

**STABILITY OF RECORDING AND NEURAL TUNING DURING INTRACORTICAL
BRAIN-COMPUTER INTERFACE ARM CONTROL**

by

John E Downey

BS in Biomedical Engineering, Johns Hopkins University, 2010

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

2017

UNIVERSITY OF PITTSBURGH
SWANSON SCHOOL OF ENGINEERING

This dissertation was presented

by

John E Downey

It was defended on

March 28th, 2017

and approved by

Andrew Schwartz Ph.D., Professor, Department of Neurobiology, University of Pittsburgh

Michael Boninger MD, Professor, Department of Physical Medicine and Rehabilitation,
University of Pittsburgh

Steven Chase Ph.D., Assistant Professor, Department of Biomedical Engineering, Carnegie
Mellon University

Dissertation Director: Jennifer Collinger Ph.D., Assistant Professor, Department of Physical
Medicine and Rehabilitation, University of Pittsburgh

Copyright © by John E Downey

2017

**STABILITY OF RECORDING AND NEURAL TUNING DURING
INTRACORTICAL BRAIN-COMPUTER INTERFACE ARM CONTROL**

John E Downey, PhD

University of Pittsburgh, 2017

For intracortical brain-computer interface (BCI) controlled neuroprosthetic arms to become a valuable assistive technology for people with upper-limb paralysis they will need to be able to adjust to a number of changes in neural activity that have not previously been well characterized. These include recording instabilities, and changes in neural tuning when interacting with objects. Here I present characterizations of these problems in two human subjects, along with a few solutions.

I quantified the rate at which recorded units become unstable within and between days to inform the design of self-recalibrating decoders. These decoders will provide BCI users with consistent performance, even as units become unstable, by updating to incorporate new units before too many original units have become unstable.

Using the quantification of stability, I also examined whether unit characteristics could predict how long a unit would be stable. I found that units with high firing rates, large peak-to-peak voltages, and more accurate tuning were most likely to remain stable. Using this result, future work should be able to create decoders that preferentially rely on stable units in order to enable high-performance BCI control for longer.

I addressed difficulties that the first subject was having using the hand to interact with objects. I identified the source of the problem as an increased firing rate across much of the population when the hand approached objects. I then developed a method to remove the increase in firing rate before decoding so that the arm kinematics became predictable when the hand approached objects.

Finally, I studied the representation of desired grasp force in primary motor cortex to enable BCI users to grasp a variety of objects, from light, fragile objects to heavy, sturdy objects. I found that primary motor cortex represents grasp force in a predictable manner during grasping, but that the tuning to grasp force is not apparent while the user carries an object. These results will enable the creation of BCI decoders that can apply the appropriate amount of force when grasping, and avoid dropping objects while transporting them.

TABLE OF CONTENTS

TABLE OF CONTENTS	vi
List of tables.....	x
Preface.....	xiii
1.0 Introduction	1
1.1 Electrophysiology.....	2
1.2 Brain-computer interface for arm control	6
1.3 Summary of work	10
2.0 General Methods	14
2.1 Electrophysiology Methods.....	14
2.1.1 Electrode implantation.....	14
2.1.2 Signal processing.....	16
2.2 Decoder Calibration	18
2.2.1 Description of decoder.....	18
2.2.2 Decoder calibration data.....	20
2.2.2.1 Calibration task and environment.....	21
2.2.2.2 Two-step calibration data collection	23
3.0 Quantifying instability	24
3.1 Introduction.....	24

3.2	Methods	26
3.2.1	Data collection and spike sorting	26
3.2.2	Unit isolation	28
3.2.3	Iterative stability determination.....	28
3.2.4	Direct comparison stability determination.....	30
3.3	Results.....	31
3.3.1	Recording quality metrics.....	31
3.3.2	Interday recording stability.....	33
3.3.3	Intraday recording stability.....	37
3.3.4	Effect of unit isolation on stability	40
3.4	Discussion	41
4.0	Predicting instability	46
4.1	Introduction.....	46
4.2	Methods	47
4.2.1	Unit characteristics.....	47
4.2.2	Cox analysis.....	50
4.2.3	Prior stability	51
4.3	Results.....	51
4.3.1	Characteristics of stable units	51
4.3.2	Cox analysis.....	54
4.3.3	Prior stability	57
4.4	Discussion	57
5.0	Object interaction.....	61

5.1	Introduction.....	61
5.2	Methods	65
5.2.1	Experiment sessions.....	65
5.2.2	Object interaction task.....	65
5.2.3	Impact of object interaction on neural activity.....	66
5.2.4	Online scaling of neural firing rates	67
5.2.5	Object transport task	68
5.3	Results.....	68
5.3.1	Object interaction task.....	68
5.3.2	Implementing scaling	71
5.3.3	Object transport	72
5.4	Discussion	72
6.0	Grasp force encoding.....	76
6.1	Introduction.....	76
6.2	Methods	78
6.2.1	Section specific recording details	78
6.2.2	Imagined grasp vs. object observation	79
6.2.3	Reach and grasp task	81
6.2.4	Force object classification.....	82
6.3	Results.....	83
6.3.1	Imagined grasp vs. object observation	83
6.3.2	Reach and grasp task	84
6.4	Discussion	88

7.0	Conclusions	91
7.1	Limitations.....	97
7.2	Future Work.....	101
	Appendix A	103
	Bibliography	132

LIST OF TABLES

Table 1. Unit stability distributions and decay rates.....	37
Table 2. Change in relative risk, given a change in each neural characteristic.	55
Table 3. ARAT performance metrics.....	121
Table 4. Multiple object task performance metrics.	126

LIST OF FIGURES

Figure 1. Illustration of intracellular and extracellular recordings.	3
Figure 2. Microelectrode array placement.	15
Figure 3. Thresholding and spike snippets.	17
Figure 4. Degrees of freedom illustration.	19
Figure 5. Calibration tasks.	22
Figure 6. Sorting examples.	27
Figure 7. Signal quality metrics.	32
Figure 8. Estimates of interday recording stability.	35
Figure 9. Estimates of intraday recording stability.	39
Figure 10. Estimates of intraday recording stability for different categories of isolation.	41
Figure 11. Illustration of baseline noise and peak-to-peak voltage.	48
Figure 12. Distributions of characteristics of units based on stability.	53
Figure 13. Predicted stability.	56
Figure 14. Object interaction tasks.	66
Figure 15. Neural response to object interaction task.	70
Figure 16. Closed-loop control of grasp force.	78
Figure 17. Grasp force methods figure.	80

Figure 18. Grasp force classification accuracy during object observation and imagined grasping.	84
Figure 19. Classification accuracy during the grasp and carry task.	85
Figure 20. Units encoding grasp force and velocity.	87
Figure 21. Array location.....	109
Figure 22. Shared control system diagram and robot testing set up.	111
Figure 23. Target positions for the multiple object task.	118
Figure 24. ARAT performance and difficulty.	120
Figure 25. Analysis of trajectory properties with and without shared control.....	123

PREFACE

My 6 years of work on brain-computer interface projects has been the most fulfilling work I have ever done. Working on this technology was a goal of mine since I was 15, and the reality of achieving that goal was everything I had hoped it would be, and more. I am looking forward to continuing my work in this field after graduating.

These opportunities were only available to me because of the work and support of numerous other people. I would like to acknowledge my advisor, Dr. Jennifer Collinger, for inviting me to work in this lab and guiding me every step of the way. I would also like to thank Drs. Michael Boninger and Andrew Schwartz who made this work a reality at University of Pittsburgh and provided valuable resources and insight. Finally, Dr. Steven Chase who provided me with much of the mathematical groundwork that underlies these projects. I have also benefitted from the assistance of other people who were not on my committee including numerous professors, post-docs, and fellow graduate students.

None of this work would have been possible without the selfless contributions of our research subjects. These people volunteer to put up with boring repetitive tasks, travel to our lab for multiple test sessions a week, and undergo multiple brain surgeries. They do all of this in the knowledge that this technology may not directly improve their own lives, but in the hope that it will benefit others who suffer similar injuries in the future. The importance of their dedication and effort cannot be overstated.

Finally, I'd like to thank my family: my parents, Lucy and Bill, who always encouraged me to pursue my interests, and provided me with a variety of opportunities to do so; and my wife, Phoebe, who has been a never ending source of support for me through personal and professional challenges during my time working on these projects.

1.0 INTRODUCTION

I have been working to develop brain-computer interfaces (BCIs) to control robotic arms to assist people with upper limb paralysis. Most people with upper limb paralysis rely on caregivers for assistance with activities of daily living (ADLs). This care can be expensive, burdensome, and deprive the individual of privacy. BCI controlled robotic arms use recordings from the intact motor cortex to allow people with upper limb paralysis to control a robotic arm by thinking about moving their own arm. We seek to make it so this robotic arm will allow them to complete many activities of daily living, restoring user's independence.

The prevalence of tetraplegia in the United States is estimated to be 151,000 people (The National SCI Statistical Center, 2015). A survey of 347 people with tetraplegia found that nearly half of them ranked arm and hand function the highest priority function for restoration (Anderson, 2004). In a different survey, approximately 75% of respondents with tetraplegia rated restored arm and hand function as very important to their quality of life (Collinger et al., 2013). BCI controlled robotic arms are being developed to address a high priority for some of the most disabled patients, and as such will provide a large improvement in quality of life for people that may live with these disabilities for decades.

1.1 ELECTROPHYSIOLOGY

Brain-computer interface control of a robotic arm uses recordings of electrical activity that neurons generate when a person thinks about moving their arm. This electrical activity is measured using electrophysiology. Electrophysiological recordings use electrodes to measure the change in voltage caused by ions moving in to and out of a neuron when it fires an action potential. These recordings can be made inside the neuron (intracellular), near the neuron (extracellular), or outside of the brain (electrocorticography). For this project, and for all intracortical BCI, extracellular recordings are used.

Hodgkin and Huxley pioneered electrophysiology in neurons beginning in the 1930's (Hodgkin & Huxley, 1939, 1945, 1952; Hodgkin, Huxley, & Katz, 1952). Using very large neurons from squid, they were able to place an electrode inside the neuron and make the first intracellular electrical recordings. Their work is the basis of our understanding of the ion channels and pumps that generate action potentials. These methods are great for studying the activity of a single neuron at a time, and have been improved so that they can be used in awake, behaving animals (Chorev, Epsztein, Houweling, Lee, & Brecht, 2009). While intracellular recordings provide very clear information about the activity of a single neuron, they do not scale well to recording many neurons at a time, making them incompatible with BCI control of robotic arms.

Action potentials occur when many small post-synaptic impulses integrate to raise the intracellular potential above a threshold, which is typically 5-25 millivolts above the resting potential. These post-synaptic impulses originate in the dendrites where other neurons form synapses and signal their activity using neurotransmitters. The action potential is integrated in the cell body, or soma, and then transmitted down the axon to the pre-synaptic sites from which the

neuron signals subsequent neurons. Both the post-synaptic impulses and the pre-synaptic action potentials can be measured using intracellular recording techniques (Figure 1). Recordings from outside the cell, called extracellular recordings, can measure the potential change due to ions entering and exiting a neuron during transmission of an action potential. The voltage changes in extracellular recordings are much smaller in magnitude (100's of microvolts vs. 100 millivolts), and opposite in sign, of those from intracellular recordings (Figure 1). The smaller magnitude of the recording means that post-synaptic impulses cannot be reliably measured and that recordings are best made in proximity to the soma or axon of a neuron.




Figure 1. Illustration of intracellular and extracellular recordings.

Extracellular recordings do not require targeting particular cells, requiring less care when placing the electrode and decreasing the chance of killing the neuron being recorded compared to intracellular recording. Because an extracellular electrode does not target an individual neuron it can record action potentials from single neurons or multiple neurons as well as field potentials relating to many neurons in the vicinity of the recording tip. Single and multi-units are defined based on the shape of the waveform, which is consistent for each action potential from a single neuron but tends to vary between neurons. Additionally, the size of the recorded action potential varies as a function of the size of the neuron and its proximity to the electrode with larger neurons and closer neurons generating higher peak-to-peak voltage recordings.

Recordings can be considered to originate from single-units when a large action potential is clearly separated from the other activity on an electrode. Recordings are considered to be multi-unit activity when multiple, largely overlapping, action potentials are recorded on an electrode, each likely originating from a different neuron. A specific form of multi-unit activity is often called “hash”. Hash units barely reach the voltage deviation threshold to be counted as firing activity and consist of activity from many different neurons that do not generate large enough changes in voltage to create clearly defined action potential traces. Hash units do contain information about planned movements, despite their indistinct origins (Perel et al., 2015). Local field potentials are the recordings of synchronized activity of many neurons in the vicinity of the recording electrode. These potentials are typically studied in the frequency domain, with changes in the power in different frequency bands conveying information about the activity of surrounding cortical tissue (Buzsáki, Anastassiou, & Koch, 2012). These three different, but related, sources of information from extracellular recordings have all been studied in the context of BCI (Perel et al., 2015).

Early extracellular recordings were made using microwires, electrodes made of steel or tungsten that is insulated above the tip and rigid enough for the point to be pushed in to the cortex far enough to record from the various layers of the brain, typically 0.5-3 mm (Schmidt, Bak, & McIntosh, 1976). Extracellular recordings also enable a much larger number of recording sites, for instance with a tetrode with multiple recording sites on a single shank (Csicsvari, 2003; Najafi, Wise, & Mochizuki, 1985). Alternatively an array of electrodes can be inserted, either by combining many microwires (Nicoletis et al., 2003) or a silicon chip etched and coated to create many recording tips (Campbell, Jones, Huber, Horch, & Normann, 1991). These methods can record from hundreds of neurons in unison while a subject is awake and moving. These large-scale recordings extract enough information from the brain to enable BCI control. Current intracortical BCI projects in human subjects all use the Blackrock microarray (shown in Figure 2A in Section 2.1.1), a silicon chip array with 100 tips, to record neural activity (Aflalo et al., 2015; Collinger et al., 2012; Gilja et al., 2015; Hochberg et al., 2006).

In our experience, we can record from approximately 300 neurons using two Blackrock microarrays implanted in motor cortex, while there are approximately 1.3 billion neurons in motor cortex (Gredal, Pakkenberg, Karlsborg, & Pakkenberg, 2000). This means that all studies of motor cortex using these electrophysiological methods immensely undersample the population. Additionally, the microarrays have 1.5 mm long electrode shanks meaning that recordings likely come from layer IV and layer V. Recordings from this layer should largely come from pyramidal neurons but likely include a smaller number of interneurons (Sherwood & Hof, 2010). Axons from pyramidal cells in M1 project down to the spinal cord and often synapse directly with motor units in the extremities. In contrast, the axons of interneurons do not leave the local area of cortex and are responsible for the local processing of information. These different types of neurons could

be differentiated based on shape of their recorded action potentials, with interneurons having narrow spike waveforms and pyramidal neurons having much wider waveforms (Bean, 2007).

Recordings from Blackrock microarrays can be unstable on the time course of hours, with recorded neurons being lost and new neurons being recorded at random intervals (Perge et al., 2013). Much of my work here (Chapters 3.0 & 4.0) is aimed at better understanding this instability to better be able to address the problems it can cause for BCI control.

Current work in extracellular electrophysiology seeks to prolong the viability of the recordings, typically by decreasing the damage the electrodes do to the brain, or by using biochemistry to attract neurons towards the recording tips (Marin & Fernández, 2010). Other approaches to recording neural activity without invasive electrodes, such as calcium imaging, functional near-infrared spectroscopy, or ultrasound recordings from piezoelectric “ neural dust”, have been proposed (Bertrand et al., 2014; Naseer & Hong, 2015; O’Shea et al., 2017) but so far none have provided the temporal and spatial resolution required to enable high-degree of freedom BCI arm control.

1.2 BRAIN-COMPUTER INTERFACE FOR ARM CONTROL

Neural control of external devices dates back to 1965 when Eberhard Fetz recorded from a single neuron in a non-human primate and taught the animal modulate the firing rate of the neuron in order to receive food rewards (Fetz, 1969). Fetz’s work proved that cortical neurons could be voluntarily controlled, presenting the possibility of using brain recordings as a control signal for computers or machines.

Recording from a single neuron requires very invasive surgery, and provides very little information at a time. The next step towards today's BCI systems was enabled by work from Apostolos Georgopolous' laboratory in the 1980s. Georgopolous' laboratory began recording from many neurons in the primary motor cortex (M1). They found that when a non-human primate moved its arm to targets in space certain neurons would consistently fire faster for movements in a particular direction and slower for movements in the opposite direction (Georgopoulos, Schwartz, & Kettner, 1986). The direction of movement that correlated with the highest firing rates was termed a neurons "preferred direction". Neurons had broadly distributed preferred directions, meaning that all types of movement would increase the firing rate of different subsets of the recorded population. By compiling the firing rates from a large enough recorded population it was possible to reconstruct the trajectory of the animals' hand as it reached to a target (Georgopoulos, Kettner, & Schwartz, 1988; Kettner, Schwartz, & Georgopoulos, 1988; Schwartz, Kettner, & Georgopoulos, 1988). The ability to reconstruct the actual movement of an animal, based solely on the recorded firing rates in cortex, was the first strong evidence that it would be possible for people without volitional control of their own arms to move robotic arms using a BCI instead.

The work from Georgopoulos' lab showed that firing rates correlated with arm movements, but for many years researchers continued to investigate what details of the movement were best encoded in the firing rates. Various labs tested whether M1 preferentially encoded movement direction, speed, velocity, joint torques, or muscle forces in non-human primate (Kakei, Hoffman, & Strick, 1999; Moran & Schwartz, 1999; Wang, Chan, Heldman, & Moran, 2007, 2010). The question about what kinematic variable M1 encodes is complicated by the high degree of correlation between variables imposed by the musculoskeletal structure of the arm. However, when these correlations were controlled for as well as possible, these studies tended to show that

velocity was the most clearly encoded kinematic variable in M1. This does not mean other information is not also encoded, and this work was done with upper-arm and wrist movement, but not individual finger movements. Current work in studying the encoding of movement, particularly hand movement in the presence of objects is discussed in section 5.1. Most research that aims to provide BCI arm control to people with tetraplegia has decoded velocity from M1 since the early 2000's (Taylor, Helms Tillery, & Schwartz, 2002).

The work discussed so far attempted to reconstruct arm trajectories from non-human primates using cortical recordings, or allowed the animals to control a cursor using neural activity resembling what was generated when they moved their own hand. The next step was to move robotic arms with the neural recordings. The first work with a BCI controlled robotic arm by a non-human primate was published by the Nicolelis laboratory. In that study, the monkey was presented with a cursor that it learned to control, but the movements of the cursor were in turn mapped to a robotic arm, enabling the monkey to unknowingly control a robotic arm (Carmena et al., 2003). A later study in the Schwartz laboratory trained a non-human primate to use a BCI controlled robotic arm that could translate in three-dimensional space and grasp objects. The animal was able to use the BCI controlled arm to eat marshmallows by reaching to them in space, grasping them, and returning them to its mouth (Velliste, Perel, Spalding, Whitford, & Schwartz, 2008). While this was the basis of the work to use a BCI controlled robotic arm discussed here, many labs continue to perform non-human primate BCI experiments both to gain a better understanding of M1 function and to improve BCI control (Ganguly, Dimitrov, Wallis, & Carmena, 2011; Koyama et al., 2010; Sadtler et al., 2014; Santhanam, Ryu, Yu, Afshar, & Shenoy, 2006).

At approximately the same time as the Velliste et. al. study, the BrainGate clinical trial was working with the first human BCI users (Hochberg et al., 2006). The BrainGate study has primarily concentrated BCI cursor control for computer use and communication (Gilja et al., 2015; Hochberg et al., 2006; B Jarosiewicz et al., 2014). These cursors could provide the last means of communication for people with ALS, allowing them to continue to interact with their friends and family, as well as communicating with caretakers to improve their quality of treatment once their disease has progressed to a locked-in state.

The BrainGate group, and the Andersen laboratory, have demonstrated low dimensional arm control in human BCI users. These studies achieved a maximum of 3 degrees of freedom (DoF) plus a binary classifier, and each enabled a subject to bring themselves a drink by using a 1 or 2 DoF decoder to trigger certain preprogrammed robotic arm movements to complete the task (Aflalo et al., 2015; Hochberg et al., 2012; Vogel et al., 2015). One study of special note allowed a man with tetraplegia to move his own arm through a combination of functional electrical stimulation and a robotic arm support (Ajiboye et al., 2017), which followed a project where a non-human primate was able to artificially activate his muscles to grasp using a BCI (Ethier, Oby, Bauman, & Miller, 2012). These studies demonstrated the possibility of BCI controlled robotic arm but provided minimal function and showed no ability to generalize to multiple tasks.

My work has shown that BCI users with tetraplegia can control a robotic arm with up to 10 degrees of freedom (DoF): 3 degrees of hand translation, 3 degrees of wrist orientation, and 4 degrees of hand shaping (Wodlinger et al., 2014). This level of control allows for sophisticated reaching and grasping that will allow users to complete a number of activities of daily living (ADLs). Based on the results of the Action Research Arm Test (ARAT), developed to monitor improvement in arm control post-stroke, and performed by our BCI robotic arm users, control is

good enough to restore a clinically relevant level of arm use (Collinger et al., 2012; Lyle, 1981; Wodlinger et al., 2014).

1.3 SUMMARY OF WORK

We have achieved high degree of freedom BCI control exclusively in a laboratory setting. In order for BCI technology to provide its intended benefits we will need to make it robust to changes in the neural recordings. In a survey of patients with tetraplegia 86% said independent operation of a BCI system was a very important design characteristic and 57% said that the daily set-up time was very important (Collinger et al., 2013). In order for this technology to be adopted for in-home use and minimize caretaker burden it will need to work reliably without the intervention of a technician for most of the day.

Intracortical BCI systems use extracellular electrophysiology (Section 1.1) to record how often the potential at each electrode crosses a threshold (the unit's firing rate, Section 2.1) and then use a decoder to convert these firing rates in to kinematic commands for the robotic arm (Section 2.2.1). The decoder is based on the relationship between the firing rates and the kinematics (neural tuning) during a decoder calibration session (Section 2.2.2). We use the neural activity to predict desired movement velocities based on many studies that found that velocity was encoded in M1 better than other kinematics or kinetics (Georgopoulos et al., 1988; Kakei et al., 1999; Koyama et al., 2010; Wang et al., 2007). Decoders are calibrated while the robotic arm moves in a predictable manner under computer control. While the arm is moving, the firing rates are recorded and the subject attempts to perform the same movements. This method is possible because, although injury

has removed the ability of the subjects to move their own arms, M1 is still capable of generating activity similar to that of someone with intact arm control moving their own arm (Kokotilo, Eng, & Curt, 2009).

Each electrode tip on the intracortical microelectrode arrays used in human intracortical BCI can record from multiple units in the local vicinity so changes in this local population can lead to changes in tuning properties (Perge et al., 2013). This can manifest as a change in baseline firing rate measured on an electrode, as well as changing the tuning to kinematics seen on that electrode. If this happens on multiple electrodes, the decoder will no longer accurately translate the neural signals into the desired movements. To date there is no published record of the time course of these changes in the recorded neural population of human subjects.

Here I quantify the instability of recorded units in two subjects (Chapter 3.0). Consistent with previous observations (Dickey, Suminski, Amit, & Hatsopoulos, 2009; Fraser & Schwartz, 2011; Perge et al., 2013; Vaidya et al., 2014), I found that approximately half of recorded units became unstable in 24 hours, but between 5 and 15% of units will remain stable for at least a month. These results indicate that well-chosen units could sustain cursor control for up to a month, increasing the independence of future BCI users.

In order to enable researchers to choose stable units, I also used the results of the quantification of instability to determine whether any unit characteristics could be used to predict how long a unit would remain stable (Chapter 4.0). I found that firing rate, peak-to-peak voltage, and strength of tuning all correlated with units that remained stable for a longer time. This predictive ability will enable the selection of units used for decoding that prolongs the usability of decoders before recalibration is necessary.

In addition to unit instability resulting from changes in the local environment at each electrode tip, we have observed that neural firing patterns also change in a context-dependent fashion based on the task a user is completing. Context-dependent changes in neural firing rates are likely to be transient, necessitating an online solution with a short response time. I found that most of the recorded units increased their firing rate when the user attempted to interact with an object (Chapter 5.0). To counteract this issue, I added a step to adjust for the correlated increase in firing rates in real-time, prior to decoding. This adjustment enabled both subjects to grasp, transport, and manipulate a variety of objects much more easily. These types of adjustments for unexpected changes in neural tuning will be a necessary step towards widespread adoption of BCI systems by people with tetraplegia.

Another context-dependent change is the encoding of desired grasp force. The BCI user only needs to issue this type of command when the robotic hand is in contact with an object. Well-controlled grasp force will be important for allowing the manipulation of a wide variety of objects. Too much force would break fragile objects, while insufficient force would allow heavy objects to slip out of the grasp. Here I present my study of how desired grasp force is represented in primary motor cortex (M1) during BCI controlled arm and hand movements to grasp different objects (Section 6.0). I found a predictable representation of desired grasp force while the BCI user attempts to grasp, but not while he tried to carry or release the object.

My work presented here covers a variety of issues presented by attempting to use a BCI arm control system for prolonged periods to accomplish a variety of tasks. Improving the robustness of decoders to recording instabilities and previously unstudied tuning characteristics in M1 are necessary steps towards a viable BCI system for people with upper-limb paralysis. The

results here suggest some solutions, and guide future work towards a clinical BCI neuroprosthetic arm system.

This document starts with a chapter on general BCI methods used for all of the projects discussed (Chapter 2.0). That is followed by 4 chapters each presenting one project each (Chapters 3.0 -6.0). Each of these chapters includes more specific background information, as well as methods and discussion to place the results in context. The document ends with general conclusions based on all of the projects (Chapter 7.0).

2.0 GENERAL METHODS

Here I describe the general methods used for the work that follows, as well as in most BCI work in our research lab. I describe the subjects, and the methods through which we recorded their neural activity. Then I describe the process through which we calibrate decoders during each test session. These methods enabled the subjects to control the robotic arm, and provide the framework within which I performed all of the analysis that follows.

2.1 ELECTROPHYSIOLOGY METHODS

2.1.1 Electrode implantation

We implanted two subjects with intracortical microelectrode arrays (Blackrock Microsystems, Salt Lake City, UT) in left primary motor cortex (M1). We used standard 10x10 platinum-coated silicon microelectrodes, but not all channels were wired to the percutaneous connector. The array's footprint was 4x4 mm, with a shank length of 1.5 cm (Figure 2A). Subject 1 was a 54-year old female diagnosed with spinocerebellar degeneration without cerebellar involvement (Boninger,

Mitchell, Tyler-Kabara, Collinger, & Schwartz, 2013). She was implanted with two 96-channel arrays for 33 months. Subject 2 was a 28-year old male with a C5 ASIA B spinal cord injury. He had two 88 channel arrays implanted in motor cortex. Subject 2 also had two 4x2.4 mm, 32 channel stimulating arrays implanted in primary sensory cortex (Figure 2B), however these arrays were not used in my work documented here. At the time of writing, the arrays had been implanted for 22 months.




Figure 2. Microelectrode array placement. A) The 10x10 electrode array. B) The locations where we implanted the arrays in Subject 1 (red) and Subject 2 (blue).

2.1.2 Signal processing

We recorded neural data during BCI test sessions, typically performed 2-3 days a week, using the NeuroPort data acquisition system (Blackrock Microsystems, Salt Lake City, UT). We recorded voltages at a 30 kHz sampling rate with a 250-4500 Hz filter. Each time the voltage signal crossed a pre-defined spike threshold (Figure 3A), the time of the crossing and a 48-sample snippet of the signal, starting 10 samples before the threshold crossing, were saved (Figure 3B). The spike threshold was set to either -5.25 (Subject 1 prior to day 565 post-implant) or -4.50 (all other test sessions) times the baseline root-mean-square voltage on each channel, and then rarely manually adjusted if necessary.




Figure 3. Thresholding and spike snippets. A) The continuous voltage trace from a single channel. The red line is the voltage threshold. B) Superimposed snippets. Everytime the voltage trace in ‘A’ crosses the threshold, a snippet of the trace is saved from the time surrounding the crossing event. The green lines illustrate how each snippet is a magnified view of the continuous voltage trace.

For use as a control signal, the number of threshold crossings were recorded every 30 ms (at 33 Hz) for Subject 1 or every 20 ms (at 50 Hz) for Subject 2. We used these binned spike counts to estimate an instantaneous firing rate by applying an exponential filter to the most recent 450 ms or 440 ms of spike counts for Subjects 1 and 2 respectively.

2.2 DECODER CALIBRATION

To enable BCI control we started each test session by calibrating a decoder. A decoder is the mathematical model that transforms the firing rates on each electrode (section 2.1.2) in to velocity commands for the robotic arm. There are many different forms that decoders can take (Gilja et al., 2012; Koyama et al., 2010; Sachs, Ruiz-Torres, Perreault, & Miller, 2016; Shanechi, Orsborn, & Carmena, 2016), but the work in our lab has exclusively used an optimal linear estimator (OLE) decoder (Collinger et al., 2012; Wodlinger et al., 2014). To calibrate the decoder for each test session, we use a two-step process where the first step is used to estimate the relationship between unit firing rates and desired kinematics in an open-loop task, which then allows a more accurate estimate by giving the subject assisted control of the robotic arm for the second step (Velliste et al., 2008). Our lab has used this decoding method to control arm translation, wrist orientation, hand shaping, and cursor movement with 1 to 10 degrees of freedom.

2.2.1 Description of decoder

The OLE decoder is based on an encoding model where the square-root transformed firing rate is a linear function of the desired kinematics:

$$\mathbf{f} = \mathbf{b}_0 + \mathbf{b}_1\mathbf{v}_1 + \mathbf{b}_2\mathbf{v}_2 + \dots + \mathbf{b}_D\mathbf{v}_D \quad \text{Equation 2.1}$$

Here, f is the square-root transformed firing rate of a single channel. The v terms represent the velocity in each numbered dimension. We calculated the b terms using linear regression with data

collected for decoder calibration. The subscripts on each term indicate which degree of freedom (DoF) it applies to (eg. translation forward-backward, wrist rotation, or grasp aperture, Figure 4). The term D indicates the number of degrees of freedom that were decoded and varied from 1 to 10 depending on the experiment being conducted.




Figure 4. Degrees of freedom illustration.

A) Diagram of the axes about which the wrist can rotate for the 3 degrees of wrist orientation. In this diagram there is 1 degree of grasp, meaning all of the fingers move in unison to grasp. B) Diagram of the axes along which the hand can translate. The origin is located at the end-point of the arm. Both diagrams represent the MPL, which was the most commonly used robotic arm in this work.

Equation 2.1 can be rewritten in matrix form as:

$$\mathbf{F} = \mathbf{B}\mathbf{V} \quad \text{Equation 2.2}$$

Here, F is the array containing the firing rate of all recorded channels. The matrix B contains the b coefficients for each channel. The array V contains the desired velocity commands. Once the encoding model is in this form we can calculate a W weights matrix such that:

$$V = FW \tag{Equation 2.3}$$

This W matrix is the decoder used to transform recorded firing rates in to velocity commands for the robotic arm. The indirect calculation of the W matrix using the B matrix from Equation 2.2 uses the following two equations, as described by Wang et. al. (Wang et al., 2007):

$$WB' = I \tag{Equation 2.4}$$

$$W = B[(B'B)^{-1}]' \tag{Equation 2.5}$$

Here I is the identity matrix and the transition from Equation 2.4 to Equation 2.5 is accomplished using the Moore-Penrose Matrix inverse.

In order to fit the B matrix, and enable the calculation of the W matrix for decoding we first had to collect a set of firing rates and arm velocities that provided a reasonable representation of how the subject would use the arm during decoding. I outline the methods used to collect this data in the section below.

2.2.2 Decoder calibration data

In order to train the OLE decoder described above (Section 2.2.1) we needed to collect firing rates recorded during attempted arm movements. To do this we used a two-step calibration procedure like the one used in non-human primates (Velliste et al., 2008), and described in our previous work (Collinger et al., 2012; Downey et al., 2016; Wodlinger et al., 2014).

2.2.2.1 Calibration task and environment

We calibrated decoders using either the Modular Prosthetic Limb (MPL) in front of a board with LEDs for presenting translation targets, or a virtual MPL (vMPL) presented in a 3 dimensional scene on a television, where virtual objects could be presented as targets (Johannes et al., 2011). For a 7 DoF decoder like the ones used in previous papers from our lab (Collinger et al., 2012; Wodlinger et al., 2014) the task used during training would have 3 phases in each trial. The first phase of a trial was for translation, the subjects were given 2 seconds to see a translation target indicated by LEDs or a virtual object (Figure 5). Once the hand moved to this translation target an audio cue would play to tell the subject how to orient the wrist. If calibration was taking place in a virtual scene the target would rotate to match the orientation target. Once the wrist was properly oriented an audio cue would instruct the subject to grasp or open the hand. In this way movements could be made in all 7 degrees of freedom to achieve targets.




Figure 5. Calibration tasks. A) The MPL completing a calibration trial including a translation phase, an orientation phase, and a hand-shaping phase. B) The vMPL completing the same three phases, with objects cueing translation instead of LEDs.

For testing that used different numbers of degrees of freedom, target sets could be expanded or contracted, or entire phases could be removed. For instance, for a 5 DoF decoder where the only wrist orientation dimension was wrist rotation, the orientation target set would be restricted to 3 positions: palm up, palm down, neutral. For a 4 DoF decoder the wrist orientation phase of the trial would be completely removed because only the translation and grasp phases were needed. For a 10 DoF decoder, like the one shown in Figure 5, the grasp phase had more targets than grasp or open to account of the extra degrees of freedom in the hand. In this way the same basic approach to collecting calibration data could be used to train decoders of varying complexity.

2.2.2.2 Two-step calibration data collection

During the first step of calibration data collection, called the observation step, we presented a task like the ones described above (Section 2.2.2.1) to the subject. The computer then moved the arm to the target while the subject attempted to execute the movements themselves. This step matched the neural firing rates generated by the subjects attempting to move the arm to the ideal kinematics generated by the computer to create a decoder for use in the second step of calibration using the equations from Section 2.2.1.

In the second step of calibration, called the orthogonal assistance step, the subjects used the decoder generated from the data collected in the first step to move the arm to achieve the presented targets. The orthogonal assistance takes the decoded velocity, then removes the components of the velocity command that were not directed along the path to the target, before implementing the command (Velliste et al., 2008). In this way, the subjects were required to provide the velocities to move the arm to the target, which ensured engagement, but were unable to make erroneous movements, maximizing the predictability of the task. We used the decoder calibrated from the firing rates and assisted velocity commands from the orthogonal assistance step to complete experiments after calibration.

3.0 QUANTIFYING INSTABILITY

3.1 INTRODUCTION

To date, intracortical BCI studies with human subjects use microelectrode arrays (Blackrock Microsystems, Salt Lake City, UT) consisting of nearly 100 recording tips inserted 1.5 mm into motor cortex (Aflalo et al., 2015; Collinger et al., 2012; Gilja et al., 2015; Hochberg et al., 2006). It has been well documented that these intracortical recordings can be unstable (Perge et al., 2013), which leads to challenges with maintaining BCI performance. BCI output is based on an assumed relationship between neural firing rates and movement commands, so if the neural activity changes, BCI performance degrades. The causes of recording instability are not well understood. The most obvious change occurs when a neuron can no longer be recorded, or when a new unit is recorded. More subtle changes, such as a change in neural tuning or behavior, can also lead to reduced BCI performance. Previous studies of the stability of individual unit recordings have been limited to a few weeks (Dickey et al., 2009; Tolias et al., 2007), with the exception of one animal that was studied for 150 days (Fraser & Schwartz, 2011). To date, there is no published record of long-term recording stability in humans with tetraplegia, which would be the target population for many BCI devices.

Previous studies have shown that it is possible to build a decoder that uses only stable units, and that performance with these decoders remains high for days to weeks. However, these studies have used a small number of units to provide two degrees of freedom of BCI control (Ganguly & Carmena, 2009). Many BCI applications will require more than two degrees of freedom. Here I seek to quantify unit stability in human BCI users, to determine if decoding methods that use stable units are feasible, or whether too few units remain stable over time to make this a viable approach. As an alternative, by understanding the rate of instability and characteristics of stable units I can provide a framework for adaptive decoders. These decoders gradually update the decoding parameters for all channels to account for changes in the recorded activity (Bishop et al., 2014; Blumberg et al., 2007; B Jarosiewicz et al., 2014; Li, Lebedev, & Nicolelis, 2011; Shpigelman, Lalazar, & Vaadia, 2008; Zhang & Chase, 2013).

The primary goal of this study was to quantify the typical length of unit stability using a previously described statistical method to determine unit stability (Fraser & Schwartz, 2011). Neural recordings from microelectrodes implanted in motor cortex were compared within and across channels to estimate stability over time, on the order of both hours and days. If individual units maintained their waveform shape, firing rate, and firing correlations, they were marked as stable. In this way, I could determine the rate at which recorded units turn over. In addition to quantifying the length of unit stability, this analysis provides the starting point for analysis of whether unit characteristics, such as firing rate, waveform size, recording noise, firing rate variability, or percent shared variance (a model-free measure of tuning (Williamson et al., 2016)), could predict how long a unit will remain stable that follows (Section 4.0).

3.2 METHODS

3.2.1 Data collection and spike sorting

We recorded 287 sessions across 33 months from Subject 1, and 51 sessions across 6 months from Subject 2 for the interday analysis. For the within day stability analysis, 5-minute long recordings were taken every hour for 7 hours. I recorded 2 sessions each 9 and 15 months after implant, and 1 session 10 months after implant from Subject 1. I recorded 4 sessions from Subject 2 at 3, 4, 5 and 6-months post-implant. Overall recording quality was tracked using two primary metrics: (1) the number of recorded units per channel and (2) the average peak-to-peak voltage of the largest unit on each channel (Figure 6 C & D) as in (Chestek et al., 2011).

Threshold crossing events were sorted offline to identify single and multi-unit activity. Here I use the term ‘unit’ to mean any sorted waveform, derived from the threshold crossing events. Units were sorted based on the spike snippets using custom sorting software written in Matlab (MathWorks, Natick, MA). This method relies on principal component analysis (PCA) to separate units based on their waveform shape. For each recording session and channel, the researcher viewed two plots of neural data prior to sorting the units. The first plot showed up to 250 overlaid spike snippets from a single channel (Figure 6A). The second plot showed all recorded snippets from that channel as points in the plane of the first two principal components (PCs) based on snippet waveform shape (Figure 6B). The researcher then specified the number of sortable units (including single-neuron, multi-neuron, and hash units) on a given channel and initialized a Gaussian mixture model expectation-maximization (EM) algorithm to create the

specified number of clusters in the PC space. Every recorded snippet was assigned to the cluster to which it had the highest probability of belonging to according to the EM algorithm. Once all snippets were assigned to a unit, any unit that had a firing rate less than 0.2 Hz was discarded prior to further analysis.

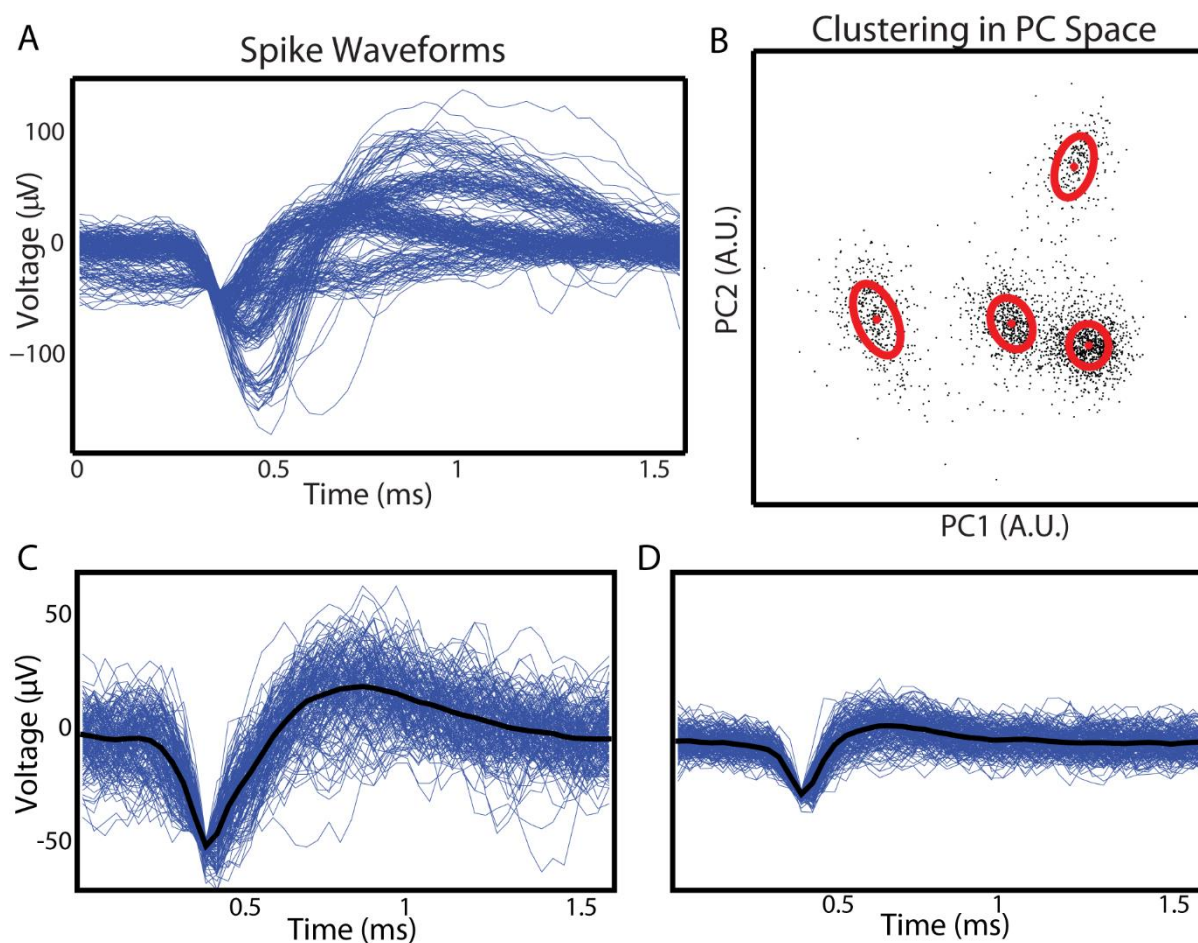


Figure 6. Sorting examples. A) The view of snippets shown during sorting. B) The same snippets in principal component space, with means (red dots) and standard deviations (red ellipsoids) for each cluster overlaid. C) 70 μV peak-to-peak unit, which was the median unit size 1-month after implant. D) 30 μV peak-to-peak unit, which was the median unit size 2-years after implant.

3.2.2 Unit isolation

After sorting units, they were assigned an isolation score. This score was calculated by taking the average of the dot products of each adjacent snippet, after normalizing them to remove effects of unit size. Higher isolation scores indicate that snippets were nearly identical, as noiseless recordings of action potentials from a single neuron would be. Lower isolation scores indicate that the unit includes multiple neurons with different shapes of action potentials, or a large amount of recording noise relative to the size of the recorded action potential. To determine what effect unit isolation had on stability I fit a two-term Gaussian model to the distribution of isolation scores and picked the local minima between the two maximums of the fit to be the threshold between well isolated and poorly isolated units.

3.2.3 Iterative stability determination

Once the snippets were sorted into units, a method developed by Fraser and Schwartz was used to determine if the units were stably recorded from one session to the next (Fraser & Schwartz, 2011) (posted to MATLAB Central as “Tracking neurons over multiple days”, identification no. 30113). A brief summary of this method is presented here. This method assumes that a neuron can only be recorded by a single electrode, and so neurons recorded on different electrodes must be unique. Based on this assumption, a null distribution of unit characteristic comparisons was computed based on neurons recorded from different channels. Although there is no way to be completely sure that two recordings come from the same neuron, a unit was considered to be stable if the unit

characteristics measured on two different days were more similar than 95% of the comparisons in the null distribution. The four unit characteristics I compared were waveform shape, auto-correlation, mean firing rate, and cross-correlation with other units. Per Fraser et al, these characteristics “are individually weak classifiers, but, because they represent independent sources of information, they can be combined into a strong classifier.”

As a first estimate of stability, pairwise comparisons were made from recordings taken on the same electrode during adjacent testing sessions. I refer to this as the iterative stability method. Pairwise comparisons continued longitudinally across all recording sessions. If a unit was stable from the first session to the second, and then again from the second to the third, it was considered stable across all three sessions. This yielded a monotonically decreasing number of stable units across time. For the intraday stability analysis, iterative pairwise comparisons were made to the recording sessions taken each hour.

The iterative stability determination method provided a conservative estimate of the number of stable neurons across three or more sessions due to its iterative nature. A false negative result during any single session comparison marked a unit as unstable for the rest of the test period, even if it was labeled as stable in all session comparisons following the false negative. While false negatives could inflate estimates of instability, a complementary feature of the iterative approach was that false positives, which occur at a 5% rate (based on the decision boundary being set at 95% of the null distribution), were rapidly corrected after only a few comparisons. Therefore, the iterative comparison method for determining stability provided a lower bound on the estimated number of stable units over the test period.

3.2.4 Direct comparison stability determination

An alternative method to determine stability is to compare each session to all following sessions directly. This direct method of comparison allowed units that may not have been recorded during some sessions to be marked as stable if they were recorded again later. The direct comparison method provided a more generous estimate of unit stability over time. With this method, each comparison has a 5% false positive rate. However, false negatives which incorrectly mark a unit as unstable, are limited to single pairwise session comparisons rather than every subsequent session. By using both the iterative comparison method and the direct comparison method, I provided a range within which the true number of stable units falls.

To perform the direct comparison, I compared each session directly to the subsequent 50 recording sessions, or all remaining sessions if there were fewer than 50, to identify stable units. I chose to only do direct comparison across 50 sessions at a time because that would always yield comparisons for at least 100 days, and continuing the comparisons beyond that point was computationally expensive while providing minimal additional information. For the intraday comparison, the first recording set in each day was directly compared to the subsequent 6 recordings. This intraday analysis was repeated separately for well-isolated and poorly-isolated units (Section 3.2.2).

3.3 RESULTS

3.3.1 Recording quality metrics

Both subjects showed the largest number of recorded units, which as defined here includes well-isolated neurons, multi-neuron waveforms, and hash, in the first two months after implant. Similarly the peak-to-peak voltage of the recorded units increased over first month of the implant before slowly decreasing. The maximum number of units per channel for Subject 1 was 1.6, while Subject 2 reached 2.3. The median peak-to-peak voltages peaked at 89 and 102 microvolts respectively. The number of recorded units and peak-to-peak voltage then declined over time (Figure 7), which is consistent with previous non-human primate and human subject studies (Chestek et al., 2011; Simeral, Kim, Black, Donoghue, & Hochberg, 2011; Suner, Fellows, Vargas-Irwin, Nakata, & Donoghue, 2005). The recording quality metrics exhibit a discontinuity at day 565 for Subject 1 due to a change in the spike threshold from -5.25 to -4.5 RMS. By moving the threshold closer to 0, more units with lower peak-to-peak voltages are recorded. This change was motivated by reports that small amplitude units carry significant movement-related information and can contribute to decoding (Fraser, Chase, Whitford, & Schwartz, 2009; Oby et al., 2016). Subject 2 had more units per recording channel compared to Subject 1. This could be attributed to using the more sensitive threshold for the duration of the Subject 2's implant, though the median peak-to-peak voltage size was similar.




Figure 7. Signal quality metrics. A) The number of sorted units per channel for each recording session. The inset shows the first 6 months for both subjects. B) The distribution of peak-to-peak voltages across all recorded channels for each recording session. Again, the inset shows the first 6 months for both subjects. The dots show the median peak-to-peak voltage for the day, with the dashed lines showing the interquartile range.

3.3.2 Interday recording stability

Data was compiled across all recording sessions to obtain the median trajectory for unit stability over the subsequent days and weeks of testing. Figure 8 shows the median and interquartile ranges for estimated unit stability rates using both the iterative (Figure 8A) and direct (Figure 8B) comparison methods. In both cases, I fit the time course of unit loss with a two-term exponential function. This implies that units do not uniformly decay at a constant rate. Instead, there is a large proportion of units (~75%) that become unstable within one week. The remaining units have longer term stability and are lost at a slower rate. The exponential function used to describe unit stability for both the iterative and direct comparison methods was:

$$\textit{Proportion of Units Still Stable} = ae^{-bt} + ce^{-dt} \quad \text{Equation 3.1}$$

where t is the number of days since an initial recordings session. As long as $b > d$, $a \times 100$ is the percentage of the recorded population that is stable in the short-term, $c \times 100$ is the percentage of the recorded population that is stable in the long-term, $(1 - e^{-b}) \times 100$ is the percentage of the short-term stable units that become unstable each day and $(1 - e^{-d}) \times 100$ is the percentage of the long-term stable units that become unstable each day.

I chose a two-term exponential model to reflect the observed bimodal clustering of decay rates within the population. It is possible that with higher resolution data we would find a continuous spectrum of decay rates instead of two discrete rates. As an approximation, we can consider the faster decaying term of the exponential fit, those represented by coefficients a and b above, to represent units that are stable only in the short term (~75% become unstable within a

week), while the slower decaying exponential term, represented by c and d above, describes the long-term stable.

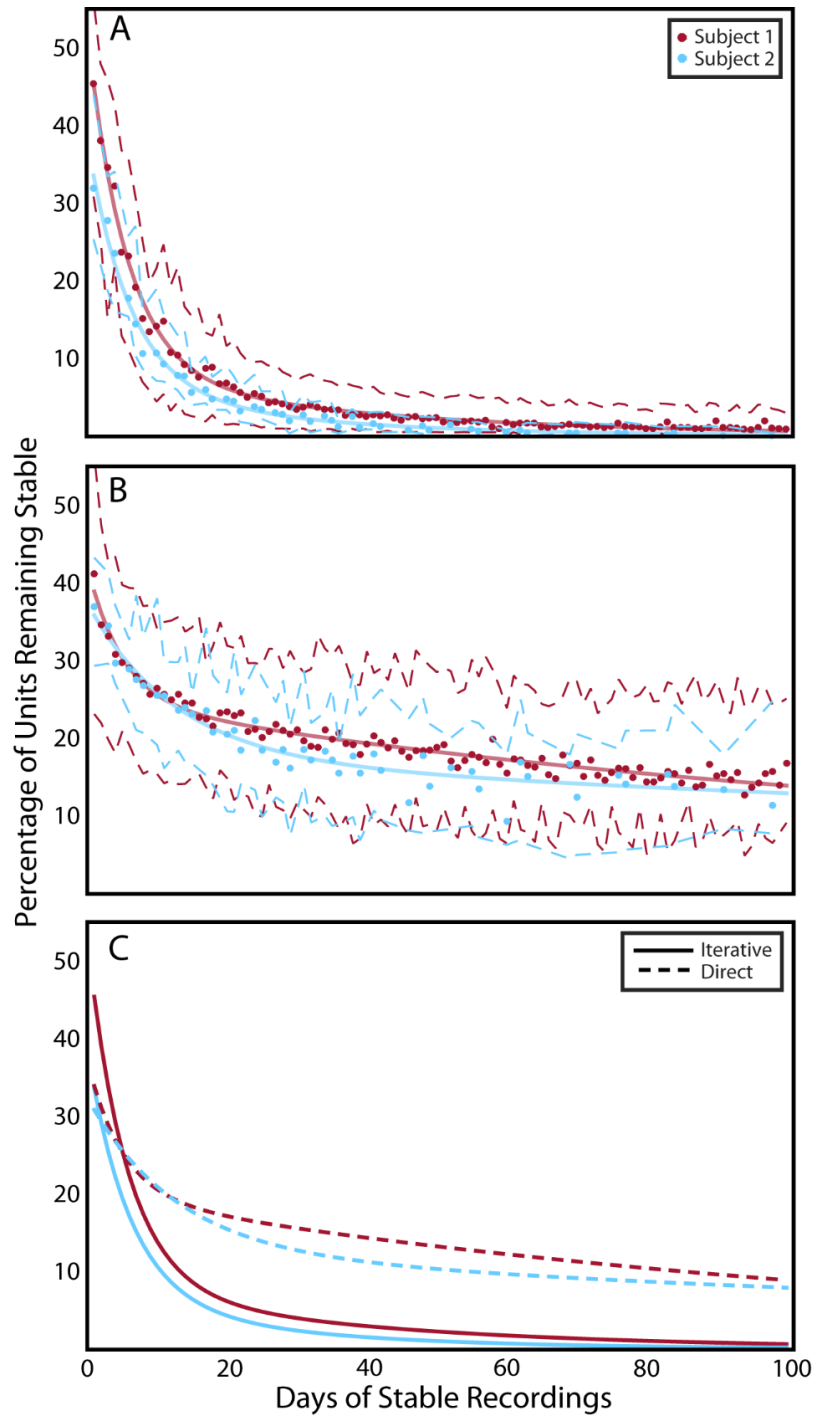


Figure 8. Estimates of interday recording stability. A) Percentage of units remaining stable calculated with the iterative-comparison method. B) Percentage of units remaining stable calculated with the direct comparison method. The data has been adjusted to account for the known 5% false positive rate. For A and B, each dot shows the median percentage of units remaining stable for each number of days, dashed lines show the interquartile range. The solid lines show the exponential fit to the medians. C) The exponential fits to estimate unit stability profiles are shown for both subjects and both comparison methods. The iterative-comparison method is more conservative and estimates fewer long-term stable units whereas the direct comparison method estimates that 10-15% of units can remain stable for months. The solid lines are for the iterative-comparison method, dashed lines for the direct-comparison method.

Figure 8C shows the two-term exponential models of unit stability for Subjects 1 and 2 using both the iterative and direct comparison methods. The participants exhibited very similar trends in the rate of unit stability.

Using the more conservative iterative-comparison method, 44.7% of units recorded from Subject 1 showed short-term stability, with 16.3% of those units becoming unstable each day. An additional 8.2% of units showed long-term stability, with only 2.6% of these units becoming unstable each day. The remaining 47.1% of recorded units did not remain stable for a full day. Subject 2 showed similar trends although a larger percentage of units (61.4%) did not remain stable for a fully day. 32% of units showed short-term stability, with 15.3% becoming unstable each day. Another 6.6% of units showed long-term stability, with 3.7% becoming unstable each day. As expected using the direct comparison approach, a larger percentage of units were considered to be long-term stable (19.7% for Subject 1 and 12.5% for Subject 2). Nearly 20% of units were stable in the short term, for 1 day to 1 week, for both subjects. Using the direct comparison approach, I estimated that approximately two-thirds of units become unstable within a single day. The estimated decay rates were lower using the direct estimation of stability as compared to the iterative approach, which is much more conservative across repeated sessions (Table 1).

Table 1. Unit stability distributions and decay rates.

Unit Stability Distributions and Decay Rates					
		Subject 1		Subject 2	
		Percentage of Units	Decay Rate (per day)	Percentage of Units	Decay Rate (per day)
Interday Iterative	Stable < 1 day	47.1%	N/A	61.4%	N/A
	Short-term Stable	44.7%	16.3%	32%	15.3%
	Long-term Stable	8.2%	2.6%	6.6%	3.7%
Interday Direct	Stable < 1 day	62.8%	N/A	67.6%	N/A
	Short-term Stable	17.5%	18.1%	19.9%	7.7%
	Long-term Stable	19.7%	0.8%	12.5%	0.5%
		Percentage of Units	Decay Rate (per hour)	Percentage of Units	Decay Rate (per hour)
Intraday Iterative	Stable < 1 hour	19.6%	N/A	20%	N/A
	Stable	80.4%	8.7%	80.0%	12.5%
Intraday Direct	Stable < 1 hour	23.2%	N/A	27.7%	N/A
	Stable	76.8%	3.6%	72.3%	5.8%

3.3.3 Intraday recording stability

The interday stability analysis revealed that half to two-thirds of the recorded units were stable for less than one day, depending on the analysis method and subject. I sought to examine unit stability over a shorter time scale by collecting 8 recordings, one every hour for 7 hours, to identify units that were stable on the order of hours, rather than days.

For both subjects, the iterative and direct comparison methods estimated that 72-80% of recorded units remained stable for at least 1 hour (Figure 9). Figure 9A and Figure 9B show the unit stability profiles for the iterative and direct comparison methods. Data was not always

collected at exact 1 hour intervals, so data is plotted using the time it was actually acquired. Overall, Subject 1 appeared to have a slightly higher percentage of stable units, which is consistent with the interday results. Both subjects showed the expected trend, in that fewer units could be recorded stably over time.

I fit the intraday stability estimates with a single-term exponential function as illustrated in Figure 9C. The exponential fits were used to estimate the rate at which stable units could no longer be recorded. Using the more conservative iterative-comparison approach, I estimated that stable units were lost at a rate of 8.7% or 12.5% per hour for Subjects 1 and 2 respectively. When I instead used the direct-comparison method, I estimated that the number of stable units decreased at a rate of 3.6% or 5.8% per hour for Subjects 1 and 2 (Table 1).

I can combine the Data from the inter- and intra-day experiments to provide a rough estimate of overall unit stability. Using the iterative-comparison method, approximately 25% of units are stable on the order of hours, 47.5% of units are stable on the order of days, and 7.5% are stable on the order of weeks. Using the direct-comparison method, approximately 45% of units are stable for hours, 17.5% for days, and 12.5% for weeks.




Figure 9. Estimates of intraday recording stability. A) Percentage of units remaining stable calculated with the iterative-comparison method. B) Percentage of units remaining stable calculated with the direct comparison method. The data has been adjusted to account for the known 5% false positive rate. For A and B, each dot shows the percentage of units remaining at the time a recording was made. The solid lines show the exponential fit to the data. C) The exponential fits to estimate unit stability profiles are shown for both subjects and both comparison methods. The iterative-comparison method is more conservative and estimates fewer stable units whereas the direct comparison method estimates that at least 50% of units can remain stable for 7 hours. The solid lines are for the iterative-comparison method, dashed lines for the direct-comparison method.

3.3.4 Effect of unit isolation on stability

When I separated the sorted units into well isolated and poorly isolated categories (Section 3.2.2) before calculating the stability rates I found that poorly-isolated units are less likely to be stable. This was expected, because poorly-isolated units are more likely to represent action potentials from multiple neurons and/or have lower signal-to-noise ratios. If a unit is composed of action potentials from multiple neurons, there is a greater chance of an individual neuron becoming unstable and making the unit as a whole unstable. Noise can vary from session to session, so units that are susceptible to noise will show greater changes between sessions, especially in firing rate.

The proportion of units that were ever stable (ie. the y-intercept of the exponential fit) was 10% higher for well-isolated units and 8% lower for poorly-isolated units, compared to the unseparated population. Similarly, the decay rate was 15% lower for well-isolated units and 11% higher for poorly isolated units, compared to the unseparated population (Figure 10). While these differences are appreciable, there are still many poorly-isolated units that remain stable throughout the day, validating their inclusion in the analysis.

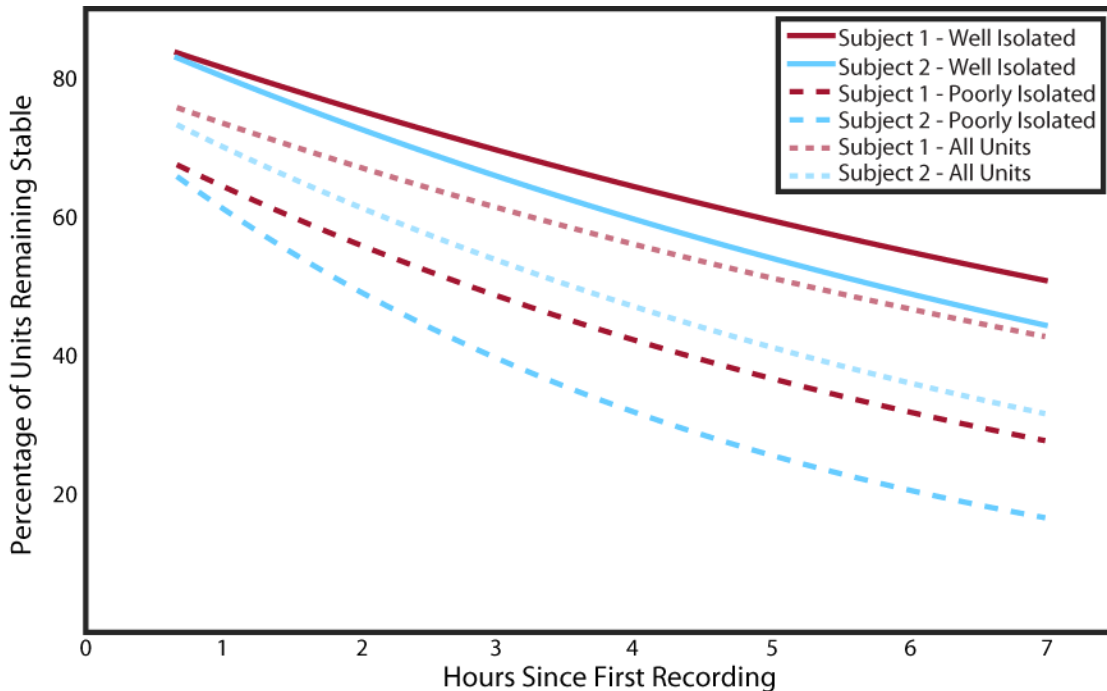


Figure 10. Estimates of intraday recording stability for different categories of isolation. The well-isolated units were the most stable, and poorly isolated units were the least stable, but all unit categories had the majority of units remain stable for at least 2 hours.

3.4 DISCUSSION

Intracortical microelectrodes arrays provide a way to measure the activity of populations of neurons in the cortex. This technology has enabled many scientific and engineering discoveries, for example in the area of motor control and brain-computer interfaces to restore movement (Collinger et al., 2012; Ganguly & Carmena, 2009; Hochberg et al., 2006; Beata Jarosiewicz et al., 2008; Sadtler et al., 2014). However, this technology also presents challenges as single unit recordings can be unstable and signal quality can degrade over time. Here I have presented a detailed look at recording stability from intracortical microelectrodes implanted in human motor cortex. Neural recordings were captured for multiple recording sessions each week for many

months, as well as multiple times during single days of testing to capture intraday stability. In general, these results confirm prior reports of signal instabilities with approximately 60% of units being lost after a single day. Even within a single day, 45% of units are unstable after 7 hours.

The long-term iterative-comparison results (Section 3.3.2) showed a slightly smaller percentage of stable units than Fraser et al. observed in the non-human primate they collected data from for months (Fraser & Schwartz, 2011). Similarly, the results Vaidya et al. observed in their non-human primate subject are comparable to the direct-comparison results I found in our human subjects (Vaidya et al., 2014). In both cases the short-term results, out to approximately 2 weeks, showed substantially higher rates of stability, but this may be due to the studies only including well-isolated units. This explanation is further supported by the unit isolation analysis (Section 3.3.4) which showed that well-isolated units are somewhat more likely to remain stable than poorly-isolated units. The similarity between previous stability analysis in non-human primates and this stability analysis in human subjects confirms the utility of further developing methods to account of instability of microelectrode array recordings in a non-human primate model.

While the results from both subjects were largely consistent, there were some small differences. It is likely that there will always be differences between any two subjects for a variety of reasons. In our study, the two subjects had different pathologies. Paralysis due to an injury, e.g. spinal cord injury (SCI) or brainstem stroke (Downey et al., 2016; Hochberg et al., 2012), is unlikely to cause a systemic change in the ability to record neurons in M1. Paralysis due to disease, e.g. spino-cerebellar degeneration or amyotrophic lateral sclerosis (ALS) (Collinger et al., 2012; Gilja et al., 2015), could result in systemic changes to the neurons in motor cortex that change the ability to record in M1. In addition to differences in pathology, biological differences such as age, amount of sleep, stress levels, and many other factors can alter the activity of neurons in cortex.

All of these factors could change the behavior of neurons, such that even if they can still be recorded they may not be recognized as stable. Gray matter volume decreases with age which could result in the electrode tips in older BCI users being deeper in cortex, changing which types of neurons are being recorded from (Seidler et al., 2010). Without a very large number of subjects, it will not be possible to control for all of these variables. It will be important to further investigate the differences between classes of likely users, currently people with ALS (Gilja et al., 2015), brain-stem stroke (Hochberg et al., 2012), or SCI (Flesher et al., 2016), so that BCI systems can be better tailored to their users if consistent differences are found in recording stability.

While the time course of instability does not provide evidence for a specific cause of instability, the smoothness of the decay curve rules out some possibilities. The consistent intraday loss of stable units while the subject has cables attached to their head rules out the possibility that interday instability is related to substantial head movements. While it is possible that rapid head movements could move the skull, and attached array wiring, relative to the brain and that the motion would move the electrode tips relative to the recorded neurons that is not the primary driver of instability. That still leaves many possibilities for causes of instability. One possibility is that micromotion of brain tissue around the electrodes due to changes in blood or intracranial pressure moves neurons relative to electrode tips. Another possibility is that neurons go through active and inactive states, in a less wide spread version of the inactivation that occurs during slow wave sleep (Volgushev, Chauvette, Mukovski, & Timofeev, 2006), drastically changing their firing rates in the inactive state and therefore being recognized as unstable. Both theories would be consistent with the markedly increased stability rate found with the direct comparison method (Section 3.2.4). Neurons that move away from electrodes could move back towards the tip if movement resembles a random walk. Neurons would also be recognized as stable again if they returned to an active state

after entering an inactive state. Either of these options, or another condition under which the same neurons are recorded multiple times discontinuously, explain why the total number of recorded neurons decreases orders of magnitude slower than the number of stable units (Figure 7 and Figure 8). This discrepancy means that instability could not be explained by neurons dying, or permanently migrating away from the electrodes.

One approach to minimizing recalibration requirements is to identify stable units and use only those units for decoding. Only a percentage of the recorded population will remain stable for any duration of time. Depending on the complexity of the task, a smaller population of units may be sufficient to enable successful task performance. For example, Ganguly & Carmena (Ganguly & Carmena, 2009) showed that a monkey could complete a cursor task using the same decoder for up to 19 days using just 15 recorded units. If we use the results from the iterative-comparison analysis we get a conservative estimate that a 96-channel array (assuming ~1 unit per channel, Figure 7A) would provide at least 15 stable units for 9 days. If we instead use the results from the direct-comparison analysis, we would expect to have at least 15 stable units for 29 days. We would therefore expect that if we could identify these stable units, cursor control could be maintained for 9-29 days.

In addition to predicting how long performance set number of neurons used in a decoder will remain stable, this stability analysis should aid in the development of unsupervised recalibrating decoders (Bishop et al., 2014; Blumberg et al., 2007; B Jarosiewicz et al., 2014; Li et al., 2011; Shpigelman et al., 2008; Zhang & Chase, 2013). These decoders adjust to the loss of units from the recorded population and the addition of the new units that replace them by recalibrating the decoder during use. The stability decay rates shown here should provide recalibration algorithms with the a priori estimates of how soon the recordings will reach that limit

and trigger a recalibration before performance decays too far. The results from the intraday comparisons are most pertinent here, showing that around half of the units turn over in the course of an 8-hour day. This indicates that updates will need to occur at least every few hours, which will require a system that is used consistently through the day. If the unsupervised recalibration can be made seamless, then a BCI user may only have to participate in supervised decoder training once and, then, be able to use the BCI for at least the rest of the day.

These results provide a clear understanding of the problems of instability. A majority of recorded units becomes unstable within a day, but those that remain stable for at least 1 day are likely to survive for multiple days with a smaller subset lasting weeks or months. There are enough stable units to control a cursor for weeks if the BCI system can identify them when calibrating a decoder. This may be sufficient for a spelling system and, for higher-DoF systems, could allow the user to issue commands for recalibration without the help of a technician. The addition of unsupervised recalibrating decoders could further improve the longevity of decoders. In order to take advantage of these stable units I next attempted to identify characteristics of units that provide evidence that unit will be one of the stable ones in the population (Section 4.0).

4.0 PREDICTING INSTABILITY

4.1 INTRODUCTION

Using the results from section 3.0 as a data set, I investigated if unit characteristics could predict how long a unit will remain stable. If a BCI system can reliably identify units that are more likely to remain stable for a long time, decoders that preferentially use the more stable units will decrease the amount of intervention that is required. These decoders would further increase the user's independence, and decrease the burden on the user due to decoder calibration.

The characteristics I investigated were unit firing rate, waveform size, recording noise, firing rate variability, or percent shared variance (a model-free measure of tuning (Williamson et al., 2016)). I measured each of these characteristics on the first day that each unit was recorded for use in predicting whether, and for how long, a unit would be stable. As an additional analysis, I also studied whether the length of time that a unit has already been stable is predictive of how much longer it will remain stable. In other words, if a unit has already been stable for one week, is it more likely to remain stable for another week than a unit that is being recorded for the first time?

This analysis shows that high firing rate, large waveform peak-to-peak voltage, and well-tuned activity are associated with units that will remain stable for weeks. Also, the longer a unit has been stably recorded, the longer you can expect it to remain stable, with diminishing returns for units that have been stable for more than a month. These results provide a promising indication that researchers will be able to develop long-term BCI decoding methods for clinical use. By identifying units that are more likely to remain stable, and preferentially using those for decoding, it should be possible to develop decoders that work for longer periods of time before needing to be retrained.

4.2 METHODS

4.2.1 Unit characteristics

The unit characteristics I studied were: peak-to-peak voltage, baseline noise, firing rate, firing rate variability, and percentage of firing rate variance explained by a factor analysis model. Unit characteristics from the first day the unit was recorded were used for analysis. These characteristics were defined as follows:

1. Baseline noise was calculated as two times the standard deviation of the first 5 and last 5 samples of each 48-sample unit snippet (Figure 11). These samples were chosen because each snippet is aligned such that the first 10 samples occur before the threshold is crossed,

and most waveforms will have returned to baseline approximately 1 ms (30 samples) after crossing threshold. Units with high baseline noise might be expected to be less stable because the recording noise could overwhelm the signal.

2. Peak-to-peak voltage was calculated as the difference in voltage from the trough to the peak (Figure 11) of the average waveform (black lines in Figure 6 C&D) for each unit.

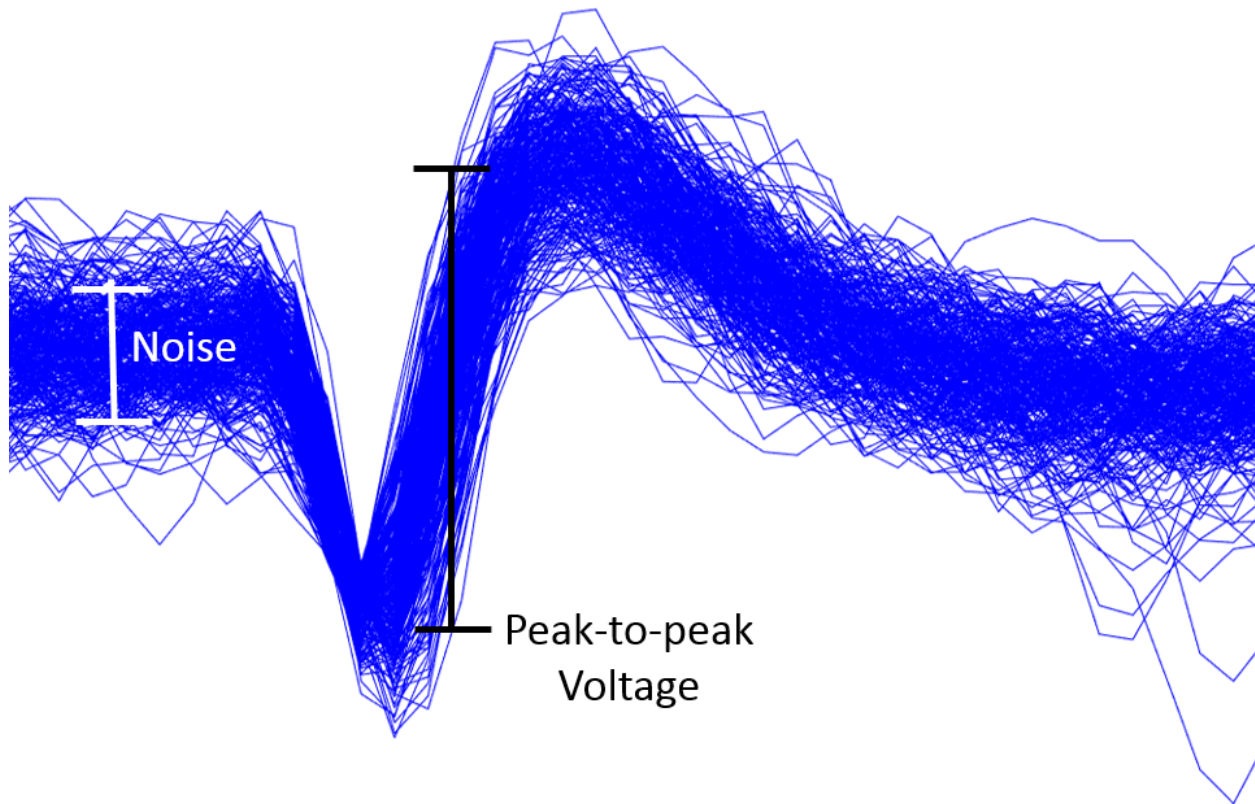


Figure 11. Illustration of baseline noise and peak-to-peak voltage.

3. Firing rate was calculated as the average number of threshold crossings per second across the 3 or 5-minute recording. Firing rate was square-root transformed to better approximate a normal distribution (Bromiley & Thacker, 2001).

4. Firing rate variability was calculated as the inverse of the variance of the inter-spike interval divided by the mean firing rate of the unit. A unit with a high level of firing rate variability might be more likely to be unstable if this is an indication that its firing behavior is more transient.

5. Model-free tuning was calculated using a factor analysis model. Since different BCI tasks were conducted across days, a 10-dimensional factor analysis model was used to identify the amount of shared firing rate variance for each unit as described in Williamson et al (Williamson et al., 2016). Factor analysis approximates the covariance matrix of the neural population activity as the sum of two components: $\Sigma \approx \Lambda + \Phi$. Here, Σ is the $n \times n$ covariance matrix measured for all n neurons in a single session. Λ is a low-rank covariance matrix that captures the 10 dimensions explaining the most shared covariance of the population. Φ is a diagonal matrix that summarizes the remaining variability not accounted for by shared fluctuations in the population activity. It represents the “private” or independent variance of each neuron that is not explained by other recorded neurons. This shared firing rate variance served as a proxy for tuning to the various BCI tasks. This value was is a ratio between how much a units firing rate is tied in to the larger network, and how much its firing rate varies in isolation. A 10-dimensional model was chosen to match the highest number of degrees-of-freedom decoded during the study (Wodlinger et al., 2014). This characteristic is called model-free tuning in this section. Higher levels of model-free tuning indicate that a neuron’s activity is better predicted by the population, which suggests this unit is less contaminated by independent noise. This could indicate a more stable unit.

I tested each characteristic to determine which were significantly different between units that were stable for at least one day, and units that were not stable between days using the Wilcoxon rank-sum test. I also tested to determine if there were differences in each characteristic for units that were stable for 1-6 days, 7-21 days, or greater than 21 days using the Kruskal-Wallis test.

4.2.2 Cox analysis

In order to determine the relative utility of the different unit characteristics for predicting the length of stability, I used the Cox method (Cox, 1972). The Cox method yields a survival function of the form:

$$S(t, i) = S_0(t) \exp(\sum_k b_k (C_k(i) - \bar{C}_k)) \quad \text{Equation 4.1}$$

where $S(t, i)$ is the survival rate for unit i at time t , $S_0(t)$ is the baseline survival rate at time t . The exponent of the equation is called the relative risk. The relative risk is calculated using the b regression coefficients for each characteristic: noise, peak-to-peak voltage, firing rate, firing rate variability, and model-free tuning. The subscript k denotes each different characteristic. $C_k(i)$ is the measured characteristic of each unit, and \bar{C}_k is the population mean for characteristic k . The higher the survival rate for a unit at a given time, the more likely it is to still be stable.

The distributions of all of the characteristics were heavily skewed. To minimize the impact of outliers on the analysis I normalized all the characteristics through log transformation before running the analysis.

4.2.3 Prior stability

I also sought to examine whether prior stability was predictive of future stability. For example, was a unit that had been stable for a week more likely to be stable for another week than a unit that had only been recorded for one day? I looked at all units that had been stable for at least x days and varied x from 1 to 60. For each of these days I then looked at the distribution for how many days longer each unit was stable. In this way I was able to compare the expected length of stability for units based on how long they had already been stably recorded. I only performed this analysis on Subject 1's recordings because Subject 2's recordings did not cover enough time. I did this analysis separately from the Cox analysis because prior stability is highly correlated with the unit characteristics that predict stability.

4.3 RESULTS

4.3.1 Characteristics of stable units

Units that were stable for at least one day had significantly different firing rate, firing rate variability, peak-to-peak voltage, noise level, and model-free tuning than units that were unstable ($p < 10^{-18}$ for all characteristics for both subjects). Similarly, in Subject 1, units that were stable for less than a week, 1-3 weeks, or greater than 3 weeks exhibited significant differences in all 5

unit characteristics ($p < 10^{-16}$, Kruskal-Wallis, Figure 12). Subject 2 showed a significant difference ($p = 0.009$ for noise, $p < 0.001$ for all others) for all characteristics other than firing rate variability ($p = 0.32$) for units that were stable for the same 3 time periods.

Unstable units had lower noise levels, lower peak-to-peak voltages, lower firing rates, less firing rate variability, and lower model-free tuning. Determining the relative predictive power of each characteristic required a more comprehensive analysis. The Cox method in the next section calculated the relative ability of each characteristic to predict a unit's survival rate.

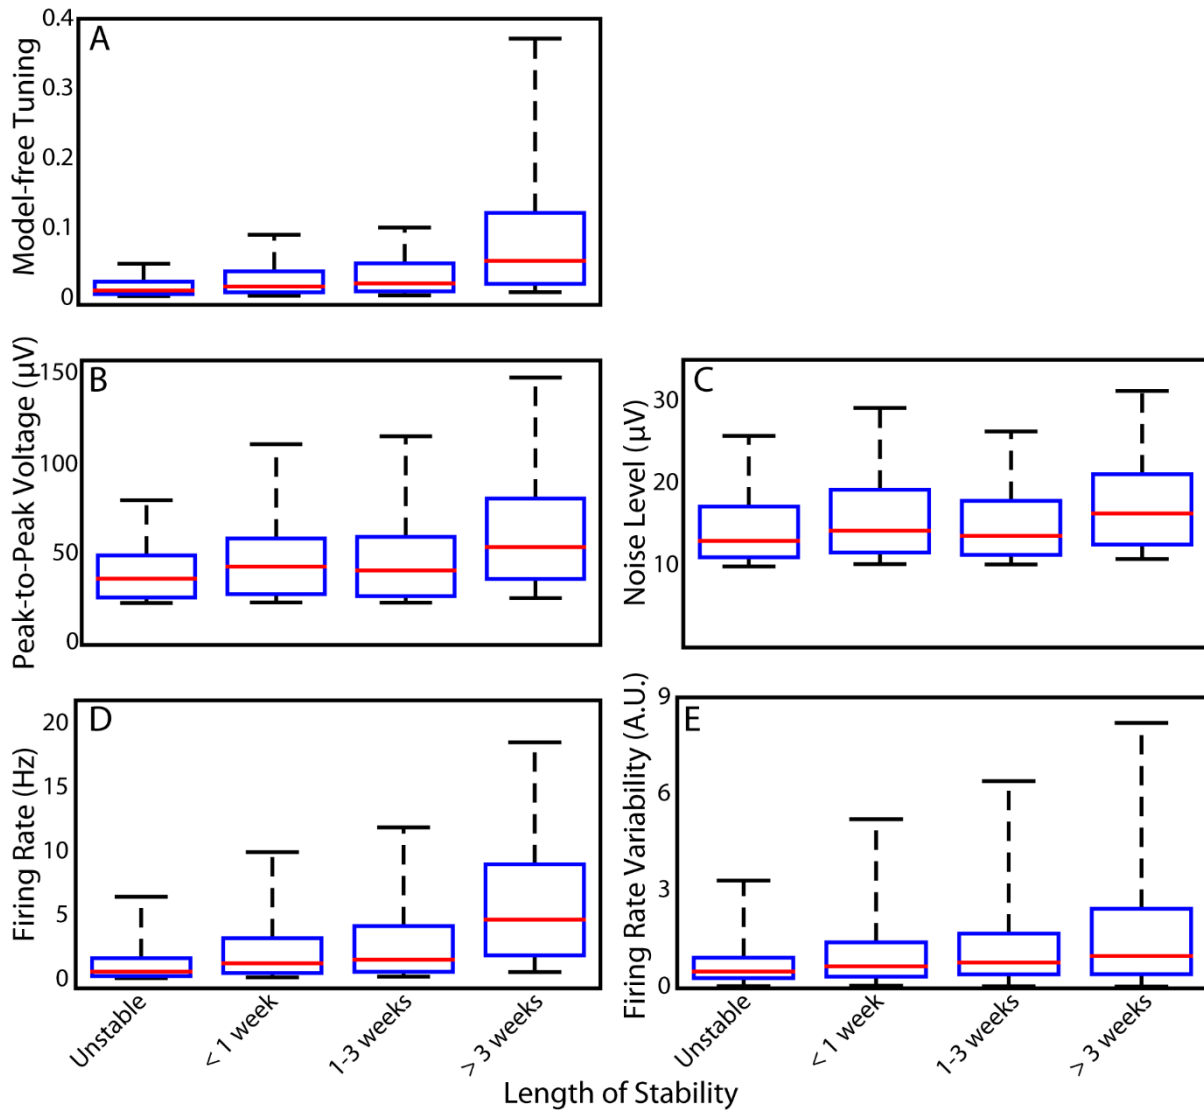


Figure 12. Distributions of characteristics of units based on stability. In all plots, the red line is the median, blue box outlines the interquartile region, and the whiskers extend from the 5th-95th percentile. A) Model-free tuning from the 10D factor analysis model. B) Peak-to-peak voltage. C) Magnitude of the noise during baseline recording. D) Average firing rate of the unit. E) Firing rate variance.

4.3.2 Cox analysis

I created a Cox model to predict the survival rate of a given unit over time based on its noise level, peak-to-peak voltage, firing rate, firing rate variability, and model-free tuning. For both subjects, higher peak-to-peak voltage, firing rate and model-free tuning all raised the survival rate. Firing rate had the largest effect, followed by peak-to-peak voltage, and then model-free tuning. Firing rate variability did not have a significant effect on the survival rate. Baseline noise had a large effect on stability in the data from Subject 1, but not Subject 2. This could be because all of Subject 2's data was collected early in the implant when signal quality is generally higher. To demonstrate the magnitude of the effect of changes in each characteristic, I calculated the relative risk for hypothetical units. The relative risk is the exponent from the survival function (Equation 4.1), which is calculated from observed characteristic values and the regression coefficients. An increase in relative risk decreases the survival rate of a unit. The hypothetical units I calculated relative risk for had the median value for each subject's entire recorded unit population in 4 of the 5 characteristics with the remaining characteristic at the 25th and 75th percentile values for the population. The change in relative risk when this is done for each characteristic is shown in Table 2.

Table 2. Change in relative risk, given a change in each neural characteristic.

Change in relative risk, given a change in each neural characteristic						
	Subject 1			Subject 2		
	25th %ile Value	75th %ile Value	Change in Relative Risk	25th %ile Value	75th %ile Value	Change in Relative Risk
Noise	11 uV	18 uV	15%*	21 uV	29 uV	1%
Peak-to-Peak Voltage	27 uV	53 uV	-13%*	53 uV	90 uV	-7%*
Firing Rate	0.50 Hz	2.4 Hz	-20%*	1.0 Hz	5.6 Hz	-13%*
Firing Rate Variability	0.28	1.0	1%	0.34	1.9	4%
Model-Free Tuning	0.006	0.028	-8%*	0.003	0.026	-10%*

The percent change in relative risk when each characteristic was changed from its 25th percentile value to its 75th percentile value while all other characteristics are held at their median. Negative changes indicated that increasing the value of a given characteristic made a unit more likely to remain stable. Starred changes in relative risk were significantly different from a 0% change.

The Cox analysis allows us to visualize the survival rates for the “best” and “worst” units. Here I define the best-case unit as one in the 95th percentile for firing rate, peak-to-peak voltage, and model-free tuning, and in the 5th percentile for noise and firing rate variability. The worst-case unit is one with the inverse characteristics. By plotting the predicted stability for these two cases and for a unit with the median value in all characteristics, we can visualize the range of possible survival rates (Figure 13A). The best-case unit had a relative risk that was 3.7x lower the worst-case unit and 2.1x lower than a unit with the median value of each characteristic. This means a best-case unit has a 36.8% chance of being stable for one week, compared to 12.8% for a median unit and only 2.4% for a worst-case unit. The best-case unit is also much more likely to be stable for 1 month, 9.5% chance, compared to the worst-case unit, 0.015%.

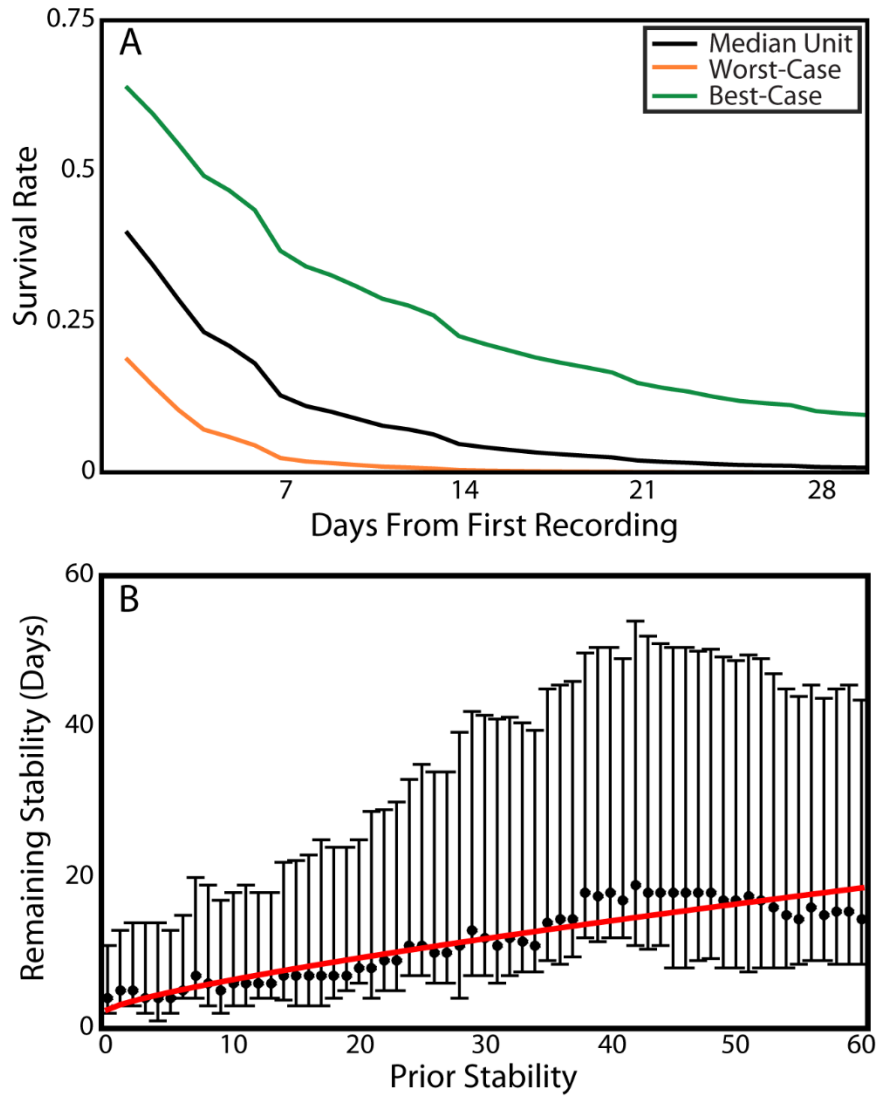


Figure 13. Predicted stability. A) Based on the Cox analysis, the best-case units (green) show far higher survival rate than units with median (black) or worst-case (orange) characteristics. B) The median and interquartile range for how long a unit will remain stable after being stable for the number of days on the x-axis. The red line shows the fit to the medians of a power function.

4.3.3 Prior stability

For Subject 1, I could also examine whether the amount of time that a unit had been stable predicted future stability. As shown in Figure 13B, the longer a unit had been stable, the longer it was likely to remain stable in the future. The relationship of prior stability to remaining stability is reasonably well fit ($R^2 = 0.85$) by the power function below:

$$\mathbf{Remaining\ Stability = 0.66t^{0.78} + 2.48} \qquad \mathbf{Equation\ 4.2}$$

Where t is the number of days a unit has already been stable. While this fit does not appear to account for the apparent plateau in the data after day 40, it provides a reasonable intuition for the importance of a unit's history.

4.4 DISCUSSION

I sought to determine whether there were particular characteristics that distinguished stable units from those units that were more likely to be unstable. I investigated unit recording noise, peak-to-peak voltage, firing rate, firing rate variability, and model-free tuning because they provide information about a unit's activity, consistency, clarity, and modulation and are independent of which task the subject is performing. In general, stable units tended to have a higher firing rate, peak-to-peak voltage, and model-free tuning value. Units that were not stable had lower values of these characteristics, likely because many of them were low-amplitude activity without clearly

defined action potentials. These indistinct units were included in analysis because they are often used for decoding (Gilja et al., 2015; Wodlinger et al., 2014) and contain significant information about movement parameters (Fraser et al., 2009; Oby et al., 2016). The unit isolation analysis (Section 3.2.2) confirmed that poorly isolated units, typically those with low-amplitude activity without clearly defined action potentials, were less likely to be stable, but enough of them remained stable to justify their inclusion in the analysis. This analysis not only provides a chance to predict which units will be stable during BCI use, but also provides evidence regarding the causes of recording instability.

When using these unit characteristics in a regression model to predict how long a unit will remain stable, I found that higher firing rates were most predictive of prolonged stability with higher peak-to-peak voltage and model-free tuning values also providing significant improvements in expected length of stability. Considering the relative importance of each of these unit characteristics could provide further evidence for some of the different hypothesized causes of unit instability.

The importance of higher firing rates for indicating units that will remain stable for a long time is particularly interesting since firing rate should be unrelated to electrode proximity. One explanation of the importance of high firing rate may be that higher firing rate units are always active during BCI tasks and never enter the inactive state proposed in section 3.4. Furthermore, since firing rate is unlikely to correlate with the proximity of the recorded neuron(s) to the electrode tip, this suggests that some mechanism other than micromotion (Section 3.4) is contributing to unit instability. However, the fact that higher peak-to-peak voltage, which correlates to the proximity of neurons to the electrode tips, is also predictive of stable units, supports the idea that micromotion could contribute to instability.

High peak-to-peak voltage could also be indicative of a larger cell, for instance a pyramidal cell vs. an interneuron. If that difference is the cause of the correlation with stability then it may support the theory of active and inactive states where interneurons, which are smaller and therefore move fewer ions during an action potential, resulting in a smaller recorded peak-to-peak voltage, are more likely to enter the inactive state and become unstable. Future work could attempt to identify whether particular types of neurons are more or less stable through waveform shape analysis (Bean, 2007) to further prove or disprove this theory.

Finally, the predictive value of model-free tuning indicates that when a unit's firing rate varies in concert with the rest of the recorded population it remains stable for longer. This may indicate that cell types that directly contribute to the motor plan (e.g. pyramidal cells) remain active during motor tasks, and are therefore more often observed to be stable during the motor task recordings used here.

When considering these 3 predictive characteristics it seems more likely that instability is related to the type of neuron being recorded, rather than motion of the neurons about the electrode tip. By this theory, certain classes of neurons are consistently in an active state while the BCI user performs tasks, while other recorded classes alternate between an active and an inactive state. Had I found that the overwhelmingly predictive characteristic was peak-to-peak voltage I would have concluded that units with low-amplitude waveforms were less stable because the neurons that generated them were further from the electrode tip. Assuming that neurons can move towards or away from the electrode, those that are far from the tip can become slightly larger or disappear respectively, while those near the electrode will just become larger or smaller. The theory of active and inactive states needs much more evidence, first by a thorough waveform analysis (Bean, 2007). Another experiment could use concurrent recording and calcium imaging of firing in non-human

primates to observe whether the recorded neurons move, or simply go silent when they become unstable (O'Shea et al., 2017).

By preferentially decoding from units with high firing rates, high peak-to-peak voltages, and higher levels of model-free tuning with the unsupervised recalibrating decoders mentioned in section 3.4, long-term stable decoders should be readily achievable. With these long-term stable decoders, BCI users will gain a degree of independence that people with upper limb paralysis currently lack. Additionally, lowering the amount of effort required on the part of users to maintain decoder performance will likely improve usage rates for the devices further justifying their implantation.

5.0 OBJECT INTERACTION

5.1 INTRODUCTION

The purpose of a BCI controlled neuroprosthetic arm is to replace lost hand and arm function in people with upper limb paralysis. In order to restore hand and arm function, the BCI needs to enable reliable grasp and manipulation of objects. Most BCI studies have focused on cursor control or reaching within a very constrained environment (Aflalo et al., 2015; Gilja et al., 2015; Hochberg et al., 2012; Simeral et al., 2011). BCI-enabled object manipulation, where the subject has simultaneous control of grasp and arm movement, has not been previously studied outside of our research group.

The laboratory of Hansjorg Scherberger has provided much of the research using extracellular recordings to study how the brain plans and executes movements to reach to and grasp different objects. The Scherberger lab has placed microarrays in anterior intraparietal cortex (AIP), ventral premotor cortex (PMv) and primary motor cortex (M1) of multiple non-human primate subjects to study the progression of information in the brain during reaching to objects. During the testing they used different object shapes and orientations as well as instructing the animals to grasp with different hand shapes (Fluet, Baumann, & Scherberger, 2010; Menz,

Schaffelhofer, & Scherberger, 2015; S. Schaffelhofer, Agudelo-Toro, & Scherberger, 2015; Stefan Schaffelhofer & Scherberger, 2016; Townsend, Subasi, & Scherberger, 2011). The Ebner laboratory also found information about whether a non-human primate was reaching towards a cylinder, cube, or polygon prism, in recordings from dorsal premotor cortex (PMd) (Hendrix, Mason, & Ebner, 2009)

They found information about the object shape and orientation were well encoded in AIP neurons when the object was presented and while the animal planned the movement, but poorly encoded during the movement (Fluet et al., 2010; S. Schaffelhofer et al., 2015; Townsend et al., 2011). Ventral premotor cortex also encodes information about object shape and orientation during object presentation and movement planning, but also during movement (Townsend et al., 2011). In addition to this tuning, PMv represents the grip type, or shape of the hand, with the encoding appearing during object presentation, but getting stronger during the movement phase of the task (Fluet et al., 2010; S. Schaffelhofer et al., 2015). Finally, Scherberger and colleagues found that M1 does not encode information about the object shape or orientation, but does encode arm movement and hand shape kinematics during the movement phase of the task (Menz et al., 2015; S. Schaffelhofer et al., 2015; Stefan Schaffelhofer & Scherberger, 2016). However, other labs report identifying information about the objects being reached towards in M1 (Hendrix et al., 2009; Mason, Gomez, & Ebner, 2002; Rouse & Schieber, 2016). It seems likely that by controlling for hand shape and object orientation in addition to the shape of the target object, the Scherberger laboratory showed that what might otherwise be identified as object tuning in M1 is actually hand shape and orientation tuning.

This work indicates that decoding from M1 during movement will provide information about the hand shape for an object, but that no direct information about the object itself would be

present. This is consistent with our previous work showing that while using a BCI, a person can continuously and simultaneously control 10 dimensions of an anthropomorphic robotic arm including positioning and orienting the hand and wrist, as well as multiple grasp postures (Wodlinger et al., 2014). However, that study also showed that calibrating the decoder in the presence of virtual objects improved the ability of the BCI user to interact with real objects. This may mean that subtle details about hand positioning in dimensions that are not decoded are present and need to be accounted for, or that human BCI users do have object information in M1 in contrast to the intact non-human primate subjects in the studies above. Without access to recordings from PMv or AIP our work focuses on how to manage the changes in neural activity in M1 during reaches to objects.

In addition to differences in which areas of cortex are being recorded from, intact non-human primates and BCI users with tetraplegia differ in a few important ways. The first, and likely most important, is that movement with an intact native arm provides readily interpretable proprioceptive and tactile sensation to reinforce the success of the movement. In contrast, a BCI user typically has only visual feedback, and even that is partially obscured when grasping objects. Our recent work stimulating in S1 to simulate sensation still does not provide proprioceptive information, and is a much poorer source of tactile information than an intact hand (Flesher et al., 2016). For these reasons, it is important to understand how M1 encodes grasping movement in proximity to objects.

Although our 10 degree of freedom BCI enabled skilled and consistent performance, we noted difficulties when the subject tried to grasp objects (Wodlinger et al., 2014). This often manifested as an unexpected velocity away from the target object, almost like a repelling force, or

the inability to execute a grasp around an object. If we removed the object from the workspace, the ability to control the kinematic dimensions of the arm was restored.

In order to calibrate the BCI, the participant had observed the robotic arm move through all 10 dimensions of control in various combinations. To mitigate the challenges with object grasping, we began to calibrate in a virtual-reality environment where the participant could observe a virtual arm grasping and transporting virtual objects (Figure 5B). Incorporating virtual objects into our calibration paradigm helped to overcome the challenges with grasping simple objects, but there remains a gap in our understanding of how grasping movements are controlled during functional task performance.

A number of non-human primate studies have investigated the encoding of movement parameters in motor cortex (Kakei et al., 1999; Moran & Schwartz, 1999; Serruya, Hatsopoulos, Paninski, Fellows, & Donoghue, 2002; Wang et al., 2013), but here I sought to specifically investigate how the presence of an object impacts motor cortex activity. I initially identified a wide-scale increase in firing of recorded M1 units as the hand reached towards an object that did not occur when the same reaching movement was made without an object. I attempted to adjust our decoding method to account for this unexpected change in neural activity. By uniformly scaling the firing rates of the recorded population so that the decoded population firing rate was constant, i.e. removing the average increase in firing rates prior to decoding, I was able to improve the performance of the two subjects in tasks requiring interaction with objects.

5.2 METHODS

5.2.1 Experiment sessions

The experiments with Subject 1 took place during 9 sessions, from 795 to 850 days after implantation. Experiments with Subject 2 took place during 3 sessions, from 661 and 673 days after implantation.

5.2.2 Object interaction task

The goal of this task was to determine how the presence of an object affected neural activity during reaching and grasping. Each subject completed two conditions: (1) Reaching to and grasping a cylindrical object and (2) Reaching to a position target and closing the hand in the absence of a physical object in the workspace (Figure 14A). The robotic arm started at the right side of the workspace for each trial. An audio cue indicated the start of the trial and the subjects moved the robotic arm to within 5cm of the target that was 35 cm to the left of the starting position, this movement was the reach phase of the trial. The same position target was used to judge the success of the reach phase regardless of whether an object was in the workspace. Once the hand was within the target area a chime prompted the subject to either close the fingers to a grasp position or to hold the fingers open for 2 seconds while remaining within 5 cm of the arm position target, based on the instructions received at the beginning of each 5 trial block. Only Subject 1 completed trials

where the hand remained open, we instructed Subject 2 to grasp for every trial. Most days the subjects controlled 5 DoF for this task, 3 degrees of translation, 1 for wrist roll, and 1 for grasp (Figure 4). Subject 1 completed 1 session without control of the wrist roll dimension, which was unnecessary for completion of the task.




Figure 14. Object interaction tasks. A) The object interaction task involved the hand to move to the same location, outlined by the dashed white line, whether there as an object present, left, or not, right. B) The object transport task required the subjects to grasp the cylinder on the left side of the table, carry it over the taped off region, and place it on the right side of the table.

5.2.3 Impact of object interaction on neural activity

The binned, filtered and square root transformed spike counts (section 2.1.2) were z-scored using the means and standard deviations from the data used to calibrate the decoder each day

(section 2.2.2.2). I did the analysis using z-scores to allow for a fair comparison between days, despite the changes in recorded units described in Section 3.0 . I aligned the z-scored firing rates on the time that the hand first reached the target region to allow for comparison between trials. I assigned the 1.5 seconds before the hand reached the target to the reach phase, and the 1.5 seconds after the hand reached the to the hold phase. In the case of Subject 1 who completed “open hand” trials in addition to “grasp” trials, I combined the data from each in to a single data set after finding that there was no significant difference during the reach phase between the two cases.

5.2.4 Online scaling of neural firing rates

In order to account for the change in neural activity when the subjects move the robotic hand to an object, I implemented a scaling factor. Because the change in activity was a broadly distributed increase in firing rate, I chose a population wide approach to address it. This also has the benefit of being immediately generalizable between days, regardless of instability in the recordings (section 3.0). The scaling factor was calculated by dividing the average firing rate during the previous 300 ms of the decoded population used by the average firing rate during decoder calibration. The firing rates of all of the decoded channels were then divided by this scaling factor before being decoded to generate kinematic commands. This maintained the relative difference in firing rates between the different channels, while lowering the decoded average population firing rate to be consistent with the data on which the decoder was calibrated.

5.2.5 Object transport task

I designed the object transport task to quickly assess the subjects' ability to pick up objects and transport them to a new location. The goal of the task was to pick up a cylindrical object (15.5 cm tall, 4.5 cm diameter, with a weighted base) on one side of a table, transport it across a 20 cm wide taped off zone on the table, and place it on the other side (Figure 14B). An experimenter immediately picked up the object and placed it back in the starting position. The goal of this task was to transport the object as many times as possible in 2 minutes. This task was modeled after the Box and Block task (Mathiowetz, Volland, Kashman, & Weber, 1985), but uses just a single, larger object. I used the task to test the effect of turning on or off the scaling factor on the functional performance of the BCI arm system by using the same decoder with or without scaling in alternating trials.

5.3 RESULTS

5.3.1 Object interaction task

During the object interaction task, Subject 1 successfully completed 88% of trials when no object was present (34/45 trials ending in a grasp, 45/45 trials ending with the hand open). In contrast, she only successfully completed the exact same movement 60% of the time if an object was present

in the target region (18/45 trials ending in a grasp, 36/45 trials ending with the hand open). Subject 2 successfully completed 79% of trials during object interaction when no object was present (55/70, all of Subject 2's trials ended with a grasp). In contrast, he only successfully completed the exact same movement 21% of the time if an object was present at the target region (15/70). The differences in performance of 28 and 58 percentage-points illustrate the scale of the object interaction problem ($p = 0.02$ and 3×10^{-5} , Wilcoxon rank-sum test).

Initially I checked to see if there was a population firing rate effect due to the subject's plan to grasp the object, or only due to the object being present. By comparing trials (45 of each) where the subject reached to grasp an object to trials where she reached to position the open hand next to the object we found no significant difference in firing rates ($p = 0.148$, Wilcoxon rank-sum test). Similarly, grasp and open-handed trials to no object did not show significantly different firing rates from each other ($p = 0.638$, Wilcoxon rank-sum test). For this reason, I did not run the open-handed case with Subject 2, and for the rest of the analysis on open-handed and grasp trials are combined in to a single data set for Subject 1.

The z-scored population firing rate shows a significant difference during the reach phase of trial, before the hand arrives at the target region. In the last 1.5 seconds of the reach phase, Subject 1's median population firing rate across 90 reaches was .44 standard deviations above expectation if a cylinder was present in the target region. In contrast, it was only .31 standard deviations above expectation if no object was present (different with $p < 10^{-100}$, Wilcoxon rank-sum). For Subject 2 these values were 0.38 and 0.27 respectively across 70 trials of each ($p < 10^{-11}$, Wilcoxon rank-sum). This change can be seen in more detail in Figure 15 where a pseudo-raster shows that the increased firing rate effects the majority of recording channels, and becomes stronger the closer the hand gets to the object.

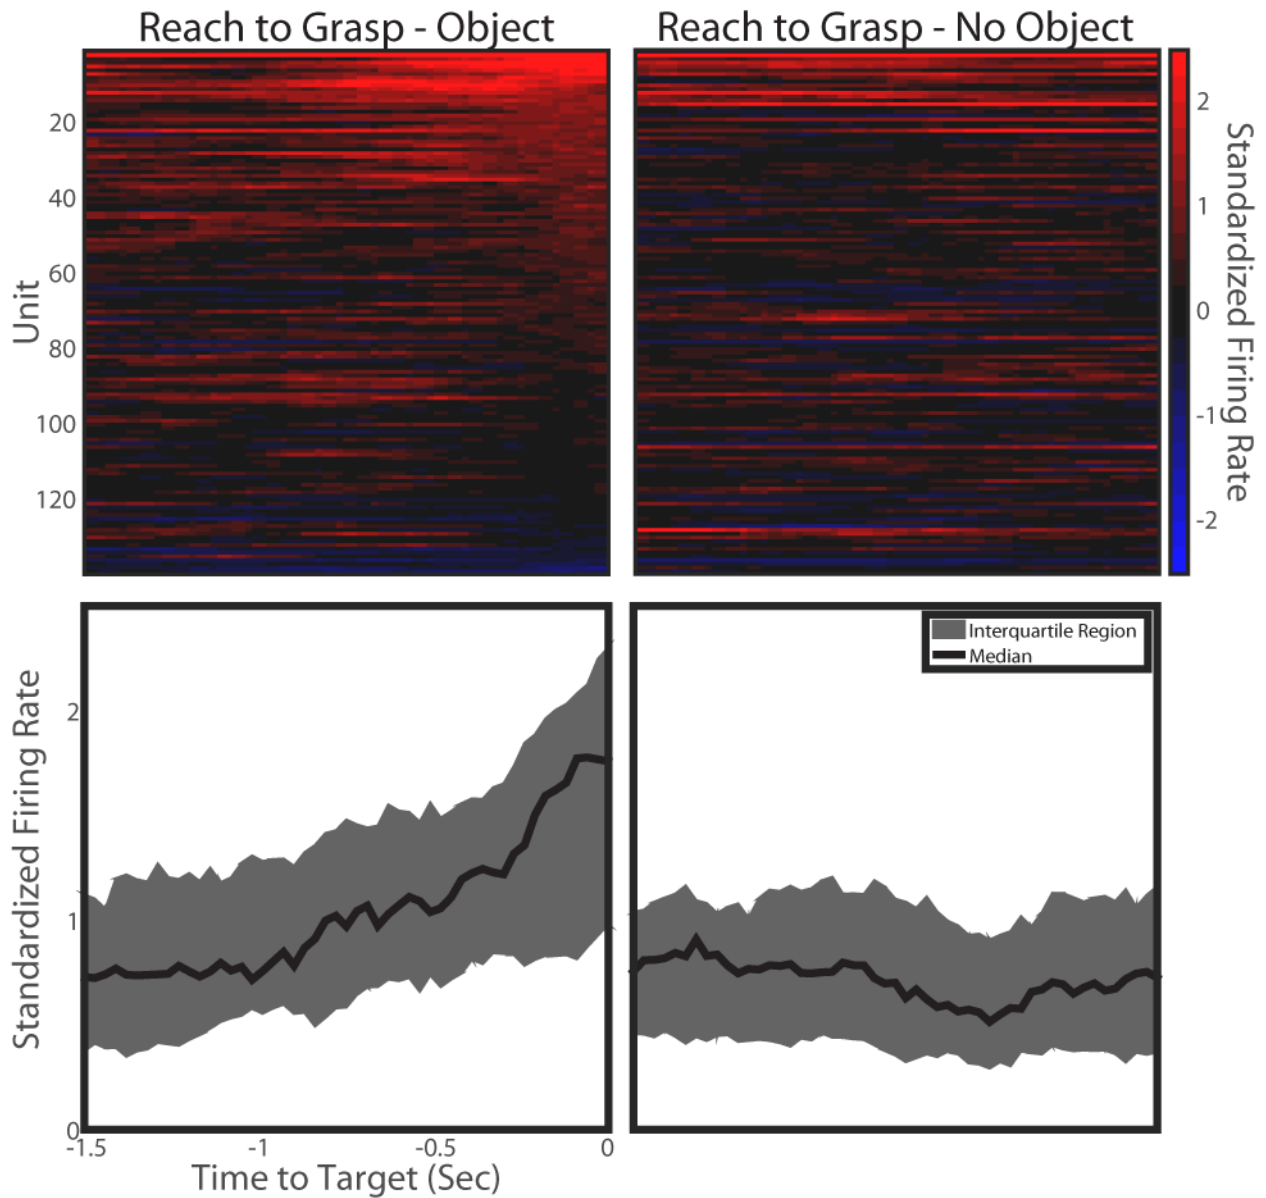


Figure 15. Neural response to object interaction task. The left side shows the last 1.5 seconds of the reach phase when reaching to an object, the right shows the same time period when reaching to no object. The top shows the average z-scored firing rates across trials for each unit in the 10 trials from a single day for Subject 1. Notice the increased brightness near the end of the reach phase across most units when an object was present, but no such increase without an object. The bottom shows the population z-scored firing rate distributions across all 9 of Subject 1's sessions.

5.3.2 Implementing scaling

The large increase in firing rate during reaches to an object provide an explanation for the poorer performance during reaches to the object. This population wide increase in firing rate violates the tuning model used to create the decoders used in this BCI system. According to the tuning model, the firing rate of each channel is a linear function of desired movement velocity (Equation 2.1). The coefficients that convert desired velocities to firing rate are typically independent (Georgopoulos et al., 1988; Wodlinger et al., 2014). Here we found that a large portion of the recorded population increased its firing rate significantly, but only when an object was present. This indicates tuning to a factor other than velocity, and that the tuning is not represented independently across the population.

In order to adjust for this unexpected change in firing rates, I implemented a scaling approach. The scaling factor was calculated by taking the average population firing rate over the most recent 300 ms and dividing that by the average population firing rate during decoder calibration. Anecdotally, scaling factors during testing tended to range between 0.9 and 1.8, with the highest values occurring in the second before the hand reached the object. Each channel's firing rate was then divided by this scaling factor before being decoded. This approach can be viewed as an implementation of divisive normalization, as seen in a number of sensory areas (Carandini & Heeger, 2012).

By implementing this scaling factor approach, we were able to raise Subject 2's success rate from 21% to 56% when a cylinder was in the target region ($p = 0.0012$, Wilcoxon rank-sum). He was successful on 79% of trials without scaling and 77% of trials with scaling when no object was present, showing no detriment to performance when objects were not present ($p = 0.94$,

Wilcoxon rank-sum). Subject 1 did not show a significant improvement in success rates with scaling on this simple task, perhaps because she had extensive practice attempting to overcome the object interaction effects on her own before scaling was implemented. The utility of scaling was more apparent in her control during the more functional task of object transport described below.

5.3.3 Object transport

I used the object transport task to provide a test of the level of functionality enabled by the scaling approach. Subject 1 was able to transport an average of 1.28 objects per minute with scaling implemented, while she was only able to transport 0.30 objects per minute without scaling across 23 trials of each. This increase in transport rate of over 400% was a significant improvement ($p = 0.00015$, Wilcoxon rank-sum). Subject 2 was able to transport 5.28 objects per minute with scaling, but only 4.65 objects per minute without scaling across 20 trials of each ($p = 0.0425$, Wilcoxon rank-sum).

5.4 DISCUSSION

These firing rate changes in M1 while interacting with objects have critical importance for BCI control. Properly identifying and addressing the issue that causes this discrepancy in control when an object is introduced will help to bring a degree of reliability and utility to BCI-controlled robotic devices that has not previously been achieved. Other papers have endeavored to approach control

discrepancies by addressing nonstationarities (Homer et al., 2013), like those quantified in section 3.0 , but these studies largely explore issues with reliability in control over time, not specific to object interactions which happen on a much shorter time scale. To my knowledge, complications with object interaction within the realm of BCI-controlled prosthetics has not been characterized prior to this work. This is because no studies have included object interaction/manipulation as a primary goal, more often focusing on 2D cursor control (Gilja et al., 2015; Hochberg et al., 2006) or control of an arm within a very constrained environment (Aflalo et al., 2015; Hochberg et al., 2012). Given that reliable grasp of objects is essential to the purpose of BCI-controlled robotics, a deeper investigation of this phenomenon was warranted.

This result also seems to be in conflict with the findings that object information is not encoded in M1 during reach to grasp movements in non-human primates (Menz et al., 2015; S. Schaffelhofer et al., 2015; Stefan Schaffelhofer & Scherberger, 2016). This discrepancy may be because those studies always reached to an object, and M1 more strongly represents that an object is present than which object is present. It could also be that the non-human primate subjects were over trained to the task and that the object interaction effect we observe has a novelty component that was extinguished in their studies. Further study in both human BCI users and non-human primates is warranted to determine the cause of this difference.

We first noticed difficulties with grasping objects with the MPL under BCI-control during the early 7D control experiments, and they became particularly problematic during 10D control experiments (Collinger et al., 2012; Wodlinger et al., 2014). This would manifest in a variety of ways: such as a velocity change in the opposite direction of the intended object, pushing the hand away from the object; or as the inability to grasp the hand once it was near an object; or the hand grasping before getting all the way to the object. The problematic movements would differ from

test session to test session, but were typically consistent within a session. When we calibrated decoders without any objects in the workspace, we often observed this phenomenon. When objects were introduced into the calibration phase the unintended movements were initially overcome. Later with our first subject, unintended movements while interacting with objects returned, even with decoders calibrated with virtual objects. It is at this point that we began the work shared here.

In this study we endeavored to explore a solution for dealing with these changes in firing patterns seen when the BCI user attempt to manipulate objects. This approach involved scaling the firing rates before decoding them to generate movement commands. This scaling removed the correlated increase in firing rates across most of the population, and improved the ability of the subjects to interact with objects.

An alternative approach to managing the difficulty of interacting with objects is a technological solution. I authored a paper, in collaboration with the National Robotics Engineering Center, showing that by combining computer-vision based robotic grasping with BCI control we could create an effective system for BCI users to reliably grasp objects with minimal effort. This paper is included as Appendix A.

While this scaling has addressed the problems caused by object interaction in the context of a BCI controlled robotic arm, many questions remain around the neurological origin of this disruption. The broadly increased firing rates in M1 during object interaction are presumably the result of inputs from other areas on the visuomotor pathway that should be traced. It is also worth investigating whether this tuning to objects is only present during neuroprosthetic arm use, perhaps due to a lack of sensory feedback, or whether it is present during native arm movements as well. Finally, it may be worth considering whether the increased firing rates are correlated with co-contractions that occur in the arm when a person interacts with an unpredictable object, in an

attempt to stabilize the arm. Answering these questions will help to better understand human motor control and further improve our ability to use neuroprosthetic arms to interact with the environment.

6.0 GRASP FORCE ENCODING

6.1 INTRODUCTION

In addition to the change in tuning observed as the BCI controlled hand approaches objects, it will be important to recognize how tuning changes while the user handles objects. In order to provide general assistance to the user, the arm will need to be able to provide a wide array of forces, as dictated by the different objects the user wants to manipulate. With the recent demonstration of the ability to use intracortical microstimulation to convey graded information about the force applied to individual fingers (Flesher et al., 2016) it is possible to return information about grasp force to the BCI user. The current method for decoding arm movement (section 2.2.1) only generates velocity commands, relying on the robot's constraints to determine how much force gets applied to objects. Information about desired wrist, and finger force has been shown in non-human primate primary motor cortex (M1) (Gaberet, 1999; Kakei et al., 1999). These studies were done with isolated wrist flexion or pinching, not while interacting with objects or making natural movements. Here I explore the representation of desired grasp force in M1 of a human BCI arm user, which could be used to enable closed-loop control of grasp force.

The Ebner laboratory has performed studies of neural tuning to grasp force in dorsal premotor cortex (PMd) and primary motor cortex (M1). Their first finding was that while M1 represented hand and arm movement plans prior to moving, force was not encoded in M1 until the hand was grasping the object (Mason et al., 2002). This provides evidence that the change in firing rate in M1 during object interaction described above (Chapter 5.0) is not due to plans to exert force on the object. Following up on their earlier work, the Ebner laboratory found that grasp force and hand shape were encoded independently in M1. The same study showed that PMd did not encode desired grasp force earlier than M1 (Hendrix et al., 2009), in contrast to the encoding of hand shape that appeared earlier in PMv than in M1 in the work in the Scherberger laboratory (Menz et al., 2015; S. Schaffelhofer et al., 2015; Stefan Schaffelhofer & Scherberger, 2016). The independence of hand shaping and grasp force, along with the apparent coincidental appearance of grasp force encoding in PMd and M1 indicates that hand shaping and grasp force planning may follow different paths in the cortex. Likely because the grasp force is readily updated after contact to account for erroneous estimations of object weight (Descoins, Danion, & Bootsma, 2006; Flanagan, Bowman, & Johansson, 2006), while hand shaping is less likely to need rapid updates.

In addition to integrating these previous findings in to a BCI system, I sought to determine whether M1 represented grasp force differently for a robotic arm. The robotic arm exerts force much differently than a biological arm, and does not provide the proprioceptive feedback that muscles would. For these reasons it is unclear if the grasping movement would be planned and executed by a BCI user as by an intact non-human primate. During the preparation for the movement, the internal models (Davidson & Wolpert, 2005) used to plan a robotic arm movement are likely to differ from a biological arm movement. However, the differences may not be drastic

as experiments have shown that numbing the fingertips only leads to a small increase in exerted force (Cole & Abbs, 1988).

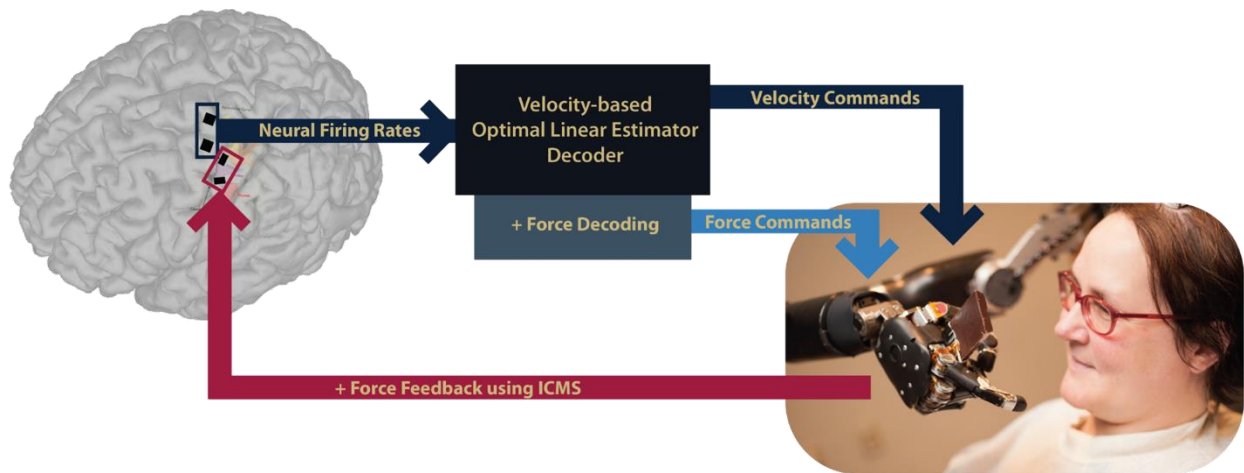


Figure 16. Closed-loop control of grasp force. The schematic shows the important components of a closed-loop grasp force BCI system. The hand would still be positioned and oriented using velocity commands. Once an object is grasped the subject will feel the amount of force exerted through intracortical microstimulation (ICMS) and be able to alter the grasp force through decoded commands.

6.2 METHODS

6.2.1 Section specific recording details

For this experiment, only Subject 2 performed the experiments. This was done because Subject 2 could receive intracortical microstimulation to simulate the sensation of grasp force, making it

possible to provide meaningful feedback of decoded grasp forces. All recording was done in M1 only, because hardware constraints currently prevent the arrays in S1 from recording and stimulating without the cables being switched (see Section 2.1 for more details about recordings).

6.2.2 Imagined grasp vs. object observation

During the imagined grasp vs object observation task the subject watched a TV screen for stimulus presentation. The task involved the presentation of 4 objects: a marshmallow, a tomato, an orange, and a can of soup (Figure 17A). I picked these objects in consultation with Subject 2 to be evenly spaced in the perceived level of grasp force required to grip them for lifting. During object observation trials, one of the four objects was presented for 4 seconds, followed by 4 seconds of a fixation cross. During imagined grasp trials, one of the objects was presented for 1.5 seconds before a chime cued the subject to imagine grasping the object with the amount of force necessary to lift it for the remaining 2.5 seconds that the object was presented. The trial finished with 4 seconds of fixation cross presentation (Figure 17B). During each of the three test sessions, the subject was first presented with 36 or 48 trials of one type, followed by the same number of the other type for a total of 132 trials of each.

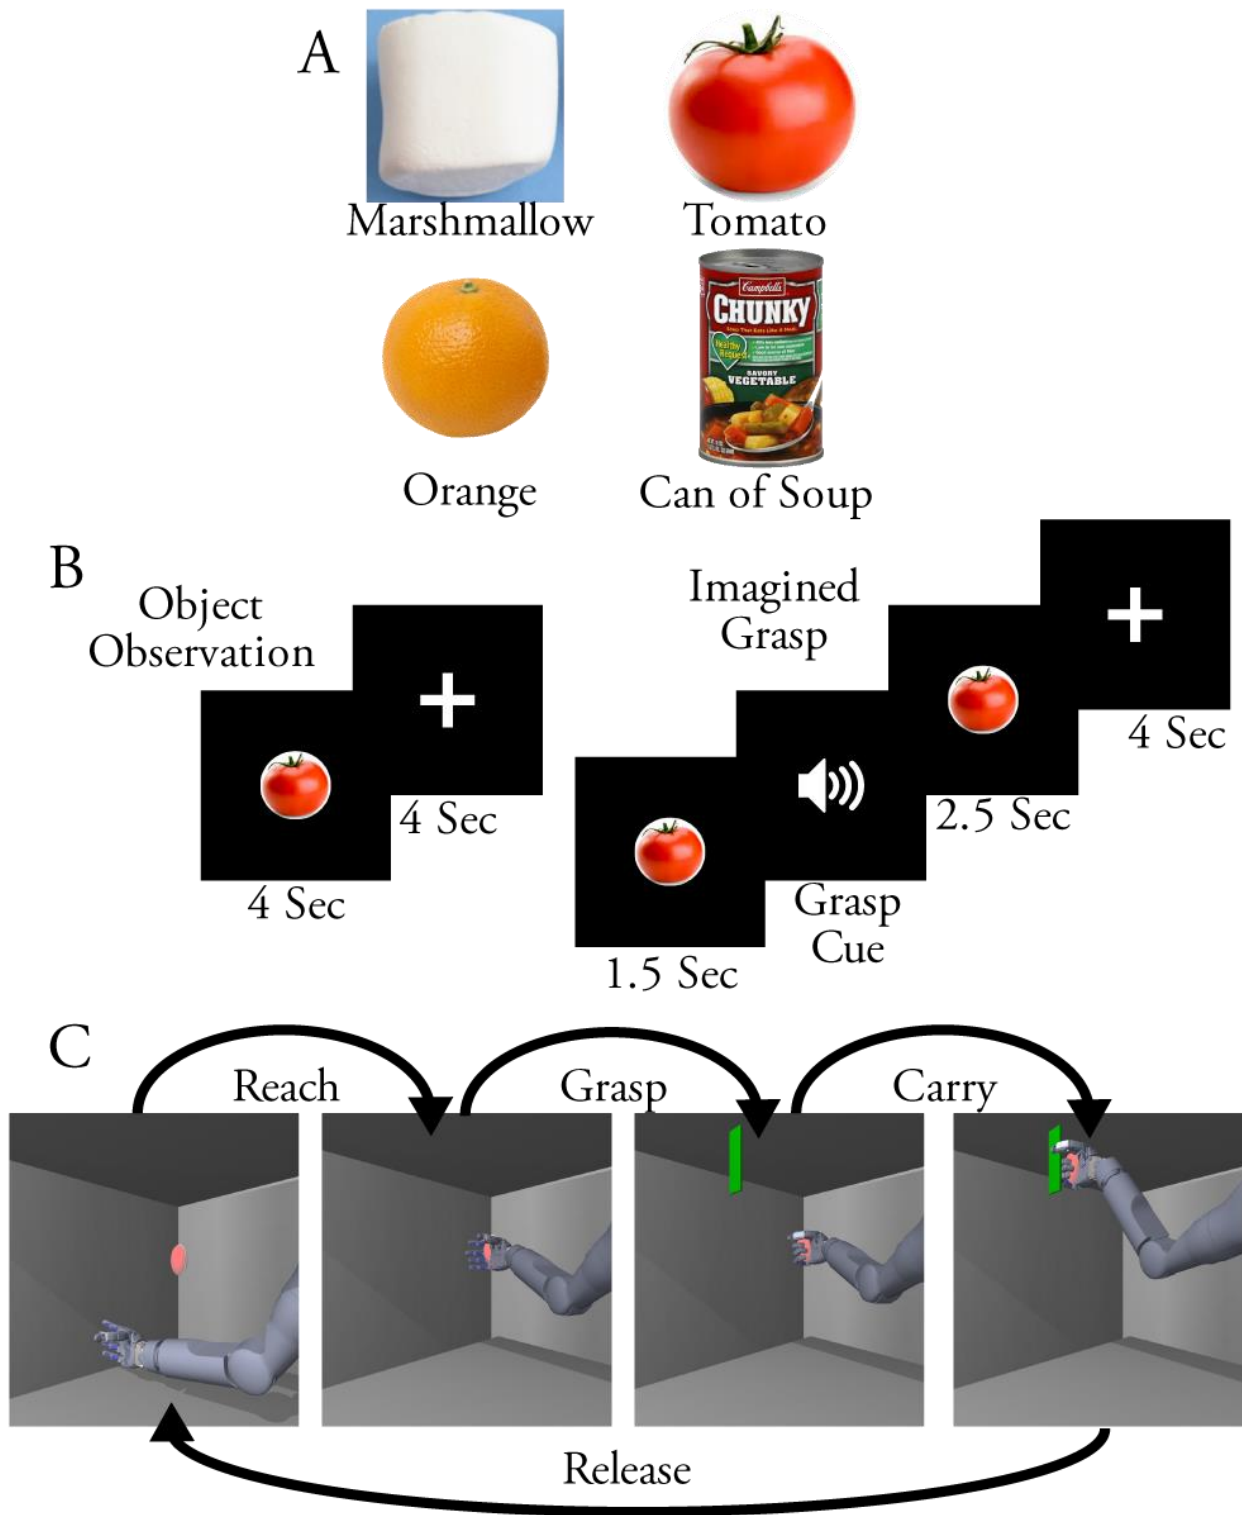


Figure 17. Grasp force methods figure. A) The 4 images used to show the force objects in both tasks. B) The order and timing of the imagined grasp trials and the object observation trials. C) An example trial of the grasp and carry task. For this task the name of one of the 4 force targets was spoken at the start of the trial.

6.2.3 Reach and grasp task

To study grasp force in the scenario of whole arm movement I adapted a virtual reality object task that we used in a previous study (Wodlinger et al., 2014) (section 2.2.2.1, Figure 5B). For this study the task required control over 3 degrees of freedom, hand translation up-down and left-right as well as grasp. For the entire task the subject was instructed to attempt to complete the task, as the computer generated the actual arm kinematics, as occurs in the first step of the 2-step calibration method we use for decoder training (Collinger et al., 2012; Downey et al., 2016; Wodlinger et al., 2014). At the start of each trial one of the 4 force objects was presented as an audio cue, to be used for the entire trial. Then a red ellipsoid appeared in the virtual scene at one of 5 translation target positions 52 cm in front of the shoulder. The targets formed a cross with one at the center, one 20 cm above the center, one 22 cm below the center, and one each 25 cm left and right of the center. Once the hand moved to the ellipsoid (the reach phase), an audio cue prompted the subject to attempt to grasp it with the amount of force appropriate for the force object presented at the start of the trial (the grasp phase). Then a green plane appeared at one of the 5 translation targets. Once the hand moved to the new target (the carry phase) an audio cue instructed the subject to release the ellipsoid (the release phase). After the release a new trial started with a new force object audio cue (Figure 17C). In both test sessions of this task 99 trials were completed.

6.2.4 Force object classification

To determine whether desired grasp force was encoded in the recordings from M1 I used naïve Bayes classification on 1 second of neural data starting .5 seconds after the grasp cue during the imagined grasp task and the corresponding time during the object observation task (2-3 seconds after the start of the trial). For cross-validation I left one random trial from each class out during classifier training, then tested the classifier on the left out trial, repeating until all trials had been left out. I repeated this process five times for each session to get multiple combinations of left out trials. Confidence intervals were determined by shuffling force object labels and repeating the classification process 500 times to determine the 95% confidence bounds on chance level.

To observe the timing of the representation of desired grasp force during both tasks I used the same classification procedure with a sliding window approach. For the imagined grasp and object observation task the entire trial was covered using 500 ms causal windows with a 100 ms step size to generate a time-series of classification accuracies. Because the reach and carry phases of the grasp and carry task were not the same length for all trials, I used the 1.6 seconds before, and 0.8 seconds after, the grasp and release cues for this analysis. If the reach or carry phase was skipped completely due to two consecutive translation targets being in the same place then only the 0.8 seconds after the grasp or release cue was used for that trial. I used these subsections of the trial to perform the same sliding window classification analysis as was done for the imagined grasp and object observation task.

6.3 RESULTS

6.3.1 Imagined grasp vs. object observation

Using the neural data from the imagined grasp vs. object observation task (section 6.2.1) I trained and tested a classifier (section 6.2.4). The classifier using the 1 second of neural data, starting 0.5 second after the grasp cue, was 69% accurate in identifying the force object presented to the subject (>33%, the 95% confidence interval calculated using label shuffling). An additional 30% of trials were misclassified as an object adjacent to the correct object on the desired force spectrum (eg. a tomato is presented but the classifier predicts a marshmallow, Figure 18). The classifier had no a priori information about the order of the objects so this pattern of mistakes indicates that the primary motor cortex is encoding the grasp forces in the order we would expect based on the chosen objects. The same analysis on data collected while the subject simply observed the objects, but did not attempt to grasp, yielded accuracies within the range of chance (29%, less than the 95% confidence threshold of 33% accuracy, Figure 18).

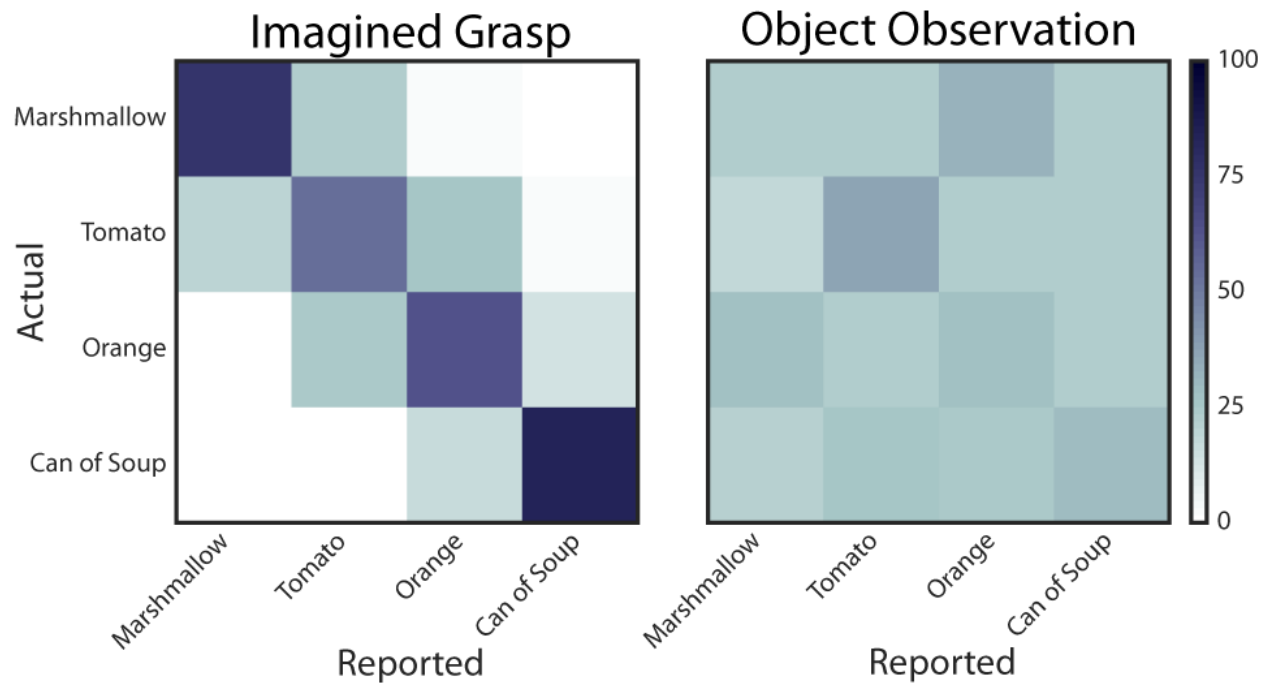


Figure 18. Grasp force classification accuracy during object observation and imagined grasping. Confusion matrices showing the accuracy of classification of the presented object from recordings in M1. On the left, the classification of the presented object is highly accurate when the subject is attempting to grasp it and errors are consistent with a continuous representation of grasp force. On the right, the classification of the presented object is random when the subject simply thinks about the object without attempting to grasp it.

6.3.2 Reach and grasp task

I used data from the more enriched reach and grasp task (section 6.2.3) to determine the timing of the presence of grasp force encoding in M1. While the arm reached towards the virtual object, the desired grasp force was not classifiable above chance level. However, consistent with the result of the imagined grasp task, the desired grasp force became classifiable at a higher than chance level during the grasp phase (47% accuracy vs. 34% chance level, Figure 19). During the carry phase however, classification accuracy returned to chance level. The subject confirmed that he continued

to attempt to grasp the object with the appropriate amount of force during the carry phase, indicating that M1 does not encode the static desired grasp force while carrying the object. Finally, during the release phase classification accuracy remained no better than chance.

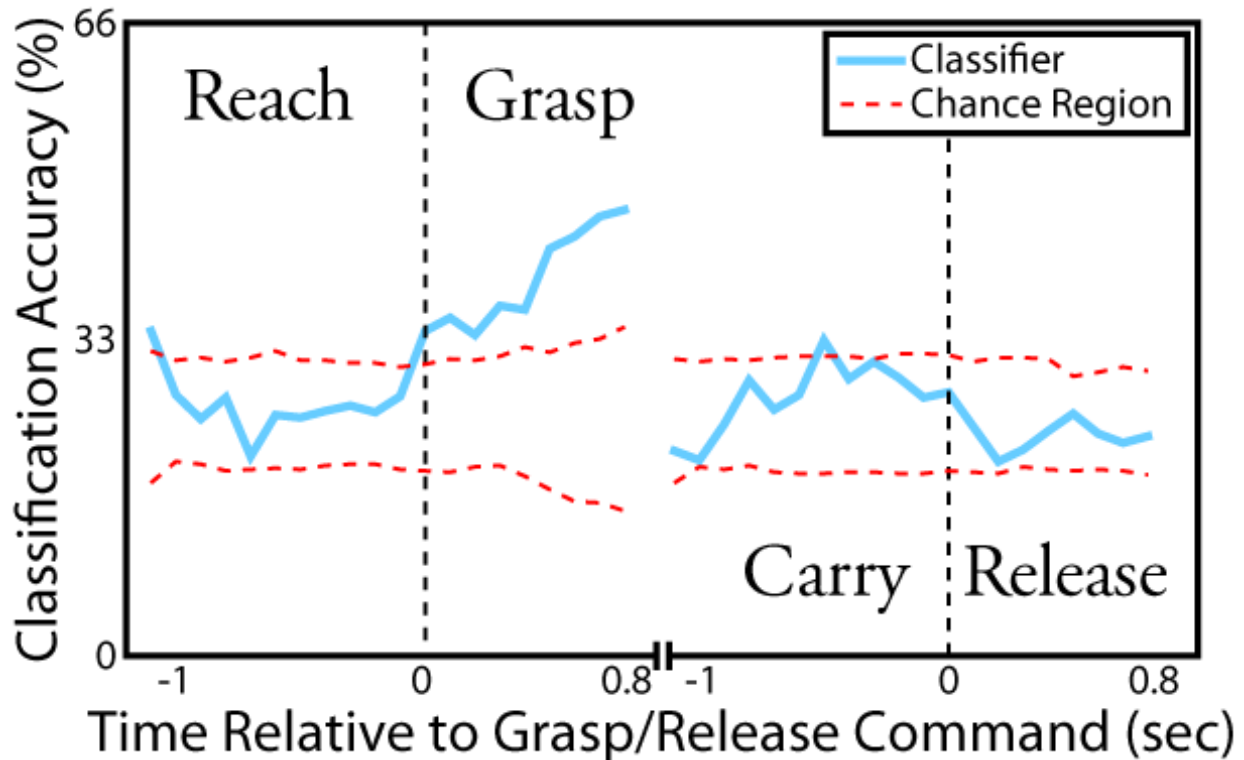


Figure 19. Classification accuracy during the grasp and carry task. The blue line tracks the classification accuracy for the force object. The red dotted lines bound the 95% confidence interval on chance performance. Only during the grasp phase does performance consistently exceed chance level performance.

I then looked for individual units that would enable this type of classification. While grasp velocity and grasp force had similar distributions of tuning accuracy across the recorded population ($p = 0.44$, K-S test), there were a few channels that showed specific tuning to these traits. Figure 20A shows one unit that is uniquely tuned to desired grasp force, with its peri-stimulus time histogram showing strong separation by presented object while the subject attempts to exert force

on the object. Figure 20B shows unit that is tuned to both grasp force and grasp velocity, ramping up its firing rate as the subject attempted to generate positive grasp velocity to close the hand, and then showing separation by trial according to which object was presented only after the hand had made contact with the object.

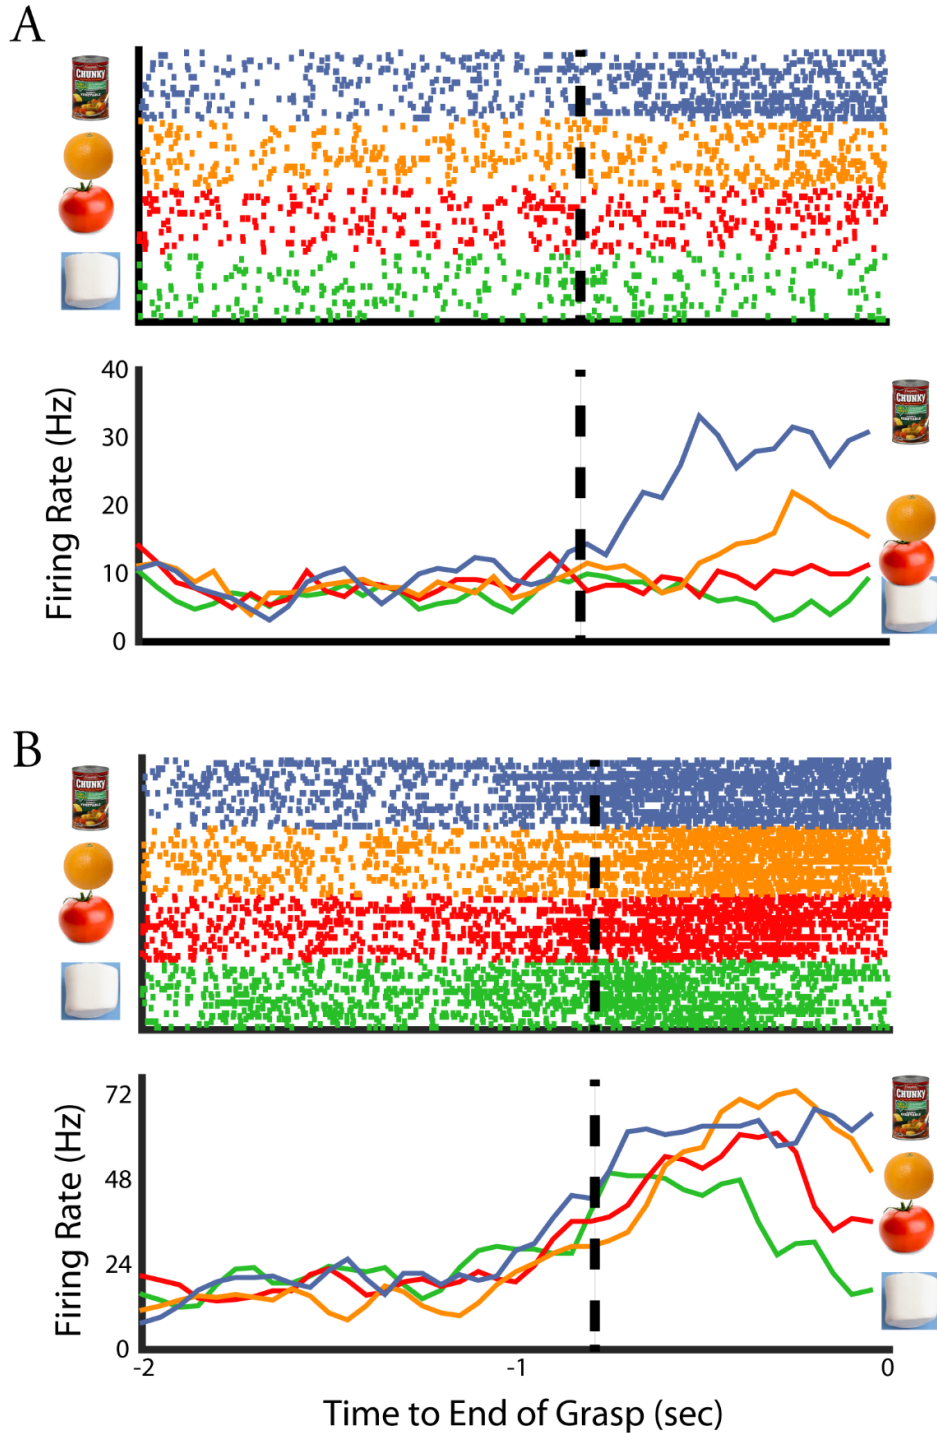


Figure 20. Units encoding grasp force and velocity. Rasters and histograms showing two units' response across 25 trials with each force object. The dashed line marks when the grasp command is issued. A) A force tuned unit shows much higher firing rate when subject grasps with the most force, with even steps down in firing rate including a decrease from baseline firing rate for the lowest force level. B) A velocity and force tuned unit. Firing rate begins increasing as the subject prepares to grasp, regardless of intended force level. Once the hand makes contact with the object, the firing rates stratify based on the desired force level.

6.4 DISCUSSION

This work showed that primary motor cortex encodes information about a BCI user's desired grasp force during active grasping, but not during the exertion of a static grasp force during hand translation. Classification of neural data collected during a task where the subject imagined grasping 4 different objects that required varying amounts of force showed was highly accurate, with errors falling almost entirely to neighboring objects, indicating an organized representation of desired grasp force in M1. When I performed the same analysis while the subject attempted to move a robotic arm to grasp and carry objects using varying grasp forces, the same organized representation was present while grasping the object, but not while carrying it to a new position or releasing the object. This indicates that the force representation in M1 is clear during the execution of the grasp, but not during periods where the force is maintained during transit or during release.

To my knowledge, no studies have reported on neural encoding of force during object transportation. The studies of force encoding from the Ebner laboratory required the animals to maintain the force level for 1.5 seconds (Hendrix et al., 2009; Mason et al., 2002), so it is also possible that the encoding of a sustained force fades in M1. This could occur if the responsibility for maintaining the force is passed along to brainstem or spinal neurons that could react more reflexively than neurons in M1. Reflexive force adjustments would be quicker, likely making them more successful for maintaining force in changing environment and as muscles fatigue.

The encoding of the increasing force, but not the static force is conceptually consistent with the standard approach of decoding movement velocity, as opposed to position (Moran & Schwartz, 1999; Wang et al., 2007). If we define arm position by the sum of the forces exerted within the arm to counteract gravity then hand translation velocity, as encoded in M1, represents the change in these forces. When these forces are constant, and the hand is at a fixed position, the encoding of that position is poor. Similarly, M1 encodes when the force against grasped objects is increasing, but not what it is when it is constant, or instantaneously released. Alternatively, it may be the case that how well grasp force is encoded in M1 is attention modulated, meaning that once the subject began attending more closely to translation than to grasp force it became difficult to identify the desired grasp force.

Using this result, BCI controlled robotic arms will need to increase their exerted force when a positive command is decoded from M1, then leave the force constant when that command is extinguished until the object is released at which point grasp force can be reset to a minimum level. This should provide users with the ability to grab a variety of objects ranging from light, delicate objects to dense, heavy objects, broadening the utility of neuroprosthetic robotic arms.

While this work lays a conceptual groundwork for the implementation of grasp force control by proving that desired grasp force is encoded in M1, there is still significant work required to enact a grasp force decoder. During these tasks, the force levels were assumed to be square waves of predetermined heights, rather than being based on real grasp movements. Grasp force exerted on real objects ramp up to the appropriate force, rather than making an instantaneous step (Johansson & Cole, 1992). Similarly, I calibrated the decoder by assuming that grasp velocity and force were independent of each other, with each controlling the hand and instantaneously switching between the two control modes. However, we know that the two traits of a grasp are

related. For instance, the hand cannot apply a negative, opening, velocity and a positive force at the same time. The BCI system will have to be able to handle contradictory commands in a way that is intuitive to the user. It remains to be seen whether the application of intracortical microstimulation-based force feedback will affect the ability to decode grasp force. We expect that to be effective, microstimulation will likely need to generate intuitive and fairly natural percepts of object contact and stability.

Our lab has begun work to attempt to address each of these issues. To address the discontinuity of assumed grasp force, we are using a virtual environment with a realistic physics engine to gather biomimetic grasp forces. We have attempted a number of heuristic ways to handle the change between controlling grasp velocity and grasp force. Implementing a type of hysteresis, which requires a strong, prolonged releasing grasp velocity, has stabilized the exertion of force on the object. Integrating stimulation to provide sensory information with the grasp force decoding has been challenging. There appear to be interaction effects of some sort, possibly consisting of attempted movements altering the sensation of stimulation rather than stimulation altering the representation of force, consistent with the fact that self-initiated forces are perceived less strongly than forces applied by an external actor (Flanagan et al., 2006). By attempting to integrate all of these factors in to a single system we will be able to make a more useful and intuitive neuroprosthetic arm for interacting with all types of objects a user might encounter in daily life.

7.0 CONCLUSIONS

The BCI controlled neuroprosthetic arms that are currently being developed have the potential to change the lives of people with upper limb paralysis, who currently rely on caregivers for assistance with activities of daily living (ADLs). An intuitive and reliable BCI system could relieve the user of the expense and lack of independence that comes with this level of assistance.

The work here has built on the years of work with non-human primates controlling cursors and then robotic arms (Carmena et al., 2003; Serruya et al., 2002; Taylor et al., 2002; Velliste et al., 2008). Our lab's work has shown that an individual with tetraplegia can control a robotic arm with up to 10 degrees of freedom (DoF): 3 degrees of hand translation, 3 degrees of hand orientation, and 4 degrees of hand shaping (Wodlinger et al., 2014). With this level of control, people with tetraplegia will be able to accomplish a number of activities of daily living (ADLs). Both subjects performed well enough on the ARAT to indicate that a clinically significant improvement in arm function is provided by the BCI (Collinger et al., 2012; Lyle, 1981).

Despite these promising results, there are a number of obstacles to clinical implementation of BCI technology. Researchers and engineers are working to make the recordings wireless (Yin, Borton, Komar, et al., 2014; Yin, Borton, Aceros, Patterson, & Nurmikko, 2014) and improve the longevity of the recordings (Kozai, Jaquins-Gerstl, Vazquez, Michael, & Cui, 2015). Other approaches to these problems would record neural activity without invasive electrodes, through

calcium imaging, functional near-infrared spectroscopy, or ultrasound recordings from piezoelectric “neural dust” (Bertrand et al., 2014; Naseer & Hong, 2015; O’Shea et al., 2017). Some of these improvements in technology are necessary for the devices to be usable in all environments and last long enough to justify implantation.

My dissertation work involved a different type of obstacle. It was known that recordings could be unstable in the short-term, with different neurons being recorded on the same electrode just hours apart (Dickey et al., 2009; Fraser & Schwartz, 2011; Perge et al., 2013; Tolias et al., 2007; Vaidya et al., 2014). I first set out to quantify the rate at which these neural units were turning over (Chapter 3.0). I then investigated how well a unit’s stability could be predicted from its noise level, peak-to-peak voltage, firing rate, firing rate variability, and model-free tuning (Chapter 4.0). One result of these studies was to better understand the potential causes of this instability. Because the peak-to-peak voltage of recorded units was not the strongest source of predictive information about whether a unit would remain stable, it is reasonable to assume that micromotion of the electrode tips relative to the recorded neurons is not the primary driver of instability. These two studies will also help to guide the development of long-term decoders that are less effected by recording instability because of unsupervised recalibration (Bishop et al., 2014; Blumberg et al., 2007; B Jarosiewicz et al., 2014; Li et al., 2011; Shpigelman et al., 2008; Zhang & Chase, 2013) or because they rely mostly on stable units (Ganguly & Carmena, 2009).

In the second half of my dissertation, I continued looking at how the mapping between recordings and desired kinematics was changing, by studying how they changed when the environment changed. I first considered how M1 responded when the BCI controlled hand approaches and object (Chapter 5.0). I then investigated how desired grasp force is encoded in M1 for use once the hand reaches the object (Chapter 6.0). This work will help as BCI controlled

neuroprosthetic arms are used in more realistic scenarios with many different types of objects. The generalizability of control is critical to providing a useful prosthetic.

By quantifying the rates of instability I enabled estimation of how often a decoder will need to be recalibrated to maintain its performance. Using the more conservative iterative-comparison method (Section 3.2.2), I showed that approximately 15% of units were not stable for even an hour, 25% of units were stable on the order of hours, 47.5% of units are stable on the order of days, and 7.5% are stable on the order of weeks. A more liberal estimate using the direct-comparison method (Section 3.2.4) showed that approximately 25% of units were not stable for an hour, 45% of units are stable for hours, 17.5% for days, and 12.5% for weeks.

Using the example from Ganguly & Carmena (Ganguly & Carmena, 2009) that a monkey could complete a cursor task using just 15 recorded units, and assuming ~1 unit per channel (Figure 7A), a 96-channel Blackrock microelectrode array would provide at least 15 stable units for 9 days based on the iterative-comparison method. If we instead use the results from the direct-comparison analysis, we would expect to have at least 15 stable units for 29 days. We would therefore expect that if we could identify these stable units, cursor control could be maintained for 9-29 days with a single electrode array.

This is promising for the ability to control a cursor for long periods of time. In addition to using the cursor to communicate (Gilja et al., 2015; B Jarosiewicz et al., 2014), or use a computer (Pandarinath et al., 2017), this would enable the BCI user to interact with a graphical user interface. This ability would allow them to initiate the recalibration of their higher degree-of-freedom neuroprosthetic arm without technician intervention. They may need to trigger this recalibration because arm control requires more stable units, and will therefore remain stable for shorter periods of time. The units that remain stable for weeks indicate that long-term decoding is possible, but

still requires the ability to identify which units are likely to remain stable and which ones will quickly become unstable.

To determine how predictable a unit's stability is I looked at 5 unit characteristics. These characteristics were baseline noise, peak-to-peak voltage, firing rate, firing rate variability, and model-free tuning (Section 4.2.1) because they provide information about a unit's activity, consistency, clarity, and modulation and are independent of the task being performed. I found that higher firing rates were most predictive of prolonged stability with higher peak-to-peak voltage and model-free tuning values also providing significant improvements in expected length of stability. The relative importance of these three characteristics to predicting a unit's stability indicates that stability is not primarily a function of micromotion, as that would lead to peak-to-peak voltage being the strongest predictor of stability. Instead, perhaps certain types of neurons alternate between active and inactive states, making them appear to be unstable because of their relative silence during the inactive state. This hypothesis requires further testing, in the form of waveform analysis (Bean, 2007), and concurrent recording and imaging of neurons in non-human primates (O'Shea et al., 2017).

This project also showed that it is possible to identify units that will become unstable much less frequently than an average unit. Using the model of probability of unit stability will allow for improved selection of stable units for the purposes of calibrating a decoder.

By combining the results from Chapters 3.0 and 4.0 future work should be able to create decoders that maintain high-performance for a long time by relying on more stable units, while also recalibrating at appropriate time points based on the rate at which units become unstable. These decoders will increase the amount of time a BCI user spends using the device between time

spent calibrating new decoders. This improved longevity will make the technology less intrusive and the user more independent.

These more reliable decoders will also need generalize to a wide variety of activities. Most activities of daily living involve interacting with objects, which has been studied in able-bodied non-human primates (Hendrix et al., 2009; Stefan Schaffelhofer & Scherberger, 2016), but not in BCI users prior to my work presented here. The first problem we encountered in the lab with object interaction was that the arm would, at times, behave erratically in proximity to objects (Wodlinger et al., 2014).

In Chapter 5.0 I found that this erratic movement was due to a sharp increase in firing rates across much of the recorded population in the second preceding contact with an object. This increase in firing rate was not accounted for in our model of neural tuning in M1 (Equation 2.1) from which our decoders are based. To counteract this shift, I introduced a scaling factor that was applied equally to all recorded units so that when the firing rates were decoded the correlated increase in firing rate had already been removed (Section 5.2.4). This increased the object transport rate for one subject by as much as 300% (Section 5.3.3). The scaling was so effective with Subject 1 that it was used for all testing with Subject 2, except for the few sessions reported here that were used to confirm its benefit for Subject 2 as well.

While the scaling made it easier for the subjects to approach and grasp objects, applying the desired amount of force to the object was still not a possibility. In Chapter 6.0 I show the results of the first steps towards decoding desired grasp force from M1 in BCI users. Since Subject 2 has arrays implanted in primary sensory cortex that allow us to stimulate and evoke the sensation of pressure (Flesher et al., 2016) it is important to allow him, and future BCI users, to be able to

control the amount of pressure exerted. This will also allow BCI users to interact with a full spectrum of objects, from light and delicate to heavy and sturdy.

I found that in a task where the subject imagined grasping each of 4 different objects, requiring different levels of grasp force, I could identify from the neural data which object he was attempting to grasp 69% of the time. Furthermore, another 30% of the time, the classifier predicted an object that was adjacent to the correct one on the scale of desired grasp force (Section 6.3.1). This result indicates that neurons in M1 encode desired grasp force along a predictable spectrum.

I then used similar methods with a more enriched task involving the virtual robotic arm (Figure 5B). This work revealed that the desired grasp force was well encoded during grasping, but not clearly available while the hand carried or released the object (Section 6.3.2). This will be important when implementing decoding for grasp force, as a loss of the ability to decode force while carrying an object could lead to consistent dropping of objects if not handled appropriately.

All of this work, and the many other projects I completed during my time as a Ph.D. student that were outside the scope of this document, is important to making BCI controlled neuroprosthetic arms a viable assistive device for people with upper-limb paralysis. Increasing the amount of time the system performs reliably before needing to be recalibrated will be critical to making the device reliable enough for anyone to use. Accounting for the neural response to interaction with objects improves the generalizability of control and enables the user to complete many tasks. Finally, implementing grasp force control will allow users to handle a variety of objects without risk of breaking or dropping them. I look forward to being able to continue the work started here to improve the utility of BCI controlled neuroprosthetic arms in the future.

7.1 LIMITATIONS

My work has taken advantage of a unique opportunity to work with the first BCI users who consistently use a robotic prosthetic arm, providing the opportunity to perform thorough, novel analysis. However, as with any study in a new field, there were a number of limitations that I could not address here. These limitations were the result of limited technology, e.g. an inability to visualize the cells I recorded from, or the amount of time and resources available, e.g. only two subjects could be studied.

One of the limitations in understanding the time course of recording stability was the relatively small percentage of time during which I could make recordings. The best way to understand when and how units change from being recorded to not being recorded would be to identify and track the firing of units 24 hours a day for weeks. I performed the intraday analysis (Section 3.3.3) to provide a better estimate of what full time recording would reveal, but there are still a number of gaps left in our understanding of stability. Continuous recording would provide a better estimate of the exact time course of unit instability and help to smooth the transition from intraday to interday stability from my work. It might also reveal changes in stability while the BCI user is in different states. For instance, M1 activity will be different while the subject is sleeping (Volgushev et al., 2006) but my current estimates of unit stability assume that hours spent sleeping are identical to those spent using the BCI.

Continuous recording would also help me better understand the source of the much higher percentage of units marked stable by the direct comparison (Section 3.2.4). If the difference is due entirely to false negatives when identifying stable units, I would expect to find that units are only marked as unstable for one comparison period at a time, and that those unstable time periods are

evenly distributed along the entire time during which units are recorded. If instead there are networks of neurons that have active and inactive states, so that units that could physically be recorded are sometimes silent, then I might expect to see a number of continuous comparison periods where multiple units are marked as unstable before returning. It is possible that units could become active and inactive on the order of a typical 3-5 minute comparison period allowing the state switching to resemble a false negative, but I would expect that this rate of activity switching would make decoding less reliable on a minute to minute basis than it currently is. This would be an interesting finding, and would support a hypothesis for which it is currently nearly impossible to find evidence in support or opposition.

Currently, long-term continuous recording is impossible because it requires the subject being connected to a large rack of equipment by two cables, and a researcher to be present to monitor everything. With the development of wireless recording technology long-term continuous recordings will soon become possible due to transmitting devices that attach to the pedestal (Yin, Borton, Komar, et al., 2014), and more thoroughly in a few years with implantable wireless transmitters (Yin, Borton, Aceros, et al., 2014).

Another limitation of the stability analysis is that there is not a gold standard for confirming that a unit truly is stable. This was mentioned in the description of the stability analysis method developed by Fraser & Schwartz (2011) (Section 3.2.3). This means that while we believe that if a unit appears much more similar between sessions than a random comparison between sessions it is the same unit, we do not have any other form of proof that this is true. Someone could confirm this belief through in vivo imaging of neuron firing (Grewe, Langer, Kasper, Kampa, & Helmchen, 2010; Hamel, Grewe, Parker, & Schnitzer, 2015). By watching a neuron fire while also recording the action potentials and then running the stability analysis it would be possible to confirm that a

single neuron was responsible for the action potentials that were marked as stable across sessions. This could be done in small mammal subjects currently, and should be developed for use in non-human primates (O'Shea et al., 2017), but is not feasible in human subjects due to the amount of time that the skull would need to be open to provide a viewing window.

A limitation in my ability to interpret the stability results comes from the possible disconnect between a unit's recorded characteristics and its functional tuning. While I expect that units that are labeled as stable in this analysis had consistent tuning characteristics based on previous work (Fraser & Schwartz, 2011), I was unable to confirm that because different tasks were completed on different days. Unit tuning can vary for different tasks (Rasmussen, Schwartz, & Chase, 2017) so these different tasks prevented a study of tuning stability through time. I expect that when BCI users begin to be able to control the system for much of the day using the wireless systems mentioned above they will complete certain tasks multiple times a day, which will allow researchers to study tuning stability thoroughly.

Additionally, decoders typically use unsorted firing rates as opposed to the sorted units studied here (Fraser et al., 2009; Perel et al., 2015; Wodlinger et al., 2014). This could mean that functional channel stability differs slightly from unit stability, however high firing rate, high amplitude units would dominate the activity recorded on channel and are most likely to be stable so I expect that stability of channels would be no worse than that of average units.

For both the object interaction and grasp force studies (Chapters 5.0 & 6.0) it is important to recognize that the results come from a compromised nervous system. Without the ability to record from M1 in an intact subject, it is not possible to know how many of the findings generalize to tuning in M1 outside of BCI arm users. I would expect the result to generalize reasonably well since functional imagining of people with spinal cord injuries attempting to move shows similar

activation to that of control subjects (Kokotilo et al., 2009). Some work with non-human primates, including and expanding upon that from the Scherberger lab (Fluet et al., 2010; Menz et al., 2015; S. Schaffelhofer et al., 2015; Stefan Schaffelhofer & Scherberger, 2016; Townsend et al., 2011), will help to close the knowledge gap, as described in Section 5.4. These studies record from multiple cortical regions while the subjects make reaches to a variety of objects to determine when and where information relevant to the reach is present in the brain. This work is very important to understanding how the brain processes reaches upstream from M1, which is where the successful BCI arm control has recorded from to date (Collinger et al., 2012; Hochberg et al., 2012; Wodlinger et al., 2014). However, non-human primate work is limited by the need to over-train subjects to a task, making it difficult to study how M1 responds to a novel object in the workspace, or to generate an appropriate amount of grip force for the wide variety of objects that can be presented to a person. For ethical reasons it is unlikely that the kind of recordings we make with a BCI system will be repeated in awake, behaving, intact humans until and unless equally high resolution, non-invasive techniques are developed.

Similarly, the fact that I did these studies with only two subjects limits their generalizability. The two subjects also had different causes of their paralysis. In some ways this is encouraging because the recording stability and object interaction results were mostly consistent between the two subjects. As more labs begin to work with human BCI users it will be important to quickly and efficiently share data and protocols to confirm findings with multiple subjects and enable quicker discovery and development.

Finally, this work was all done with the robotic arm in a laboratory or in a virtual reality environment. These very controlled environments are necessary for us to establish our understanding of how M1 represents movement in the context of controlling a robotic arm. Using

the arm in a real-world environment will undoubtedly introduce unexpected changes to the recorded activity, in the same way that introducing objects to the work space did (Chapter 5.0). This is why it is important to begin emulating a natural environment for BCI testing, and transition to testing in the real world soon. This will be made easier by the development of wireless systems (Yin, Borton, Aceros, et al., 2014), and intelligent robotic systems to help adapt to the environment (Downey et al., 2016; K. Muelling et al., 2017). By moving BCI development out of the controlled environment soon, we will avoid “over-fitting” our technology, by creating additional difficulties for use in the real world in order to make minor improvements in performance in the lab.

7.2 FUTURE WORK

While the work in this document takes important steps towards a reliable, generalizable BCI controlled neuroprosthetic arm there is still much work to be done. The work on unit stability will inform the further development of self-recalibrating decoders that is already underway (Bishop et al., 2014; Blumberg et al., 2007; B Jarosiewicz et al., 2014; Li et al., 2011; Shpigelman et al., 2008; Zhang & Chase, 2013). The integration of the ability to recognize units that are more likely to remain stable will help these methods by providing information about which units to rely on when looking for changes in the overall population. These decoders will be increasingly important for the continued development of BCI as wireless systems (Yin, Borton, Komar, et al., 2014) begin to allow testing for more than a handful of hours at a time, and in a variety of contexts. Ultimately they will also enable clinical BCI systems to operate without the regular technical intervention that is still available in the research setting.

The current solution for the problems with object interaction seems to be sufficient, but the ultimate source of the observed increase in firing rate is still open to investigation. Whether the same response is observed in monkeys reaching with their native arms will help determine if this is a BCI specific phenomena. If it is not, it may be that the response is analogous to, or possible concurrent with, muscle co-contraction to stabilize the hand when unpredictable disturbances are possible. Whatever the source of the change in firing rate, it could be a useful signal for a system like the one in Appendix A to trigger intelligent robotic responses to objects.

Our current and planned work to build on the information about grasp force encoding in M1 is discussed in Section 6.4. Briefly, continuously decoding desired grasp force and arbitrating between that and decoded grasp velocity while integrating stimulation for feedback of the enacted grasp force will create a system that allows nuanced grasping of objects. This type of decoder will allow for testing of a variety of activities of daily living that involve delicate objects.

The future work in these areas will require a combination of improved understanding of how primary motor cortex, and possibly other implanted regions, encode desired movements and engineering solutions to a variety of problems. This field of work will continue to evolve as we push the technology closer to its final use case and eventual clinical translation.

APPENDIX

BLENDING OF BRAIN-MACHINE INTERFACE AND VISION-GUIDED

AUTONOMOUS ROBOTICS IMPROVES NEUROPROSTHETIC ARM

PERFORMANCE DURING GRASPING

This paper was published in the Journal of NeuroEngineering and Rehabilitation by: John E. Downey, Jeffrey Weiss, Katharina Muelling, Arun Venkataraman, Jean-Sebastien Valois, Martial Hebert, J. Andrew Bagnell, Andrew B. Schwartz, and Jennifer L. Collinger.

A.1 ABSTRACT

A.1.1 Background

Recent studies have shown that brain-machine interfaces (BMIs) offer great potential for restoring upper limb function. However, grasping objects is a complicated task and the signals extracted

from the brain may not always be capable of driving these movements reliably. Vision-guided robotic assistance is one possible way to improve BMI performance. We describe a method of shared control where the user controls a prosthetic arm using a BMI and receives assistance with positioning the hand when it approaches an object.

A.1.2 Methods

Two human subjects with tetraplegia used a robotic arm to complete object transport tasks with and without shared control. The shared control system was designed to provide a balance between BMI-derived intention and computer assistance. An autonomous robotic grasping system identified and tracked objects and defined stable grasp positions for these objects. The system identified when the user intended to interact with an object based on the BMI-controlled movements of the robotic arm. Using shared control, BMI controlled movements and autonomous grasping commands were blended to ensure secure grasps.

A.1.3 Results

Both subjects were more successful on object transfer tasks when using shared control compared to BMI control alone. Movements made using shared control were more accurate, more efficient, and less difficult. One participant attempted a task with multiple objects and successfully lifted one of two closely spaced objects in 92% of trials, demonstrating the potential for users to accurately execute their intention while using shared control.

A.1.4 Conclusions

Integration of BMI control with vision-guided robotic assistance led to improved performance on object transfer tasks. Providing assistance while maintaining generalizability will make BMI systems more attractive to potential users.

A.2 BACKGROUND

Recent brain-machine interface (BMI) work has shown that people with tetraplegia can control robotic arms using signals recorded by intracortical electrodes (Collinger et al., 2012; Hochberg et al., 2012; Wodlinger et al., 2014). In order to make this technology broadly available to people with upper limb impairment it will need to be reliable under a variety of non-ideal conditions. Intracortical BMIs suffer from limitations that can negatively impact performance including the small number of simultaneously recorded neurons (Collinger et al., 2012; H. K. Kim, Park, & Srinivasan, 2009), the degradation of recorded signal quality over time (Chestek et al., 2011), and intraday changes in the recorded units (Perge et al., 2013). We have recently shown that motor cortex signaling is context-dependent as the extracted signal changes between object grasping and free movement (Wodlinger et al., 2014). If this change is not taken into account, the BMI user has limited ability to control the robotic arm near an object. As with natural reaching, the user must determine how to optimally position the hand to grasp the object for the intended action (Johansson & Flanagan, 2009). Currently, BMIs for arm control do not provide somatosensory feedback for the user (Collinger et al., 2012; Hochberg et al., 2012; Wodlinger et al., 2014), which may impair the normal grasping process (Bensmaia & Miller, 2014). Finally, another potential barrier to

optimal performance is that the visual feedback that a BMI user receives is of a robotic arm rather than their own hand, which may introduce sensory conflicts (Cusack, Patterson, Thach, Kistenberg, & Wheaton, 2014).

Intelligent, vision-guided robotic assistance is one way to improve BMI performance during grasping. Specifically, the system could identify objects in the workspace, define stable grasp positions, and stabilize the hand during grasping (Bagnell et al., 2012; Grest, Woetzel, Koch, & Kiel, 2005; Toshev, Makadia, & Daniilidis, 2009). Previous work towards shared control between a BMI and a computer vision-based system has been limited to specific tasks that were automatically executed by the robot once the user identified a target (Bell, Shenoy, Chalodhorn, & Rao, 2008; McMullen et al., 2014). Other shared control work with intracortical BMI used state switching, such that either the BMI user or the robotic system had control during specific phases (Katyal et al., 2014; Vogel et al., 2015). Instead of this ‘hand off’ from volitional to automatic control, a more ideal system would identify the subject’s intention and optimally blend autonomous assistance with the subject’s volition in a seamless fashion. In order to allow the system to generalize across many tasks, it is important that the volitional contribution remains high at all times. Preservation of BMI control over high-level functions such as gross arm movements and target selection would maintain the user’s agency, while automation could handle low-level functions like hand orientation. A generalizable shared control-based BMI should enable the system to operate robustly over a wide variety of functions in different contexts.

The ability to operate a BMI independently for many activities is considered very important by potential BMI users with tetraplegia (Collinger et al., 2013; Huggins, Moinuddin, Chiodo, & Wren, 2015). Additionally, users of other shared control systems report frustration if it becomes obvious that the system is providing autonomous control, reducing the sense of agency (D. J. Kim

et al., 2012). We propose an alternative approach; the computer provides assistance with low level control tasks such as accurately aligning the hand with an optimal grasp position defined in real-time. The assistance becomes stronger as the system becomes more confident in its identification of the user's intent. This method of shared control allows the user to make unconstrained movements in the workspace, and then provides assistance once the system deciphers the user's intent. In this way, arm control can be made more accurate during tasks that require a high degree of accuracy while still allowing the user to directly control the majority of the movement.

This paper describes a shared control system for grasping that improved the ability of 2 subjects with tetraplegia to transfer objects using a BMI to control a robotic arm. The use of shared control also lowered the perceived difficulty of the task. Analysis of the arm movement kinematics during reaches to an object shows that although shared control led to slower peak movement speeds the resulting trajectories were more stable and efficient. Additionally, we demonstrated the ability of the user to specify their intention through successful completion of a task requiring object selection.

A.3 METHODS

A.3.1 Study design

This study was completed as part of two clinical trials of intracortical BMIs conducted under Investigational Device Exemptions at the University of Pittsburgh (Boninger, n.d.-a, n.d.-b). One subject with tetraplegia from each clinical trial participated in this study.

The primary objectives of the study were to determine the extent to which shared control improves the functional performance of a BMI prosthetic arm. Performance was evaluated using reaching and grasping tasks including a clinical assessment, the Action Research Arm Test (ARAT) (Lyle, 1981), and a task in which the user was instructed to pick up one of two possible objects, called the multiple object task. The order of shared control and unassisted test blocks were randomized each day, and subjects were blinded to order of the blocks.

A.3.2 Subjects:

Both subjects provided informed consent prior to participation in any study procedures, which included implantation of intracortical microelectrode arrays (Blackrock Microsystems, Inc., Salt Lake City, Utah). Informed consent was obtained prior to participation in any study procedures. Subject 1 was a 54-year old female diagnosed with spinocerebellar degeneration without cerebellar involvement (Boninger et al., 2013). Her injury was motor complete at the C4 level, but sensation was generally intact with some hypersensitivity. She had two 4x4 mm, 96-channel arrays implanted in her left primary motor cortex (M1; Figure 21) 31 months prior to the 3 test sessions reported here. Subject 2 was a 28 year old male with a cervical spinal cord injury. Although a complete examination was not performed, we estimate his injury to be C5 motor and C4 sensory AIS B (www.asia-spinalinjury.org). Subject 2 had two 4x4 mm, 88-channel arrays implanted in left somatosensory cortex (S1) and two 4x2.4 mm, 32-channel arrays implanted in left parietal cortex (Figure 21). Subject 2's arrays were placed relative to anatomical landmarks resulting in placement in somatosensory cortex rather than motor cortex. While this likely contributed to differences in performance between subjects, this placement provided the opportunity to study the benefits of shared control for a BMI using sub-optimal neural control. One session was conducted

7 weeks post-implant and the other two were 14 weeks post-implant. These studies were conducted under Investigational Device Exemptions from the Food and Drug Administration and with approval from the University of Pittsburgh Institutional Review Board.




Figure 21. Array location. The approximate location of the microelectrode recording arrays for both subjects on a template brain. Subject 1 had 2 96-channel arrays implanted in M1 (green squares). Subject 2 had 2 88-channel arrays implanted in S1 (yellow squares) and 2 32-channel arrays implanted more posterior (yellow rectangles).

A.3.3 Neural recording

Neural data was acquired with the Neuroport Neural Signal Processor (Blackrock Microsystems). At the beginning of each test session a threshold was set for all recorded channels at -4.5 times the root-mean-square voltage (RMS) for Subject 1 or -5.25 times RMS for Subject 2. Firing rates for

Subject 1 were estimated for each channel by binning the number of recorded threshold crossings every 30 ms (33 Hz update rate). For Subject 2, a 20 ms bin size was used (50 Hz update rate). Firing rates were low-pass filtered using an exponential smoothing function with a 450 ms window for Subject 1, and a 440 ms window for Subject 2. Each channel was considered to be a neural unit, though many channels recorded multi-unit activity.

A.3.4 BMI decoding:

The goal of the experiment was to enable BMI control of a robotic arm (WAM Arm, Barrett Technology, Inc., Newton, MA) during reaching and grasping tasks. The WAM Arm is a 7 degree of freedom robot with a 4 degree of freedom 3-fingered Barrett Hand (Figure 22B). Each day a neural decoder was created to transform neural firing rates into continuous 3 dimensional endpoint translation and 1 dimensional grasp velocity commands for the WAM Arm. The wrist was maintained in a neutral position by computer control during calibration and throughout testing. Custom code was used to compute joint torque commands to execute the decoded endpoint velocity commands.

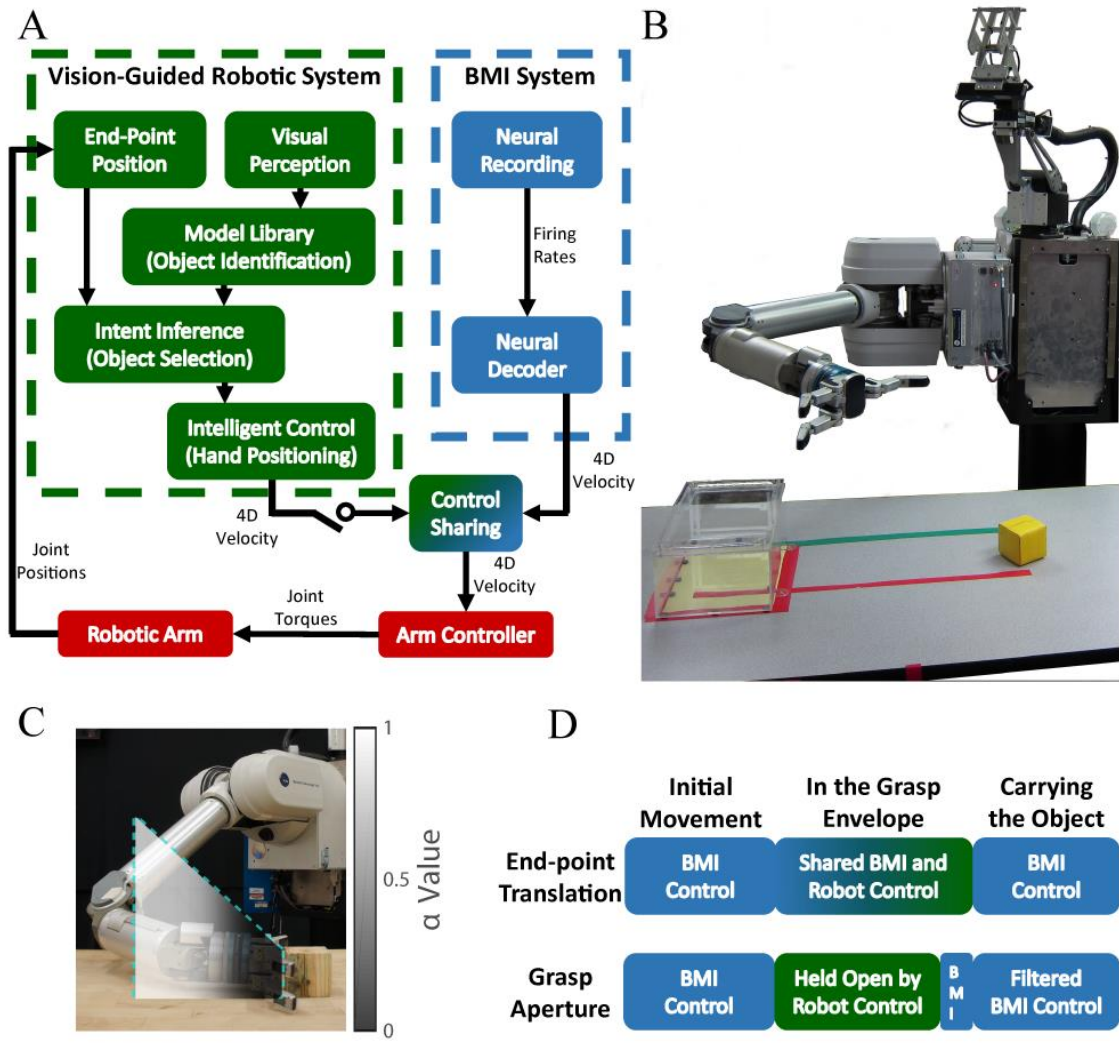


Figure 22. Shared control system diagram and robot testing set up. System diagram for the vision-guided shared control. The blue boxes show the BMI system decoding endpoint translational and grasp velocity. The green boxes show the components of the vision-guided robotic system for grasping. If shared control was not in use, only the output of the BMI system was used to send commands to the arm, but with shared control, the control signal of the vision-guided system was blended with that of the BMI system to create the final robot command. (B) The 7.5 cm cube (yellow) and the target box (clear box) were positioned on the table, as shown, to start the ARAT trials. The subject sat approximately 1 meter to the left of the robot. (C) An example of the central cross-section of the grasp envelope for a stable grasp position on a 7.5 cm cube is outlined by the blue dotted line. The shading shows the gradient of shared control (α value), with white areas being completely controlled by the BMI user and darker areas having more robot control. (D) A trial progression schematic showing when translation and grasp control are under BMI control (blue) or robot control (green). Wrist orientation was always maintained in a neutral posture under computer control.

A two-step calibration method was used to create an optimal linear estimation (OLE) decoder as previously described (Collinger et al., 2012; Wodlinger et al., 2014). The subjects first observed and attempted to control the WAM Arm as the computer commanded it to move. Each trial started with translation to one of 10 possible targets (points in space) at which point the hand grasped or released. Subject 1 completed 40 trials, while Subject 2 completed 60 trials, taking approximately 10 or 15 minutes respectively. Once this step was completed, an OLE decoder was derived based on an encoding model relating recorded firing rates to the computer generated arm and hand kinematics. The encoding model relating unit firing rate to arm kinematics is:

$$f = b_0 + b_x v_x + b_y v_y + b_z v_z + b_g v_g$$

where f is a unit's square root transformed firing rate, v is the kinematic velocity and b is the calculated coefficient relating the velocity and firing rate for each controlled dimension. The dimensions are x , y , and z translation and g grasp aperture. Optimal linear estimation with ridge regression was used to transform calculated b coefficients in to decoding weights that were applied to recorded firing rates to generate kinematics for the arm (Collinger et al., 2012).

The decoder trained in this first step was then used during a second step to give the subjects control of the arm in order to execute the same task as in the first stage of calibration. During this second step the computer restricted decoded brain control commands to those that moved the hand directly towards or away from the specified target during the translation phase of the trial as demonstrated by Velliste et al (Velliste et al., 2008). The arm position was held still while the user issued grasp velocity commands during the grasp phase of each trial. A new OLE decoder was trained using the recorded firing rates and kinematics generated during 40 (Subject 1) or 60

(Subject 2) reach and grasp movements under brain control. The gains for translation and grasp velocity commands were tuned with feedback from the subjects to achieve what they felt was the best balance of speed and accuracy when using the BMI alone (i.e. without shared control). This new decoder was then used for BMI control during the rest of the testing session (Figure 22A, blue blocks).

A.3.5 Vision-guided robotic shared control:

Shared control was provided by utilizing a computer vision system that monitored the work environment with a RGB-D camera mounted above the arm base (Figure 22B). The vision system identified objects by matching depth-image templates from a model library. This library also contains a set of pre-defined hand positions relative to the object that would result in a stable grasp and each is associated with a grasp envelope. The stable grasp positions for each test object were manually determined prior to testing. The grasp envelope was a 25 cm long truncated cone with the small end being the size of the object. The grasp envelope cone opened at an 80 degree angle from the stable grasp position and was oriented along an ideal approach path for a given grasp position (Figure 22D).

During the experiment, the subjects controlled the robotic arm and hand to perform reaching and grasping tasks. Once the subject directed the hand into the envelope of a stable grasp of an object, the shared control system inferred that the user intended to grasp that object (Figure 22A, green blocks). The certainty of this intention was updated for every velocity command from the BMI system (33 Hz for Subject 1, 50 Hz for subject 2). Once the hand was inside the grasp envelope, the system assisted the user in positioning the robot hand to grasp the object by blending the translation commands of the vision-guided robotic system and the BMI system. While the hand

was in the grasp envelope, but not yet at the stable grasp position, the hand was maintained in an open position and user generated grasp velocities that would otherwise open and close the hand were ignored. As the hand moved closer to the stable grasp position and the certainty of intention to grasp increased the robotic system's commands gained more weight, however if the user directed the hand outside the envelope at any point they regained complete control of the arm. Essentially the shared control system regulated the amount of assistance based on the certainty of the BMI user's intention (Figure 22A, Control Sharing). The final translation velocity commands sent to the robot were calculated using a linear blending of user and robot-generated commands:

$$C = (1 - \alpha)R + \alpha B$$

where C is the velocity command sent to the robot, R is the robotic system's velocity command, and B is the BMI system's velocity command derived from the BMI user. The arbitration factor α defines the amount of control given to the user and is computed as:

$$\alpha = (1 + e^{-aD+o})^{-1}$$

where D is the scaled distance between the hand and the optimal position for grasping projected along the central axis of the grasp envelope. D is 0 when the hand is at the stable grasp position and 1 when the hand is at the furthest point in the grasp envelope (Figure 22C). The constants a and o are parameters that are set to ensure α is in the range of $[\alpha_{min}, 1]$. For this study they were manually set to $a = 11.6$ and $o = 6.9$, making $\alpha_{min} = 0.001$. Outside the grasping envelope the user has full control of the robot arm ($\alpha = 1.0$). Once the hand reached the stable grasp position, the

robotic system had nearly complete control of hand position ($\alpha = 0.001$). At the grasping pose the user was required to issue hand closing-velocity commands to grasp the object. The object was grasped with a pre-programmed constant finger torque until a release command was issued.

Once a successful grasp was assured by the system, the user regained unassisted control of endpoint translation while the assist system applied a low-pass filter to the grasp velocity commands, so that transient release commands did not cause the object to drop. The low-pass filter was applied until the hand reached the pre-programmed release area or the subject issued a prolonged hand-opening velocity command to release the object. Figure 22D provides a timeline view of when robotic assistance or BMI control were being responsible for translation and grasp velocities. A more detailed description of the vision-guided robotic shared control system is given in Muelling et al (Katharina Muelling et al., 2015). This previous conference paper focused on the technical details of the system and included high level performance metrics, but no detailed kinematic analysis, for 2 of the 5 ARAT sessions in this paper. Muelling et al. (Katharina Muelling et al., 2015) also described the results of more unstructured tasks that demonstrated the abilities of the system but did not allow for comparison to unassisted control.

A.3.6 Action Research Arm Test

Functional control of the robotic arm was tested using a subset of the Action Research Arm Test (ARAT) that was developed for measuring arm and hand function during recovery from stroke (Lyle, 1981). The ARAT involves moving cubes of various sizes from one location to another. For each ARAT trial, the cube (2.5, 5, 7.5, or 10 cm) was placed on the left side of the table, approximately 40 cm to the left of the hand's starting position which was 30 cm above the table, similar to (Collinger et al., 2012; Wodlinger et al., 2014) (Figure 22B). The subject was instructed

to pick up the cube and place it on top of a box positioned on the right side of the workspace. The subject could regrasp an object if it was dropped or moved. Completion times were recorded for successful trials. A trial was marked as out of bounds if the object was pushed or dropped outside the workspace of the arm. All other trials ended after 2 minutes and were marked as timed out. Subjects were informed that all 3 trials would be counted towards their performance score, in contrast to previous studies which instructed participants that only the best trial would count, allowing for different approach strategies to be tested (Collinger et al., 2012; Wodlinger et al., 2014). Trials were presented in blocks of 3 with the same cube size and assistance conditions (with or without shared control). The subject was not told which condition was being tested and the order of the test conditions was randomized. After each block the subjects were asked to rate the difficulty of the task on a scale of 1 to 10 (1=easiest, 10=hardest task imaginable). Subject 1 completed 3 sessions of ARAT testing, and Subject 2 completed 2 sessions.

A.3.7 Trajectory comparisons

To identify the effect of shared control on the kinematics of the ARAT we analyzed the path lengths of successful trials. Here we consider path lengths for the whole trial, for before the first grasp attempt, and for after the object was successfully grasped for the last time. The first grasp attempt was identified as the initial decrease in aperture to less than 10% of the full range. The path length before this point was labeled as pre-grasp path length. In order to calculate post-grasp path length, we identified the last grasp (10% of the minimum aperture) prior to a successful release of the object. This allowed us to compute a post-grasp path length even when the object was dropped during the transport phase. When only one grasp was required to complete the trial

the path length for the whole trial was equal to the pre-grasp path length plus the post-grasp path length, but in all other cases the sum of the two would be smaller than the whole trial path length.

A.3.8 Multiple object task:

After completing both sessions of the ARAT task, Subject 2 was asked to perform a task that required him to correctly select one of two objects placed on a table. Subject 1 was no longer participating in the study by the time this task was developed. At the beginning of each trial, two 7.5 cm cubes were placed 10 cm apart (from inside edge to inside edge) in one of 3 possible orientations (Fig 3). The subject was told which of the two objects to grasp immediately prior to the trial starting, but the robotic system was not informed of the target. The goal of the task was to select the correct cube and lift it at least 7.5 cm off the table (i.e. above the other cube). This task was in some ways easier than the ARAT since the object only had to be grasped and lifted, but not transported and released at a new target. The subject had 60 seconds to complete the trial and completion times were recorded for successful trials. In addition to timing out, trials were considered failures if the cube was pushed beyond the reach of the arm, or if the wrong cube was lifted. The test session consisted of 24 trials with assistance and 24 trials without assistance, in randomized blocks of 4-5 consecutive trials under one condition. At the end of each block, the subject reported a 1-10 difficulty score as in the ARAT.

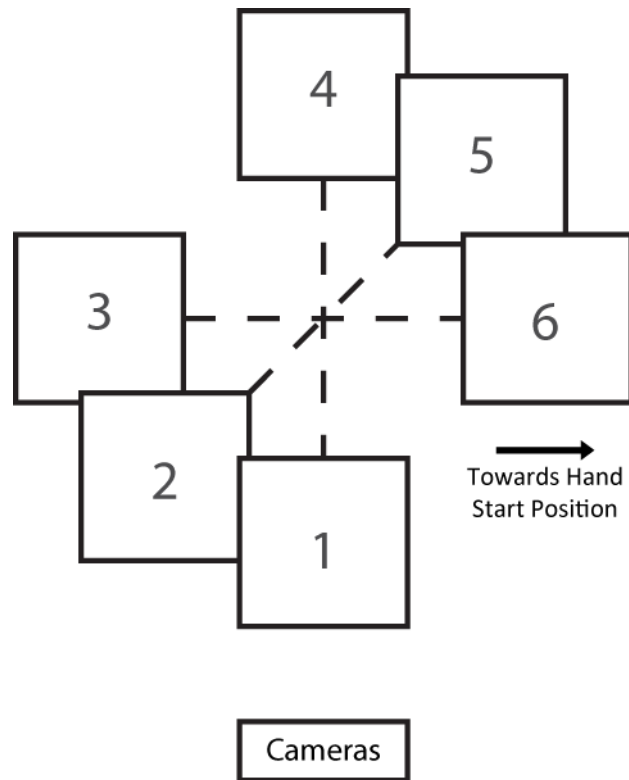


Figure 23. Target positions for the multiple object task. The 7.5 cm target cubes filled the squares in the diagram and were separated by 20 cm. For a single trial, the cubes were placed at 2 positions connected by dashed lines, and the subject was instructed to pick up 1 of the 2 cubes. The position numbers correspond to the target numbers in Table 2. The cube in Figure 2B is at the same point on the table as the intersection of the dashed lines here. The “Cameras” box and hand position arrow indicate the location of those components of the robot at the start of the trial.

A.4 RESULTS

A.4.1 Action Research Arm Test performance comparison

Both subjects performed significantly better on the ARAT tasks with shared control (both subjects: $p < 0.001$, Fisher’s Test). Subject 1 successfully completed the tasks in 78% of the trials with shared control while only succeeding in 22% of trials without. Subject 2 successfully completed

the tasks in 46% of trials with shared control, but failed all unassisted trials. Figure 24A shows the distribution of completion times for successful trials, as well as the percentage of trials that failed by timing out, or moving the object out of bounds. The median completion time for Subject 1 for trials with shared control was 17.5 seconds, while unassisted trials had a median of 31.5 seconds. With only 8 successful unassisted trials, this difference did not reach significance ($p = 0.31$, Wilcoxon rank-sum test). The median completion time for successful shared control trials for Subject 2 was 35 seconds.




Figure 24. ARAT performance and difficulty. (A) The frequency of each trial result for Subject 1 (left) and Subject 2 (right). Completion times are shown for successful trials and the failure mode (time out or out of bounds) is noted for failed trials. Assisted (blue bars) and unassisted (red bars) trials are shown separately. (B) The frequency of each reported difficulty score for assisted and unassisted trial sets (1= extremely easy, 10= extremely difficult). Both subjects were more successful and reported that the task was easier during the trials with shared control.

The purpose of using shared control is to make tasks easier, make it possible to perform a wider range of tasks, and to alleviate user frustration. Participants reported a 1-10 difficulty score after each set of three trials with a single cube size and test condition (with or without shared control) combination. Both subjects had an easier experience during shared control trials. Subjects 1 and 2 reported an average difficulty score of 4.1 and 4.4 for assisted trials and 7.7 and 8.1 respectively

for unassisted trials (both subjects: $p < 0.02$, Wilcoxon signed-rank test; Figure 24; Table 1).

Success rates, completion times, and difficulty ratings for each cube size are listed in Table 3.

Table 3. ARAT performance metrics.

Subject 1 completed 9 trials of each cube size, both with and without shared control assistance. Subject 2 completed 6 trials of each. The time to complete the task is averaged across all successful trials. If there were no successful trials the cell was left blank. Mean difficulty is on a 1 (extremely easy)-10 (extremely difficult) scale.

Sessions	Cube	Success Rate		Mean Completion Time (Sec)		Mean Difficulty	
		w/ Assist	w/o Assist	w/ Assist	w/o Assist	w/ Assist	w/o Assist
Subject 1	10 cm	67%	0%	25	-	3.7	8.3
	7.5 cm	100%	44%	22	23	3.0	6.3
	5 cm	67%	11%	16	14	4.7	7.7
	2.5 cm	78%	33%	48	64	5.0	8.3
	Total	78%	22%	28	37	4.1	7.7
Subject 2	10 cm	0%	0%	-	-	7.0	7.0
	7.5 cm	50%	0%	29	-	4.5	7.5
	5 cm	50%	0%	20	-	4.0	9.0
	2.5 cm	83%	0%	42	-	2.0	9.0
	Total	46%	0%	33	-	4.4	8.1

A.4.2 Trajectory comparisons

To identify how the use of shared control affected the execution of the ARAT task, we computed the speed profile and total path length for each trial. Since Subject 1 was the only subject to complete ARAT trials without assistance, her data forms the primary basis of comparison shown in Figure 25. Endpoint speed for Subject 2's assisted trials is shown for reference. Specifically, we show the distribution of endpoint speeds when the hand was within 10 cm of the table during successful trials for both subjects (Figure 25A). We used this distance criterion to identify hand approach, even if the subject had difficulty grasping, or bumped the object before successfully grasping it. The distributions of endpoint speed show that lower hand speeds are maintained during trials with shared control as compared to unassisted trials.

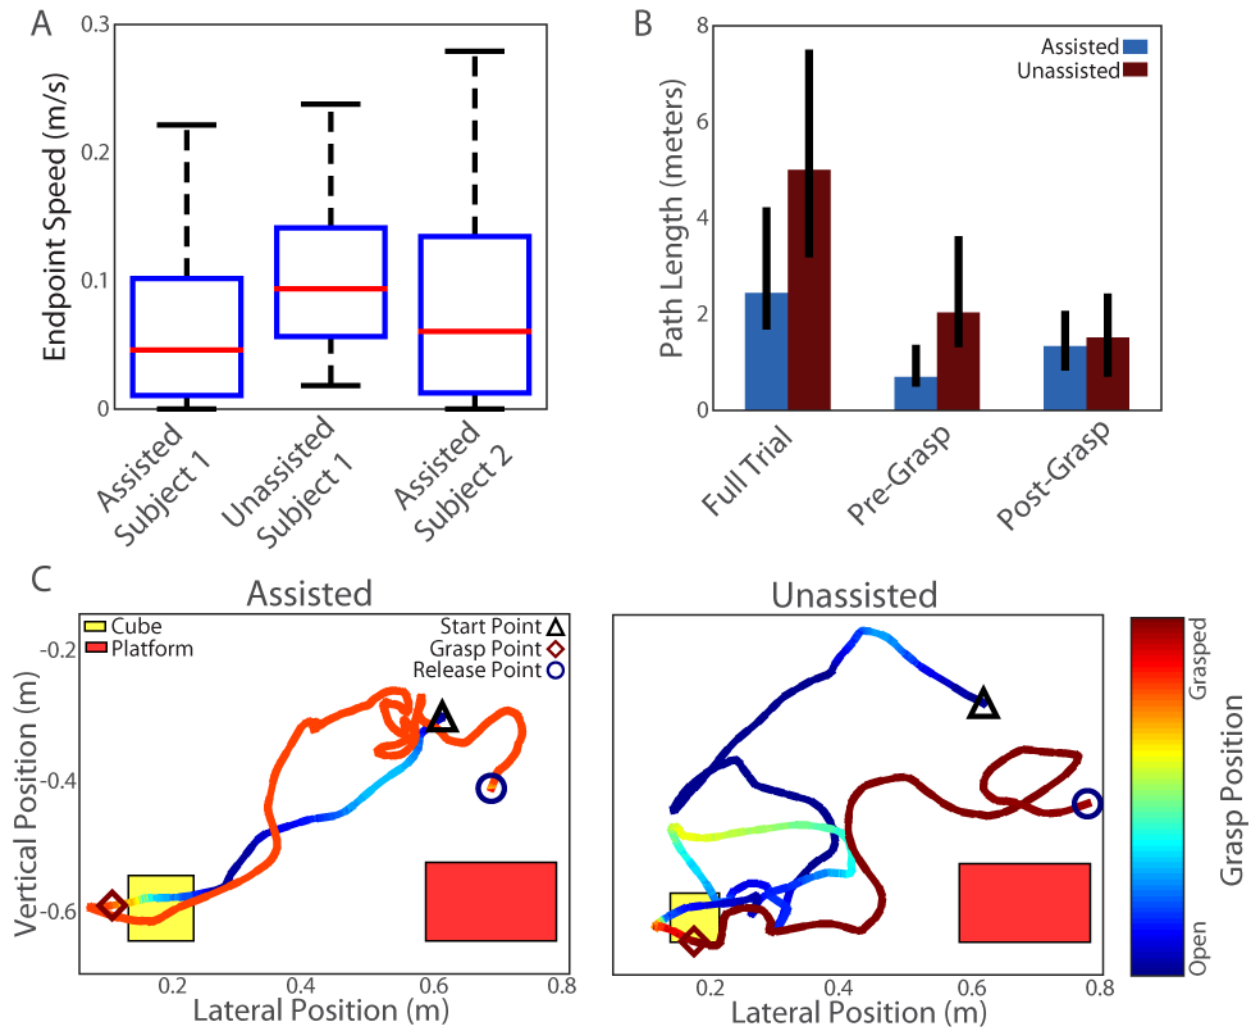


Figure 25. Analysis of trajectory properties with and without shared control. (A) A box plot distribution of hand translation speeds across all time bins while the hand was less than 10 cm above the table during successful trials. The red line is the median speed, the blue box show the interquartile region, and the whiskers span the 5th-95th percentile. The speed distribution for assisted trials for both subjects skews low indicating that the hand was steadier when approaching the object. (B) Subject 1’s path lengths during successful trials, first for the full trials, then separated by the path length before the first grasp attempt and the path length after the object was grasped. The assisted trials benefit the most during the pre-grasp portion of the trial. (C) Subject 1’s hand trajectories with median path lengths for their assistance condition. The color shows the grasp aperture. The release point is marked where the hand opened to allow the object to drop onto the platform. We did not specify to the subject how the object had to be placed, or released, onto the platform. Movie S2 shows both trials.

As shown in Figure 24 and Table 3, Subject 1 tended to have faster trial completion times when using shared control (but did not reach significance). The faster completion times during shared control trials, which also had slower movement speeds, indicate that movements were more direct.

The direct movements are apparent from Subject 1's significantly shorter overall path lengths during trials with shared control (median: 2.44 vs. 5.00 meters, $p = 0.026$ Wilcoxon rank-sum test, Figure 25). While variability in task execution was present between trials, Figure 25C shows the endpoint and grasp trajectories for the trials nearest the median in overall path length to provide a representative comparison of performance between the two conditions.

To identify which phase of the task was most improved by shared control, we examined the path length before and after the first grasp attempt of each trial, defined as when the grasp aperture first reached 10% of its minimum in each trial. Note that the subject missed the block on the first try and had to grasp again on 7% of assisted and 25% of unassisted successful trials. On these trials with missed grasps 2-4 total grasp attempts were made. She also dropped the block after grasping it but before reaching the release point on 21% of assisted and 13% of unassisted successful trials and needed to grasp it again to finish the trial. On trials in which the cube was dropped there were typically only 2 total grasp attempts. In both cases, movements after the first grasp attempt were not included in the calculation of pre-grasp path length. Likewise, movements before the final grasp attempt were not included in the calculation of post-grasp path length. Shared control trials had much shorter path lengths prior to grasp (median: 0.70 vs. 2.04 meters, $p < 0.005$ Wilcoxon rank-sum test), while the difference was less pronounced during the carry phase, which occurred after the object was grasped for the last time (median: 1.34 vs. 1.51 meters, $p = 0.86$ Wilcoxon rank-sum test). The shorter path lengths and slower hand movements prior to grasping indicate that the shared control improved task performance primarily by stabilizing the movement near the object to eliminate unintended or inaccurate movements that could interfere with the user's ability to quickly complete the task.

During successful shared control trials, some level of computer assistance was active for 49% of the trial on average (Subject 1 – 53%, Subject 2 – 38%). This assistance was primarily rendered while the user was attempting to grasp the object, which takes a great deal of precision, making up 34% of the total trial time. After grasping, the subjects regained complete control of translation and only grasp aperture was assisted to aid with carrying of the object. The object was grasped for transport for 15% of the trial time on average including the multiple transport attempts if the object was dropped prematurely. This amount of time is small relative to the time to grasp the object because releasing the object required much less precision than grasping it. The subjects were in complete control of the arm for 51% of the trial on average.

A.4.3 Selection from multiple objects

Subject 2 successfully lifted the correct object in 92% of trials with shared control compared to only 46% of unassisted trials during the multiple object task (Table 4; significantly different with $p < 0.001$, Fisher's test). Only 2 of the 24 trials with shared control were failures, one for lifting the wrong object, and one for moving the target out of the workspace without lifting it off the table. In contrast, 13 of the 24 trials without shared control were failures, 8 because the targeted object was moved outside of the arm's workspace, and 5 because the targeted object was not lifted within the 1 minute time limit.

Table 4. Multiple object task performance metrics.

Target	Success Rate		Median Completion Time (Sec)	
	w/ Assist	w/o Assist	w/ Assist	w/o Assist
1	100%	50%	8.7	19.3
2	100%	50%	8.4	17.8
3	75%	67%	7.7	30.5
4	100%	80%	8.1	20.7
5	100%	25%	8.3	28.7
6	75%	0%	9.6	-
Total	92%	46%	8.3	26.3

In addition to increasing the frequency of successful trials, shared control decreased the median completion time of successful trials from 26.3 seconds to 8.3 seconds (Table 2; $p = 0.03$, Wilcoxon rank-sum test). Consistent with the clear difference in performance, the subject gave the shared control trials an average difficulty rating of 1.4 and the unassisted trials a difficulty rating of 5.6.

A.5 DISCUSSION

In this study, we showed that the combination of a BMI system blended with vision-guided autonomous robotic control, improved the operation of a robotic arm during moderately complex

reach-and-grasp tasks. Once computer vision had identified objects in the workspace, the shared control algorithm used the observed BMI-generated arm trajectory, to infer intention and computes a new command signal with the appropriate contribution from BMI and autonomous control. Previous studies that combined BMI control with a computer vision-based autonomous robotic system relied on visual attention toward a specific object on a computer screen. One study identified targeted objects on a computer display using gaze tracking (McMullen et al., 2014) and another, relied on the EEG P300-evoked response to flashes on the screen near objects (Bell et al., 2008). In both cases, once the object was selected, the task was completed automatically. The user was unable to intervene. Rather than relying on triggered pre-programmed arm movements, our system continuously tracked objects while the user maintained high-level control of the arm from which the shared control system could infer intent. Shared control allowed the system to continuously correct for user errors, such as dropping an object in a new location, and allowed the user to correct system errors, such as incorrectly identifying the desired target. Our objective was to maintain as much volitional control as possible, while providing assistance for the most difficult parts of the task. Once the shared control system detected that the user had selected an object, the hand was stabilized for grasping. Closure of the hand around the object was triggered by the subject. This method left the users in control of the task progression at all times and allowed them to override the assistance by moving the arm away from the object if the robotic system misidentified the user's intent.

There have been two published shared control systems involving intracortical BMI in a human subject (Katyal et al., 2014; Vogel et al., 2015). Katyal et al. presented a system that gave a BMI user translational control until the hand approached a visually identified object at which point entirely autonomous translation and grasping was triggered (Katyal et al., 2014). This system

only dealt with one object, and automatically grasped the identified object based on hand proximity with no attempt to identify intent and did not allow the subject to abort the grasp. In a different study by Vogel et al., the subject controlled 2 degrees of arm translation and a separately decoded discrete state switch command (Vogel et al., 2015). When the state switch command was detected, the system executed a pre-programmed drinking task (i.e., grasp and lift bottle, tilt bottle, or lower and release bottle) and waited for the next state switch command. When attempting to grasp or move the bottle, the user had unassisted control of the 2 degrees of translation. While these systems worked well for specific pre-programmed tasks, they could not generalize to novel settings or adapt to real-world use.

For the ARAT trials, the users maintained complete control of the arm for the majority of the trial. While computer assistance was active for less than half of the ARAT trial on average, it shortened the hand's path length, especially prior to grasping. This demonstrates the significant improvement in the reach-to-grasp movement with shared control. This improvement results in a decreased endpoint velocity when the hand is near the table, indicating that the hand is stable while grasping. Grasping the object while the hand moves quickly might be the ideal result to minimize trial time, however from observation we saw that if the hand moved quickly while near the object it often knocked the object out of position rather than grasping it, making lower velocities near the object more beneficial. We also included a task with multiple objects in order to demonstrate the system's flexibility in a more realistic environment where the user had to choose the correct target. Subject 2 was much more successful on the multiple object tasks with shared control than without. This real-time shared control system accomplished the goals of maintaining user autonomy when possible, limiting user frustration due to undesired movements, and decreasing the overall difficulty of using the system.

Shared control can extend the limits of current BMI technology. One problem with the chronic microelectrode recording technique is the degradation of recorded signal quality over time (Chestek et al., 2011). This was a factor for Subject 1, who had been able to control the robotic arm with up to 10 degrees of freedom (DoF) (Wodlinger et al., 2014). In general, the number of DoF a user can control is limited by the relevant information that can be extracted from a limited sample size of recorded units. Because of the worsening signal quality, 31 months into the study Subject 1 could only control 4 DoF (3D translation and grasp) with variable levels of performance. Shared control compensated for this decline in performance and allowed her to more reliably and easily perform object-transfer tasks with 4 DoF. Similarly, Subject 2 was limited to unreliable 4 DoF control at the time of this study. Despite the limited DoF, both subjects were able to achieve reliable functional control with the shared control system. If control is limited even further than it was for the subjects in our study due to signal quality degradation or limited recordable information, assistance may need to be increased. It may be possible to increase the size of the grasp envelope or to automate grasping completely so that the system is less reliant on the user's unreliable input. However, this would come at the cost of generalizability to novel tasks. The system would need to be customized to provide optimal performance while maintaining a desirable level of independence for the user.

Another general problem is that BMI users currently rely exclusively on visual feedback, which may contribute to ineffective and unstable grasping (Johansson & Flanagan, 2009). The increased stability from shared control resulted in low endpoint speeds when the hand was in the grasp envelope (Figure 25). While slowing down or smoothing the translation commands might create similar stability, this would likely come with a tradeoff of slowing other phases of movements or limiting the ability to make corrective movements. To compensate for the lack of

somatosensation, the shared control system biased the hand aperture toward a closed configuration once an object was grasped, decreasing the likelihood of the object being dropped prematurely. Shared control is one way to maximize the function of a low dimensional control signal and/or limited feedback.

While this prototype shared control system already demonstrated how the addition of autonomous control can enhance BMI performance, a number of advances can be made in the near future. The object library in these sessions was primarily composed of simple geometric objects, but could be expanded to include a large variety of objects, and together with machine-learning algorithms, used to identify graspable portions of larger objects. Muelling et al. (Katharina Muelling et al., 2015) also showed the potential to add pre-programmed actions for a subset of identified objects, such as turning a door knob or pouring from a can, that could provide assistance for higher level tasks. This is similar to the work completed by Vogel et al. (Vogel et al., 2015) with the addition of continuous object identification and shared control. Ideally, this pre-defined information could be replaced by brain-derived signals encoding object identity and intended use. Further, this system could be modified to allow the user to turn the assistance on and off using brain-derived, or external (e.g. switches), control signals. This would allow the user to more easily interact with objects with different intentions, such as pushing an object versus grasping an object.” Additionally, while the release zone for the ARAT task in the current study was specified explicitly in the system, release of an object could be assisted based on proximity to a surface like a shelf or table, allowing generalization to a wider variety of tasks. In future applications, the camera could be mounted to the user’s wheelchair along with the robotic arm to allow for maximum utility and portability.

The work here is a proof-of-concept as an initial step to create a more flexible system to make neuroprosthetic arms more functional for future users. In this study the subjects reported that using the arm was significantly easier with shared control than with BMI alone, and while they recognized when the assistance was active, they never commented that it interfered with their intended actions. Balancing the control between the user and the automated system will be important to provide high performance while ensuring that the user feels that they can use the device reliably in many different situations. As both technologies continue to improve, robotic prosthetic control should become both easier and more useful for the people who need it.

A.6 CONCLUSIONS

The combination of BMI and computer vision-based grasping creates a system that can allow people without use of their arms to control a robotic prosthetic to perform functional tasks in cases where neither technology would be sufficient on its own. The BMI provides the user with high-level control of the pace and goals of the arm movements. The computer vision system helps with the details of the movement, ensuring a secure grasp in the presented cases, but also by identifying how to act on a specific object based on its shape (Katharina Muelling et al., 2015). Balancing the control between the user and the automated system will be important to provide high performance while ensuring that the user feels the device is reliable and responsive to their commands in a variety of situations. As both technologies continue to improve, robotic prosthetic control should become both easier and more useful for the people who need it.

BIBLIOGRAPHY

- Aflalo, T., Kellis, S., Klaes, C., Lee, B., Shi, Y., Pejsa, K., ... Andersen, R. (2015). Decoding motor imagery from the posterior parietal cortex of a tetraplegic human. *Science*, 348(6237), 906–910. <http://doi.org/10.7910/DVN/GJDUTV>
- Ajiboye, A. B., Willett, F. R., Young, D. R., Memberg, W. D., Murphy, B. A., Miller, J. P., ... Kirsch, R. F. (2017). Restoration of reaching and grasping movements through brain-controlled muscle stimulation in a person with tetraplegia: A proof-of-concept demonstration. *The Lancet*, 389(10081), 1821–1830. [http://doi.org/10.1016/S0140-6736\(17\)30601-3](http://doi.org/10.1016/S0140-6736(17)30601-3)
- Anderson, K. D. (2004). Targeting recovery: priorities of the spinal cord-injured population. *Journal of Neurotrauma*, 21(10), 1371–1383. <http://doi.org/10.1089/neu.2004.21.1371>
- Bagnell, J. A., Cavalcanti, F., Cui, L., Galluzzo, T., Hebert, M., Kazemi, M., ... Zhu, R. (2012). An integrated system for autonomous robotics manipulation. *IEEE International Conference on Intelligent Robots and Systems*, 2955–2962. <http://doi.org/10.1109/IROS.2012.6385888>
- Bean, B. P. (2007). The action potential in mammalian central neurons. *Nature Reviews Neuroscience*, 8(6), 451–465. <http://doi.org/10.1038/nrn2148>
- Bell, C. J., Shenoy, P., Chalodhorn, R., & Rao, R. P. N. (2008). Control of a humanoid robot by a noninvasive brain-computer interface in humans. *Journal of Neural Engineering*, 5, 214–220. <http://doi.org/10.1088/1741-2560/5/2/012>
- Bensmaia, S. J., & Miller, L. E. (2014). Restoring sensorimotor function through intracortical interfaces: progress and looming challenges. *Nature Reviews Neuroscience*, 15(5), 313–25. <http://doi.org/10.1038/nrn3724>
- Bertrand, A., Seo, D., Maksimovic, F., Carmena, J. M., Maharbiz, M. M., Alon, E., & Rabaey, J. M. (2014). Beamforming approaches for untethered, ultrasonic neural dust motes for cortical recording: a simulation study. *Conference Proceedings : ... Annual International Conference of the IEEE Engineering in Medicine and Biology Society. IEEE Engineering in Medicine and Biology Society. Annual Conference*, 2625–2628. <http://doi.org/10.1109/EMBC.2014.6944161>
- Bishop, W., Chestek, C. C., Gilja, V., Nuyujukian, P., Foster, J. D., Ryu, S. I., ... Yu, B. M. (2014). Self-recalibrating classifiers for intracortical brain-computer interfaces. *Journal of Neural Engineering*, 11(2), 26001. <http://doi.org/10.1088/1741-2560/11/2/026001>

- Blumberg, J., Rickert, J., Waldert, S., Schulze-Bonhage, A., Aertsen, A., & Mehring, C. (2007). Adaptive classification for brain computer interfaces. *Annual International Conference of the IEEE Engineering in Medicine and Biology - Proceedings*, (November 2014), 2536–2539. <http://doi.org/10.1109/IEMBS.2007.4352845>
- Boninger, M. (n.d.-a). Cortical Recording and Stimulating Array Brain-Machine Interface. Retrieved June 23, 2015, from <https://clinicaltrials.gov/ct2/show/NCT01894802?term=boninger&rank=2>
- Boninger, M. (n.d.-b). Microelectrode Brain-Machine Interface for Individuals With Tetraplegia. Retrieved June 23, 2015, from <https://clinicaltrials.gov/ct2/show/NCT01364480?term=boninger&rank=3>
- Boninger, M., Mitchell, G., Tyler-Kabara, E., Collinger, J., & Schwartz, A. B. (2013). Neuroprosthetic control and tetraplegia--authors' reply. *Lancet*, 381(9881), 1900–1. [http://doi.org/10.1016/S0140-6736\(13\)61154-X](http://doi.org/10.1016/S0140-6736(13)61154-X)
- Bromiley, P. A., & Thacker, N. A. (2001). *The Effects of a Square Root Transform on a Poisson Distributed Quantity. Technical Report.*
- Buzsáki, G., Anastassiou, C. a, & Koch, C. (2012). The origin of extracellular fields and currents--EEG, ECoG, LFP and spikes. *Nature Reviews. Neuroscience*, 13(6), 407–20. <http://doi.org/10.1038/nrn3241>
- Campbell, P. K., Jones, K. E., Huber, R. J., Horch, K. W., & Normann, R. A. (1991). A silicon-based, three-dimensional neural interface: Manufacturing processes for an intracortical electrode array. *IEEE Transactions on Biomedical Engineering*, 38(8), 758–768. <http://doi.org/10.1109/10.83588>
- Carandini, M., & Heeger, D. (2012). Normalization as a canonical neural computation. *Nature Reviews Neuroscience*, (November), 1–12. <http://doi.org/10.1038/nrn3136>
- Carmena, J. M., Lebedev, M. A., Crist, R. E., O'Doherty, J. E., Santucci, D. M., Dimitrov, D. F., ... Nicolelis, M. A. L. (2003). Learning to control a brain-machine interface for reaching and grasping by primates. *PLoS Biology*, 1(2), 193–208. <http://doi.org/10.1371/journal.pbio.0000042>
- Chestek, C. A., Gilja, V., Nuyujukian, P., Foster, J. D., Fan, J. M., Kaufman, M. T., ... Shenoy, K. V. (2011). Long-term stability of neural prosthetic control signals from silicon cortical arrays in rhesus macaque motor cortex. *Journal of Neural Engineering*, 8(4), 45005. <http://doi.org/10.1088/1741-2560/8/4/045005>
- Chorev, E., Epsztein, J., Houweling, A. R., Lee, A. K., & Brecht, M. (2009). Electrophysiological recordings from behaving animals--going beyond spikes. *Current Opinion in Neurobiology*, 19(5), 513–519. <http://doi.org/10.1016/j.conb.2009.08.005>

- Cole, K. J., & Abbs, J. H. (1988). Grip force adjustments evoked by load force perturbations of a grasped object. *Journal of Neurophysiology*, *60*(4), 1513–1522. Retrieved from <http://jn.physiology.org/content/jn/60/4/1513.full.pdf>
- Collinger, J. L., Boninger, M. L., Bruns, T. M., Curley, K., Wang, W., & Weber, D. J. (2013). Functional priorities, assistive technology, and brain-computer interfaces after spinal cord injury. *Journal of Rehabilitation Research & Development*, *50*(2), 145–160. <http://doi.org/10.1016/j.biotechadv.2011.08.021>. Secreted
- Collinger, J. L., Wodlinger, B., Downey, J. E., Wang, W., Tyler-Kabara, E. C., Weber, D. J., ... Schwartz, A. B. (2012). High-performance neuroprosthetic control by an individual with tetraplegia. *The Lancet*, *6736*(12), 1–8. [http://doi.org/10.1016/S0140-6736\(12\)61816-9](http://doi.org/10.1016/S0140-6736(12)61816-9)
- Cox, D. R. (1972). Regression Models and Life-Tables. *Journal of the Royal Statistical Society*, *34*(2), 187–220.
- Csicsvari, J. (2003). Massively Parallel Recording of Unit and Local Field Potentials With Silicon-Based Electrodes. *Journal of Neurophysiology*, *90*(2), 1314–1323. <http://doi.org/10.1152/jn.00116.2003>
- Cusack, W. F., Patterson, R., Thach, S., Kistenberg, R. S., & Wheaton, L. A. (2014). Motor performance benefits of matched limb imitation in prosthesis users. *Experimental Brain Research*, *232*(7), 2143–2154. <http://doi.org/10.1007/s00221-014-3904-2>
- Davidson, P. R., & Wolpert, D. M. (2005). Widespread access to predictive models in the motor system: a short review. *Journal of Neural Engineering*, *2*(3), S313–S319. <http://doi.org/10.1088/1741-2560/2/3/S11>
- Descoins, M., Danion, F., & Bootsma, R. J. (2006). Predictive control of grip force when moving object with an elastic load applied on the arm. *Experimental Brain Research*, *172*(3), 331–342. <http://doi.org/10.1007/s00221-005-0340-3>
- Dickey, A. S., Suminski, A., Amit, Y., & Hatsopoulos, N. G. (2009). Single-unit stability using chronically implanted multielectrode arrays. *Journal of Neurophysiology*, *102*(2), 1331–9. <http://doi.org/10.1152/jn.90920.2008>
- Downey, J. E., Weiss, J. M., Muelling, K., Venkatraman, A., Valois, J.-S., Hebert, M., ... Collinger, J. L. (2016). Blending of brain-machine interface and vision-guided autonomous robotics improves neuroprosthetic arm performance during grasping. *Journal of NeuroEngineering and Rehabilitation*, *13*(1), 28. <http://doi.org/10.1186/s12984-016-0134-9>
- Ethier, C., Oby, E. R., Bauman, M. J., & Miller, L. E. (2012). Restoration of grasp following paralysis through brain-controlled stimulation of muscles. *Nature*, *485*(7398), 368–371. <http://doi.org/10.1038/nature10987>
- Fetz, E. E. (1969). Operant Conditioning of Cortical Unit Activity. *Science*, *163*(3870), 955–958.

- Flanagan, J. R., Bowman, M. C., & Johansson, R. S. (2006). Control strategies in object manipulation tasks. *Current Opinion in Neurobiology*, *16*(6), 650–659. <http://doi.org/10.1016/j.conb.2006.10.005>
- Flesher, S. N., Collinger, J. L., Foldes, S. T., Weiss, J. M., Downey, J. E., Tyler-Kabara, E. C., ... Gaunt, R. A. (2016). Intracortical microstimulation of human somatosensory cortex, 1–11. <http://doi.org/10.1126/scitranslmed.aaf8083>
- Fluet, M.-C., Baumann, M. A., & Scherberger, H. (2010). Context-Specific Grasp Movement Representation in Macaque Ventral Premotor Cortex. *Journal of Neuroscience*, *30*(45), 15175–15184. <http://doi.org/10.1523/JNEUROSCI.3343-10.2010>
- Fraser, G. W., Chase, S. M., Whitford, A., & Schwartz, A. B. (2009). Control of a brain-computer interface without spike sorting. *Journal of Neural Engineering*, *6*(5), 55004. <http://doi.org/10.1088/1741-2560/6/5/055004>
- Fraser, G. W., & Schwartz, A. B. (2011). Recording from the same neurons chronically in motor cortex. *Journal of Neurophysiology*. <http://doi.org/10.1152/jn.01012.2010>
- Gabernet, K. L. (1999). Context-dependent force coding in motor and premotor cortical areas, 123–133.
- Ganguly, K., & Carmena, J. M. (2009). Emergence of a stable cortical map for neuroprosthetic control. *PLoS Biology*, *7*(7), e1000153. <http://doi.org/10.1371/journal.pbio.1000153>
- Ganguly, K., Dimitrov, D. F., Wallis, J. D., & Carmena, J. M. (2011). Reversible large-scale modification of cortical networks during neuroprosthetic control. *Nature Neuroscience*, *14*(5), 662–7. <http://doi.org/10.1038/nn.2797>
- Georgopoulos, A. P., Kettner, R. E., & Schwartz, A. B. (1988). Primate motor cortex and free arm movements to visual targets in three-dimensional space. II. Coding of the direction of movement by a neuronal population. *The Journal of Neuroscience*, *8*(8), 2928–2937. <http://doi.org/10.1089/scd.2011.0674>
- Georgopoulos, A. P., Schwartz, A. B., & Kettner, R. E. (1986). Neuronal Population Coding of Movement Direction. *Science*, *233*(233), 1416–1419.
- Gilja, V., Nuyujukian, P., Chestek, C. a, Cunningham, J. P., Yu, B. M., Fan, J. M., ... Shenoy, K. V. (2012). A high-performance neural prosthesis enabled by control algorithm design. *Nature Neuroscience*, *15*(12), 1752–7. <http://doi.org/10.1038/nn.3265>
- Gilja, V., Pandarinath, C., Blabe, C. H., Nuyujukian, P., Simeral, J. D., Sarma, A. A., ... Henderson, J. M. (2015). Clinical translation of a high-performance neural prosthesis. *Nature Medicine*, *21*(10), 6–8. <http://doi.org/10.1038/nm.3953>
- Gredal, O., Pakkenberg, H., Karlsborg, M., & Pakkenberg, B. (2000). Unchanged total number of neurons in motor cortex and neocortex in amyotrophic lateral sclerosis: A stereological study.

Journal of Neuroscience Methods, 95(2), 171–176. [http://doi.org/10.1016/S0165-0270\(99\)00175-2](http://doi.org/10.1016/S0165-0270(99)00175-2)

- Grest, D., Woetzel, J., Koch, R., & Kiel, C. (2005). Nonlinear Body Pose Estimation from Depth Images, 285–292.
- Grewe, B. F., Langer, D., Kasper, H., Kampa, B. M., & Helmchen, F. (2010). High-speed in vivo calcium imaging reveals neuronal network activity with near-millisecond precision. *Nature Methods*, 7(5), 399–405. <http://doi.org/10.1038/nmeth.1453>
- Hamel, E. J. O., Grewe, B. F., Parker, J. G., & Schnitzer, M. J. (2015). Cellular level brain imaging in behaving mammals: An engineering approach. *Neuron*, 86(1), 140–159. <http://doi.org/10.1016/j.neuron.2015.03.055>
- Hendrix, C. M., Mason, C. R., & Ebner, T. J. (2009). Signaling of grasp dimension and grasp force in dorsal premotor cortex and primary motor cortex neurons during reach to grasp in the monkey. *Journal of Neurophysiology*, 102(1), 132–45. <http://doi.org/10.1152/jn.00016.2009>
- Hochberg, L. R., Bacher, D., Jarosiewicz, B., Masse, N. Y., Simeral, J. D., Vogel, J., ... Donoghue, J. P. (2012). Reach and grasp by people with tetraplegia using a neurally controlled robotic arm. *Nature*, 485(7398), 372–5. <http://doi.org/10.1038/nature11076>
- Hochberg, L. R., Serruya, M. D., Friehs, G. M., Mukand, J. A., Saleh, M., Caplan, A. H., ... Donoghue, J. P. (2006). Neuronal ensemble control of prosthetic devices by a human with tetraplegia. *Nature*, 442(7099), 164–71. <http://doi.org/10.1038/nature04970>
- Hodgkin, A. L., & Huxley, A. F. (1939). Action potentials recorded from inside a nerve fibre. *Nature*, 144(3651), 710–711.
- Hodgkin, A. L., & Huxley, A. F. (1945). Resting and action potentials in single nerve fibres. *The Journal of Physiology*, 104(2), 176–195. <http://doi.org/10.1113/jphysiol.1945.sp004114>
- Hodgkin, A. L., & Huxley, A. F. (1952). Currents carried by sodium and potassium ions through the membrane of the giant axon of *Loligo*. *The Journal of Physiology*, 116(4), 449–472. <http://doi.org/10.1113/jphysiol.1952.sp004717>
- Hodgkin, A. L., Huxley, A. F., & Katz, B. (1952). The effect of sodium ions on the electrical activity of the giant axon of the squid. *The Journal of Physiology*, 108(1), 37–77. <http://doi.org/10.1113/jphysiol.1949.sp004310>
- Homer, M., Perge, J., Black, M., Harrison, M., Cash, S., & Hochberg, L. (2013). Adaptive Offset Correction for Intracortical Brain Computer Interfaces. *IEEE Transactions on Neural Systems and Rehabilitation Engineering: A Publication of the IEEE Engineering in Medicine and Biology Society*, 22(2), 239–248. <http://doi.org/10.1109/TNSRE.2013.2287768>
- Huggins, J. E., Moinuddin, A. A., Chiodo, A. E., & Wren, P. A. (2015). What Would Brain-Computer Interface Users Want: Opinions and Priorities of Potential Users With Spinal Cord

- Injury. *Archives of Physical Medicine and Rehabilitation*, 96(3), S38–S45.e5. <http://doi.org/10.1016/j.apmr.2014.05.028>
- Jarosiewicz, B., Bacher, D., Sarma, A. A., Masse, N. Y., Simeral, J. D., Sorice, B., ... Hochberg, L. R. (2014). Virtual typing by people with tetraplegia using a stabilized, self-calibrating intracortical brain-computer interface. *IEEE BRAIN Grand Challenges Conference, Washington, DC*, 7(313), 1–11. <http://doi.org/10.1126/scitranslmed.aac7328>
- Jarosiewicz, B., Chase, S. M., Fraser, G. W., Velliste, M., Kass, R. E., & 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. <http://doi.org/10.1073/pnas.0808113105>
- Johannes, M. S., Bigelow, J. D., Burck, J. M., Harshbarger, S. D., Kozlowski, M. V., & Van Doren, T. (2011). An overview of the developmental process for the modular prosthetic limb. *Johns Hopkins APL Technical Digest (Applied Physics Laboratory)*, 30(3), 207–216.
- Johansson, R. S., & Cole, K. J. (1992). Sensory-motor coordination during grasping and manipulative actions. *Current Opinion in Neurobiology*, 2(6), 815–823. [http://doi.org/10.1016/0959-4388\(92\)90139-C](http://doi.org/10.1016/0959-4388(92)90139-C)
- Johansson, R. S., & Flanagan, J. R. (2009). Coding and use of tactile signals from the fingertips in object manipulation tasks. *Nature Reviews. Neuroscience*, 10(5), 345–359. <http://doi.org/10.1038/nrn2621>
- Takei, S., Hoffman, D. S., & Strick, P. L. (1999). Muscle and Movement Representations in the Primary Motor Cortex . Muscle and Movement Representations in the Primary Motor Cortex ., 2136(1999), 1–5. <http://doi.org/10.1126/science.285.5436.2136>
- Katyal, K. D., Johannes, M. S., Kellis, S., Aflalo, T., Klaes, C., Mcgee, T. G., ... Mcloughlin, M. P. (2014). A Collaborative BCI Approach to Autonomous Control of a Prosthetic Limb System. In *IEEE International Conference on Systems, Man, and Cybernetics* (pp. 1479–1482).
- Kettner, R. E., Schwartz, A. B., & Georgopoulos, A. P. (1988). Primate motor cortex and free arm movements to visual targets in three-dimensional space. III. Positional gradients and population coding of movement direction from various movement origins. *J.Neurosci.*, 8(8), 2938–2947. <http://doi.org/10.1089/scd.2011.0674>
- Kim, D. J., Hazlett-Knudsen, R., Culver-Godfrey, H., Rucks, G., Cunningham, T., Portée, D., ... Behal, A. (2012). How autonomy impacts performance and satisfaction: Results from a study with spinal cord injured subjects using an assistive robot. *IEEE Transactions on Systems, Man, and Cybernetics: Systems*, 42(1), 2–14. <http://doi.org/10.1109/TSMCA.2011.2159589>
- Kim, H. K., Park, S., & Srinivasan, M. A. (2009). Developments in brain-machine interfaces from the perspective of robotics. *Human Movement Science*, 28(2), 191–203. <http://doi.org/10.1016/j.humov.2008.12.001>

- Kokotilo, K. J., Eng, J. J., & Curt, A. (2009). Reorganization and preservation of motor control of the brain in spinal cord injury: a systematic review. *Journal of Neurotrauma*, 26(11), 2113–2126. <http://doi.org/10.1089/neu.2008.0688.Reorganization>
- Koyama, S., Chase, S. M., Whitford, A. S., Velliste, M., Schwartz, A. B., & Kass, R. E. (2010). Comparison of brain-computer interface decoding algorithms in open-loop and closed-loop control. *Journal of Computational Neuroscience*, 29(1–2), 73–87. <http://doi.org/10.1007/s10827-009-0196-9>
- Kozai, T. D. Y., Jaquins-Gerstl, A. S., Vazquez, A. L., Michael, A. C., & Cui, X. T. (2015). Brain tissue responses to neural implants impact signal sensitivity and intervention strategies. *ACS Chemical Neuroscience*, 6(1), 48–67. <http://doi.org/10.1021/cn500256e>
- Li, Z., Lebedev, M. A., & Nicolelis, M. A. L. (2011). Adaptive Decoding for Brain-Machine Interfaces Through. *Neural Computation*, 3204, 3162–3204.
- Lyle, R. C. (1981). A performance test for assessment of upper limb function in physical rehabilitation treatment and research. *International Journal of Rehabilitation Research*, 4(4), 483–492. <http://doi.org/10.1097/MRR.0b013e3282fc0f93>
- Marin, C., & Fernández, E. (2010). Biocompatibility of intracortical microelectrodes: current status and future prospects. *Frontiers in Neuroengineering*, 3(May), 8. <http://doi.org/10.3389/fneng.2010.00008>
- Mason, C. R., Gomez, J. E., & Ebner, T. J. (2002). Primary motor cortex neuronal discharge during reach-to-grasp: controlling the hand as a unit. *Archives Italiennes de Biologie*, 140, 229–236.
- Mathiowetz, V., Volland, G., Kashman, N., & Weber, K. (1985). Adult Norms for the Box and Block Test of Manual Dexterity. *American Journal of Occupational Therapy*, 39(6), 386–391.
- McMullen, D. P., Hotson, G., Katyal, K. D., Wester, B. A., Fifer, M. S., McGee, T. G., ... Crone, N. E. (2014). Demonstration of a semi-autonomous hybrid brain-machine interface using human intracranial EEG, eye tracking, and computer vision to control a robotic upper limb prosthetic. *IEEE Transactions on Neural Systems and Rehabilitation Engineering*, 22(4), 784–796. <http://doi.org/10.1109/TNSRE.2013.2294685>
- Menz, V. K., Schaffelhofer, S., & Scherberger, H. (2015). Representation of continuous hand and arm movements in macaque areas M1, F5, and AIP: a comparative decoding study. *Journal of Neural Engineering*, 12(5), 56016. <http://doi.org/10.1088/1741-2560/12/5/056016>
- Moran, D. W., & Schwartz, A. B. (1999). Motor cortical representation of speed and direction during reaching. *J Neurophysiol*, 82, 2676–2692.
- Muelling, K., Venkatraman, A., Valois, J.-S., Downey, J. E., Weiss, J., Javdani, S., ... Bagnell, J. A. (2015). Autonomy Infused Teleoperation with Application to BCI Manipulation. In *Robotics: Science and Systems* (Vol. 11). <http://doi.org/10.15607/RSS.2015.XI.039>

- Muelling, K., Venkatraman, A., Valois, J.-S., Downey, J. E., Weiss, J., Javdani, S., ... Bagnell, J. A. (2017). Autonomy infused teleoperation with application to brain computer interface controlled manipulation. *Autonomous Robots*. <http://doi.org/10.1007/s10514-017-9622-4>
- Najafi, K., Wise, K. D., & Mochizuki, T. (1985). A High-Yield IC-Cfompatible Multichannel Array Reco rding. *Electron Devices IEEE Transactions on*, *32*(7), 1206–1211. <http://doi.org/10.1109/T-ED.1985.22102>
- Naseer, N., & Hong, K.-S. (2015). fNIRS-based brain-computer interfaces: a review. *Frontiers in Human Neuroscience*, *9*(January), 1–15. <http://doi.org/10.3389/fnhum.2015.00003>
- Nicolelis, M. A. L., Dimitrov, D., Carmena, J. M., Crist, R., Lehew, G., Kralik, J. D., & Wise, S. P. (2003). Chronic, multisite, multielectrode recordings in macaque monkeys. *Proceedings of the National Academy of Sciences*, *100*(19), 11041–11046. <http://doi.org/10.1073/pnas.1934665100>
- O’Shea, D. J., Trautmann, E., Chandrasekaran, C., Stavisky, S., Kao, J. C., Sahani, M., ... Shenoy, K. V. (2017). The need for calcium imaging in nonhuman primates: New motor neuroscience and brain-machine interfaces. *Experimental Neurology*, *287*, 437–451. <http://doi.org/10.1016/j.expneurol.2016.08.003>
- Oby, E. R., Perel, S., Sadtler, P. T., Ruff, D. A., Mischel, J. L., Montez, D. F., ... Chase, S. M. (2016). Extracellular voltage threshold settings can be tuned for optimal encoding of movement and stimulus parameters. *Journal of Neural Engineering*, *13*(3), 36009. <http://doi.org/10.1088/1741-2560/13/3/036009>
- Pandarath, C., Nuyujukian, P., Blabe, C. H., Sorice, B. L., Saab, J., Willett, F. R., ... Henderson, J. M. (2017). High performance communication by people with paralysis using an intracortical brain-computer interface. *eLife*, *6*, 1–27. <http://doi.org/10.7554/eLife.18554>
- Perel, S., Sadtler, P. T., Oby, E. R., Ryu, S. I., Tyler-Kabara, E. C., Batista, A. P., & Chase, S. (2015). Single-Unit Activity, Threshold Crossings, and Local Field Potentials in Motor Cortex Differentially Encode Reach Kinematics. *Journal of Neurophysiology*, jn.00293.2014. <http://doi.org/10.1152/jn.00293.2014>
- Perge, J. A., Homer, M. L., Malik, W. Q., Cash, S., Eskandar, E., Friehs, G., ... Hochberg, L. R. (2013). Intra-day signal instabilities affect decoding performance in an intracortical neural interface system. *Journal of Neural Engineering*, *10*(3), 36004. <http://doi.org/10.1088/1741-2560/10/3/036004>
- Rasmussen, R. G., Schwartz, A., & Chase, S. M. (2017). Dynamic range adaptation in primary motor cortical populations. *eLife*, *6*, 1–20. <http://doi.org/10.7554/eLife.21409>
- Rouse, A. G., & Schieber, M. H. (2016). Spatiotemporal Distribution of Location and Object Effects in Primary Motor Cortex Neurons during Reach-to-Grasp. *Journal of Neuroscience*, *36*(41), 10640–10653. <http://doi.org/10.1523/JNEUROSCI.1716-16.2016>

- Sachs, N. A., Ruiz-Torres, R., Perreault, E. J., & Miller, L. E. (2016). Brain-state classification and a dual-state decoder dramatically improve the control of cursor movement through a brain-machine interface. *Journal of Neural Engineering*, *13*, 12.
- Sadtler, P. T., Quick, K. M., Golub, M. D., Chase, S. M., Ryu, S. I., Tyler-Kabara, E. C., ... Batista, A. P. (2014). Neural constraints on learning. *Nature*, *512*(7515), 423–426. <http://doi.org/10.1038/nature13665>
- Santhanam, G., Ryu, S. I., Yu, B. M., Afshar, A., & Shenoy, K. V. (2006). A high-performance brain-computer interface. *Nature*, *442*(7099), 195–8. <http://doi.org/10.1038/nature04968>
- Schaffelhofer, S., Agudelo-Toro, A., & Scherberger, H. (2015). Decoding a Wide Range of Hand Configurations from Macaque Motor, Premotor, and Parietal Cortices. *Journal of Neuroscience*, *35*(3), 1068–1081. <http://doi.org/10.1523/JNEUROSCI.3594-14.2015>
- Schaffelhofer, S., & Scherberger, H. (2016). Object vision to hand action in macaque parietal, premotor, and motor cortices. *eLife*, *5*(2016JULY), 1–24. <http://doi.org/10.7554/eLife.15278>
- Schmidt, E. M., Bak, M. J., & McIntosh, J. S. (1976). Long-term chronic recording from cortical neurons. *Experimental Neurology*, *52*(3), 496–506.
- Schwartz, A. B., Kettner, R. E., & Georgopoulos, A. P. (1988). Primate motor cortex and free arm movements to visual targets in three-dimensional space. I. Relations between single cell discharge and direction of movement. *J.Neurosci.*, *8*(August), 2913–2927.
- Seidler, R. D., Bernard, J. A., Burutolu, T. B., Fling, B. W., Gordon, M. T., Gwin, J. T., ... Lipps, D. B. (2010). Motor control and aging: Links to age-related brain structural, functional, and biochemical effects. *Neuroscience and Biobehavioral Reviews*, *34*(5), 721–733. <http://doi.org/10.1016/j.neubiorev.2009.10.005>
- Serruya, M. D., Hatsopoulos, N. G., Paninski, L., Fellows, M. R., & Donoghue, J. P. (2002). Instant neural control of a movement signal. *Nature*, *416*(6877), 141–2. <http://doi.org/10.1038/416141a>
- Shanechi, M. M., Orsborn, A. L., & Carmena, J. M. (2016). Robust Brain-Machine Interface Design Using Optimal Feedback Control Modeling and Adaptive Point Process Filtering. *PLoS Computational Biology*, *12*(4), 1–29. <http://doi.org/10.1371/journal.pcbi.1004730>
- Sherwood, C. C., & Hof, P. R. (2010). The evolution of neuron types and cortical histology in apes and humans. *Evolution of Nervous Systems*, *4*, 355–378. <http://doi.org/10.1016/B0-12-370878-8/00022-7>
- Shpigelman, L., Lalazar, H., & Vaadia, E. (2008). Kernel-ARMA for Hand Tracking and Brain-Machine Interfacing During 3D Motor Control. *Neural Information Processing Systems*, 1489–1496. <http://doi.org/citeulike-article-id:6002127>
- Simeral, J. D., Kim, S.-P., Black, M. J., Donoghue, J. P., & Hochberg, L. R. (2011). Neural control of cursor trajectory and click by a human with tetraplegia 1000 days after implant of an

- intracortical microelectrode array. *Journal of Neural Engineering*, 8(2), 25027. <http://doi.org/10.1088/1741-2560/8/2/025027>
- Suner, S., Fellows, M. R., Vargas-Irwin, C., Nakata, G. K., & Donoghue, J. P. (2005). Reliability of signals from a chronically implanted, silicon-based electrode array in non-human primate primary motor cortex. *IEEE Transactions on Neural Systems and Rehabilitation Engineering : A Publication of the IEEE Engineering in Medicine and Biology Society*, 13(4), 524–41. <http://doi.org/10.1109/TNSRE.2005.857687>
- Taylor, D. M., Helms Tillery, S. I., & Schwartz, A. B. (2002). Direct Cortical Control of 3D Neuroprosthetic Devices. *Science*, 296(5574), 1829–1832. <http://doi.org/10.1126/science.1070291>
- The National SCI Statistical Center. (2015). Spinal Cord Injury (SCI) Facts and Figures at a Glance. Retrieved January 1, 2015, from <https://www.nscisc.uab.edu/Public/Facts2015.pdf>
- Tolias, A. S., Ecker, A. S., Siapas, A. G., Hoenselaar, A., Keliris, G. A., & Logothetis, N. K. (2007). Recording chronically from the same neurons in awake, behaving primates. *Journal of Neurophysiology*, 98(6), 3780–90. <http://doi.org/10.1152/jn.00260.2007>
- Toshev, A., Makadia, A., & Daniilidis, K. (2009). Shape-based object recognition in videos using 3D synthetic object models. *2009 IEEE Computer Society Conference on Computer Vision and Pattern Recognition Workshops, CVPR Workshops 2009*, 288–295. <http://doi.org/10.1109/CVPRW.2009.5206803>
- Townsend, B. R., Subasi, E., & Scherberger, H. (2011). Grasp Movement Decoding from Premotor and Parietal Cortex. *Journal of Neuroscience*, 31(40), 14386–14398. <http://doi.org/10.1523/JNEUROSCI.2451-11.2011>
- Vaidya, M., Dickey, A., Best, M. D., Coles, J., Balasubramanian, K., Suminski, A. J., & Hatsopoulos, N. G. (2014). Ultra-long term stability of single units using chronically implanted multielectrode arrays. *2014 36th Annual International Conference of the IEEE Engineering in Medicine and Biology Society, EMBC 2014*, 4872–4875. <http://doi.org/10.1109/EMBC.2014.6944715>
- Velliste, M., Perel, S., Spalding, M. C., Whitford, A. S., & Schwartz, A. B. (2008). Cortical control of a prosthetic arm for self-feeding. *Nature*, 453(7198), 1098–101. <http://doi.org/10.1038/nature06996>
- Vogel, J., Haddadin, S., Jarosiewicz, B., Simeral, J. D., Bacher, D., Hochberg, L. R., ... van der Smagt, P. (2015). An assistive decision-and-control architecture for force-sensitive hand-arm systems driven by human-machine interfaces. *The International Journal of Robotics Research*, 1–18. <http://doi.org/10.1177/0278364914561535>
- Volgushev, M., Chauvette, S., Mukovski, M., & Timofeev, I. (2006). Precise long-range synchronization of activity and silence in neocortical neurons during slow-wave oscillations [corrected]. *The Journal of Neuroscience: The Official Journal of the Society for Neuroscience*, 26(21), 5665–72. <http://doi.org/10.1523/JNEUROSCI.0279-06.2006>

- Wang, W., Chan, S. S., Heldman, D. A., & Moran, D. W. (2007). Motor cortical representation of position and velocity during reaching. *Journal of Neurophysiology*, 97(6), 4258–4270. <http://doi.org/10.1152/jn.01180.2006>.
- Wang, W., Chan, S. S., Heldman, D. a, & Moran, D. W. (2010). Motor cortical representation of hand translation and rotation during reaching. *The Journal of Neuroscience : The Official Journal of the Society for Neuroscience*, 30(3), 958–62. <http://doi.org/10.1523/JNEUROSCI.3742-09.2010>
- Wang, W., Collinger, J. L., Degenhart, A. D., Tyler-Kabara, E. C., Schwartz, A. B., Moran, D. W., ... Boninger, M. L. (2013). An electrocorticographic brain interface in an individual with tetraplegia. *PloS One*, 8(2), e55344. <http://doi.org/10.1371/journal.pone.0055344>
- Williamson, R. C., Cowley, B. R., Litwin-Kumar, A., Doiron, B., Kohn, A., Smith, M. A., & Yu, B. M. (2016). Scaling Properties of Dimensionality Reduction for Neural Populations and Network Models. *PLoS Computational Biology*, 12(12), e1005141. <http://doi.org/10.1371/journal.pcbi.1005141>
- Wodlinger, B., Downey, J. E., Tyler-Kabara, E. C., Schwartz, A. B., Boninger, M. L., & Collinger, J. L. (2014). Ten-dimensional anthropomorphic arm control in a human brain-machine interface: difficulties, solutions, and limitations. *Journal of Neural Engineering*, 12(1), 16011. <http://doi.org/10.1088/1741-2560/12/1/016011>
- Yin, M., Borton, D. A., Aceros, J., Patterson, W. R., & Nurmikko, A. V. (2014). A 100-Channel Hermetically Sealed Implantable Device for Chronic Wireless Neurosensing Applications. *IEEE Trans. Biomed. Eng.*, 7(2), 115–128. <http://doi.org/10.1109/TBCAS.2013.2255874.A>
- Yin, M., Borton, D. A., Komar, J., Agha, N., Lu, Y., Li, H., ... Nurmikko, A. V. (2014). Wireless neurosensor for full-spectrum electrophysiology recordings during free behavior. *Neuron*, 84(6), 1170–1182. <http://doi.org/10.1016/j.neuron.2014.11.010>
- Zhang, Y., & Chase, S. M. (2013). A stabilized dual Kalman filter for adaptive tracking of brain-computer interface decoding parameters. *Conference Proceedings : ... Annual International Conference of the IEEE Engineering in Medicine and Biology Society. IEEE Engineering in Medicine and Biology Society. Annual Conference, 2013*, 7100–7103. <http://doi.org/10.1109/EMBC.2013.6611194>