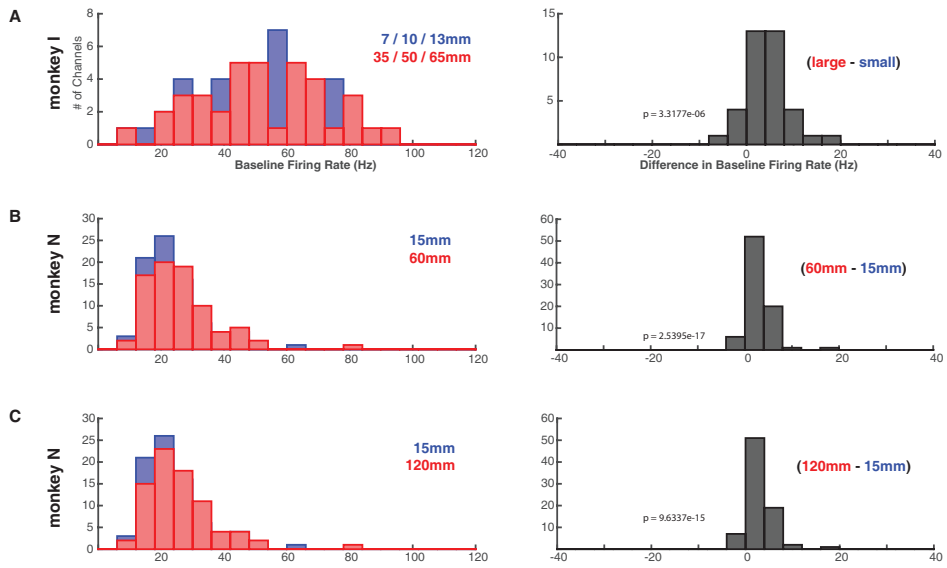

Figure 5.1: Population level changes in baseline firing rates during natural reaching. Red bars indicate the longer distance reaches, the blue bars indicate the shorter distance reaches, and the gray bars indicate the channel-by-channel firing rate difference between the two distances. (A) Comparison of monkey I's baseline firing rates between the grouping of short distance (7mm, 10mm, 13mm) and long distance (35mm, 50mm, and 65mm) reaches. (B) Comparison of monkey N's baseline firing rates between the 15mm reaches and 60mm reaches. (C) Comparison of monkey N's baseline firing rates between the 15mm reaches and 120mm reaches.

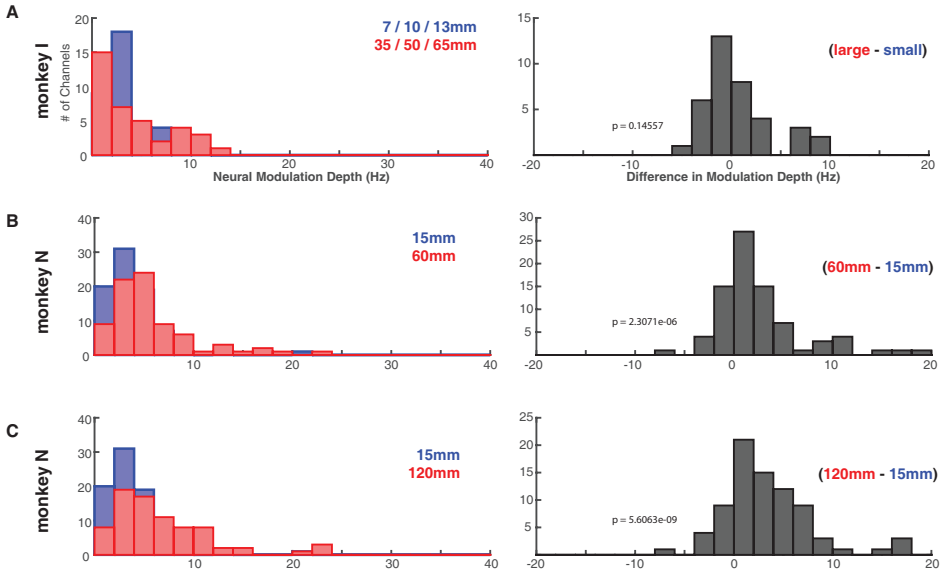

Figure 5.2: Population level changes in modulation depths during natural reaching. Red bars indicate the longer distance reaches, the blue bars indicate the shorter distance reaches, and the gray bars indicate the channel-by-channel firing rate difference between the two distances. (A) Comparison of monkey I's modulation depths between the grouping of short distance (7mm, 10mm, 13mm) and long distance (35mm, 50mm, and 65mm) reaches. (B) Comparison of monkey N's modulation depths between the 15mm reaches and 60mm reaches. (C) Comparison of monkey N's modulation depths between the 15mm reaches and 120mm reaches.

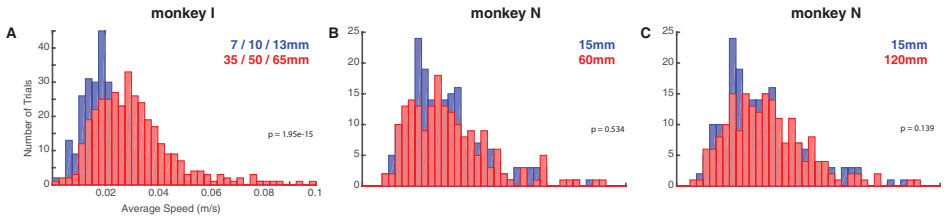

Figure 5.3: Changes in reach speed during short and long distance reaches. Average reach speed for each trial, lasting 200ms from target onset. (A) Monkey I, reach speeds for the grouping of short distance reaches (blue) and long distance reaches (red). Long distance reaches were significantly faster. (B) Monkey N, reach speeds for the 15mm reach distance (blue) and the 60mm reach distance (red). There was not a significant difference in speed. (C) Monkey N, reach speeds for the 15mm reach distance (blue) and the 120mm reach distance (red). There was not a significant difference in speed.

monkey I, we did find a significant difference in average speed between the long distance reaches and the short distance reaches 200ms after target onset ($p = 1.95e\text{-}15$), Figure 5.3A. Interestingly, this finding did not hold for monkey N, Figure 5.3B-C. For both comparisons in monkey N, there was no significant difference in movement speed between the 15mm reaches and the 60mm, or the 15mm reaches and the 120mm. Looking between the three subplots in Figure 5.3, it seems like this is due to monkey N using a faster speed for short reaches than monkey I used for short reaches. Now, let's take this knowledge and use it to understand the tuning curve changes that we saw.

With monkey I, we did observe a change in the baseline of the neural units' tuning curves between the short reach distances and the long reach distances. In monkey N, we saw similarly sized changes in the baseline firing rate between the short reach distance and the two longer reach distances. It is interesting, however, that only monkey I changed his speed between the two reach distances, inside the analysis time window. Monkey N had the

same speed between 15mm reach and 60mm reach, and between 15mm reaches and 120mm reaches. Therefore, the change in baseline signal the upcoming reach distance, regardless of the whether the monkey has actually changed his speed.

The findings become more difficult to interpret as we switch to modulation depth. Monkey I did not show a change in modulation depth between short and long reach distances; however, he did show a difference in reach speed. Monkey N did show a change in modulation depth between short and long reach directions; however, he did not show a difference in reach speed. How can these contradictory findings be so? The answer may lie in the brain area from which the neural recordings were taken. In monkey I we recorded from PMd and in monkey N we recorded from M1. From our data, it would show that on average, PMd does not change it's modulation depth as a function of reach distance. From monkey N, our data show that M1 does change it's modulation depth as a function of reach distance. Both of these findings are potentially problematic. They highlight the lack of a systematic relationship between neural activity and reach distance and reach speed. As the neural activity shows, reaching more slowly is not neurally equivalent to reaching a short distance. In order to justify these strong claims, more neural data would need to be collected from PMd and from M1.

5.3.2 Examining online BCI control during large and small movements

For monkey L, we trained three different types of decoders. One decoder was trained on short distance movements, one decoder was trained on long distance movements, and one decoder was trained on short and long distance movements. Monkey L then used decoder trained on short distance movements to reach short distances during two different experimental sessions, Figure 5.4A. The first day plotted was the first day that monkey L had ever attempted at making short distance movements. Even so, he did pretty well, attaining a 79% success rate. On a following session, be used the decoder to reach to short distance targets in two different blocks. His performance was improved on this day, attaining a 88% and 91% success rate on the respective blocks. These results clearly demonstrate that a BCI decoder can be successfully used to reach to short distance targets.

Next, Monkey L then used decoder trained on long distance movements to reach long distances during three different experimental sessions, Figure 5.4B. Again, monkey L was able to use his decoder to successfully reach to these targets. His performance improved slightly across days. Even though he was well trained in BCI control, he had not previously experienced the 50ms hold requirement. This requirement took some practice, but he was able to improve. Two blocks of long distance reaches were run on each of the three experimental sessions. The monkey's performance was usually slightly worse during the second block. We believe this was partially due to the fatigue from the BCI control block(s) which occurred in between the two long distance reach blocks. Overall, monkey L did well at performing long distance reaches with the long distance decoder.

Finally, Monkey L used the decoder trained on short and long distance movements, Figure 5.4C. To test this decoder, we interleaved trials of long distance reaches and short distance reaches. His combined long and short distance performance is shown in the purple bars in figure 1C. His overall performance was mediocre, around 70% success. To determine the source of the poor control, we separated the trials into long distance and short distance trials. When looking at only the long distance trials, monkey L did very well. He did just as well with the decoder trained on both distances as he did with the decoder trained on only long distances. However, we see a different story when we look at the short distance reaches. Monkey L did poorly when trying to acquire the short distance targets. He attained a 48% success on the first experimental session and 40% success on the second experimental sessions. This result shows that simply training on the entire velocity to be used will not result in a better decoder. The neural activity does not linearly map between short and long distance reaches. What was intriguing what that decoder trained on both distances did not detriment the monkey's ability to acquire the long distance targets. Only the short distance targets were affected.

The detriment to only the short distance targets could have been for one of two reasons. During the training process, we noticed that the monkey was having a difficult time appropriately reaching to the short distance targets, but fairly good at reaching to the long distance targets. He would move to quickly pass the short distance targets to successfully acquire them; he did not slow down. This could have been for one of two reasons. First,

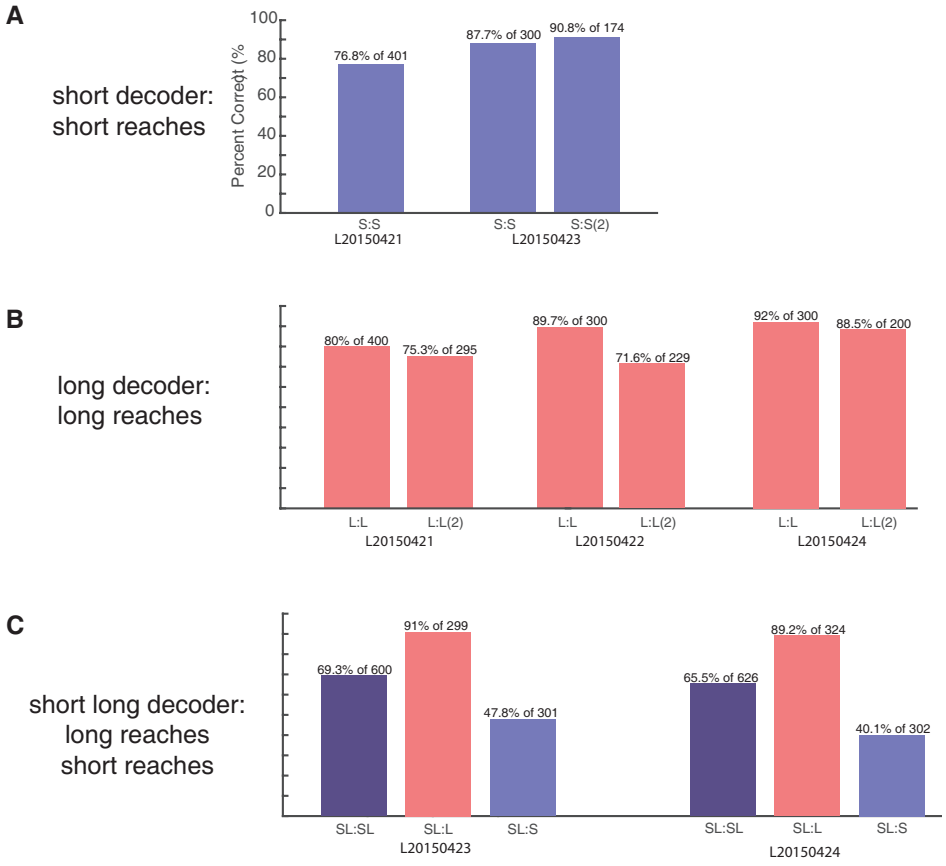

Figure 5.4: Online BCI performance using decoders trained on three types reaches. Blue bars indicate BCI control while reaching to short distance targets. Red bars indicate BCI control while reaching to long distance targets. Purple bars indicate the combined performance of reaching to short and long distance targets. (A) Monkey L using a decoder trained on short distance reaches. (B) Monkey L using a decoder trained on long distance reaches. (C) Monkey L using a decoder trained on both short and long distance reaches. The first bar in each day indicates the combined performance across long and short distance reaches. The subsequent two bars break the performance into the long and short distance trials.

maybe he was not attempting to slow down. One possible strategy would have been to just use the same intention to hit the short and long distance targets. Maybe because in the monkey's previous gradual training experience, he only needed to acquire targets at a single distance. He was not practiced at modulating his neural activity for different distances. It is possible that with further, sustained practice, the monkey may have been able to differently modulate his neural activity for short and long distance reaches. Alternatively, his neural activity could have been responding naturally to the short and long distance reaches, and the decoder being trained was unable to recognize the neural signature for both reach distances. We hypothesize, that the second reason is true - that the decoder could not sufficiently understand the neural signatures for long and short distance reaches.

The ideal way for monkey to use the linear velocity Kalman filter decoder would be to keep a consistent preferred direction and baseline firing rate between the two conditions, but increase the neural modulation depth when attempting to reach higher velocities to long distance targets. On one experimental day, we compared the neural tuning properties while the monkey used the decoder trained on small distances to reach to small distances to the neural properties of using the decoder trained on long distances to reach to long distances, Figure 5.5A-C. We found a significant difference between the two distance conditions for all three tuning properties. Comparing the short distance to the long distance, we saw an average preferred direction change of 22°, modulation depths increase by 5.18Hz, and baseline firing rates increase by 2.1Hz.

The changes in baseline and preferred direction do not bode well for using a single velocity Kalman filter decoder. However, we went ahead and tried to build one to monkey L to use. On a second experimental day, we compared the neural tuning properties while the monkey used the decoder trained on both distances to reach to small distances to the neural properties of using the decoder trained on both distances to reach to long distances, Figure 5.5D-F. We see that the monkey did attempt to further differentiate his modulation depths between short and long distance reaches. He increased the difference in modulation depth from 5.18Hz to 12.76Hz. Additionally, (or consequentially), the difference between his preferred directions and baseline firing rates also changed, (41.8°, 3.47Hz). Unfortunately, even with this now large difference in modulation depth, the monkey still had difficulty in

slowing the decoder down enough to acquire the short distance targets. The tuning curves generated from using an ideal BCI decoder would have shown the monkey using the same neural activity patterns to between the decoder trained on short distances, and a decoder trained on short and long distances. The fact that the monkeys neural activity changed between the two decoders shows that the decoder trained on short and long distance reaches does not accurately capture the neural encoding of short distances.

5.4 DISCUSSION

Our interest in the motor cortical representation of distance stemmed from our desire to improve better brain-computer interfaces (BCIs). We know that BCIs have trouble 'honing-in' on the target of interest. They circle or pass-through the target, unable to hold still. We wondered if there was a different neural representation for small movements, one that was not being understood by the decoder.

Through our experiments with natural and BCI movements, we looked for changes in neural tuning parameters. First, we looked into preferred direction changes. During natural arm movements, the preferred direction changes were variable and more analyses are needed to determine if these changes are real. During BCI movements, while we noted a definite change in preferred direction, but the interpretation is unclear. Monkey L was an experience BCI user, and his preferred directions had converged to approximately the same preferred direction. The cause of this neural clustering has been under investigation by our lab for a number of years, but a clear conclusion has yet to be made. Since we do not know the cause of neural clustering, it is difficult to interpret the meaning of monkey L's preferred direction changes. When looking to previous literature, we found contradictory results. Fu et al. found that preferred directions varied little to different movement distances (1993). Moran and Schwartz found that preferred directions were invariant movement speeds (1999). However, Churchland et al. saw preferred directions which varied greatly with movement distance and movement speed (2006). At this point in time, we can conclude only that preferred directions do seem to change, however, more investigation is needed.

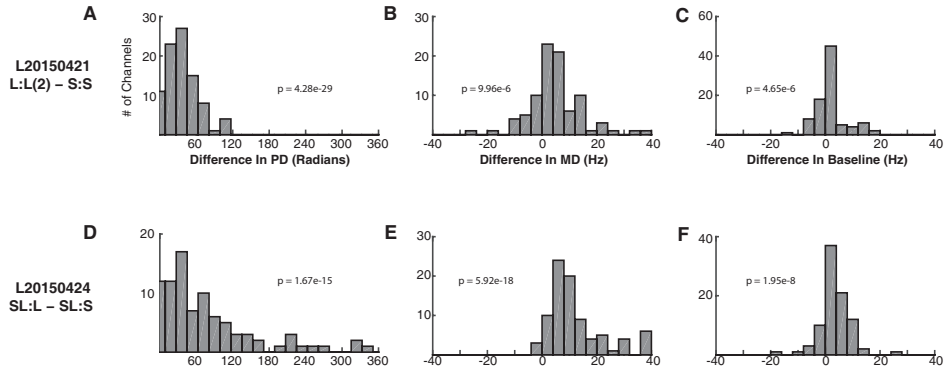

Figure 5.5: Tuning curves changes between short and long BCI reaches. First row: Differences in tuning curves between using a decoder trained and used on long reaches (L:L(2)) and a decoder trained and used on short reaches (S:S). (A) Between the long and short reaches, the preferred directions changed 22°. (B) Between the long and short reaches, the modulation depths changed 5.18Hz. (C) Between the long and short reaches, the baseline firing rate changed 2.1Hz. Second row: Differences in tuning curves between using a decoder trained on short and long reaches, and tested on long reaches (SL:L(2)) and a decoder trained on short and long reaches, and tested on short reaches. (D) Between the long and short reaches, the preferred directions changed 41.8°. (E) Between the long and short reaches, the modulation depths changed 12.76Hz. (F) Between the long and short reaches, the baseline firing rate changed 3.47Hz.

Next, we looked into the baseline firing rate as a function of movement distance. We consistently saw a small, but significant, shift in the baseline firing rate. When the monkeys made short distance movements, their neural tuning curves on average lowered 2-4Hz. This was also true for monkey L, while he was making short distance BCI movements. Additionally, we looked into the modulation depth as a function of movement distance. In monkey I, which was implanted in PMd, we did not see a change in modulation depth. In monkeys N and L, which were implanted in M1, we saw that the modulation depth decreased for short distance movements. Again, this change was small but significant. With monkey L, we tried to build a single representation which incorporated short and long movement distances. When we used this single representation to make BCI movements to long distance targets, the BCI representation worked well. This single representation did not work well for trying to make BCI movements to short distance targets. Even though monkey L tried to use the single representation by increasing the firing rate differences between large and small movements, he was still had difficulty obtaining the short distance targets. While other studies have highlighted the significant changes in modulation depth (Fu et al. 1993, Churchland et al. 2006) between different movement distances, we would like to highlight that these are small differences. For example, Fu et al. found that the majority of neurons showed modulation to at least one movement direction. These carefully chosen words highlight the non-linear nature of the neural representation of distance.

Others have also noted this observation, except instead, with a lack of speed modulation. Golub et al. found that there was a much greater prevalence of direction tuning, rather than speed tuning, in motor cortex (2014). While movement distance and movement speed are not the same variable, they are usually positively correlated. The idea that M1 'lacks speed tuning' is becoming more prevalent in the BCI community as a rule of thumb. However, if most researchers were asked whether they thought neurons would fire differently for very small and for very large movements, I believe the overwhelming majority would say yes. With such a large disparity, surely speed and/or amplitude tuning would manifest for such a large kinematic difference? We are finding, however, that kinematic differences are larger than the neural differences. This is most evident with the disparity in BCI performance with the different decoders.

We believe that some form of neural rescaling is occurring as a function of movement speed and/or amplitude. This is not such a strange or new idea. Neural rescaling is commonly seen in the sensory system (Kuffler 1953, Hubel and Wiesel 1962, Rees and Palmer 1988, Simons 1978). Additionally, in 1999, Hepp-Reymond et al. noted that the neural representation of pinch force could be changed based on the context of the task. If the monkey needed to acquire three different pinch force levels, a neuron's firing rate would take on three, evenly spaced firing rate levels inside of it's dynamic range. If the monkey needed to acquire only two pinch force levels, the neuron's firing would take on two, evenly spaced firing rate levels inside of it's dynamic range. The firing rate - to - pinch force representation changed based on the context of the task.

If neural rescaling is context dependent, must a neuron decide in which reference frame the rescaling should occur? Should the neuron rescale it's firing rate to the movement's distance, horizontal or vertical location, speed, reach duration, or multiple parameters simultaneously? This question falls in line with a discussion point in Churchland et al. 2006. They authors hypothesized that perhaps there is no fundamental reference frame. Perhaps, certain movement parameters are specified first, specified by the context of the immediate task. Once the context has been decided, maybe then a task-specific, reference frame emerges.

While their suggestions might be true, a consistent representation is needed in order to build a BCI decoder. However, we need to do a better job of identifying situations where the representation might quickly switch to a different representation. In dire need of research is the fast context change from a large straight movement to a short fast corrective movement. In our data from natural arm reaches and BCI reaches, we have seen large and small movements use different, but overlapping, neural representations. In order to allow BCI decoders to make small corrective movements as well as they make large straight reaches, we will need a way to quickly determine when a context change is imminent. In Golub et al., their strategy was to look for large angular deviations in the cursor velocity (2014). This served as a reliable indicator that the monkey wanted to 'hone-in' and reduce the speed of his decoder. More decoding strategies like this are needed. We should not view these strategies

as 'hacks' or giving up on the attempt to find the 'real' invariant tuning. Such a concept may not exist. If we wish to improve decoders, we will need more carefully for clues in the kinematics or neural data to signal context changes (Shenoy et al. 2003).

One avenue for pursuing these contextual changes will be to look early in the reach. Such as we saw in our data, there were tuning curve changes very early in the reach, between 0-200ms after the target's onset. It has been previously noted that neurons seem to be modulated by speed information early in the reach, and direction information later in the reach (Johnson et al. 1999). One possible explanation maybe that the motor cortex is initializing the reach; one component of that initialization maybe the reach speed. Once the reach has been properly initialized, the reach can unfold more passively, along the initialized reach dynamics (Churchland et al. 2010, Churchland et al. 2012, Mante et al. 2013).

6.0 NEURAL CONSTRAINTS ON LEARNING

Figures and text in this chapter are from Sadtler et al. 2014. In this chapter, I describe my research analyzing the neural constraints on learning. The primary contributor to this work was Patrick Sadtler. As the second author, I have reproduced here the research components to which I contributed design direction, insight, and problem solving skills. My efforts included the background experiments and persistence to convince the team to restrain the monkeys' arms during BCI control, training the intended second monkey, many discussions on the best way to measure the difficulty of each perturbed decoder, and pushing team to systematically analyze the after-effects of using the perturbed decoder.

Assessment tools are as important for BCI control as they are for natural motor control. In Chapters 3 and 4, I discuss my efforts to assess natural sensorimotor control in humans and monkeys. Additionally in Chapter 4, I begin to assess BCI motor control in monkeys. BCI motor control was poor when compared to natural motor control. In try to understand why, we looked into the neural representation of large and small movements in Chapter 5. However, it is unlikely that the representation of large and small movements is the only constraint on poor BCI performance. Here, in Chapter 6, we further look into constraints on neural activity, particularly, the intrinsic neural network. In order to improve a person's ability to use a BCI decoder, it is important to understand the constraints on person's ability to use the decoder. In understanding these constraints, we will be able to work around, or with, them to design an easier, and more learnable, BCI decoder.

6.1 INTRODUCTION

People excel at learning new motor tasks. Starting at a young age, we learn how to walk, ride a bike, and drink from a glass. Unfortunately, there are some movements we cannot make. Even for the trained contortionist or Olympic athlete, the scope of our motor repertoire is constrained by our physiology. The shape and size of our bodies, the geometry of our joints, and the dynamics of our muscles constrain the movements we make. If the physical constraints could be removed, are there other non-physical constraints on our ability to learn new movements? In order to perform a specific movement, our brains must produce a specific pattern of neural activity. Perhaps, our neural activity patterns also constrain our ability to learn new movements.

Neurons modulate their activity together, rather than independently. While performing a specific action, such as reaching to the right, some neurons will tend to increase their firing rates while others will decrease their firing rates. This co-modulation of neural activity creates a framework for studying the structure of the neural network. If a subject performs a variety of motor actions while we record from a population of neurons, we can find the most common neural co-modulation patterns, which we term the intrinsic manifold. Our question becomes, are we constrained by our co-modulation patterns, by our intrinsic manifold? Must our neural activity lie on the intrinsic manifold, or can we learn to move our neural activity patterns away from the intrinsic manifold?

To test whether we can learn to move our neural activity away from the intrinsic manifold, we used a brain computer interface. With brain computer interfaces, the experimenter can completely describe the entire system being studied. We are recording from all of the inputs - the neurons - which describe the cursor movements. With actual physical movements, the movements are generated by entire brain, not just the subset of neurons being recorded. In a BCI paradigm, the user controls a cursor on a computer screen by generating activity patterns across a population of neurons. A BCI offers advantages for studying learning because we can observe all of the neurons that directly control an action, and we can fully specify the mapping from neural activity to action. This allowed us to define which activity patterns would lead to task success and to test whether subjects were capable of generating

them. Previous studies have shown that BCI learning can be remarkably extensive (Ganguly and Carmena 2009, Fetz 1969, Jarosiewicz et al. 2008, Hwang et al. 2013, Rouse et al. 2013, Engelhard et al. 2013), raising the intriguing possibility that most (or all) novel BCI mappings are learnable.

Here we show that the animals could readily learn to proficiently control the cursor using neural activity patterns that were within the intrinsic manifold. However, animals were less able to learn to proficiently control the cursor using activity patterns that were outside of the intrinsic manifold. These results suggest that the existing structure of a network can shape learning. On a timescale of hours, it seems to be difficult to learn to generate neural activity patterns that are not consistent with the existing network structure. These findings offer a network-level explanation for the observation that we are more readily able to learn new skills when they are related to the skills that we already possess.

6.2 METHODS

Two male Rhesus macaques (aged 7 and 8 years) were trained to move a cursor from the center of a screen to one of eight radially arranged targets by modulating the activity of 85-91 neural units (that is, threshold crossings on each electrode) recorded in the primary motor cortex (Figure 6.1). To represent the activity of the neural population, we defined a high- dimensional space called the neural space)where each axis corresponds to the activity of one neural unit. The activity of all neural units during a short time period is represented as a point in this space (Figure 6.2A). At each time step, the neural activity (a green point in Figure 6.2A) is mapped onto a control space (black line in Figure 6.2A); two-dimensional plane in the actual experiments, corresponding to horizontal and vertical cursor velocity) to specify cursor velocity. The control space is the geometrical representation of a BCImapping. At the start of each day, we calibrated an 'intuitive mapping' by specifying a control space that the monkey used to move the cursor proficiently.

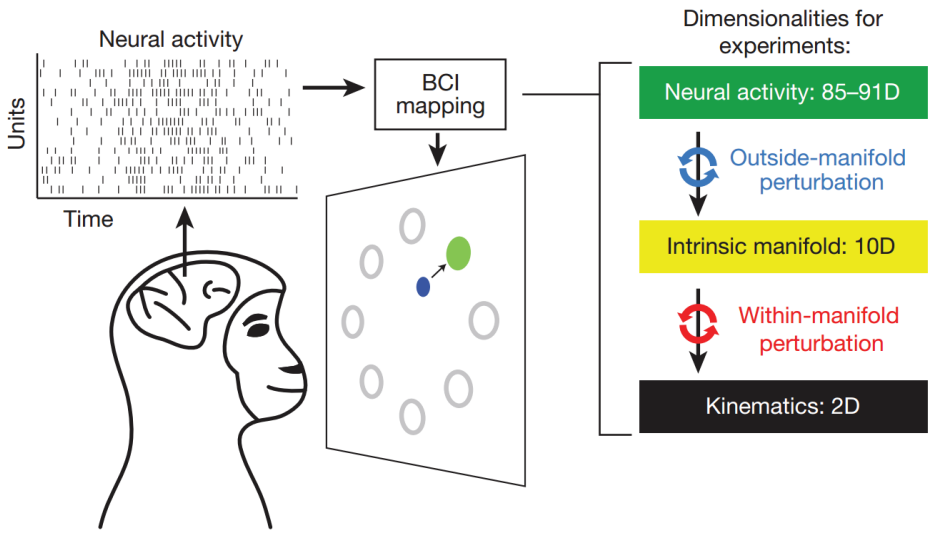

Figure 6.1: Using a brain-computer interface to study learning. Monkeys moved the BCI cursor (blue circle) to acquire targets (green circle) by modulating their neural activity. The BCI mapping consisted of first mapping the population neural activity to the intrinsic manifold using factor analysis, then from the intrinsic manifold to cursor kinematics using a Kalman filter. This two-step procedure allowed us to perform outside-manifold perturbations (blue arrows) and within-manifold perturbations (red arrows). D, dimensions.

At the beginning of each day we also characterized how each neural unit changed its activity relative to the other neural units (that is, how the neural units co-modulated). In the simplified network represented in Figure 6.2A, neurons 1 and 3 positively co-modulate due to common input, whereas neurons 1 and 2 negatively co-modulate due to an indirect inhibitory connection. Such co-modulations among neurons mean that neural activity does not uniformly populate the neural space (Cunningham and Yu 2014, Mazor and Laurent 2005, Mante et al. 2013, Rigotti et al. 2013, Churchland et al. 2012, Luczak et al. 2009). We identified the low-dimensional space that captured the natural patterns of co-modulation among the recorded neurons. We refer to this space as the intrinsic manifold (yellow plane in Figure 6.2A-B). By construction, the intuitive mapping lies within the intrinsic manifold. Our key experimental manipulation was to change the BCI mapping so that the control space was either within or outside of the intrinsic manifold. A within-manifold perturbation was created by re-orienting the intuitive control space but keeping it within the intrinsic manifold (depicted as the red line in Figure 6.2B). This preserved the relationship between neural units and co- modulation patterns, but it altered the way in which co-modulation patterns affected cursor kinematics (red arrows, Figure 6.1). An outside-manifold perturbation was created by re-orienting the intuitive control space and allowing it to depart from the intrinsic manifold (depicted as the blue line in Figure 6.2B). This altered the way in which neural units contributed to co-modulation patterns, but it preserved the way in which co-modulation patterns affected cursor kinematics (blue arrows, Figure 6.1).

6.2.1 Electrophysiology and behavioral monitoring

We recorded from the proximal arm region of the primary motor cortex in two male Rhesus macaques (Macaca mulatta, aged 7 and 8 years) using 96-channel microelectrode arrays (Blackrock Microsystems) as the monkeys sat head-fixed in a primate chair. All animal handling procedures were approved by the University of Pittsburgh Institutional Animal Care and Use Committee. At the beginning of each session, we estimated the root-mean-square voltage of the signal on each electrode while the monkeys sat calmly in a darkened room. We then set the spike threshold at 3.0 times the root-mean-square value for each channel. Spike

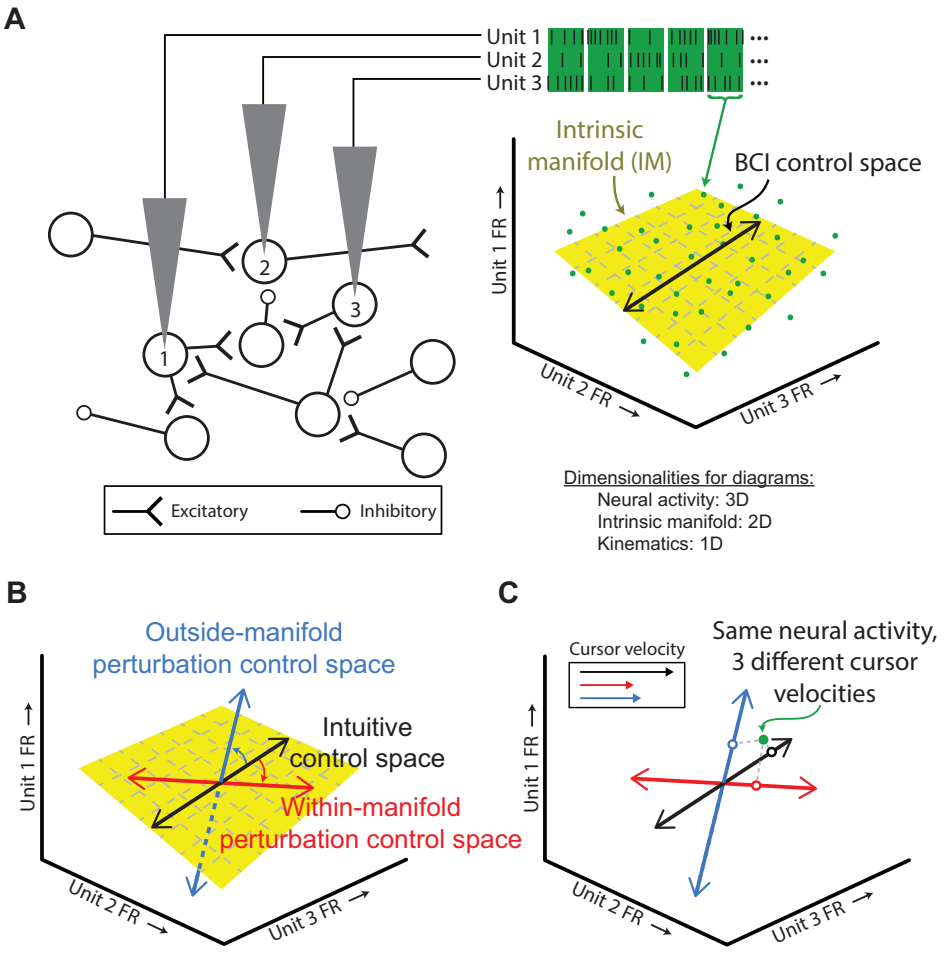

Figure 6.2: The intrinsic manifold and BCI cursor perturbations. (A) A simplified, conceptual illustration using three electrodes. The firing rate (FR) observed on each electrode in a brief epoch define a point (green dots) in the neural space. The intrinsic manifold (yellow plane) characterizes the prominent patterns of co-modulation. Neural activity maps onto the control space (black line) to specify cursor velocity. (B) Control spaces for an intuitive mapping (black arrow), within-manifold perturbation (red arrow) and outside-manifold perturbation (blue arrow). (C) Neural activity (green dot) elicits different cursor velocities (open circles and inset) under different mappings. Arrow colors as in B.

counts used for BCI control were determined from the times at which the voltage crossed this threshold. We refer to the threshold crossings recorded on one electrode as one neural unit. We used 85-91 neural units each day. We did not use an electrode if the threshold crossing waveforms did not resemble action potentials or if the electrode was electrically shorted to another electrode. The data were recorded approximately 19-24 months after array implantation for monkey J and approximately 8-9 months after array implantation for monkey L. We monitored hand movements using an LED marker (PhaseSpace Inc.)on the hand contralateral to the recording array. The monkeys' arms were loosely restrained. The monkeys could have moved their forearms by approximately 5cm from their armrests, and there were no restrictions on wrist movement. The hand movements during the BCI trials were minimal, and we observed that the monkeys' movements did not approach the limits of the restraints.

6.2.2 Task flow

Each day began with a calibration block during which we determined the parameters of the intuitive mapping. The monkeys then used the intuitive mapping for 400 trials (monkey J) or 250 trials (monkey L) during the baseline block. We then switched to the perturbed mapping for 600 trials (monkey J) or 400 trials (monkey L) for the perturbation block. This was followed by a 200-trial washout block with the intuitive mapping. Together, the perturbation and washout blocks comprised a perturbation session. The transitions between blocks were made seamlessly, without an additional delay between trials. We gave the monkey no indication which type of perturbation would be presented. On most days, we completed one perturbation session (monkey J, 50 of 58 days; monkey L, 29 of 30 days). On nine days, we completed multiple perturbation sessions.

6.2.3 Experimental sessions

We conducted 78 (30 within-manifold perturbations; 48 outside-manifold perturbations) sessions with monkey J. We conducted 31 sessions (16 within-manifold perturbations; 15 outside-manifold perturbations) with monkey L. For both monkeys, we did not analyze a

session if the monkey attempted fewer than 100 trials with the perturbed mapping. For monkey J, we did not analyze 11 sessions (2 within-manifold perturbations; 9 outside-manifold perturbations). For monkey L, we did not analyze 3 sessions (2 within-manifold perturbations; 1 outside-manifold perturbation).

6.2.4 BCI calibration procedures

Each day began with a calibration block of trials. The data that we recorded during these blocks were used to estimate the intrinsic manifold and to calibrate the parameters of the intuitive mappings. For monkey J, we used two calibration methods (only one on a given day), and for monkey L, we used one method for all days. The following describes the BCI calibration procedures for monkey J. The first method for this monkey relied on the neural signals being fairly stable across days. At the beginning of each day, the monkey was typically able to control the cursor proficiently using the previous days intuitive mapping. We collected data for calibration by having the monkey use the previous days intuitive mapping for 80 trials (10 per target).

We designed the second method because we were concerned about the potential for carry-over effects across days. This method relied on passive observation of cursor movement. The monkey observed the cursor automatically complete the center-out task for 80 trials (10 per target). At the beginning of each trial, the cursor appeared in the center of the monkeys workspace for 300 ms. Then, the cursor moved at a constant velocity (0.15m/s) to the pseudo-randomly chosen target for each trial. When the cursor reached the target, the monkey received a juice reward. After each trial, there was a blank screen for 200ms before the next trial. For both methods for monkey J, we used the neural activity recorded 300ms after the start of each trial until the cursor reached the peripheral target for BCI calibration.

The following describes the BCI calibration procedure for monkey L. We observed that neural activity for this monkey was not as stable from day to day as it was for monkey J. As a result, we could not use the calibration procedure relying on the previous days intuitive mapping. Additionally, the observation-based calibration procedure was not as effective at generating an intuitive decoder for monkey L as it had been for monkey J. Therefore, we

used a a closed-loop calibration procedure(similar to reference 32) to generate the intuitive decoder. The procedure began with 16 trials (2 to each target) of the observation task. We calibrated a decoder from these 16 trials in the same manner as the first method for monkey J.We then switched to the BCI center-out task, and the monkey controlled the velocity of the cursor using the decoder calibrated on the 16 observation trials. We restricted movement of the cursor so that it moved in a straight line towards the target (that is,any cursor movement perpendicular to the straight path to the target was scaled by a factor of 0). After 8 trials (1 to each target), we calibrated another decoder from those 8 trials. The monkey then controlled the cursor for 8 more trials with this newly calibrated decoder with perpendicular movements scaled by a factor of 0.125. We then calibrated a new decoder using all 16 closed-loop trials. We repeated this procedure over a total of 80 trials until the monkey was in full control of the cursor (perpendicular velocity scale factor). We calibrated the intuitive mapping using the 80 trials during which the monkey had full or partial control of the cursor. For each of those trials, we used the neural activity recorded 300ms after the start of the trial until the cursor reached the peripheral target.

6.2.5 BCI center-out task

The same closed-loop BCI control task was used during the baseline, perturbation and washout blocks. At the beginning of each trial, the cursor (circle, radius 18mm) appeared in the center of the workspace. One of eight possible peripheral targets (chosen pseudo-randomly) was presented (circle, radius 20mm; 150mm (monkey J) or 125mm (monkey L) from center of workspace, separated by 45°. A 300ms freeze period ensued when the cursor did not move. After the freeze period, the cursor's velocity was controlled by the monkey through the BCI mapping. The monkey had 7.5s to move the cursor into the peripheral target. If the cursor acquired the target within the time limit, the monkey received a juice reward. After 200ms, the next trial began. With the intuitive mappings, the monkeys' movement times were near 1,000ms, but the monkeys sometimes exceeded the 7.5s acquisition time limit with the perturbed mappings. If the cursor did not acquire the target within the time limit, there was a 1.5s time-out before the start of the next trial.

6.2.6 Estimation of the intrinsic manifold

We identified the intrinsic manifold from the population activity recorded during the calibration session using the dimensionality reduction technique factor analysis (Santhanam et al. 2009, Yu et al. 2009). The central idea is to describe the high-dimensional population activity u in terms of a low-dimensional set of factors z. Each factor is distributed according to the standard normal distribution N. This can be written in vector form as:

$$z \sim \mathrm{N}(0, I) \tag{6.1}$$

where I is the identity matrix. The neural activity is related to those factors by:

$$\boldsymbol{u}|z \sim \mathrm{N}(\Lambda z + \boldsymbol{\mu}, \Psi) \tag{6.2}$$

where $\boldsymbol{u} \in \mathbb{R}^{q \times 1}$ is a vector of z-scored spike counts (z-scoring was performed separately for each neural unit) taken in non-overlapping 45 ms bins across the q-neural units, and $z \in \mathbb{R}^{10 \times 1}$ contains the ten factors. That is, the neural activity $\boldsymbol{u}$ given a set of factors z is distributed according to a normal distribution with mean $\Lambda z + \boldsymbol{\mu}$ and diagonal covariance Ψ. The intrinsic manifold is defined as the column space of Λ. Each factor, or latent dimension, is represented by a column of Λ. We estimated Λ, $\boldsymbol{\mu}$ and Ψ using the expectation-maximization algorithm. The data collected during the calibration sessions had 1,470 $\pm$ 325 (monkeyJ, mean $\pm$ standard deviation) and 1,379 $\pm$ 157 (monkey L) samples.

6.2.7 Intuitive mappings

The intuitive mapping was a modified version of the standard Kalman filter (Wu et al. 2006). A key component of the experimental design was to use the Kalman filter to relate factors z to cursor kinematics rather than to relate neural activity directly to the cursor kinematics. This modification allowed us to perform the two different types of perturbation. We observed that performance with our modified Kalman filter is qualitatively similar to performance with a standard Kalman filter (data not shown).

The first step in the construction of the intuitive mapping was to estimate the factors using factor analysis (equations (1) and (2)). For each z-scored spike count vector $\boldsymbol{u}_t$, we computed the posterior mean of the factors $\hat{z}_t = \mathrm{E}[z_t|\boldsymbol{u}_t]$. We then z-scored each factor (that is, each element of $\hat{z}_t$ separately. The second step was to estimate the horizontal and vertical velocity of the cursor from the z-scored factors using a Kalman filter:

$$\boldsymbol{x}_t|\boldsymbol{x}_{t-1} \sim \mathrm{N}(A\boldsymbol{x}_{t-1} + \boldsymbol{b}, Q) \tag{6.3}$$

$$\hat{z}_t|\boldsymbol{x}_t \sim \mathrm{N}(C\boldsymbol{x}_t + \boldsymbol{d}, R) \tag{6.4}$$

where $\boldsymbol{x}_t \in \mathbb{R}^{2\times 1}$ is a vector of horizontal and vertical cursor velocity at time step t. We fitted the parameters A, $\boldsymbol{b}$, Q, C, $\boldsymbol{d}$ and R using maximum likelihood by relating the factors to an estimate of the monkeys' intended velocity during the calibration sessions. At each time point, this intended velocity vector either pointed straight from the current cursor position to the target with a speed equal to the current cursor speed (monkey J, first calibration task) or pointed straight from the center of the workspace to the target with a constant speed (0.15m/s, monkey L and monkey J, second calibration task).

Because spike counts were z-scored before factor analysis, $\boldsymbol{\mu} = 0$. Because factors were z-scored before decoding into cursor velocity, $\boldsymbol{d} = 0$. Because calibration kinematics were centered about the center of the workspace, $\boldsymbol{b} = 0$. The decoded velocity used to move the cursor at time step t was $\hat{\boldsymbol{x}}_t = \mathrm{E}[\boldsymbol{x}_t|\hat{z}_1, \ldots, \hat{z}_t]$. We can express $\hat{\boldsymbol{x}}_t$ in terms of the decoded velocity at the previous time step $\hat{\boldsymbol{x}}_{t-1}$ and the current z-scored spike count vector $\boldsymbol{u}_t$:

$$\hat{\boldsymbol{x}}_t = M_1\hat{\boldsymbol{x}}_{t-1} + M_2\boldsymbol{u}_t \tag{6.5}$$

$$M_1 = A - KCA \tag{6.6}$$

$$M_2 = K\Sigma_z\beta \tag{6.7}$$

$$\beta = \Lambda^T(\Lambda\Lambda^T + \Psi)^{-1} \tag{6.8}$$

As part of the procedure for z-scoring factors, Σ_z is a diagonal matrix where the (p, p) element is the inverse of the standard deviation of the pth factor. K is the steady-state Kalman gain matrix. We z-scored the spike counts and the factors in the intuitive mappings so that the perturbed mappings (which were based on the intuitive mappings) would not require a neural unit to fire outside of its observed spike count range.

6.2.8 Perturbed mappings

The perturbed mappings were modified versions of the intuitive mapping. Within-manifold perturbations altered the relationship between factors and cursor kinematics. The elements of the vector $\hat{z}_t$ were permuted before being passed into the Kalman filter. This preserves the relationship between neural units and the intrinsic manifold, but changes the relationship between dimensions of the intrinsic manifold and cursor velocity. Geometrically, this corresponds to re-orienting the control space within the intrinsic manifold. The following equations describe the within-manifold perturbations:

$$\hat{\boldsymbol{x}}_t = M_1\hat{\boldsymbol{x}}_{t-1} + M_{2,WM}\boldsymbol{u}_t \tag{6.9}$$

$$M_{2,WM} = K\eta_{WM}\Sigma_z\beta \tag{6.10}$$

where η_{WM} is a 10×10 permutation matrix defining the within-manifold perturbation matrix. Each element of a perturbation matrix is either a 1 or a 0. In each column an d in each row of a permutation matrix, one element is 1 and the other elements are 0. Outside-manifold perturbations altered the relationship between neural units and factors.

The elements of $\boldsymbol{u}_t$ were permuted before being passed into the factor analysis model (blue arrows, Figure 6.1). This preserves the relationship between factors and cursor velocity, but changes the relationship between neural units and factors. Geometrically, this corresponds to re-orienting the control space within the neural space and outside of the intrinsic manifold. The following equations describe outside-manifold perturbations:

$$\hat{\boldsymbol{x}}_t = M_1 \hat{\boldsymbol{x}}_{t-1} + M_{2,OM} \boldsymbol{u}_t \tag{6.11}$$

$$M_{2,OM} = K \Sigma_z \beta \eta_{OM} \tag{6.12}$$

where η_{OM} is a $q \times q$ permutation matrix defining the outside-manifold perturbation matrix.

6.2.9 Amount of learning

This section corresponds to Figure 6.3C. For each session, we computed the amount of learning during perturbation blocks as a single, scalar value that incorporated both changes in success rate (percent of trials for which the peripheral target was acquired successfully) and target acquisition time. We sought to use a metric that captured how much the monkeys' performance improved throughout the perturbation block relative to how much it was impaired at the beginning of the perturbation block. Having a single value for each session allowed us to more easily compare learning across sessions and to relate the amount of learning to a variety of properties of each perturbation . We also analyzed each performance criterion individually for each monkey without any normalization. We saw consistent differences in learnability. Thus, our results do not rely on the precise form of our learning metric, but the form we used provides a scalar value as a convenient summary metric.

As success rate and target acquisition time are expressed in different units, we first normalized each metric. We found the mean and standard deviation of the success rates and target acquisition times across all non-overlapping 50-trial bins in the baseline, perturbation and washout blocks for each monkey. We then z-scored the success rates and target acquisi-

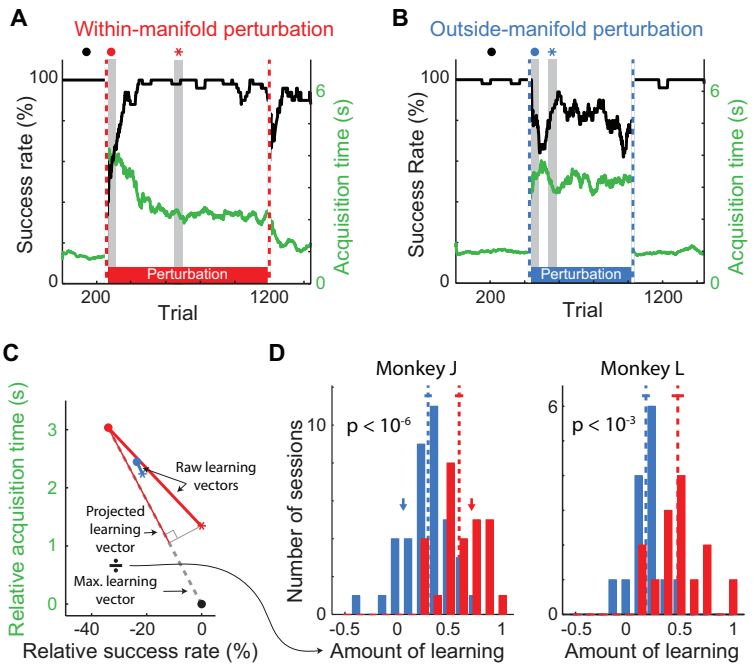

Figure 6.3: Better learning for within-manifold perturbations than outside-manifold perturbations. (a)(b) Task performance during one representative within-manifold perturbation session (A) and one representative outside-manifold perturbation session (b). Black trace, success rate; green trace, target acquisition time. Dashed vertical lines indicate when the BCI mapping changed. Grey vertical bands represent 50-trial bins used to determine initial (red and blue dots) and best (red and blue asterisks) performance with the perturbed mapping. (C) Quantifying the amount of learning. Black dot, performance with the intuitive mappings; red and blue dots, performance (success rate and acquisition time are relative to performance with intuitive mapping) just after the perturbation was introduced for sessions in Figure 6.3A and Figure 6.3B; red and blue asterisks, best per projected onto the maximum learning vector, normalized by the length of the maximum learning vector. This is the ratio of the length of the thin red line to the length of the dashed line. (D) Amount of learning for all sessions. A value of 1 indicates complete learning of the relationship between neural activity and kinematics, and 0 indicates no learning. Learning is significantly better for within-manifold perturbations than for outside-manifold perturbations. Arrows indicate the sessions shown in Figure 6.3A (red) and Figure 6.3B (blue). Dashed lines, means of distributions; solid lines, mean ± standard error of the mean (s.e.m.).

tion times separately for each monkey. Figure 6.3C shows normalized performance projected onto veridical units. For each session, we computed the average z-scored success rate and the average z-scored target acquisition time across all bins in the baseline block.

$$\boldsymbol{P}_B = \begin{bmatrix} s_B \\ a_B \end{bmatrix} \tag{6.13}$$

where $\boldsymbol{P}_B$ is the performance, s_B is the average normalized success rate and a_B is the average normalized acquisition time during the baseline block (monkey J, 386.0 $\pm$ 82.5 trials; monkey L, 292.1 pm 43.5 trials). We also compared the normalized success rates and acquisition times for all bins in the perturbation blocks.

$$\boldsymbol{P}_P(j) = \begin{bmatrix} s_P(j) \\ a_P(j) \end{bmatrix} \tag{6.14}$$

where $\boldsymbol{P}_P(j)$ is the performance, $s_B(j)$ is the normalized success rate, and $a_B(j)$ is the average normalized acquisition time during the jth 50-trial bin of the perturbation block. Empirically, we observed that the monkeys' performance during the perturbation blocks did not exceed the performance during the baseline blocks. Therefore, we define a maximum learning vector $\boldsymbol{L}_{max}$ as a vector that extends from the performance in the first bin with the perturbed mapping to the point corresponding to baseline performance (Figure 6.3c).

$$\boldsymbol{L}_{max} = \boldsymbol{P}_B - \boldsymbol{P}_P(1) \tag{6.15}$$

The length of this vector is the initial performance impairment because it describes the drop in performance that resulted when we switched from the baseline block to the perturbation block (shown in Figure 6.4A). For each bin j within the perturbation blocks, we defined a raw learning vector $\boldsymbol{L}_{raw}(j)$. This vector extended from the point corresponding to initial performance during the perturbation block to the point corresponding to performance during each bin.

$$\boldsymbol{L}_{raw}(j) = \boldsymbol{P}_P(j) - \boldsymbol{P}_P(1) \tag{6.16}$$

We projected the raw learning vectors onto the maximum learning vector. These were termed the projected learning vectors $\boldsymbol{L}_{proj}(j)$.

$$\boldsymbol{L}_{proj}(j) = \left(\boldsymbol{L}_{raw}(j) \cdot \frac{\boldsymbol{L}_{max}}{\|\boldsymbol{L}_{max}\|}\right)\left(\frac{\boldsymbol{L}_{max}}{\|\boldsymbol{L}_{max}\|}\right) \tag{6.17}$$

The lengths of the projected learning vectors relative to the lengths of the maximum learning vectors define the amount of learning in each 50-trial bin $\boldsymbol{L}_{bin}(j)$.

$$L_{bin}(j) = \frac{\|\boldsymbol{L}_{proj}(j)\|}{\|\boldsymbol{L}_{max}\|} \tag{6.18}$$

An amount of learning of 0 indicates that the monkey did not improve performance, and a value of 1 indicates that the monkey fully improved (up to the level during the baseline block). For each session, we computed the amount of learning for all bins, and we selected the largest one as the amount of learning for that session.

$$L_{session} = \max_j(L_{bin}(j)) \tag{6.19}$$

Figure 6.3C shows the raw learning vectors for one bin in each of two sessions (thick blue and red lines), along with the projected learning vector (thin red line) and the maximum learning vector (dashed gray line) for one of those sessions.

6.2.10 Statistics

For the histograms in Figures 6.3D and 6.4, the significances of the differences in distributions between within-manifold perturbation samples and outside-manifold perturbation samples were determined with two tailed Students t-tests assuming unequal variances of the two samples. We ensured that each histogram followed a normal distribution (Kolmogorov-Smirnov test).

6.3 RESULTS

To regain proficient control of the cursor under a within-manifold perturbation, the animals had to learn new associations between the natural co-modulation patterns and the cursor kinematics (Figure 6.2C). To restore proficient control of the cursor under an outside-manifold perturbation, the animals had to learn to generate new co-modulation patterns among the recorded neurons. Our hypothesis predicted that within-manifold perturbations would be more readily learnable than outside-manifold perturbations.

Just after the perturbed mappings were introduced, BCI performance was impaired (Figure 6.3A-B, first gray vertical band). Performance improved for the within-manifold perturbation (Figure 6.3A), showing that the animal learned to control the cursor under that mapping. In contrast, performance remained impaired for the outside-manifold perturbation (Figure 6.3B), showing that learning did not occur. We quantified the amount of learning as the extent to which BCI performance recovered from its initial impairment to the level attained while using the intuitive mapping (Figure 6.3C). For within-manifold perturbations, the animals regained proficient control of the cursor (red histograms in Figure 6.3D), indicating that they could learn new associations between natural co-modulation patterns and cursor kinematics. For outside-manifold perturbations, BCI performance remained impaired (blue histograms in Figure 6.3D), indicating that it was difficult to learn to generate new co-modulation patterns, even when those patterns would have led to improved performance in the task. These results support our hypothesis that the structure of a network determines which patterns of neural activity (and corresponding behaviors) a subject can readily learn to generate.

Just after the perturbed mappings were introduced, BCI performance was impaired (Figure 6.3A-B, first gray vertical band). Performance improved for the within-manifold perturbation (Figure 6.3A), showing that the animal learned to control the cursor under that mapping. In contrast, performance remained impaired for the outside-manifold perturbation (Figure 6.3B), showing that learning did not occur. We quantified the amount of learning as the extent to which BCI performance recovered from its initial impairment to the level attained while using the intuitive mapping (Figure 6.3C). For within-manifold perturbations,

the animals regained proficient control of the cursor (red histograms in Figure 6.3D, $n = 28$, monkey J; $n = 14$, monkey L), indicating that they could learn new associations between natural co-modulation patterns and cursor kinematics. For outside-manifold perturbations, BCI performance remained impaired (blue histograms in Figure 6.3D, $n = 39$, monkey J; $n = 15$, monkey L), indicating that it was difficult to learn to generate new co-modulation patterns, even when those patterns would have led to improved performance in the task. This finding is supported by the fact that there was a significant difference in the amount of learning between the two types of perturbations. The significance was tested using a two-tailed Students t-test. These results support our hypothesis that the structure of a network determines which patterns of neural activity (and corresponding behaviors) a subject can readily learn to generate.

Two additional lines of evidence show that BCI control was more learnable when using within-manifold perturbations than outside-manifold perturbations. First, perturbation types differed in their after-effects. After a lengthy exposure to the perturbed mapping, we again presented the intuitive mapping (the second dashed vertical line in Figure 6.3A-B). Following within-manifold perturbations, performance was impaired briefly, indicating that learning had occurred (Shadmehr et al. 2010). Following outside-manifold perturbations, performance was not impaired, which is consistent with little, if any, learning having occurred. Second, the difference in learnability between the two types of perturbation was present from the earliest sessions, and over the course of the study the monkeys did not improve at learning.

These results show that the intrinsic manifold was a reliable predictor of the learnability of a BCI mapping: new BCI mappings that were within the intrinsic manifold were more learnable than those outside of it. We considered five alternative explanations for the difference in learnability. First, we considered the possibility that mappings which were more difficult to use initially might be more difficult to learn. We ensured that the initial performance impairments were equivalent for the two perturbation types (Figure 6.4A).

Second, we posited that the animals must search through neural space for the new control space following the perturbation. If the control spaces for one type of perturbation tended to be farther from the intuitive control space, then they might be harder to find, and thus,

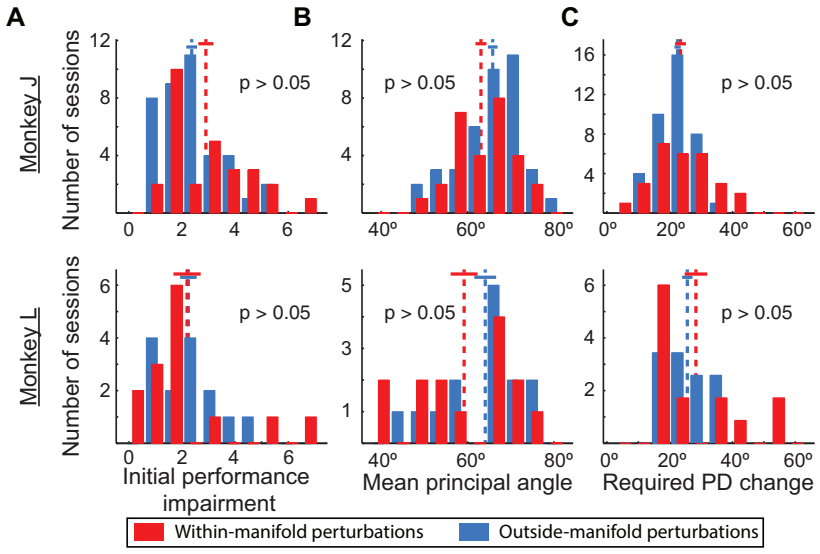

Figure 6.4: Alternate explanations do not explain the difference in learnability between the two types of perturbations. (A) Performance impairment immediately following within-manifold and outside-manifold perturbations. (B) Mean principal angles between intuitive and perturbed mappings. (C) Mean required change in preferred direction (PD) for individual neural units. For all panels: red, within-manifold perturbations; blue, outside-manifold perturbations; dashed lines, means of distributions; solid lines, mean ± s.e.m.; P values are for two-tailed Students t-tests; same number of sessions as in Figure 6.3D.

learning would be reduced. We ensured that the angles between the intuitive and perturbed control spaces did not differ between the two perturbation types (Figure 6.4B). Incidentally, Figure 6.4B also shows that the perturbations were not pure workspace rotations. If that were the case, the angles between control spaces would have been zero, not in the range of 40-80u as shown.

Third, we considered how much of an impact the perturbations exerted on the activity of each neural unit. Learning is manifested (at least in part) as changes in the preferred direction (that is, the direction of movement for which a neuron is most active) of individual neurons (Jarosiewicz et al. 2008, Li et al. 2001). If learning one type of perturbation required larger changes in preferred directions of neural units, then those perturbations might be harder to learn. We predicted the changes in preferred directions that would be required to learn each perturbation while minimizing changes in activity. We ensured that learning the two perturbation types required comparable preferred direction changes (Figure 6.4C).

Fourth, for one monkey (L), we ensured that the sizes of the search spaces for finding a strategy to proficiently control the cursor were the same for both perturbation types (see Methods). Fifth, hand movements were comparable and nearly non-existent for both perturbation types and should therefore have had no impact on learnability. We conclude from these analyses that the parsimonious explanation for BCI learning is whether or not the new control space is within the intrinsic manifold. These alternative explanations did reveal interesting secondary aspects of the data; they partially explained within-category differences in learnability, albeit in an idiosyncratic manner between the two monkeys.

A key step in these experiments was the identification of an intrinsic manifold using dimensionality reduction (Cunningham and Yu 2014). Although our estimate of the intrinsic manifold can depend on several methodological factors, the critical property of such a manifold is that it captures the prominent patterns of co-modulation among the recorded neurons, which presumably reflect underlying network constraints. For consistency, we estimated a linear, ten-dimensional intrinsic manifold each day. In retrospect, we considered whether our choice of ten dimensions had been appropriate (Figure 4). We estimated the intrinsic dimensionality of the neural activity for each day (Figure 6.5A); the average dimensionality was about ten (Figure 6.5B). Even though the estimated dimensionalities ranged

from 4-16, the selection of ten dimensions still provided a model that was nearly as good as the best model (Figure 6.5C). Because the top few dimensions captured the majority of the co-modulation among the neural units (Figure 6.5D), we probably could have selected a different dimensionality within the range of near-optimal dimensionalities and still attained similar results. We note that we cannot make claims about the 'true' dimensionality of the primary motor cortex, in part because it probably depends on considerations such as the behaviors the animal is performing and, perhaps, its level of skill.

6.4 DISCUSSION

Sensorimotor learning probably encompasses a variety of neural mechanisms, operating at diverse timescales and levels of organization. We posit that learning a within-manifold perturbation harnesses the fast timescale learning mechanisms that underlie adaptation (Salinas 2004), whereas learning an outside-manifold perturbation engages the neural mechanisms required for skill learning (Picard et al. 2013,Rioult-Pedotti et al. 2000). This suggests that learning outside-manifold perturbations could benefit from multi-day use (Ganguly and Carmena 2009, Peters et al. 2014). Such learning might require the intrinsic manifold to expand or change orientation.

Other studies have employed dimensionality-reduction techniques to interpret how networks of neurons encode information (Cunningham and Yu 2014, Mazor and Laurent 2005, Mante et al. 2013, Rigotti et al. 2013, Churchland et al. 2012, Luczak et al. 2009) and change their activity during learning (Paz et al. 2005, Durstewitz et al. 2010). Our findings strengthen those discoveries by showing that low-dimensional projections of neural data are not only visualization toolsthey can reveal causal constraints on the activity attainable by networks of neurons. Our study also indicates that the low-dimensional patterns present among a population of neurons may better reflect the elemental units of volitional control than do individual neurons.

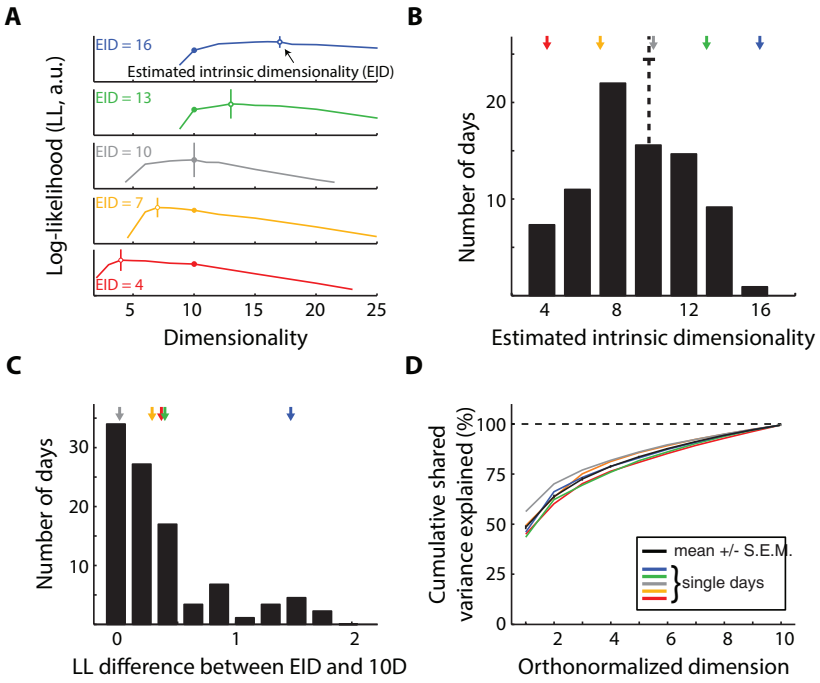

Figure 6.5: Properties of the intrinsic manifold. (A) Cross-validated log likelihoods (LL; arbitrary units) of the population activity for different days. The peaks (open circles) indicate the estimated intrinsic dimensionality (EID). Vertical bars indicate the standard error of log-likelihoods, computed across four cross-validation folds. A ten-dimensional intrinsic manifold was used for all experiments (solid circles). (B) EID across all days and both monkeys (mean ± s.e.m., 9.81 ± 0.31). (C) Difference between the LL for the ten dimensional (10D) model and the EID model. Units are the number of standard errors of LL for the EID model. For 89% (78 of 88) of the days, the LL for the ten-dimensional model was within one standard error of the EID model. All sessions were less than two standard errors away. d, Cumulative shared variance explained by the ten-dimensional intrinsic manifold used during the experiment. Colored curves correspond to the experimental days shown in Figure 6.5A. The black curve shows the mean ± s.e.m. across all days (n = 88; monkey J, 58; monkey L, 30).

In summary, a BCI paradigm enabled us to reveal neural constraints on learning. The principles we observed may govern other forms of learning (Braun et al. 2010, Jeanne et al. 2013, Gu et al. 2011, Ingvalson et al. 2012, Park et al. 2014) and perhaps even cognitive processes. For example, combinatorial creativity (Boden 1998), which involves re-combining cognitive elements in new ways, might involve the generation of new neural activity patterns that are within the intrinsic manifold of relevant brain areas. Transformational creativity, which involves creating new cognitive elements, may result from generating neural activity patterns outside of the relevant intrinsic manifold. More broadly, our results help to provide a neural explanation for the balance we possess between adaptability and persistence in our actions and thoughts (Ajemian et al. 2013).

7.0 SECONDARY PROJECTS AND CONTRIBUTIONS

In addition to the larger projects described in Chapters 3 - 6, I was also involved with several smaller projects and science outreach to middle school and high school students. The additional projects in eye-hand coordination, electrocorticography, monkey training, as well as my outreach, is described in this chapter.

7.1 ADDITIONAL PROJECTS

During my first years during my graduate training, I experienced different levels of the motor system. My endeavors have spanned single unit recordings in the brain stem, electrocorticography (ECoG) recordings in human patients with epilepsy, and additional monkey training during wherein I witnessed the powerful effects of neural plasticity.

7.1.1 Eye-hand coordination

In my first year, I attempted to perform an experiment to look into the interaction between fast eye movements, known as saccades, and reaching movements. This was performed with ample help from Dr. Neeraj Gandhi. Typically, if some one wants to reach to a target location, they first acquire that target location with eye movements. The person then guides their reach by looking at their intended target. We wished to disrupt this normal function. Our goal was to slip a single electrode into a monkey's brain stem and briefly microstimulate the omnipause neurons (OPN). These neurons fire at a high rate to 'lock' the position of the eyes in place. When the neurons pause their firing, the pause allows an eye

movement to happen. Thus, we intended to microstimulate the OPN when the eyes were mid-saccade. This would help us determine temporal linkages between coordinated eye and hand movements.

The experiments had progressed far enough that we began driving single microelectrodes into the brain stem. Two guide tubes were used, the first to penetrate the dura mater, and the second to penetrate the tentorium. During our first few attempts, we were not able to locate the OPN, which would have been easily identified by their unique firing rate patterns. We were, however, we believe we located one of the cranial nerves. We did several rounds of microstimulation, and we were able to illicit eye movements.

Unfortunately, our experiment was never completed. During our microstimulation tests, we learned that the monkey had developed cancer. We attempted steroid treatments, to help spur his appetite. Several weeks later, his condition was not improving and the animal was euthanized. If in the future, the opportunity arises, it would be interesting to learn the results of this experience.

7.1.2 Electrocorticography

During my beginning years as a graduate student, I was interested in learning about all possible sources of neural signals which could enable brain-computer interface control. In my primary lab with Dr. Aaron Batista, we primarily used chronically implanted microelectrode arrays to record neural signals from the brain. This allowed us to obtain a high-frequency, largely independent channels of neural activity. By recording extracellular neural signals, we could record the richest signal content and create highly controllable brain-computer interfaces. However, there are several downsides to using extracellular neural recordings. First, the microelectrode arrays must be chronically implanted into the brain itself. This is the most invasive of all neural recording techniques, and can lead to brain swelling, bleeding, and infection. Second, the microelectrode arrays illicit a foreign-body response and become encapsulated in tissue after a period of six months to five years. The tissue encapsulated the

microelectrodes so that they can no longer record the extracellular signals from the neurons. Third, the microelectrode arrays slowly shift in the brain, over the course of a single days to larger shifts over a period of weeks. These neural instabilities mean that brain-computer interfaces must frequently be calibrated, on a daily or weekly basis.

After learning of all of these issues, I realized many years of additional research will be required before intracortical brain-computer interfaces could become a clinically-viable option. Around the same period of time, I started learning about electrocorticography (ECoG), or recording electrical potentials from the surface of the brain. ECoG seemed like a fairly reasonable recording technology for brain-computer interfaces. The electrodes are closer to the surface of the brain than electroencephalography (EEG) electrodes, and thus able to sample signals with higher spatial and temporal complexity than EEG. However, because ECoG electrodes are less invasive than intracortical arrays, and we believed they would not initiate the same foreign-body response. And since each electrode records from larger cortical region, the electrical recordings would be more stable over time, and require less constant calibrations.

With this intriguing recording technology, I partially transitioned to Dr. Wei Wang's research group. Dr. Wang's group worked with human patients suffering from intractable epilepsy. This meant that doctors were not able to well control the patients' epilepsy using medication. The final option to help bring relief to a patient is by removing the portion of the brain generating the seizures. Before removing a portion of the brain, doctors must verify that they will be removing the correct portion of tissue. They also wish to make sure this tissue does not contain a primary motor or language region, which would leave the patient severely disabled if that tissue were retracted. To accomplish all of these verifications, doctor acutely implant ECoG arrays over the area of cortex where they believe the seizures are being generated. Then, over a week, the patient remains in the hospital to be monitored for seizures. I also had the valuable opportunity to watch how they tested for the motor or language region. To test for a motor region, the doctors would stimulate each electrode with varying levels of current. If that electrode was centered over a motor area, the stimulation

would cause the corresponding part of the body to move. To test for a language region, the doctors had the person read from a book while they were stimulating the electrodes. If that electrode was centered over a language region, the person's speech would stop, be disrupted, or altered.

For the week that the doctors were waiting and performing these tests, Dr. Wang's group asked the patients if they would like to perform brain-computer interface research. If they consented, we hooked our neural recording equipment to the hospital's equipment so that we could simultaneously record from the ECoG electrodes. Once all of our computers were connected, the patients were asked to move portions of their body, such as their arm, tongue, or foot. We wished to see if any of these movements caused increased neural activity on one of the electrodes. If certain electrodes did become activated, we also tested subjects to see if the same channels could be activated through imagined, rather than overt, movements. Once certain movements, electrodes, and frequency bands had been identified, a simple two-dimensional decoder was built. In most cases, one of the overt movements, such as raising your arm, would be specified to make the BCI cursor travel upwards. Then, a different movement, such as tongue motion, would be specified to make the BCI cursor travel to the right. Then the absence of each of the movements would cause the BCI cursor to move in the opposite direction.

From my several observations with the epilepsy patients, these BCI cursors seemed difficult to control. If was difficult to make the BCI cursor move in the intended direction, especially diagonally, which required two movements which humans to not usually make in a coordinated fashion. Also, subjects had a difficult time simultaneously not moving other parts of their body. Similar to people playing video or computer games, they desire to move other parts of the body to help get the object moving in the 'right' direction. However, these extraneous movements only generated more neural noise and made it harder for the patient to control the decoder.

After spending my first year in graduate school working with highly constrained monkey experiments, it was frustrating to perform experiments which had many additional, unaccounted for variables. Additionally, I witnessed the strong somatotopy of the brain. Each cortical area was strongly activated by movement by a particular portion of the body. With

ECoG recordings, it was easy to see when a particular limb was in motion. However, it was difficult to discern the fine details of that movement. In order to properly discern the details, the other sources of variability would need to be accounted for, and removed.

Because I was still interested in ECoG, but wanted better accounting of the variables, I helped initiated a research partnership between Dr. Aaron Batista and Dr. Wei Wang's lab. Dr. Wang's group was interested in testing out a custom micro-ECoG array design in a monkey subject. Their plan was to implant the same type of array in a human subject later in the year. My role was to create the partnership between the two groups, and setup the everything needed for the implant surgery. Because the ECoG array was going to be covering such a large portion of monkey K's brain, we wanted to get a good idea of how the array would lay over cortex. To get a clearer picture, I made a three-printed version of monkey K's brain. This was not a straightforward process. We got the shape and size of the brain from monkey K's MRI images. However, the MRI image file types were not directly convertible to the standard 3D printing '*.stl' file type. First, I tried using the skull stripping function in MRIcro. This did yield a helpful view of the cortical surface, however, the file type was still not correct. I contacted a 3D printing company, and using a software called Mimics, he was able to separate the cortical surface from the remainder of the the MRI files. He then printed a 3D model of monkey K's brain. Unfortunately, there was one variable for which I had not accounted–the MRIs showed a mirror image of the brain. What was printed as the right side of the brain was actually the left side of monkey K's brain. We were still able to use the model for surgical planning purposes.

After the implant, we started receiving signals from the electrodes. The original purpose of the implant was to verify the integrity and usefulness of the ECoG grid design. However, after seeing the signals, the project's scope expanded. It was desired that we collect hand kinematics while recording ECoG. I tried undertook this assignment for several weeks. Afterwards, when the monkey was not progressing well in his training, I removed myself from the project so that I could focus on my primary thesis work. Other students in the Batista and Wang labs continued to work with the monkey and collect ECoG data over a period of many months. After monkey K was sacrificed, the ECoG array was removed and a histological study was performed on the tissue which surrounded the array (Degenhart et al. 2015).

This study was important because it was one of first to look into the long-term effects of ECoG implants. Many people assumed that there would be less, to no, tissue encapsulation. However, when the array was explanted 666 days later, the array was heavily encapsulated in the stereotypical foreign body giant cells. Even so, reach-related activity was still recorded 562 days after implantation. This finding bodes well for continued research into using ECoG as a method for obtaining stable-long term neural recordings for neurophysiological research (Degenhart et al. 2015).

I gained much knowledge through this experience. The most valuable experience was interacting with the clinicians and patients, gaining a more personal point of view of those suffering from neurological disorders. Additionally, the broad somatotopy of the brain was astounding. From my classes, I already knew that sensorimotor cortex interacted with specific parts of the body. However, witnessing the microstimulation process and partaking in the BCI decoder training process, it became clearer to me the strength and importance of this organization. With intracortical recordings, there is usually limited somatotopy within the level of the array. Additionally, at least for BCI decoding, the original somatotopy was not extremely important. With training, both the neurons in the hand region of PMd, or the trunk region of M1, could be trained to move a BCI decoder in the horizontal and vertical directions. At the level of ECoG recordings, however, the somatotopy was more of a limiting factor.

7.1.3 Monkey H training and neural plasticity

During my time working on the experiments related to Chapter 6, I was training monkey H to make center-out reaching movements. Additionally, I placed focus on teaching the monkey to reach through virtual tubes (Sadtler et al. 2015), a task which came before the experiments in Chapter 6. A tube was a virtual constraint on a reach trajectory. For example, in a standard center-out reach, the monkey would reach straight from the center to the peripheral target. With a tube constraint, the tube placed a boundary perpendicular to the reach direction. We could then constrain the amount of perpendicular error the monkey could have for a

successful reach. I also trained monkey H to make curved reached trajectories. During those reaches, the tube was curved so that the monkey first moved away from the target and then curved back to the target.

After the training process, I scheduled and planned the the surgery to implant a microelectrode array into monkey H's right hemisphere motor cortex. Several days after the surgery, we were unable to record an single unit activity. We found and treated a large bacterial infection caused by the array implant. After the infection subsided, we were still not able to record any single unit activity, we performed a follow-up surgery to investigate the cause. Upon reopening the skull, we found there to be a large hole in the brain where the array had been implanted. The array was suspended inside of the hole. We removed the array, cleaned up the the tissue surrounding the hole, and waited to see how monkey H would do. For several months, monkey H's left arm was paralyzed. Miraculously, over time, monkey H did regain almost full use of his left arm. It was one of the most remarkable events I have witnessed. For the amount of brain tissue that was lost, I fully believed that monkey H would never regain use of his arm. However, with time, neural plasticity enabled him to make an almost complete recovery. I believe this may be in part due to the monkey's drive to be able to use his arm. Unlike human patients, there was no one who told monkey H the extent of his injury. No one told him that it would be unlikely, near impossible, for him to regain movement in his arm. But through persistence and determination, monkey H's brain was able to regain control of his left arm, even with premotor, motor, and portions of sensory cortex removed. Neural plasticity is able to over come great challenges when given enough time and will-power.

7.2 MENTORING AND OUTREACH

A substantial portion of my graduate student career was spent mentoring undergraduate and junior graduate students. It was a real joy to me to be able to help bring younger students into the sciences, while additionally having others support my projects. Within my time in the Batista lab, I mentored: Frank Petraglia, Ivan Smalianchuk, Stephen Whaite, Nicholas

Card (undergraduate students), and Robert Morrison, Jeffery Chiou, Bridget Endler, and Jessica Mischel (graduate students). They provided me an immense learning opportunity, how to manage several eager, hard-working people so that they are able to produce high quality work, that is both meaningful to their future goals, and to my current needs. I will treasure my experiences with each of the people.

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. It helped spur a lot of interest in research and the sciences. At the end of the program, we also talked with the students about how we reached our current career path. This was particularly helpful for juniors or seniors in high school trying to decide which majors to pursue in college. Another important aspect of the program was to highlight the humane use of animals in a research setting. I talked about how my lab performs research on rhesus monkeys and implants a chip in their brain so that we can study their neural signals. This speech was difficult to give to students the first several times, because there were always concerned looks on students' or teachers' faces. However, I then highlight how our goal is to help people who are severely paralyzed from the below the neck. We can help give them people more independence by allowing them to control a computer cursor, or robotic arm with their mind. Most students, and teachers, did immediately understand this long-term research goal. And the finally, I put this goal back in line with animal research. I would gently inform the students that it would be extremely unethical to implant these types of devices in human patients without understanding their short-term and long-term implications. Each time I gave the presentation, I witnessed it's impact on students. Using animals for research is not a clear-cut problem. If we want to advance and improve our medical procedures, animal sacrifice is a necessary component. Even though it is necessary, we still need to be good stewards of these animals and provide them the best life possible considering their circumstances.

8.0 CONCLUSIONS AND FUTURE DIRECTIONS

In Chapter 3, human subjects performed the Critical Stability Task (CST). In this set of experiments, the finding which surprised me the most was the differential effectiveness of the vibrotactile feedback. Human subjects were better able to use the vibrotactile feedback in conjunction with the pinch force control than with hand position control. How is it that the same feedback signal could be use better for one control method over another? The answer seemed to come in the form of matching the dynamics of the feedback signal with the dynamics of the effector. The 'bursty' vibrotactile feedback was better used during pinch force control, which was itself a 'bursty' control method. The sharp on/off nature of the feedback signal matched the on/off nature of pressing on the force transducer.

I believe an overlooked area of research has been to design the dynamics of the feedback signal so that they are congruent with the dynamics of the controller. In vibrotactiles research with humans, many research studies have attempted to determine whether frequency or amplitude modulation is better at conveying information to the subject. These studies have come to different conclusions. This may be in large part due to the different tasks being performed, and the different control methods being tested. It is entirely feasible that frequency modulation might be a better match for one control method, and amplitude modulation might be a better match for another. A similar idea has been pursued by Loughlin et al. for balance prostheses. The central idea is that each person's balance control strategy is slightly different; people are different heights, weights, and have learned to balance differently. To account for these differences, the vibrotactile feedback could be customized to each person to amplify or minimize specific information.

Similar to customizing the feedback signal to each person, we also need to better understand how subjects are using particular feedback signals. We will need to develop more detailed models of the feedback component in the CST. The original Critical Tracking Task (CTT) developed by Jex et al. assumed that the subject perfectly received, or perfectly understood, the feedback signal. In order to better study alternate forms of feedback using the CST, we will need to develop a generic model for learning to use a new feedback signal. When first introduced to the vibrotactile feedback signals, both humans and monkeys poorly used the feedback. With extensive training, the monkeys were able to use the vibrotactile feedback nearly as well as the visual feedback. What changed to allow the monkeys to use the feedback? Did they improve at perceiving the fine details of the signal, or did they improve at making the correct motor response to the vibrotactile signal? From my experiences, I would argue that the majority of the learning occurred the the sensorimotor transformation level. The monkeys most likely also improved in their ability to differentiate minute vibrotactile feedback differences. However, the majority of their performance gains likely occurred by learning to make the right motor action. How would this hypothesis be tested? One avenue would be to build two components of a vibrotactile-motor transfer function. First would be a perceptual transformation. This would convey the accuracy with which the subject understands the system state being rendered. Second would be a perceptual-motor transfer function. Early in vibrotactile feedback training, the perceptual-motor mapping is more similar to a cognitive strategy, such as a fuzzy logic controller (Zedeh 1968). From my experiences looking at the human subject data, the subjects respond in a characteristic manner to a broad range of the vibrotactile signal. This simplistic strategy is reminiscent of a non-linear fuzzy logic controller taking in a general piece of information and responding with a set movement. It isn't until later in vibrotactile feedback training, seen in the monkey subjects, that the perceptual-motor mapping becomes more linear. To understand how one learns a new form of feedback, it will be important to understand how the sensory-motor transformation changes with learning. Additionally, once the sensory-motor transformation has been characterized, it is likely that it will provide insight as to how the new feedback signal could be improved in order to better convey information.

Before more widespread adoption of the CST, several more task characterizations should be performed. These characterizations include testing the affect of the boundary locations and additional time delays. In the original CTT, Jex et al. found that the task was insensitive to the location of the failure boundaries. Once the subject could no longer stabilize the instability, the system's position grew exponentially – the location of the failure boundary was not very important. With the modified version of the task however, the monkeys stabilize the system for only a short period of time, six seconds. Because 'stability' is now being assessed within a period of time, the failure boundaries play a more important role. It is possible for the monkey to lose control of the system, right before the trial is over. Even though the system's position would be moving towards the boundary at an exponential speed, the trial might still end successfully. It is for this reason that a more thorough assessment of failure boundaries should be performed.

A second test that should be performed is adding an additional delay between the system's position and the rendered feedback. For example, a delay of 50ms could be added from when the updated system position was calculated to when the position was rendered to the subject. This could be done for both versions of the CST, where the instability increases within a trial or across different trials. The objective would be to find how the added delay alters the subject's critical instability score. Ideally, there would be a systematic shift between a certain length delay and a certain decrease in critical instability score. This would provide additional validation that a subject's critical instability score corresponds to a specific sensorimotor delay.

Additional CST modifications that could be performed include incorporating intracortical microstimulation feedback, using a second-order rather than first-order system, making the CST into a two-dimensional rather than one-dimensional task, or adding white noise or a sum of sinusoids to the unstable system equation. By adding noise to the unstable system, it would be possible to determine the frequency response of the human controller.

In Chapter 5, I investigated the neural representation of large and small amplitude movements. I believe this issue will be of growing importance in the coming years. In order for brain-computer interfaces to becomes clinically relevant, they will need to smoothly transition between gross and fine motor control. The brain-computer interface community

is now looking at this problem through the lens of speed representation in motor cortex. Since the 1980's, it has been found that movement direction is well-represented in motor cortex (Georgopoulos et al. 1986). BCI decoders since that time have used the principle of a single preferred movement direction. Movement speed was found to be present in the motor cortex within a reach (Schwartz 1993, Schwartz 1994, Moran and Schwartz 1999). Since that time, decoders based off of cursor velocity have been developed and utilized (Wu et al. 2006). However, it is coming to light that the speed tuning is probably insufficient to yield exceptional BCI performance (Golub et al. 2014). I have found that there are significant tuning differences between large amplitude reaches and small amplitude reaches. While they do manifest as a small change in modulation depth, it is not sufficiently large to say that the neurons are well speed-tuned. Additionally, there is a baseline shift in the tuning curve. Baseline shifts cannot be accounted for within the same decoding model. When trying to use the same decoding model for large and small amplitude reaches, the decoder performs poorly for the small amplitude reaches.

All of these observations demonstrate that neural tuning does not conform to our decoding models. Our decoders work well for a very specific range of movements, but they do not generalize across tasks. In order to generalize across tasks, there will need to be a way to model contextual changes in baseline firing rates, preferred direction changes, and changes in modulation depths. Our current decoding models are too rigid to account for the fluid mapping between neural activity and movement.

The fluid nature of neural tuning was tested in Chapter 6. One of the motivations for performing the experiments in Chapter 6 was to test the idea that neural tuning could re-configure itself to the tuning constraints specified by a neural decoder. In 2009, Ganguly and Carmena found that it was possible to re-configure the weights of their decoder, and over time, the monkey could learn to alter his neural activity to control the re-configured decoder. We further tested this hypothesis by creating two different types perturbed decoders: one that maintained the original correlations present in the intrinsic manifold, and one that broke the correlations present in the intrinsic manifold. We found that within a day, monkeys had an easier time learning a new decoder if it maintained the original correlation structure. Our experiment tested the fluidity of neural representations by enforcing a new structure

between the neurons. However, I am curious as to whether there was a contextual switch between controlling the original intuitive decoder and the perturbed decoders. If there was a contextual switch, it is possible that the baseline firing rate of the neurons changed. However, because we normalized the firing rates based on activity while controlling the intuitive decoder, the normalization may not be accurate for the perturbed decoder. in the future, it would be interesting to investigate the non-normalized firing rates between the intuitive and perturbed decoding sessions.

In conclusion, in order to make improvements in brain-computer interfaces, we will need to identify weaknesses in the sensorimotor control loop. Weaknesses can be done through psychomotor experiments and in brain-computer interface experiments. In order to actually discover weaknesses, it is imperative to use tests that can fundamentally test the limits of control. One of these limits is the time between sensing feedback and responding with a motor action. The Critical Stability Task will be an important tool for assessing weaknesses present in the sensory or motor components of the loop. Finally, BCI decoders will need to adopt a more generalized framework to account for contextual task changes. The combination of context-dependent decoders, more difficult tasks, and additional forms of feedback will allow brain-computer interfaces to truly transform people's lives.

APPENDIX

ADDITIONAL METHODS FOR TESTING NEURAL CONSTRAINTS ON LEARNING

A.1 CHOOSING A PERTURBED MAPPING

We used data from the first 200 trials (monkey J) or 150 trials (monkey L) of closed-loop control during the baseline blocks to determine the perturbation matrix that we would use for the session. The procedure we used had three steps (detailed below). First, we defined a set of candidate perturbations. Second, we predicted the open-loop cursor velocities for each candidate perturbation. Third, we selected one candidate perturbation. We aimed to choose a perturbation such that the perturbed mapping would not be too difficult for the monkeys to use nor so easy that no learning was needed to achieve proficient performance.

For monkey J, we often alternated perturbation types across consecutive days. For monkey L, we determined which type of perturbation we would use each day before the first experiment. That order was set randomly by a computer. We did this in order to avoid a detectable pattern of perturbation types. The following describes the first step in choosing a perturbed mapping: defining the candidate perturbations. For within-manifold perturbations, η_{WM} is a 10×10 permutation matrix. The total number possible η_{WM} is 10 factorial (3,628,800). We considered all of these candidate within-manifold perturbations.

For outside-manifold perturbations, η_{OM} is a $q \times q$ permutation matrix, where q is the number of neural units. For a population of 90 neural units, there are 90 factorial ($> 10^{100}$) possible values of η_{OM}. Due to computational constraints, we were unable to consider every possible η_{OM} as a candidate perturbation. We used slightly different procedures to determine the candidate outside-manifold perturbations for the two monkeys.

The procedure we used for monkey J is as follows. We permuted the neural units independently. We chose to permute only the neural units with the largest modulation depths (mean number of units permuted, 39 ± 18). Permuting the units with larger modulation depths impacted the monkeys ability to proficiently control the cursor more than would permuting units with smaller modulation depths. For each session, we randomly chose 6 million η_{OM} that permuted only the specified units. This formed the set of candidate outside-manifold perturbations.

The procedure we used for monkey L is as follows. To motivate it, note that the two perturbation types altered the intuitive mapping control space with in a different number of dimensions of the neural space for monkey J. Within-manifold perturbations were confined to N dimensions of the neural space, but outside-manifold perturbations were confined to N dimensions of the neural space (where N is the number of permuted units, 39 on average). Thus, the dimensionality of the space through which the monkey would have to search to find the perturbed control space was different for the two types of perturbed mappings; it was larger for the outside-manifold perturbations than it was for the within-manifold perturbations. We recognized that this difference may have affected the monkeys ability to learn outside-manifold perturbations. For monkey L, we reduced the size of the search space for the outside-manifold perturbations, thereby equalizing the size of the search space for the two perturbation types. We did this by constraining η_{OM} so that the number of possible η_{OM} was equal to the number of candidate within-manifold perturbations. We then considered all η_{OM} to be candidate outside-manifold perturbations. To construct outside-manifold perturbations, we assigned each neural unit to one of eleven groups. The first ten groups had an equal number of neural units. The eleventh group had the remaining neural units. We specifically put the neural units with the lowest modulation depths in the eleventh group. The $10m$ (where m is the number of neural units per group) neural units

with the highest modulation depths were randomly assigned to the first ten groups. We created outside-manifold perturbations by permuting the first ten groups, keeping all the neural units within a group together. Thus, the number of possible η_{OM} is 10 factorial, all of which were considered as candidate outside-manifold perturbations.

We attempted to keep these groupings as constant as possible across days. On some days, one electrode would become unusable (relative to the previous day) as evident from the threshold crossing waveforms. When this occurred, we kept all of the groupings fixed that did not involve that electrode. If an electrode in one of the first ten groups became unusable, we would substitute it with a neural unit from the eleventh group.

The following describes the second step in choosing a perturbed mapping: estimating the open-loop velocities of each candidate perturbation. The open-loop velocity can be thought of as a coarse approximation to how the cursor would move if the monkey did not learn. The open-loop velocity measurement captures how the neural activity updates the velocity of the cursor from the previous time step, whereas the closed-loop decoder (equation (5)) also includes contributions from the decoded velocity at the previous timestep ($M_1\hat{\boldsymbol{x}}_{t-1}$) as well as from the neural activity at the current time step ($M_2\boldsymbol{u}_t$). To compute the open-loop velocity, we first computed the average z-scored spike counts of every neural unit in the first 200 (monkey J) or 150 (monkey L) trials of the baseline block. We binned the spike counts from 300ms to 1,300ms (monkey J) or 1,100ms (monkey L) after the beginning of each trial, and then averaged the spike counts for all trials to the same target. Together, these comprised 8 spike count vectors (one per target). For each of the spike count vectors, we computed the open-loop velocity for the candidate perturbations:

$$\boldsymbol{x}_{OL}^i = M_{2,P}\boldsymbol{u}_B^i \tag{.1}$$

where $\boldsymbol{u}_B^i$ is the mean z-scored spike count vector for the ith target. $M_{2,P}$ is $M_{2,WM}$ for within-manifold perturbations and $M_{2,OM}$ for outside-manifold perturbations.

The following describes the third step in choosing a perturbation: selecting a candidate perturbation. For each candidate perturbation, we compared the open-loop velocities under the perturbed mapping to the open-loop velocities under the intuitive mapping on a per-target basis. We needed the velocities to be dissimilar (to induce learning) but not so

different that the animal could not control the cursor. For each target, we measured the angles between the 2D open-loop velocity vectors. We also measured the magnitude of the open-loop velocity for the perturbed mapping. For each session, we defined a range of angles (average minimum of range across sessions: mean ± s.e.m, 19.7°± 7.0°; average maximum of range across sessions: 44.4°± 8.9°)and a range of velocity magnitudes (average minimum of range across sessions, 0.7mm/s ± 0.4mm/s; average maximum of range across sessions, 5.5mm/s ± 4.0mm/s). Note that when the monkey controlled the cursor in closed-loop (equation (5)), the cursor speeds were much greater than these ranges of open-loop velocities. This is because M_1 was nearly an identity matrix for our experiments. Thus, the term $M_1\hat{\boldsymbol{x}}_{t-1}$ is expected to be larger than the term $M_2\boldsymbol{u}_t$. We found all candidate perturbations for which the angles and magnitudes for all targets were within the designated ranges. From the candidate perturbations that remained after applying these criteria, we arbitrarily chose one to use as the perturbation for that session.

A.2 PRINCIPAL ANGLES BETWEEN INTUITIVE AND PERTURBED CONTROL SPACES

The control spaces for the intuitive and perturbed BCI mappings in our experiments were spanned by the rows of M_2 for the intuitive mapping, $M_{2,WM}$ for within-manifold perturbations and $M_{2,OM}$ for outside-manifold perturbations. Because we z-scored spike counts in advance, the control spaces for each day intersected at the origin of the neural space. The two principal angles between the intuitive and perturbed control spaces defined the maximum and minimum angles of separation between the control spaces (Figure 3b).

A.3 REQUIRED PREFERRED DIRECTION CHANGES

One way in which learning is manifested is by changes in how individual neurons are tuned to the parameters of the movement, in particular the preferred direction. For each session, we

sought to compute the required changes in preferred direction for each neural unit that would lead to proficient control of the cursor under the perturbed mapping. One possibility would be to examine the columns of M_2 and $M_{2,P}$. Each column can be thought of as representing the pushing direction and pushing magnitude of one unit (that is, the contribution of each neural unit to the velocity of the cursor). We could simply estimate the required change in preferred direction by measuring the change in pushing directions for each unit between the intuitive and perturbed mappings. However, this method is not suitable for the following reason. For outside-manifold perturbations for monkey J, we permuted only a subset of the neural units. As a result, the columns of $M_{2,OM}$ corresponding to the non-permuted units were the same as in M_2. By estimating the required changed in preferred direction as the difference in directional components of M_2 and $M_{2,OM}$, we would be implicitly assuming that the monkey is capable of identifying which units we perturbed and changing only their preferred directions, which appears to be difficult to achieve in the time frame of a few hours7. Therefore, we sought a more biologically plausible method of computing the required preferred direction changes.

Using a minimal set of assumptions, we computed the firing rates that each unit should show under one particular learning strategy. Then, we computed the preferred direction of each unit using those firing rates and compared them to the preferred directions during the baseline block. The following were the assumptions used to compute the firing rates:

1. We assumed the monkeys would intend to move the cursor to each target at the same velocity it exhibited under the intuitive mapping. Fitts' Law predicts that movement speed depends on movement amplitude and target size, and these were always the same in our experiments.
2. The firing rates for the perturbed mapping should be as close as possible to the recorded firing rates when the monkeys used the intuitive mapping. This keeps the predicted firing rates within a physiological range.

We used the following procedure to compute the required preferred direction changes. First, we found the average normalized spike count vector $\boldsymbol{u}_B^i$ across time points (300-1,000 ms after the start of the trial) and all trials to each target (i) during the baseline blocks. We

minimized the Euclidean distance between $\boldsymbol{u}_B^i$ and $\boldsymbol{u}_P^i$, the normalized spike count vector for the perturbed mapping (assumption 2), subject to $M_2\boldsymbol{u}_B^i = M_{2,P}\boldsymbol{u}_P^i$ (assumption 1). $M_2\boldsymbol{u}_B^i$ (the open-loop velocity for the intuitive mapping) is known from the baseline block. For a given perturbed mapping (with $M_{2,P}$), we sought to find $\boldsymbol{u}_P^i$ that would lead to the same open-loop velocity, which has a closed-form solution:

$$\boldsymbol{u}_P^i = \boldsymbol{u}_B^i + M_{2,P}^T \left(M_{2,P} M_{2,P}^T\right)^{-1} (M_2 - M_{2,P})\boldsymbol{u}_B^i \tag{.2}$$

For each neural unit k, we computed its preferred direction $\theta_B(k)$ with the intuitive mapping by fitting a standard cosine tuning model.

$$u_B^i(k) = m_k \cos(\theta_i - \theta_B(k)) + b_k \tag{.3}$$

where $u_B^i(k)$ is the kth element of $\boldsymbol{u}_B^i$, m_k is the depth of modulation, b_k is the model offset of unit k, and θ_i is the direction of the ith target. We also computed the preferred direction of each unit for the perturbed mapping $\theta_P(k)$ in the same way. Figure 3c shows histograms of the equation below, averaged across all units for each session:

$$|\theta_P(k) - \theta_B(k)| \tag{.4}$$

A.4 ESTIMATION OF INTRINSIC DIMENSIONALITY

This section accompanies Figure 6.5A-C. During all experiments, we identified a ten dimensional (factors) intrinsic manifold. Offline, we confirmed this was a reasonable choice by estimating the intrinsic dimensionality of the data recorded in each calibration block. For each day, we performed a standard model selection procedure to compare factor analysis models with dimensionalities ranging from 2 to 30. For each candidate dimensionality, we used four fold cross-validation. For each fold, we estimated the factor analysis model parameters using 75% of the calibration data. We then computed the likelihood of the remaining

25% of the calibration data with the factor analysis model. For each dimensionality, we averaged the likelihoods across all folds. Each days intrinsic dimensionality was defined as the dimensionality corresponding to the largest cross-validated data likelihood of the calibration data for that day.

A.5 MEASURING THE CUMULATIVE SHARED VARIANCE EXPLAINED

This section corresponds to Figure 4d. Factor analysis partitions the sample covariance of the population activity $\text{cov}(\boldsymbol{u})$ into a shared component $\Lambda\Lambda^T$ and an independent component Ψ. In offline analyses, we sought to characterize the amount of shared variance along orthogonal directions within the intrinsic manifold (akin to measuring the lengths of the major and minor axes of an ellipse). These shared variance values are given by the eigenvalues of $\Lambda\Lambda^T$, which can be ordered from largest to smallest. Each eigenvalue corresponds to an orthonormalized latent dimension, which refers to identifying orthonormal axes that span the intrinsic manifold. Each orthonormalized dimension is a linear combination of the original ten dimensions. The cumulative shared variance curve is thus informative of how oblong-shaped the shared variance is within the manifold, and it can be compared across days. By definition, the cumulative shared variance explained reaches 100% using all ten dimensions, and none of the independent variance Ψ is explained by those latent dimensions.

BIBLIOGRAPHY

Ajemian, R., D'Ausilio, A., Moorman, H., and Bizzi, E. (2013). A theory for how sensorimotor skills are learned and retained in noisy and nonstationary neural circuits. *Proceedings of the National Academy of Sciences of the United States of America*, 110(52):E5078–87.

Andrews, T. J. and Coppola, D. M. (1999). Idiosyncratic characteristics of saccadic eye movements when viewing different visual environments. *Vision Research*, 39(17):2947–2953.

Bacher, D. H. (2009). *Real-time somatosensory feedback for neural prosthesis control: system development and experimental validation*. Master of science, University of Pittsburgh.

Bark, K., Wheeler, J. W., Premakumar, S., and Cutkosky, M. R. (2008). Comparison of skin stretch and vibrotactile stimulation for feedback of proprioceptive information. *2008 Symposium on Haptic Interfaces for Virtual Environment and Teleoperator Systems*, pages 71–78.

Batista, A. P., Buneo, C. A., Snyder, L. H., and Andersen, R. A. (1999). Reach plans in eye-centered coordinates. *Science*, 285(5425):257–260.

Birznieks, I., Jenmalm, P., Goodwin, a. W., and Johansson, R. S. (2001). Encoding of direction of fingertip forces by human tactile afferents. *The Journal of Neuroscience*, 21(20):8222–8237.

Boden, M. a. (1998). Creativity and artificial intelligence. *Artificial Intelligence*, 103(1-2):347–356.

Braun, D. a., Mehring, C., and Wolpert, D. M. (2010). Structure learning in action. *Behavioural Brain Research*, 206(2):157–65.

Buneo, C. a., Jarvis, M. R., Batista, A. P., and Andersen, R. a. (2002). Direct visuomotor transformations for reaching. *Nature*, 416(6881):632–636.

Burns, M. and Moskowitz, H. (1980). Effects of diphenhydramine and alcohol on skills performance. *European Journal of Clinical Pharmacology*, 17(4):259–66.

Callaway, E. M. (1998). Local circuits in primary visual cortex of the macaque monkey. *Annual Review of Neuroscience*, 21:47–74.

Carmena, J. M., Lebedev, M. A., Crist, R. E., O'Doherty, J. E., Santucci, D. M., Dimitrov, D. F., Patil, P. G., Henriquez, C. S., and Nicolelis, M. A. L. (2003). Learning to control a brain-machine interface for reaching and grasping by primates. *PLoS Biology*, 1(2):193–208.

Chatterjee, A., Chaubey, P., Martin, J., and Thakor, N. (2008). Testing a Prosthetic Haptic Feedback Simulator With an Interactive Force Matching Task. *Journal of Prosthetics and Orthotics*, 20(2):27–34.

Cheng, L., Kazman, R., and Robinson, J. (1996). Vibrotactile feedback in delicate virtual reality operations. *ACM Multimedia*, pages 243–251.

Cherian, A., Krucoff, M. O., and Miller, L. E. (2011). Motor cortical prediction of EMG: evidence that a kinetic brain-machine interface may be robust across altered movement dynamics. *Journal of Neurophysiology*, 106(2):564–75.

Churchland, M. M., Cunningham, J. P., Kaufman, M. T., Foster, J. D., Nuyujukian, P., Ryu, S. I., and Shenoy, K. V. (2012). Neural population dynamics during reaching. *Nature*, 487:51–56.

Churchland, M. M., Cunningham, J. P., Kaufman, M. T., Ryu, S. I., and Shenoy, K. V. (2010). Cortical preparatory activity: representation of movement or first cog in a dynamical machine? *Neuron*, 68(3):387–400.

Churchland, M. M., Santhanam, G., and Shenoy, K. V. (2006). Preparatory activity in premotor and motor cortex reflects the speed of the upcoming reach. *Journal of Neurophysiology*, 96(6):3130–3146.

Cincotti, F., Kauhanen, L., Aloise, F., Palomäki, T., Caporusso, N., Jylänki, P., Mattia, D., Babiloni, F., Vanacker, G., Nuttin, M., Marciani, M. G., and Del R Millán, J. (2007). Vibrotactile feedback for brain-computer interface operation. *Computational Intelligence and Neuroscience*, 2007:48937.

Collinger, J. L., Wodlinger, B., and Schwartz, A. B. (2013). 7 Degree-of-Freedom Neuroprosthetic Control By an Individual With Tetraplegia. *Lancet*, 381(9866):557–564.

Crevecoeur, F. and Scott, S. H. (2013). Priors Engaged in Long-Latency Responses to Mechanical Perturbations Suggest a Rapid Update in State Estimation. *PLoS Computational Biology*, 9(8):e1003177.

Cunningham, J. P. and Yu, B. M. (2014). Dimensionality reduction for large-scale neural recordings. *Nature Neuroscience*, 17(11).

Dadarlat, M. C., Doherty, J. E. O., and Sabes, P. N. (2015). A learning-based approach to artificial sensory feedback leads to optimal integration. *Nature Neuroscience*, 18(1):138–144.

Damian, D. D., Hernandez Arieta, A., Martinez, H., and Pfeifer, R. (2012). Slip speed feedback for grip force control. *IEEE Transactions on Biomedical Engineering*, 59(8):2200–2210.

Degenhart, A. D., Mischel, J., Endler, B., Smalianchuk, I., Ashmore, R. C., Cui, T., Dum, R., Elizabeth, C., Wang, W., and Batista, A. P. (2015). Evaluation of a chronically-implanted electrocorticographic electrode grid in a nonhuman primate. *To be submitted*, pages 1–24.

Dhillon, G. S. and Horch, K. W. (2005). Direct neural sensory feedback and control of a prosthetic arm. *IEEE Transactions on Neural Systems and Rehabilitation Engineering*, 13(4):468–72.

Diederich, A. and Colonius, H. (2004). Bimodal and trimodal multisensory enhancement: effects of stimulus onset and intensity on reaction time. *Perception & Psychophysics*, 66(8):1388–1404.

Dum, R. P. and Strick, P. L. (2002). Motor areas in the frontal lobe of the primate. *Physiology and Behavior*, 77:677–682.

Durstewitz, D., Vittoz, N. M., Floresco, S. B., and Seamans, J. K. (2010). Abrupt Transitions between Prefrontal Neural Ensemble States Accompany Behavioral Transitions during Rule Learning. *Neuron*, 66(3):438–448.

Engelhard, B., Ozeri, N., Israel, Z., Bergman, H., and Vaadia, E. (2013). Inducing Gamma Oscillations and Precise Spike Synchrony by Operant Conditioning via Brain-Machine Interface. *Neuron*, 77(2):361–375.

Fetz, E. E. (1969). Operant Conditioning of Cortical Unit Activity. *Science*, 163(3870):955–958.

Fitzsimmons, N. A., Drake, W., Hanson, T. L., Lebedev, M. A., and Nicolelis, M. A. L. (2007). Primate reaching cued by multichannel spatiotemporal cortical microstimulation. *The Journal of Neuroscience*, 27(21):5593–602.

Fogassi, L., Raos, V., Franchi, G., Gallese, V., Luppino, G., and Matelli, M. (1999). Visual responses in the dorsal premotor area F2 of the macaque monkey. *Experimental Brain Research*, 128:194–199.

Fu, Q.-G., Suarez, J. I., and Ebner, T. (1993). Neuronal specification of direction and distance during reaching movements in the superior precentral premotor area and primary motor cortex of monkeys. *Journal of Neurophysiology*, 70(5):2097–2116.

Gail, A. and Andersen, R. a. (2006). Neural dynamics in monkey parietal reach region reflect context-specific sensorimotor transformations. *The Journal of Neuroscience*, 26(37):9376–84.

Gandhi, N. and Katnani, H. (2011). Motor Functions of the Superior Colliculus. *Annual Review of Neuroscience*, 34:205–231.

Ganguly, K. and Carmena, J. M. (2009). Emergence of a stable cortical map for neuroprosthetic control. *PLoS Biology*, 7(7):e1000153.

Georgopoulos, A. P., Schwartz, A. B., and Kettner, R. E. (1986). Neuronal population coding of movement direction. *Science*, 233(4771):1416.

Ghez, C., Gordon, J., and Ghilardi, M. F. (1995). Impairments of reaching movements in patients without proprioception. II. Effects of visual information on accuracy. *Journal of Neurophysiology*, 73(1):361–372.

Gilja, V., Nuyujukian, P., Chestek, C. A., Cunningham, J. P., Yu, B. M., Fan, J. M., Churchland, M. M., Kaufman, M. T., Kao, J. C., Ryu, S. I., and Shenoy, K. V. (2012). A high-performance neural prosthesis enabled by control algorithm design. *Nature Neuroscience*, 15(12):7–10.

Godlove, J. M., Whaite, E. O., and Batista, A. P. (2014). Comparing temporal aspects of visual, tactile, and microstimulation feedback for motor control. *Journal of Neural Engineering*, 11(4):046025.

Golub, M. D., Yu, B. M., Schwartz, A. B., and Chase, S. M. (2014). Motor cortical control of movement speed with implications for brain-machine interface control. *Journal of Neurophysiology*, 112(2):411–29.

Gordon, J., Ghilardi, M. F., and Ghez, C. (1995). Impairments of reaching movements in patients without proprioception. I. Spatial errors. *Journal of Neurophysiology*, 73(1):347–60.

Gu, Y., Liu, S., Fetsch, C. R., Yang, Y., Fok, S., Sunkara, A., DeAngelis, G. C., and Angelaki, D. E. (2011). Perceptual learning reduces interneuronal correlations in macaque visual cortex. *Neuron*, 71(4):750–761.

Hall, W. C. and Moschovakis, A. K., editors (2004). *The Superior Colliculus: New Approaches for Studying Sensorimotor Integration*. CRC Press.

Hepp-Reymond, M., Kirkpatrick-Tanner, M., Gabernet, L., Qi, H.-X., and Weber, B. (1999). Context-dependent force coding in motor and premotor cortical areas. *Experimental Brain Research*, 128(1-2):123–133.

Hochberg, L. R., Bacher, D., Jarosiewicz, B., Masse, N. Y., Simeral, J. D., Vogel, J., Haddadin, S., Liu, J., Cash, S. S., van der Smagt, P., and Donoghue, J. P. (2012). Reach and grasp by people with tetraplegia using a neurally controlled robotic arm. *Nature*, 485(7398):372–375.

Horch, K., Meek, S., Taylor, T. G., and Hutchinson, D. T. (2011). Object discrimination with an artificial hand using electrical stimulation of peripheral tactile and proprioceptive pathways with intrafascicular electrodes. *IEEE Transactions on Neural Systems and Rehabilitation Engineering*, 19(5):483–489.

Hubel, D. N. and Wiesel, T. N. (1962). Receptive fields, binocular interaction and functional architecture in the cat's visual cortex. *Journal of Physiology*, 160:106–154.

Hulliger, M., Nordh, E., Thelin, A. E., and Vallbo, A. B. (1979). The responses of afferent fibres from the glabrous skin of the hand during voluntary finger movements in man. *Journal of Physiology*, 291:233–249.

Hwang, E. J., Bailey, P. M., and Andersen, R. a. (2013). Volitional Control of Neural Activity Relies on the Natural Motor Repertoire. *Current Biology*, pages 1–9.

Ingvalson, E. M., Holt, L. L., and McClelland, J. L. (2012). Can native Japanese listeners learn to differentiate /r-l/ on the basis of F3 onset frequency? *Bilingualism: Language and Cognition*, 29(6):997–1003.

Jarosiewicz, B., Chase, S. M., Fraser, G. W., Velliste, M., Kass, R. E., and Schwartz, A. B. (2008). Functional network reorganization during learning in a brain-computer interface paradigm. *Proceedings of the National Academy of Sciences of the United States of America*, 105(49):19486–91.

Jeanne, J. M., Sharpee, T. O., and Gentner, T. Q. (2013). Associative learning enhances population coding by inverting interneuronal correlation patterns. *Neuron*, 78(2):352–363.

Jenmalm, P., Birznieks, I., Goodwin, A. W., and Johansson, R. S. (2003). Influence of object shape on responses of human tactile afferents under conditions characteristic of manipulation. *European Journal of Neuroscience*, 18(1):164–176.

Jex, H., McDonnell, J., and Phatak, A. (1966). A "Critical" Tracking Task for Manual Control Research. *IEEE Transactions on Human Factors in Electronics*, 7(4):138–145.

Jiang, L., Cutkosky, M., Ruutiainen, J., and Raisamo, R. (2009). Using Haptic Feedback to Improve Grasp Force Control in Multiple Sclerosis Patients. *IEEE Transactions on Robotics*, 25(3):593–601.

Johansson, R. S. (1978). Tactile sensibility in the human hand: receptive field characteristics of mechanoreceptive units in the glabrous skin area. *Journal of Physiology*, 281:101–125.

Johansson, R. S. and Cole, K. (1992). Sensory-motor coordination during grasping and manipulative actions. *Current Opinion in Neurobiology*, 2:815–23.

Johansson, R. S. and Flanagan, J. R. (2009). Coding and use of tactile signals from the fingertips in object manipulation tasks. *Nature Reviews Neuroscience*, 10(5):345–59.

Johansson, R. S. and Westling, G. (1987). Signals in tactile afferents from the fingers eliciting adaptive motor responses during precision grip. *Experimental Brain Research*, 66:141–154.

Johnson, M. T. V., Coltz, J. D., and Ebner, T. J. (1999). Encoding of target direction and speed during visual instruction and arm tracking in dorsal premotor and primary motor cortical neurons. *European Journal of Neuroscience*, 11(12):4433–4445.

Kaczmarek, K. A., Webster, J. G., Bach-y Rita, P., and Tompkins, W. J. (1991). Electrotactile and vibrotactile displays for sensory substitution systems. *IEEE Transactions on Biomedical Engineering*, 38(1):1–16.

Kadkade, P. P., Benda, B. J., Schmidt, P. B., and Wall, C. (2003). Vibrotactile Display Coding for a Balance Prosthesis. *IEEE Transactions on Neural Systems and Rehabilitation Engineering*, 11(4):392–9.

Kandel, E. R., Schwartz, J. H., and Jessell, T. M. (2000). The Bodily Senses. In *The Principles of Neural Science*, pages 430–471. McGraw-Hill, New York, fourth edi edition.

Kuffler, S. W. S. W. (1953). Discharge Patterns and Fnctional Organization of Mammalian Retina. *Journal of Neurophysiology*, 16(1):37–68.

Ledbetter, N. M., Ethier, C., Oby, E. R., Hiatt, S. D., Wilder, a. M., Ko, J. H., Agnew, S. P., Miller, L. E., and Clark, G. a. (2013). Intrafascicular stimulation of monkey arm nerves evokes coordinated grasp and sensory responses. *Journal of Neurophysiology*, 109:580–590.

Leeb, R., Gwak, K., Kim, D.-S., and Millan, J. D. R. (2013). Freeing the visual channel by exploiting vibrotactile BCI feedback. *35th Annual International Conference of the IEEE EMBS*, 2013:3093–6.

Li, C. S., Padoa-Schioppa, C., and 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, C.-C., Whitney, S. L., Loughlin, P., Furman, J. M., Redfern, M. S., Sienko, K. H., and Sparto, P. J. (2015). The effect of age on postural and cognitive task performance while using vibrotactile feedback. *Journal of Neurophysiology*, 113(7):2127–2136.

London, B. M., Jordan, L. R., Jackson, C. R., and Miller, L. E. (2008). Electrical stimulation of the proprioceptive cortex (area 3a) used to instruct a behaving monkey. *IEEE Transactions on Neural Systems and Rehabilitation Engineering*, 16(1):32–6.

London, B. M. and Miller, L. E. (2013). Responses of somatosensory area 2 neurons to actively and passively generated limb movements. *Journal of Neurophysiology*, 109:1505–1513.

Loughlin, P., Mahboobin, A., and Furman, J. (2011). Designing vibrotactile balance feedback for desired body sway reductions. *33rd Annual International Conference of the IEEE Engineering in Medicine and Biology Society*, 2011:1310–3.

Luczak, A., Barthó, P., and Harris, K. D. (2009). Spontaneous events outline the realm of possible sensory responses in neocortical populations. *Neuron*, 62(3):413–25.

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

Mazor, O. and Laurent, G. (2005). Transient dynamics versus fixed points in odor representations by locust antennal lobe projection neurons. *Neuron*, 48(4):661–673.

McAuley, J. H., Rothwell, J. C., and Marsden, C. D. (1999). Human anticipatory eye movements may reflect rhythmic central nervous activity. *Neuroscience*, 94(2):339–350.

Mehta, B. and Schaal, S. (2002). Forward models in visuomotor control. *Journal of Neurophysiology*, 88:942–953.

Mennie, N., Hayhoe, M., and Sullivan, B. (2007). Look-ahead fixations: Anticipatory eye movements in natural tasks. *Experimental Brain Research*, 179(3):427–442.

Miall, R. C. and Wolpert, D. M. (1996). Forward Models for Physiological Motor Control. *Neural Networks*, 9(8):1265–1279.

Miller, J. O. and Low, K. (2001). Motor processes in simple, go/no-go, and choice reaction time tasks: a psychophysiological analysis. *Journal of Experimental Psychology*, 27(2):266.

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

Moran, D. W. and Schwartz, A. B. (1999). Motor cortical representation of speed and direction during reaching. *Journal of Neurophysiology*, 82:2676–2692.

Mountcastle, V. B., Steinmetz, M. A., and Romo, R. (1990). Frequency Discrimination in the Sense of Flutter: Psychophysical Measurements Correlated with Postcentral Events in Behaving Monkeys. *The Journal of Neuroscience*, 10(9):3032–44.

Nowak, L. G., Munk, M. H., Girard, P., and Bullier, J. (1995). Visual latencies in areas V1 and V2 of the macaque monkey. *Visual Neuroscience*, 12:371–384.

O'Doherty, J. E., Lebedev, M. A., Ifft, P. J., Zhuang, K. Z., Shokur, S., Bleuler, H., and Nicolelis, M. A. L. (2011). Active tactile exploration using a brainmachinebrain interface. *Nature*, 479:228–31.

Paillard, J. and Brouchon, M. (1968). Active and passive movements in the calibration of position sense. In Freedman, S., editor, *The Neuropsychology of Spatially Oriented Behavior*, pages 37–55. Dorsey Press, Homewood, IL.

Park, D. C., Lodi-Smith, J., Drew, L., Haber, S., Hebrank, A., Bischof, G. N., and Aamodt, W. (2014). The impact of sustained engagement on cognitive function in older adults: the Synapse Project. *Psychological Science*, 25(1):103–12.

Paz, R., Natan, C., Boraud, T., Bergman, H., and Vaadia, E. (2005). Emerging patterns of neuronal responses in supplementary and primary motor areas during sensorimotor adaptation. *The Journal of Neuroscience*, 25(47):10941–10951.

Pesaran, B., Nelson, M. J., and Andersen, R. a. (2010). A relative position code for saccades in dorsal premotor cortex. *The Journal of Neuroscience*, 30(19):6527–37.

Peterka, R. J. (2002). Sensorimotor Integration in Human Postural Control. *Journal of Neurophysiology*, 88:1097–1118.

Peters, A. J., Chen, S. X., and Komiyama, T. (2014). Emergence of reproducible spatiotemporal activity during motor learning. *Nature*, 510(7504):263–7.

Petzoldt, T., Bellem, H., and Krems, J. F. (2013). The Critical Tracking Task: A Potentially Useful Method to Assess Driver Distraction? *Human Factors: The Journal of the Human Factors and Ergonomics Society*, 56(4):789–808.

Picard, N., Matsuzaka, Y., and Strick, P. L. (2013). Extended practice of a motor skill is associated with reduced metabolic activity in M1. *Nature Neuroscience*, 16(9):1340–7.

Potvin, A. R., Doerr, J. A., Estes, J. T., and Tourtellotte, W. W. (1977). Portable clinical tracking-task instrument. *Medical & Bological Engineering & Computing*, 15(4):391–7.

Ptito, M., Moesgaard, S. M., Gjedde, A., and Kupers, R. (2005). Cross-modal plasticity revealed by electrotactile stimulation of the tongue in the congenitally blind. *Brain*, 128(3):606–614.

Pylatiuk, C., Mounier, S., Kargov, A., Schulz, S., and Bretthauer, G. (2004). Progress in the development of a multifunctional hand prosthesis. In *Engineering in Medicine and Biology Society, 2004. IEMBS'04. 26th Annual International Conference of the IEEE*, pages 4260–4263.

Quick, K. M., Card, N. S., Whaite, S. M., Mischel, J., Loughlin, P., and Batista, A. P. (2014). Assessing Vibrotactile Feedback Strategies by Controlling a Cursor with Unstable Dynamics. In *Engineering in Medicine and Biology Society (EMBC), 2014 36th Annual International Conference of the IEEE*, pages 2589–2592.

Quick, K. M., Mischel, J., Loughlin, P., and Batista, A. P. (2015). A task to quantify the effectiveness of feedback for motor control. *To be submitted*.

Ramaekers, J. G., Kauert, G., van Ruitenbeek, P., Theunissen, E. L., Schneider, E., and Moeller, M. R. (2006a). High-potency marijuana impairs executive function and inhibitory motor control. *Neuropsychopharmacology*, 31(10):2296–303.

Ramaekers, J. G., Moeller, M. R., van Ruitenbeek, P., Theunissen, E. L., Schneider, E., and Kauert, G. (2006b). Cognition and motor control as a function of Delta9-THC concentration in serum and oral fluid: limits of impairment. *Drug and Alcohol Dependence*, 85(2):114–22.

Rayner, K. and McConkie, G. W. (1976). What guides a reader's eye movements? *Vision Research*, 16(8):829–837.

Rees, A. and Palmer, A. R. (1988). Rate-intensity functions and their modification by broadband noise for neurons in the guinea pig inferior colliculus. *The Journal of the Acoustical Society of America*, 83(4):1488–1498.

Rigotti, M., Barak, O., Warden, M. R., Wang, X.-J., Daw, N. D., Miller, E. K., and Fusi, S. (2013). The importance of mixed selectivity in complex cognitive tasks. *Nature*, 497(7451):585–90.

Rioult-Pedotti, M. S., Friedman, D., and Donoghue, J. P. (2000). Learning-induced LTP in neocortex. *Science*, 290(5491):533–536.

Rombokas, E., Stepp, C. E., Chang, C., Malhotra, M., and Matsuoka, Y. (2013). Vibrotactile sensory substitution for electromyographic control of object manipulation. *IEEE Transactions on Biomedical Engineering*, 60(8):2226–32.

Romo, R., Hernández, A., Zainos, A., and Salinas, E. (1998). Somatosensory discrimination based on cortical microstimulation. *Nature*, 392(March):387–390.

Rouse, A. G., Williams, J. J., Wheeler, J. J., and Moran, D. W. (2013). Cortical Adaptation to a Chronic Micro-Electrocorticographic Brain Computer Interface. *Journal of Neuroscience*, 33(4):1326–1330.

Sadtler, P. T., Quick, K. M., Golub, M. D., Chase, S. M., Ryu, S. I., Tyler-Kabara, E. C., Yu, B. M., and Batista, A. P. (2014). Neural constraints on learning. *Nature*, 512(7515):423–426.

Sadtler, P. T., Ryu, S. I., Tyler-Kabara, E. C., Yu, B. M., and Batista, A. (2015). Brain computer interface control along instructed paths. *Journal of Neural Engineering*, 12(016015):1–13.

Salinas, E. (2004). Fast remapping of sensory stimuli onto motor actions on the basis of contextual modulation. *The Journal of Neuroscience*, 24(5):1113–1118.

Santhanam, G., Yu, B. M., Gilja, V., Ryu, S. I., Afshar, A., Sahani, M., and Shenoy, K. V. (2009). Factor-Analysis Methods for Higher-Performance Neural Prostheses. *Journal of Neurophysiology*, 102:1315–1330.

Schroeder, C. E., Wilson, D. A., Radman, T., Scharfman, H., and Lakatos, P. (2010). Dynamics of Active Sensing and perceptual selection. *Current Opinion in Neurobiology*, 20(2):172–6.

Schwartz, A. B. (1993). Motor cortical activity during drawing movements: population representation during sinusoid tracing. *Journal of Neurophysiology*, 70(1):28–36.

Schwartz, A. B. (1994). Direct cortical representation of drawing. *Science*, 265(5171):540–542.

Scott, S. H. and Loeb, G. E. (1994). The computation of position sense from spindles in mono- and multiarticular muscles. *The Journal of Neuroscience*, 14(12):7529–7540.

Shadmehr, R., Smith, M. a., and Krakauer, J. W. (2010). Error correction, sensory prediction, and adaptation in motor control. *Annual Review of Neuroscience*, 33:89–108.

Shenoy, K., Churchland, M., Santhanam, G., Yu, B., and Ryu, S. (2003). Influence of movement speed on plan activity in monkey pre-motor cortex and implications for high-performance neural prosthetic system design. *Proceedings of the 25th Annual International Conference of the IEEE Engineering in Medicine and Biology Society*, 2:1897–1900.

Simons, D. J. (1978). Response properties of vibrissa units in rat SI somatosensory neocortex. *Journal of Neurophysiology*, 41(3):798–820.

Skoglund, S. (1973). Joint Receptors and Kinaesthesia. In *Somatosensory System*, pages 111–136. Springer.

Sober, S. J. and Sabes, P. N. (2003). Multisensory integration during motor planning. *The Journal of Neuroscience*, 23(18):6982–92.

Sparks, D. L. and Mays, L. E. (1990). Signal transformations required for the generation of saccadic eye movements. *Annual review of neuroscience*, 13:309–336.

Suminski, A. J., Tkach, D. C., Fagg, A. H., and Hatsopoulos, N. G. (2010). Incorporating Feedback from Multiple Sensory Modalities Enhances Brain-Machine Interface Control. *Journal of Neuroscience*, 30(50):16777–16787.

Surwillo, W. W. (1973). Choice reaction time and speed of information processing in old age. *Perceptual and motor skills*, 36(1):321–322.

Taylor, D. M., Tillery, S. I. H., and Schwartz, A. B. (2002). Direct cortical control of 3D neuroprosthetic devices. *Science*, 296(5574):1829–32.

Tyler, M., Danilov, Y., and Bach-Y-Rita, P. (2003). Closing an open-loop control system: vestibular substitution through the tongue. *Journal of Integrative Neuroscience*, 2(2):159–164.

Vallbo, A. B., Hagbarth, K., Torebjork, H., and Wallin, B. (1979). Somatosensory, proprioceptive, and sympathetic activity in human peripheral nerves. *Physiological Reviews*, 59(4):919–957.

Van Essen, D. C. and Maunsell, J. H. R. (1983). Hierarchical organization and functional steams in the visual cortex. *Trends Neurosciences*, 6(9):370–375.

Velliste, M., Perel, S., Spalding, M. C., Whitford, A. S., and Schwartz, A. B. (2008). Cortical control of a prosthetic arm for self-feeding. *Nature*, 453(7198):1098–101.

Venkatraman, S. and Carmena, J. M. (2011). Active sensing of target location encoded by cortical microstimulation. *IEEE Transactions on Neural Systems and Rehabilitation Engineering*, 19(3):317–24.

Weber, D. J., Friesen, R., and Miller, L. E. (2012). Interfacing the somatosensory system to restore touch and proprioception: essential considerations. *Journal of Motor Behavior*, 44(6):403–18.

Weber, D. J., London, B. M., Hokanson, J. a., Ayers, C. a., Gaunt, R. a., Torres, R. R., Zaaimi, B., and Miller, L. E. (2011). Limb-state information encoded by peripheral and central somatosensory neurons: implications for an afferent interface. *IEEE Transactions on Neural Systems and Rehabilitation Engineering*, 19(5):501–13.

Westling, G. and Johansson, R. S. (1987). Responses in glabrous skin mechanoreceptors during precision grip in humans. *Experimental Brain Research*, 66(1):128–140.

Wheatstone, C. (1838). Contributions to the physiology of vision - Part the first: On some remarkable and hitherto unobserved phenomena of binocular vision. *Philosophical Transactions of the Royal Society of London*, 128:371–394.

Wheatstone, C. (1852). The Bakerian Lecture: Contributions to the physiology of vision - Part the second: On some remarkable and hitherto unobserved phenomena of binocular vision. *Philosophical Transactions of the Royal Society of London*, 142:1–17.

Wise, S. P., Boussaoud, D., Johnson, P. B., and Caminiti, R. (1997). Premotor and parietal cortex: corticocortical connectivity and combinatorial computations. *Annual Review of Neuroscience*, 20:25–42.

Wu, W., Gao, Y., Bienenstock, E., Donoghue, J. P., and Black, M. J. (2006). Bayesian population decoding of motor cortical activity using a Kalman filter. *Neural Computation*, 18(1):80–118.

Yu, B. M., Cunningham, J. P., Santhanam, G., Ryu, S. I., Shenoy, K. V., and Sahani, M. (2009). Gaussian-process factor analysis for low-dimensional single-trial analysis of neural population activity. *Journal of Neurophysiology*, 102(1):614–35.

Zedeh, L. (1968). Fuzzy Algorithms. *Information and Control*, 12:94–102.

Zhai, S. and Milgram, P. (1993). Human performance evaluation of isometric and elastic rate controllers in a 6 dof tracking task. *Proceedings of SPIE Vol . 2057 Telemanipulator Technology, 1993, Boston*, 2057:1–12.