

Neural correlates of learning and intent during human brain-computer interface use

Jeremiah D. Wander

A dissertation

Submitted in partial fulfillment of the

Requirements for the degree of

Doctor of Philosophy

University of Washington

2015

Reading Committee:

Jeffrey Ojemann, Co-chair

Rajesh Rao, Co-chair

Eric Chudler

Program Authorized to Offer Degree:

Department of Bioengineering

©Copyright 2015
Jeremiah D. Wander

University of Washington

Abstract

Neural correlates of learning and intent during human brain computer-interface use

Jeremiah D. Wander

Co-chairs of the Supervisory Committee:

Associate Professor Rajesh Rao

Computer Science and Engineering

Professor and Division Chief Jeffrey Ojemann

Neurological Surgery

Brain-computer interface (BCI) technologies can potentially be used restore function in patients with severe motor disorders, however, BCI devices currently do not perform well enough to warrant the risk or expense relative to other treatments. This is due in part to limitations of current BCI architectures, which utilize signals from a relatively small portion of the cortex and exclusively rely on a fixed mapping between neural activity and the output device. In contrast, when executing native motor function, the nervous system invokes a sophisticated series of bottom-up and top-down modulators that dynamically change the relationship between cortical function and motor output, based on an individual's capability, task demands, attentional focus and numerous other factors. Characterization of the neural correlates of these higher-order cognitive facets of BCI use is a critical first step in the development of such systems. **The work below focuses on identifying the distributed neural correlates of BCI skill acquisition and goal-oriented task execution, and leveraging these signals to improve BCI performance.**

We identified multiple cortical regions that become very active during novice BCI use, and are subsequently less active in the experienced user. Activity changes within these regions suggest distributed cortical processing, but could also be explained by nonspecific co-activation, so it was then necessary to show that these regions also interact in a meaningful way. To this end, we demonstrated that during BCI use there are high-frequency amplitude-amplitude interactions taking place on local spatial scales and non-linear low-frequency to high-frequency phase-phase interactions covering larger cortical distances. Lastly, we characterized neural correlates of a user's intended action immediately before and during BCI use, and trained a machine learning-based system to identify these activity patterns and leverage them in a hierarchical BCI.

These findings are directly applicable to BCI design. Next-generation BCI architectures will include signals from multiple cortical regions to allow for dynamic device control strategies. Furthermore, by leveraging BCI as a platform for scientific inquiry, we have been able to develop our understanding of the networks involved in acquisition and execution of the BCI skill, and the neural mechanisms of interaction enabling communication across these networks. Understanding these relationships is at the core of understanding the tremendous adaptive capability of the nervous system, and successfully translating brain activity into action.

1 Acknowledgements

There are a number of people whom I would like to acknowledge for the vital role they played in my graduate education and the generation of this document.

First, I want to thank Jeffrey Ojemann and Rajesh Rao as well as the other members of my supervisory committee – Eberhard Fetz, Eric Chudler, Colin Studholme and Adrienne Fairhall – for their guidance throughout the graduate school process.

I also want to thank the members of the Gridlab that I was fortunate enough to work with: Lise Johnson, Tim Blakely, Kurt Weaver, Jared Olson, Felix Darvas, Devapratim Sarma, Kai Miller, Hai Sun, David Su, Andrew Ko, Kaitlyn Casimo and James Wu. In particular I would like to recognize Tim, Lise, Kurt, Jared and Felix for the time they each took to serve as mentors for me.

My work is inherently collaborative and spans many disciplines, so I would also like to thank Stavros Zanos, Charlie Matlack, Brian Mogen, and the members of the University of Washington Regional Epilepsy Center and the Center for Sensorimotor Neural Engineering.

Most importantly, I would like to thank my family, and especially my wife, Leslie, for the unwavering and vital support they offered me throughout this process.

2 Contents

2.1 Table of contents

1	Acknowledgements.....	4
2	Contents.....	5
2.1	Table of contents	5
2.2	List of figures.....	7
3	Introduction	10
3.1	The clinical need	10
3.2	Brain-computer interfaces.....	11
3.3	Relevant motor physiology and neurophysiology	15
3.3.1	Primary motor cortex.....	15
3.3.2	Non-M1 frontal cortices.....	17
3.3.3	Parietal and temporal cortices.....	18
3.3.4	Fronto-parietal connectivity	19
3.4	Electrophysiological models for BCI.....	20
3.4.1	ECoG BCI.....	23
3.4.2	Disadvantages of ECoG BCI	24
3.4.3	Advantages of ECoG BCI.....	25
3.5	Distributed cortical processing during BCI use	25
4	General materials and methods	27

4.1	The standard BCI model.....	27
4.1.1	Components of a BCI.....	27
4.2	Data collection and signal processing.....	29
5	Distributed cortical activity patterns during BCI use.....	33
5.1	Introduction and background.....	33
5.2	Materials and methods.....	35
5.3	Results.....	39
5.4	Discussion.....	53
5.5	Related publications and presentations.....	56
6	Multi-site cortical interactions during BCI use.....	57
6.1	Introduction and background.....	57
6.2	Materials and methods.....	60
6.3	Results.....	67
6.4	Discussion.....	75
6.5	Related publications and presentations.....	81
7	Neural representation of intention during BCI use.....	83
7.1	Introduction and background.....	83
7.2	Materials and methods.....	88
7.2.1	Retrospective RJB studies.....	89
7.2.2	GoalBCI studies.....	94

7.3	Results.....	102
7.3.1	Retrospective RJB studies	102
7.3.2	GoalBCI Task.....	111
7.4	Discussion.....	126
7.5	Related publications and presentations	132
8	Conclusions	133
9	References	140

2.2 List of figures

Figure 1	- Lateral view of the human brain.....	16
Figure 2	- Overview of the RJB task	31
Figure 3	- Spatial distribution of electrode coverage.....	36
Figure 4	- Behavioral performance	40
Figure 5	- Behavioral analysis of cursor trajectories measured against inferred intended trajectories	41
Figure 6	- Run-by-Run task performance	42
Figure 7	- Cortex-wide activity during BCI use and modeling of early vs. late activity patterns	45
Figure 8	- Time-by-trial HG activation.....	48
Figure 9	- Time-by-trial HG activation for individual subjects.....	49
Figure 10	- Cortex-wide activation changes.....	50
Figure 11	- Changes in mu-beta activity patterns	52
Figure 12	- Changes in HG activity patterns observed with alternative data segmentation method	53
Figure 13	- Post-hoc analysis workflow.....	62

Figure 14 - Behavioral performance for all subjects included in interaction analysis	67
Figure 15 - Exemplar plots (from two subjects S2 and S5) showing alignment of up trials by HG onset... 68	
Figure 16 - Exemplar response-locked STWC maps.....	69
Figure 17 - Spatial distribution of significant STWC interactions.....	70
Figure 18 - Spatial comparison of non-CTL electrodes involved in HG STWC interactions for response-locked and cue-locked trials	71
Figure 19 - Group-average of significant bPLV interactions for up and down targets	73
Figure 20 - Changes in STWC and bPLV interactions over time.....	74
Figure 21 - Comparison of spatial distribution of significant STWC and bPLV interactions	75
Figure 22 - Proposed linked-oscillator model to generate cross-frequency coupling.....	81
Figure 23 - Depiction of the GoalBCI task and block randomized design	96
Figure 24 - Depiction of the eight possible targets presented to the user during the GoalBCI task.....	97
Figure 25 - Individual task performance in the retrospective analysis of the RJB data for the purposes of goal inference	103
Figure 26 - Average Strength of intent representation in the RJB data set as a function of time, stratified across frequency	104
Figure 27 - Individual spectrotemporal significance maps for the RJB data set.....	105
Figure 28 - Spatial distribution of electrodes showing significant representation of intent during the cue period.....	107
Figure 29 - Average classification performance across all subjects for the instantaneous and aggregate goal inference approaches.....	108
Figure 30 - Individual classification performance for both the instantaneous (a) and aggregate (b) approaches.....	109
Figure 31 - Summary of post-hoc trial biasing analyses	111

Figure 32 - GoalBCI behavioral performance as measured by (a) hitrate and (b) ISE	113
Figure 33 – Individual spectrotemporal significance maps for the GoalBCI data set	115
Figure 34 - time course of the strength of direction representation across all nine GoalBCI subjects during the late cue and early feedback phases	116
Figure 35 – Impact of real-time goal inference on Hit rate and ISE.....	117
Figure 36 – Impact of classification correctness on behavioral outcomes in biased trials.....	117
Figure 37 – Classification results in 10-fold cross-validated post-hoc classification analyses.....	118
Figure 38 – Classification results in 10-fold cross-validated post-hoc classification analyses for the single visual control subject	119
Figure 39 - Velocity results from multiple regression model.....	121
Figure 40 – Error results from multiple regression model.....	123
Figure 41 – Interaction (velocity x error) results from multiple regression model	124
Figure 42 – Target location results from multiple regression model	125
Figure 43 - Representation of significant relationships in the multiple regression model by Brodmann area	126

3 Introduction

There exist a number of neuromuscular disorders and physical impairments that impact an individual's ability to perform motor functions. This impairment is manifest as a loss of ability to dexterously move about in or manipulate one's environment, a loss of ability to communicate, or both. Because of ongoing advancements in robotics, machine learning, wearable computing, biosignal acquisition and processing and electrode design, the field of neural engineering has a tremendous opportunity to develop brain-computer interface (BCI) technologies that can be used to restore these lost functions. However, BCI devices currently do not perform well enough to replace the current state of the art treatments. For BCI devices that are used to continuously control an end effector, such as a computer cursor or a robotic arm, this limitation is due in part to the ubiquitous primary motor cortex-centric BCI design patterns and an incomplete understanding of the distributed cortical processing that enables learning and execution of the neuroprosthetic skill.

3.1 The clinical need

Neuromuscular disorders are a broad class of ailments that cause loss of function through one or more of the following means: (1) by disrupting the central nervous system's capability to generate an appropriate motor command, (2) by disrupting the flow of information from the central to the peripheral nervous system, or (3) by disrupting the peripheral nervous system's ability to effectively carry out that motor command. In its advanced stages, amyotrophic lateral sclerosis (ALS) may be one of the most severe of these disorders, but there are numerous others including traumatic brain injury, stroke and spinal cord injury. These disorders affect nearly two million people in the United States, and many more worldwide (Murray & Lopez, 1996). Severity varies from patient to patient, ranging from mild impairment of motor function to complete loss of voluntary muscular control, an affliction that is referred to as locked-in syndrome (Smith & Delargy, 2005).

Current treatments for neuromuscular disorders vary widely in concordance with the severity of the ailment, but can be separated into two primary groups. In the event that limited muscular function remains (e.g., some cases of ischemic stroke), *rehabilitative* physical therapy (Pascual-Leone et al., 2005) is employed to restore as much of the lost function as is possible. Experimental enhancement of physical therapy via transcranial magnetic stimulation (TMS) (Harvey & Nudo, 2007; Plautz et al., 2003), transcranial direct-current stimulation (tDCS) (Schlaug et al., 2008), and direct cortical stimulation (DCS) (Adkins et al., 2006, 2008; Adkins-Muir & Jones, 2003; Huang et al., 2008; Levy et al., 2008; Plow et al., 2009) is also being explored.

In cases where rehabilitative therapies are not an option, such as in the advanced stages of ALS, other means have been explored to replace lost function such as eye-trackers, blink interfaces (Doble et al., 2003), and assistant-mediated communication (Wu & Voda, 1985). Though these strategies are effective at providing patients with some portion of their original function, they leave much to be desired in terms of complete restoration.

3.2 Brain-computer interfaces

In response to this clinical need, over the past two decades, many researchers have been investigating direct interface to the nervous system as a means to create rehabilitative and assistive devices that provide patients with a significantly greater portion of their original function than the current standard of care (for review see Green & Kalaska, 2011; Moran, 2010; Wander & Rao, 2014; Wolpaw et al., 2002). Researchers are attempting to interface to the nervous system at a number of different levels, including the cerebral cortex, the spinal cord, and at the neuromuscular junction; the version of these devices that interface to the cortex are called BCIs or alternatively brain-machine interfaces. In the context of restoration of lost motor output, the proposed idea of an assistive BCI is that by recording neural activity ‘upstream’ of where the ailment has impacted the nervous system, and decoding that activity to

determine the patient's intended action, an external device can be driven to carry out that action. Three primary testbeds for BCI research are cursor control and communication in tetraplegic and locked-in patients (Hochberg et al., 2006; Kubler & Kotchoubey, 2001) and neural control of prosthetic arms (Carmena et al., 2003; Pfurtscheller et al., 2000; Velliste et al., 2008). This patient population has been concentrated on by many researchers as an appropriate short-term target as the prospect of restoration of capability to communicate to locked-in patients may soon outweigh the risk associated with the current invasiveness of many BCI electrode implantation procedures (Gilja et al., 2011; Moran, 2010).

Though BCI is a non-specific term that can be applied to a variety of system architectures ranging from BCIs for motor control to deep brain stimulators for the treatment of clinical depression, in this section we provide a brief review of the state of the art of the former, as it is relevant to the clinical need described above. In discussing BCIs for motor control, we will consider only BCI architectures that leverage volitional modulation of neural activity for the manipulation of an end effector.

In foundational work, Fetz and colleagues demonstrated that, when given feedback, the brain could learn to volitionally modulate the activity of single neurons (Fetz, 1969). In this experiment, the estimates of individual firing rates of a small number of cortical neurons were mapped directly to a simple visual feedback device, and the subject developed the capacity to modulate these firing rates, effectively controlling a simple, one-dimensional (1-D) BCI. A similar feat was accomplished a number of years later by Wolpaw and colleagues using a non-invasive approach, mapping spectral (μ -beta) changes recorded at the scalp to vertical control of a cursor on a computer monitor (Wolpaw et al., 1991). Since these initial BCI experiments, the primary investigative and engineering push in the field has been to increase the dexterity, robustness, and clinical viability of these devices. The current state of the art for BCI depends greatly on the electrophysiological signal being used to run the BCI. Aggregate activity from large numbers of neurons (field potentials) have been used to control end effectors in up to 3 dimensions (McFarland et al., 2010; Wang et al., 2013) and the high-bar is currently set by Collinger

and colleagues who trained a human user to control a seven degree of freedom (DOF) robotic arm using neural signals recorded with a microelectrode implant (Collinger et al., 2013). Unfortunately, direct comparison across these and other experiments is extremely difficult because of inconsistencies in task paradigms and evaluation metrics (Thomas et al., 2013).

There are a number of impediments to the long-term clinical deployment of brain-computer interfaces, including host-response to implanted electrodes, long-term power requirements, and wireless communication of neural signals to remote devices. One of the largest road-blocks in the development of these devices is simply that their performance relative to the current standard of care is not good enough to warrant the risk and cost to the patient (Gilja et al., 2011).

As a result, there is a collective push in the field to increase BCI performance in applications that restore a patient's ability to communicate or to control external devices. There are a number of approaches being investigated; the two most common being an increase of coverage density in motor cortex (i.e. higher number of input channels per unit area) and development of more sophisticated decoding algorithms (i.e. extracting more robust signals from the input channels). Recent work by the Schwartz (Wang et al., 2013) and Shenoy (Gilja et al., 2012) groups, respectively, provides examples of developments in these two directions. Other groups are working in a different, but potentially complimentary direction, constructing BCI systems that maximize the performance gains made possible by the tremendous adaptive capability of the brain (Orsborn et al., 2012).

An important point to note is that the vast majority of BCIs harness signals recorded solely from the primary motor cortex (M1) or premotor cortex (PMv / PMd) (Blakely et al., 2009; Simeral et al., 2011; Wang et al., 2013; Wolpaw et al., 1991). A noteworthy exception to this is the line of research being pursued by Andersen and colleagues that utilizes activity patterns in posterior parietal cortex (PPC) (Mulliken et al., 2008), though this approach still utilizes firing rate changes in a single cortical area to

achieve BCI control. We know, however, from a vast body of electrophysiological and imaging motor studies, that the generation of dexterous, goal-directed movement is a coordinated effort on the part of multiple cortical areas, and that the functional roles being carried out in these areas are diverse. Correspondingly, we posit that BCI architectures may benefit significantly by incorporating multiple streams of task-relevant information from a variety of cortical sources. An example of this would be a BCI that simultaneously decodes both the higher-level goal an individual is trying to achieve from one set of cortical structures and the low-level BCI motor commands that they are using and utilizes both channels of information to improve BCI performance. This notion of hybrid BCI that employs control signals at multiple levels of abstraction has been implemented previously (Cheung et al., 2012; Shanechi et al., 2013), but has yet to reach its full potential.

Controlling a BCI with volitional modulation of neural activity is intriguing in that it is arguably a non-motor or quasi-motor task. When the BCI is being driven by signals recorded from M1, one might expect that the same cortical (Jenkins et al., 1994; Schlaug et al., 1994) and subcortical areas (Hikosaka et al., 2002) that are involved in visuomotor tasks to be involved in the control of a BCI as well. However, the two tasks are significantly different in that, with the exception of visual feedback, the afferent activity that is critical to successful acquisition of traditional motor tasks is not present during the use of a BCI. Furthermore, it has been previously shown that humans can control a BCI using expressly non-motor imagery (Vansteensel et al., 2010) and non-human primates (NHPs) can do so using 'arbitrary' neurons (Fetz, 1969). Correspondingly, the first step in the process in the development of BCIs that leverage distributed activity patterns is to understand how the various areas in the brain work together to enable learning and continued execution of the neuroprosthetic skill.

3.3 Relevant motor physiology and neurophysiology

An extensive review of neurophysiology and motor physiology is beyond the scope of this document. For additional information, the reader is encouraged to review the seminal text by Kandel and Schwartz (Kandel et al., 2012). However, in an effort to provide context for the experimental designs, results, and discussions that follow, we will provide a brief summary of the cortical structures relevant to BCI use, and the underlying neurophysiology that is leveraged in the electrophysiological recording models described below.

3.3.1 Primary motor cortex

The cerebral cortex is spatially organized into areas that have been linked with specific functions, initially through post-mortem lesion studies and electrical stimulation, and now through *in vivo* functional imaging and electrophysiology. When locating cortical areas associated with motor outputs from the body, one of the most important cortical landmarks to identify is the central sulcus, which starts at the lateral fissure and travels superiorly and posteriorly. It separates the frontal lobe and the parietal lobe. Canonically, the gyrus anterior to the central sulcus, called the pre-central gyrus, is referred to as M1, though cytoarchitectonic delineations specify that M1 comprises only the posterior half of the lateral portion of the pre-central gyrus (M1 corresponds to Brodmann Area [BA] 4). The majority of descending afferent motor tracts that carry neural activity that will eventually result in motor output originate in M1.




Figure 1 - Lateral view of the human brain, the central sulcus is highlighted in red. Image is public domain, source: <http://commons.wikimedia.org>

As was determined in the 1930s (Penfield & Boldrey, 1937), motor cortex itself is somatotopically organized such that the motor outputs of the body are represented in an orderly fashion across M1. This can be seen in the image of the motor homunculus, below. Another feature of the motor representation on M1 that is important to the use of these neural populations for BCIs is that the amount of cortical area dedicated to a specific part of the body is related to the degree and precision of motor control executed in that part. Consequently, the motor cortical representation of face and hand movements is very large in humans and non-human primates (NHPs).

This large representation of hand and face movement, combined with the somatotopic organization of M1 is one of the primary reasons that it is often selected as a target site for BCI, however there are other critical reasons that should also be noted. First, layer V of M1 contains the cell bodies of many large pyramidal neurons, from which extracellular recordings can be readily taken, and which (because of the orientation of their dendritic arbor and their columnar organization) generate relatively discernible cortical and scalp potentials. Second, an extensive foundation of motor electrophysiology studies has demonstrated strong directional tuning between M1 pyramidal neurons and position and velocity of distal limbs, allowing for BCI decoders to be seeded with parameters derived from neural recordings taken during overt motor movement.

3.3.2 Non-M1 frontal cortices

Directly anterior to M1 and extending from the lateral sulcus to the superior frontal sulcus is the premotor cortex (BA 6), which is further divided into dorsal and ventral areas (PMd and PMv). Superior to this and wrapping around to the medial aspect of the brain, stopping at the cingulate sulcus is the supplementary motor area (SMA); coverage of this portion of the brain is exceedingly rare in an electrocorticographic (ECoG) epilepsy model and is thus not discussed extensively in the remainder of this document. Premotor areas, on the other hand, are commonly covered during epilepsy monitoring and thus provide an excellent first candidate region for BCI-relevant extra-M1 neural activity.

PMv has been studied extensively in NHPs, and is traditionally held to be involved in preparatory motor activity related to hand motion (Rizzolatti & Luppino, 2001; Weinrich & Wise, 1982) , and during the movement for both reaching and grasping actions (Xiao et al., 2006). Additionally, so-named “mirror neurons” in PMv have been seen to respond to observation of relevant movements (Rizzolatti et al., 1996).

Though also active during preparation (Mauritz & Wise, 1986), PMd has been demonstrated previously to be more associated with trajectory planning and control of the more proximal musculature involved in reach motions (Pesaran et al., 2006). In an NHP model, cells have been found in PMd that respond to changes in characteristics of the target of a reach (Shen & Alexander, 1997), which, when taken with the above, implies a role for PMd in the translation of a reach target to a motor command through the use of internal models.

Collectively, the area of brain anterior to the premotor areas is called prefrontal cortex (PFC), further subdivided into a number of regions that vary, to some degree depending on the body of literature in which they are being discussed. In general PFC is considered the cortical center for executive control (Fuster, 2000) and a critical region for orchestrating coordinated complex behaviors to achieve higher-

level goals (Miller & Cohen, 2001). Various regions of the PFC have demonstrated involvement in complex, goal-driven behavior in both human (Rudorf & Hare, 2014) and NHP (Kobayashi et al., 2007) models. The dorsolateral PFC (dlPFC; BAs 8,9 & 46) is a specific sub-region of PFC of particular relevance to BCI control, based on previous demonstrations of involvement in higher order motor control (for review see Damasio et al., 1996) and spatial working memory (Barbey et al., 2013; Funahashi et al., 1989; Goldman-Rakic, 1996).

3.3.3 Parietal and temporal cortices

Posterior to the central sulcus is primary somatosensory cortex (SC; BAs 1-3). Though it is not discussed extensively in the context of this document, it is one of the sensory processing areas – along with auditory and visual processing areas – that provides sensory information to its more caudal neighbor, the posterior parietal cortex (PPC; BAs 5 and 7).

PPC is a group of cortical regions, located, as the name implies, in the posterior portion of the parietal lobe. It is bisected by the intraparietal sulcus (IPS) into the superior and inferior parietal lobules (SPL and IPL, respectively). The PPC on the whole is an associative cortical region implicated in planning and online control of visually guided movements (Buneo & Andersen, 2006a; Mulliken et al., 2008). A significant amount of work has gone into understanding the subregions of PPC and their particular functional roles in the variety of motor movements carried out by primates (reviewed in Vesia & Crawford, 2012). It is worth noting that the terminology between the two species differs and can be confusing. In NHPs, we commonly discuss three major regions the PPC: the parietal reach region (PRR), the lateral intraparietal area (LIP), and the anterior intraparietal area (AIP). These areas are classically associated with reaching, saccades, and grasping respectively (Rizzolatti et al., 1998). The human homologues are slightly less well defined, but they are the mid-posterior IPS (mIPS), mIPS and the parieto-occipital sulcus (SPOC), and the anterior IPS (aIPS), respectively. Recent imaging studies have

begun to demonstrate similar roles for these regions, though the functional separations do not appear to be as strict (Gallivan et al., 2011, 2013) .

Inferior to the IPL, surrounding the supramarginal gyrus, where the temporal and parietal lobes meet is a region called the temporoparietal junction (TPJ; BAs 22, 39 and 40). Though not explicitly involved in reach and grasp planning, the TPJ is implicated in a number of cognitive functions. Of particular interest to its potential role during BCI execution is evidence of involvement in bottom-up modulation of sensory inputs relevant for target selection (Geng & Mangun, 2011).

Anterior and inferior to this is another cortical region of interest, namely the posterior portion of the superior temporal gyrus (STG). Classically thought to be involved primarily in emotional, social and language-related processing (Friederici & Rueschemeyer, 2003; Radua et al., 2010), the STG has recently been demonstrated to be involved in multiple cognitive processes that may be highly relevant to neuroprosthetic control. More specifically, the posterior portion of the STG has been demonstrated to be activated during selective processing of visual stimuli (Hopfinger et al., 2000) and during observation of geometric shapes that were following goal-directed trajectories (Schultz et al., 2004).

3.3.4 Fronto-parietal connectivity

The classical viewpoint on parieto-frontal networks is that in both areas different functions are spatially segregated and the pathways between parietal regions and their frontal counterparts are discrete and separate (Rizzolatti et al., 1998), an idea that was supported through the study of anatomical connectivity between frontal and parietal regions (Tanné-Gariépy et al., 2002; Wise & Boussaoud, 1997). In these studies the authors demonstrated two segregated structural pathways connecting fronto-parietal networks; specifically noting pathways between PMv and AIP, and between PMd and more posterior parietal areas. This structural model supports Rizzolatti *et al.*'s functional model of area specific specialization, with PRR-to-PMd connectivity playing a role in reaching movements (Wise &

Boussaoud, 1997) and AIP-to-PMv connectivity playing a role in grasping (Rizzolatti & Luppino, 2001). Despite structural evidence for separate fronto-parietal networks, this notion has since given way, to some degree, to a more overlapping or even multi-functional model of parietal and frontal visuomotor processing (Davare et al., 2011). As ever-improving capability of functional MRI (fMRI) allows for study of the human correlates of macaque PPC areas, we now have evidence of end-effector independent activity changes in a number of frontal and parietal regions (Gallivan et al., 2011). Such findings are critical to future BCI study as the BCI end-effector represents a completely novel end effector to be used in spatial manipulation tasks; to be able to do so successfully, the brain will likely need to be capable of leveraging pre-existing visuomotor processing circuitry. For further information, the reader is referred to a number of excellent reviews of parietal processing and fronto-parietal networks (Buneo & Andersen, 2006a; Rizzolatti et al., 1998; Vesia & Crawford, 2012).

3.4 Electrophysiological models for BCI

There are two primary classes of physiologic signals used to drive BCIs: electrophysiological recordings such as single-unit activity (Carmena et al., 2003; Jackson, Baker, et al., 2006; Lebedev et al., 2005; Moritz et al., 2008), local field potentials (LFP) (Moran, 2010), ECoG (Blakely et al., 2009; Felton et al., 2007; Krusienski & Shih, 2010; Schalk et al., 2008), the electroencephalogram (EEG) (Moran, 2010), and the magnetoencephalogram (MEG) (Mellinger et al., 2007); hemodynamic responses such blood-oxygen level dependent (BOLD) changes as measured in fMRI (Sitaram et al., 2008) and functional near infra-red spectroscopy. There are alternatives (e.g., optogenetic and ultrasound-based imaging methods), though these are much less common. Instead, we will focus on the four common classes of electrophysiological recordings and provide a discussion of their relative merits and demerits in the context of real-time BCI for control of an end effector, leading to an explanation of why ECoG is the most appropriate signal modality for the work described below.

The methods used to acquire these signals range from very invasive (single unit activity and LFP) to minimally invasive (ECoG) to non-invasive (EEG). In general, there is a formidable sacrifice of spatial resolution, temporal resolution, or both as less invasive signal acquisition methods are used (Gilja et al., 2011; Moran, 2010).

The majority of BCI experiments as well as pre-clinical BCI applications (Birbaumer et al., 2000; Flor et al., 1995; McFarland et al., 1997; Pfurtscheller et al., 2000; Wolpaw & McFarland, 2004) use EEG as a signal source because it is non-invasive, healthy subjects are readily available, the necessary equipment is comparatively inexpensive, signal quality typically does not degrade over time, and the use of EEG in humans does not impose the same regulatory burdens as ECoG or microelectrode recordings. Within the EEG modality, BCIs have been constructed that use evoked potentials such as the visual evoked potentials in the P300 speller, volitional changes in slow cortical potentials, and volitional event-related desynchronization and/or synchronization of the sensorimotor rhythms (SMR) mu and beta (Wolpaw et al., 2002). Regrettably, there are a number of factors of EEG that currently make it poorly suited for clinical deployment and for the purposes of the study described below. First, the dura mater, cerebrospinal fluid (CSF), skull and scalp collectively act to spatially mix and low-pass filter the true cortical potentials, such that the potentials recorded at the surface of the scalp have lost tremendous spatial specificity and typically do not contain frequency content above approximately 60 Hz. Second, EEG is easily contaminated by surface electromyographic (EMG) potentials as well as ambient electromagnetic noise (e.g., 50/60 Hz line noise) resulting in a relatively poor signal to noise ratio (SNR). As a result, EEG-based BCI typically use classification-based decoder architectures (selecting from a discrete set of control parameters) as opposed to regression-based decoder architectures (mapping to a continuous space of control parameters), and require longer time intervals over which to average observations to collect a statistically robust assessment of the underlying neural activity.

At the other end of the spectrum are electrophysiological recordings made using invasive microelectrodes, such as the Utah electrode array (UEA), a microelectrode array consisting of 100, 1-1.5 mm long electrodes configured in a 10x10 grid with an inter-electrode spacing of 300-400 μm . The electrodes themselves are shielded along the shank and conductive only at the tip. It is immediately apparent that there is a tremendous difference in spatial specificity between the UEA and an EEG system. However, it is important to note that the UEA itself records from an area of less than 13 square millimeters, making correct placement over anatomical areas a critical component of experimental design and execution; microelectrode recordings from distributed cortical sites require multiple implants. The UEA and other microelectrode recording devices are designed to record extracellular action potentials and LFPs from nearby cell bodies, thus they are typically placed such that the conductive tip sits as near to the layer V pyramidal neurons as possible. Action potentials recorded from these neurons provide an extremely high-fidelity signal that has been utilized in a BCI for continuous end effector control on numerous occasions (e.g., Fetz, 1969; Simeral et al., 2011). Though these signals provide excellent fidelity and have been demonstrated to be highly capable of adaptation to task requirements (Fetz & Baker, 1973; Fetz, 1969; Ganguly & Carmena, 2009; Ganguly et al., 2011), the long-term recording of activity from single units is a difficult technical task that has yet to be solved completely. Though there are limited instances of microelectrode implants that can still record sufficient neural activity to control a BCI up to 3 years post-implant (Simeral et al., 2011), in most cases, either due to movement of the array or underlying physiological changes, intracortical implants are functionally limited to shorter operating time-frames (Kipke et al., 2008).

A signal modality that strikes a compromise between these two is ECoG. ECoG is referred to as a minimally invasive recording technique as it still requires surgery to implant the necessary electrodes, but the electrodes sit on (instead of penetrating) the pial surface. ECoG has even been demonstrated to allow for successful acquisition of spectral components up to 200 Hz when recorded epidurally (Gomez-

Rodriguez et al., 2010), potentially further lessening its invasiveness. ECoG signals are almost exclusively acquired in the context of clinical treatment for intractable epilepsy wherein patients undergo long-term monitoring (approx. 7d) of ECoG activity during ictal events for the identification and eventual resection of a seizure focus.

Within the context of motor function and cognitive processing, there are two frequency ranges of the ECoG signal that have been focused on heavily in the literature. They are a band-limited, low frequency feature (12-25 Hz), which is the ECoG correlate of the mu-beta rhythm discussed in EEG literature, and a broadband, high frequency feature (70-200 Hz), referred to as high-gamma (HG). Changes in HG activity have been postulated to reflect changes in the overall firing rate and/or firing synchrony of underlying neural populations (Miller et al., 2007; Ray et al., 2008).

Changes in HG activity are more spatially focal than changes in the mu/beta band (Miller et al., 2007), and have subsequently been concentrated on in research that uses the ECoG signal modality as a temporally and spatially local indicator of underlying cortical processing. This signal has been applied to a number of ECoG BCI paradigms, typically with changes in spectral estimates of HG power being directly mapped to end effector control parameters (Blakely et al., 2009; Schalk et al., 2008).

3.4.1 ECoG BCI

Ten years ago, ECoG cortical potentials were first leveraged as control signals for a BCI application (Leuthardt et al., 2004). Since that time, a number of subsequent studies have demonstrated not only that ECoG BCI users are capable of performing multiple simultaneous types of motor imagery to achieve multiple dimensions of control (Schalk et al., 2008; Wang et al., 2013), but also that ECoG BCI can be successfully used over multiple days without the need for classifier retraining (Blakely et al., 2009). ECoG BCIs are generally controlled through volitional modulation of HG activity at one or more electrodes; in

humans, initial execution of the task is done through the use of a motor (e.g., hand motor imagery) or cognitive (e.g., mental arithmetic) task that modulates activity in the controlling electrodes.

Because clinical-scale ECoG grids typically cover approximately 64 cm² of cortical tissue, and are only controlled by a subset of the electrodes in the grid, ECoG BCIs provide an excellent opportunity to probe the nervous system for additional physiological details relevant to BCI use. One example of this is a recent study demonstrating a spatially focal change in sleep spindle density correlated with training on BCI task, a finding suggestive of offline learning taking place in the brain after having the opportunity to perform the highly novel BCI task (Johnson et al., 2012).

3.4.2 Disadvantages of ECoG BCI

It is important to note that ECoG data are collected in an opportunistic recording model. With few exceptions, ECoG subjects are undergoing long-term monitoring for epileptic focus identification and resection. Depending on the etiology of their epilepsy, these subjects can have a host of neural complications, including but not limited to inter-ictal activity and cortical reorganization. The full extent of how epilepsy impacts the nervous system is still largely unknown and likely varies across individuals; conclusions drawn from studies based on this patient population must be accompanied by this consideration.

In addition to this, the median duration of observation for epileptics undergoing this procedure is 7d. Typically subjects are neither able nor willing to participate in research studies until the third or fourth post-operative day, which limits the amount of training time a given subject can receive on the BCI to at most four days. Extensive studies on cortical adaptation associated with BCI (e.g. responses to perturbation after the user achieves a learned state) use are extremely difficult under these circumstances.

3.4.3 Advantages of ECoG BCI

These caveats to the research use of ECoG recordings notwithstanding, ECoG-based BCIs possess several advantages relative to SUA and EEG-based BCIs that make them the appropriate tool for the inquiry discussed in this document. First, because the electrodes are not typically moved from day-to-day, and the cortical potentials are fairly stationary, ECoG BCIs allow for robust control over the course of multiple days without retraining of the classifier (Blakely et al., 2009). ECoG recordings easily remain consistent over multiple days, and have been shown to perform well for up to 30 days (Wang et al., 2013) with daily classifier updates.

Relative to scalp surface potentials, ECoG presents an excellent opportunity for BCI research because, as of yet, there are no robust computational algorithms for extracting spatially focal HG activity from EEG on a single trial basis. Work is being performed to construct an inverse model that will permit us to back project likely cortical sources of EEG data (Darvas et al., 2004, 2010), thus increasing SNR and potentially providing access to information in the HG range at the scalp, however, these techniques are still under development.

As was mentioned above, ECoG provides an excellent compromise (relative to EEG and microelectrodes) of spatial distribution of coverage and signal fidelity, making it an appropriate tool for the investigation of activity patterns in diverse cortical areas during BCI use.

3.5 Distributed cortical processing during BCI use

As was outlined above, BCIs show great promise for changing how we interact with the world. The field of brain-computer interfacing has demonstrated the tremendous adaptive potential of brain and machine by showing BCIs can be based on activity from one to millions of neurons, with response latencies from tens to thousands of milliseconds. However, the performance of these devices is not currently sufficient to warrant the risk and expense to their target patient population. We posit that the

relative single-mindedness of current architectures is one significant obstacle to the performance improvements necessary to make these devices clinically viable. We propose the development of BCI architectures that recognize and leverage the spatially and temporally heterogeneous patterns of activity observable across the brain to overcome this obstacle. Though we have, as a field, repeatedly demonstrated the capability of the brain to develop control over these novel interfaces, we have done so with little attention to the adaptive processes taking place in neural populations that are not directly linked to control of the BCI. An important first step is to develop an understanding of which cortical structures are involved in BCI skill acquisition and task execution and to characterize the relationships between these regions. From there, specific relationships between neural activity and task demands can be extracted and provided to computational systems as additional channels of highly task-relevant information.

In the studies described in the remainder of this document, we demonstrate two important points that advance our understanding of the way brain and machine interact during BCI use. First, we demonstrate that execution of neuroprosthetic control in a 1-D ECoG BCI is accompanied by changes in neural activity in a variety of functionally heterogeneous cortical structures, and that these distributed regions interact with M1 in meaningful ways that are indicative of underlying patterns of structural and functional connectivity. Second, we provide evidence for extra-M1 coding of the higher-level goal both immediately prior to and during BCI task execution, and validate a potential framework for leveraging these intentional signals.

4 General materials and methods

Whether intended as a therapeutic device or as an experimental tool for fundamental research, many of the constituent components of a BCI are the same. Additionally, many of the neuroscientific methods associated with processing of the data that drive a BCI are common across different analyses.

Accordingly, this section outlines the standard BCI architecture and a number of the experimental techniques and analytical tools that are used throughout the remainder of this document.

4.1 The standard BCI model

Though the implementation of the individual subunits vary widely from research group to research group, and are generally focused on engineering efforts to improve BCI performance, the typical architecture for all BCIs is the same. Systems include the following: a means of recording neural signals, a computational algorithm to extract features of interest from those signals, a decoder to transform those features into one or more control signals, and a device, virtual or realized, that carries out the actions dictated by the control signals and provides feedback to the user via one or more sensory modalities.

4.1.1 Components of a BCI

Data acquisition hardware. Research in the area of fully implantable data acquisition hardware is moving quickly and will be critical in the development of viable and dependable BCIs. For the time being, however, the majority of human BCI studies use general purpose biosignal amplification systems. Most common electrophysiological data acquisition systems record between 16 and 256 channels simultaneously, at sampling rates specific to the signal being recorded. ECoG signals are typically sampled at > 1000 Hz using either AC or DC coupled instrumentation amplifiers, depending on the intended use of the signals being recorded.

Feature selection algorithm. There exist a wide variety of methods to extract features of interest from neural signals. Many of these are highly specific to the type of signal being recorded, such as the isolation of action potentials from individual neurons using a window discriminator-based spike sorter, or the spatial unmixing of EEG signals using independent components analysis. Often, feature selection algorithms are chosen to effectively leverage the strengths of the decoder that will be fed the output of the feature selection algorithm, with specific attention to reducing or increasing the dimensionality of the feature set as appropriate. Typically feature selection also attempts transform the neural data in such a way that the neural signals being discriminated become linearly separable. Typically, real-time feature extraction in an ECoG BCI involves re-referencing one or more previously selected channels to reduce common mode signal, and using a spectral estimator to determine a time variant estimate of power in the HG range recorded from the electrodes of interest.

Decoding. Decoding architecture options are as numerous as feature selection algorithms if not more so. Decoders are also typically highly application specific, based not only on the dimensionality and nature of the decoder outputs, but also on engineering tolerances specific to the application for which the BCI will be used. One trend in the field of BCI research has been toward building more intelligent decoders that are capable of robustly mapping motor-based neural features to BCI control, however the work discussed in this document provides an alternative view of decoder enhancement. We posit that decoders will benefit from additional channels of information that are more cognitive in nature. Such signals could provide decoders with an understanding of the user's current state, their intended goal in a multi-step movement, or knowledge of when the decoder appears to be configured incorrectly and is in need of updating.

Applications. BCI applications are highly specific to their intended use. They range from EEG-based consumer-grade fashion devices to deep brain stimulators and neurally-controlled robotic limbs. The

content of this document focuses specifically on the last of these due to their clinical relevance to the target population of patients with motor disorders.

4.2 Data collection and signal processing

ECoG Grids. In all studies described below, subjects were implanted with platinum sub-dural ECoG electrodes for the purpose of seizure focus localization at Harborview Medical Center and Children’s Hospital in Seattle, Washington. The physical makeup (number and arrangement of electrodes) and implant location of all grids were based on clinical indication. Arrays were either 8x8, 6x8, 4x8, or 2x8 grids or 1x8, 1x6, or 1x4 strips with 2.4mm diameter exposed recording surface and a 1cm inter-electrode distance.

Institutional approval. Subjects provided informed consent in accordance with the Institutional Review Board’s direction and patient data were anonymized in accordance with HIPAA mandate.

All procedures were carried out within the University of Washington Regional Epilepsy Center, either at Harborview Medical Center or Seattle Children’s Hospital after informed consent was obtained. For children under age 18 parental consent was obtained along with consent from the child (age 14 or above) or assent of the child (age 7-13). The protocol was approved by the Institutional Review Board at both institutes.

Recordings. Experimental recordings were performed at the patient’s bedside without interrupting the clinical recording systems. Either Synamps2 (Neuroscan, El Paso, TX, USA) or g.USBamps (GugerTec, Graz, Austria) sampled at 1000 Hz and 1200 Hz respectively were used for recording. Cortical potentials referenced against a scalp electrode, and were digitized and processed using the BCI2000 software suite (Schalk et al., 2004) which provided real-time feedback to the user.

Motor screening. Prior to online control, subjects performed overt motor screening to determine candidate electrodes for BCI use. Depending on each subject's coverage, they performed gross hand motor movements (of the hand contralateral to the implant site), mouth motor movements, or both. Visual cues were presented for 3 sec followed by a 3 sec inter-trial interval. This process was repeated 10-30 times for each of the two motions. They were then asked to repeat this screening process but with imagined movement. Electrodes that demonstrated statistically significant change in HG power during activity as compared to rest in either or both of these tasks were chosen as candidate controlling electrodes. In cases where there was more than one candidate controlling electrode, the electrode used for online control was chosen based on the magnitude of change between activity and rest and/or neuroanatomical relevance.

BCI Tasks. Both BCI paradigms were driven by spectral power changes in a portion of the HG frequency band of a single electrode determined to be modulated by motor imagery. Only a subset of the HG range was used during online control (approx. 75-100 Hz) for computational tractability and to eliminate the need for real-time notch filtering to reduce line noise harmonics. HG activity was chosen as the control feature as it has been previously postulated that HG activity is a correlate of underlying population level firing rates or coherence in firing (Ray et al., 2008).

Right-justified box task. The standard right justified box task (RJB) (Wolpaw et al., 2003) is depicted in Figure 2. During execution of the BCI task, the subject is presented with one of two targets, occupying either the top half or the bottom half of the right-most edge of the screen. After a cue interval of two sec, the cursor appears on the left edge of the monitor and travels to the right at a constant horizontal velocity, such that the duration of the feedback period is fixed (typically 3 sec). The subject controls the vertical velocity of the cursor by modulating HG activity at the previously selected controlling electrode (CTL); performance of motor imagery causes the cursor to travel up and remaining at rest causes the cursor to travel down. Their objective is to complete each trial with the cursor in the specified target

area for that trial. HG activity recorded at CTL is mapped to vertical cursor velocity using a simple linear decoder that was trained in the first set of trials. The task consists of four phases: rest, cue, feedback, and reward. Throughout the remainder of this document – when discussing the RJB task – targets occupying the top half and bottom half of the screen are referred to as “up-targets” and “down-targets,” respectively.

The cursor’s vertical velocity was updated every 40ms and controlled by changes in HG activity at the controlling electrode as calculated by an auto-regressive filter using the previous 500ms of data. This time-variant estimate of HG activity was z-normalized against 12 seconds of stored data (6 for each target type) and then mapped to cursor velocity. The normalizer was typically adapting (collecting reference data and updating normalization parameters) only during the first run (18 trials), however in cases where non-stationarity of the signals showed obvious bias, the normalizer was allowed to recalibrate. Subjects participated in the experiment over the course of multiple days; duration of the recording sessions was dictated by the subjects’ willingness and capability to participate.




Figure 2 - Overview of the RJB task depicting the spatial scale of the ECoG grids, as well as the phases and timing of the task. Subjects were presented with a target occupying either the upper (up target; shown) or lower half (down target; not shown) of the right-most edge of the screen and had 3 sec to control the vertical position of the feedback cursor such that it ended the trial in the target area.

Cortical reconstructions and anatomical labeling. Pre-operative magnetic resonance images (MRI) were co-registered with post-operative Computed Tomography (CT) scans using the Statistical Parametric Mapping software package (Penny et al., 2006). Three-dimensional reconstructions of the pial surface were generated using Freesurfer (freely available for download at <http://surfer.nmr.mgh.harvard.edu/>) and custom code implemented in Matlab (The MathWorks, Natick, MA). Electrode positions estimated

from post-operative CT were projected to the reconstructed pial surface using the method outlined by Hermes and colleagues (Hermes et al., 2010). MRI images and projected electrode locations were normalized to Talairach coordinates using Freesurfer.

Anatomical labels were estimated using the Human Motor Area Template (HMAT) (Mayka et al., 2006), and the Talairach Daemon (Lancaster & Rainey, 1997; Lancaster et al., 2000). The HMAT atlas is based on the meta-analysis of 126 motor-based fMRI studies, thus it does not include posterior parietal cortex or prefrontal cortex, but does subdivide BA6 into PMv and PMd.

5 Distributed cortical activity patterns during BCI use

The majority of subjects who attempt to learn control of a brain–computer interface (BCI) can do so with adequate training. Much like when one learns to type or ride a bicycle, BCI users report transitioning from a deliberate, cognitively-focused mindset to near automatic control as training progresses. What are the neural correlates of this process of BCI skill acquisition? Seven subjects were implanted with electrocorticography (ECoG) electrodes and had multiple opportunities to practice a 1-D BCI task. **As subjects became proficient, strong initial task-related activation was followed by lessening of activation in pre-frontal cortex, dorsal pre-motor cortex, and posterior parietal cortex, areas that have previously been implicated in the cognitive phase of motor sequence learning and abstract task learning. These results demonstrate that though the use of a BCI only requires modulation of a local population of neurons, a distributed network of cortical areas is involved in the acquisition of BCI proficiency.**

5.1 Introduction and background

Over the past 50 years, it has been demonstrated that, when given feedback, the brain can learn to volitionally modulate the activity of single neurons (Chapin et al., 1999; Fetz & Baker, 1973; Fetz, 1969; Kennedy & Bakay, 1998; Koralek et al., 2012) and populations of neurons (Fabiani et al., 2004; Leuthardt et al., 2004; Miller, Schalk, et al., 2010; Wolpaw et al., 1991). This modulation can occur in the absence of overt movement (Leuthardt et al., 2004; Miller, Schalk, et al., 2010), or even when overt movement is not possible (Hochberg et al., 2006; Kennedy & Bakay, 1998). The use of these signals to control external devices for restoration of lost function or as potential feedback signals for rehabilitation holds great promise for the future; however, our understanding of the underlying processes that drive this activity modulation is incomplete.

The ability to voluntarily modulate neural activity to control a BCI appears to be a learned skill, similar to learning to ride a bike or swing a golf club. Investigators have repeatedly documented that task performance typically increases over the course of practice (Carmena et al., 2003; Ganguly & Carmena, 2009; Moritz et al., 2008; Schalk et al., 2008; Wolpaw & McFarland, 2004), which is indicative of a learning process taking place in the brain. In many cases, human BCI users have anecdotally reported transitioning from a very deliberate, cognitive approach to nearly automatic execution. Whereas they were initially focused on the cognitive task (e.g. motor imagery) that drives neural activity and in turn controls the end effector, after training, they report a goal-directed approach focused directly on the end effector itself.

BCI use is interesting in that it shares characteristics with both concrete motor tasks as well as abstract cognitive tasks (Green & Kalaska, 2011). Users of BCIs are manipulating their physical environment, but doing so without necessarily moving their bodies, thus depriving their brains of the multiple modalities of sensory feedback that have been demonstrated to be vital to learning a motor skill (Wolpaw, 2007).

The transition from a cognitive to an automatic phase, combined with the fact that the majority of single-unit and field-potential BCIs are driven by signals recorded from motor areas, suggests the hypothesis that BCI skill learning shares many commonalities with motor sequence learning. Indeed, in most cases, large initial gains in performance are followed by gradually diminishing improvements, congruent with Fitts' well-known law of motor performance during skill learning (Fitts, 1954). Following results on neural activity patterns during motor skill learning (Doyon & Ungerleider, 2002; Mier et al., 1998), one would therefore expect significant involvement of motor, prefrontal, and parietal cortical areas, as well as the striatum and cerebellum during the cognitive phase of BCI skill acquisition.

To investigate the role of distributed cortical networks in BCI skill acquisition, we recorded population-scale neural activity simultaneously from various locations across the cortex using ECoG while subjects

performed a BCI task. The task required subjects to modulate a sub range (approx. 70–100 Hz) of the high-gamma band (HG, 70–200 Hz) in ECoG surface potentials to control the position of a cursor on a computer monitor. HG activity was used because it has been previously shown to correspond to an increase in the firing rate and/or coherence of underlying local neural populations (Ray et al., 2008). **In this chapter, we investigated two hypotheses: (a) that motor learning networks across the cortex participate in BCI learning even though the task only requires modulation of a localized neural population, and (b) that activity in these distributed cortical networks changes with increasing BCI task proficiency.**

5.2 Materials and methods

Subjects. This study was completed as a retroactive analysis of previously recorded ECoG data. Inclusion criteria for subjects were as follows: (1) subjects needed to have participated in the 1-D, right justified box task; (2) activity changes in the controlling electrode needed to be driven by overt motor movement or motor imagery; (3) subjects needed to participate in 50 or more trials. Note that subjects were not chosen for specific electrode coverage (other than the motor coverage necessary to participate in the task), thus not all subjects had coverage in all areas discussed. Seven subjects (2 female, mean age 25.6 [range 18–32]) met these inclusion criteria. Figure 3 summarizes the spectrum of coverage for the seven subjects.

Table 1 - Demographic information for subjects included in the analysis of distributed cortical activity patterns during BCI use. Abbreviations: R right, L left, F frontal, P parietal, T temporal, O occipital, M Act refers to the motor task associated with the controlling electrode, M/MI motor/motor imagery.

ID	SID	Gender	Age	M Act	M/MI Type	Coverage
S1	fc9643	F	26	Tongue	Overt	R-F/P/T
S2	26cb98	M	22	Tongue	Imagined	R-F/P/T
S3	38e116	M	18	Hand	Overt	R-F/P/T
S4	4568f4	M	27	Tongue	Overt	R-F/P/T
S5	30052b	M	29	Tongue	Imagined	R-F/T
S6	mg	M	32	Tongue	Imagined	L-F
S7	04b3d5	F	25	Tongue	Imagined	L-F/P/T

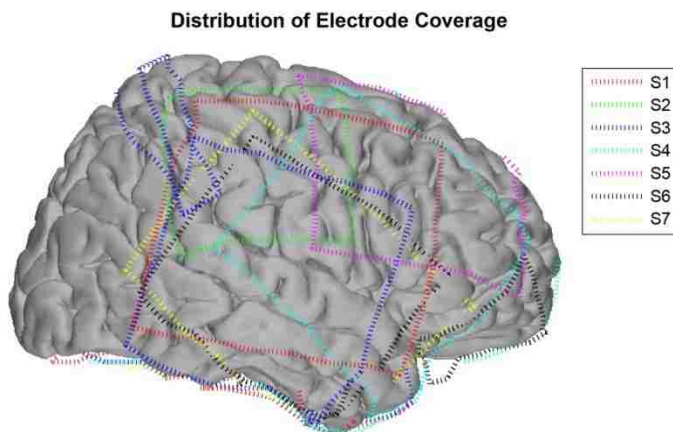


Figure 3 - Spatial distribution of electrode coverage – Shows the differences in electrode coverage for the 7 study subjects. Dashed lines forming a rectangle or square represent the approximate outlines of grids (8x8, 6x8, 4x8, or 2x8). Single dashed lines represent the approximate centerline of strips (1x8, 1x6, 1x4). Left hemispheric coverage is translated to the right hemisphere to allow for direct comparison.

Recordings. See section 4.2.

BCI Task. All subjects in this study performed the RJB task. See section 4.2 for additional detail.

Anatomical labeling. See section 4.2.

Preprocessing. Signals were common average re-referenced by subtraction of the average signal recorded at all electrodes within an implanted grid or strip. This was done to eliminate any common noise introduced by activity recorded at the reference electrode. Re-referenced signals were band pass filtered for the HG range (70-200 Hz) using a fourth-order Butterworth filter and a time variant estimate of band-power was calculated using the square of the magnitude of the Hilbert transform. Data were then log transformed to become approximately normally distributed such that the assumptions made by statistical tests employed in further analyses would be met. To account for changes in recording characteristics from session to session the resulting log-power estimate was z-normalized with respect to rest samples only per channel and session. Throughout the remainder of this section, the signal resulting from the pre-processing steps will be referred to as HG activation.

Quantification of average responses. Trials were divided into four periods: rest, cue, feedback, and reward. These periods were 1, 2, 3, and 1 second, respectively, resulting in a trial length of 7 seconds. The HG activation for each of these periods was calculated by averaging all HG activation samples within a given period. Average activation for a given electrode across all trials of a given type was calculated as the mean HG activation during all feedback periods for trials of a given type minus the mean HG activation during all rest periods.

Estimation of HG activation separability and learning states. In order to compare “unlearned” and “learned” states, it is necessary to define a trial or trials that delineate these two states. Alternatives exist, such as assuming a continuous learning process or selecting a portion of early trials and late trials to be representative of the learned and unlearned states. However, as performance increases in task learning typically follow asymptotic trends, and there is no guarantee that users of a task will learn at the same rate, we thought it necessary to use a data-driven approach to determine the transition trial. As behavioral performance on the RJB task saturates quickly, assessment of trends in activation patterns at the controlling electrode itself is a reasonable method for determining this transition trial.

Our algorithm for determining the transition trial makes the following assumptions: (a) that there are in fact two states (this does not preclude the existence of more than two states, but our algorithm will only detect a single transition and will effectively group sub-states), and that (b) the two states can be differentiated by observing changes in a subject’s ability to differentially modulate activity at the controlling electrode for up targets as compared to down targets.

The transition trial was determined as follows:

- (1) A running estimate of each subject’s ability to separate HG activation in up targets as compared to down targets was calculated.

- a. The mean HG activations for up and down targets were separately smoothed using a Gaussian kernel (5 trials full width at half max [FWHM]).
- b. Because up and down targets were presented in random sequence, a running difference between smoothed up target and down target activations was calculated by linearly interpolating between observation points.
- c. From these operations, an estimate for the difference between activation in up targets and down targets as a function of trial number was established.

(2) A model of two distinct Gaussian distributions was fit to the difference estimate such that the difference between these two distributions (measured using the statistic explained below) was maximized. The single free parameter in this model was the estimated transition trial.

Distance between the distributions was calculated using the following equation to account not only for differences in the means of the two distributions but also for the variances of the two distributions.

$$D_{el} = \frac{(\bar{e} - \bar{l})^3}{|\bar{e} - \bar{l}| \sigma_{e:l}^2} \frac{N_e - N_l}{N_{e:l}^2} \quad (1)$$

where D_{el} represents the separability of the early and late trials, e and l represent the early and late trials themselves, N_e and N_l represent the number of samples in each of these sets and $e:l$ represents the joint set of all trials.

When comparing two distributions, this measure represents the proportion of the variance of the joint distribution that can be explained by the difference in the means of the two sub-distributions. It has been used previously in the assessment of ECoG signals (Blakely et al., 2009; Miller et al., 2007; Miller, Schalk, et al., 2010).

Generation of time course activations. The average time course of HG activation for all trials of a given type were constructed by averaging the time course of HG activation for a given target and smoothing

with a Gaussian kernel (250 ms FWHM). Standard errors were calculated before smoothing and were also smoothed using a Gaussian kernel (250 ms FWHM). Trial-by-trial progressions of the time course of HG activation were generated by aligning all trials into an $M \times T$ matrix where M was the number of trials of a given type and T was the length of a trial in samples. This matrix was then smoothed simultaneously in both dimensions using a 2-D Gaussian kernel (300 ms and 7 trials FWHM).

Comparison of early and late activations. Shifts in activation from early to late learning states were calculated in much the same way that average responses were calculated except that instead of comparing activation during the feedback period to the activation during rest, activation during the feedback period of early trials was compared to activation during the feedback period of late trials.

5.3 Results

Behavioral performance. Overall performance, depicted in Figure 4, increased significantly as subjects gained experience with the task ($p = 0.035$; right-sided t -test, $N = 7$). Behavioral performance typically began above chance levels and saturated quickly. Cursor trajectories could potentially be used to infer non-saturating performance metrics, but would require an assumption be made regarding optimal trajectory (see below for further discussion). Thus, we used activity patterns at the controlling electrode to assess learning-related changes and define a transition from a learned to an unlearned state.




Figure 4 - Behavioral performance – (a) Mean task performance for all users. SEM shown in gray around trend line. Chance performance and the 95th percentile of chance performance are shown as horizontal dashed and dotted lines, respectively. Note that mean performance was above chance beginning with the first run. The number of runs performed by each user varied, thus all seven users contributed to data for the first three runs, six users contributed through the first five runs, and four contributed to the sixth run. (b) Comparison of performance for the first and least runs performed by each subject. Significant differences ($p < 0.05$) denoted with an asterisk ().**

Behavioral results assuming an optimal cursor trajectory. The solution space for how a given subject could achieve success in the RJB is extremely large. Subjects were not asked to report their intended cursor trajectory. If one makes the assumption that the intended cursor trajectory is constant throughout all trials, and that the subject asymptotically approaches that trajectory, the intended trajectory can be approximately inferred by looking at the last few trials conducted by each subject. Each individual trial can then be compared to this inferred trajectory, giving an estimate of how well a given subject’s performance approached this intended trajectory as they gained experience with the task. In an effort to obtain an understanding of some of the behavioral parameters of the RJB, for the five of seven subjects for whom target trajectories were recorded, we performed this analysis. In three of these five subjects we saw a statistically significant decrease in mean-squared error (MSE) over the course of all trials performed, compared against an intended trajectory estimated from the last 20% of trials of a given target type. Individual results are shown in Figure 5.

It is important to note that this method assumes that performance gains are made only by the subject improving their ability to precisely execute an indented trajectory, not by fine-tuning of the intended

trajectory itself. We find the assumptions made in this analysis less robust than those used in the primary analysis that uses power changes in the controlling electrode as an assessment of task performance, but they do provide insight in to the fact that as subjects gain task experience, and exhibit changes in dynamics of neural control, there are concomitant changes in some behavioral aspects of the task.

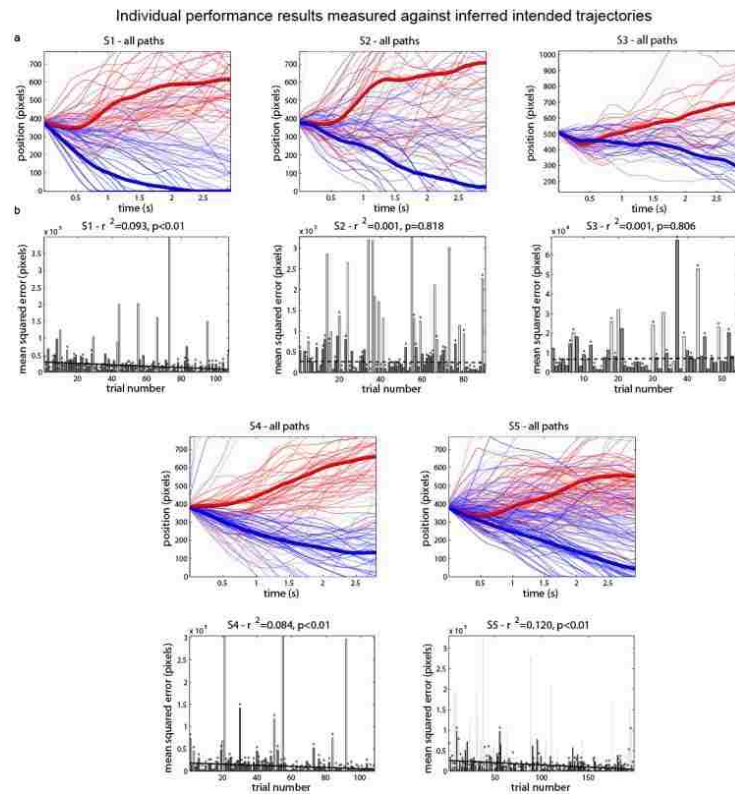


Figure 5 - Behavioral analysis of cursor trajectories measured against inferred intended trajectories – (a) Recorded trajectories for individual subjects shown as cursor position as a function of time during the feedback phase of the task. Trajectories for up targets are shown in red and for down targets are shown in blue. Earlier and later trials are depicted in lighter and darker shades, respectively. The inferred intended trajectory is shown as a thick red or blue line. **(b)** MSE for each individual trial relative to the inferred intended up and down trajectories. Successful trials shown in dark gray and failed trials are shown in light gray. Up targets are notated with a diamond (\diamond) above each bar. To assess performance trends we performed logistic regression of MSE values as a function of trial number. Significant trends are shown as solid regression lines, non-significant trends are shown as dashed lines.

Run-by-Run performance. During the feedback period, the cursor travelled from left to right across the screen over the course of 3 sec. In this time period, the subject was tasked with causing the cursor to ‘hit’ the indicated target on the right side of the screen, meaning, that at the end of the 3 sec feedback

period, the cursor needed to be within the vertical area defined by that target. In terms of whether or not subjects were able to achieve this requirement, task performance quickly saturated, making assessment of improvement as a function of time impossible. Based on simulated, randomized replay of the task using previously recorded ECoG data that were phase scrambled and random target sequences (144 runs of 17 trials each, totaling 2448 trials), chance task performance for a given run was 48.8% with a 95th percentile of 64.7%. Accordingly, if actual performance for a given run exceeded the 95th percentile, then performance on that run can be considered above chance with statistical significance. Figure 6 shows individual performance trends for each subject on a run-by-run basis. Note that 5 of the 7 subjects were performing above chance by the end of their first run and performance often saturated quickly, necessitating the use of alternative methods to assess changes in aptitude.

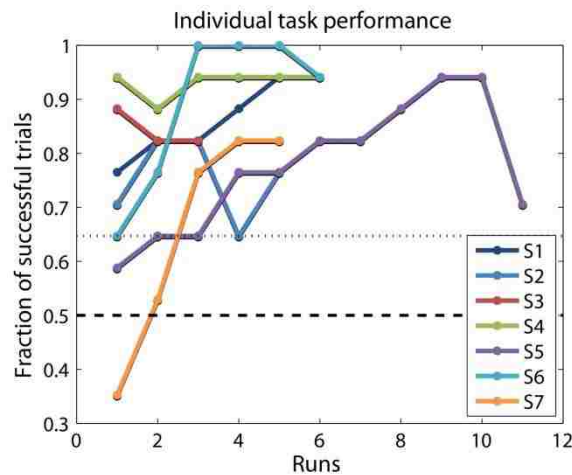


Figure 6 - Run-by-Run task performance – Task performance for each user shown on a run-by-run basis. The 95th percentile of chance performance is represented in gray. Note that for five of the seven subjects task performance was above chance by the end of the first run, and for many subjects performance quickly saturated.

Volitional modulation of activity at the controlling electrode. In six of seven subjects (all subjects except for 04b3d5), we found a statistically significant increase in HG activity during the feedback period of up-targets as compared to all rest periods (right-sided two-sample *t*-test, Bonferonni corrected, $27 < N_1 < 98$, $51 < N_2 < 187$; $p < 0.0001$). The seventh subject ($N_1 = 45$, $N_2 = 85$, $p = 0.271$) did not demonstrate

a significant difference between up-targets and rest; however, in this subject we found a statistically significant decrease in HG activity between down-targets and rest (left-sided two-sample t -test, Bonferonni corrected, $N_1 = 45$, $N_2 = 85$, $p = 0.0028$), suggesting that though activity suppression below baseline is generally not employed as a control strategy, it may have been the strategy used by this subject. These same activity decreases during down targets were not observed in the other six subjects (left-sided two-sample t -test, Bonferonni corrected, $27 < N_1 < 100$, $51 < N_2 < 187$; $0.075 < p < 1$). See Table 2 for details.

Table 2 - Summary of subject population, data collected, and basic statistical analysis. Abbreviations: SID, subject identifier; N number of trials performed; M/MI overt motor or motor imagery; CTL, controlling electrode; U, up; D, down; R, rest. D_{el} is the separability metric, and p_{el} is the p-value associated with separability of the two learning states.

ID	SID	Trials			HG at CTL.		Trans. Trial	D_{el}	p_{el}
		N	Up	Down	p U vs R	p D vs R			
S1	fc9643	108	54	54	< 0.0001	0.075	48	0.6422	< 0.0001
S2	26cb98	90	45	45	< 0.0001	0.994	75	0.3403	< 0.0001
S3	38e116	54	27	27	< 0.0001	1	22	0.8689	< 0.0001
S4	4568f4	108	54	54	< 0.0001	0.983	53	0.6097	< 0.0001
S5	30052b	198	98	100	< 0.0001	0.986	32	0.4076	< 0.0001
S6	Mg	108	54	54	< 0.0001	1	20	0.2039	< 0.0001
S7	04b3d5	90	45	45	0.271	0.0028	40	0.7695	< 0.0001

Task-modulated activity throughout cortex. We recorded ECoG data from 652 electrodes across the seven subjects. Of these electrodes, 83 were excluded from analyses because they contained non-physiologic artifacts (resulting from poor contact, placement over scar tissue, etc). In the remaining 569, we found 152 electrodes showing statistically significant increases in HG activity between feedback during up targets as compared to rest (right-sided two-sample t -test, Bonferonni corrected, $27 < N_1 < 98$, $51 < N_2 < 187$; $p < 8.787 \times 10^{-5}$). Further, of that same 569 electrodes, we found 125 electrodes showing

statistically significant increases in HG activity between all targets and rest (right-sided two-sample t -test, Bonferroni corrected, $54 < N_1 < 198$, $51 < N_2 < 187$; $p < 8.787 \times 10^{-5}$). Electrodes showing significant activity increase were concentrated in cortical areas previously known to be associated with motor learning: M1, primary somatosensory cortex, dlPFC, PMd, and PMv (Jenkins et al., 1994; Schlaug et al., 1994). Additionally, electrodes showing increases were found in posterior parietal cortex (PPC), an area associated with sensorimotor tasks involving visual feedback (Buneo & Andersen, 2006a). Additional task-modulated electrodes were found in other cortical areas, such as the temporal parietal junction and the inferior temporal gyrus, though to a lesser extent. Task-modulated activity in SMA was noted in the only subject with SMA coverage, but this finding was not included in our analyses as it could not be verified in multiple subjects. Figure 7 illustrates the spatial distribution of electrodes showing significant activation for up targets relative to rest.

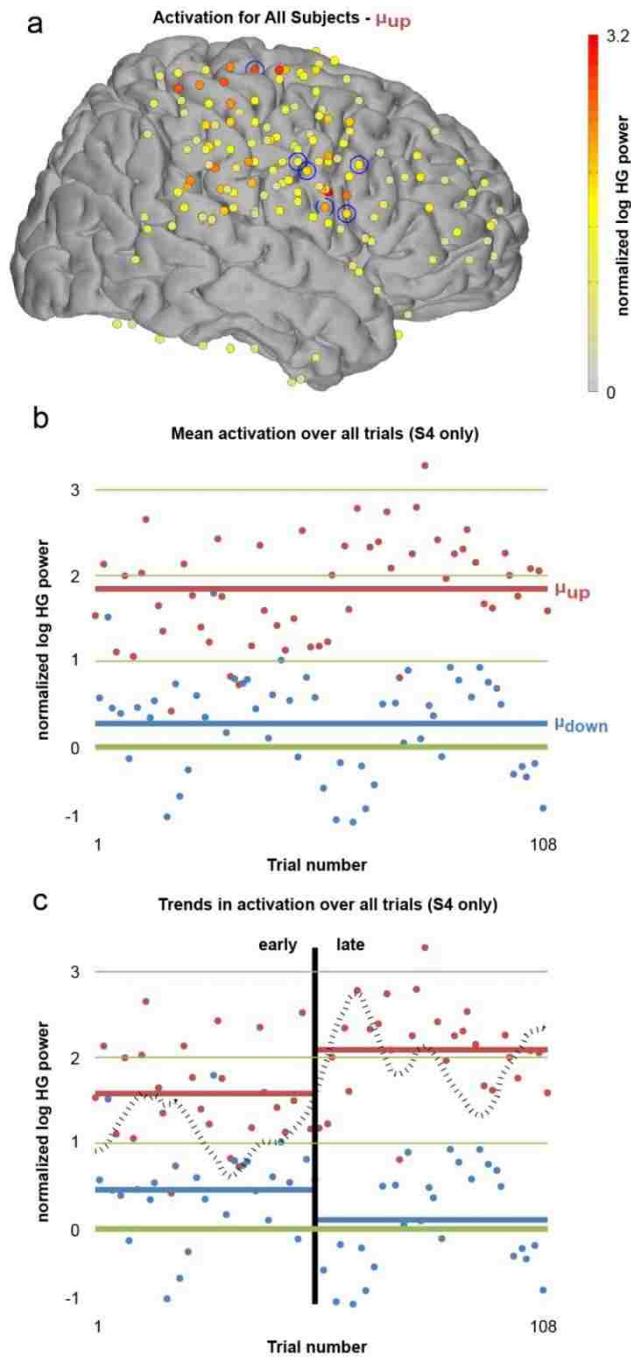


Figure 7 - Cortex-wide activity during BCI use and modeling of early vs. late activity patterns – (a) Activation during up targets for all lateral electrodes for all subjects (left coverage projected to right hemisphere) shown on the Talairach brain. Note widespread cortical activation including frontal, middle-parietal, and posterior-parietal areas. Controlling electrodes are circled in blue. (b) Activation for an example subject (S4 / 4568f4) during each feedback period normalized against log HG power during rest. Each dot represents one trial, up and down targets are shown in red and blue, respectively. Thick red and blue lines represent mean activation for all trials, respectively. (c) Early-late trial division shown for subject S4. Separability of high-gamma activity during up and down trials was used as a measurement of task proficiency. A separability measure (black dotted line) was modeled as two Gaussian distributions and fit to the data such that the distance between the distributions was maximized.

Determination of unlearned vs. learned states. Motor skill learning has been characterized as having multiple distinct learning phases: cognitive, associative, and autonomous (Fitts, 1954). These loosely correspond to understanding what actions a certain skill requires, optimization of performance at a skill through repeated action, and development of an automatic capacity to carry out the skill, respectively. Recent work has developed a neural model describing the involvement of various cortical and sub-cortical structures during motor learning (Doyon & Benali, 2005; Doyon & Ungerleider, 2002). This model postulates involvement of pre-frontal and pre-motor as well as parietal areas during the initial, cognitive phase of motor sequence learning.

With this motivation in mind we sought to capture differences in HG activity at the controlling electrode that changed with this transition from an “unlearned” to a “learned” state. We employed the data-driven approach described above in the methods. In comparing unlearned and learned distributions, distance value magnitudes close to one implied that the majority of the variance in the joint distribution can be explained by the difference in the means of the two sub-distributions. The sign of the resulting distance measure demonstrates the direction of shift in the means of the two sub-distributions, where a positive value implies that the mean of the learned distribution is greater than the mean of the unlearned distribution. Within the context of this task, a large, positive distance value for a given transition trial implies that the subject’s ability to differentially modulate activity during up targets as compared to down targets was much greater after the transition trial as compared to before. Furthermore, a relatively large distance value implies that one of the solution strategies employed was to increase the difference in activity at the controlling electrode in up targets relative to down targets. This is compared to the alternative possibility that the user only reduces the variability of their control signal as they gain experience. Were this the case we would have observed relatively low distances between the unlearned and learned distributions.

In all seven subjects, relatively large distances between the distributions of power separation for learned and unlearned states demonstrated that the applied model effectively separated those two states ($0.2039 < D_{el} < 0.8689$). Accordingly, the differences in distributions of power separation for unlearned and learned states were highly significant (both-sided two-sample *t*-test, Bonferonni corrected, $p < 0.0001$, see Table 2 for detail).

Cortex-wide changes in activation from unlearned to learned states. HG activation of all non-controlling electrodes was visualized on a trial-by-trial basis to observe activity dynamics at these electrodes. Figure 8 shows example trial-by-trial plots for electrodes located throughout motor-learning-associated cortical areas (see Figure 9 for similar plots separated by subject). M1/SC and PMd/PMv electrodes demonstrated activation primarily during up targets, and PFC and PPC demonstrated activation during both up and down targets. As can be seen these figures, changes in activation in these electrodes over the course of many trials were often, though not always, well aligned with the transition from unlearned to learned state as defined solely by HG power at the controlling electrode. It is notable that in some subjects, transitions in HG activity patterns were approximately temporally aligned with breaks in experimental sessions, suggestive that offline learning was taking place during these periods. This is complimentary to previous findings of increased sleep spindle density local to the controlling electrode correlated with training on an ECoG BCI (Johnson et al., 2012).




Figure 8 - Time-by-trial HG activation – (a) Distribution of electrodes shown in (b-e) on the Talairach brain. The subject from which each electrode was taken is noted; not all subjects had coverage in all areas (see Figure 3). (b-e, subplot 1) average HG activation in a given electrode for all trials separated by up and down targets. Subject is specified in the subplot title. Phases of the task (ordered from L to R: rest, cue, feedback, reward) are separated by vertical bars. Dotted line represents SEM. (b-e, subplots 2 & 3) trial-by-trial HG activation for all trials, separated by up (subplot 2) and down (subplot 3) targets. Trial count is shown on the vertical axis. Breaks in the experimental session of more than 8 hours are denoted with an asterisk (*). Time, as described for subplot 1, is shown on the horizontal axis. The black horizontal bar represents the transition trial, and the early and late trials relative to this point are denoted with black arrows. (b) M1/SC electrode shows continued increase in activation for up targets, congruent with changes at the controlling electrode. (c) PFC electrode shows activation for both up and down targets, decreasing near the point of model separation. (d) PMd electrode shows activation for up targets, decreasing near the point of model separation. (e) PPC electrode shows subtle activation during up and down targets, decreasing near the point of model separation.

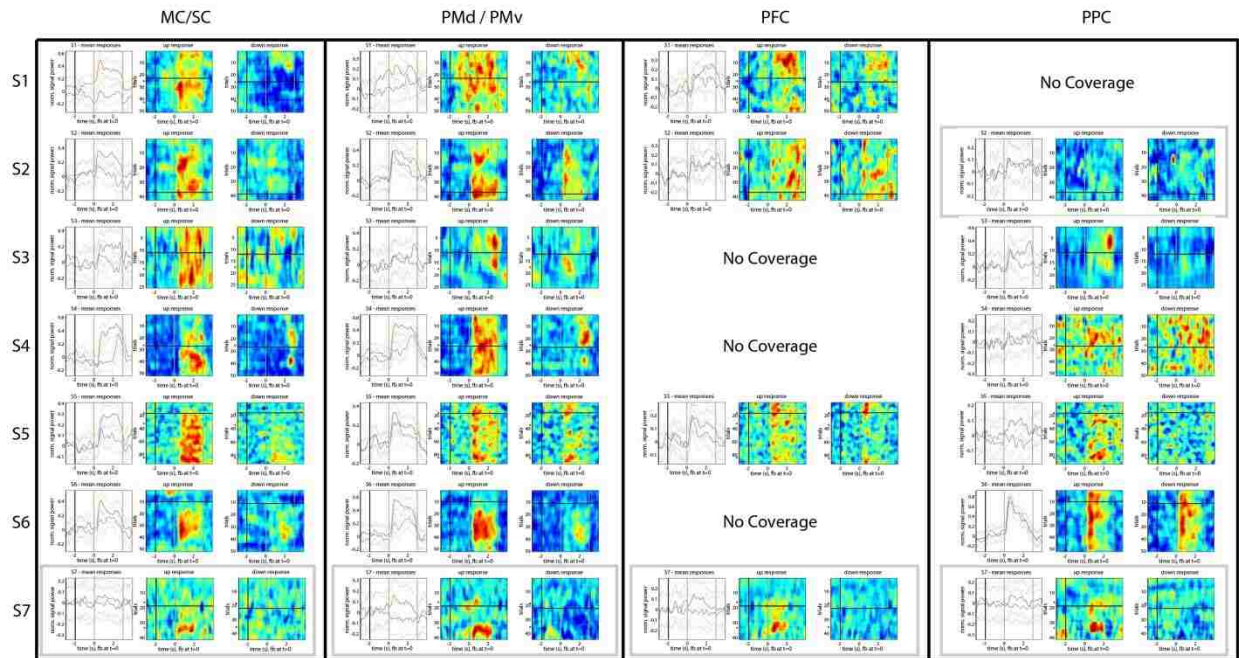


Figure 9 – Time-by-trial HG activation for individual subjects – (a) average HG activation for all trials separated by up and down targets. Phases of the task (ordered from L to R: rest, cue, feedback, reward) are separated by vertical bars. Dotted line represents SEM. Up and down target activations shown in red and blue respectively. (b & c) trial-by-trial high-gamma activation for all trials, separated by up (b) and down (c) targets. Trial count is shown on the vertical axis, time, as described for subplot 1 is shown on the horizontal axis. The black horizontal bar represents the model separation point derived solely from the controlling electrode. The pair of values in each subplot title denotes the minimal and maximal values represented by the heat map, where blue corresponds to the minimum and red the maximum. Plots surrounded in gray boxes are for electrodes that did not exhibit significant task-modulation.

Dynamics of activity in remote electrodes were quantitatively evaluated by comparing mean activity during feedback in these electrodes in the unlearned and learned states. 67 electrodes showed a significant change in HG activity during all targets from the unlearned to learned states (both-sided two-sample *t*-test, Bonferroni corrected, $p < 8.787 \times 10^{-5}$, $20 < N_1 < 75$, $15 < N_2 < 166$). A large portion of these changes corresponded to significant lessening in activation in the frontal cortex and posterior parietal cortex, as depicted in Figure 10a. A smaller portion corresponded to significant increases of activation in areas surrounding the controlling electrode. In order to determine anatomically relevant patterns of activity dynamics, electrodes were assigned to cortical areas using the HMAT atlas (Mayka et al., 2006) and the Talairach daemon (Lancaster & Rainey, 1997; Lancaster et al., 2000). Significant

lessening of activation was found in 31 electrodes corresponding to labels of PMd, PFC, and PPC, as shown in Figure 10b. M1/SC and PMv exhibited ambivalent trends in HG from early to late trials.

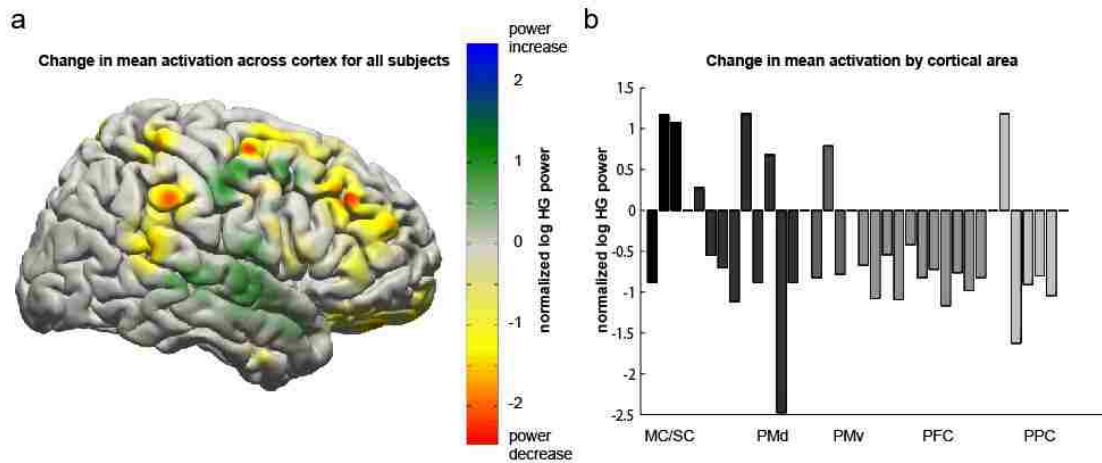


Figure 10 - Cortex-wide activation changes – (a) Spatial distribution of change in mean activation comparing early to late trials for all targets for all subjects. Activations for individual electrodes are normalized against rest periods from the same electrode for a given run. Activation change values are blurred using a 12cm FWHM Gaussian filter. Frontal areas and posterior parietal areas exhibited lessening in task-related activation over the course of BCI use. (b) Change in mean activation for all electrodes showing significant change from early to late trials, classified into approximate cortical areas.

Evaluation of lower frequency cortical activity. HG is the range of frequencies that often receives the primary focus of attention in ECoG studies, partially because until recently (Darvas et al., 2010), it had not been demonstrated that high frequency signals could be obtained using non-invasive methods. However, the volitional modulation of lower frequency cortical rhythms such as the mu-beta rhythm (12-30 Hz) is often used as a control signal in non-invasive BCIs (Fabiani et al., 2004; Lemieux et al., 1997; McFarland et al., 1997; Wolpaw et al., 1991) and is thus worth discussing as a complementary analysis to that of HG activity.

Of the 569 electrodes investigated from the seven subjects, we found 128 electrodes showing statistically significant decreases in mu-beta (12-30 Hz) activity between feedback during up targets as compared to rest (left-sided two-sample t-test, Bonferroni corrected, $27 < N_1 < 98$, $51 < N_2 < 187$; $p < 8.631 \times 10^{-5}$). These findings are shown in Figure 11. Further, of that same 569 electrodes, we found 123

electrodes showing a statistically significant decrease in mu-beta activity between all targets and rest (left-sided two-sample t-test, Bonferroni corrected, $54 < N_1 < 198$, $51 < N_2 < 187$; $p < 6.284 \times 10^{-5}$).

Electrodes showing significant activity decrease were again distributed throughout frontal and parietal cortices. This spatially diffuse, task-related decrease of mu-beta activity is consistent with previous findings that motor-related desynchronization of mu-beta activity is more spatially widespread than increases in HG activity (Miller et al., 2007). Trial-by-trial patterns showing this activity decrease can be seen in Figure 11.

Using the transition trials defined by HG activity at the controlling electrode, we performed a similar analysis to assess whether activation patterns in the mu-beta range changed over the course of transitioning from a “learned” to an “unlearned” state. Though there were a few electrodes exhibiting significant changes between these two states, as is illustrated in Figure 11, these changes were less spatially organized and lower in magnitude than HG changes over a similar time period.




Figure 11 - Changes in mu-beta activity patterns – (a) Significant decreases (normalized mu-beta power) during up targets for all lateral electrodes for all subjects (left hemispheric coverage projected to right hemisphere) projected on to the Talairach brain. Note widespread cortical desynchronization in frontal and parietal areas. (b - subplot 1) Average mu-beta activity patterns in a given electrode from a specific subject for all trials separated by up and down targets. Subject is specified in the subplot title. Phases of the task (ordered from L to R: rest, cue, feedback, reward) are separated by vertical bars. Dotted line represents SEM. (b - subplots 2 & 3) trial-by-trial mu-beta activity patterns for all trials, separated by up (subplot 2) and down (subplot 3) targets. Trial count is shown on the vertical axis, and breaks in the experimental session of more than 8 hours are denoted with an asterisk (*). Time, as described for subplot 1 is shown on the horizontal axis. The black horizontal bar represents the model separation point derived solely from HG activity in the controlling electrode. (c) Spatial distribution of change in mean mu-beta activation comparing early to late trials for all targets for all subjects. Activations for individual electrodes are normalized against rest periods from the same electrode for a given run to eliminate gain differences between electrodes and non-stationarities within a given electrode over time. Activation change values are blurred using a 12cm FWHM Gaussian filter. (d) Change in mean mu-beta activation for all electrodes showing significant change from early to late trials, classified in to approximate cortical areas.

Alternative methods for assessment of learning states. To assess the robustness of findings with respect to our method for determination of “early” and “late” learning states, we performed additional analyses of distributed changes in HG activity comparing the first 30% and the last 30% of all trials

recorded for each subject, an approach that has been used previously (Koralek et al., 2012). Other than the definition of the transition period, the method for these analyses was identical to what has been described previously in this chapter. As can be seen in Figure 12, we found similar spatial patterns showing less HG activation in PFC and PPC in the last 30% of trials as compared to the first 30%. This suggests that the observed spatiotemporal patterns in HG activity are robust to alternative methods of segmenting the data in to learning states.

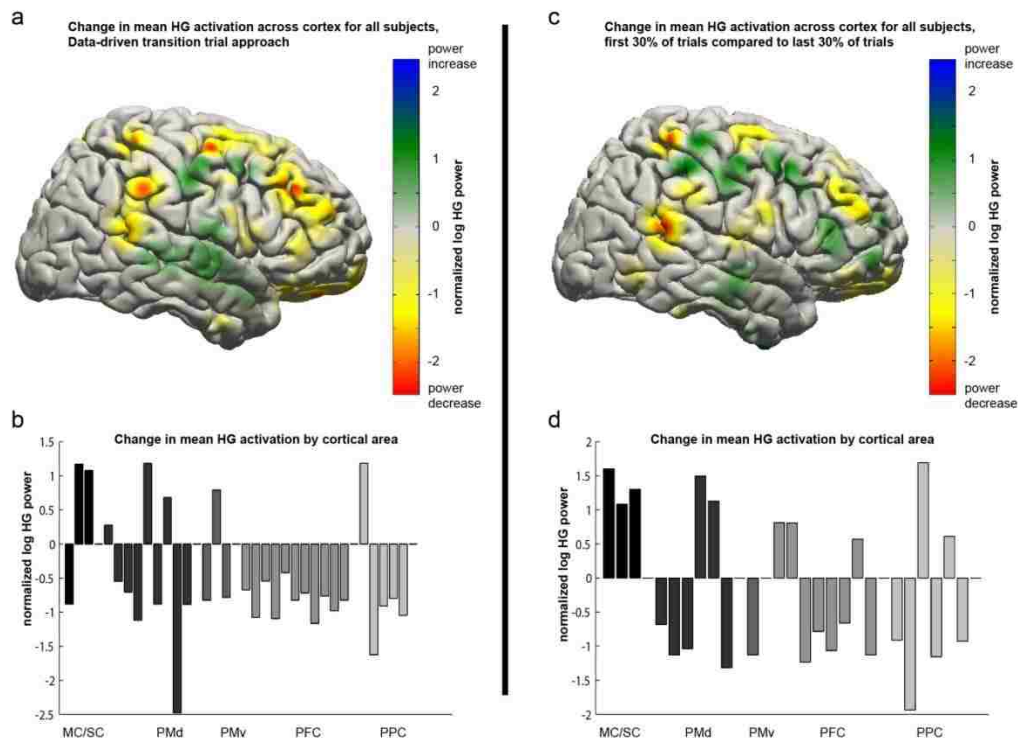


Figure 12 - Changes in HG activity patterns observed with alternative data segmentation method – (a,b) repetition of Figure 4, given for direct comparison with the subsequent two panels. (c) Spatial distribution of change in mean HG activation comparing the first 30% to the last 30% of all trials recorded for each subject. Note similar spatial patterns as compared to panel a. (d) Change in HG activation for all electrodes showing significant change from early to late trials, classified in to approximate cortical areas.

5.4 Discussion

Our results demonstrate that when learning a BCI cursor-control task based on signals recorded from a single electrode over motor cortex, there is a distributed network of cortical areas that is involved in acquisition of skill in the task. This network includes, but is not limited to surrounding sensorimotor and

visuomotor areas: PMd/PMv, PFC, and PPC. We further demonstrated that as a user develops proficiency with the BCI, activity in parts of this network (PMd, PFC, and PPC) decreases, a finding that is indicative of the mental shift from cognitive to automatic task execution often anecdotally reported in human BCI studies.

In this study, we sought to investigate changes taking place in motor networks as users developed both their control strategy and execution proficiency. To do so, we employed a data-driven approach that relied solely on the signal that each subject was using to control the BCI. We found that dorsal premotor, prefrontal and posterior parietal cortices exhibited decreased task-modulated activity as users transitioned from a naive to more experienced state. This transition occurred quickly and was contemporaneous with changes taking place in the BCI control signal.

It is interesting to note the patterns of distributed activity observed across multiple subjects. The time course of activation in pre-motor electrodes parallels tightly that of corresponding M1/SC electrodes. Further, premotor areas were primarily more activated for up targets than down targets, instances in which the controlling electrodes were required to exhibit the same activity. These observations suggest a direct relationship between premotor areas and the controlling electrode. This can be contrasted to activity patterns observed in prefrontal and posterior parietal areas, which exhibited fairly equivalent activation for both up and down targets. These activity patterns are a reflection of effort or engagement on the part of the subject, and correspondingly decrease as the subject gains task automaticity.

In contrast to ECoG BCIs, non-invasive, EEG-based BCIs typically harness power changes occurring in lower frequencies to achieve control. When EEG subjects are performing motor imagery, the typical frequency band chosen is in the mu-beta range (12-30 Hz) as power changes in this band have been demonstrated to be negatively correlated with movement and motor imagery (McFarland et al., 2000). When using the methods described above to assess whether similar changes were taking place in the

mu-beta range as subjects develop experience with an ECoG BCI, we found that though activity in the mu-beta band was strongly task-modulated, it did not undergo changes during learning to the same degree that we observed them in HG activity. It is important to note, however, that as the cursor was not being driven by mu-beta activity changes, the impact of this observation to non-invasive BCIs will require further investigation.

A logical and necessary extension of these findings is to investigate the roles that subcortical networks play in this same learning process. Previous studies have demonstrated the vital and differential roles of the basal ganglia and cerebellum during motor sequence learning and motor adaptation (Doyon & Benali, 2005; Hikosaka et al., 2002), but the involvement of these sub-cortical networks in the process of BCI skill acquisition remains an open question, with implications for both fundamental neuroscience and the incremental improvement of BCI frameworks. Recent work performed by Koralek and colleagues (Koralek et al., 2012) has demonstrated that the striatum is involved in and critical to development of proficiency with a BCI in a rat model. A notable finding, given that effective use of the BCI in that study did not explicitly require recruitment of the motor system outside of motor cortex, yet task performance was degraded in subjects with impaired cortico-striatal interaction. Investigation of the role of these structures in humans will have to be left to other recording modalities, as ECoG provides information only regarding activity near the cortical surface, and is subject to spatial undersampling based on the distribution of electrodes.

The distributed dynamics in cortical activity that we have demonstrated here have significant implications in development of co-adaptive BCI frameworks. BCI investigators have recognized the need for BCI architectures that accommodate the dynamic nature of the neural signals used as inputs (Vidaurre et al., 2010, 2011), but our findings suggest that a number of the most commonly used methods (e.g., common spatial patterns - Müller-Gerking et al., 1999) may require updating to handle both the spatial and temporal variability in input signals.

The ability to detect correlates of cognitive load during BCI task learning and execution holds great potential for expanding the current limitations of BCIs. Continuous control of a BCI has been demonstrated in two dimensions using ECoG (Schalk et al., 2008) and EEG (Wolpaw & McFarland, 2004), with degrading performance as users attempt control in more dimensions. Real-time monitoring of cognitive load during BCI skill acquisition would allow for titrated increases in task complexity, potentially facilitating an increase in the total number of dimensions that could be simultaneously controlled, thus allowing operation of more complex devices.

Our results also demonstrate the potential that BCI holds as a technique for probing neural systems *in vivo* (Moran, 2010). By applying specific task requirements but allowing users to employ native learning strategies we observed that distributed cortical networks are involved in the cognitive phase of BCI skill acquisition, but the degree of involvement of these networks lessens as users transition to automatic execution of the task. Future work will allow us to probe whether activity in this network is renewed when task dynamics are perturbed and the user is required to adapt to novel task conditions.

5.5 Related publications and presentations

[1] **JD Wander**, T Blakely, LA Johnson, F Darvas, KJ Miller, RPN Rao, JG Ojemann. Dynamics of distributed cortical activity demonstrated over the course of learning to use a brain-computer interface.

Proceedings of the 42nd Annual Society for Neuroscience conference, New Orleans, LA, October, 2012.

[2] **JD Wander**, et al. Distributed cortical adaptation during learning of a brain-computer interface task.

Proceedings of the National Academy of Sciences (2013).

6 Multi-site cortical interactions during BCI use

Use of a brain-computer interface (BCI) has been demonstrated to be a learned skill that involves recruitment of neural populations that are directly linked to BCI control as well as those that are not. The nature of interactions between these populations, however, remains largely unknown. In this chapter, we employed a data-driven approach to assess the interaction between both local and remote cortical areas during the use of an electrocorticographic BCI. Comparing the controlling area with remote areas, we evaluated relationships between the amplitude envelopes of band limited powers as well as non-linear phase-phase interactions. **We found amplitude-amplitude interactions in the high gamma (HG, 70-150 Hz) range that were primarily located in the posterior portion of the frontal lobe, near the controlling site, and non-linear phase-phase interactions involving multiple frequencies (cross-frequency coupling was observed between 8-11 Hz and 70-90 Hz) taking place over larger cortical distances.** Further strength of the amplitude-amplitude interactions decreased with time, whereas the phase-phase interactions did not. **These findings suggest multiple modes of cortical communication taking place during BCI use that are specialized for function and depend on interaction distance.**

6.1 Introduction and background

Direct communication between brain and machine provides a powerful platform for both the development of clinical therapies and scientific inquiry. By providing the brain with a completely novel output pathway, experimentalists have an opportunity to observe the ways in which the brain responds to and develops control over this new output mechanism. A number of studies have demonstrated that the use of a brain-computer interface (BCI) is a learned skill (Carmena et al., 2003; Ganguly & Carmena, 2009; Moritz et al., 2008; Schalk et al., 2008; Wolpaw & McFarland, 2004), and that the brain can learn this skill more effectively when the transformation that maps neural activity to BCI control is consistent

(Ganguly & Carmena, 2010). Further, it has been demonstrated that the nature of the neural signals being used to drive the BCI change with practice (Ganguly & Carmena, 2009) and that there are also changes in neural activity in populations that are not directly linked to BCI control [both local to the controlling site (Ganguly et al., 2011); and at more remote sites (see Chapter 5)]. The mechanisms underlying learning of BCI control have many similarities to those for learning motor control (Green & Kalaska, 2011). Repeated BCI training can have lasting effects on motor networks, altering functional connectivity in cortico-thalamic networks during execution of a finger-tapping task (Young et al., 2014). To date, there have been no systematic studies of cortico-cortical interaction during BCI use. Other than a recent study demonstrating the need for corticostriatal interaction during the BCI learning process in a rodent model (Koralek et al., 2012), we have little understanding of the networks involved in acquisition of the neuroprosthetic skill.

The brain is a vastly distributed and parallelized system, requiring effective and efficient communication between both neighboring and distant neural populations (Buzsáki et al., 2013). Correspondingly, of equal interest to within-region changes in synchrony of neural activity are changes in interactivity between regions. Cortical, cortico-subcortical, and cortico-muscular coherence have all been observed in the mu range (8-12 Hz) during slow movements (Gross et al., 2002). Similar observations have been made regarding long-distance synchrony in the beta range, both in cortico-cortical interactions between M1, SC and PPC during a visual discrimination task (Brovelli et al., 2004). Another form of phase-phase synchrony, the phase locking value (PLV) (Lachaux et al., 1999) has been used to quantify linear interactions during execution of a cognitive task (Doesburg et al., 2008). While these examples are restricted to within-frequency phase-phase interactions, it has been suggested that cross-frequency (i.e. non-linear) interactions could reflect much richer cortical interactivity (Buzsáki & Draguhn, 2004).

Interaction between neurons and neural populations encompasses a variety of neural mechanisms, including coordinated increases in firing rates, periodic synchrony, and complex feedback loops (see

Salinas & Sejnowski, 2001, for review). The link between these mechanisms and their corresponding signatures in population-based physiological signals is incompletely understood. Each of these mechanisms may manifest differently in population-scale neural recordings as anything from changes in raw covariances to detectable differences in the complex non-linear coupling of spectral components. It has been hypothesized that long-distance cortical communication is mediated by a relatively small number of direct connections, because direct connection of all communicating cells across these long distances would be biologically infeasible (Bullmore & Sporns, 2012; Buzsáki et al., 2004). Further, Buzsáki and colleagues argued that oscillatory activity is central to the maintenance of efficient cortical information flow within increasingly large and complex mammalian cortices (Buzsáki et al., 2013). Such a network model would be well served by the use of oscillatory synchrony, or rhythmic interactions, to allow for maximal efficiency in processing (Schnitzler & Gross, 2005).

However, there are various ways in which cortical field potentials can be related. Whether these different relationships play differing roles in cortical processing, or whether they are indicative of a single underlying network of connectivity remains an open question. It has recently been shown that though high-frequency, amplitude-amplitude correlations in cortical field potentials are predictive of underlying local structural connectivity, this relationship deteriorates over longer distances (Keller et al., 2014). Coupling this with the theory that oscillatory synchrony is critical to long-range cortical communication leads to a testable hypothesis of distance-specificity by interaction type: when observing simultaneous amplitude-amplitude and phase-phase interactions taking place during a cognitive task such as BCI use, the former will be observed over shorter distances and the latter over longer ones.

In chapter 5, we demonstrated frontal and parietal regions that were active during the initial use of BCI using ECoG signals from motor cortex. These areas became less active with repeated use. Here, we examine the interactions between areas outside of the site used for BCI control with reference to the signal from this electrode. We hypothesized that there exist task-driven amplitude-amplitude and

phase-phase interactions observable in the ECoG field potential between the controlling electrode and remote cortical structures and that these two interaction types are present on differing spatial scales.

6.2 Materials and methods

Subjects and motor screening. Ten human subjects (1 female, mean age 26.9y) with intractable epilepsy were implanted with platinum sub-dural ECoG grids (AdTech, Racine, WI) for the clinical purpose of seizure focus localization and resection. These subjects were monitored for between four and ten days before removal of the arrays and surgical resection of the seizure focus. During this time the subjects participated in multiple recording sessions, spread over one to three days.

This study was a retrospective analysis of previously collected ECoG data. To determine which subjects were eligible for inclusion in this study, the following criteria were applied: (1) subjects needed to have participated in the 1-D RJB task; and (2) subjects needed to perform the task above chance levels in order to demonstrate intentional control. Of the 11 subjects originally eligible per these inclusion criteria, one subject was eliminated from this study based on extreme cortical distortion due to a previously resected peri-central cavernous malformation.

For detail on motor screening see section 4.2. For the remainder of this chapter, the single electrode selected during motor screening will subsequently be referred to as CTL.

Table 3 - Demographic information for subjects included in the analysis of multi-site cortical interactions observed during BCI use. Abbreviations: R right, L left, F frontal, P parietal, T temporal, O occipital, M Act refers to the motor task associated with the controlling electrode. Note, S9 was removed from the study due to extreme cortical distortion.

ID	SID	Gender	Age	M Act	Coverage
S1	30052b	M	29	Tongue	R-F/T
S2	4568f4	M	27	Tongue	R-F/P/T
S3	3745d1	M	14	Tongue	L-F/T
S4	26cb98	M	22	Tongue	R-F/P/T
S5	fc9643	F	26	Tongue	R-F/P/T
S6	58411c	M	54	Hand	L-T
S7	0dd118	M	11	Hand	L-F
S8	7ee6bc	M	29	Hand	R-F/P/T
S10	f83dbb	M	19	Hand	R-T
S11	7662c2	M	38	Tongue	R-F/T

ECoG data collection and the BCI task. See section 4.2.

Cortical reconstructions and anatomical labeling. See section 4.2.

Evaluation of behavioral performance. Subjects' individual performance levels were calculated as the fraction of completed trials wherein the subject successfully ended the trial in the target area. Though theoretical chance performance on this task is 0.5, the behavioral performance necessary to be significantly greater than chance was dependent on the number of trials performed and thus varied from subject to subject. Confidence intervals on chance performance were evaluated on an individual basis by characterizing the distribution of average task performance under the null hypothesis that success and failure were equally likely outcomes on any given trial. To synthesize chance performance data we drew N random samples from the binomial distribution where N was the number of trials conducted by each subject; the average of these samples was one sample of chance performance under the null hypothesis. The distribution of chance performance was characterized by repeating this process 1000 times.

Data pre-processing. Data were first manually inspected for any channels or time periods that contained obvious non-physiologic artifact or substantial inter-ictal activity. For each subject, the data were re-referenced by subtracting the common average among all good channels. Signals were then notch-filtered to remove line noise using 4th-order Butterworth filters at 60 and 120 Hz. For the purposes of short-time windowed correlation (STWC) analyses, time-variant spectral estimates were extracted by bandpass filtering the signals using 4th-order Butterworth filters and then taking the magnitude of the Hilbert transform. Spectral estimates were derived for the canonical frequency bands, mu/alpha (μ/α ; 8-12 Hz) and beta (β ; 12-24 Hz), as well as for the high-gamma range (HG; 70-150 Hz). These spectral estimates were then temporally smoothed using a 47 msec FWHM Gaussian window. Finally, for computational tractability of remaining analyses, signals were then resampled to 400 Hz.

See Figure 13 for an overview of the post hoc signal processing pathway.

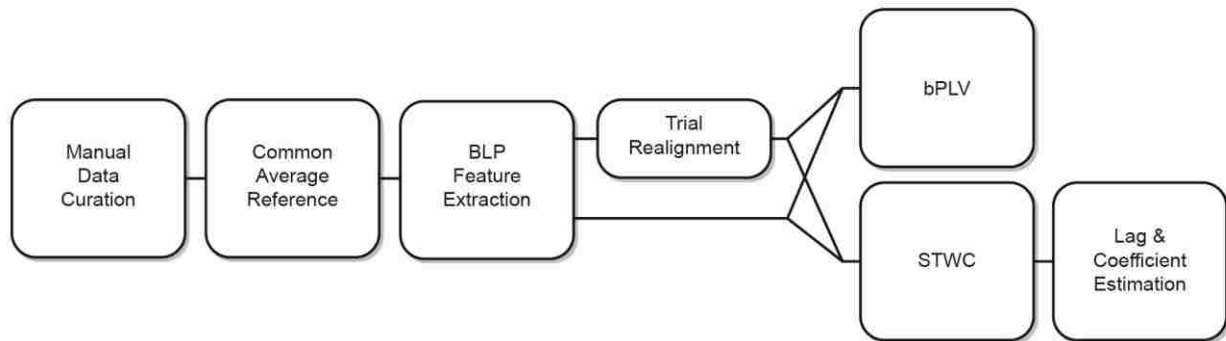


Figure 13 - Post-hoc analysis workflow. Abbreviations: band-limited power (BLP), bi-phase locking value (bPLV), and short-time windowed covariance (STWC).

Reaction-time estimation and trial realignment. Because of the potential for trial-to-trial variability in response time to the task, in addition to performing interaction analyses on trials aligned on cue presentation (cue-locked), we also performed these analyses on trials that were realigned based on initial onset of HG activity at CTL (response-locked). This allowed us to investigate both cue- and response-related interactions. Identification of this onset was performed as follows: First, each trial was temporally smoothed using a 470 msec FWHM Gaussian window. Then a pre-onset baseline value was

defined as the lowest value in the smoothed HG that occurred in the first second of the feedback period. The time at which this baseline value occurred was also noted. Next, the maximum value that occurred after the pre-onset baseline and before two seconds into the feedback period was also determined. The HG onset was defined as the first point after the pre-onset baseline when half the distance between the baseline and the maximum was crossed.

It is noteworthy that this approach is only suitable when activity at the controlling electrode changes during feedback relative to rest. In the previous chapter we demonstrated that subjects typically only modulate HG above baseline for up-targets, thus all interaction analyses computed on response-locked trials are based solely on activity changes during up-target trials.

STWC Analyses. Several methods exist to test for cross-frequency coupling: phase-amplitude coupling (PAC) (Canolty et al., 2006; Miller et al., 2012; Miller, Hermes, et al., 2010), bispectrum- and bicoherence-based measures (Nikias & Mendel, 1993; Nolte et al., 2004) which involve both phase and amplitude, pure phase-based measures (i.e. the bi-phase locking measure) (Darvas, Ojemann, et al., 2009) and amplitude-amplitude coupled measures [e.g. dynamic causal modeling (Friston et al., 2003) and short-time windowed covariance (STWC) (Blakely et al., 2014)]. While all these methods interrogate neural data for some form of cross-frequency interaction, they are each sensitive to different mechanisms that produce this cross-frequency coupling. While PAC and bi-spectral/coherence have been extensively discussed in the literature, in this chapter we focused on pure phase- and pure amplitude-based measures, which can be seen as testing for large ensemble to large ensemble interaction (amplitude-amplitude) and highly synchronized ensemble to highly synchronized ensemble interaction (phase-phase), where the groups of neurons involved can be small.

We assessed transient temporal amplitude-amplitude correlations in HG activity between CTL and remote electrodes using the normalized form of the STWC measure (Blakely, Ojemann, & Rao, 2014).

$$C(x, y, t, \tau, \delta) = \frac{1}{\sigma_{x,t,\tau}\sigma_{y,t+\delta,\tau}} \sum_{i=t-\frac{\tau}{2}}^{t+\frac{\tau}{2}} \frac{(x_i - \bar{x})(y_{\delta+i} - \bar{y}_{\delta})}{(\tau + 1)} \quad (2)$$

where $t \in [1, T]$ and $\delta \in [-\Delta, \Delta]$ and x and y are the two signals being considered, τ is the window size over which the correlation is being calculated, t is the time (or sample) within the signal x , and δ is the lag of the window from y with respect to the window from x . $\sigma_{x,t,\tau}$ and $\sigma_{y,t+\delta,\tau}$ are the sample standard deviations from the two data windows. This method is specifically suited to teasing out amplitude-amplitude interactions in neural signals that are not only transient (e.g., event-driven), but also potentially occur at slightly different points in time in each of the two signals.

Individual STWC maps were calculated for each trial using a window width of 500 msec and a maximum lag of 300 msec. Average STWC maps were then generated separately for cue-locked and response-locked trials. We evaluated HG-HG, HG- β , and HG- μ/α interactions in order to observe both within-frequency and cross-frequency amplitude-amplitude effects.

To isolate interactions relevant to task execution, we only evaluated interactions occurring within the first second of the feedback period (cue-locked trials) or ± 500 msec from HG onset at CTL (response-locked trials).

From each significant interaction, we extracted both a maximal STWC coefficient from the average STWC map as well as the corresponding lag at which this coefficient occurred. The former provides information regarding the relative strength of the interaction whereas the latter provides information as to the relative timing of the activity changes between the two areas.

Significance of STWC interactions was evaluated using a bootstrap approach on surrogate neural data. Using 100 iterations per channel pair, average STWC maps were calculated on trial-shuffled, phase-randomized neural signals. Phase randomization was used to destroy any temporal interaction between

the two channels while preserving the individual power spectral characteristics of each channel. STWC maps were calculated as above and for each interaction the maximal coefficient from the average maps for all channel pairs was retained to characterize the multiple comparison-corrected distribution of maximal STWC coefficients one would expect to see under the null hypothesis of no interaction between electrodes. Only STWC coefficients greater than 95% of this distribution were considered significant; reported p values are the probability of seeing the observed STWC coefficient under the null hypothesis of no interaction.

Bi-phase coupling. The bi-phase coupling value (bPLV) is a non-linear measure of cortical interaction.

The bPLV can be computed from the time varying phase of the signal for a pair of frequencies as:

$$B_{XYZ}(t, f_1, f_2) = \left| \frac{1}{N} \sum_{j=1}^N e^{i(\phi_X^j(t, f_1) + \phi_Y^j(t, f_2) + \phi_Z^j(t, f_1 + f_2))} \right| \quad (3)$$

Here $\phi_X^j(t, f_1)$ is the phase of signal X at frequency f_1 and time t for the j^{th} trial, $\phi_Y^j(t, f_2)$ is the phase of signal Y at frequency f_2 and $\phi_Z^j(t, f_1 + f_2)$ is the phase of the coupled signal Z (Darvas, Ojemann, et al., 2009). The three signals (i.e. the sources X and Y and the target Z) can either be located in different positions, or in any combination be distributed across one to three electrodes. Here we choose a configuration, where the source signals, X and Y , reside in one location and the target is located at a different electrode. This configuration of the bPLV is similar to the one used in our earlier studies of bi-phase coupling in the motor system (Darvas, Miller, et al., 2009) and allows for an interpretation of causal directionality in the Granger sense, as here the phase at the target location is predicted by the phases of the source location, but not vice versa. Similar to our earlier studies, we compute trial-wise phase coupling, but we limit the source signals X and Y in this study to CTL and test for interactions to all other electrodes in the montage.

For a single pair of electrodes, we compute a frequency by frequency by time bPLV map, for interaction frequencies limited to coupling from 7-25 Hz to 70-100 Hz to a resulting target frequency ranging from 77 to 105 Hz. This range is motivated by our earlier studies (Darvas, Miller, et al., 2009), which found alpha/beta range coupling to high gamma frequencies during overt movement, as well as by the fact that this particular frequency range avoids the power line frequency at 60 Hz and its harmonics.

We use the continuous complex Morlet wavelet to compute a time varying phase for each frequency pair and the corresponding interaction frequency with a frequency resolution of 1 Hz. The resulting bPLV map per pair thus contains 589 time series, each 2200 samples long, covering the time from -3 s prior to the beginning of BCI control to approx. 2.5 s post control onset, when the trial ends. Maps are computed for all subjects across all channels in the montage over all trials.

We integrate the bPLV time series for each frequency pair from the onset of BCI control to 1 s post onset to test our hypothesis of a task specific increase in bPLV during execution of BCI control.

Even after integrating the bPLV time series over a time interval, we are left with a large number of potential interactions and in the absence of a specific hypothesis about which frequencies and channel pairs should increase during BCI control, we use a statistical threshold to identify significant bPLV changes. While for individual bPLV values, an analytical expression for the null-distribution exists (Darvas, Ojemann, et al., 2009), we have no such description for the time integrated bPLV and thus must resort to non-parametric tests. We employed a similar maximum statistic approach as the one applied to STWC coefficients. We use trial shuffling (Lachaux et al., 1999) to generate new samples of the time integrated bPLV, where we randomly shift trials between the controlling electrode and the target electrode. Here we assume that there exists no coupling on the time scale of full trials (which last >5s) and thus this method will generate an appropriate null-hypothesis. We generate 10,000 resamples per

channel pair, but for each resample, after computing the integrated frequency by frequency map, we only retain the maximum value across all frequencies.

We then compute a p -value by comparing the original time integrated bPLV against the histogram of maximum values. This way we avoid having to control for multiple comparisons across the whole frequency by frequency map, which has a variable resolution, which would in case of a simple Bonferroni correction lead to a too conservative threshold.

Since channel pairs can be considered independent, we Bonferroni correct the resulting p -value from the maximum statistic by the total number of pairs examined.

6.3 Results

Behavioral performance. As is necessitated by our study inclusion criteria, all subjects performed above chance levels on the BCI task (N varies by subject, *binomial test*, $p < 0.05$). This is important to subsequent analyses as it serves to demonstrate that subjects had intentional control of the neural signal being used to control the BCI. Figure 14 shows that performance levels were above 95% chance performance confidence intervals for each subject; see Table 4 for additional detailed behavioral data.




Figure 14 - Behavioral performance for all subjects included in interaction analysis. Chance task performance denoted as black horizontal line. 95% Confidence intervals on chance performance differ by subject and are shown as a dashed line.

Reaction time and trial realignment. For up trials, where HG onsets could be calculated, we determined the average reaction time across the ten subjects to be 775 msec (\pm 140 msec std). Though not statistically significant there was a subtle correlation between an individual subject's reaction time and task performance (Pearson's $\rho = -0.39$, $p = 0.26$) that may have been related to attentional vigilance or other factors. Figure 15 shows response-locked HG activity from CTL for two subjects. Individual subject reaction times and standard deviations are reported in Table 4.




Figure 15 - Exemplar plots (from two subjects S2 and S5) showing alignment of up trials by HG onset . Color-coded normalized HG activity from the CTL electrode is shown for all trials. Trials were realigned based on onset of HG activity, such that HG onset occurs at $t=0$. Black dots denote the actual moment of cue presentation (beginning of the feedback period) for each trial.

Table 4 - Individual behavioral results. All represents the total number of valid BCI trials performed; performance is the fraction of those trials that were successful. In the second section we have listed the total number valid electrodes recorded from each subject. Lastly, response time is the mean time, in msec, between the start of the BCI trial and an increase in HG activity at CTL.

SID	BCI Trials		Electrodes	Response Time
	All	Performance	<i>N</i>	Mean \pm SD msec
30052b	84	0.8	82	908 \pm 471
4568f4	108	0.78	85	753 \pm 385
3745d1	39	0.72	63	886 \pm 500
26cb98	97	0.74	44	743 \pm 482
fc9643	164	0.9	94	788 \pm 469
58411c	120	0.63	63	821 \pm 412
Odd118	68	0.76	63	733 \pm 521
7ee6bc	110	0.66	56	816 \pm 484
f83dbb	89	0.64	61	986 \pm 480
7662c2	50	0.78	82	596 \pm 370

STWC Interactions. When performing STWC interactions on response-locked trials, we identified 31 total electrodes, from a total of 9 of the 10 subjects that exhibited significant STWC interactions with the CTL electrode ($p < 0.05$; see bootstrap methods for detail). Exemplar interactions are shown in Figure 16.




Figure 16 - Exemplar response-locked STWC maps from four subjects with significant STWC interactions showing a remote electrode (a) coactivated with, (b) leading, (c) lagging, or (d) both leading and lagging the CTL electrode. Significant interactions are circled in a white boundary. The solid black horizontal line depicts a lag of zero and dashed horizontal and vertical lines intersect at the peak STWC coefficient that was extracted and used in subsequent analyses.

We note that though the electrodes considered in the STWC analyses come from a large number of cortical areas, as is depicted in Figure 17b, and that there are a number of areas outside of traditional motor regions that are task-modulated during BCI, the areas interacting with the control electrode are

almost exclusively contained to the posterior portion of the frontal lobe. Nine of the 31 electrodes interacting with CTL were found in ventral PMv meaning that 29% of observed significant interactions occurred within PMv, though less than 5% of all electrodes considered were over that area. Further, the interactions between this area and CTL showed larger overall STWC coefficients ($N_1 = 9, N_2 = 22$, two-sample t -test, $p = 0.0017$) and more of a tendency to lead CTL ($N_1 = 9, N_2 = 22$, two-sample t -test, $p = 0.027$) than interactions involving other regions. We also identified significant interactions between CTL and M1, SC, PMd and a small number of extra-motor areas.

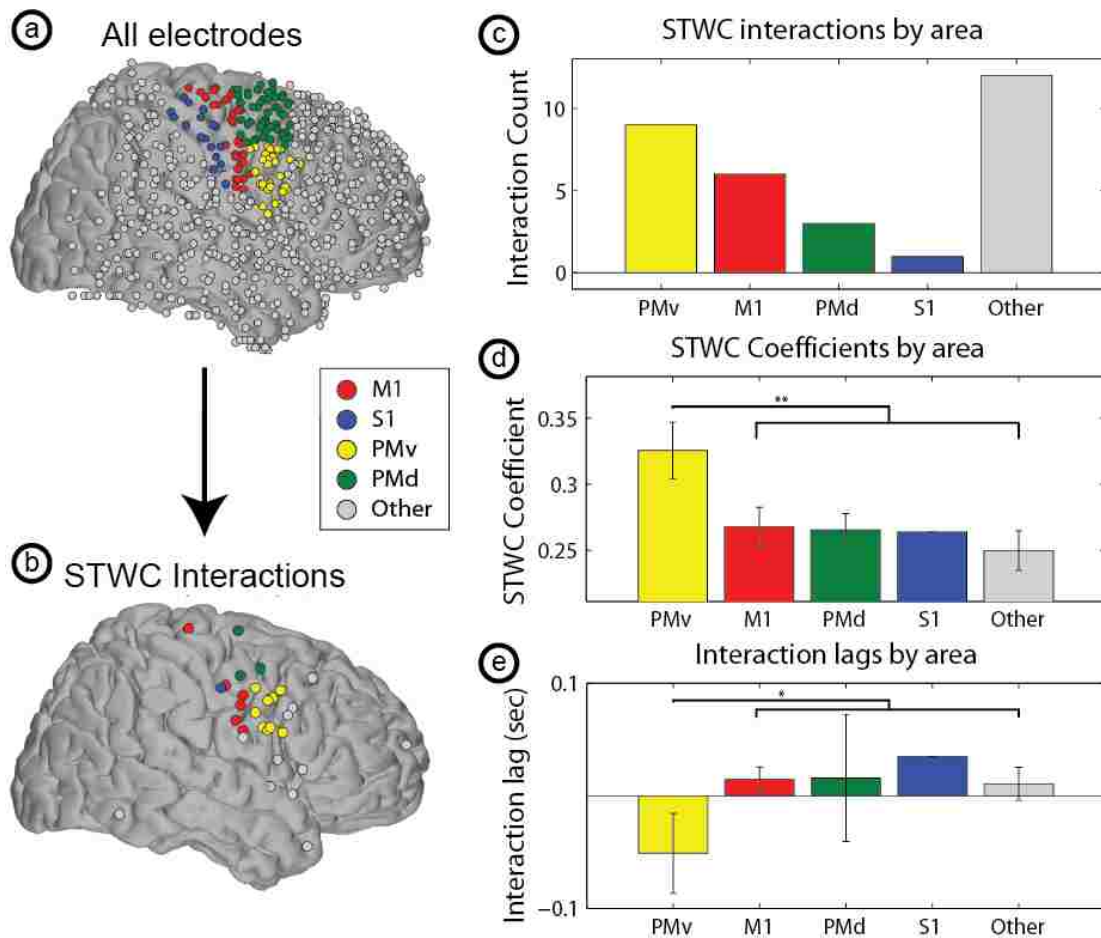



Figure 17 - Spatial distribution of significant STWC interactions. Subplot (a) depicts the locations of all non-CTL electrodes across all subjects whereas (b) shows the subset of these electrodes that were involved in significant response-locked STWC interactions. Note that the majority of significant interactions were seen in the posterior portion of the frontal lobe. Subplot (c) shows the frequency of interactions in various cortical regions as defined by the HMAT atlas. Subplots (d and e) show average STWC coefficients and lags (respectively) across those same regions. One star (**) denotes $p < 0.05$ and two stars (***) denote $p < 0.01$.

STWC Analyses on cue-locked trials. STWC analyses reported above were performed on trials that had been realigned based on the onset of HG activity at CTL. Additionally, in order to determine whether the observed effect was tightly bound to HG onset (and not to cue presentation), we re-performed STWC analyses on cue-locked trials. Other than the alignment of individual trials during creation of average STWC maps, the procedure for this was identical to what was used for response-locked trials.

We identified 23 total electrodes, from a total of 7 of 10 subjects that exhibited significant interactions with the CTL electrode ($p < 0.05$; bootstrap approach described above). Again these electrodes were primarily located in the posterior portion of the frontal lobe; the spatial distribution of electrodes identified when applying STWC to unaligned trials is qualitatively similar to what was described for response-locked trials (See Figure 18). Though slight, the decreased number of significant interactions on cue-locked trials relative to aligned trials (23 significant interactions as opposed to 31), combined with the fact that on average, significant STWC peaks calculated on unaligned trials occurred $478 (\pm 36$ SEM) msec after cue presentation suggests that the observed amplitude-amplitude interactions are involved more with the execution of motor imagery than the immediate response to the cue.

Response-locked STWC Interactions



Cue-locked STWC Interactions




Figure 18 - Spatial comparison of non-CTL electrodes involved in HG STWC interactions for response-locked and cue-locked trials

In addition to performing STWC analyses on HG amplitude-amplitude interactions, we also calculated STWC maps for expressly non-linear (i.e. cross-frequency) amplitude-amplitude interactions using response-locked trials. For each electrode pair (including CTL and a single remote electrode – denoted with a subscript of R), we considered the following cross-frequency couplings: $HG_{CTL} \leftrightarrow \beta_R$, $HG_{CTL} \leftrightarrow \alpha_R$, $\beta_{CTL} \leftrightarrow HG_R$, and $\alpha_{CTL} \leftrightarrow HG_R$. Across all comparisons, of which there were 2732, we found only two significant cross-frequency STWC interactions. We interpret this to indicate that there are minimal cross-frequency amplitude-amplitude interactions, and that phase-phase relationships better capture the interactions between these frequency bands.

bPLV Coupling. We observed significant cue-locked bPLV interactions in 8 of the 10 subjects. All significant interactions were ‘outgoing’ from CTL, meaning that the phases of an alpha and HG frequency at CTL were predictive of the phase at the sum of those two frequencies at a remote site. Results of the grand average of the bPLV for individually significant alpha-HG interactions based on cue-locked trials are shown in Figure 19. Since bPLV significance for individuals was computed as an integrated value over a time interval, and the bPLV measures phase synchrony on a time scale of <10 ms, the individual time series over a scale of more than 5s can appear noisy and highly variable over time. However, on the group level, a stable trend emerges: an overall increase in bPLV between 0.5 s and 1 s post task onset, when comparing up vs. down targets, or up targets vs. baseline.

We applied the same analysis (see methods) to response-locked trials, but under these conditions we observed no significant bi-phase coupling. Additionally, we evaluated linear phase-phase interactions from 1 to 200 Hz using the standard PLV, and observed no robust trends across subjects in within-frequency coupling.




Figure 19 - Group-average of significant bPLV interactions for up and down targets . Dark is the average for up-targets, light for down targets. The shaded areas show the 84% confidence interval (1 standard deviation) of the group average. Vertical black bar at t=0 represents the onset of the feedback period.

Changes in interactions across skill acquisition. In order to begin to understand whether either of these interactions are indicative of the skill acquisition process, or are present during task execution in general, we re-evaluated all significant STWC and bPLV interactions using subsets of the trials. Comparing interactions from early trials (the first half of trials executed by a subject) and late trials (the second half of trials executed by a subject) we saw no statistical difference in bPLV interactions between remote electrodes and the CTL electrode (see Figure 20). This is in contrast with STWC interactions where we saw a significant decrease in median (per subject) interaction strength (paired t -test, $N=9$, $p = 0.016$) from early to late trials.

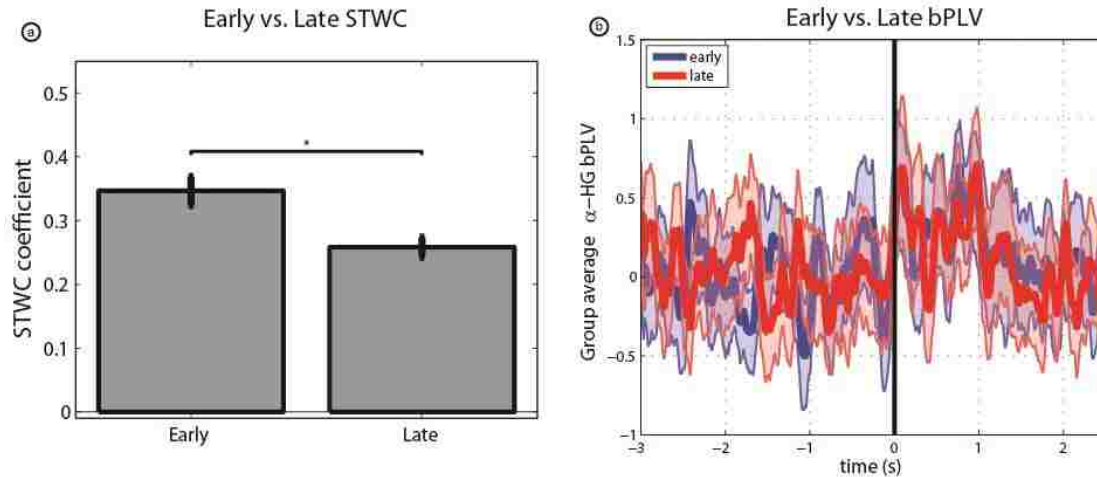


Figure 20 - Changes in STWC and bPLV interactions over time. (a) Depicts change in median (per subject) response-locked STWC coefficients from early to late trials across all significant interactions. (b) Demonstrates that there was no significant change in grand-average bPLV from early to late trials. Black vertical bar at t=0 represents the onset of the feedback period. The shaded areas show the 84% confidence interval (1 standard deviation) of the group average. One star (*) denotes $p < 0.05$.

One potential criticism of this finding is that though STWC is normalized for differences in amplitude, like any correlation measure, it *is* sensitive to changes in signal-to-noise (SNR). As we observed in chapter 5, a number of cortical areas exhibit significant changes in task-driven HG activity that may – assuming a stationary noise floor – impact STWC strength over the course of skill acquisition. To control for this possibility we calculated SNR for all electrodes over the early and late trial periods, and repeated the above analyses excluding all interactions that involved electrodes with a significant decrease in SNR from early to late trials. In this case we still observed a significant decrease in median STWC strength ($N=8$, $p=0.001$) from early to late trials.

Comparison of spatial distribution of STWC and bPLV interactions. There are a number of electrophysiological studies that suggest that the very nature of amplitude-amplitude and phase-phase interactions are different (for review see Canolty & Knight, 2010; Schnitzler & Gross, 2005). Though both may be indicative of information flow between cortical areas, the spatial and/or temporal scales over which these interactions take place may be quite distinct. With respect to spatial extent, we found this to be the case. Though we considered interactions from all cortical areas with electrode coverage, we

primarily found significant STWC interactions close to primary motor cortex, often in PMv, or other nearby cortical regions. bPLV interactions, on the other hand, were much more spatially distributed, extending to PMd, the STG and prefrontal cortex. Considering the median STWC-to-CTL distance for each subject as a single observation to adjust for repeated measures within subjects, we found that the distance covered by STWC interactions was, on average, 23.22 mm, whereas the mean distance spanned by bPLV interactions was 38.35 mm, and that the distributions of these distances were significantly different (two-sample *t*-test, $N_1 = 9$, $N_2 = 8$, $p = 0.04$). Figure 21 provides additional detail as to the spatial distributions of these two interaction types.

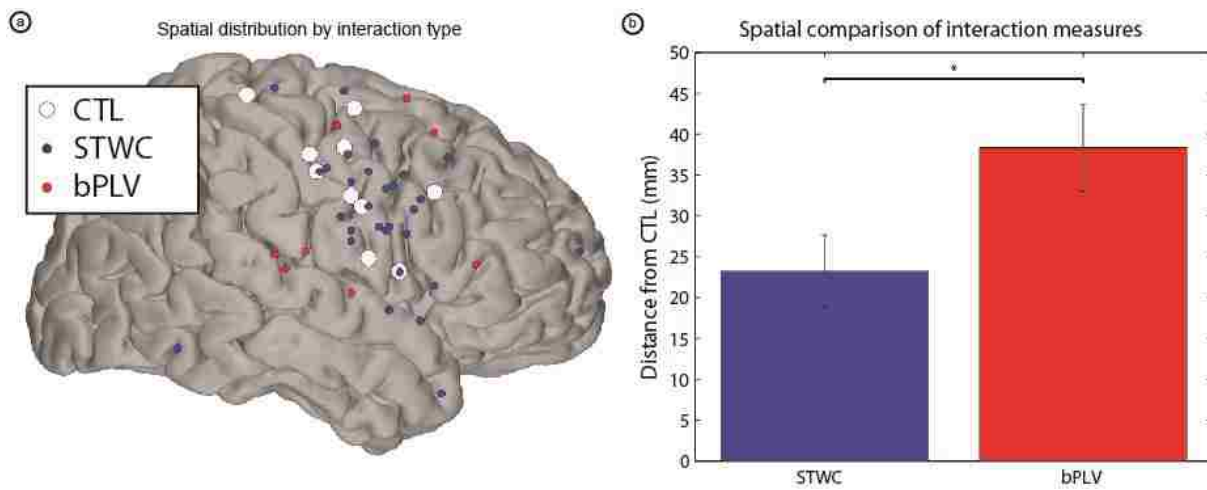


Figure 21 - Comparison of spatial distribution of significant STWC and bPLV interactions. (a) Shows all electrodes involved in significant STWC or bPLV interactions across all subjects. Note the spatial localization of electrodes participating in STWC interactions to the posterior portion of the frontal lobe, and the slightly broader distribution of bPLV electrodes. (b) Provides quantification of this effect, comparing median STWC-CTL distances (per subject) with bPLV distances. One star (“*”) denotes $p < 0.05$.

6.4 Discussion

In this chapter we have demonstrated the presence of both HG amplitude-amplitude (STWC) and cross-frequency phase-phase (bPLV) interactions between cortical areas during BCI use. Coupling this with the fact that a number of cortical and subcortical regions have been shown to be active during the BCI task (see chapter 5, as well as Koralek et al., 2012), it appears that execution of the neuroprosthetic skill is a coordinated effort involving multiple cortical areas. Additionally, we have demonstrated that cross-

frequency phase-phase and within-frequency amplitude-amplitude interactions occur simultaneously during a cognitive task, but exhibit distinctly different spatial scales, which is indicative of at least two different modes of trans-cortical communication during BCI use. Lastly, we have demonstrated that the nature of these amplitude-amplitude interactions changes over the course of skill acquisition for a subgroup of the communicating regions. Our findings suggest the presence of multiple mechanisms of cortico-cortical communication that may play differing roles in task execution.

Though our analyses included every implanted electrode for each subject, we found that the vast majority of significant STWC interactions occurred in or near the posterior portion of the frontal lobe and cover relatively short cortical distances (~2 cm). This not only implicates these regions in successful execution of a BCI driven by M1-derived control signals, but speaks to the relative temporal consistency of interactions between these areas and the controlling area and the possibility that these regions are responsible for similar facets of task execution. Given the tendency of neural circuits to optimize with BCI training (Wodlinger et al., 2014), this may in part explain the observed decrease in STWC over the course of skill acquisition. Along those same lines, it is interesting to note the relative absence of significant STWC interactions between CTL and cortical areas further upstream in the action portion of the perception-action cycle (Fuster, 2000). It is noteworthy that the prefrontal cortex, which we demonstrated in chapter 5 to be active during BCI task execution, participated in relatively few significant STWC interactions with CTL. There are a few potential explanations for this: because there are few direct PFC-to-M1 connections (Miller & Cohen, 2001), there may be a lack of temporal consistency in interactions between these areas that would render them statistically insignificant in our model-free analytical approach. Additionally, it is very likely that the task-relevant information represented in PFC is more related to goal-direction (Kobayashi et al., 2007; Rudolf & Hare, 2014) and working memory (Barbey et al., 2013; Funahashi et al., 1989; Goldman-Rakic, 1996) than direct BCI

control, thus we would not necessarily expect to observe tight temporal correlations between PFC and M1.

The question remains of what function STWC interactions are playing to assist in BCI task execution. The fact that we observed temporal structure in STWC interactions that both followed and that preceded activity in the control electrode speaks to a potential feedback role for these regions. Signals that followed the control electrode activity could carry information about just-completed task performance whereas activity that precedes control electrode performance could carry information about motor planning that might result in improved task performance, though these mechanistic interpretations remain speculative.

Of particular interest is our observation that these high-frequency STWC interactions generally covered shorter cortical distances than bPLV interactions. This is in agreement with recent findings that inter-area correlations in the low-pass filtered HG envelope are more predictive of local than distant structural connectivity (Keller et al., 2014) and consistent with previous hypotheses regarding phase-phase interactions as an appropriate means for long-distance information transfer (Darvas, Miller, et al., 2009; Schnitzler & Gross, 2005).

Oscillatory cortical activity has been studied extensively at multiple spatial scales (Schnitzler & Gross, 2005), and, as a field, we have resounding evidence demonstrating that specific cortical oscillations respond reliably sensorimotor and cognitive events (Kahana, 2006). One prevalent hypothesis about inter-area oscillatory coupling is that it facilitates communication between two regions, however, when such coupling is expressly linear (e.g. classical PLV), it comes at the expense of independent computation occurring in those two regions at those frequencies, which may be problematic for cortical areas responsible for different functions (i.e. distant cortical sites). Biphase coupling, on the other hand, is a proposed mechanism for information transfer between regions that preserves functional independence

between the two regions. We note, that though we observed significant bPLV coupling between a number of remote regions during the BCI task, no significant linear PLV was observed.

Cortical hubs proposed by Buzsaki et al. (2004) and evidenced by Keller et al. (2014) are intrinsically tasked with selectively processing information that comes from multiple streams simultaneously, attending to relevant information and muting the rest. Low-frequency oscillatory synchrony is one proposed mechanism for this gating (Fries, 2005; van Elswijk et al., 2010; Womelsdorf et al., 2007). However, if two streams of information are being integrated that occur on intrinsically different timescales (e.g., response to a visual stimulus and internal motor imagery state) the mechanism of linear phase-phase coupling may be insufficient and cross-frequency coupling may be a viable alternative.

Another possibility is that the role of information generated in one region may be different from its role in a distant region. If narrow-band HG changes are representative of selective activation of a subnetwork within an area (Canolty & Knight, 2010) then biphasic coupling presents one potential mechanism for transferring information from one such network to another, either within or between cortical regions. This is conceptually similar to the hypothesized role of PAC: that it serves to transfer information across spatial and temporal scales, from distributed low-frequency oscillatory networks to local high-frequency ones (Canolty & Knight, 2010).

As was described above, bPLV can readily be interpreted as a measure of effective (i.e. directed) functional connectivity. This is because the phases of the two multiplicative frequencies are predictive of the phase at the sum of their frequencies, but not vice versa. Whether STWC can be interpreted as a measure of functional (i.e. undirected) or effective connectivity depends, in part, on the lag at which correlative relationships occur. Though the potential existence of a hidden third source, exerting influence over the two visible nodes means that any conclusions as to causal influence of one of the visible nodes on the other must be tempered, at the very least, significant STWC coefficients at non-zero

lags are indicative of information flow of some kind. Our finding of significant STWC lag relationships indicates that we are not merely observing non-specific co-activation of neighboring neural populations. The fact that we observed PMv typically leading CTL in significant PMv-to-CTL interactions makes sense considering traditional models of pre-motor influence on M1, but we note that this finding was based on extraction of a single lag value from each interaction. In reality we expect that there is likely bi-directional information flow between these and other cortical regions during BCI task execution.

The presence of bPLV synchronization exclusively during cue-locked trials is potentially indicative of the role that this synchrony may play in distributed processing of the task demands and subsequent execution plan. It may be that widespread cortical synchrony is necessary to develop attentional focus or to create the appropriate state associated with task execution. The precise timing of task performance does not appear to impact bPLV measures: when the data were realigned to response times, the of bPLV change with the task largely vanished. The increase in bPLV lasted only over the initial portion of the task meaning that bPLV changes had largely returned to baseline by the time the task was completed. This further supports that the cortical synchronization indicated by bPLV changes represents coordinated information flow related to the anticipation and state related to task performance rather than execution of the task itself.

Though there is evidence for cross-frequency coupling in both cortical (Canolty et al., 2006; Miller et al., 2012) and cortico-subcortical networks (Belluscio et al. 2012), we have only limited data relating these oscillatory phenomena to activity changes in underlying neuronal networks (Murthy & Fetz, 1996). Regardless of whether or not synchrony in oscillatory components of field potentials are directly involved in neural computation or are simply epiphenomena of underlying activity patterns, through a better understanding of the mechanisms that generate them we will be able to extend our interpretations of these coupling motifs when they are observed.

From the simplest paired coupled oscillator model (Mirolo & Strogatz, 1990) to much more complex models of networks of interneurons (Buzsáki & Chrobak, 1995), periodic temporal structure in firing and, correspondingly, oscillatory behavior in extracellular field potentials, can be generated via activity in recurrently-connected networks. It has been argued that the complexity (i.e. number of neurons) and spatial extent of these networks give rise to their characteristic oscillatory frequency (Buzsáki & Draguhn, 2004), which leads to the resultant hypotheses that higher-frequency oscillatory phenomena are indicative of more local computation and that cross-frequency coupling may be a mechanism for communication between networks of different spatio-spectral scales.

One could generate, then, a model of individual neurons to test this possibility. This model would include two separate oscillatory networks (see Figure 22), with varying population sizes and modeled axonal lengths to generate different (and physiologically feasible) oscillatory frequencies when driven at a basal input level. By then connecting these two sub-networks and varying both the strength of the connection between them and the strength of the respective inputs to each subnetwork, one would expect to see activity from each sub-network impact activity in the other. This would render not only in the local field potential, modeled as the sum of post-synaptic potentials between neurons in the sub-network, but also in the timing of firing in individual neurons.

In addition to probing a potential model for how cross-frequency coupling is generated in the brain, this would allow us to test whether a single mechanism could generate the bi-phase coupling observed above and previously (Darvas, Miller, et al., 2009) as well as phase entrainment of HG activity to lower-frequency oscillatory activity (Canolty et al., 2006; Miller et al., 2012) – a distinct possibility based on the intrinsic link between phase and amplitude in narrow-band oscillatory signals (Aydore et al., 2013).

Ⓐ Uncoupled Networks Ⓑ Coupled Networks

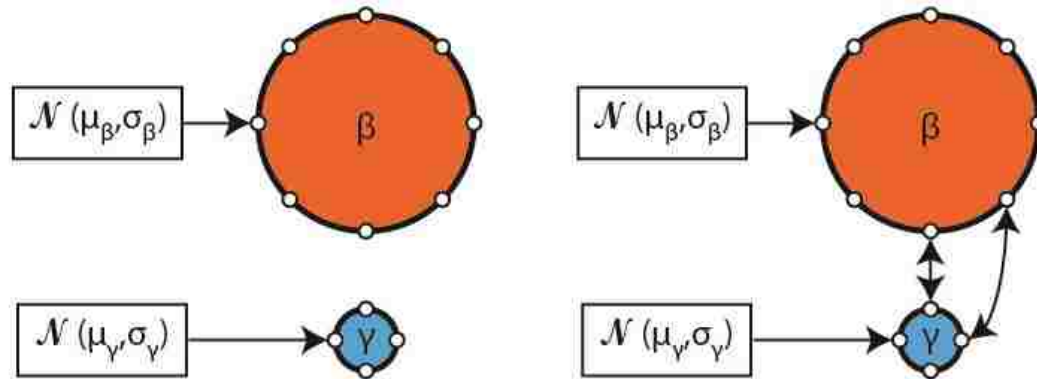


Figure 22 - Proposed linked-oscillator model to generate cross-frequency coupling. (a) Two externally-driven recurrent networks will oscillate at a frequency corresponding to the number of neurons in those networks and the modeled transmission time between them. (b) By coupling those networks, one can demonstrate the potential role of cross-frequency coupling to allow for information transfer between them. The spectral properties of activity in each separate sub-network will be that of the other, both in terms of modeled field potentials around those networks and firing patterns of individual neurons within them.

Structural interconnectivity and functional interactivity between populations of cortical neurons are at the core of human cognition. These interactions are dynamic and render in various ways in population-level cortical signals. Though in this study we have simultaneously demonstrated two different interactions in the ECoG surface potential, studies such as this will benefit greatly from an increased understanding of the anatomical mechanisms and network architectures that underlie the various forms of interactions observed in electrophysiologic recordings.

6.5 Related publications and presentations

- [1] **JD Wander**, RPN Rao, JG Ojemann. Multiple roles of ventral premotor cortex in BCI task learning and execution. Proceedings of the 6th International Brain-Computer Interface Conference, Graz, Austria, 2014.
- [2] **JD Wander**, D Sarma, K Weaver, RPN Rao, JG Ojemann, F Darvas. Non-linear cortical interactions during brain-computer interface use. Proceedings of the 44th Annual Society for Neuroscience conference, Washington, DC November, 2014.

[3] **JD Wander**, D Sarma, LA Johnson, EE Fetz, RPN Rao, JG Ojemann, F Darvas. Short- and long-distance cortical interactions during brain-computer interface task execution. *In Review*.

7 Neural representation of intention during BCI use

7.1 Introduction and background

Traditional BCI paradigms have made great strides toward dexterous control of end-effectors such as robotic arms, allowing users manipulate devices with as many as ten DOF (Wodlinger et al., 2014). Direct control paradigms are limited, however, in that they require constant cognitive vigilance on the part of the user who is volitionally modulating the multiple neural signals required to directly control such a large number of degrees of freedom. An alternative to this methodology is to simultaneously extract higher-level goal-related information from remote cortical sites to provide the BCI with additional channels of task-relevant information that can enhance its ability to carry out the movement intended by the user.

In this chapter, we address this end from two directions. **First, we determine the existence of and quantify the neural representation of the higher-level goal in a pair of one-dimensional BCI tasks. Second, we evaluate a closed-loop, machine learning-based approach to including the inference of this higher-level goal in BCI task execution.**

Neural representation of intention during motor movement. When executing natural motor movements, the brain seamlessly and rapidly transforms sensory information into the motor plan necessary to achieve its goal (Buneo & Andersen, 2006a; Fuster, 2000). Astoundingly, this multi-stage process of sensory integration, goal identification, motor plan formulation and finally motor execution can occur in less than a second. As it is central to goal-driven behavior, numerous studies have investigated the neurophysiological correlates of this transformation and the subsequent motor execution (see Fuster, 2000 for review).

Single unit activity during movements (most often reach movements and saccades) has been correlated with a wide variety of movement parameters: target location in various reference frames, joint angles, force, velocity, and position (Evarts, 1968; Fetz & Cheney, 1980; Georgopoulos et al., 1982). These activity changes have been observed not only in M1, but also in premotor areas (Mushiake et al., 1991), SMA (Mushiake et al., 1991), and PPC (Mulliken et al., 2008). In addition, both behavioral (Prablanc & Martin, 1992) and neurophysiologic studies (discussed individually below) have demonstrated the presence of a significant feedback component in online motor execution, indicative of the fact that the brain is constantly updating an error signal that is representative of both the ongoing task state and the higher-level goal (Desmurget & Grafton, 2000). Mackay provided indirect evidence for this by showing that individual cell firing rates in NHP area 7a (near the IPS) changed over the course of reach execution (1992). The role of PPC in online reach error tracking was substantiated through the demonstration that interruption of PPC processing via TMS during a two-step reaching task resulted in increased correction errors (Desmurget et al., 1999).

Motor execution is not instantaneous, thus correspondingly, there are also numerous studies that investigate representations of parameters of pending movement *before* that movement has been initiated (Batista et al., 1999; Gallivan et al., 2011, 2013; Scherberger et al., 2005). Of particular interest to the field of BCI research is the pre-movement neural representation of the higher-level goal that will be achieved by a particular movement because such a signal could be readily utilized in multi-stage BCI architectures. A number of cortical regions, primarily in fronto-parietal networks have been implicated in this process of goal formulation and storage.

A thorough series of studies by Andersen and colleagues have demonstrated the role of PPC, specifically the parietal reach region (PRR) in transformation of visual input to an internal representation of the task involved in reach planning in monkeys (Batista et al., 1999; Buneo et al., 2002; Hwang et al., 2012; Mulliken et al., 2008; Scherberger et al., 2005). These patterns of activity are so robustly correlated with

kinematic parameters that they have been successfully used to reconstruct kinematics of the reaching motion (Scherberger et al., 2005). Human homologues of these regions are not as well understood, limited in-part by restriction to primarily non-invasive recording techniques. Recent imaging evidence supports specific participation of more dorsal portions of the posterior parietal cortex (specifically areas surrounding IPS) in both reach and saccade planning (Gallivan et al., 2011) as well as voluntary shifts in spatial attention (Yantis et al., 2002). This is in contrast to evidence for bottom-up attentional modulation of activity in the more ventral TPJ, where BOLD responses are elicited through the presentation of auditory or visual stimuli that are informative of a spatial target (Geng & Mangun, 2011). Current imaging evidence supports roles for both dorsal and ventral portions of the parietal cortex in spatial target discrimination (Gallivan et al., 2011; Geng & Mangun, 2011).

Such signals are not at all exclusive to PPC, Additional work has demonstrated single unit activity in M1 (Tanji & Evarts, 1976) and premotor cortices (Weinrich & Wise, 1982) immediately preceding movement onset. For the former of those two studies, differential activity patterns were observed preceding flexion and extension of the arm. Pellegrino and Wise (1993) expanded on this by demonstrating that firing changes in PMd reflected more than just sensory processing or resultant motor command but were additionally modulated by higher-level task-related factors. More recently, imaging studies performed in humans have corroborated these findings of encoding of task-relevant information in frontal cortical areas, demonstrating effector-independent (reach vs. saccade) changes in PMd BOLD activity that was sufficient for significant predictive cross-effector classification of movement direction. Additionally, they demonstrated changes, albeit weaker, in PMv BOLD signals only before reaching actions that were sufficient to allow for classification of pending movement direction (Gallivan et al., 2011).

This finding is reasonable, as we classically consider fronto-parietal networks between and AIP to PMv to be involved in the formulation of grasping motions and hand postures and between PRR & LIP and PMd to be involved in reach and saccade processing (depending on whether the frontal eye fields are

considered as a part of or neighbor to PMd) (For review see Rizzolatti & Luppino, 2001; Wise & Boussaoud, 1997). However, more recent studies are beginning to question the historical distinction between reach and grasp areas (Davare et al., 2011), especially as we translate knowledge from primate electrophysiological studies to human imaging work (Vesia & Crawford, 2012). This distinction may be particularly relevant for the field of invasive BCI because NHP SUA studies typically involve training animals and decoders during overt reaching motions (Carmena et al., 2003; Mulliken et al., 2008) whereas human field-potential-based studies often train subjects to utilize grasp-related imagery to develop control (Blakely et al., 2009; Leuthardt et al., 2004).

Another candidate frontal region for encoding of higher-level spatial goals is dIPFC, traditionally thought to be involved in working memory processes (Barbey et al., 2013; Funahashi et al., 1989; Goldman-Rakic, 1996) and goal-driven motor behavior (Kobayashi et al., 2007; Rudorf & Hare, 2014). In the same study discussed above, Gallivan et al. (2011) demonstrated sufficient coding of intent in dIPFC to perform classification of the pending reach direction.

As a field, we have yet to see evidence on the same scale of spatial target coding in far-field potentials (ECoG/EEG) recorded in humans. A number of invasive (Gomez-Rodriguez et al., 2010; Wang et al., 2012; Williams et al., 2013) and non-invasive (Kornhuber & Deecke, 1965; Lew et al., 2012; Salvaris & Haggard, 2014) studies have utilized cortical field potentials as an indicator of pending movement, without respect to movement direction, though instances of decoding pending movement direction are much more rare (Schalk et al., 2007; Wang et al., 2012). Interestingly, both of these studies cite low-frequency voltage deviations (termed the local motor potential [LMP] in these studies) (Kornhuber & Deecke, 1965) as a primary predictor of movement direction, and neither (though this may be due to limited coverage of the area) cite PPC as a cortical region where movement direction is coded. The degree to which clinical scale ECoG can detect differential activity patterns related to goal-directed motor behavior in fronto-parietal networks remains an important open question. This is the case, not

only for practical BCI research, but also in terms of our fundamental understanding of how these cortical regions interact to enable goal-driven motor behavior.

Pre-trial and in-trial decoding of task demands in BCI. The notion of prescient inference of a higher-level goal has been a subject of considerable interest in BCI research, as it this information can be utilized in a variety of ways. Firstly, when detecting higher-level goals, the most basic of these is whether to do anything at all (movement vs. rest). Though the mechanisms of inhibition are not completely understood, it is well established that motor cortical activity does not always result in motor outputs (Kaufman et al., 2014). The responsibility then falls, at least in part, on a BCI decoder to tell the difference between motor cortical activity that is intended to drive movement and that which is not. Since long before non-invasive neural signals were being thought of as a means to manipulate external devices, researchers have been using them as potential predictors of the intent to move (Gomez-Rodriguez et al., 2010; Kornhuber & Deecke, 1965; Wang et al., 2012). One potential application of these signals to BCI systems is simply to treat them as a gating function, determining, moment to moment, whether neural activity is meant to be driving end-effector output. The ability to decode this higher-level intent to move has been a significant focus of intent estimation in field potential-based BCIs (Lew et al., 2012; Williams et al., 2013), typically relying on changes in low-frequency oscillatory activity.

On the other hand, if the signal modality and decoding architecture are sufficient to allow for inference as to specifically what the intended goal is, then there are two potential strategies for how to utilize this information. The first possibility is the direct execution of the inferred goal by an automated (i.e. robotic) system; this removes the cognitive load associated with direct control from the user, but also prevents any real-time adjustments to the execution strategy. This strategy has been employed in non-invasive recording techniques (Bell et al., 2008; Cheung et al., 2012) due to their relatively low bandwidth that makes kinematic control of an end effector via sensorimotor activity changes more difficult, though robust evoked visual responses allow a user to quickly scan a space of potential goals

(Bell et al., 2008). It has also been demonstrated using invasive recordings from PPC (Musallam et al., 2004) and PMd (Santhanam et al., 2006), where substantially higher data rates have been demonstrated (see Santhanam et al., 2006 for discussion).

Alternatively the transform implemented by the decoder and/or the structure of the task could be modified to assist the user in achieving the inferred goal, while still granting the ability for online corrective action. This approach is technically challenging and correspondingly has been explored less thoroughly. During post-hoc analyses of a traditional (kinematic cursor control only) BCI, Mulliken et al. (2008) demonstrated the ability to recreate kinematic parameters of the cursor as well as target locations in two-dimensional space using a modified Kalman filter and microelectrode recordings from PPC in an NHP. Shanechi and colleagues then extended this work, constructing a BCI system that utilized information recorded from PMd and SMA about inferred target position from the cue period to bias control during the feedback period, demonstrating that this hybrid approach outperformed either approach separately (Shanechi et al., 2013).

The purpose of this chapter is to demonstrate whether there exists a neural representation of similar intention signals in clinical-scale ECoG data *before* and/or *during* use of an abstract BCI, and to investigate whether single-trial estimation of intent can be leveraged to improve BCI performance in real time.

7.2 Materials and methods

This chapter of the dissertation is logically organized into two separate experiments: (1) analysis of the RJB dataset including quantification of pre-trial and in-trial indicators of intent and assessment of the predictive power of these activity patterns; and (2) collection and analysis of data using a novel BCI task, again investigating intent indicators, validating a potential real-time approach to use of these neural

signals, and examining neural activity patterns during closed-loop BCI execution. Both the methods and the following results section reflect this organization scheme.

Because one may expect to see neural activity changes in so many cortical areas, and in the ECoG observation model we do not always have coverage of each of these areas, in the following studies, we opted for a purely data-driven approach to identification of goal-correlated neural activity patterns. This approach considers all channels recorded, regardless of their location on cortex, as candidate sources of information coding for intention during BCI execution.

7.2.1 Retrospective RJB studies

Subjects. As this was a retrospective analysis, principled selection of subjects to include in the study was done by selecting only subjects that met the following two inclusion criteria: (a) the subject must have performed at least three runs of the RJB task, and (b) the subject must have performed above chance levels.

Table 5 - Demographic and behavioral information from the post-hoc intent analysis performed on RJB data. Abbreviations: R right, L left, F frontal, P parietal, T temporal, O occipital, M Act refers to the motor task associated with the controlling electrode.

ID	SID	Gender	Age	M Act	Coverage	N Trials	N Bad	Hit rate	Chance CI
S1	7662c2	M	38	Hand	R-F/T	51	1	78%	64%
S2	30052b	M	29	Tongue	R-F/T	198	114	80%	61%
S3	4568f4	M	27	Tongue	R-F/P/T	117	9	78%	59%
S4	3745d1	M	14	Tongue	L-F/P	41	2	72%	67%
S5	26cb98	M	22	Tongue	R-F/P/T	126	26	74%	60%
S6	fc9643	F	26	Tongue	R-F/P/T	210	46	90%	57%
S7	58411c	M	54	Hand	L-F/P/T	135	15	63%	59%
S8	Odd118	M	11	Hand	L-F/T	74	6	76%	62%
S9	7ee6bc	M	29	Hand	R-F/P/T	128	18	66%	59%
S10	38e116	M	18	Hand	R-F/P	51	9	86%	64%
S11	f83dbb	M	19	Hand	R-F/P/T	92	3	64%	61%

Recordings. See section 4.2.

BCI Task. See section 4.2.

Quantification of behavioral performance. Calculation of chance performance was identical to the approach outlined in section 6.2.

Preprocessing. Preprocessing steps were similar to what has been used previously (See section 6.2), with a few exceptions. Correspondingly, they are listed out here completely as to prevent confusion and to allow for direct recreation of the analyses. Data were first manually inspected for any channels or time periods that contained obvious non-physiologic artifact or substantial inter-ictal activity. For each subject, the data were re-referenced by subtracting the common average among all good channels. Signals were then notch-filtered to remove line noise using 4th-order Butterworth filters at 60 and 120 Hz. Time-variant spectral estimates were extracted by bandpass filtering the signals using 4th-order Butterworth filters and then taking the magnitude of the Hilbert transform and squaring it. Spectral estimates were derived for the canonical frequency bands, delta (δ ; 1-4 Hz), theta (θ ; 4-7 Hz), mu/alpha (μ/α ; 8-15 Hz) and beta (β ; 16-31 Hz), as well as for the high-gamma range (HG; 70-150 Hz) and the LMP (1-10 Hz). Finally, band-limited powers were resampled down to 100 Hz.

Identification of spatial/temporal/spectral representation of intent. In order to understand the spatial, temporal, and spectral distribution of neural responses that correlate with intended movement direction, we evaluated, on a per-channel and per-frequency basis, the time course of statistically significant differences in average neural activity between the two target types. Because this requires a very large number of statistical comparisons, we further downsampled neural activity patterns under the assumption that spectral activity patterns would not vary greatly over relatively short timescales. Using a window width of 200 msec and a step size of 100 msec, we calculated the average value of each spectral feature (for a given channel and frequency) over successive windows. This resulted in an effective final sampling rate of 10 Hz. To control for differences between trials in baseline activity of

each feature, the mean value of each channel/spectral feature from the rest phase of each epoch was subtracted from all samples within that epoch. Next, we performed individual two-sample t-tests for each time/frequency/channel combination.

To correct for the relatively large number of comparisons performed in this analysis, we utilized a spatio-temporal cluster size maximum statistic approach. Treating each frequency separately, one can conceptualize the resultant time-by-channel significance map as an image. In turn, respecting the spatial geometry of the grid (8x8 as opposed to 1x64), this significance map becomes a volume. Within each of these volumes, we calculated the sizes of all contiguous regions (clusters) of significant voxels, with the rationale of identifying regions that are larger than the largest region that would be expected by chance. To non-parametrically characterize the distribution of chance region sizes, we shuffled the label of trial types and recalculated the sizes of all significant regions within the randomized map. After recording the size of the largest significant region in this chance map, this process was repeated 100 times. The 95th percentile of this chance distribution represents a conservative estimate for the minimum number of contiguous voxels in a region from the true map necessary to deem that region significant. Because this approach leverages the spatial relationship between electrodes, these analyses were restricted only to electrodes from the primary grid.

To compare the relative strength of the five spectral features, we then calculated the fraction of electrodes within each subject's primary grid that exhibited a significant difference in a given spectral feature across the two target types as a function of time. Because the feedback period for subject 58411c was 2s as opposed to 3s, their results on this and subsequent analyses in this chapter are presented individually, but not included in group averages.

Classification analyses. These studies present an opportunity to understand the neural representation of higher-level goals during a BCI task. An important consideration for the utilization of these signals in a

BCI architecture is to understand whether these activity patterns can successfully be used to train a classifier capable of inferring the user's goal. We specifically wanted to understand two questions: (1) whether neural activity from any particular moment during task execution allowed for particularly robust classification of intent, and (2) at what time point during the trial, if any, are the aggregate neural activity features from all previous time points in the trial sufficient to perform successful inference. To address these two questions, we utilized neural activity from consecutive 100 msec windows as predictive features and performed instantaneous and aggregate forms of our classification analyses. In the former, activity features from each time-step within the trial are considered separately. In the latter, all statistically meaningful features from previous time-steps are included as potential features in the training of the classifier for a given time-step. Other than the treatment of the features, the methodological approach to these two analyses is the same. In both cases, we performed 10-fold cross validation training and testing of a support vector machine (SVM) using a radial basis function (RBF) kernel to allow for non-linear relationships between predictors (Chang & Lin, 2011) (available for download from www.csie.ntu.edu.tw/~cjlin/libsvm). The two free parameters in this model, C and λ were fixed to 1 and $1/N$, where N is the number of features used for training. Before training the SVM, feature selection was performed on training data from each fold using regularized least squares regression using the lasso algorithm (Tibshirani, 1996). To prevent over fitting to the training data, lasso was conducted using 4-fold cross validation and 10 monte-carlo repetitions per fold. The algorithm was configured to include no more than 10 non-zero weights in the final regression model.

In the event that lasso identified a best model fit with no non-zero weights, we selected the three features with the highest pearson's correlation coefficient (R^2) for SVM training.

Our classification analyses were computationally intensive. In order to maintain reasonable computational complexity and execution time, we restricted classification analyses to use only HG

activity from the six frequency bands extracted from the neural data. As can be seen below in section 7.3.1, HG activity is the most robust indicator of both cue-phase and feedback-phase intention.

Post-hoc trial performance biasing. We opted for an assistive approach to hierarchical BCI that modifies one or more task parameters on a per-trial basis in an effort to assist in task execution. In our post-hoc analysis of the RJB data, we modified the starting y-position (y_0) of the cursor on the screen. Because a portion of our classification analyses utilized HG activity from the feedback phase of the trial, after the cursor would have already begun moving, this theoretical analysis violates the causality necessary for a real-time application. However, in an effort to characterize the full-time course of intent representation, we included neural activity from both the cue and feedback phases.

In the case of the two-target RJB, the target can take one of two positions on the screen, as described previously. Assigning position values to those two targets as $y_{up} = 1$ and $y_{down} = 0$, and assuming that the target direction A has already been classified as a specific direction a , the start position of the cursor was modified according to the following formula:

$$y'_0 = (y_a - y_0) * 2\alpha(p_{A=a|X} - 0.5) + y_0 \quad (4)$$

Where y'_0 is the biased cursor start position, y_a is the y-position of the target in the direction of a from the origin, y_0 is the original cursor start position, α is a bias magnitude term ranging from 0 (no confidence in inference) to 1 (complete confidence), X is the observed neural data and $p_{A=a|X}$ is the posterior probability of classification of the intended target direction being a given the set of neural features X .

We performed post-hoc biasing of RJB task execution using the results from the classification analyses at a spectrum of α values from 0 to 1 in 0.01 unit steps.

7.2.2 GoalBCI studies

Subjects. To be eligible to participate in this study, subjects needed coverage of the precentral gyrus, specifically with one or more electrodes that demonstrated significant activity correlated with hand or tongue movement. After the initial screening task, subjects were given the opportunity to perform the GoalBCI task as many times as they wanted, but the subjects had to complete at least 3 runs to be considered for subsequent analyses.

Table 6 – GoalBCI Subject information and behavioral performance summary. Subjects performing significantly better than chance on either hit rate or ISE are shown in bold face. The asterisk (*) denotes a subject that was used as a purely visual control who was presented the visual stimuli associated with the task without being granted BCI control. Abbreviations: R right, L left, F frontal, P parietal, T temporal, O occipital, M Act refers to the motor task associated with the controlling electrode.

ID	SID	Gender	Age	M Act	Coverage	N Trials	N Bad	Chance		Chance	
								Hit rate	(95% CI)	ISE	(95% CI)
S1	d6c834	M	25	Hand	LR-F/P	129	10	18.4%	24.5%	1.31	1.36
S2	6cc87c	F	11	Hand	L-F/P	134	1	45.4%	22.8%	0.704	1.35
S3	ada1ab	M	16	Tongue	R-F/T/P	54	9	16.7%	20.5%	1.44	1.36
S4	6b68ef	F	50	Hand	R-F/T/P	78	0	14.5%	19.3%	1.40	1.35
S5	8adc5c	F	42	Tongue	R-F	49	2	22.5%	26.7%	1.69	1.56
S6	5050b0	M	16	Hand	R-P	63	0	17.9%	8.16%	1.25	1.26
S7	a9952e	F	13	Hand	L-F	106	6	45.3%	25.8%	1.14	1.39
S8	d5cd55	F	35	Hand	L-F/T/P	112	1	25.9%	25.3%	1.50	1.57
S9	9d10c8	M	13	Tongue	L-F/T/P	125	17	38.4%	26.7%	1.20	1.46
S10	979eab*	M	20	N/A	L-F/T/O	32	1	N/A	N/A	N/A	N/A

BCI Task. As an alternative to the RJB, the design of the GoalBCI task is tailored to allow us to simultaneously probe the cortex for representation higher-level intent while also providing a more behaviorally rich task that breaks the perfect correlation between direction of the target from the cursor starting position and correct direction of travel given the current cursor position. The task is a one-dimensional variant of the center-out task that has been used in numerous invasive and non-invasive BCI studies (e.g., Simeral et al., 2011). Subjects are presented with the target for 2-3s (the 'cue' period), then the cursor appears and the subject must modulate neural activity of the controlling electrode to

move the cursor into the space defined by the target (the 'feedback' period). To succeed in any given trial, the subject must move the cursor into the space defined by the target and maintain it within that region for a specified dwell time (1s). If the cursor leaves the target before this dwell duration has completed, any subsequent entry into the target area will still be required to dwell for the entire duration of the dwell time. The dwell time requirement was chosen in conjunction with the overall control gain of the system such that the subject is unlikely to acquire a target simply by 'drifting' across it; instead they can overshoot and must subsequently correct for that error. As was the case with the RJB, the cursor's vertical velocity was driven by spectral power changes in a sub-band (~75-100 Hz) of the HG range from the electrode determined during motor screening. Subjects were trained to modulate this activity feature using motor imagery; however there was no restriction on the subject continuing to explicitly use motor imagery as they developed task proficiency. Subjects were instructed not to perform overt movements to drive the cursor.

Whereas the two-dimensional center-out task typically has four or eight target locations, all equidistant from the cursor origin and placed radially every 90° or 45°, respectively, the task used in this chapter restricts movement of the cursor to the vertical dimension and places targets along the vertical line that passes through the cursor origin. Also unlike the standard center-out task, targets in this task have properties other than location that are varied on each trial. The targets were placed either above or below the origin at one of two distances from the origin (20% or 35% of the total screen height). They were also one of two diameters (8% or 16% of total screen height). Task flow and target direction/distance/sizes are illustrated in Figure 23 and Figure 24.

As was the case with the RJB, the cursor's vertical velocity was updated every 40msec and controlled by changes in HG activity at the controlling electrode as calculated by an auto-regressive filter using the previous 500msec of data. This time-variant estimate of HG activity was normalized against six seconds of stored data for each target type and then mapped to cursor velocity. The normalizer was typically

allowed to adapt (collecting reference data and updating normalization parameters) only during the first run; however in cases where non-stationarity of the signals showed obvious decoder bias, it was recalibrated.

We utilized a block-randomized trial design, to present the user with an approximately equal number of occurrences of each target type, even in runs that were aborted before completion, while controlling for potential order effects between trial types. Each block consisted of a single presentation of each of the eight potential targets. Each run consisted of three consecutive blocks. Depending on the subject's capability, each run could last between 2.25 and 6.75 minutes.

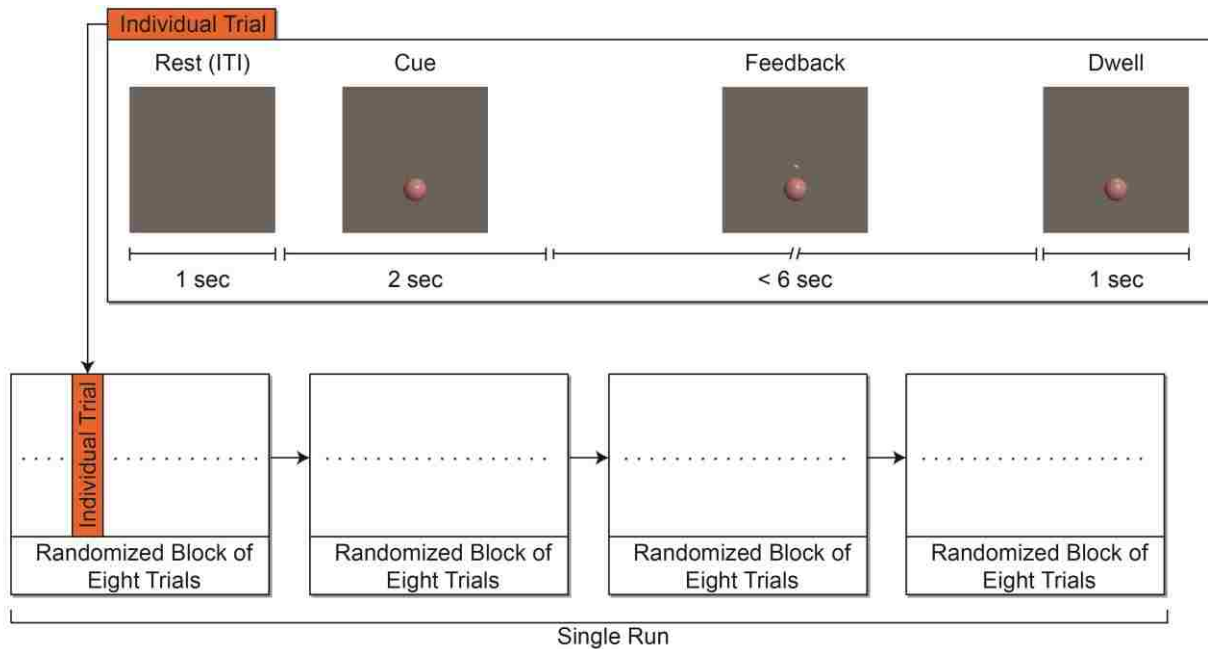


Figure 23 - Depiction of the GoalBCI task and block randomized design

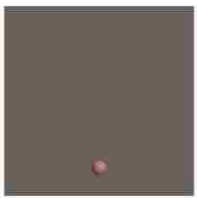


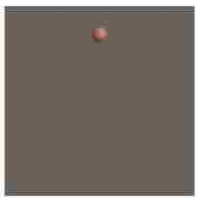




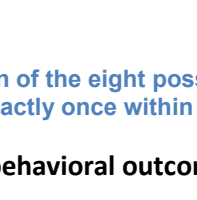
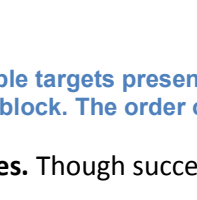
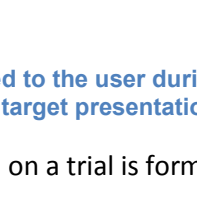
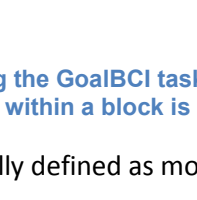
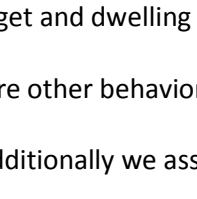
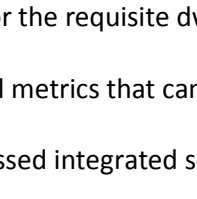
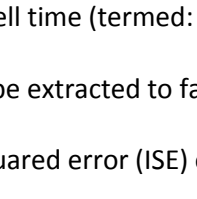
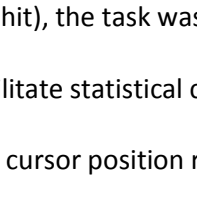
		Down		Up	
		Far	Near	Far	Near
Small					
					
Large					
					

Figure 24 - Depiction of the eight possible targets presented to the user during the GoalBCI task. Each target will be presented exactly once within a block. The order of target presentation within a block is randomized.

Quantification of behavioral outcomes. Though success on a trial is formally defined as moving the cursor over the target and dwelling for the requisite dwell time (termed: a hit), the task was designed such that there were other behavioral metrics that can be extracted to facilitate statistical comparisons of performance. Additionally we assessed integrated squared error (ISE) of cursor position relative to the nearest edge of the target on each trial.

Preprocessing for offline analysis. As the reader will see demonstrated in the results of the post-hoc analysis of the RJB task, HG activity is by far the most robust indicator of intent during BCI task execution, thus all preprocessing steps were as described above, with the exception that the sole spectral feature extracted from the raw ECoG signal was the HG power envelope (70-150 Hz). To maintain computational tractability and to allow for direct comparison of ongoing task state (cursor position, etc.), this spectral feature was downsampled to the update rate of the BCI2000 system (40msec / 25 Hz).

Identification of neural correlates of intended movement direction. Identification of HG activity patterns associated with different target directions was performed similarly to the approach described previously for the RJB data, with one notable difference. Instead of using the conservative maximum statistic approach to determine the significance threshold, we characterized the complete distribution of cluster sizes in the 100 randomized bootstrap repetitions. Given that subjects in this study did not participate in as many trials of the BCI task and did not expressly have the requirement of above chance task performance, this was done to minimize the risk of type II error.

Classification of intended movement direction. Post hoc classification analyses were conducted on all nine subjects. The approach was conceptually similar to the RJB classification analyses performed earlier in this chapter, though there were a number of noteworthy differences, so the method is explained in full below.

The neural activity features used in GoalBCI classification analyses were the average HG power in each recorded electrode during the cue phase. We performed 10-fold cross-validated training and testing again using an SVM with an RBF kernel. In an effort to maximize classifier performance on relatively little training data, in this analysis we added the additional step of performing a parameter sweep for C , the SVM penalty parameter and γ , the single free parameter in the RBF kernel.

This was done using a standard train-test-validate approach, with a 2-stage iterative grid search for optimum values of C and γ . In the outer stage of this process candidate values of C were $2^x, x \in [-5, -3, -1 \dots 15]$ and candidate values of γ were $2^x, x \in [-15, -13, -11 \dots 3]$. In the second stage, given maximal exponents from stage one as C_+ and γ_+ candidate values for both C and γ were calculated as $2^x, x \in [y - 2, y - 1.5, y - 1 \dots y + 2]$ where y is C_+ and γ_+ as is appropriate.

During the training phase of each fold, the training data were broken into nine additional folds (each containing $8/9^{\text{th}}$ of the training data) that were used to train multiple SVMs across a spectrum of

parameter values. Each of these SVMs was then used to classify on the remaining $1/9^{\text{th}}$ of data from that fold. From these classification results, the values for C and γ that produced the best classification accuracy over all sub-folds were then used to train a single SVM on the entire set of training data. This single SVM was used to evaluate classification accuracy for that fold. The reader is referred to LIBSVM – A practical guide to SVM classification for additional methodological detail (unpublished, available freely for download at <http://www.csie.ntu.edu.tw/~cjlin/papers/guide/guide.pdf>). Because of the real-time requirement of low computational complexity, feature selection was performed on a per-fold basis using the computationally efficient mRMR feature selection algorithm (Peng et al., 2005), modified to be appropriate for continuous data (utilizing correlation as opposed to mutual information as the measure of relevance). The mRMR algorithm maximizes the relevance of all selected features while minimizing the redundancy of any two features in the retained feature set, and has a single free parameter of the number of features to be selected. This parameter was set to 10.

Implementation of a real-time BCI using inferred intended movement direction. In addition to the post-hoc analyses, we developed a version of this pathway that could run in real-time in the BCI2000 signal processing pathway. The feature and parameter selection portions of this pathway were conducted using custom Matlab software and were identical to what is described above. The final inference was not performed using a cross-validation approach, however, but instead was deployed as a real-time system. At the patient's bedside, we collected data from approximately two full runs of the GoalBCI task and took a short break (< 2 minutes) to train the classifier that would be used for real-time biasing in subsequent runs.

The classifier was configured to report the posterior probability of classification of intended movement direction, such that the starting position of the cursor can be dynamically changed on a trial-to-trial basis. During this closed-loop evaluation, one-half of the performed trials were catch trials where the

output of the classifier will not be used to bias the trial. This is done to allow for direct comparison of the added value of inclusion of goal inference in our BCI architecture.

We biased the starting position of the cursor by moving it toward the mean position of all targets in a given direction, proportionally to the posterior probability of the estimation. Assuming that the target direction A , was classified as a specific direction a , the biasing of start position was be done according to the following equation:

$$y'_0 = (y_{\mu|a} - y_0) * 2(p_{A=a|X} - 0.5) + y_0 \quad (5)$$

Where y'_0 represents the biased cursor start position, $y_{\mu|a}$ is the mean y-position of all targets in the direction of a from the origin, y_0 is the original cursor start position, X is the observed neural data and $p_{A=a|X}$ is the posterior probability of classification of the intended target direction being a given the neural data X . Note that this is similar to equation 4 above, however, now targets can have one of multiple vertical positions and the confidence term α has been set to one.

Multiple regression model of ongoing representation of task state. All previous discussion of the neural representation of intention has treated intent as a static value (e.g. coding for target position) that could be inferred before or during task execution. In addition to this representation of intent, we wanted to consider the possibility that neural activity patterns during *ongoing* task execution could reveal information about both the higher-level goal being attempted by the subject and the closed-loop computation being carried out by the brain during the feedback period to achieve that goal.

To that end, we employed a channel-wise lagged multiple regression analysis relating HG activity patterns to ongoing task state during the feedback period. Because the subject had direct control over cursor velocity and was expressly trying to minimize the distance between the cursor and the target, we chose these two quantities and their interaction as primary predictors of neural activity. Because we were simultaneously interested in higher-level goal signals, we included the position of the target in this

regression model, recognizing that it is not time variant within a given trial. The simple formulation of the regression model is then as follows:

$$x(t) \rightarrow v(t) + e(t) + ve(t) + p_{target} + \epsilon \quad (6)$$

where $x(t)$, the HG activity from a single ECoG channel, can be approximately modeled as a linear combination of cursor velocity [$v(t)$], the distance between the cursor and the target [$e(t)$], the interaction between the two [$ve(t)$], and the position of the target [p_{target}]. Respecting the fact that both the brain and the BCI decoder introduce delays between task state parameters and their corresponding neural representations, this model can then be expanded as follows:

$$x(t) \rightarrow v(t + \tau_v) + e(t + \tau_e) + ve(t + \tau_{ve}) + p_{target} + \epsilon \quad (7)$$

where each time-variant predictor now has a single lag that best captures the relationship between that predictor and the neural activity.

After the preprocessing steps described above, the first step in this analysis is then to identify and remove all trials where the cursor became ‘stuck’ on the top or bottom of the workspace. In these cases, the nature of the relationship between task state and neural activity changes drastically (especially at the controlling electrode, where ECoG-to-velocity mapping is strictly defined). Then, on a channel-by-channel basis, we selected the appropriate lags for each of the time-variant predictors. This was done using a cross-correlation approach, with a maximum lag between neural activity and a given predictor of $lag \in [-1, 1]$ seconds. Because task state (e.g. cursor velocity) is undefined outside of the feedback period, these regression analyses only included neural activity from time points where behavioral state was defined at all lags (from 1s after feedback begins to 1s before it ends). Lagged predictor matrices were generated by shifting predictor vectors by up to 1s in either direction (30 samples) and concatenating them, then dropping all time periods where the predictor was not defined. We then calculated correlation coefficients and corresponding p-values for each predictor-to-ECoG lag, and

Bonferonni corrected the p-values for the number of lags considered. The significant (post-correction) lag that resulted in the highest correlation coefficient was retained as the lag that best captured the relationship for that particular ECoG channel/predictor pair. In the case there were no significant correlations for that pair, that predictor was not included in subsequent regression analyses for that specific ECoG channel.

With lags for each channel/predictor pair in hand, we then performed channel-wise multiple regression analyses, relating the multiple factors of task state to the HG activity patterns recorded in each electrode. From this we extracted and subsequently Bonferonni corrected (based on the number of channels for a given subject) p-values indicating whether a given predictor (at its corresponding lag) was significantly predictive of HG neural activity for a given ECoG channel.

7.3 Results

7.3.1 Retrospective RJB studies

Subject population and behavioral results. This study was performed on 11 individuals (1 female; Mean age 26.1 ± 12.0 std). The mean number of trials performed by each subject was 111 ± 57 (std). Average task performance, measured in terms of successful target acquisition was $75 \pm 8.6\%$ (std). Under the assumption of an equal probability of success or failure on any given trial, the 95% CI of chance performance varied across subjects, depending on the total number of trials performed by that subject. Correspondingly, individual trial counts, performance, and CI's are included in Table 5. These results are summarized in Figure 25.




Figure 25 - Individual task performance in the retrospective analysis of the RJB data for the purposes of goal inference. Chance performance is at 50% for all subjects, and the 95% CI of chance performance (dotted line) varies from subject to subject.

Spectro-temporal features correlated with intent. The primary goal of this analysis was to determine which spectral components from the ECoG signal are most predictive of the intended target of the subject and when during task execution they occur.

Far and away, we found the strongest representation of the intended target in HG signals during the ongoing task execution; in 10 of 11 subjects we observed significant differences in HG activity across trial types (two-sample t-test, $p < 0.05$; cluster-size corrected), with mean t-statistic magnitudes (considered across all electrodes) of over 0.6 during the feedback period. Spectro-temporal significance maps are shown for each subject in Figure 27. HG was not the only frequency band in which information predictive of intent was coded. During the feedback period, we observed significant differences in beta, alpha, theta, and delta (in 9, 2, 3, and 3 of the 11 subjects, respectively) activity as a function of target direction. We also evaluated the LMP for predictive power but observed significant differences across trial types in only 1/11 subjects. Excluding the LMP, the relative strength of significant activity differences in the different frequency bands as a function of time within the trial is summarized in Figure 26. As can be seen individually in these two figures, differences in HG activity most commonly occurred during the feedback period after approximately 500 msec, indicative of the fact that intent is most

strongly represented after a behavioral response latency and during ongoing task execution. Note that subject 58411c was excluded from Figure 26 because their feedback duration was only 2s.

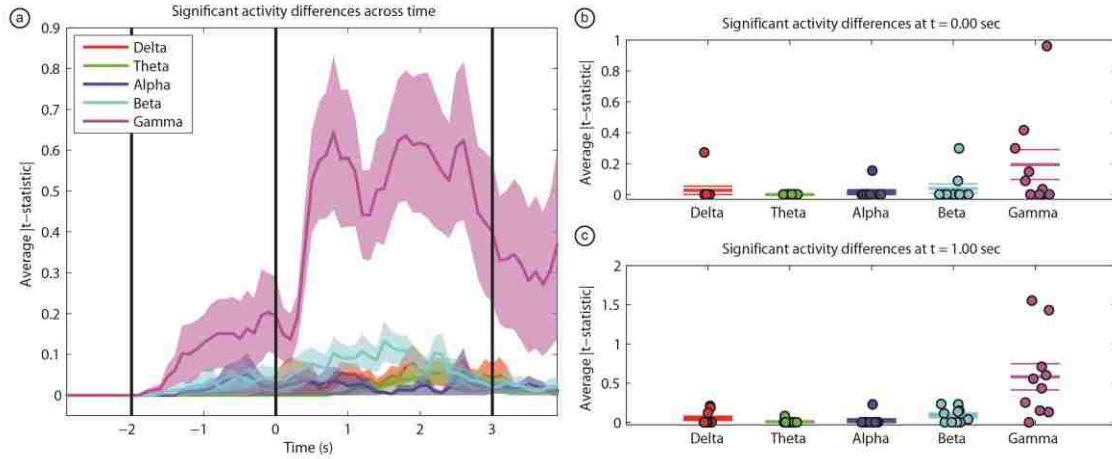


Figure 26 - Average Strength of intent representation in the RJB data set as a function of time, stratified across frequency. (a) Note markedly higher average absolute t-statistics for HG than the lower frequencies. (b-c) Individual time slices taken from (a). These serve to demonstrate the distributions and individual contributions of data points for each curve from the subjects. Thick horizontal bar represents the mean value of all data points at that given time for the specific frequency. Thin horizontal bars represent the SEM.

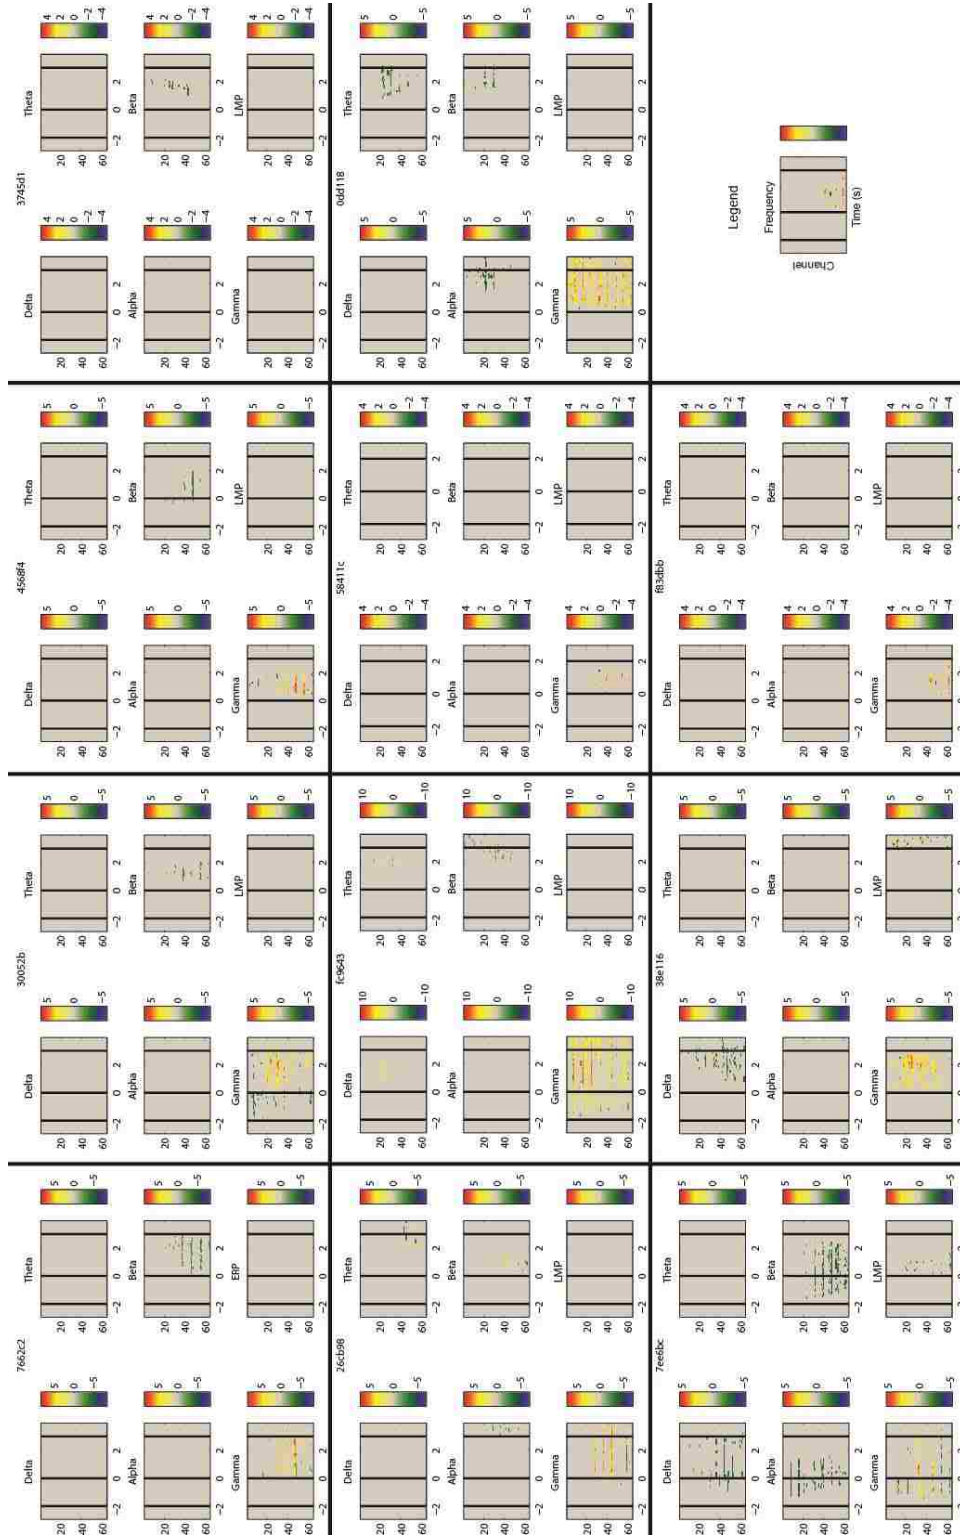


Figure 27 - Individual spectrotemporal significance maps for the RJB data set. Maps are shown for the four canonical oscillatory frequencies (δ , θ , α & β) as well as HG and the LMP. The legend provides axis and colormap labels.

Cue phase representation of intent. Though Figure 26 demonstrates that the majority of subjects only demonstrated intent representation during the feedback period, we note three subjects who had significant intent representation during the cue period as well: 30052b, fc9643, and 7ee6bc. This can be seen individually in each of their spectro-temporal feature maps, but is also readily apparent in Figure 26b, which shows the individual contributions of each subject to the various curves at $t = 0$ and $t = 1$. These three subjects appear to account for most of the significant, goal-related HG activity during the cue phase.

Looking specifically at the spatial representation of electrodes with strong target representation during the final second of the cue phase ($-1.5 < t < 0$; $\text{abs. mean}(t\text{-statistic}) > 1.96$), we observe an interesting spatial distribution of electrodes. This is depicted in Figure 28. Specifically, in subject 30052b, we see representation in BAs 22 and 6, corresponding to the TPJ and PMd respectively. For fc9643, we see representation in BAs 4, 6, 10, and 43, corresponding to M1, PMv/d, PFC, and TPJ, respectively. Lastly, for subject 7ee6bc, we observe representation in BAs 4, 5, 6, and 40, corresponding to M1, PPC, PMd, and lateral parietal cortex. It appears, at least in this small sample size, that fronto-parietal networks play a role in pre-trial representation of intent in BCI task execution. Why we do not observe similar activity patterns in the other eight subjects is curious, and may be related to strategic or internal representation differences across users, neither of which were quantified in this study.




Figure 28 - Spatial distribution of electrodes showing significant representation of intent during the cue period. *t*-statistics were averaged over the majority of the cue period ($-1 < t < 0$), and thresholded to exclude all average values between -1.96 and 1.96. Note representation of intent in M1, PMv, PFC, TPJ, and PPC. Electrodes not exhibiting significant changes are shown as black dots to provide information about subject-specific coverage.

Inference of intent. Using the approach described above, we assessed the instantaneous and aggregate classification accuracy that could be achieved using HG activity features. In this case, aggregate classification gives a sense of the relative time/accuracy tradeoff for inference applications.

The trend of relatively stronger representation of intent during the feedback period carried over to classification analyses, with mean instantaneous classification accuracies of $55.4 \pm 2.0\%$ (mean \pm SEM typical throughout this paragraph) and $65.7 \pm 3.2\%$ during the cue ($-1 < t < 0$) and feedback ($0.5 < t < 3$) periods, respectively. Using the aggregate approach, predictive power roughly equivalent during both task phases with mean classification accuracies of $57.3 \pm 3.3\%$ and $71.4 \pm 2.8\%$, again for cue and feedback periods, respectively. As an experimental control, we utilized the same neural activity and classification approach attempting to classify shuffled trial labels. In the control case we observed mean instantaneous classification accuracies of $50.8 \pm 2.1\%$ and $50.0 \pm 2.1\%$ (cue and feedback) and mean aggregate classification accuracies of $47.5 \pm 2.7\%$ and $45.0 \pm 3.6\%$ (cue and feedback). For the instantaneous classification approach, classification during the feedback phase (paired *t*-test, typical throughout this paragraph, $p = 8.82 \times 10^{-4}$) significantly outperformed the randomized control, but classification during the cue phase did not ($p = 0.12$). This same trend was observed in the aggregate case, where we observed significantly better performance relative to control

during the feedback phase ($p = 4.66 * 10^{-4}$), but statistically equivalent performance during the cue phase ($p = 0.061$).

These results are summarized in Figure 29, with individual classification results included for each subject separately in Figure 30. Again we note that subject 58411c was not included in the group averages of classification accuracy because of a feedback period duration of only 2s.

Looking at Figure 29, we note that average instantaneous classification accuracy increases drastically over the first 500 msec of task execution, consistent with our observation of an increased presence of HG features that code for differences in intended target over this same time period.

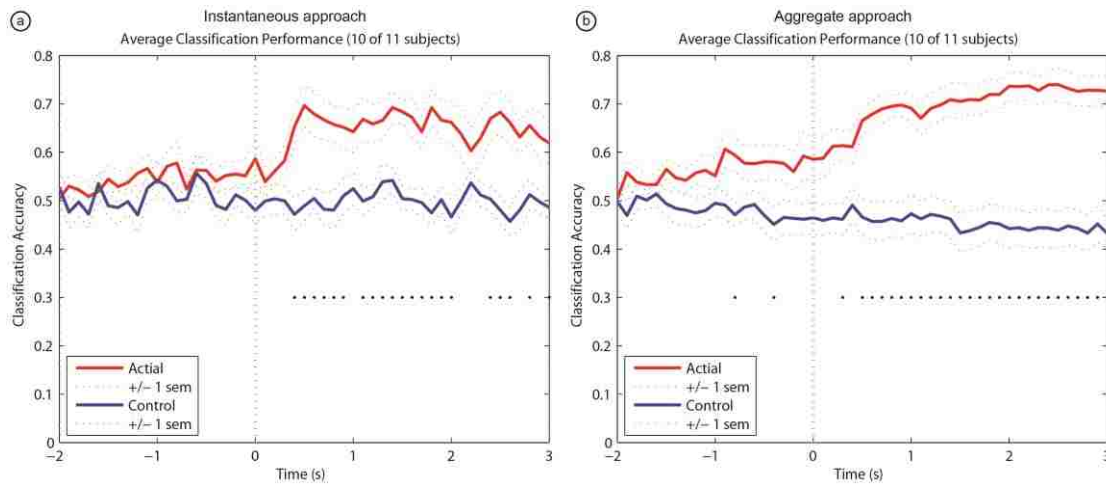


Figure 29 - Average classification performance across all subjects for the instantaneous and aggregate goal inference approaches. Actual classifier performance is shown in red (dashed line represents ± 1 SEM) and a control using randomized target labels is shown in blue. Any time points where classification performance using true trial labels was significantly different from the control condition (paired t-test, FDR corrected) are notated with black dots.

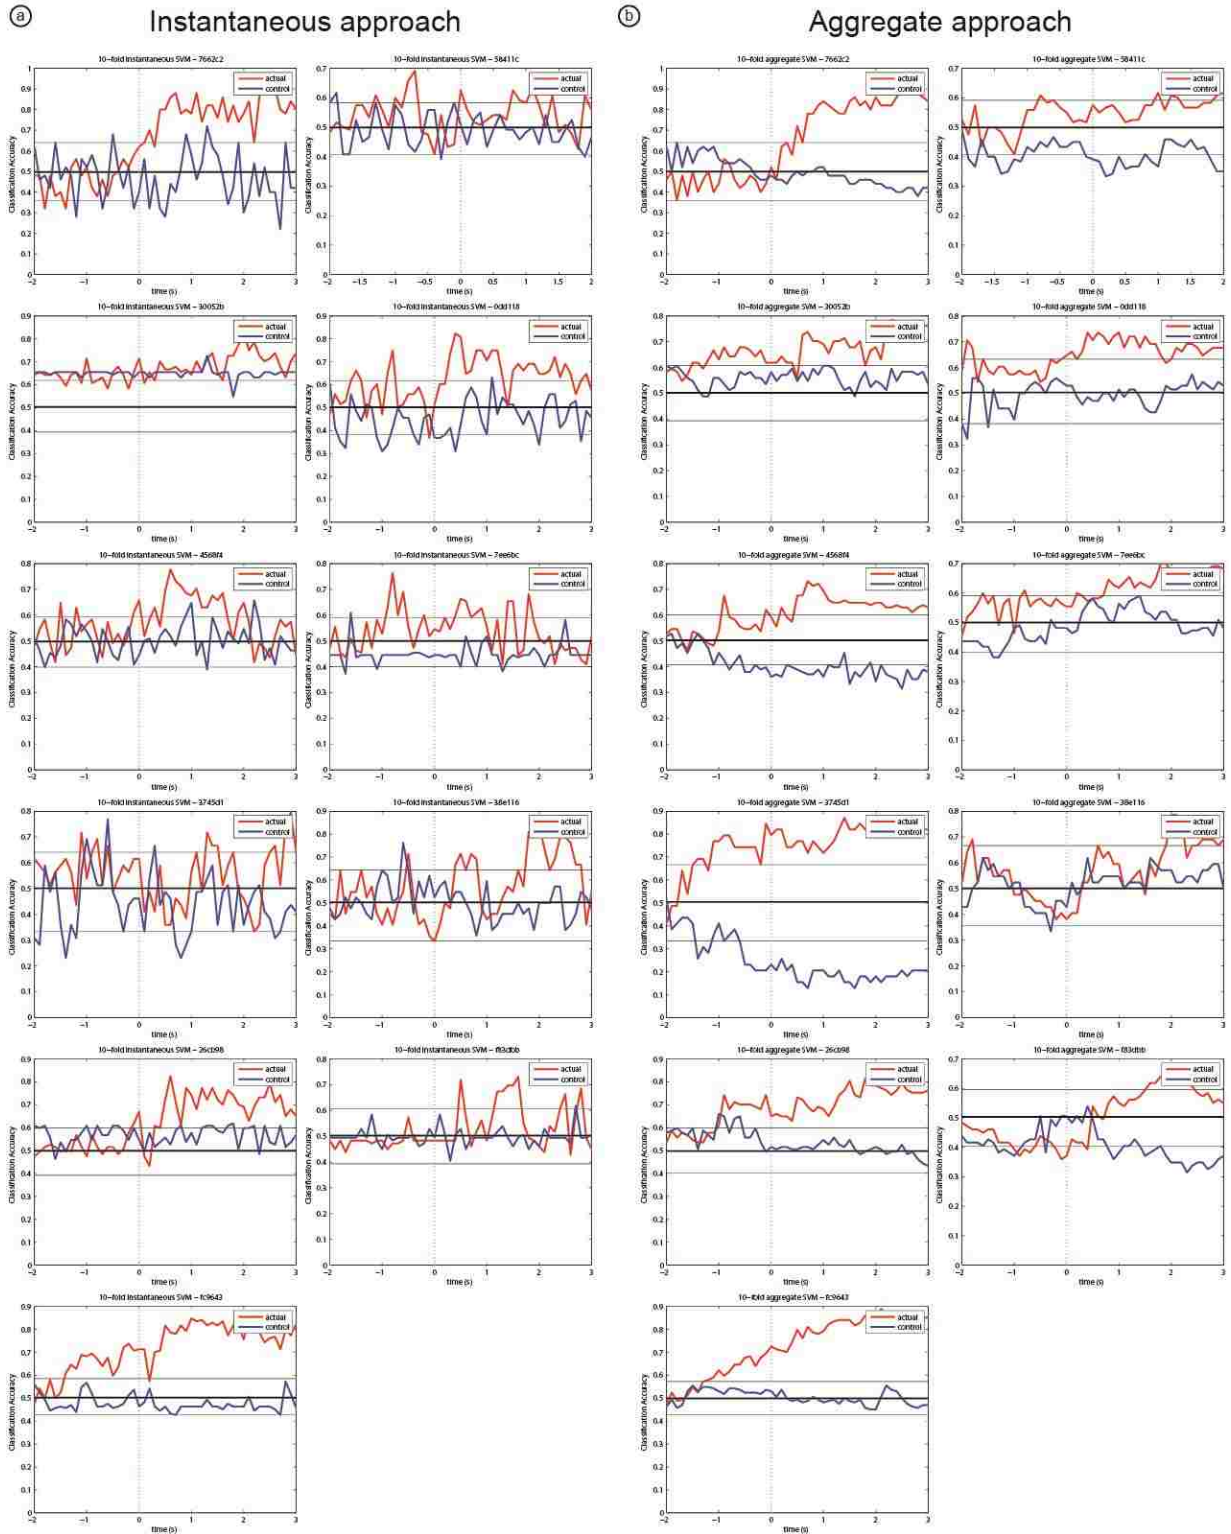


Figure 30 - Individual classification performance for both the instantaneous (a) and aggregate (b) approaches. Classification based on true trial labels shown in red and for the randomized control condition shown in blue.

Task biasing. In a post-hoc fashion, we investigated whether utilization of goal inference during the task execution would have modified overall task performance. Because these experimental modifications were theoretical as opposed to being executed in real-time, we had the luxury of performing multiple instantiations of the modification, testing for the relative effect of inference weight (the α value in equation 4).

For instantaneous classification of intent, we observed weak behavioral performance increases when inference was done on neural data from the cue phase (up to 1.7% average performance increase across all subjects occurred at an alpha value of 0.21). This was specifically true for low α values, below approximately 0.4. During the feedback period, performance increases were higher (up to 6.0%, alpha of 0.57), and this enhancement was much less sensitive to changes in α . See Figure 31 for more detail.

When comparing this to aggregate intent classification, the general trend is the same: better classification of intent when the learner has access to data from the feedback period as opposed to data only from the cue period. We observed peak average performance increases of only 0.9% ($\alpha = 0.13$) during the cue phase, and performance increases of 5.7% during the feedback phase ($\alpha = 0.46$).




Figure 31 - Summary of post-hoc trial biasing analyses. Subplots (a - instantaneous) and (d - aggregate) represent the average performance gains across all subjects (10 of 11) utilizing inference performed at each time point (x-axis) at various values of alpha (x-axis). Subplots (b) and (e) represent the same for the randomized control condition. Subplots (c) and (f) present the peak average change in behavioral performance as a function of alpha across the cue ($-1.5 < t < 0$) and feedback periods ($0.5 < t < 3$), in red and blue respectively. Corresponding control conditions are shown in dashed lines.

7.3.2 GoalBCI Task

In an effort to test the viability of the pre-trial goal inference approach in a real-time setting, and to overcome some of the behavioral shortcomings of the RJB task for dissociating low-level imagery-related neural activity from high-level goal-related neural activity, we designed the GoalBCI task, as described above.

Subject population and behavioral results. This task was performed by 9 subjects, all of whom were included in the analyses. Because there were no inclusion criteria, other than performance of the task, some subjects performed at chance levels and did not demonstrate adept control of the BCI.

For the 9 subjects (5 female, mean age 25 ± 14 std), the mean number of trials (after eliminating corrupted data) performed by each subject was 89 ± 32 (std). Because this task was substantially more difficult than the RJB, task performance was measured using two metrics. The first and most straightforward performance metric (hit rate) was the fraction of trials in which subjects successfully

moved the cursor to the target and held it within the target for 1s, meeting the successful end condition for the trial. The average hit rate across the 9 subjects was $27.2 \pm 12.5\%$ (std). We also evaluated subject performance using ISE, calculated as the instantaneous distance between the cursor center and nearest target edge. ISE is normalized by the number of pixels across the vertical extent of the workspace such that the final units are screen-seconds. Average ISE across the 9 subjects was 1.3 ± 0.29 screen-seconds.

Similar to the RJB, individual chance performance levels and 95% CI's on those performance levels varied from subject to subject based on based on the number of trials performed and differences in cursor kinematics due to variability in the BCI decoder. In contrast to the RJB, however, neither hit rate nor ISE could be simulated with random draws from a known distribution. To characterize chance task performance under the null hypothesis that subjects did not have volitional cursor control, we replayed all recorded cursor trajectories with a randomized target location and recalculated hit rate and ISE. This process was repeated 1000 times per subject to characterize the distributions for hit rate and ISE under the null hypothesis.

Demographic information, behavioral performance, and chance performance levels are given in detail in Table 6. Behavioral performance is also summarized in Figure 32.

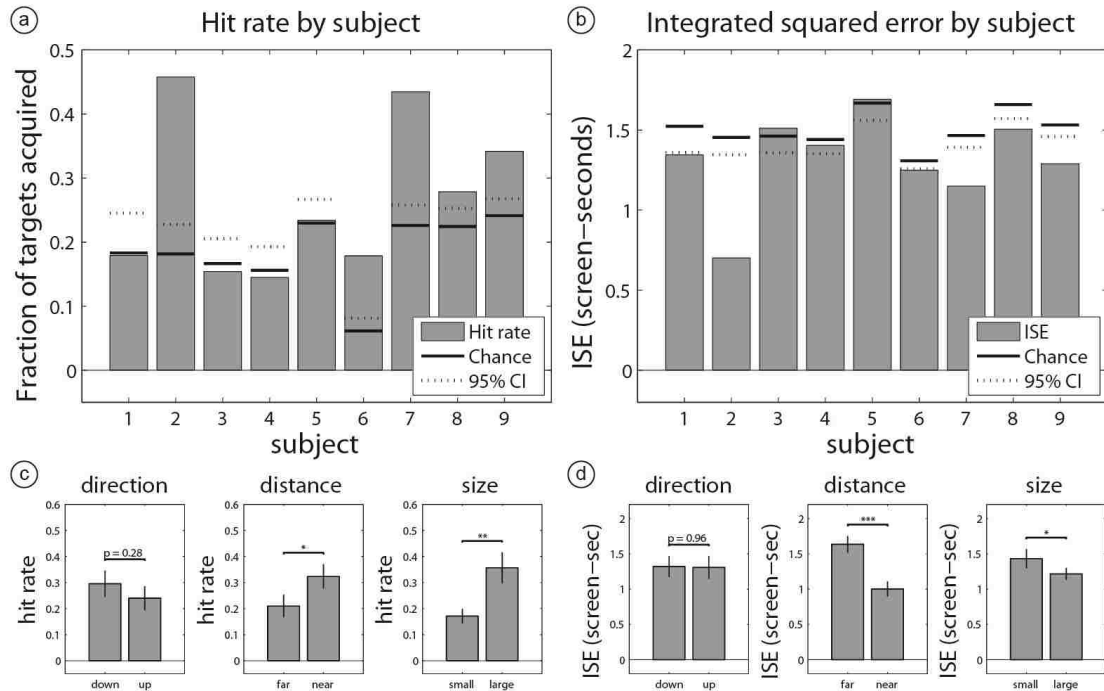


Figure 32 - GoalBCI behavioral performance as measured by (a) hitrate and (b) ISE. Chance performance was calculated on a per-subject basis using synthetically recreated trajectories and random target locations. Mean chance performance shown as a solid horizontal bar, with the 95% CI (right-sided [a] / left-sided [b]) shown as a dashed horizontal bar. (c and d) Hit rate and ISE across different target types. One star (*) signifies $0.01 < p < 0.05$, two stars () signifies $0.001 < p < 0.01$, three stars (***) signifies $p < 0.001$.**

In this task, we also took the opportunity to evaluate the relative influence of changes in target size and target distance from the origin on behavioral performance. Targets each possessed three characteristics that could be varied independently: direction, distance, and size. We found that direction (up vs. down) did not have a significant impact on hit rate (paired t-test, $N=9$, $p = 0.28$) nor ISE ($p = 0.96$). Targets that were farther away from the origin were significantly harder to hit, than those that were near, resulting in a significant impact on both hit rate ($p = 0.01$) and ISE ($p < 0.001$). Similarly, smaller targets were harder to hit than large ones, impacting both hit rate ($p = 0.002$) and ISE ($p = 0.04$). See Figure 32c and d for detail.

Comparing the impact of these target characteristics on hit rate simultaneously using a two-way ANOVA with target size and distance from origin as factors, we see that through they both have significant effects, target size appears to have a stronger impact on this measure of behavioral performance

($F(1,32) = 14.68, p = 0.0006$) than distance from the origin ($F(1,32) = 5.48, p = 0.026$). This suggests that it is easier for subjects to generate large cursor movements as compared to precise ones.

HG representation of target direction. As was the case for the neural activity from the RJB task, we sought to understand the time-course of neural activity patterns that were representative of each subject's higher-level goal. Specifically, it was of interest to know whether this goal was represented in HG activity during the cue phase and/or the feedback phase of task execution. We observed significant HG activity during each of these phases in seven of nine subjects, though we note that in general representation during the feedback phase was more consistent and stronger. In contrast to the similar results from the RJB task, we note that the strength of intent representation is not as strong for the GoalBCI subjects as compared to the RJB subjects. This is partially explainable by the fact that the GoalBCI task can require subjects to generate both positive and negative cursor velocities to reach either an up or down target, but is also possibly due to the fact that in the RJB studies we selected specifically for subjects that could perform the task at above chance levels, whereas in this study we did not. HG activity difference maps for the nine individual subjects are shown in Figure 33.

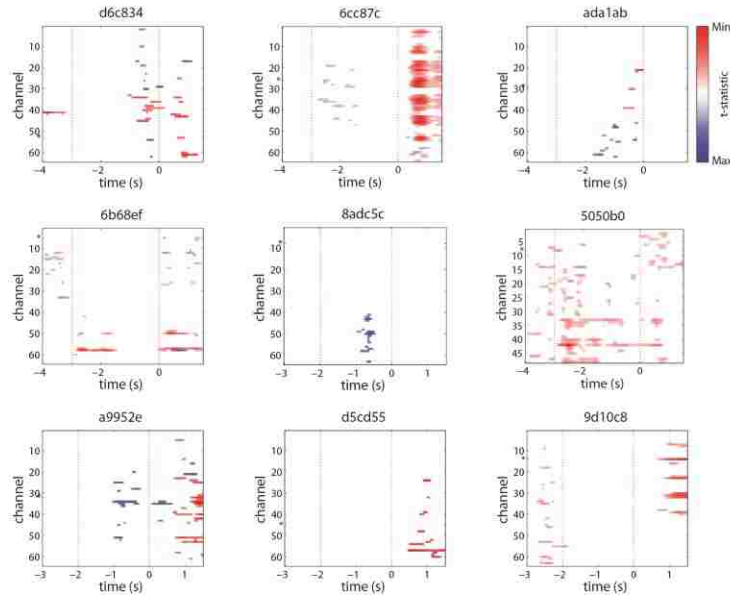


Figure 33 – Individual spectrotemporal significance maps for the GoalBCI data set. Maps were calculated for HG only. Red represents the maximum t-statistic on a given map and blue represents the minimum. See materials and methods for statistical thresholding approach.

To understand at which time periods during task execution there was a robust HG indicator of intended movement direction across subjects, we calculated the average absolute value of the t-statistic shown in Figure 33. Looking at the population average for this measure in Figure 34, we see that once again, the strongest representation of intended movement direction occurs during the initial portion of the feedback phase, and can, at list partially simply be attributed to the performance of motor imagery. Two important exceptions to this are subjects 6b68ef and 5050b0, both of whom exhibited robust differences in HG activity across trial conditions; these activity differences occurred in the posterior temporal lobe, the medial and lateral parietal lobe, and the medial occipital lobe. These are areas that are rarely covered by ECoG electrodes in epileptic focus resection procedures so it may be that similar activity patterns occurred in other subjects, but by nature of the observation model, we did not have an opportunity to record them.




Figure 34 - time course of the strength of direction representation across all nine GoalBCI subjects during the late cue and early feedback phases. Note minimal cue-phase representation of intent is visible in the grand average. Again, the strongest representation of intent is visible after ~500 msec in to the feedback period.

Inference of intent based on neural activity during the cue phase. Using the real-time system described above, we performed inference and subsequent trial biasing for a portion of the BCI trials performed by three of the BCI subjects (a9952e, d5cd55, and 9d10c8). We note that all three of these subjects naturally performed the BCI task at above chance levels, both in terms of hit rate and ISE). Classification accuracy on biased trials for these three subjects was at chance levels (58.4%, 43.75%, and 44.8% respectively). In general the approach previously applied to more invasive neural recordings (Mulliken et al., 2008; Shanechi et al., 2013) did not generalize to macro-scale ECoG. Hit rate was not significantly different in biased trials from catch trials (χ^2 tests, $p = 1, 1,$ and 0.33 for S7, S8, and S9, respectively) nor was ISE (two-sample t-tests, $p=0.9, 0.35, 0.42$). Correspondingly, behavioral outcomes were not significantly improved in biased trials relative to catch trials; see Figure 35 for detail. We do note, however that on the occasions when classification was correct, the biasing of the start position had a noteworthy impact on task performance for two of the three subjects (hit rate: $p = 1, 0.06, 1$; ISE: $p = 0.0071, 0.16, 0.08$). This is shown in Figure 36 and indicates that the assistive approach to biasing may be viable in cases where inference of the goal is possible.

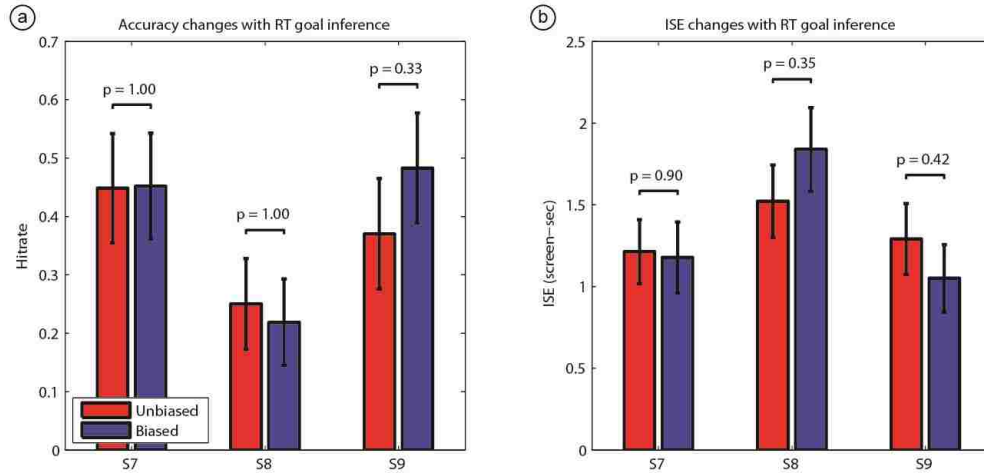


Figure 35 – Impact of real-time goal inference on Hit rate and ISE. Results for individual subjects are shown separately. Unbiased trials are shown in red, biased trials are shown in blue. Classification of intended target direction was at chance levels, thus overall performance trends were not significantly impacted.




Figure 36 – Impact of classification correctness on behavioral outcomes in biased trials. Performance metrics in instances of incorrect intent classification are shown in red, and those of correct intent classification are shown in blue. Marked performance increases in two of three subjects suggest that cursor start position biasing would be a viable approach in cases where cue-phase inference was possible. Two asterisks (*) denote p values of < 0.01 but > 0.001.**

We applied the same signal processing and inference pathway post-hoc, in a 10-fold cross-validation approach to see if, like the previous RJB results, pre-trial goal inference would be successful for a subset (N=3) of subjects. *A priori* we were specifically interested in subjects S4 (6b68ef) and S6 (5050b0) because of their strong pre-trial target-specific response. As one might expect from the relative strengths of cue phase direction representation, post-hoc classification for these two subjects was much

better (82.6% and 92.9%, respectively) as compared to the other seven subjects ($54.1 \pm 2.8\%$ [SEM]) and a target-randomized experimental control ($51.1 \pm 2.1\%$ [SEM]). We note that S5 (8adc5c) also exhibited classification performance significantly above chance though they did not possess as strong of a cue-phase intent representation.

Figure 37 summarizes the individual post-hoc classification accuracies relative to chance classification performance (calculated from the binomial distribution), and also presents the spatial locations of the electrodes from subjects 6b68ef and 5050b0 that contributed features that were significantly correlated with classification labels after FDR correction. Again, we see the localization of these electrodes to the posterior temporal, lateral parietal, medial parietal and medial occipital lobes.

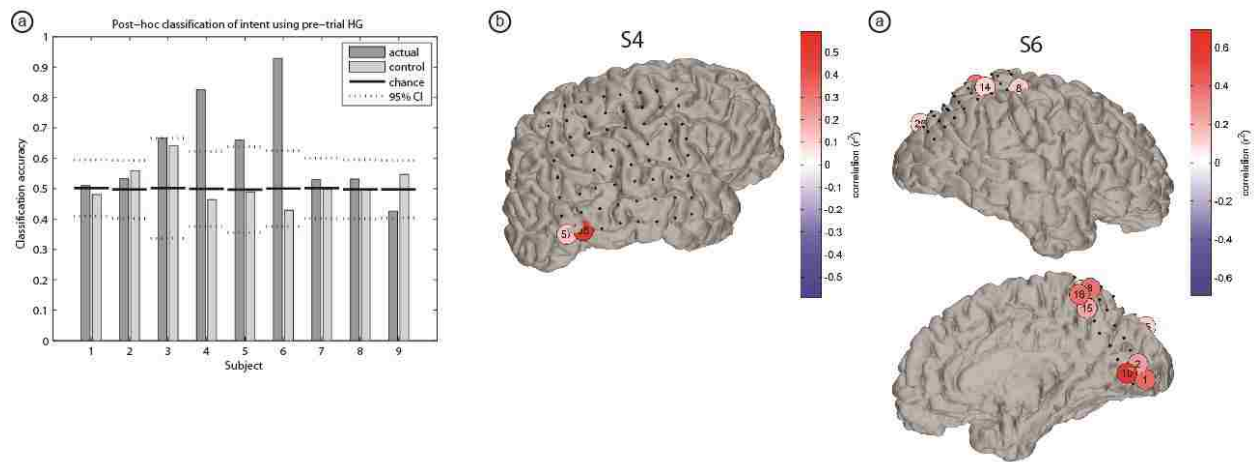


Figure 37 – Classification results in 10-fold cross-validated post-hoc classification analyses. Chance classification performance for each subject is shown as a black horizontal line, with two-sided 95% CI's shown as dashed lines. Classification based on true trial ordering is shown as a dark bar, and a control based on randomized trial ordering is shown as a light bar. (b and c) For two subjects with above chance classification, spatial distribution of contributing electrodes is shown. All electrodes with significant correlations ($p < 0.05$, FDR corrected) with target direction are shown on individual subject brains. Uncorrelated electrodes are shown as solid black dots. Though S5 had above chance classification accuracy, they had no electrodes with significant correlations after FDR correction.

Visual control for cue phase representation. As a control to determine whether (a) significant visual representation of the target's location in the workspace exists outside of primary and secondary visual cortices in the absence of task salience and (b) whether any representation under these circumstances were sufficient to classify the intended target direction, we presented all of the visual stimuli associated

with the BCI task to a novice user. During the false feedback period, the cursor did not move from the starting position, and the subject was not instructed to perform any behavioral task other than to observe the visual stimuli. After multiple comparison correction (FDR), we found no electrodes that exhibited significant correlations between target direction and mean HG activity during the cue phase (two-sample t-test). There were three electrodes in occipital (BA 19) and lateral parietal (BA 39) with correlations significant at an uncorrected significance threshold of $p < 0.01$, though this representation was not sufficient to allow for classification of target direction above chance levels (classification accuracy of 45.2%). This suggests that though representation of target location can be driven in visual processing areas in the absence of task salience, it is not sufficient to explain task-related activity changes seen in PPC (where the control subject also had coverage). See Figure 38 for detail.

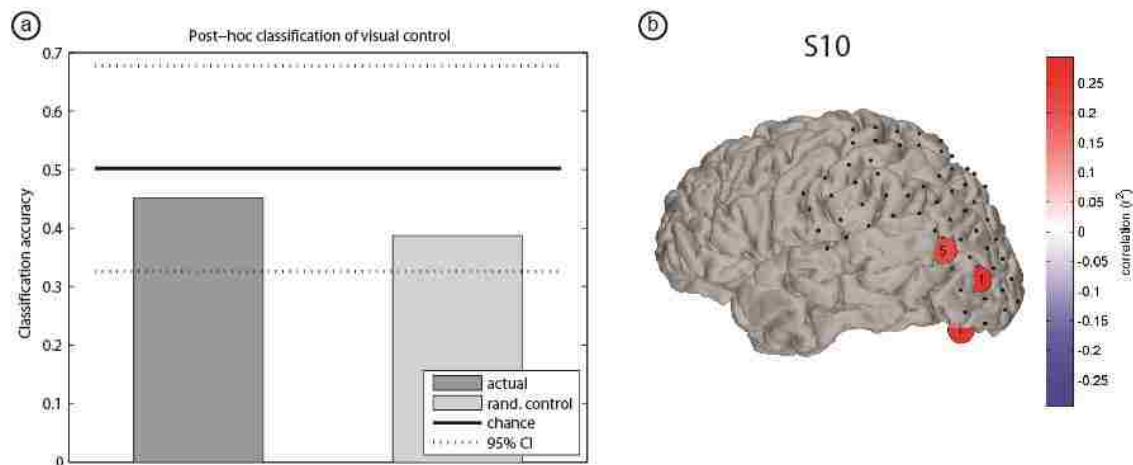


Figure 38 – Classification results in 10-fold cross-validated post-hoc classification analyses for the single visual control subject. All depicted values are as described in Figure 37.

Regression analyses reveal ongoing representation of task demands. It is apparent that there is not a consistent, strong representation of intended movement direction during the cue phase of the two BCI tasks described above. However, in both cases, we observe consistent activity patterns across multiple subjects during the feedback period that are correlated with target direction. We utilized an electrode-

wise multiple-regression analysis approach to determine how these activity changes relate to ongoing changes in the task state, specifically cursor kinematics, momentary error, and the overall goal.

We found fairly widespread correlations between cursor velocity and HG activity. Significant correlations were observed between lagged velocity and HG in all nine subjects. Due to the nature of the task design itself (linking HG activity at the controlling channel to cursor velocity), we would expect to see a strong correlation in at least one channel for each subject, however we note that multiple electrodes from eight of nine subjects were significantly correlated with velocity. The distribution of weighted average lags (see methods) was significantly less than zero (-278 ± 88 msec, mean \pm SEM; one-sample t-test, $N=9$, $p = 0.014$). We note that though the majority of lag relationships implied that HG activity was leading cursor velocity, there were significant relationships in many subjects in which changes in HG lagged changes in velocity. This implies that these regions are responding to the velocity of the cursor as it moves along the screen, and may play a potential role in feedback processing. Looking at Figure 39, we observe no particular spatial patterning associated with electrodes that lead or lag cursor velocity.

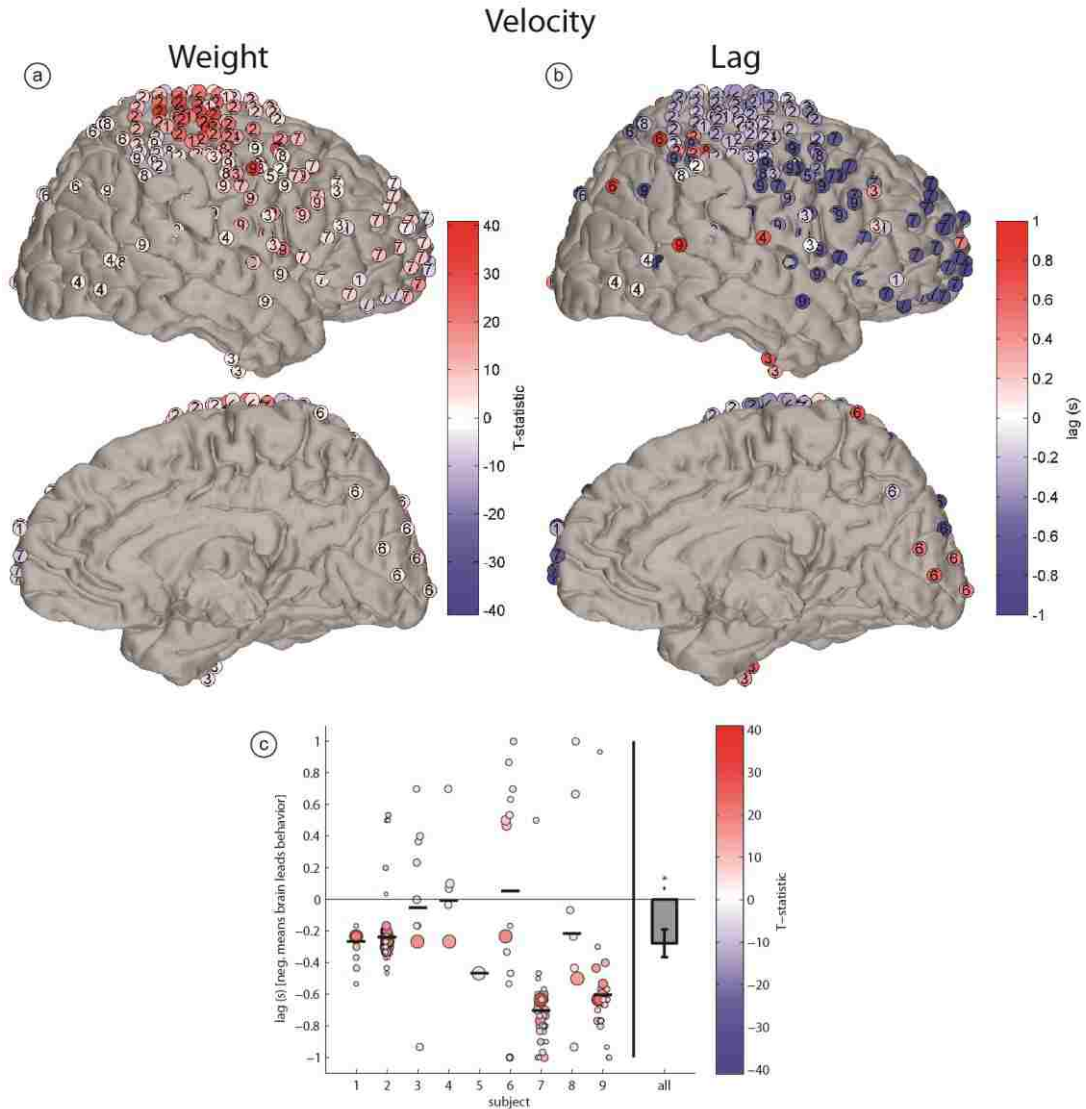


Figure 39 - Velocity results from multiple regression model. Subplots (a) and (b) depict the weights and lags of all significantly velocity-related electrodes, respectively. Both lateral and medial cortical surfaces are shown. For subjects with left-hemispheric coverage, electrode positions were projected onto the right hemisphere. Contributing subject ID's are shown in each electrode. Color of each electrode depicts the t-statistic (a) and lag (b) of the electrode-velocity relationship in the multiple regression analysis. Subplot (c) provides the same information, separated by subject. Bubble diameter and color are representative of the t-statistic in the multiple regression analysis. Horizontal black bars represent the weighted means of lags, weighted by the absolute value of t-statistic. Bar plot on the far right represents the mean and SEM of these weighted means. One star (***) denotes a p-value of < 0.05.

The second predictor included in our regression model was the signed distance between the cursor and the target. Error was represented substantially less commonly than velocity, with significant relationships in eight of nine subjects, and a median occurrence in only two electrodes per subject.

Subject S2 (6cc87c) is a notable exception to this, as they exhibited significant correlations between HG and error in 34 electrodes. Figure 40 shows the spatial distribution of these error relationships, where the dominance of subject S2 in terms of this effect can be readily observed. Though the representation of error was somewhat less prevalent than velocity, the lag at which the significant HG-error relationships occurred was significantly greater than zero (443 ± 157 msec, mean \pm SEM; one-sample t-test, $N=8$, $p = 0.039$) indicating that neural activity in the regions demonstrating significant correlations is responding to the error state of the BCI, as opposed to predicting it.

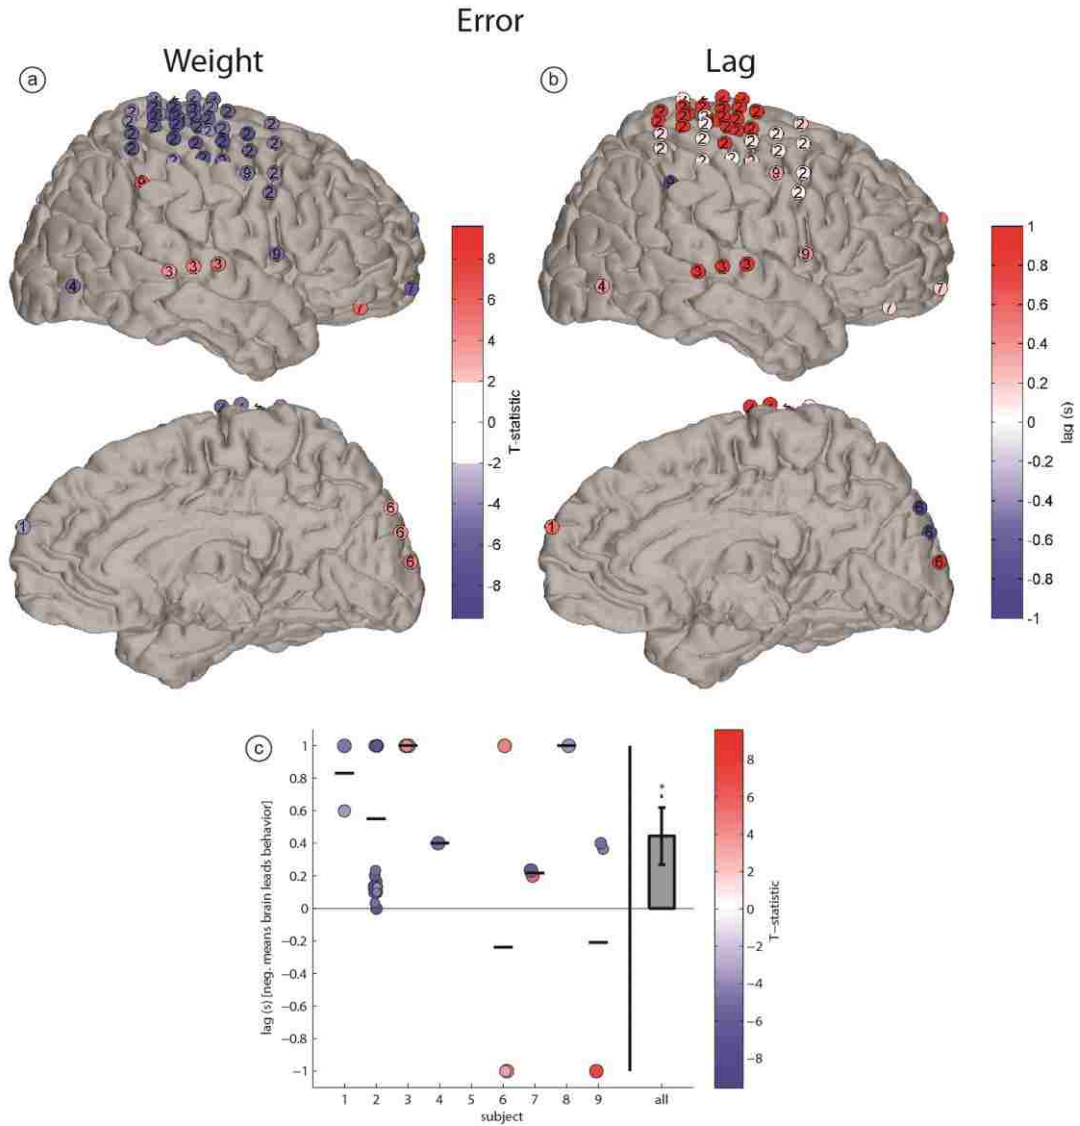


Figure 40 – Error results from multiple regression model. All depicted values are as described in Figure 39.

The interaction between error and velocity appeared to capture a more meaningful relationship between task state and neural activity. Again there were significant relationships observed in eight of nine subjects, however in this case, the median number of electrodes exhibiting significant relationships with the interaction term was 10. Similar to error, the lag at which the significant HG-interaction relationships occurred was significantly greater than zero (283 ± 113 msec, mean \pm SEM; one-sample t-test, $N=8$, $p = 0.04$) indicating that neural activity is primarily responding to this task state. Similar to the

observation of the velocity relationships, we note that in the majority of subjects (seven of nine) there are significant HG-interaction relationships at negative lags, suggesting that the brain may also be predicting the interaction between velocity and error during ongoing task execution. Spatial distribution of these relationships, as well as their strengths and lags are included in Figure 41.

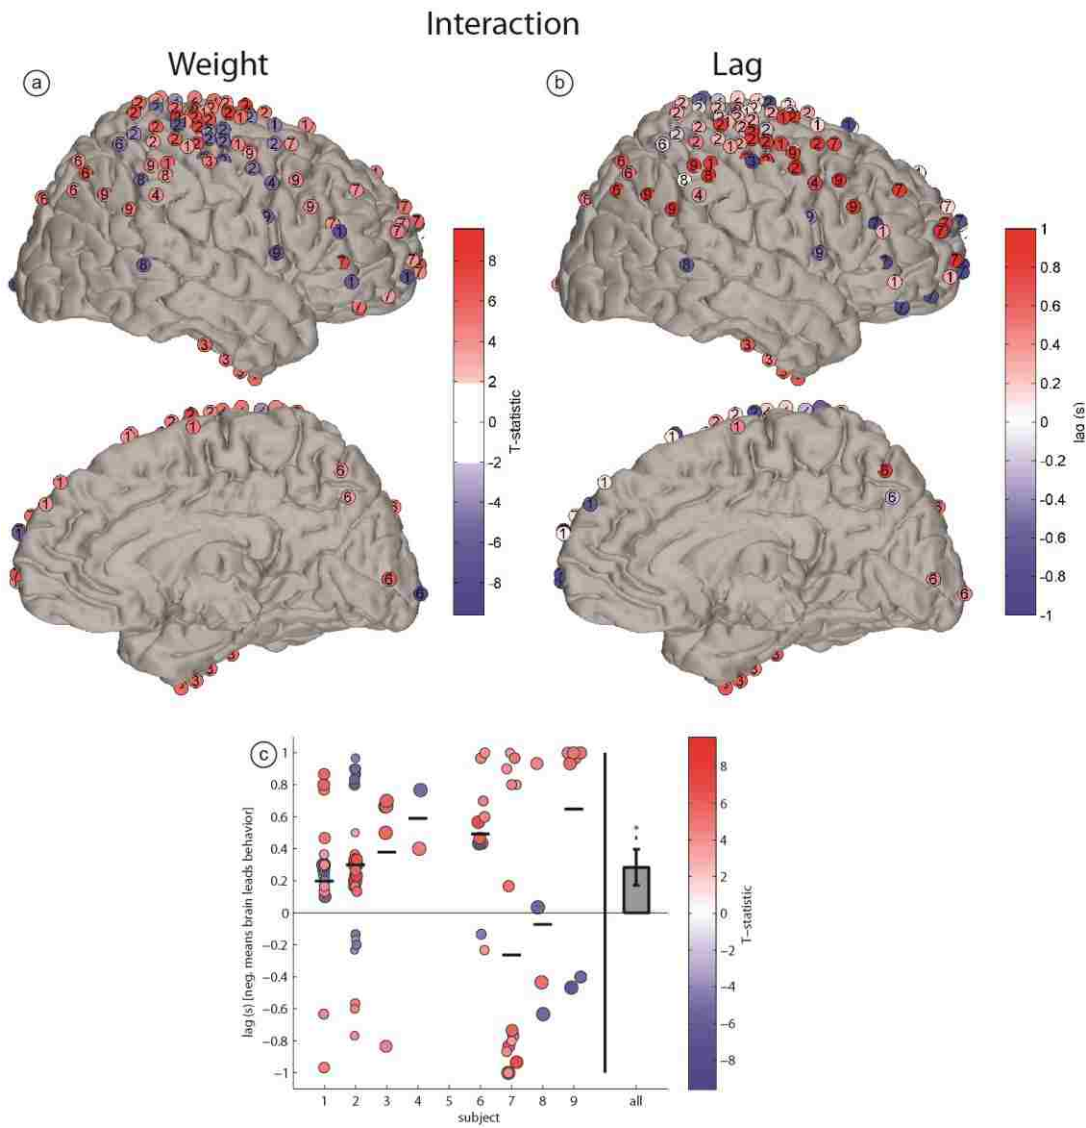


Figure 41 – Interaction (velocity x error) results from multiple regression model. All depicted values are as described in Figure 39.

The final component in our regression model was the position of the target on the screen, which is assumed to be representative of the higher-level goal that the user is attempting to reach. One benefit

of the multiple regression model is that it allows us to dissociate changes in neural activity attributable to continuously variable task state (e.g. velocity) from changes in neural activity attributable to the higher-level goal. We found a significant representation of target position during the feedback period that was primarily located outside of M1 (BA 4). The strongest representations (see Figure 42) were found in premotor (BA 6) and frontopolar and dorsal prefrontal areas (BAs 10, 11, 46).




Figure 42 – Target location results from multiple regression model. As target location is constant throughout each trial, lag is not meaningful and was not calculated. Correspondingly, only target location-related weights are shown. Both lateral and medial cortical surfaces are shown. For subjects with left-hemispheric coverage, electrode positions were projected onto the right hemisphere. Contributing subject ID's are shown in each electrode. Color of each electrode depicts the t-statistic of the electrode-target location relationship in the multiple regression analysis.

In three of eleven subjects from the legacy RJB dataset, we observed significant representation of the intended movement direction during the cue phase. Electrode coverage across these subjects was variable, based on differing clinical indication, but we observed representation of intent in cortical areas classically associated reach/grasp planning (PPC, PMv), attention (TPJ), and executive control (PFC). In ten subjects we observed significant differences in HG activity during the feedback phase as a function of intended target. In the RJB task, however, intent and execution may be tightly coupled as there is perfect correlation between target direction relative to the cursor's starting position and the correct direction of travel at any given time within the trial (see chapter 5 for additional discussion). Therefore, any activity changes observed during the feedback period could be attributable to either higher-level intention or low-level execution and dissociating the two is difficult.

To address this shortcoming, we developed a more behaviorally rich 1-D BCI task that weakened this coupling. This task, termed the GoalBCI task, is the 1-D equivalent of the center-out task and allowed for subjects to overshoot the target, requiring subsequent trajectory correction to successfully hit it.

Because we observed notable cue-phase representation of intent in a portion of the subjects from the RJB analysis, we constructed the GoalBCI task to evaluate the possibility of an ECoG-based real-time system that combined estimation of higher level intent and direct execution control, as had been demonstrated previously using single unit activity in NHP's (Shanечи et al., 2013). Testing this hierarchical approach in three subjects, we found that it did not appear to generalize to a human ECoG model, though this negative finding could partially be explained by inopportune coverage, limited classifier training data, and or limited subject sample size. We expect that successful goal inference is possible using human ECoG given application-based electrode sizing/placement and longer training periods for both the subject and the decoder. To serve as a contrast, Shanечи et al. utilized recordings from PMd and SMA, the latter of which is rarely covered in clinical human ECoG studies and the inference model in these studies was trained on approximately 3 times as many trials.

To support this hypothesis, in a post-hoc analysis using the same classifier structure but a larger subject population (all nine instead of the three who attempted the real-time approach) and more training data (10-fold cross-validation allows for training on 90% of data recorded per subject) we could successfully predict intent in three of the nine subjects. Prediction in these subjects was primarily based in posterior parietal and posterior temporal areas, suggestive that it was based on evoked visual responses, trajectory planning in parietal regions or both. Though it would not generalize to the base case of control of a mechanical end-effector in diverse environments, a system that relies on visual responses for intent classification could provide tremendous performance gains for communication-based BCIs for severely motor-impaired individuals (Santhanam et al., 2006). As has been argued previously, this is arguably one of the specific domains where BCI will first reach clinical viability (Gilja et al., 2011; Moran, 2010).

It is important to consider our findings with respect to previous approaches to infer intentional signals during BCI control. As one putative source for such signals is PPC, the Andersen group has focused heavily on building BCI systems based on activity from single neuron recordings in this region (Mulliken et al., 2008). They report the ability to recreate cursor trajectories during manual (joystick-based) cursor control from PPC ensemble activity and to utilize BCI decoders trained during manual control for real-time trajectory decoding. This work demonstrates directional tuning of PPC activity at a single-unit level, but whether the online control based on this activity is representative of successful inference of intention or direct operant conditioning is uncertain. As an extension to this, Shanechi and colleagues built a multi-stage BCI system that expressly incorporates firing activity *before* the feedback period as the prior for a Bayesian framework for trajectory decoding. They also report behavioral performance significantly above chance levels, and note that performance declines if the predictive portion of their decoder is not used. We note a few important distinctions between this study and the work described in this chapter that are valuable points of consideration when designing future ECoG-based attempts at

direct inference of higher-level intention signals. First, in both studies mentioned above, subjects were trained to perform a task that expressly involved reaching, a motor activity known to elicit activity changes in PMd and PPC (Pesaran et al., 2006; Rizzolatti et al., 1998), the two cortical regions where electrode arrays were placed. Neural activity patterns during *true* reaching movements were utilized to train decoding architectures. This is in contrast to the approach described in this document, where the BCI decoder was trained during the initial stages of BCI task execution and the cognitive task being conducted by the subject did not expressly involve reaching movements. It may be that training the subject and decoder using overt motor movements that expressly activate these two cortical regions provides more robust, directionally-tuned indicators of the higher-level goal. Second, there is a tremendous difference in spatial resolution of a intra-cortical micro array and clinical-scale ECoG. While small neuronal ensembles in PPC and PMd may exhibit directional tuning sufficient for classification of intended movement direction, these ensembles may overlap and become indistinguishable from each other at the spatial scale of macro-ECoG. Research-specific, μ ECoG arrays with inter-electrode spacing of 1-3 mm may be necessary to resolve directionally-tuned far field potentials.

Both studies found that there is a strong correlation between neural activity during the feedback phase and the position of the target that the subject was trying to reach. As we note above, this relationship (especially in the case of the RJB) may be partially attributable to the strong correlation between target position and the velocity command (and hence HG activity at the controlling electrode) necessary reach that target. In an effort to disentangle these phenomena, we employed an electrode-wise lagged multiple regression approach to characterizing the simultaneous relationships between neural activity and BCI control, feedback signals, and the higher-level goal. In these analyses we provided evidence that higher-level goal is represented during the feedback period primarily in premotor (BA6) prefrontal (BAs 10, 11 and 46), parietal (BAs 5 and 7), and superior temporal (BAs 41 and 42) areas. Activation of fronto-parietal networks during ongoing task execution is reasonable based on their demonstrated roles during

visuomotor task processing (Buneo & Andersen, 2006b). Additionally, activation of more anterior frontal regions may be indicative of a putative role in goal formation (Kobayashi et al., 2007; Rudolf & Hare, 2014) and working memory processes (Barbey et al., 2013; Funahashi et al., 1989; Goldman-Rakic, 1996). The observed activity changes in the STG are more unexpected, but may be indicative of evoked activity in networks involved in attentional visual processing of goal-directed behavior (Schultz et al., 2004), especially when considered in conjunction with the observation from the previous chapter of M1-to-STG coupling during task execution.

Inherent to all the studies described in this chapter is the underlying assumption that the roles carried out by various cortical structures during reaching and grasping motions will carry over to trajectory-based control of a BCI. Gallivan and colleagues (2011) demonstrated effector-independent responses during reach and saccade planning in parietal subregions (superior parieto-occipital cortex and IPS), as well as dorsomedial PFC and M1. However, both of these end-effectors have well established places within an individual's body schema; whether such a representation exists for a completely novel and disembodied end-effector (i.e. BCI) has not yet been demonstrated. If we consider the nervous system's ability to forward-model reach trajectories and their corresponding errors (Desmurget & Grafton, 2000) as evidence for internal representation of embodied end-effectors, then our observation of changes in HG activity before errors occur during BCI task execution provides what may be initial observations of an internal representation of a virtual end-effector.

Even if neural circuitry associated with reach and saccade planning can be quickly co-opted for trajectory planning of a BCI, we reiterate that BCI task execution is not expressly a motor task in nature. Both the biomechanics of the 'movement' and the non-visual feedback (or lack thereof) are quite different than what they would be during execution of a native motor task. BCI use can be likened as much to a cognitive task as a motor one; in some cases studies even cognitive strategies to generate activity changes in controlling areas (Friedrich et al., 2013; Vansteensel et al., 2010). Correspondingly,

the 'goal' may or may not be spatial, especially in human BCI studies, where the cognitive tool used to train the users is typically gross motor imagery (Blakely et al., 2009; Leuthardt et al., 2004) as opposed to the initially overt execution of a specific grasping or reaching motion, as is commonly used in primate studies (Carmena et al., 2003; Mulliken et al., 2008). In fact, for human users, the internal representation of the goal may not be spatial at all, it could be strategic – something more abstract. An interesting extension to the work above would be to build a BCI that could provide either auditory, tactile, or visual feedback, all presenting the same information to the user but in a variety of sensory domains (separately or simultaneously). One might expect then, that for a proficient BCI user, specific sensory association areas would be activated by the various feedback modalities, but information about task execution and intention would eventually converge, most likely in premotor areas.

Once an estimate for the subject's higher-level goal has been successfully inferred, there are two potential strategies for how to utilize that information. The first possibility is that generation and execution of the motor commands necessary to achieve that goal could be offloaded onto a robotic/automated system (Bell et al., 2008; Santhanam et al., 2006). Alternatively, the structure of the task could be modified to assist the user in achieving the inferred goal (Shanechi et al., 2013). In the studies above, we favored the latter approach, as it maintained direct participation by the BCI subject, however we recognize that as capability to perform inference improves, overall performance from the BCI system could be improved through direct execution. An interesting fusion of these two models is the concept of a hierarchical BCI (Cheung et al., 2012), that permits direct training of lower level 'BCI motor primitives' that can then be invoked through higher-level commands. System architectures like this show great long-term promise as capable and flexible BCI systems that maintain user agency, but also mitigate the requirement of constant cognitive vigilance currently necessary for successful direct BCI control.

7.5 Related publications and presentations

- [1] **JD Wander**, RPN Rao, JG Ojemann (2014). Multi-region goal inference improves performance in an invasive brain-computer interface task. 1st Annual Neurofutures conference, Seattle, WA, June, 2014.

8 Conclusions

Though BCI technology represents a powerful potential therapeutic tool for patients with severe motor disorders and, in the long-term, a potentially disruptive technology for how we interact with the world, systems based on this technology have not yet reached clinical viability. There are a number of impediments to this milestone, ranging from electrode rejection by the body to limitations in wireless power transmission. In this work, we addressed one limitation in particular: the lack of utilization of neural activity from a variety of functionally diverse cortical structures. Current BCI architectures utilize a *constant* and *direct* mapping between neural activity patterns and end-effector control. This activity is typically derived from a relatively small cortical area (e.g., M1). Such systems are relatively inflexible, require constant attentional vigilance on the part of the user, allow for little abstraction of function and do not leverage other potentially informative signals that can be extracted from elsewhere in the brain.

Utilizing a human ECoG model, performing a 1-D BCI task, we initially demonstrated that multiple cortical regions are active during BCI task execution (M1, PMv, PMd, PPC), and that a portion of these regions (primarily PMd and PFC) exhibit less activity as subjects develop task proficiency. Interestingly, this finding is consistent with observed activity patterns (both in terms of areas activated and changes in activation over time) seen during execution of a true motor task (Chein & Schneider, 2005). This suggests that in an effort to develop neuroprosthetic control, the brain harnesses pre-existing cortical learning scaffolds and provides a potentially informative target for future co-adaptive BCI architectures that are 'aware' of the user's learning process.

We then extended these findings, utilizing the same dataset, by demonstrating meaningful transcortical interactions between a number of these remote areas and the portion of the brain directly involved in BCI control. Utilizing two distinct methods – namely STWC and bPLV – we found evidence for spatial selectivity by interaction type; amplitude-amplitude interactions occurred over significantly shorter

cortical distances than phase-phase interactions. This finding is consistent with recent work predicting structural connectivity based on amplitude correlations between regions (Keller et al., 2014), and taken together they provide evidence in favor of slightly more complicated models of long-distance cortical interactions (Buzsáki & Draguhn, 2004).

Lastly, we provided evidence for functional heterogeneity of these cortical regions, demonstrating in two separate studies that activity in multiple cortical regions is representative of the intention of the BCI user, and that such intention signals are primarily present during ongoing task execution. In a subset of the study subjects, we implemented and validated a real-time system that attempted to utilize preparatory changes in HG activity to bias execution in order to improve task performance. Though these subjects did not demonstrate significant behavioral improvements, we showed in a post-hoc analysis on all subjects that behavioral improvement was possible in subjects with parieto-occipital coverage.

Interestingly, though execution of our BCI task involves no true motor execution and lacks both somatosensory and proprioceptive feedback, we found that many regions traditionally involved in visuomotor skill execution were also involved in BCI use (Chein & Schneider, 2005), and tended to represent similar facets of task execution as they would during execution of a native motor task (e.g., target location in visual processing areas and PPC, initial task learning in PFC). This implies that even though execution of neuroprosthetic control is expressly non-motor in practice, and lacks the afferent feedback that is so critical to successful motor skill execution, the brain harnesses pre-existing networks for skill acquisition and execution. One opportunity for future BCI research will be to utilize it as a scientific tool for evaluating the functionality of motor-related networks in the absence of select feedback modalities, or with experimentally modified feedback (see Wander & Rao, 2014 for further discussion).

One higher-order question that this work helps us to address is: “Which areas are important to BCI control?” The series of studies described above provide, to our knowledge, the first evidence of widespread cortical task-modulation during BCI task execution. However, in contrast to local activity changes observed by Carmena and colleagues (Ganguly et al., 2011), the explanation of simple correlation with cursor dynamics appears insufficient. We observe strong task-related activity patterns in remote regions that do not simply mirror what is taking place in M1 (e.g., target direction non-specific activity changes in dlPFC). One potential model for the activity relationships observed across these regions is that during novice task execution, prefrontal areas play an active role in coordinating information flow from sensory integration to premotor areas (Miller & Cohen, 2001), helping to build a model that maps intended target direction to necessary motor plan. As subjects develop proficiency with the task, this cognitive approach to task execution gives way to a more automated one that does not require as much active participation by prefrontal areas. This would explain both the target non-specific changes in PFC as well as the decrease in these activity patterns over the course of learning. Additionally, it could explain the lack of learning-related changes in PMv, as this region may continue to play a role in successful task execution.

As it does in execution of native visuomotor tasks, PPC appears to be involved in this process of sensory integration and target localization during a BCI task. However, it is by no means the only cortical region that represents target-specific activity changes, which is consistent with fMRI-based findings in humans performing reach, grasp and saccade tasks where direction is represented in both parietal and frontal cortices (Gallivan et al., 2011, 2013). We found though, that coverage of PPC was not sufficient to guarantee an observable representation of higher-order goal. We pose two potential explanations of this: first, it may be that spatial scale of clinical-grade ECoG grids is not fine enough to differentiate populations of parietal neurons that respond to targets in specific portions of the visual field. Second, it is quite possible that an individual’s particular task execution strategy does not activate pre-existing

visuomotor networks. In future work, strong consideration should be given to how the neuroprosthetic skill is taught and what types of motor imagery are being utilized by individual subjects.

This raises an interesting question about the teaching of the neuroprosthetic skill. As we described above, there is significant variability in how different research groups train study subjects to control a BCI (Lotte et al., 2013); in our studies alone subjects controlled the BCI based on either tongue or hand motor imagery, depending on their electrode coverage. Whether PMv, PMd or both are involved in a BCI task may depend largely on the way the individual was trained to perform the task. As an example, NHPs that are trained to perform a reaching task that is later converted to a BCI task (Mulliken et al., 2008) may preferentially involve PMd neurons, whereas human subjects that begin their BCI training with gross grasp performance (e.g., a power grip) (Blakely et al., 2009) may be more likely to recruit cortical networks involved in grasp control.

When considering our results, it is important to consider the limitations inherent with human ECoG studies. The work discussed here exclusively involved humans with intractable epilepsy and study subjects were undergoing clinical treatment at the time of their participation in this study. The nature of this subject population should be kept in mind when making extensive cross-subject generalizations and generalizations to healthy populations. Electrode placement was driven by clinical need and thus not all subjects had coverage of all cortical areas that were discussed above, though in all cases where it was possible, we made efforts to evaluate effects in multiple subjects and restrict findings observed in single individuals. Additionally, epilepsy is a disease with widespread effects, the full extent of which is still largely unknown. Depending on the root cause of an individual's epilepsy, they may have undergone significant cortical reorganization and cortical field potentials may be contaminated by inter-ictal activity. Findings from studies based on epileptic patients should bear in mind these considerations.

Use of a BCI was completely novel to the subjects in the above studies. Improved performance and accompanying changes in neural activity are to be expected throughout the entire duration of participation in the studies (approx. 1-3 d) and would probably continue long after if we had an opportunity to observe the subjects for longer time periods (Wodlinger et al., 2014). The development of robust, co-adaptive hierarchical BCI architectures will require opportunities to collect human data for long periods (weeks to months). Fortunately, with the recent advent of the Neuropace® device as a long-term ECoG-based treatment for refractory epilepsy, opportunities for truly long-term ECoG studies may be more feasible.

Another limitation to consider is the likely lack of generalizability to the BCI systems described above; this consideration is of particular importance as it is endemic to nearly all BCI designs. The hospital environment where data collection takes place and the BCI tasks that subjects are conducting are both tremendously simplistic environments when compared to the eventual reality of a deployed BCI. As Blakely and colleagues noted, BCI decoding parameters could be held constant over a few days with no appreciable performance decrease on a simple BCI task (Blakely et al., 2009), but we have no evidence yet that a static mapping would suffice when confronted with varied task demands. Within the context of the goal inference studies described above, we only had the opportunity to present subjects with a small space of potential visual targets. It is reasonable to expect that systems trained on such limited data would likely not generalize well when confronted with different stimuli that evoke a similar higher-level goal. To account for this, we believe it will be necessary to implement BCI systems that continuously adapt and retrain (over long time periods) to adjust for changes in device requirements as well as underlying patterns of neural activity related to intention and execution signals.

The conclusions of this work must also be tempered by the fact that these are observational studies. The functional roles of each of these cortical areas are inferred based on correlative evidence. Interventional studies that either temporarily (e.g., electrical stimulation, neurotransmitter inhibition) or permanently

(e.g., lesion studies in appropriate models) disrupt function in select cortical areas will be necessary to confirm specific roles of regions and networks. In contrast, facilitation of activity in these cortical areas through enhancement of local plasticity during or between sessions may provide an opportunity to begin to translate previous studies in cortical plasticity (Jackson, Mavoori, et al., 2006) to clinically functional outcomes.

In line with some of the criticisms described above, we think a highly relevant line of inquiry to be pursued in the future will be to further our understanding of how different cortical networks are selectively recruited in response to changes in the BCI paradigm and learning environment. One example of this that has been mentioned previously is to explicitly train human BCI subjects using grasp vs reach motor imagery to determine whether there is any selectivity (either initial or continuous) in recruitment of specific fronto-parietal networks. Such a study may be more appropriate for either simultaneous EEG-fMRI or micro-scale (2-3 mm inter-electrode spacing) human ECoG based on the small spatial scale of the specific subregions of PPC that one would want to observe. Furthermore, in the work above, we investigated cortical networks involved in learning of a BCI task *de novo*. Perturbation of this learned system, either through manipulation of the decoder or the specific rules of the task, will provide us an opportunity to investigate the role of carry-over learning and the BCI equivalent of motor-skill adaptation, both of which are of particular relevance to generalizable BCI frameworks.

Though it was only touched on briefly in Chapter 7, neural representation of error has been well studied previously (Falkenstein et al., 2000; Gemba et al., 1986; Milekovic et al., 2012), even in the context of BCI task execution (Buttfield et al., 2006; Ferrez & Millán, 2008; Wander et al., 2013). In terms of the latter, event-related error potentials have been employed extensively as a means to perform quasi-supervised BCI decoder co-adaptation (Buttfield et al., 2006; Ferrez & Millán, 2008; Gürel & Mehring, 2012). Continuous error tracking, on the other hand, has not been explored as thoroughly and may be

an excellent candidate for an additional channel of information to pass to multi-level BCI architectures. At the very least, such a signal may be predictive of pending failure or even danger to a BCI user.

While we are working toward the necessary advancements to make motor-based BCIs a clinically viable technology, it is important to keep in mind that the term BCI encompasses many different types of systems, ranging from stimulators for sensory substitution (e.g., cochlear implant) and symptomatic treatment (e.g., deep brain stimulation for Parkinson's disease) to recording devices that will someday allow locked-in patients to communicate. Advances made in any one of these sub-domains can and should benefit ongoing investigation in the others, and at every step along the way where it is feasible, we should be leveraging BCI technology not only as a potential therapeutic medical device, but also as an enabling technology to better understand the inner workings of the nervous system.

~

The primary contribution of the work described above is that it demonstrates the fact that volitional modulation of activity in a small portion of the brain to control a BCI is, in fact, a concerted effort on the part of multiple cortical structures. This, in and of itself, is a novel finding, however it is intended to serve as a springboard for a new way of thinking about BCI architectures. We hope that it will encourage future scientific inquiry and systems development that leverages signals from multiple cortical areas, on multiple spatial scales to develop highly functional, hierarchical BCI systems.

9 References

- Adkins, D. L., Campos, P., Quach, D., Borromeo, M., Schallert, K., & Jones, T. a. (2006). Epidural cortical stimulation enhances motor function after sensorimotor cortical infarcts in rats. *Experimental Neurology*, 200(2), 356–70.
- Adkins, D. L., Hsu, J. E., & Jones, T. a. (2008). Motor cortical stimulation promotes synaptic plasticity and behavioral improvements following sensorimotor cortex lesions. *Experimental Neurology*, 212(1), 14–28.
- Adkins-Muir, D. L., & Jones, T. a. (2003). Cortical electrical stimulation combined with rehabilitative training: enhanced functional recovery and dendritic plasticity following focal cortical ischemia in rats. *Neurological Research*, 25(8), 780–8.
- Aydore, S., Pantazis, D., & Leahy, R. M. (2013). A note on the phase locking value and its properties. *NeuroImage*, 74, 231–44.
- Barbey, A. K., Koenigs, M., & Grafman, J. (2013). Dorsolateral prefrontal contributions to human working memory. *Cortex*, 49(5), 1195–205.
- Batista, a P., Buneo, C. a, Snyder, L. H., & Andersen, R. a. (1999). Reach plans in eye-centered coordinates. *Science*, 285(5425), 257–60.
- Bell, C. J., Shenoy, P., Chalodhorn, R., & Rao, R. P. N. (2008). Control of a humanoid robot by a noninvasive brain-computer interface in humans. *Journal of Neural Engineering*, 5(2), 214–20.
- Birbaumer, N., Kubler, A., Ghanayim, N., Hinterberger, T., Perelmouter, J., Kaiser, J., ... Flor, H. (2000). The thought translation device (TTD) for completely paralyzed patients. *IEEE Transactions on Rehabilitation Engineering*, 8(2), 190–193.
- Blakely, T., Miller, K., Zanos, S., Rao, R. P. N., & Ojemann, J. (2009). Robust, long-term control of an electrocorticographic brain-computer interface with fixed parameters. *Neurosurgical Focus*, 27(1), E13.
- Blakely, T., Ojemann, J. G., & Rao, R. P. N. (2014). Short-time windowed covariance : A metric for identifying non-stationary , event-related covariant cortical sites. *Journal of Neuroscience Methods*, 222, 24–33.
- Brovelli, A., Ding, M., Ledberg, A., Chen, Y., Nakamura, R., & Bressler, S. L. (2004). Beta oscillations in a large-scale sensorimotor cortical network: directional influences revealed by Granger causality. *Proceedings of the National Academy of Sciences of the United States of America*, 101(26), 9849–54.
- Bullmore, E., & Sporns, O. (2012). The economy of brain network organization. *Nature Reviews. Neuroscience*, 13(5), 336–49.

- Buneo, C. a, & Andersen, R. A. (2006a). The posterior parietal cortex: sensorimotor interface for the planning and online control of visually guided movements. *Neuropsychologia*, 44(13), 2594–606.
- Buneo, C. a, Jarvis, M. R., Batista, A. P., & Andersen, R. a. (2002). Direct visuomotor transformations for reaching. *Nature*, 416(6881), 632–6.
- Buneo, C. A., & Andersen, R. A. (2006b). The posterior parietal cortex: sensorimotor interface for the planning and online control of visually guided movements. *Neuropsychologia*, 44(13), 2594–606.
- Buttfield, A., Ferrez, P., & Millán, J. (2006). Towards a robust BCI: error potentials and online learning. *IEEE TNSRE*, 14(2), 164–8.
- Buzsaki, G., & Chrobak, J. J. (1995). Temporal structure in spatially organized neuronal ensembles: A role for interneuronal networks. *Current Opinion in Neurobiology*, 5, 504–510.
- Buzsáki, G., & Draguhn, A. (2004). Neuronal oscillations in cortical networks. *Science (New York, N.Y.)*, 304(5679), 1926–9.
- Buzsáki, G., Geisler, C., Henze, D. a, & Wang, X.-J. (2004). Interneuron Diversity series: Circuit complexity and axon wiring economy of cortical interneurons. *Trends in Neurosciences*, 27(4), 186–93.
- Buzsáki, G., Logothetis, N., & Singer, W. (2013). Scaling brain size, keeping timing: evolutionary preservation of brain rhythms. *Neuron*, 80(3), 751–64.
- Canolty, R., Edwards, E., Dalal, S. S., Soltani, M., Nagarajan, S. S., Kirsch, H. E., ... Knight, R. T. (2006). High gamma power is phase-locked to theta oscillations in human neocortex. *Science*, 313(5793), 1626–8.
- Canolty, R., & Knight, R. (2010). The functional role of cross-frequency coupling. *Trends in Cognitive Sciences*, 14(11), 506–515.
- Carmena, J. M., Lebedev, M. A., Crist, R. E., O’Doherty, J. E., Santucci, D. M., Dimitrov, D. F., ... Nicolelis, M. A. L. (2003). Learning to control a brain-machine interface for reaching and grasping by primates. *PLoS Biology*, 1(2), E42.
- Chang, C.-C., & Lin, C.-J. (2011). Libsvm: A Library for Support Vector Machines. *ACM Transactions on Intelligent Systems and Technology*, 2(3), 1–27.
- Chapin, J. K., Moxon, K. A., Markowitz, R. S., & Nicolelis, M. A. L. (1999). Real-time control of a robot arm using simultaneously recorded neurons in the motor cortex. *Nature Neuroscience*, 2(7), 664–70.
- Chein, J. M., & Schneider, W. (2005). Neuroimaging studies of practice-related change: fMRI and meta-analytic evidence of a domain-general control network for learning. *Brain Research*, 25(3), 607–23.
- Cheung, W., Sarma, D., Scherer, R., & Rao, R. P. N. (2012). Simultaneous brain-computer interfacing and motor control: expanding the reach of non-invasive BCIs. *IEEE EMBS*, 2012, 6715–8.

- Collinger, J. L., Wodlinger, B., Downey, J. E., Wang, W., Tyler-Kabara, E. C., Weber, D. J., ... Schwartz, A. B. (2013). High-performance neuroprosthetic control by an individual with tetraplegia. *Lancet*, 381(9866), 557–64.
- Damasio, A. R., Damasio, H., & Christen, Y. (Eds.). (1996). *Neurobiology of Decision-Making* (p. 219). Berlin, Heidelberg: Springer.
- Darvas, F., Miller, K., Rao, R. P. N., & Ojemann, J. G. (2009). Nonlinear phase-phase cross-frequency coupling mediates communication between distant sites in human neocortex. *The Journal of Neuroscience*, 29(2), 426–35.
- Darvas, F., Ojemann, J. G., & Sorensen, L. B. (2009). Bi-phase locking - a tool for probing non-linear interaction in the human brain. *NeuroImage*, 46(1), 123–32.
- Darvas, F., Pantazis, D., Kucukaltun-Yildirim, E., & Leahy, R. M. (2004). Mapping human brain function with MEG and EEG: methods and validation. *NeuroImage*, 23 Suppl 1, S289–99.
- Darvas, F., Scherer, R., Ojemann, J. G., Rao, R. P. N., Miller, K., & Sorensen, L. B. (2010). High gamma mapping using EEG. *NeuroImage*, 49(1), 930–8.
- Davare, M., Kraskov, A., Rothwell, J. C., & Lemon, R. N. (2011). Interactions between areas of the cortical grasping network. *Current Opinion in Neurobiology*, 21(4), 565–70.
- Desmurget, M., Epstein, C. M., Turner, R. S., Prablanc, C., Alexander, G. E., & Grafton, S. T. (1999). Role of the posterior parietal cortex in updating reaching movements to a visual target. *Nature Neuroscience*, 2(6), 563–7.
- Desmurget, M., & Grafton, S. (2000). Forward modeling allows feedback control for fast reaching movements. *Trends in Cognitive Sciences*, 4(11), 423–431.
- Doble, J., Haig, A., & Anderson, C. (2003). Impairment, activity, participation, life satisfaction, and survival in persons with locked-in syndrome for over a decade: follow-up on a previously reported cohort. *J Head Trauma Rehabil*, 18(5), 435–444.
- Doesburg, S. M., Roggeveen, A. B., Kitajo, K., & Ward, L. M. (2008). Large-scale gamma-band phase synchronization and selective attention. *Cerebral Cortex (New York, N.Y. : 1991)*, 18(2), 386–96.
- Doyon, J., & Benali, H. (2005). Reorganization and plasticity in the adult brain during learning of motor skills. *Current Opinion in Neurobiology*, 15(2), 161–7.
- Doyon, J., & Ungerleider, L. (2002). Functional anatomy of motor skill learning. In *Neuropsychology of Memory, Third Edition* (p. 519). New York: Guilford Press.
- Evarts, E. V. (1968). Relation of pyramidal tract activity to force exerted during voluntary movement. *Journal of Neurophysiology*, 31(1), 14–27.

- Fabiani, G. E., McFarland, D. J., Wolpaw, J. R., & Pfurtscheller, G. (2004). Conversion of EEG activity into cursor movement by a brain-computer interface (BCI). *IEEE TNSRE*, 12(3), 331–8.
- Falkenstein, M., Hoormann, J., Christ, S., & Hohnsbein, J. (2000). ERP components on reaction errors and their functional significance: a tutorial. *Biological Psychology*, 51(2-3), 87–107.
- Felton, E. a, Wilson, J. A., Williams, J. C., & 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*, 106(3), 495–500.
- Ferrez, P., & Millán, J. (2008). Simultaneous real-time detection of motor imagery and error-related potentials for improved BCI accuracy. *Proceedings of the 4th International Brain-Computer Interface Meeting*, 3–8.
- Fetz, E. E. (1969). Operant conditioning of cortical unit activity. *Science*, 163(February), 955–958.
- Fetz, E. E., & Baker, M. (1973). Operantly Conditioned Patterns of Precentral Unit Activity and Correlated Responses in Adjacent Cells and Contralateral Muscles. *Journal of Neurophysiology*, 36, 179–204.
- Fetz, E. E., & Cheney, P. D. (1980). Postspike facilitation of forelimb muscle activity by primate corticomotoneuronal cells. *Journal of Neurophysiology*, 44, 751–772.
- Fitts, P. (1954). The information capacity of the human motor system in controlling the amplitude of movement. *Journal of Experimental Psychology*, 47(6), 381.
- Flor, H., Elbert, T., Knecht, S., Wienbruch, C., Pantev, C., Birbaumer, N., ... Taub, E. (1995). Phantom-limb pain as a perceptual correlate of cortical reorganization following arm amputation. *Nature*, 375(6531), 482–4.
- Friederici, A., & Rueschemeyer, S. (2003). The role of left inferior frontal and superior temporal cortex in sentence comprehension: localizing syntactic and semantic processes. *Cerebral ...*, 170–177.
- Friedrich, E. V. C., Neuper, C., & Scherer, R. (2013). Whatever works: a systematic user-centered training protocol to optimize brain-computer interfacing individually. *PloS One*, 8(9), e76214.
- Fries, P. (2005). A mechanism for cognitive dynamics: neuronal communication through neuronal coherence. *Trends in Cognitive Sciences*, 9(10), 474–80.
- Friston, K. J., Harrison, L., & Penny, W. (2003). Dynamic causal modelling. *NeuroImage*, 19(4), 1273–1302.
- Funahashi, S., Bruce, C., & Goldman-Rakic, P. (1989). Mnemonic coding of visual space in the monkey's dorsolateral prefrontal cortex. *J Neurophysiol*, 6(2).
- Fuster, J. M. (2000). Executive frontal functions. *Experimental Brain Research*, 133(1), 66–70.

- Gallivan, J., McLean, D. A., Flanagan, J. R., & Culham, J. C. (2013). Where one hand meets the other: limb-specific and action-dependent movement plans decoded from preparatory signals in single human frontoparietal brain areas. *The Journal of Neuroscience*, *33*(5), 1991–2008.
- Gallivan, J., McLean, D., Smith, F., & Culham, J. (2011). Decoding effector-dependent and effector-independent movement intentions from human parieto-frontal brain activity. *The Journal of Neuroscience*, *31*(47), 17149–68.
- Ganguly, K., & Carmena, J. M. (2009). Emergence of a stable cortical map for neuroprosthetic control. *PLoS Biology*, *7*(7).
- Ganguly, K., & Carmena, J. M. (2010). Neural Correlates of Skill Acquisition with a Cortical Brain–Machine Interface. *Journal of Motor Behavior*, *42*(6), 355–360.
- Ganguly, K., Dimitrov, D. F., Wallis, J. D., & Carmena, J. M. (2011). Reversible large-scale modification of cortical networks during neuroprosthetic control. *Nature Neuroscience*, *14*(5), 662–7.
- Gemba, H., Sasaki, K., & Brooks, V. B. (1986). “Error” potentials in limbic cortex (anterior cingulate area 24) of monkeys during motor learning. *Neuroscience Letters*, *70*(2), 223–7.
- Geng, J. J., & Mangun, G. R. (2011). Right temporoparietal junction activation by a salient contextual cue facilitates target discrimination. *NeuroImage*, *54*(1), 594–601.
- Georgopoulos, A. P., Kalaska, J. F., Caminiti, R., & Massey, J. T. (1982). On the relations between the direction of two-dimensional arm movements and cell discharge in primate motor cortex. *The Journal of Neuroscience*, *2*(11), 1527–37.
- Gilja, V., Chestek, C. a, Diester, I., Henderson, J. M., Deisseroth, K., & Shenoy, K. V. (2011). Challenges and opportunities for next-generation intracortically based neural prostheses. *IEEE Transactions on Bio-Medical Engineering*, *58*(7), 1891–9.
- Gilja, V., Nuyujukian, P., Chestek, C. a, Cunningham, J. P., Yu, B. M., Fan, J. M., ... Shenoy, K. V. (2012). A high-performance neural prosthesis enabled by control algorithm design. *Nature Neuroscience*, *15*(12), 1752–7.
- Goldman-Rakic, P. (1996). The prefrontal landscape: implications of functional architecture for understanding human mentation and the central executive. *Phil. Trans. R. Soc. Lond.*
- Gomez-Rodriguez, M., Grosse-Wentrup, M., Peters, J., Naros, G., Hill, J., Scholkopf, B., & Gharabaghi, A. (2010). Epidural ECoG Online Decoding of Arm Movement Intention in Hemiparesis. *First Workshop on Brain Decoding: Pattern Recognition Challenges in Neuroimaging*, 36–39.
- Green, A. M., & Kalaska, J. F. (2011). Learning to move machines with the mind. *Trends in Neurosciences*, *34*(2), 61–75.

- Gross, J., Timmermann, L., Kujala, J., Dirks, M., Schmitz, F., Salmelin, R., & Schnitzler, a. (2002). The neural basis of intermittent motor control in humans. *Proceedings of the National Academy of Sciences of the United States of America*, 99(4), 2299–302.
- Gürel, T., & Mehring, C. (2012). Unsupervised adaptation of brain-machine interface decoders. *Frontiers in Neuroscience*, 6(November), 164.
- Harvey, R. L., & Nudo, R. J. (2007). Cortical brain stimulation: a potential therapeutic agent for upper limb motor recovery following stroke. *Topics in Stroke Rehabilitation*, 14(6), 54–67.
- Hermes, D., Miller, K., Noordmans, H. J., Vansteensel, M. J., & Ramsey, N. F. (2010). Automated electrocorticographic electrode localization on individually rendered brain surfaces. *Journal of Neuroscience Methods*, 185(2), 293–8.
- Hikosaka, O., Nakamura, K., Sakai, K., & Nakahara, H. (2002). Central mechanisms of motor skill learning. *Current Opinion in Neurobiology*, 12(2), 217–22.
- Hochberg, L. R., Serruya, M. D., Friebs, G. M., Mukand, J. A., Saleh, M., Caplan, A. H., ... Donoghue, J. P. (2006). Neuronal ensemble control of prosthetic devices by a human with tetraplegia. *Nature*, 442(7099), 164–71.
- Hopfinger, J., Buonocore, M., & Mangun, G. (2000). The neural mechanisms of top-down attentional control. *Nature Neuroscience*, 3(3).
- Huang, M., Harvey, R. L., Stoykov, M. E., Ruland, S., Weinand, M., Lowry, D., & Levy, R. (2008). Cortical stimulation for upper limb recovery following ischemic stroke: a small phase II pilot study of a fully implanted stimulator. *Topics in Stroke Rehabilitation*, 15(2), 160–72.
- Hwang, E. J., Hauschild, M., Wilke, M., & Andersen, R. a. (2012). Inactivation of the parietal reach region causes optic ataxia, impairing reaches but not saccades. *Neuron*, 76(5), 1021–9.
- Jackson, A., Baker, S. N., & Fetz, E. E. (2006). Tests for presynaptic modulation of corticospinal terminals from peripheral afferents and pyramidal tract in the macaque. *The Journal of Physiology*, 573(Pt 1), 107–20.
- Jackson, A., Mavoori, J., & Fetz, E. E. (2006). Long-term motor cortex plasticity induced by an electronic neural implant. *Nature*, 444(7115), 56–60.
- Jenkins, I. H., Brooks, D. J., Frackowiak, R. S. J., Passingham, F. E., Nixon, P., & Passingham, R. (1994). Motor sequence learning: a study with positron emission tomography. *The Journal of Neuroscience*, 14(6), 3775–3790.
- Johnson, L. a, Blakely, T., Hermes, D., Hakimian, S., Ramsey, N. F., & Ojemann, J. G. (2012). Sleep spindles are locally modulated by training on a brain-computer interface. *Proceedings of the National Academy of Sciences of the United States of America*, 109(45), 18583–8.

- Kahana, M. J. (2006). The cognitive correlates of human brain oscillations. *The Journal of Neuroscience*, 26(6), 1669–72.
- Kandel, E. R., Schwartz, J. H., & Jessell, T. M. (2012). *Principles of Neural Science*. (E. R. Kandel, J. H. Schwartz, & T. M. Jessell, Eds.) *Neurology* (5th Editio., Vol. 4, p. 1414). McGraw-Hill.
- Kaufman, M. T., Churchland, M. M., Ryu, S. I., & Shenoy, K. V. (2014). Cortical activity in the null space: permitting preparation without movement. *Nature Neuroscience*, 17(3), 440–448.
- Keller, C. J., Honey, C. J., Entz, L., Bickel, S., Groppe, D. M., Toth, E., ... Mehta, A. D. (2014). Corticocortical evoked potentials reveal projectors and integrators in human brain networks. *The Journal of Neuroscience*, 34(27), 9152–63.
- Kennedy, P. R., & Bakay, R. A. (1998). Restoration of neural output from a paralyzed patient by a direct brain connection. *Neuroreport*, 9(8), 1707–11.
- Kipke, D. R., Shain, W., Buzsáki, G., Fetz, E. E., Henderson, J. M., Hetke, J. F., & Schalk, G. (2008). Advanced neurotechnologies for chronic neural interfaces: new horizons and clinical opportunities. *The Journal of Neuroscience*, 28(46), 11830–8.
- Kobayashi, S., Kawagoe, R., Takikawa, Y., Koizumi, M., Sakagami, M., & Hikosaka, O. (2007). Functional differences between macaque prefrontal cortex and caudate nucleus during eye movements with and without reward. *Experimental Brain Research*, 176(2), 341–55.
- Koralek, A. C., Jin, X., Long II, J. D., Costa, R. M., & Carmena, J. M. (2012). Corticostriatal plasticity is necessary for learning intentional neuroprosthetic skills. *Nature*, 483, 331–335.
- Kornhuber, H., & Deecke, L. (1965). Brain potential in voluntary movements and passive movements in man: readiness potential and reafferent potentials. *Pflugers Archiv*, (284), 1–17.
- Krusienski, D. J., & Shih, J. J. (2010). Control of a Visual Keyboard Using an Electrographic Brain-Computer Interface. *Neurorehabilitation and Neural Repair*, 25(4), 323–31.
- Kubler, A., & Kotchoubey, B. (2001). Brain-computer communication: unlocking the locked in. *Psychological Bulletin*, 127(3), 358–375.
- Lachaux, J. P., Rodriguez, E., Martinerie, J., & Varela, F. J. (1999). Measuring phase synchrony in brain signals. *Human Brain Mapping*, 8(4), 194–208.
- Lancaster, J., & Rainey, L. (1997). Automated labeling of the human brain: a preliminary report on the development and evaluation of a forward-transform method. *Human Brain Mapping*, 5(4), 238–242.
- Lancaster, J., Woldorff, M. G., Parsons, L. M., Liotti, M., Freitas, C. S., Rainey, L., ... Fox, P. T. (2000). Automated Talairach atlas labels for functional brain mapping. *Human Brain Mapping*, 10(3), 120–31.

- Lebedev, M. a, Carmena, J. M., O'Doherty, J. E., Zacksenhouse, M., Henriquez, C. S., Principe, J. C., & Nicolelis, M. a L. (2005). Cortical ensemble adaptation to represent velocity of an artificial actuator controlled by a brain-machine interface. *The Journal of Neuroscience*, 25(19), 4681–93.
- Lemieux, L., Allen, P. J., Franconi, F., Symms, M. R., & Fish, D. R. (1997). Recording of EEG during fMRI experiments: patient safety. *Magnetic Resonance in Medicine*, 38(6), 943–52.
- Leuthardt, E. C., Schalk, G., Wolpaw, J. R., Ojemann, J. G., & Moran, D. W. (2004). A brain-computer interface using electrocorticographic signals in humans. *Journal of Neural Engineering*, 1(2), 63–71.
- Levy, R., Ruland, S., Weinand, M., Lowry, D., Dafer, R., & Bakay, R. A. (2008). Cortical stimulation for the rehabilitation of patients with hemiparetic stroke: a multicenter feasibility study of safety and efficacy. *Journal of Neurosurgery*, 108(4), 707–14.
- Lew, E., Chavarriaga, R., Silvoni, S., & Millán, J. (2012). Detection of self-paced reaching movement intention from EEG signals. *Frontiers in Neuroengineering*, 5(July), 13.
- Lotte, F., Larrue, F., & Mühl, C. (2013). Flaws in current human training protocols for spontaneous Brain-Computer Interfaces: lessons learned from instructional design. *Frontiers in Human Neuroscience*, 7(September), 568.
- Mackay, A. (1992). Properties of Reach-Related Neuronal Activity in Cortical Area 7A, 67(5).
- Mauritz, K., & Wise, S. (1986). Premotor cortex of the rhesus monkey: neuronal activity in anticipation of predictable environmental events. *Experimental Brain Research*, 61(2), 229–244.
- Mayka, M., Corcos, D., Leurgans, S., & Vaillancourt, D. (2006). Three-dimensional locations and boundaries of motor and premotor cortices as defined by functional brain imaging: a meta-analysis. *Neuroimage*, 31(4), 1453–1474.
- McFarland, D. J., Lefkowitz, A. T., & Wolpaw, J. R. (1997). Design and operation of an EEG-based brain-computer interface with digital signal processing technology. *Behavior Research Methods*, 29(3), 337–345.
- McFarland, D. J., Miner, L. A., Vaughan, T. M., & Wolpaw, J. R. (2000). Mu and beta rhythm topographies during motor imagery and actual movements. *Brain Topography*, 12(3), 177–86.
- McFarland, D. J., Sarnacki, W. a, & Wolpaw, J. R. (2010). Electroencephalographic (EEG) control of three-dimensional movement. *Journal of Neural Engineering*, 7(3), 036007.
- Mellinger, J., Schalk, G., Braun, C., Preissl, H., Rosenstiel, W., Birbaumer, N., & Kübler, A. (2007). An MEG-based brain-computer interface (BCI). *NeuroImage*, 36(3), 581–93.
- Mier, H. van, Tempel, L. W., Perlmutter, J. S., Raichle, M. E., Peterson, S. E., van Mier, H., & Petersen, S. E. (1998). Changes in brain activity during motor learning measured with PET: effects of hand of performance and practice. *Journal of Neurophysiology*, 80(4), 2177–99.

- Milekovic, T., Ball, T., Schulze-Bonhage, A., Aertsen, A., & Mehring, C. (2012). Error-related electrocorticographic activity in humans during continuous movements. *Journal of Neural Engineering*, 9(2), 026007.
- Miller, E., & Cohen, J. (2001). An integrative theory of prefrontal cortex function. *Annual Review of Neuroscience*, 167–202.
- Miller, K., Hermes, D., Honey, C. J., Hebb, A. O., Ramsey, N. F., Knight, R. T., ... Fetz, E. E. (2012). Human motor cortical activity is selectively phase-entrained on underlying rhythms. *PLoS Computational Biology*, 8(9), e1002655.
- Miller, K., Hermes, D., Honey, C. J., Sharma, M., Rao, R. P. N., Nijs, M. Den, ... Leuthardt, E. C. (2010). Dynamic Modulation of Local Population Activity by Rhythm Phase in Human Occipital Cortex During a Visual Search Task. *Frontiers in Human Neuroscience*, 4(October), 1–16.
- Miller, K., Leuthardt, E. C., Schalk, G., Rao, R. P. N., Anderson, N. R., Moran, D. W., ... Ojemann, J. G. (2007). Spectral changes in cortical surface potentials during motor movement. *The Journal of Neuroscience*, 27(9), 2424–32.
- Miller, K., Schalk, G., Fetz, E. E., den Nijs, M., Ojemann, J. G., Rao, R. P. N., & Den, M. (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*, 107(9), 4430–5.
- Mirollo, R. E., & Strogatz, S. H. (1990). Synchronization of Pulse-Coupled Biological Oscillators. *SIAM Journal on Applied Mathematics*.
- Moran, D. (2010). Evolution of brain-computer interface: action potentials, local field potentials and electrocorticograms. *Current Opinion in Neurobiology*, 20(6), 741–745.
- Moritz, C. T., Perlmutter, S. I., & Fetz, E. E. (2008). Direct control of paralysed muscles by cortical neurons. *Nature*, 456(7222), 639–42.
- Müller-Gerking, J., Pfurtscheller, G., & Flyvbjerg, H. (1999). Designing optimal spatial filters for single-trial EEG classification in a movement task. *Clinical Neurophysiology*, 110, 787–798.
- Mulliken, G. H., Musallam, S., & Andersen, R. a. (2008). Decoding trajectories from posterior parietal cortex ensembles. *The Journal of Neuroscience*, 28(48), 12913–26.
- Murray, C., & Lopez, A. (1996). *The global burden of disease: a comprehensive assessment of mortality and disability from diseases, injuries and risk factors in 1990 and projected to 2020*. Cambridge, MA: Harvard University Press.
- Murthy, V., & Fetz, E. (1996). Synchronization of neurons during local field potential oscillations in sensorimotor cortex of awake monkeys. *Journal of Neurophysiology*, 76(6).
- Musallam, S., Corneil, B., & Greger, B. (2004). Cognitive control signals for neural prosthetics. *Science*, 305(July).

- Mushiake, H., Inase, M., & Tanji, J. (1991). Neuronal activity in the primate premotor, supplementary, and precentral motor cortex during visually guided and internally determined sequential movements. *Journal of Neurophysiology*, 66, 705–718.
- Nikias, C., & Mendel, J. (1993). Signal processing with higher-order spectra. *IEEE Signal Processing Magazine*, July.
- Nolte, G., Bai, O., Wheaton, L., Mari, Z., Vorbach, S., & Hallett, M. (2004). Identifying true brain interaction from EEG data using the imaginary part of coherency. *Clinical Neurophysiology : Official Journal of the International Federation of Clinical Neurophysiology*, 115(10), 2292–307.
- Orsborn, A. L., Dangi, S., Moorman, H. G., & Carmena, J. M. (2012). Closed-Loop Decoder Adaptation on Intermediate Time-Scales Facilitates Rapid BMI Performance Improvements Independent of Decoder Initialization Conditions. *IEEE TNSRE*, 20(4), 468–77.
- Pascual-Leone, A., Amedi, A., Fregni, F., & Merabet, L. B. (2005). The plastic human brain cortex. *Annual Review of Neuroscience*, 28, 377–401.
- Pellegrino, G. di, & Wise, S. (1993). Visuospatial versus visuomotor activity in the premotor and prefrontal cortex of a primate. *The Journal of Neuroscience*, 13(3), 1227–1243.
- Penfield, W., & Boldrey, E. (1937). Somatic motor and sensory representation in the cerebral cortex of man as studied by electrical stimulation. *Brain*, 60, 389–443.
- Peng, H., Long, F., & Ding, C. (2005). Feature selection based on mutual information criteria of max-dependency, max-relevance, and min-redundancy. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 27(8), 1226–1238.
- Penny, W., Friston, K., Ashburner, J., Kiebel, S., & Nichols, T. (Eds.). (2006). *Statistical Parametric Mapping: The Analysis of Functional Brain Images* (1st Editio.). Academic Press.
- Pesaran, B., Nelson, M. J., & Andersen, R. a. (2006). Dorsal premotor neurons encode the relative position of the hand, eye, and goal during reach planning. *Neuron*, 51(1), 125–34.
- Pfurtscheller, G., Guger, C., Müller, G., Krausz, G., & Neuper, C. (2000). Brain oscillations control hand orthosis in a tetraplegic. *Neuroscience Letters*, 292(3), 211–4.
- Plautz, E. J., Barbay, S., Frost, S. B., Friel, K. M., Dancause, N., Zoubina, E. V, ... Nudo, R. J. (2003). Post-infarct cortical plasticity and behavioral recovery using concurrent cortical stimulation and rehabilitative training: a feasibility study in primates. *Neurological Research*, 25(8), 801–10.
- Plow, E. B., Carey, J. R., Nudo, R. J., & Pascual-Leone, A. (2009). Invasive cortical stimulation to promote recovery of function after stroke: a critical appraisal. *Stroke; a Journal of Cerebral Circulation*, 40(5), 1926–31.
- Prablanc, C., & Martin, O. (1992). Automatic control during hand reaching at undetected two-dimensional target displacements. *Journal of Neurophysiology*, 455–469.

- Radua, J., Phillips, M. L., Russell, T., Lawrence, N., Marshall, N., Kalidindi, S., ... Surguladze, S. a. (2010). Neural response to specific components of fearful faces in healthy and schizophrenic adults. *NeuroImage*, 49(1), 939–46.
- Ray, S., Crone, N. E., Niebur, E., Franaszczuk, P. J., & Hsiao, S. S. (2008). Neural correlates of high-gamma oscillations (60-200 Hz) in macaque local field potentials and their potential implications in electrocorticography. *The Journal of Neuroscience*, 28(45), 11526–36.
- Rizzolatti, G., Fadiga, L., Gallese, V., & Fogassi, L. (1996). Premotor cortex and the recognition of motor actions. *Brain Research. Cognitive Brain Research*, 3(2), 131–41.
- Rizzolatti, G., & Luppino, G. (2001). The cortical motor system. *Neuron*, 31, 889–901.
- Rizzolatti, G., Luppino, G., & Matelli, M. (1998). The organization of the cortical motor system: new concepts. *Electroencephalography and Clinical Neurophysiology*, 106(4), 283–296.
- Rudorf, S., & Hare, T. a. (2014). Interactions between Dorsolateral and Ventromedial Prefrontal Cortex Underlie Context-Dependent Stimulus Valuation in Goal-Directed Choice. *Journal of Neuroscience*, 34(48), 15988–15996.
- Salinas, E., & Sejnowski, T. (2001). Correlated neuronal activity and the flow of neural information. *Nature Reviews Neuroscience*, 2, 539–550.
- Salvaris, M., & Haggard, P. (2014). Decoding intention at sensorimotor timescales. *PloS One*, 9(2), e85100.
- Santhanam, G., Ryu, S. I., Yu, B. M., Afshar, A., & Shenoy, K. V. (2006). A high-performance brain-computer interface. *Nature*, 442(7099), 195–8.
- Schalk, G., Kubánek, J., Miller, K., Anderson, N. R., Leuthardt, E. C., Ojemann, J. G., ... Wolpaw, J. R. (2007). Decoding two-dimensional movement trajectories using electrocorticographic signals in humans. *Journal of Neural Engineering*, 4(3), 264–75.
- Schalk, G., McFarland, D. J., Hinterberger, T., Birbaumer, N., & Wolpaw, J. R. (2004). BCI2000: a general-purpose brain-computer interface (BCI) system. *IEEE Transactions on Biomedical Engineering*, 51(6), 1034–43.
- Schalk, G., Miller, K., Anderson, N. R., Wilson, J. A., Smyth, M. D., Ojemann, J. G., ... Leuthardt, E. C. (2008). Two-dimensional movement control using electrocorticographic signals in humans. *Journal of Neural Engineering*, 5(1), 75–84.
- Scherberger, H., Jarvis, M. R., & Andersen, R. a. (2005). Cortical local field potential encodes movement intentions in the posterior parietal cortex. *Neuron*, 46(2), 347–54.
- Schlaug, G., Knorr, U., & Seitz, R. (1994). Inter-subject variability of cerebral activations in acquiring a motor skill: a study with positron emission tomography. *Experimental Brain Research*, 98(3), 523–34.

- Schlaug, G., Renga, V., & Nair, D. (2008). Transcranial direct current stimulation in stroke recovery. *Archives of Neurology*, 65(12), 1571–6.
- Schnitzler, A., & Gross, J. (2005). Normal and pathological oscillatory communication in the brain. *Nature Reviews. Neuroscience*, 6(4), 285–96.
- Schultz, J., Imamizu, H., Kawato, M., & Frith, C. D. (2004). Activation of the human superior temporal gyrus during observation of goal attribution by intentional objects. *Journal of Cognitive Neuroscience*, 16(10), 1695–705.
- Shanechi, M. M., Williams, Z. M., Wornell, G. W., Hu, R. C., Powers, M., & Brown, E. N. (2013). A Real-Time Brain-Machine Interface Combining Motor Target and Trajectory Intent Using an Optimal Feedback Control Design. *PLoS ONE*, 8(4), e59049.
- Shen, L., & Alexander, G. (1997). Preferential representation of instructed target location versus limb trajectory in dorsal premotor area. *Journal of Neurophysiology*, 1195–1212.
- Simeral, J. D., Kim, S. P., Black, M. J., Donoghue, J. P., & Hochberg, L. R. (2011). Neural control of cursor trajectory and click by a human with tetraplegia 1000 days after implant of an intracortical microelectrode array. *Journal of Neural Engineering*, 8(2), 025027.
- Sitaram, R., Weiskopf, N., Caria, A., Veit, R., Erb, M., & Birbaumer, N. (2008). fMRI Brain-Computer Interfaces. *IEEE Signal Processing Magazine*, 25(1), 95–106.
- Smith, E., & Delargy, M. (2005). Locked-in syndrome. *BMJ*, 330(7488), 406.
- Tanji, J., & Evarts, E. (1976). Anticipatory activity of motor cortex neurons in relation to direction of an intended movement. *Journal of Neurophysiology*, 1062–1068.
- Tanné-Gariépy, J., Rouiller, E. M., & Boussaoud, D. (2002). Parietal inputs to dorsal versus ventral premotor areas in the macaque monkey: evidence for largely segregated visuomotor pathways. *Experimental Brain Research*, 145(1), 91–103.
- Thomas, E., Dyson, M., & Clerc, M. (2013). An analysis of performance evaluation for motor-imagery based BCI. *Journal of Neural Engineering*, 10(3), 031001.
- Tibshirani, R. (1996). Regression shrinkage and selection via the lasso. *Journal of the Royal Statistical Society*, 58(1), 267–288.
- Van Elswijk, G., Maij, F., Schoffelen, J.-M., Overeem, S., Stegeman, D. F., & Fries, P. (2010). Corticospinal beta-band synchronization entails rhythmic gain modulation. *The Journal of Neuroscience*, 30(12), 4481–8.
- Vansteensel, M. J., Hermes, D., Aarnoutse, E. J., Bleichner, M. G., Schalk, G., van Rijen, P. C., ... Ramsey, N. F. (2010). Brain-computer interfacing based on cognitive control. *Annals of Neurology*, 67(6), 809–16.

- Velliste, M., Perel, S., Spalding, M. C., Whitford, A. S., & Schwartz, A. B. (2008). Cortical control of a prosthetic arm for self-feeding. *Nature*, *453*(7198), 1098–101.
- Vesia, M., & Crawford, J. D. (2012). Specialization of reach function in human posterior parietal cortex. *Experimental Brain Research*, *221*(1), 1–18.
- Vidaurre, C., Sannelli, C., Müller, K.-R., & Blankertz, B. (2010). Machine-Learning-Based Coadaptive Calibration for Brain-Computer Interfaces. *Neural Computation*, *23*, 791–816.
- Vidaurre, C., Sannelli, C., Müller, K.-R., & Blankertz, B. (2011). Co-adaptive calibration to improve BCI efficiency. *Journal of Neural Engineering*, *8*(2), 025009.
- Wander, J. D., Olson, J. D., Ojemann, J. G., & Rao, R. P. N. (2013). Cortically-Derived Error-Signals During BCI Use. In *Proceedings of the 5th International Brain-Computer Interface Meeting*.
- Wander, J. D., & Rao, R. P. N. (2014). Brain-computer interfaces: a powerful tool for scientific inquiry. *Current Opinion in Neurobiology*, *25*, 70–5.
- Wang, W., Collinger, J. L., Degenhart, A. D., Tyler-Kabara, E. C., Schwartz, A. B., Moran, D. W., ... Boninger, M. L. (2013). An Electrographic Brain Interface in an Individual with Tetraplegia. *PloS One*, *8*(2), e55344.
- Wang, Z., Gunduz, A., Brunner, P., Ritaccio, A. L., Ji, Q., & Schalk, G. (2012). Decoding onset and direction of movements using Electrographic (ECoG) signals in humans. *Frontiers in Neuroengineering*, *5*(August), 15.
- Weinrich, M., & Wise, S. (1982). The premotor cortex of the monkey. *The Journal of Neuroscience*, *2*(9), 1329–45.
- Williams, J. J., Rouse, A. G., Thongpang, S., Williams, J. C., & Moran, D. W. (2013). Differentiating closed-loop cortical intention from rest: building an asynchronous electrographic BCI. *Journal of Neural Engineering*, *10*(4), 046001.
- Wise, S., & Boussaoud, D. (1997). Premotor and parietal cortex: Corticocortical connectivity and combinatorial computations. *Annual Review of Neuroscience*, *(20)*, 25–42.
- Wodlinger, B., Downey, J. E., Tyler-Kabara, E. C., Schwartz, a B., Boninger, M. L., & Collinger, J. L. (2014). Ten-dimensional anthropomorphic arm control in a human brain-machine interface: difficulties, solutions, and limitations. *Journal of Neural Engineering*, *12*(1), 016011.
- Wolpaw, J. R. (2007). Brain-computer interfaces as new brain output pathways. *The Journal of Physiology*, *579*(Pt 3), 613–9.
- Wolpaw, J. R., Birbaumer, N., McFarland, D. J., Pfurtscheller, G., & Vaughan, T. M. (2002). Brain-computer interfaces for communication and control. *Clinical Neurophysiology*, *113*(6), 767–791.

- Wolpaw, J. R., & 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*, 101(51), 17849–54.
- Wolpaw, J. R., McFarland, D. J., Neat, G. W., & Forneris, C. A. (1991). An EEG-based brain-computer interface for cursor control. *Electroencephalography and Clinical Neurophysiology*, 78(3), 252–9.
- Wolpaw, J. R., McFarland, D. J., Vaughan, T. M., & Schalk, G. (2003). The Wadsworth Center brain-computer interface (BCI) research and development program. *IEEE TNSRE*, 11(2), 204–7.
- Womelsdorf, T., Schoffelen, J.-M., Oostenveld, R., Singer, W., Desimone, R., Engel, A. K., & Fries, P. (2007). Modulation of neuronal interactions through neuronal synchronization. *Science*, 316(5831), 1609–12.
- Wu, Y. C., & Voda, J. A. (1985). User-friendly communication board for nonverbal, severely physically disabled individuals. *Archives of Physical Medicine and Rehabilitation*, 66(12), 827–8.
- Xiao, J., Padoa-Schioppa, C., & Bizzi, E. (2006). Neuronal correlates of movement dynamics in the dorsal and ventral premotor area in the monkey. *Experimental Brain Research*, 168(1-2), 106–19.
- Yantis, S., Schwarzbach, J., Serences, J. T., Carlson, R. L., Steinmetz, M. a, Pekar, J. J., & Courtney, S. M. (2002). Transient neural activity in human parietal cortex during spatial attention shifts. *Nature Neuroscience*, 5(10), 995–1002.
- Young, B. M., Nigogosyan, Z., Remsik, A., Walton, L. M., Song, J., Nair, V. a, ... Prabhakaran, V. (2014). Changes in functional connectivity correlate with behavioral gains in stroke patients after therapy using a brain-computer interface device. *Frontiers in Neuroengineering*, 7(July), 25.