EVALUATION AND ADVANCEMENT OF ELECTROCORTICOGRAPHIC BRAIN-MACHINE INTERFACES FOR INDIVIDUALS WITH UPPER-LIMB PARALYSIS

by

Alan D. Degenhart

B.S., The Ohio State University, 2005

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
2014

UNIVERSITY OF PITTSBURGH

SWANSON SCHOOL OF ENGINEERING

This dissertation was presented

by

Alan D. Degenhart

It was defended on

October 30, 2014

and approved by

Wei Wang, MD, PhD, Department of Physical Medicine and Rehabilitation, University of Pittsburgh

Elizabeth Tyler-Kabara, MD, PhD, Department of Neurological Surgery, University of Pittsburgh

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

Byron Yu, PhD, Department of Electrical and Computer Engineering, Carnegie Mellon University

Dissertation Director: Wei Wang, MD, PhD, Department of Physical Medicine and Rehabilitation, University of Pittsburgh

Dissertation Co-Director: Elizabeth Tyler-Kabara, MD, PhD, Department of Neurological Surgery, University of Pittsburgh

EVALUATION AND ADVANCEMENT OF ELECTROCORTICOGRAPHIC BRAIN-MACHINE INTERFACES FOR INDIVIDUALS WITH UPPER-LIMB PARALYSIS

Alan D. Degenhart, PhD

University of Pittsburgh, 2014

Brain-machine interface (BMI) technology aims to provide individuals with movement paralysis a natural and intuitive means for the restoration of function. Electrocorticography (ECoG), in which disc electrodes are placed on either the surface of the dura or the cortex to record field potential activity, has been proposed as a viable neural recording modality for BMI systems, potentially providing stable, long-term recordings of cortical activity with high spatial and temporal resolution. Previous demonstrations of BMI control using ECoG have consisted of short-term periods of control by able-bodied subjects utilizing basic processing and decoding techniques. This dissertation presents work seeking to advance the current state of ECoG BMIs through an assessment of the ability of individuals with movement paralysis to control an ECoG BMI, an investigation into adaptation during BMI skill acquisition, an evaluation of chronic implantation of an ECoG electrode grid, and improved extraction of BMI command signals from ECoG recordings.

Two individuals with upper-limb paralysis were implanted with high-density ECoG electrode grids over sensorimotor cortical areas for up to 30 days, with both subjects found to be capable of voluntarily modulating their cortical activity to control movement of a computer cursor with up to three degrees of freedom. Analysis of control signal angular error and the tuning characteristics of ECoG spectral features during the acquisition of brain control revealed that both decoder calibration and fixed-decoder training could facilitate performance improvements. In addition, to better understand the capability of ECoG to provide robust, long-term recordings, work was conducted assessing the effects of chronic implantation of an ECoG electrode grid in a non-human primate, demonstrating that movement-related modulation could be recorded from electrode nearly two years post-implantation despite the presence of substantial fibrotic encapsulation. Finally, it was found that the extraction of command signals from ECoG recordings could be improved through the use of a decoding method incorporating weight-space priors accounting for the expected correlation structure of electrical field potentials. Combined, this work both demonstrates the feasibility of ECoG-based BMI systems as well as addresses some of key challenges that must be overcome before such systems are translated to the clinical realm.

TABLE OF CONTENTS

LIST OF TABLES

LIST OF FIGURES

PREFACE

This dissertation contains the culmination of over 7 years of work. As with many undertakings of such scope, it could not have been accomplished without the help of a great number of people. I would like to thank my advisors, Drs. Wei Wang and Elizabeth Tyler-Kabara, for their mentorship throughout my graduate training and for giving me the opportunity to work on such an exciting project, as well as Drs. Douglas Weber and Aaron Batista, who have served as unofficial mentors through the years and have provided me with invaluable advice on my professional career. I would also like to thank the past and present members of the Human Rehabilitation and Neural Engineering Laboratory (hRNEL), the Rehabilitation and Neural Engineering Laboratory (RNEL), Sensory Motor Integration Laboratory and Engineering (SMILE), SCABBY journal club, and the subjects contributing to the data presented in this dissertation.

Finally, this work would have not been possible without the continual support and understanding of my family: my wife Jenelle, my parents Doug and Wendy, my sister Lindsey, and Watson. This work is dedicated to them.

1.0 INTRODUCTION

Individuals with movement disorders due to spinal cord injury (SCI) and amyotrophic lateral sclerosis (AlS) are often left with severely limited ability to interact with their environment. For these individuals, a number of devices exist which seek to restore functionality, including myoelectric-controlled protheses and eye and head tracking systems for computer control. Unfortunately, these devices are limited by both their reliability and capability, as they typically provide only limited functional restoration while requiring a significant amount of maintenance effort on the part of either users or qualified technicians.

Brain-machine interfaces (BMIs) seek to alleviate these shortcomings by attempting to infer assistive device commands from signals recorded from the cerebral cortex. By inferring user intent from cortical activity, BMIs can theoretically provide high-fidelity command signals for assistive devices and a more naturalistic and intuitive experience for the user. Advances in neural recording technology, the development of anthropomorphic prosthetic devices, and increased computational power of computing platforms have made clinical BMI systems increasingly viable, enabling a number of studies evaluating BMI technology in end-user populations.

This dissertation presents work seeking to advance the current state of electrocorticographic (ECoG) brain-machine interfaces for continuous and proportional control of assistive devices. It is organized in the following manner: Chapter 2 provides a background on brain-machine interface systems, reviewing the components common to most BMIs and motivating the following work. Chapter 3 introduces electrocorticography and its applicability for BMIs. Chapter 4 presents results of a study evaluating an ECoG BMI system for individuals with upper-limb paralysis. Chapter 5 investigates the cortical and computational processes underlying acquisition of control of an ECoG BMI system. Chapter 6 presents the results investigating the effects of chronic implantation of a ECoG electrode grid in a non-human primate. Chapter 7 describes a novel decoding algorithm for the extraction of BMI commands from ECoG signals. Finally, Chapter 8 summarizes the contents of this dissertation, outlines directions for future research, and discusses the future of the BMI field.

2.0 BRAIN-MACHINE INTERFACES

A brain-machine interface is a complicated system involving multiple components which work in concert to translate a user's intent into action. Figure 2.1 illustrates the components of a typical BMI system. The user generates cortical activity which reflects their intent, which is recorded using one of several *neural recording modalities*. Task-relevant activity is then extracted from raw neural signals using a *feature extraction* process, the output from which is then passed through a *decoder* which converts them into command signals for an assistive device. These components can be summarized as follows:

Neural recording modality and control strategy. The neural recording modality is the interface used to record cortical activity. In most cases this process involves recording a time-varying voltage signal using some type of electrode. Cortical activity reflects the BMI *control strategy*, which determines the relationship between user intent and cortical activity.

Feature extraction and decoding. Raw signals recorded from the brain must be converted into control commands for an assistive device. This is a two-stage process: neural signals are first converted into a set of features which contain information about the intent of the user this process is commonly referred to as *feature extraction*. These features are then translated into device command signals using a *neural decoder*. Though the particulars of this process varies with decoder type, in general a decoder consists of a set of equations which translate extracted features into command signals.

Assistive device. The device which is controlled by the user. Though the type of assistive device is dictated by the needs of the user, most current BMI systems involve *continuous* and *proportional* control of computer cursors or prosthetic limbs[1].

Numerous choices in neural recording modalities, extraction algorithms, and assistive devices are available for for use in BMIs. Researchers developing BMI systems must therefore balance the relative merits of these options to design a system that appropriately meets the needs of a target clinical population. In order to do this, an in-depth understanding of the advantages and disadvantages of the available design choices is required; the following sections describe the components of a BMI in more detail and discuss the currently-available methods and technologies for each.

---

[1] Though a great deal of research has gone into the development of BMI-based communication aids (e.g. P300 spellers [Donchin et al., 2000]), the work contained in this dissertation focuses on the development BMIs for restoration of motor function for individuals with movement paralysis. As such, BMI communication aids and discrete systems will not be discussed further.

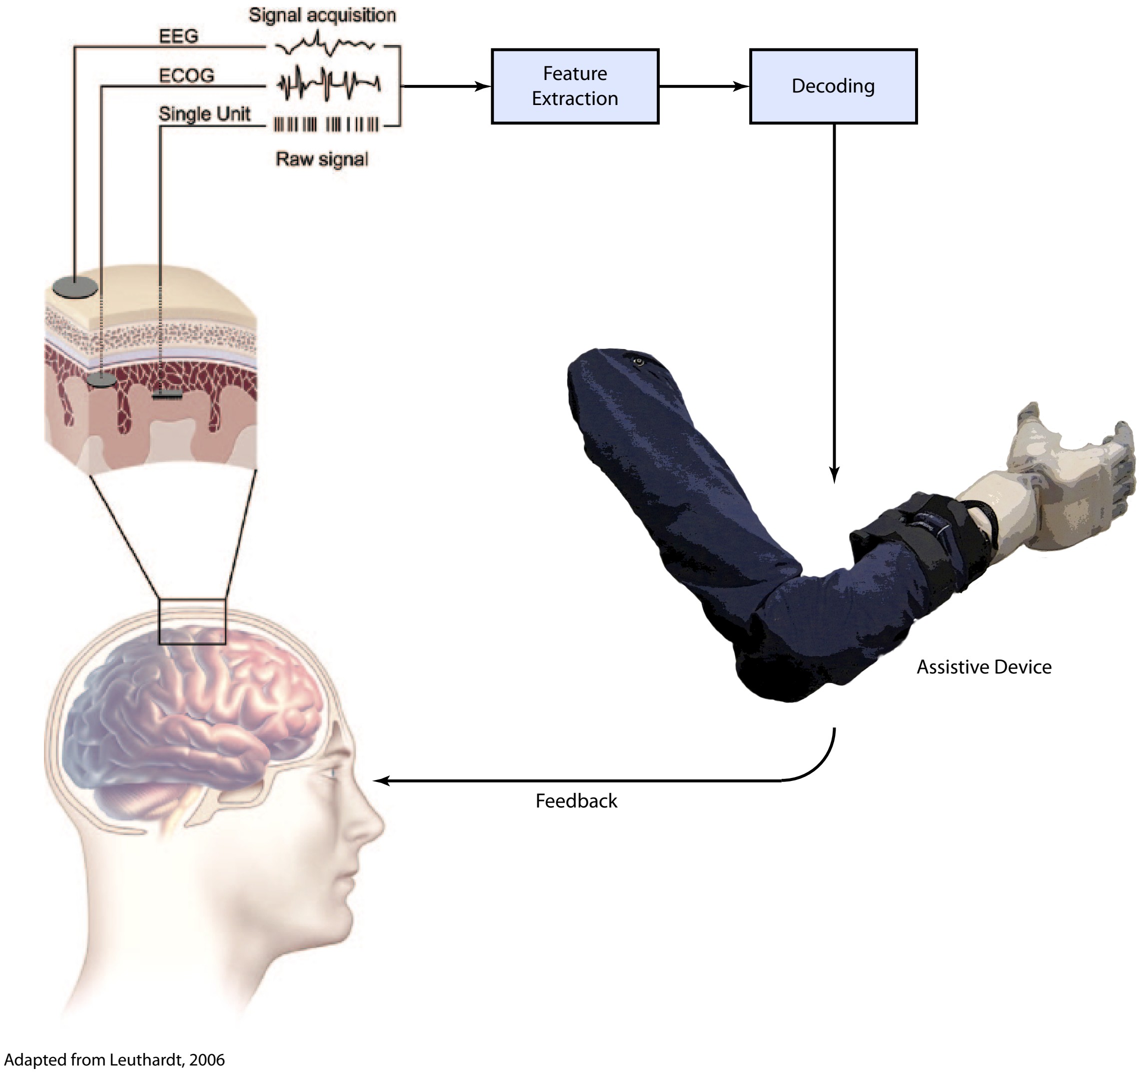

Figure 2.1: BMI system diagram. Neural signals are first recorded from the cortex using one of a number of available neural recording modalities. Salient information is obtained from these recordings using a feature extraction process. Extracted features are then passed through a decoder which converts them to command signals, which are used to control an assistive device. The user then receives feedback about the current state of the assistive device. Image adapted from [Leuthardt et al., 2006b].

2.1 NEURAL RECORDING MODALITIES

A wide range of neural recording modalities are available for BMI use, including electroencephalography (EEG), electrocorticography (ECoG), local field potentials (LFPs), single/multi-unit activity (SUA/MUA), functional magnetic resonance imaging (fMRI), magnetoencephalography (MEG). These methods can be differentiated by the type of neural signal recorded, then type of electrode (or sensor) used to record the activity, the spatial and temporal resolution of the recordings, invasiveness of the technique, and portability. Several of these modalities, including fMRI and MEG, involve the use of a large scanner and are not portable as a result. Though such modalities have been used successfully for brain-machine interfaces [Sudre et al., 2010], the non-portability of these systems makes them impractical for clinical BMI use; these methods will not be discussed further.

The chief modalities used in clinically-viable BMI systems are EEG, ECoG, LFPs, and MUA/SUA. These methods are common in their use of electrodes to record electrical activity from the cortex, but vary drastically in electrode design and placement. These differences have implications on the characteristics of the neural activity recorded by the modality, and can be characterized by their performance, decoding stability, longevity, and invasiveness.

Performance. BMI performance can be defined as the quality of control which can be achieved using a specific modality. This capability can be dependent upon a number of factors, including the number of neural features provided, their independence, encoding of relevant kinematic parameters, and signal-to-noise ratio (SNR). Modalities providing large numbers of independent, high-SNR features which encode kinematic parameters relevant to the task will typically be capable of supporting a high-performance BMI system.

Stability. Stability can defined as the ability of the BMI system to maintain consistent performance without frequent experimenter intervention (e.g., neural decoder updates), and is dependent upon the stability of the neural signals recorded by the modality. If the recording modality can record activity from the same neural population over time, and the encoding of information that neural population does not change with time, a neural recording modality will enable stable BMI operation.

Longevity. The longevity of a neural recording modality refers to the ability of the neural interface to record signals from the cortex for long periods of time. Longevity is a critical feature for a chronic BMI system, as it can facilitate long-term BMI performance without the need for future surgical intervention to repair or update the system. It should be noted that the concept of longevity is distinct from that of stability; a neural recording modality may provide low stability yet high longevity (or vice-versa). For example, the neural population recorded by a modality may change drastically from day-to-day, but the ability of the modality to record signals from the cortex may remain unchanged over time.

Invasiveness. The invasiveness of a recording modality may be defined as the risk the subject is placed at as a result of use of a particular recording modality. For example, those modalities requiring surgical implantation of electrodes can be considered highly invasive, while those modalities placing subjects at low-risk are considered non-invasive or minimally-invasive.

An “ideal” brain-machine interface would be capable of high-performance, stable control over long periods of time at low risk to the subject. Unfortunately, in reality this is not the case for currently-available recording modalities. Generally, a tradeoff exists between the aforementioned characteristics; the choice of a neural recording modality for a particular BMI application must balance these characteristics to achieve the desired level of performance. The following sections briefly describe EEG, ECoG, LFP, and MUA/SUA in additional detail, including the ability of these modalities to satisfy the requirements of a clinical BMI system.

2.1.1 Electroencephalography (EEG)

Electroencephalography (EEG) is a technique in which disc electrodes are placed on the scalp to record electrical field potentials from the brain (Figure 2.2). As such, EEG is considered a minimally invasive recording technology. This has led to widespread use of EEG for BMIs, including two and three-dimensional cursor control [Wolpaw and McFarland, 2004, McFarland et al., 2010, Foldes and Taylor, 2013]. Typically, EEG BMIs are based on volitional changes in the spectral power of electrical field potentials in the μ (8 - 12 Hz) or β (18 - 26 Hz) frequency bands, activity which as been shown to be related to motor cortical activity [McFarland et al., 2000] (see [Wolpaw et al., 2002] for a review).

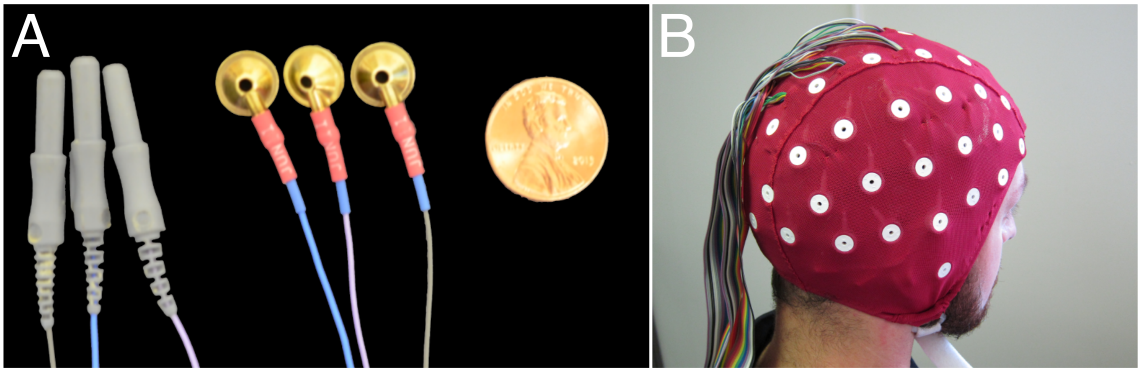

Figure 2.2: (*A*) Standard EEG cup electrodes. (*B*) Example of subject wearing an EEG electrode cap.

EEG electrodes are typically affixed to the scalp daily using common electrode montages (e.g. the *10-20 sytem*), which allow for relatively consistent placement of electrodes. Despite this, EEG recordings are still prone to instability issues arising from inconsistent electrode placement across days and intra-day changes in electrode impedance. However, EEG recordings may be considered to have high longevity, as the minimally-invasive nature of the method mean EEG is not susceptible to the brain tissue reaction believed to result in diminished recording quality over time for more invasive implanted microelectrodes [Polikov et al., 2005].

Unfortunately, the distance of recording electrodes from the cortex results in both low spatial resolution and low signal-to-noise ratio for EEG signals. Synchronous cortical activation of $> 6cm^2$ is required to generate measurable field potentials at the scalp [COOPER et al., 1965, Ebersole, 1997]. In general, this drastically limits the performance of EEG-based BMI systems, requiring synchronization of large cortical populations ($\sim 60,000,000$ neurons) to generate field potentials and reducing the number of independent sources which can be measured using EEG.

2.1.2 Electrocorticography (ECoG)

Similar to EEG, electrocorticographic (ECoG) signals originate from electrical field potentials in the brain. ECoG electrodes, typically grids of platinum disc electrodes several millimeters in diameter, are placed beneath the skull either epidurally (above the dura) or subdurally (beneath the dura) as shown in Figure 2.3. Placement of ECoG electrodes is considered an invasive procedure, requiring a craniotomy and, in the case of subdural implantation, retraction of the dura. ECoG has been shown to enable two-dimensional closed loop control of cursors or a prosthetic limb [Leuthardt et al., 2004, Wilson et al., 2006, Schalk et al., 2008c, Leuthardt et al., 2011, Yanagisawa et al., 2012].

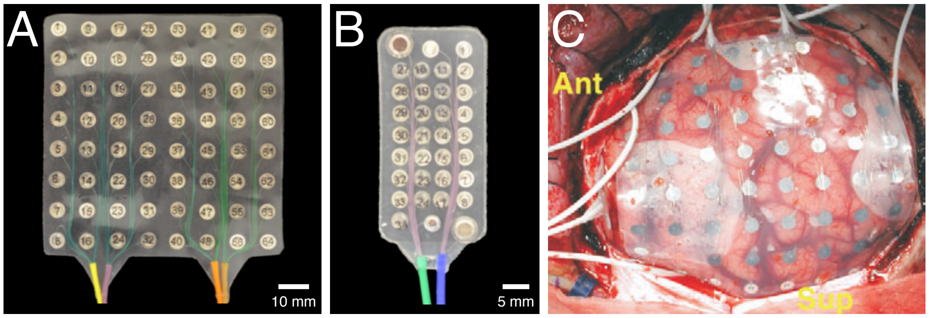

Figure 2.3: (*A*) Standard 64-contact ECoG electrode grid used for seizure monitoring in individuals with epilepsy. (*B*) Custom high-density ECoG electrode grid. (*C*) Intraoperative image showing placement of a standard ECoG electrode grid for epilepsy monitoring. Image reproduced from [Leuthardt et al., 2004].

As a result of the proximity of electrodes to the brain, ECoG recordings have increased spatial resolution compared to that of EEG recordings, on the order of ~ 1 mm for subdural ECoG [Freeman et al., 2000, Slutzky et al., 2010]. This increase in spatial resolution is accompanied by an increase in the spectral content of ECoG recordings. In addition to the μ and β bands, task-related ECoG activity is commonly observed in the low-γ (30 – 60 Hz) and high-γ ($>$ 60 Hz) frequency ranges. Based on the size of common electrode sizes, modulation of ECoG activity involves synchronization of $3 - 7mm^2$ regions of cortical tissue, an order of magnitude smaller than of EEG.

ECoG recordings generally considered to have high longevity and stability, as placement of electrodes does not compromise the integrity of cortical tissue. Furthermore, the dependence of ECoG on the activity of large populations of neurons is believed to prevent these recordings from being sensitive to changes in the activity of individual neurons [Moran, 2010, Leuthardt et al., 2004]. ECoG has been shown to provide stable, offline decoding of arm movement kinematics over a period of ~ 9 months [Chao et al., 2010]. Whether such longevity enables long-term closed-loop control of an ECoG BCI system, however, remains to be seen.

2.1.3 Local field potentials (LFPs)

Local field potentials (LFPs) are extracellular field potentials measured with penetrating microelectrodes (Figure 2.4). Similarly to EEG and ECoG, LFPs are commonly analyzed in the frequency domain, with the frequency range of interest overlapping with that of ECoG (< 250 Hz). Placement of penetrating electrodes is considered an invasive procedure, as electrodes are inserted into cortical tissue. While a number of studies have investigated the encoding of movement-related information in LFPs (e.g., [Heldman et al., 2006]), to-date the only examples of closed-loop BMI control using LFPs have focused on two-dimensional cursor control by non-human primates [So et al., 2014, Flint et al., 2013].

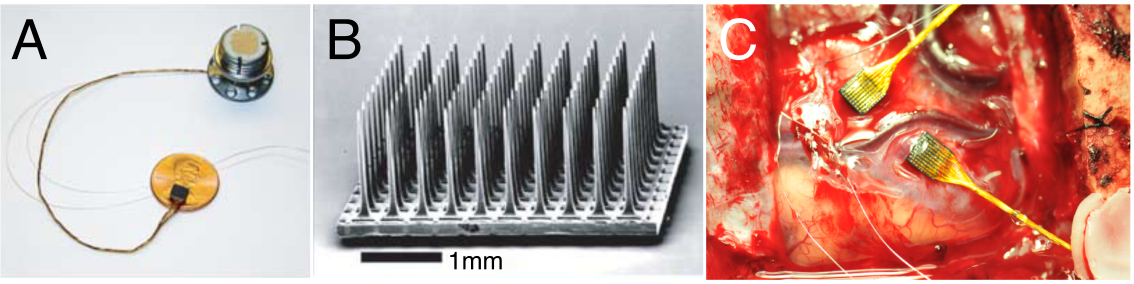

Figure 2.4: The "Utah" penetrating microelectrode array. (*A*) Utah array attached to a NeuroPort pedestal. Image reproduced from [Hochberg et al., 2006]. (*B*) Scanning electron micrograph of a Utah array showing electrode details. Image reproduced from [Hochberg et al., 2006]. (*C*) Arrays inserted into cortical tissue. Image courtesy of the Human Rehabilitation and Neural Engineering Laboratory (hRNEL).

LFP recordings are believed to represent neural activity within several millimeters of the tip of the recording electrode [Logothetis et al., 2001, Juergens et al., 1999]. It has been shown that high-gamma-band LFP activity is highly correlated with the average firing rate recorded by microelectrodes [Ray et al., 2008a], and that the spectral tuning of LFP activity is similar to that of individual neurons recorded on the same electrode [Heldman et al., 2006]. Thus, LFP recordings have a greater spatial resolution that that of either EEG or ECoG, allowing electrical field potentials to be recorded from a larger number of independent neural populations.

Recordings from penetrating microelectrodes have been shown to suffer from a decline in signal quality commonly attributed to a foreign body reaction to the implanted electrodes [Collinger et al., 2012, Simeral et al., 2011, Chestek et al., 2011]. Despite this, it has been postulated that it may be possible to record high-gamma-band LFPs over years [Moran, 2010]. This belief has been supported by recent work showing stable closed-loop control of a computer cursor for nearly 12 months using LFP recordings [Flint et al., 2013], suggesting that LFPs may provide sufficiently high stability and longevity to enable clinical BMI systems.

2.1.4 Single-unit/multi-unit activity (SU/MUA)

Lastly, single and multi-unit activity (SU/MUA) represent the activity of individual neurons as recorded by penetrating microelectrodes (Figure 2.4). As opposed to LFPs, SU/MUA recordings are obtained by band-pass filtering recorded electrode voltages over the $300 - 5,000$ Hz range, allowing individual action potentials to be identified [Schwartz et al., 2006]. This activity may be further broken down into the action potentials of isolated neurons (single-unit activity), or the activity of multiple neurons measured on a single electrode (multi-unit activity). In either case, SU/MUA action potential events are typically converted to time-varying firing rates [Dayan and Abbott, 2001]. As SU/MUA utilizes the same electrodes as that used for LFP recordings, it is also considered an invasive recording modality.

Single and multi-unit activity, as recorded by penetrating mircoelectrodes, is sensitive to neural activity within several hundred microns of the electrode [Logothetis et al., 2001], making it the recording modality with the highest spatial resolution currently used in BMI systems. This has made SU/MUA the recording modality of choice for use in BMI systems, both as a result of the encoding of movement direction kinematics in single-unit activity [Georgopoulos et al., 1982, Schwartz et al., 1988] and from early work showing that movement kinematics can be reconstructed from populations of neurons [Georgopoulos et al., 1986]. Combined with the advent of multi-electrode recording arrays [Campbell et al., 1991, Vetter et al., 2004, Musallam et al., 2007], SU/MUA recording has enabled high-performance BMI control. SU/MUA has been the prominent recording modality for BMI research, having been shown to enable closed-loop control of computer cursors [Flint et al., 2013, Gilja et al., 2012, Simeral et al., 2011, Kim et al., 2011, Ganguly and Carmena, 2009, Santhanam et al., 2006, Hochberg et al., 2006, Taylor et al., 2002, Wessberg et al., 2000], as well as robotic arms [Collinger et al., 2012, Hochberg et al., 2012, Velliste et al., 2008] in both humans and non-human primates.

Unfortunately, the access to the activity of individual neurons comes at the cost of decreased longevity for SU/MUA-based BMIs. Penetrating microelectrodes have been shown to be prone to brain tissue responses encapsulating implanted electrodes [Turner et al., 1999, Carter and Houk, 1993, Polikov et al., 2005]. Such encapsulation can lead to neural degeneration around the implant site [Polikov et al., 2005, McConnell et al., 2009, Biran et al., 2005], potentially leading to the reduction in signal quality observed in some intracortical BMI studies [Collinger et al., 2012, Simeral et al., 2011, Chestek et al., 2011]. Conversely, some researchers attribute such degradation in signal quality to a gradual failure of electrode array insulation [Barrese et al., 2013]. Despite such findings, several groups have found that SU/MUA recording can yield stable and robust BMI control [Ganguly and Carmena, 2009, Flint et al., 2013].

2.2 THE NEURAL BASIS OF CLOSED-LOOP BMI CONTROL

The link between user intent and recorded neural activity is critical to the successful operation of a BMI. Ideally, neural activity should directly reflect a user's desired BMI commands. However, in practice this not always as straight forward as it might seem; the BMI *control strategy*, or the cognitive process employed by the user to generate task-modulated neural activity, can differ drastically depending on the quality and type of recording modality. For individuals with movement paralysis, BMI control strategies are further complicated by the inability of the users to generate overt movements to serve as the basis for control. In these cases, congruence between neural activity during overt, attempted, and/or imagined movement can be leveraged to obtain closed-loop control. The following section will briefly discuss the use of motor imagery and observation-related activity to enable BMI control, as well as two of the prominent control strategies for clinical BMI systems: *naturalistic* and *somatotopic* control strategies.

2.2.1 Cortical activity during motor imagery, attempted movement, and movement execution.

Much of the work investigating the encoding of movement-related information in the cortex has utilized able-bodied subjects. While such studies serve to justify the feasibility of BMI systems, the potential candidates for BMIs are no longer able to make overt movements to generate neural activity to serve as the basis for closed-loop control. Some method of reliably generating motor cortical activity in the absence of overt movement must therefore be used.

The use of action observation has been proposed as a method for generating neural activation in the absence of overt movement [Collinger et al., 2014]. This approach is based on the concept of *mirror neurons* in the cortex, which have been shown to exhibit similar activity during the execution and observation of specific movements [Gallese et al., 1996]. Congruent activity has been observed between action observation and action execution in MEG [Hari et al., 1998, Caetano et al., 2007, Press et al., 2011], EEG [Muthukumaraswamy et al., 2004, Perry and Bentin, 2009], ECoG [Collinger et al., 2014], and penetrating microelectrodes [Tkach et al., 2008, Dushanova and Donoghue, 2010]. An action-observation-based approach typically will begin with the user observing automated movements of the BMI effector. Neural data collected during this task can then be used to calibrate the neural decoder used for closed-loop control (see Section 2.3). However, it is important to point out that pure observation-related activity is not sufficient to operate a BMI system; at some point the user's intent, rather than passive observation, must be used to drive cortical activity.

A closely related method for eliciting cortical activity in individuals with movement paralysis is the use of motor imagery or attempted movement. Here, subjects are instructed to either imagine performing or attempt to perform a particular movement, sometimes while simultaneously observing the same movement being performed by the BMI effector. As opposed to the action observation paradigm, using motor imagery or attempted movement to drive neural activity requires intent on the part of the user. Cortical activation during imagined movement has been shown in individuals with movement paralysis, with individual neurons exhibiting cosine tuning to imagined movement direction [Truccolo et al., 2008]. While it has been shown that there are distinct differences in neural activity during motor imagery and attempted movement [Hotz-Boendermaker et al., 2008], both approaches provide BMI users with a volitional means with which cortical activity can be generated. In practice, the line between these two types of activity becomes blurred during decoder calibration (see Section 2.3.3), where a BMI based initially on observation-related activity is eventually controlled by neural activity elicited by attempted movement.

2.2.2 The "naturalistic" control strategy

Perhaps the most straightforward approach to achieving closed-loop BMI control is the use of intended effector movement to generate task-related neural activity. The term "naturalistic" will be used to refer to this approach, as it attempts to utilize the natural encoding of the desired BMI movement commands in the neural population. In theory, such a control strategy is intuitive for the user; one simply thinks of the desired effector movement to control the system.

Early brain control efforts in non-human primates focused on neural control of computer cursors based on overt arm and hand movements [Serruya et al., 2002, Taylor et al., 2002, Carmena et al., 2003, Ganguly and Carmena, 2009]. In these studies, animals were first trained to perform a task under hand control of the computer cursor. Movement of the cursor was then switched over to neural control. In some cases, animals were permitted to continue to make arm movements during BMI control [Serruya et al., 2002], while in others the animal's arm was restrained [Taylor et al., 2002, Ganguly and Carmena, 2009]. Other non-human primate studies have shown closed-loop control of prosthetic limbs [Velliste et al., 2008, Velliste et al., 2014]; here again animals were initially trained to perform the control task using overt movements and subsequently transitioned to BMI control. Unfortunately, the use of the non-human primate model in these studies makes it difficult to ascertain the control strategy being employed by the subject as a consequence of the inability of experimenters to provide explicit instruction to the subjects. However, given that subjects initially perform the BMI task under overt movement control, it is likely that some form of a naturalistic strategy is being employed by the subjects in these studies.

Clinical BMI studies, alternatively, have the benefit of allowing explicit instructions to be provided to subjects. Such studies have shown that naturalistic control strategies using imagined [Hochberg et al., 2006, Truccolo et al., 2008, Hochberg et al., 2012] or observed [Collinger et al., 2012] arm and hand movements can serve as the basis for BMI control. Though effective, these approaches are dependent upon the detailed encoding of movement-related information in cortical activity. When such information is not available, a naturalistic-based approach may not be capable of providing the richness of information required to obtain control.

2.2.3 The "somatotopic" control strategy

EEG and, to a lesser extent, ECoG recordings, do not provide sufficient encoding of detailed movement kinematic information to enable BMI control using a naturalistic control strategy[2]. In these cases, an alternative approach must be used to generate task-modulated neural activity. One commonly-used approach is to assign imagined or attempted movements of different body parts to control different degrees of freedom of the effector. This approach is referred to as *somatotopic* control strategy.

The somatotopic control strategy has been successfully used to obtain closed-loop computer cursor control using EEG [McFarland et al., 2010] and ECoG [Leuthardt et al., 2004, Schalk et al., 2008c]. Potential movements for closed-loop control are typically identified using some form of movement screening task, during which neural responses are observed during a set attempted or imagined movements such as tongue protrusion or hand movement [Wolpaw and McFarland, 1994]. Movements eliciting strong cortical modulation can then be mapped onto the effector. For example, for a two-dimensional cursor task, one movement (e.g. hand grasp) could be assigned to control the velocity of the cursor in the X-dimension, while a different movement (e.g. tongue protrusion) could be used to control the velocity of the cursor in the Y-dimension.

Unfortunately, this approach typically results in a conflict between the intent of the user (e.g. to move the cursor) and the imagined or attempted movements made by the user to generate cortical activity. This conflict may serve to increase the cognitive load placed on the user, making the BMI system fatiguing to control over long periods of time. This may particularly be the case during somatotopic-based control of a prosthetic limb, where there will likely be a conflict between the attempted arm movements eliciting cortical modulation and the intended movement of the arm. Whether or not the cognitive load of a somatotopy-based BMI system can be reduced through long-term training remains to be seen.

2.3 FEATURE EXTRACTION, DECODING ALGORITHMS, AND DECODER CALIBRATION.

In order to control an external device using a BMI system, recorded neural signals must somehow be translated into device command signals. This translation occurs through a multi-step process in which relevant information, or "features", are first extracted from neural signals and are then passed through a *decoding algorithm* which maps the extracted features onto the output command signals. The specifics of this mapping are determined in the *decoder calibration* process, during which parameters of the decoding algorithm are

---

[2] The encoding of movement-related activity in ECoG recordings will be discussed in more detail in Chapter 3.

learned. The following sections will describe the feature extraction, decoding, and calibration processes in more detail.

2.3.1 Feature extraction

In most cases it is extremely difficult to extract meaningful information directly from raw neural signals. Typically, once neural signals have been recorded from the cortex using one of the aforementioned methods, a set of time-varying “neural features” must be extracted from the data which are informative about the intent of the user. This *feature extraction* process is highly dependent on the type of recording modality used. For SU/MUA recording this commonly involves the identification of action potential events (“spikes”) using a combination of thresholding and/or spike-sorting, followed by the binning of spikes in larger time windows [Dayan and Abbott, 2001]. For electrical field potentials (EEG, ECoG, LFPs), this commonly involves the transformation of time-domain signals into the time-frequency domain and possibly the averaging of these signals across specific frequency bands [Schalk et al., 2008c].

2.3.2 Decoding algorithms

Once a set of relevant features has been extracted from the recorded neural signals, a *neural decoder* is used to convert these features into a command signal for an external device. In contrast to feature extraction methods, decoding algorithms are less dependent upon the choice of neural recording modality, and can vary from extremely simple to exceedingly complex. In general, decoding methods take the following form:

$$\mathbf{y} = f(\mathbf{x}) \tag{2.1}$$

where $\mathbf{y} = [y_1, \dots, y_M]^\top$ is a M-dimensional vector representing the decoded output of the system (e.g. cursor movement velocity), and $\mathbf{x} = [x_1, \dots, x_D]^\top$ is a D-dimensional vector of neural features. The functional form of $f(\mathbf{x})$ depends on the particulars of a given method, though nearly all involve a set of *decoding weights* which map neural activity onto desired BMI commands; these weights are determined during the *decoder calibration* process. A few of the most common BMI decoding methods are briefly described in the following sections.

2.3.2.1 Manual weight assignment Perhaps the simplest decoding algorithm used for BMI systems involves the manual assignment of neural features to command signals by the experimenter. Such an approach has been used in both EEG [Wolpaw and McFarland, 1994] and ECoG [Leuthardt et al., 2004, Schalk et al., 2008c] BMIs, with advanced implementations utilizing adaptive weighting of the selected neural features to improve control [Wolpaw and McFarland, 2004]. This may be viewed as an extreme version of *feature selection* where a small subset of neural features is chosen to serve as the basis for BMI control. Though sometimes effective, these methods have limited utility, as manual assignment of decoding weights can prove intractable as the number of neural features and output command signals increases.

2.3.2.2 Linear-regression-based methods Linear-regression-based methods utilize a set of decoding weights to linearly map neural features onto desired BMI command signals. Such methods take the form

$$y_i = \mathbf{w}_i^\top \mathbf{x} + \epsilon_i \tag{2.2}$$

where $\mathbf{w}_i$ is a D-dimensional weight vector mapping the neural feature vector $\mathbf{x}$ onto the ith BMI command signal and ϵ_i is an additive noise term. Though simple in concept, linear regression can become problematic for large values of D (i.e. for large numbers of neural features), particularly in cases where calibration data is limited. In these situations, a number of variants of regularized linear regression, such as L_1 and L_2-regularized linear regression, can be utilized. Such methods seek to minimizes the L_1 or L_2 norm of $\mathbf{w}_i$, respectively [Bishop, 2006, Ganguli and Sompolinsky, 2012]. Linear regression has been used for closed-loop control of a computer cursor by an individual the spinal cord injury [Hochberg et al., 2006]. One commonly used variant of linear regression commonly used for BMIs is the Weiner filter, which takes the following form:

$$y_i = w_0 + \sum_{u=-m}^{N} \mathbf{w}_i(u)\mathbf{x}(t-u) + \epsilon_i \tag{2.3}$$

Here, weights $\mathbf{w}_i(u)$ are determined for a predefined set of time lags (u), with $\mathbf{x}(t-u)$ representing the neural activity at time $t-u$. Weiner filters have been successfully used in a number of BMI studies [Carmena et al., 2003, Ganguly and Carmena, 2009].

2.3.2.3 Neuron-specific methods Several decoding algorithms utilized in BMI systems have been developed specifically to take advantage of the characteristics of neuronal populations. Perhaps the most well-known of these methods is the Population Vector Algorithm (PVA) [Georgopoulos et al., 1986]. The PVA first defines an encoding model for each neuron of the form

$$f_i = b_{0,i} + \mathbf{b}_i^\top \mathbf{d} + \epsilon_i \tag{2.4}$$

where f_i is the firing rate of the ith neuron, $\mathbf{d}$ is a M-dimensional direction (or velocity) vector, $b_{0,i}$ is a firing rate offset term, and $\mathbf{b}_i = [b_1, \ldots, b_M]^\top$ is the *preferred direction* of the ith neuron (i.e. the direction which the ith neuron exhibits its highest firing rate)[3]. This model describes the dependence of the neural firing rate on movement kinematics, with this functional form of the encoding model is equivalent to the cosine tuning model that has been shown to well-characterize the response of primary motor cortical neurons to two and three-dimensional arm movements [Georgopoulos et al., 1982, Schwartz et al., 1988]. The firing rates of individual neurons are normalized according to Equation 2.5, and the kinematic output of the population is predicted from normalized firing rates according to Equation 2.6 [Chase et al., 2009].

$$r_i = \frac{f_i - b_{0,i}}{\|\mathbf{b}_i\|} \tag{2.5}$$

---

[3] The notation here has been changed from that of the previous section to reflect the dependence of these methods on assumptions specific to the characteristics of SU/MUA recordings.

$$\hat{\mathbf{d}} = k_s \frac{M}{N} \sum_{i=1}^{N} r_i \mathbf{b}_i \tag{2.6}$$

Here, $\hat{\mathbf{d}}$ is the predicted kinematic output and k_s is a normalization constant. While the PVA has been successfully used to obtain closed-loop control of a computer cursor [Taylor et al., 2002] and prosthetic limb [Velliste et al., 2008], it assumes a uniform distribution of preferred direction vectors across the neuronal population. The PVA is biased when this assumption is violated, potentially leading to decreased performance [Chase et al., 2009].

The optimal linear estimator (OLE) [Salinas and Abbott, 1994, Kass et al., 2005] attempts to correct for this bias by finding a new set of preferred direction vectors which optimally predict the kinematic output from neuronal firing rates. Here, preferred direction vectors from all neurons are collected in a single matrix $\mathbf{B}$:

$$\mathbf{B} = \left[\frac{\mathbf{b}_1}{\|\mathbf{b}_1\|}, \ldots, \frac{\mathbf{b}_1}{\|\mathbf{b}_1\|} \right]^\top \tag{2.7}$$

The instantaneous normalized firing rate of the neuronal population, $\mathbf{r} = [r_1, \ldots, r_N]^\top$, can then be expressed as:

$$\mathbf{r} = \mathbf{B}\mathbf{d} \tag{2.8}$$

If the kinematic prediction is formulated according to Equation 2.9, it can be shown that the optimal set of decoding weights, $\mathbf{W} = [\mathbf{w}_1, \ldots, \mathbf{w}_M]$, is found by Equation 2.10 [Chase et al., 2009].

$$\mathbf{d} = \mathbf{W}\mathbf{r} \tag{2.9}$$

$$\mathbf{W} = (\mathbf{B}^\top \mathbf{B})^{-1} \mathbf{B}^\top \tag{2.10}$$

The OLE has been shown to outperform the PVA in offline decoding of hand trajectories [Chase et al., 2009], and has been used for online control of a prosthetic limb [Collinger et al., 2012]. However, in online control situations subjects can compensate for biases in the PVA such that the performance of the PVA and OLE are equivalent [Koyama et al., 2010]. Finally, it is important to note that the OLE assumes independence between neurons [Kass et al., 2005], which may limit the applicability of this method in instances where this assumption is not valid.

2.3.2.4 State-space methods Another common class of decoders used in BMI studies are *state-space* methods. In comparison to previously-discussed methods, state-space models assume that the underlying *state* of a system (e.g. intended BMI command signal) evolves smoothly over time, and that noisy measurements of this state are observed (e.g. recorded neural data). Perhaps the most well-known of these models is the Kalman Filter (KF) [Kalman, 1960]. The state model for the KF may be expressed by Equation 2.11.

$$\mathbf{y}_{t+1} = \mathbf{A}_t \mathbf{y}_t + \mathbf{w}_t \quad , \quad \mathbf{w}_t \sim \mathcal{N}(\mathbf{0}, \mathbf{W}_t) \tag{2.11}$$

where $\mathbf{y}_{t+1} \in \mathbb{R}^{M \times 1}$ is the state of the system at time $t+1$, $\mathbf{y}_t$ is the state of the system at time (t), and $\mathbf{A} \in \mathbb{R}^{M \times M}$ is a matrix relating the state of the system at time $t+1$ to time t. Vector $\mathbf{w}_t$ is an additive noise term. The observation model relating the kinematic state of the system to the observed neural data is

$$\mathbf{x}_t = \mathbf{C}_t \mathbf{y}_t + \mathbf{q}_t \quad , \quad \mathbf{q}_t \sim \mathcal{N}(\mathbf{0}, \mathbf{Q}_t) \tag{2.12}$$

where $\mathbf{x}_t \in \mathbb{R}^{N \times 1}$ is a vector of neural firing rates, $\mathbf{C}_t \in \mathbb{R}^{N \times M}$ is a matrix relating the kinematic state to the neural data, and $\mathbf{q}_t$ represents the observation noise in the neural data. A detailed derivation of the KF, particularly as applied to neural data, is provided in [Wu et al., 2006].

The form of the state model (Equation 2.11), particularly that of $\mathbf{A}$, determines the smoothness with which the state evolves. For example, for $\mathbf{A} = \mathbf{I}$, the state model takes the form of a random walk model which encourages gradual changes in $\mathbf{y}_t$ [Brockwell, 2004]. While smoothness can be imposed on previously-discussed decoding methods by pre-filtering prior to decoding, state-based models such as the KF have the advantage of combining state filtering and estimation in a single probabilistic framework [Koyama et al., 2010].

State-space models such as the KF and the Particle Filter have been shown to outperform the PVA and OLE in offline decoding of arm movement velocities from motor cortical data [Brockwell, 2004, Wu et al., 2006], while the Kalman Filter has also been used successfully in real-time BMI control [Hochberg et al., 2012, Kim et al., 2011, Kim et al., 2008]. Furthermore, several variants on this method have been developed which seek to improve performance of the KF, such as in the extraction of speed information from neural data [Golub et al., 2014] or incorporation into a closed-loop calibration framework [Gilja et al., 2012]. The latter study is particularly noteworthy for the development of BMIs, as it seeks to integrate aspects of the closed-loop decoder calibration process into the assumptions of the decoding method in order to improve performance as well as generalizability across BMI tasks.

2.3.3 Decoder calibration

Regardless of the type of decoder used in a BMI system, the parameters of the decoder mapping extracted neural features onto effector command signals must be determined. This occurs during the *decoder calibration* process. Decoder calibration begins with the collection of neural data (*calibration data*), from which decoding parameters are learned. This set of parameters is commonly referred to simply as the *decoder*.

Decoder calibration commonly begins with the collection of a set of neural data during a movement observation or attempted/imagined movement task [Hochberg et al., 2006, Hochberg et al., 2012, Collinger et al., 2012]. Decoding parameters are learned from this data, after which the user can be given closed-loop control of the effector. However, it may be the case that performance with this decoder is poor, particularly in case of those decoders trained on observation-related activity. Such decoders may suffer from differences between neural activity during movement observation and that of attempted or imagined movement. In this situation, additional rounds of decoder calibration based on closed-loop control can be used to improve performance. Several studies have formalized this process, referring to it as a "coadaptive" prediction process [Taylor et al., 2002] or closed-loop decoder adaptation [Orsborn et al., 2012]. Other work has sought to further improve the calibration process by making assumptions about the intent of the user during brain control [Gilja et al., 2012].

In some cases, it is necessary to provide computer assistance to the user during decoder calibration. Such assistance may come in the form of attenuation of cursor movement errors [Velliste et al., 2008], or the addition of a control signal bias towards presented targets [Rouse et al., 2013]. Initial brain control (i.e. immediately following decoder calibration) is aided by high levels of computer assistance; the level of assistance provided to the user is incrementally reduced as the user becomes proficient with closed-loop control. Additional rounds of decoder calibration may be performed in conjunction with assist level reduction so that the neural decoder reflects the user's improved ability to control the BMI.

Finally, the amount of time required to adequately perform decoder calibration time must be given consideration. From a clinical perspective, it is beneficial to the user to reduce this time as much as possible, as a lengthy calibration process would be a substantial burden for users. However, the quality of the decoder obtained can depend on the amount of calibration data used to learn the decoding parameters. Typically, the accuracy of decoding parameters learned from calibration data increases with the amount of data. Furthermore, as the number of degrees of freedom (DoF) of the effector increase, the number of kinematic conditions sampled must also increase; ultimately this leads to an increase in the amount of calibration data required as the number of DoF increases.

Ultimately, decoder calibration is a critical step in achieving control of a BMI. Just as decoder parameters are learned from the user's neural data during calibration, the user learns to control the BMI system. This interaction between the BMI user and the system, manifested as learning during closed-loop control, is discussed in the following section.

2.4 LEARNING DURING BRAIN-MACHINE-INTERFACE CONTROL

Perhaps the most important aspect of a BMI system is the interaction of the user with the device. During operation, the user receives some form of feedback about the current state of the system. Typically, this is visual feedback of the state of the effector, though other forms such as vibrotactile [Godlove et al., 2014], peripheral stimulation [Horch et al., 2011], or cortical stimulation [Venkatraman and Carmena, 2011] are being investigated. Comparing resultant of the BMI effector with the intended movement commands allows the user to make adjustments to their intended movements in order to improve performance. This process will be referred to as *BMI learning*, and is a topic which is beginning to receive increased attention as BMI systems have become more prevalent.

Some of the earliest work investigating learning showed that non-human primates could learn to modulate the firing rates of individual neurons in primary motor cortex [Fetz, 1969]. Though not specifically in a BMI context, this work nonetheless laid the groundwork for future BMI studies, with recent work showing such learning can be used to control a one-dimensional BMI system [Moritz and Fetz, 2011]. Similarly, non-human primates can learn to differentially modulate the γ-band amplitude of epidural ECoG signals to control a one-dimensional BMI [Rouse et al., 2013].

Correlates of learning have been observed in a number of BMI-specific studies. Increases in control performance have been observed across multi-day exposure to a BMI system in non-human primates [Taylor et al., 2002, Carmena et al., 2003, Musallam et al., 2004] as well as humans [Collinger et al., 2012], though the cause of this presumed learning effect is unclear. In one such study it has been shown that the preferred directions of neurons during hand control and brain control diverge during prolonged BMI use, and that this divergence is accompanied by an increase in the tuning of neurons to the model assumed by the neural decoder [Taylor et al., 2002].

A number of research groups have begun going beyond demonstrations of learning during BMI skill acquisition to probe the learning which can occur during the acquisition of BMI control. Subjects have been shown to be able to compensate for perturbations in the preferred directions of a subset of neurons during PVA-based BMI control [Jarosiewicz et al., 2008]. It has also been found that subjects can learn to control a fixed-decoder BMI system over the course of several weeks, and successful control persists across days without re-calibration of the decoder [Ganguly and Carmena, 2009]. By exploring the ability subjects to modify their cortical activity to gain control of BMI systems, studies such as these are beginning to shed light on the mechanisms underlying the acquisition of BMI control. This work may ultimately lead to new training and calibration methods to improve performance of BMI systems.

While most of this research has focused on learning for the perspective of demonstrating or improving closed-loop control, BMIs are gaining increasing attention as a tool to address basic-neuroscience questions about sensory and motor control and learning [Wander and Rao, 2014]. Commonly-used motor learning paradigms, such as visuomotor rotations, have been applied to closed-loop BMI control in order to investigate cortical adaptation [Chase et al., 2012], while other researchers have used BMI paradigms to investigate the mechanisms underlying skill learning [Koralek et al., 2012]. While such studies may not be of immediately applicable to current BMI systems, the development of a deeper understanding of the mechanisms of learning in the cortex could ultimately lead to the development of more sophisticated BMI systems.

2.5 CONCLUSIONS: CHOOSING THE APPROPRIATE BMI SYSTEM

This chapter has highlighted the essential components and considerations in the development of a clinical BMI system, as well as provided an overview of the currently-available technologies and methods (e.g.

recording modalities and neural decoding algorithms) which are available. Ultimately, the relative merits of these options must be weighted in order to develop a BMI which best satisfies the needs of a given clinical population. For individuals with upper-limb paralysis, an ECoG-based system would appear to strike a good balance between the aforementioned requirements, providing access population-level activity with relatively higher spatial resolution than EEG while avoiding the consequences of implantation of microelectrodes. The following chapter will discuss electrocorticography in further detail and provide information that is critical to the appropriate choice of feature extraction, decoding, and decoding methods for an ECoG-based BMI system.

3.0 ELECTROCORTICOGRAPHY

Electrocorticography (ECoG), sometimes referred to as intracranial electroencephalography (iEEG), is an invasive neural recording modality in which electrodes are placed on the surface of the cortex (*subdural* ECoG) or dura (*epidural* ECoG) to record cortical field potentials. Used as early as 1939 [Penfield, 1939], this technique has historically been used in the treatment of pharmacologically-intractable epilepsy to identify epileptogenic foci prior to surgical resection. More recently, ECoG has emerged as a neuroscientific tool used to study human cortical activity, and as a potential recording modality for clinical brain-machine interfaces. This chapter will provide an overview of electrocorticography, including the physiological basis of the electrocorticographic signal, the encoding of cortical activity in ECoG, ECoG processing techniques, and ECoG-based BMI systems.

3.1 PHYSIOLOGICAL BASIS OF THE ELECTROCORTICOGRAPHIC SIGNAL

As ECoG records extracellular voltages from the cortex, these signals share much in common with local field potentials (LFPs) and electroencephalography (EEG), techniques which record cortical field potentials using penetrating microelectrodes and scalp macroelectrodes, respectively. Drawing heavily upon classical EEG techniques, ECoG activity is commonly characterized by changes in in the power of specific frequency bands [Miller et al., 2007a]. Though the exact ranges used may vary from study-to-study, field potential activity is generally broken down into the δ (1 – 3 Hz), θ (4 – 8 Hz), α (9 – 12 Hz, sometimes referred to as the μ rhythm when recording from sensorimotor cortex [Crone et al., 1998b]), β (12 – 30 Hz), γ (30 – 80 Hz), and high-γ ($>$ 80 Hz) bands [Buzsáki and Draguhn, 2004].

During movements, potentials recorded from motor cortex have been traditionally characterized by event-related desynchronization (ERD; an decrease in band power) of the α and β ranges [Crone et al., 1998b], or event-related synchronization (ERS; an increase in band power) of the γ and high-γ ranges [Crone et al., 1998a]. Though the correlation of the γ and high-γ bands during simple movements may appear to suggest they stem from similar sources, a number of studies have provided evidence that low-frequency (α and β) and high-frequency (γ and high-γ) stem from distinct, albeit related, sources. Investigations of the spatial distribution of motor cortical activity has shown low-frequency activity to be more spatially diffuse than high-frequency activity [Crone et al., 1998a, Szurhaj et al., 2006, Miller et al., 2012]. Consistent with this, modeling of the ECoG frequency spectrum has revealed that these spectra could arise from broadband modulation related to local neuronal processing obscured by activity in the α and β bands [Miller et al., 2009a]. Though this suggests uniformity of the high-frequency band, significant stimulus-related responses have been observed in narrow bands as high as 500 Hz [Gaona et al., 2011]. In addition, distinction has been drawn between the γ and high-γ bands, where stimulus manipulations have shown that γ and high-γ field potentials recorded from visual cortex can be decoupled, with high-γ modulation in particular believed to reflect aggregate action potential activity [Ray and Maunsell, 2011].

Though such studies have provided strong evidence for the independence of high and low-frequency field potential activity, the underlying sources contributing to this activity are still somewhat unclear. It is commonly believed that β activity is associated with thalamocortical circuits [Pfurtscheller and Lopes da Silva, 1999], while γ and high-γ activity reflects localized neuronal processing [Miller et al., 2007a]. This has led some to believe that low-frequency oscillations, which are decreased during movements, are indicative of resting-state activity [Miller et al., 2012]. However, others have suggested that β activity may act as a suppressive mechanism for gating motor function on account of coupling of the β and high-frequency bands specifically during non-movement states [Miller et al., 2012].

Unfortunately, development of a detailed model for the origin of electrical field potentials is hindered by the 'inverse problem', a term used to refer to the process of inference of cortical sources from multi-electrodes [Jewett and WILLISTON, 1971, Nunez and Srinivasan, 2006, Buzsáki et al., 2012]. As multiple sources undoubtedly contribute to the extracellular voltage measured by a single electrode, attribution of measured activity to any one source is problematic. Several studies have begun to address this through the use of volume-conductor and finite element models in an effort to characterize the sensitivity of ECoG electrodes to simulated cortical sources [Slutzky et al., 2010, Wodlinger et al., 2011]. However, much work is needed before the relationship between localized network activity and electrical field potentials is fully understood. Nevertheless, as the following section will discuss, this has not prevented ECoG from being utilized to study a broad number of cortical processes.

3.2 ENCODING OF CORTICAL ACTIVITY IN ECOG

Though the exact origins of ECoG activity may not be fully understood, access to ECoG recordings obtained from individuals undergoing monitoring for intractable epilepsy has provided researchers with a unique opportunity to study cortical activity in humans. ECoG has been used to investigate a number of cortical processes, including the auditory system [Edwards et al., 2005, Trautner et al., 2006], the visual system [Lachaux et al., 2005], language [Crone et al., 2001, Mainy et al., 2007, Kellis et al., 2010, Wang et al., 2011a, Pei et al., 2011], attention [Tallon-Baudry et al., 2005, Jung et al., 2008, Ray et al., 2008b], and the motor system. The encoding of motor cortical activity in ECoG is of particular interest to BMI researchers, as it provides insight into how activity occurring during natural arm and hand movements can be utilized in the development of brain-machine interfaces.

Early studies of motor cortical activity with ECoG were limited to characterizing responses to simple movements. These efforts demonstrated the existence of both ERD (i.e., a decrease in the power of the α and β bands) and ERS (i.e., an increase in the power of the γ band) during movements such as tongue protrusion or fist-clenching [Arroyo et al., 1993, Miller et al., 2007a, Crone et al., 1993, Crone et al., 1998b, Crone et al., 1998a]. Furthermore, characterization of ECoG responses during overt movements has shown that low-frequency activity is more spatially diffuse than high-frequency responses, consistent with theories of high-frequency field potentials reflecting the activity of local neuronal populations [Crone et al., 1998a, Miller et al., 2007a]. Importantly, these responses have been shown to be *somatotopically* organized, consistent with the traditional cortical homunculus [Toro et al., 1994, Miller et al., 2007a]. Based on such organization, activation arising from movement of distal body parts (e.g., the foot or leg) is expected to occur medially on the pre-central gyrus, while that arising from the movement of more proximal body parts (e.g., the face or tongue) should be located laterally on the pre-central gyrus.

Building upon these findings, researchers have investigated the encoding of detailed kinematic and kinetic parameters in ECoG recordings. It has been shown that gross movements, such as leftward and rightward reaches, elicit stereotypical responses in ECoG recordings [Leuthardt et al., 2004] (Figure 3.1A), and are sufficiently robust to enable the classification of reach direction [Chin et al., 2007]. Several groups have attempted to decode continuous arm position during two-dimensional reaching tasks from human ECoG activity with some success, showing the low-frequency component and high-γ band activity to be most informative about movements of the arm [Schalk et al., 2007, Pistohl et al., 2008, Nakanishi et al., 2013]. These findings have been replicated in non-human primate studies, where it has been found that three-dimensional movement trajectories can be predicted from both subdural and epidural ECoG activity [Shimoda et al., 2012, Chao et al., 2010]. Additionally, cosine tuning of ECoG spectral features has been observed during both circle-drawing and center-out reaching tasks [Schalk et al., 2007, Ball et al., 2009, Anderson et al., 2012]. Aside from predictions of arm movement kinematics, ECoG activity has been shown to allow for the prediction of muscle activity during a reach-to-grasp task [Shin et al., 2012].

More recently, studies of the encoding of motor cortical activity recorded with ECoG have focused on movement of the hand. Characteristic ECoG modulation has been observed in response to finger movements (Figure 3.1B), with the activity of electrodes found to be preferentially responsive to individual finger movements [Miller et al., 2009b, Kubánek et al., 2009]. This has enabled the movement of individual fingers to be classified from ECoG recordings with high accuracy [Kubánek et al., 2009, Wang et al., 2009, Chestek et al., 2013], as well as the reconstruction of the trajectories of individual finger movements [Kubánek et al., 2009, Acharya et al., 2010]. Perhaps not surprisingly, this has enabled the prediction of hand posture from recordings of ECoG activity [Degenhart et al., 2011a, Pistohl et al., 2012, Chestek et al., 2013]. Encoding of grasp force has also been observed in ECoG activity (Figure 3.1C) [Flint et al., 2014, Degenhart et al., 2011a], though whether such activity is distinguishable from responses to hand posture or kinematics is unknown.

Though the aforementioned studies provide evidence for the encoding of motor cortical activity in ECoG, it is still unclear how such encoding relates to that of single-neuron activity. Preferred directions of individual motor cortical neurons have been shown to be well-distributed throughout 3D space during reaching [Schwartz et al., 1988]; studies demonstrating cosine tuning in ECoG have not performed similar analyses to fully characterize responses to arm movements. Additionally, it has been shown that the preferred directions of neurons in primary motor cortex exhibit columnar organization in the sub-millimeter domain [Amirikian and Georgopoulos, 2003, Georgopoulos et al., 2007]. How such structure may influence field potential recordings obtained from electrodes several millimeters in diameter is currently unknown. Furthermore, the discrepancy between the mixed finger movement movement representation found in single-unit activity [Schieber and Hibbard, 1993] and observations of observations of individual-finger representation in ECoG recordings [Miller et al., 2009b, Kubánek et al., 2009] underscores the need for a unifying framework relating neuronal activity and field potential recordings. Ultimately, additional investigation is needed into how the somatotopic representation of movement in ECoG recordings is related to the tuning properties of the neuronal population in motor cortex.

3.3 ECOG RECORDING TECHNIQUES

3.3.1 Electrodes and electrode design

As shown by Figure 3.2, ECoG electrodes can vary widely in electrode size, inter-electrode distance, and fabrication process. Traditional electrode grids, commonly used for pre-surgical mapping with individual with intractable epilepsy, consist of platinum-iridium discs embedded in a silicon sheet with 4 mm electrode diameters and 10 mm inter-electrode (center-to-center) spacing, and are typically arranged in 4 to 8 electrode strips or large (8 × 8) grids (Figure 3.2A) [Crone et al., 1998b]. Variants of these grids with reduced inter-electrode distance (4 mm) have been used in an effort to increase spatial sampling of the cortex (Figure 3.2B) [Wang et al., 2009]. In order to further reduce electrode diameter and inter-electrode distance, some researches have utilized platinum-iridium mirco-wire grids with electrode diameters of 300μm (Figure 3.2C) [Rouse et al., 2013], while others have used lithographic techniques to fabricate thin-film grids with electrode diameters of 300μm and inter-electrode distances of 1 mm (Figure 3.2D) [Schendel et al., 2013, Thongpang et al., 2011]. In perhaps the biggest departure from traditional electrode designs, flexible, high-density grids have been fabricated utilizing nanomembrane transistors with thousands of electrodes and sub-millimeter inter-electrode spacing (Figure 3.2D) [Viventi et al., 2011].

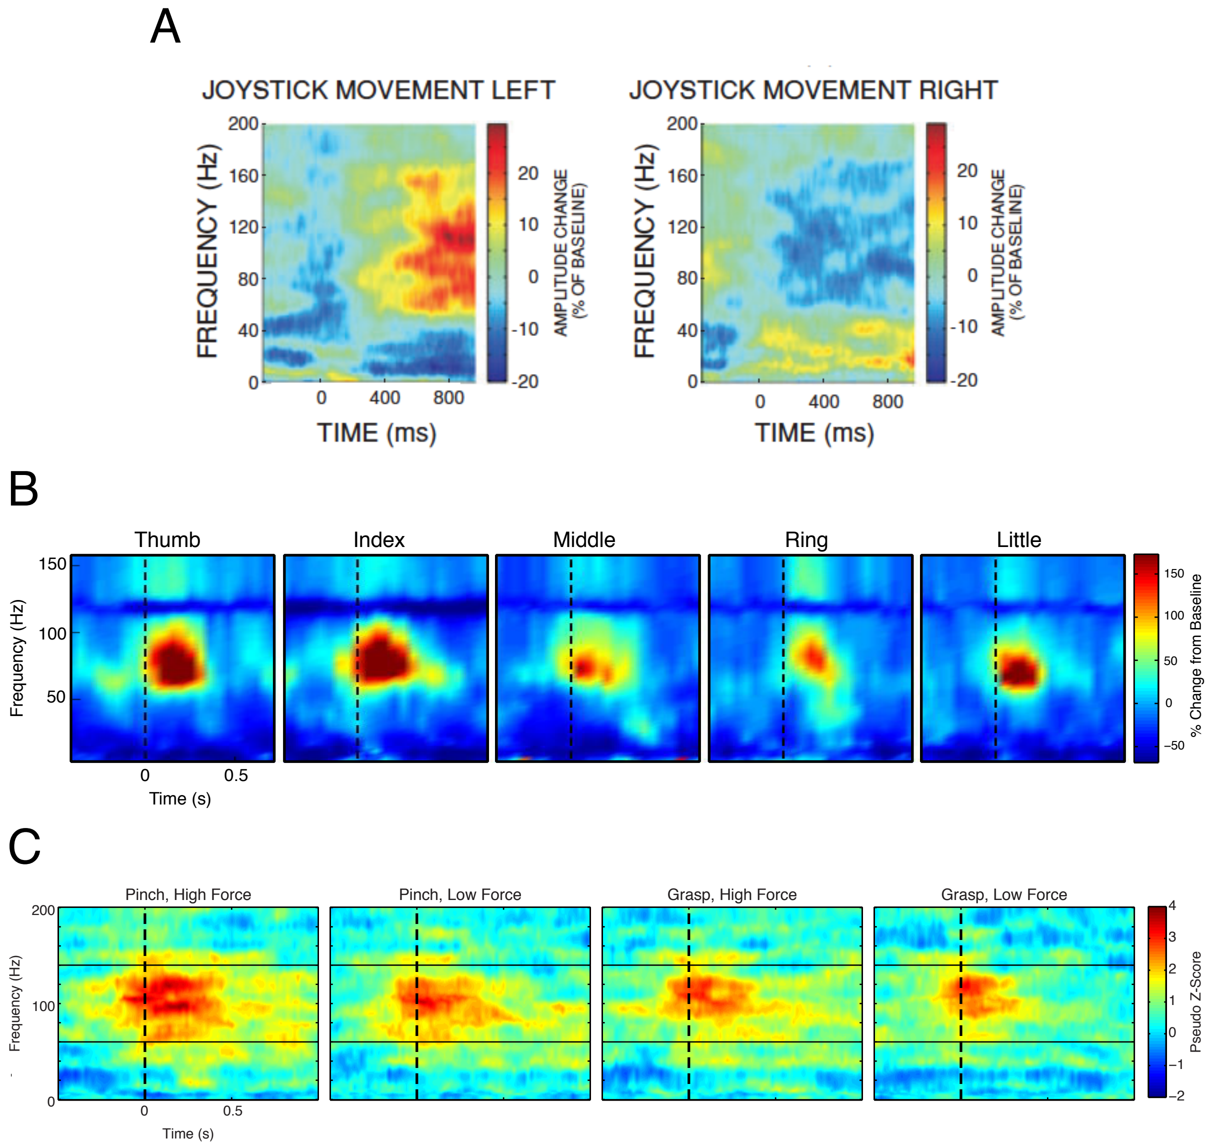

Figure 3.1: *Examples of ECoG modulation during arm and hand movement tasks.* (*A*) Modulation during leftwards and rightwards arm-directed joystick movements. Figure reproduced from [Leuthardt et al., 2004]. (*B*) Modulation during individual finger movements. Panels show time-frequency responses for a single electrode for movement of different fingers. Figure reproduced from [Wang et al., 2009]. (*C*) Modulation during a two-force, two-grasp hand posture task. Panels show time-frequency responses for a single electrode for different task conditions. Figure reproduced from [Degenhart et al., 2011a].

Examination of the effect of electrode diameter and spacing on recording quality has provided insight into how ECoG grid designs may be improved. Analysis of spontaneous ECoG activity in human subjects has found that inter-electrode distances of 1.25mm are required to effectively sample cortical tissue and avoid under-sampling and aliasing [Freeman et al., 2000]. Using finite element modeling, optimal inter-electrode spacing values of 0.6 mm (subdural), 0.6 - 0.9mm (epidural), and 2 mm (scalp) have been suggested, with results indicating that the main difference between subdural and epidural electrodes is the influence of a cerebral spinal fluid (CSF) layer [Slutzky et al., 2010]. Additionally, volume conductor modeling of cortical surface electrodes has found minimal benefit to electrode sizes smaller than 1 mm, with smaller electrode diameters exhibiting decreased signal-to-noise values [Wodlinger et al., 2011]. Collectively, these studies suggest that a tradeoff exists between the quality and specificity of recordings obtained with surface electrodes; while smaller electrodes may enable recording from increasingly independent neural populations, this independence comes at the cost of decreased signal-to-noise ration. Ultimately, determination of the optimal ECoG electrode design must balance the encoding of information in field potentials as a function of distance with the electrical characteristics of electrodes.

3.3.2 Signal processing and feature extraction

Before meaningful information can be extracted from ECoG recordings, these signals must undergo a number of processing steps, including re-referencing, filtering, and transformation into the frequency domain. In order to remove common-mode artifacts, a common average reference (CAR) is typically used [Schalk et al., 2007]:

$$s'_i(t) = s_i(t) - \sum_{j=1}^{N} s_j(t), \tag{3.1}$$

where $s_i(t)$ is the time-domain signal of the ith electrode at time t, N is the number of electrodes, and $s'_i(t)$ is the re-referenced signal. In addition to the standard CAR, adaptive filtering methods have been developed which are capable of adapting to channel-dependent amplitude or polarity of the common-mode signal [Kelly et al., 2013]. Re-referenced signals can then be notch-filtered at harmonics of 60 Hz to remove power-line noise contamination.

Once filtered, signals are then converted into a neural feature set containing task-relevant information. Though some studies have utilized the time-domain ECoG signal, termed the *low-pass filtered component* (LFC) [Pistohl et al., 2012] or *local motor potential* (LMP) [Schalk et al., 2007], feature extraction typically occurs through a spectral estimation process, where time-domain field potentials are transformed into the frequency domain. Numerous spectral estimation methods have been employed for the extraction of frequency content from field potential signals, including bandpass filtering [Rouse et al., 2013], the fast fourier transform (FFT) [Blakely et al., 2009, Edwards et al., 2005], wavelet [Chao et al., 2010, Gaona et al., 2011, Miller et al., 2009b], multitaper [Ball et al., 2009, Zhuang et al., 2010] autoregressive (AR) [Leuthardt et al., 2011, Leuthardt et al., 2004, McFarland et al., 2010, Pei et al., 2011, Wang et al., 2013a], and matching pursuit

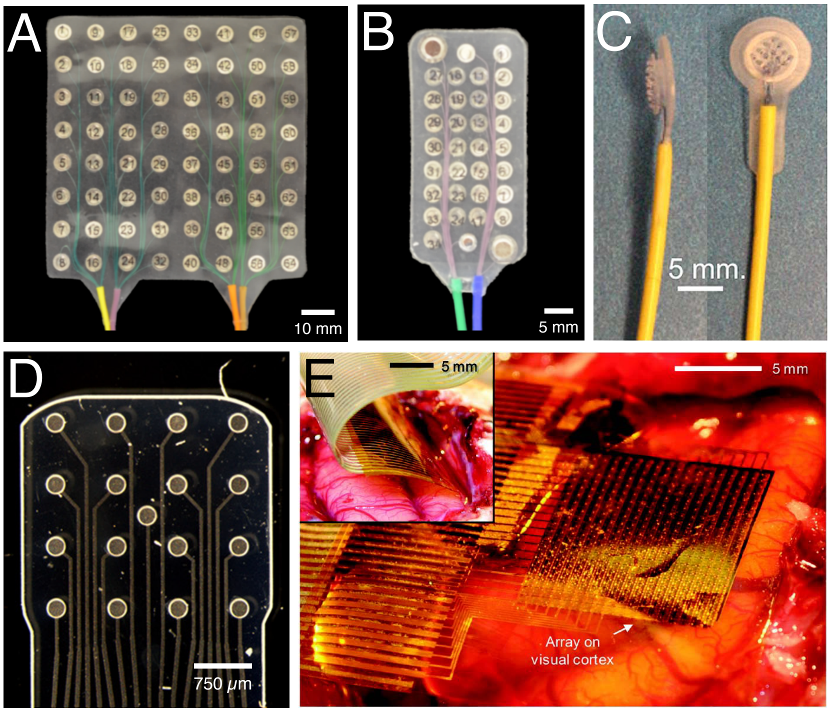

Figure 3.2: *Examples of ECoG electrode grids.* (*A*) Standard 64-electrode grid. (*B*) High-density 32-contact electrode grid. (*C*) 16-electrode microwire grid. Image reproduced from [Schalk and Leuthardt, 2011]. (*D*) Micro-ECoG electrode grid fabricated using lithography. Image reproduced from [Schendel et al., 2014]. (*E*) Flexible, high-density electrode grid. Image reproduced from [Viventi et al., 2011].

[Ray et al., 2008b] algorithms. These methods range from simple (e.g., bandpass filtering) to exceedingly complex (e.g, matching pursuit); the choice of spectral estimation method used for the analysis of ECoG recordings must be chosen to appropriately reveal task-relevant information while meeting experimental requirements. For example, though the matching pursuit algorithm allows for high-resolution time-frequency analysis, this method is only suited for offline analyses, as it is both non-causal and requires substantial computation time [Ray et al., 2008b].

For ECoG-based BMI control, the spectral estimation method used must be both causal and computationally efficient enough to be run in real-time. Thus, among the numerous spectral estimation algorithms available, only the bandpass filter [Rouse et al., 2013], FFT [Yanagisawa et al., 2011, Yanagisawa et al., 2012], and autoregressive [Leuthardt et al., 2004, Schalk et al., 2008c, Hinterberger et al., 2008, Blakely et al., 2009, Leuthardt et al., 2011, Wang et al., 2013a] methods have been used during real-time control. Of these, AR-based methods, such as the *maximum-entropy* (MEM) [Marple Jr, 1987] and Burg [Kay and Marple, 1981] algorithms, are by far the most common, in part on account of their ability to reproduce both narrowband and broadband frequency components [Kay and Marple, 1981]. With these methods, conversion of time-domain signals into the time-frequency domain is performed using a windowed spectral estimation process in which the frequency content of a short (*windowed*) segment of data is estimated, the data window is then stepped temporally, and the process is repeated to generate a time-varying estimate of the spectral content of the signal. It is important to note that choices of the length of the spectral estimation window can effect the characteristics of the resultant time-frequency data. As shown by Figure 3.3, the *window length* determines the influence of the most recent neural data on the spectral estimate, with shorter window lengths increasing the influence of the most recent neural data at the expense of an increase in the variability of the estimate. Determination of the appropriate window size for real-time BMI operation must therefore balance the desired responsiveness of the system with the stability of the spectral estimation process.

Once transformed into the time-frequency domain, ECoG signals frequently are normalized prior to analysis. ECoG field potentials exhibit a $1/f$ falloff in spectral power with frequency [Freeman et al., 2000], which can be problematic during analysis for several reasons. First, this makes visualization of time-frequency data difficult, as power in the in the high-frequency band is extremely small when compared to that of the low-frequency band. Secondly, this creates difficulty when averaging spectral power across frequency bands, resulting in a disproportionally high contribution of lower-frequency components to the average. To account for this, spectral estimates can be normalized through transformation into change-from-baseline values [Heldman et al., 2006] (Equation 3.2) or log-transformed and converted to pseudo-Z-scores [Tallon-Baudry et al., 2005, Edwards et al., 2009] (Equation 3.3) relative to a “baseline” condition:

$$f'_i = \frac{f_i - \bar{f}_i}{\bar{f}_i}, \tag{3.2}$$

$$f'_i = \frac{f_i - \bar{f}_i}{\sigma_i}, \tag{3.3}$$

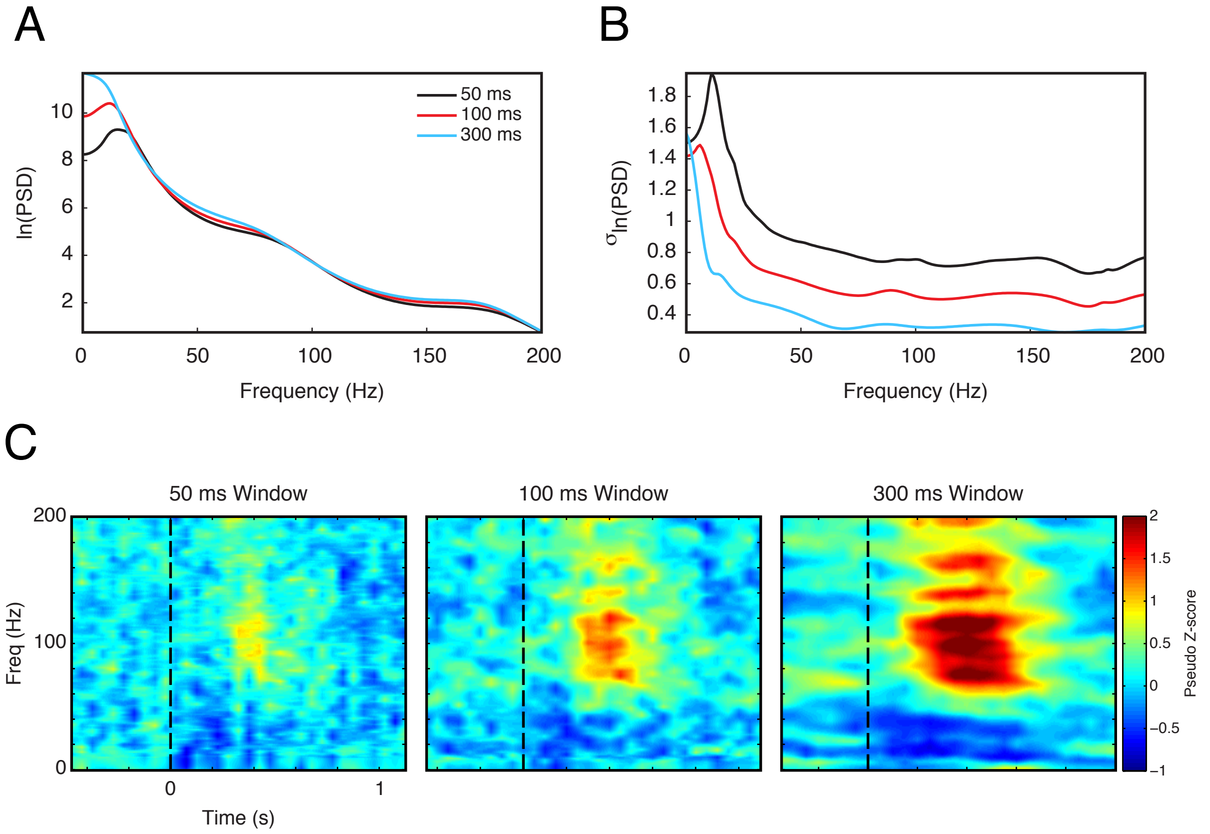

Figure 3.3: *Effect of spectral estimation window length on time-frequency responses.* (*A*) Average log-power for ECoG activity estimated with 50 ms, 100 ms, and 300 ms window lengths. (*B*) Standard deviation of log-spectral-power for different window lengths. (*C*) Time-frequency responses during hand grasping for different window lengths. Dashed lines at $t = 0$ marks the onset of the grasp cue. A fixed step size of 50 ms was used in all conditions.

where f_i is the unnormalized power of the ith spectral feature (i.e., power for a particular frequency band on a single electrode), $\bar{f}_i$ is average power of the ith spectral feature during the baseline condition, σ_i is the standard deviation of the raw band power of the ith feature during the baseline condition, and f_i' is the normalized activity of the ith feature. Baseline data is commonly taken from the inter-trial period preceding individual trials [Heldman et al., 2006], or from a continuous segment of resting-state data at the beginning of a testing session [Wang et al., 2013a]. While both change-from-baseline and pseudo-Z-score normalization methods have been utilized in the analysis of field potential data, the pseudo-Z-score method has the distinct advantage of accounting for the variance of the baseline feature activity.

Following normalization, processed time-frequency data consists of normalized power estimates for each frequency band for each electrode. Depending on the frequency resolution and range used during spectral estimation, the number of features per electrode may be quite large. Furthermore, as previously described, ECoG activity in specific frequency ranges (e.g., the μ and γ bands) is believed to reflect distinct cortical processes. Thus, as a final pre-processing step prior to analysis, normalized spectral estimates are sometimes averaged across larger frequency ranges, such as the $8 - 12$ Hz, $18 - 24$ Hz, $75 - 115$ Hz, $125 - 195$ Hz, and $159 - 175$ Hz bands [Kubánek et al., 2009]. Benefits of such averaging include *feature reduction*, which can help to prevent overfitting when decoding information from ECoG activity, and an increase in the signal-to-noise ratio (SNR) gained by averaging out independent noise of the smaller frequency bands. However, it is important to note that the boundaries of such bands are somewhat arbitrary, and may detrimentally affect analyses if improperly chosen. For example, it has been shown that the upper limit of the high-frequency range is typically determined by the noise floor of the amplifier [Miller et al., 2009a]; inclusion of frequency bands beyond this upper bound when averaging could inadvertently reduce SNR. Furthermore, fine-grained structure in ECoG spectra, which has been observed over narrow frequency bands [Gaona et al., 2011], may be lost as a consequence of averaging over large frequency ranges.

3.3.3 ECoG decoding techniques

A wide range of decoding algorithms have been employed in the extraction of information from ECoG data. These methods vary from discrete classifiers such as Naïve Bayes [Chestek et al., 2013], linear discriminant analysis [Pistohl et al., 2012], and support vector machines [Wang et al., 2009, Kanas et al., 2014], to continuous decoders such as linear regression [Kelly et al., 2012, Wang et al., 2013a], partial least squares [Chao et al., 2010, Shimoda et al., 2012], Wiener filters [Nakanishi et al., 2013, Flint et al., 2014], the kalman filter [Pistohl et al., 2008], and recurrent neural networks [Gunduz et al., 2007]. While this suggests that most machine learning techniques can be applied to ECoG recordings, there are several characteristics of these data which should be considered when choosing a particular decoding method. First, ECoG signals are high-dimensional; it is not uncommon for spectral estimation to be performed with 5 Hz or smaller frequency resolution, resulting in more than 600 features for a 32-electrode grid. Secondly, the nature of many human ECoG experiments, particularly those performed pre-surgical clinical mapping, dictate that data must be collected in a short amount of time, leading to datasets with few observations with which decoders can be trained. Combined, these characteristics can lead to *over-fitting* decoder parameters, resulting in poor performance.

To combat this, a number of techniques, including *feature selection*, *dimensionality reduction*, and *regularization* can be employed. Feature selection, which refers to the removal of uninformative components from the feature set used for decoding, has been employed for the prediction of arm movement trajectories using partial least squares [Chao et al., 2010, Shimoda et al., 2012]. Additionally, the aforementioned feature averaging process, in which the total feature set is reduced to the activity of a few pre-determined frequency bands [Kubánek et al., 2009], can also be viewed as a type of feature selection. Dimensionality reduction, including techniques such as Principal Components Analysis (PCA), and Factor Analysis (FA), reduces the feature set size by projecting the original high-dimensional feature set space into a lower-dimensional space and has been used in the prediction of finger movements from ECoG [Miller et al., 2009b, Degenhart et al., 2011a]. Finally, regularization techniques, which impose a penalty on decoding weights solutions based on characteristics of the weights themselves, have been utilized in the prediction of arm movement trajectories [Nakanishi et al., 2013], grasp [Pistohl et al., 2012, Flint et al., 2014], and cursor movement [Kelly et al., 2012].

3.4 ELECTROCORTICOGRAPHIC BRAIN-MACHINE INTERFACES

Based on the encoding of movement-related information in ECoG recordings, the higher spatial resolution and reduced susceptibility to artifact contamination compared to non-invasive recording modalities such as EEG [Freeman et al., 2000, Slutzky et al., 2010], and access to human subjects undergoing pre-surgical monitoring with subdural electrodes, researchers have investigated the potential for ECoG to support brain-machine interface systems. While the majority of ECoG BMI studies have been conducted with human subjects, non-human primate models have also been utilized to perform studies not feasible in a human subject population. The following sections will provide a brief overview of some of the key findings of both human and non-human primate BMI literature.

3.4.1 Human BMI studies[1]

Early ECoG BMI studies primarily focused on demonstrations of closed-loop control by human subjects. By using motor cortical activity during real or imagined movements, it has been shown that human subjects are capable of controlling computer cursors to perform one-dimensional [Leuthardt et al., 2004, Leuthardt et al., 2006a, Felton et al., 2007] or two-dimensional [Schalk et al., 2008c] tasks. Consistent with the encoding

---

[1] Though several studies, notably [Hinterberger et al., 2008], have investigated ECoG BMIs as communication aids, this work falls outside of the scope of BMI systems for the restoration of movement and will not be discussed further.

of movement-related activity in ECoG recordings, these studies utilized somatotopic control strategies to guide computer cursor movement. For one-dimensional control, this consisted of performing an overt or imagined movement, such as opening and closing of the hand, to drive cursor movement in one direction (e.g., the positive Y-direction) and relaxation to move the cursor in the opposing direction (e.g., the negative Y-direction). For two-dimensional control, activity from a second overt or imagined movement was used to guide cursor movement in the added dimension, with pairs of movement used to obtain 2D control taken from somoatotopically-distant locations such as the tongue and the hand [Schalk et al., 2008c]. More recently, it has been shown that ECoG can be used to control simple movements of a prosthetic limb, including by an individual with moderate motor impairment (spasticity) [Yanagisawa et al., 2012]. However, it is worth noting that in this study all subjects utilized overt arm and hand movements as the basis for closed-loop control.

ECoG brain-machine interface studies in humans have also provided valuable insight into the cortical processes underlying closed-loop brain control. Examination of activity during imagined somatotopic movements has revealed that closed-loop feedback can enhance cortical modulation [Miller et al., 2010], suggesting that subject learning may occur during the acquisition of ECoG-based BMI control. Human subjects have also been found to be capable of learning closed-loop cursor control using auditory imagery [Wilson et al., 2006, Felton et al., 2007], demonstrating the potential for non-motor cortical activity to serve as the neural substrate for BMI systems. Finally, it has been shown that stable, multi-day control of a 1D ECoG BMI system can be achieved using fixed decoding parameters, providing some evidence for the theorized stability of ECoG activity [Blakely et al., 2009].

Unfortunately, most human ECoG studies, including work investigating the encoding of movement-related information in ECoG, are conducted with subjects undergoing monitoring for intractable epilepsy. While this provides access to a much larger subject population, the nature of studies performed in the epilepsy monitoring unit (EMU) have a number of distinct disadvantages. First and foremost, studies in the EMU are hindered by their reliance upon a non-target subject population. As individuals with epilepsy are able to make overt movements, something that individuals with movement paralysis are incapable of doing, BMI studies in the EMU typically rely on imagined movements to serve as the basis for closed-loop control. As differences between imagined and attempted movements have been observed in individuals with paralysis [Hotz-Boendermaker et al., 2008], it is unclear how brain-control results obtained with individuals with epilepsy will generalize to the target clinical population for BMI systems. Secondly, research performed in the EMU is limited by short study durations and low subject motivation on account of post-operative pain and/or medication, making the study of BMI control over longer timescales difficult. In order to fully determine the feasibility of ECoG BMI systems, an investigation into the ability of individuals with movement paralysis to control such systems must be conducted.

3.4.2 Non-human primate BMI studies

Though limited, a number of studies have used non-human primate models to conduct BMI studies not feasible with human subjects. Foremost amongst these are investigations into learning during brain-machine interface control. Such studies are typically not feasible with human subjects, as they require substantial amounts of training time. It has been shown that through the course of multi-day BMI training, animals are capable of de-correlating gamma-band ECoG activity between pairs of electrodes [Rouse and Moran, 2009, Rouse et al., 2013]. Such work is particularly important to the development of ECoG-based BMI systems, as the lack of detailed kinematic information in ECoG suggests that some form of cortical adaptation may ultimately be necessary to achieve high-performance closed-loop control.

In addition to enabling the study of learning and adaptation during ECoG BMI control, the ability to conduct chronic studies with non-human primates has enabled the investigation of long-term stability of ECoG signals. As ECoG recordings presumably reflect the activity of thousands of neurons, it is believed they may be less sensitive to changes in activity from any individual neuron [Moran, 2010, Leuthardt et al., 2004]. This has led some to postulate that ECoG may be capable of providing robust and stable recordings, potentially eliminating the need for daily updating of decoding parameters. Several studies have begun to shed light on this, demonstrating long-term offline decoding of movement kinematics from subdural [Chao et al., 2010] and epidural [Shimoda et al., 2012] ECoG recordings using fixed decoding parameters. However, to-date, long-term closed-loop control of an ECoG BMI system with a fixed neural decoder has not been demonstrated.

3.5 CONCLUSIONS: ADVANCING THE CURRENT STATE OF ECOG BMI SYSTEMS

The capability of electrocorticography to record robust movement-related activity from the cortex, combined with demonstrations of closed-loop BMI control, provides key evidence for the potential of ECoG to support brain-machine interfaces. However, there are still several unanswered questions which must be addressed before ECoG BMI technology can be fully translated to the clinical realm. The following chapters will present work attempting to address several of these shortcomings and advance the current state of ECoG brain-machine interfaces. First, the ability of individuals with movement paralysis to control an ECoG-based BMI system will be assessed (Chapter 4). Second, the factors contributing to adaptation during the acquisition of ECoG-based BMI control will be investigated (Chapter 5). Third, the capability of ECoG to sustain robust, long-term closed-loop brain control will be evaluated (Chapter 6). Finally, a novel decoding algorithm will be presented which attempts to improve the extraction of BMI command signals from ECoG recordings through the incorporation of field-potential-specific correlation structure (Chapter 7).

4.0 AN ELECTROCORTICOGRAPHIC BRAIN-MACHINE INTERFACE FOR INDIVIDUALS WITH UPPER-LIMB PARALYSIS[1]

As mentioned in the previous chapter, while work with patients undergoing clinical brain mapping for treatment of epilepsy or chronic pain has demonstrated that BMI control signals can be extracted from ECoG recordings [Schalk et al., 2008c, Leuthardt et al., 2004, Acharya et al., 2010, Chao et al., 2010, Wang et al., 2009, Kellis et al., 2010, Miller et al., 2009b, Yanagisawa et al., 2012], to-date successful BMI control has not been demonstrated in individuals with movement paralysis. This chapter presents an investigation into the feasibility of an ECoG-based BMI system in two individuals with upper-limb paralysis. High-density ECoG grids were implanted subdurally over sensorimotor cortical areas of subjects for up to 28 days, during which they were trained to control 2D and 3D cursor movement using ECoG signals. We show that subjects were able to voluntarily modulate their cortical activity to gain successful closed-loop control, further validating the use of ECoG as a neural recording modality capable of supporting clinically-viable BMI systems.

4.1 METHODS

4.1.1 Subjects and surgical procedures

All experimental procedures were approved by the Institutional Review Board at the University of Pittsburgh and followed all guidelines for human subject research. Written informed consent was obtained before initiating any research procedures.

Subject 1 (S1) was a 30-year-old right-handed male with tetraplegia caused by a complete C4 level spinal cord injury seven years prior to the study. Subject 2 (S2) was a 54-year-old right-handed male with amyotrophic lateral sclerosis (ALS) diagnosed 9 years prior to enrollment in the study. Both subjects were capable of neck and shoulder control but could not initiate voluntary arm or hand movement. In addition, subject S2 was ambulatory at the time of the study.

In order to guide placement of ECoG electrode grids, functional magnetic resonance imaging (fMRI)

---

[1] A version of this chapter has been published presenting the results of closed-loop BMI control by an individual with spinal cord injury [Wang et al., 2013a]. This chapter expands upon this publication by adding results from a second subject.

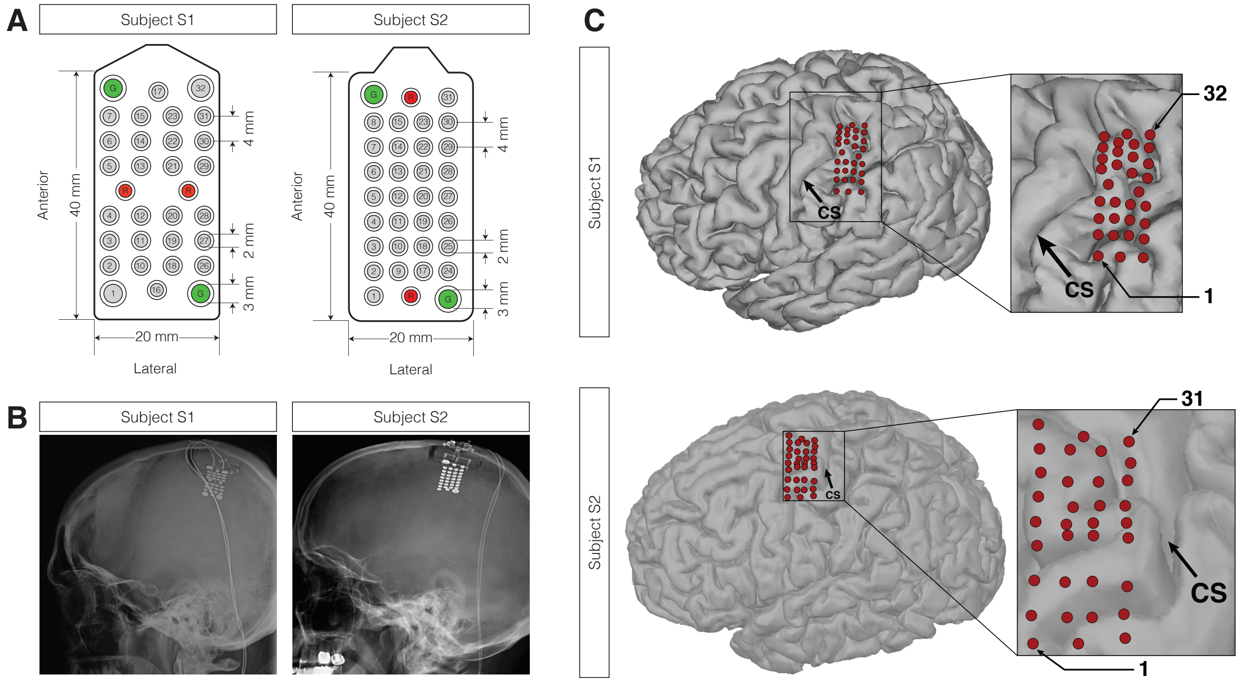

Figure 4.1: *ECoG grid design and location. A.* ECoG grid design for subjects S1 and S2. Gray circles represent recording electrodes. Red and green circles represent represent upside-down (skull-facing) reference and ground electrodes, respectively. *B.* Post-operative x-ray imaging showing implanted ECoG grids. *C.* Location of ECoG electrodes on the cortical surface. Electrode locations were determined using post-operative head x-ray, computed tomography (CT), structural MRI images, and intra-operative images. In addition, electrode localization for Subject S1 also utilized the coordinates of exposed electrodes recorded by the surgical navigation system (Brainlab AG, Feldkirchen, Germany) during the grid implantation surgery [Hermes et al., 2010, Miller et al., 2007b]

was conducted prior to the implantation surgery in order to localize subjects' sensorimotor cortex. Subjects were presented with videos of isolated arm and hand movements (e.g. elbow flexion/extension, wrist flexion/extension, hand grasp/release) interleaved with visual fixation (rest) blocks and were instructed to attempt to make the movements depicted by the stimuli. Functional images were collected using a T2*-weighted echoplanar imaging (EPI) pulse sequence (31 oblique axial slices, in-plane resolution $2mm \times 2mm$, $3mm$ slice thickness, no gap, $TR = 2000ms$, echo time $TE = 29ms$, $FA = 90°$, $GRAPPA = 2$, matrix size = 96 × 96, field of view $FOV = 192mm$). Stimulus presentation and synchronization with the MRI scanner was performed using E-Prime (Psychology Software Tools, Inc., Sharpsburg, PA, USA). Raw blood-oxygen-level dependent (BOLD) responses were convolved with a hemodynamic response kernel and converted to statistical t-maps comparing BOLD activity between attempted movement and rest epochs for each stimuli using the SPM8 toolbox (Wellcome Department of Imaging Neuroscience, London, United Kingdom). In order to compare the cortical activity across movements, these t-maps were then coregistered with the anatomical MRI scan for each subject, thresholded, and rendered on the cortical surface using the Freesurfer image analysis suite (`http://surfer.nmr.mgh.harvard.edu/`) and the SUMA toolbox for AFNI [Cox, 1996]. Target electrode grid implantation locations were chosen based on the correspondence of this fMRI activity with known anatomical landmarks in order to maximize coverage of sensorimotor cortex.

Subjects were implanted with high-density ECoG grids (PMT Corp, Chanhassen, MN USA) consisting of 32 (subject S1) or 34 (subject S2) platinum disc electrodes embedded in a $2cm \times 4cm$ silicone sheet (Figure 4.1A). Electrodes were either 2mm or 3mm in diameter and were spaced $4mm$ apart. Electrode grids were implanted subdurally over the hand and arm areas of left sensorimotor cortex through a small craniotomy approximately $3cm \times 3cm$ in size. Following placement of the electrode grid, the dura was reapproximated and the bone flap was replaced and affixed to the skull using titanium straps. Electrode leads were tunneled subcutaneously to the chest and exited the skin below the left clavicle. Figures 4.1B and 4.1C show post-operative x-ray images of the implanted electrode grids and the approximate location of the electrodes on the cortical surface[2]. Per U.S. FDA 510(K) regulations, electrode grids were explanted after no longer than 30 days.

4.1.2 Neural recording and preprocessing

Neural signals were recorded and digitized at 1200Hz using the g.USBamp biosignal amplification system (Guger Technologies, Austria) and processed in 33ms blocks, resulting in a 30Hz system update rate. Upside-down (skull-facing) electrodes served as reference and ground electrodes for all recordings (Figure 4.1A). Raw time-domain signals were notch-filtered at 60Hz, 120Hz, and 180Hz to remove power line noise artifacts. Spectral power of the filtered signals was the computed using the Burg autoregressive method [Kay and

---

[2] While electrodes were located primarily over the pre-central gyrus for Subject S2, for Subject S1 electrodes were placed over the post-central gyrus. Though motor cortical areas were initially targeted for this subject, the presence of a venous lake over a substantial region of the the target implantation site necessitated more posterior placement of the electrode grid than originally planned.

Marple, 1981] over the [0Hz – 200Hz] frequency range (25th order, 10Hz frequency bands) using 300ms (subject S1) and 100ms (subject S2) sliding windows. Instantaneous power estimates for each feature were log-transformed and then converted to pseudo Z-scores relative to a baseline resting condition typically collected at the beginning of each testing session [Edwards et al., 2009, Ray et al., 2008b]. All signal processing, neural decoding, and experiment control was performed using Craniux, a LabVIEW-based open-source BMI software suite [Degenhart et al., 2011b] (See Appendix A).

4.1.3 Experimental tasks

Subjects performed two types of experimental tasks: movement screening tasks and closed-loop brain control tasks. Movement screening tasks were performed following grid implantation in order to characterize cortical modulation in response to attempted movements and to identify the command strategy to be used during closed-loop control. Following this, subjects performed brain-control tasks in which they controlled the movement of either a computer cursor or a prosthetic limb using control signals derived from their cortical activity. The following sections describe these tasks in detail.

4.1.3.1 Motor screening task Motor screening tasks were used to identify attempted movements eliciting strong cortical modulation which would ultimately serve as the basis for closed-loop BMI control. Subjects were situated in front of a display and presented with an approximately first-person view of isolated, planar movements performed by either an avatar (Subject S1) or an experimenter (Subject S2). Stimuli consisted of movements of the shoulder, elbow, wrist, and hand; the full stimulus set consisted of shoulder ab/adduction, shoulder flexion/extension, shoulder internal/external rotation, elbow flexion/extension, wrist flexion/extension, wrist pronation/supination, whole-hand grasp, and flexion/extension of individual fingers. Subjects were instructed to attempt to make the movement depicted by the stimulus.

For Subject S1, movements of the avatar were presented in a continuous manner, with joint angle position driven by a 0.5 Hz sinusoid such that each movement phase (e.g. flexion) was 1s in duration. A single movement sequence (stimulus) consisted of 5 repetitions (cycles) of this movement. For Subject S2, individual movements consisted of an initial hold period (2s) the first movement phase (e.g. flexion, 1s in duration), a second hold period (2s), and the second movement phase (e.g. extension, 1s in duration), with an entire movement sequence consisting of 5 repetitions of this. As opposed to the movement sequence used for Subject S1, this movement sequence allowed cortical responses to the individual movement phases to be isolated. For both subjects, individual movement sequences for selected stimuli were presented in a pseudorandom order interleaved with 2s inter-trial interval periods.

4.1.3.2 Cursor control task The majority of brain-control sessions consisted of a center-out cursor control task in a virtual environment [Taylor et al., 2002, Moran and Schwartz, 1999]. During this task, subjects were given control of a “cursor”, rendered as a sphere in a three-dimensional workspace, and instructed to guide this cursor towards spherical targets. The virtual environment utilized a right-handed Cartesian coordinate system where the X-axis pointed to the subjects right, the Y-axis pointed upward, and the Z-axis pointed toward the subject. The center of the cursor was constrained to remain within the workspace boundary at all times.

Subjects performed the cursor control task with both 2D and 3D target configurations (i.e. 2D and 3D tasks), beginning with the 2D task and progressing to the 3D task once satisfactory 2D performance had been achieved. The 2D target configuration consisted of a set of 4 (subjects S1, S2) or 8 (Subject S1 only) uniformly-distributed targets presented in the X-Y plane, while the 3D target configuration (subjects S1, S2) consisted of two 4-target planes at situated different Z-axis coordinates. Cursor movement was constrained to remain in the X-Y plane for the 2D task. Trials consisted of presentation of one pseudo-randomly selected target in the workspace; the subject was required to acquire this target with the cursor in order to complete the trial; trials were considered successful if the cursor overlapped with the target at any point (i.e. no target hold time was enforced). Maximum trial durations of 5s/7s (2D/3D) and 2s/3s were enforced for subjects S1 and S2, respectively, with trials in which the subject was not able to acquire the target before the end of the trial considered as failed trials. The size of the cursor and targets were adjusted in order to control task difficulty.

4.1.3.3 Prosthetic arm control task

Towards the end of the implantation period, subjects performed a prosthetic arm task where they attempted to control the movement of the Modular Prosthetic Limb (Subject S1, The Applied Physics Laboratory, Laurel, MD) [Harris et al., 2011], or the DEKA Arm (Subject S2, DEKA Research and Development Corporation, Manchester, NH). This was intended only as a brief demonstration since a more extensive study was precluded by the limited duration of the protocol.

Subjects controlled the endpoint velocity of the hand, with joint-angle feedback information used to determine the position of the arm in the workspace. For the Modular Prosthetic Limb (MPL), joint angle feedback was converted to 3D position by proprietary software provided by the manufacturer, with the endpoint of the limb defined as a point 2 cm from the center of the palm. For the DEKA Arm, joint angle positions were converted to 3D endpoint position using a 2-link forward kinematics model [Waldron and Schmiedeler, 2008]. Wrist position was either controlled automatically with the hand in a open configuration (Subject S1), or fixed in a neutral position with the hand in a closed-fist configuration (Subject S2).

Targets were arranged in either a 4-target 2D (Subject S1) or 8-target 3D (Subject S2) configuration, and consisted of virtual cubes defined within the workspace of the arm. For both subjects, foam objects were used to indicate the approximate location of targets. For Subject S1, foam blocks were placed behind the virtual targets (e.g. away from the subject) and the subject was instructed to move the palm of the arm such that it was in front of the blocks. For Subject S2, foam balls were presented in the workspace using a custom automatic target presentation system, which used a set of linear stages to advance targets into the workspace. These targets were arranged such that the location of the target ball approximately coincided with that of the virtual target. The position of the arm in the workspace, rather than the collision of the arm with the presented targets, was used to assess task completion.

4.1.4 Neural signal decoding and calibration

Normalized time-frequency data were used as the basis for real-time brain control. Dura-facing electrodes, as well as those containing high amounts of noise indicative of poor electrical connectivity, were removed from the set of electrodes used for control. In addition, the [0–40] Hz frequency range was removed for all electrodes prior to decoding. The resultant feature sets consisted of 448 (28 electrodes × 16 frequency bands) and 432 (27 electrodes × 16 frequency bands) neural features for subjects S1 and S2, respectively.

Intended velocity command signals were predicted from instantaneous feature activities in real-time using Equation 4.1,

$$\hat{\mathbf{v}} = \mathbf{W}\mathbf{f} \tag{4.1}$$

where $\mathbf{W} \in \mathbb{R}^{M \times D}$ is a decoding weight matrix mapping the D-dimensional feature vector $\mathbf{f} \in \mathbb{R}^{D \times 1}$ onto the M-dimensional command velocity vector $\hat{\mathbf{v}} \in \mathbb{R}^{M \times 1}$. For Subject S1, linear regression was used to find $\mathbf{W}$ using Equation 4.2:

$$\mathbf{W} = \mathbf{V}\mathbf{F}^{\dagger} \tag{4.2}$$

where $\mathbf{V} = [\mathbf{v}_1, \ldots, \mathbf{v}_N]$ and $\mathbf{F} = [\mathbf{f}_1, \ldots, \mathbf{f}_N]^\top$ are concatenated matrices of the N time-averaged single-trial observations of the desired movement direction and associated neural feature activity during decoder calibration (see below). The superscript "†" denotes the Moore-Penrose pseudoinverse.

For Subject S2, the Optimal Linear Estimator (OLE) [Salinas and Abbott, 1994, Kass et al., 2005] was used to find weight matrix $\mathbf{W}$. This process began by first fitting an encoding model for each neural feature of the form:

$$f_i = b_0 + \mathbf{b}_i^\top \mathbf{d} + \epsilon \tag{4.3}$$

where f_i is the instantaneous feature activity for the ith neural feature, b_0 is a constant offset term, $\mathbf{d} \in \mathbb{R}^{D \times 1}$ is the intended movement direction, and $\mathbf{b}_i \in \mathbb{R}^{D \times 1}$ is the *preferred direction vector* relating movement direction to neural activity. Preferred direction vectors from all features were collected into a single matrix $\mathbf{B} = [\mathbf{b}_1, \ldots, \mathbf{b}_N]$, from which the decoding weight matrix was found according to Equation 4.4.

$$\mathbf{W} = (\mathbf{B}^\top \mathbf{B})^{-1} \mathbf{B}^\top \tag{4.4}$$

Decoder calibration was performed in different manners for both subjects. For Subject S1, initial decoder calibration was performed using 40 trials in which the cursor was automatically moved to the target under computer control (observation-based calibration). Subsequently, calibration sessions consisted of 16-trial blocks of closed-loop cursor control, during which the subject was instructed to direct the cursor to the target as quickly as possible without correcting for movement errors. In order to ensure neural decoding weights changed gradually during calibration, updated decoding weights were calculated as a weighted sum of the old decoding weights and those estimated from the newly-acquired calibration data (Equation 4.5):

$$\mathbf{W}_{new} = (1 - \alpha)\mathbf{W}_{old} + \alpha\mathbf{W}_{calib} \tag{4.5}$$

where $\mathbf{W}_{old}$ were the set of decoding weights used for control during calibration, $\mathbf{W}_{calib}$ were the set of weights resulting from the current decoder calibration block (Equation 4.4), and $\mathbf{W}_{new}$ were the resultant set of weights to be used for subsequent brain control trials. Parameter α specifies the proportion of the resultant decoder weights contributed by the calibration data ($\mathbf{W}_{calib}$); a value of $\alpha = 0.2$ was used throughout the study. Updating weights in this manner allowed progression from 2D to 3D control tasks to occur seamlessly: decoding weights for the added (e.g. third) dimension were initialized to zero, and the calibration procedure proceeded as previously described.

Calibration was typically performed in 5-block sequences (“rounds”), with decoding weights updated after each block. Following this, the subject performed the closed-loop cursor control task until performance (as determined by success rate) was deemed to plateau, at which point an additional round of decoder calibration was performed. By interleaving decoder calibration and closed-loop control, the subject was trained to use the BMI in a “coadaptive” manner. This allowed adaptation on the part of both the neural decoder and the subject to occur gradually, and attempted to allow these adaptation processes to reinforce, rather than interfere with, one another. Such coadaptative approaches to BMI training have been successfully used in other BMI studies [Taylor et al., 2002, Orsborn et al., 2012].

For Subject S2, decoder calibration occurred more frequently and was commonly performed at the start of each experimental testing session. Calibration sessions were was run in 80-trial (2D) or 96-trial (3D) blocks. As with Subject S1, initial decoder calibration was performed using an observation-based calibration procedure. Subsequent decoder calibration was performed under closed-loop control when possible, though observation-based control was used in cases where closed-loop performance using the most recent decoder was poor. In order to optimize the coadaptation process, we attempted to identify the optimal value of α in Equation 4.5 by minimizing the cross-validated sum-of-squared error over the full range of possible α values. However, in practice this always resulted in an optimal α value of 1.0; thus, updated decoding weights for Subject S2 were solely based on the data collected during the most recent decoder calibration block.

4.1.5 Computer assistance during closed-loop control

Computer assistance was used to facilitate brain control training by attenuating the component of the cursor control signal perpendicular to the vector from the cursor to the target by an experimenter-controlled assist factor [Velliste et al., 2008]. To accomplish this, the instantaneous movement control signal was decomposed into two components: one pointing towards the target and a second, perpendicular to the target direction, representing the instantaneous movement error (see Equations 4.6 and 4.7). The updated (assisted) control signal is then found according to Equation 4.8.

$$\mathbf{v}_c = \mathbf{v}_t + \mathbf{v}_e \quad , \quad \mathbf{v}_t \perp \mathbf{v}_e \tag{4.6}$$

$$\mathbf{v}_t = \frac{\mathbf{p}_t - \mathbf{p}_c}{\|\mathbf{p}_t - \mathbf{p}_c\|} \left(\mathbf{v}_t^\top \frac{\mathbf{p}_t - \mathbf{p}_c}{\|\mathbf{p}_t - \mathbf{p}_c\|} \right) \tag{4.7}$$

$$\hat{\mathbf{v}}_c = \mathbf{v}_t + \gamma(\mathbf{v}_c - \mathbf{v}_t) \tag{4.8}$$

Here, $\mathbf{v}_c$ is the instantaneous (unassisted) control signal, $\mathbf{v}_t$ and $\mathbf{v}_e$ are target and error components of the control signal, $\mathbf{p}_c$ and $\mathbf{p}_t$ are the positions of the cursor and target, γ is the assist level, and $\hat{\mathbf{v}}_c$ is the updated control signal. An conceptual illustration of this process is provided by Figure 4.2. At full computer assist ($\gamma = 0$), the cursor is constrained to remain on a line from the starting position to the target, while at $\gamma = 1$, the cursor is under full brain control (i.e. no constraints on cursor movement).

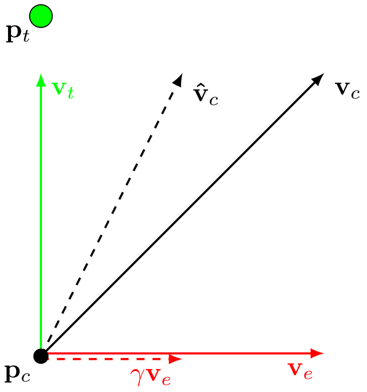

Figure 4.2: *Illustration of computer assistance of cursor movement.* Instantaneous cursor and target positions are indicated by the black and green circles, respectively. The instantaneous control signal vector $\mathbf{v}_c$ is decomposed into components pointing towards the target ($\mathbf{v}_t$) and a perpendicular error vector ($\mathbf{v}_e$). The error vector is attenuated by the assist level γ and added back to the target vector to produce the assisted control signal vector $\hat{\mathbf{v}}_c$.

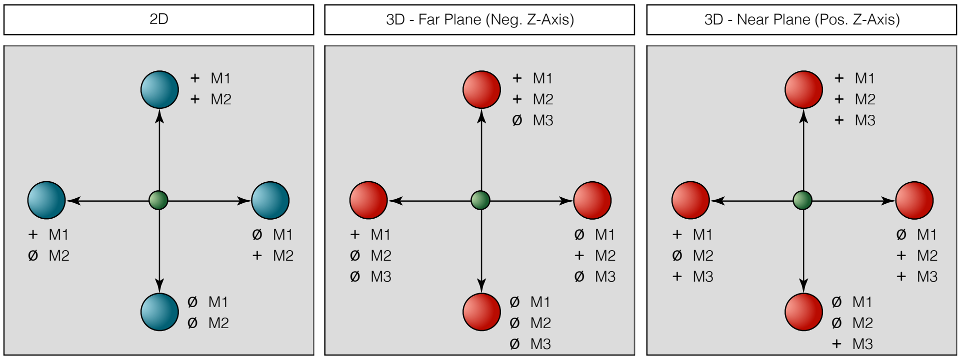

Figure 4.3: *Illustration of the somatotopic control strategy.* Circles represent target directions, with the associated attempted movement shown next to the target (e.g. 'M1', 'M2', 'M3'). '+': attempted movement, 'ø': relaxation. Strategies are shown for both two-dimensional (*left*) and three-dimensional (*middle, right*) cursor tasks. For the sake of clarity, control strategies for the 3D cursor task has been broken into that used for targets in the far plane (away from the subject, negative Z-axis), and that used for targets in the near plane (towards the subject, positive Z-axis).

The level of computer assistance provided to the subjects was adjusted depending on performance. For Subject S1, assistance was used only during initial attempts at brain control, while for Subject S2, computer assistance was used more regularly throughout brain control sessions. In general, for Subject S2 the assist level was set to 1.0 immediately following decoder calibration, and was reduced by a set amount (typically 0.1 or 0.25) if the success rate over a block of trials was at least 70%. Controlling the assist level in this manner attempted to keep the subject motivated while still providing sufficient visual feedback of the subject's actual performance to allow for improvements in performance.

4.1.6 Determination of BMI control strategies

In order to generate directionally-modulated cortical activity to serve as the basis for closed-loop control, subjects were instructed to use a somatotopic control strategy in which they associated attempted arm and hand movement with desired cursor movement direction. Time-frequency responses during the motor screening task (see Section 4.1.3.1) were examined to identify those movements eliciting robust gamma-band modulation. A subset of movements were then chosen which exhibited spatially-distinct patterns of gamma-band activity. Once three such movements were identified, they were mapped onto the cursor movement workspace as shown by Figure 4.3.

The first two movements were arranged in a "push–pull" configuration to control movement of the cursor along the X-axis such that attempting Movement 1 would move the cursor in the positive X-direction and attempting Movement 2 would move the cursor in the negative X-direction. Movement in the positive Y-direction was generated by attempting Movements 1 and 2 simultaneously, while movement in the negative Y-direction movement was generated by relaxing (i.e. no movement). For three-dimensional control, Movement 3 was directly mapped to the Z-axis such that attempting Movement 3 would move the cursor in the positive Z-direction. Alternative somatotopic mappings have been used for EEG control which utilize a one-to-one correspondence between attempted movements and cardinal movement directions (e.g. Movement 1: Positive X-direction, Movement 2: Positive Y-direction) [Foldes and Taylor, 2013]. However, such mappings result in complicated movement combinations for movements along the cardinal axis directions, potentially increasing the cognitive burden on the subject. By orienting attempted movements in the manner described by Figure 4.3, it was hoped that movement along the cardinal axis directions placed the lowest cognitive burden on the subject possible.

4.1.7 Characterization of brain-controlled cursor movement

A number of performance metrics were used to characterize different aspects of brain control performance, including success rate, corrected success rate, time to target, path efficiency, movement error, and fraction of time spent at the boundary of the workspace. Success rate was defined as the number of successful trials divided by the total number of trials. The corrected success rate was calculated in a similar manner, but with successful trials defined as the number of trials in which the presented target was acquired before the cursor would have made contact with any of the other possible targets would all targets have been presented simultaneously. Only the success rate metrics were calculated using all trials; the remaining metrics were calculated using data from successful trials only.

Time to target was defined as the average time from the onset of brain control until first contact with the target. Path efficiency (also referred to as the *distance ratio* [Simeral et al., 2011]) was defined according to Equation 4.9,

$$PE = \frac{1}{N}\sum_{i=1}^{N}\frac{1}{\|\mathbf{p}_t^i - \mathbf{c}\| - r_t}\sum_{j=1}^{M_i-1}\|\mathbf{p}_c^{j+1} - \mathbf{p}_c^j\| \tag{4.9}$$

where PE is the path efficiency, N is the number of trials, M_i is the number of time points in the ith trial, $\mathbf{p}_c^j$ is the position of the cursor at the jth time point, $\mathbf{p}_t^i$ is the position of the target for the ith trial, $\mathbf{c}$ is the center of the workspace, and r_t is the radius of the target. Lower path efficiency scores indicate more accurate trajectories; a path efficiency score of 1 indicates a perfectly straight trajectory. Movement error was defined as the mean perpendicular distance of the cursor position to the ideal straight-line trajectory normalized by the distance to the target, and was calculated according to Equation 4.10,

$$ME = \frac{1}{N}\sum_{i=1}^{N}\frac{1}{M_i(\|\mathbf{p}_t^i - \mathbf{c}\| - r_t)}\sum_{j=1}^{M_i}\left\{\|\mathbf{p}_c^j - \mathbf{c}\|^2 - \left[(\mathbf{p}_c^j - \mathbf{c})^\top \frac{\mathbf{p}_t^i - \mathbf{c}}{\|\mathbf{p}_t^i - \mathbf{c}\|}\right]^2\right\}^{\frac{1}{2}} \tag{4.10}$$

Both path efficiency and movement error metrics seek to quantify the deviation of cursor trajectories from an ideal, straight-line path. Finally, the boundary fraction, indicating the fraction of time in which the cursor was at the edge of the workspace boundary, was calculated as the number of time points in which the cursor was at a workspace boundary divided by the total number of time points. This metric provided a means with which to assess subjects' reliance upon the workspace boundary constraint to successfully complete the task.

4.2 RESULTS

4.2.1 Cortical Activity during motor screening

ECoG signals recorded from the sensorimotor cortex of both subject demonstrated robust modulation during attempted arm and hand movement, even in the absence of overt movements. Figures 4.4 and 4.5 show time-frequency responses for all recording electrodes during selected attempted movements for subjects S1 and S2, respectively. Typical event-related synchronization (ERS, increase in spectral power) of the gamma/high-gamma frequency bands and event-related desynchronization (ERD, decrease in spectral power) for the sensorimotor rhythm (10–30 Hz), tightly time-locked to stimulus onset, was observed for a number of attempted movements. Spatial patterns of gamma/high-gamma modulation followed the expected somatotopic organization of motor (Subject S2) and sensory (Subject S1) cortices, with the centroids of activity for distal movements medial and those of proximal movements lateral on the pre-central (S2) and post-central (S1) gyri.

Once cortical responses to attempted arm and hand movements were examined, a subset of movements were chosen to serve as the basis for cursor control using a somatotopic control strategy as described in Section 4.1.6. For Subject S1, attempted hand flexion and elbow flexion were chosen for two-dimensional control (Figure 4.6, top row, first column) and attempted hand flexion, elbow flexion, and wrist flexion were chosen to serve as the basis of three-dimensional control (Figure 4.6, top row, second and third columns). For Subject S2, attempted thumb flexion and middle finger flexion were chosen for two-dimensional control (Figure 4.6, bottom row, first column) and attempted thumb flexion, little finger flexion, and elbow flexion were chosen to serve as the basis of three-dimensional control (Figure 4.6, bottom row, second and third columns). For Subject S1, the control strategy used for 2D control was preserved when transitioning to 3D control; though this was attempted for Subject S2 poor 3D performance necessitated using a new set of attempted movements for 3D control (see Section 4.2.2).

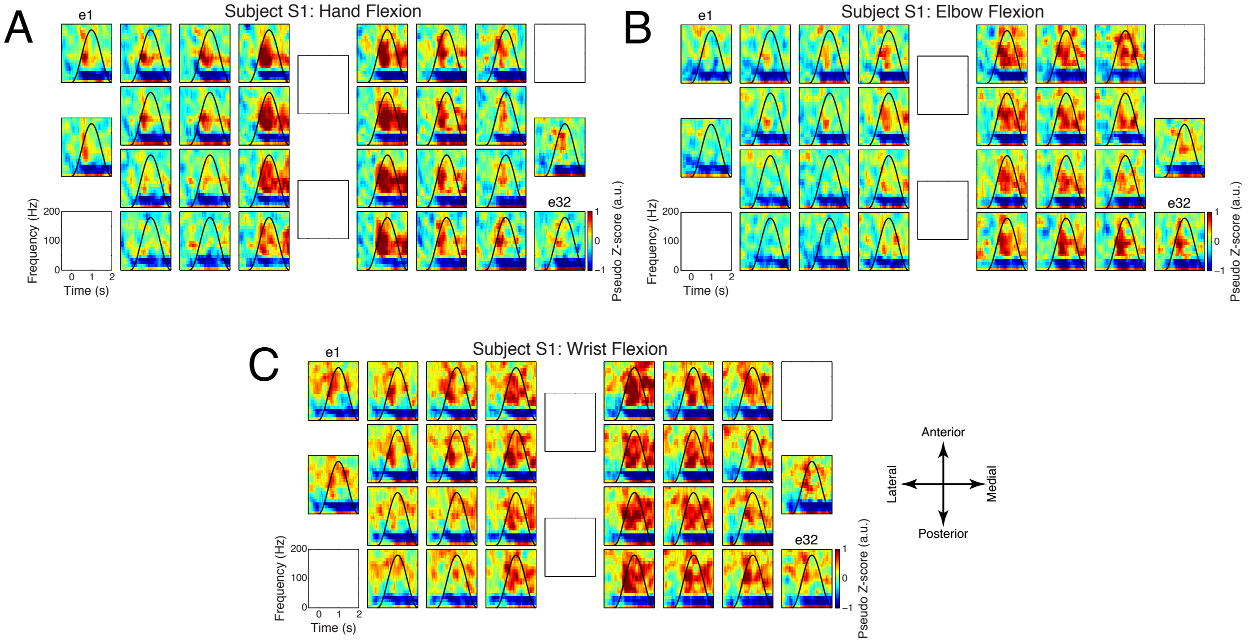

Figure 4.4: *Averaged electrode grid time-frequency responses across all electrodes for Subject S1 for selected attempted movements.* Each plot shows the time-frequency response for a single electrode averaged across repetitions of one attempted movement. Instructed kinematic profiles are indicated by the *black* line in each plot. (*A*) Hand flexion/extension. (*B*) Elbow flexion/extension. (*C*) Wrist flexion/extension. Electrode layout and numbering are as depicted in Figure 4.1.

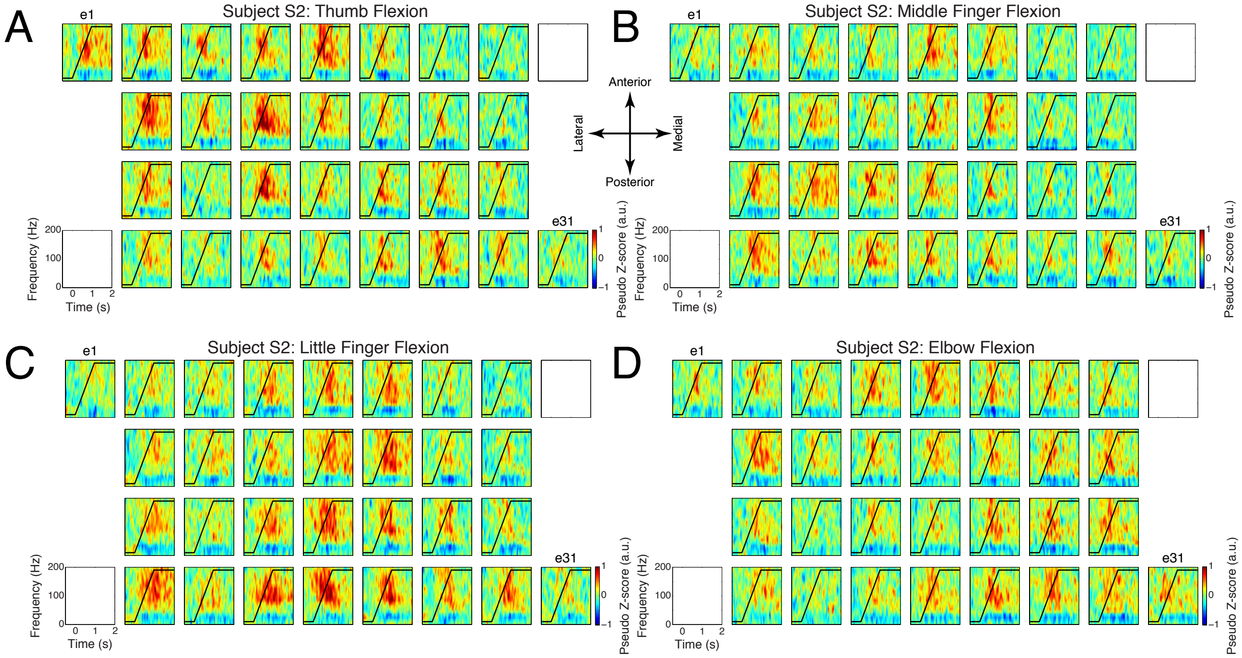

Figure 4.5: *Averaged electrode grid time-frequency responses across all electrodes for Subject S2 for selected attempted movements.* Each plot shows the time-frequency response for a single electrode averaged across repetitions of one attempted movement. Instructed kinematic profiles are indicated by the *black* line in each plot. (*A*) Thumb flexion/extension. (*B*) Middle finger flexion/extension. (*C*) Little finger flexion/extension. (*D*) Elbow flexion/extension. Electrode layout and numbering are as depicted in Figure 4.1. Note that the qualitative differences in time-frequency responses for Subject S2 as compared to Subject S1 (Figure 4.4), particularly the time course and strength of modulation are in large part due to the different spectral estimation window lengths used for the two subjects.

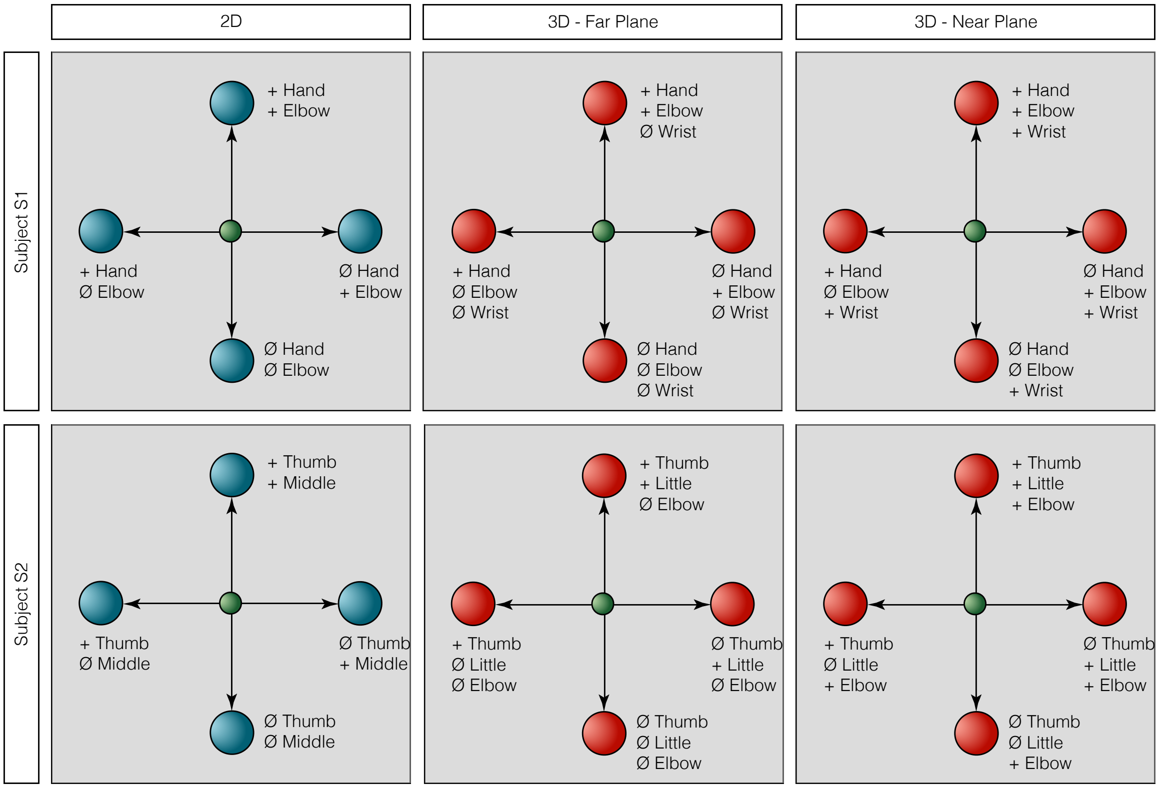

Figure 4.6: *Control strategies utilized to obtain closed-loop BMI control.* Circles represent target directions, with the associated attempted movement shown next to the target. '+': attempted movement, 'ø': relaxation. Strategies are shown for both two-dimensional and three-dimensional cursor tasks. For the sake of clarity, control strategies for the 3D cursor task has been broken into that used for targets in the far plane (away from the subject, *middle* column), and that used for targets in the near plane (towards the subject, *right* column).

4.2.2 Cortical control of cursor movement

The general progression of the cursor control task, including success rate, is shown by Figure 4.7. Subjects began with a 4-target 2D task and progressed to an 8-target 3D task once satisfactory 2D control was achieved. Subject S1 also performed an 8-target 2D task between 4-target 2D and 8-target 3D control; this was not performed for Subject S2 in favor of transitioning to 3D control as quickly as possible[3].

For Subject S1, control performance using the decoder calibrated on the first day of brain control (day 15) decreased as the computer assist level was decreased (Figure 4.7A). Only after removing the assist altogether and re-calibrating the decoder on Day 19 were improvements in performance observed. Additional calibration sessions on days 20 and 24 facilitated further improvements in performance. On day 24, the subject transitioned to 3D cursor control. Three-dimensional brain control was built upon the subject's existing 2D control capability in two ways. First, the existing association between attempted movement and 2D cursor movement direction was preserved, while a third attempted movement, wrist flexion/extension, was added to control cursor movement along the Z-axis. Second, the existing set of decoding weights used during 2D control was used during decoder calibration for the 3D task (see Equation 4.5). This attempted to prevent the decoding weights from changing rapidly during initial periods of 3D control, and ensure the control strategy used for 2D control was preserved during 3D control. It is worth noting that for Subject S1, 3D cursor control performance declined after testing day 25, likely due to a change in the control strategy employed by the participant. On testing day 26, the subject reported that they had begun using cursor movement imagery (i.e. imagining the size of the cursor increase/decrease) to drive movement of the cursor in the Z-dimension; performance during this period was poor. Though efforts were made to transition back to the original 3D control strategy, the subject was not able to re-establish satisfactory control.

Subject S2 also began with a 4-target 2D cursor control task. Unlike Subject S1, the assist level was reduced much more rapidly and the decoder re-calibrated more often for Subject S2 in an effort to increase the rate at which the subject acquired control (Figure 4.7B). Early efforts at control utilized a thumb/index finger control strategy; after 4 days of mediocre performance this was changed to a thumb/middle finger strategy on day 11, which yielded improved performance. The subject transitioned to a 3D task on Day 13, but required two additional changes in control strategy before peak 3D performance was achieved using a thumb/little finger/elbow control strategy on day 25. Unlike Subject S1, initial plans for closed-loop control with subject S1 were to use three different finger movements as the basis for 3D control, with attempted elbow movement serving as the basis for the grasp dimension of a 4 degree-of-freedom reach-and-grasp prosthetic arm task. However, attempts at achieving satisfactory 3D control using only attempted finger movement were ultimately unsuccessful, and this strategy was abandoned in favor of the thumb/little finger/elbow control strategy used at the end of the 3D control testing.

---

[3] As Subject S1 was the first subject participating in the study, we were unsure whether 3D cursor control would be possible using ECoG. Thus, we elected to fully evaluate Subject S1's 2D cursor control ability using an 8-target task in the event that 3D control was unsuccessful. For Subject S2, we were confident of the ability of the subject to achieve 3D control (based on the results from Subject S1) and thus chose to forgo the 8-target 2D task.

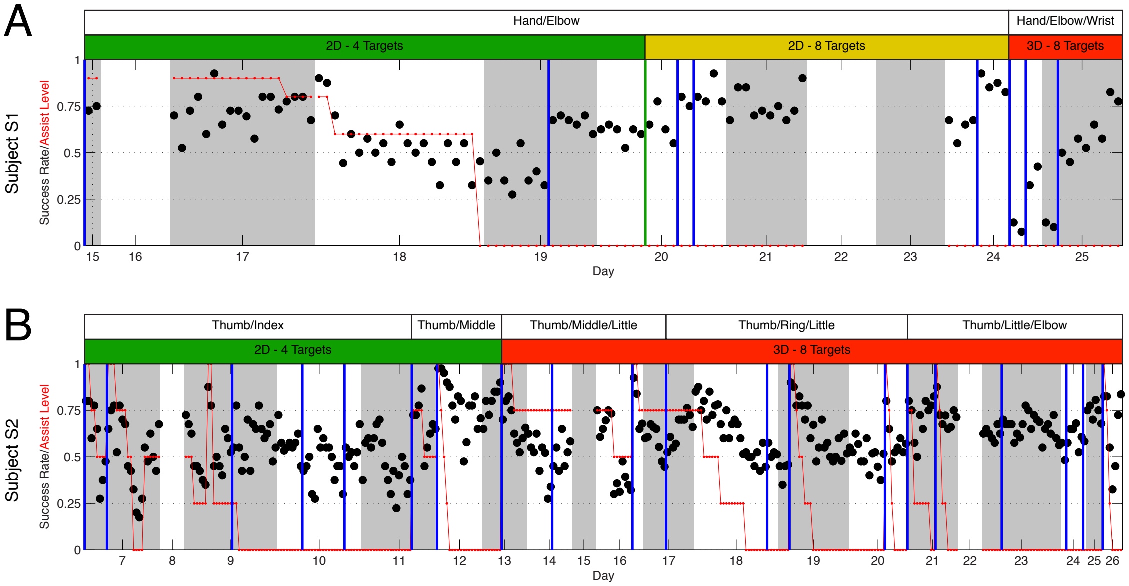

Figure 4.7: *BMI control performance across days for subjects S1 (A) and S2 (B).* BMI control success rate (*black*) and computer assist level (*red*) is plotted as a function of testing day. Each point represents a single "block" of closed-loop control consisting of 40 trials. Alternating white and gray regions mark individual days, while vertical blue lines mark the occurrence of neural decoder calibration. White bars above each panel indicate the somatotopic control strategy used for that period, while the green, yellow, and red bars indicate the task difficulty. Days without success rate data (days 16, 22, and 23 for Subject S1, days 8, 15, 22 for Subject S2) were planned days off.

Averaged cursor movement trajectories for subjects S1 and S2 are shown by Figure 4.8. Qualitatively, 2D trajectories were relatively straight for both subjects, with the exception of those towards the left target for Subject S2, which exhibited a curved trajectory. Curvature in the movement trajectories was more pronounced during 3D control. In particular, cursor movements during 3D control generally began with movement in the X–Y plane before moving in the Z-dimension. This suggests that subjects utilized a "step-wise" strategy in which they first attempted to move the cursor to the appropriate location in the X–Y plane before the Z-dimension. This may have been due to subjects' greater familiarity with the 2D control strategy, or an effort on the part of the subjects to reduce the cognitive burden of the task by performing movement in the X–Y plane and Z-dimension in a sequential manner.

As expected, increases in the spectral power of the gamma/high-gamma range and decreases in the sensorimotor rhythm were observed surrounding the onset of cursor movement. Time-frequency responses during closed-loop cursor control reflected the instructed control strategies; several representative examples of such responses are shown by Figure 4.9 for selected electrodes from subjects S1 and S2. For example, time-frequency responses for electrode 15 exhibited increases in gamma-band activity for targets in the upper-right quadrant of the workspace (Figure 4.9B). This is consistent with the response of electrode 15 during attempted elbow flexion/extension (Figure 4.4B), and the instructed control strategy (Figure 4.6, top-left), which associated attempted elbow flexion with movements to the upward and right targets. Consistent with the subjects' ability to successfully control the cursor, pairs of electrodes were observed which exhibited modulation for different target directions. For Subject S1, time-frequency data from electrode 5 (Figure 4.9A) and electrode 15 (Figure 4.9B) were preferentially modulated for the upper-left and upper-right quadrants of the workspace, respectively. For Subject S2, electrodes 5 (Figure 4.9C) and 27 (Figure 4.9D) showed exhibited similar responses.

Performance metrics for 2D and 3D cursor control are summarized in Table 4.1. Though Subject S1 achieved marginally higher success rates than Subject S2, Subject S2 performed substantially more trials. Success rates for Subject S1 were 0.87 (2D, calculated over 176 trials) and 0.71 (3D, 160 trials), while those for Subject S2 were 0.78 (2D, 680 trials) and 0.68 (3D, $2,160$ trials). Corrected success rates for Subject S1 (0.82/0.70, 2D/3D) and S2 (0.77/0.48) suggest that Subject S1 was capable of more accurate cursor control; path efficiency, movement error, and boundary fraction values corroborate this, as all metrics are higher for Subject S2 than S1, indicating greater deviation from straight-line trajectories. However, time-to-target was substantially higher for Subject S1 ($2.27s/3.25s$, 2D/3D) than for S1 (0.72/0.86). Figure 4.10 shows success rate, time-to-target, and path efficiency broken down by target. Non-uniformity across targets was observed across these metrics, indicative of difficulty by the subjects to reach specific targets. For example, Subject S1 had difficulty acquiring the rightmost target in the workspace (Target 1) during 2D control, as evidenced by a lower success rate, longer time-to-target, and higher path efficiency for this target relative to others. Similarly, Subject S2 had difficulty acquiring the leftmost target (Target 3) during 2D control; this is also evidenced by the curved trajectory towards this target (Figure 4.8, second row, first column).

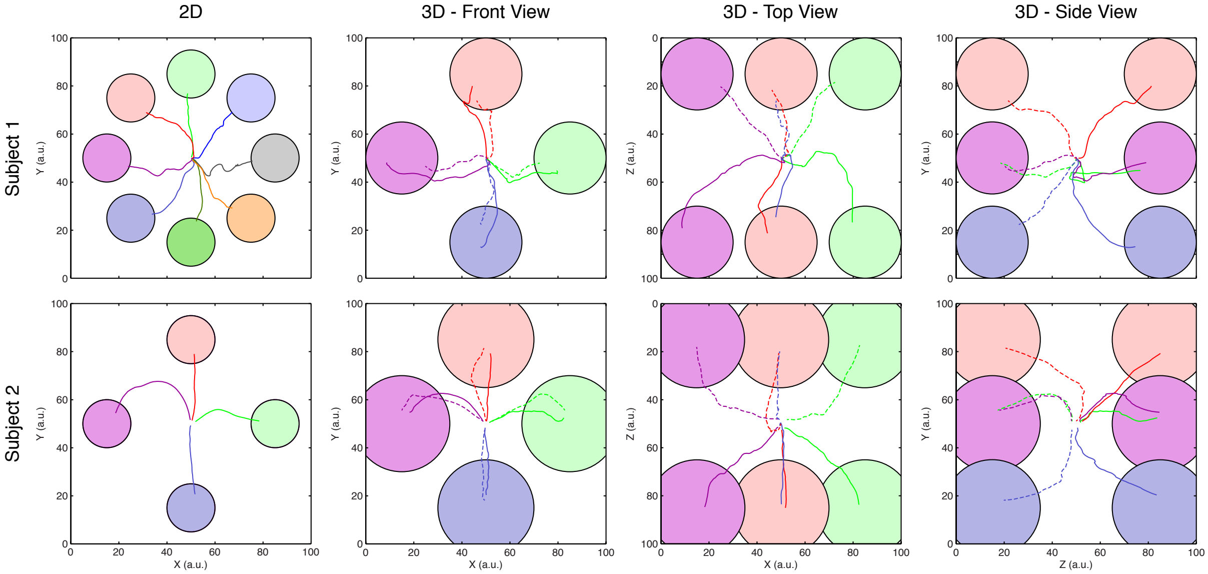

Figure 4.8: *Average trajectories during two and three-dimensional cursor control.* Averaged trajectories (successful trials only) are shown for subjects S1 (top row) and S2 (bottom row) for the 2D (left column) and 3D (middle and right columns) tasks. Colors of the individual trajectories correspond to their respective targets. Trajectories for the 3D task have been separated into three views (front, top, and side) for the sake of clarity. Trajectories towards targets in the far plane of the workspace are indicated by dashed lines. Note that each 3D view results in the obstruction of one or more targets; while the targets themselves are obscured the trajectories are plotted as normal. Thus, some pairs of trajectories appear to terminate at the same target (e.g. all pairs of trajectories toward each target in the *middle-left* column); one of each of the pairs of these trajectories correspond to the obstructed (non-visible) target.

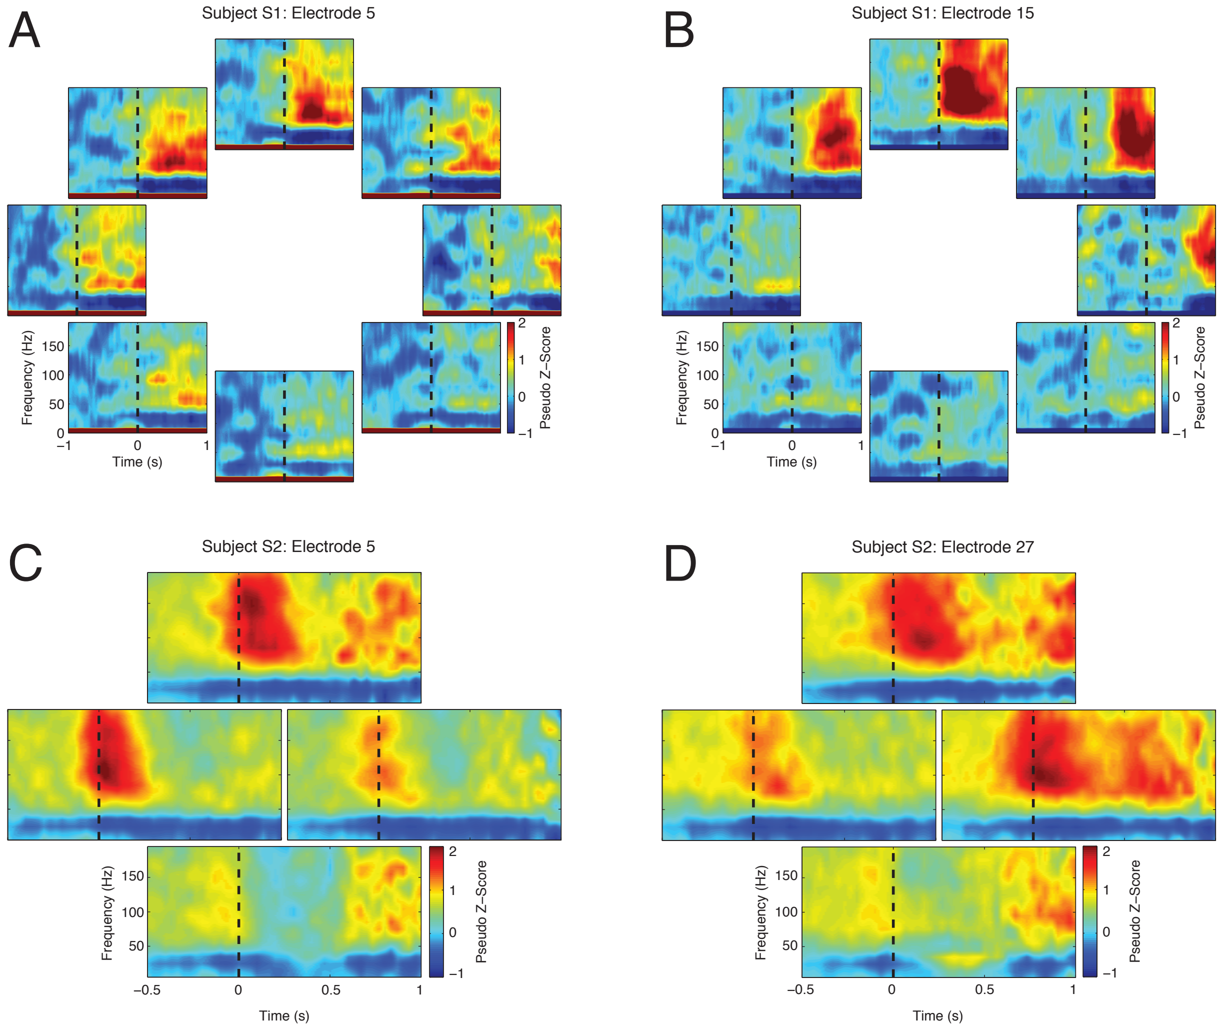

Figure 4.9: *Averaged time-frequency data for selected electrodes during two-dimensional cursor control.* Individual panels represent time-frequency data for a single electrode, averaged across trials, towards a single target. Responses are shown for the following: (*A*) Subject S1, electrode 5, (*B*) Subject S1, electrode 15, (*C*) Subject S2, electrode 5, (*D*) Subject S2, electrode 27. The layout of the time-frequency plots corresponds to the position of the targets in the workspace (see Figure 4.8, *left* column). Dashed black lines indicate the onset of cursor control.

Table 4.1: *Cursor control performance metrics.* Number of trials, success rate, corrected success rate, time-to-target, path efficiency, movement error, and boundary fraction are provided for selected peak performance periods for 2D and 3D control. Data for time-to-target, path efficiency, movement error, and boundary fraction are presented as mean ± standard deviation.

	Subject S1		**Subject S2**	
	2D	**3D**	**2D**	**3D**
Trials	176	160	680	2160
Success rate	0.87	0.71	0.78	0.68
Success rate (corrected)	0.82	0.70	0.77	0.48
Time to target (s)	2.27 ± 1.02	3.25 ± 1.31	0.72 ± 0.45	0.86 ± 0.71
Path efficiency	1.95 ± 0.74	2.45 ± 0.95	1.97 ± 0.98	5.11 ± 3.96
Movement error	0.20 ± 0.09	0.30 ± 0.15	0.28 ± 0.22	0.70 ± 0.38
Boundary fraction	0.00 ± 0.00	0.01 ± 0.04	0.01 ± 0.03	0.17 ± 0.21

4.2.3 Prosthetic limb control

After achieving satisfactory 3D cursor control, subjects performed closed-loop control of either the MPL (Subject S1) or the DEKA Arm (Subject S2). For Subject S1, arm control was performed on the last testing day prior to explantation of the electrode grid (day 27), and was limited to a brief demonstration consisting of 15 trials of a two-dimensional center-out task. Successful trajectories from these trials are shown by Figure 4.11; the subject was able to achieve a success rate of 0.467 (7/15 trials). Though performance was poor, 3D cursor control during testing days 26 and 27, the decoding weights from which were used for arm control, was also poor as a result of a change in the control strategy by the participant[4].

Subject S2 performed DEKA Arm control sessions on testing days 24, 25, and 26. For these three days, the subject began with a 3D cursor control task; once satisfactory control had been achieved (typically after $200 - 300$ trials), the subject was transitioned to arm control. As with Subject S1, decoding weights obtained during 3D cursor control were used for arm control. Computer assistance was used to aid the acquisition of arm control. Though performance was relatively high during computer assisted trials, the subject was only able to achieve a maximum unassisted success rate of 0.417 over 48 trials; trajectories for these trials are shown in Figure 4.12. Though performance for targets in the near target plane was relatively poor, acquisition of of the right, left, and downward targets in the far target plane was fairly reliable.

---

[4] See Section 4.2.2.

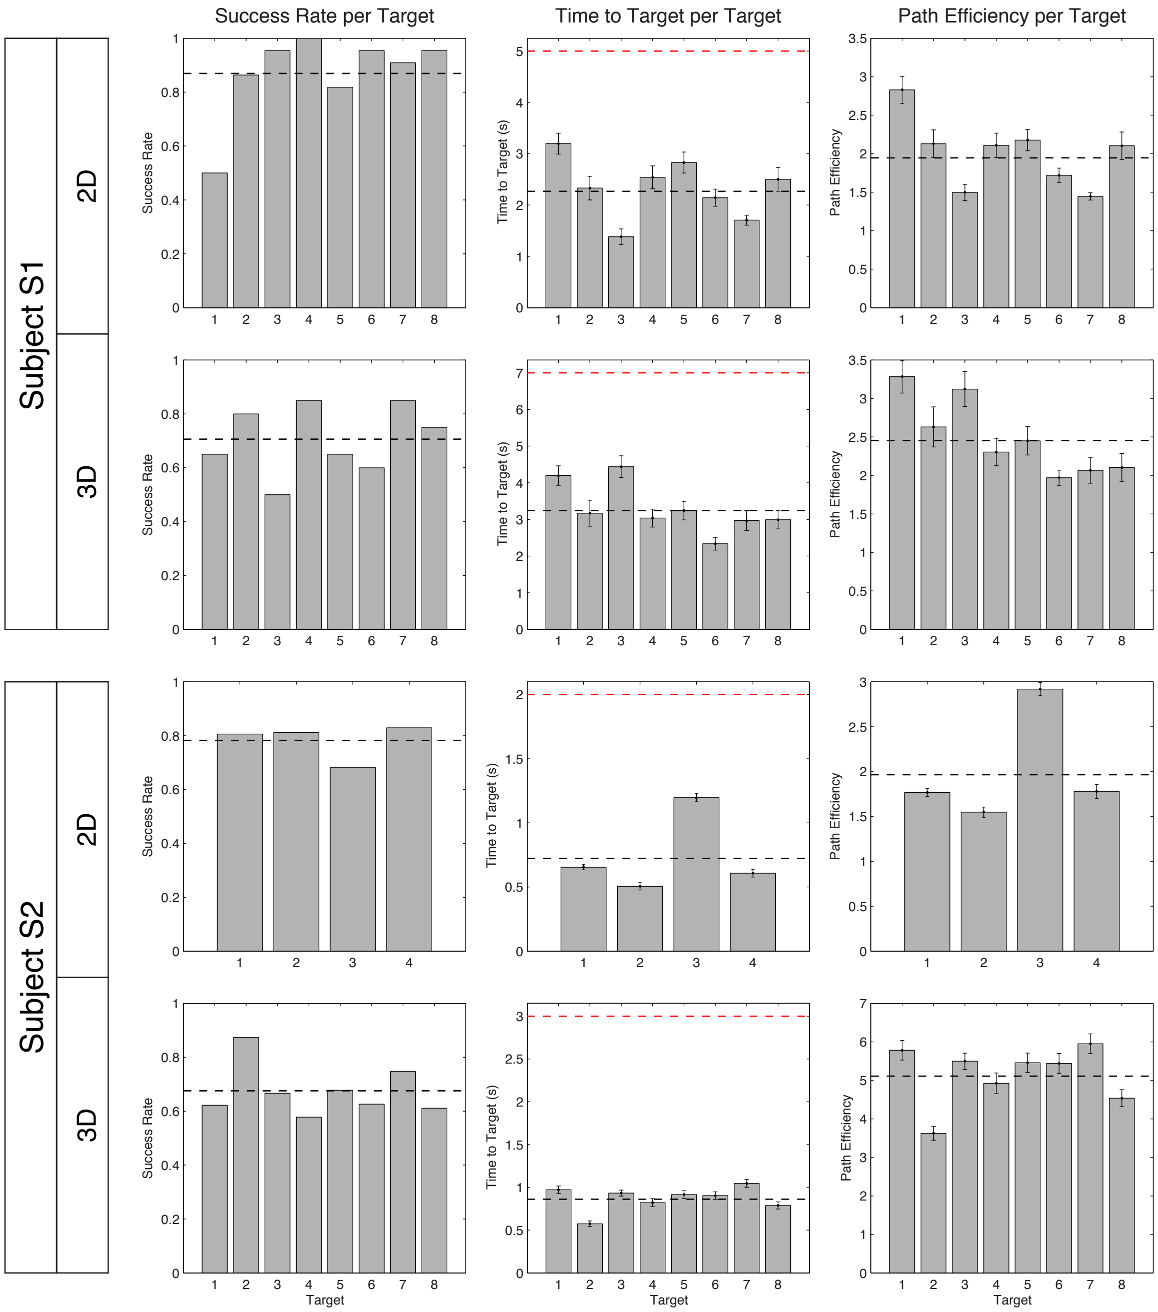

Figure 4.10: *Performance metrics by target location during closed-loop cursor control.* Success rate (*left* column), time to target (*middle* column), and path efficiency per target (*right* column) are shown for subjects S1 and S2 during 2D and 3D control. Dashed black lines indicate the average across all targets (all plots), while dashed red lines (time-to-target plots only) indicate the maximum trial duration. Error bars (time-to-target and path efficiency plots) indicate mean ± SEM. Differences in time to target and path efficiency metrics were statistically significant for all conditions (one-way ANOVA, $p < 0.01$).

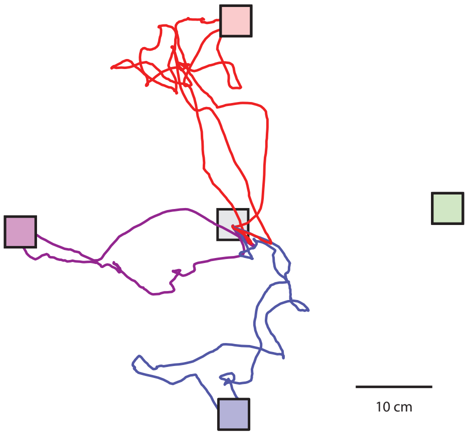

Figure 4.11: *Arm control trajectories for Subject S1.* The subject was able to achieve a success rate of 0.467 over 15 trials of unassisted arm control.

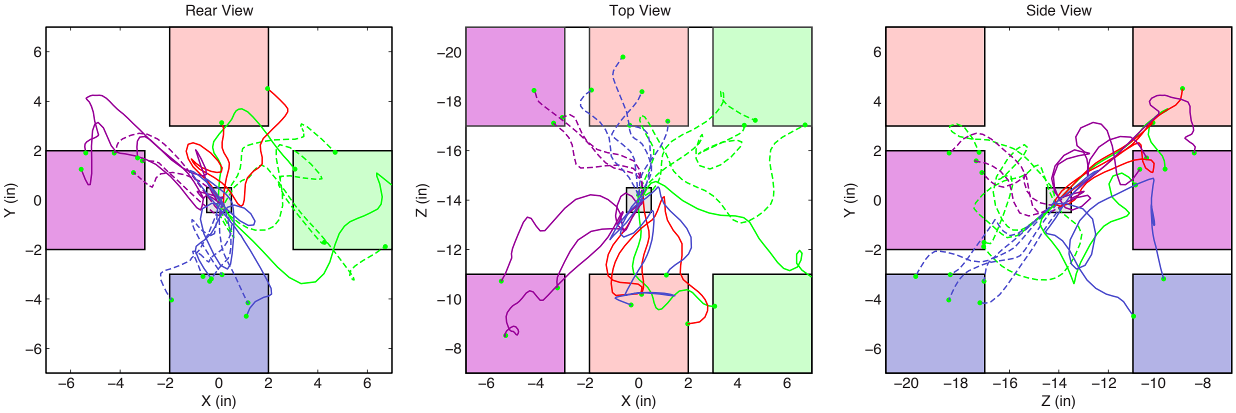

Figure 4.12: *Arm control trajectories for Subject S2.* The subject was able to achieve a success rate of 0.417 over 48 trials. Trajectories towards the far plane of targets (i.e. away from the subject) are indicated by dashed lines.

4.3 DISCUSSION

The presented study investigated the feasibility of an ECoG-based BMI in two individuals with upper limb paralysis. We show that subjects were capable of voluntarily modulating cortical activity in motor and somatosensory cortex during attempted movement, and that, consistent with previous fMRI studies, the somatotopic orgainzation of this activity was preserved from that exhibited by able-bodied individuals [Cramer et al., 2005, Shoham et al., 2001]. High-gamma band activity, believed to represent local neuronal population activity [Crone et al., 2006, Miller et al., 2009a], was found to be tightly coupled to attempted arm and hand movement, as expected from previous reports of motor cortical neuronal activity recorded with intracortical microelectrode arrays in individuals with tetraplegia [Hochberg et al., 2006, Hochberg et al., 2012, Truccolo et al., 2008, Collinger et al., 2012]. Furthermore, we have shown that subjects were able to modulate their cortical activity using a somatotopic control strategy in order to achieve three-dimensional cursor control. Finally, and perhaps most importantly, the entirety of the study, including implantation and explantation of the electrode grid, was performed without complications.

When comparing control performance between subjects, we found that Subject S1 was more accurate (lower path efficiency, movement error, and boundary fraction) while Subject S2 was faster (shorter time-to-target). While it may be tempting to attribute such differences in performance to methodological differences between subjects, it is difficult to draw conclusions from such a comparison. Performance differences may be the result of any number of factors, including grid placement, decoder calibration, target size, cursor gain, subject instruction or motivation. Future work is needed in order to determine the effect of these factors on BMI control quality.

Interestingly, we find that both somatosensory (Subject S1) and motor (Subject S2) cortical activity can be used to control an ECoG-based BMI system. The ability of Subject S1 to voluntarily modulate somatosensory cortical activity is particularly notable, if not surprising. Activation of both pre and post-central gyri is often observed in individuals with chronic spinal cord injury during attempted movement [Cramer et al., 2005, Shoham et al., 2001, Hotz-Boendermaker et al., 2008] and in able-bodied individuals during motor imagery in the absence of overt movement [Miller et al., 2010, Christensen et al., 2007, Lacourse et al., 2005, Porro et al., 1996]. Such somatosensory cortical activity may represent efferent copies of motor control signals [Christensen et al., 2007, Crapse and Sommer, 2008, Gritsenko et al., 2007], or reflect engagement of sensory imagery [Hotz-Boendermaker et al., 2008]. Somatotopic organization of sensory cortex has been shown in both humans [Penfield and Rasmussen, 1950, Penfield and Jasper, 1954, Penfield and Boldrey, 1937] and non-human primates [Merzenich et al., 1978]. More recent work has shown that functional magnetic resonance imaging reveals somatotopic organization of both primary motor and somatosensory cortex during hand movements, with somatosensory cortex exhibiting less overlap between cortical volumes responsive to different finger movements [Hluštík et al., 2001]. This suggests that somatosensory cortex may actually be a more ideal source of cortical activity for a BMI based on a somatotopic control strategy than primary motor cortex, though more work is needed in characterizing the relative organization and discriminability of activity in these regions during attempted movement.

Previous studies have demonstrated ECoG-based BMI control by able-bodied individuals undergoing presurgical brain mapping [Schalk et al., 2008c, Leuthardt et al., 2004, Miller et al., 2010]. To our knowledge, the brain control results presented here are the first demonstration of continuous and porportional three-dimensional BMI control using ECoG. We believe that our results can be attributed to several factors. First, we utilized custom "high-density" electrode grids with both smaller electrode diameter and inter-electrode spacing than standard ECoG grids. This design allows for the sampling of more localized neuronal populations on account of the smaller electrode area and greater spatial resolution on account of the decreased inter-electrode distance. We believe that combined, these characteristics allow for better sampling of the somatotopic organization of sensorimotor cortex. Such sampling is likely critical for the development of high degree-of-freedom BMIs utilizing somatotopic control strategies, where cortical activity related to a number of different attempted movements must be discriminated.

Second, the decoding algorithms and decoder calibration process used in the presented work likely contributed to the ability of subjects to obtain closed-loop control. Whereas other ECoG BMI studies have utilized manual weight assignment to map neural features onto BMI command signals [Leuthardt et al., 2004, Schalk et al., 2008c], we took advantage of population-level decoding algorithms, such as the Optimal Linear Estimator (OLE). Such algorithms are able to learn decoding parameters from calibration data without the need for experimenter intervention. As these methods utilize all available neural features, they are able to take advantage of the redundant encoding of attempted movements in the population, generating output control signals as weighted summations of all input features and thereby increasing the signal-to-noise ratio of the predicted output. In addition, the co-adaptive training process used, which alternated between periods of decoder calibration and fixed-decoder BMI practice, likely facilitated incremental learning by the participant [Ganguly and Carmena, 2009] and gradual refinement of decoder weights. Previous work has shown that incremental updates to decoding parameters over short time scales results in both a convergence of parameters and improvements in performance [Orsborn et al., 2012]. Though the co-adaptative process we utilized differed from this in that we updated parameters less frequently and over longer time scales, we nonetheless observe that performance improvements are driven by both decoder calibration and fixed-decoder closed-loop training. A more thorough examination of the contributions of decoder calibration and fixed-decoder training to improvements in ECoG BMI performance is presented in Chapter 5.

Finally, the duration of the study likely played a substantial role in the ability of subjects to obtain closed-loop control. Previous demonstrations of two-dimensional cursor control using ECoG in patients undergoing presurgical mapping have been limited to less than an hour of training time [Schalk et al., 2008c]. In contrast, subjects in the presented study performed BMI control experiments for as many as 17 testing days, increasing their exposure to the BMI task and providing an opportunity for performance improvements resulting from the aforementioned co-adaptive training process to be retained across experimental sessions. Though our study was not specifically designed to assess the stability of ECoG-based BMI control over time, we observed that decoding parameters could often be retained from one day to the next with no discernible change in performance. This finding is in agreement with other work showing stable offline prediction of movement kinematics from ECoG recordings for up to 9 months with fixed decoding parameters [Chao et al., 2010]. However, further work is needed in order to characterize the nature of day-to-day instabilities in chronic ECoG recordings.

Unfortunately, our efforts at translating closed-loop cursor control to the control of a prosthetic limb yielded limited success. Increases in cognitive load have been shown to reduce BMI control performance [Foldes and Taylor, 2013]. It may be the case that the presence of the physical limb and targets presented additional distractions, potentially increasing cognitive burden placed on subjects. However, a more likely cause for the poor performance may be a conflict between cortical activity during attempted movement and that occurring while observing movements of the prosthetic limb. Congruent activity between action observation and action execution has been demonstrated in ECoG recordings [Collinger et al., 2014], and well as in a number of different recording modalities (see Chapter 2.2.1). When using a somatotopic control strategy based on attempted arm and hand movements to control a prosthetic limb, attempted movements used to control the limb will inevitably differ from the resultant movements which are observed by the subject. Further work is required investigating the effect of such discrepancies on cortical activity and its implications for somatotopic BMIs.

The difficulties encountered during prosthetic limb experiments allude to problems which may arise in the development of a generalizable BMI system. Performance of a clinically-viable BMI, particularly those seeking to restore arm and hand function, must be invariant to differences in application and/or context. Differences in motor cortical activity have been observed in goal-directed and non-goal-directed movements [Nishitani and Hari, 2000, Järveläinen et al., 2004]; such differences suggest that a BMI system trained on one task may not necessarily perform well on a different task. It may be the case that a BMI based on a somatotopic control strategy, which breaks the association between observation-related and attempted-movement-related cortical activity, could result in a system which is more generalizable across applications. The extent to which ECoG-based somatotopic BMI control can be generalized through training, however, remains to be seen.

Unfortunately, our findings are not without their limitations. Though the duration of electrode implantation in the presented work is longer that that of previous human ECoG BMI studies, the necessity that electrode grids be removed after 30 days precluded the us from answering a number of questions of interest. As the goal of the study was to demonstrate reliable closed-loop control, in-depth study of the stability of ECoG recordings or a detailed investigation of the learning processes underlying ECoG control were forgone in favor of attempting to maximize the degrees-of-freedom which could be obtained with an ECoG BMI. Despite this, the three degrees-of-freedom obtained in the presented work are lower than that which has been demonstrated using penetrating microelectrodes [Collinger et al., 2012]; further study into the maximum number of independent control signals which can be extracted from motor cortex using ECoG is required. Such efforts may be aided by the use of new electrode technology which increases the spatial resolution of cortical surface recordings [Viventi et al., 2011]. Finally, the capability of ECoG to sustain long-term, stable BMI performance, as has recently been demonstrated for up to a year using multi-unit activity and local field potentials [Flint et al., 2013], must be evaluated.

It is difficult to draw conclusions from a direct comparison between the results obtained from subjects S1 and S2. Such comparisons are limited by the numerous methodological differences between the subjects, including electrode implant location (sensory cortex versus motor cortex), decoder calibration procedure (closed-loop versus observation), and decoding algorithm (linear regression versus OLE). In addition, the cause of limb paralysis also differed between the two subjects (SCI for Subject S1; ALS for Subject S2). It is possible that any one of these differences could have effects on the characteristics of the BMI control achieved by the subjects; additional carefully-controlled studies are required in order to assess the effects of these differences on closed-loop control. Despite these differences, however, the fact that both subjects were capable of successfully achieving closed-loop control provides some evidence for the robustness of ECoG BMI systems.

In conclusion, we have shown that individuals with upper-limb paralysis can successfully control a BMI using ECoG with up to three degrees of freedom. Our results, combined with the promise of ECoG to provide robust, long-term recordings [Chao et al., 2010, Blakely et al., 2009] with relatively low hardware and software requirements suggest that an ECoG BMI system is a viable solution for the restoration of function for individuals with movement paralysis. We believe that further development into novel decoding algorithms, BMI user training approaches, and fully-implantable devices with telemetry [Rouse et al., 2011], will ultimately allow for longer studies with more subjects, further facilitating the translation of this technology to the clinical realm.

5.0 CORTICAL AND DECODER ADAPTATION EFFECTS DURING ELECTROCORTICOGRAPHIC BRAIN-MACHINE INTERFACE SKILL ACQUISITION

In the previous chapter it was shown that individuals with upper-limb paralysis were able to successfully control a brain-machine interface with control signals derived from ECoG recordings from sensorimotor cortex. While this serves as an important demonstration of the feasibility of an ECoG BMI system, we nonetheless desire to advance such systems through the identification of means by which control performance can be improved. In this chapter, we investigate the contribution of decoder calibration and fixed-decoder training to changes in closed-loop control performance, and attempt to characterize the extent of neural adaptation which can occur during acquisition of ECoG-based BMI control.

5.1 BACKGROUND

An understanding of the learning processes underlying the acquisition of closed-loop brain control is essential to advancing the state-of-the-art of BMI systems. Though limited, some studies are beginning to address the concept of learning during BMI skill acquisition. Early work by Fetz and colleagues has demonstrated the ability of non-human primates to learn to control the firing rate of motor cortical neurons when provided with visual or auditory feedback about their instantaneous activity [Fetz and Finocchio, 1975, Fetz and Baker, 1973, Fetz, 1969]. More recently, it has been shown that subjects can learn to control three-dimensional cursor movement using population SU/MUA activity [Taylor et al., 2002]. In this study, improvements in performance were observed both within and across testing days, and were facilitated both by iterative neural decoder updates and the adaptation of the tuning characteristics of individual neurons. Moreover, performance was found to continue to improve even after fixation of decoding parameters. The ability of subjects to learn fixed decoding parameters was further demonstrated by Ganguly and Carmena, where it was found that stable BMI control could be obtained over a 19-day period using a fixed decoder, with prominent learning effects observed during initial brain control sessions [Ganguly and Carmena, 2009]. This work has been extended to the realm of ECoG, where closed-loop BMI training has been found to allow non-human primates to de-correlate micro-ECoG signals from neighboring electrodes [Rouse and Moran, 2009, Rouse et al., 2013].

Taken together, these studies provide strong evidence that improvements in brain control performance are driven by both subject adaptation and by refinement of the neural decoder. However, little is known about the relationship of these two processes to improvements in control performance; if the decoder calibration process can be tailored to work with the learning process of the subject, it is possible that both the subject and the decoder can quickly converge upon a stable BMI control solution. Therefore, it is important to understand the contributions of each of these agents (decoder and user) on improvements in BMI control.

Unfortunately, these two processes are inextricably related, as neural decoder improvements are driven by the user's ability to consistently modulate their neural activity and user learning is driven by feedback provided by the neural decoder. To improve performance the user must interact with the decoder calibration process, stabilizing their neural responses during calibration to allow optimal decoding parameters to be estimated, then potentially adapting these parameters during closed-loop control. Recently, a number of groups have begun formalizing such "co-adaptive" BMI training algorithms, showing that iterative updates of decoding parameters based off of closed-loop control can be used to achieve successful cursor control by non-human primates [Gage et al., 2005, Mahmoudi and Sanchez, 2011, Gilja et al., 2012, Orsborn et al., 2012, Dangi et al., 2014].

Our work demonstrating closed-loop ECoG-based brain control by individuals with upper-limb paralysis presented in the preceding chapter utilized a two-agent co-adaptation training process which alternated between periods of decoder calibration and participant training. During decoder calibration, the human participants were instructed to adopt a consistent cursor control strategy while not correcting for errors in the resultant cursor movement. During training, the participant was then allowed to correct for errors in the resultant cursor movement to maximize performance. This approach aims to improve the BMI training process by attempting to ensure that only one agent (i.e. the neural decoder or participant) is adapting at a given time. However, in order for this to be effective, the human subject must be capable of modulating their neural activity in a consistent manner during decoder calibration. Furthermore, the participant must also be capable of adapting their neural activity during fixed decoder trials in order to maximize performance. To better understand the effects of decoder calibration and subject learning on brain-control performance, we analyzed data from subjects S1 and S2 during periods of BMI skill acquisition.

5.2 METHODS

We utilized two approaches to investigate the factors contributing to the acquisition of closed-loop brain control. First, we evaluated the contributions of fixed-decoder training and decoder calibration to improvements in performance by identifying changes in the angular error of BMI command signals. Secondly, we sought to uncover evidence of neural adaptation during brain control acquisition by tracking changes in the preferred directions of ECoG spectral features over time. The following section will describe these analyses in detail.

5.2.1 Angular error estimation during closed-loop cursor control

5.2.1.1 Datasets In order to assess the contribution of fixed-decoder training and decoder calibration on changes in control signal angular error, we analyzed data from subjects S1 and S2[1] during closed-loop brain control. Figure 5.1 depicts the data sets chosen for analysis for both subjects. For each of the analysis conditions (fixed-decoder and calibration), we selected pairs of 80-trial blocks which either surrounded a period of fixed decoder training or were immediately before and after decoder calibration. For the sake of simplicity, we refer to these as *pre-adaptation* and *post-adaptation* trial blocks for each condition, where the adapting agent is either the subject (fixed-decoder training) or the decoder (decoder calibration).

Trials in which computer assistance was provided to subjects were excluded from analysis in order to eliminate effects resulting from the masking of the true (unassisted) cursor movement. Furthermore, we restricted our analysis to data collected during two-dimensional cursor control, both due to the small amount of 3D control data available for Subject S1 as well as the difficulties in resolving bias for 3D control signals. In total, we analyzed 3 fixed-decoder and 4 calibration data sets for Subject S1, and 3 fixed-decoder and 2 calibration data sets for Subject S2. All data sets analyzed for Subject S1 utilized a consistent "hand/elbow" control strategy, but varied in the number of targets (either 4-target or 8-target). For Subject S2, data sets contained trials utilizing either "thumb/index" or "thumb/middle" control strategies during a 4-target task.

5.2.1.2 Analysis Control signal angular error was estimated by computing the angle between the instantaneous control signal vector and the idealized target direction for each trial. To do this, we first computed normalized target directions according to Equation 5.1:

$$\hat{\mathbf{p}}_t = \frac{\mathbf{p}_t - \mathbf{p}_c}{\|\mathbf{p}_t - \mathbf{p}_c\|}, \tag{5.1}$$

where $\hat{\mathbf{p}}_t$ is the normalized target direction and $\mathbf{p}_t$ and $\mathbf{p}_c$ are normalized target and center positions, respectively. Using these direction vectors, we calculated the target angle as:

$$\theta_t = \text{atan2}(p_{t,y}, p_{t,x}), \tag{5.2}$$

---

[1] See Chapter 4 for detailed subject descriptions.

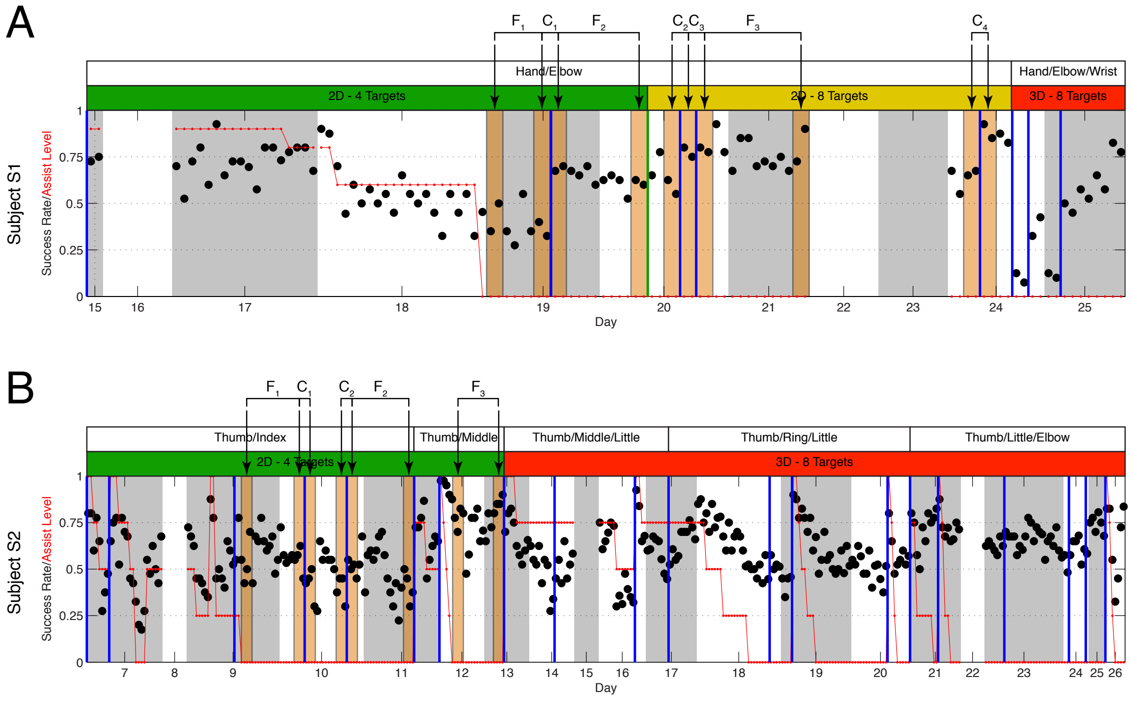

Figure 5.1: *Data sets used in angular error analysis.* Trial sets used in angular error analysis are shown for subjects S1 (*A*) and S2 (*B*). *Black* dots indicate the brain control success rate calculated over 40-trial blocks, while the shaded *orange* regions marked by black arrows indicate the trial blocks used in analysis. '*F*': Fixed-decoder training. '*C*': Decoder calibration. Subscripts indicate the set number for a given condition (fixed-decoder training or decoder calibration).

where $\operatorname{atan2}(y, x)$ is the quadrant-resolved arctangent of the vector defined by $[x, y]^\top$:

$$\operatorname{atan2}(y, x) = \begin{cases} \arctan \frac{y}{x} & x > 0 \\ \arctan \frac{y}{x} + \pi & y \geq 0,\ x < 0 \\ \arctan \frac{y}{x} - \pi & y < 0,\ x < 0 \\ +\frac{\pi}{2} & y > 0,\ x = 0 \\ -\frac{\pi}{2} & y < 0,\ x = 0 \end{cases} \tag{5.3}$$

A rotation matrix $\mathbf{R}$ for each target direction could then be calculated using Equation 5.4:

$$\mathbf{R}_t = \begin{pmatrix} \cos\theta_t & -\sin\theta_t \\ \sin\theta_t & \cos\theta_t \end{pmatrix}, \tag{5.4}$$

which was then used to rotate instantaneous control signal vectors $\mathbf{y}_i$ towards the origin:

$$\hat{\mathbf{y}}_i = \mathbf{R}\mathbf{y}_i. \tag{5.5}$$

Instantaneous angular error estimates, θ_e, were then calculated according to Equation 5.6:

$$\theta_e = \operatorname{atan2}(y_{i,y}, y_{i,x}), \tag{5.6}$$

where $y_{i,x}$ and $y_{i,x}$ are the x and y components of the instantaneous control signal vector. Calculating angular error in this manner preserved the sign of the angle between $\hat{\mathbf{p}}_t$ and $\mathbf{y}_i$, allowing biases in error to be resolved.

For each data set, instantaneous angular error was calculated over the $[0 - 500$ ms$]$ interval relative to the onset of cursor control and merged across trials to generate pre and post-adaptation angular error distributions. Pairs of distributions were generated for all trials, irrespective of target direction, as well as for individual targets. For purposes of comparison, we characterized angular error distributions by both their mean (μ_{pre}, μ_{post}) and standard deviation ($\sigma_{pre}, \sigma_{post}$), and used equations 5.7 and 5.8 to evaluate the change in the mean and standard deviation of angular error distributions:

$$\Delta\mu = |\mu_{post}| - |\mu_{pre}|, \tag{5.7}$$

$$\Delta\sigma = \sigma_{post} - \sigma_{pre}. \tag{5.8}$$

Two-tailed t-tests were used to determine the presence of bias for each distribution (i.e. $\mu_{pre} = 0$, $\mu_{post} = 0$), while two-sample t-tests assuming unequal variances were used to determine the significance of differences in the means between pre and post-adaptation distributions (i.e. $\mu_{pre} = \mu_{post}$). A two-sample F-test was used to determine the significance of differences in variance between the pre and post-adaptation distributions (i.e. $\sigma^2_{pre} = \sigma^2_{post}$). All statistical comparisons were considered significant at $p < 0.05$.

5.2.2 Estimation of preferred directions for ECoG spectral features

5.2.2.1 Dataset To investigate cortical adaptation during acquisition of closed-loop brain control, we examined data from Subject S1[2] during two-dimensional brain control. As shown by Figure 5.1A, 2D brain-control training for Subject S1 encompassed 7 testing days spanning 10 calendar days. All 2D brain control trials (excluding decoder calibration trials), $3,640$ in total, were used in our analysis.

We investigated two electrode subsets in an effort to better understand the potential influences of control strategy on cortical adaptation. As Subject S1 used a consistent “hand/elbow” somatotopic strategy during all 2D brain control trials, we defined two electrode subsets on the basis of their responses during attempted hand and elbow movements (see Figure 4.4). The “elbow” electrode subset included electrodes 7, 15, 23, and 31, while the “hand” subset included electrodes 4, 12, 20, and 28. The analyses described in the following section were preformed on these two subsets in addition to the full set of all electrodes.

5.2.2.2 Analysis Cortical adaptation during closed-loop brain control was evaluated by estimating the preferred directions (PDs) of ECoG spectral features during blocks of brain control trials. Normalized time-frequency data for each trial was averaged over the $[0-500 \text{ ms}]$ time window relative the the onset of cursor control. Data from ground and reference electrodes, as well as the activity of the $[0-10 \text{ Hz}]$ frequency band for each electrode, was excluded from analysis. Averaged time-frequency data for each of the remaining features were Z-scored and fit to a 2D cosine tuning model:

$$x_i = b_{i,0} + b_{i,x} d_x + b_{i,y} d_y + \epsilon, \tag{5.9}$$

where x_i is the normalized activity for the ith neural feature, d_x and d_y are the x and y components of the normalized target direction vector $\mathbf{d} = [d_x, d_y]^\top$, $b_{i,0}$ is a constant offset, and $\mathbf{b}_i = [b_{i,x}, b_{i,y}]^\top$ is the preferred direction of the ith feature. All neural features not found to be significantly tuned to target direction at $p < 0.05$ were excluded from further analysis. To compare PDs across all features, we converted preferred direction vectors to angles using Equation 5.10:

$$\theta_i = \text{atan2}(b_y, b_x). \tag{5.10}$$

Distributions of preferred directions were compared for selected 2D brain control trial sets to identify changes in the encoding of intended movement direction as a result of closed-loop training. On account of differences in the spatial specificity of the low ($f < 40$ Hz) and high ($f \geq 40$ Hz) frequency bands, we analyzed PD distributions for these bands independently. A two-sided rank sum test was used to determine the significance of differences in the medians of PD distributions for early and late brain control trials, while a two-sample F-test was used to compare the variance of these distributions.

To investigate changes in the encoding of intended movement direction in the neural population, we also computed PDs for non-overlapping 40-trial blocks for all 2D brain control trials. Preferred directions

---

[2] We did not analyze brain control data from Subject S2 due the relatively short duration of 2D training combined with frequent changes in the brain control strategy employed by the subject.

across trial blocks were compared by calculating the average PD vector for each electrode from the set of significantly-tuned features independently for both the low and high frequency bands. Averaged PD vectors for each electrode were then converted to angles using Equation 5.10 and re-referenced to the PD for the first 40-trial block:

$$\Delta\theta_i(k) = \theta_i(k) - \theta_i(1), \tag{5.11}$$

where $\theta_i(k)$ is the PD of the ith electrode for trial block k, $\theta_i(1)$ is the PD of the ith electrode for the first trial block, and $\Delta\theta_i(k)$ is the change in PD of the ith electrode between the first and kth blocks. Calculated $\Delta\theta_i(k)$ values were then averaged across all electrodes for each trial block and regressed against block number to determine if the PD distributions for the low and high frequency bands changed in a systematic manner over the course of brain control training. This was accomplished by fitting PD changes calculated using Equation 5.11 to the following model:

$$\Delta\theta(k) = mk + \Delta\theta_0, \tag{5.12}$$

where $\Delta\theta(k)$ is the predicted mean change in preferred direction for the kth trial block, m is the rate of change of the preferred direction, and $\Delta\theta_0$ is a PD offset term. We regressed the average number of tuned features for each trial block against the block number in order to determine if the number of significantly-tuned neural features changed as a result of training in a similar manner.

5.3 RESULTS

5.3.1 Changes in angular error during fixed-decoder training and decoder calibration

In order to identify the effects of fixed-decoder training and decoder calibration on brain-control performance, we examined both cursor command signals and distributions of angular error surrounding periods of subject adaptation (fixed-decoder training) and decoder adaptation (decoder calibration). Figures 5.2 and 5.3 depict control signal and angular error distributions for Subject S1 during fixed-decoder training and decoder calibration, respectively, while Figures 5.4 and 5.5 depict control signal and angular error distribution for Subject S2. A summary of all angular error distributions and p-values from associated statistical tests is provided by Table 5.1.

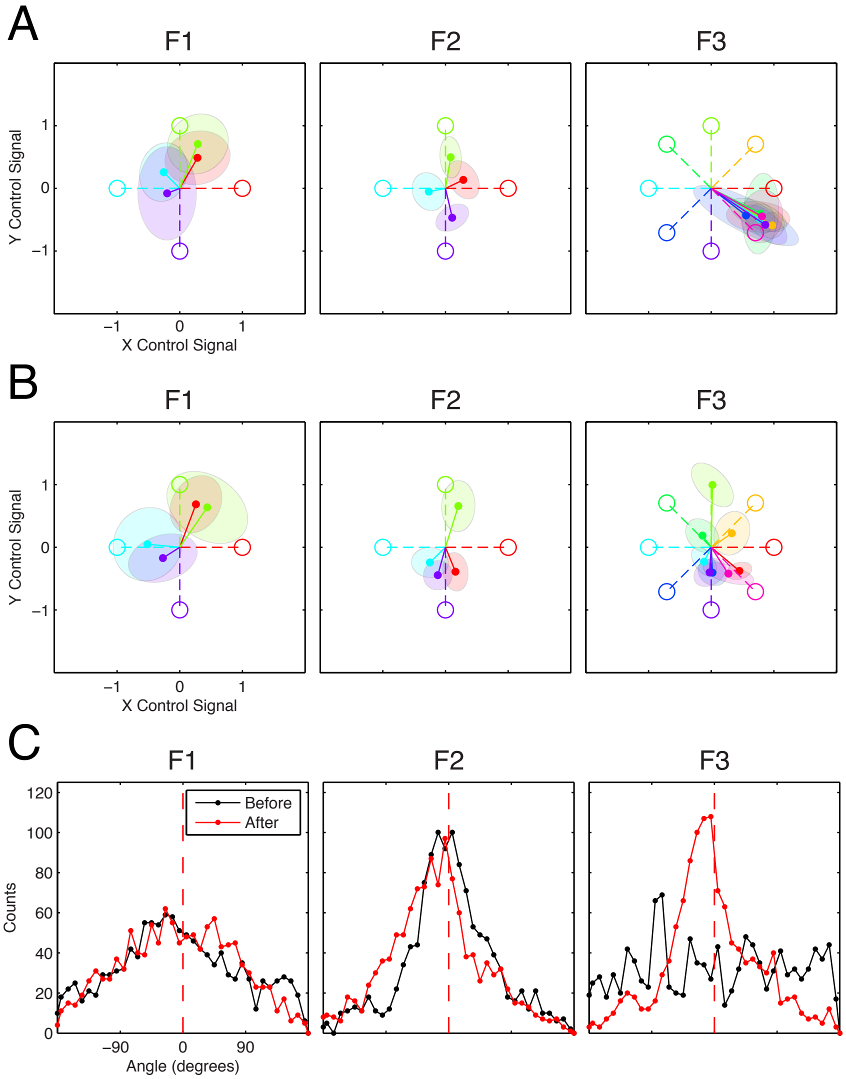

Figure 5.2: *Control signal error for Subject S1 before and after fixed-decoder training.* (*A*) Mean control signal vectors (solid lines) for each target before fixed-decoder training. Dashed lines indicate ideal (straight-line) control signal vectors towards each target. Shaded regions around mean control signal vector endpoints represent one standard deviation covariance ellipses. (*B*) Mean control signal vectors after fixed-decoder training. (*C*) Histograms of control signal angular error before (*black*) and after (*red*) fixed-decoder training. Results are provided for fixed-decoder data sets *F1*, *F2*, and *F3* (see Figure 5.1).

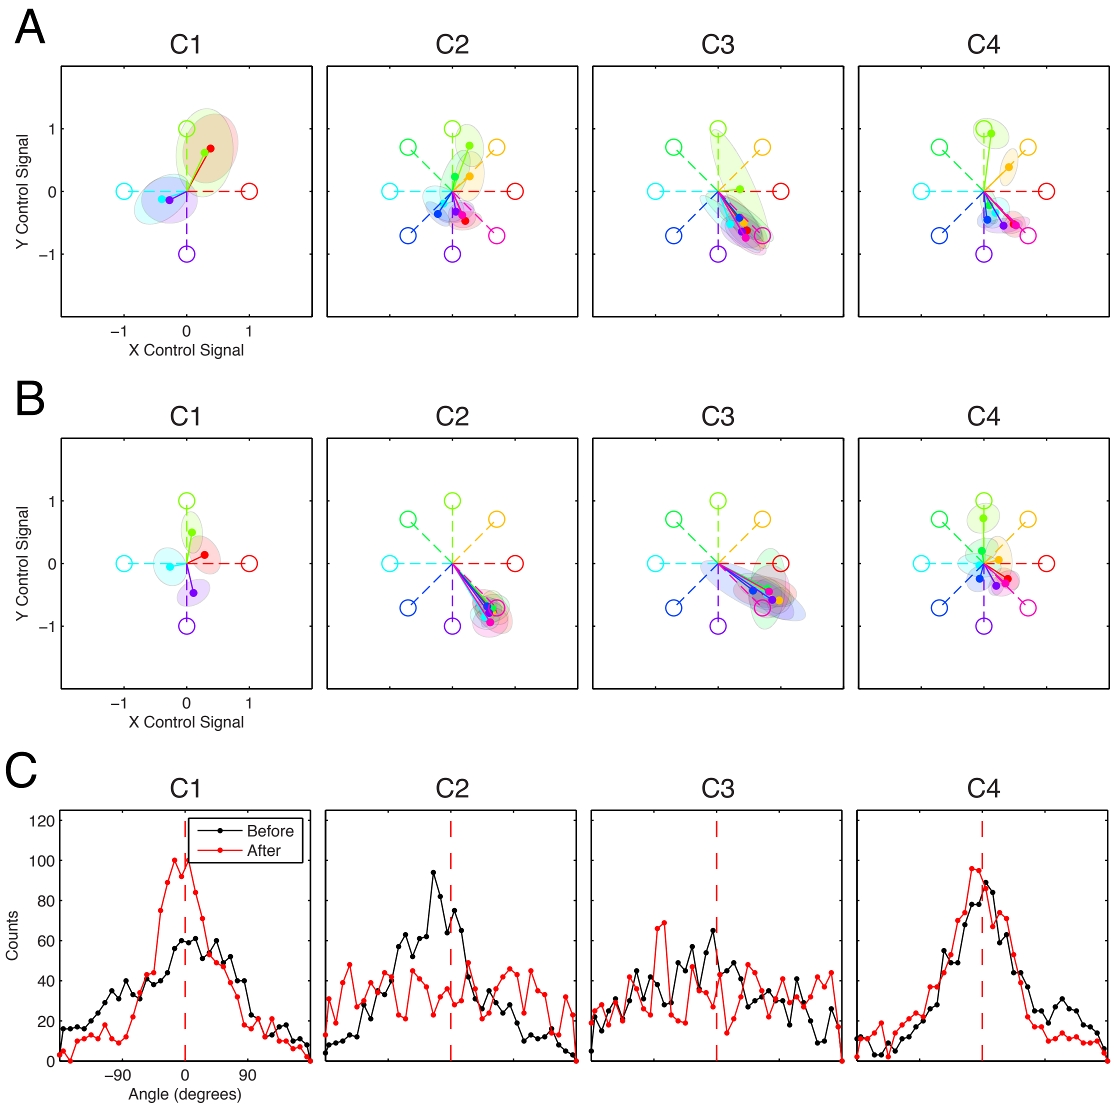

Figure 5.3: *Control signal error for Subject S1 before and after decoder calibration.* (*A*) Mean control signal vectors for each target before decoder calibration. (*B*) Mean control signal vectors after decoder calibration. (*C*) Histograms of control signal angular error before and after decoder calibration. Results are provided for calibration data sets *C1*, *C2*, *C3*, and *C4* (see Figure 5.1).

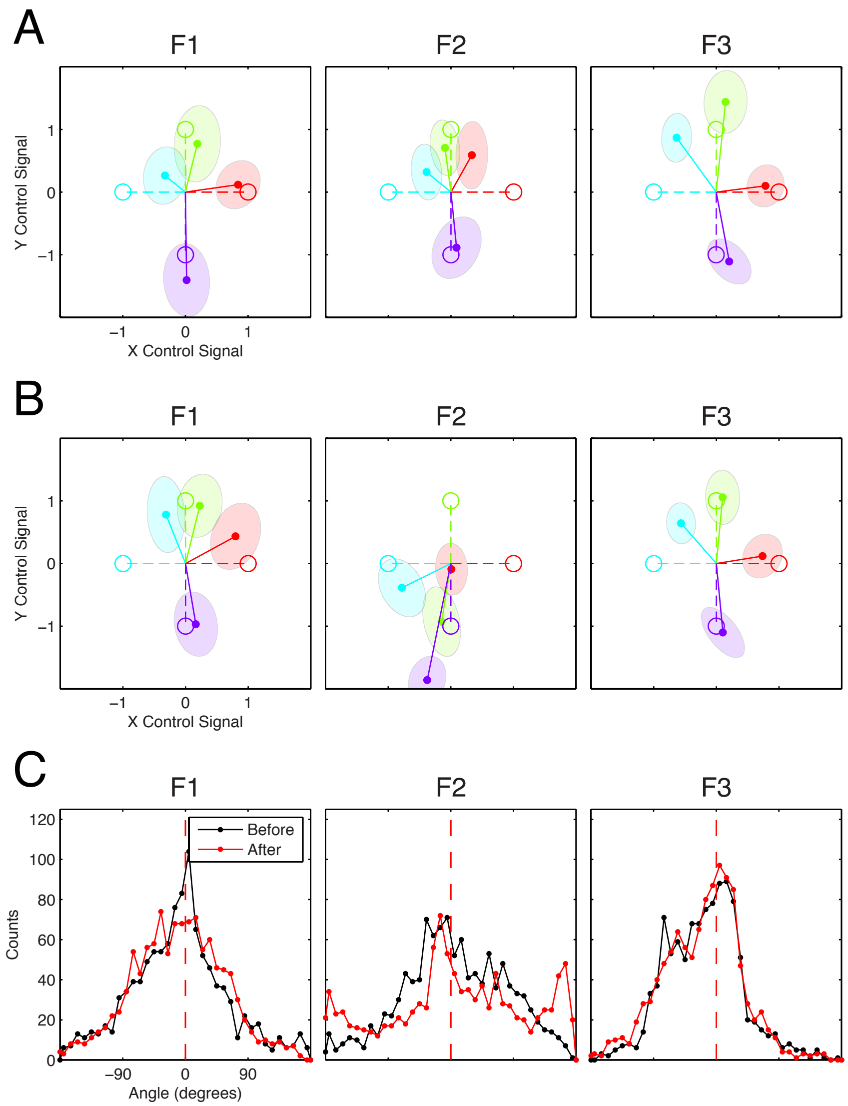

Figure 5.4: *Control signal error for Subject S2 before and after fixed-decoder training.* (*A*) Mean control signal vectors for each target before fixed-decoder training. (*B*) Mean control signal vectors after fixed-decoder training. (*C*) Histograms of control signal angular error before and after fixed-decoder training. Results are provided for calibration data sets *F1*, *F2*, and *F3* (see Figure 5.1).

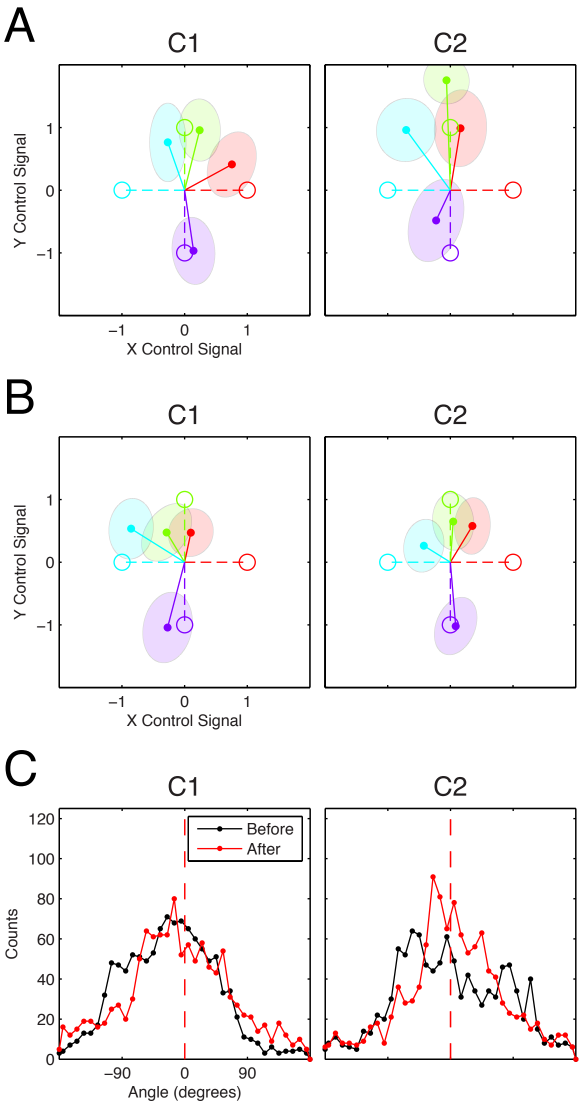

Figure 5.5: *Control signal error for Subject S2 before and after decoder calibration.* (*A*) Mean control signal vectors for each target before decoder calibration. (*B*) Mean control signal vectors after decoder calibration. (*C*) Histograms of control signal angular error before and after decoder calibration. Results are provided for calibration data sets *C1* and *C2* (see Figure 5.1).

Table 5.1: *Summary of angular error distributions.* Calculated mean and standard deviation values are provided for all angular error distributions. In addition, p-values for all statistical tests comparing pre and post-adaptation distributions are shown. Values listed in bold indicate those statistical comparisons considered significant at $p < 0.05$.

Subject	**Dataset**	μ_{pre}	μ_{post}	σ_{pre}	σ_{post}	$p(\mu_{pre} = 0)$	$p(\mu_{post} = 0)$	$p(\mu_{pre} = \mu_{post})$	$p(\sigma^2_{pre} < \sigma^2_{post})$	$p(\sigma^2_{pre} > \sigma^2_{post})$
S1	F1	−1.92	−2.30	88.33	80.84	0.453	0.326	0.912	0.999	**0.001**
	F2	6.47	−13.03	63.36	67.35	**4.22e−4**	**3.16e−11**	**3.79e−11**	**0.0174**	0.983
	F3	4.72	0.17	103.20	65.37	0.114	0.929	0.197	1.000	**1.14e−54**
S1	C1	0.34	6.47	81.15	63.36	0.887	**4.22e−4**	**0.0398**	1.000	**9.09e−18**
	C2	−11.46	−2.13	70.87	101.61	**2.67e−8**	0.469	**0.0092**	**2.77e−35**	1.000
	C3	−4.18	4.72	90.28	103.20	0.109	0.114	**0.0246**	**1.91e−6**	1.000
	C4	14.48	−2.94	73.47	67.71	**1.42e−11**	0.133	**1.82e−9**	0.998	**0.0024**
S2	F1	−3.96	−5.02	68.19	65.23	0.0539	0.0098	0.706	0.931	0.0686
	F2	10.22	12.15	73.90	99.40	**4.23e−6**	**8.12e−5**	0.610	**1.49e−22**	1.000
	F3	−11.53	−14.80	52.97	55.45	**2.83e−12**	**2.94e−18**	0.161	0.0674	0.932
S2	C1	−17.01	−5.63	65.43	75.92	**5.98e−18**	**0.0120**	**1.23e−4**	**2.55e−7**	1.000
	C2	4.01	7.92	76.64	70.54	0.0882	**1.88e−4**	0.216	0.997	**0.0031**

Upon examination of these data, several characteristics become apparent. First, we observe instances where both fixed-decoder training and decoder calibration result in decreases in the variability and overall error of the angular error distributions. For example, average control signal vectors for dataset *F3* from Subject S1 exhibit a strong bias towards the lower-right quadrant of the workspace before decoder training (Figure 5.2A); after training these vectors are much more closely aligned with their respective targets (Figure 5.2B). This is confirmed by an examination of the angular error distributions before and after fixed-decoder training for this dataset (Figure 5.2C). Here, we find that the distribution of angular error after fixed-decoder training exhibits a prominent peak near $\theta = 0$ which was not present in the pre-training distribution, indicative of a significant decrease in the variance of the angular error distribution as a result of fixed-decoder training (see Table 5.1). Similar changes in error distributions were observed for dataset *C1* from Subject S1 (Figure 5.3) and dataset *C2* from Subject S2 (Figure 5.5).

Surprisingly, in several instances decoder calibration resulted in an increase in overall angular error for Subject S1. For datasets *C2* and *C3* for Subject S1, control signals after calibration exhibited a strong bias towards the lower-right quadrant of the workspace (Figure 5.3B). This is also evident in the pre and post-calibration angular error distributions for these datasets (Figure 5.3C), with the post-calibration error distributions found to be more variable than the pre-calibration distributions. This is in contrast to all other calibration datasets, where angular error was either decreased or unaffected following decoder calibration. While the cause for the increase in angular error post-calibration for these datasets is unclear, it is important to note that these calibration sessions were performed shortly following a change from a 4-target to an 8-target cursor control task. It may be the case that this change in task difficulty adversely affected the ability of Subject S1 to control the computer cursor, resulting in poorly-estimated decoding parameters.

In order to identify if consistent changes in angular error were evident across datasets, we examined the change in the mean and standard deviation of angular error distributions for individual target directions before and after decoder calibration and fixed-decoder training for each subject; these results are shown by Figure 5.6. By plotting the mean and standard deviation of the pre and post-adaptation distributions as a two-dimensional vector, we can compare the observed angular error changes to those expected by pure bias correction (Figure 5.6A) or variance reduction (Figure 5.6B). In general, we find that angular error exhibited greater changes from the pre to post-adaptation distributions for Subject S1, as indicated by the larger magnitude of angular error change vectors for Subject S1 (Figure 5.6C,D) than Subject S2 (Figure 5.6E,F). However, we did not observe consistent changes in the angular error distributions for either subject which were consistent with that expected from bias correction or variance reduction.

To better investigate changes in angular error during adaptation periods, we plotted the change in the standard deviation between pre and post-adaptation distributions ($\Delta\sigma$, Equation 5.8) against that of the mean of the same distributions ($\Delta\mu$, Equation 5.7) for individual targets; these results are shown by Figure 5.7. By charactering angular error distributions in this manner, systematic reductions in bias and/or variance could be identified by the specific quadrant of the $\Delta\mu - \Delta\sigma$ space in which data points for specific data sets

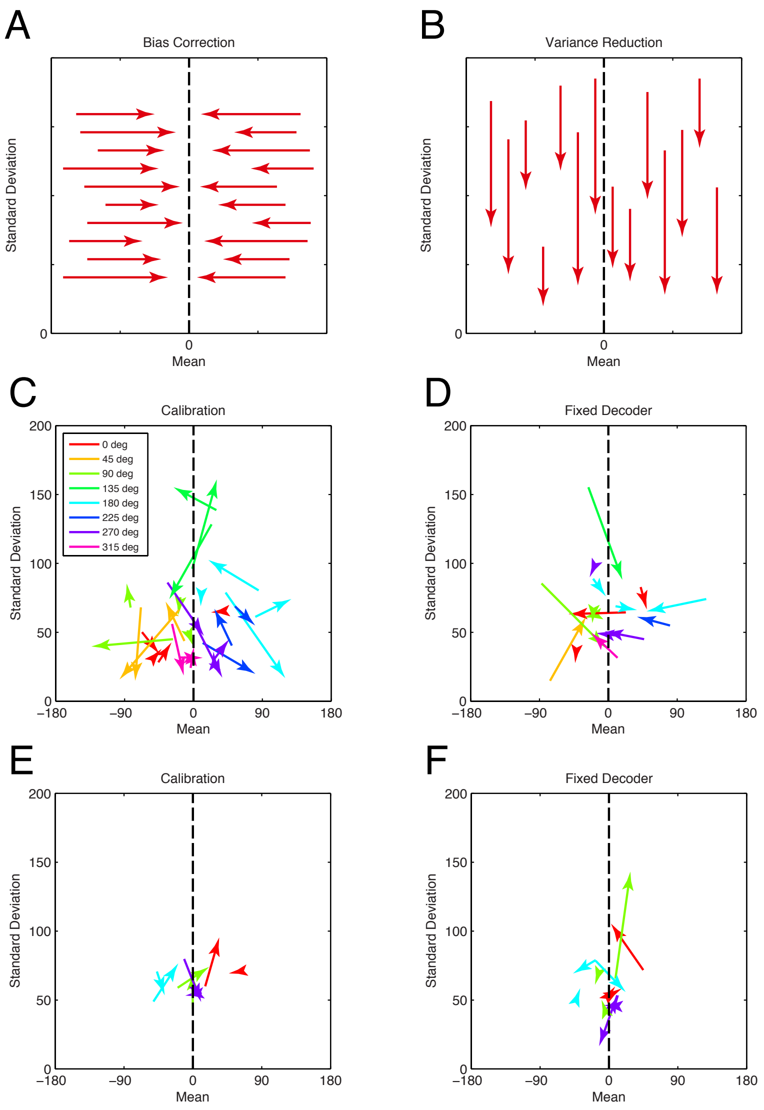

Figure 5.6: *Change in angular error distributions during adaptation.* Pre and post-adaptation angular error distributions are represented as a single vector for each target direction, with the tail of each vector representing the mean and standard deviation of the pre-adaptation distribution and the head representing mean and standard deviation for the post-adaptation distribution. Colors represent angular error change vectors for individual target directions. (*A*) Conceptual example of angular error change vectors expected as a result of a reduction in control signal bias. (*B*) Conceptual example of angular error change vectors expected as a result of a reduction in control variance. (*C*) Angular error change vectors resulting from decoder calibration for Subject S1. (*D*) Angular error change vectors resulting from fixed-decoder training for Subject S2. (*E*) Angular error change vectors resulting from decoder calibration for Subject S1. (*F*) Angular error change vectors resulting from fixed-decoder training for Subject S2.

were found to reside (Figure 5.7A). While we find that this analysis confirms our observations for specific data sets, such as the reduction in variance for dataset *C1* for Subject S1 (Figure 5.7B) and the bias correction observed during adaptation period *F3* for Subject S1 (Figure 5.7C), neither decoder calibration or fixed decoder training resulted in systematic changes in either the mean or variability of control signal angular error across all datasets.

5.3.2 Cortical adaptation during closed-loop training

Analysis of changes in angular error during closed-loop brain control provides a means by which control performance improvements can be investigated in more detail. However, as demonstrated by the findings in Section 5.3, investigating changes in BMI command signals alone makes dissociating the effects of decoder calibration and subject training on performance changes difficult. In an effort to overcome this limitations, we investigated the neural encoding of intended movement direction during closed-loop brain control by tracking changes in the preferred directions of ECoG spectral features over the course of 2D brain control training.

To accomplish this, changes in preferred direction were calculated for non-overlapping 40-trial blocks for all 2D brain control trials. Figure 5.8A shows the trend in average change in preferred direction for both the low and high-frequency bands as a function of trial block. A significant linear trend for both the low and high-frequency bands was found ($m_L = 0.539°/\text{block}$, $p_L = 2.76e{-}7$; $m_H = 0.539°/\text{block}$, $p_H = 3.00e{-}6$), corresponding to predicted counter-clockwise rotation of 49° for both the low and high-frequency bands. We also examined whether the number of significantly-tuned features significantly changed as a function of trial block (Figure 5.8B). The number of significantly-tuned low-frequency spectral features was not found to vary significantly as a function of trial block ($m_L = 0.062$ features/block, $p_L = 0.327$), while the number of high-frequency spectral features was found to vary significantly as a function of trial block ($m_H = 0.318$ features/block, $p_H = 0.0212$).

Trends in preferred direction change and the number of significantly-tuned neural features were also examined for "elbow" and "hand" electrode subsets. As shown by Figure 5.8C, significant linear relationships between preferred direction change and trial block were found for both the low and high-frequency bands ($m_L = 0.747°/\text{block}$, $p_L = 9.16e{-}6$; $m_H = 0.656°/\text{block}$, $p_H = 8.96e{-}22$) for the subset of electrodes preferentially responsive to attempted elbow movement, though a significant relationship between the number of significantly-tuned features and trial block was not found for either frequency band (Figure 5.8D: $m_L = 0.000°/\text{block}$, $p_L = 0.986$; $m_H = 0.089°/\text{block}$, $p_H = 0.054$). Interestingly, when the preferred direction analysis was restricted to this electrode subset, the seemingly gradual change in preferred direction shown by Figure 5.8A appeared to be a sudden shift occurring at trial block 56, coinciding with the decoder calibration performed on testing day 19 (see Figure 5.1A). When the change in preferred direction was examined for the subset of electrodes preferentially responsive to attempted hand movement, we observed that while a significant linear trend was found to exist for the low-frequency band, the high-frequency band

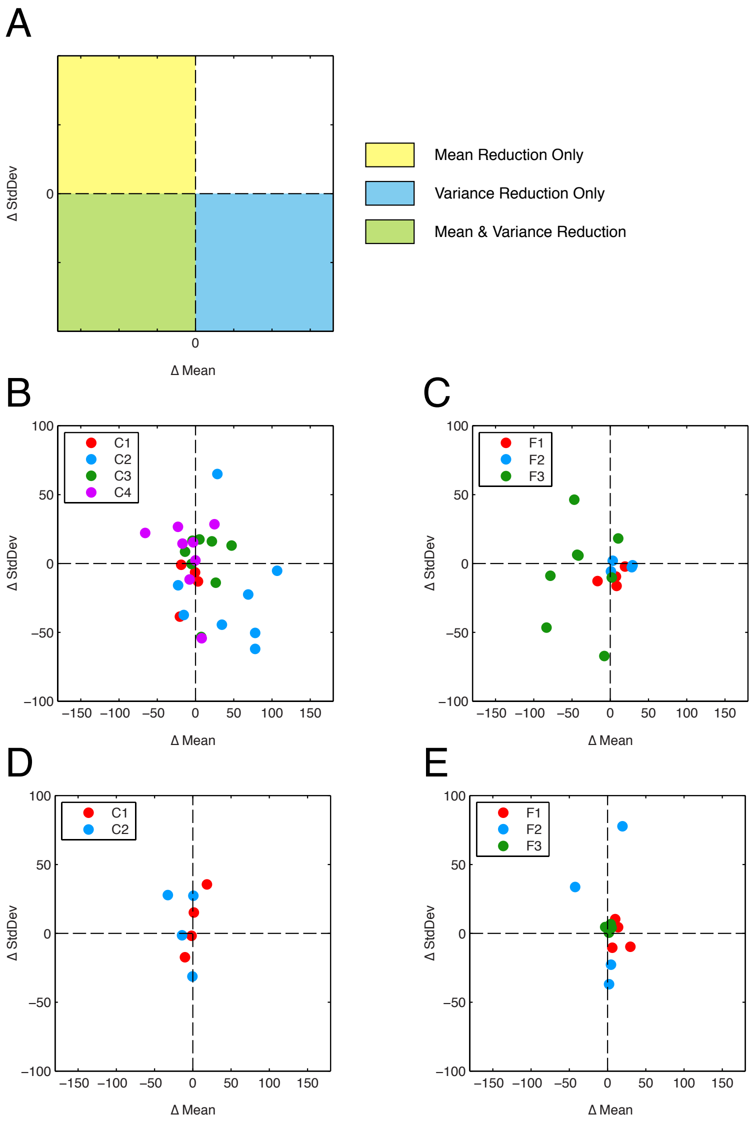

Figure 5.7: *Effect of decoder and subject adaptation on angular error distributions.* Shown are changes in the mean and standard deviation of angular error distributions as a result of adaptation. Individual points represent the change in the standard deviation ($\Delta\sigma$) plotted agains the change in absolute mean error ($\Delta\mu$) for a single target direction and dataset. Colors represent changes in angular error distributions for the indicated dataset. (*A*) Conceptual illustration indicating the regions corresponding to reductions in the mean (*yellow*), variance (*blue*), and mean and variance (*green*) of the angular error distributions. (*B*) Change in angular error distributions as a result of decoder calibration for Subject S1. (*C*) Change in angular error distributions as a result of fixed-decoder training for Subject S1. (*D*) Change in angular error distributions as a result of decoder calibration for Subject S2. (*E*) Change in angular error distributions as a result of fixed-decoder training for Subject S2.

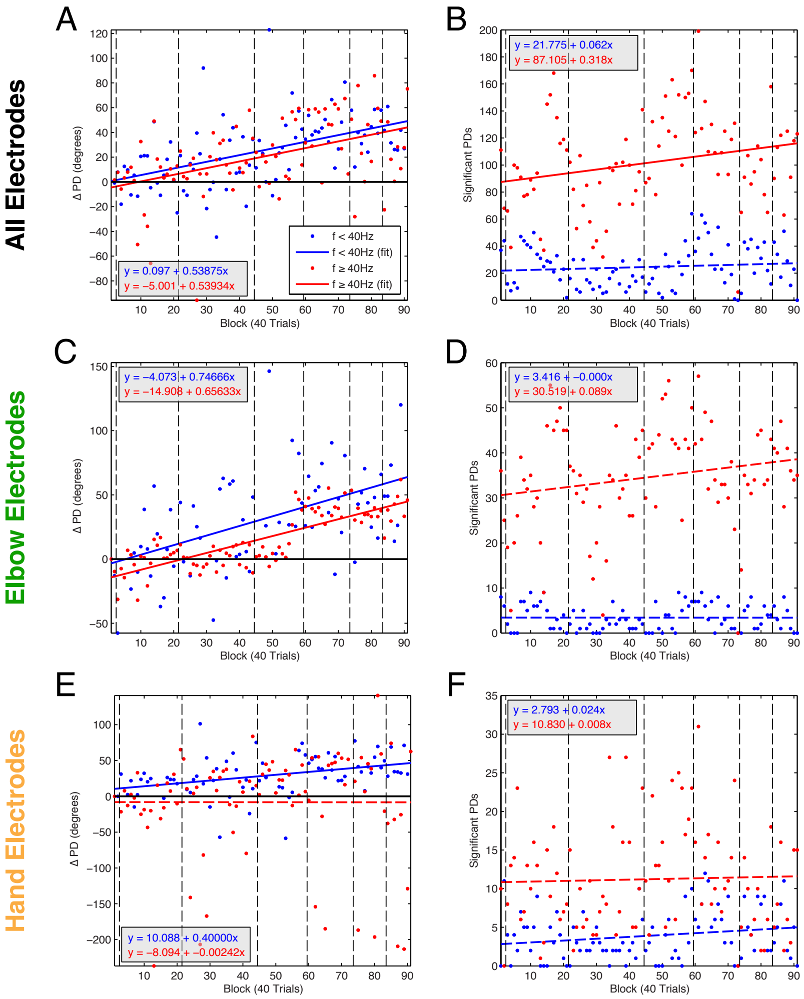

Figure 5.8: *Trends in preferred direction changes during closed-loop training.* Shown are trends in the change in preferred direction (*left* column) and the number of significantly-tuned features (*right* column) over the course of 2D brain-control training for Subject S1. Data are shown for all electrodes (*top* row), electrodes responsive to attempted elbow movement (*middle* row), and electrodes responsive to attempted hand movement (*middle* row). Dots indicate values calculated over 40-trial blocks, with lines indicating the results of a linear regression fit. Solid lines indicate those fits with a slope significantly different than zero at $p < 0.05$; dashed lines indicate non-significant fits. *Blue* and *red* colors indicate results for low-frequency ($f < 40$ Hz) and high-frequency ($f \geq 40$ Hz) features, respectively. (*A*) Change in preferred direction for all electrodes. (*B*) Change in number of tuned features for all electrodes. (*C*) Change in preferred direction for the elbow electrode subset. (*D*) Change in number of tuned features for the elbow electrode subset. (*E*) Change in preferred direction for the hand electrode subset. (*F*) Change in number of tuned features for the hand electrode subset.

no longer exhibited a significant relationship between preferred direction change and trial block (Figure 5.8E: $m_L = 0.400°$/block, $p_L = 8.57e{-}4$; $m_H = -0.002°$/block, $p_H = 0.994$). As with the subset of elbow electrodes, a significant change in the number of tuned features as a function of trial block was not found for the hand electrode subset (Figure 5.8D: $m_L = 0.024°$/block, $p_L = 0.062$; $m_H = 0.008°$/block, $p_H = 0.762$).

Noting the apparent shift in preferred direction occurring near decoder calibration on testing day 19, we examined the distribution of preferred directions of all significantly-tuned features for two specific sets of trials: (1) the initial and final 80 trials of brain control, and (2) the 480 trials surrounding (240 before/240 after) decoder calibration on day 19. Figure 5.9A shows the distribution of preferred directions of the significantly-tuned features for the low frequency ($f < 40$ Hz) and high frequency ($f \geq 40$ Hz) bands for the first 80 trials of 2D brain control, while Figure 5.9B shows the PD distributions for the final 80 trial of brain control; Figure 5.9C provides a histogram of preferred direction angles for the same data. For initial brain control trials, the distribution of the preferred direction for the high-frequency features was found to be centered in the upper-right quadrant of the workspace ($\bar{\theta}_{H,pre} = 62.8°$), while that of the low-frequency features was found to be approximately centered in the negative-Y direction ($\bar{\theta}_{L,pre} = 254.7°$). However, after 7 days of 2D cursor control, the preferred directions for all features was found to rotate counter-clockwise, with the distributions of the high and low-frequency features now centered in the upper-left ($\bar{\theta}_{H,post} = 124.9°$) and lower-right ($\bar{\theta}_{L,post} = 315.4°$) quadrants of the workspace, respectively. This rotation corresponded to changes in the means of the preferred direction distributions of 60.7° and 62.1° for the low and high-frequency bands. Additionally, we observed a significant increase in the variance of the preferred direction distribution for the high-frequency band after 2D brain control ($p = 1.48e{-}5$); a similar effect was not found for the low-frequency band ($p = 0.51$).

When preferred directions were examined surrounding decoder calibration on testing day 19, we observe a similar shift between the pre-calibration (Figure 5.9D) and post-calibration (Figure 5.9E) distributions. As seen in Figure 5.9F, the shift in the preferred direction distribution for the low-frequency band ($\bar{\theta}_{L,pre} = 244.4°$, $\bar{\theta}_{L,post} = 297.4°$, $\Delta\bar{\theta}_L = 53.1°$) was similar to that occurring between the initial and final brain control trials, while that of the high-frequency band ($\bar{\theta}_{H,pre} = 88.5°$, $\bar{\theta}_{H,post} = 115.1°$, $\Delta\bar{\theta}_H = 32.1°$) was found to be less than that occurring throughout the course of 2D brain control training. However, preferred direction distributions for the post-calibration (Figure 5.9E) and final brain control (Figure 5.9B) trials were found to be similar to one another, suggesting little cortical adaptation occurred between decoder calibration on day 19 and the final 2D brain control trials on day 24.

We next investigated the spatial distribution of preferred directions for both the initial/final and pre/post-calibration sets of 2D brain control trials. As shown by Figure 5.10B, electrodes in the medial half of the electrode grid exhibited similar preferred directions during early brain control trials, while few electrodes in the lateral half of the grid were found to be significantly-tuned to target direction. However, by the end of 2D brain control training, preferred directions for medial electrodes were found to have rotated counter-clockwise, with electrodes in the lateral aspect of the electrode grid now exhibiting a substantial number

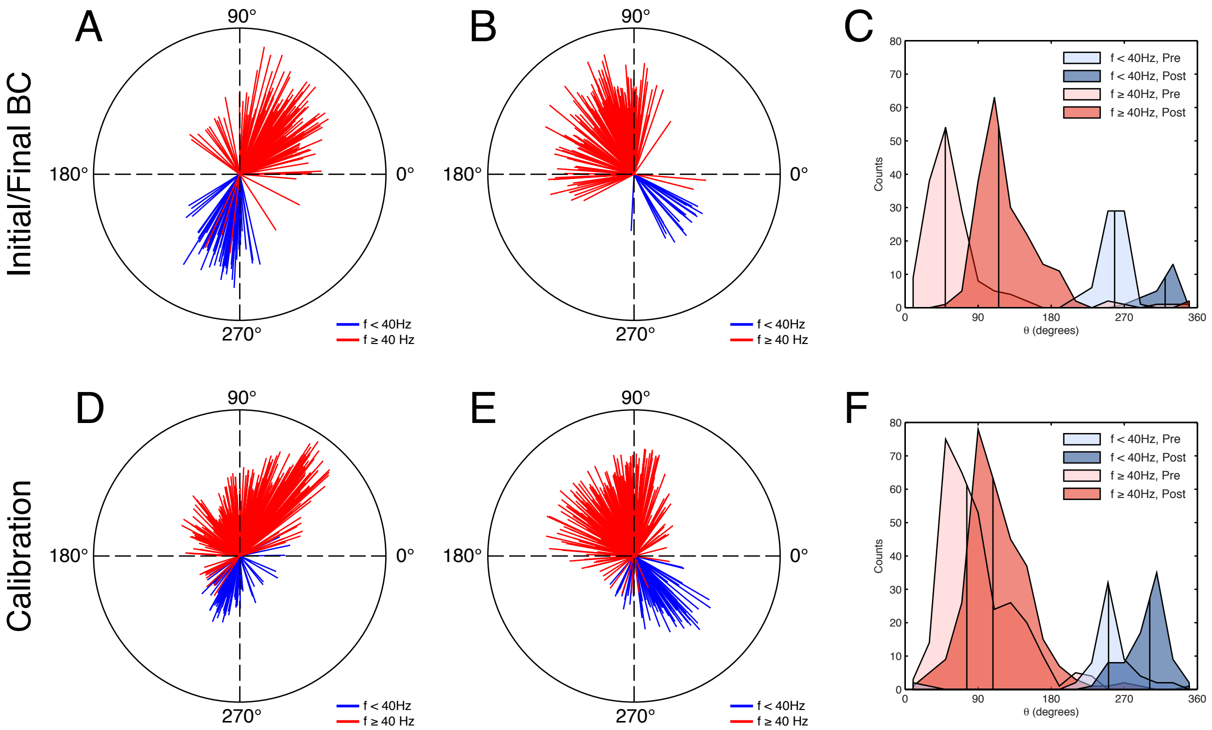

Figure 5.9: *Change in distributions of preferred directions during 2D computer cursor control.* (*A*) Preferred directions for significantly-tuned spectral features for the first 80 trials of 2D brain control. *Blue* vectors indicate preferred directions for low-frequency ($f < 40$ Hz) features, while *red* vectors indicate preferred directions for high-frequency ($f \geq 40$ Hz) features. (*B*) Preferred directions for the last 80 trials of 2D brain control. (*C*) Histograms of preferred directions for early and late brain control trials. *Blue*-shaded regions indicate preferred direction distributions for low-frequency ($f < 40$ Hz) features, while *red*-shaded regions vectors indicate preferred direction distributions for high-frequency ($f \geq 40$ Hz) features. For each frequency range, light-shaded regions indicate early brain control trials, while dark-shaded regions indicate late brain control trials. Vertical black lines indicate the mean of each preferred direction distribution. (*D*) Preferred directions for significantly-tuned spectral features for the 240 trials preceding decoder calibration on testing day 19. (*E*) Preferred directions for the 240 trials following decoder calibration on testing day 19. (*F*) Histograms of preferred directions for pre and post-calibration trials.

of significantly-tuned spectral features. When the spatial distribution of preferred directions was examined surrounding decoder calibration (Figure 5.10C), we observe that while the pre-calibration distributions for the medial electrodes are similar to those of initial brain control trials, there was less consistency between the PD distributions of the initial brain control and pre-calibration trials for the medial half of electrodes. Following decoder calibration, however, the spatial distribution of preferred directions was found to closely match that of the final brain control trials. Furthermore, for the post-calibration and final brain control trial sets we observe that high-frequency features exhibit a preferred direction gradient, with medial electrodes having preferred directions near $\theta = 90°$, transitioning to $\theta = 180°$ for lateral electrodes.

Finally, we compared the preferred direction vectors of the elbow and hand electrode subsets to those expected from the instructed control strategy. Based on this strategy (Figure 5.11A), elbow and hand electrodes would exhibit preferred directions vectors towards the upper-right and upper-left quadrants of the workspace, respectively (Figure 5.11B). During initial brain control trials, the preferred direction vectors for elbow electrodes were found to align with that expected from the instructed strategy, while the subset of hand electrodes exhibited similar tuning to that of the set of elbow electrodes (Figure 5.11C). However, by day 19, prior to decoder calibration, we find that the distribution of PDs for the hand subset had rotated nearly 90° (Figure 5.11D). Following decoder calibration, the preferred directions of the hand electrode subset were now found to align closely with that expected from the instructed control strategy, while the PDs of the elbow subset had rotated to approximately $\theta = 90°$ (Figure 5.11E). A slight counter-clockwise rotation of the PDs for both electrode subsets was observed from the post-calibration trials on day 19 to the final 2D brain control trials (Figure 5.11F). The preferred direction vectors for the elbow and hand electrode subsets on day 24, rotated approximately 45° from that expected from the instructed control strategy, was consistent with a rotated control strategy in which attempted elbow movement was used to move the cursor to the upper-right, attempted hand movement was used to move the cursor the the lower-left, attempted hand and elbow movement was used to move the cursor to the upper-left, and relaxation was used to move the cursor to the lower-right (Figure 5.11H).

5.4 DISCUSSION

We have examined the influences of adaptation, both on the part of the decoder and of the neural population, to improvements in ECoG BMI performance. To accomplish this, we compared angular error before and after decoder calibration and fixed-decoder training. In addition, we tracked changes in the neural encoding of attempted movement directions during the course of 2D brain control training. The following sections will discuss our findings in additional detail, highlight the limitations of the presented work, and propose additional experiments which could be performed to address these shortcomings.

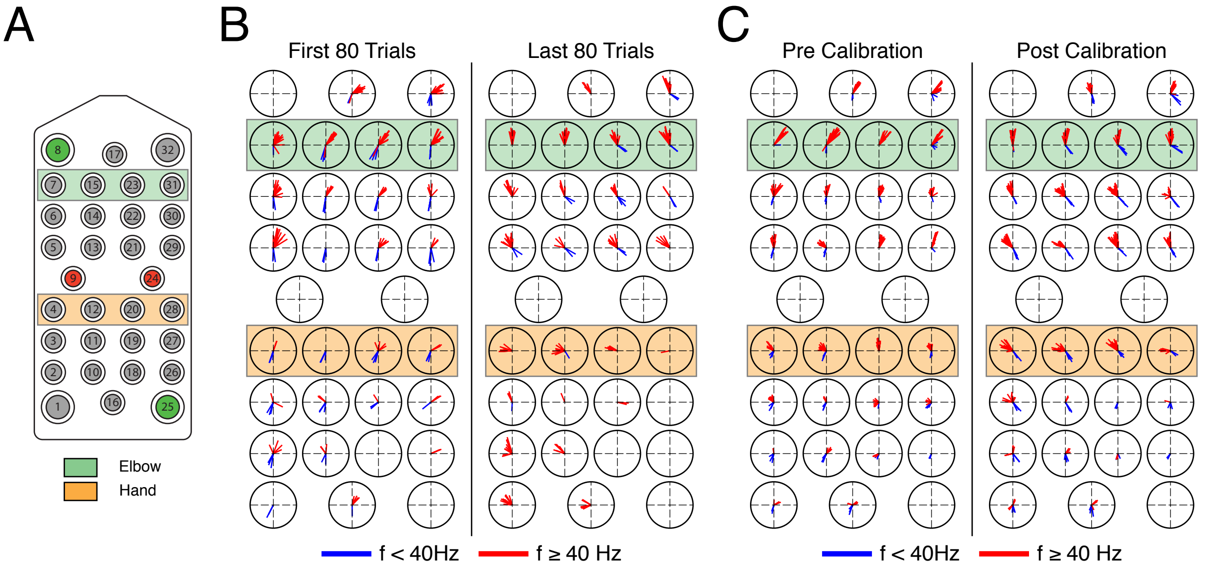

Figure 5.10: *Spatial distribution of preferred directions during closed-loop cursor control for Subject S1.* Shown are preferred direction distributions for all electrodes for the initial and final brain control trials, as well as those trials surrounding decoder calibration on testing day 19. (*A*) Electrode schematic for Subject S1. (*B*) Spatial distribution of preferred directions for the first and last 80 trials of two-dimensional brain control training. Each circular plot shows the preferred directions of significantly-tuned spectral features for an individual electrode. Plots for individual electrodes are arranged as depicted by the electrode schematic. *Blue* vectors indicate preferred directions for low-frequency ($f < 40$ Hz) features, while *red* vectors indicate preferred directions for high-frequency ($f \geq 40$ Hz) features. (*C*) Spatial distribution of preferred directions before and after decoder calibration on testing day 19. *Green* and *orange* regions in all panels indicate those electrodes deemed to be responsive to attempted elbow and hand movements, respectively.

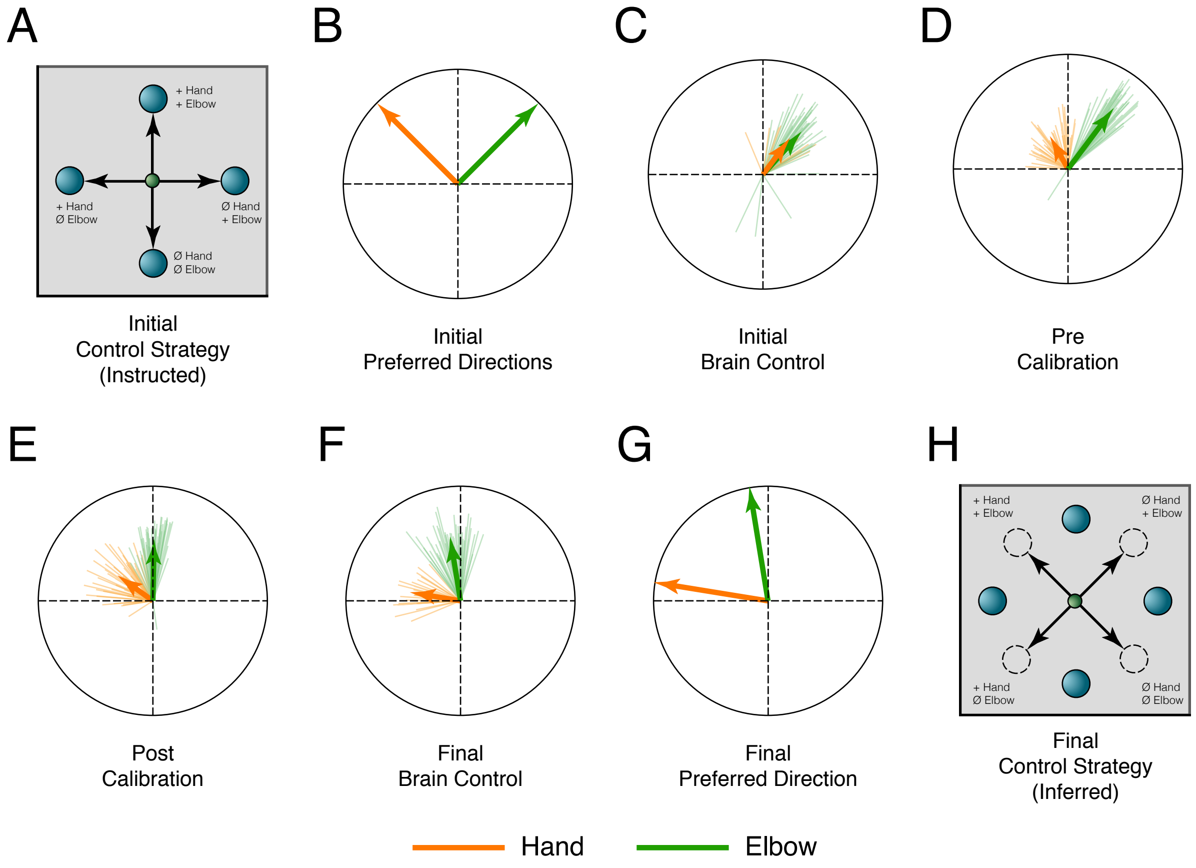

Figure 5.11: *Control strategy adaptation during acquisition of closed-loop brain control for Subject S1.* Shown is a comparison of the expected and actual preferred direction vectors for subsets of elbow and hand electrodes over the course of 2D cursor control. (*A*) Instructed 2D control strategy for Subject S1. (*B*) Expected preferred directions of elbow (*green*) and hand (*orange*) electrode subsets. (*C*) Preferred directions calculated during the first 80 trials of closed-loop brain control. Thin lines indicate the preferred directions of individual features, while thick lines indicate the mean preferred direction vector for each electrode subset. (*D*) Preferred direction vectors for the 240 trials preceding decoder calibration on testing day 19. (*E*) Preferred direction vectors for the 240 trials following decoder calibration on testing day 19. (*F*) Preferred direction vectors for the final 80 trials of closed-loop cursor control on day 24. (*G*) Final normalize preferred direction vectors for elbow and hand electrode subsets. (*H*) Final control strategy inferred from final preferred direction vectors.

5.4.1 Changes in angular error during decoder calibration and fixed-decoder training

By comparing the angular error between ideal and actual BMI control signals surrounding decoder calibration events and periods of fixed-decoder training, we are able to characterize changes in brain control performance beyond that which is possible by simply examining target acquisition success rate. For example, as shown by Figure 5.1A, performance, as judged by overall cursor control success rate, during fixed-decoder dataset *F2* for Subject S1 is relatively consistent (success rate: ∼75%). However, as shown by Figure 5.2, angular error distributions for early trials for this dataset are characterized by a strong control signal bias, which is reduced through closed-loop training. The fact that this reduction in control signal angular error is not evident in the overall success rate highlights the sensitivity of our approach for identifying changes in cursor control performance.

Not surprisingly, we find cases where decoder calibration significantly improves brain control performance (Subject S1, datasets *C1* and *C4*; Subject S2, dataset *C2*). Over time, changes in neural activity, either resulting from non-stationarities in the neural population [Perge et al., 2013] or learning on the part of the subject [Taylor et al., 2002, Ganguly and Carmena, 2009], may change the patterns of neural activity expressed during closed-loop control. By re-estimating decoder parameters using newly-acquired neural data, these changes can be accounted for and control performance can be improved. Unfortunately, as demonstrated by calibration datasets *C2* and *C3* for Subject S1, we find that decoder calibration can also lead to dramatic increases in angular error, manifested as strong biases in BMI control signals. Interestingly, the success rate following these calibration events does not reflect the high degree of angular error present in the instantaneous control signal; in fact, the success rate following calibration for dataset *C2* shows a substantial increase in performance. As control signal angular error was calculated over the first 500 ms of each trial, this indicates that Subject S1 was able to rapidly correct for biases in the instantaneous control signal to successfully complete the brain control task.

It is tempting to attribute control signal biases during the first 500 ms of cursor control[3] for each trial to slow reaction times on the part of the subject. With the somatotopic control strategy used for both subjects, cortical activity during rest is used to drive cursor movement in the downward direction. This results in a confound between resting-state activity during the inter-trial interval of the task and cortical activity during intended downwards movement of the cursor. As a result, slow reaction times would be manifested as downward biases in the control signal during the early period of each trial. However, as evidenced by the low degree of control signal angular error for other datasets for this subject (e.g., dataset *C1*, post-calibration; dataset *C2*, pre-calibration), Subject S1 was capable of generating target-aligned control signal vectors during this cursor control epoch. Thus, it is unlikely that the control signal biases observed post-calibration for Subject S1 are solely due to slow reaction times on the part of the subject and are likely at least in part due to poor estimation of decoding parameters during calibration.

---

[3] As described in Chapter 4, subjects were presented with the target for 500 ms prior to the start of cursor control for each trial.

In addition to changes in angular error resulting from decoder calibration, we find that fixed-decoder training can result in a decreases both angular error and angular error variance. Specifically, for dataset *F3*, Subject S1 was able to correct for a strong control signal bias resulting from decoder calibration over the course of nearly 600 trials spanning two days (Figure 5.2). This is consistent with previous findings of performance increases resulting from fixed-decoder training in non-human primates [Ganguly and Carmena, 2009, Sadtler et al., 2014] and during short-term ECoG BMI control by individuals undergoing pre-surgical monitoring for epilepsy management [Leuthardt et al., 2004]. Perhaps surprisingly, we do not observe reductions in angular error resulting from fixed-decoder training for any of the other datasets analyzed. Considering the low cursor control success rate for several of these datasets (notably, *F1* and *F2* for Subject S1), based on previous studies it might be expected that performance improvements could be gained through fixed-decoder training. Though it is difficult to ascertain the reason for the lack of fixed-decoder adaptation in these cases, it may be the case that the somatotopic brain control strategy constrains the type of adaptation possible for ECoG BMI systems. We will discuss this possibility further in the following section.

5.4.2 Cortical adaptation during acquisition of brain control

To assess the extent of cortical adaptation during the acquisition of ECoG BMI control, we calculated the preferred direction of all neural features over the course of 2D brain control training. In general, we find that the preferred directions of both the low-frequency ($f < 40$ Hz) and high-frequency ($f \geq 40$ Hz) features exhibited a significant change in preferred direction of approximately $60°$. Furthermore, this rotation was accompanied by a significant increase in the number of significantly-tuned features for the high-frequency band across all electrodes. An examination of changes in preferred direction for specific somatotopic electrode subsets revealed a sudden shift in preferred direction for the subset of electrodes found to be active during attempted elbow movement; no such change was observed for the subset of electrodes active during attempted hand movement. This shift in tuning was observed for both low and high frequency features, despite the exclusion of low frequency features from the decoder used for cursor control. In addition, the observed change in preferred directions was found to coincide with the decoder calibration performed on testing day 19, with the distribution of preferred directions for post-calibration trials closely resembling that of the final brain control trials on testing day 24. Furthermore, this shift in preferred direction was accompanied by a dramatic increase in brain-control performance post-calibration.

Recent work has shown that co-adaptive BMI training can facilitate improvements in performance characterized by a convergence of decoder parameters [Orsborn et al., 2012]. Similarly, a number of other BMI studies have shown that decoder calibration based on periods of closed-loop control can result in successful BMI operation [Gage et al., 2005, Mahmoudi and Sanchez, 2011, Gilja et al., 2012]. The performance improvements observed after decoder calibration on day 19 are consistent with these finding, as the combination of closed-loop training and re-estimation of decoder parameters facilitated an increase in cursor control success rate not possible with the initial decoder, despite over 3 days of fixed-decoder training. Furthermore, we also show that decoder calibration resulted in a change in the tuning of the neural population, suggesting that co-adaptive training may facilitate cortical adaptation as well as refinement of decoding parameters.

Interestingly, we observed a gradient of preferred directions across electrodes for the high-frequency range, particularly for post-calibration (day 19) and final brain control trial blocks. Perhaps not surprisingly, the preferred directions of the low-frequency band did not show this behavior, consistent with the differing degrees of spatial correlation expected from the low and high frequency bands [Brunner et al., 2009, Schalk and Leuthardt, 2011]. For the high-frequency band, we found that the preferred directions exhibited a gradual counter-clockwise rotation along the medial-lateral axis, and that the spatial correlation between preferred directions appears to be preserved throughout the course of closed-loop brain control. While the reason for this gradient is unclear, we speculate that it is a consequence of the spatial correlation between the elbow and hand-related ECoG activity. As found during BMI screening tasks (Chapter 4, Figure 4.4), attempted movements of the elbow and hand elicited gamma-band activity across large numbers of electrodes. Though this activity was sufficiently independent to enable closed-loop cursor control, a number of electrodes were found to be responsive for both elbow and hand movements. During brain control these electrodes would be activated during both attempted elbow hand movements, resulting in preferred directions between those of "elbow-only" or "hand-only" electrodes.

The spatial correlation of preferred directions highlights a key question in the development of ECoG-based BMI systems, namely, whether or not ECoG can provide the requisite signal independence to support high degree-of-freedom (DoF) control. Recent demonstrations of high-dimensional prosthetic limb control [Collinger et al., 2012, Hochberg et al., 2012] require highly-independent population activity; our finding of a preferred direction gradient across electrodes suggests that ECoG activity during attempted movements may not naturally be sufficiently independent to enable high DoF control. However, analysis human ECoG recordings has shown that intermittent synchronization of gamma-band activity does not exist beyond inter-electrode distances of 2 cm [Menon et al., 1996]. In addition, volume conductor analysis of ECoG electrode characteristics has shown that, for electrode diameters similar to those used in the presented work, sensitivity drops by more than 80% for sources at radial distances greater than 3 mm from the recording site [Wodlinger et al., 2011]. These findings provide evidence that the observed spatial correlation of ECoG recordings is primarily a cortical, rather than electrical, phenomena. Thus, through closed-loop training it may be possible to sufficiently de-correlate these sources to enable high-DoF BMI control. Recently, it has been shown that non-human primates are capable of de-correlating epidural ECoG signals between randomly-chosen pairs of electrodes [Rouse et al., 2013]. Additional work investigating the extent to which such de-correlation is possible, and how best to facilitate it, is likely needed before the full potential of ECoG BMI systems can be realized.

Perhaps our most intriguing finding is the rotation of preferred directions observed over the course of closed-loop training. Preferred direction vectors during initial brain control trials were found to be highly correlated for elbow and hand electrode subsets. Through closed-loop training, Subject S1 was able to de-correlate this activity, resulting in a rotation of the preferred direction for the hand subset of electrodes away from that of the elbow electrodes. Global rotation of the preferred directions for the majority of electrodes was then observed after decoder re-calibration, resulting in an increase in success rate. Based on these results, we postulate that Subject S1 covertly employed changes in their BMI control strategy in order to facilitate improvements in closed-loop control. Previously, adaptation of the encoding models of individual neurons has been observed during closed-loop control [Taylor et al., 2002], as well as in response to various perturbations [Jarosiewicz et al., 2008, Chase et al., 2012]. Unfortunately, a comparison of single-unit and ECoG adaptation is complicated by the use of a somatotopic control strategy for ECoG BMI control. It is likely that the type of adaptation possible during ECoG BMI control is constrained by subjects' reliance upon such abstract strategies. Thus, while it is possible to independently modify the tuning characteristics of individual neurons during closed-loop control, adaptation during ECoG control may be limited to gross modifications of the somatotopic control strategy (e.g., global rotations of PDs). Though researchers have begun investigating adaptation of ECoG signals in a BMI framework [Rouse et al., 2013], future work is needed to fully understand the limitations on such adaptation, specifically in the context of closed-loop BMI control by human subjects.

5.4.3 Limitations

There are a number of limitations to the presented work that make it difficult to draw conclusions from our findings. Perhaps most obvious of these is the experimental design. The goal of the experiments generating the data analyzed in this chapter was to assess the ability of individuals with upper-limb paralysis to control an ECoG-based BMI system. As such, decoder calibration and fixed-decoder training were used in an attempt to facilitate brain-control improvements, and not to directly assess the effects of these events. It may be the case that our inability to identify effects of decoder calibration or fixed-decoder training on is a consequence of manner in which this training was performed. For example, decoder calibration and fixed-decoder training could potentially facilitate improvements in brain-control performance at different points in the acquisition of closed-loop control. To adequately investigate this, it is likely that experiments systematically varying the BMI training procedure would have to be performed across a large subject population; such an experiment is outside of the goals of the work described in Chapter 4.

Additionally, our findings are complicated by the fact that data from Subject S2 could not be utilized to assess cortical adaptation during the acquisition of closed-loop brain control. For this subject, the goal of brain control experiments was to maximize the degrees of freedom extracted from ECoG recordings. As a consequence of this, the control strategy used for 2D brain control changed multiple times over the course of BMI experiments. Furthermore, once satisfactory brain control was achieved, the subject progressed from 2D brain control to 3D brain control; this prevented the analysis methods utilized for data from Subject S1 from being applied to data from Subject S1.

5.4.4 Additional experiments

As described in the preceding section, a number of factors prevented us from fully assessing the effects of decoder and cortical adaptation on ECoG BMI performance. In this section, we outline several additional experiments which could potentially address these shortcomings.

5.4.4.1 Effects of cortical adaptation and decoder calibration on ECoG BMI performance. In this chapter, we attempted to assess the contributions of decoder calibration and cortical adaptation during fixed-decoder training on improvements in BMI control. However, interpretation of our findings was difficult on account of the nature in which BMI experiments with subjects S1 and S2 were performed. To remedy this, we propose an experiment in which 2D BMI performance is assessed during acquisition of brain control using a somatotopic control strategy.

As opposed to individuals with upper-limb paralysis, a limited subject population, adaptation during brain control will be assessed in individuals with intractable epilepsy undergoing pre-surgical mapping. By focusing on this subject population, which should allow for an increased number of subjects, it is expected that inter-subject variability can be accounted for. The somatotopic screening procedure described in Chapter 4 will be utilized to identify those electrodes active for selected overt arm and hand movements. These movements will be mapped on to the cursor workspace in a similar manner as that utilized for subjects S1 and S2. Subjects will then perform closed-loop brain control experiments under the following conditions:

1. *Fixed decoding parameters.* Activity for individual electrodes will be classified as being preferentially active for specific arm and hand movements based on the results of the previously-described somatotopic screening procedure. Using the desired control strategy, decoding weights will be directly assigned for specific electrodes. Blocks of closed-loop cursor control will be performed with these fixed decoding parameters.
2. *Initial calibration only.* An initial decoder will be trained using a set of fully-assisted trials. Subjects will then proceed to perform blocks of cursor control, with no additional updates to the decoder performed.
3. *Co-adaptive decoder calibration with fixed updates.* Similar to [Orsborn et al., 2012], BMI training will proceed in a manner consisting of alternating periods of fixed-decoder training and decoder calibration. The number of trials between decoder calibration events will remain fixed.
4. *Co-adaptive decoder calibration with varying updates.* Rather than update decoding parameters over a fixed timescale, decoding calibration will be performed only after brain control performance has been found to plateau.

It is expected that both the rate of acquisition of closed-loop control, as well as the maximum level of performance achieved, will vary with the type of closed-loop training utilized. By identifying the training and calibration procedures which allow subjects to most rapidly acquire robust closed-loop control, we can optimize the BMI training process to maximize performance and reduce training time.

5.4.4.2 Decorrelation of ECoG signals through closed-loop BMI training. The spatial structure of the correlation between preferred directions throughout the course of closed-loop BMI training raises a key question regarding the nature of ECoG recordings, namely, how does the ability of subjects to de-correlate field potential signals from different electrodes vary as a function of inter-electrode distance? The results presented in this chapter suggest that inter-electrode correlation is preserved during closed-loop training. To better understand this, we propose an experiment in which the ability of subjects to modify the mapping for a somatotopic control strategy is assessed.

Individuals with upper-limb paralysis will be implanted with ECoG electrode in a similar manner as described in Chapter 4. Electrodes will be classified by their responses during attempted arm and hand movements; those exhibiting preferential activity for individual movements will be used as the basis for closed-loop cursor control. Similar to BMI experiments with subjects S1 and S2, these movements will be mapped onto the cursor control space using a somatotopic control strategy. However, rather than fit a set of decoding parameters during a calibration procedure, a modified version of the population vector algorithm (PVA) will be used, with the preferred directions for electrodes assigned based on the desired somatotopic mapping. Brain control experiments will then be performed under the following conditions:

1. *Baseline control.* Subjects will first perform a set of closed-loop experiments using the standard somatotopic control strategy to character baseline control performance.
2. *Visuomotor rotation of the PVA decoder.* Once satisfactory control has been obtained using the standard somatotopic mapping, a visuomotor rotation will be applied to the PVA decoder such that all preferred directions are rotated in a similar manner. Subjects will then be required to re-establish closed-loop control. As the target directions will remain unchanged from the baseline control condition, the visuomotor rotation will assess the ability of subjects to generalize the somatotopic mapping in order to generate cursor movements in arbitrary directions.
3. *Rotation of preferred directions for individual movements.* Finally, the preferred directions of electrodes corresponding to individual arm and hand movements will be rotated in the PVA decoder. As opposed to a pure visuomotor rotation, this will require subjects to generate a new somatotopic mapping in order to re-establish control.

In addition to assessing the quality of closed-loop control achieved for these conditions, the effect of the various experimental conditions on inter-electrode correlation will be determined. By providing a better understanding of the constraints on the adaptability of ECoG activity, it is believed that this work will provide insight into the capability of ECoG to support high degree-of-freedom BMI systems.

5.4.5 Conclusions

Though interpretation of our analysis investigating changes in angular error as well as cortical adaptation during closed-loop ECoG control are complicated by the particulars of the experimental design used, we nonetheless believe the presented work provides crucial insight into the potential for adaptation during electrocorticographic BMI control. By better understanding the interplay between decoder calibration and subject training, better BMI training procedures may be developed which shorten subject training time and maximize control performance. Furthermore, improved understanding the ability of users to modify the tuning properties of neural features combined with the stability of ECoG recordings may allow for enhancement in the degree of control capable possible with ECoG-based BMI systems in a chronic implantation scenario. Ultimately, we believe that these findings, as well as the additional experimental work proposed, lay the groundwork for the development of high-performance, clinically-viable brain-machine interface systems using ECoG.

6.0 EVALUATION OF A CHRONICALLY-IMPLANTED ELECTROCORTICOGRAPHIC ELECTRODE GRID IN A NON-HUMAN PRIMATE[1]

In Chapter 4, it was shown that individuals with upper-limb paralysis can learn to control an ECoG-based BMI system over a the course of 28 days. Though this serves as an essential demonstration of the potential of an ECoG-based brain-machine interface, a clinical BMI system must be capable of sustaining performance over a much longer period of time, preferably without the need for experimenter interventions such as decoder re-calibration. Unfortunately, apart from limited work investigating the stability of offline decoding of arm movement trajectories from ECoG recordings [Chao et al., 2010], relatively little is known about the capability of electrocorticography to support a chronically-implanted BMI system. In this chapter, we explore BMI control capability and longevity of ECoG implants for BMI applications through an evaluation of a chronically-implanted ECoG electrode grid in a non-human primate. First, BMI control capability using an ECoG array implanted over premotor, primary motor, and sensory cortex is demonstrated. Next, recording quality over time is examined by freezing the neural decoder used for real-time control and observing performance over multiple days. Finally, the viability of ECoG implants for long term applications is investigated by examining recording quality obtained during overt hand movement and histological data.

6.1 BACKGROUND

ECoG recordings presumably reflect the activity of thousands of neurons, and as a result may be less sensitive to changes in activity from any individual neuron [Moran, 2010, Leuthardt et al., 2004]. It has been postulated that ECoG may be capable of providing robust and stable recordings, thus eliminating the need for daily updating or retraining of the decoder which translates multi-channel neural recordings into BMI control signals. Several studies have begun to shed light on this, demonstrating long-term offline decoding of movement kinematics from subdural [Chao et al., 2010] and epidural [Shimoda et al., 2012] ECoG recordings using fixed decoding parameters. To-date, long-term closed-loop control of an ECoG BMI system with a fixed neural decoder has not been demonstrated.

---

[1] This chapter consists of a modified version of a manuscript currently in review: Degenhart AD, Eles J, et al., *Histological evaluation of a chronically-implanted electrocorticographic electrode grid in a non-human primate.* The version presented here incorporates the results of a fixed-decoder BMI experiment which has been excluded from the manuscript in preparation based on the apparent contribution of artifacts to closed-loop control.

Such stability is in contrast to SU/MUA activity, where recordings from single electrodes are much more sensitive to small changes in the local environment surrounding the recording site. Several research groups have demonstrated recording variability over time using single and multi-unit activity [Schwartz et al., 2006, Rousche and Normann, 1998, Schwartz, 2004, Flint et al., 2013]. However, it has also been shown that SU/MUA as well as LFP activity can provide long-term stable recordings, enabling closed-loop control of a computer cursor for up to a year with a fixed decoder [Ganguly and Carmena, 2009, Flint et al., 2013].

Single and multi-unit recordings can suffer from signal deterioration, manifested as a reduction in the number of electrodes recording individual neurons or a decrease in signal amplitude [Collinger et al., 2012, Simeral et al., 2011, Chestek et al., 2011]. This degradation has been attributed to a host tissue response that begins with trauma from initial probe penetration. The mechanical disruption of tissue from implantation develops into a chronic inflammatory response hallmarked by an aggregation of astrocytes and microglia in a glial scar around the probe, as well as progressive neurodegeneration within a $100\mu m$ radius around the implantation site [Polikov et al., 2005, McConnell et al., 2009, Biran et al., 2005]. Glial ensheathment is theorized to act as an electrical insulator for the electrodes, and along with neurodegeneration, is believed to contribute to signal decay [Turner et al., 1999]. Neurodegeneration is associated with the presence of neurotoxic cytokines in the "kill zone" around the implant, which could be produced by the reactive astrocytes and microglia of the glial scar [Biran et al., 2005]. It has been theorized that mismatch between the stiffness of electrode devices and the brain produce shear force during natural head movements; this chronic mechanical stress may underlie chronic inflammation [Biran et al., 2007]. Additionally, blood-brain barrier disruption following implantation leads to a build-up of neurotoxic species in the vicinity of the electrodes [Saxena et al., 2013, McConnell et al., 2009].

Since ECoG arrays do not penetrate the cortex, they avoid intracortical shear and blood-brain barrier disruption, potentially making encapsulation less likely to occur during long-term implantation [Szarowski et al., 2003, Bjornsson et al., 2006, Biran et al., 2005]. Despite the promise of ECoG in BMI and neuroscience applications, very few studies evaluating long-term host-tissue response to either epidural or subdural grids have been conducted. While arrays have been implanted for over 1 year in humans and in non-human primates with viable neural recording [Shimoda et al., 2012, Morrell and RNS System in Epilepsy Study Group, 2011], connective tissue overgrowth has been observed in epidural implants after only one week in rats, and leptomeningeal inflammation has been seen after 6 months of sub-dural implantation in the same model [Schendel et al., 2013, Schendel et al., 2014, Henle et al., 2011]. It has been shown that tissue growth around epidural ECoG grids implanted in rats exhibits dural thickening and scar tissue encapsulation over the top of the array as early as one month following implantation [Schendel et al., 2014]. While the relationship between connective tissue proliferation and electrical impedance change is unclear, it is possible that the tissue acts as insulation between electrode sites and the brain, reducing signal quality. While no evidence of tissue fibrosis has been noted in epilepsy monitoring arrays implanted in humans for up to 30 days [Fountas and Smith, 2007, Van Gompel et al., 2008, Wong et al., 2009, Wang et al., 2013a], the long-term outlook remains unclear. Further, while qualitative evidence exists showing that epidural ECoG grids do not result in inflammatory microglia or astrocyte activation in the cortical tissue beneath the implant, there appear to be no studies exploring the impact of subdural grids on cortical tissue health.

6.2 METHODS

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

The major events in the study were as follows: after implantation, a series of closed-loop brain control experiments were performed to evaluate the performance of an ECoG-based BMI system utilizing a static decoder. Shortly before explantation of the electrode grid, 15 sessions of a standard center-out reaching task were conducted in order to characterize the ECoG signal modulation during overt hand movements. Following explantation, a histological analysis of the electrode and cortical tissue was performed to evaluate the nature of the foreign body response to the electrode grid.

6.2.1 ECoG grid implantation surgery

A male Rhesus macaque (*mucaca mulatta*) was anesthetized, and a craniotomy was performed over the left motor and premotor cortex. The dura was retracted to expose an area approximately 2x2 cm between the arcuate and central sulci (Figure 6.1B). A custom-built 15-channel ECoG grid (Figure 6.1A, PMT Corp, Chanhassen, MN, USA) was placed directly on the exposed brain surface (Figure 6.1C), and the dura and the bone were reapproximated. Wires from the grid were connected to a Cereport pedestal connector (Blackrock Microsystems, Salt Lake City, UT, USA) affixed to the skull.

6.2.2 Neural recording and task control

Signals from the ECoG grid were recorded with a g.USBamp Biosignal Amplifier (g.tec Medical Engineering), and sampled at 1200 Hz. All recording, online processing, decoding, task control and presentation was performed using the Craniux Brain Computer Interface system [Degenhart et al., 2011b]. Dura-facing electrodes 4 and 13 were used as reference and ground electrodes for all recordings (Figure 6.1A). Visual stimuli were presented on a 22-inch computer monitor placed approximately 2.5 feet from the animal.

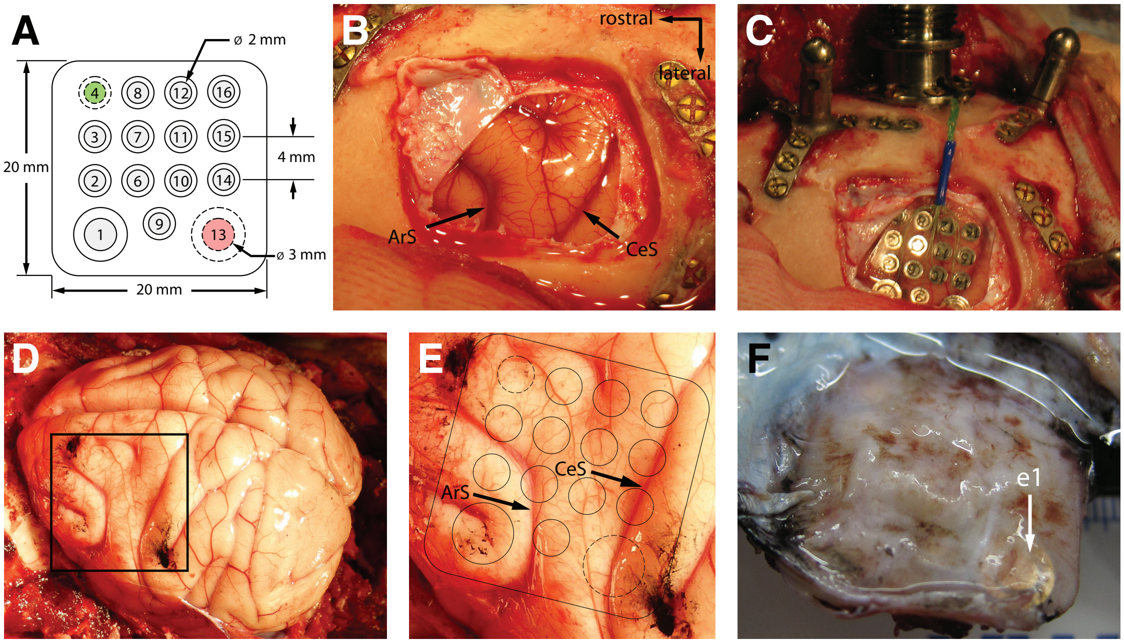

Figure 6.1: (*A*). Schematic of the electrode grid. Shaded gray areas indicate the exposed area for each electrode. Electrodes 4 (green) and 13 (red) are upside-down (i.e. dura-facing) and serve as the reference and ground, respectively, for neural recordings. A top view of grid is shown (i.e. the dura-facing side). (*B*). Craniotomy detail showing the exposed cortex (ArS: arcuate sulcus, CeS: central sulcus). (*C*). Placement of the electrode grid. (*D*). Brain immediately following removal from the skull during postmortem explantation. Black marks indicate the rostral-medial and caudal-lateral corners of the grid. (*E*). Cortical tissue underneath the ECoG grid indicated by the black box in (*D*). (*F*). Underside of the encapsulated grid following explantation. The location of electrode 1 (e1) is indicated by the white arrow.

6.2.3 Brain control task

Brain control experimental sessions began on day 208 post-implant and ended on day 299 (51 testing sessions spanning 85 calendar days)[2].

In order to derive neural features for closed-loop control, spectral estimation was performed in real-time using the Burg Autoregressive (AR) method [Robinson, 1982] over the 40 to 180 Hz range (25th order, 10 Hz bands). These gamma and high gamma bands were chosen because they are typically spatially localized and informative about underlying neural processes [Crone et al., 1998a, Crone et al., 2006, Miller et al., 2007a, Heldman et al., 2006]. Spectral estimates were calculated every 33 ms using a sliding window of 300 ms of raw data. AR data were log-transformed, then normalized to pseudo Z-scores relative to a baseline condition [Degenhart et al., 2011b]. Normalized spectral estimates were used as the neural features for closed-loop BMI control. Twelve of the 15 channels available on the grid were used for control, excluding those electrodes providing reference and ground signals (e4 and e13), and a third (e2) to which electrical connectivity was lost before experiments began[3]

Linear regression was used to determine a mapping between neural features and cursor movement. This mapping took the form of a matrix of weights ($\mathbf{W}$) to be applied to the neural features ($\mathbf{f}$) for decoding such that the predicted cursor movement velocity ($\hat{\mathbf{v}}$) was a weighted sum of neural feature input as described by Equation 6.1. The decoding weights were calculated using Equation 6.2:

$$\hat{\mathbf{v}} = \mathbf{W}\mathbf{f} \tag{6.1}$$

$$\mathbf{W} = \mathbf{V}\mathbf{F}^{\dagger} \tag{6.2}$$

where $\mathbf{V} = [\mathbf{v}_1, \ldots, \mathbf{v}_N]$ and $\mathbf{F} = [\mathbf{f}_1, \ldots, \mathbf{f}_N]^\top$ are concatenated matrices of observations of the desired cursor movement direction and associated neural features, respectively, collected during decoder calibration. For calibration data, the desired cursor movement direction was defined as the instantaneous unit vector pointing from the cursor to the target, averaged over a trial. The superscript "†" denotes the Moore-Penrose pseudoinverse. Initial calibration data was acquired from one block of trials during which the animal observed automated movements of the cursor, with subsequent updates of the decoder using all trials from a single block (see below) of closed-loop brain control data.

The monkey was trained to perform a two-dimensional 8-target center-out cursor task. At the beginning of each trial, a cursor appeared at the center of a computer screen. Simultaneously, a pseudo-randomly selected target appeared at one of 8 possible targets approximately 10cm from the starting position of the cursor. During a 500 ms "hold" window immediately after target onset, the cursor was held fixed at the

---

[2] Delays resulting from connector damage and exploratory BMI sessions prevented initiation of experimental procedures until 208 days post-implantation. Additional closed-loop control experiments were conducted between days 300 and 541 post-implantation. Unfortunately, complications arising from the presence of artifacts during closed-loop control (see Section 6.3.2) limited the utility of these experimental sessions. As a result, these sessions were excluded from further analysis

[3] It was later found that this was an issue with the cable connecting the Cereport pedestal and the neural amplifiers and not that of the electrode grid specifically.

center of the visual display. Afterwards, the cursor moved under brain control. The monkey was required to move the cursor to the displayed target within 3 seconds. If the cursor overlapped with the target for 100-200 ms (randomly determined each trial), the trial was considered successful and the monkey was immediately given a water reward.

The task was run in 40-trial blocks, with multiple blocks run each testing day. In the initial testing sessions, the task software guided training for brain control by actively moving the cursor directly to the target ("active assist"). Following one fully assisted observation block (100% active assist) an initial decoder was trained, and the amount of assistance provided to the animal was reduced. This assistance was incrementally lowered between blocks until the cursor was under full brain control (0% active assist). The decision to lower the assist level was based on the success rate, such that high success rates over multiple blocks prompted a decrease in the assist level. Computer assistance was not provided to the animal after testing day 7, thus success rates reported for subsequent sessions represent proficiency with full brain control.

During the initial period of training, the decoding weights were updated periodically between blocks. Updates consisted of a weighted average of the old and new decoding weight parameters, such that changes in the decoder were gradual [Wang et al., 2013a]. Once satisfactory closed-loop performance was achieved under full brain control (testing day 6) the neural decoder was frozen, and remained constant for the remainder of the brain-control experiment.

6.2.4 Hand control task

Hand control experiments began on day 542 post-implant and ended on day 562 (15 testing days spanning 21 calendar days). During these experiments, the animal performed a standard 2D center-out task in a virtual environment, with the position of the hand controlling the location of a computer cursor in a two-dimensional plane. Hand position was tracked in real-time using the Phasespace optical tracking system (Phasespace, San Leandro, CA) and rendered on a computer screen as a sphere in a virtual workspace. Trials began with the appearance of the cursor and central target; the animal was then required to move the cursor to the central target, holding the cursor over the target for 400–600 ms. One of eight peripheral targets would then appear, to which the animal was required to reach. A target hold time of 200 ms was enforced. The animal was provided with water rewards immediately following successful completion of a trial.

Time domain data was transformed into the time-frequency domain using the Burg autoregressive method (0 - 200Hz range, 2Hz frequency bands, 100th order, 100ms non-causal window, 33ms step size)[4], log-transformed, then normalized to pseudo Z-scores relative to the spectral power during the central target hold period. Trials were then manually aligned to movement onset using the cursor speed profile for each trial.

---

[4] As opposed to closed-loop control, where spectral estimates parameters were constrained by real-time processing requirements, spectral estimates performed for offline analysis were performed using parameters more suitable for a detailed examination of the spectral content of recorded signals.

6.2.5 Artifact identification

Hand and brain control trials were visually examined for artifacts in both the time and time-frequency domains. These artifacts were characterized by large-amplitude, broadband transients across the majority recording electrodes, and were believed to be the result of jaw movements based on their consistent appearance during the reward period of each trial [Kelly et al., 2013]. Time-domain ECoG signals were band-pass filtered from 5 - 250 Hz (4th order butterworth filter) and plotted along with time-frequency data for all electrodes. Trials were marked as containing artifacts if they exhibited both large-amplitude voltage transients and broadband frequency modulation (> 40Hz) across all electrodes. Trials containing artifacts during the central hold or target acquisition periods were excluded from further analysis (hand control data only).

6.2.6 Cosine tuning analysis

ECoG signals were fit to a standard cosine tuning model [Georgopoulos et al., 1982] in order to characterize the directional modulation of ECoG signals during both brain and hand control tasks. Normalized spectral data was first averaged over the $[70 - 110]$ Hz frequency band and from movement onset to target acquisition (brain control) or the $[-100ms, 100ms]$ interval relative to movement onset (hand control) for each trial. Cosine tuning curves were fit to this averaged data using equation 6.3.

$$f = b_0 + b_x x + b_y y + \epsilon \tag{6.3}$$

6.2.7 Explant

The total implant duration was 666 days. Surgical complications negated the possibility of perfusing the animal before removing the brain. After exposure of the skull, the original bone flap was removed to expose the dura. The skull, dura, and encapsulated electrode grid were then removed as a single piece and the entire brain was extracted. The brain and the encapsulated array were then placed in a 10% formalin + 10% glycerin solution for 8 days followed by 10% formalin + 20% glycerin for 26 days to fixate the tissue. Fixated tissue was then frozen and sectioned into $50\mu m$ sections for staining. Sections were cut perpendicular to the central sulcus. The electrode grid was carefully removed from the encapsulation “envelope,” which was similarly fixed for 6 days and then stored in phosphate tris azide (PTA) solution until it was cut into $50\mu m$ sections for staining. Encapsulation tissue was cut perpendicular to the placement of the grid.

6.2.8 Immunohistochemistry

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

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

6.2.9 Confocal imaging

Confocal imaging was performed with an Olympus Fluoview 1000 confocal scanning microscope (Olympus). All images were taken with a 20X or 40X objective to optimize cellular resolution and image frame size, and were taken at multiple focal depths for each frame as to image the full depth of a tissue slice. This ensured that image analysis was not biased by choice of a single image depth. Confocal laser power, photomultiplier tube voltage (the sensitivity of the image detector), and photomultiplier offset (background level of image detector) were selected to ensure that image pixels did not exceed upper or lower detection limits. Images ($n = 5$ per stain) were collected from cortical regions directly under randomly selected electrode sites on the ECoG array or random sites of both the dura-facing and cortex-facing sides encapsulation from the middle of the tissue envelope. For cortical imaging, images from the contralateral hemisphere were collected for comparison. Images were matched to the same sagittal slice depth and anterior-posterior region of interest as the ipsilateral hemisphere. Tissue encapsulation images were compared to images from random regions of interest of control dura mater retrieved from $> 2cm$ from the tissue encapsulation.

To determine cortical layers and cortical thickness, disconnected images of cortex were stitched to create continuous high-resolution images of the entire cortical depth using Fiji, an Image-J (NIH) plug-in [Preibisch et al., 2009]. Layers I/II-III were discerned from layer V by the location of layer V giant pyramidal cells [Matelli et al., 1991]. Stitched images were used to measure cortical depth ($n = 5$) between conditions.

Neuronal and microglial cell densities were determined for layers I/II-III and V with hand counting facilitated by Image J Cell Counter ($n = 5$). Because GFAP labels extensive networks of astrocytic processes, discerning individual cell bodies for cell counting is inconclusive. Thus, the proportion of cortex occupied by reactive astrocytic signal (% GFAP signal) was determined by first setting a pixel intensity threshold to the mean pixel value of layer I/II-III, where the most intense signal was localized. Because the majority of pixels in a given image are not GFAP-signal, the pixels below the mean can be discounted as noise. Once thresholded, the GFAP signal was determined by automating a count of the non-zero pixels ($n = 5$). Implanted cortex and contralateral cortex were compared for all metrics by t-tests with significance defined to be $p < 0.05$.

Encapsulation and dura mater images were qualitatively analyzed to determine the identities and morphologies of encapsulation cells. Based on the cell specificity of the antibodies, vimentin(+)/Iba-1(+) and vimentin(-)/Iba-1(+) cells were considered to be macrophages/microglia, while vimentin(+)/Iba-1(-) cells were considered to be fibroblasts. Multi-nucleated cells were considered cells that contained more than one Hoescht 33342 labeled nuclei in a single cell body.

6.2.10 Collagen-I imaging

Collagen I, a key component of tissue encapsulation, can be visualized using second-harmonic generation (SHG) imaging. SHG imaging takes advantage of a second order non-linear optical property of collagen type I to visualize an intrinsically generated optical signal that can be used to locate and quantify collagen I in tissue. This is preferred to traditional histological staining protocols, which have been shown to have less signal specificity and require chemical processing that may alter the tissue quality [Strupler et al., 2007].

SHG images of tissue encapsulation and dura mater were captured using a laser through a Nikon A1Plus multiphoton scanning confocal microscope and Nikon NIS-Elements Microscope Imaging Software. SHG was generated at an $830nm$ wavelength, and signal was collected via a bandpass filter that isolated tissue auto-fluorescence ($435 - 700nm$) and SHG signal ($380 - 400nm$). Images were taken with a 25X objective to maximize signal resolution and imaging frame; stitching software (EIS-Elements Microscope Imaging Software, Nikon) was used to consolidate disconnected images to make a seamless, high-resolution image of the encapsulation through the dorsal-ventral plane.

Encapsulation and dura mater thickness was determined by measuring average thickness of tissue extent as denoted by auto-fluorescence. Because SHG signal was confined within an uninterrupted, fibrous area, percent SHG-signal was measured by dividing the average thickness of SHG area divided by the total tissue thickness. Such measures were generated for encapsulation tissue from the cortex-facing and dura-facing sides, as well as for control dura mater ($n = 5$ for all groups). Encapsulation and dura matter thickness and percent SHG signal were compared between cortex-facing encapsulation, skull-facing encapsulation, and control dura mater groups by one-way ANOVA tests with Tukeys post-hoc tests. Significance for all comparisons was defined to be $p < 0.05$.

6.3 RESULTS

6.3.1 Closed-loop brain control with a static decoder

Across 46 testing days (77 calendar days) of BMI control with a static decoder, performance was often over 70% correct (Figure 6.2). This would appear to be impressive control, considering not only that the decoder was held constant, but also the tight timing constraints in the task: the animal had three seconds in which to acquire the presented targets. Of particular interest is a gradual upward trend in performance between testing days 16 and 33; with a static decoder, performance improvements can most likely be attributed to learning on the part of the animal. Note that the decoder was being calibrated for the first six days; we include these data for the sake of completeness.

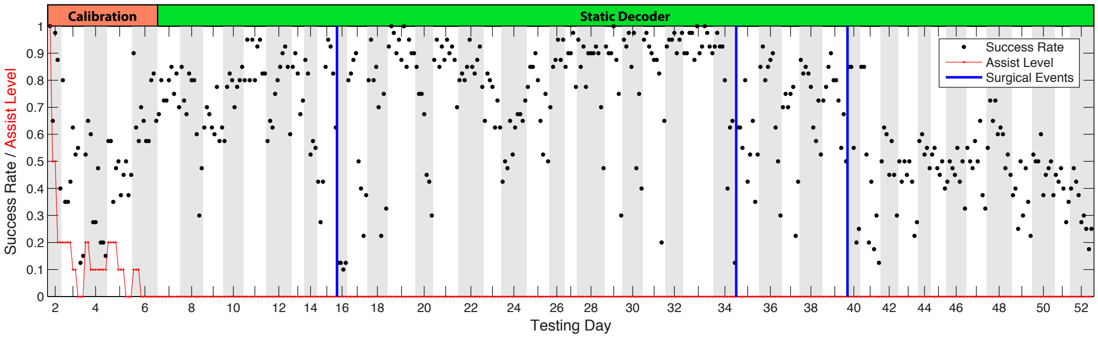

Figure 6.2: *Task performance is high during closed-loop cursor control with a static decoder.* Success rate for 40-trial blocks are indicated by black dots. Assist level and surgical events are indicated by the red and blue lines, respectively. Alternating gray and white bands indicate individual testing days.

Daily brain-control sessions were typically characterized by higher success rates during early trials and decreasing success rates at the end of each session; this reduction in performance was attributed to a decrease in motivation during the day. Intermittent days with abnormally low performance were observed (testing day 24), though subsequent days showed performance return to a high level. Figure 6.3 shows the maximum and mean success rates per testing day as a function of day post-implantation.

Surgical procedures were performed following experimental testing days 15, 34, and 39. These procedures involved attempts to repair exposure of the wire bundle running form the electrode grid to the pedestal. Drastic decreases in performances were observed following each of these interventions. Following the last of these procedures performance was observed to continue to decrease until testing day 52, when brain-control with the static decoder was halted. Reductions in performance during this period were characterized by an inability of the animal to move the cursor towards the lower-right quadrant of the workspace.

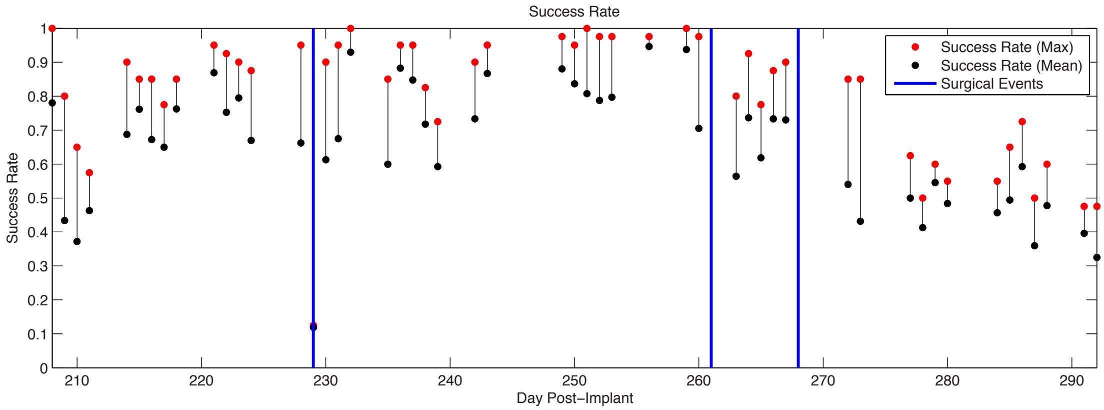

Figure 6.3: *Task performance across testing days during fixed-decoder brain control.* Maximum (red) and mean (black) success rate are shown for each testing day. Blue lines indicate days in which surgical interventions were performed.

6.3.2 Influence of artifacts on brain control

Unfortunately, offline examination of data collected during closed-loop brain control revealed the presence of artifacts presumed to be the result of jaw movement. Artifacts were identified during the reward period of the majority of trials, but were also observed during the target acquisition period of a number of trials. Figure 6.4 shows an example of a trial where artifacts were identified during ITI (inter-trial interval), center-hold, and reward trial states, while Figure 6.5 shows an example of a trial containing artifacts during all trial states.

In order to determine if the presence of artifacts influenced the animal's ability to control the computer cursor, trials were examined for artifacts for four days spanning the entirety of the fixed-decoder experiment. Table 6.1 presents data summarizing the presence of successful trials containing artifacts for testing days 2 (brain control calibration), 7 (early brain control performance), 33 (peak brain control performance), and 52 (final brain control performance). While relatively few successful trials contained artifacts during decoder calibration (9%), the presence of artifacts increased during early brain control (22%), peak brain control (41%), and final brain control (23%).

When the relationship between artifact rate and target direction was examined, it was found that artifacts were predominately present for the bottom-left and bottommost targets. Figure 6.6 depicts artifact rate as a function of target direction for the selected brain control sessions. Artifact rates were particularly biased with target direction for later brain control sessions (peak brain control performance and final brain control performance), suggesting that artifacts likely contributed to the animal's ability to control the computer cursor.

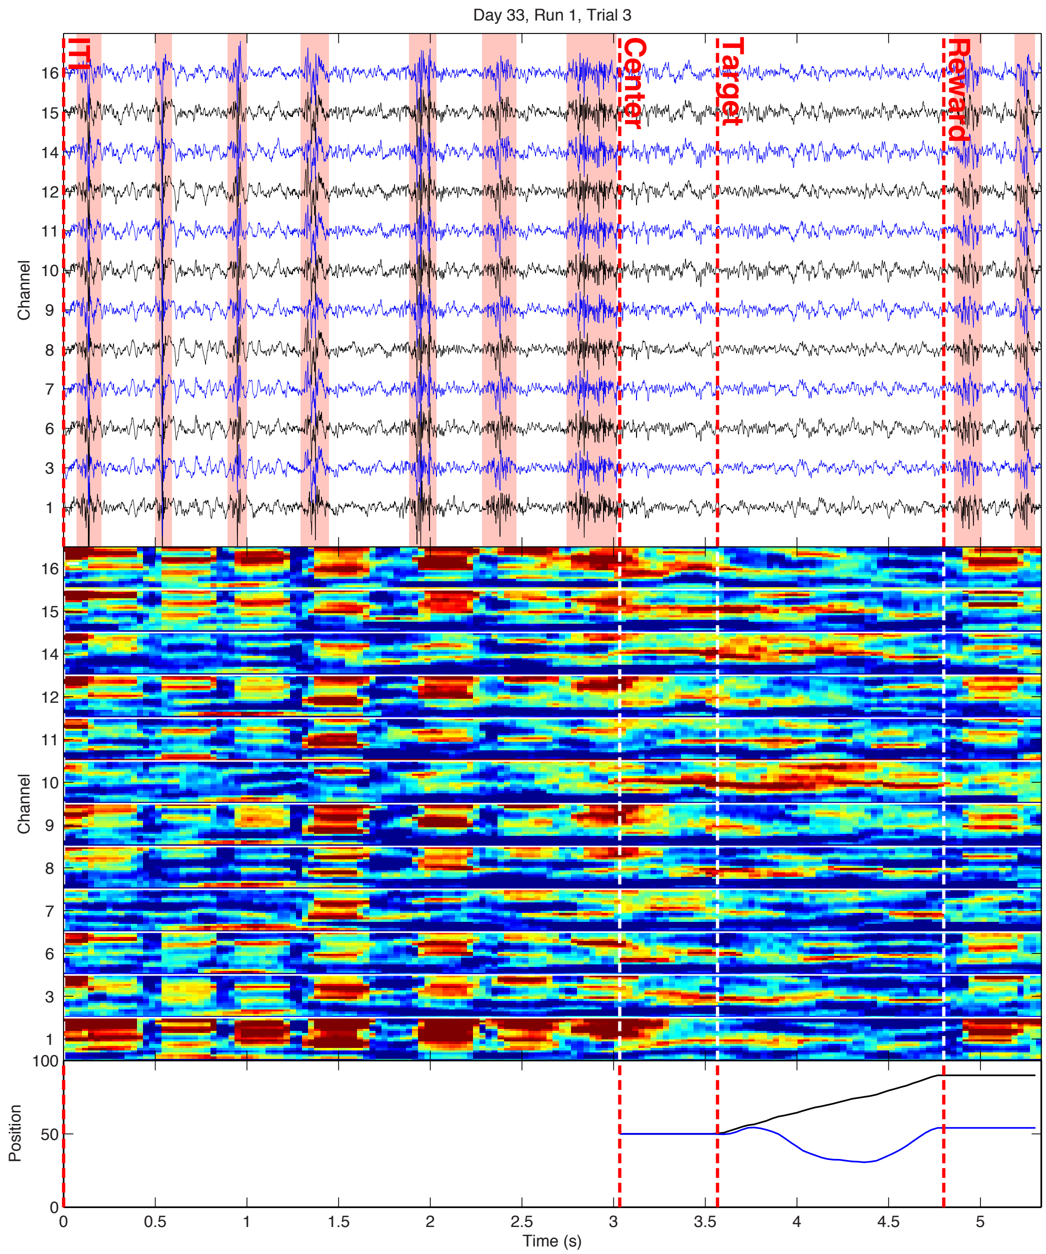

Figure 6.4: *Example of a brain-control trial without artifacts during the target acquisition period.* Time-domain ECoG voltages (*top*), time-frequency data (*middle*), and cursor position (*bottom*) for a representative trial without artifacts during the target acquisition trial state. Vertical dashed lines indicate onset of trial states. Red-shaded regions on the top plot indicate manually-identified artifacts. X and Y cursor position is indicated by the black and blue lines, respectively, in the bottom plot.

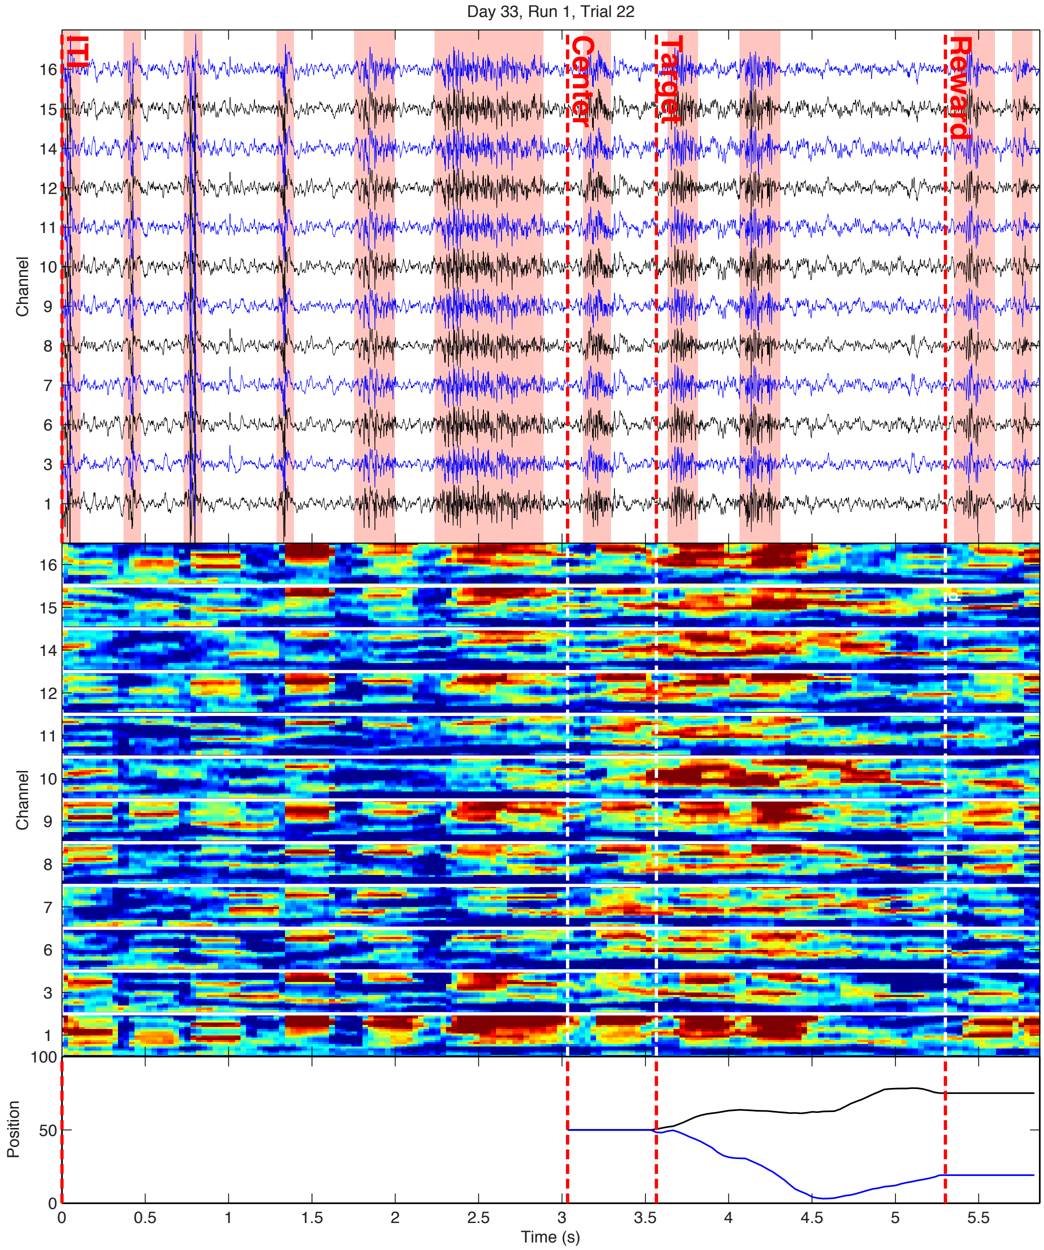

Figure 6.5: *Example of a brain-control trial with artifacts during the target acquisition period.* Time-domain ECoG voltages (*top*), time-frequency data (*middle*), and cursor position (*bottom*) for a representative trial with artifacts during the target acquisition trial state. Vertical dashed lines indicate onset of trial states. Red-shaded regions on the top plot indicate manually-identified artifacts. X and Y cursor position is indicated by the black and blue lines, respectively, in the bottom plot.

Table 6.1: *Summary of artifacts in brain-control data for selected testing days.* Data are provided for four exemplar testing days spanning the duration of closed-loop cursor control experiments. Artifact rates are calculated based on the number of successful trials where at least one artifact was identified during the target acquisition (control) trial epoch. ST: Successful trials, AT: Artifact trials, AR: Artifact rate.

				Target Direction AR (ST)							
Day	ST	AT	AR	1	2	3	4	5	6	7	8
2	156	14	0.09	0.15 (20)	0.00 (20)	0.05 (21)	0.05 (19)	0.05 (19)	0.12 (17)	0.30 (20)	0.00 (20)
7	274	59	0.22	0.06 (35)	0.09 (35)	0.06 (33)	0.31 (36)	0.29 (34)	0.44 (36)	0.19 (26)	0.26 (39)
33	450	185	0.41	0.33 (55)	0.05 (56)	0.07 (59)	0.20 (55)	0.45 (55)	0.83 (58)	0.75 (55)	0.61 (57)
52	117	27	0.23	N/A (0)	0.10 (10)	0.00 (26)	0.06 (32)	0.35 (31)	0.69 (16)	1.00 (2)	N/A (0)

6.3.3 ECoG modulation during brain and hand-controlled cursor tasks

Figure 6.7 shows time-frequency data for one selected electrode (e10) at various points during the lifetime of the implant[5]. Early calibration trials were characterized by weak modulation of spectral power with target direction (Figure 6.7A). However, by the beginning of the static decoder experimental sessions modulation was observed to increase (Figure 6.7B), with electrodes remaining strongly modulated by the task throughout periods of high performance (Figure 6.7C). This characteristic pattern of modulation was found to disappear by the end of fixed decoder sessions (Figure 6.7D), with electrodes exhibiting drastically reduced modulation by the end of closed-loop experiments (testing day 52).

Normalized spectral data for each target and electrode during peak brain control performance is shown by Figure 6.8. While many of the electrodes (e.g. e10, e14, e15) are modulated strongest by targets in the down-right direction, a few electrodes (e12, e16) exhibit strongest modulation for targets in the down-left direction. Based on the prevalence of artifacts during these trials (Figure 6.6C), it appears that modulation for these electrodes is likely driven by the presence of artifacts.

Figure 6.9 shows the results of 2D cosine tuning analysis performed on brain control data. Tuning curves during initial brain control training did not exhibit significant modulation with target direction. After initial brain control training, all electrodes were found to become significantly tuned, with preferred directions tightly clustered in the lower-right quadrant of the workspace. The distribution of preferred directions was found to become less clustered during peak brain control performance. In addition, a slight increase in depth of modulation was also observed from initial to peak brain control sessions. By the end of the fixed decoder experiment preferred directions were found to have shifted dramatically, with a decrease in depth

---

[5] As shown by Figures 6.5 and 6.6, it appears that the presence of artifacts contributed to the animal's ability to control the computer cursor. Thus, these data contain artifacts which resulted in broadband increases in spectral power which overlap with the gamma-band modulation expected in ECoG recordings. As the intent of this (and subsequent) figures is to demonstrate spectral modulation underlying closed-loop cursor control, trials containing artifacts have not been removed from these data.

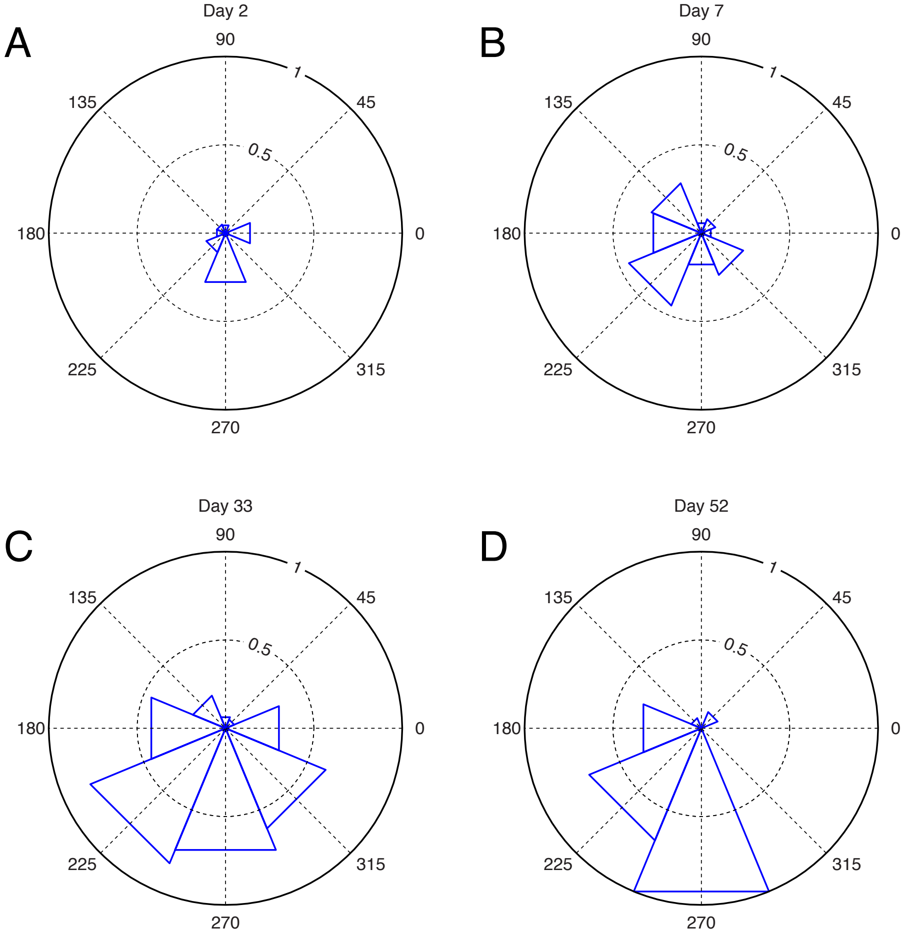

Figure 6.6: *Artifact presence rate per target direction for selected brain-control testing days.* Fraction of successful trials with artifacts during the target acquisition state is shown for testing days 2, 7, 33, and 52 for each of the 8 target directions.

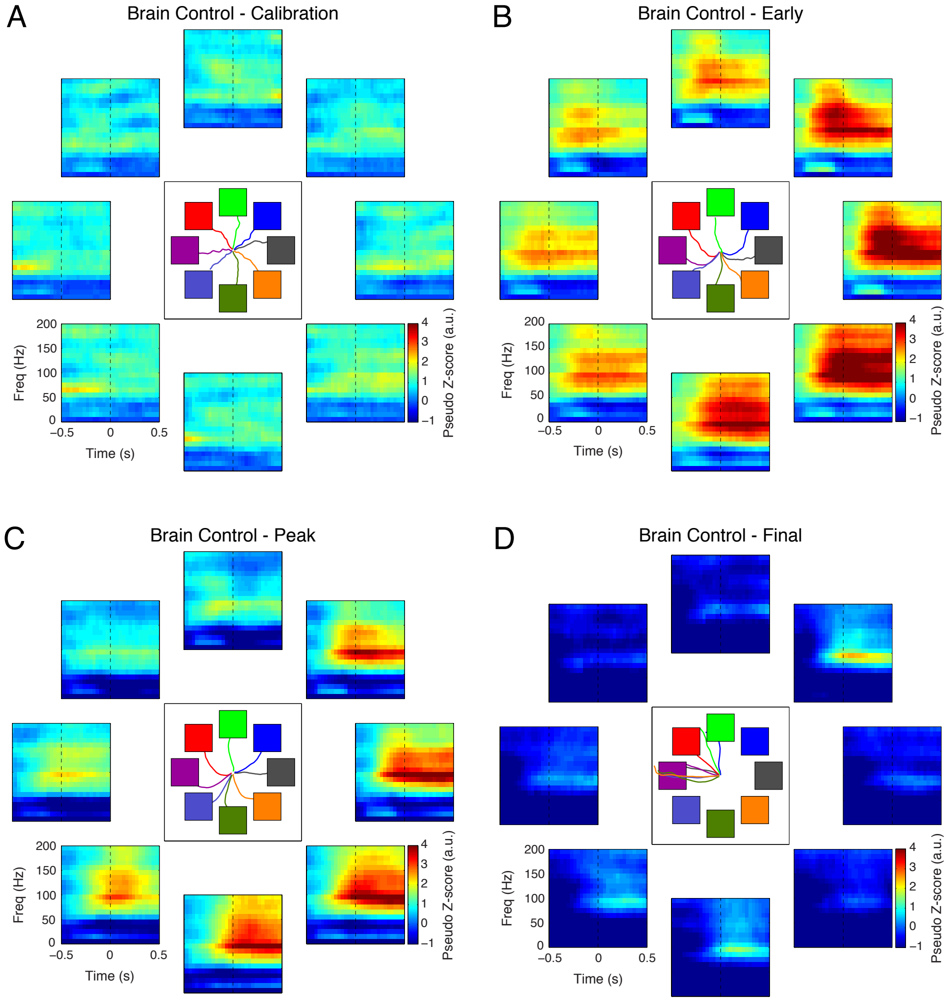

Figure 6.7: *ECoG signal modulation is directionally-tuned during 8-target center-out brain-control.* Averaged time-frequency data is shown for a single electrode (e10) for all tasks. Averaged trajectories for each target are shown in the center plot of each panel. (A). Initial brain control training (testing days 2 and 3, $N = 328$ trials). (B). Initial brain control performance (testing days 7 and 8, $N = 720$). (C). Peak brain control performance (testing days 32 and 33, $N = 760$). (D) Final brain control performance (testing days 51 and 52, $N = 600$). Time-frequency data was normalized with respect to a fixed set of "baseline" parameters collected during the first brain control session (i.e. converted to pseudo-Z-scores). Note that the color scale for panels (A) – (D) is identical.

of modulation also observed. The similarity between the distribution of movement artifacts (Figure 6.6) and the spatial pattern of gamma-band modulation suggest that artifacts contributed substantially to the observed tuning of ECoG modulation.

ECoG signals recorded during hand control experiments, conducted between day 542 and 562 post-implantation, also exhibited clear modulation with target direction (Figures 6.10, 6.11). Characteristic decreases in the mu and beta frequency bands, in conjunction with increases in the gamma band, were observed. Gamma band modulation was found to be strongest over the $70 - 100$ Hz frequency range, and was tightly locked to movement onset. As opposed to data collected during brain control, trials containing artifacts were removed prior to analysis of hand control data. Thus, the modulation observed during hand control can only be attributed to arm movements on the part of the animal.

Figure 6.11 shows normalized spectral data for each target and electrode during center-out hand control trials. As opposed to brain control trials, nearly all electrodes show preferential modulation for upper-left targets. This is confirmed by the results of the cosine tuning analysis shown in Figure 6.12. The preferred directions of high-gamma band activity from all electrodes was tightly clustered during the hand control task, as would be expected by the highly correlated time-frequency data shown by Figure 6.11. Though the distribution of preferred directions observed during hand control differ from those seen during brain control, it is not necessarily the case that similar distributions would be expected given that no constraints were imposed on hand movement position during the brain control task. In addition, as previously mentioned, the presence of artifacts during brain control likely was responsible for the tuning of gamma-band activity to target direction; thus, it should not be expected that brain control and hand control data should exhibit similar preferred direction distributions.

6.3.4 Cortical architecture

Chronic ECoG implantation did not significantly affect cortical cytoarchitecture. Using Nissl staining and immunohistochemistry with antibodies targeted to neurons (NeuN), microglia (Iba-1), and reactive astrocytes (GFAP), we evaluated cortical thickness and cell densities in layers I/II-III and V in cortical regions under the ECoG array and in the same regions in the contralateral hemisphere. These results are shown in Figure 6.13.

Noting a mechanical depression of the brain under the ECoG grid post-explantation, we measured and compared cortical thickness between hemispheres. Despite the gross morphological depression, thickness of the cortical tissue under the grid ($2.8 \pm 0.04mm$) was not statistically different from that of the contralateral cortex (Figure 6.13D; $2.7 \pm 0.09mm$; t-test: $p = 0.19$). Further investigation of changes in cell densities of neurons and microglia in layers I/II–III and V between cortex under the grid versus contralateral cortex showed no significant change in neuronal density between hemispheres, and significant increases only for the microglia of the superficial layers (I/II–III; $p = 0.027$). There were no significant differences in percentage of reactive astrocyte signal for either layer I/II-III or V between groups (summarized in Table 6.2).

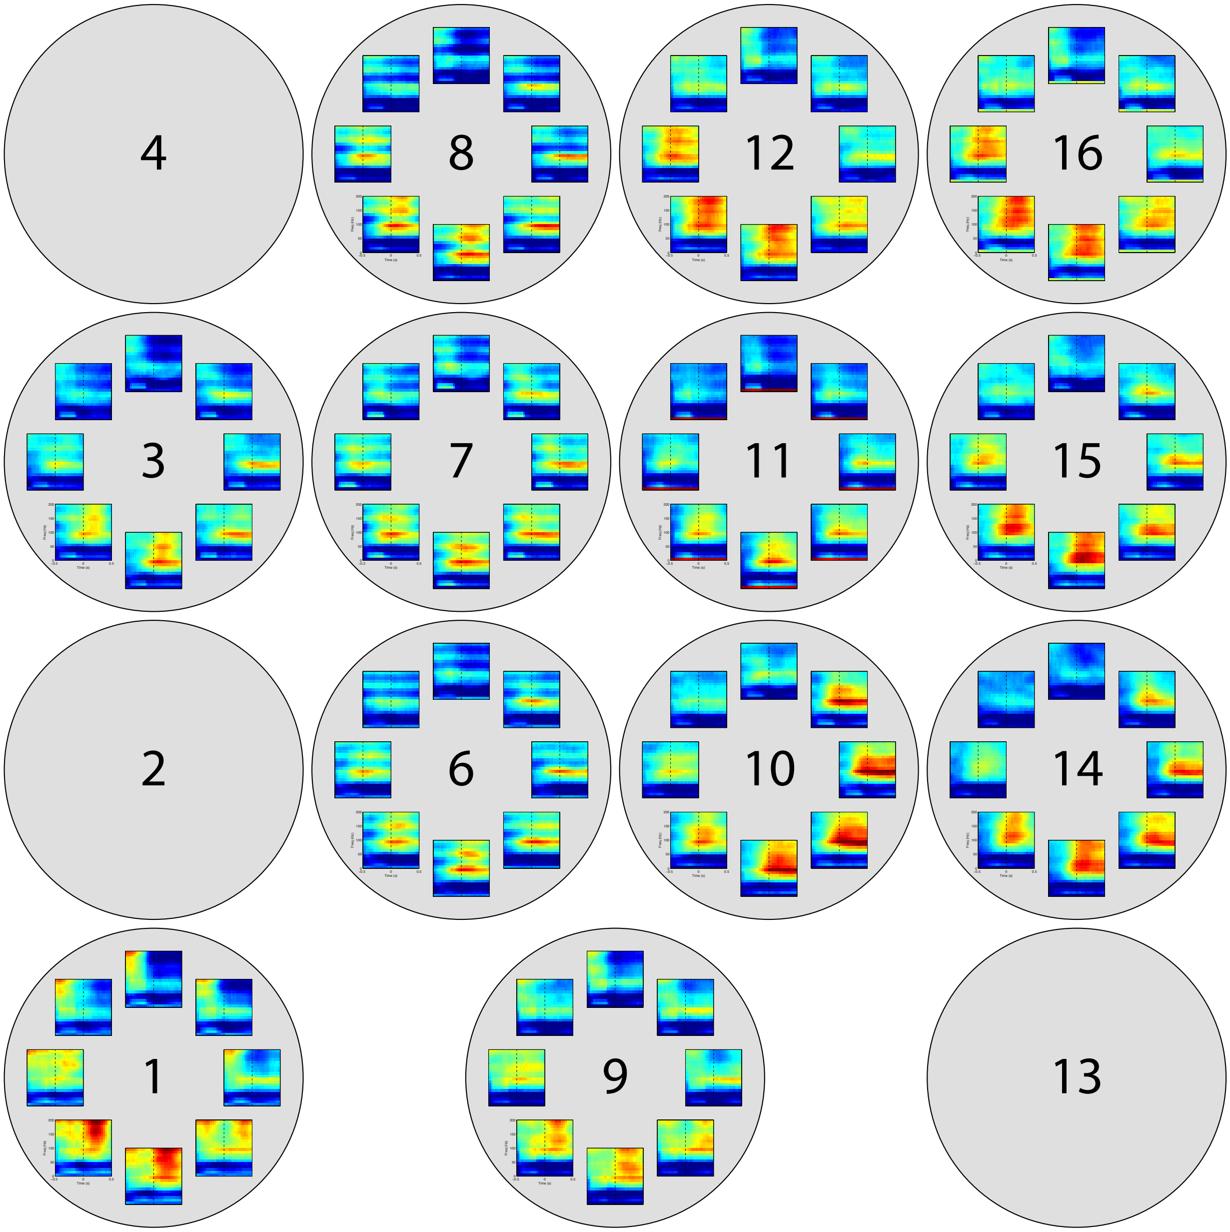

Figure 6.8: *Directional modulation of ECoG signals during brain control.* Data is shown for the peak brain control performance epoch (sessions 32 and 33, $N = 760$). Movement onset is indicated by the dashed black line for each plot. Data for electrodes 4 and 13 are not available on account of these electrodes serving as ground and reference electrodes, respectively. Data for electrode 2 is not available on account of connectivity to this electrode being lost shortly after grid implantation.

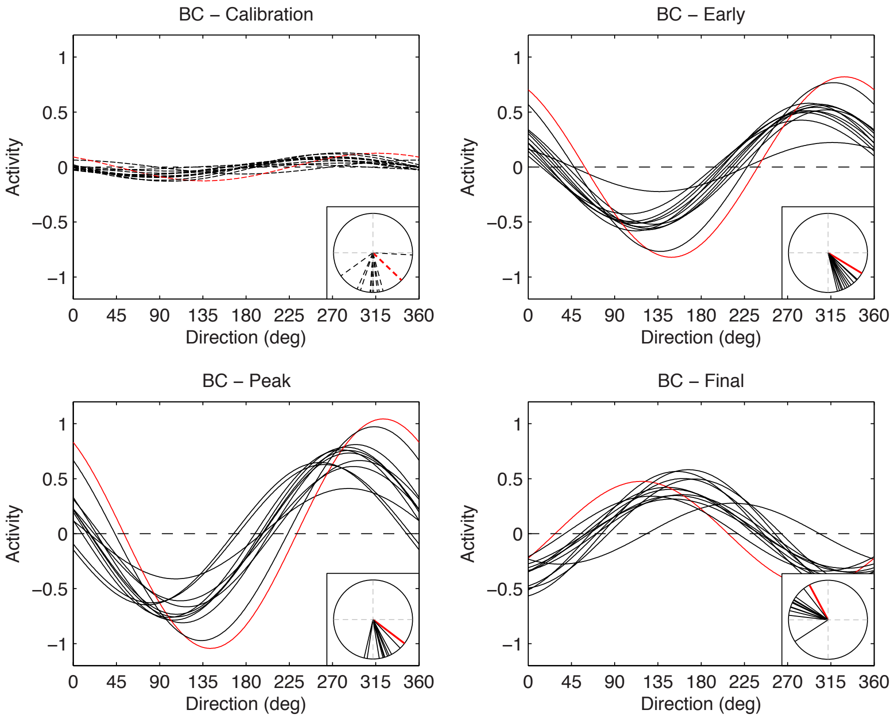

Figure 6.9: *Cosine tuning of ECoG signals during brain control.* Cosine tuning curves are shown for brain control calibration, early, peak, and final epochs (see Figure 6.8). Tuning curves for electrode 10 are shown in red; all other tuning curves are shown in black. Solid lines indicate electrodes exhibiting significant tuning to target direction ($p \leq 0.05$), while dashed lines indicate non-significant fits. *Inset.* Normalized preferred directions for each channel plotted on a unit circle. The mean of each curve (b_0) has been removed prior to plotting.

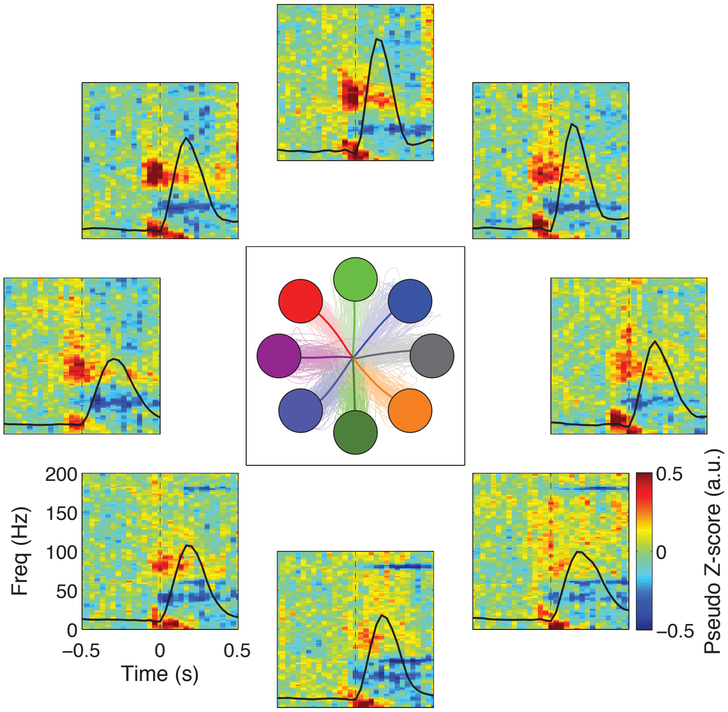

Figure 6.10: *ECoG signal modulation during 8-target center-out hand-controlled cursor tasks.* Averaged time-frequency data are shown for a single electrode (e10) for all tasks. Averaged (thick lines) and individual trajectories (thin lines) for each target are shown in the center panel. Time-frequency data was normalized with respect to the spectral data during a central hold period preceding each trial. Black lines show average speed profiles for each target.

Qualitatively, microglia in both cortex under the array and the contralateral cortex exhibited a ramified morphology [Stence et al., 2001].

6.3.5 Fibrotic encapsulation

Chronic subdural ECoG implantation resulted in fibrotic encapsulation of the grid. The grid was removed by making an incision along the anterior portion of the encapsulation and pulling the grid with forceps. Surprisingly, the grid offered little resistance to removal, indicating that mechanical tethering between the grid and encapsulation tissue was minimal. Using second-harmonic generation (SHG) imaging, we detected collagen I in sections of the tissue encapsulation and control dura mater ($> 2cm$ from implantation site). Using filters to simultaneously image second-harmonic signals as well as tissue autofluorescence, we quantified both the thickness of encapsulation tissue and the percentage of encapsulation tissue that was collagen I-positive (Figure 6.14). Because the distal portion of the encapsulation (dura-facing) was the original,

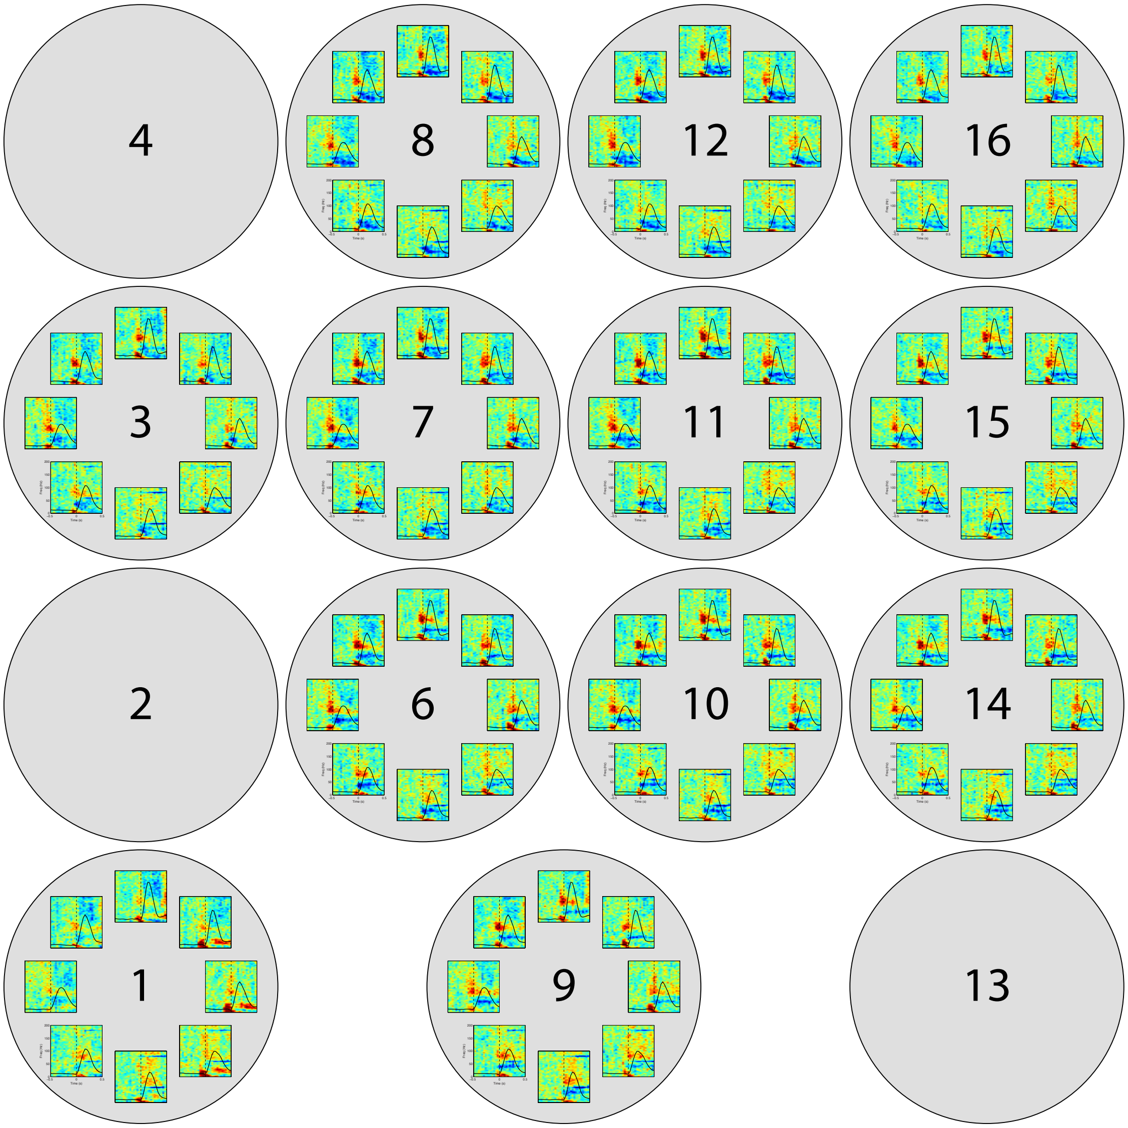

Figure 6.11: *Directional modulation of ECoG signals during hand control.* Data is shown for all successful hand control trials ($N = 1,145$). Movement onset is indicated by the dashed black line for each plot. Data for electrodes 4 and 13 are not available on account of these electrodes serving as ground and reference electrodes, respectively. Data for electrode 2 is not available on account of connectivity to this electrode being lost shortly after grid implantation.

autografted dura mater, we analyzed it separately from the cortex-facing portion of the encapsulation, which grew de novo following initial implantation. Both aspects of the tissue encapsulation were compared to control dura mater taken more than $2cm$ from the implantation site. There were statistically significant differences in the thicknesses of the tissues (one-way ANOVA: $F(2,14) = 136.13$, $p < 0.001$), with both dura-side encapsulation ($0.82 \pm 0.04mm$) and cortex-side encapsulation ($1.76 \pm 0.09mm$) being thicker than control dura mater ($0.36 \pm 0.03mm$; Tukeys post-test: $p < 0.001$). The cortex-side encapsulation was also significantly thicker than dura-side encapsulation ($p < 0.001$).

SHG imaging revealed encapsulation tissue to be comprised of a cellular region that did not express strong SHG signal and a collagenous region that was strongly SHG(+) (schematic: Figure 6.14B,E). Using the tissue thickness derived above, we were able to assess the relative proportions of cellular and collagenous regions by measuring the area of collagenous region (SHG(+) region) and divide it by tissue thickness. This showed that the proportion of collagenous region was significantly different between the tissues (one-way ANOVA: $F(2,14) = 44.33$; $p < 0.001$). Control dura mater had a significantly higher percentage of collagenous

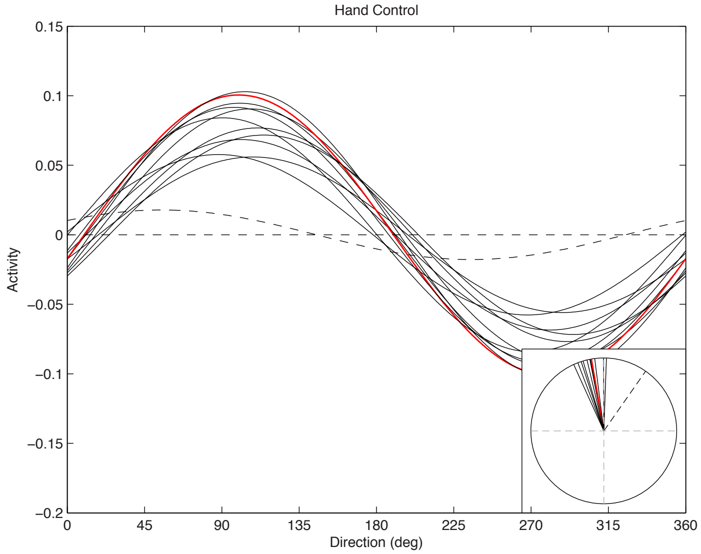

Figure 6.12: *Cosine tuning of ECoG signals during hand control.* The tuning curve for electrode 10 is shown in red; all other tuning curves are shown in black. Solid lines indicate electrodes exhibiting significant tuning to target direction ($p \leq 0.05$), while dashed lines indicate non-significant fits. *Inset.* Normalized preferred directions for each channel plotted on a unit circle. The mean of each curve (b_0) has been removed prior to plotting. Note that spectral data during hand control was normalized with respect to different baseline parameters than brain control, making comparisons of the depth of modulation between this figure and Figure 6.9 invalid.

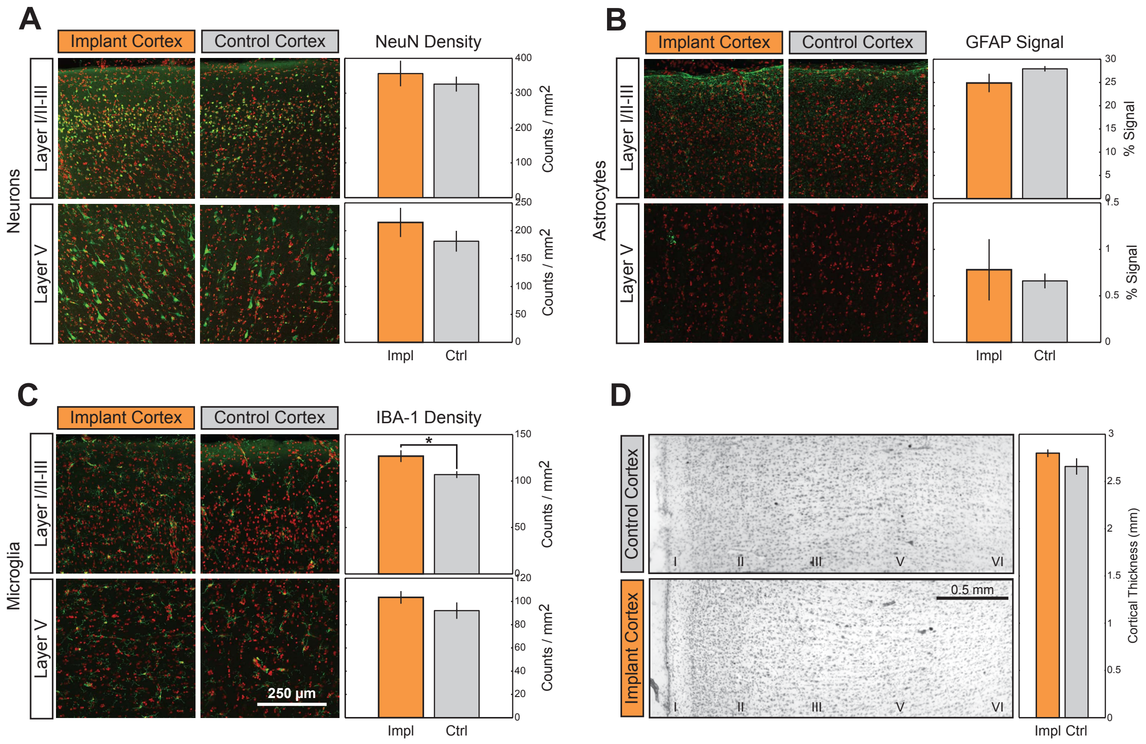

Figure 6.13: *Cortical histology shows minimal damage resulting from long-term ECoG grid implantation.* All sections were counterstained with Hoescht 33342 to label cell nuclei (red). (*A–B*). Neither neurons (A; green) nor astrocytes (B; green) were significantly affected by implantation in either layer I/II–III or layer V. (*C*). Microglial signal was significantly increased in layer I/II–III following implantation, but not in layer V. Data presented as mean ± SEM; '*' denotes significant difference from control ($p < 0.05$). (*D*). Comparison of Nissl-stained motor cortex between implanted and control hemispheres. Cortical layers are indicated by roman numerals I - VI.

tissue ($96.4 \pm 0.33\%$) than either brain-side encapsulation ($82.5 \pm 2.3\%$; Tukeys post-test: $p < 0.01$) or dura-side encapsulation ($58.9 \pm 4.5\%$; $p < 0.001$). The percentage of collagenous tissue in the cortex-facing encapsulation was also significantly greater than that of the dura-facing encapsulation ($p < 0.001$).

In order to determine the cellular composition of the encapsulation we used immunohistochemistry (described in Section 6.2.8). We identified fibroblasts (vimentin(+)/Iba-1(-)) and macrophages (vimentin(+/-)/Iba-1(+)) in all tissue groups. Control dura mater was largely composed of fibroblasts, many of which exhibited elongated nuclei (Figure 6.15C), consistent with previous literature [Adeeb et al., 2012]. Macrophages were sparsely distributed. This closely resembled the collagenous region ($> 300\mu m$ from the array; Figure 6.15B) of the tissue encapsulation, which also contained elongated fibroblasts and macrophages. The "cellular region" of encapsulation ($< 300\mu m$ from the array; Figure 6.15A) was highly cell dense with round,

Table 6.2: Cell density (neurons and microglia, counts/mm^2) and % signal (astrocytes) in cortical layers I/II–III and layer V. Data presented as mean ± SEM, significance ($p < 0.05$) denoted by (*).

Cortical Layer		Implanted Cortex	Control Cortex	p-value
Layer I/II–III	Neurons (counts/mm^2)	356.39 ± 36.76	326.02 ± 21.37	0.5
	Astrocytes (% signal)	24.88 ± 1.99	27.94 ± 0.58	0.18
	Microglia (counts/mm^2)	126.63 ± 6.33	106.93 ± 3.7	0.027*
Layer V	Neurons (counts/mm^2)	214.85 ± 26.19	181.19 ± 18.65	0.32
	Astrocytes (% signal)	0.74 ± 0.33	0.18 ± 0.08	0.73
	Microglia (counts/mm^2)	103.47 ± 5.45	92.08 ± 7.07	0.24

mononuclear macrophages as well as multinuclear, foreign body giant cells (vimentin(+)/Iba-1(+)). Mononuclear/multinuclear distinction was based on nuclei count (Figure 6.15A, inset).

6.4 DISCUSSION

We have shown that a chronically-implanted subdural ECoG electrode grid is capable of recording physiological signals from the cortex for nearly two years, with no electrode failure occurring during the implant lifetime. Furthermore, signals were sufficiently stable over time to allow closed-loop brain control performance with a static decoder to be sustained for a period of 78 days, though it is likely that performance was aided by activity generated as a result of jaw movement artifacts. Upon removal of the electrode grid, it was found that the grid was encapsulated in fibrotic tissue on both the cortex-facing and dura-facing aspects. Histological analysis of explanted cortical tissue showed no damage resulting from the electrode implant, as neuronal and glia density and cortical thickness were found to be comparable to control (contralateral hemisphere) tissue.

6.4.1 Closed-loop brain control

The ability of the monkey to perform closed-loop control of a computer cursor with a fixed decoder appears to provide evidence that ECoG may facilitate stable, long-term control of BMI systems. Previous ECoG studies have shown offline hand reconstruction from subdural ECoG recordings can be sustained for up to approximately 150 days [Chao et al., 2010], with epidural signals yielding stable decoding for up to 10 days [Shimoda et al., 2012]. In comparison, closed-loop BMI control has been demonstrated for up to 19 days using single-unit recordings [Ganguly and Carmena, 2009]. Notably, recent work has shown multi-unit

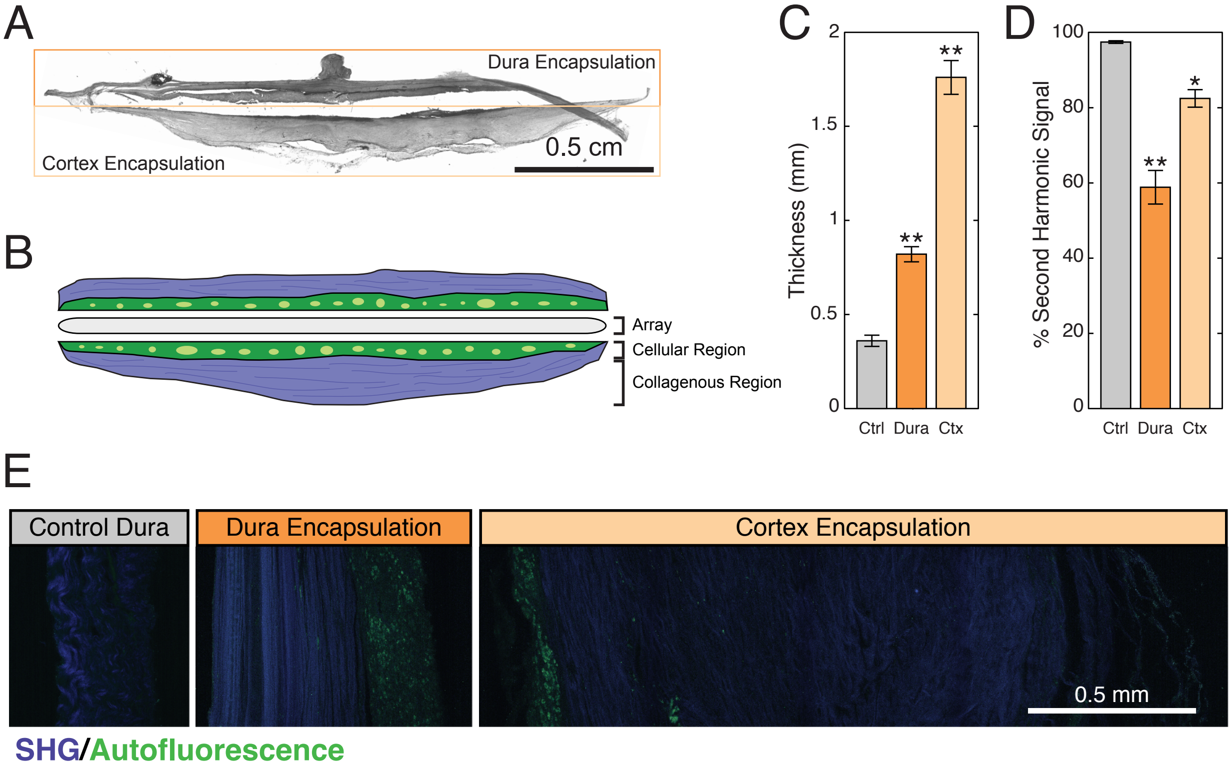

Figure 6.14: *(A)* Full section of encapsulation tissue. *(B)* Schematic of encapsulation tissue composition. Both dura-facing and brain-facing aspects of the encapsulation tissue consisted of a cellular region (green) and a collagenous region (purple). *(C)* Comparison of thickness of dura-facing and brain-facing aspects of encapsulation tissue to control dura. (*D*) The percentage of SHG(+) tissue was significantly reduced in encapsulation tissue. *(E)* Sample images of SHG imaging of control dura mater (*left*), the skull-facing aspect of the encapsulation (*middle*), and the cortex-facing aspect of the encapsulation (*right*). SHG signal is shown in blue while tissue autofluorescence is shown in green. Inner-facing aspects of tissue encapsulation were SHG(-), while distal portions consisted of a thick sheet of fibril SHG(+) elements, likely collagen I. Data in *(C, D)* are presented as mean ± SEM, with asterisks (*) and (**) denoting significant difference from control at $p < 0.01$ and $p < 0.001$, respectively.

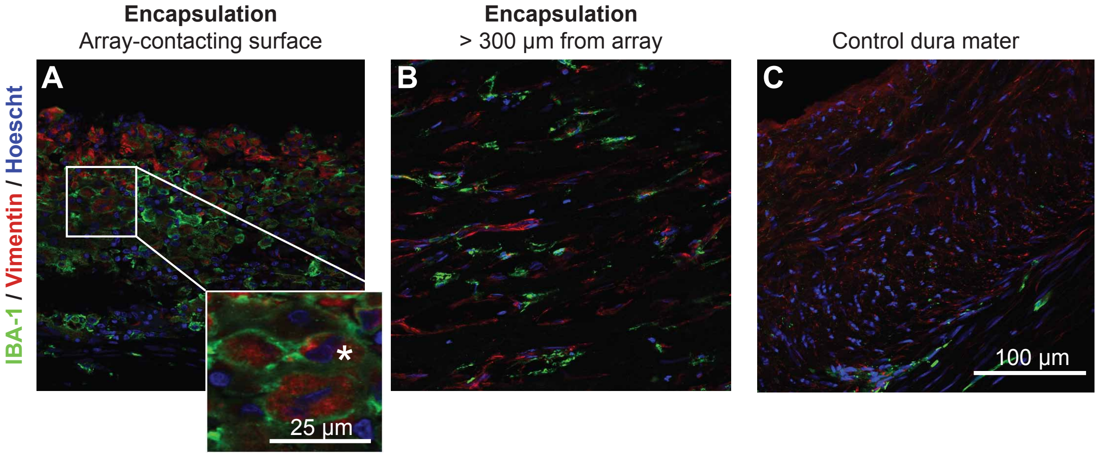

Figure 6.15: *Immunohistochemical staining of encapsulation tissue.* Tissue was stained with Hoescht 33342 (blue; nuclei) and antibodies directed to Iba-1 (macrophages/microglia; green) or vimentin (macrophages/fibroblasts; red). (*A*). Array-contacting aspects of the encapsulation were highly cell dense, populated with vimentin(+)/Iba-1(+) and vimentin(-)/Iba-1(+) macropahges as well as vimentin(+)/Iba-1(-) fibroblasts. *Inset.* Mono and multi-nucleated (*) macrophages. *B.* Distal portions of encapsulation were hallmarked by elongated vimentin(+)/Iba-1(-) fibroblasts and vimentin(-)/Iba-1(+) macrophages. *C.* Control dura mater is largely elongated vimetin(+)/Iba-1(-) fibroblasts with infrequent vimentin(-)/Iba-1(+) macrophages.

and LFP recordings can sustain stable closed-loop BMI performance for up to 203 and 348 days, respectively, using penetrating microelectrode arrays [Flint et al., 2013].

Our results show a marked decline in performance beginning approximately 50 days after the start of fixed-decoder experimental sessions. The reason for such a decline is unclear, though several factors may be responsible. It is tempting to attribute this to surgical interventions (covering of an exposed wire bundle) given the proximity of the performance decline to that of the surgical events. As these surgeries did not involve opening the skull, it is difficult to imagine they could be responsible for the observed changes in performance. Similarly, it is unlikely that motivational issues could be solely responsible for the decline in performance, as the time-frequency data from late fixed-decoder experimental sessions showed a noticeable decline in amplitude from that of peak brain control performance, more indicative of a baseline shift in spectral power.

The encapsulation seems a natural cause for the decline in performance, increasing the distance of the recording electrodes from the cortical source and, as a result, decreasing the overall signal amplitude. However, the time course of the decline does not seem to be consistent with that expected from a typical foreign body reaction, which may plateau at 3 months post-implant (see Section 6.4.3).

It has recently been shown that performance of a skilled task over a period of months results in decreased metabolic activity in primary motor cortex, likely due to a decrease in presynaptic activity [Picard et al., 2013]. In the present study, we observe a decrease in ECoG activity during a brain control task utilizing a fixed decoder over a period where brain control performance appeared to either improved or remained static. It may be the case that increased metabolic efficiency resulted in a decrease in the baseline ECoG spectrum as a result of continued practice, i.e. development of a BMI "skill". Such a decrease, when combined with a fixed neural decoder (including fixed baseline normalization parameters), could be responsible for the decline in performance observed towards the end of the brain control experiment.

Unfortunately, the presence of artifacts in these data make it difficult to attribute the observed changes in the ECoG power spectrum to one particular source. Artifacts were present even during periods of declining brain control performance, suggesting that these artifacts may not have contributed to the observed decrease in spectral amplitude observed over the course of the fixed-decoder brain control experiment. Additional experiments are needed in order to appropriately assess the factors contributing to changes in ECoG spectral amplitude during prolonged BMI experiments.

6.4.2 Influence of artifacts on closed-loop brain control

Recording artifacts detrimentally influenced our closed-loop BMI experiments. Post-hoc examination of neural recordings during closed-loop BMI control revealed the presence of a large number of artifacts. These artifacts have been described previously [Kelly et al., 2013], and we believe they are related to jaw movements based on their timing coinciding with that of the reward period during the brain control task. We found that the number of artifacts increased with improvements in closed-loop performance, and that the presence of artifacts during successful trials was biased towards the lower-left quadrant of the workspace. The presence of such a bias strongly suggests that the animal adopted a strategy of utilizing artifact-related activity to drive movement of the computer cursor.

It is likely that the animal was able to learn to use artifacts to control the BMI as a result of the task timing. The animal was provided with a liquid reward immediately following successful acquisition of the target during brain control. This task sequence may have resulted in anticipation of the liquid reward, and resultantly an incorporation of artifacts into the data used for decoder calibration. Subsequent decoder calibrations may have then unintentionally reinforced the animal's use of artifacts for brain control. As suggested by the increase in the number of artifacts with increasing task performance, the animal likely learned over the course of brain control experiments that these artifacts resulted in stereotyped movement of the cursor and that they could be used to improve performance in the task.

Following cessation of the fixed-decoder brain control experiment and the discovery of the prevalence of artifacts in data collected during this experiment, we performed additional brain control sessions where we attempted to prevent the animal from using jaw movement artifacts during brain control. These efforts involved the incorporation of common-average referencing (CAR) of neural signals [Schalk et al., 2007], improved task sequencing separating the brain control and reward periods of trials, and root-mean-square threshold detection for invalidation of trials containing artifacts. Unfortunately, none of these efforts were successful in preventing the animal from generating artifacts during closed-loop brain control.

Though it has been suggested that ECoG is less sensitive to artifacts than EEG [Schalk and Leuthardt, 2011], we nonetheless observed a high prevalence of artifacts in the neural recordings obtained in the presented study. It may be the case that artifacts in ECoG recordings are more common than previously believed, and may be particularly problematic for clinical BMI systems where control must be robust in the presence of mouth, head, and neck movements. In many BMI experiments, such movements are minimized though subject instruction (in the case of human experiments) or head fixation (for non-human primates); such interventions are impractical for a clinical BMI.

Finally, it is unclear whether the presence of encapsulation tissue surround the electrode grid influence the presence of artifacts. It has also been shown that chewing artifacts may be larger for epidural ECoG recordings than for subdural recordings [Shimoda et al., 2012]. Though the time course of the encapsulation growth in our study is unknown, it is possible that the presence of encapsulation tissue exacerbated the severity of the artifacts observed in this study. It is likely that better electrode grid and amplification hardware designs, innovations in signal filter algorithms such as automated common-mode rejection [Kelly et al., 2013], and careful decoder calibration paradigms will be necessary to limit the influence of artifacts on ECoG-based BMI systems.

6.4.3 Histological findings

During the grid implantation surgery, we resected the dura mater, replaced it over the ECoG array, and sutured it in place. After 666 days of implantation, dura mater/fibrous encapsulation tissue was found in a contiguous piece surrounding the top of the ECoG array where it was placed during surgery as well as under the ECoG array. In order to examine the foreign body response to the chronically-implanted electrode, we examined both the cortical tissue underlying the grid and the fibrotic encapsulation tissue surrounding the grid.

The characteristics of the cerebral cortex underneath the ECoG grid was consistent with that of the tissue of the control contralateral hemisphere. Most importantly, cortical thickness and neuronal density of the tissue under the array were statistically indistinguishable from the contralateral tissue, with no morphological differences apparent at any spatial scale. Measurements of thickness and neuronal density for both hemispheres agreed with those of previous anatomical studies of primate frontal cortex [Matelli et al., 1991, Gittins and Harrison, 2004]. Only the microglial density of the superficial cortical layers was significantly different between the implanted and non-implanted hemispheres. Despite their increased density, these cells were in a “resting” microglia morphology, which indicates that these microglia were not actively responding to trauma or other noxious stimuli [Stence et al., 2001]. It is possible that the persistent, unactivated microglia population is part of the foreign body response to the implanted ECoG array, where increased macrophage density persists in the vicinity of the implant for its lifetime [Anderson, 2001, Sanders et al., 2000]. We observed a concavity of the brain under the electrode array; it is also conceivable that this mechanical deformation incited an independent tissue response that resulted in tissue changes at the locus of mechanical stress [Ding et al., 2008]. The astrocytic GFAP expression between control and implanted hemispheres was not different. Healthy cortical tissue shows low levels of GFAP expression in gray matter, protoplasmic astrocytes, and strong expression at the surface of the brain in the glia limitans. Expression becomes drastically higher in pathological and traumatic situations [Eng et al., 2000]. Qualitatively, our finding of low levels of gray-matter GFAP(+) cells under the array suggests that the array was not actively causing trauma to the cortex.

The cellular distribution in the encapsulation tissue implies that the wound-healing response to implantation consisted of a stereotypic foreign body response, which involved aggregation of mononuclear macrophages and multinucleated foreign body giant cells to the implant site and encapsulation of the device in a collagenous envelope. Aggregated cells and tissue encapsulation generally persist through the lifetime of the implant, with pro-inflammatory cytokine expression diminishing within the first month as anti-inflammatory/pro-wound healing cytokines are expressed [Lynn et al., 2011, Brodbeck et al., 2003, Anderson et al., 2008]. The fibrous encapsulation demarcates the final stage of wound healing in which the tissue disrupted by implantation is either regenerated from cells of the original cell type, or replaced with fibrous connective tissue. Given that dura mater is already largely fibrous connective tissue and mesenchymally derived fibroblasts, it was unclear to us the extent to which the tissue encapsulation was fibrous encapsulation or regrown dura mater [Adeeb et al., 2012, Anderson, 2001, Anderson et al., 2008]. We observed a gradient where tissue proximal to the implant more closely resembled fibrous encapsulation, and tissue distal to the implant more closely resembled control dura mater.

We found that both the dura-facing encapsulation and cortex-facing encapsulation were thicker than the control dura mater, which would be expected of a foreign body tissue encapsulation. Dural thickening may be an inevitable consequence of craniotomy and/or durotomy that is simply exacerbated by the presence of a foreign body. Previously, it has been shown that merely performing a craniotomy in New Zealand white rabbits resulted in a 3.8 fold increase in dural thickness at 3 weeks, with a reduction to a 2.6 fold increase at 3 months [Nunamaker and Kipke, 2010]. Replacing dura with an alginate hydrogel resulted in a 2.8 fold increase of dural thickness at 3 weeks and a 3.1 fold increase at 3 months. Dural thickening of 2mm 8 weeks after a 2cm dural incision has been observed in coonhound dogs; following application of a poly(ethylene) glycol based dural sealant, the healed dura was found to have thickened as much as 4mm [Preul et al., 2003]. Furthermore, it has been found that mengineal cells almost double collagen production following injury [Sajanti et al., 1999], with computational models of collagen I fibrosis following biomaterial implantation corroborating this finding [Su et al., 2011]. In these experiment-validated models, increasing numbers of fibroblasts at the implant site results in significantly increased collagen deposition. Since the predominating cell type of dura mater is the fibroblast, it is plausible that we might expect pronounced collagen I production following implantation.

While we have shown that the foreign body response to chronic ECoG grid implantation can result in grid encapsulation, the time course for this response is still unknown. ECoG electrodes implanted up to 30 days clinically for epilepsy monitoring do not exhibit such encapsulation [Fountas and Smith, 2007, Van Gompel et al., 2008, Wong et al., 2009, Wang et al., 2013a], while it has recently shown that encapsulation for epidural ECoG grids can occur as early as one month post implantation [Schendel et al., 2013, Schendel et al., 2014]. Mild chronic inflammation, restricted to the leptomeninges, has been observed after 25 weeks (6 months) in mirco-ECoG electrodes implanted subdurally in a rat model [Henle et al., 2011] while other dural manipulations resulting in fibrosis may plateau within 3 months [Nunamaker and Kipke, 2010, Preul et al., 2003, Su et al., 2011].

6.4.4 Study limitations

Unfortunately, the nature of the closed-loop brain-control experiments we conducted prevented cortical responses from being observed under constant experimental conditions. Evaluation of ECoG activity during a stereotypical motor task (e.g. a center-out reaching task) at regular intervals during periods of closed-loop brain control would likely have shed light on the cause of declining performance. However, such experiments were not able to be performed for practical reasons. Ultimately, the decline in closed-loop brain control performance with a static decoder may be the result of a combination of factors, including the presence of the encapsulation, declining motivation, and perhaps even cortical changes resulting from BMI skill learning.

Furthermore, a more thorough analysis of signal quality over time is difficult given the nature of the brain-control experiment. Changes in ECoG signal characteristics associated with BMI learning are likely occurring concurrently with those changes associated with electrode encapsulation, making dissociation of these two effects problematic. The time course of such changes will need to be studied in detail, as these will likely affect the stability of BMI control with static decoding parameters. Nevertheless, the fact that chronic implantation of the ECoG electrode grid did not result in cortical damage provides evidence for the utility of ECoG as a recording modality for a BMI system. This, combined with the ease with which the electrode grid could be removed from the encapsulating tissue, suggest that a chronic ECoG BMI system could in theory be upgraded if a more promising electrode technology becomes available.

The effect of encapsulation tissue on ECoG signal independence also warrants further investigation. While it has been shown that ECoG can be used to obtain three-dimensional BMI control [Wang et al., 2013a], it is unknown if ECoG signals are sufficiently independent to enable higher degree-of-freedom (DoF) control. It has previously been shown through finite element modeling that human epidural electrical field potentials exhibit lower spatial resolution than those recorded subdurally [Slutzky et al., 2010], suggesting that the presence of encapsulation is likely to decrease the spatial resolution of ECoG recordings due to an increased distance of the recording electrodes from the cortex. Considering this, any encapsulation-related decreases in electrode independence could decrease the utility of an ECoG BMI system for higher degree-of-freedom control.

6.4.5 Implications and future directions

Given the current interest in both epidural and subdural ECoG recordings for BMI and other clinical applications, it is worth considering the implications of our findings on the utility of subdurally-implanted ECoG grids. Epidural implantation of ECoG grids is argued to result in a lower risk of infection than subdural ECoG, while an acute study in humans showed subdural implantation provides higher amplitude neural signals [Bundy et al., 2014]. The current work suggests that long-term subdural implantation could result in a tissue encapsulation, though it is unclear how this encapsulation may affect recording quality. Furthermore, encapsulation has also been observed in epidural implantations [Schendel et al., 2013, Schendel et al., 2014]. Because these studies were conducted in different animal models with different ECoG arrays, direct comparison is not possible. A long-term study to evaluate the effect of implantation location on tissue response and recording quality is necessary to fully understand any potential tradeoffs between locations. Regardless, the fact that a similar reaction type is observed despite implantation location suggests that cutting the dura is not solely responsible for the response observed in this study. Further, the fact that profound dura thickening, increased collagen deposition, and microglial and astrocytic activation is also observed following a craniotomy without interruption of the dura mater [Nunamaker and Kipke, 2010, Sajanti et al., 1999, Xu et al., 2007] suggests the extent of tissue encapsulation may be exacerbated by other factors beyond cutting the dura. Subdural and epidural implants also have the potential to mechanically irritate either dura or neural tissue based on the relative micromotion of the brain/dura and implant, which could exacerbate fibrosis [Schendel et al., 2013, Gilletti and Muthuswamy, 2006].

Many strategies can and have been pursued to minimize the foreign body response to ECoG electrodes. To mitigate dural thickening from craniotomy and durotomy, minimally invasive ECoG insertion techniques can be used, such as the expandable shape memory alloy guidewires [Yamakawa et al., 2010]. It is postulated that micromotion of the ECoG array at the surface of the brain/dura may irritate surrounding tissue, causing inflammation and fibrosis [Schendel et al., 2013]. Using flexible materials such as polyimide as an array substrate may mitigate this inflammation [Yeager et al., 2008, Rubehn et al., 2009]. Additionally, dissolvable silk fibroin shuttles have been used to introduce ultracompliant, thin polyimide mesh ECoG arrays to cortex. Such arrays can conform to the brains topography upon contact which could reduce micromotion, and have been shown to mitigate some of the tissue proliferation under the ECoG grid when compared to traditional grids [Kim et al., 2010, Schendel et al., 2013, Schendel et al., 2014]. Many biomaterials applications achieve reduced biofouling with hydrophilic polymer coating that inhibit serum protein binding [Brodbeck et al., 2003, Collier et al., 2004]. Additionally, biomolecule surface immobilization has been used to “camouflage” implants against the bodys defense mechanisms [Azemi et al., 2011, Kolarcik et al., 2012]. Further, controlled release of anti-inflammatory or other therapeutics can be achieved by trapping the drug in either degradable or conductive polymer surface modification. In the former case, drug is released as hydrolytic or enzymatic degradation of the polymer exposes encapsulated drug [Norton et al., 2005, Rujitanaroj et al., 2013, Bridges and Garcia, 2008, Wang et al., 2013b, Wadhwa et al., 2012]. In the latter case, electrical stimulation of the conductive polymer electrostatically releases ensnared drug, allowing it to diffuse out of the coating and enabling on-demand release [Svirskis et al., 2010, Weaver et al., 2014, Luo et al., 2011]. Anti-mitotic drugs applied directly to a chronic durotomy site have been shown to safely prevent dura regrowth and could be a viable candidate for controlled release at the ECoG-brain interface [Spinks et al., 2003]. These strategies should be assessed in the context of long-term ECoG implantation and neural recording to determine the optimal strategy for more biocompatible arrays.

Clearly, brain-surface dwelling implants present different challenges from cortex-penetrating neural electrodes where there is little fibrous encapsulation, but slow signal decay likely due to neuronal degeneration [Barrese et al., 2013, Freire et al., 2011, Collinger et al., 2012, McConnell et al., 2009, Saxena et al., 2013]. While it is predicted that intracortical devices will last 8 years [Barrese et al., 2013], the lifespans of surface dwelling grids are unknown. To our knowledge, our 666-day implant is the longest reported to date. Our results suggest that these devices do not affect neuronal density even after nearly two years of implantation. Thus, if ECoG arrays do fail over time, we do not expect it to be from the same mode as intracortical electrodes. Rather, our results suggest that the only likely failure mode would be signal loss due to excessive fibrous encapsulation. However, even after 666 days and 1.7mm of fibrous encapsulation, our ECoG grid showed robust movement-related modulation during hand-control tasks. Because we have only explored one time point, it is unclear whether the encapsulation is stagnant, growing, or shrinking. Other studies of sub-dural implants found that fibrous encapsulation plateaued by 3 months post-implantation [Nunamaker and Kipke, 2010, Preul et al., 2003, Su et al., 2011]. Further experiments with more time points must be conducted to understand the progression of fibrous encapsulation in the sub-dural ECoG context. Furthermore, the effect of fibrous encapsulation on electrode independence must also be evaluated. Though we observe significant modulation of ECoG activity during overt hand movements, nearly all electrodes exhibit similar preferred directions, indicative of a high degree of spatial correlation between electrodes. What effect fibrotic encapsulation has on this correlation is unclear, and requires further investigation.

We believe our results also have implications for the viability of ECoG for long-term high-resolution brain recording. ECoG has increasingly become a neuroimaging method of choice in a variety of neuroscience fields, including the study of speech and language processes [Bouchard et al., 2013, Pasley et al., 2012, Kellis et al., 2010, Wang et al., 2011a, Cervenka et al., 2011, Pei et al., 2011]. The fact that subdural ECoG grids can remain implanted for extended periods of time without damaging the cortex will facilitate the study of cognitive processes over long timescales. The ease with which ECoG electrode grids can be removed from encapsulation tissue may provide for repeated implants, thereby enabling long-term studies using ECoG in appropriate patient populations.

Our results may also inform other fields that use implantable brain surface electrodes for non-BMI neural recording and neuromodulation applications, such as pain management and the treatment of epilepsy and mental disorders. For example, work is currently in progress investigating the efficacy of chronically-implanted depth and subdural electrodes for purposes of epilepsy management through responsive cortical stimulation [Morrell and RNS System in Epilepsy Study Group, 2011]. It may be the case that stimulation parameters would need to be updated over time in order to account for the effects of encapsulation. Long-term studies of epilepsy or other neurological disorders, specifically those tracking disease progression or recovery, may also be enabled through the use of chronically-implanted of ECoG electrodes.

7.0 SPECTRALLY AND SPATIALLY-CONSTRAINED DECODING OF ELECTRICAL FIELD POTENTIALS USING EMPIRICAL BAYES

We have previously discussed both the efficacy of ECoG BMI systems (Chapter 4) and the chronically-implanted ECoG BMI systems (Chapter 6). While these results provide an essential evidence for the feasibility of a clinical BMI system, we ultimately seek the realization of a high-degree-of-freedom, high-performance BMI. Integral to this is maximizing the extraction of information from the cortex with the *neural decoder*, the component of the brain-machine interface which is responsible for the translation of cortical activity into BMI command signals. In this chapter, we present a novel decoding algorithm which uses empirical Bayes in conjunction with weight-space covariance priors to find spectrally and spatially-constrained decoding weight solutions. Using both simulated data and ECoG recordings from individuals with upper-limb paralysis, we show that by imposing prior distributions over decoding weights which encourage spectral and spatial correlation, decoding performance can be achieved which surpasses that of standard least-squares linear regression, and in some cases, Optimal Linear Estimation. Additionally, we find that the benefits of this method are greatest in data-limited regimes, making it particularly appropriate for ECoG BMI systems.

7.1 BACKGROUND

ECoG signals are commonly recorded from between 16 and 128 electrodes simultaneously and are analyzed in the time-frequency domain. As a result, the number of potential features (i.e., spectral power in a particular frequency band for a single electrode) available for use in analysis can easily exceed several thousand. For example, 128 channels of data analyzed with 5 Hz frequency resolution results in over 5,000 potential features. With growing interest in high-density and micro ECoG grids [Degenhart et al., 2011a, Rouse and Moran, 2009, Wang et al., 2009], local field potentials [Zhuang et al., 2010], and expansion of the frequency range of interest [Gaona et al., 2011], this number is only expected to increase. Such large feature set sizes can be problematic for two reasons. First, large feature sets can lead to *ad-hoc* feature selection methods which may fail to identify important features. Second, ECoG features are often highly correlated temporally, spectrally, and spatially due to the biophysics of field potentials and choices in signal preprocessing. This can result in large numbers of highly dependent features, problematic for certain decoding methods (e.g. Naïve Bayes classification) as well as for making inferences about the structure of ECoG signals, such as the functional specificity of broadband spectral modulation. Finally, the high-dimensional space is typically undersampled, as the number of features (the number of channels times the number of frequency bands) often exceeds the number of data points. If not properly handled, this can easily lead to overfitting of neural decoder parameters and consequently a reduction in performance for BMI systems.

Several solutions are available for handling the problem of overfitting, including dimensionality reduction, feature selection, or sparse decoding methods. *Ad hoc* feature selection has been widely employed in ECoG literature as a means of controlling the size of feature sets. Frequently this takes the form of averaging the frequency spectra for individual channels across larger bands [Kubánek et al., 2009], or the use of movement screening paradigms to restrict the feature space to task-modulated channels and frequency bands [Schalk et al., 2008c]. Other, more principled feature selection approaches, such as "greedy" (forward) feature selection [Zhuang et al., 2010] have been applied to the decoding of field potential signals. Principal components analysis (PCA) has been applied to extract low-dimensional sets of features from ECoG data, enabling the prediction of individual finger movements [Miller et al., 2009b, Wang et al., 2009], as well as classification of hand posture and force levels [Degenhart et al., 2011a]. We have recently shown that sparse decoding methods, such as L_1 and elastic net regularized linear regression, can be used offline to decode BMI cursor control trajectories and overt hand posture from large (>1000) ECoG feature sets with greater accuracy than non-regularized methods, with elastic net regularization allowing a sparse set of decoding parameters to be found which retains correlations between important features [Kelly et al., 2012].

The preservation of the influence of inter-feature correlations, particularly for frequency spectra calculated with fine spectral resolution, is likely to be important for the realization of high-performance ECoG-based BMI systems. It has been hypothesized that the ECoG frequency spectrum modulation occurs in a broadband manner for frequencies between 80 and 500 Hz [Miller et al., 2009a], though other evidence suggests that functional modulation of ECoG signals may exist within sub-bands of the high gamma (>60 Hz) range [Gaona et al., 2011]. Furthermore, it has been shown that non-human primates are capable of de-correlating the spectral modulation of neighboring micro-ECoG electrodes in narrow frequency bands through closed-loop BMI training [Rouse and Moran, 2009]. The use of such *a priori* information may prove beneficial within the context of ECoG BMI decoding, allowing for an increase in the signal-to-noise ratio of BMI control commands through the identification and inclusion of functionally-distinct patterns of spectral and spatial modulation into neural decoding schemes.

The incorporation of prior knowledge into decoding algorithms has been utilized with multiple neural recording modalities, including magnetoencephalography (MEG), functional magnetic imaging (fMRI), ECoG, and LFP. Prior information about the constraints governing finger movements has been used to improve the accuracy of finger movement prediction from ECoG signals, though in this case *ad hoc* feature selection was still used [Wang et al., 2011b]. In the field of neuroimaging, decoding of joystick movements using MEG has been shown to be improved by the incorporation of a prior constraint on the cortical region of interest over which movement-related information is expected to be recorded [Zhang et al., 2011]. A similar approach has been shown to aid in the prediction of brain state from BOLD responses, where prior spatial information about the cortical regions expected to be activated by a given set of stimuli has been used to enhance prediction [Chu et al., 2011]. Finally, the encoding of visual stimuli by LFP recordings has been uncovered using implicit assumptions of the functional similarity of neighboring frequency bands to divide a frequency range of interest into sub-bands which maximize information content [Magri et al., 2011]. This suggests that a similar treatment of *a priori* information about the spectral and spatial correlation structure of field potentials may aid in the decoding of information from these signals.

7.2 MOTIVATION AND APPROACH

We motivate the proposed approach by first identifying the key characteristics of ECoG time-frequency responses to naturalistic stimuli. An example of ECoG modulation during an attempted movement task is shown in Figure 7.1. There are several notable characteristics present in these data. First, time-frequency responses for individual electrodes exhibit a substantial degree of correlation across frequency, with neighboring frequencies having similar activity. Secondly, ECoG responses are highly correlated spatially, with neighboring electrodes showing similar responses.

Further insight into the nature of ECoG features can be gained by examining the correlation structure of recordings directly. Figure 7.2 shows covariance matrices for ECoG brain-control subjects S1 and S2[1]. As seen in these matrices, the spectral modulation of neighboring frequency bands on individual electrodes tend to co-vary with one another, as indicated by the high degree of covariance along the diagonal of Figures 7.2C and 7.2F. In addition, neighboring electrodes also exhibit a substantial degree of covariance, as indicated by the high covariance of the off-diagonal elements in Figures 7.2B and 7.2E. For example the covariance of neighboring electrodes 7 and 15 is high quite high for subject S1 (Figure 7.2B), as expected by the close proximity of these electrodes to one another on the electrode grid (Figure 7.2A).

These signal characteristics allow us to define our decoding approach. Given the high degree of correlation of spectrally and spatially-similar ECoG features, we make the assumption that the decoding weights should exhibit the same correlation structure as the neural features. Such structure can be imposed on the set of decoding weights by using *a priori* information of the correlation of ECoG features to define a prior distribution over the set of decoding weights. As will be shown in the following section, this approach can be formalized as a special case of Bayesian linear regression.

---

[1] A thorough description of these subjects can be found in Chapter 4.

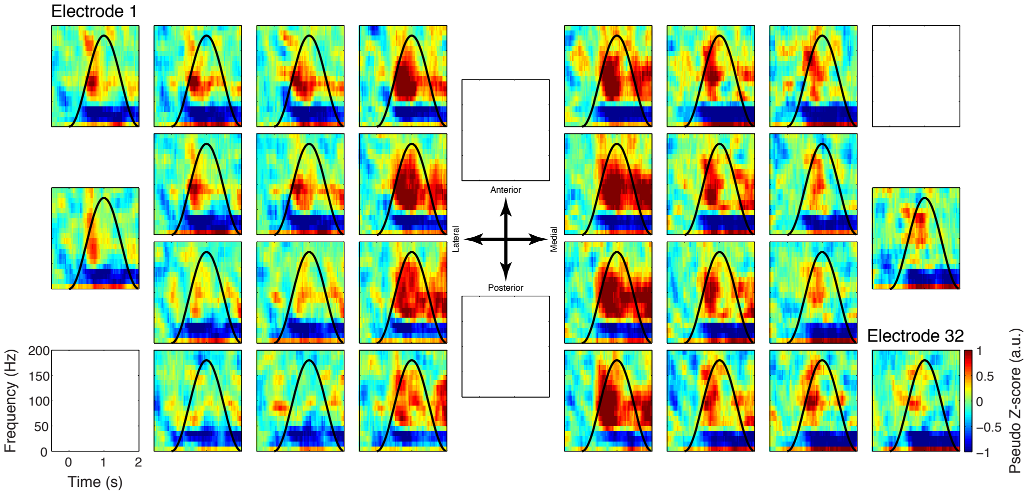

Figure 7.1: *ECoG time-frequency responses during attempted hand flexion.* Time-frequency data are shown for an individual with upper-limb paralysis implanted with a 32-electrode ECoG grid. Individual panels show the time-frequency responses for individual electrodes, with the instructed attempted hand aperature indicated by the *black* line. Refer to Chapter 4, Figure 4.1 for grid layout and placement details.

7.3 METHODS

7.3.1 Bayesian linear regression

Incorporation of prior knowledge of the correlation structure of ECoG recordings into the decoding method begins by first defining the functional form of the model used. We first assume that kinematic variables are a linear function of neural activity, as defined the standard linear regression model shown by Equation 7.1:

$$y = \mathbf{w}^\top \mathbf{x} + \epsilon \quad , \quad \epsilon \sim \mathcal{N}(0, \sigma_n^2), \tag{7.1}$$

where y is a single kinematic variable (e.g. X-axis velocity), $\mathbf{x} \in \mathbb{R}^{D\times 1}$ is a vector of instantaneous neural activity, $\mathbf{w} \in \mathbb{R}^{D\times 1}$ is the vector of decoding weights relating neural activity to kinematics, and ϵ is zero-mean additive observation noise with variance σ_n^2. Such models have been used successfully in BMI systems (see Chapter 2), including the closed-loop control results presented in Chapter 4.

Based on the observation that the neural feature vector $\mathbf{x}$ exhibits a characteristic pattern of covariance, we make the assumption that the elements of the decoding weights vector $\mathbf{w}$ share this correlation structure. This is accomplished by setting a prior distribution over $\mathbf{w}$. By specifying the form of this prior such that it matches our *a priori* knowledge about the expected correlation structure of the neural feature set, we can attempt to find a solution to Equation 7.1 which reflects this structure.

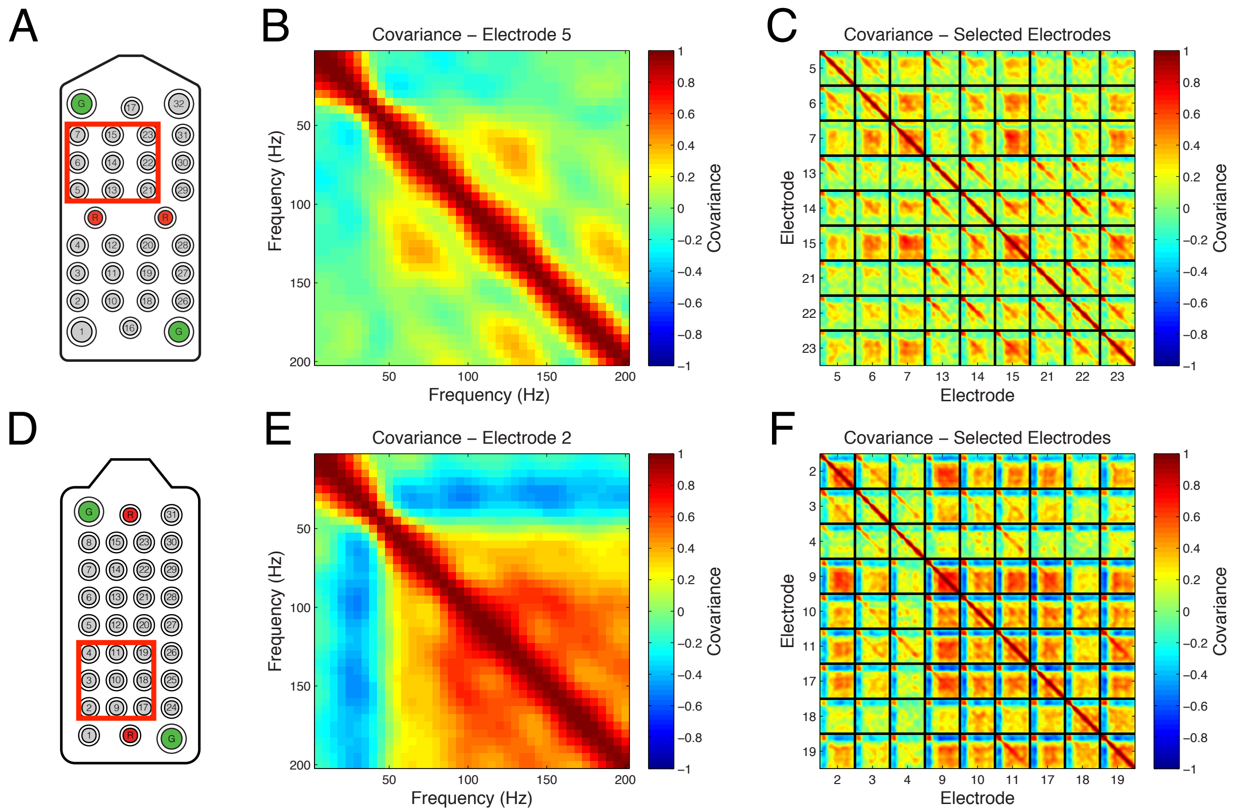

Figure 7.2: *Example of covariance structure of human ECoG recordings.* Covariance matrices are shown for selected electrode subsets for human ECoG subjects S1 (*top* row) and S2 (*bottom* row). *(A)* Electrode schematic for Subject S1. The subset of electrodes for which the covariance matrix is provided is shown by the *red* box. *(B)* Detail of covariance for a single electrode (e5) for Subject S1. *(C)* Covariance matrix for selected electrodes for Subject S1. Divisions between electrodes are indicated by *black* lines. *(D)* Electrode schematic for Subject S2. *(E)* Detail of covariance for a single electrode (e2) for Subject S2. *(F)*. Covariance matrix for selected electrodes for Subject S2.

We define the prior distribution over $\mathbf{w}$ is defined as a zero-mean Gaussian distribution with covariance $\mathbf{\Sigma}_p$ as shown by Equation 7.2:

$$p(\mathbf{w}) \sim \mathcal{N}(\mathbf{0}, \mathbf{\Sigma}_p). \tag{7.2}$$

Under this prior, the decoding weights are calculated according to Equation 7.3: [Rasmussen and Williams, 2006]:

$$\mathbf{w} = \sigma_n^{-2}\mathbf{A}^{-1}\mathbf{X}\mathbf{y}, \tag{7.3}$$

where $\mathbf{A} = \sigma_n^{-2}\mathbf{X}\mathbf{X}^\top + \mathbf{\Sigma}_p^{-1}$, $\mathbf{X} = [\mathbf{x}_1, \ldots, \mathbf{x}_N] \in \mathbb{R}^{D\times N}$ is an matrix of neural feature observations, and $\mathbf{y} = [y_1, \ldots, y_N]^\top$ is a vector of kinematic observations. A detailed derivation of these equations is provided in Appendix C.

7.3.2 Parametric prior covariance matrices

The functional form of the covariance prior over the decoding weights, $\mathbf{\Sigma}_p$, specifies the desired correlation structure of the decoding weights vector $\mathbf{w}$. We specify the general form of $\mathbf{\Sigma}_p$ to be:

$$\mathbf{\Sigma}_p = \sigma_{w,f}^2\mathbf{K} + \sigma_{w,n}^2\mathbf{I}, \tag{7.4}$$

where $\sigma_{w,f}^2$ and $\sigma_{w,n}^2$ are the functional and noise variance of the prior over $\mathbf{w}$, respectively, and $\mathbf{K}$ is a matrix which specifies the functional form of the prior. The diagonal term $\sigma_{w,n}^2\mathbf{I}$ ensures that $\mathbf{\Sigma}_p$ is full-rank and thus invertible. Parametric covariance matrices are defined by specifying the elements of $\mathbf{K}$ according to a particular *kernel* with the general form:

$$\mathbf{K}(x_i, x_j) = g(x_i, x_j, \theta), \tag{7.5}$$

which indicates that the kernel specifying the covariance prior over the decoding weights is a function of features x_i and x_j as well as the set of *hyperparameters* of the prior θ.

We investigate three specific covariance kernels. The first of these, referred to here as the *diagonal* kernel, is the simplest of the three and is defined as:

$$\mathbf{K}_d(x_i, x_j) = \delta(x_i, x_j), \tag{7.6}$$

where $\delta(x_i, x_j) = 1$ if $i = j$ and 0 otherwise. As this implies that $\mathbf{\Sigma}_p$ is proportional to the identify matrix, use of the diagonal prior makes the assumption that the elements of $\mathbf{w}$ do not co-vary and are thus independent from one another. In this case, the probabilistic linear regression model is equivalent to L_2-regularized linear regression ("ridge regression"), and prevents overfitting by penalizing solutions for $\mathbf{w}$ by the L^2 norm $|\mathbf{x}|$ [Park and Pillow, 2011].

The two other covariance kernels, the *frequency-block-diagonal* and *frequency-spatial* kernels, attempt to incorporate the known correlation structure of ECoG recordings into $\mathbf{\Sigma}_p$. The *frequency-block-diagonal* kernel specifies correlations between features for a single electrode as a function of the distance in frequency between the features and assumes independence between electrodes:

$$\mathbf{K}_f(x_i, x_j) = \exp\left\{-\frac{1}{2\ell_f^2}\left[f(x_i) - f(x_j)\right]^2\right\}\delta_{el}(x_i, x_j) \tag{7.7}$$

where $f(x_i)$ is the frequency band corresponding to feature x_i, hyperparameter ℓ_f^2 is the "spectral length constant" determining the correlation between elements of $\mathbf{w}$ as a function of the difference in frequency between their corresponding features, and delta function $\delta_{el}(x_i, x_j)$ is 1 if features x_i and x_j are on the same electrode, and 0 otherwise. The *frequency-spatial* kernel builds upon the *frequency block-diagonal* kernel by incorporating covariance as a function of the spatial distance between electrodes:

$$\mathbf{K}_{f,s}(x_i, x_j) = \mathbf{K}_f(x_i, x_j)\mathbf{K}_s(x_i, x_j) = \exp\left\{-\frac{\left[f(x_i) - f(x_j)\right]^2}{2\ell_f^2}\right\}\exp\left\{-\frac{d^2(x_i, x_j)}{2\ell_s^2}\right\}, \tag{7.8}$$

where $d^2(x_i, x_j)$ is the square of the spatial distance between electrodes for features x_i and x_j and ℓ_s is the "spectral length constant" hyperparameter determining the correlation between elements of $\mathbf{w}$ as a function of the distance between the electrodes for their corresponding features.

Figure 7.3 shows example prior covariance matrices for the diagonal, frequency-block-diagonal, and frequency-spatial kernels for a simulated 3 × 3 electrode grid with inter-electrode distance of 3 mm. As the length constant hyperparameters (ℓ_f^2 and ℓ_s^2) increase, the degree of inter-feature correlation incorporated into the prior increases. Furthermore, as $\ell_s^2 \to 0$, the frequency-spatial kernel becomes equivalent to the frequency block-diagonal kernel, and as $\ell_f^2 \to 0$, the frequency-block-diagonal kernel becomes equivalent to the diagonal kernel. Thus, the frequency-spatial kernel can be considered the most general case of the three, with the diagonal and frequency block-diagonal special cases of the frequency-spatial kernel for specific values of ℓ_f^2 and ℓ_s^2. For the sake of brevity, the prior covariance matrices for the diagonal, frequency block-diagonal, and frequency-spatial kernels will be referred to as the diagonal, frequency block-diagonal, and frequency-spatial *priors* for the remainder of this chapter, with the decoders for these priors abbreviated by EB_d, EB_f, and EB_{fs}, respectively.

7.3.3 Evidence-based maximization of parameters and hyperparameters

The previous sections have shown how decoding weights can be calculated provided training data and the set of hyperparameters governing the prior over the covariance matrix $\mathbf{\Sigma}_p$. While it is possible to specify the values of the hyperparameters, we choose to learn these values directly from the training data. This is accomplished by maximizing the likelihood of the kinematic observations given the neural observations, $p(\mathbf{y}|\mathbf{X})$, which can be expressed as:

$$p(\mathbf{y}|\mathbf{X}) = \int p(\mathbf{y}|\mathbf{X}, \mathbf{w})p(\mathbf{w})\ \mathrm{d}\mathbf{w}. \tag{7.9}$$

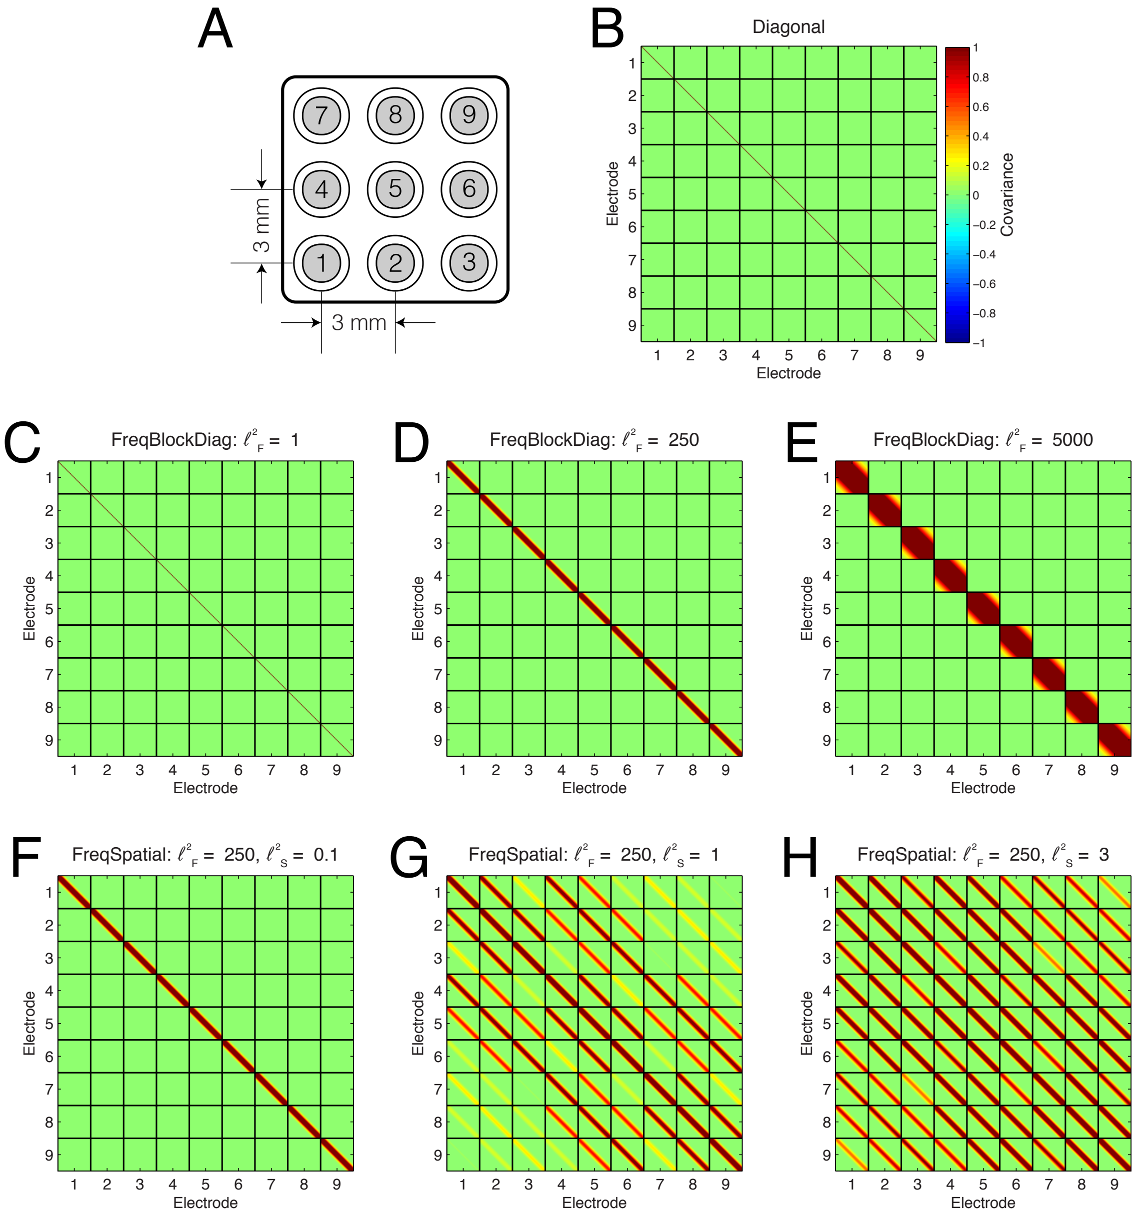

Figure 7.3: *Example of covariance priors.* Example covariance matrix priors ($\mathbf{\Sigma}_p$) are shown for three classes of covariance priors. 'Diagonal': diagonal prior. 'FreqBlockDiag': frequency-block-diagonal prior. 'FreqSpatial': frequency-spatial prior. (A) Schematic of electrode grid used for generating covariance priors ('frequency-spatial' prior only). (B) Diagonal prior. (C) Frequency block diagonal prior, $\ell_f^2 = 1\ Hz^{-2}$. (D) Frequency block diagonal prior, $\ell_f^2 = 250\ Hz^{-2}$. (E) Frequency block diagonal prior, $\ell_f^2 = 5000\ Hz^{-2}$. (F) Frequency-spatial prior, $\ell_f^2 = 250\ Hz^{-2}$, $\ell_s^2 = 0.1\ mm^{-2}$. (G) Frequency-spatial prior, $\ell_f^2 = 250\ Hz^{-2}$, $\ell_s = 1\ mm^{-2}$. (H) Frequency-spatial prior, $\ell_f^2 = 250\ Hz^{-2}$, $\ell_s^2 = 3\ mm^{-2}$.

This quantity is commonly referred to as the marginal likelihood or *evidence* [Linden and Sahani, 2003]. Replacing the probabilities in Equation 7.9 with their respective distributions allows us to express the log-evidence as a function of known quantities $\mathbf{y}$, $\mathbf{w}$, $\mathbf{X}$, $\mathbf{A}$, $\mathbf{\Sigma}_p$, σ_n^2 [Linden and Sahani, 2003]:

$$\ln \mathcal{E}(\sigma_n^2, \theta) = \frac{1}{2}\ln|2\pi\mathbf{A}| - \frac{N}{2}\ln(2\pi\sigma_n^2) - \frac{1}{2}\ln|2\pi\mathbf{\Sigma}_p| - \frac{1}{2}\mathbf{y}^\top\left(\frac{\mathbf{I}}{\sigma_n^2} - \frac{\mathbf{X}^\top\mathbf{A}\mathbf{X}}{\sigma_n^4}\right)\mathbf{y}. \tag{7.10}$$

By taking the derivative of Equation 7.10 with respect to σ_n^2 and the set of hyperparameters of the prior θ, the optimal values of these (hyper)parameters can be found using gradient ascent. Appendix C provides a detailed description of this process, including derivation of the evidence function and the partial derivatives with respect to the each of the (hyper)parameters. Estimating the prior distribution over the set of decoding weights from the data in this manner is known as *empirical Bayes*, and has been used previously to characterize the receptive fields in visual and auditory cortex [Linden and Sahani, 2003, Park and Pillow, 2011].

7.4 CHARACTERIZATION AND VALIDATION

Simulations were performed in order to validate the implementation of the method and characterize its performance under specific conditions. These simulations included assessing the ability of our implementation of to recover known (hyper)parameters, as well as evaluating the performance of empirical Bayes relative to that of standard least-squares as a function of characteristics of the training set.

7.4.1 Simulation

We generated simulated data according to the linear regression model provided by Equation 7.1 using decoding weights randomly drawn from a Gaussian distribution with zero-mean and covariance $\mathbf{\Sigma}_p$. This procedure is outlined by Algorithm 1.

Briefly, simulation of data began by choosing the form of the covariance prior, specifying the set of hyperparameters governing the prior, and then evaluating $\mathbf{\Sigma}_p$. A single set of weights $\mathbf{w}$ were then drawn at random from $\mathcal{N}(\mathbf{0}, \mathbf{\Sigma}_p)$, the prior distribution over the decoding weights. A set of simulated neural responses were generated by drawing at random from a zero-mean uniform distribution; these responses were then transformed into noiseless kinematic observations through multiplication with the decoding weights vector. Zero-mean noise with variance σ_n^2 was then added to these observations to generate the final set of noisy kinematic data.

Hyperparameter fitting was performed by minimizing the negative-log-evidence function through nonlinear optimization using the conjugate gradient method [Rasmussen, 2006]. As the optimization procedure is only guaranteed to find the local minimum of the negative-log-evidence function, multiple random restarts of the optimization procedure were used in an attempt to find the best set of (hyper)parameters. Once

Algorithm 1 Generation of simulated data (Frequency block-diagonal covariance prior)

```
Input
    $\mathbf{y}_f = [y_{f,1}, \ldots, y_{f,N}]^\top$            ▷ known function values
    $N$                                                          ▷ number of desired observations
    $D$                                                          ▷ number of features
    $\sigma_n^2$                                                 ▷ observation model noise variance
    $\theta = \{\sigma_{p,f}^2, \sigma_{p,n}^2, \ell_f^2\}$      ▷ covariance prior hyperparameters

for $i = 1 : D$ do
    for $j = 1 : D$ do
        $\Sigma_p(i,j) = \sigma_{p,f}^2 \exp\left[-\frac{1}{2\ell_f^2}(f_i - f_j)^2\right]\delta_{el}(i,j) + \sigma_{p,n}^2\delta_{ij}$
    end for
end for
Draw $\mathbf{w}$ from $\mathcal{N}(\mathbf{0}_D, \Sigma_p)$
Draw $N$ instances of $\mathbf{x}_i$ at random; $X = [\mathbf{x}_1, \ldots, \mathbf{x}_N]$
$\mathbf{y}_f = (\mathbf{w}^\top X)^\top$
$\mathbf{y}_n = [y_{n,1}, \ldots, y_{n,N}]^\top$, where $y_i \sim \mathcal{N}(0, \sigma_n^2)$
$\mathbf{y} = \mathbf{y}_f + \mathbf{y}_n$
```

the optimal parameter values were obtained from the simulated training data, the resultant set of decoding weights were obtained using Equation 7.3. For comparison, we also computed the standard least-squares solution (LS):

$$\mathbf{w}_{ls} = (\mathbf{X}\mathbf{X}^\top)^{-1}\mathbf{X}\mathbf{y} = \mathbf{X}^\dagger\mathbf{y}, \tag{7.11}$$

where $\mathbf{X}^\dagger$ is the Moore-Penrose pseudoinverse of $\mathbf{X}$. Estimated weights (both least-squares and empirical Bayes solutions) were compared to the known weights vector generated during the simulation procedure, with the the quality of the fit assessed using mean-squared-error (MSE):

$$MSE = \frac{1}{D}\mathbf{w}_{actual}^\top\mathbf{w}_{fit}, \tag{7.12}$$

where D is the number of features, $\mathbf{w}_{actual}^\top$ is the set of known decoding weights, and $\mathbf{w}_{fit}^\top$ is the set of decoding weights fit with either EB or LS.

Figure 7.4 shows the results of a single example simulation comparing least squares to empirical Bayes with the frequency-spatial prior for a 3×3 electrode grid with inter-electrode distance of 3mm. Figure 7.4A, 7.4B, and 7.4C show the actual, empirical-Bayes-fit, and least-squares-fit decoding weights vectors, respectively, while Figure 7.4D and 7.4E show the covariance priors calculated from the known and fit hyperparameters. Finally, Figure 7.4F depicts the mean-squared error for empirical Bayes and least-squares. As expected, estimating the decoding weights using empirical Bayes does a better job of recovering the known decoding

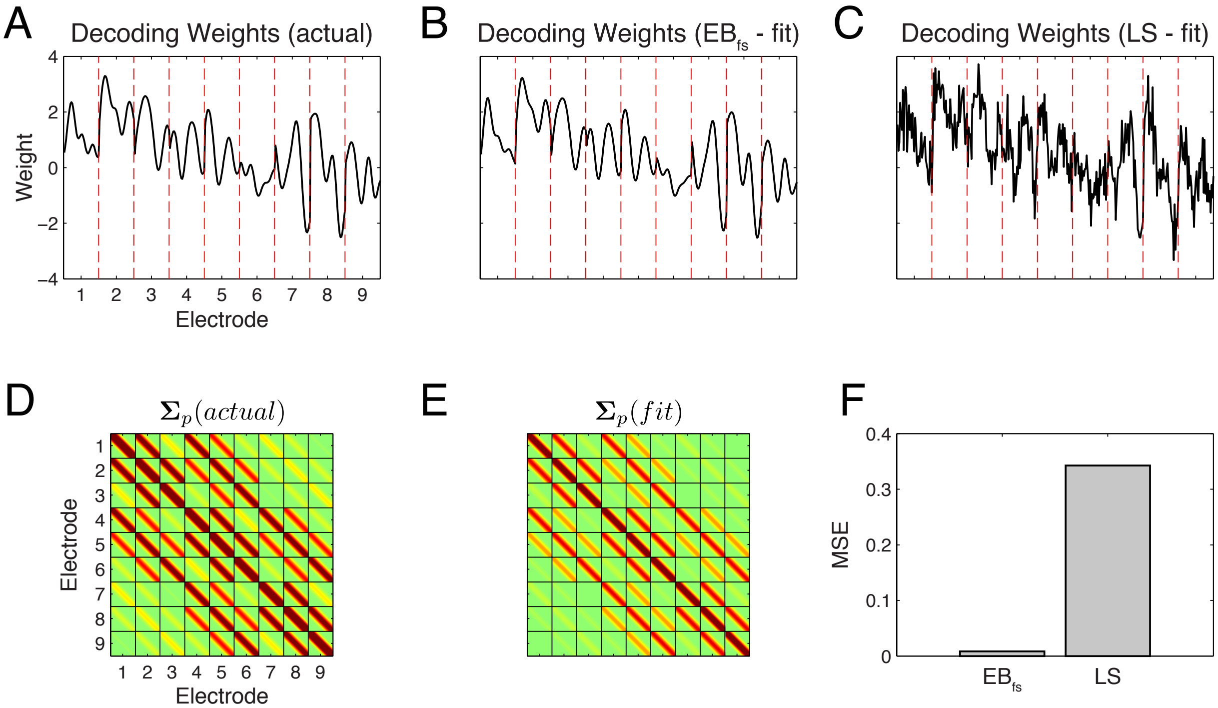

Figure 7.4: *Example of Empirical Bayes simulation results for the frequency-spatial prior.* An example of the validation procedure for the Empirical Bayes method is shown for the frequency-spatial covariance prior with parameters $\sigma_n^2 = 1$, $\sigma_{p,f}^2 = 2$, $\ell_f^2 = 1000\ Hz^{-2}$, $\ell_s^2 = 1\ mm^{-2}$, $N = 400$, $\Delta f = 5\ Hz$ for the 9-electrode grid shown in Figure 7.2A. (A) Simulated decoding weights. *Red* lines demarcate weights for simulated individual electrodes. (B) Weights recovered using Empirical Bayes (frequency-spatial prior). (C) Weights recovered using least-squares. (D) Actual covariance prior. (E) Covariance prior calculated using estimated hyperparameters. (F) Mean-squared error of recovered decoding weights for Empirical Bayes (EB_{fs}) and least squares (LS) decoders.

weights vector than least-squares. Furthermore, the similarity of the actual and fit covariance priors (Figure 7.4D and 7.4E) show that the evidence-optimization method is successful at approximating the known set of hyperparameters.

7.4.2 Recovery of parameters

In order to verify our implementation of empirical Bayes, we evaluated the ability of the evidence maximization process to recover known (hyper)parameter values using the simulation procedure described in the previous section. For each of the three covariance priors investigated, we compared known and recovered (hyper)parameters over a range of known values chosen to encompass those values likely for ECoG recordings. The values of all other (hyper)parameters were held constant when sweeping a particular (hyper)parameter through its range of likely values, with 50 simulations performed for each known value.

The recovery of (hyper)parameters for each of the 3 evaluated covariance priors is shown in Figure 7.5. As expected, the median values of the fit parameters match the known values. In general, parameters of the model (i.e., σ_n^2) were recovered more accurately than those of the prior (i.e., ℓ_f^2). Estimates of ℓ_s^2, the spatial length constant of the frequency-spatial prior, were found to be quite variable, particularly for small values of the hyperparameter. Furthermore, we observe that as the model becomes increasingly complex through the incorporation of additional hyperparameters in the prior, the ability to accurately recover the set of (hyper)parameters decreases.

In order to investigate this further, we compared the error in length constant hyperparameter recovery to the log ratio of EB to LS decoding weight recovery error. Hyperparameter recovery error was calculated according to Equation 7.13, while decoding weight recovery error ratio was calculated according Equation 7.14:

$$E_\ell = \log |\ell_{actual} - \ell_{fit}| \tag{7.13}$$

$$ER_{\mathbf{w}} = \log \left(\frac{MSE_{EB}}{MSE_{LS}} \right), \tag{7.14}$$

where ℓ_{actual} and ℓ_{fit} are the actual and fit length constants, and MSE_{EB} and MSE_{LS} are the mean-squared-error for the EB and LS decoders, respectively. Figure 7.6 shows $ER_{\mathbf{w}}$ plotted against E_ℓ for the frequency-block-diagonal (Figure 7.6A) and frequency-spatial (Figure 7.6B and 7.6C) priors. For the frequency-block-diagonal prior, nearly all of the data points lie below the $R_{\mathbf{w}} = 0$ line, indicating that the empirical Bayes method does a better job of recovering the true decoding weights than least-squares for this prior, regardless of the spectral length constant recovery error. This was not found to be the case for the frequency-spatial prior, however. Here, hyperparameter recovery was more error-prone than for frequency-block-diagonal prior. Furthermore, in cases where hyperparameter recovery error was high, the EB_{fs} decoder tended to do a poorer job recovering the true decoding weights than the LS decoder. Additionally, we observe that when length constant recovery error is high for one hyperparameter of the EB_{fs} decoder, it is typically

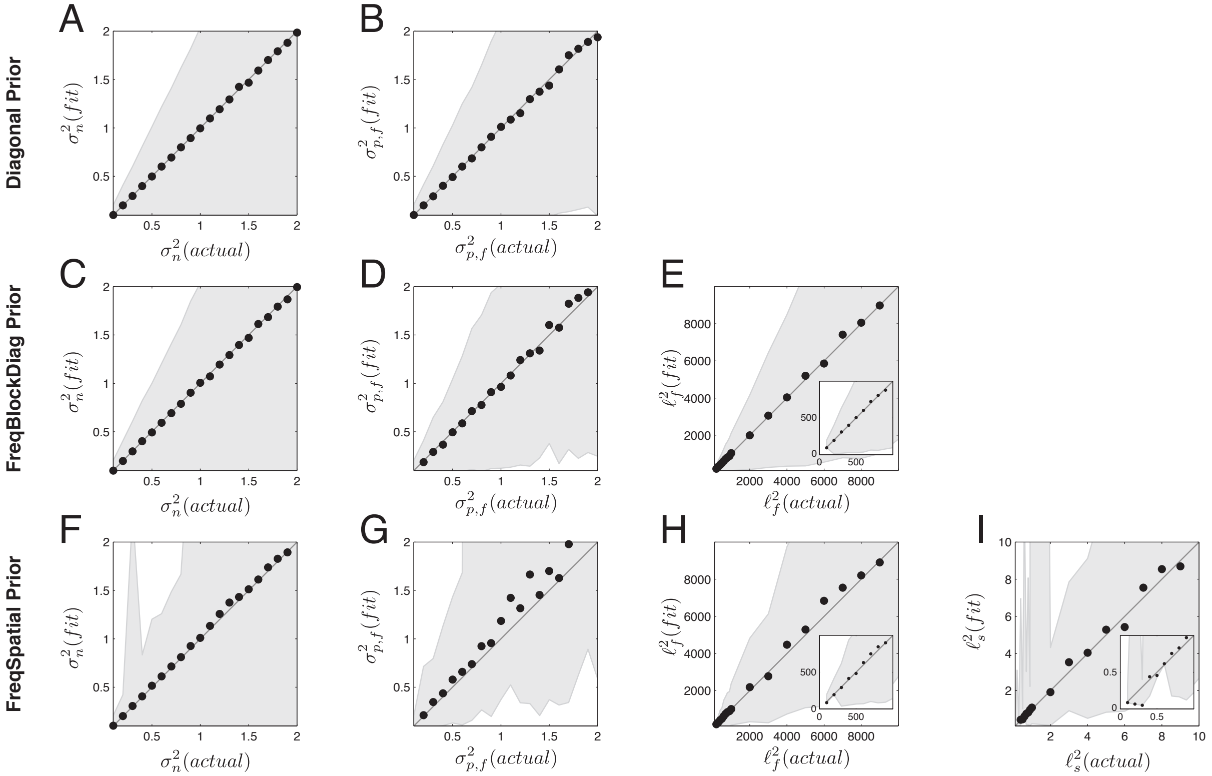

Figure 7.5: *Validation of ability Empirical Bayes to recover known (hyper)parameters.* Plots of actual (hyper)parameters versus recovered (fit) hyperparameters are shown for simulated data generated using the diagonal (*top* row), frequency-block-diagonal (*middle* row), and frequency-spatial (*bottom* row) covariance priors. (*A*) Diagonal prior, parameter σ_n^2. (*B*) Diagonal prior, hyperparameter $\sigma_{p,f}^2$. (*C*) Frequency-block-diagional prior, parameter σ_n^2. (*D*) Frequency-block-diagional prior, hyperparameter $\sigma_{p,f}^2$. (*E*) Frequency-block-diagional prior, hyperparameter ℓ_f^2. (*F*) Frequency-spatial prior, parameter σ_n^2. (*G*) Frequency-spatial prior, hyperparameter $\sigma_{p,f}^2$. (*H*) Frequency-spatial, hyperparameter ℓ_f^2. (*I*) Frequency-spatial, hyperparameter ℓ_s^2. *Black* dots indicate the median fit parameter value, while shaded gray regions indicates the middle 50% quantile range of fit values. Each panel shows the results obtained by sweeping the value of the specified (hyper)parameter over the indicated range, with simulations were repeated 50 times for each condition ($N = 1000$, $\Delta f = 5Hz$) for a 3×3 electrode configuration with an inter-electrode distance of $4mm$. Unless swept, constant values of $\sigma_n^2 = 1$, $\sigma_{p,f}^2 = 2$, $\ell_f^2 = 1000\ Hz^{-2}$, $\ell_s^2 = 1\ mm^{-2}$ were used for all simulations.

high for the other hyperparameter as well. In general, these results indicate that hyperparameter estimation is more accuate for the frequency-block-diagonal prior than for the frequency-spatial prior, with inaccurate length constant estimates for the frequency-spatial prior leading to decreased performance of the EB_{fs} decoder relative to the LS decoder.

7.4.3 Characterization of performance

In order to better understand the potential advantages of the EB_f and EB_{fs} decoders, we characterized the behavior of empirical Bayes with the frequency-block-diagonal covariance prior as a function of both the number of training observations and the number of electrodes. To accomplish this, simulations were performed in which we varied either the number of observations or the number of electrodes, holding all other (hyper)parameters constant. The results of these simulations are shown by Figure 7.7.

When varying the number of observations (Figure 7.7, *top* row), we find that the error in the recover of the decoding weights for the EB_f decoder decreases as the number of observations is increased. Apart from small values of N_{obs}, where the two methods are equivalent, the recovery of the decoding weights using the EB_f decoder is more accurate than using least-squares. Performance of the LS decoder exhibits a peak in weight recovery error where the number of observations is equal to the number of features. The region to the left of this peak corresponds to those conditions where the number of observations is less than the number of features (the *under-determined* region), while the region to the right corresponds to those conditions where the number of observations exceeds the number of features (the *over-determined* region). For the under-determined region, the least-squares solution is computed using the minimum-norm estimate of pseudoinverse of $\mathbf{X}$ in Equation 7.11[2], with the standard pseudoinverse used for the over-determined region. As the number of observations is increased (i.e. as $N_{obs} \to \infty$), the error of the two methods begins to converge, as evidenced by the increase in the error ratio (MSE_{EB}/MSE_{LS}) for increasingly large values of N_{obs}. Furthermore, the variability in the recovery of (hyper)parameters σ_n^2 and ℓ_f^2 is reduced for larger numbers of training observations.

On the other hand, we find that the performance of both methods decreases as the number of electrodes, and consequently the number of simulated neural features, is increased (Figure 7.7, *bottom* row). In addition, the ratio of the error in the recovery of the simulated decoding weights decreases with increasing numbers of electrodes. This indicates that while performance of both the EB_f and LS decoders should be expected to decrease for increasing numbers of electrodes, the performance decrease for EB_f should be less than that of LS. The ability of empirical Bayes to estimate observation noise parameter σ_n^2 shows little change with increasing numbers of electrodes, while estimates of hyperparameter ℓ_f^2 become increasingly more accurate. This is to be expected, as increasing the number of electrodes (and consequently, the number of features) provides additional pairwise comparisons of the neural feature activity which are used to estimate the spectral length constant characterizing the correlation between features as a function of frequency.

---

[2] See Appendix B.

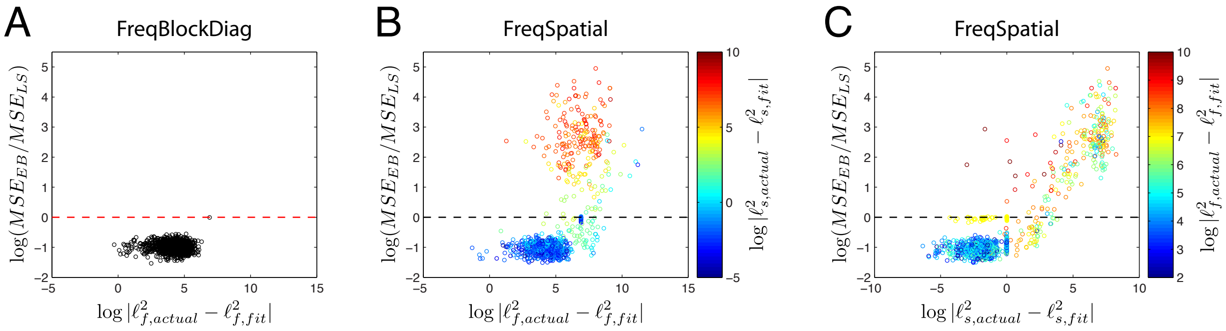

Figure 7.6: *Effect of hyperparameter fit on recovery of decoding weights.* The ratio of decoding weight recovery error of empirical Bayes to that of least-squares regression is plotted as a function of hyperparameter recovery error for the EB_f and EB_{fs} decoders. (A) Hyperparameter recovery error for ℓ_f^2 for the frequency-block-diagonal covariance prior. (B) Hyperparameter recovery error of ℓ_f^2 for the frequency-spatial covariance prior. Colors indicate the hyperparameter recovery error of ℓ_s. (C) Hyperparameter recovery error of ℓ_s^2 for the frequency-spatial covariance prior. Colors indicate the hyperparameter recovery error of ℓ_f^2. Dashed lines indicates equivalent performance of Empirical Bayes and least-squares decoders. Results are shown for the simulation results from the "sweep σ_n" condition for the frequency-block-diagonal and frequency-spatial covariance priors (Figure 7.5C and 7.5F, respectively).

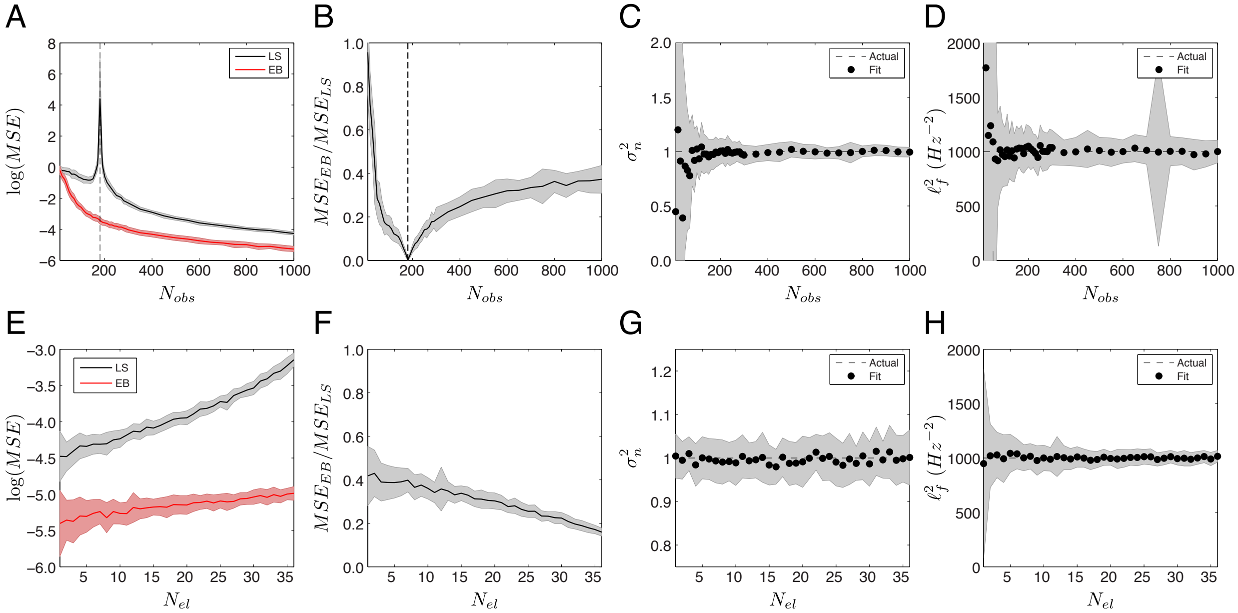

Figure 7.7: *Characterization of the performance of the* EB_f *decoder as a function of the number of observations and features.* EB_f performance and (hyper)parameter recovery is shown as a function of the number of observations (*top* row) and the number of electrodes (*bottom* row). (A) Log-mean-squared error (log MSE) of fit decoding weights for the LS and EB_f decoders as a function of the number of observations N_{obs}. (B) Ratio of EB to LS MSE as a function of N_{obs}. (C) Recovery of parameter σ_n^2 as a function of N_{obs}. (D) Recovery of spectral length constant ℓ_f^2 as a function of N_{obs}. (E) Log-mean-squared error of fit decoding weights for the LS and EB_f decoders as a function of the number of electrodes N_{el}. (F) Ratio of EB to LS MSE as a function of N_{el}. (G) Recovery of parameter σ_n^2 as a function of N_{el}. (H) Recovery of spectral length constant ℓ_f^2 as a function of N_{el}. (Hyper)parameters $\sigma_n^2 = 1$ and $\ell_f^2 = 1000\ Hz^{-2}$ were fixed for all simulations. Simulations of performance as a function of N_{obs} were performed for a 9-electrode grid with $\Delta f = 10\ Hz$, while simulations of performance as a function of N_{el} were performed with $N_{obs} = 1000$ and $\Delta f = 10\ Hz$. Results are presented as the mean ± standard deviation over 50 repetitions at each condition. The dashed *black* line in panels (A) and (B) indicate the point at which N_{obs} is equal to the number of features.

Taken together, the simulations summarized by Figure 7.7 suggest that the performance benefits of the empirical Bayesian decoders with the frequency-block-diagonal and frequency-spatial priors should be greatest when estimating decoding weights for large numbers of neural features using few training observations. As this is commonly the case for many ECoG BMI experiments, these results illustrate the potential advantage of the EB_f and EB_{fs} methods.

7.5 APPLICATION TO ELECTROPHYSIOLOGICAL DATA

In order to evaluate the performance of empirical Bayes in decoding movement-related activity from ECoG recordings, we assessed the ability of the EB_f and EB_{fs} decoders in predicting BMI command signals from human ECoG recordings obtained during closed-loop cursor control. The following sections will describe the datasets and methods used in this comparison, as well as the performance of empirical Bayesian decoding compared to several commonly-used methods.

7.5.1 Datasets

Electrocorticographic recordings from two human subjects (subjects S1 and S2) with upper limb paralysis were used to assess the performance of proposed empirical Bayes decoding method. Detailed descriptions of these subjects are provided in Chapter 4. Data sets for each subject consisted of 160 trials of two-dimensional closed-loop computer cursor control. Subject S1 performed an 8-target variant of the 2D cursor task, while Subject S2 performed a 4-target task. Specific trial sets were selected from the peak performance period for each subject.

7.5.2 Evaluation of performance

We compared the performance of empirical Bayes with the three previously-described priors (diagonal, frequency-block-diagonal, and frequency-spatial) to two methods commonly utilized in brain-machine interface systems: linear-least-squares regression and Optimal Linear Estimation. Least-squares linear regression was performed according to Equation 7.11, while OLE was implemented as described in Chapter 4.

Time-domain ECoG recordings were transformed to the time-frequency domain using the Burg method [Kay and Marple, 1981] over the $0 - 200$ Hz frequency range (25th order, 100ms window size) at frequency resolutions of 5Hz, 10Hz, and 20Hz. Dura-facing electrodes, as well as those exhibiting a high degree of noise, were excluded from the decoding analysis, yielding 26 and 27 electrodes for Subjects S1 and S2, respectively. Time-frequency data were normalized to pseduo-Z-scores relative to responses during the target presentation epoch of the task [Edwards et al., 2009, Ray et al., 2008b]. Normalized time-frequency data were averaged over the $0 - 500$ms time window relative to the onset of cursor control in an attempt to reduce the influence of error correction on the decoding results. The resultant feature sets consisted of 260, 520, and 1040 features (20Hz, 10Hz, and 5Hz frequency resolution, respectively) for Subject S1 and 270, 540, and 1080 features for Subject S2.

Least-squares (LS), optimal linear estimation (OLE), and empirical Bayesian decoders with diagonal (EB_d), frequency-block-diagonal (EB_f), and frequency-spatial (EB_{fs}) priors were used to predict 2D target direction from the neural feature sets. In order to assess the influence of the amount of training data on performance, 2, 4, and 20-fold cross validation was used (80, 120, and 152 training trials, respectively). Both neural and kinematic data were normalized prior to decoding by subtracting the mean of the training set of data. Decoding accuracy was evaluated using squared-error (SE) and was averaged over the two kinematic dimensions. Performance between decoding conditions was compared using one-tailed t-tests, with differences in mean squared error considered significant at $p < 0.05$. In all, we examined performance of the 5 decoding methods at a total of 9 different conditions (3 frequency resolutions × 3 cross validation splits).

7.5.3 Decoding results

Figure 7.8 depicts the mean squared error of all for all decoders across all evaluated conditions for subject S1 and S2. Generally, we find that the mean squared error of the predicted kinematics for the EB_f and EB_{fs} decodes is lower than that of the LS, OLE, and EB_d decoders for Subject S1, and lower than the LS and EB_d decoders for Subject S2. We observe several interesting characteristics when the MSE of the predicted kinematics is examined across conditions. First, we observe that the error for empirical Bayesian decoders tend to decrease as the amount of training data is increased (i.e. the *fraction of training data, FTD, is increased*), consistent with simulations showing that the accuracy of the recovered weights increases as the number of training observations is increased. Secondly, decoding accuracy of the EB decoders does not appear to be substantially affected by the number of features, regardless of the amount of training data. This is intriguing, as the number of features is increased fourfold from the 20Hz frequency resolution to the 5Hz frequency resolution with no decline in performance.

Interestingly, we find that performance using the LS decoder decreases as the number of observations increases or the number of features decreases. While this may seem counter-intuitive, it is important to note that in all the conditions evaluated, the number of features (260 – 1080) exceeded the number of observations (80 – 152). In this case, the system of equations $\mathbf{y} = \mathbf{X}^\top \mathbf{w}$ is under-determined, and the minimum norm estimate[3] was used to estimate the pseudoinverse $\mathbf{X}^\dagger = (\mathbf{X}\mathbf{X}^\top)^{-1}\mathbf{X}$ in the least-squares solution for decoding weights vector $\mathbf{w}$.

An example of the comparison of mean squared error for a single experimental condition is provided by Figure 7.9, while the p-values for all statistical comparisons are provided by Figure 7.10. One tailed t-tests were used to identify the conditions for which EB_f and EB_{fs} out performed standard LS and OLE

---

[3] See Appendix XX

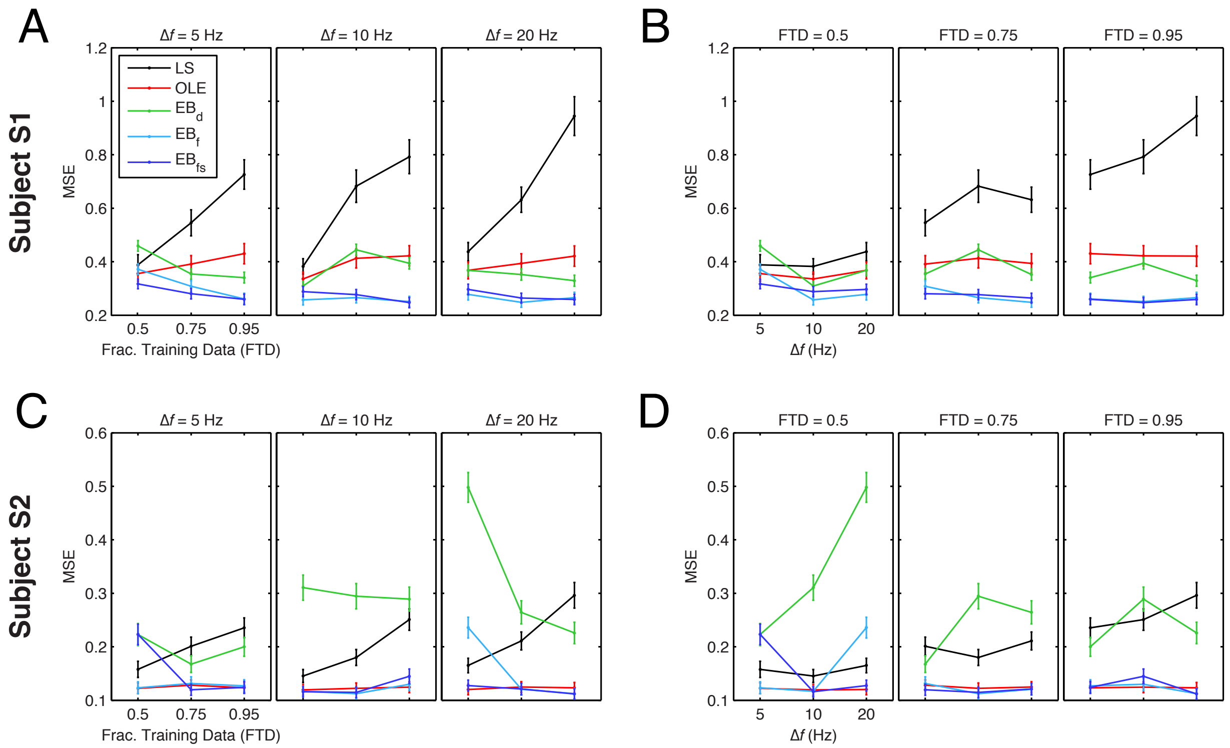

Figure 7.8: *Decoding performance across experimental conditions.* A comparison of mean squared error (MSE) for all decoding methods and evaluation conditions is shown for subjects S1 and S2. (A) Performance evaluated on data from Subject S1, plotted as a function of fraction of the training data (FTD). (B). Subject S1, plotted as a function of frequency resolution (Δf). (C) Subject S2, plotted as a function of FTD. (D) Subject S2, plotted as a function of Δf. *Black*, *red*, *green*, *light-blue*, and *dark-blue* lines show performance for the LS, OLE, EB_d, EB_f, and EB_{fs} decoders, respectively. Error bars indicate SEM. Note that the pairs of plots for each subject (i.e. A and B) contain the same data, but are plotted differently in order to show decoding accuracy as a function of both FTD and Δf.

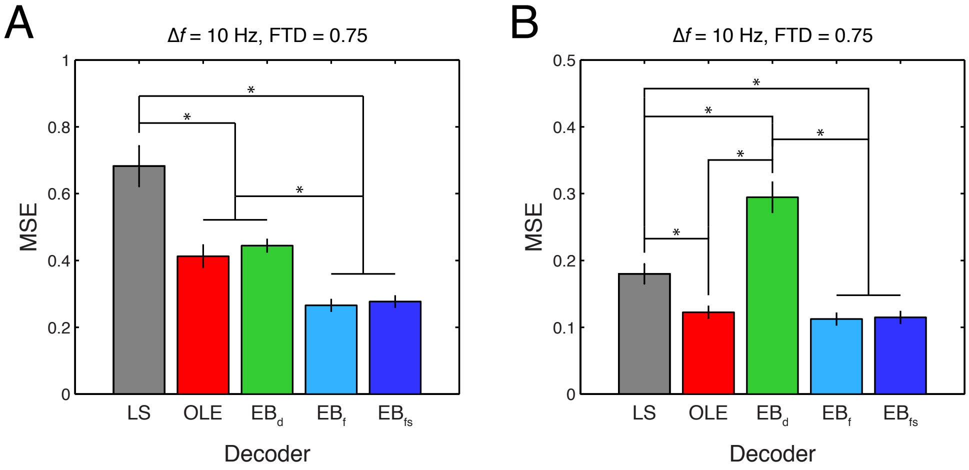

Figure 7.9: *Decoding accuracy for a single experimental condition.* Comparison of mean squared error (MSE) for all evaluated decoding methods is shown for subjects S1 (A) and S2 (B) for the $\Delta f = 10Hz$, $FTD = 0.75$ condition. Error bars indicate SEM, while asterisks indicate significance at $p < 0.05$.

decoders, as well as those conditions where the standard decoders outperformed the empirical Bayesian decoders. Finally, a summary of the comparisons between the three empirical Bayes decoders, least-squares, and OLE decoders is provided by Figure 7.11. Together, these figures confirm the qualitative comparisons apparent in Figures 7.8: EB_f and EB_{fs} decoders outperform LS for data from both subjects S1 and S2, with EB_f and EB_{fs} also outperforming OLE when evaluated on data from Subject S1. Performance for the EB_f, EB_{fs}, and OLE decoders was found to be comparable for data from Subject S2.

Finally, we examined the hyperparameters of the prior covariance matrix $\mathbf{\Sigma}_p$ fit using evidence maximization for both the EB_f and EB_{fs} decoders; a summary of these values is provided by Figure 7.12. While there is a degree of overlap in the value of the spectral length constant ℓ_f fit by EB_f and EB_{fs}, the values found by the EB_{fs} decoder are noticeably higher than those found by the EB_f decoder for both subjects S1 and S2. In addition, we a tendency of the length constants (both ℓ_f and ℓ_s) fit by the EB_{fs} decoder to approach zero, though as shown by Figure 7.12D and 7.12H, both ℓ_f and ℓ_s were not fit to near-zero values simultaneously.

7.6 DISCUSSION

We have developed a empirical Bayesian framework which seeks to improve the decoding of information from ECoG recordings. We propose two parametric covariance matrix priors, the *frequency-block-diagonal* and

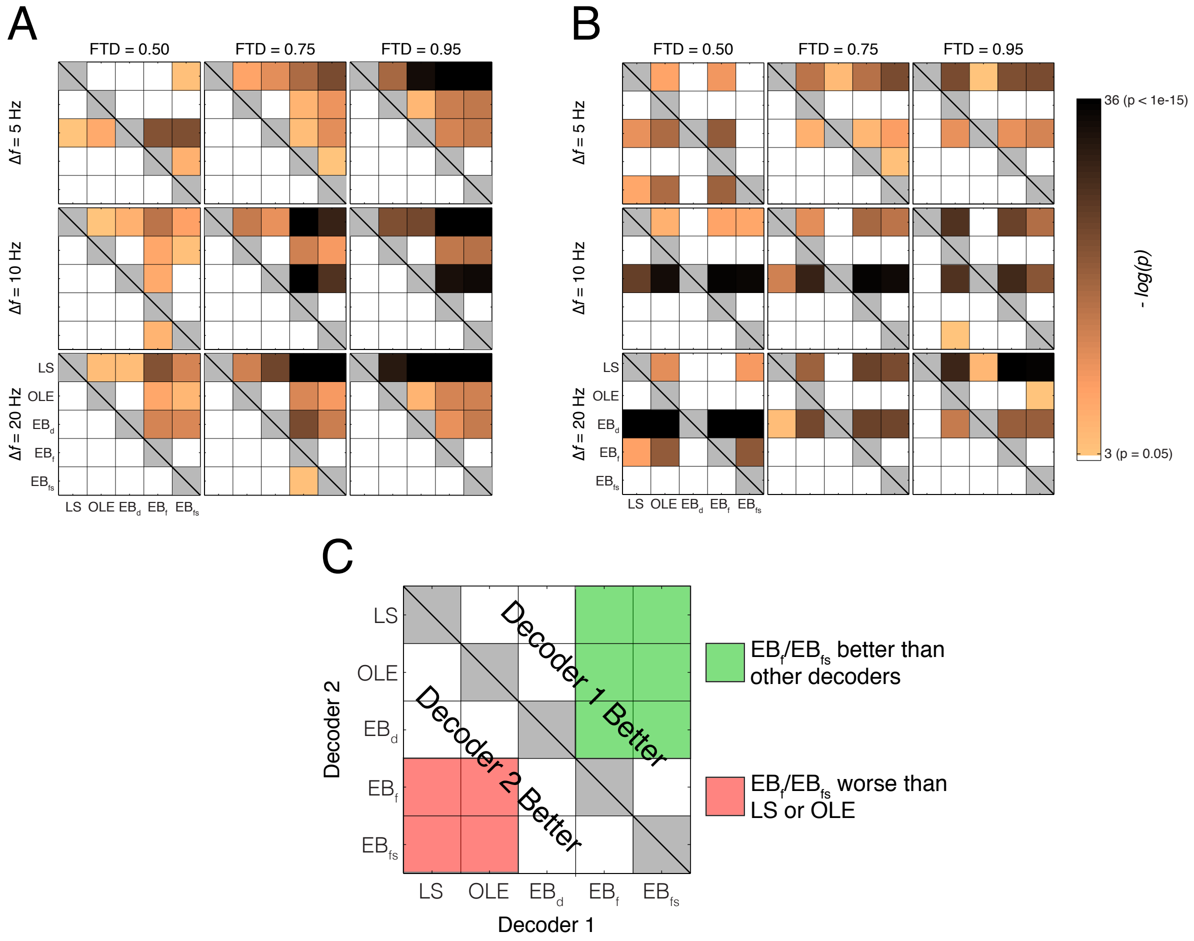

Figure 7.10: *Comparison of decoding accuracy across decoders.* Negative-log-p-values resulting from one-sided t-tests comparing squared error distributions between decoders are shown for subjects S1 (A) and S2 (B). Non-significant comparisons are indicated by *white* squares, while colored squares indicate comparisons which were significant at $p < 0.05$. Comparisons of decoding accuracy for a single condition (i.e. frequency resolution and fraction of training data) are represented by a single 5×5 square as shown by the schematic in (C). The upper triangular half of each 5×5 square indicates the results of a left-tailed test evaluating whether the squared-error distribution of the decoder on the X-axis (Decoder 1) is significantly less than that of the decoder on the Y-axis (Decoder 2). The lower triangular half of each 5×5 square indicates the results of a right-tailed test evaluating whether the squared-error distribution of Decoder 1 is significantly greater than that of the Decoder 2. The *green*-shaded regions in (C) indicate statistical comparisons which, if significant, indicate that the Empirical Bayesian decoders (EB_f and EB_{fs}) out-perform the other evaluated decoders. The *red*-shaded regions in (C) indicate those comparisons which the EB_f and EB_{fs} were significantly worse than the least-squares or OLE decoders.

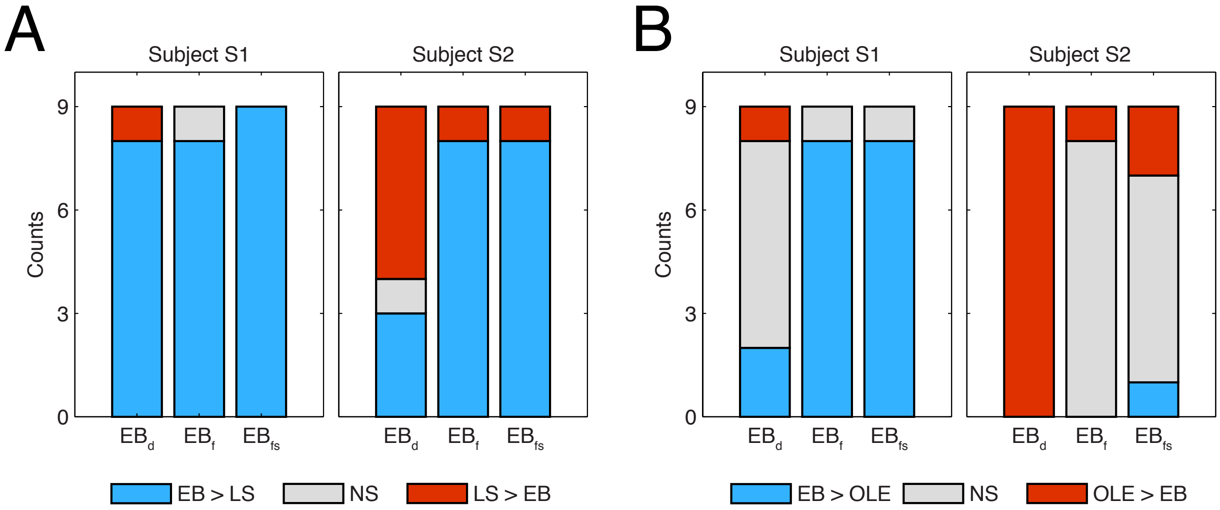

Figure 7.11: *Summary of statistical comparisions between Empirical Bayes decoders and traditional decoders.* Counts of significant differences in decoding accuracy for all conditions are shown for subject S1 and S2. (A) Empirical Bayesian decoders compared to least-squares. (B) Empirical Bayesian decoders compared to OLE. *Blue* bars indicate counts of conditions where the empirical Bayesian decoder outperformed the traditional decoder (either LS or OLE), *red* bars indicate conditions where the traditional decoder outperformed the empirical Bayesian decoder, and *gray* bars indicate experimental conditions where there was no significant difference in decoding accuracy between the two decoders.

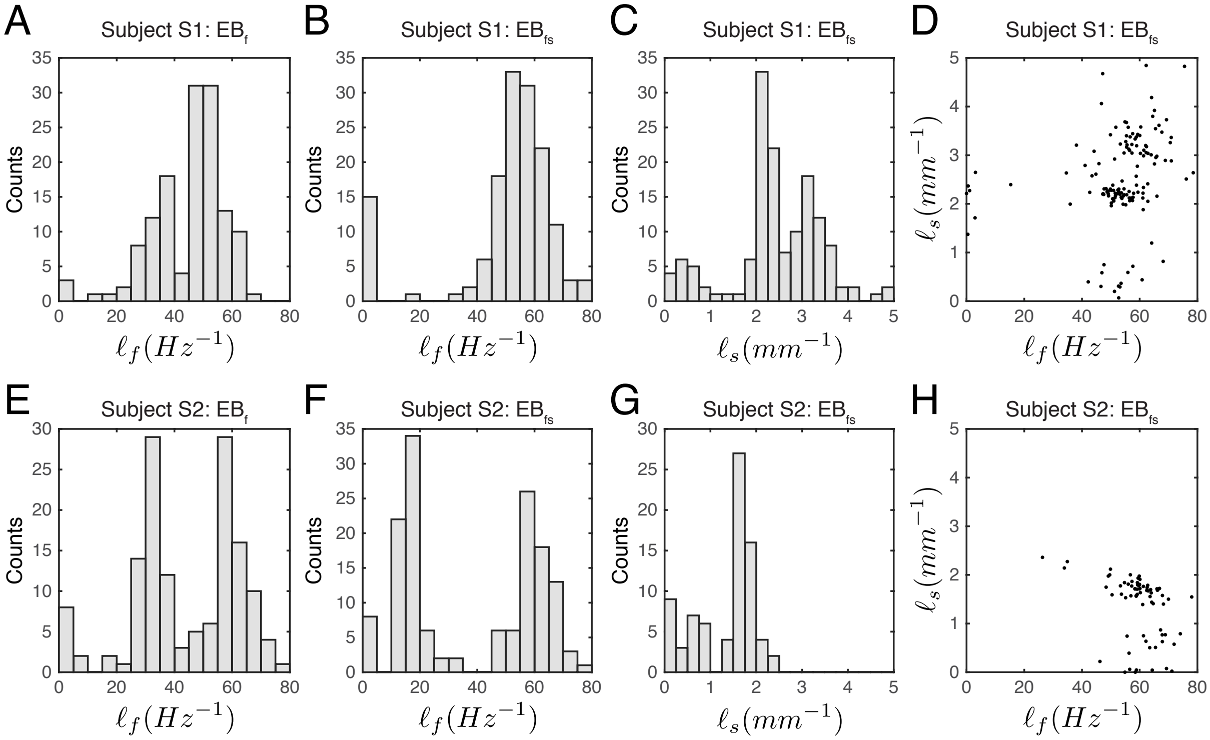

Figure 7.12: *Hyperparameters fit to data from electrophysiological recordings.* Shown are distributions of hyperparameters ℓ_f and ℓ_s fit to data from subjects S1 (*top* row) and S2 *bottom* row. (A) Subject S1: frequency-block-diagonal prior, hyperparameter ℓ_f. (B) Subject S1: frequency-spatial prior, hyperparameter ℓ_f. (C) Subject S1: frequency-spatial prior, hyperparameter ℓ_s. (D) Scatter plot of ℓ_f versus ℓ_s for Subject S1. (E) Subject S2: frequency-block-diagonal prior, hyperparameter ℓ_f. (F) Subject S2: frequency-spatial prior, hyperparameter ℓ_f. (G) Subject S2: frequency-spatial prior, hyperparameter ℓ_s. (H) Scatter plot of ℓ_f versus ℓ_s for Subject S2.

frequency-spatial priors, which impose specific spectral and spatial correlation structure on decoding weights fit using probabilistic linear regression. Using simulated data, we show that hyperparameters of these priors can be fit using evidence maximization, and that recovery of decoding weights is improved relative to linear least-squares regression when simulated weights are randomly drawn from these priors. This framework was found to out-perform least-squares regression, and in some cases Optimal Linear Estimation, when predicting intended target direction from ECoG recordings obtained during closed-loop cursor control by individuals with upper-limb paralysis.

By placing a prior distribution over the decoding weights, probabilistic linear regression constrains the resultant set of weights by effectively penalizing those solutions which do not match the correlation structure of the prior. As such, it is important to ensure that the form of the priors chosen match those expected from the data. The frequency-block-diagonal and frequency-spatial priors employed here were chosen to capture two characteristics of ECoG recordings: correlation amongst neighboring frequency bands (spectral correlation), and correlation amongst neighboring electrodes (spatial correlation). As shown by Figures 7.2 and 7.2, the forms of the covariance matrix prior chosen reflect the prominent correlation structure of ECoG signals.

The performance benefits of the EB_f and EB_{fs} decoders could possibly attributed to regularization of the decoding weights solution. As the hyperparameters of the covariance priors are fit to training data using evidence maximization, empirical Bayes with the frequency-block-diagonal and frequency-spatial covariance priors can potentially reduce to L_2-regularized linear regression (ridge regression) for specific length constant values. This provides a check against overfitting by preventing magnitude of the decoding weights from growing too large, and has been used to control overfitting when decoding kinetic information from ECoG signals [Flint et al., 2014]. Furthermore, this allows for the estimation of the weights in cases where the number of training observations is less than the number of neural features; this is commonly the case for BMI studies, where the number of training observations is limited on account of time and motivational considerations. To determine whether the performance benefits of EB_f and EB_{fs} were solely an effect of regularization, we compared the EB_f and EB_{fs} decoders to empirical Bayes with a diagonal prior, which is equivalent to ridge regression [Park and Pillow, 2011]. In addition, we assessed performance of all decoders in a data-limited regime (i.e., the number of observations is less than the number of features), where the least-squares solution is that which minimizes the norm of the decoding weights (see Appendix XX). We find that the EB_f and EB_{fs} decoders outperform both the minimum-norm least-squares estimate and ridge regression in nearly all conditions evaluated. Thus, we attribute the performance benefits of EB_f and EB_{fs} to something beyond those conveyed by regularization.

By imposing spectrally and spatially-constrained priors when solving for the decoding weights, solutions can be obtained in which neighboring features have similar weight values. This allows cortical activity from functionally-similar frequency bands to be combined, potentially increasing the signal-to-noise ratio of predicted kinematics by averaging out independent observation noise across neural features. Previously this has been performed by averaging features in an *ad-hoc* manner [Acharya et al., 2010]. We find that decoding performance using the proposed empirical Bayesian approach does not appreciably change when the frequency band resolution of ECoG recordings is changed from 20Hz to 5Hz, despite the resultant 4-fold increase in the number of features. While not a direct comparison to prediction using *ad-hoc* feature averaging, this suggests that the proposed EB decoder may provide the de-noising effect of averaging across larger frequency bands while simultaneously allowing for the de-correlation of frequency bands which may occur during closed-loop training [Rouse and Moran, 2009].

It is important to note that the framework presented here has the potential to be extended through the development of more descriptive covariance priors. Such priors could theoretically account for the anti-correlated nature of the low frequency ($f < 40Hz$) and high frequency ($f > 40Hz$) bands, as well as allow for spatial correlation to vary as a function of frequency. However, as implied by the differences in hyperparameter recovery between the frequency-block-diagonal and frequency-spatial priors, increasing the complexity of the covariance prior through the incorporation of additional hyperparameters may make maximization of the evidence function more difficult. Furturemore, the probabilistic nature of empirical Bayes also may allow this work to be extended to more advanced decoding methods, such as Kalman filters which utilize probabilistic linear regression to track parameter updates over time [Li et al., 2011].

There are several limitations to the presented work. First, we observe that the values of the hyperparameters fit to simulated data, particularly for the frequency-spatial covariance prior, can be prone to error. Ultimately this resulted in an inability of the method to recover the true decoding weight parameters, and consequently, diminished performance of the method. We also find that the distributions of length constant hyperparameter values fit to human ECoG recordings using the EB_{fs} method exhibited exhibited prominent peaks near $\ell^2 = 0$. These findings appear to be indicative of an inability of the evidence maximization procedure to find the global maximum of the evidence function. It may be the case that more advanced optimization methods will improve the ability of the EB_{fs} method to recover reasonable estimates of length constant hyperparameters.

Interestingly, we find that OLE performs comparably to the EB_f and EB_{fs} decoders when applied to data from Subject S2. It is worth noting that OLE was used to obtain closed-loop brain-control with this subject. It is possible that brain control training resulted in neural adaptation which improved the fit of neural activity to the cosine-tuning model assumed by OLE; such adaptation has been observed during closed-loop brain control by non-human primates [Taylor et al., 2002]. An evaluation of decoding performance on electrical field potential data collected during an open-loop task, such as a center-out reaching task, will prevent closed-loop adaptation effects from influencing decoding results.

Furthermore, performance of OLE when evaluated on data from Subject S2 calls into question one of the assumptions of the proposed empirical Bayes approach: namely, that correlation between the neural population implies correlation between decoding weights. In the case of independent sets of features, it may be preferable to de-correlate weights for those features exhibiting significant correlation. In fact, such de-correlation allows the Optimal Linear Estimator to correct for biases in the distribution of the encoding models of the neural population; this allows OLE to improve decoding performance over traditional methods such as the Population Vector Algorithm [Chase et al., 2009]. It is possible that improvements in decoding performance over the proposed empirical Bayesian approach may be achieved by placing a prior distribution over the encoding parameters of the population (i.e. the preferred directions of the individual neural features) rather than the decoding parameters themselves. Such an approach is not possible within the empirical Bayesian framework applied here, as the weights are calculated independently for each kinematic dimension.

Despite these potential shortcomings, we show that an empirical Bayesian framework utilizing covariance priors encouraging spectral and spatial correlation amongst decoding weights has the potential to improve the decoding of information from electrocorticographic signals. We believe that this method is particularly well-suited for use in decoding command signals for ECoG-based BMI systems given the data-limited nature of BMI experiments. However, the approach developed here may be equally applicable for the decoding of semantic [Wang et al., 2011a] or language [Pei et al., 2011, Pasley et al., 2012] information from field potentials recorded from the cortex using a number of different recording modalities. Ultimately, we hope that the use the empirical Bayesian decoding framework presented here can yield further insight into the information contained in large-scale cortical recordings.

8.0 CONCLUSIONS

This dissertation has presented work attempting to advance the state of electrocorticographic brain-machine interfaces. In In Chapter 4, it was shown that individuals with upper-limb paralysis are capable of modulating sensorimotor cortical activity to successfully control a brain-machine interface with three degrees of freedom. This work is the first demonstration of successful ECoG-based BMI control by individuals with upper-limb paralysis, as well as the first instance of three-dimensional BMI control using ECoG. Though these results are encouraging, there are several key limitations worthy of mention. First, though satisfactory three-dimensional control was demonstrated, the level of performance achieved was limited compared to recent work showing high-dimensional control using intracortical electrodes [Hochberg et al., 2012, Collinger et al., 2012]. Second, it was found that translation of three-dimensional cursor control to three-dimensional robotic arm control was difficult, suggesting the generalization of somatotopic-based closed-loop control may be problematic. Such shortcomings may be overcome through long-term closed-loop training, which could facilitate learning on the part of the subject, or through improved extraction of BMI command signals from ECoG recordings. These shortcomings were addressed through investigation into cortical adaptation during acquisition of ECoG-based BMI control, an evaluation of a chronically-implanted BMI system in a non-human primate, and the development of an improved decoding method for ECoG signals.

To better understand how subject learning and decoder calibration could improve control performance, both cortical and decoder adaptation were investigated during the acquisition of somatotopic BMI control, with it found that decoder calibration as well as cortical adaptation during fixed-decoder training could facilitate performance improvements (Chapter 5). In order to better understand the capability of ECoG to provide robust, long-term recordings, work was conducted assessing the effects of chronic implantation of an ECoG electrode grid, demonstrating that movement-related modulation could be recorded from a ECoG electrode grid nearly two years post-implantation despite the presence of substantial fibrotic encapsulation (Chapter 6). Finally, to improve the extraction of information command signals from ECoG, a decoding method was developed which improved BMI decoding performance through the incorporation of weight-space priors accounting for the expected correlation structure of electrical field potentials (Chapter 7). Combined, this work both demonstrates the feasibility of ECoG-based BMI systems as well as addresses some of key challenges which must be overcome before such systems are translated to the clinical realm. We will conclude by briefly discussing the implications of this work as well as the future of electrocorticographic brain-machine interfaces.

8.1 IMPLICATIONS AND FUTURE DIRECTIONS

In order for a brain-machine interface system to truly improve the quality of life for individuals with movement disorders, it must be capable of restoration of function while being both robust and easy to use. The work presented in this dissertation provides evidence that an electrocorticography-based BMI could potentially meet these requirements. The demonstration of closed-loop cursor control by two subjects with upper-limb paralysis shows that voluntary modulation of sensorimotor cortex is possible long after spinal cord injury or onset of ALS. Furthermore, it was found that a somatotopic control strategy can enable closed-loop control with at least three degrees of freedom, despite the increased cognitive burden likely placed on subjects as a consequence of this strategy. Future studies are required to determine the upper bound on the number of independent command signals which could potentially extracted from ECoG recordings using this strategy. Though the quality of closed-loop control demonstrated here may be considered by some to be mediocre compared to that shown by studies utilizing penetrating microelectrodes, it is important to note that the goal of this work was to demonstrate the feasibility of an ECoG-based BMI system, and not to maximize control performance. Such efforts are better suited for long-term studies providing sufficient training time to fully evaluate the quality of closed-loop control. In addition, long-term studies will provide an opportunity to determine if the difficulties encountered when transitioning from cursor control to prosthetic limb control can be overcome through prolonged training.

Such efforts will likely be enhanced by the investigation into a chronically-implanted BMI system presented in Chapter 6. By showing that recordings can be obtained from subdural electrodes implanted for nearly 2 years, evidence for the likely failure mode for current subdural ECoG electrode technology, namely, fibrotic grid encapsulation, is provided. The characterization of the nature of this encapsulation provided by the histological analyses presented will aid in the development of novel technologies seeking to reduce the foreign body response to implanted electrode grids and enable a robust, long-term brain interface. While such an interface is necessary for BMI systems, this also enables long-term studies of cortical activity using ECoG. Such work would inevitably include an evaluation of long-term ECoG BMI performance, necessary to confirm the often-theorized stability of ECoG suggested by offline decoding of electrical field potentials [Chao et al., 2010, Flint et al., 2013].

The ability to obtain stable, long-term ECoG recordings will also allow for the improved study of cortical adaptation during acquisition of BMI control, including additional investigation into the results presented in Chapter 5. As evidenced by the preservation of the relationship between the preferred directions of individual neural features during cortical adaptation, use of the the somatotopic control strategy could constrain the cortical adaptation possible during closed-loop control. By studying such adaptation over longer timescales, as well as in response to perturbations requiring modification of the instructed control strategy, the existence of such constraints can be determined. This could allow for the development of novel training strategies specifically developed for the somatotopic control strategy, potentially providing for both an enhancement in control performance and a reduction in training time.

The seeming reliance of ECoG BMI systems on the somatotopic strategy highlights a key problem for the realization of high degree-of-freedom BMI control using ECoG, namely, whether a somatotopic-based ECoG BMI system can provide the signal independence required for assistive devices with many DoF. Control of a prosthetic limb for reaching and grasping will requires a minimum of 4 independent command signals, with more required for fine wrist and hand control. As demonstrated by the spatial correlation of preferred directions during periods of closed-loop brain control, the number of independent neural populations which can be activated using a somatotopic control strategy may be limited. Thus, in order to achieve high DoF control using ECoG, some form of cortical adaptation may be required. Recent efforts at decorrelating ECoG signals through closed-loop training [Rouse et al., 2013] provide a crucial demonstration of the capacity of subjects to increase the independence of ECoG signals. However, additional work is needed to fully assess the capacity of human subjects to adapt their cortical activity over time. Whether the adaptation mechanisms observed in Chapter 5 persist over longer timescales, and how cortical adaptation is constrained by somatotopic BMI control strategies, is an area worthy of further study. In addition, the cognitive burden imposed by somatotopic-based BMI control must be better understood. BMI systems based on SU/MUA recordings seemingly enable intuitive control for users, as such systems directly translate movement intention into action without the need for abstract control strategies. In order for an ECoG BMI system to truly approach clinical relevancy, a better understanding of the potential for the internalization of the somatotopic control strategy through closed-loop training, and how such internalization may be affected by adaptation, will be required.

Such efforts could potentially be avoided altogether through the use of a naturalistic control strategy. Such a strategy would likely result in a more intuitive control interface for the user, and is seemingly supported by work demonstrating that arm movement kinematics can be predicted from ECoG recordings [Nakanishi et al., 2013, Shimoda et al., 2012, Chao et al., 2010]. Unfortunately, the potential for these findings to be translated to online control is currently unknown, as successful closed-loop control of a virtual or prosthetic limb during natural reaching movements has not yet been demonstrated. Additional studies are required to fully understand the encoding of arm and hand movements in field potentials. Such efforts should include theoretical modeling work to determine how single-neuron tuning properties are reflected in population-level ECoG activity to predict which spatial scales, if any, tuning to extrinsic movement might be expected at. It may be possible that by increasing the spatial sampling of cortex through the use of high-density electrode grids [Viventi et al., 2011] combined with the use of decoding algorithms tailored to high-dimensional field potential signals, such as that presented in Chapter 7, some degree of ECoG-based BMI control using a naturalistic control strategy can be achieved.

8.2 THE FUTURE OF ELECTROCORTICOGRAPHIC BRAIN-MACHINE INTERFACES

Ultimately, the work presented in this dissertation provides an important demonstration of the potential of ECoG to support brain-machine interfaces for functional restoration for individuals with movement paralysis, and provides a foundation for additional research seeking to fully translate ECoG BMI technology to the clinical realm. At this point in time, however, it is still unclear whether such a system will ever truly be realized, as a number of significant scientific and technological hurdles must first be overcome. Considering the success of recent demonstrations of high-dimensional control using SU/MUA recordings, the case against ECoG-based BMI systems appear quite strong: the encoding of movement parameters in ECoG is extremely coarse compared to that of single and multi-unit activity, ECoG signals appear lack the sufficient spatial resolution and independence to support high degree of freedom control, and ECoG-based control currently relies heavily upon abstract control strategies which likely impose a heavy cognitive burden on users.

However, based on the work presented here as well as previous investigation into ECoG-based BMI systems, we can speculate as to what the ideal ECoG brain-machine interface might look like. High-density electrode grids, utilizing surface treatments to reduce the foreign body response and prolong electrode life, will enable the recording of field potentials from the surface of the cortex with high spatial resolution over long periods of time. Though the use of a somatotopic control strategy and training procedures, individuals with movement disorders will gradually learn to independently modulate the activity of small electrode subsets to achieve high-dimensional control, with control becoming increasingly internalized and natural over the course of training. Sophisticated co-adaptive decoding algorithms, specifically designed to handle the high-dimensional and correlated nature of electrical field potential signals, will facilitate the process, resulting in a system that successfully provides functional restoration, is stable and robust, and easy to use.

While it may seem unlikely to some that such a system will ever come to pass, it is important to consider all BMI research as part of a much larger enterprise in the development of assistive and restorative technology. At this point in time, we can only guess as to what a clinical-viable BMI system might look like. It is not unreasonable to expect that systems may vary widely in design depending on the nature of the movement disorder and the restoration of function desired, potentially incorporating signals from multiple cortical recording modalities. Though much work still remains, the collective efforts of researchers, clinicians, and subject volunteers continue to move us closer to the realization of a brain-machine interface system which will make a profound impact in the lives of individuals with movement disorders.

APPENDIX A

CRANIUX: A LABVIEW-BASED MODULAR SOFTWARE FRAMEWORK FOR BRAIN-MACHINE INTERFACE RESEARCH[1]

A.1 INTRODUCTION

Brain-machine interface (BMI) technology aims to establish a direct link for transmitting information between the brain and external devices. It offers a rich and natural assistive device control interface for individuals with disabilities [Schwartz et al., 2006, Wang et al., 2010], and is a rapidly-progressing, extremely active research area in the field of neuroscience and neural engineering. Various neural signal modalities, including electroencephalography (EEG) [McFarland et al., 2010], magnetoencephalography (MEG) [Mellinger et al., 2007], electrocorticography (ECoG) [Schalk et al., 2008c], intracortical local field potentials (LFPs) [Heldman et al., 2006], and neuronal firing rates [Hochberg et al., 2006, Velliste et al., 2008, Ganguly and Carmena, 2009], have been used for BMI research. Regardless of the input modality, all BMI systems require an essential suite of software capable of acquiring neural signals continuously and converting them in real-time or near real-time into specific BMI control commands for an external device, such as a prosthetic hand, in order to accomplish a specific task.

To conduct innovative and unique BMI studies, researchers very often need to implement new signal processing techniques, neural decoding algorithms, or experimental paradigms in a BMI software package. Given the rapid progression of the field, it is desirable to reduce the time it takes from the conception of a new idea to software implementation, data collection, and data analysis. However, the increasing complexity of BMI systems has made this problematic. For example, sophisticated neural decoding algorithms previously studied in offline analysis are now being investigated for real-time BMI control [Koyama et al., 2010]. Additionally, more advanced external devices are being controlled by BMI systems, such as the dexterous prosthetic arm and hand system developed by the Revolutionizing Prosthetics project [Adee, 2008, Adee, 2009]. These advancements call for an open-source software framework that enables BMI researchers to better focus on the essential engineering and scientific questions they are investigating and to develop advanced

---

[1] The contents of this appendix have been published [Degenhart et al., 2011b] and have been reproduced here as the version accepted for publication.

BMI features more efficiently. This framework should be able to manage the basic software operations common to many BMI studies, and should be easily extendable in a high-level programming environment that offers the ease and flexibility for programming new BMI modules.

One successful open-source general purpose BMI software package is BCI2000, a modular C++-based system for neural signal acquisition, data saving, stimulus presentation, and more [Schalk et al., 2004], which has been widely distributed among academic institutions and used in numerous research studies [Schalk, 2014]. Source code as well as binary executable files are freely available for download, allowing end users to either use the software as-is or modify it to suit their own needs. One of the greatest advantages of BCI2000 is its modular, lightweight, and portable design, making it extremely popular and successful in the BMI research community. Recently, the BCPy2000 open-source framework [Hill, 2008] has been made available as a user contribution package to BCI2000. This framework follows the same system architecture as BCI2000, but it allows BMI researchers to develop new modules in Python, a high-level language that greatly reduces software programming complexity for fast prototyping of new BMI software.

This paper presents an open-source open-access real-time BMI software framework inspired by BCI2000 termed "Craniux," developed using LabVIEW (National Instruments, Inc.), a high-level multi-platform graphical development environment. Craniux implements a core framework for BMI operation, including modular architecture, network communications between modules, data flow control, data visualization, data storage, and graphical user interfaces. Craniux offers a unique set of advantages that can greatly facilitate BMI software development and research. First, it enables BMI researchers to develop and share new BMI modules in the LabVIEW development environment and take full advantage of many features inherent to this environment, such as:

- High-level graphical programming for fast development and run-time debugging.
- A rich set of data visualization options and graphical user interface elements.
- Ease of multi-threading and parallel processing programming, including automatic parallelism and multi-core processor support.
- A large number of high-quality LabVIEW function libraries for signal processing and streamlined integration with a wide range of engineering hardware (e.g. National Instruments controller cards).
- Re-use and sharing of custom-made LabVIEW modules as sub-VI (virtual instrument) blocks.

Second, facilitated by the above-mentioned LabVIEW features, we have further implemented functionality critical for BMI research:

- Real-time operation in which the system is capable of acquiring a block of neural data, processing this data, and generating an output before the next block of data is received [Schalk et al., 2004, Wilson et al., 2010].
- Online neural decoder training capability accomplished through data sharing between real-time operations and parallel decoder training.

- 'On-the-fly' data visualization and online experiment parameter control.
- Deterministic control of system execution, including parameter updates and display of visualization data.
- Streaming and storage of raw neural data, various intermediate processing data, and experimental parameters to disk for offline analysis.
- Distribution of BMI modules across computers using well-defined generic network communication protocols optimized for data transmission between software modules.

Finally, Craniux has been developed to be a lightweight, extendable, and portable software framework. Its modular architecture, well-defined user interfaces, and generic network communication protocol make it very easy to maintain and develop BMI engines. The existing engines and standard template engines provide a starting point for new engine development.

In the following sections, we will first introduce the basic system architecture of the Craniux software. We will then provide system performance testing results based on both simulated and real experimental data. The last section will further discuss the uniqueness of this software framework as compared to other existing BMI software tools, its advantages and limitations, and future directions.

A.2 SYSTEM ARCHITECTURE

The Craniux software package has been designed to be a highly modularized system, capable of operating across both a distributed network of computers and on a single computer. To accomplish this, and to make data transfer between engines as reliable as possible, all data communication is conducted using the TCP/IP protocol. Data saving is implemented using the LabVIEW TDMS (Technical Data Management Streaming) framework [National Instruments, 2014], ensuring all system data are streamed to disk as quickly as possible in order to maximize system performance. The following sections describe the system framework, engine execution, GUI operation, communication protocols, and data saving operation in further detail.

A.2.1 Distributed engine framework

Figure A1 depicts the design of the Craniux system. Inspired by the BCI2000 framework, this system consists of five distinct components: the System Launcher, Acquisition engine, Signal Processing engine, Application engine, and data saving manager, and may be distributed across as many as four computers. Furthermore, each engine has an associated graphical user interface (GUI), through which the user interacts with the engine. The main system components perform the following functions:

System Launcher. The System Launcher is the initial interface the user is presented with when running the software, and allows the user to specify system-level parameters at runtime. It is here that the specific engines, their network locations, and high-level experimental parameters (e.g. subject ID, date, investigators, and session number) are specified. Additionally, the System Launcher controls the start and stop of execution, though it itself is not a part of the real-time operation of the system.

Acquisition engine. Acquisition engines are responsible for the acquisition and initial pre-processing (e.g. spectral estimation) of neural data from some signal source such as an amplifier or user datagram protocol (UDP) connection.

Signal Processing engine. Signal Processing engines receive data from the Acquisition engine and are responsible for the processing of this data, such as the generation of a control signal.

Application engine. Application engines receive data from the Signal processing engine and are responsible for the control of interaction between the subject and the BMI.

Data saving manager. The data saving manager is responsible for the saving of Craniux data and receives input from the Acquisition, Signal Processing, and Application engines.

In order to ensure sequential processing of data through each engine, system execution proceeds from the Acquisition engine to the Signal Processing engine, then to the Application engine, and finally from the Application engine back to the Acquisition engine; only one of each type of engine may be running at a given time. This cyclical data flow guarantees that each block of data received by the Acquisition engine is processed and a system output is generated before processing of the next block of data begins.

At any point in operation, the system may be suspended and any of the engines replaced with another of the same type, preserving the state of those engines that remain running. This is desirable for BMI operation, as system parameters such as neural decoder weights obtained during operation with a specific application (e.g. a center-out computer cursor task) may be retained and immediately used for a new application (e.g. the control of a robotic arm). Table 1 provides a list of the current engines available in the Craniux system.

A.2.2 Engine execution

Each engine in the system operates in a basic sequence, first receiving data from a previous engine, processing the received data, and sending the relevant results of that processing to the next engine in the signal chain. Figure A2 outlines the basic flow of the execution of an individual engine. Engine execution first begins with the initialization of all parameters, including the loading of user-specified parameter files, and the identification of those engine-specific parameters to be saved. From here, execution proceeds into the main sequence of the engine, where the engine: (1) waits for data from the previous engine in the signal chain, (2) processes the received data (or performs some other action), and (3) sends the results of this processing to the next engine in the signal chain. Execution then proceeds back to (1), where the engine waits for the next block of input data. Operating in parallel to this main sequence are a number of additional threads, such as data saving, engine-specific processes not capable of or not requiring real-time operation (e.g. neural decoder training), and communication with the engine's GUI. A detailed description of engine

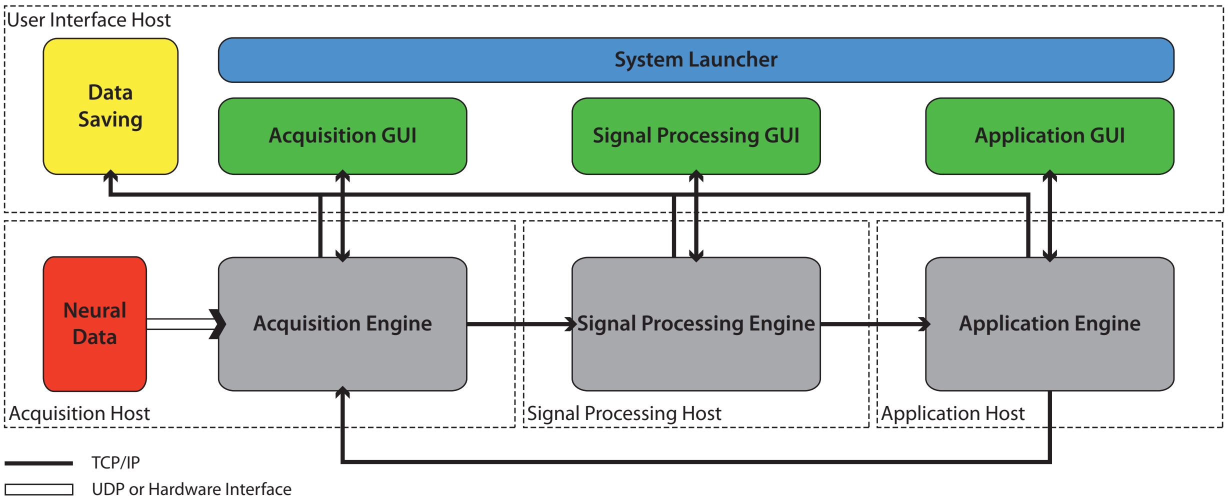

Figure A1: *Craniux system framework.* The Craniux system is comprised of the Acquisition, Signal Processing, and Application engines, their associated GUIs, the System Launcher, and the data saving manager. Engines and user interface elements are spread across four network hosts: the Acquisition Host, the Signal Processing Host, the Application Host, and the User Interface Host, though the same computer may serve as multiple hosts. Network communication between system engines, as well as communication between engines and GUIs, is performed using the TCP/IP protocol. A block of neural data enters the system through the Acquisition engine, which sends preprocessed data to the Signal Processing engine. The Signal Processing engine generates a control signal, which is then sent to the Application engine. The Application engine then communicates any relevant application-specific data (e.g. target information used for neural decoder training) back to the Acquisition engine, which reads the next block of neural data. Bidirectional data transfer occurs between engine-specific GUIs and their associated engines, with system parameters transferred from the GUI to the engine and visualization data transferred from the engine to the GUI. Finally, the System Launcher is responsible for loading the desired engines, tracking general experimental parameters, and experimental control.

Table A1: List of current Acquisition, Signal Processing, and Application engines.

Module Name	**Typical Location**	**Description**
gUSBamp	Acquisition	Reads signals from the gUSBamp amplification system.
SimECoG	Acquisition	Generates simulated ECoG signals modulated by the computer cursor position.
SignalFiltering	Chain	Performs filtering (band-pass, notch) of time-domain signals, as well as visualization.
SpectralEstimation	Chain	Calculates the power spectra of input signals.
MATLABDecoder	Ring	Performs linear decoding of input signals, as well as buffering of feedback data and decoder training in MATLAB.
StimulusPresentation	Application	Controls sequencing and display of text, image, and video stimuli.
CursorTask	Application	StateMachine-based cursor control task in 3-dimensional space.
VirtualArmControl	Application	StateMachine-based control of a virtual arm.
DEKAArmControl	Application	StateMachine-based control of the DEKA prosthetic arm, including presentation of targets using the Automatic Target Presentation (ATP) system.

execution, including the enforcement of deterministic execution within engine components, is provided in the Supplemental Materials.

A.2.3 Graphical user interface (GUI) elements

The GUI for each engine is responsible for both on-the-fly control of engine-specific parameters as well as the visualization of engine-specific data. Permitting on-the-fly control is essential to successful BMI operation, as during real-time closed-loop BMI operation it is often necessary to dynamically adjust parameters such as the computer assist level or computer cursor speed [Velliste et al., 2008]. As opposed to a traditional graphical user interface, which simply serves as a front end user interface for a LabVIEW application, GUIs in the Craniux system exist as stand-alone applications. It is through these applications that the user interacts with each engine. GUIs and their associated engines maintain reciprocal two-way communication; parameter value change events are monitored by the GUI and transmitted to its associated engine via TCP/IP, while data to be visualized is transmitted from the engine to the GUI. It should be noted that parameter value changes are instantaneously transmitted from the GUI to the engine and are accessed by the engine at the beginning of its main sequence in order to ensure the consistency of all parameter values throughout the processing of a single block of data. Parameter value changes are also index-stamped and saved to disk, allowing the complete reconstruction or replay of the full system state during offline analysis. In order to allow for data visualization on-the-fly, all data elements are transmitted from the engine to the GUI, providing the experimenter with the most accurate representation of the state of the engine. This occurs in parallel with the real-time main sequence execution so that this communication does not interfere with the timing of the execution of the main sequence of the engine.

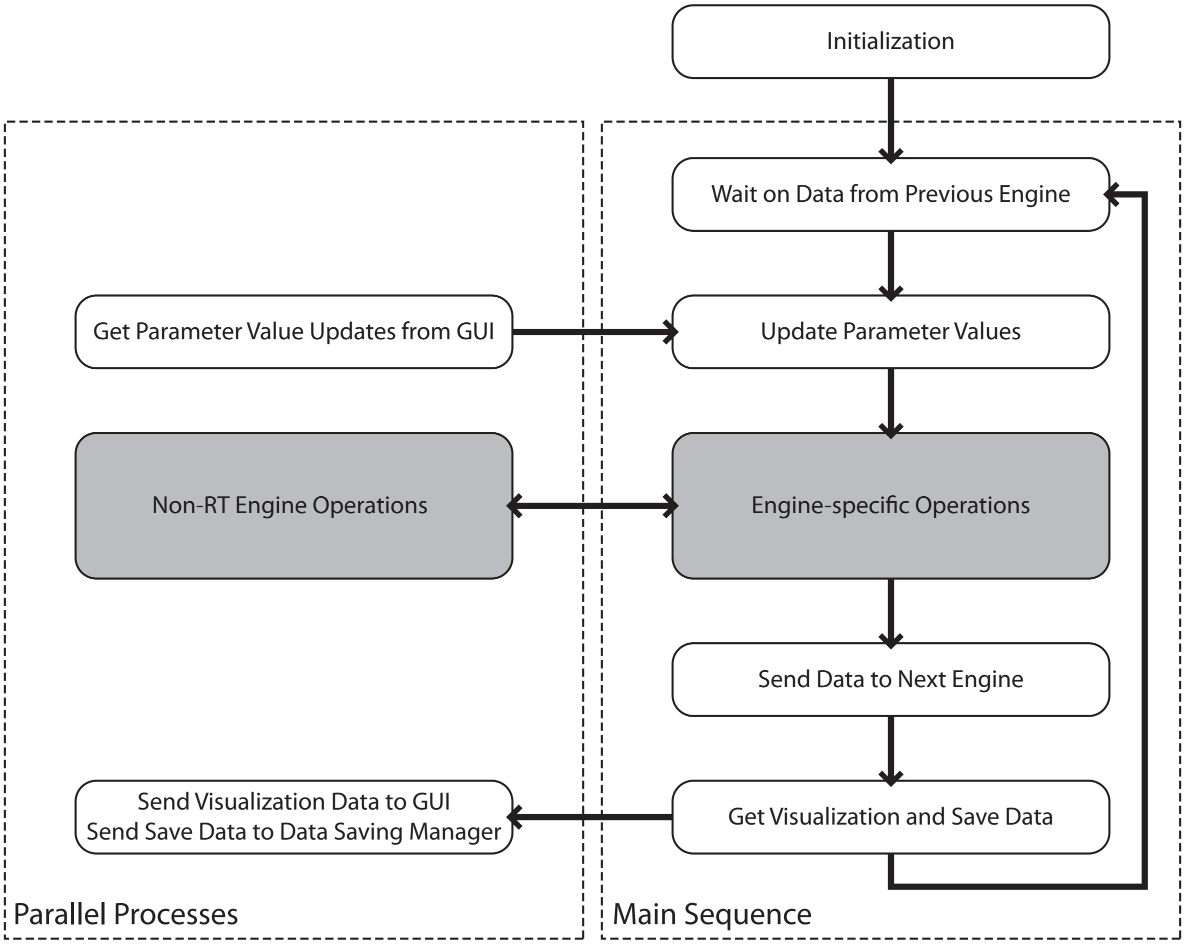

Figure A2: *Engine execution.* After initialization, each engine proceeds into the 'main sequence' loop, in which core engine processes are executed sequentially. Data are first received from the previous engine in the signal chain, and any parameter value changes received from the engine's GUI are updated locally. The system next performs any actions specific to the individual engine (e.g. calculation of a control signal or updating of a display), and sends the results of these actions to the next engine in the signal chain. Current values of any data items to be visualized are placed in a queue, and data is sent to the data saving manager. The engine then proceeds to the beginning of the main sequence loop to await the arrival of the next input. Parallel to the main sequence loop are any parallel processes designed to operate asynchronously. These processes will always include receiving parameter value updates from the GUI and sending visualization data to the GUI, and may include individual engine-specific operation such as decoder training or monitoring for events. Shaded blocks represent those areas to be modified by the developer during the creation of new engines, while white blocks represent sections of code providing core functionality.

A.2.4 Communication between components

Communication between Craniux components utilizes self-establishing and self-repairing network connections that provide efficient, reliable data flow robust to any data type or combination of variables that is sent over them. For communication between engines, these connections take the form of a ring that maintains data flow and controls program execution. For communication between engines and their GUIs or the data saving manager a single TCP connection is established. When creating new engines, this ring requires no input and single connections only require the developer to provide a network host name. The only user input necessary is the IP address of each engine, which is specified on the System Launcher. Available ports are automatically selected for each connection.

All network connections use the TCP protocol. TCP was chosen over UDP because its superior reliability is important in a ring structure responsible for the control of program execution; a dropped packet between engines would break the ring and leave each engine waiting for data that will never arrive. It is also important to note that Nagle's algorithm [Nagel, 1984] was disabled for all connections used in the Craniux system. The Nagle algorithm attempts to reduce TCP packet overhead and bandwidth usage by intentionally delaying transmission so that multiple packets can be combined before being sent. Here, the latency introduced by this algorithm is unacceptable and bandwidth usage is not a concern. The concept behind Nagle's algorithm is retained in our system, however, as all data to be sent simultaneously is combined into a single packet before transmission.

To send variables over the network, the developer must only create a list of the variable names on the sending side of the connection. No information on variable type or size is needed. The provided variable names are packed together with their values into a single variable of LabVIEW's 'variant' data type, which is then sent over the network. On the receiving side of the connection, the data is read and parsed into the correct values, which are written to those existing variables on the receiving side with the same name and data type as the sent variables. Additional information on the transfer of information between components has been provided in the Supplemental Materials.

A.2.5 Data saving

The Craniux framework for saving data is a reliable process that minimizes latency introduced by saving and creates highly accessible data. Data saving is conducted by an independent data saving manager, which receives data from all engines. This data is initially saved in LabVIEW's TDMS format, which was specifically created for quickly and continuously streaming large amounts of data to the hard drive to help eliminate data-saving bottlenecks in speed normally introduced by slow writes to disk [National Instruments, 2014]. When saving data, a packet containing all the variable values to be saved and the data packet index is placed into a first-in-first-out (FIFO) buffer. Parallel to the main execution of Craniux, these packets are removed from the buffer and sent to the data saving manager, located on the user interface host, via the communications framework described in Section A.2.4. Upon receiving a packet, the data saving manager streams the data to a TDMS file. When creating new engines it is only necessary to provide a list of variable names to be saved; these items will be automatically identified and their values saved accordingly.

A single TDMS file is saved for each experimental run; stopping or suspending system execution closes all references to the current data file. Within each data file, data saved by each engine is separated into two groups: sampled variables and controls. Sampled variables are data sampled continuously at each update of the BMI system, such as cursor position during a brain-controlled cursor movement task. As controls are normally parameter settings that are infrequently updated (e.g. the number of targets), these values are only saved when changed. The current data packet number is included in every save operation so that the experiment can be reconstructed afterwards with the data properly aligned in time. A separate LabVIEW VI has been created to convert Craniux TDMS files into the MATLAB (Mathworks, Inc.) MAT format. These MAT files contain structures for each engine paired with each data type (sampled variables and controls). The saved values for each variable are stored in an array, with the data packet number array providing the time index for each element. Array variables are stored as cell arrays, allowing them to be aligned with their associated data packet numbers and enabling the data structure to handle dynamic changing of array sizes during a BMI session.

A.3 SYSTEM VALIDATION

A.3.1 Closed-loop cursor movement control using simulated ECoG signals

Validation of the operation of the Craniux system was initially conducted using an electrocorticographic (ECoG) signal simulator in which experimenter-controlled mouse cursor movement was used to modulate the high gamma band activity of a number of synthetic signals. This simulator is capable of generating 32 channels of analog signals with directionally tuned high-gamma band (70-120Hz) activity emulating ECoG signals recorded from human subjects [Leuthardt et al., 2004, Wang et al., 2010] according to Equation A.1 [Georgopoulos et al., 1986, Heldman et al., 2006]:

$$S = S_1 + d\cos(\theta)S_2 \tag{A.1}$$

where S is a single simulated directionally modulated ECoG signal, S_1 is a pink noise signal with a $1/frequency$ power falloff [Keshner, 1982]. S_2 is a second pink noise signal band-pass filtered between 70-120Hz, d controls the depth of modulation of the high-gamma band, and θ is the angle between the preferred direction of the simulated ECoG signal and the vector pointing from the center of the computer screen to the current mouse cursor position on the computer screen. The preferred directions of the 32 simulated signals were uniformly distributed over two-dimensional (2D) space. Simulated signals were generated at 2400Hz using a National Instruments NI PCI-6723 32 channel analog output board on a simulation computer

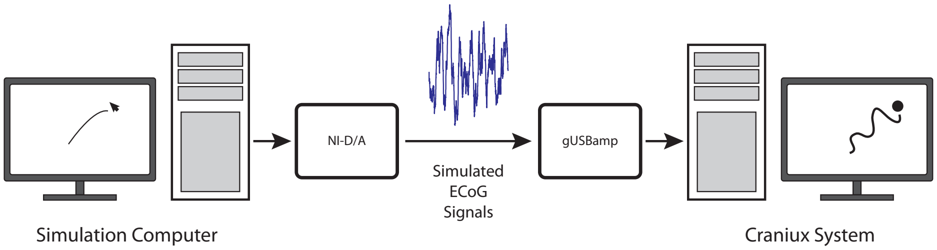

Figure A3: *Simulated ECoG experimental setup.* Experimenter-controlled mouse position on the simulation computer modulates the high-gamma power of simulated directionally tuned ECoG signals. These signals are output at 2400Hz using a National Instruments D/A card and are read into the Craniux system using the g.USBamp amplification system and BCI2000. The Craniux system then decodes the desired cursor position from the simulated signals using the Population Vector Algorithm.

(Windows XP x86 operating system, AMD Athlon 64 FX-62 Dual Core CPU @ 2.81 GHz, 3.5 GB RAM, NVIDIA GeForce 7900 GS video card) and then stepped down to match the amplitude of typical ECoG signals recorded from human subjects.

Simulated ECoG signals were then sampled at 1200Hz using the g.USBamp amplification system (Guger Technologies, OEG) on a separate computer (Windows XP x86 operating system, Intel Core i7 CPU 920 @ 2.67 GHz, 2.49 GB RAM, 2 NVIDIA GeForce 9800 GT video cards) and sent to the Craniux system as binary UDP packets using a simplified version of the BCI2000 software package. BCI2000 was used in this case due to its reliability and efficiency in interfacing with the g.USBamp amplification system. These raw time-domain signals entering the Craniux system were first converted into the frequency domain using LabVIEW's built-in autoregressive (AR) spectral estimation function (10Hz bins, 500ms window) in the Read UDP Binary Acquisition engine and then passed to the Population Vector Signal Processing engine. Here, signals were normalized to pseudo Z-scores based on Equation A.2 [Tallon-Baudry et al., 2005, Edwards et al., 2009]:

$$f_{norm,i,j} = \frac{f_{i,j} - \bar{f}_{i,j}}{\sigma_{i,j}} \tag{A.2}$$

where $f_{norm,i,j}$, $f_{i,j}$, and $\bar{f}_{i,j}$ are the normalized, raw, and mean power of the ith channel and jth frequency band, respectively, and $\sigma_{i,j}$ is the standard deviation of the raw band power of the ith channel and jth frequency band. Mean and standard deviation values were calculated based on data collected during a baseline condition in which the computer cursor on the simulation computer remained in the center of the screen (i.e. no modulation of high-gamma band activity).

The brain control task used was a typical 2-D center-out design, with the movement direction of a cursor controlled by multiple ECoG signal features across 32 channels according to the Population Vector Algorithm [Georgopoulos et al., 1986]:

$$f_i = b_{0,i} + b_{x,i}m_x + b_{y,i}m_y \tag{A.3}$$

$$P = \sum_{i}^{N}(d_i - b_{0,i})\mathbf{C}_i \tag{A.4}$$

where f_i is the activity of individual feature i, m_x and m_y are the desired movement in the x and y direction, $b_{0,i}$, $b_{x,i}$ and $b_{y,i}$ are coefficients found using linear regression relating desired movement to the activity of feature i, P_i is the trajectory vector predicted by the activity of feature i, d_i is the instantaneous activity of feature i, and $\mathbf{C}_i = [b_{x,i}\ b_{y,i}]/(b_{x,i}^2 + b_{y,i}^2)^{1/2}$ is a vector representing the preferred direction of feature i.

The standard workflow used to achieve ECoG-controlled 2D cursor movement with the Craniux framework is described below. Though simulated ECoG signals were used here to validate the system, this workflow will be similar for real neural signals.

1. *Collection of baseline data.* Once the Craniux system is started, approximately 3 minutes of baseline data is collected, from which the Craniux system will calculate feature mean and standard deviation values. These will then be used in the calculation of pseudo-Z scores for all ECoG signal features in real-time.
2. *Collection of training data for the neural decoder.* During this period, the experimenter will use the ECoG signal simulator to generate modulated ECoG signals based on the target position (i.e. desired cursor movement direction). The ECoG data along with target position are automatically buffered by Craniux for neural decoder training.
3. *Training of the neural decoder.* During this period, the buffered data is used to train the neural decoder. A multiple linear regression procedure is used to determine the degree of directional tuning and preferred direction for each ECoG signal feature as mentioned above [Schwartz et al., 1988, Wang et al., 2007]. The resulting R-squared values and preferred directions are displayed by the Population Vector GUI, allowing experimenters to visualize the results on-the-fly and interactively select a subset of directionally tuned ECoG signal features for brain control.
4. *Real-time brain control.* Activities of ECoG signal features selected during step (3) are then used to generate the population vector, a 2D velocity control signal that drives the cursor. Figure A4 shows the Population Vector GUI during closed-loop brain control, illustrating the user interface elements provided to the user during this process.

It is worth noting that all the above procedures are conducted in a continuous BMI session without stopping and restarting the Craniux system. This streamlined workflow allows BMI studies to be conducted smoothly and efficiently. Furthermore, steps (2) and (3) can be conducted at any time during a BMI session in parallel with step (4). This allows the neural decoder to be re-calibrated on-the-fly to adapt to any

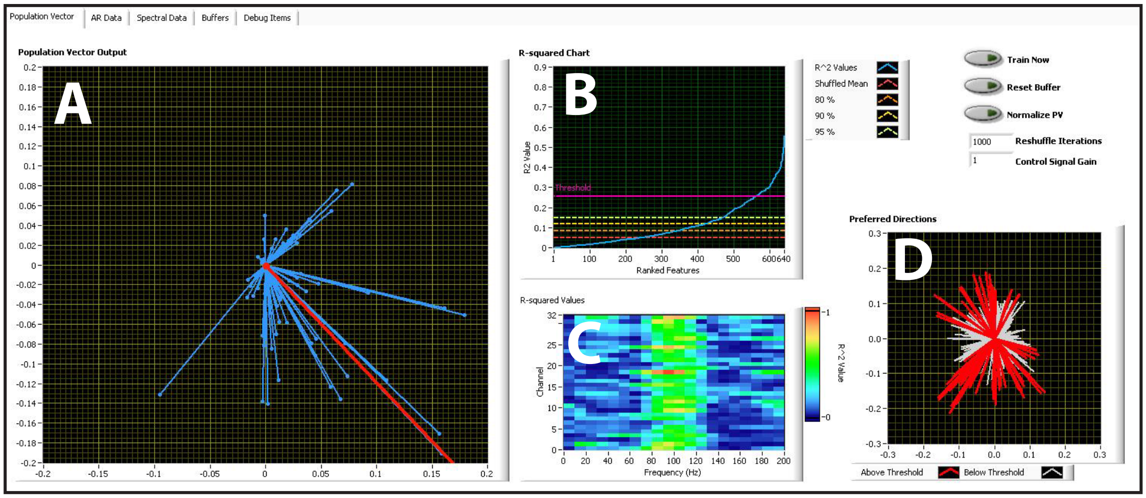

1000 Reshuffle Iterations

1 Control Signal Gain

Figure A4: *Craniux system screenshot during Population Vector-based control. A.* Plot of the instantaneous activity of each feature used for cursor control along its preferred direction (blue) and the resultant population vector (red). *B.* R^2-value plot indicating the distribution of R^2 values obtained during Population Vector training (blue) compared to the mean, 80th, 90th, and 95th percentile R^2 values obtained after training on 1000 iterations of target-shuffled data (red, dark orange, light orange, and yellow lines). The threshold above which features are chosen for use in the decoder is shown by the pink line. *C.* R^2 values obtained during Population Vector training arranged by channel and frequency band. Note that the 70-120Hz frequency band features show high R^2 values across all channels, consistent with the method used to generate the simulated ECoG signals. *D.* The preferred direction distribution of all features. Red lines correspond to those features with R^2 values above the user-determined threshold, while white lines are those features falling below the threshold.

potential changes or non-stationarities of input neural signals, a key element for achieving and maintaining reliable brain control [Taylor et al., 2002]. Figure A5 shows an example of directionally modulated normalized time-frequency data for one ECoG signal saved by the Craniux system, as well as trajectories of the cursor during real-time brain control.

A.3.2 Brain-controlled cursor movement using real ECoG signals recorded from a human subject

Further validation of the Craniux system was conducted in a human subject undergoing subdural epilepsy monitoring. Informed consent was obtained from the subject prior to testing; all experimental procedures were approved by the University of Pittsburgh Institutional Review Board and followed all guidelines for human subject research. Experimental methods used were similar to those presented in [Wang et al., 2009], with the exception that the Craniux system was used for data collection and brain control. Standard ECoG electrodes exhibiting high-gamma band modulation in response to overt movement screening tasks were chosen for use in closed-loop control. High-gamma band power (70-110Hz) of two neighboring ECoG electrodes were used to control the vertical movement of a cursor with a push-pull scheme, with the cursor control signal calculated according to Equation A.5:

$$c_y = a(s_1 - s_2) - b \tag{A.5}$$

Where c_y is the one-dimensional control signal, s_1 and s_2 are the high-gamma band power of the two neighboring electrodes used for control, and a and b are gain and offset terms used to normalize the control signal to zero-mean and unit-variance. Thus, in order to achieve satisfactory brain control, the subject had to de-correlate the activity of the two electrodes to generate the desired cursor control signal. Brain control sessions began with the collection of baseline data for normalization purposes as described in the previous section. Individual trials began with the placement of the cursor at the center of the computer screen along with the presentation of one of two peripheral targets located in the vertical plane of the workspace (e.g. a "center-out" task). Trials in which the subject was able to hit the presented target within the maximum trial length of 10 seconds were deemed successful; failure to do so resulted in an unsuccessful trial. All trials were followed by an inter-trial interval of 2 seconds in which neither the cursor nor the target were visible. Figure A6 shows the results of one brain control session, during which the subject was able to achieve an 88% success rate.

A.3.3 System timing

To evaluate the consistency of system performance, timing characteristics were analyzed for a typical Craniux setup using 15th order autoregressive (AR) spectral estimation of 10 Hz frequency bins over 0.5 second windows of simulated neural data (see Section A.3.1), the Linear Decoder Signal Processing engine, and the

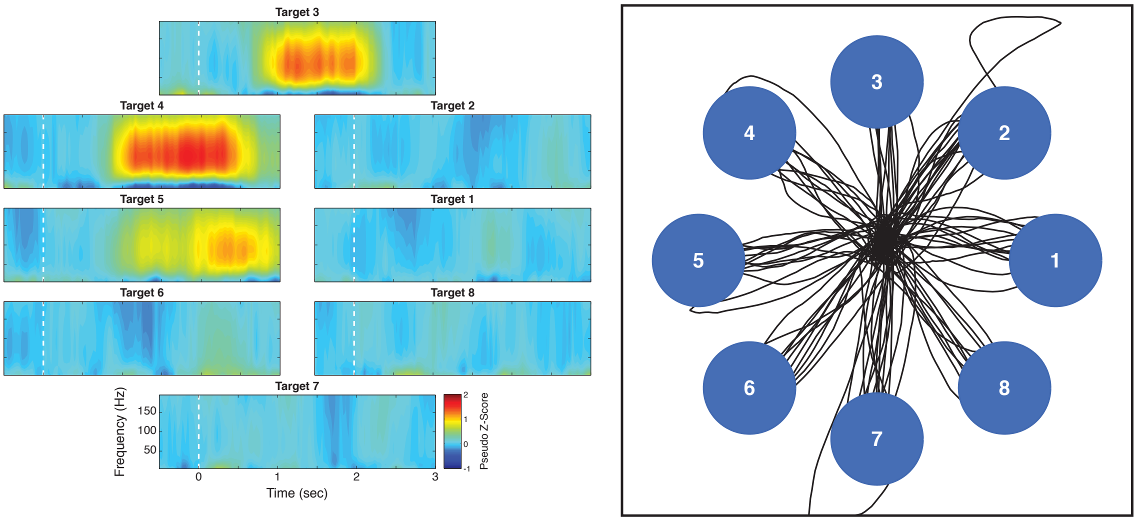

Figure A5: *Closed-loop brain control using simulated ECoG signals. Left.* Time-frequency plots of a single simulated ECoG signal averaged across all repetitions of an 8-target center-out cursor control task. Plots are aligned to target presentation at time $t = 0$ (dashed white line). In all, a total of 32 channels of simulated directionally tuned ECoG signals were generated. *Right.* Real-time cursor trajectories controlled by simulated ECoG signals using the Population Vector Algorithm.

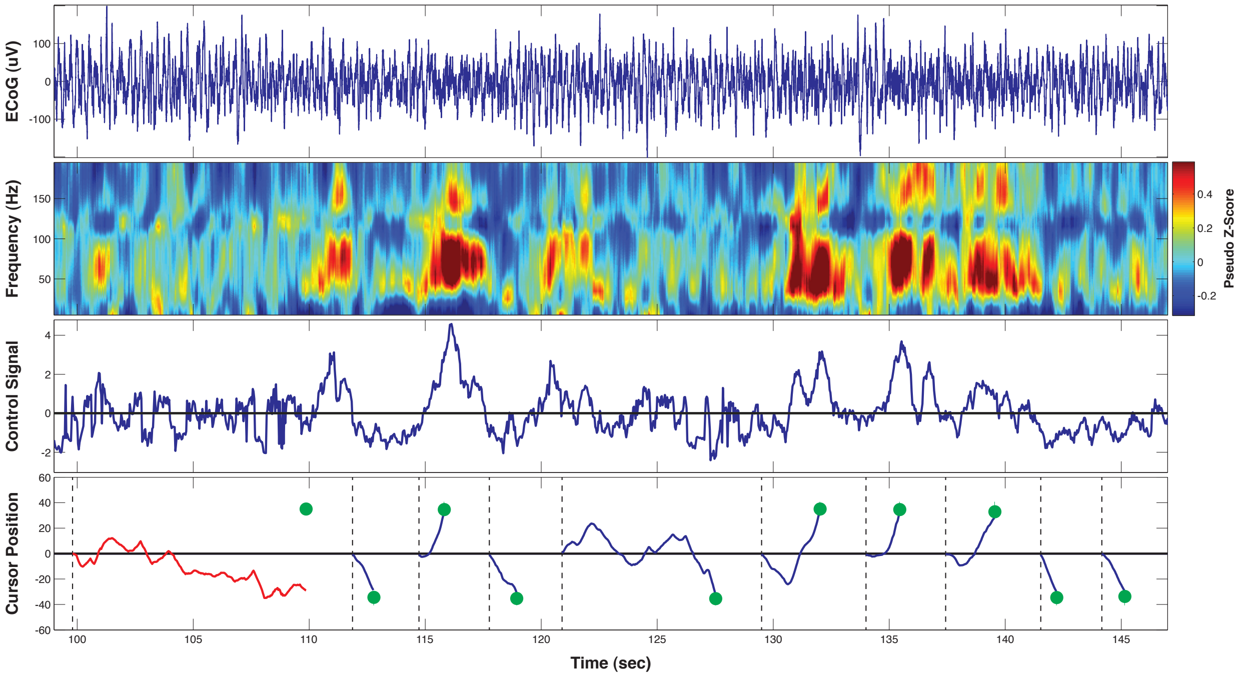

Figure A6: *Closed loop ECoG-based computer cursor control.* One-dimensional computer cursor control using the Craniux system in a subject implanted with ECoG electrodes. *Top.* Raw time-domain ECoG signal for one of two electrodes used for cursor control. *Top-middle.* Time-frequency data saved by the Craniux system for the same electrode. *Bottom-middle.* Control signal generated by the Craniux system to control computer cursor movement. Positive control signal values move the cursor in the up direction, while negative control signal values will move the computer cursor down. Note that control signal values are unitless as they have been normalized to zero-mean and unit variance. *Bottom.* Vertical cursor positions generated by the neural control signal. Dashed black lines represent target onset, green circles indicate the position of presented targets, blue lines indicate cursor trajectories for successful trials, and red lines indicate cursor trajectories for unsuccessful trials.

Center-out Cursor Control Application engine. Analog and digital data were sampled by the g.USBamp amplification system at 1200 Hz and acquired directly by Craniux at a 33.3 ms frame rate. To trigger timing events, a digital signal was sent from Craniux back to the digital input of the amplifiers so that the timing events could be acquired precisely and synchronously with the raw analog input signal at 1200 Hz.

Three different timing tests were conducted: a system processing test, a display update rate test, and an overall system latency test. The first two tests (system processing and display update rate) were performed on both a single computer (Windows XP x86 operating system, Intel Core i7 CPU 920 @ 2.67 GHz, 2.49 GB RAM, 2 NVIDIA GeForce 9800 GT video cards) and with Craniux distributed across the network so that data acquisition, spectral estimation, and GUIs were hosted on one computer (the same as that used for local timing test, see above) while Signal Processing and Application engines, including the 3D render window, were hosted on a separate computer (Windows XP x86 operating system, AMD Athlon 64 FX-62 Dual Core CPU @ 2.81 GHz, 3.5 GB RAM, NVIDIA GeForce 7900 GS video card). For both configurations tests were conducted using 16, 32, and 64 channels of data. The third test (system latency) was run only on the single-computer configuration with processing performed on 32 channels of data.

The first test used 5,000 consecutive frames of collected data to measure the system processing time, the time between the arrival of a block of data from the amplifier and the time when the Craniux system had finished all processing on the data and begun waiting on the next block. These results are shown in the second column of Table 2. As expected, processing time was found to increase with the number of processed channels, but remained below the 33.3 ms time required to maintain a consistent frame rate and prevent the loss of data. Distributing Craniux across the network showed improvements in processing time for all channel configurations. Since processing time is only required to remain below the frame rate, running Craniux as a distributed system is not necessary unless the system is under a heavy load. AR spectral estimation was found to require the most processing time, especially as the number of channels increased. These results indicate the extra processing time made available when Craniux is run as a distributed system could easily be utilized to run more complex signal processing algorithms or to decrease the frame rate.

The second test also used 5,000 consecutive frames of data, but now measured the refresh rate, the amount of time between consecutive display updates on the Center-out Cursor Control engine. The results are shown in the third column of Table 2. The refresh time was found to be 33.3 ms for all configurations, precisely what would be expected given the system frame rate. Furthermore, the low variability of this timing indicates that the user would experience a consistent cursor update with no noticeable jitter.

The final test measured system latency, the elapsed time between a neural signal event and the point in time when the Craniux system can generate an action in response to this event. A 10 Hz sine wave with zero offset was input to 1 channel of the amplifier; this channel of data was fed through the Craniux system to the point at which the display was updated in the Center-out Cursor Control engine, occurring just before processing fully completes and the system begins waiting on the next data block. At this point, if a zero-crossing was detected on the sine wave, the digital output bit being written back to the amplifier was

Table A2: *Characterization of system timing.* Craniux system processing time, refresh rate, and latency for local and network system configurations under various processing loads. Values shown are mean timing values plus or minus one standard deviation from the mean.

System Configuration	**Processing Time (ms)**	**Refresh Rate (ms)**	**Latency (ms)**
Local, 16 channels	12.8 ± 0.7	33.3 ± 0.5	N/A
Network, 16 channels	9.9 ± 0.6	33.3 ± 0.5	N/A
Local, 32 channels	17.8 ± 0.8	33.3 ± 0.7	33.2 ± 9.6
Network, 32 channels	15.2 ± 1.0	33.3 ± 0.4	N/A
Local, 64 channels	28.0 ± 1.3	33.3 ± 0.7	N/A
Network, 64 channels	24.9 ± 0.9	33.3 ± 0.4	N/A

flipped. In this case, the elapsed time between a zero-crossing of the sine wave (a simulated neural event) and the bit value change (the time of the system response) indicates the system latency. Data was collected for 5,000 consecutive sine wave zero-crossings, with zero-crossing events symmetrically distributed about the center of each 33.3 ms data frame. In distributing the zero-crossings in this way it is known that the latency should have an average of slightly less than half the frame length plus the mean processing time (33.9 ms for this configuration) and a range nearly equal to the frame length. The latency was found to have a mean and standard deviation of 33.2 ± 9.6 ms, meeting all expectations.

A.4 DISCUSSION

Craniux is a powerful, yet simple and easily extendable, open-source framework for BMI studies that require high-performance real-time BMI software. Currently, a number of open-source software solutions for BMI research are available for academic use. These software packages include extremely specialized, high-overhead systems used in non-human primate BMI research [Taylor et al., 2002, Bacher et al., 2008, Velliste et al., 2008], highly-modular, visual-programming based software platforms such as OpenViBE [Renard et al., 2010], as well as portable, lightweight systems for human BMI research [Schalk et al., 2004]. Software tools for more specific BMI research applications have also been made available, from toolboxes allowing for the interfacing of MEG systems in real-time for BMI use [Sudre et al., 2010, Oostenveld, 2014] to real-time brain mapping software capable of quickly identifying signals from electrocorticographic electrodes related to cortical activity corresponding to overt movement, speech, and sensory stimulation [Schalk et al., 2008b, Schalk et al., 2008a].

The Craniux framework is inspired by the system architecture design of BCI2000, and we believe that it takes advantage of several unique features of the LabVIEW graphical development environment for developing real-time BMI software. By making it an open-access and open-source software framework, we hope to serve the research community on at least two fronts. First, at the basic level, Craniux is a BMI software solution with an easy-to-use graphical user interface. Those researchers interested in BMIs can use this software to conduct research without writing custom software. Second and most importantly, we hope this framework will facilitate the development of new BMI paradigms and signal processing algorithms by the research community through providing the basic functionalities of BMI system operation, allowing researchers to focus on the development of their specific research questions. Finally, as this framework is set up in the LabVIEW environment, it naturally inherits the many advantages offered by the high-level graphical nature of LabVIEW programming.

In its current form, the Craniux framework demonstrates the benefits and ease with which it can be used and modified to develop new BMI paradigms and algorithms. The simplicity of the LabVIEW programming language makes the creation of new BMI engines accessible to individuals who may not be familiar with object-oriented programming. Would-be developers can simply take one of the provided engine templates, implement their desired operation, and save the engine under a new name (this process is described in greater detail in the Supplemental Materials). This new engine will then be available for use in the Craniux framework, without the need for the compiling of code down to executables as required by programming languages such as C/C++. The debugging of newly created engines can also be easily performed during run-time through the use of LabVIEW's built-in debugging tools. The dataflow-driven nature of Craniux further simplifies debugging, allowing system execution to be halted and resumed at any point during operation without the loss of the current state of the system. These tools, along with advanced data visualization options, make the rapid prototyping of highly sophisticated neural signal processing techniques possible.

Craniux currently offers a rich set of options to visualize BMI data on-the-fly at multiple processing stages in various formats. Neural signals, such as EEG, MEG, or ECoG, can be viewed as scrolling time-frequency plots or dynamic spatiotemporal plots in the frequency domain. This is beneficial for online examination of neural signal quality, as certain features may be difficult to view in a simple plot of time-domain raw neural signals. The results of calculations performed during the training and application of neural decoding algorithms can also be visualized on-the-fly, providing researchers with the opportunity to select neural signal features, visualize decoding weights, and examine decoder outputs without suspending operation of the system. For example, our implementation of the Population Vector Algorithm allows researchers to dynamically change the value of the R-squared threshold used for feature selection, view the preferred direction distribution of the currently selected features, as well as view the instantaneous contribution of all features to the control signal output by the algorithm. This visualization capability is of particular importance when using and developing sophisticated decoding algorithms, as it allows BMI researchers to judge the validity of the decoding weights on-the-fly and make adjustments of neural signal processing and other BMI experiment parameters accordingly.

We have also shown the potential for enhancement in Craniux performance through its distribution across multiple network hosts, as assigning individual engines to separate computers eliminates the possibility of competition between engines for system resources. The separation of graphical user interface elements from real-time engines further improves system performance by ensuring the real-time engine execution is not affected by user interface interaction events or data visualization. Furthermore, the capability to distribute the Craniux system across multiple network hosts could prove especially useful in long-term human BMI studies. Experimental sessions could be run remotely on a daily basis, eliminating the need for either subjects or investigators to travel to participate in these sessions. This will become important as BMI technology moves into preclinical and clinical trials.

Craniux also offers a streamlined workflow for BMI research. It allows for on-the-fly control of specific experimental parameters, offering experimenters great flexibility for BMI user training. For example, an experimenter can quickly adjust the output gain of a neural decoder if it is deemed that a brain-controlled cursor is moving in the correct direction but with a very low speed. In our experience, this flexibility is critical for effective BMI training. Meanwhile, the Craniux system is capable of capturing all changes in experimental parameters along with BMI data, allowing researchers to perform offline analysis of BMI sessions. Furthermore, various BMI procedures, including the collection of baseline data for the normalization of neural signals and the training of neural decoding algorithms may be performed without the cessation of system operation. This provides both BMI researchers and experimental subjects with a seamless experience in which system parameters can be continuously updated to improve BMI performance.

It should be noted that the timing of the Craniux system is dependent on the timing of the Acquisition engine, which currently can be driven by UDP packets sent from neural acquisition hardware or controlled explicitly by the Acquisition engine itself (e.g. the ‘SimECoG’ engine). Any number of neural recording hardware solutions may be used for BMI operation provided that data recorded by these devices can be packaged and transmitted via UDP. Additionally, hardware-specific acquisition engines can also be created within the Craniux framework.

In addition, it is important to mention that editing or developing new BMI engines in the Craniux framework requires the purchase of LabVIEW. However, it is not uncommon for open-source research tools to be built upon commercial software; two such examples are the EEGLAB [Delorme and Makeig, 2004] and FieldTrip packages for neural data analysis. These packages are both built upon MATLAB, an extremely powerful commercial data analysis software package. Just as many researchers are now using MATLAB instead of custom-written C programs for data analysis, we believe that the time and effort saved by the use of the Craniux system in BMI software development will outweigh the cost of the LabVIEW software. Furthermore, if the Craniux software is to be used as a self-contained out-of-the-box software package, all engines can be compiled down to binary executable files and run using the freely available LabVIEW Runtime Engine, eliminating the need for the LabVIEW software. Finally, as demonstrated in Section A.3, given the computing capability of current personal computers and the code optimization inherently performed by the LabVIEW environment, the overall performance of the Craniux system is comparable to BMI systems developed using other programming languages. Hence, the gain from using the high-level LabVIEW programming environment does not come at the expense of significant sacrifices in system performance.

The Craniux software package, including in-depth documentation and detailed operation instructions for all engines, has been made available free of charge to academic institutions and can be accessed at http://www.engr.pitt.edu/rnel/hRNEL/software.html. The Craniux software package can be downloaded as a library of LabVIEW virtual instruments (VIs), and all stable system updates will be made available for download.

A.5 CONCLUSIONS

While other open-access open-source BMI software solutions are currently available, we feel that the Craniux software package fills a specific need in the realm of BMI research. Powerful yet lightweight, this system allows experimenters to rapidly develop and test cutting-edge technology in an online environment, whether it be new neural signal processing techniques, new neural decoders, or advanced prosthetic devices. This system offers an easy-to-use "out-of-the-box" solution for BMI research as well as other neural data visualization and processing purposes. Additionally, the Craniux system provides an extendable framework through the provision of template engines. The provided framework possesses the basic fundamental architecture for running closed-loop BMI experiments, and enables other researchers to take advantage of LabVIEW functionality to design and conduct novel experimental paradigms without the need to implement their own core system framework. It is also worth noting that functionality offered by the Craniux framework also lends itself useful for other neuroscience research and even neurorehabilitation applications that could benefit from real-time processing and visualization of neural data, such as cortical source imaging using EEG or MEG recordings. It is with these characteristics in mind that we feel the Craniux software package will prove an important addition to the BMI research community.

APPENDIX B

MINIMUM NORM ESTIMATION

B.1 THE STANDARD LEAST-SQUARES SOLUTION

Considering the standard linear regression problem:

$$y = \mathbf{w}^\top \mathbf{x},$$

where $\mathbf{w} \in \mathbb{R}^{M \times 1}$ is a vector of decoding weights mapping the feature observation vector $\mathbf{x} \in \mathbb{R}^{M \times 1}$ to the predicted kinematic output. Given the set of N kinematic observations $\mathbf{y} = [y_1, \ldots, y_N]^\top$ and feature observations $\mathbf{X} = [\mathbf{x}_1, \ldots, \mathbf{x}_N]$, this can be expressed as:

$$\mathbf{y} = \mathbf{X}^\top \mathbf{w}.$$

Solving for $\mathbf{w}$, the standard least-squares solution is obtained:

$$\mathbf{w} = (\mathbf{X}\mathbf{X}^\top)^{-1}\mathbf{X}\mathbf{y}.$$

However, this solution requires $\mathbf{X}\mathbf{X}^\top \in \mathbb{R}^{M \times M}$ to be invertible. As the rank of $\mathbf{X}\mathbf{X}^\top$ is the minimum of M and N, for instances where the number of observations is greater than the number of features (e.g., $N > M$), $\mathbf{X}\mathbf{X}^\top$ will be full-rank and thus invertible. However, in cases where the number of features exceeds the number of observations (e.g., $N < M$), $\mathbf{X}\mathbf{X}^\top$ will not be full-rank and thus cannot be inverted. In such cases, an alternative solution for $\mathbf{w}$ must be found.

B.2 THE MINIMUM NORM SOLUTION

One particular solution for $\mathbf{w}$ in the underdetermined regime (e.g., $N < M$) is the *minimum norm* solution. This solution is defined as that which has the smallest norm of all possible solutions. Finding the minimum norm solution can be formulated as the constrained minimization problem [Boyd, 2007]:

$$\text{minimize} \quad \mathbf{w}^\top \mathbf{w}$$

$$\text{subject to} \quad \mathbf{y} = \mathbf{X}^\top \mathbf{w},$$

which can be solved using Lagrange multipliers. The Lagrangian is:

$$\Lambda(\mathbf{w}, \lambda) = \mathbf{w}^\top \mathbf{w} + \lambda(\mathbf{X}^\top \mathbf{w} - \mathbf{y}),$$

which gives gradients:

$$\nabla_{\mathbf{w}} \Lambda(\mathbf{w}, \lambda) = 2\mathbf{w} + \mathbf{X}\mathbf{w} = 0 \tag{B.1}$$

$$\nabla_{\lambda} \Lambda(\mathbf{w}, \lambda) = \mathbf{X}^\top \mathbf{w} - \mathbf{y} = 0. \tag{B.2}$$

Solving for $\mathbf{w}$ in Equation B.1:

$$\mathbf{w} = -\frac{1}{2}\mathbf{X}\lambda. \tag{B.3}$$

Replacing $\mathbf{w}$ in Equation B.2, the expression for λ may be found:

$$\mathbf{X}^\top(-\frac{1}{2}\mathbf{X}\lambda) - \mathbf{y} = 0$$

$$-\frac{1}{2}\mathbf{X}^\top \mathbf{X}\lambda = \mathbf{y}$$

$$\lambda = -2(\mathbf{X}^\top \mathbf{X})^{-1}\mathbf{y}.$$

Inserting this into Equation B.3 gives the final decoding weights solution:

$$\mathbf{w} = \mathbf{X}(\mathbf{X}^\top \mathbf{X})^{-1}\mathbf{y} \tag{B.4}$$

The minimum norm solution is that which is returned in the `MATLAB` programming environment when using the `pinv` function to find the pseudoinverse of a non-square matrix where $N < M$.

APPENDIX C

DERIVATION OF EMPIRICAL BAYES

The following appendix will outline the derivation of the Empirical Bayes decoding method presented in Chapter 7. First, the terminology and notation used in the derivation of this method is provided in Section C.1. A derivation of probabilistic liner regression is provided in Section C.2. Finally, the evidence maximization method and resultant partial differential equations used during parameter optimization is presented in Section C.3.

C.1 TERMINOLOGY AND NOTATION

The notation included below will be used in the subsequent sections. Non-bold font (e.g. x) is used to indicate scalar values, lower-case bold font (e.g. $\mathbf{x}$) is used to indicate vectors, and capital bold font (e.g. $\mathbf{X}$) is used to indicate matrices. All vectors are assumed to be column vectors.

N	Number of observations.
D	Number of parameters (features).
$\mathbf{y}$	$N \times 1$ vector of kinematic observations.
$\mathbf{w}$	$D \times 1$ vector of parameters (“decoding weights”).
X	$D \times N$ design matrix of feature observations. Columns represent independent observations.
Σ_p	Covariance prior.
ℓ_x^2	Length constant. Subscript x specifies the type of length constant.

C.2 PROBABILISTIC LINEAR REGRESSION

The following presents a derivation of probabilistic linear regression based upon that presented by [Rasmussen and Williams, 2006], with the expansion of selected steps for the sake of completeness. In order to enforce a prior distribution over $\mathbf{w}$, the standard linear regression model must be formulated in a probabilistic framework. Considering the previously-described regression model:

$$y = \mathbf{w}^\top \mathbf{x} + \epsilon \quad , \quad \epsilon \sim \mathcal{N}(0, \sigma_n^2),$$

the likelihood of a single kinematic observation, $p(y_i|\mathbf{w}, \mathbf{x})$ is defined as:

$$p(y_i|\mathbf{w}, \mathbf{x}_i) = \frac{1}{\sqrt{2\pi}\sigma_n} \exp\left[-\frac{(y_i - \mathbf{w}^\top \mathbf{x}_i)^2}{2\sigma_n^2}\right].$$

The likelihood of a set of N independent kinematic observations $\mathbf{y} = [y_1, \ldots, y_N]^\top$ given the set of neural feature observations $\mathbf{X} = [\mathbf{x}_1, \ldots, \mathbf{x}_N]$ can be expressed as the product of the individual probabilities as follows:

$$\begin{aligned}
p(\mathbf{y}|\mathbf{w}, \mathbf{X}) &= \prod_{i=1}^{N} p(y_i|\mathbf{w}, \mathbf{x}_i) \\
&= \prod_{i=1}^{N} \frac{1}{\sqrt{2\pi}\sigma_n} \exp\left[-\frac{(y_i - \mathbf{w}^\top \mathbf{x}_i)^2}{2\sigma_n^2}\right] \\
&= \frac{1}{(2\pi\sigma_n^2)^{N/2}} \exp\left[-\frac{1}{2\sigma_n^2}\sum_{i=1}^{N}(y_i - \mathbf{w}^\top \mathbf{x}_i)^2\right] \\
&= \frac{1}{(2\pi\sigma_n^2)^{N/2}} \exp\left[-\frac{1}{2\sigma_n^2}(\mathbf{y} - \mathbf{X}^\top \mathbf{w})^\top(\mathbf{y} - \mathbf{X}^\top \mathbf{w})\right] \\
&= \mathcal{N}(\mathbf{w}^\top \mathbf{X}, \sigma_n^2 \mathbf{I}).
\end{aligned}$$

We seek to find the set of decoding weights $\mathbf{w}$ that is most likely given both the kinematic and neural observations ($\mathbf{y}$ and $\mathbf{X}$, respectively), which is is equivalent to maximizing the probability distribution $p(\mathbf{w}|\mathbf{y}, \mathbf{X})$, the posterior distribution over the decoding weights. This can be expressed as a function of the likelihood using Bayes' rule:

$$p(\mathbf{y}, \mathbf{w}|\mathbf{X}) = p(\mathbf{y}|\mathbf{X}, \mathbf{w})p(\mathbf{w}|\mathbf{X}) = p(\mathbf{w}|\mathbf{X}, \mathbf{y})p(\mathbf{y}|\mathbf{X})$$

$$p(\mathbf{w}|\mathbf{y}, \mathbf{X}) = \frac{p(\mathbf{y}|\mathbf{X}, \mathbf{w})p(\mathbf{w}|\mathbf{X})}{p(\mathbf{y}|\mathbf{X})}. \tag{C.1}$$

Noting that the prior distribution over $\mathbf{w}$ is independent of $\mathbf{X}$, and that the marginal likelihood is a constant and independent of $\mathbf{w}$, Equation C.1 reduces to

$$p(\mathbf{w}|\mathbf{y}, \mathbf{X}) \propto p(\mathbf{y}|\mathbf{X}, \mathbf{w})p(\mathbf{w}). \tag{C.2}$$

The prior distribution over $\mathbf{w}$ is defined as a zero-mean Gaussian distribution with covariance $\mathbf{\Sigma}_p$:

$$p(\mathbf{w}) \sim \mathcal{N}(\mathbf{0}, \mathbf{\Sigma}_p). \tag{C.3}$$

Combining Equations C.1 and C.2 allows the posterior distribution over $\mathbf{w}$ to be re-written as:

$$
\begin{aligned}
p(\mathbf{w}|\mathbf{y},\mathbf{X}) &\propto p(\mathbf{y}|\mathbf{X},\mathbf{w})p(\mathbf{w}) \\
&\propto \exp\left[-\frac{1}{2\sigma_n^2}(\mathbf{y}-\mathbf{X}^\top\mathbf{w})^\top(\mathbf{y}-\mathbf{X}^\top\mathbf{w})\right]\exp\left[-\frac{1}{2}\mathbf{w}^\top\mathbf{\Sigma}_p^{-1}\mathbf{w}\right] \\
&\propto \exp\left[-\frac{1}{2\sigma_n^2}(\mathbf{y}-\mathbf{X}^\top\mathbf{w})^\top(\mathbf{y}-\mathbf{X}^\top\mathbf{w})-\frac{1}{2}\mathbf{w}^\top\mathbf{\Sigma}_p^{-1}\mathbf{w}\right] \\
&\propto \exp\left[-\frac{1}{2\sigma_n^2}(\mathbf{y}^\top\mathbf{y}-\mathbf{y}^\top\mathbf{X}^\top\mathbf{w}-\mathbf{w}^\top\mathbf{X}\mathbf{y}+\mathbf{w}^\top\mathbf{X}\mathbf{X}^\top\mathbf{w})-\frac{1}{2}\mathbf{w}^\top\mathbf{\Sigma}_p^{-1}\mathbf{w}\right] \\
&\propto \exp\left[-\frac{1}{2\sigma_n^2}\mathbf{w}^\top\mathbf{X}\mathbf{X}^\top\mathbf{w}-\frac{1}{2}\mathbf{w}^\top\mathbf{\Sigma}_p^{-1}\mathbf{w}+\frac{1}{2\sigma_n^2}\mathbf{y}^\top\mathbf{X}^\top\mathbf{w}+\frac{1}{2\sigma_n^2}\mathbf{w}^\top\mathbf{X}\mathbf{y}\right]\exp\left[-\frac{1}{2\sigma_n^2}\mathbf{y}^\top\mathbf{y}\right]
\end{aligned}
$$

Disregarding those terms which are independent of $\mathbf{w}$ (which are thus constant and do not change the proportionality) and collecting terms:

$$
\begin{aligned}
p(\mathbf{w}|\mathbf{y},\mathbf{X}) &\propto \exp\left[-\frac{1}{2}\mathbf{w}^\top(\sigma_n^{-2}\mathbf{X}\mathbf{X}^\top+\mathbf{\Sigma}_p^{-1})\mathbf{w}+\frac{1}{2\sigma_n^2}\mathbf{y}^\top\mathbf{X}^\top\mathbf{w}+\frac{1}{2\sigma_n^2}\mathbf{w}^\top\mathbf{X}\mathbf{y}\right] \\
&\propto \exp\left\{-\frac{1}{2}\left[\mathbf{w}^\top(\sigma_n^{-2}\mathbf{X}\mathbf{X}^\top+\mathbf{\Sigma}_p^{-1})\mathbf{w}-\sigma_n^{-2}\mathbf{w}^\top\mathbf{X}\mathbf{y}-\sigma_n^{-2}\mathbf{y}^\top\mathbf{X}^\top\mathbf{w}\right]\right\}
\end{aligned}
$$

In this form it becomes apparent that the first term in the exponential is that of a Gaussian over $\mathbf{w}$ with covariance $(\sigma_n^{-2}\mathbf{X}\mathbf{X}^\top+\mathbf{\Sigma}_p^{-1})^{-1}$. Letting $A=\sigma_n^{-2}\mathbf{X}\mathbf{X}^\top+\mathbf{\Sigma}_p^{-1}$, the posterior probability can be written as:

$$
p(\mathbf{w}|\mathbf{y},\mathbf{X}) \propto \exp\left[-\frac{1}{2}\left(\mathbf{w}^\top\mathbf{A}\mathbf{w}-\sigma_n^{-2}\mathbf{w}^\top\mathbf{A}\mathbf{A}^{-1}\mathbf{X}\mathbf{y}-\sigma_n^{-2}\mathbf{y}^\top\mathbf{X}^\top\mathbf{A}^{-\top}\mathbf{A}^\top\mathbf{w}\right)\right].
$$

Letting $\bar{\mathbf{w}}=\sigma_n^{-2}\mathbf{A}^{-1}\mathbf{X}\mathbf{y}$, and noting that $\mathbf{A}^\top=\mathbf{A}$ due to symmetry, this reduces to:

$$
p(\mathbf{w}|\mathbf{y},\mathbf{X}) \propto \exp\left[-\frac{1}{2}\left(\mathbf{w}^\top\mathbf{A}\mathbf{w}-\mathbf{w}^\top\mathbf{A}\bar{\mathbf{w}}-\bar{\mathbf{w}}^\top\mathbf{A}\mathbf{w}\right)\right],
$$

which can be further simplified by completing the square through the addition and subtraction of the term $\left(-\frac{1}{2}\bar{\mathbf{w}}^\top\mathbf{A}\bar{\mathbf{w}}\right)$ in the exponential:

$$
\begin{aligned}
p(\mathbf{w}|\mathbf{y},\mathbf{X}) &\propto \exp\left[-\frac{1}{2}\left(\mathbf{w}^\top\mathbf{A}\mathbf{w}-\mathbf{w}^\top\mathbf{A}\bar{\mathbf{w}}-\bar{\mathbf{w}}^\top\mathbf{A}\mathbf{w}+\bar{\mathbf{w}}^\top\mathbf{A}\bar{\mathbf{w}}-\bar{\mathbf{w}}^\top\mathbf{A}\bar{\mathbf{w}}\right)\right] \\
&\propto \exp\left[-\frac{1}{2}(\mathbf{w}-\bar{\mathbf{w}})^\top\mathbf{A}(\mathbf{w}-\bar{\mathbf{w}})\right]\exp\left[-\frac{1}{2}\bar{\mathbf{w}}^\top\mathbf{A}\bar{\mathbf{w}}\right].
\end{aligned}
$$

Finally, as the second exponential term does not depend on $\mathbf{w}$, it may be safely disregarded to yield:

$$
p(\mathbf{w}|\mathbf{y},\mathbf{X}) \propto \exp\left[-\frac{1}{2}(\mathbf{w}-\bar{\mathbf{w}})^\top\mathbf{A}(\mathbf{w}-\bar{\mathbf{w}})\right]
$$

$$
p(\mathbf{w}|\mathbf{X},\mathbf{y}) \sim \mathcal{N}(\bar{\mathbf{w}},\mathbf{A}^{-1}),
$$

Given a new feature observation $\mathbf{x}_*$, the kinematic prediction y_* is obtained by maximizing the probability distribution:

$$
p(y_*|\mathbf{x}_*,\mathbf{X},\mathbf{y}),
$$

the probability of the kinematic prediction given the new feature observation, the set of kinematic training data $\mathbf{y}$, and the set of neural training data $\mathbf{X}$. This is equivalent to the joint probability distribution of y_* and $\mathbf{w}$, integrated over $\mathbf{w}$:

$$p(y_*|\mathbf{x}_*, \mathbf{X}, \mathbf{y}) = \int p(y_*, \mathbf{w}|\mathbf{x}_*, \mathbf{X}, \mathbf{y}) d\mathbf{w}$$

. By factoring $\mathbf{w}$ out of the joint probability distribution $p(y_*, \mathbf{w}|\mathbf{x}_*, \mathbf{X}, \mathbf{y})$, this reduces to:

$$p(y_*|\mathbf{x}_*, \mathbf{X}, \mathbf{y}) = \int p(y_*|\mathbf{x}_*, \mathbf{X}, \mathbf{y}, \mathbf{w}) p(\mathbf{w}|\mathbf{x}_*, \mathbf{X}, \mathbf{y}) d\mathbf{w}.$$

This can be further reduced by noting that $\mathbf{w}$ is independent of $\mathbf{x}_*$, and that y_* is independent of $\mathbf{X}$ and $\mathbf{y}$ given $\mathbf{w}$:

$$\begin{aligned} p(y_*|\mathbf{x}_*, \mathbf{X}, \mathbf{y}) &= \int p(y_*|\mathbf{x}_*, \mathbf{X}, \mathbf{y}) p(\mathbf{w}|\mathbf{x}_*, \mathbf{X}, \mathbf{y}) d\mathbf{w} \\ &= \mathcal{N}\left(\frac{1}{\sigma_n^2}\mathbf{x}_*^\top \mathbf{A}^{-1}\mathbf{X}\mathbf{y}, \mathbf{x}_*^\top \mathbf{A}^{-1}\mathbf{x}_*\right). \end{aligned}$$

As the maximum of a Gaussian distribution occurs at its mean, this give the final equation for kinematic predictions y_* given observed feature activity $\mathbf{x}_*$:

$$\begin{aligned} y_* &= \frac{1}{\sigma_n^2}\mathbf{x}_*^\top \mathbf{A}^{-1}\mathbf{X}\mathbf{y} \\ &= \bar{\mathbf{w}}^\top \mathbf{x}_*. \end{aligned}$$

C.3 EVIDENCE MAXIMIZATION

The previous section described how decoding weights and kinematic predictions are obtained for probabilistic linear regression. However, this is dependent upon knowing $\mathbf{\Sigma}_p$, the covariance prior over the decoding weights. As $\mathbf{\Sigma}_p$ can be determined by a number of hyperparameters, the values for these hyperparameters must be found before the final decoding weights solution can be calculated. One way in which this can be accomplished is through *evidence maximization*, where the probability distribution for the evidence $p(\mathbf{y}|\mathbf{X}, \theta)$, the probability of the data given hyperparameters, is maximized with respect to the set of hyperparameters. This can be expressed using known probability distributions as:

$$\mathcal{E}(\theta) = p(\mathbf{y}|\mathbf{X}) = \int p(\mathbf{y}|\mathbf{X}, \mathbf{w}) p(\mathbf{w}) d\mathbf{w}, \tag{C.4}$$

where θ is used to indicate the set of hyperparameters, $p(\mathbf{y}|\mathbf{X}, \mathbf{w}) \sim \mathcal{N}(\mathbf{X}^\top \mathbf{w}, \sigma_n^2 \mathbf{I})$, and $p(\mathbf{w}) \sim \mathcal{N}(\mathbf{0}, \mathbf{\Sigma}_p)$. Replacing the probability distributions in Equation C.4 with their known values, this becomes [Linden and

Sahani, 2003]:

$$\begin{aligned}
\mathcal{E}(\theta) &= \int \mathcal{N}(\mathbf{X}^\top \mathbf{w}, \sigma_n^2 \mathbf{I}) \mathcal{N}(\mathbf{0}, \mathbf{\Sigma}_p) d\mathbf{w} \\
&= \int \frac{1}{(2\pi\sigma_n^2)^{M/2}} \exp\left[-\frac{1}{2\sigma_n^2}(\mathbf{y} - \mathbf{X}^\top \mathbf{w})^\top (\mathbf{y} - \mathbf{X}^\top \mathbf{w})\right] \frac{1}{(2\pi)^{M/2}|\mathbf{\Sigma}_p|^{1/2}} \exp\left[-\frac{1}{2}\mathbf{w}\mathbf{\Sigma}_p^{-1}\mathbf{w}^\top\right] d\mathbf{w} \\
&= (2\pi\sigma_n^2)^{-\frac{M}{2}} (2\pi)^{-\frac{M}{2}} |\mathbf{\Sigma}_p|^{\frac{-1}{2}} \int \exp\left[-\frac{1}{2}(\mathbf{w} - \bar{\mathbf{w}})^\top \mathbf{A} (\mathbf{w} - \bar{\mathbf{w}})\right] d\mathbf{w} \\
&= (2\pi\sigma_n^2)^{-\frac{M}{2}} |\mathbf{\Sigma}_p|^{-\frac{1}{2}} |\mathbf{A}^{-1}|^{\frac{1}{2}} \exp\left(-\frac{1}{2}\bar{\mathbf{w}}^\top \mathbf{A} \bar{\mathbf{w}}\right)
\end{aligned}$$

This equation must be maximized with respect to noise variance parameter σ_n^2 and hyperparameters of Σ_p (indicated by θ), which is equivalent to maximizing the log-evidence:

$$\ln \mathcal{E}(\sigma_n^2, \theta) = \frac{1}{2} \ln |2\pi \mathbf{A}| - \frac{N}{2} \ln(2\pi\sigma_n^2) - \frac{1}{2} \ln |2\pi \mathbf{\Sigma_p}| - \frac{1}{2} \mathbf{y} \left(\sigma_n^{-2} \mathbf{I} - \sigma_n^{-4} \mathbf{X}^\top \mathbf{A} \mathbf{X}\right) \mathbf{y}^\top.$$

As σ_n is required to remain positive, constrained maximization of $\ln \mathcal{E}(\sigma_n^2, \theta)$ can be re-formulated under the change of variable $\sigma_n^2 = e^{\alpha_n}$ as an unconstrained maximization of:

$$\ln \mathcal{E}(\alpha, \theta) = \frac{1}{2} \ln |2\pi \mathbf{A}| - \frac{N}{2} \ln(2\pi e^{\alpha}) - \frac{1}{2} \ln |2\pi \mathbf{\Sigma_p}| - \frac{1}{2} \mathbf{y} \left(e^{-\alpha} \mathbf{I} - e^{-2\alpha} \mathbf{X}^\top \mathbf{A} \mathbf{X}\right) \mathbf{y}^\top. \tag{C.5}$$

In order to maximize this using gradient ascent, the partial derivatives with respect to all of the parameters must be computed. Taking partial derivative of Equation C.5 with respect to α_n yields the following:

$$\begin{aligned}
\frac{\partial}{\partial \alpha} \ln \mathcal{E}(\alpha, \theta) &= -N + \mathrm{Tr}(\mathbf{I} - \mathbf{A}\Sigma_p^{-1}) + e^{-\alpha}(\mathbf{y} - \mathbf{w}^\top \mathbf{X})(\mathbf{y} - \mathbf{w}^\top \mathbf{X})^\top \\
&= -N + \mathrm{Tr}(\mathbf{I} - \mathbf{A}\Sigma_p^{-1}) + e^{-\alpha}(\mathbf{y}\mathbf{y}^\top - \mathbf{y}\mathbf{X}^\top \mathbf{w} - \mathbf{w}^\top \mathbf{X} \mathbf{y}^\top + \mathbf{w}^\top \mathbf{X}\mathbf{X}^\top \mathbf{w}).
\end{aligned}$$

In order to compute the partial derivatives with respect to θ, the set of hyperparameters of $\mathbf{\Sigma}_p$, the partial derivative of the log-evidence function can be expressed as a function of the partial derivative of $\mathbf{\Sigma}_p$:

$$\frac{\partial}{\partial \theta} \ln \mathcal{E}(\alpha, \theta) = \frac{1}{2} \mathrm{Tr}\left[\mathbf{\Sigma_p}^{-1} \left(\frac{\partial}{\partial \theta} \mathbf{\Sigma_p}\right) \mathbf{\Sigma_p}^{-1} (\Sigma_p - \mathbf{A} - \mathbf{w}\mathbf{w}^\top)\right]. \tag{C.6}$$

The following sections will outline calculation of the partial derivatives for the various covariance prior forms.

C.3.1 General covariance prior

The general form of the covariance prior is expressed as:

$$\mathbf{\Sigma}_p = \sigma_{w,f}^2 \mathbf{K} + \sigma_{w,n}^2 \mathbf{I}.$$

Under the change of variables $\sigma_{w,f}^2 = e^{-\alpha_{w,f}}$ and $\sigma_{w,n}^2 = e^{-\alpha_{w,n}}$, this becomes:

$$\mathbf{\Sigma}_p = e^{-\alpha_{w,f}} \mathbf{K} + e^{-\alpha_{w,n}} \mathbf{I},$$

which leads to expressions for the partial derivatives with respect to $\alpha_{w,f}$ and $\alpha_{w,n}$

$$\frac{\partial}{\partial \alpha_{w,f}} \mathbf{\Sigma}_p = -e^{-\alpha_{w,f}} \mathbf{K}, \tag{C.7}$$

$$\frac{\partial}{\partial \alpha_{w,n}} \mathbf{\Sigma}_p = -e^{-\alpha_{w,n}} \mathbf{I}. \tag{C.8}$$

In order to compute partial derivatives with respect to hyperparameters of the kernel $\mathbf{K}$, the partial derivative of the covariance prior is expressed as a function of the partial derivative of the kernel:

$$\frac{\partial}{\partial \theta} \mathbf{\Sigma}_p = e^{-\alpha_{w,f}} \frac{\partial}{\partial \theta} \mathbf{K} \tag{C.9}$$

C.3.2 L_2-regularized linear regression

For the L_2-regularized linear regression, $\mathbf{K} = \mathbf{I}$. As this kernel does not contain any parameters, no additional partial derivatives are required.

C.3.3 Frequency-block-diagonal prior

For the frequency-block-diagonal prior, the kernel is defined as:

$$\mathbf{K}_f(x_i, x_j) = \exp\left\{ -\frac{1}{2\ell_f^2} \left[f(x_i) - f(x_j) \right]^2 \right\} \delta_{el}(x_i, x_j)$$

where $f(x_i)$ is the frequency band corresponding to feature x_i, hyperparameter ℓ_f^2 is the "spectral length constant", and delta function $\delta_{el}(x_i, x_j)$ is 1 if features x_i and x_j are on the same electrode, and 0 otherwise. Under the change of variable $\ell_f^2 = e_f^\gamma$, the partial derivative of the covariance kernel with respect to γ_f is:

$$\frac{\partial}{\partial \gamma_f} \mathbf{K}_f(x_i, x_j) = \frac{1}{2} e^{-\gamma_f} \left[f(x_i) - f(x_j) \right] \mathbf{K}_f(x_i, x_j). \tag{C.10}$$

C.3.4 Frequency-spatial prior

For the frequency-spatial prior, the kernel is defined as:

$$\mathbf{K}_{f,s}(x_i, x_j) = \mathbf{K}_f(x_i, x_j) \mathbf{K}_s(x_i, x_j) = \exp\left\{ -\frac{\left[f(x_i) - f(x_j) \right]^2}{2\ell_f^2} \right\} \exp\left\{ -\frac{d^2(x_i, x_j)}{2\ell_s^2} \right\},$$

where $d^2(x_i, x_j)$ is the square of the spatial distance between electrodes for features x_i and x_j and ℓ_s is the "spectral length constant" hyperparameter determining the correlation between elements of $\mathbf{w}$ as a function of the distance between the electrodes for their corresponding features. Under the change of variable $\ell_f^2 = e_f^\gamma$ and $\ell_s^2 = e_s^\gamma$, the partial derivatives of the covariance kernel with respect to γ_f and γ_s are:

$$\frac{\partial}{\partial \gamma_f} \mathbf{K}_{f,s}(x_i, x_j) = \frac{1}{2} e^{-\gamma_f} \left[f(x_i) - f(x_j) \right] \mathbf{K}_f(x_i, x_j) \mathbf{K}_s(x_i, x_j), \tag{C.11}$$

$$\frac{\partial}{\partial \gamma_s} \mathbf{K}_{f,s}(x_i, x_j) = \frac{1}{2} e^{-\gamma_s} d^2(x_i, x_j) \mathbf{K}_f(x_i, x_j) \mathbf{K}_s(x_i, x_j), \tag{C.12}$$

C.4 IMPLEMENTATION

Evidence maximization was preformed through minimization of the negative-log-evidence function via. gradient descent. Gradient descent was performed using the `minimize.m` function by Carl Edward Rasmussen (`http://learning.eng.cam.ac.uk/carl/code/minimize/minimize.m`), which uses Polack-Ribiere conjugate gradients to compute search directions, a line search using quadratic and cubic polynomial approximations, and the Wolfe-Powell stopping criteria used together with the slope ratio method for guessing initial step sizes [Rasmussen, 2006]. Random restarts ($n = 5$) were used to define initial parameter values, with the final parameter set chosen as that minimizing the negative log-evidence of the data. All methods were implemented in the MATLAB programming environment.

BIBLIOGRAPHY

[Acharya et al., 2010] Acharya, S., Fifer, M. S., Benz, H. L., Crone, N. E., and Thakor, N. V. (2010). Electrocorticographic amplitude predicts finger positions during slow grasping motions of the hand. *Journal of Neural Engineering.*

[Adee, 2008] Adee, S. (2008). Dean Kamen's" Luke Arm" Prosthesis Readies for Clinical Trials. *IEEE Spectrum.*

[Adee, 2009] Adee, S. (2009). Winner: the revolution will be prosthetized. *IEEE Spectrum.*

[Adeeb et al., 2012] Adeeb, N., Mortazavi, M. M., Tubbs, R. S., and Cohen-Gadol, A. A. (2012). The cranial dura mater: a review of its history, embryology, and anatomy. *Child's Nervous System.*

[Amirikian and Georgopoulos, 2003] Amirikian, B. and Georgopoulos, A. P. (2003). Modular organization of directionally tuned cells in the motor cortex: Is there a short-range order? *Proceedings of the National Academy of Sciences of the United States of America.*

[Anderson, 2001] Anderson, J. M. (2001). Biological responses to materials. *Annual Review of Materials Research.*

[Anderson et al., 2008] Anderson, J. M., Rodriguez, A., and Chang, D. T. (2008). Foreign body reaction to biomaterials. *Seminars in immunology.*

[Anderson et al., 2012] Anderson, N. R., Blakely, T., Schalk, G., Leuthardt, E. C., and Moran, D. W. (2012). Electrocorticographic (ECoG) correlates of human arm movements. *Experimental brain research Experimentelle Hirnforschung Expérimentation cérébrale.*

[Arroyo et al., 1993] Arroyo, S., Lesser, R. P., Gordon, B., Uematsu, S., Jackson, D., and Webber, R. (1993). Functional significance of the mu rhythm of human cortex: an electrophysiologic study with subdural electrodes. *Electroencephalography and clinical Neurophysiology.*

[Azemi et al., 2011] Azemi, E., Lagenaur, C. F., and Cui, X. T. (2011). The surface immobilization of the neural adhesion molecule L1 on neural probes and its effect on neuronal density and gliosis at the probe/tissue interface. *Biomaterials.*

[Bacher et al., 2008] Bacher, D., McFerron, J., Krishnamurthy, N., and Batista, A. (2008). An Expimental Rig for Closed-Loop Neuroprosthetics. In *Poster presented as part of the Society for Neuroscience Conference. Washington D.C.*

[Ball et al., 2009] Ball, T., Schulze-Bonhage, A., Aertsen, A., and Mehring, C. (2009). Differential representation of arm movement direction in relation to cortical anatomy and function. *Journal of Neural Engineering.*

[Barrese et al., 2013] Barrese, J. C., Rao, N., Paroo, K., Triebwasser, C., Vargas-Irwin, C., Franquemont, L., and Donoghue, J. P. (2013). Failure mode analysis of silicon-based intracortical microelectrode arrays in non-human primates. *Journal of Neural Engineering.*

[Biran et al., 2005] Biran, R., Martin, D. C., and Tresco, P. A. (2005). Neuronal cell loss accompanies the brain tissue response to chronically implanted silicon microelectrode arrays. *Experimental neurology.*

[Biran et al., 2007] Biran, R., Martin, D. C., and Tresco, P. A. (2007). The brain tissue response to implanted silicon microelectrode arrays is increased when the device is tethered to the skull. *Journal of biomedical materials research. Part A.*

[Bishop, 2006] Bishop, C. M. (2006). *Pattern Recognition and Machine Learning.* Springer Verlag.

[Bjornsson et al., 2006] Bjornsson, C. S., Oh, S. J., Al-Kofahi, Y. A., Lim, Y. J., Smith, K. L., Turner, J. N., De, S., Roysam, B., Shain, W., and Kim, S. J. (2006). Effects of insertion conditions on tissue strain and vascular damage during neuroprosthetic device insertion. *Journal of Neural Engineering.*

[Blakely et al., 2009] Blakely, T., Miller, K. J., Zanos, S. P., Rao, R. P. N., and Ojemann, J. G. (2009). Robust, long-term control of an electrocorticographic brain-computer interface with fixed parameters. *Neurosurgical focus.*

[Bouchard et al., 2013] Bouchard, K. E., Mesgarani, N., Johnson, K., and Chang, E. F. (2013). Functional organization of human sensorimotor cortex for speech articulation. *Nature.*

[Boyd, 2007] Boyd, S. (2007). Least-norm solutions of undetermined equations. Retrieved from `http://see.stanford.edu/materials/lsoeldsee263/08-min-norm.pdf`.

[Bridges and Garcia, 2008] Bridges, A. W. and Garcia, A. J. (2008). Anti-inflammatory polymeric coatings for implantable biomaterials and devices. *Journal of diabetes science and technology.*

[Brockwell, 2004] Brockwell, A. E. (2004). Recursive Bayesian Decoding of Motor Cortical Signals by Particle Filtering.

[Brodbeck et al., 2003] Brodbeck, W. G., Voskerician, G., Ziats, N. P., Nakayama, Y., Matsuda, T., and Anderson, J. M. (2003). In vivo leukocyte cytokine mRNA responses to biomaterials are dependent on surface chemistry. *Journal of biomedical materials research. Part A.*

[Brunner et al., 2009] Brunner, P., Ritaccio, A. L., Lynch, T. M., Emrich, J. F., Wilson, J. A., Williams, J. C., Aarnoutse, E. J., Ramsey, N. F., Leuthardt, E. C., Bischof, H., and Schalk, G. (2009). A practical procedure for real-time functional mapping of eloquent cortex using electrocorticographic signals in humans. *Epilepsy & behavior : E&B.*

[Bundy et al., 2014] Bundy, D. T., Zellmer, E., Gaona, C. M., Sharma, M., Szrama, N., Hacker, C., Freudenburg, Z. V., Daitch, A., Moran, D. W., and Leuthardt, E. C. (2014). Characterization of the effects of the human dura on macro- and micro-electrocorticographic recordings. *Journal of Neural Engineering.*

[Buzsáki et al., 2012] Buzsáki, G., Anastassiou, C. A., and Koch, C. (2012). The origin of extracellular fields and currents–EEG, ECoG, LFP and spikes. *Nature Reviews Neuroscience.*

[Buzsáki and Draguhn, 2004] Buzsáki, G. and Draguhn, A. (2004). Neuronal oscillations in cortical networks. *Science (New York, NY).*

[Caetano et al., 2007] Caetano, G., Jousmäki, V., and Hari, R. (2007). Actor's and observer's primary motor cortices stabilize similarly after seen or heard motor actions. *Proceedings of the National Academy of Sciences of the United States of America.*

[Campbell et al., 1991] Campbell, P. K., Jones, K. E., Huber, R. J., Horch, K. W., and Normann, R. A. (1991). A silicon-based, three-dimensional neural interface: manufacturing processes for an intracortical electrode array. *IEEE transactions on bio-medical engineering.*

[Carmena et al., 2003] Carmena, J. M., Lebedev, M. A., Crist, R. E., O'doherty, J. E., Santucci, D. M., Dimitrov, D. F., Patil, P. G., Henriquez, C. S., and Nicolelis, M. A. L. (2003). Learning to control a brain-machine interface for reaching and grasping by primates. *PLoS Biology.*

[Carter and Houk, 1993] Carter, R. R. and Houk, J. C. (1993). Multiple single-unit recordings from the CNS using thin-film electrode arrays. *IEEE transactions on rehabilitation engineering : a publication of the IEEE Engineering in Medicine and Biology Society.*

[Cervenka et al., 2011] Cervenka, M. C., Boatman-Reich, D. F., Ward, J., Franaszczuk, P. J., and Crone, N. E. (2011). Language mapping in multilingual patients: electrocorticography and cortical stimulation during naming. *Frontiers in Human Neuroscience.*

[Chao et al., 2010] Chao, Z. C., Nagasaka, Y., and Fujii, N. (2010). Long-term asynchronous decoding of arm motion using electrocorticographic signals in monkeys. *Frontiers in Neuroengineering.*

[Chase et al., 2012] Chase, S. M., Kass, R. E., and Schwartz, A. B. (2012). Behavioral and neural correlates of visuomotor adaptation observed through a brain-computer interface in primary motor cortex. *Journal of Neurophysiology.*

[Chase et al., 2009] Chase, S. M., Schwartz, A. B., and Kass, R. E. (2009). Bias, optimal linear estimation, and the differences between open-loop simulation and closed-loop performance of spiking-based brain-computer interface algorithms. *Neural networks : the official journal of the International Neural Network Society.*

[Chestek et al., 2013] Chestek, C. A., Gilja, V., Blabe, C. H., Foster, B. L., Shenoy, K. V., Parvizi, J., and Henderson, J. M. (2013). Hand posture classification using electrocorticography signals in the gamma band over human sensorimotor brain areas. *Journal of Neural Engineering.*

[Chestek et al., 2011] Chestek, C. A., Gilja, V., Nuyujukian, P., Foster, J. D., Fan, J. M., Kaufman, M. T., Churchland, M. M., Rivera-Alvidrez, Z., Cunningham, J. P., Ryu, S. I., and 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.*

[Chin et al., 2007] Chin, C. M., Popovic, M. R., Thrasher, A., Cameron, T., Lozano, A., and Chen, R. (2007). Identification of arm movements using correlation of electrocorticographic spectral components and kinematic recordings. *Journal of Neural Engineering.*

[Christensen et al., 2007] Christensen, M. S., Lundbye-Jensen, J., Geertsen, S. S., Petersen, T. H., Paulson, O. B., and Nielsen, J. B. (2007). Premotor cortex modulates somatosensory cortex during voluntary movements without proprioceptive feedback. *Nature Neuroscience.*

[Chu et al., 2011] Chu, C., Ni, Y., Tan, G., Saunders, C. J., and Ashburner, J. (2011). Kernel regression for fMRI pattern prediction. *NeuroImage.*

[Collier et al., 2004] Collier, T. O., Anderson, J. M., Brodbeck, W. G., Barber, T., and Healy, K. E. (2004). Inhibition of macrophage development and foreign body giant cell formation by hydrophilic interpenetrating polymer network. *Journal of biomedical materials research. Part A.*

[Collinger et al., 2014] Collinger, J. L., Vinjamuri, R., Degenhart, A. D., Weber, D. J., Sudre, G. P., Boninger, M. L., Tyler-Kabara, E. C., and Wang, W. (2014). Motor-related brain activity during action observation: a neural substrate for electrocorticographic brain-computer interfaces after spinal cord injury.

[Collinger et al., 2012] Collinger, J. L., Wodlinger, B., Downey, J. E., Wang, W., Tyler-Kabara, E. C., Weber, D. J., McMorland, A. J., Velliste, M., Boninger, M. L., and Schwartz, A. B. (2012). High-performance neuroprosthetic control by an individual with tetraplegia. *Lancet.*

[COOPER et al., 1965] COOPER, R., WINTER, A. L., CROW, H. J., and WALTER, W. G. (1965). Comparison of subcortical, cortical and scalp activity using chronically indwelling electrodes in man. *... and clinical neurophysiology.*

[Cox, 1996] Cox, R. W. (1996). AFNI: software for analysis and visualization of functional magnetic resonance neuroimages. *Computers and biomedical research, an international journal.*

[Cramer et al., 2005] Cramer, S. C., Lastra, L., Lacourse, M. G., and Cohen, M. J. (2005). Brain motor system function after chronic, complete spinal cord injury. *Brain : a journal of neurology.*

[Crapse and Sommer, 2008] Crapse, T. B. and Sommer, M. A. (2008). Corollary discharge circuits in the primate brain. *Current opinion in neurobiology.*

[Crone et al., 2001] Crone, N. E., Hao, L., Hart, J., Boatman, D., Lesser, R. P., Irizarry, R., and Gordon, B. (2001). Electrocorticographic gamma activity during word production in spoken and sign language. *Neurology.*

[Crone et al., 1993] Crone, N. E., Lesser, R. P., Krauss, G. L., and Nathan, S. S. (1993). *Topographic mapping of human sensorimotor cortex with electrocortical spectra.* Epilepsia.

[Crone et al., 1998a] Crone, N. E., Miglioretti, D. L., Gordon, B., and Lesser, R. P. (1998a). Functional mapping of human sensorimotor cortex with electrocorticographic spectral analysis. II. Event-related synchronization in the gamma band. *Brain : a journal of neurology.*

[Crone et al., 1998b] Crone, N. E., Miglioretti, D. L., Gordon, B., Sieracki, J. M., Wilson, M. T., Uematsu, S., and Lesser, R. P. (1998b). Functional mapping of human sensorimotor cortex with electrocorticographic spectral analysis. I. Alpha and beta event-related desynchronization. *Brain : a journal of neurology.*

[Crone et al., 2006] Crone, N. E., Sinai, A., and Korzeniewska, A. (2006). High-frequency gamma oscillations and human brain mapping with electrocorticography. *Progress in brain research.*

[Dangi et al., 2014] Dangi, S., Gowda, S., Moorman, H. G., Orsborn, A. L., So, K., Shanechi, M., and Carmena, J. M. (2014). Continuous closed-loop decoder adaptation with a recursive maximum likelihood algorithm allows for rapid performance acquisition in brain-machine interfaces. *Neural computation.*

[Dayan and Abbott, 2001] Dayan, P. and Abbott, L. F. (2001). *Theoretical Neuroscience.* MIT Press.

[Degenhart et al., 2011a] Degenhart, A. D., Collinger, J. L., Vinjamuri, R., Kelly, J., Tyler-Kabara, E. C., and Wang, W. (2011a). Classification of hand posture from electrocorticographic signals recorded during varying force conditions. In *Engineering in Medicine and Biology Society,EMBC, 2011 Annual International Conference of the IEEE.*

[Degenhart et al., 2011b] Degenhart, A. D., Kelly, J. W., Ashmore, R. C., Collinger, J. L., Tyler-Kabara, E. C., Weber, D. J., and Wang, W. (2011b). Craniux: A LabVIEW-Based Modular Software Framework for Brain-Machine Interface Research. *Computational Intelligence and Neuroscience.*

[Delorme and Makeig, 2004] Delorme, A. and Makeig, S. (2004). EEGLAB: an open source toolbox for analysis of single-trial EEG dynamics including independent component analysis. *Journal of Neuroscience Methods.*

[Ding et al., 2008] Ding, M. C., Lo, E. H., and Stanley, G. B. (2008). Sustained focal cortical compression reduces electrically-induced seizure threshold. *Neuroscience.*

[Donchin et al., 2000] Donchin, E., Spencer, K. M., and Wijesinghe, R. (2000). The mental prosthesis: assessing the speed of a P300-based brain-computer interface. *IEEE transactions on rehabilitation engineering : a publication of the IEEE Engineering in Medicine and Biology Society.*

[Dushanova and Donoghue, 2010] Dushanova, J. and Donoghue, J. (2010). Neurons in primary motor cortex engaged during action observation. *European Journal of Neuroscience*.

[Ebersole, 1997] Ebersole, J. S. (1997). Defining epileptogenic foci: past, present, future. *Journal of clinical neurophysiology*.

[Edwards et al., 2009] Edwards, E., Nagarajan, S. S., Dalal, S. S., Canolty, R. T., Kirsch, H. E., Barbaro, N. M., and Knight, R. T. (2009). Spatiotemporal imaging of cortical activation during verb generation and picture naming. *NeuroImage*.

[Edwards et al., 2005] Edwards, E., Soltani, M., Deouell, L. Y., Berger, M. S., and Knight, R. T. (2005). High gamma activity in response to deviant auditory stimuli recorded directly from human cortex. *Journal of Neurophysiology*.

[Eng et al., 2000] Eng, L. F., Ghirnikar, R. S., and Lee, Y. L. (2000). Glial fibrillary acidic protein: GFAP-thirty-one years (1969-2000). *Neurochemical research*.

[Felton et al., 2007] Felton, E. A., Wilson, J. A., Williams, J. C., and Garell, P. C. (2007). Electrocorticographically controlled brain-computer interfaces using motor and sensory imagery in patients with temporary subdural electrode implants. Report of four cases. *Journal of Neurosurgery*.

[Fetz, 1969] Fetz, E. E. (1969). Operant conditioning of cortical unit activity. *Science (New York, NY)*.

[Fetz and Baker, 1973] Fetz, E. E. and Baker, M. A. (1973). Operantly conditioned patterns on precentral unit activity and correlated responses in adjacent cells and contralateral muscles. *Journal of Neurophysiology*.

[Fetz and Finocchio, 1975] Fetz, E. E. and Finocchio, D. V. (1975). Correlations between activity of motor cortex cells and arm muscles during operantly conditioned response patterns. *Experimental brain research Experimentelle Hirnforschung Expérimentation cérébrale*.

[Flint et al., 2014] Flint, R. D., Wang, P. T., Wright, Z. A., King, C. E., Krucoff, M. O., Schuele, S. U., Rosenow, J. M., Hsu, F. P. K., Liu, C. Y., Lin, J. J., Sazgar, M., Millett, D. E., Shaw, S. J., Nenadic, Z., Do, A. H., and Slutzky, M. W. (2014). Extracting kinetic information from human motor cortical signals. *NeuroImage*.

[Flint et al., 2013] Flint, R. D., Wright, Z. A., Scheid, M. R., and Slutzky, M. W. (2013). Long term, stable brain machine interface performance using local field potentials and multiunit spikes. *Journal of Neural Engineering*.

[Foldes and Taylor, 2013] Foldes, S. T. and Taylor, D. M. (2013). Speaking and cognitive distractions during EEG-based brain control of a virtual neuroprosthesis-arm. *Journal of neuroengineering and rehabilitation*.

[Fountas and Smith, 2007] Fountas, K. N. and Smith, J. R. (2007). Subdural Electrode-Associated Complications: A 20-Year Experience. *Stereotactic and functional neurosurgery*.

[Freeman et al., 2000] Freeman, W. J., Rogers, L. J., and Holmes, M. D. (2000). Spatial spectral analysis of human electrocorticograms including the alpha and gamma bands. *Journal of Neuroscience*.

[Freire et al., 2011] Freire, M. A. M., Morya, E., Faber, J., Santos, J. R., Guimaraes, J. S., Lemos, N. A. M., Sameshima, K., Pereira, A., Ribeiro, S., and Nicolelis, M. A. L. (2011). Comprehensive analysis of tissue preservation and recording quality from chronic multielectrode implants. *PLoS ONE*.

[Gage et al., 2005] Gage, G. J., Ludwig, K. A., Otto, K. J., Ionides, E. L., and Kipke, D. R. (2005). Naïve coadaptive cortical control. *Journal of Neural Engineering*.

[Gallese et al., 1996] Gallese, V., Fadiga, L., Fogassi, L., and Rizzolatti, G. (1996). Action recognition in the premotor cortex. *Brain : a journal of neurology*.

[Ganguli and Sompolinsky, 2012] Ganguli, S. and Sompolinsky, H. (2012). Compressed sensing, sparsity, and dimensionality in neuronal information processing and data analysis. *Annual Review of Neuroscience.*

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

[Gaona et al., 2011] Gaona, C. M., Sharma, M., Freudenburg, Z. V., Breshears, J. D., Bundy, D. T., Roland, J., Barbour, D. L., Schalk, G., and Leuthardt, E. C. (2011). Nonuniform high-gamma (60-500 Hz) power changes dissociate cognitive task and anatomy in human cortex. *The Journal of neuroscience : the official journal of the Society for Neuroscience.*

[Georgopoulos et al., 1982] Georgopoulos, A. P., Kalaska, J. F., Caminiti, R., and Massey, J. T. (1982). On the relations between the direction of two-dimensional arm movements and cell discharge in primate motor cortex. *The Journal of neuroscience : the official journal of the Society for Neuroscience.*

[Georgopoulos et al., 2007] Georgopoulos, A. P., Merchant, H., Naselaris, T., and Amirikian, B. (2007). Mapping of the preferred direction in the motor cortex. *Proceedings of the National Academy of Sciences.*

[Georgopoulos et al., 1986] Georgopoulos, A. P., Schwartz, A. B., and Kettner, R. E. (1986). Neuronal population coding of movement direction. *Science (New York, NY).*

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

[Gilletti and Muthuswamy, 2006] Gilletti, A. and Muthuswamy, J. (2006). Brain micromotion around implants in the rodent somatosensory cortex. *Journal of Neural Engineering.*

[Gittins and Harrison, 2004] Gittins, R. and Harrison, P. J. (2004). Neuronal density, size and shape in the human anterior cingulate cortex: a comparison of Nissl and NeuN staining. *Brain research bulletin.*

[Godlove et al., 2014] Godlove, J. M., Whaite, E. O., and Batista, A. P. (2014). Comparing temporal aspects of visual, tactile, and microstimulation feedback for motor control. *Journal of Neural Engineering.*

[Golub et al., 2014] Golub, M. D., Yu, B. M., Schwartz, A. B., and Chase, S. M. (2014). Motor cortical control of movement speed with implications for brain-machine interface control.

[Gritsenko et al., 2007] Gritsenko, V., Krouchev, N. I., and Kalaska, J. F. (2007). Afferent Input, Efference Copy, Signal Noise, and Biases in Perception of Joint Angle During Active Versus Passive Elbow Movements. *Journal of Neurophysiology.*

[Gunduz et al., 2007] Gunduz, A., Ozturk, M., Sanchez, J., and Principe, J. (2007). Echo State Networks for Motor Control of Human ECoG Neuroprosthetics. *Neural Engineering, 2007. CNE '07. 3rd International IEEE/EMBS Conference on.*

[Hari et al., 1998] Hari, R., Forss, N., Avikainen, S., Kirveskari, E., Salenius, S., and Rizzolatti, G. (1998). Activation of human primary motor cortex during action observation: a neuromagnetic study. *Proceedings of the National Academy of Sciences of the United States of America.*

[Harris et al., 2011] Harris, A., Katyal, K., Para, M., and Thomas, J. (2011). Revolutionizing Prosthetics software technology. In *Systems, Man, and Cybernetics (SMC), 2011 IEEE International Conference on.*

[Heldman et al., 2006] Heldman, D. A., Wang, W., Chan, S. S., and Moran, D. W. (2006). Local field potential spectral tuning in motor cortex during reaching. *IEEE transactions on neural systems and rehabilitation engineering : a publication of the IEEE Engineering in Medicine and Biology Society.*

[Henle et al., 2011] Henle, C., Raab, M., Cordeiro, J. G., Doostkam, S., Schulze-Bonhage, A., Stieglitz, T., and Rickert, J. (2011). First long term in vivo study on subdurally implanted micro-ECoG electrodes, manufactured with a novel laser technology. *Biomedical microdevices*.

[Hermes et al., 2010] Hermes, D., Miller, K. J., Noordmans, H. J., Vansteensel, M. J., and Ramsey, N. F. (2010). Automated electrocorticographic electrode localization on individually rendered brain surfaces. *Journal of Neuroscience Methods*.

[Hill, 2008] Hill, J. (2008). BCPy2000. Retrieved from `http://bci2000.org/downloads/BCPy2000/BCPy2000.html`.

[Hinterberger et al., 2008] Hinterberger, T., Widman, G., Lal, T. N., Hill, J., Tangermann, M., Rosenstiel, W., Schölkopf, B., Elger, C., and Birbaumer, N. (2008). Voluntary brain regulation and communication with electrocorticogram signals. *Epilepsy & behavior : E&B*.

[Hluštík et al., 2001] Hluštík, P., Solodkin, A., Gullapalli, R. P., and Noll, D. C. (2001). Somatotopy in human primary motor and somatosensory hand representations revisited. *Cerebral . . .*.

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

[Hochberg et al., 2006] Hochberg, L. R., Serruya, M. D., Friehs, G. M., Mukand, J. A., Saleh, M., Caplan, A. H., Branner, A., Chen, D., Penn, R. D., and Donoghue, J. P. (2006). Neuronal ensemble control of prosthetic devices by a human with tetraplegia. *Nature*.

[Horch et al., 2011] Horch, K., Meek, S., Taylor, T. G., and Hutchinson, D. T. (2011). Object Discrimination With an Artificial Hand Using Electrical Stimulation of Peripheral Tactile and Proprioceptive Pathways With Intrafascicular Electrodes. *IEEE transactions on neural systems and rehabilitation engineering : a publication of the IEEE Engineering in Medicine and Biology Society*.

[Hotz-Boendermaker et al., 2008] Hotz-Boendermaker, S., Funk, M., Summers, P., Brugger, P., Hepp-Reymond, M.-C., Curt, A., and Kollias, S. S. (2008). Preservation of motor programs in paraplegics as demonstrated by attempted and imagined foot movements. *NeuroImage*.

[Jarosiewicz et al., 2008] Jarosiewicz, B., Chase, S. M., Fraser, G. W., Velliste, M., Kass, R. E., and Schwartz, A. B. (2008). Functional network reorganization during learning in a brain-computer interface paradigm. *Proceedings of the National Academy of Sciences*.

[Järveläinen et al., 2004] Järveläinen, J., Schürmann, M., and Hari, R. (2004). Activation of the human primary motor cortex during observation of tool use. *NeuroImage*.

[Jewett and WILLISTON, 1971] Jewett, D. L. and WILLISTON, J. S. (1971). Auditory-evoked far fields averaged from the scalp of humans. *Brain : a journal of neurology*.

[Juergens et al., 1999] Juergens, E., Guettler, A., and Eckhorn, R. (1999). Visual stimulation elicits locked and induced gamma oscillations in monkey intracortical-and EEG-potentials, but not in human EEG. *Experimental Brain Research*.

[Jung et al., 2008] Jung, J., Mainy, N., Kahane, P., Minotti, L., Hoffmann, D., Bertrand, O., and Lachaux, J.-P. (2008). The neural bases of attentive reading. *Human brain mapping*.

[Kalman, 1960] Kalman, R. E. (1960). A New Approach to Linear Filtering and Prediction Problems. *Journal of Fluids Engineering*.

[Kanas et al., 2014] Kanas, V. G., Mporas, I., Benz, H. L., Sgarbas, K. N., Bezerianos, A., and Crone, N. E. (2014). Joint Spatial-Spectral Feature Space Clustering for Speech Activity Detection from ECoG Signals. *Biomedical Engineering, IEEE Transactions on*.

[Kass et al., 2005] Kass, R. E., Ventura, V., and Brown, E. N. (2005). Statistical issues in the analysis of neuronal data. *Journal of Neurophysiology.*

[Kay and Marple, 1981] Kay, S. and Marple, S. J. (1981). Spectrum analysis—A modern perspective. In *Proceedings of the IEEE.*

[Kellis et al., 2010] Kellis, S., Miller, K., Thomson, K., Brown, R., House, P., and Greger, B. (2010). Decoding spoken words using local field potentials recorded from the cortical surface. *Journal of Neural Engineering.*

[Kelly et al., 2012] Kelly, J. W., Degenhart, A. D., Siewiorek, D. P., Smailagic, A., and Wang, W. (2012). Sparse linear regression with elastic net regularization for brain-computer interfaces. *Conference proceedings : Annual International Conference of the IEEE Engineering in Medicine and Biology Society IEEE Engineering in Medicine and Biology Society Conference.*

[Kelly et al., 2013] Kelly, J. W., Siewiorek, D. P., Smailagic, A., and Wang, W. (2013). Automated filtering of common-mode artifacts in multichannel physiological recordings. *IEEE transactions on bio-medical engineering.*

[Keshner, 1982] Keshner, M. V. (1982). 1/f noise. *Proceedings of the IEEE.*

[Kim et al., 2010] Kim, D. H., Viventi, J., Amsden, J. J., Xiao, J., and Vigeland, L. (2010). Dissolvable films of silk fibroin for ultrathin conformal bio-integrated electronics. *Nature materials.*

[Kim et al., 2008] Kim, S.-P., Simeral, J. D., Hochberg, L. R., Donoghue, J. P., and Black, M. J. (2008). Neural control of computer cursor velocity by decoding motor cortical spiking activity in humans with tetraplegia. *Journal of Neural Engineering.*

[Kim et al., 2011] Kim, S.-P., Simeral, J. D., Hochberg, L. R., Donoghue, J. P., Friehs, G. M., and Black, M. J. (2011). Point-and-Click Cursor Control With an Intracortical Neural Interface System by Humans With Tetraplegia. *IEEE transactions on neural systems and rehabilitation engineering : a publication of the IEEE Engineering in Medicine and Biology Society.*

[Kolarcik et al., 2012] Kolarcik, C. L., Bourbeau, D., Azemi, E., Rost, E., Zhang, L., Lagenaur, C. F., Weber, D. J., and Cui, X. T. (2012). In vivo effects of L1 coating on inflammation and neuronal health at the electrode-tissue interface in rat spinal cord and dorsal root ganglion. *Acta biomaterialia.*

[Koralek et al., 2012] Koralek, A. C., Jin, X., Long, J. D., Costa, R. M., and Carmena, J. M. (2012). Corticostriatal plasticity is necessary for learning intentional neuroprosthetic skills. *Nature.*

[Koyama et al., 2010] Koyama, S., Chase, S. M., Whitford, A. S., Velliste, M., Schwartz, A. B., and Kass, R. E. (2010). Comparison of brain-computer interface decoding algorithms in open-loop and closed-loop control. *Journal of Computational Neuroscience.*

[Kubánek et al., 2009] Kubánek, J., Miller, K. J., Ojemann, J. G., Wolpaw, J. R., and Schalk, G. (2009). Decoding flexion of individual fingers using electrocorticographic signals in humans. *Journal of Neural Engineering.*

[Lachaux et al., 2005] Lachaux, J.-P., George, N., Tallon-Baudry, C., Martinerie, J., Hugueville, L., Minotti, L., Kahane, P., and Renault, B. (2005). The many faces of the gamma band response to complex visual stimuli. *NeuroImage.*

[Lacourse et al., 2005] Lacourse, M. G., Orr, E. L. R., Cramer, S. C., and Cohen, M. J. (2005). Brain activation during execution and motor imagery of novel and skilled sequential hand movements. *NeuroImage.*

[Leuthardt et al., 2011] Leuthardt, E. C., Gaona, C., Sharma, M., Szrama, N., Roland, J., Freudenberg, Z., Solis, J., Breshears, J., and Schalk, G. (2011). Using the electrocorticographic speech network to control a brain-computer interface in humans. *Journal of Neural Engineering.*

[Leuthardt et al., 2006a] Leuthardt, E. C., Miller, K. J., Schalk, G., Rao, R. P. N., and Ojemann, J. G. (2006a). Electrocorticography-based brain computer interface–the Seattle experience. *IEEE transactions on neural systems and rehabilitation engineering : a publication of the IEEE Engineering in Medicine and Biology Society.*

[Leuthardt et al., 2006b] Leuthardt, E. C., Schalk, G., Moran, D., and Ojemann, J. G. (2006b). The emerging world of motor neuroprosthetics: a neurosurgical perspective. *Neurosurgery.*

[Leuthardt et al., 2004] Leuthardt, E. C., Schalk, G., Wolpaw, J. R., Ojemann, J. G., and Moran, D. W. (2004). A brain-computer interface using electrocorticographic signals in humans. *Journal of Neural Engineering.*

[Li et al., 2011] Li, Z., O'doherty, J. E., Lebedev, M. A., and Nicolelis, M. A. L. (2011). Adaptive decoding for brain-machine interfaces through Bayesian parameter updates. *Neural computation.*

[Linden and Sahani, 2003] Linden, J. and Sahani, M. (2003). Evidence Optimization Techniques for Estimating Stimulus-Response Functions.

[Logothetis et al., 2001] Logothetis, N. K., Pauls, J., Augath, M., Trinath, T., and Oeltermann, A. (2001). Neurophysiological investigation of the basis of the fMRI signal. *Nature.*

[Luo et al., 2011] Luo, X., Matranga, C., Tan, S., Alba, N., and Cui, X. T. (2011). Carbon nanotube nanoreservior for controlled release of anti-inflammatory dexamethasone. *Biomaterials.*

[Lynn et al., 2011] Lynn, A. D., Blakney, A. K., Kyriakides, T. R., and Bryant, S. J. (2011). Temporal progression of the host response to implanted poly(ethylene glycol)-based hydrogels. *Journal of biomedical materials research. Part A.*

[Magri et al., 2011] Magri, C., Mazzoni, A., Logothetis, N. K., and Panzeri, S. (2011). Optimal band separation of extracellular field potentials. *Journal of Neuroscience Methods.*

[Mahmoudi and Sanchez, 2011] Mahmoudi, B. and Sanchez, J. C. (2011). A Symbiotic Brain-Machine Interface through Value-Based Decision Making. *PLoS ONE.*

[Mainy et al., 2007] Mainy, N., Kahane, P., Minotti, L., Hoffmann, D., Bertrand, O., and Lachaux, J.-P. (2007). Neural correlates of consolidation in working memory. *Human brain mapping.*

[Marple Jr, 1987] Marple Jr, S. L. (1987). Digital spectral analysis with applications. *Englewood Cliffs.*

[Matelli et al., 1991] Matelli, M., Luppino, G., and Rizzolatti, G. (1991). Architecture of superior and mesial area 6 and the adjacent cingulate cortex in the macaque monkey. *Journal of Comparative Neurology.*

[McConnell et al., 2009] McConnell, G. C., Rees, H. D., Levey, A. I., Gutekunst, C.-A., Gross, R. E., and Bellamkonda, R. V. (2009). Implanted neural electrodes cause chronic, local inflammation that is correlated with local neurodegeneration. *Journal of Neural Engineering.*

[McFarland et al., 2000] McFarland, D. J., Miner, L. A., Vaughan, T. M., and Wolpaw, J. R. (2000). Mu and Beta Rhythm Topographies During Motor Imagery and Actual Movements. *Brain Topography.*

[McFarland et al., 2010] McFarland, D. J., Sarnacki, W. A., and Wolpaw, J. R. (2010). Electroencephalographic (EEG) control of three-dimensional movement. *Journal of Neural Engineering.*

[Mellinger et al., 2007] Mellinger, J., Schalk, G., Braun, C., Preissl, H., Rosenstiel, W., Birbaumer, N., and Kübler, A. (2007). An MEG-based brain-computer interface (BCI). *NeuroImage.*

[Menon et al., 1996] Menon, V., Freeman, W. J., Cutillo, B. A., Desmond, J. E., Ward, M. F., Bressler, S. L., Laxer, K. D., Barbaro, N., and Gevins, A. S. (1996). Spatio-temporal correlations in human gamma band electrocorticograms. *Electroencephalography and clinical Neurophysiology.*

[Merzenich et al., 1978] Merzenich, M. M., Kaas, J. H., Sur, M., and Lin, C. S. (1978). Double representation of the body surface within cytoarchitectonic areas 3b and 1 in "SI" in the owl monkey (Aotus trivirgatus). *The Journal of comparative neurology.*

[Miller et al., 2012] Miller, K. J., Hermes, D., Honey, C. J., Hebb, A. O., Ramsey, N. F., Knight, R. T., Ojemann, J. G., and Fetz, E. E. (2012). Human motor cortical activity is selectively phase-entrained on underlying rhythms. *PLoS computational biology.*

[Miller et al., 2007a] Miller, K. J., Leuthardt, E. C., Schalk, G., Rao, R. P. N., Anderson, N. R., Moran, D. W., Miller, J. W., and Ojemann, J. G. (2007a). Spectral changes in cortical surface potentials during motor movement. *The Journal of neuroscience : the official journal of the Society for Neuroscience.*

[Miller et al., 2007b] Miller, K. J., Makeig, S., Hebb, A. O., Rao, R. P. N., denNijs, M., and Ojemann, J. G. (2007b). Cortical electrode localization from X-rays and simple mapping for electrocorticographic research: The "Location on Cortex" (LOC) package for MATLAB. *Journal of Neuroscience Methods.*

[Miller et al., 2010] Miller, K. J., Schalk, G., Fetz, E. E., den Nijs, M., Ojemann, J. G., and Rao, R. P. N. (2010). Cortical activity during motor execution, motor imagery, and imagery-based online feedback. *Proceedings of the National Academy of Sciences of the United States of America.*

[Miller et al., 2009a] Miller, K. J., Sorensen, L. B., Ojemann, J. G., and den Nijs, M. (2009a). Power-law scaling in the brain surface electric potential. *PLoS computational biology.*

[Miller et al., 2009b] Miller, K. J., Zanos, S., Fetz, E. E., Den Nijs, M., and Ojemann, J. G. (2009b). Decoupling the Cortical Power Spectrum Reveals Real-Time Representation of Individual Finger Movements in Humans. *Journal of Neuroscience.*

[Moran, 2010] Moran, D. (2010). Evolution of brain-computer interface: action potentials, local field potentials and electrocorticograms. *Current opinion in neurobiology.*

[Moran and Schwartz, 1999] Moran, D. W. and Schwartz, A. B. (1999). Motor cortical representation of speed and direction during reaching. *Journal of Neurophysiology.*

[Moritz and Fetz, 2011] Moritz, C. T. and Fetz, E. E. (2011). Volitional control of single cortical neurons in a brain-machine interface. *Journal of Neural Engineering.*

[Morrell and RNS System in Epilepsy Study Group, 2011] Morrell, M. J. and RNS System in Epilepsy Study Group (2011). Responsive cortical stimulation for the treatment of medically intractable partial epilepsy. *Neurology.*

[Musallam et al., 2007] Musallam, S., Bak, M. J., Troyk, P. R., and Andersen, R. A. (2007). A floating metal microelectrode array for chronic implantation. *Journal of Neuroscience Methods.*

[Musallam et al., 2004] Musallam, S., Corneil, B. D., Greger, B., Scherberger, H., and Andersen, R. A. (2004). Cognitive control signals for neural prosthetics. *Science (New York, NY).*

[Muthukumaraswamy et al., 2004] Muthukumaraswamy, S. D., Johnson, B. W., and McNair, N. A. (2004). Mu rhythm modulation during observation of an object-directed grasp. *Brain research. Cognitive brain research.*

[Nagel, 1984] Nagel, J. (1984). Congestion control in IP/TCP internetworks. Retrieved from `http://tools.ietf.org/html/rfc896`.

[Nakanishi et al., 2013] Nakanishi, Y., Yanagisawa, T., Shin, D., Fukuma, R., Chen, C., Kambara, H., Yoshimura, N., Hirata, M., Yoshimine, T., and Koike, Y. (2013). Prediction of Three-Dimensional Arm Trajectories Based on ECoG Signals Recorded from Human Sensorimotor Cortex. *PLoS ONE.*

[National Instruments, 2014] National Instruments (2014). The NI TDMS file format.

[Nishitani and Hari, 2000] Nishitani, N. and Hari, R. (2000). Temporal dynamics of cortical representation for action. *Proceedings of the National Academy of Sciences of the United States of America*.

[Norton et al., 2005] Norton, L. W., Tegnell, E., Toporek, S. S., and Reichert, W. M. (2005). In vitro characterization of vascular endothelial growth factor and dexamethasone releasing hydrogels for implantable probe coatings. *Biomaterials*.

[Nunamaker and Kipke, 2010] Nunamaker, E. A. and Kipke, D. R. (2010). An alginate hydrogel dura mater replacement for use with intracortical electrodes. *Journal of biomedical materials research*.

[Nunez and Srinivasan, 2006] Nunez, P. L. and Srinivasan, R. (2006). *Electric Fields of the Brain*. Oxford University Press.

[Oostenveld, 2014] Oostenveld, R. (2014). The FieldTrip buffer for real-time access to EEG/MEG data. Retrieved from `http://fieldtrip.fcdonders.nl/development/realtime/buffer_overview`.

[Orsborn et al., 2012] Orsborn, A. L., Dangi, S., Moorman, H. G., and Carmena, J. M. (2012). Closed-loop decoder adaptation on intermediate time-scales facilitates rapid BMI performance improvements independent of decoder initialization conditions. *IEEE transactions on neural systems and rehabilitation engineering : a publication of the IEEE Engineering in Medicine and Biology Society*.

[Park and Pillow, 2011] Park, M. and Pillow, J. W. (2011). Receptive field inference with localized priors. *PLoS computational biology*.

[Pasley et al., 2012] Pasley, B. N., David, S. V., Mesgarani, N., Flinker, A., Shamma, S. A., Crone, N. E., Knight, R. T., and Chang, E. F. (2012). Reconstructing speech from human auditory cortex. *PLoS Biology*.

[Pei et al., 2011] Pei, X., Barbour, D. L., Leuthardt, E. C., and Schalk, G. (2011). Decoding vowels and consonants in spoken and imagined words using electrocorticographic signals in humans. *Journal of Neural Engineering*.

[Penfield, 1939] Penfield, W. (1939). *The Epilepsies: With a Note on Radical Therapy*. New England Journal of Medicine.

[Penfield and Boldrey, 1937] Penfield, W. and Boldrey, E. (1937). Somatic motor and sensory representation in the cerebral cortex of man as studied by electrical stimulation. *Brain : a journal of neurology*.

[Penfield and Jasper, 1954] Penfield, W. and Jasper, H. (1954). Epilepsy and the functional anatomy of the human brain.

[Penfield and Rasmussen, 1950] Penfield, W. and Rasmussen, T. (1950). The cerebral cortex of man; a clinical study of localization of function.

[Perge et al., 2013] Perge, J. A., Homer, M. L., Malik, W. Q., Cash, S., Eskandar, E., Friehs, G., Donoghue, J. P., and Hochberg, L. R. (2013). Intra-day signal instabilities affect decoding performance in an intracortical neural interface system. *Journal of Neural Engineering*.

[Perry and Bentin, 2009] Perry, A. and Bentin, S. (2009). Mirror activity in the human brain while observing hand movements: a comparison between EEG desynchronization in the mu-range and previous fMRI results. *Brain research*.

[Pfurtscheller and Lopes da Silva, 1999] Pfurtscheller, G. and Lopes da Silva, F. H. (1999). Event-related EEG/MEG synchronization and desynchronization: basic principles. *Clinical neurophysiology : official journal of the International Federation of Clinical Neurophysiology*.

[Picard et al., 2013] Picard, N., Matsuzaka, Y., and Strick, P. L. (2013). Extended practice of a motor skill is associated with reduced metabolic activity in M1. *Nature Neuroscience*.

[Pistohl et al., 2008] Pistohl, T., Ball, T., Schulze-Bonhage, A., Aertsen, A., and Mehring, C. (2008). Prediction of arm movement trajectories from ECoG-recordings in humans. *Journal of Neuroscience Methods*.

[Pistohl et al., 2012] Pistohl, T., Schulze-Bonhage, A., Aertsen, A., Mehring, C., and Ball, T. (2012). Decoding natural grasp types from human ECoG. *NeuroImage*.

[Polikov et al., 2005] Polikov, V. S., Tresco, P. A., and Reichert, W. M. (2005). Response of brain tissue to chronically implanted neural electrodes. *Journal of Neuroscience Methods*.

[Porro et al., 1996] Porro, C. A., Francescato, M. P., Cettolo, V., Diamond, M. E., Baraldi, P., Zuiani, C., Bazzocchi, M., and di Prampero, P. E. (1996). Primary motor and sensory cortex activation during motor performance and motor imagery: a functional magnetic resonance imaging study. *The Journal of neuroscience : the official journal of the Society for Neuroscience*.

[Preibisch et al., 2009] Preibisch, S., Saalfeld, S., and Tomancak, P. (2009). Globally optimal stitching of tiled 3D microscopic image acquisitions. *Bioinformatics*.

[Press et al., 2011] Press, C., Cook, J., Blakemore, S.-J., and Kilner, J. (2011). Dynamic modulation of human motor activity when observing actions. *Journal of Neuroscience*.

[Preul et al., 2003] Preul, M. C., Bichard, W. D., and Spetzler, R. F. (2003). Toward optimal tissue sealants for neurosurgery: use of a novel hydrogel sealant in a canine durotomy repair model. *Neurosurgery*.

[Rasmussen, 2006] Rasmussen, C. E. (2006). Mimimize.m. Retrieved from `http://learning.eng.cam.ac.uk/carl/code/minimize/minimize.m`.

[Rasmussen and Williams, 2006] Rasmussen, C. E. and Williams, C. K. I. (2006). *Gaussian Processes for Machine Learning*. The MIT Press.

[Ray et al., 2008a] Ray, S., Crone, N. E., Niebur, E., Franaszczuk, P. J., and Hsiao, S. S. (2008a). Neural correlates of high-gamma oscillations (60-200 Hz) in macaque local field potentials and their potential implications in electrocorticography. *The Journal of neuroscience : the official journal of the Society for Neuroscience*.

[Ray and Maunsell, 2011] Ray, S. and Maunsell, J. H. R. (2011). Different Origins of Gamma Rhythm and High-Gamma Activity in Macaque Visual Cortex. *PLoS Biology*.

[Ray et al., 2008b] Ray, S., Niebur, E., Hsiao, S. S., Sinai, A., and Crone, N. E. (2008b). High-frequency gamma activity (80-150Hz) is increased in human cortex during selective attention. *Clinical neurophysiology : official journal of the International Federation of Clinical Neurophysiology*.

[Renard et al., 2010] Renard, Y., Lotte, F., Gibert, G., Congedo, M., Maby, E., Delannoy, V., Bertrand, O., and L 233 cuyer, A. (2010). OpenViBE: An Open-Source Software Platform to Design, Test, and Use Brain–Computer Interfaces in Real and Virtual Environments. *Presence*.

[Robinson, 1982] Robinson (1982). A historical perspective of spectrum estimation. *Proceedings of the IEEE*.

[Rousche and Normann, 1998] Rousche, P. J. and Normann, R. A. (1998). Chronic recording capability of the Utah Intracortical Electrode Array in cat sensory cortex. *Journal of Neuroscience Methods*.

[Rouse and Moran, 2009] Rouse, A. G. and Moran, D. W. (2009). Neural adaptation of epidural electrocorticographic (EECoG) signals during closed-loop brain computer interface (BCI) tasks. *Conference proceedings : Annual International Conference of the IEEE Engineering in Medicine and Biology Society IEEE Engineering in Medicine and Biology Society Conference*.

[Rouse et al., 2011] Rouse, A. G., Stanslaski, S. R., Cong, P., Jensen, R. M., Afshar, P., Ullestad, D., Gupta, R., Molnar, G. F., Moran, D. W., and Denison, T. J. (2011). A chronic generalized bi-directional brain–machine interface. *Journal of Neural Engineering*.

[Rouse et al., 2013] Rouse, A. G., Williams, J. J., Wheeler, J. J., and Moran, D. W. (2013). Cortical adaptation to a chronic micro-electrocorticographic brain computer interface. *Journal of Neuroscience.*

[Rubehn et al., 2009] Rubehn, B., Bosman, C., Oostenveld, R., Fries, P., and Stieglitz, T. (2009). A MEMS-based flexible multichannel ECoG-electrode array. *Journal of Neural Engineering.*

[Rujitanaroj et al., 2013] Rujitanaroj, P.-o., Jao, B., Yang, J., Wang, F., Anderson, J. M., Wang, J., and Chew, S. Y. (2013). Controlling fibrous capsule formation through long-term down-regulation of collagen type I (COL1A1) expression by nanofiber-mediated siRNA gene silencing. *Acta biomaterialia.*

[Sadtler et al., 2014] Sadtler, P. T., Quick, K. M., Golub, M. D., Chase, S. M., Ryu, S. I., Tyler-Kabara, E. C., Yu, B. M., and Batista, A. P. (2014). Neural constraints on learning. *Nature.*

[Sajanti et al., 1999] Sajanti, J., Björkstrand, A. S., Finnilä, S., Heikkinen, E., Peltonen, J., and Majamaa, K. (1999). Increase of collagen synthesis and deposition in the arachnoid and the dura following subarachnoid hemorrhage in the rat. *Biochimica et biophysica acta.*

[Salinas and Abbott, 1994] Salinas, E. and Abbott, L. F. (1994). Vector reconstruction from firing rates. *Journal of Computational Neuroscience.*

[Sanders et al., 2000] Sanders, J. E., Stiles, C. E., and Hayes, C. L. (2000). Tissue response to single-polymer fibers of varying diameters: evaluation of fibrous encapsulation and macrophage density. *Journal of biomedical materials research.*

[Santhanam et al., 2006] Santhanam, G., Ryu, S. I., Yu, B. M., Afshar, A., and Shenoy, K. V. (2006). A high-performance brain-computer interface. *Nature.*

[Saxena et al., 2013] Saxena, T., Karumbaiah, L., Gaupp, E. A., Patkar, R., Patil, K., Betancur, M., Stanley, G. B., and Bellamkonda, R. V. (2013). The impact of chronic blood-brain barrier breach on intracortical electrode function. *Biomaterials.*

[Schalk, 2014] Schalk, G. (2014). BCI2000 Publications. Retrieved from `http://www.schalklab.org/publications`.

[Schalk et al., 2008a] Schalk, G., Brunner, P., Gerhardt, L. A., Bischof, H., and Wolpaw, J. R. (2008a). Brain-computer interfaces (BCIs): detection instead of classification. *Journal of Neuroscience Methods.*

[Schalk et al., 2007] Schalk, G., Kubánek, J., Miller, K. J., Anderson, N. R., Leuthardt, E. C., Ojemann, J. G., Limbrick, D., Moran, D., Gerhardt, L. A., and Wolpaw, J. R. (2007). Decoding two-dimensional movement trajectories using electrocorticographic signals in humans. *Journal of Neural Engineering.*

[Schalk and Leuthardt, 2011] Schalk, G. and Leuthardt, E. C. (2011). Brain-computer interfaces using electrocorticographic signals. *IEEE reviews in biomedical engineering.*

[Schalk et al., 2008b] Schalk, G., Leuthardt, E. C., Brunner, P., Ojemann, J. G., Gerhardt, L. A., and Wolpaw, J. R. (2008b). Real-time detection of event-related brain activity. *NeuroImage.*

[Schalk et al., 2004] Schalk, G., McFarland, D. J., Hinterberger, T., Birbaumer, N., and Wolpaw, J. R. (2004). BCI2000: a general-purpose brain-computer interface (BCI) system. *IEEE transactions on biomedical engineering.*

[Schalk et al., 2008c] Schalk, G., Miller, K. J., Anderson, N. R., Wilson, J. A., Smyth, M. D., Ojemann, J. G., Moran, D. W., Wolpaw, J. R., and Leuthardt, E. C. (2008c). Two-dimensional movement control using electrocorticographic signals in humans. *Journal of Neural Engineering.*

[Schendel et al., 2014] Schendel, A. A., Nonte, M. W., and Vokoun, C. (2014). The effect of micro-ECoG substrate footprint on the meningeal tissue response. *Journal of Neural . . . .*

[Schendel et al., 2013] Schendel, A. A., Thongpang, S., Brodnick, S. K., Richner, T. J., Lindevig, B. D. B., Krugner-Higby, L., and Williams, J. C. (2013). A cranial window imaging method for monitoring vascular growth around chronically implanted micro-ECoG devices. *Journal of Neuroscience Methods.*

[Schieber and Hibbard, 1993] Schieber, M. and Hibbard, L. (1993). How somatotopic is the motor cortex hand area? *Science (New York, NY).*

[Schwartz, 2004] Schwartz, A. B. (2004). Cortical neural prosthetics. *Annual Review of Neuroscience.*

[Schwartz et al., 2006] Schwartz, A. B., Cui, X. T., Weber, D. J., and Moran, D. W. (2006). Brain-controlled interfaces: movement restoration with neural prosthetics. *Neuron.*

[Schwartz et al., 1988] Schwartz, A. B., Kettner, R. E., and 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. *The Journal of neuroscience : the official journal of the Society for Neuroscience.*

[Serruya et al., 2002] Serruya, M. D., Hatsopoulos, N. G., Paninski, L., Fellows, M. R., and Donoghue, J. P. (2002). Instant neural control of a movement signal. *Nature.*

[Shimoda et al., 2012] Shimoda, K., Nagasaka, Y., Chao, Z. C., and Fujii, N. (2012). Decoding continuous three-dimensional hand trajectories from epidural electrocorticographic signals in Japanese macaques. *Journal of Neural Engineering.*

[Shin et al., 2012] Shin, D., Watanabe, H., Kambara, H., Nambu, A., and Isa, T. (2012). PLOS ONE: Prediction of Muscle Activities from Electrocorticograms in Primary Motor Cortex of Primates. *PLoS ONE.*

[Shoham et al., 2001] Shoham, S., Halgren, E., Maynard, E. M., and Normann, R. A. (2001). Motor-cortical activity in tetraplegics. *Nature.*

[Simeral et al., 2011] Simeral, J. D., Kim, S.-P., Black, M. J., Donoghue, J. P., and 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.*

[Slutzky et al., 2010] Slutzky, M. W., Jordan, L. R., Krieg, T., Chen, M., Mogul, D. J., and Miller, L. E. (2010). Optimal spacing of surface electrode arrays for brain-machine interface applications. *Journal of Neural Engineering.*

[So et al., 2014] So, K., Dangi, S., Orsborn, A. L., Gastpar, M. C., and Carmena, J. M. (2014). Subject-specific modulation of local field potential spectral power during brain–machine interface control in primates. *Journal of Neural Engineering.*

[Spinks et al., 2003] Spinks, R. L., Baker, S. N., Jackson, A., Khaw, P. T., and Lemon, R. N. (2003). Problem of dural scarring in recording from awake, behaving monkeys: a solution using 5-fluorouracil. *Journal of Neurophysiology.*

[Stence et al., 2001] Stence, N., Waite, M., and Dailey, M. E. (2001). Dynamics of microglial activation: A confocal time-lapse analysis in hippocampal slices. *Glia.*

[Strupler et al., 2007] Strupler, M., Pena, A. M., Hernest, M., Tharaux, P. L., Martin, J. L., Beaurepaire, E., and Schanne-Klein, M. C. (2007). Second harmonic imaging and scoring of collagen in fibrotic tissues. *Optics Express.*

[Su et al., 2011] Su, J., Todorov, M., Gonzales, H. P., Perkins, L., Kojouharov, H., Weng, H., and Tang, L. (2011). A Predictive Tool for Foreign Body Fibrotic Reactions Using 2-Dimensional Computational Model. *Open access bioinformatics.*

[Sudre et al., 2010] Sudre, G., Wang, W., Song, T., Kajola, M., Vinjamuri, R., Collinger, J. L., Degenhart, A. D., Bagic, A. I., and Weber, D. J. (2010). rtMEG: A Real-Time Software Toolbox for Brain-Machine Interfaces Using Magnetoencephelography. *17th International Conference on Biomagnetism Advances in Biomagnetism – Biomag2010.*

[Svirskis et al., 2010] Svirskis, D., Travas-Sejdic, J., Rodgers, A., and Garg, S. (2010). Electrochemically controlled drug delivery based on intrinsically conducting polymers. *Journal of controlled release : official journal of the Controlled Release Society.*

[Szarowski et al., 2003] Szarowski, D. H., Andersen, M. D., Retterer, S., Spence, A. J., Isaacson, M., Craighead, H. G., Turner, J. N., and Shain, W. (2003). Brain responses to micro-machined silicon devices. *Brain research.*

[Szurhaj et al., 2006] Szurhaj, W., Labyt, E., Bourriez, J.-L., Kahane, P., Chauvel, P., Mauguière, F., and Derambure, P. (2006). Relationship between intracerebral gamma oscillations and slow potentials in the human sensorimotor cortex. *European Journal of Neuroscience.*

[Tallon-Baudry et al., 2005] Tallon-Baudry, C., Bertrand, O., Hénaff, M.-A., Isnard, J., and Fischer, C. (2005). Attention modulates gamma-band oscillations differently in the human lateral occipital cortex and fusiform gyrus. *Cerebral cortex (New York, NY : 1991).*

[Taylor et al., 2002] Taylor, D. M., Tillery, S. I. H., and Schwartz, A. B. (2002). Direct cortical control of 3D neuroprosthetic devices. *Science (New York, NY).*

[Thongpang et al., 2011] Thongpang, S., Richner, T. J., Brodnick, S. K., Schendel, A., Kim, J., Wilson, J. A., Hippensteel, J., Krugner-Higby, L., Moran, D., Ahmed, A. S., Neimann, D., Sillay, K., and Williams, J. C. (2011). A micro-electrocorticography platform and deployment strategies for chronic BCI applications. *Clinical EEG and neuroscience.*

[Tkach et al., 2008] Tkach, D., Reimer, J., and Hatsopoulos, N. G. (2008). Observation-based learning for brain–machine interfaces. *Current opinion in neurobiology.*

[Toro et al., 1994] Toro, C., Deuschl, G., Thatcher, R., Sato, S., Kufta, C., and Hallett, M. (1994). Event-related desynchronization and movement-related cortical potentials on the ECoG and EEG. *Electroencephalography and Clinical Neurophysiology/Evoked Potentials Section.*

[Trautner et al., 2006] Trautner, P., Rosburg, T., Dietl, T., Fell, J., Korzyukov, O. A., Kurthen, M., Schaller, C., Elger, C. E., and Boutros, N. N. (2006). Sensory gating of auditory evoked and induced gamma band activity in intracranial recordings. *NeuroImage.*

[Truccolo et al., 2008] Truccolo, W., Friehs, G. M., Donoghue, J. P., and Hochberg, L. R. (2008). Primary motor cortex tuning to intended movement kinematics in humans with tetraplegia. *Journal of Neuroscience.*

[Turner et al., 1999] Turner, J. N., Shain, W., Szarowski, D. H., Andersen, M., Martins, S., Isaacson, M., and Craighead, H. (1999). Cerebral astrocyte response to micromachined silicon implants. *Experimental neurology.*

[Van Gompel et al., 2008] Van Gompel, J. J., Worrell, G. A., Bell, M. L., Patrick, T. A., Cascino, G. D., Raffel, C., Marsh, W. R., and Meyer, F. B. (2008). Intracranial electroencephalography with subdural grid electrodes: techniques, complications, and outcomes. *Neurosurgery.*

[Velliste et al., 2014] Velliste, M., Kennedy, S. D., Schwartz, A. B., Whitford, A. S., Sohn, J.-W., and McMorland, A. J. C. (2014). Motor cortical correlates of arm resting in the context of a reaching task and implications for prosthetic control. *Journal of Neuroscience.*

[Velliste et al., 2008] Velliste, M., Perel, S., Spalding, M. C., Whitford, A. S., and Schwartz, A. B. (2008). Cortical control of a prosthetic arm for self-feeding. *Nature*.

[Venkatraman and Carmena, 2011] Venkatraman, S. and Carmena, J. M. (2011). Active Sensing of Target Location Encoded by Cortical Microstimulation. *IEEE transactions on neural systems and rehabilitation engineering : a publication of the IEEE Engineering in Medicine and Biology Society*.

[Vetter et al., 2004] Vetter, R. J., Williams, J. C., Hetke, J. F., Nunamaker, E. A., and Kipke, D. R. (2004). Chronic neural recording using silicon-substrate microelectrode arrays implanted in cerebral cortex. *IEEE transactions on bio-medical engineering*.

[Viventi et al., 2011] Viventi, J., Kim, D.-H., Vigeland, L., Frechette, E. S., Blanco, J. A., Kim, Y.-S., Avrin, A. E., Tiruvadi, V. R., Hwang, S.-W., Vanleer, A. C., Wulsin, D. F., Davis, K., Gelber, C. E., Palmer, L., Van der Spiegel, J., Wu, J., Xiao, J., Huang, Y., Contreras, D., Rogers, J. A., and Litt, B. (2011). Flexible, foldable, actively multiplexed, high-density electrode array for mapping brain activity in vivo. *Nature Neuroscience*.

[Wadhwa et al., 2012] Wadhwa, J., Nair, A., and Kumria, R. (2012). Emulsion forming drug delivery system for lipophilic drugs. *Acta poloniae pharmaceutica*.

[Waldron and Schmiedeler, 2008] Waldron, K. and Schmiedeler, J. (2008). Kinematics. In *Springer Handbook of Robotics*.

[Wander and Rao, 2014] Wander, J. D. and Rao, R. P. N. (2014). Brain-computer interfaces: a powerful tool for scientific inquiry. *Current opinion in neurobiology*.

[Wang et al., 2007] Wang, W., Chan, S. S., Heldman, D. A., and Moran, D. W. (2007). Motor cortical representation of position and velocity during reaching. *Journal of Neurophysiology*.

[Wang et al., 2013a] Wang, W., Collinger, J. L., Degenhart, A. D., Tyler-Kabara, E. C., Schwartz, A. B., Moran, D. W., Weber, D. J., Wodlinger, B., Vinjamuri, R. K., Ashmore, R. C., Kelly, J. W., and Boninger, M. L. (2013a). An electrocorticographic brain interface in an individual with tetraplegia. *PLoS ONE*.

[Wang et al., 2010] Wang, W., Collinger, J. L., Perez, M. A., Tyler-Kabara, E. C., Cohen, L. G., Birbaumer, N., Brose, S. W., Schwartz, A. B., Boninger, M. L., and Weber, D. J. (2010). Neural interface technology for rehabilitation: exploiting and promoting neuroplasticity. *Physical medicine and rehabilitation clinics of North America*.

[Wang et al., 2009] Wang, W., Degenhart, A. D., Collinger, J. L., Vinjamuri, R., Sudre, G. P., Adelson, P. D., Holder, D. L., Leuthardt, E. C., Moran, D. W., Boninger, M. L., Schwartz, A. B., Crammond, D. J., Tyler-Kabara, E. C., and Weber, D. J. (2009). Human motor cortical activity recorded with Micro-ECoG electrodes, during individual finger movements. *Conference proceedings : Annual International Conference of the IEEE Engineering in Medicine and Biology Society IEEE Engineering in Medicine and Biology Society Conference*.

[Wang et al., 2011a] Wang, W., Degenhart, A. D., Sudre, G. P., Pomerleau, D., and Tyler-Kabara, E. C. (2011a). Decoding semantic information from human electrocorticographic (ECoG) signals. In *Engineering in Medicine and Biology Society,EMBC, 2011 Annual International Conference of the IEEE*.

[Wang et al., 2013b] Wang, Y., Papadimitrakopoulos, F., and Burgess, D. J. (2013b). Polymeric "smart" coatings to prevent foreign body response to implantable biosensors. *Journal of Controlled Release*.

[Wang et al., 2011b] Wang, Z., Ji, Q., Miller, K. J., and Schalk, G. (2011b). Prior knowledge improves decoding of finger flexion from electrocorticographic signals. *Frontiers in Neuroscience*.

[Weaver et al., 2014] Weaver, C. L., LaRosa, J. M., Luo, X., and Cui, X. T. (2014). Electrically controlled drug delivery from graphene oxide nanocomposite films. *ACS nano*.

[Wessberg et al., 2000] Wessberg, J., Stambaugh, C. R., Kralik, J. D., Beck, P. D., Laubach, M., Chapin, J. K., Kim, J., Biggs, S. J., Srinivasan, M. A., and Nicolelis, M. A. (2000). Real-time prediction of hand trajectory by ensembles of cortical neurons in primates. *Nature*.

[Wilson et al., 2006] Wilson, J. A., Felton, E. A., Garell, P. C., Schalk, G., and Williams, J. C. (2006). ECoG factors underlying multimodal control of a brain-computer interface. *IEEE transactions on neural systems and rehabilitation engineering : a publication of the IEEE Engineering in Medicine and Biology Society*.

[Wilson et al., 2010] Wilson, J. A., Mellinger, J., Schalk, G., and Williams, J. (2010). A procedure for measuring latencies in brain-computer interfaces. *IEEE transactions on bio-medical engineering*.

[Wodlinger et al., 2011] Wodlinger, B., Degenhart, A. D., Collinger, J. L., Tyler-Kabara, E. C., and Wang, W. (2011). The impact of electrode characteristics on electrocorticography (ECoG). In *Engineering in Medicine and Biology Society,EMBC, 2011 Annual International Conference of the IEEE*.

[Wolpaw et al., 2002] Wolpaw, J. R., Birbaumer, N., McFarland, D. J., Pfurtscheller, G., and Vaughan, T. M. (2002). Brain-computer interfaces for communication and control. *Clinical neurophysiology : official journal of the International Federation of Clinical Neurophysiology*.

[Wolpaw and McFarland, 1994] Wolpaw, J. R. and McFarland, D. J. (1994). Multichannel EEG-based brain-computer communication. *Electroencephalography and clinical Neurophysiology*.

[Wolpaw and McFarland, 2004] Wolpaw, J. R. and McFarland, D. J. (2004). Control of a two-dimensional movement signal by a noninvasive brain-computer interface in humans. *Proceedings of the National Academy of Sciences of the United States of America*.

[Wong et al., 2009] Wong, C. H., Birkett, J., Byth, K., Dexter, M., Somerville, E., Gill, D., Chaseling, R., Fearnside, M., and Bleasel, A. (2009). Risk factors for complications during intracranial electrode recording in presurgical evaluation of drug resistant partial epilepsy. *Acta Neurochirurgica*.

[Wu et al., 2006] Wu, W., Gao, Y., Bienenstock, E., Donoghue, J. P., and Black, M. J. (2006). Bayesian population decoding of motor cortical activity using a Kalman filter. *Neural computation*.

[Xu et al., 2007] Xu, H.-T., Pan, F., Yang, G., and Gan, W.-B. (2007). Choice of cranial window type for in vivo imaging affects dendritic spine turnover in the cortex. *Nature Neuroscience*.

[Yamakawa et al., 2010] Yamakawa, T., Aou, S., Ishizuka, S., Suzuki, M., and Fujii, M. (2010). Subdural electrode array manipulated by a shape memory alloy guidewire for minimally-invasive electrocorticogram recording. *World Automation Congress (WAC), 2010*.

[Yanagisawa et al., 2011] Yanagisawa, T., Hirata, M., Saitoh, Y., Goto, T., Kishima, H., Fukuma, R., Yokoi, H., Kamitani, Y., and Yoshimine, T. (2011). Real-time control of a prosthetic hand using human electrocorticography signals. *Journal of Neurosurgery*.

[Yanagisawa et al., 2012] Yanagisawa, T., Hirata, M., Saitoh, Y., Kishima, H., Matsushita, K., Goto, T., Fukuma, R., Yokoi, H., Kamitani, Y., and Yoshimine, T. (2012). Electrocorticographic control of a prosthetic arm in paralyzed patients. *Annals of Neurology*.

[Yeager et al., 2008] Yeager, J. D., Phillips, D. J., Rector, D. M., and Bahr, D. F. (2008). Characterization of flexible ECoG electrode arrays for chronic recording in awake rats. *Journal of Neuroscience Methods*.

[Zhang et al., 2011] Zhang, J., Sudre, G., Li, X., Wang, W., Weber, D. J., and Bagic, A. (2011). Task-related MEG source localization via discriminant analysis. *Conference proceedings : Annual International Conference of the IEEE Engineering in Medicine and Biology Society IEEE Engineering in Medicine and Biology Society Conference*.

[Zhuang et al., 2010] Zhuang, J., Truccolo, W., Vargas-Irwin, C., and Donoghue, J. P. (2010). Decoding 3-D reach and grasp kinematics from high-frequency local field potentials in primate primary motor cortex. *IEEE transactions on bio-medical engineering*.