10-10-2013

ECoG correlates of visuomotor transformation, neural plasticity, and application to a force-based brain computer interface

Jordan John Williams
Washington University in St. Louis

Follow this and additional works at: https://openscholarship.wustl.edu/etd
 Part of the Biomedical Engineering and Bioengineering Commons

WASHINGTON UNIVERSITY IN ST. LOUIS

School of Engineering and Applied Science

Department of Biomedical Engineering

Dissertation Examination Committee:
Daniel W. Moran, Chair
Dennis L. Barbour
John P. Cunningham
Timothy E. Holy
Lawrence H. Snyder
Kilian Q. Weinberger

ECoG Correlates of Visuomotor Transformation, Neural Plasticity, and Application to a

Force-Based Brain Computer Interface

by

Jordan John Williams

A dissertation presented to the Graduate School of Arts and Sciences
of Washington University in partial fulfillment of the
requirements for the degree of

DOCTOR OF PHILOSOPHY

December 2013
Saint Louis, Missouri

Contents

List of Tables

List of Figures

Acknowledgments

At the culmination of my graduate school career, I would like to thank those who helped me reach this point and traveled with me on the journey. First, I would like to thank my parents, Darlys and Gary, for being my biggest supporters in every single activity I have participated in over the years and for being the biggest proponents of higher education. I would like to thank my older sister and brother, Kristi and Justin, for inspiring me to become an engineer and for letting their little brother tag along at a young age to see what science is all about.

Finally, I need to thank my loving wife, Rebel, for bearing with me through this entire process and reminding me that there is life outside of lab, as well as our baby daughter, Lyric, for the smile she can bring to my face after the most frustrating of days as well as her endless sense of wonder, amazement, and curiosity with the world around her that reminds us what science is all about.

<div align="right">

Jordan John Williams

</div>

Washington University in St. Louis
December 2013

For Rebel, Lyric, Mocha, and Java.

ABSTRACT OF THE DISSERTATION

ECoG correlates of visuomotor transformation, neural plasticity, and application to a force-based brain computer interface

by

Jordan John Williams

Doctor of Philosophy in Biomedical Engineering

Washington University in St. Louis, 2013

Professor Daniel W. Moran, Chair

Electrocorticography (ECoG) has gained increased notoriety over the past decade as a possible recording modality for Brain-Computer Interface (BCI) applications that offers a balance of minimal invasiveness to the patient in addition to robust spectral information over time. More recently, the scale of ECoG devices has begun to shrink to the order of micrometer diameter contacts and millimeter spacings with the intent of extracting more independent signals for BCI control within less cortical real-estate. However, most control signals to date, whether within the field of ECoG or any of the more seasoned recording techniques, have translated their control signals to kinematic control parameters (*i.e.* position or velocity of an object) which may not be practical for certain BCI applications such as functional neuromuscular stimulation (FNS). Thus, the purpose of this dissertation was to present a novel application of ECoG signals to a force-based control algorithm and address its feasibility for such a BCI system. Micro-ECoG arrays constructed from thin-film polyimide were implanted epidurally over areas spanning premotor, primary motor, and

parietal cortical areas of two monkeys (three hemispheres, three arrays). Monkeys first learned to perform a classic center-out task using a brain signal-to-velocity mapping for control of a computer cursor. The BCI algorithm utilized day-to-day adaptation of the decoding model to match the task intention of the monkeys with no need for pre-screeening of movement-related ECoG signals. Using this strategy, subjects showed notable 2-D task profiency and increased task-related modulation of ECoG features within five training sessions. After fixing the last model trained for velocity control of the cursor, the monkeys then utilized this decoding model to control the acceleration of the cursor in the same center-out task. Cursor movement profiles under this mapping paralleled those demonstrated using velocity control, and neural control signal profiles revealed the monkeys actively accelerated and decelerated the cursor within a limited time window (1-1.5 seconds). The fixed BCI decoding model was recast once again to control the force on a virtual cursor in a novel "mass-grab" task. This task required targets not only to reach to peripheral targets but also account for an additional virtual mass as they "grabbed" each target and moved it to a second target location in the presence of the external force of gravity. Examination of the ensemble control signals showed neural adaptation to variations in the perceived mass of the target as well as the presence or absence of gravity. Finally, short rest periods were interleaved within blocks of each task type to elucidate differences between active BCI intention and rest. Using a *post-hoc* state-decoder model, periods of active BCI task control could be distinguished from periods of rest with a very high degree of accuracy (~99%). Taken together, the results from these experiments present a first step toward the design of a dynamics-based BCI system suitable for FNS applications as well as a framework for implementation of an asyncrhonous ECoG BCI.

1 Introduction

1.1 Motivation

Over the past two decades, a number of approaches to implementing a brain-computer interface (BCI) for clinical applications have been developed. From non-invasive electroencephalogram (EEG) based systems used to type by mentally selecting letters on a computer screen to using highly invasive intracortical electrode signals for control of a robotic arm, the degree of task complexity that these devices can handle is continually evolving. Goals for these applications often include restoring some degree of motor control to patients with conditions such as spinal cord injury or amyotrophic lateral sclerosis. Motor restoration might be something as simple as controlling a cursor on a computer screen or as complex as controlling an arm, whether natural or prosthetic.

In either case, current BCIs rely primarily on kinematic parameter based control signals. Indeed, our lab has recently demonstrated accurate control of a computer cursor using a velocity-based command signal extracted from an intermediately invasive recording modality known as epidural electrocorticography (EECoG). This kinematic-based control is completely sufficient for control of virtual objects such as those found in computer environments. However, this type of control signal may not be ideal for controlling real-world objects with mass (i.e. functional electrical stimulation of a paralyzed limb) that are subject to external and unexpected forces. With this in mind, we wish to expand on our current line of experiments by developing a force-based BCI that may be more suitable for prosthetic applications. Based on evidence in our previous studies that demonstrated rapid neural adaptation on the order of days when a non-human primate learned a novel velocity-based BCI algorithm, we hypothesize that a non-human primate will also be capable of learning a force-based control algorithm with comparable ease. In addition to its implications for future BCI applications, our results may shed additional insight into the functional organization of motor cortex and its encoding strategies.

1.2 Specific Aims

Given the previous motivation, three specific aims were focused upon for examination in this dissertation. Aim 1 represents a series of open-loop experiments designed to examine possible representations of visuomotor transformations present in EECoG over motor related cortical areas. Aim 2 consists of a series of closed-loop experiments geared toward examining the feasibility of acceleration or force-based BCI control signals. Finally, the experiments and analyses in Aim 3 are directed toward evaluating evidence of neural adaptation on the part of the subject toward a BCI algorithm and examining machine learning techniques to adapt a BCI algorithm to the user in order to improve performance. Specifically, the three aims are:

Aim 1: Identify cortical areas that may be amenable to modulating a force-based control signal. In searching for a suitable brain area to control a force-based BCI, we desire one of two qualities: either an area that intrinsically represents force or acceleration during normal reaching tasks or an area that exhibits plasticity in its movement parameter representation. Thus, we will:

> a) Assess intrinsic position, velocity, and acceleration representations in the EECoG activity of motor cortical areas (primary motor (M1), dorsal premotor (PMd), and parietal area 5) during joystick "reaching" tasks in which joystick position is directly translated to a virtual cursor's position. In these tasks, we will use joystick position as a proxy measure for hand position that will also allow for easy remapping to other movement parameters.
>
> b) Identify shifts in neural activity representation toward cursor movement parameters by training monkeys on joystick tasks in which joystick position is translated to a cursor's velocity or acceleration. For example, a neuron normally tuned to hand velocity may shift to cursor velocity after remapping movement parameters, thus indicating a degree of flexibility/plasticity in an area's encoding strategy.

Aim 2: Determine whether subjects can learn to modulate a force-based control signal in closed-loop, EECoG BCI tasks. We wish to address whether a force-based BCI is a viable control signal for applications such as functional neuromuscular stimulation of paralyzed limbs. Thus, we will:

a) Train monkeys on acceleration and force-based BCI tasks in a virtual environment simulating simple real-world physics.

b) Quantify movement times, trajectories, and accuracy during acceleration and force-based BCI tasks and compare with results of the conjugate joystick task to assess the obtainable degree of control.

Aim 3: Optimize current EECoG decoding strategies for multidimensional BCI control by examining adaptation on the part of the subject and control algorithm. Past experiments have demonstrated that neural plasticity on the subject's part can account for a large degree of the control obtained in simple 1-D and 2-D tasks with arbitrarily assigned control features. We wish to determine the degree to which incorporating machine learning techniques with larger feature sets can improve subject performance. Thus, we will:

a) Use post-hoc analysis of BCI trials to identify neural adaptation to a given algorithm as well as additional features that may have resulted in better performance.

b) Employ day-to-day adaptation of our control algorithm to the monkey's neural strategy to examine how these additional features translate to performance in practice.

1.3 Dissertation Organization

This thesis is divided into eight chapters. Chapter 2 provides relevant background material on current BCI technology, recent developments in ECoG-based research, and current approaches to implementing FNS. Chapter 3 presents the general experimental setups and methodology used throughout this thesis. Chapter 4 addresses

the goals outlined in Aim 1 through a series of open-loop joystick experiments, while Chapters 5 and 6 present a series of closed-loop experiments geared toward addressing Aim 2 through evaluation of an acceleration-based control signal in a center-out task (Chapter 5) and a force-based control signal in a novel "mass-grab" task (Chapter 6). Chapters 7 and 8 are geared toward investigation of Aim 3. Specifically, Chapter 7 examines differences between periods of active BCI control and rest with the goal of selectively pruning BCI training data to improve model performance. Chapter 8 goes on to demonstrate evidence of adaptation on part of both the user and the decoding algorithm using machine learning techniques and how these translate to performance. Finally, Chapter 9 summarizes the conclusions and ramifications of these experiments and presents suggestions for directions of future experiments along these lines.

2 Background

Motor impairment resulting from damage to the brain or spinal cord can have a dramatic impact on a patient's quality of life. Because these conditions often hinder the patient from performing what most would consider simple, everyday tasks, an additional burden is transferred to the patient's family as well as society. Assistive devices that rely on residual physical capabilities (e.g. head, mouth, or eye movements) have been developed to restore some degree of motor functionality and independence. These devices often require unnatural transformations from the behavior to the desired output, and they may offer limited utility for severely disabled individuals such as those with late-stage amyotrophic lateral sclerosis (ALS) in which motor function is virtually absent. However, many disease states that affect the motor system, including spinal cord injury and ALS, disrupt the body's motor output signals but leave the brain from which these motor control signals originate fully intact. As a result, this leaves an ideal opportunity to bypass the damaged area and reroute the brain's electrical activity using a direct brain-computer interface (BCI) with external devices. This redirection can use brain signals to directly control devices such as a computer mouse, a prosthetic limb, or a muscle stimulator, and could thus significantly increase a patient's independence through an expanded array of tasks they would normally be unable to perform.

2.1 BCI Recording Modalities

A variety of methods for recording brain signals in humans and non-human primates have been investigated as possible output control signals in a brain-computer interface. In general, these modalities measure extracellular potentials on the order of microvolts from individual neurons or collections of neurons. However, one major design consideration that is encountered with each modality is the trade-off between the invasiveness of a device in acquiring neural signals and the robustness of those signals. At one extreme, much work has been invested in developing BCIs based on electroencephalography (EEG) [1], [2], which is a non-invasive approach that detects neural signals using electrodes placed on the scalp. While the non-invasiveness is an

attractive characteristic, the relatively low spatial and spectral frequency content has limited the degree of accuracy that subjects can attain compared to other modalities. On the other end of the spectrum, BCIs that employ penetrating intracortical electrodes for recording single unit activity have the capacity for much greater spatial and temporal resolution and consequently have demonstrated a higher degree of accuracy [3], [4]. However, since they are highly invasive, these devices not only cause local neural and vascular damage but also increase the chance for CNS infections [5]. In addition, these implantable devices are prone to immunological tissue encapsulation that increases electrical impedance and degrades signal quality over time [6].

As a result of the two extremes these approaches represent, recent research in the BCI domain has begun to explore the efficacy of electrocorticography (ECoG) as a signal acquisition modality. The hypothesis is that ECoG may offer a much less invasive alternative to intracortical electrodes since the brain parenchyma is not penetrated while producing a control signal with much better temporal and spatial resolution than EEG. As with EEG, the signal conditioning algorithm most often involves a time-frequency analysis of the electrical activity. However, because the electrodes are placed much closer to the brain than in EEG recordings and do not experience the low-pass filtering effects of the skull/scalp, ECoG signals have better spatial resolution on the order of mm [7] and better spectral frequency content ranging up to 200 Hz [8], [9]. This approach has proven to be successful in a number of human trials [9–11] which have employed imagined motor, speech, and auditory tasks for brain control of a computer cursor. These studies have shown that the higher gamma frequencies present in ECoG (60-200 Hz) can be rapidly modified through biofeedback to meet the demands of the BCI task, and, consequently, significant control can be achieved by subjects with only minutes of training.

The ECoG arrays used in the experiments above were typically clinical arrays implanted acutely in human patients for the purpose of epilepsy monitoring and seizure foci localization. These clinical arrays generally employ configurations with millimeter order electrode diameters and inter-electrode spacing on the order of centimeters. More recent experiments in non-human primates have demonstrated asynchronous

6

decoding of arm movements over a chronic time period using similarly sized electrodes [12]. However, some experimenters have questioned whether independent information in the ECoG signal may be available on a finer scale than is apparent on standard clinical arrays, and have thus begun investigating the use of ECoG arrays with sub-millimeter electrode diameters and millimeter order spacings, *i.e.* micro-ECoG arrays [13–16]. These studies have shown promise in decoding motor actions such as speech [13], individual finger movements [14], and wrist movements [15]. In addition to their acute decoding capabilities, micro-ECoG arrays have also shown to be viable recording modalities on a chronic time scale of months in term of impedance and signal characteristics [17] as well as BCI control [18], [19]. While the total longevity of ECoG array implants have not matched that of penetrating intracortical electrodes [20] due to the relative immaturity of chronic experiments, the existing studies to date do suggest that this longer time scale is realistically attainable.

Finally, while a majority of micro- and macro-ECoG studies have examined arrays implanted subdurally, our group has recently focused on the performance of epidurally implanted arrays [18], [19]. During these experiments, monkeys demonstrated chronic performance of and neural adaptation to a BCI task in which spectral power within a 75-105 Hz band on arbitrarily chosen control channels was used to control a virtual cursor. Since these arrays do not penetrate the blood-brain barrier, they theoretically leave one more line of defense against possible infection in tact as opposed to subdural arrays. Provided that epidural arrays can offer signal fidelity on par with subdural arrays, it is our hypothesis that this type of array could provide an ideal balance of signal fidelity, robustness, and device invasiveness for a BCI platform.

2.2 Motor Cortical Representations

If the goal of a BCI system is to reproduce the complex movements that are natural to the human body in a prosthetic device, then it seems intuitive that the system would perform best if the output control signal was appropriately matched to its natural counterpart in the brain, i.e. if a robotic arm is to perform reaching tasks, the brain might adapt most naturally to a BCI system that utilizes brain signals which normally

encode arm movement parameters. Accordingly, much research has been devoted to understanding how motor-associated areas of cortex naturally act to produce volitional movements.

One of the premier debates in motor neurophysiology centers around whether primary motor cortex encodes low-level commands such as individual muscle activations or forces or whether it encodes higher-level commands such as a desired position or velocity with distal structures in the motor system (i.e. spinal cord) performing the appropriate transformations to muscle activations. Initial studies examining wrist movements [21–23] and precision gripping [24] in monkeys suggested that M1 neurons may encode various low-level parameters of movement including the magnitude and rate of change of force as well as dynamic and static components of force (force resulting/not resulting in movement, respectively). A later study by Georgopoulos et al examined 2-D representations of force production using a manipulandum [25]. These experiments suggested that neuronal firing in M1 correlated with the direction of the dynamic component of force and not the static or total force exerted by the monkey. The sum of these early experiments suggests the presence of low-level parameter encoding in motor cortical structures.

In addition to the previous experiments examining neural responses to load-bearing movements, many researchers have investigated representations of higher-order movement parameters during free-reaching and drawing tasks. Results from initial studies of this nature by Georgopoulos et al. suggested that cells in motor cortex are directionally tuned for movement, i.e. the cell's activity is highest for a given direction of movement (the cell's "preferred direction") and decreases in a cosine-tuned fashion as the movement direction progresses farther away from the preferred direction [26–28]. The activity of a population of these directionally-tuned cells can be used to predict intended movements using a population vector algorithm (PVA). Subsequent studies aimed to correlate this directional-tuning with other parameters such as position, velocity, acceleration, as well as effects of posture. While neuronal activity was found to be influenced by all of these parameters, the most well represented high-level parameters in motor cortex appear to be position and velocity, and these also appear to

be cosine tuned. Even between position and velocity, both variables may be represented in the same neuron simultaneously, and the relative weighting of each parameter's representation in a neuron's firing rate (i.e. 90% of a neuron's activity relates to encoding velocity, 10% to position) may shift throughout the course of a movement [29], [30]. For example, velocity is well represented in the middle of a reaching task when speed is at its maximum while position appears to be more represented at the end of a reach when posture maintenance is required. Overall, these studies suggest that there is a prominent representation of high-level movement parameters in motor cortex. However, these parameters are not necessarily statically represented by one neural population or another, suggesting a degree of neural plasticity in what parameters these areas may encode.

While primary cortex has been extensively studied to reveal possible encoding parameters, other motor areas have also been analyzed for their role in sensorimotor transformations, and these unique characteristics of representation may also prove useful for BCI applications. Many of the same encoding schemes that have been identified in primary motor cortex have also been observed in a number of premotor areas [31], [32]. One notable difference is that the delay periods between neural activity and associated movement in premotor cortex differs from those seen in primary motor cortex with delays being seen both earlier than M1 activity as well as after movement has been initiated. This might suggest that premotor areas are involved in both developing a motor plan to pass on to primary motor cortex as well as receiving perceptual feedback about how the intended action is actually proceeding. For example, when visual feedback of a monkey's reach is distorted such that it appears to be drawing a circle when it is physically drawing an ellipse or vice versa, ventral premotor (PMv) neurons reflect the perceived path of the hand while M1 neurons reflect the physical path [33].

In summary of these observations, we suspect that motor and premotor areas may be ideal candidates for employing a force-based BCI. The wide variety of parameter representations, both low- and high-level, as well as the ability to shift representations within a movement suggest that these areas could suitably adapt to a

number of control algorithms. In addition, the flexible representation of perceived endpoint effectors by premotor areas may be especially useful for BCI applications where movement of the desired endpoint should be dissociated from physical movements.

2.3 BCI Control Signals

Despite the variety of low- and high-level movement parameters that have been correlated with neural activity in motor cortical areas, mainly position and velocity have been used as control signal outputs to drive movement (*i.e.* non-classification) BCI devices. Part of this selective usage is driven by the ease of implementing these parameters. A position output signal from the brain can be directly mapped to the position of a cursor on a screen, while a velocity control signal can simply be integrated over time to produce a position on a screen. As force or acceleration terms are added, the equations of motion become more complex, especially as the number of degrees of freedom is increased. However, using a simple brain control output signal such as position can also have its drawbacks. If a brain-derived position signal is directly mapped to a cursor position, then any neural noise and inherent variability is also mapped directly to the screen, resulting in jittery movements. When higher order time-derivative parameters are used as control signals, the successive integrations used to transform these signals into a cursor position essentially act to low-pass filter the signal such that the net intention of movement is passed through, and smoother cursor trajectories may result. To examine this concept in more detail, consider a sinusoidal component of a control signal, *c(t,f)*, with fundamental frequency *f*:

$$c(t, f) = \sin(2\pi ft) \qquad (2.1)$$

If this control signal is used as a velocity control signal, the component of position, *p(t,f)*, due to this control signal component at frequency *f* would be:

$$p(t, f) = \int \sin(2\pi ft)\, dt \qquad (2.2)$$

$$p(t, f) = \frac{1}{2\pi f} \cos(2\pi ft) \qquad (2.3)$$

If *c(t,f)* was instead used as an acceleration control signal requiring two integrations

before translating to the final position signal, the resulting position would be:

$$p(t, f) = \iint \sin(2\pi ft)\, dt \quad (2.4)$$

$$p(t, f) = \int \frac{1}{2\pi f} \cos(2\pi ft) \quad (2.5)$$

$$p(t, f) = \frac{-1}{(2\pi f)^2} \sin(2\pi ft) \quad (2.6)$$

If we consider a hypothetical control signal in which high frequency components consist mostly of noise while lower frequency components include volitional intent, then Equation (2.6) shows that the contribution from these noisy high frequency components would be attenuated relative to low-frequency volitional components.

Several of these theoretical control scenarios were examined by Marathe and Taylor in healthy human subjects using transformations of physical reaching characteristics to control a computer cursor [34]. In these experiments, the hand position, goal (a discrete form of the end-point or target of a reaching movement, in
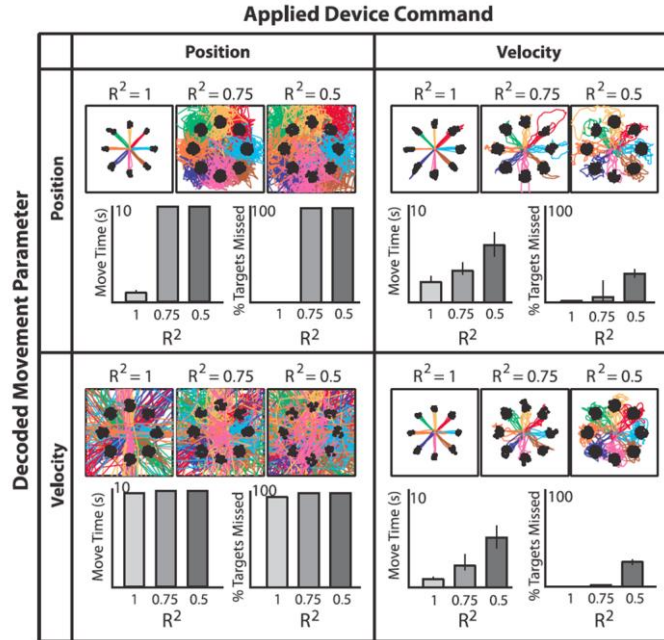
Figure 2.1 Simulated BCI performance when various levels of noise are added to position or velocity control signals. Both decoded position and velocity signals perform well when no noise is present and they are remapped to control the virtual version of themselves. However, performance degrades quickly if noise present in the control signal is remapped to position. Remapping to a velocity control signal, conversely, maintains a viable degree of performance. (From Marathe and Taylor, 2011)

these experiments simply the hand position binned into four quadrants), and velocity were remapped to the same or remaining movement characteristics of a computer cursor with various amounts of added noise. Subjects had to use these remapped physical movements in order to complete a center-out task. As expected, remapping a given noisy control signal to cursor velocity resulted in much better performance than remapping the same signal to the cursor position or target (see Figure 2.1). However, this study did not examine whether this result extended to even higher-order time derivatives such as acceleration or force. Based on the observations outlined here, we suspect that this trend will continue to these control parameters and could potentially offset the advantage in ease of translating control signals to higher-level parameters such as position and velocity.

Another factor for why position and velocity control signals are so prevalent in current BCI systems is their corresponding prevalence in motor cortex research over the past two decades as mentioned previously. A number of studies have successfully extended movement reconstruction algorithms previously used to reconstruct physical movements from neural activity to BCI control of a computer cursor [3], [4], [35]. Anecdotal evidence from these studies suggests that the monkeys would initially still make arm movements to try to complete the BCI task. After a period of training, though, the monkeys would eventually put their arms down and rely entirely on brain control for the task. Thus, the neural activity in motor cortex became dissociated from its original behavior which suggests a degree of neural plasticity that might make pre-screening an area for a particular parameter representation unnecessary. Further signs of neural plasticity include changes in the tuning profiles of neurons during closed-loop BCI experiments. Many cells used for closed-loop control become even more cosine tuned to direction in order to better match the control algorithm [36]. Cells also show changes in their preferred direction, but this adaptation is not uniform across a population (i.e. some neurons show great changes while others show hardly any change in preferred direction) [37]. Recent experiments from our laboratory have extended these notions of neural plasticity to the ECoG realm. Rouse et al. showed that monkeys can adapt their neural activity to better fit an arbitrarily chosen control algorithm as

opposed to basing a decoding algorithm off of neural activity pre-screened for movement correlations [19]. Overall, these results suggest that the brain is fairly plastic in terms of adapting to the control algorithm that is implemented, and this may aid in developing a force-based control system.

Finally, while computer cursor control has been a staple of BCI research, some studies have shown brain-control extended to more real-life applications such as controlling a robotic arm [38]. Since a robotic arm has real mass and inertia, an inverse dynamic model is necessary to transform a desired velocity signal into the appropriate forces to produce that trajectory. To date, these calculations have been computed by an external controller in the robotic arm using a known physical model of the arm and any external forces on it. In this light, control of the robotic arm is still purely kinematic from the monkey's point of view. However, for functional neuromuscular stimulation (FNS) in which a patient's muscles are externally stimulated in an attempt to produce natural volitional movements, kinematic based control signals might be an inappropriate choice for efficient control due to variability in each individual's limb characteristics. Any attempt to model an individual's arm and its environment would most likely be prohibitively complicated or incomplete, and this process would have to be repeated for each patient. In this case, it might be more appropriate to use a force-based control signal and leave any inverse model computations to the brain to learn over time. It is our hypothesis that the numerous instances of neural plasticity mentioned above make motor cortical structures viable candidates for performing these inverse model computations, and thus a force-based control algorithm may also be the most appropriate choice for more complex BCI applications such as FNS where inter-individual variability might have a significant effect.

2.4 Functional Neuromuscular Stimulation

Functional neuromuscular stimulation (FNS), also referred to as functional electrical stimulation (FES), has often been mentioned as a likely target application for BCI control. For conditions such as spinal cord injury in which the motor chain is broken but the remaining nervous system components on either side of the damaged

section remain mostly intact, FNS would offer a solution for rerouting motor commands to healthy muscles if the original path could not be repaired. Several methods for eliciting muscular activity by an external device have been investigated. These range from activating multiple muscle groups or coordinated movement primitives through microstimulation of the spinal cord [39] to stimulation of individual muscle groups through intrafasicular stimulation [40]. Clinically approved systems for cervical spinal cord injuries have also been developed that link stimulation of desired muscle groups normally innervated by spinal neurons below the lesion to movement patterns (e.g. wrist movements) generated by higher cervical roots whose function have been spared [41]. While these systems do provide a significant improvement in a patient's quality of life and independence, they also require a sharing of physical resources (e.g. must extend the wrist to stimulate grasp closing) that may limit the number of independent movements that can be performed.

More recently, however, Moritz et al. demonstrated that non-human primates can directly control the functional stimulation of a transiently paralyzed arm using the activity of neurons in motor cortex [42]. Along with establishing one of the first successful couplings of a BCI decoding system to FNS of a subject's own limbs, this study also showed that a neuron's previous movement associations did not have a significant bearing on the neuron's final control performance under BCI conditions. Consistent with the previously mentioned results by Rouse et al. [19], this finding has significant implications on BCI design considerations, especially of the dynamics-based type discussed in this thesis, as it substantially expands the potential candidate control signals.

While the study by Moritz et al. did demonstrate BCI control of select muscle groups to generate bidirectional wrist torques, the task could be considered analogous to a one-dimensional position-based task as the monkey had to match a series of static torque "targets" in order to receive a reward. Other investigators have constructed dynamic models as a theoretical test bed for more complex FNS systems. Blana et al. developed a simulated controller for a simple six-muscle, two-joint dynamic musculoskeletal arm model [43]. This controller employed an artificial neural network

14

(ANN) for the inverse-dynamic model of the feed-forward plan muscle stimulations, and a sequential proportional-integral-derivative (PID) controller and ANN as a feedback model for correcting errors due to muscle fatigue and external disturbances. An expansion of this study by Chadwick et al. showed that a human subject with tetraplegia could control a virtual version of this arm constrained to the horizontal plane with no gravity [44]. However, in both of these studies, the input from the subject to the controller was an endpoint velocity command. As a result, control of this "dynamic" system was still strictly kinematic from the subject's viewpoint as it has been for brain control of robotic arms [38], [45].

2.6 Asynchronous BCIs

Although much of the research in motor BCI development, both for ECoG and other modalities, seeks to stretch the limits of existing technologies to develop faster and more accurate control of a device, another goal within the field is to develop asynchronous systems – *i.e.* real-time decoding of intention with self-paced operation of a prosthetic device. Several experiments in non-human primates have demonstrated either asynchronous decoding of arm movements [12] or BCI control of a robotic arm [38] during feeding in which explicit task cues were not presented. Even in BCI experiments with more structured behavioral tasks such as the common center-out task as in [46], continuous control of a computer cursor could still be considered asynchronous as the subject chooses when to start a trial by moving to the center target, what path to take to the target, etc., as long as these movements fall within the parameters for completing a "correct" trial. However, as noted by Ludwig [47], these experiments often assume that the subject is intentionally controlling the device at all times. This assumption can lead to unwanted movements when the user is not intending to control the device (*e.g.* daydreaming, sleeping, distractions from the task), or it can introduce faulty training data if a subject is not actively attending to performing the BCI task.

Due to these considerations, many asynchronous designs have focused on detecting transitions between intended and unintended periods of control. In the EEG

literature, one common methodology for addressing this goal is the use of a brain controlled "switch" [48–51]. These may operate by simply detecting low-frequency imagined movement-related potentials within normal EEG activity [48], or by combining different extraction modalities such as event-related desynchronization (ERD) as a switch to engage BCI control and steady-state evoked potentials (SSEP) as the device's operative control signal [50], [51]. Achtman *et al* utilized single and multi-unit activity in PMd related to movement planning signals in a state estimator for detecting the presence of plan activity, and then used a separate model to evaluate neural activity during this planning period to classify the subject's intended target [52]. These particular EEG and single-unit examples employed models based on neural activity pre-screened for correlations with a real or imagined movement. Ludwig et al sought to circumvent this requirement for *a priori* knowledge of movement correlates by allowing rodent subjects to build their own neural output states to accomplish an auditory version of the center-out task [47]. In this study, baseline neural recordings were first mapped to the center position for the neural cursor. Then, neurally separable states were gradually identified for each target as the subject learned to generate different neural responses for each target. A Bayesian maximum-likelihood classifier was used to select the most likely intended target, and the neural response during each trial was used to update the history and probabilities associated with each target. In this manner, no *a priori* training of the decoder was necessary to start brain-control experiments, and the classifier could adapt and recalibrate to the subject over time.

After consideration of the asynchronous studies to date mentioned above, several observations arise. First, the majority of designs have incorporated a baseline or switching element into a discrete classification scheme. Theoretically, these systems would be useful in applications such as a brain-controlled keyboard to prevent undesired characters from being selected during idle states. However, to the best of our knowledge, asynchronous designs of this nature have not been applied to continuous valued decoding systems such as would be necessary for smooth control of a robotic arm. The implementation of the "No Control" (NC) state also varies considerably between designs. Several of the EEG systems mentioned previously used salient neural

features (*e.g.* SSVEPs from observing a flashing LED) to switch on and off the control state. Conversely, Ludwig *et al* used the baseline neural state as a control state for returning a cursor to the center position, making no distinction between transitioning to this state to begin a trial and naturally transitioning to it during periods of inattention to the task. Finally, asynchronous experiments to date have been limited to EEG and single/multi-unit studies as these have been the major modalities for BCI control. With the exception of several open-loop decoding experiments [12], [53], which do not rely on task events but also do not focus on detecting intermittent periods of control within a closed-loop brain control experiment, experiments examining asynchronous BCI control using ECoG have not been performed to date.

2.7 Summary

Based on the studies and principles outlined here, there are several avenues of BCI research that warrant further exploration. First and foremost for the direction of this thesis, dynamics-based control in BCI systems has yet to be thoroughly investigated as a possible control scheme, especially using ECoG as a recording modality. Control of robotic arms, single muscle groups, and simulated FNS models has been demonstrated using single-unit recordings, but control on the part of user has been purely kinematic. Due to possible advantages offered by a dynamics-based control signal over higher-level kinematic variables as well as numerous evidence of neural plasticity observed in BCI settings, a force-based control signal may indeed be a feasible alternative strategy for applications such as FNS. With recent advances in ECoG technology, epidurally implanted arrays with sub-millimeter electrode sites may offer a suitable platform for this type of system with high spatial content, long-term stability, and a relatively decreased risk to the patient compared to modalities that cross the blood-brain barrier. Many BCI studies utilizing ECoG signals to date have used pre-screened features for control, so developing algorithms to exploit neural plasticity at this scale will likely be a key component for implementing a dynamics-based control scheme. Finally, while accurately deciphering the intent of a BCI user while actively

using a neural prosthetic is very important for the utility of the device, just as important may be the ability to detect when the user is actively intending to use the device or not. This concept of asynchronous operation of a BCI system has received little attention in the ECoG arena, as has detecting these transitions in other recording modalities without pre-screening for control features. Overall, addressing these questions may not only layout important design considerations for future BCI systems, but may also illuminate important neurophysiological phenomena related to BCI operation.

3 General Experimental Methods

The experiments in this dissertation focus on the use of epidural ECoG and its application to a BCI. This chapter describes the animal protocols, device technology, behavioral and recording setups, and basic signal processing used in the extraction of cortical signals. Since these features were used in a wide variety of BCI settings and analyses, later chapters will describe the behavioral tasks and analyses specific to those experimental paradigms in more detail.

3.1 Subjects

The majority of experiments and analyses described in this dissertation utilize recordings obtained from two male, 6-10 kg Rhesus macaques (*Macaca mulatta*, monkeys K and P) chronically implanted with epidural ECoG arrays. A set of analyses and experiments in Chapter 7 utilizes data from previous experiments in our lab and an extension of these experiments. These experiments utilized two additional monkeys (one *Macaca mulatta*, monkey J, and one *Macaca fasicularis*, monkey M) also chronically implanted with epidural ECoG arrays. The design and anatomical coverage of the arrays used in the previous studies differ from the arrays implanted in monkeys K and P, but many major design considerations (electrode site diameter, spacing, material) remain similar. In addition, the behavioral setup, recording hardware, and signal acquisition and processing were identical. Thus, unless noted otherwise, fundamental findings from one array type will be assumed to cross over well to the other type. The array design and implantation of monkeys J and M will be summarized as necessary in this chapter, but further details can be found in their original publications [19].

3.2 Behavioral Setup

Each monkey learned to perform a variety of behavioral tasks through operant conditioning with a liquid reward. The monkey sat in a custom primate chair approximately 20" from an LCD monitor (Dell Inc.) displaying visual feedback of the behavioral task. The monkey's arms and legs were completely enclosed within the chair

but were free to move during BCI experiments. For training and the joystick experiments in Chapter 4, a joystick (APEM 9000 Series Joystick, APEM Components Inc.) was attached to the front of the chair with the handle extending into the interior of the chair such that the monkey could move the joystick in the frontal plane at approximately chest level. Maximum deflection of the handle in any direction was approximately 4 cm. Each monkey's left arm was restrained during joystick experiments so that the right arm controlled the joystick while left hemisphere cortical activity was recorded.

The basic behavioral task template used in this dissertation was for each monkey to interact within a virtual two-dimensional environment displayed on the computer screen through control of a small spherical cursor. This cursor was controlled using either position information relayed from the joystick or a processed control signal extracted from the monkey's ECoG signals.

3.3 ECoG Arrays

Monkeys K and P were implanted bilaterally with 32-channel thin-film polyimide arrays in the epidural space. The arrays were designed and manufactured with collaboration from the lab of Dr. Justin Williams at the University of Wisconsin (Madison, WI). Each array consisted of circular electrode sites ranging in diameter from 300-600 μm (four distinct sizes with surface areas 2, 3, or 4 times that of the smallest size) and 3 mm center-to-center spacing facing the dura (see Figure 3.1). Electrode sites were made up of gold contacts coated with platinum-black powder. Sites were arranged in a 4-by-8 rectangle with each row consisting of a single electrode size (two rows of each size). These sites were printed on one half of a flexible polyimide ribbon. On the other half, signals were referenced to a large ground pad and similarly sized reference sites sitting above the electrode sites, facing toward the skull. This referencing configuration was accomplished on a single planar printed surface by folding the polyimide ribbon in half before connecting both ends to a ZIF connector (Tucker Davis Technologies, Alachua, FL). This configuration was chosen such that skull facing ground and references would remain relatively quiet from cortical signals but still pick

up noise or artifact in the local recording area above recording sites for rejection by differential amplifiers. The base of the ZIF connector was secured in acrylic within a recording chamber with protective cap, both constructed from Delrin® plastic (Dupont). In addition, the polyimide neck of the array was coated with silicone rubber to provide cushioned relief from movement of the neck relative to the surrounding skull.

Monkeys J and M were implanted unilaterally with 28 channel epidural ECoG arrays fabricated from 300 μm platinum-iridium (90%-10%) microwires. The wires were embedded in an acrylic mold within a titanium recording chamber, and the recording ends of the microwires were ground flush with the flat acrylic surface to rest on the dura. Recording sites were arranged in a hexagonal pattern such that each electrode was 3 mm from its closest six neighbors, and the boundary of the overall pattern was roughly circular. Signals were referenced to an identical reference electrode
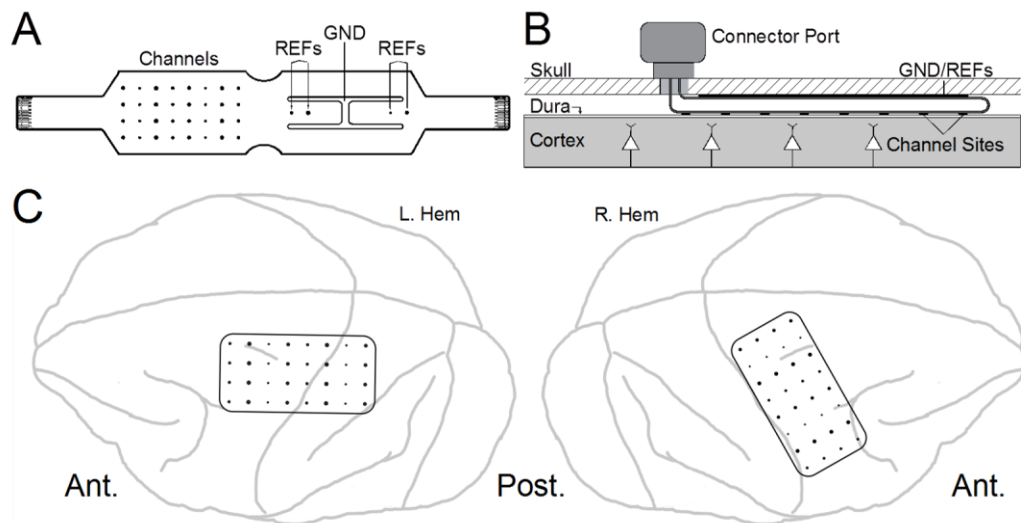
Figure 3.1 Epidural ECoG array design and placement in monkeys K and P. A) 32 channel polyimide thin-film ECoG grid layout with ground and references. B) Schematic of ECoG array placement in relation to dura, skull, and external connector port. The planar layout in (A) is folded in half before being connected to a ZIF connector within the connector port. After folding and insertion underneath the skull, channels face downward toward the brain while ground and reference sites face upward toward the skull. C) Array placement over left and right hemispheres. Left hemisphere arrays in both monkeys were oriented anterior to posterior from approximately premotor cortex to sensory and parietal areas. Right hemisphere arrays were oriented medial to lateral, anterior to central sulcus to follow primary and premotor cortical areas.

site facing the dura, and the titanium chamber sitting on the skull served as the recording ground.

3.4 Surgical Implantation

The surgical procedure for implanting monkeys K and P under general anesthesia with chronic bilateral ECoG arrays is summarized as follows: 1) A 19 mm craniotomy was manually drilled with a hand trephine while leaving the dura intact. The circular bone plug from the crantiotomy was retained for later replacement over the array. 2) The folded polyimide array was slid through the craniotomy and underneath the skull. Two surgical spatulas were used push the array and gently depress the brain from the skull as necessary. Insertion stopped once approximately 2/3 of the array had been slid underneath the skull. At this point, the silicone-coated neck of the array was situated against the far edge of the craniotomy, and half of the recording chamber rested on top of the skull. 3) The bone plug was replaced underneath the recording chamber into the craniotomy above the exposed array. The bone plug was pressed against the silicone-coated neck, and any gaps between the bone plug and skull were filled with moistened Gelfoam® (Pfizer) to provide a scaffold for skull regrowth. 4) The bone plug was secured to the rest of the skull with a 2 cm titanium strap. 5) The recording chamber, titanium strap, and bone plug were secured in place to the surrounding skull with dental acrylic. 6) Steps 1-5 were repeated on the opposite hemisphere. 7) 4-5 additional 4 cm titanium straps were attached to peripheral sections of skull near temporal level with screws to provide extra stability for an acrylic head cap. 8) The acrylic head cap was fabricated by laying thin layers of dental acrylic on top of the exposed skull, titanium straps, and bases of the recording chambers.

The left hemisphere array of each monkey was oriented in the rostral-caudal direction spanning approximately premotor areas to S1/Parietal Area 5 (see Figure 3.1c). Stereotaxic coordinates of the craniotomy center were (Anterior 15 mm, Lateral 15 mm), corresponding to the anterior most row of electrodes being centered at approximately (Anterior 20, Lateral 15). This placement and orientation was chosen to provide coverage of arm and shoulder motor areas spanning from abstract planning

(parietal areas near the posterior boundary of the array) to areas closer to physical movement execution (premotor and primary motor cortex at the anterior half of the array).

Right hemisphere arrays were oriented medial-laterally to approximately follow the pre- and primary motor strips. The craniotomy center resided at stereotaxic coordinates (Anterior 15 mm, Lateral 9 mm), placing the medial most row of electrodes at approximately (Anterior 15 mm, Lateral 4 mm). The rationale for this orientation centered around providing an alternative BCI strategy for the monkey and whether this might affect the quality of BCI control. In the left hemisphere array orientation, the array followed roughly one anatomical area (i.e. proximal arm of the motor humunculus) through multiple cortical movement processing stages. The orientation of the right hemisphere arrays followed multiple anatomical areas in mainly one or two movement processing stages. Thus, it is possible that these orientations could provide insight into whether including coverage of more diverse anatomical areas or more diverse processing areas might provide more independent signals for BCI control.

The left hemisphere was recorded from during visuomotor transformation experiments examined in Chapter 4 as well as BCI experiments in Chapters 5, 6, and 7. The right hemisphere of each monkey was intended for use in BCI experiments after left hemisphere experiments were completed, and Monkey K did complete BCI experiments with both hemispheres. However, Monkey P only completed BCI experiments with the left hemisphere due to breakage of the right hemisphere connector.

3.5 ECoG Recordings and Signal Processing

ECoG signals were recorded using Tucker Davis Technology (TDT, Alachua, FL) neurophysiology recording equipment. Raw signal waveforms were band limited to 3-500 Hz and amplified to a total of 50X gain in low-impedance headstage and pre-amplifier hardware before digitization. Digitized signals were transmitted from battery powered pre-amplifiers to TDT base stations over fiber optic cables to avoid 60 Hz contamination. Signals were sampled and processed on the base station DSPs at 6 kHz

and downsampled to 2 kHz for storage and offline analysis.

The real-time brain control schemes and offline analyses in this dissertation primarily focus on trends in spectral amplitude or power within various frequency bands. Specifically, BCI experiments utilized five frequency bands: 8-15 Hz, 15-24 Hz, 70-90 Hz, 90-110 Hz, and 130-150 Hz. The lowest two frequency bands are commonly referred to as the alpha and beta bands, respectively, and have been frequently used in EEG BCI experiments [1], [54], [55]. The highest three bands are typically referred to as residing in the high gamma frequency range, and have been used more frequently for brain control in ECoG experiments [11], [18], [56]. Spectral amplitudes in these ranges were estimated using an envelope detection method more commonly known for its use in AM radio demodulation.

$$b_{i,j}(t) = f_j[s_i(t)] \qquad (3.1)$$

$$a_{i,j}(t) = f_{<3Hz}(|b_i(t)|) \qquad (3.2)$$

The first step in envelope detection requires band pass filtering the raw signal, $s(t)$, on channel i within the frequency band j (Eqn. (3.1)). Band-pass filtering was accomplished by passing each channel through a series of two 16th order digital Butterworth filters (one high-pass filter, one low-pass filter at the band boundaries). The band limited signal, $b(t)$, was full wave rectified and low-pass filtered below 3 Hz (1st order Butterworth) to produce the spectral amplitude estimate, $a(t)$.

The spectral estimates calculated in the manner above for a given feature produce a distribution of amplitudes skewed toward higher amplitudes (see Figure 3.2). If a single feature with this characteristic was used for one-dimensional control of a computer cursor's velocity by calculating the amplitude difference from the mean (i.e. feature amplitudes above the mean would produce movement in an arbitrary "positive" direction and amplitudes below the mean would produce movement in the opposite "negative" direction), the subject may experience difficulty in moving in the negative direction as greater speeds would be possible in the positive direction. Because the integral of velocities below the mean is, by definition, equal to that of velocities above the mean, the subject should still be able to reach the same position in either case.

However, because the range of positive velocities in this skewed distribution is greater than that of negative velocities, the time needed to reach a target in the positive direction would theoretically be less than that needed to reach an equally distanced negative target. Due to this consideration, a natural log transformation (Eqn (3.3)) was performed on each feature to produce a more symmetrical distribution about the mean.

$$c_{i,j}(t) = ln(a_{i,j}) \qquad \textbf{(3.3)}$$

The Pearson 2 skewness coefficient, Sk_2, is a simple measure of a distribution's skewness and symmetry that works well for unimodal distributions [57]. This coefficient is related to a distribution's mean ($\bar{x}$), median (m), and standard deviation (σ) by:

$$Sk_2 = \frac{3(\bar{x} - m)}{\sigma} \qquad \textbf{(3.4)}$$

Positive and negative coefficient values indicate rightward and leftward skewed distributions, respectively, while a coefficient value of zero indicates a symmetric distribution. As seen in Figure 3.2, the log transformation significantly decreases the skewness of the amplitude distribution.
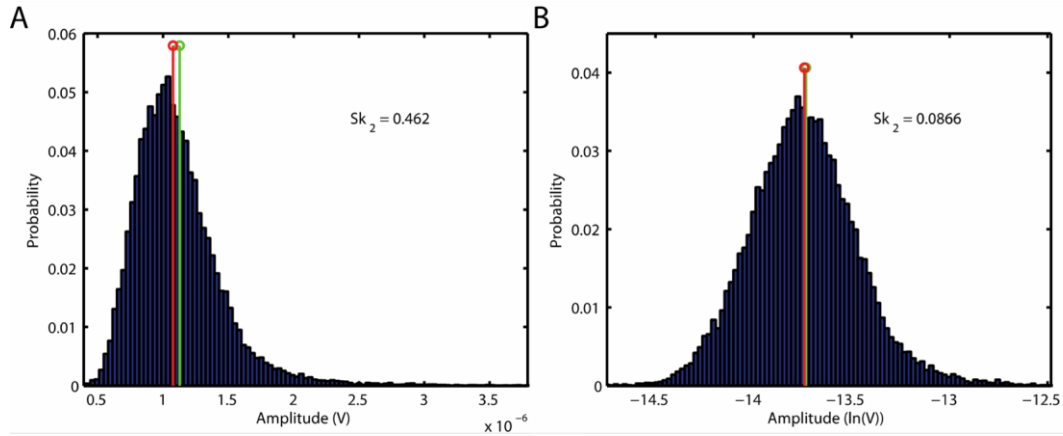
Figure 3.2. Histogram of (A) raw and (B) log transformed amplitudes. Green and red stem plots indicate the mean and median of each distribution, respectively. The Pearson 2 skewness coefficients, Sk$_2$, are indicated for each distribution.

Because the spectral power in a given ECoG signal typically falls off at a rate inversely proportional to the frequency, the log transformed features $c(t)$ were further normalized as follows:

$$x(t)_{i,j} = \frac{c(t)_{i,j} - \bar{c}_{i,j}}{\sigma_{i,j}} \qquad (3.5)$$

$$\sigma_{i,j} = \sqrt{\frac{1}{T}\Sigma\left(c(t)_{i,j} - \bar{c}_{i,j}\right)^2} \qquad (3.6)$$

Features were normalized by subtracting baseline values, $\bar{c}$, from each feature, $c(t)$, and then dividing by σ, the RMS value of the difference between the feature and its baseline. Normalization was performed to prevent a possible bias in our BCI decoding models toward lower frequency features with higher powers. In addition, the physiological significance of the magnitude of weights in the decoding model should be easier to interpret than if features were not normalized.

At the beginning of each recording day, 5 minutes of baseline activity was recorded with no behavioral task displayed to the monkey. The baseline and scaling terms, $\bar{c}$ and σ, were computed over this time period and held constant for the rest of the recording day. This static baseline period is in contrast with past strategies in our lab of using a running baseline during BCI experiments. The rationale for the static baseline strategy used in the current experiments stems from stark differences that can arise in a running baseline between periods of active performance of a BCI task and rest. These differences will be examined in further detail in Chapter 7.

3.6 BCI Control Scheme

The BCI experiments examined in this dissertation utilized ECoG derived control signals to move a spherical cursor in two dimensions on a computer screen. Although one of the main goals of this work was to demonstrate the feasibility of a force-based BCI for interactions with real-world objects and forces, translating ECoG signals to movement within a virtual environment allowed us complete control of the virtual world's physics, and thus our general BCI scheme could be tested within a wider variety of BCI settings. In addition, examining BCI control of a virtual device allowed us to implement a learning curve more individually tailored to the progress of each subject.

A schematic for the closed-loop processing from raw ECoG signals to a cursor

position on the screen is depicted in Figure 3.3. First, raw signals, $s(t)$, are band-pass filtered, rectified, and low-pass filtered in TDT hardware to produce an amplitude estimate on each channel with five frequency bands. These amplitude features are relayed to a host computer and then normalized as in Section 3.5 to the spectral features, $x(t)$. These control features are vectorized at each time point to produce a 160 x 1 feature vector, $X(t)$. In other words, the i^{th} channel and j^{th} frequency band are mapped to the $(32*(j-1)+i)^{th}$ feature:

$$x_{i,j}(t) \rightarrow X_{32(j-1)+i}(t) \qquad (3.7)$$

An instantaneous control signal, $y(t)$, is produced by a linear transformation of X(t) by a
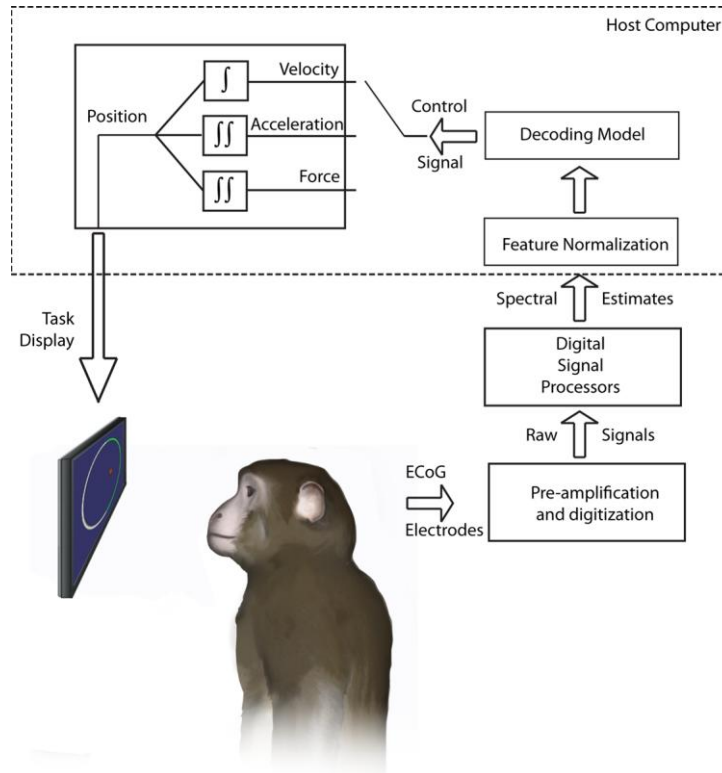
Figure 3.3. BCI Control Schematic. Raw signals are first recorded from the monkey's epidural ECoG electrodes, amplified, and digitized. Digital signal processors produce spectral estimates for each channel of the amplitude within five frequency bands using an envelope detection method. Spectral estimates are sent to a host computer for feature normalization. A decoding model transforms the normalized features into a two-dimensional control signal, which is subsequently transformed into either a velocity, acceleration, or force variable depending on the experiment the monkey is performing (experiments and control schemes lasted several weeks before switching). The kinematic or dynamic variable is then integrated appropriately to a cursor position, which is displayed on an LCD screen along with the rest of the behavioral task.

27

decoding matrix, w, and a constant offset vector, b_0:

$$y(t) = w^T X(t) + b_0 \qquad (3.8)$$

$y(t)$ and b_0 are both two-dimensional, and w is a 160 x 2 dimensional matrix.

To examine the feasibility of a dynamics based control signal compared with the "gold standard" velocity control signal that has typically been used in movement-focused BCIs, we applied the control signal, $y(t)$, to a number of kinematic and dynamic variables before integrating these appropriately to the displayed cursor position on the screen. The highest level variable examined in these experiments was the transformation from control signal to the two-dimensional cursor velocity, $v(t)$:

$$v(t) = g_v \cdot y(t) \qquad (3.9)$$

where g_v is a velocity gain constant. The position, $p(t)$, is simply calculated for the continuous case as the integral:

$$p(t) = \int_0^t v(t)dt + p(0) \qquad (3.10)$$

Since the position is calculated on a host computer at distinct time steps, the discrete version is more applicable:

$$p(t + dt) = p(t) + v(t)dt \quad (3.11)$$

The next kinematic parameter explored was the acceleration of the cursor, $a(t)$, with acceleration gain, g_a:

$$a(t) = g_a \cdot y(t) \qquad (3.12)$$

A similar discrete integration results in the cursor velocity in an undamped system:

$$v(t + dt) = v(t) + a(t)dt \quad (3.13)$$

The velocity term is then integrated as in Eqn. (3.11) to produce the visual position feedback. Technically, equations (3.12) and (3.13) represent brain control of an object's acceleration regardless of any external forces acting on it. Through preliminary trials of this control system with human subjects controlling cursor acceleration with a joystick's position, however, this undamped system is very difficult to control in practice. This control scheme often resulted in oscillatory behavior as subjects would fail to decelerate sufficiently to stop at a target, overshoot the target, overcompensate in accelerating toward the target, overshoot the target again, etc. To make this control scheme more

tractable to the subject and also provide a mean of gradually training toward "pure" acceleration control, a damping term was added to Eqn. (3.12):

$$a(t) = g_a \cdot y(t) - k \cdot v(t) \quad (3.14)$$

where k is a viscous damping constant, producing a force proportional in magnitude and opposite in direction to the cursor's velocity. Thus, in the absence of a driving control signal, the cursor velocity will exponentially go to zero over time. If k is large enough such that $k \cdot dt \approx 1$, Eqn. (3.13) reduces to Eqn. (3.9) and the control scheme resembles velocity control. As the damping constant is decreased, control of the cursor becomes more and more "acceleration-like". In this manner, we can decrease the learning curve and ease the transition to acceleration control for the subject. Since we have included an external force, Eqn (3.14) technically represents dynamic control of the cursor. However, since we have only included a simple viscosity term that shifts the degree of control between velocity-like and acceleration-like without accounting for the mass of the cursor or any objects it may act on, we will refer to this control scheme as acceleration control throughout this dissertation.

Using a simple expansion of Eqn (3.14), the last brain control type examined in this work was dynamic (*i.e.* force) control. By definition, the acceleration of an object is the vector sum of forces acting on the object divided by its mass. Previously, we assumed a constant mass that could be absorbed into the gain term, g_a. For the set of experiments explained in more detail in Chapter 6, however, the total mass under control of the monkey, m_{total}, was allowed to vary. In addition, this combined mass was also subjected to a downward force, G, simulating gravity. Thus, the relationship from the ECoG control signal, $y(t)$, to cursor acceleration becomes:

$$a(t) = \frac{g_a \cdot y(t) - k \cdot v(t) - G}{m_{total}} \quad (3.15)$$

Two notes are important to keep in mind from this relationship. First, G is a constant force that acts only on one dimension, while the inertial control force and viscous force act in both dimensions. Second, G is a static force, while the first two force terms that are mainly responsible for producing the desired movement profile are time varying. Thus, the monkey must not only account for both static and dynamic forces, but also

account for these differently between the two dimensions of the task. Similar to the acceleration control scheme, g_a, k, G, and m_{total} all represent settings that can be adjusted by the experimenter. These variables can not only be used to aid in gradually training the subject with a force-based control scheme, but also to tease apart the static and dynamic force components needed to interact in this virtual environment.

After translating the ECoG signal to the desired control signal and integrating appropriately on the host computer, the cursor position and the rest of the virtual task scene is relayed to the monkey's monitor for display. The monkey observes the visual feedback of his cursor's movement, attempts to adjust his cortical output to produce a desired movement toward completing the BCI task, and the cycle repeats.

3.7 Behavioral Tasks

Three behavioral tasks were used during the course of the experiments examined in this thesis: a radial-choice task, a classic center-out task, and a novel mass-grab task. The radial-choice task is a ballistic control task useful for assessing directional decoding information and theoretical data transfer rates for a discrete classification BCI [18], [19] as well as for initial training on a particular control modality (i.e. joystick or brain control). The center-out task is a well-known behavioral task often used in both non-human primate and human reaching studies to elucidate possible movement parameter representations in neural data [58–60] as well as act as a test bed for BCI applications [11], [46], [61]. As it used in this thesis, the center-out task requires more precision and fine control than the radial-choice task, and will be used to assess BCI performance as would be applied for reaching movements. Finally, the mass-grab is an extension of the center-out task designed for examination of force-based control in which the subject must account for the masses of the cursor and targets as well as any external forces acting on them. Between blocks of these active behavioral tasks, rest blocks were interleaved to assess active BCI usage and resting states.

3.7.1 Radial-choice Task

First, each subject learned a radial choice task in which a highlighted arc on a

circular ring served as the target for the task (see Figure 3.4). At the beginning of each trial, the cursor was automatically reset to the center of the ring. One hundred milliseconds later, an arc of the ring was highlighted to signify the desired target. For the 2-D brain control experiments in this thesis, the ring was divided into four possible target arcs for a chance rate of 25%, while Figure 3.4 depicts a version with eight targets used for joystick training. The subject would then move the cursor toward the highlighted portion of the ring. Moving the cursor through the correct arc of the ring resulted in a water reward followed by a new trial. Moving through a non-highlighted portion ended the trial with no reward or opportunity to correct the mistake before a new trial would begin. Failure to move through any segment of the ring within 5 seconds resulted in a time-out error and a new trial. Peripheral targets were chosen in a pseudorandom block design such that a predefined number of trial repetitions (8 trials for joystick experiments, 12-16 for BCI experiments) to each target must have been successfully completed before a block was complete. Control mappings for this task
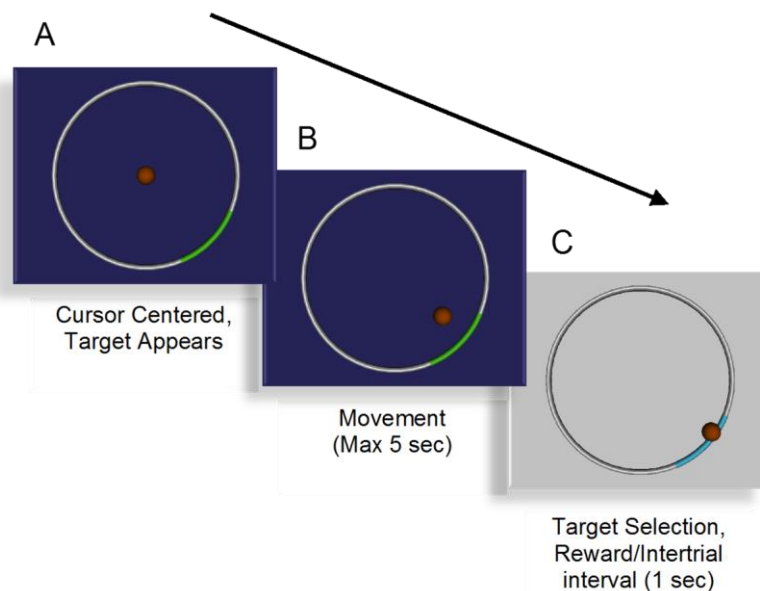
Figure 3.4. Two-dimensional radial-choice task. A) At the start of the trial, the cursor is moved to the center and a highlighted peripheral targets appears. B) The monkey then has five seconds to select the correct target. C) Once a target is selected or the maximum movement time has been reached, the trial is over, a one second inter-trial interval occurs, and the monkey is rewarded if the correct target was chosen.

were always to cursor velocity (i.e. joystick position to velocity or brain signal to cursor velocity). Since this task does not require deceleration to stop on a particular target and was to be primarily used for training and BCI learning assessment, the use of acceleration control was deemed unnecessary for these experiments.

Initial training under brain-control was assisted through the use of a cursor bias velocity that was added to the brain-derived velocity signal. This bias velocity always pointed in the direction of the target. During early training sessions, the magnitude of the bias signal was set relatively high compared to the brain control signal such that the cursor would drift toward the target, but the monkey would still receive visual feedback about the effect of his brain activity on the movement of the cursor. The bias velocity was also used such that the cursor would eventually (but slowly) drift to the target and provide a reward to keep the monkey from becoming frustrated, but the monkey also had the opportunity to identify a neural state that made the cursor travel in the direction of the target faster and therefore be rewarded more often. As the monkey identified productive cortical modulations over days, the bias velocity magnitude was gradually decreased until the cursor was under complete brain control.

3.7.2 Center-out Task

After becoming proficient at the radial choice task and demonstrating unassisted performance well above chance, each subject learned to use the joystick- or brain-controlled cursor to complete a classic center-out task [46], [58], [60–62] (see Figure 3.5 A-D). Subjects would first move the virtual cursor to a center target sphere and maintain contact with it for a predetermined amount of time (Hold A period). After satisfying the Hold A time, the center target would disappear, and a peripheral target sphere would appear at one of a set of possible locations. For joystick tasks, sixteen target locations equally spaced around the center of the virtual workspace were possible, while for BCI tasks, four targets were possible (two each on the horizontal and vertical axes). The subject would move the cursor to the peripheral target and hold it there for another period of time (Hold B), and a water reward would be administered to end the trial. Failure to complete either hold time, to reach the center target within 10 seconds

of its appearance to start the trial, or to reach the peripheral target within a designated period from its appearance (Max Movement Time) would end the trial with no reward. Hold A and B times were gradually increased as the monkey improved at slowing down the cursor on approach to a target. The Max Movement Time was also gradually decreased over the course of days to encourage the monkey to make smooth accelerations and decelerations of the cursor similar to normal reaching movement profiles. Peripheral targets in each task block were chosen in the same pseudorandom manner as for the radial-choice task. Monkeys utilized control mappings to position, velocity, and acceleration during joystick experiments, and mappings to velocity and acceleration during brain control trials.
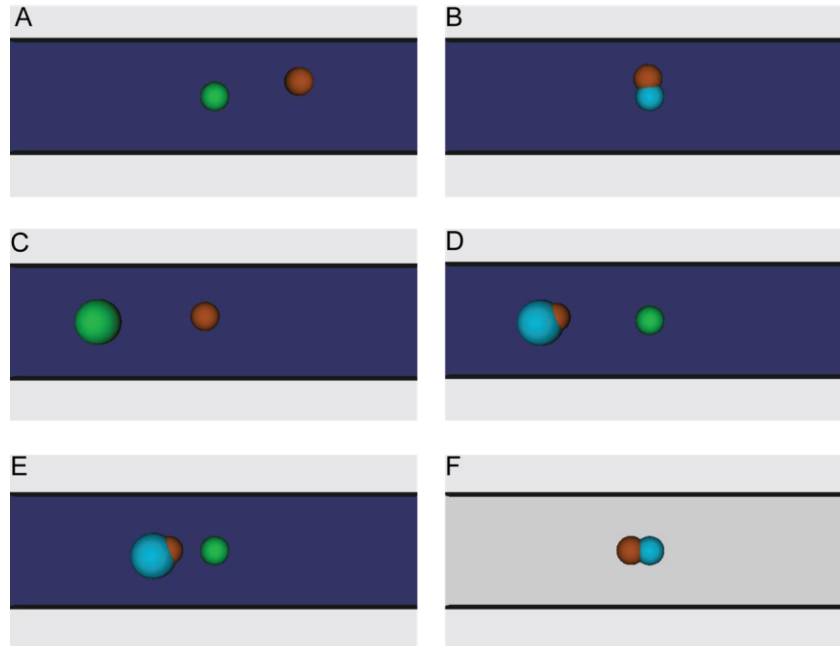
Figure 3.5. Two-dimensional center-out and mass-grab tasks. A) The monkey moves a cursor (orange sphere) toward the center target (green sphere) at the start of a trial. B) Upon capturing the center target by maintaining contact for a hold period (Hold A), the center target disappears and a peripheral target appears (panel C). D) The monkey moves the cursor to capture the peripheral target (Hold B). Panels A-D constitute the standard center-out task used in all joystick tasks and initial BCI tasks. E) For mass grab trials under brain control, upon completion of the Hold B time in panel D, the cursor "grabs" the peripheral target and must be returned to the center for a third hold time (Hold C, panel F) to complete the trial. During mass grab trials, the cursor and peripheral targets assume different masses such that different virtual forces must be applied on the outgoing (panels A-D, cursor mass alone) and return (E-F, cursor+target mass) paths.

3.7.3 Mass-grab Task

For further examination of the efficacy of a force-based BCI, monkeys were challenged with a modified center-out task (i.e. "mass-grab" task) that incorporated simple, simulated real-world physics analogous to those one might experience with everyday reaching tasks (see Figure 3.5). For this task, both the cursor and peripheral targets were assigned virtual masses that correlated visually with the size of each sphere. The first part of the trial proceeded exactly as the previously described center-out trials with a force-based control signal producing a cursor acceleration magnitude inversely proportional to the cursor's mass (see Eqn. (3.15)). Upon successful completion of the hold time at the peripheral target (Hold B), the cursor "grabbed" the peripheral target. The monkey's force control signal then had to accelerate the combined mass of the cursor and peripheral target back toward the center target again for a final hold period (Hold C). For each trial, one of three mass sizes was randomly selected for the peripheral targets while the cursor mass remained constant between trials.

In addition to the task complexity introduced by using different mass terms for the outgoing and return cursor trajectories as well as variable peripheral target masses, a non-zero gravity term in Eqn. (3.15) was included such that the monkey had to maintain a static vertical force in order to maintain a horizontal cursor path. A virtual floor and ceiling with error penalties for contact were also used to encourage horizontal trajectories during horizontal target trials. With the use of variable peripheral masses and virtual arena constraints, the monkey was required to adjust this static force level between the outgoing and return path as well as between trials for different peripheral masses. Again, a pseudorandom block design was used, and four trials to each target mass/target combination had to be successfully completed before finishing a block (48 correct trials per block).

3.7.4 Rest Trials

At the end of a block of any of the aforementioned task types, a block of "rest trials" was initiated in which the background of the subject's virtual arena remained constant, but visual feedback of the ECoG controlled cursor was removed and no

rewards were provided. With no visual feedback or reward motivation, these trials were intended to disengage the subject from actively modulating his ECoG signal for BCI control. Since no task related events were present during these trials, individual trial lengths were fixed at 3 seconds. No indication of transition from rest trial to rest trial was presented. After 20 trials (~60 seconds), the rest block ended and the beginning of a new task block was signaled by the appearance of the cursor and target ring/sphere.

3.8 Experimental Time Course

Following surgical implantation of the ECoG arrays, each monkey was allowed approximately one week for recovery before beginning joystick experiments. Since each monkey had been training on acceleration-control with a joystick just prior to surgery, joystick experiments with ECoG recordings began with acceleration-control, followed by velocity-control, and the finally position-control mapping. Monkey K regained prior proficiency levels and completed joystick mapping experiments in 3-4 days per mapping. Monkey P was slower to recover from surgery and initially seemed reluctant to perform the behavioral task. Thus, joystick acceleration-control mapping took monkey P 10 days to achieve pre-surgical proficiency at the task, while he regained proficiency at the velocity and position mappings in 4 days each.

Upon completing the joystick mapping experiments, each monkey began brain controlled radial-choice task experiments using velocity-control with signals from the left hemisphere. Monkey K demonstrated unassisted proficiency consistently above chance after 6 training days, and moved on to the center-out task after approximately 3½ weeks. Monkey P required two weeks before demonstrating unassisted performance above chance on the radial-choice task. However, this monkey was also observed to frequently take breaks within a task block which may have confounded the adaptive decoding algorithm described in Chapter 5 and may have also introduced a delay into his learning curve. He moved onto the center-out task after three weeks of total training on the radial-choice task. Moving onto the center-out task also aided the adaptive training algorithm for monkey P as the subject must move the cursor to the center target to start a trial, thus offering better insight into which trials the monkey was

actively attempting to complete.

Upon moving onto the center-out task, monkey K performed the center-out task with velocity-control for one week before moving on to acceleration-control for three weeks, after which he performed the mass-grab task with force-control for two weeks. Monkey P spent three weeks on the center-out task with velocity-control and two weeks with acceleration control, and finished on the mass-grab task for four weeks. For both monkeys during mass-grab training weeks, each training day began with a shorter block of acceleration-control center-out trials to serve as a control with no gravity forces or additional masses for comparison with mass-grab trials.

After finishing the initial round of control schemes with his left hemisphere, monkey K repeated all brain-control experiments with signals from only his right hemisphere array. The second round of brain-control experiments proceeded much faster than the first as monkey K was completing the radial-choice task above chance performance with no assistance by the end of the first day. He spent one week on the radial-choice task, one week with velocity-control during center-out trials, one week with acceleration-control during center-out trials, and three weeks on mass-grab experiments (additional iterations of gain, masses, and gravity were investigated during this time period compared to the first round).

4 ECoG Correlates of Visuomotor Transformation: Open-Loop Analyses

4.1 Introduction

In general, experiments geared toward the development of efferent or motor BCI systems come in one of two common flavors. On one hand, the neurophysiology of movement is investigated to see how well stereotyped movements can be reproduced offline using only neural signals (*i.e.* "open-loop" analysis). On the other hand, experimenters may investigate how well a subject can control a prosthetic device in practice using only neural signals regardless of any physical movements that may be correlated with those signals (*i.e.* "closed-loop" analysis). These two experimental paradigms may also be combined by examining the adaptation of neurophysiological phenomena found in open-loop experiments when directly applied to closed-loop control of a prosthetic device.

In this chapter, we investigate possible ECoG correlates of visuomotor transformation, *i.e.* how processing chains in the brain transform a high-level goal oriented command ("I want to place this cup on that shelf.") to the lower-level commands necessary to realize the goal ("Activate anterior deltoid and biceps, then triceps…"). Past open-loop experiments have often used a center-out paradigm with non-human primates to identify correlations between movement parameters during free arm reaching motions (*i.e.* hand position, velocity, acceleration, force, EMG) and neural activity (single unit activity, LFPs, ECoG spectra, etc.). For example, in past experiments in our lab, monkeys were first trained in a virtual reality simulator in which infrared markers placed on the arm were used to track the 3-D position of the arm and hand in space and map that position directly to the position of a 3-D computer cursor on a stereoscopic monitor. After mastering these tasks, the monkeys would then undergo surgery for implantation of an acute recording chamber, and the animals would

continue to perform these tasks while neural activity was recorded. In this scenario, the visual goal-oriented effector or "tool" of the task (*i.e.* the virtual cursor) is directly mapped from the physical effector for the task (*i.e.* the hand) with a certain gain term.

Building upon these studies, some experiments have examined the roles of different cortical areas in movement planning and execution when physical movement and perception of those movements are not as tightly coupled. As mentioned in Chapter 2, Schwartz *et al.* studied the differential representation of motor perception and execution in motor and premotor cortices by changing the gain of a visually displayed cursor in one dimension while the monkey performed a circle or ellipse drawing task [33]. The findings from this study suggested that pre-motor cortex is more closely correlated with a movement plan or the perceived end result of movement while primary motor cortex is more concerned with the actual physical movements that take place. However, as in other studies, the relationship between physical movements and the perceived end-effector of the task is still very closely related, *i.e.* one differs from the other by only a gain term.

In the experiments presented in this chapter, we sought to dissociate the relationship between physical movement and movement of the goal-oriented effector of a task to a greater degree. Toward this goal, monkeys were trained on a set of center-out tasks in 2-D space similar to those used in our prior research. However, whereas in our prior research the monkeys used free arm movements in space, the current experiment subjects completed tasks through movement of a joystick. The use of a joystick toward the first aim in this thesis allowed us to dissociate between physical arm movements and cursor movements on a screen, and assess a given area's possible relationship with either physical hand or perceived end-effector movement parameters.

4.2 Methods

4.2.1 Joystick Mapping

All open-loop experiments in this chapter employed the use of a joystick in

order to translate physical hand movements into cursor movements on the monkey's screen. The monkey's left hand was restrained while the right hand was free to operate the joystick. As described in Chapter 3, the 2-D position of the joystick, *y(t)*, was remapped to either the 2-D position (*p(t)*), velocity (*v(t)*), or acceleration (*a(t)*) of the cursor displayed to the monkey on a computer monitor as follows:

$$p(t) = g_p \cdot y(t) \qquad (4.1)$$

$$v(t) = g_v \cdot y(t) \qquad (4.2)$$

$$a(t) = g_a \cdot y(t) - k \cdot v(t) \quad (4.3)$$

where g is a gain term for the corresponding mapping modality and k is a viscous damping term used in the position-to-acceleration mapping. Acceleration and velocity were integrated to position using the following discrete update equations:

$$v(t + dt) = v(t) + a(t)dt \quad (4.4)$$

$$p(t + dt) = p(t) + v(t)dt \quad (4.5)$$

Only one joystick mapping was used on a given experiment day. Monkeys were fully trained on each joystick mapping prior to surgical implantation of the ECoG arrays. During training on the acceleration mapping, the damping term was gradually decreased and the gain term was increased to more closely resemble pure acceleration-based control rather than velocity-based control. However, to make the task more feasible for the monkey, minimal damping was still included under this control scheme. This inclusion of damping likely more closely resembles real-life control scenarios as it is very rare to encounter a completely frictionless or undamped system. For example, the viscoelastic properties of muscles factor considerably into the joint torques they produce.

4.2.2 Behavioral Task

Monkeys K and P performed a standard 2-D center-out task as described in section 3.7.2 with the use of a joystick. Sixteen radially spaced targets were possible for

selection in a pseudorandom block design. For final goal experiments from which psychophysical and neural data in this chapter was analyzed, Hold A and B times were randomly selected from between 700-800 ms for each trial. Maximum movement times were set at 750 ms for position-to-position mapping experiments, 1100 ms for position-to-velocity mapping, and 1500 ms for position-to-acceleration mappings. These hold and movement time values were determined from pre-implantation training on the mapping tasks at which point the monkeys did not complete the task consistently with longer hold times or shorter movement times.

4.2.3 Joystick/Cursor Psychophysics

Correct trials that met the above hold and movement time criteria were included for further analyses. Since the absolute speed component of high gamma tuning using micro-ECoG has been shown to dominate any directional modulation [63], trials within a given joystick mapping were collapsed across all target directions. The joystick and cursor position data collected during the experiments were low-pass filtered below 10 Hz. Joystick/cursor velocity estimates were calculated by taking the discrete derivative of the low-pass filtered position signals and low-pass filtering the result. Accelerations were calculated in the same manner from the corresponding velocity signals.

For binned analyses, trial data was binned based on the velocity of the displayed cursor in two different manners. For the first method, the movement period was defined by finding the maximum cursor velocity projected along a straight line from the center to the target during the task-defined movement period (after Hold A and before the beginning of Hold B). The start and end of the binned movement period were then defined as the last time point before the maximum velocity and the first time point after, respectively, at which the cursor velocity was above 15% of the maximum velocity. This movement period was then binned by downsampling the joystick data into 40 equal time bins. Using the same bin width, 30 bin samples preceding the start of this period and 30 bin samples following this period were included for analysis. Although this binning method removes some temporal variability that could be introduced by delayed reaction times on the part of the monkey, it could also artificially

reduce the variability in velocity at the beginning and end points of the movement period. Because of this consideration, a second binning method was employed for comparison. The start of the binned movement period was defined this time as the first time at which the cursor position projected along a straight line from center to target was greater than 15% of the distance from center to target. The end of this period was defined as the first point at which the projected cursor distance from the target was less than 15% of the center-to-target distance. Binning of this redefined movement period then proceeded as before. Using either binning method, joystick and cursor movement parameters during trials across all targets for a given experimental day (one joystick mapping) were averaged together for comparison with ECoG data. As no appreciable difference in the results obtained using both methods emerged, results in this chapter are limited to the first method for simplicity. However, the second method will be retained for BCI binned analyses for the reasons mentioned above.

For non-binned analyses, several simple performance metrics were calculated. First, the average Time to Acquire Target (TAT) was calculated as the average time between the end of the Hold A period and the beginning of the Hold B period. While this metric gives an idea of how fast a subject could touch a peripheral target after it was presented, it does not give any additional information about the cursor trajectory or level of control. A second metric, $R_{max/B}$ was calculated as the ratio of the maximum speed during the movement period to the speed at the end of the Hold B period:

$$R_{max/B} = \frac{|v_{max}|}{|v_{Hold\ B}|} \qquad (4.6)$$

This ratio provides a picture of the subject's ability to accelerate and stop at a target. Low values near 1 indicate several possibilities, either: 1) the subject accelerates the cursor to a fast speed but does not slow down substantially and speeds through the target, or 2) the subject accelerates the cursor at a relatively slow, consistent speed in order to meet the Hold B criteria. A sufficiently long Hold B criterion precludes the first possibility, while a sufficiently short Max Movement Time criterion lessens the likelihood of the second possibility. Since these two speeds are not necessarily independent (*i.e.* a trial with a higher maximum velocity is likely to have a higher Hold B

velocity as it is requires more deceleration to come to a complete stop), the mean ratio statistic was calculated as the mean of sample ratios calculated on a trial-by-trial basis as opposed to calculating the ratio of the means.

Finally, several metrics were used to quantify the cursor trajectory itself. A normalized path length metric, d_{path}, was calculated as the path-dependent distance traveled by the cursor over the movement period by integrating the absolute speed and dividing by the distance, d_{line}, between the start and end cursor positions over the same period:

$$d_{path} = \frac{\int_{move}|v(t)|dt}{d_{line}} \qquad (4.7)$$

As the shortest path from the center target to the peripheral target is a straight line, path lengths greater than this value indicate the degree of deviation from a straight line trajectory. Along these lines, a dot product metric, $dp_{vel-targ}$, between the instantaneous cursor velocity vector and the cursor-to-target vector was calculated:

$$dp_{vel-targ} = \frac{1}{T_{move}} \int_{T_{move}} \frac{\vec{v}(t) \cdot \vec{u}(t)}{\|\vec{v}(t)\| \|\vec{u}(t)\|} dt \qquad (4.8)$$

where $\vec{v}(t)$ is the instantaneous velocity vector and $\vec{u}(t)$ is the instantaneous vector pointing from the cursor to the target. This metric has an advantage in that it can assess the subject's ability to correct the cursor's trajectory as a function of time.

4.2.4 ECoG Recordings and Signal Processing

Epidural ECoG signals from the left hemisphere array of each monkey (contralateral to the joystick hand) were recorded during these experiments. Spectral estimates for the amplitude within the 75-105 Hz band on each channel were obtained and normalized to units of standard deviations in the same manner as in the BCI control scheme described in Chapter 3.5. This high gamma frequency range was chosen as it has previously shown correlation with physical movement parameters in LFP experiments from our lab [64] as well as a sound control feature for ECoG BCI control [18], [65]. Since we were interested in the actual time course of ECoG correlations with movement parameters and not how a particular decoding scheme might work in

practice, the "filtfilt" function in Matlab was used for band-pass and low-pass filtering of neural and movement data in order to remove filtering delays. For binned analyses, binning time points derived from the cursor psychophysics were used to sample the normalized 75-105 Hz spectral amplitude on a given channel and trial, and this data was then averaged across correct trials to all targets on a given day and channel. For non-binned analyses, neural and movement data were downsampled to 50 Hz.

4.3 Results

4.3.1 Psychophysics and Performance

Figure 4.1 shows the cursor paths to all 16 targets for each monkey and experimental day used in further analyses (one mapping per day) along with the number of trials included. As seen from the figure, the trajectories appear consistent across each mapping experiment, except the position-to-position mapping day for monkey P
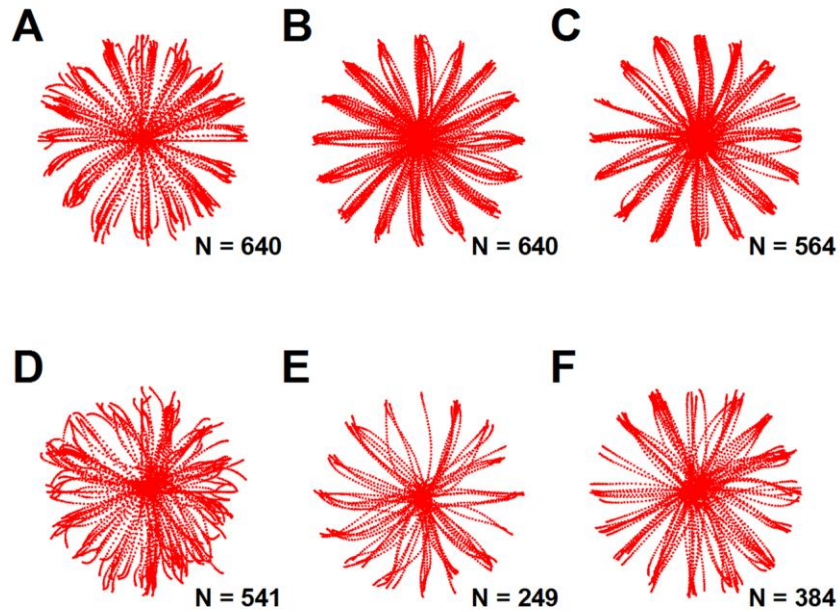
Figure 4.1. Cursor trajectories for joystick mapping experiments. Panels A-C show cursor trajectories to all 16 targets for monkey K on position-to-position, position-to-velocity, and position-to-acceleration mappings, respectively. Panels D-F show cursor trajectories for similar mapping experiments performed by monkey P. Total trial numbers for a given mapping experiment day are indicated by N. (For plotting purposes, trials were downsampled by a factor of 3.)

appeared to show the most variability in cursor path.

An example of response times for one monkey (K) to react to the appearance of the peripheral target within a position-to-position mapping experiment is illustrated in Figure 4.2. The average response time upon appearance of the peripheral target for the cursor to leave the area previously occupied by the center target was approximately 350 ms. Basing the response time on the joystick velocity reaching 15% of the trial's maximum velocity resulted in a response time of 300 (±42) ms. This response time for the displayed cursor increased with the other two mappings due to a delay associated with the integration(s) from joystick position to cursor position. However, examination of the joystick speed under these mappings revealed similar response times, indicating that visual response to and physical initiation of the task did not differ significantly between mappings.
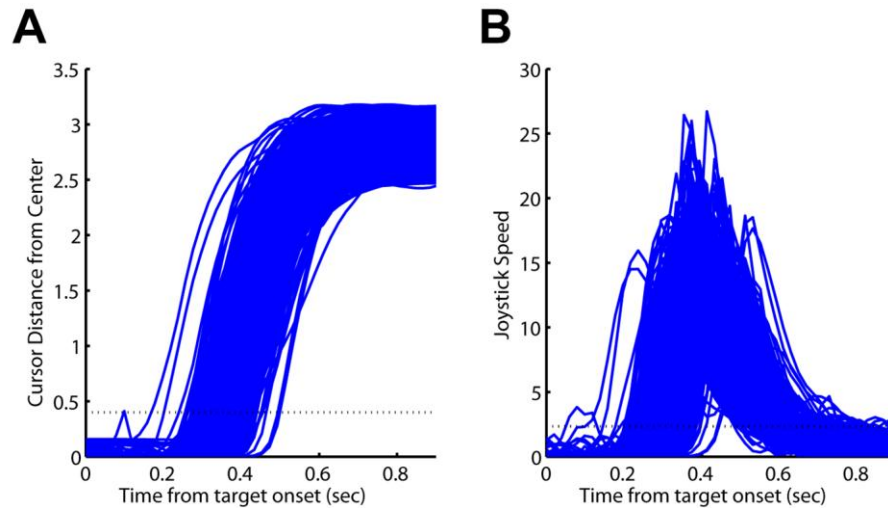
Figure 4.2. Position-to-position mapping task response times. Traces depict the trial-by-trial (A) cursor distance from the center target in the direction of the displayed target and (B) joystick velocity as a function of time. At t=0, the peripheral target is displayed to the monkey. The dotted line indicates the threshold used for calculating the response time for each trial (leaving center target for (A) or speed > 0.15 max speed for (B)).

Table 4.1 summarizes the "reaching" characteristics and performance of each monkey and mapping experiment. As seen from the statistics, the TAT increased across the position, velocity, and acceleration mapping tasks due to the complexity of the required joystick movements. For example, the position mapping task required one outward movement with the joystick, the velocity task required two movements (out

and back to center), and the acceleration task required 3-4 movement phases (out to accelerate, back to center and in the opposite direction to decelerate (possibly two distinct phases), and back to center again once). Conversely, d_{path}, was similar across tasks although statistically different on average with the large number of trials analyzed. The dot product metric, $dp_{vel-targ}$, between target direction and cursor velocity was appreciably lower for the acceleration task than the other two mappings. With the exception of the position mapping experiment for monkey P, Hold B speeds were similar across tasks while maximum speeds consistently decreased across tasks (position mappings with the highest speeds). The resulting speed ratios were highest for the position mapping. However, this statistic is especially sensitive to the denominator (especially small values) so assigning too much value to small differences between the velocity and acceleration mappings might be premature without considering the overall context of the other statistics.

Table 4.1. Summary of psychophysical results from joystick mapping experiments. Values in parentheses represent standard deviations. The ratio, $R_{max/B}$, was calculated from the trial-by-trial ratio of the maximum speed to the Hold B speed, and standard deviations were calculated from this distribution as well.

		Joystick-to-Cursor Mapping		
		Pos-to-Pos	Pos-to-Vel	Pos-to-Acc
TAT (sec)	K	537 (67)	967 (49)	1255 (78)
	P	541 (98)	1033 (71)	1330 (88)
d_{path}	K	1.030 (0.018)	1.036 (0.037)	1.064 (0.052)
	P	1.143 (0.143)	1.054 (0.045)	1.044 (0.030)
dp_{vel_targ}	K	0.85 (0.14)	0.81 (0.10)	0.77 (0.09)
	P	0.88 (0.11)	0.89 (0.10)	0.79 (0.13)
Max Speed (units/s)	K	15.70 (2.32)	4.56 (0.40)	3.99 (0.34)
	P	12.73 (2.97)	4.16 (0.37)	3.65 (0.31)
Hold B Speed (units/s)	K	0.41 (0.42)	0.38 (0.59)	0.53 (0.33)
	P	1.69 (1.20)	0.67 (0.55)	0.58 (0.40)
$R_{max/B}$	K	19.0 (2.6)	3.50 (0.14)	4.80 (0.53)
	P	5.35 (0.53)	4.00 (0.43)	4.08 (0.20)

While these statistics are not overly exciting within their own context, they do provide information about how the individual mapping tasks compare in difficulty. Additionally, they provide a standard for comparison for the brain-control versions of the velocity and acceleration tasks in Chapter 5.

4.3.2 Task-related ECoG Signatures

As an attempt to deduce any preferential correlation of neural features with either physical or end-effector movement characteristics, the mean joystick and cursor movement parameters (position, velocity, and acceleration) for each mapping experiment were examined alongside the corresponding mean 75-105 Hz neural profiles as shown in Figures 4.3 and 4.4 (monkeys K and P, respectively). The profiles in these figures were aligned based on the shape of the cursor velocity profile during each trial. Thus, the general cursor movement profiles were held roughly constant across mapping experiments and resembled those seen in Figures 4.3A and 4.4A except scaled horizontally with time. Given this setup, if a channel's high gamma activity is coupled to a perceived movement characteristic of the cursor (*i.e.* cursor position, velocity, or acceleration), its neural profile should remain unchanged across joystick mappings. Conversely, if a neural feature follows a physical movement feature of the monkey's joystick/hand that reveals itself in panel D of either of those figures, then its profile should change shape over the other two mappings and continue to follow its correlated joystick variable in those tasks.

As seen in panels A and D of Figures 4.3-4.4, a significant number of channels on the anterior half of each monkey's array appeared to correlate well with either cursor or joystick velocity. Upon examining the results from the other mapping tasks, these neural profiles appear to follow the physical speed characteristics of the joystick. An important consideration to note when examining these profiles is that previous experiments in non-human primate experiments have shown correlations between high gamma ECoG activity and the absolute speed of reaching movements, but have argued against directional tuning of this activity (*i.e.* true velocity tuning) [63]. Thus, while
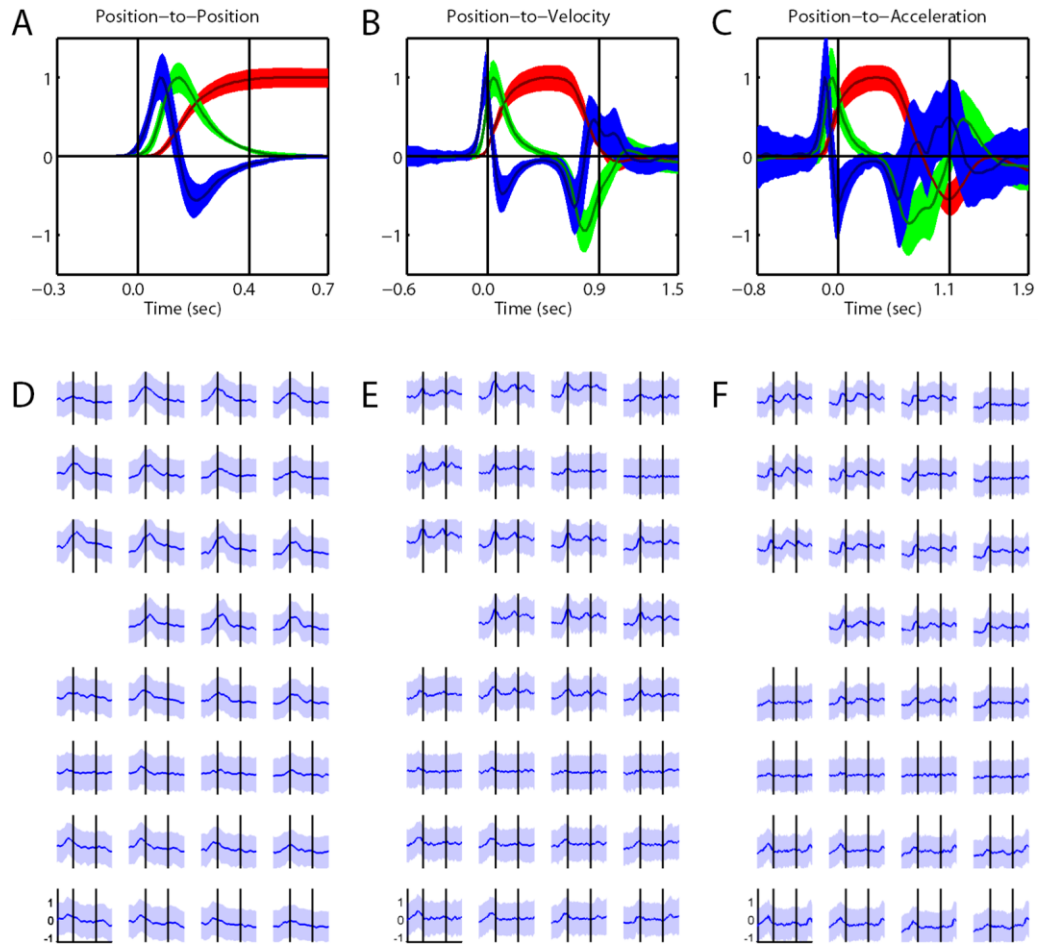
Figure 4.3. Joystick movement and neural profiles binned by cursor velocity across mapping experiments, monkey K. <u>Top Panels</u>: Red, blue, and green traces indicate joystick position, velocity, and acceleration, respectively, projected along a vector from center to target. A-C show mean joystick parameters during the position, velocity, and acceleration mapping tasks, respectively. Vertical black bars indicate the edges of the binned movement period defined by cursor speeds at 15% of maximum cursor speed. Time values indicate rounded mean movement period duration. Actual time span distributions were 0.40 ± 0.01, 0.89 ± 0.07, and 1.13 ± 0.14 seconds for A-C. <u>Bottom Panels</u>: D-F show normalized mean ECoG amplitude profiles in the 75-105 Hz band on each channel in the array's spatial arrangement (top rows - anterior, left columns - lateral) for the corresponding mapping tasks in A-C. Black vertical bars correspond to the same time markers in A-C. Shaded regions on all plots indicate ± 1 standard deviation.
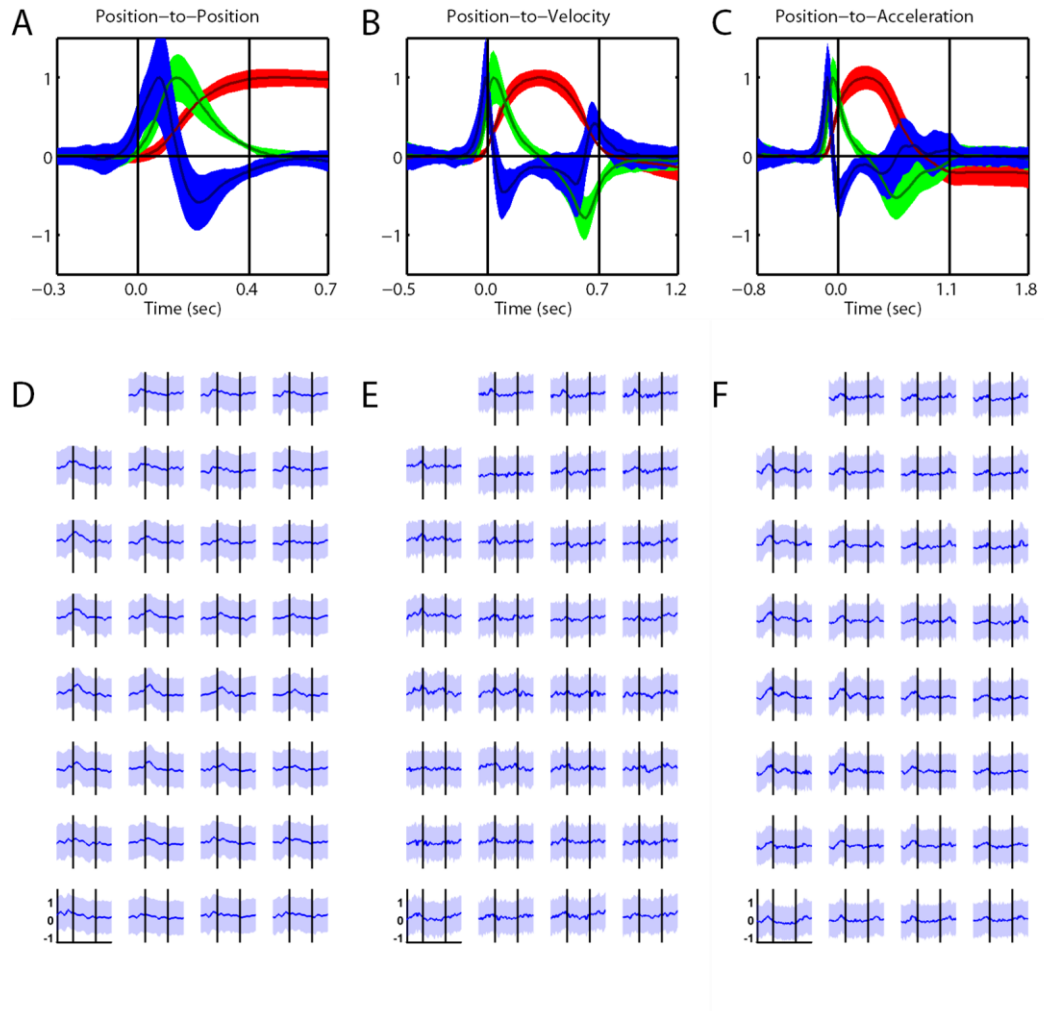
Figure 4.4. Joystick movement and neural profiles binned by cursor velocity across mapping experiments, monkey P. Panel descriptions are the same as in Figure 4.3. Time span distributions between black lines were 0.42±0.06, 0.67±0.28, and 1.06±0.41 seconds for A-C, respectively, and the corresponding channel ECoG profiles in D-F.

positive and negative deflections in joystick velocity are plotted to illustrate the movements necessary to complete the mapping tasks, the absolute value of this parameter (and for acceleration as well) should be considered against neural activity as well. If the absolute value of the velocity traces in each mapping are considered, then a single positive speed deflection should be present during the position mapping task, two positive deflections for the velocity mapping task, and three positive deflections for the

acceleration mapping task. Accordingly, one, two, and three similarly spaced positive deflections emerge in the mean neural traces for the anterior channels (first 4-5 rows) that appear to align (with delay) to the corresponding joystick speed. A strong example channel of this is shown in Figure 4.5.
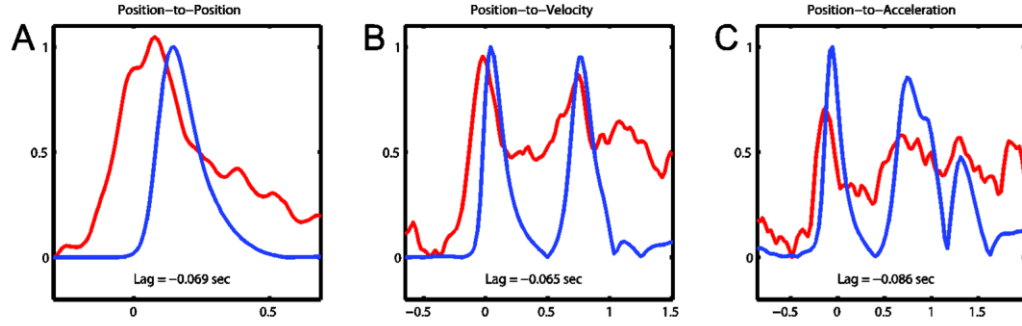
Figure 4.5. Example channel with correlation to joystick velocity. Blue traces indicate normalized joystick velocity projected onto the center-to-target vector while normalized neural traces are shown in red. t=0 on the horizontal time axes indicate cursor speed exceeding 15% of its maximum value. The lag between the first velocity and neural peaks are provided for each task. Negative values indicate neural traces lead velocity traces.

Admittedly, it is difficult to discount whether these neural profiles could alternatively be correlated with absolute joystick acceleration profiles. Taking the absolute value of the acceleration profiles would produce closely spaced positive deflections that would overlap significantly with the speed traces, and the temporal resolution of the neural signals obtained using the envelope estimation method used throughout this thesis may not be able to finely distinguish such peaks. However, perhaps more importantly than whether these channels strictly represent joystick speed or acceleration is the fact that their neural profiles change shape from the position-to-position task, suggesting that these signals are more closely related to the physical movement aspects of the task as opposed to the perceived cursor characteristics.

The posterior-most rows exhibit a single bell-shaped curve near the beginning of each task in Figures 4.3-4.4. Since the shape of these profiles remain consistent across tasks, it is unlikely they are as closely tied to physical movement characteristics as more anterior channels and could possibly follow cursor characteristics. However, as discussed in the next section, the timing of these deflections discount the likelihood of a purely cursor-movement representation in these channels.

4.3.3 Neural Correlation Lags

While a rudimentary examination of the neural profile shape on each ECoG channel provides a great deal of insight into possible movement parameter representation, a more formalized time-series analysis of correlation between neural and movement features offers further insight into an area's function as well as possible confirmation of cortical anatomy. To measure the correlation between neural and movement parameters, neural features were first shifted relative to a particular movement feature by aligning the prominent peaks of both. The correlation coefficient of the two features was then computed at the peak-aligned lag time. Figure 4.6 depicts a spatial representation of both the time lag and correlation between ECoG features and joystick velocity in each mapping task for monkey K (similar spatial results were observed for monkey P). As seen in the position-to-position mappings of this figure, a majority of channels show strong correlation with the joystick/cursor speed (indistinguishable in this task). At the anterior end of the array, the high gamma ECoG deflection leads the speed profile by approximately 100-150 ms, and this lag time decreases to approximately 50-100 ms by the fourth row. These lag times are consistent with lead times between premotor and primary motor activity and hand velocity results published in previous single-unit studies [59]. During the velocity and acceleration mapping tasks, the lag times of the anterior rows appears to decrease slightly to 50-100 ms, but the relative trend appears to generally conserved (greater lag at most anterior channels).

As mentioned in the previous section, the posterior-most rows in Figures 4.3-4.4 could appear at first glance to correlate well with the perceived cursor velocity across all mappings as a single prominent deflection is seen on each mapping task. This hypothesis could be anatomically plausible as these channels were intended to be implanted near sensory and high-level motor planning areas. However, examination of the time course of these deflections argues against this hypothesis as the apparent lag between the ECoG deflection and the cursor velocity trace increases across tasks from approximately 220 ms for the position task to over 900 ms during the acceleration mapping task. On the other hand, when referenced to the appearance of the peripheral

target, the peak of the ECoG deflections on these channels appears to be consistent across tasks at approximately 350 ms.
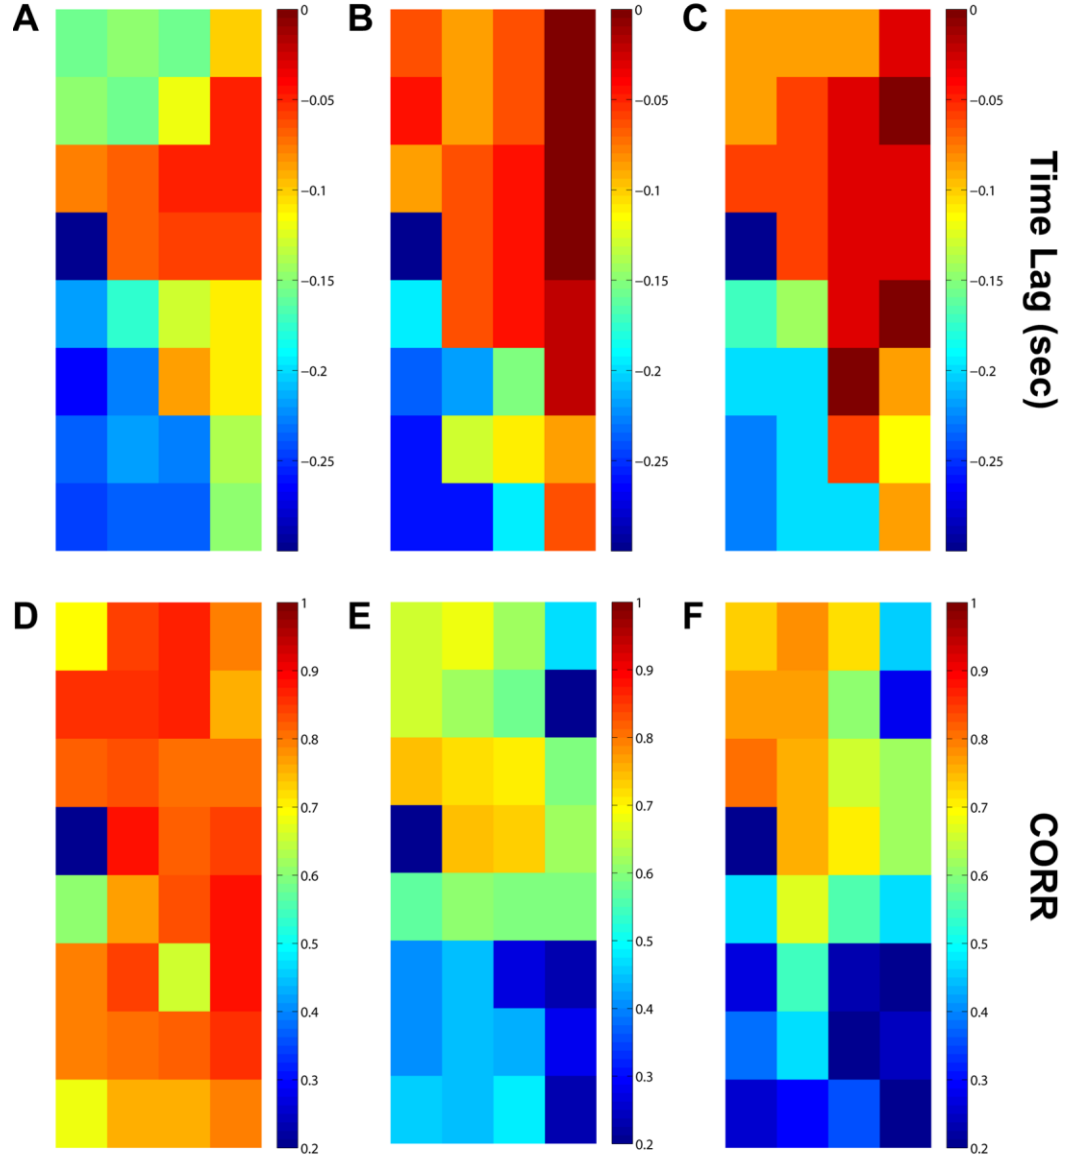
Figure 4.6. Time lags and correlation between ECoG features and speed profiles, monkey K. Top panels depict the time lag shift of ECoG array features that produces maximum correlation with joystick speed for (A) position, (B) velocity, and (C) acceleration mapping tasks. Corresponding correlation coefficients (D-F) are shown in the bottom panels. (The channel in the 1st column, 4th row is a noisy channel excluded from analyses as in Figure 4.3.)

This event-related depolarization could be caused by several hypothetical phenomena. First, the neural profiles on these channels could simply be a visual sensory reaction to the target stimulus. Areas in the Posterior Parietal Cortex (PPC) receive strong visual input and often show activity related to saccades or reaches [66], [67]. As our goal for the surgical implantation of these arrays was to span from premotor areas to parietal areas, it is possible that the posterior rows of these arrays provide some coverage over cortical areas of this visual or saccade related nature. It is also possible that these electrode sites cover more reaching related areas such as the Parietal Reach Region (PPR), and the activity on these channels represents some more high-level reaching plan such as movement of the cursor. However, PPC regions such as PPR that show more directly reaching-related behavior typically exhibit sustained, elevated activity throughout any delay periods and through the actual reach. In the cases discussed here, these channels only display brief, transient depolarizations whose timing remains consistent with the target stimulus and not the cursor or joystick movements. While it is difficult to speculate the exact role of these channels, their time course at least suggests a more visual or early abstract planning role over lower-level instructions closer to the motor output.

4.4 Conclusions

The results from this chapter provide both interesting neurophysiological insights as well as practical information for the BCI experiments discussed later. From a purely scientific standpoint, the velocity and remapping experimental data suggest that classic "motor" areas such as premotor and primary motor cortex are primarily concerned with coordinating the desired physical movement profiles needed to accomplish a task and not necessarily the movement profile of any secondary tool or effector coupled to those physical movements. So while somewhere in the brain of a person playing a racing video game a neuron might be saying, "Accelerate forward, turn right, slow down ...", the motor cortical areas examined in this chapter are most likely relaying something along the lines of, "Move thumb up, move thumb right, move thumb down ...". Conversely, the posterior edge of channels in both monkey arrays

appeared to be more related to target cues than physical movement. However, there was little evidence for consistent correlations with cursor movement parameters. It is not implausible that this conceptual relation exists in the brain, but just less likely that it exists in the cortical areas covered in these experiments.

Although these findings may not be groundbreaking, they do offer several practical benefits that can be exploited. First, the analyses in this chapter provide some anatomical evidence toward individual channel locations without sacrificing an animal for histology. As one goal for our ongoing BCI research is to demonstrate the long-term viability of these chronic implants, it may be many years after the initial experiments before the precise array placement can be confirmed through histology. Furthermore, localization through imaging techniques may be hampered by the electrode design (large metal ground bars sitting directly above electrode sites as shown in Figure 3.1 along with numerous titanium straps anchored to the monkey's skull). Thus, the correlation and lag results (along with our pre-surgical stereotaxic planning) can provide an educated estimate of where premotor, primary motor, posterior parietal, etc. cortices may reside in relation to the ECoG grid until further histological data can be obtained. These approximate channel locations can then be used as background information for how these channels perform in a BCI context in later experiments.

In addition, the movement parameter representations suggested here also provide a backdrop for comparison in BCI experiments. Initially, we suggested that a cortical area found to have flexible movement parameter representation in these experiments (*e.g.* coupled with perceived cursor movement regardless of physical movement as opposed to being tied to one physical movement parameter) might be suitable for a wider range of control schemes. However, in light of the results presented here, the opposite may very well be true. If a particular area is consistently related to the perceived end-effect of movement (*e.g.* the displayed cursor position or velocity), it may be difficult to remap that representation to a different parameter of the same effector (*e.g.* cursor acceleration) as the desired and perceived parameters would likely cause interference. On the other hand, using an area with lower-level representations closer to the motor output may offer easier translation to other control schemes. The

psychophysical results alone from this study as well as others [34] support this idea as monkeys and humans can be very adept at using one motor command (hand/joystick position) to control another (cursor position, velocity, acceleration, etc.). The following chapters will explore whether this finding translates to BCI control as well, and the psychophysical joystick results from this chapter can provide a benchmark for which to compare BCI versions of these tasks.

5 ECoG Application to a 2-D Kinematic BCI

5.1 Introduction

The experiments in this chapter are used to assess the use of microECoG signals as applied to 2-D kinematic BCI control. To date, BCI control schemes have employed translations from neural signals to kinematic variables such as the position or velocity of a virtual cursor or robotic arm. Thus, we will first examine performance metrics and ECoG characteristics when applied to a velocity control signal for use in a standard center-out task as our control for comparison. We will then examine these measures when similar experiments are performed with an acceleration-based control signal. Although still technically a kinematic variable, acceleration is very close to a force-based control signal (force scaled by a constant mass) and represents a first step toward addressing the feasibility of dynamics-based BCI control.

5.2 Methods

Subjects learned to perform two closed-loop tasks during the course of these experiments based on kinematic control. Subjects first learned brain control of the cursor using a velocity-based control signal in the radial-choice task. After initial training on this task, subjects transferred proficiency with the velocity control signal to the center-out task. After further refining their control in this task, the monkeys finally learned to perform the center-out task using an acceleration control signal. Movement parameter and control signal profiles as well performance metrics were then compared between the two control signal modalities.

5.2.1 Radial-choice Task

As in the joystick experiments described in Chapter 4, a radial-choice task (see section 3.7.1) was used to train the monkeys on brain-control of the cursor as well as provide neural training data for initial decoding algorithm models. As in the joystick training, the cursor would initially appear at the center of the ring, followed 100 ms later

by the appearance of the highlighted target arc on the ring. The monkey was allowed 5 seconds to move through the highlighted arc before the trial would time-out (no reward was given and a new trial was started). For these experiments, the circular arc was divided into four possible target locations for a chance rate of 25%.

In this task, monkeys controlled the velocity of the cursor:

$$v(t) = g_v \cdot y(t) \qquad (5.1)$$

where $v(t)_{2x1}$ is the two-dimensional cursor velocity, $y(t)_{2x1}$ is the ECoG-derived control signal, and g_v is a constant velocity gain term for scaling the control signals to an appropriate range for movements on the screen. Initial training under brain-control was assisted through the use of a cursor bias velocity that was added to the brain-derived velocity signal. This bias velocity always pointed in the direction of the target. During early training sessions, the magnitude of the bias signal was set relatively high compared to the brain control signal such that the cursor would drift toward the target, but the monkey would still receive visual feedback about the effect of his brain activity on the movement of the cursor. The bias velocity was also used such that the cursor would eventually (but slowly) drift to the target and provide a reward to keep the monkey from becoming frustrated, but the monkey also had the opportunity to identify a neural state that made the cursor travel in the direction of the target faster and therefore be rewarded more often. As the monkey identified productive cortical modulations over days, the bias velocity magnitude was gradually decreased until the cursor was under complete brain control.

5.2.2 Center-out Task

After becoming proficient at the radial choice task and demonstrating BCI performance well above chance with no assistance, each monkey learned to complete a classic center-out task (see section 3.7.2) using the brain-derived velocity control signal. As in the radial-choice task, one of four peripheral target locations residing directly on the horizontal and vertical axes were pseudorandomly presented as the target for a given trial. Failure to complete either the Hold A or Hold B time, to reach the center target within 10 seconds of the start of the trial, or to reach the peripheral target within the

maximum movement time of its appearance would end the trial with no reward. Hold A and B times were gradually increased from 100 to 500 ms as the monkey improved at slowing down the cursor on approach to a target. The maximum movement time to reach the peripheral target was also decreased from 5 seconds to 2 seconds, although this time limit was scarcely reached under the velocity control mapping.

Following training on the center-out task with velocity control, each monkey performed the same task with an acceleration-based control signal:

$$a(t) = g_a \cdot y(t) - k \cdot v(t) \qquad (5.2)$$

where $a(t)$ is the cursor acceleration, $y(t)$ is the neural control signal, and k is the viscous damping constant. Initially, the damping constant was set high such that the equations of motion essentially reduce to a velocity control scheme (see section 3.6). The damping constant was then gradually reduced from experimental session to session as the monkey learned that manual "braking" of the cursor with his control signal was necessary to slow down the cursor in a timely manner for hold periods. At the end point of these experiments, the damping constant held a non-negligible value such that the monkey needed to apply his control signal in the opposite direction of the target to slow the cursor, but not to the extent of the peak acceleration signal. Hold A and B times were again gradually increased to 500 ms while the maximum movement time was decreased to a minimum value of 1600 ms.

5.2.3 Decoding Algorithm and Training

The normalized spectral ECoG features, $X(t)_{160x1}$, were transformed to the two-dimensional BCI control signal, $y(t)_{2x1}$, through the following linear equation:

$$y(t) = w^T X(t) + b_0 \qquad (5.3)$$

where w_{160x2} is the decoding weight vector (160 features x 2 dimensions) and $b_{0_{2x1}}$ is a constant bias term for the control signal. The resulting control signal was then mapped to either the two-dimensional velocity or acceleration as in equations (5.1-5.2) and integrated appropriately to the displayed cursor position.

The weight vector and constant terms were adapted to the monkey's neural strategy during velocity control experiments on a daily basis by observing the displayed

target positions for a given set of trials and regressing against the accompanying ECoG feature responses, **X**. No *a priori* correlations between neural and movement features were used in the training of any decoding model. Instead, each monkey's initial decoding model consisted of random weights paired with a high bias velocity and used strictly neural and cursor data to train future models.

Training data consisted of 50 consecutive trials to each of the four targets. Since the trials were pseudo-randomly balanced across targets, the 200 total trials were chosen from the block of time with the highest average percentage correct across targets. We chose to train on a limited subset each day under the assumption that the monkey was performing with maximum effort during this period, thus reducing the risk of introducing noise by training on trials during which the monkey was not actively attending to the task.

Using this training set, an ℓ_1-regularized least-squares regression model was formulated using a MATLAB toolbox (SLEP, Arizona State University [68]) to learn the weight vector *w* and bias b_0 within the following optimization problem:

$$min_{(w,b_0)} \sum_{t=1}^{T} \left((w^T X(t) + b_0) - \hat{y}(t) \right)^2 + \lambda |w|_1 \qquad (5.4)$$

In this context, $\hat{y}(t)$ is the desired control vector pointing from the cursor to the desired target, and $\lambda \geq 0$ is the ℓ_1 regularization constant that trades off training accuracy with model complexity. (For simplicity, the magnitude of the desired vector was normalized to a unit vector such that we were primarily interested in correctly modeling directional intent.) Increasing values of λ drive *w* toward a sparse decoding matrix selective for the most relevant features [69]. This regularization parameter was chosen using either leave-one-out cross-validation (LOOCV) with trial-averaged data from radial choice task experiments, or 10-fold cross-validation with real-time data from center-out task experiments. The regularization constant with the highest cross-validation accuracy was used to train the day's final model over the entire 200 trial subset. This post-hoc decoding model was then employed in the next training day's closed-loop experiments.

For velocity control, a directional unit vector from the cursor to the target

should be an accurate estimate of the subject's desired cursor direction for model training. This estimate might become noisy as the cursor closely approaches the target due to the radii of the cursor and target spheres. For example, if the cursor's trajectory is slightly off course from the center of the target but still likely to result in overlap of the two radii, the monkey may desire to simply continue on a course nearly tangential to the target. This plan would result in an estimated desired vector (cursor-to-target) that is nearly perpendicular to the true desired vector and a dot product between the two vectors near zero.

However, for an acceleration control signal, estimating the directional intent of the subject is more difficult as discrepancies between the estimated and actual desired control vectors can be much greater than in the velocity control scheme. These discrepancies arise due to limited knowledge of the subject's favored time course for accelerating/decelerating the cursor. An analogous situation to this might be the tendencies of drivers as they approach a red stoplight. One person might tend to keep his/her foot on the gas until the last second before firmly pressing the brake to stop. Another person might let off the gas pedal half a block before reaching the stoplight, begin feathering the brake early on, and only need to lightly depress the brake further to stop at the end. If each person's desired acceleration model was wrongly applied to the other's, the dot product between the estimated and true desired control vectors would be negative for a significant period of time before reaching the stoplight. When applied to our BCI model, this scenario would result in a significant proportion of noisy or grossly incorrect training data.

For these reasons, we chose not to use neural and cursor data from acceleration control experiments to update our decoding models. The decoding models used to control acceleration for each monkey were trained on data from the last day of velocity control. The monkeys were then expected to adapt their neural modulation under these static decoding models to produce appropriate acceleration control profiles for completing the center-out task.

5.2.4 BCI Psychophysics

For binned data analyses, neural and cursor data were binned based on cursor velocity characteristics as in Chapter 4. Since neural features were directly causal for cursor movement and any delays between neural and cursor features were due to associated signal processing delays and not neurophysiological delays, the relative number of bins preceding, during, and following the movement period were changed from the joystick experiments. Each movement period (based on cursor velocity first rising above and then below 15% of max velocity) was divided into 70 equal time bins. This period was then preceded and followed by 15 equal length time bins for a total of 100 bins per trial. For non-binned analyses, neural and movement data were downsampled to 50 Hz.

The performance metrics introduced in Chapter 4 were also calculated for BCI center-out experiments. These metrics included the Time to Acquire Target (TAT), maximum-to-Hold B velocity ratio ($R_{max/B}$), normalized path length (d_{path}), and dot product between the velocity and cursor-target vectors ($dp_{vel-targ}$). These metrics were calculated only for center-out trials in which the peripheral target was touched (either no error or Hold B error), but the percentage of trials with each type of error (Start error, Hold A error, Movement error, Hold B error, or none) were noted for each day to give more context to the performance metrics.

5.3 Results

5.3.1 General Task Performance

Figure 5.1 shows center-out cursor trajectories during example training days late in velocity and acceleration control experiments for each monkey hemisphere. From gross inspection of the figure, BCI "reaches" to the target were fairly similar between control schemes for each monkey. Reaches were not nearly as tight and consistent as those seen using the same control schemes during the joystick mapping experiments (Figure 4.1).

Along with this figure, Table 5.1 shows the relative proportion of errors for

each control scheme. From these example days, it appears that one of the more difficult aspects of either control variable experiment was successfully completing the Hold A time to reveal the peripheral target. In the context of these experiments, the period from the appearance of the center target to fulfillment of Hold A is the least constrained and most variable since the cursor can start at any position on the screen depending on where the last trial ended. Conversely, following the Hold A period, the cursor always starts at approximately the same point and has one of four desired trajectories to finish the trial. Anecdotally, the Hold A errors tended to occur in groups. For example, the cursor might be in a far corner of the screen when the center target appeared. With the increased distance to the target, the monkey would over-accelerate past the target for one Hold A error, over-compensate and overshoot the center target again for a second error, and so on until slowing the cursor enough to fulfill the first hold time. In essence, this task setup allows for two "reaches" per trial: one variable in length and direction and one static in length and discretely variable in direction. Other examples of BCI center-out experiments commonly omit the first reach or its error type by either automatically returning the cursor to the center at the beginning of a trial or by allowing unlimited time and attempts at holding the cursor in the center before officially starting a trial.

Given that the Hold A time was fulfilled, the monkey reached the peripheral target within the allotted movement time (either no error or a Hold B error) more often than not (Movement Time error). Within trials where the peripheral target was reached, correct trials were more frequent than Hold B errors. Between cursor mappings, the proportion of correct trials compared to Hold B errors was generally higher for velocity mapping days.

Start errors (failure to reach the center target within 10 seconds) were relatively uncommon in most cases. These errors usually occurred when the monkey was clearly not attending to the task. Monkey P was often noted to take multiple breaks of this nature within a training day, consistent with his higher percentage of start errors.

Table 5.1 Trial error summary for example BCI days. Data indicates the number and relative percentage of trials for each error type within a day's session.

Monkey (Hemisphere)	Control Variable	Correct Trials (%)	Start Errors (%)	Hold A Errors (%)	Movement Time Errors (%)	Hold B Errors (%)
K (L Hem)	Velocity	291 (34.4)	4 (0.4)	341 (40.3)	50 (5.9)	161 (19.0)
	Accel	192 (30.5)	0 (0)	172 (27.3)	167 (26.6)	98 (15.6)
K (R Hem)	Velocity	567 (52.9)	27 (2.5)	356 (33.2)	25 (2.3)	97 (9.0)
	Accel	306 (20.3)	1 (0.1)	685 (45.5)	247 (16.4)	267 (17.7)
P (L Hem)	Velocity	329 (22.8)	371 (25.6)	474 (32.8)	86 (5.9)	186 (12.9)
	Accel	225 (16.3)	101 (7.3)	599 (43.3)	296 (21.4)	162 (11.7)
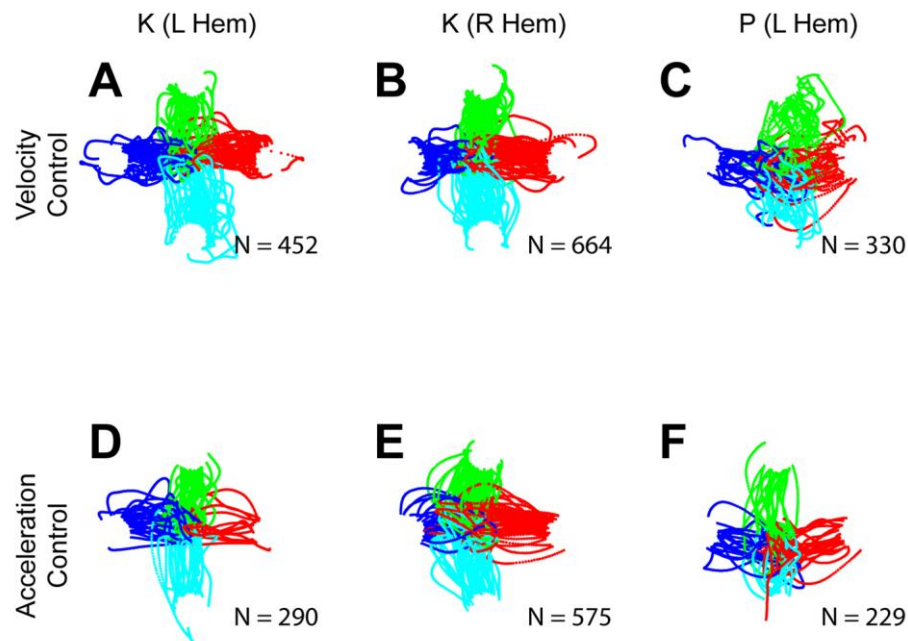
Figure 5.1. Cursor trajectories for BCI center-out experiments. Panels A-C show cursor trajectories using the velocity control scheme to all 4 targets for all three monkey hemispheres. Panels D-F show cursor trajectories for acceleration control experiments. Total trial numbers for a given mapping experiment day are indicated by *N*, and targets are grouped by color. (For plotting purposes, trials were downsampled to display every third trial.)

5.3.2 BCI Psychophysics

The psychophysical metrics introduced in Chapter 4 are shown for the example

BCI center-out experiments in Table 5.2. As seen from comparing mapping task results within a given monkey hemisphere, these performance metrics were generally similar between the two task versions. Although many of these measures exhibit mean statistical differences between mappings based on the standard error of the mean (SEM), trial-by-trial comparison yields substantial overlap of parameter distributions based on each metric's standard deviation. Mean differences between metrics were also not necessarily consistent across monkey hemispheres. For example, monkey P had a higher $R_{max/B}$ for late acceleration-mapped experiments, but monkey K demonstrated the opposite trend for both hemispheres.

Table 5.2 Summary of psychophysical metrics from BCI center-out experiments. (Standard deviations indicated in parentheses)

		BCI Mapping	
		Velocity	Acceleration
TAT (sec)	K (L)	1.80 (0.57)	1.46 (0.27)
	K (R)	1.26 (0.27)	1.20 (0.23)
	P (L)	1.36 (0.47)	1.36 (0.27)
d_{path}	K (L)	1.49 (0.31)	1.32 (0.21)
	K (R)	1.37 (0.20)	1.48 (0.32)
	P (L)	1.61 (0.37)	1.42 (0.29)
dp_{vel_targ}	K (L)	0.70 (0.20)	0.69 (0.19)
	K (R)	0.72 (0.18)	0.68 (0.17)
	P (L)	0.69 (0.18)	0.67 (0.19)
Max Speed (units/s)	K (L)	4.38 (1.13)	4.09 (0.89)
	K (R)	4.64 (0.93)	5.33 (1.13)
	P (L)	5.51 (1.23)	4.45 (0.98)
Hold B Speed (units/s)	K (L)	1.88 (1.02)	1.74 (1.00)
	K (R)	1.90 (0.86)	2.26 (1.24)
	P (L)	1.96 (1.04)	1.57 (0.79)

$R_{max/B}$	K (L)	1.87 (0.17)	1.80 (0.08)
	K (R)	2.09 (0.08)	1.94 (0.28)
	P (L)	2.21 (0.11)	2.42 (0.20)

Comparing Table 5.2 with the joystick results in Table 4.1, several observations emerge. Movement times for BCI tasks were slower than their joystick counterparts, although differences between velocity and acceleration mappings were greater for joystick tasks than BCI tasks. The path length and dot product metrics were longer and smaller, respectively, than their joystick counterparts, consistent with the trajectories in Figure 5.1 showing more deviation from a straight line trajectory. Average trial max speeds were actually higher for BCI experiments under both mappings. However, average Hold B speeds were also higher, resulting in a lower velocity ratio, $R_{max/B}$. In general, it appears that the brain-to-control signal gain for each mapping task was set appropriately high such that max speeds were similar or even slightly higher than in the joystick tasks. However, the tradeoff for this characteristic was accuracy in the cursory trajectory as well as the ability to stop the cursor on the target.

5.3.3 Movement and Control Signal Profiles

Although the metrics in the previous section were compared across control signal mappings as well as between joystick and BCI experiments, the goal of this thesis is not necessarily to demonstrate superiority of one control variable over another for BCI use. Indeed, velocity and force variables could be considered tools toward different tasks in the context of BCI systems (i.e. velocity for pointing/selecting, force for interacting with objects). Rather, we wish to address the feasibility of a force-based BCI. Toward this aim, we investigated the ability of subjects to modulate neural signals in a manner amenable to properly controlling the acceleration or forces on an object.

In the center-out task and its respective task constraints for acceleration control, monkeys had to both accelerate the cursor at a sufficient magnitude in order to reach the cursor in a limited time period as well as apply a decelerating control signal in order to stop the cursor on the target. Figures 5.2-5.4 shows cursor movement parameters for

under velocity and acceleration control for each monkey hemisphere. As seen from the figures, the movement parameter profiles for each mapping look considerably similar to each other. Comparing the movement profiles across targets, it is apparent that each monkey had preferred targets for which they had more consistent control in that direction.

It is interesting to note which movement parameters present the most visible "noise" under each mapping condition. For velocity control, the cursor velocity – the variable directly containing neural variability – contains a significant amount of visible variability. This variance decreases when integrated to position but is amplified for the acceleration derivative. For acceleration control, however, variability assumed to originate from neural input arises in the acceleration parameter while integration to velocity and position drives this variability down. For the final integration, the position trace demonstrates less variance under the acceleration mapping than the velocity mapping in many cases.

Also seen in these figures, each monkey is clearly decelerating the cursor at the end of the movement under both mappings. While the statistics from Table 5.2 suggest that the cursor is not completely stopped at the end of the Hold B period, the velocity and acceleration traces demonstrate that a significant deceleration below zero is present in each case. However, because a viscous damping term was included with the acceleration mapping, examining the acceleration profiles alone does not necessitate that each monkey's control signal followed suit. To further investigate this point, Figure 5.5 depicts the mean neural control signal projected along the center-to-target line under each mapping for monkey P. As seen from panel A, the neural control signal profile matches that of the corresponding velocity profile in panel A of Figure 5.4. In panel B, however, there is a visible decrease in the amplitude of the deceleration phase of the control signal. The difference during this phase between the acceleration and control
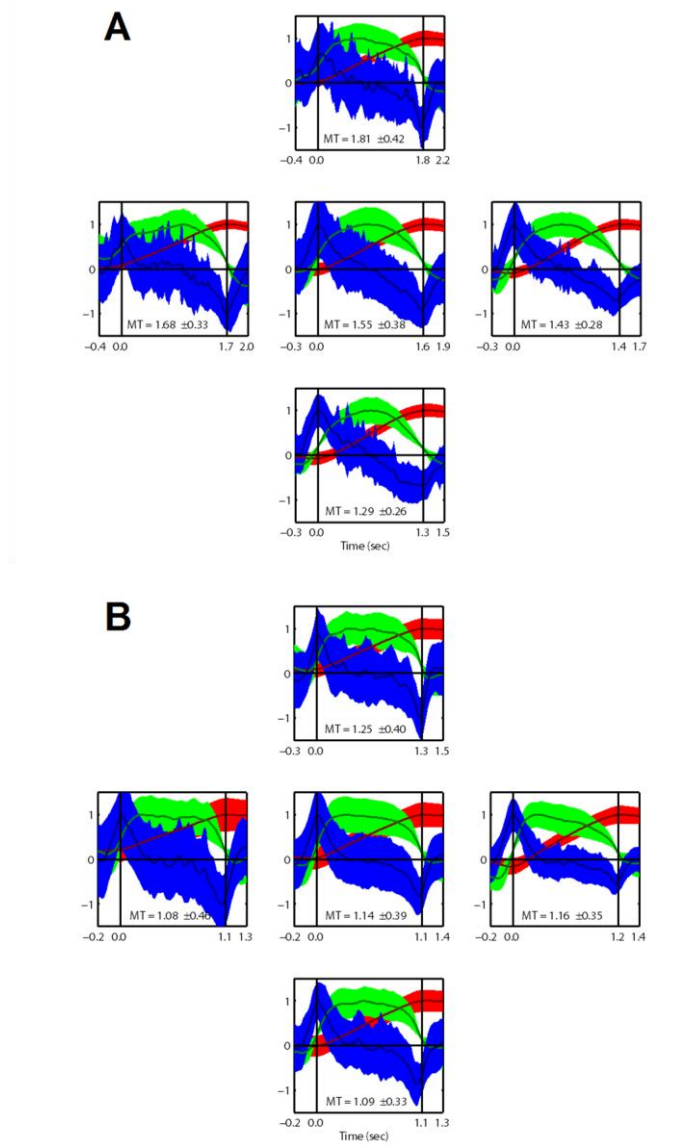
Figure 5.2. BCI cursor movement profiles for monkey K, left hemisphere. Mean cursor movement parameters per target (peripheral plots) and overall means (center plots) are plotted for **(A)** velocity control and **(B)** acceleration control. Red, green, and blue traces indicate mean cursor position, velocity, and acceleration projected along the center-to-target line, and shaded regions indicate standard deviations. Black bars depict the period used to bin data by cursor velocity, and MT text indicates the average time between these points.
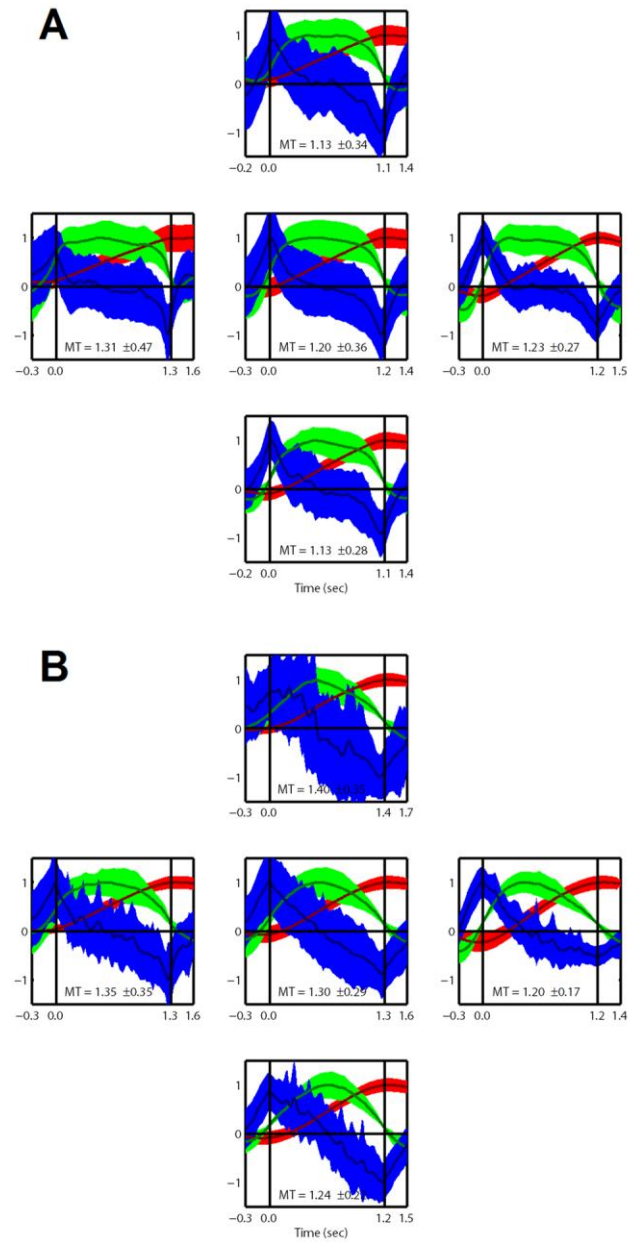
Figure 5.3. BCI cursor movement profiles for monkey K, right hemisphere. (A) Profiles using velocity control. (B) Profiles using acceleration control. See Figure 5.2 for further description.

67
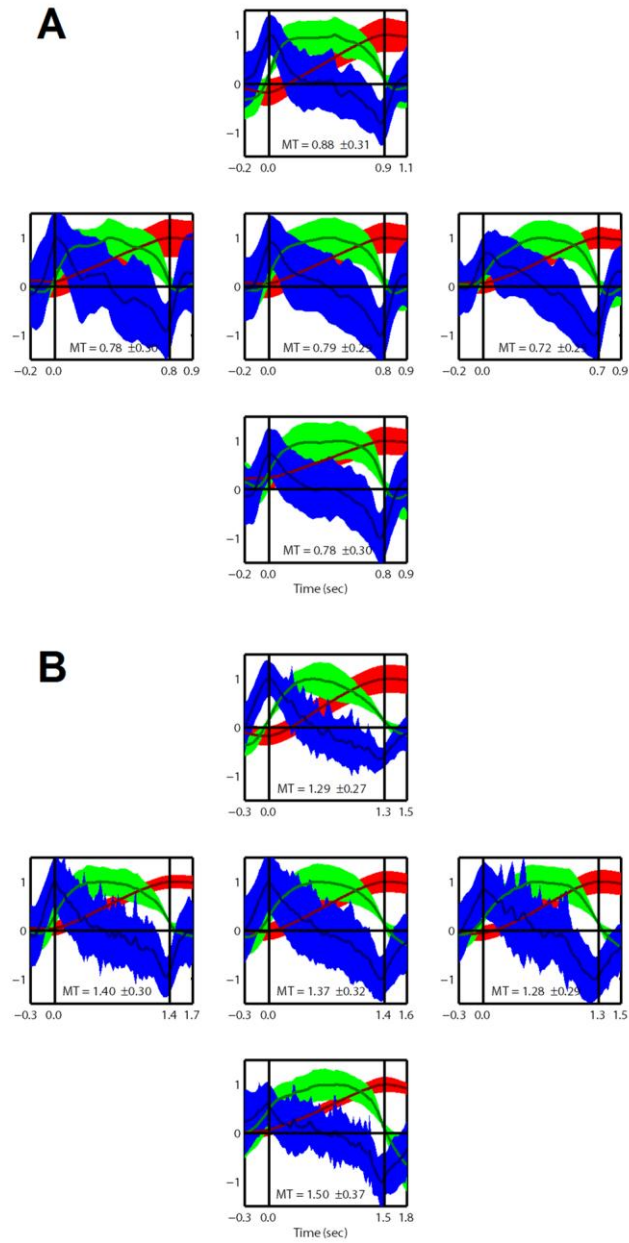
Figure 5.4. BCI cursor movement profiles for monkey P, left hemisphere. (A) Profiles using velocity control. (B) Profiles using acceleration control. See Figure 5.2 for further description.
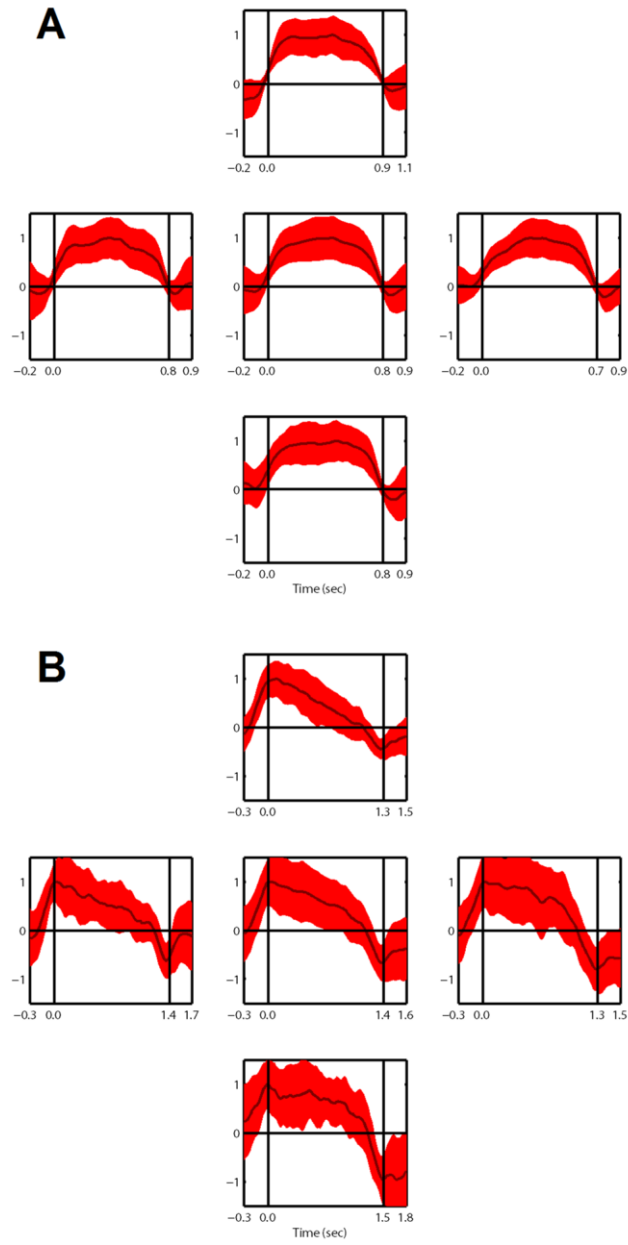
Figure 5.5. BCI control signal profiles for monkey P, left hemisphere. (A) Profiles under velocity control. (B) Profiles under acceleration control. Dark lines indicate mean profile while shaded areas indicate ± one standard deviation. Time distributions and plot layout descriptions are identical to those in Figure 5.4.

signal profiles is due to the damping term included in Eqn (5.2). If this damping term was high enough, the monkey could use a neural control signal profile similar to those under the velocity control mapping to accelerate the cursor while letting the viscous force decelerate the cursor. However, it is clear that the neural control signal was actively modulated in the opposite direction of the target to slow the cursor rather than simply returning to baseline. Neural profiles under acceleration control for monkey K (both hemispheres) showed similar results.

5.3.4 Control Signals vs. Cursor Position

The choice to bin cursor and neural data based on cursor velocity characteristics may have biased the shape and variance of the resulting profiles. By aligning data on points where cursor velocity at the beginning and end of the movement period was small (15% of maximum velocity for the trial), the variance of the velocity profiles at these points was minimized while shifting that variance to other sources (*e.g.* time length of the movement period, profiles at times away from the alignment points).

To verify that the trends in cursor and control signal profiles were not merely an artifact of a particular binning process, similar analyses were carried out with data binned by cursor position. As before, all cursor and control signal variables were projected onto the center-to-target vector. Percentiles for cursor positions were then calculated at 5% increments. Cursor position data were binned by these percentiles along with corresponding velocity, acceleration, or control signal data. Figure 5.6 shows an example of neural control signal profiles as a function of cursor position. As seen from the figure, the projected control signal modulates away from the target as the cursor draws within the final 25% of the center-to-target distance. These results are in agreement with the profiles binned by velocity and also exhibit the asymmetric control modulation due to the damping force.
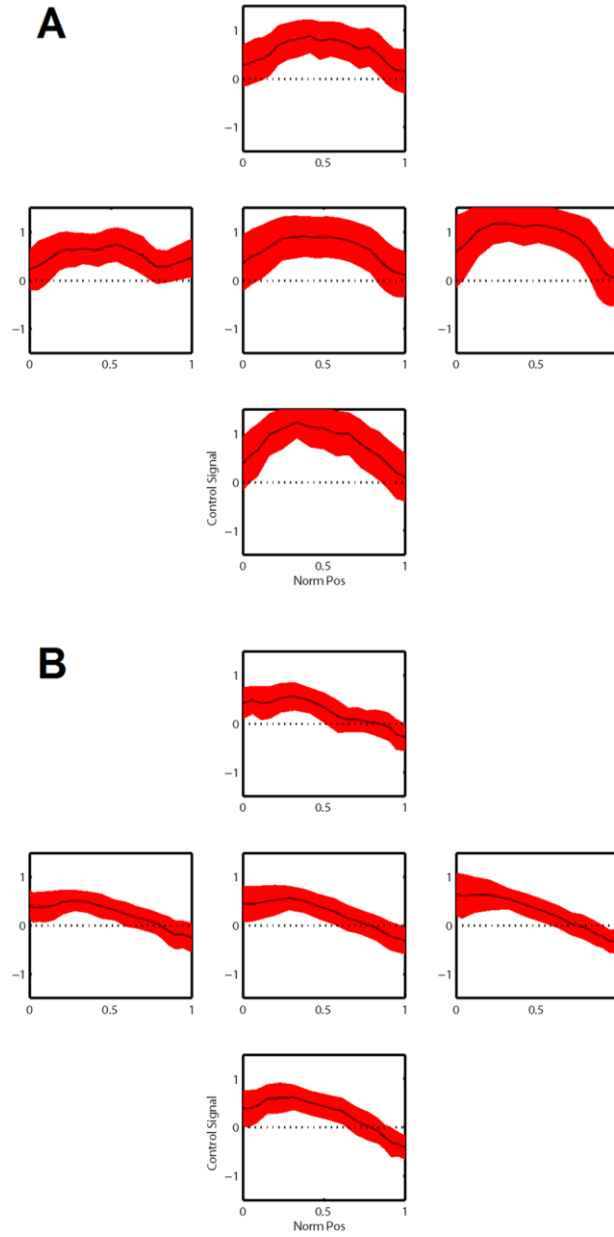
Figure 5.6. BCI control signal vs. cursor position. **(A)** Profiles under velocity control. **(B)** Profiles under acceleration control. Dark lines depict the mean profile while shaded areas indicate ± one standard deviation. Cursor limits were normalized from [center target position, peripheral target position] to [0,1]. Plot locations indicate averaging over individual targets (peripheral plots) or all targets (center plots).

5.3.5 Decoding Weights vs. Control Modality

As mentioned in Chapter 2, a number of neural correlations with different movement parameters, both kinematic and dynamic, have been suggested in motor-related cortical areas through open-loop experiments such as those examined in Chapter 4. The breadth of these representations have not been as well studied in BCI contexts, however, as output signals have been primarily constrained to kinematic variables such as position or velocity. In the BCI experiments in this thesis, we did not assume any prior correlations between ECoG signals and movements, and instead allowed decoding models to develop solely from BCI trial data and adaptation on the part of the subject and algorithm. Thus, using the decoding framework used in these experiments, the following question arises: Do different ECoG/cortical features lend themselves more or less to a specific BCI control variable? Conversely, one could ask whether it is matching neural features with their predilection for a specific movement variable or the intrinsic modulatory capacity of recorded neural features that is more important to the success of a BCI system.

To address these questions, the neural control signals for a given day were reanalyzed alongside their causal ECoG features to evaluate which features were most responsible for driving the velocity or acceleration control signal output. Two approaches were used in these analyses. For the first approach, the actual model weights for a given day were used to compare a neural feature's actual modulation with its expected modulation based on the observed control signal. For each feature, the two-dimensional control signal was projected onto the feature's preferred direction vector. The resulting one-dimensional control signal time series was then correlated with the time series of the ECoG feature. Intuitively, features with high correlations in this analysis should contribute constructively to the control signal output while features with small correlations might contribute noise or less substantial input to the overall control signal. Likewise, features with negative correlations might contribute negatively to the control signal, although this could theoretically be used to balance out other feature contributions.

Figure 5.7 depicts correlation values for monkey K (left hemisphere) using each

mapping. The figure suggests the most causative features for producing the control signal reside primarily in the high gamma bands. From visual inspection, it also appears there is some overlap in the most correlated signals between control mappings. To quantify the degree of overlap between the two mapping conditions, correlations were sorted in descending order. The top 20 features with the highest correlation values for each mapping were considered, and the percentage of features that appeared in each mapping set was calculated. Percentages for monkey K's left hemisphere, monkey K's right hemisphere, and monkey P's left hemisphere were 70%, 55%, and 30%, respectively. Doubling the number of top features included to 40, the overlap percentages only increased marginally to 72.5%, 62.5%, and 32.5%.

As an alternative approach to identifying the most influential features for BCI control, the decoding weights were ignored and only the neural features and control signal were considered. The rationale for a secondary approach was that the regularization scheme we employed was designed to produce relatively sparse decoding matrices, i.e. two neural features that provide roughly the same information would be pruned down to the feature providing slightly more consistent information. Because the models were trained on data from the previous day (a relatively slow adaptation), the monkey could hypothetically be trying to recruit additional features for a given purpose, but these features may have no weighted influence on the control signal due to previous regularization. In addition (and perhaps more importantly), since models were not retrained during acceleration control experiments, the monkey may have tried to modulate unweighted features toward the new task, but these features would also show no correlation to the control signal using the previous analysis.

Given these considerations, the same algorithm used to train the decoding models on a daily basis (Eqns (5.3-5.4)) was employed while considering the control signal to be the "desired" signal. From this, new models were trained to reproduce the original control signal. Whereas the previous goal of the training algorithm was to produce a decoding model with the highest validation accuracy for desired direction, the goal of this simulation was to reproduce the original control signal with as few features
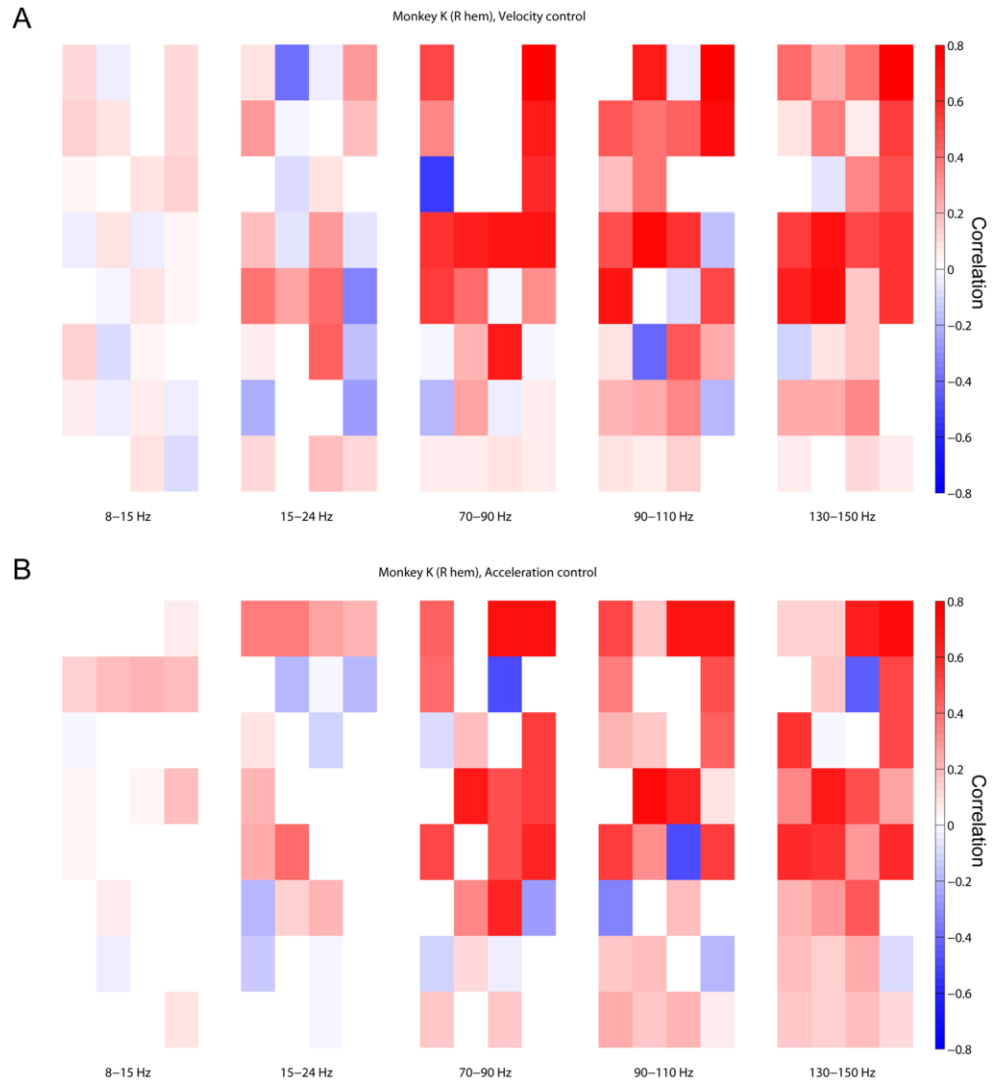
Figure 5.7. Feature correlation with projected control signals. (A) Correlations during a velocity control day. (B) Correlations using acceleration control during an acceleration control day. Data taken from monkey K, right hemisphere.

as possible. Thus, we aimed to identify the ECoG features most correlated with our estimation of the monkey's intent (*i.e.* control signal) without directly biasing the results by the original decoding model. For regularization to over-constrain the decoding matrix, a normalized regularized constant, λ_{norm}, of 0.25 was used ($\lambda_{norm} = 0$ indicates no regularization/all features used; $\lambda_{norm} = 1$ indicates the zero solution model where all model weights are set to 0). This value of λ_{norm} generally produced a number of non-zero weight features between 18-32. Figure 5.8 shows an example of the magnitude of weight matrices for the same experiments shown in Figure 5.7. As in the previous approach, the degree of overlap between decoding models was calculated by finding the fraction of non-zero features from one set (the smaller set if non-zero features were not equal after regularization) also present in the other set. Using this approach, the overlap percentages for monkey K's left hemisphere, monkey K's right hemisphere, and monkey P were 68%, 79%, and 32%, respectively.

Both of the above analyses considered spectral features independent of their spatial location. If just the channel locations of the top correlations/weights were considered, the channel overlap for monkey hemispheres K-L, K-R, and P-L became 85%, 90%, and 80% using the first correlation approach, and 71%, 75%, and 38% using the over-constrained model approach.

In general, with the exception of several spurious results from monkey P, the overlap results suggest that the monkeys were using primarily the same features for control under either mapping condition. This overlap increased when features were grouped by spatial location in most cases.
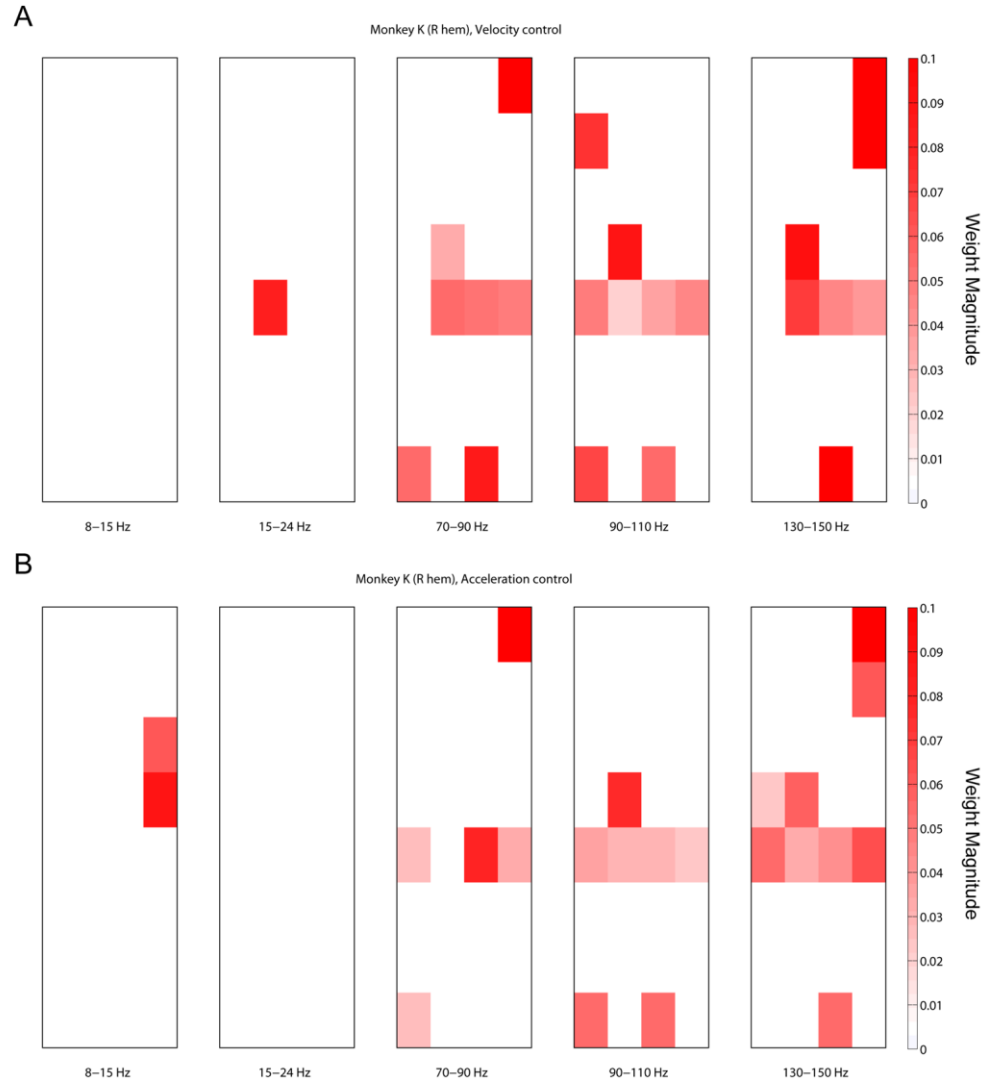
Figure 5.8. Over-constrained model weights. (A) Absolute weight magnitude during a velocity control day. (B) Absolute weight magnitude using acceleration control during an acceleration control day. Data taken from monkey K, right hemisphere.

5.4 Conclusions

The results from the experiments in this chapter represent a first step toward the design of a force-based BCI system. In these studies, it was shown that monkeys can consistently modulate a neural control signal above and below baseline in a relatively short time period (<1.5 seconds) compared to many past BCI center-out

experiments in order to meet the demands of the task. As noted before, acceleration control is not truly dynamics-based control, but it is a step further in that direction than has previously been reported. This experimental setup would be akin to monkey's producing the appropriate forces on a constant-valued mass to complete the task. Experiments in Chapter 6 will expand on these results by forcing the monkey to account for various possible masses as well as external forces with which to interact.

While the inclusion of a viscous force may appear to be "cheating" at first, forces of this nature would likely be present in a BCI-FNS application due to the visco-elastic properties of muscles [70], [71]. We did not exhaustively attempt to match these properties in our simulations, but merely presented a simple, biologically relevant model as a first-pass test at evaluating the feasibility of such a system. As work along this line of BCI research progresses, however, it will be important to incorporate more realistic muscle models into such simulations.

The error rates for different portions of the task reflect the difficulty of the task as well as inexperience at controlling a BCI system. Making the subjects manually return the cursor to the center target at the beginning of a trial appeared to affect the complexity of the task as evidenced by the Hold A error rate. In addition, this added complexity to the rest of the movement period by increasing variance in the initial conditions of the cursor (position, velocity) as well as the resulting analyses. In real-life situations, complex movements are often initiated from a set resting position. For example, before a baseball pitcher makes a complicated throw, he rests in a set position on the mound for several seconds before initiating his movement sequence. If he were required to make the same throwing movement within 0.5 seconds of touching the pitching rubber, his success rate and accuracy would likely suffer tremendously. In our experiments, perhaps trajectories during the movement period would have been more consistent had the computer reset the cursor to the center or had the monkey been allowed to exit the center target briefly as long as he reentered and maintained contact for a longer period of time (~2-3 seconds) before revealing the peripheral target. Even with the drawbacks of the complexity of our task design, however, each subject demonstrated remarkable control given that the initial hold period was satisfied.

It is also important to note that performance within each mapping task was measured at a time point only a few weeks to months after each monkey had started to learn brain-control. It takes newborn humans months to years to develop and learn how to use their innate nervous and muscular systems to move around independently in our world. Likewise, novel tasks like swinging a golf club take weeks to learn the basics of and years (or lifetimes) to master. Given this, it is likely subjects studies would continue to improve at these tasks with practice over time.

Finally, many of the early results presented in this chapter were aimed at a fairly binary question: Can subjects control an acceleration-based BCI? However, results presented later in this chapter offered more neurophysiologic insight that could be relevant to BCI design. In these experiments, we allowed the animals to develop their own strategies for controlling the cursor through biofeedback and attempted to infer and adapt to these strategies in our decoding models. In agreement with previous ECoG BCI studies [9], [65], the high gamma bands examined in this study appeared to be utilized heavily for control. Additionally, we queried whether the subjects might use different areas for control under velocity-based control than acceleration-based control. Most of the results presented here suggest that they did not. However, these results may be biased by the fact that new models were not updated for acceleration control experiments. The fact that the monkeys were still able to utilize existing models trained on velocity control data and adapt it for acceleration-based control presents a novel training protocol for future BCI systems. Using this protocol, the system would adapt to the user as he/she learns basic velocity-based brain control. Following sufficient training, the system would then fix the model and allow the user to adapt to the new context for the decoding model of acceleration/force control. This may indeed be a practical approach to developing a force-based BCI should an accurate model of intended force be unavailable.

6 ECoG Application to a 2-D Dynamic BCI

6.1 Introduction

The results from the previous chapter provide a valuable initial assessment of the feasibility of a dynamics-based BCI system. The ability of subjects to cognitively accelerate and decelerate an object using an acceleration-based control signal is an important pre-requisite to realize before more complex dynamics can be incorporated. This type of control would be similar to controlling an arm (albeit with more degrees of freedom) to making pointing movements in one plane with no other masses with which to interact or forces for which to account. However, this control scheme would be of limited use in the physical world in which we live and interact. Even an act as seemingly simple as taking a drink from a bottle of soda requires estimation of the mass of the bottle, exerting appropriate horizontal forces to smoothly accelerate/decelerate the combined mass of your arm and the bottle to your mouth as well as vertical forces to account for gravity, and using continuous visual and/or proprioceptive feedback to correct for errors in your initial estimate.

Given this, we sought to expand on the results of the previous chapter by applying micro-ECoG signals as a force-based control signal in a virtual environment with simple, real-world physics. In this chapter, we examine whether monkeys can adapt this control scheme to a novel BCI reaching task in which the subject must dynamically move virtual objects with various masses while accounting for a static virtual gravity force. In this manner, we sought to test the feasibility of a force-based control system under a simplified simulation of reaching movements that we perform every day.

6.2 Methods

6.2.1 Mass-grab Task

After demonstrating proficiency at the center-out task using an acceleration control signal as described in Chapter 5, monkeys K and P learned to perform the mass-grab BCI task detailed in section 3.7.3 using a force-based neural control signal. To summarize this task, after moving and holding the cursor at a center target (Hold A), the monkey used his neural control signal to apply a force on the cursor in order to accelerate it from the center of the screen to a peripheral target. This portion of the task was identical to the center-out task in Chapters 4 and 5 and will be referred to as "Movement A". Upon touching the presented peripheral target with the cursor and holding for a brief period (Hold B), the monkey "grabbed" the target, assumed control of the combined mass of his cursor and the target's mass, and accelerated both to the center target for another hold period (Movement B and Hold C).

Although the BCI version of this task immediately followed the acceleration-based BCI experiments in the last chapter, both monkeys had previous experience with this task using the joystick. During this training, the joystick position was mapped to the force on the cursor (i.e. joystick position was used as a surrogate to the neural control signals used for control in this chapter). Joystick and cursor data were recorded for comparison with the control signal and movement characteristics under brain control.

While one emphasis of the BCI center-out experiments in the previous chapter was to demonstrate relatively quick movement profiles similar to natural reaching movements (e.g. a bell-shaped speed profile), the focus of analyses for experiments in this chapter was on adaptation to different loads controlled by the monkey on the outbound and return trajectories. During each experimental day, three different target masses were possible. The smallest peripheral target presented possessed a mass negligible compared to that of the cursor mass. This mass condition was used as a control as the outbound and return path force magnitudes and trajectories should roughly approximate each other. The remaining two masses were non-negligible in

comparison to the cursor mass, and the heaviest mass was normally double that of the second heaviest mass. However, the mass of the heavier targets in relation to the cursor mass was varied on separate sessions to examine the dynamic range of the neural control signal as applied to this task. For example, on one day, the cursor mass would be 100 units, Mass-1 would be 5 units, Mass 2 would be 50 units, and Mass 3 would be 100 units. On another day, Masses 2 and 3 would be 100 and 200 units, respectively. The radius of the displayed target was proportional to the cubic root of its mass in order to provide a visual cue for its mass. As in the previous radial-choice and center-out tasks, targets and masses were chosen in a pseudo-random manner such that 4 trials to each target location/mass combination had to be successfully completed before finishing a task block.

In addition to varying the masses of the peripheral targets, the gravitational force was also occasionally varied between sessions. Varying these two variables allowed us to probe adaptation by the monkey to account for both inertial and static forces. By dynamic or inertial forces, we refer to the force profile required to produce a desired acceleration profile for a given mass, while the static force component refers to the constant downward force due to gravity. To illustrate the differences between these forces, consider a mass-grab trial to a target on the horizontal axis with no gravity present. The amplitude of the horizontal force component on the return path should be greater than that of the outbound path due to the additional force needed to accelerate the extra mass of the target. No vertical force should be present as the trajectory should be horizontal and gravity is not present. If gravity is introduced during a trial with the same cursor and target masses, vertical forces that differ between the outbound and return paths should arise to account for gravitational pull on the combined cursor/target mass and produce a fairly horizontal trajectory. However, the horizontal (inertial) force profiles should remain the same as in the previous trial since the masses of each object did not change and no other horizontal forces were introduced.

6.2.2 Decoding Algorithm

The same decoding model used for each monkey in acceleration control experiments (see section 5.2.3) was used to decode the force-based control signal in these experiments. As in the acceleration control experiments, this model was trained based on data from the end of velocity-control BCI experiments and remained static throughout acceleration and force control experiments. To reiterate the ECoG-to-control signal mapping, the normalized spectral ECoG features, $X(t)_{160x1}$, were transformed to the two-dimensional BCI control signal, $y(t)_{2x1}$, through the following linear equation:

$$y(t) = w^T X(t) + b_0 \qquad (6.1)$$

where w_{160x2} is the decoding weight vector (160 features x 2 dimensions) and $b_{0_{2x1}}$ is a constant bias term for the control signal. The resulting control signal was mapped to a two-dimensional force, F_{BCI}, acting on the cursor or cursor/target mass:

$$F_{BCI}(t) = g_F \cdot y(t) \qquad (6.2)$$

where g_F is the constant force gain term. The force gain was set identical to the acceleration gain in Chapter 5 so that the outbound control profile should resemble those seen in the acceleration control experiments when no gravity was present. The control signal-derived force, F_{BCI}, was translated to the cursor's acceleration, $a(t)$:

$$a(t) = \frac{g_a \cdot y(t) - k \cdot v(t) - G}{m_{total}} \qquad (6.3)$$

where k is the viscous damping constant, $v(t)$ is the cursor velocity, G is the gravitational force (acts only in the vertical dimension), and m_{total} is the current mass being acted upon (either the cursor mass or the combined cursor/target mass).

6.2.3 Movement Parameter and Control Signal Analysis

As in Chapter 5, neural responses to task variables were evaluated by examining the ECoG-derived control signal. Trial data was binned separately for each movement phase since the lengths of the two periods were seldom equal. The time between the beginning of Movement A and the end of Hold B was divided into 100 equal bins, and the time between the beginning of Movement B and the end of Hold C was divided

into 100 equal bins. Binned data was then averaged across trials to a given target location and target mass.

Data was not averaged across target locations as the external force of gravity affected horizontal axis targets differently than vertical axis targets. For example, the horizontal component of the control signal would only need to accelerate and decelerate the cursor/target masses (in a fashion similar to the profiles seen in center-out tasks in Chapter 5) for targets on the horizontal axes. Meanwhile, the vertical control component should produce a static positive force to maintain the vertical position of the cursor/target against gravity. In this case, the dynamic and static forces are segregated to separate dimensions. However, for vertical targets, the vertical control component would be responsible for both static and dynamic force considerations, while the horizontal control component should show little modulation. In addition, gravity would assist movement of the combined cursor/target mass during Movement B for the upper target but hinder it for the lower target.

Finally, a discriminability index, d', was calculated for control signals and movement parameters between different mass conditions to examine how changes in the mass under brain control affected the distribution of the signal of interest. This metric measures the separation between two distributions and provides an indication of how well one could distinguish between two conditions given a signal of interest. The metric is calculated by:

$$d' = \frac{\mu_2 - \mu_1}{\sqrt{p_1\sigma_1^2 + p_1\sigma_1^2}} \qquad (6.4)$$

where μ is the mean of the signal of interest under a given condition (1 or 2 in this case), σ is the corresponding standard deviation of the conditional distribution, and p is the relative proportion of samples for a given condition ($p_1 + p_2 = 1$). The index is presented in units of standard deviation. For example, a d' of 2 would indicate the means of the two distributions are separated by two pooled standard deviations. The d' metric will be utilized more heavily in Chapters 7 and 8 but is introduced here for the motivation stated above.

6.3 Results

6.3.1 Target Mass vs. Control Signal

To provide a "gold standard" for comparison, Figure 6.1 shows examples of the velocity characteristics (projected onto the center-to-target line), the joystick-based control signal profiles to each target/mass combination, and the vertical position of the cursor for each target/mass combination. The data shown are from a training session for monkey K, but monkey P showed similar trends. Several observations concerning the monkey's task performance and strategy under "ideal" conditions emerge from this figure. First, the velocity and control profiles during Movement A for trials to each mass for a given target are nearly identical. This control period confirms that the monkey is only accounting for the mass of the cursor (identical during this period for trials to all three target masses) and is not prematurely adjusting for the additional mass of the displayed target. Second, the velocity profiles during the Movement B period are more flat and spread out over a longer period of time with heavier target masses. Ideally, we would hope to see the velocity profiles during this period to be identical across mass conditions as well if the monkey was able to completely adjust for the additional mass. However, this observation is related to the dynamic range limits of the control signal. For example, as seen from the horizontal control signals (Control X) to targets 0 and 2 and the vertical control signals (Control Y) to target 3, the control signals are limited to ±1 due to the mechanical constraints of the joystick. The implications of this finding will be discussed further in section 6.3.4. Third, the vertical control signals during trials to the vertical-most target, target 1, were very similar for each mass condition. This special case is reflective of the monkey's strategy for that particular target as he would accelerate the cursor upward in a similar manner, grab the target, and essentially let gravity do most of the work to accelerate the cursor/target mass downward. The monkey would then produce an upward force at the last second to slow the combined mass enough to satisfy the Hold C period. In contrast, both the horizontal and vertical control components to other targets show evidence of adaptation to account for the differences in target mass as well as the resulting effect of

84

gravity.

On sessions using similar settings for gravity, cursor mass (100 units), and target masses (10, 100, 200 units) to those used in the previous figure, Figure 6.2 shows the same profiles for monkeys K and P under brain control. As seen from the figure, similar profile characteristics from the joystick experiments reappear during BCI operation. Movement A velocity and control profiles were similarly consistent across mass conditions to a given target. The relatively small range in the vertical cursor position demonstrates that the differences seen in the corresponding vertical control signals were used to produce flat trajectories for each mass. Again, both monkeys generally required more time to complete trials with heavier target masses. The relationship between total mass and trial time was not directly proportional, however, and was not always consistent across targets. For example, as seen for target 2 on Figure 6.2A, trial Movement B times for all three masses were similar around 2 seconds give or take approximately 0.25 seconds. If the monkey was simply to apply the same force level for each mass, the Movement B times for the heaviest two masses should be roughly 2x and 3x that of the smallest target, respectively. The relative proximity of the actual period lengths suggest that the monkey was accounting for the inertial properties of the cursor and targets, yet perhaps somewhat incompletely. Profiles for the other horizontal target (target 0) in the same subfigure show some of the same target preferences seen in Chapter 5 in which the monkey appears to struggle with consistent kinematic profiles to some targets more than others. The consistent movement times and vertical control signals for target 1 are consistent with the strategy observed in the previous paragraph in which the monkey simply has "put on the brakes" against gravity long enough to satisfy the Hold C period. The long trials for target 3 appear to be due to the dynamic range limitations also mentioned previously. Regardless of these nuances, there is visible evidence in both monkeys' control signals of adaptation both to the variation in the total mass under their control and the gravity field in which they operate.
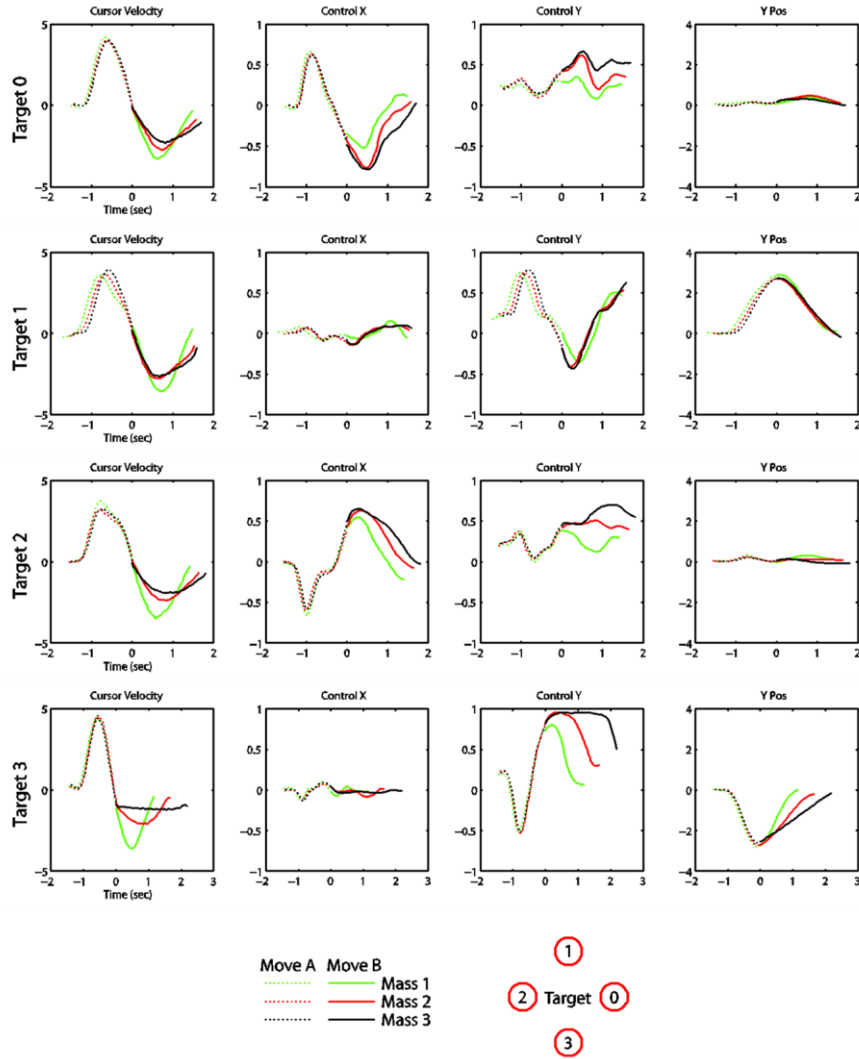
Figure 6.1 Mass-grab control characteristics under joystick control. Monkeys K and P initially trained on the mass-grab task using a joystick position-to-force control signal. Each row displays mean profiles for velocity, each control signal dimension (normalized joystick position), and the vertical position of the cursor as functions of time for each target location. Colored lines indicate the mass of the peripheral target, and dashed/solid lines indicate movement during Movement A and Movement B periods, respectively (t=0 indicates capture of the peripheral target and the beginning of Movement B). Velocity profiles are projected onto the center-to-target vector such that positive values indicate movement toward the peripheral target and negative values away from it. (Data from monkey K).

A
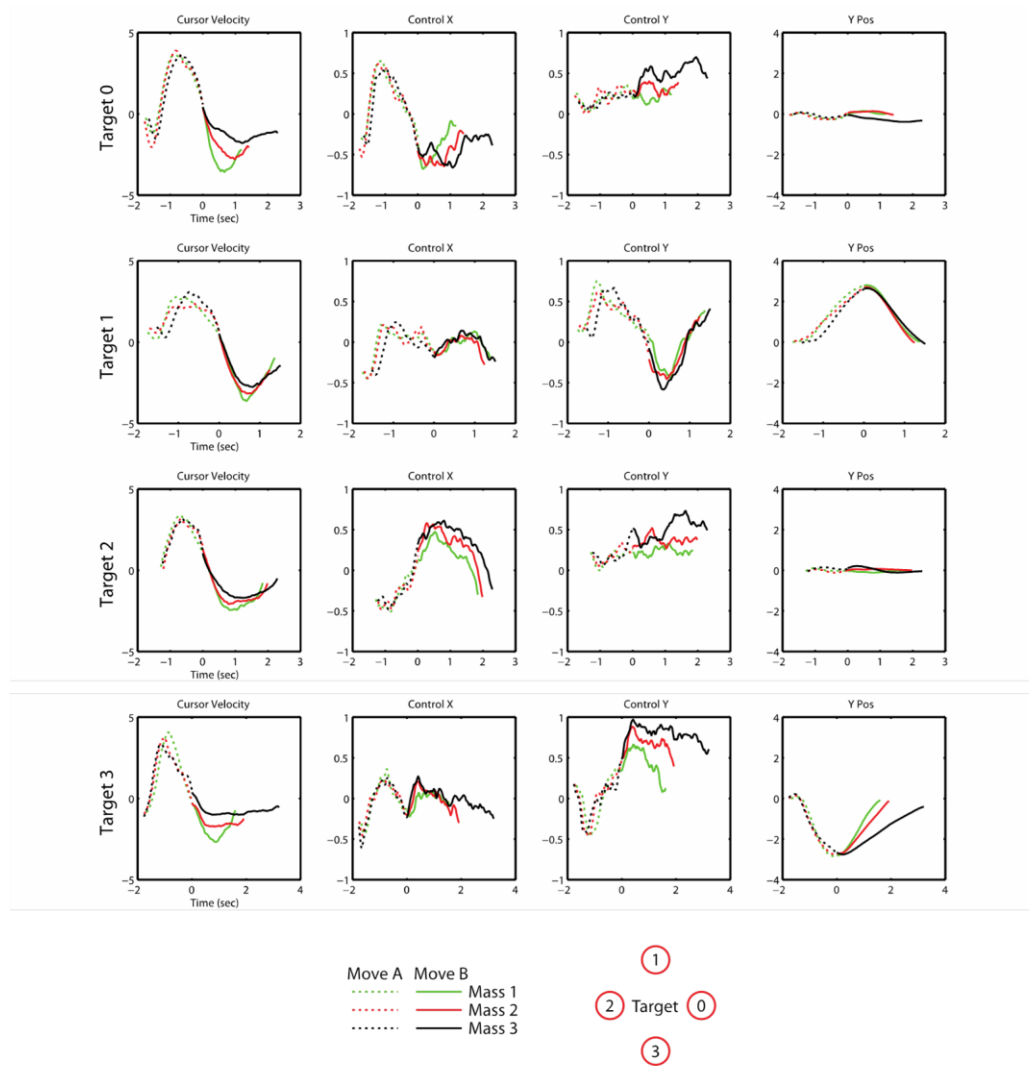
Figure 6.2 Mass-grab control characteristics under brain control. Using an ECoG-derived control signal mapped to a force on the cursor or combined cursor+target mass, each monkey performed the mass-grab task with gravity and mass conditions identical to those used for the joystick example in Figure 6.1. Plot descriptions are the same as in Figure 6.1. A) Control characteristics for monkey K. (continued on next page)

B
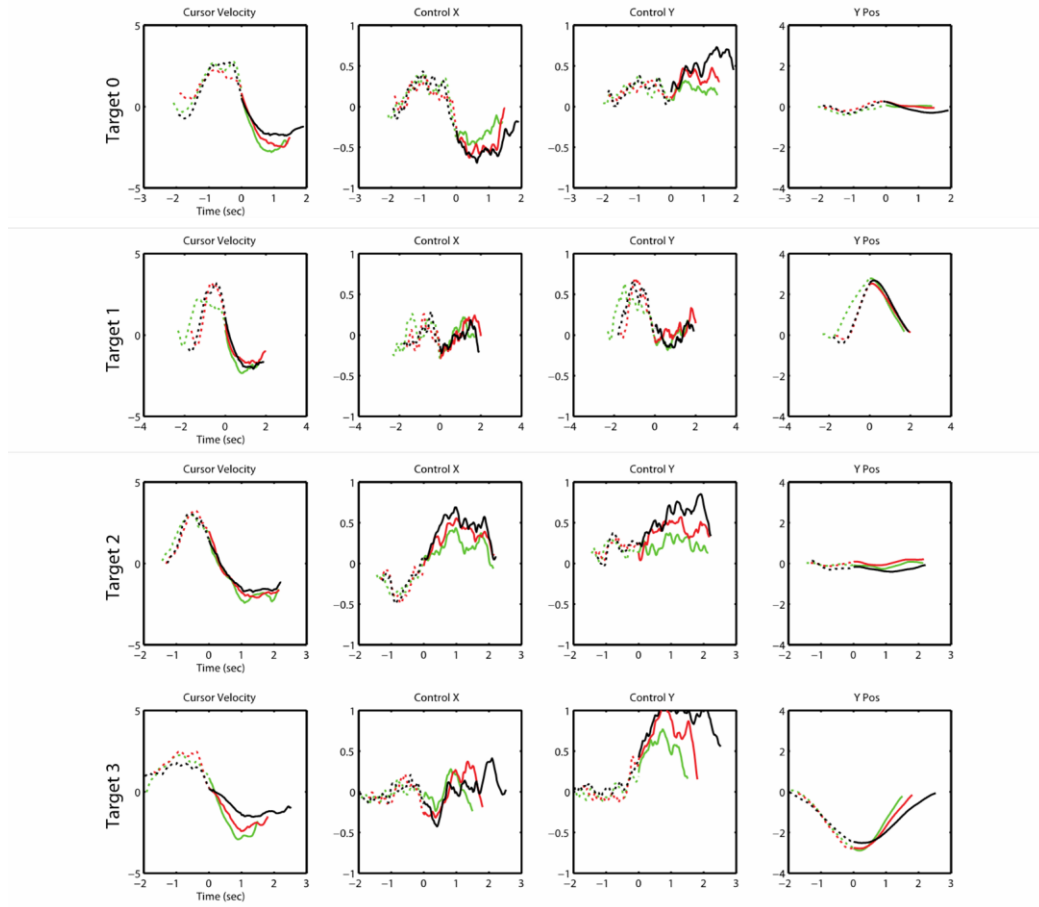
Figure 6.2 (continued). B) Control characteristics for monkey P.

The traces in Figures 6.1 and 6.2 show mean profiles over time. They do not provide an indication of how the distributions for each mass overlap, however. To examine the degree to which target masses could be distinguished from one another using only the neural control signal, the d' statistic was calculated between target mass permutations. Because the time length of trials varied between target masses, the d' metric was calculated on trial data binned by position. Figure 6.3 illustrates the d' metric as a function of position and mass for a horizontal target during joystick and BCI experiments. As seen from the figure, d' remains relatively close to zero during Movement A as the cursor mass is the same in each case. During Movement B, however, differences in the vertical control signal arise between target mass conditions

to account for the additional gravitational force. Between the two control modalities, joystick control demonstrated the most consistent differences in the vertical control signal with d' values near 1.5, 2, and 3 for combinations Mass 3/Mass 2, Mass 2/Mass 1, and Mass 3/Mass 1, respectively (A positive value indicates that the mean of the signal of interest was greater for the first mass listed than the second). d' values were smaller under brain control but were still generally greater during Movement B than Movement A. These smaller d' values were due to variance in the control signal as the differences in the mean control signal profiles between target masses appear similar across control modalities (see Figures 6.1 and 6.2A).
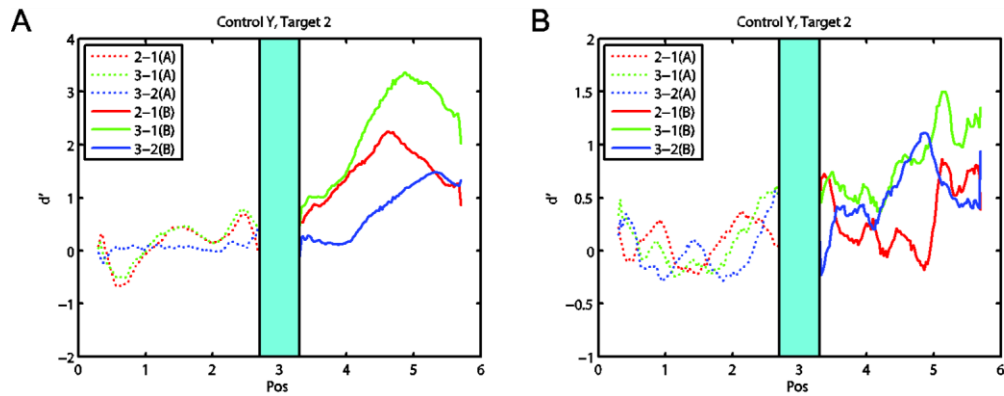
Figure 6.3. **Discriminability between mass conditions for vertical control signals.** **Discriminability indices, d', are plotted as a function of total cursor position between the three target mass pair combinations for a horizontal target. Colored traces indicate the mass pairing (i.e. 3-1 indicates comparing Mass 3 with Mass 1 with positive d' values indicating Mass 3 had a higher mean value) and the movement period (A - Movement A, B - Movement B). Position on the x-axis is total distance projected along the center-to-target vector. Cursor positions from 0-3 indicate the outward movement toward the peripheral target during Movement A, and positions 3-6 indicate the return movement to the center with 6 residing back at the center of the screen. Cyan boxes roughly indicate the Hold B period when the cursor touches the peripheral target. A) d' values using joystick control. B) d' values using brain control under conditions similar to those in (A).**

Despite these smaller d' values, it is still apparent that the subject is adapting his vertical control signal to compensate for deviations in the cursor's vertical position from zero. Figure 6.4A shows an example of the BCI control signal as a function of vertical cursor position while Figure 6.4B shows the regression lines of this type of plot for each target mass (all to target 2 as an example). The negative correlations observed in the figure suggest that the subject increases his vertical control signal to compensate for drops in cursor position and decreases the control signal in response to the cursor

rising. The constant offset of the lines from the origin are consistent with the presence of gravity, and the relative distance of these lines is consistent with the increasing target loads.
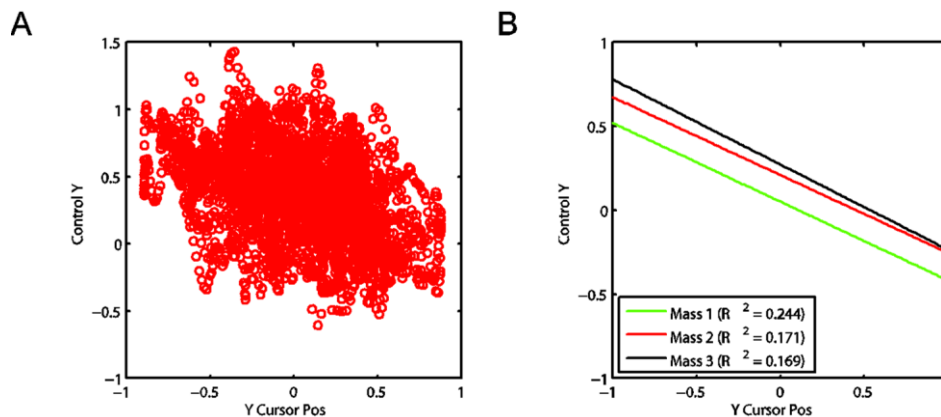
Figure 6.4. Control signal vs. cursor position during mass-grab task. A) Example plot of sampled vertical control signal data as a function of the cursor's vertical position. Data in this example was taken from trials from monkey K, target 2, Mass 2. B) Regression lines for the vertical control signal vs. vertical cursor position (monkey K, target 2). Colored lines indicate the target mass for the data fit, and R^2 values are provided for each fit. (All p-values = 0).

6.3.2 Gravity vs. Control Signal

Figure 6.5 displays the differences in control profiles for similar target mass and gain settings in the presence and absence of gravity for a horizontal target. As seen from this figure and previous figures where gravity was present, a separation in vertical neural control signals can be observed during the Movement B period. In the absence of gravity, these separations disappear, and the vertical control signal remains relatively constant near zero during both movement periods. However, examination of the horizontal control signals shows a similarity between the two panels in that a degree of adaptation is observed with increasing target mass. These results demonstrate that the subjects can adapt to both dynamic (inertial) and static (gravity) force loads. In addition, compensation for these forces was accomplished separately from each other using only visual feedback of the cursor. In other words, adaptation of the horizontal control signal to increased target loads during horizontal target trials did not accompany adaptation of the vertical control signal when no gravity was present. Although these

findings were discussed in the context of different types of forces likely to be present in a force-based BCI system, this decoupling/decorrelation of control axes is an important characteristic to develop for BCI control of multiple degrees of freedom regardless of the framework in which it is employed.
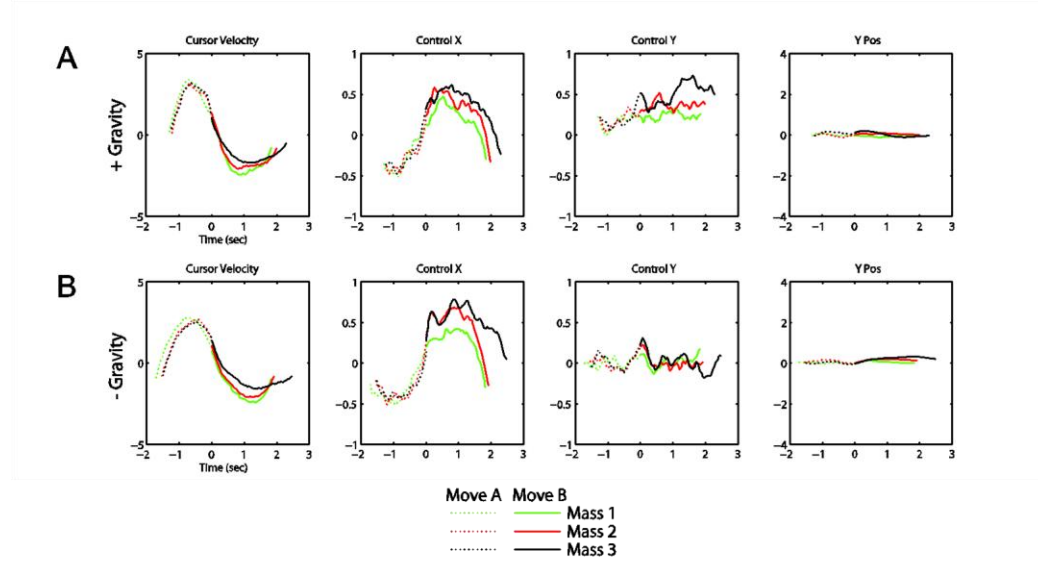
Figure 6.5.   BCI control characteristics with and without gravity.   Mean movement parameter and BCI control signal profiles similar to those described in Figures 6.1 and 6.2 to a horizontal target (target 2). A) Profiles with simulated gravity present. B) Profiles with no gravity included.

6.3.3 Dynamic Range of Control Signal

One important design consideration brought to light during the course of these experiments was choosing an appropriate operating point for the ECoG control signal. The profiles for monkey K in Figure 6.2A were extracted from a training day in which the gravity for the session was set to a value approximately twice that of previous days' settings. For the initial session on that day, the monkey's normal settings for the force gain term, g_F, were used. Figure 6.6 depicts profiles from that session similar to those shown in Figures 6.2. From the figure, it can be seen that the peak of the control signals for each mass during the Movement B period often reside in the range between 0.8-1.0. During this session, it was obvious that the monkey struggled more with accelerating the heavier masses, maintaining vertical cursor position for horizontal

targets, and returning the cursor from the bottom target to the center. In fact, the monkey was unable to finish the block from which this figure was taken due to his inability to complete trials to target 3 with the largest target mass. After this failed session, g_F was increased by approximately 40%. The monkey was then able to consistently complete trials to all targets as shown in Figure 6.2A. Referring back to that figure, increasing the gain term set the vertical control signals for horizontal targets into a range near 0.2-0.7, allowing the monkey to more accurately compensate for gravity. Trials to target 3 still demonstrated vertical control signals maxed out near 0.8-1.0, but this control signal level was still sufficient to move the heaviest targets upwards, even if at a reduced speed.

The previous results suggest an absolute range within which the monkey could operate. Reexamining the statistics for the profiles in Figures 6.2 also lends useful insight into the dynamic range of the system. For a given mass and target condition, the mean standard deviation of each control signal profile (horizontal or vertical) averaged across a given movement period ranged from 0.25 to 0.31 units. This variance in control signal was consistent across targets, target mass, control dimension, and movement period, suggesting it could be an inherent property of the control signal (i.e. possible the "noise floor" of the control signal). The mean vertical control signals for each mass on horizontal targets were separated from the next closest target mass by 0.15 to 0.20 units.

Combining the range of standard deviations with the observed range of the control signal ([-1 1]) would imply that this ECoG-derived control signal can support a range of 3-4 standard deviations per dimension (6-8 if negative and positive values are considered together) and similarly a maximum d' value of 3-4 (6-8) for discrimination of binary force values (0 or ± F1). Thus, this information could be used to adjust the gain of the system to achieve a desired discriminability for given gradations in force. For example, one might desire to have a signal-to-noise ratio (i.e. d') of 1.0 for increments in
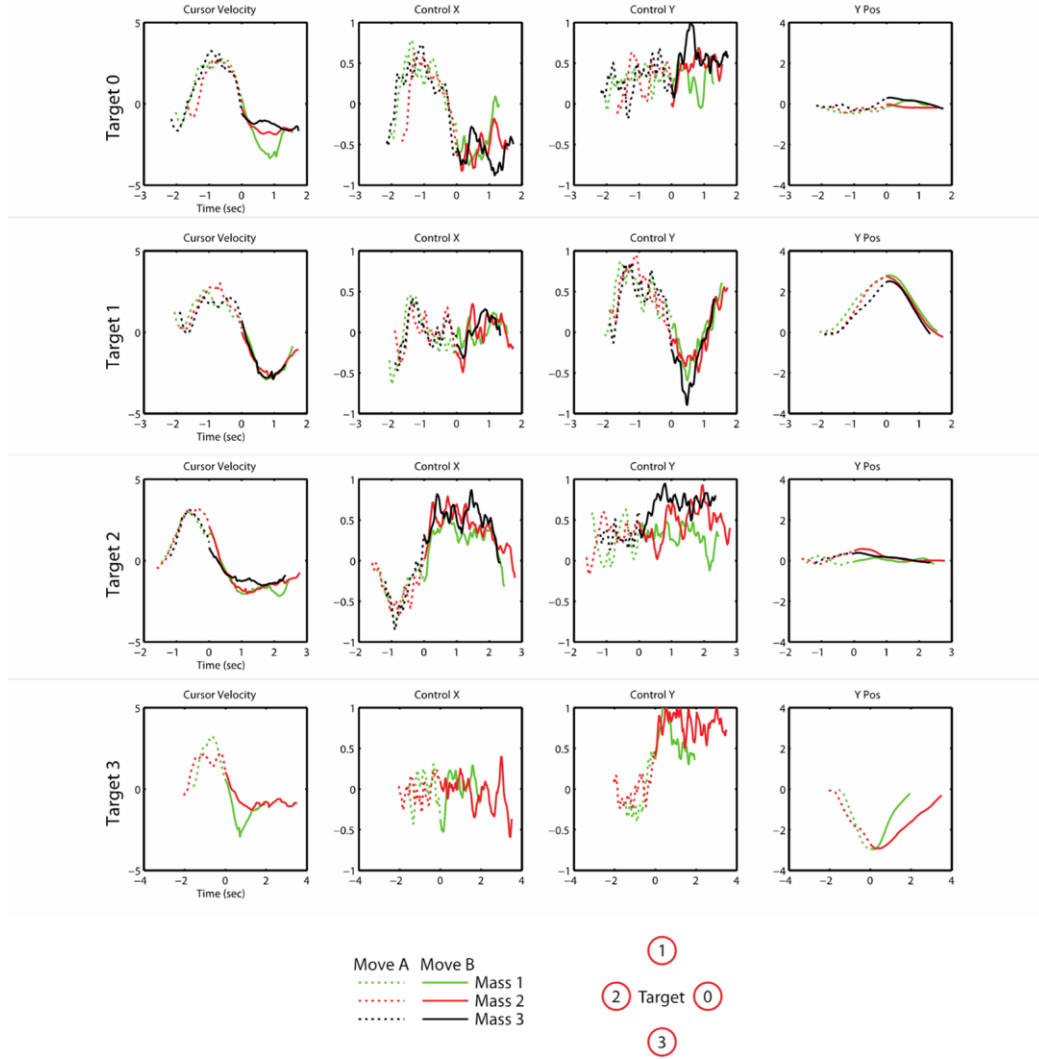
Figure 6.6 Control characteristics under low gain conditions. Monkey performed a mass-grab task using gravity and mass settings identical to those used in Figure 6.2A but with a lower control signal-to-force gain term. Profiles for Mass 3 in target 3 are absent due to the monkey's inability to successfully complete trials to that target/mass combination with the gain and gravity settings. Comparison with Figure 6.2A reveals that a higher gain constant produced a better dynamic range for the given gravity and target mass settings.

force of 1 N. Consequently, this would limit the range of the system to ± 4 N. Obviously, satisfaction of both the dynamic range and sensitivity specifications are not guaranteed. Increasing the dynamic range for decoding of larger forces likely comes at the cost of decreased sensitivity in the production of smaller forces. Analogous real-life systems for these differing force requirements could be a hand/finger grip system requiring fine and accurate changes in force for handling small or delicate objects versus

a system that may require production of larger forces but may also have larger acceptable variation in force such as the proximal arm (i.e. carrying a basketball to lifting a heavy suitcase).

6.4 Conclusions

The results presented here provide a novel demonstration of a force-based BCI control signal. Rather than using brain-derived signals to simply point to or select a target as has been commonly demonstrated in most BCI studies to date, subjects in this study were forced to use a similar brain signal to interact with and account for varying environmental force conditions. Although the task and range of conditions was relatively simple compared to those experienced in real-life, the results do present a promising first step in that direction.

The mass-grab task examined in this chapter was initially discussed in the context of applying it toward FNS of a natural arm for reaching and grasping. However, these results present broader implications for BCI applications. The decoding model used by both monkeys was initially trained to produce a control signal suitable for velocity control of a cursor. Once this model was fixed for the acceleration mapping task in Chapter 5 and the mass-grab task in this chapter, the monkeys successfully adapted these base control signals to match the context of the BCI task. Thus, as alluded to in the last chapter, the particular variable to be decoded may not be the deciding factor in the success of a BCI system. Rather, basic properties of the control signal itself such as d' measures are likely more important to optimize.

Continuing along this ideological framework, a wider array of possible BCI applications begins to materialize. For example, imagine designing a BCI system for a person paralyzed below the waist who wishes to drive an automobile once again. Using conventional BCI wisdom, one might design a decoder to decipher the angle or velocity of the ankle joint as the foot presses on the gas pedal for a desired acceleration. An inverse model would then be used to compute the appropriate forces to realize the decoded foot movements, and these foot movements would then be translated through the gas pedal and the car's engine to the car's acceleration. The first two steps of this

system could be eliminated, however, if BCI force control signals were relayed directly to the car in lieu of the gas and brake pedals.

Along these lines, BCI signals might not only be applied to restore lost function as has been traditionally championed, but also to augment natural function. Inferring intent directly in this manner from cortical sources might decrease delays associated with transmission through the central and peripheral nervous systems. One example of this augmentation proposed recently would be to decrease reaction braking time by inferring this intent directly from EEG signals [72]. Utilizing brain signals in this manner might greatly expand the current breadth of BCI research and present exciting new possibilities for human-computer interaction.

7 Differentiating BCI Intention from Rest

7.1 Introduction

In this study, we examine the spectral differences in epidural micro-ECoG signals between active BCI control and resting states. As in the experiments by Ludwig et al described in Chapter 2 [47], no pre-screening for movement related features is performed beforehand, and our adaptive control strategy allows the monkeys in these experiments to develop their own neural states for completing a BCI task. However, we do not force these neural states to be distinct from baseline spectral characteristics. Thus, we are able to examine differences that might arise naturally between control and resting states when no *a priori* movement features are assumed. We also assess the degree to which these states can be decoded on both a trial-by-trial and real-time basis, as well as the robustness of these models over time. The results presented here could provide valuable insight into the neurophysiology of BCI control as well as important considerations toward the design of an asynchronous BCI system.

7.2 Methods

7.2.1 BCI Tasks

Subjects performed two closed-loop tasks during the course of these experiments. Monkeys K and P used ECoG-derived control signals to complete the radial-choice and center-out tasks as described in Chapter 5. The analyses in this chapter focus on data sets from which the monkey used a control signal-to-cursor velocity mapping. These data sets were also taken from sessions after initial brain control training had been completed such that no cursor bias was used to assist the subjects. For each task, peripheral targets were chosen in a pseudorandom block design such that 15 trial repetitions to each target must have been successfully completed before a block was complete.

At the end of a task block, a block of "rest trials" was initiated in which the

background of the subject's virtual arena remained constant, but visual feedback of the ECoG controlled cursor was removed and no rewards were provided. With no visual feedback or reward motivation, these trials were intended to disengage the subject from actively modulating his ECoG signal for BCI control. Since no task related events were present during these trials, individual trial lengths were fixed at 3 seconds. No indication of transition from rest trial to rest trial was presented. After 20 trials (~60 seconds), the rest block ended and the beginning of a new task block was signaled by the appearance of the cursor and target ring/sphere.

7.2.2 Task vs. Rest Amplitude Differences

To examine the gross differences between the active state of participation in a BCI task and a resting state, we inspected the amplitude differences of the spectral features used for BCI control. Trial averages of the band amplitude estimates on each channel were compared between task and rest states using the discriminability index metric, d':

$$d'_{task,rest} = \frac{\bar{A}_{task} - \bar{A}_{rest}}{\sqrt{p_{task}\sigma^2_{task} + p_{rest}\sigma^2_{rest}}} \qquad (7.1)$$

where A is the mean band amplitude for a given state (task or rest), σ^2 is the state's variance in spectral amplitude, and p is the proportion of trials for that state. The d' metric is a classic measure of signal detection theory that describes how well an observer could discriminate between two signal states using a given feature.

7.2.3 Intra-task Amplitude Modulation

To compare and contrast the characteristic ECoG modulations found between task and rest trials with modulation patterns used within the BCI task itself, we calculated a modified version of the d' statistic based on a feature's modulation with target direction. Normally, d' is used to assess the ability to discriminate a signal from noise or two distinct signal distributions from each other. For signals related to target direction, however, the variable of interest is two-dimensional and continuous. In order to adapt this metric, we first regressed each feature's spectral band amplitude against

target direction to fit the following model:

$$X_{band,chan} = 0.5A_{pp}\cos(\theta_{target} - \theta_0) + b_0 \qquad (7.2)$$

where X is the predicted feature amplitude, A_{pp} is the peak-to-peak amplitude modulation of the feature, θ_{target} and θ_0 are the presented and preferred target directions, respectively, and b_0 is a constant offset term. If the spectral feature X is well fit by this cosine tuning model, A_{pp} can be approximated by the mean difference in amplitude between the preferred and anti-preferred target directions. A_{pp} can then be applied to the compute the directional or target discriminability index within task trials:

$$d'_{target} = \frac{A_{pp}}{\sqrt{\sum_{j=0}^{N-1} p_j \sigma_j^2}} \qquad (7.3)$$

where N is the total number of possible targets, σ_j^2 is the variance in amplitude for target j, and p_j is the proportion of trials to that target.

7.2.4 Task vs. Rest Decoding Model

After examining differences in modulation between BCI task and rest states within individual features, we sought to reveal how well the working intention of the subjects could be decoded on both a trial-by-trial as well as a real-time basis using the entire feature ensemble. Spectral features were normalized in the same manner as before for cursor control. Task and rest trials/samples were labeled with ground truth values of y={+1, -1}, respectively. The model used previously for directional decoding was recast for the purposes of binary classification. Raw decision values for predicting whether task or rest periods occurred were given by:

$$y_{raw}(t) = w^T X(t) + b_0 \qquad (7.4)$$

and final classification labels as:

$$y_{task/rest}(t) = sign(y_{raw}(t)) \qquad (7.5)$$

Again, w_{160x1} is the task/rest state decoding matrix, $X(t)_{160x1}$ is the normalized feature set, and $b_{0_{1x1}}$ is a constant term that serves as a decision threshold in this classification setting. The model was regularized and validated in the same manner as for the cursor movement model (leave-one-out cross-validation (LOOCV) for trial-averaged data or

10-fold cross-validation with real-time center-out task data). In addition, each model was tested on a set of trials from the same day exclusive of those used in the model's construction.

7.2.5 Model Performance Evaluation

The threshold term, b_0, in equation (7.4) can be adjusted to tune certain accuracy metrics toward the desired usage of a BCI device. For example, lowering b_0 should result in a lower false negative rate (fewer task periods incorrectly classified as rest periods), while raising b_0 should result in a lower false positive rate (rest periods incorrectly classified as task periods). The first scenario might be preferable in situations where turning off a device prematurely might present safety concerns such as possible falls, and the latter might be preferable for preventing undesired movements during sleep. In either case, final accuracy metrics such as false positive and negative rates become more akin to design considerations and do not necessarily encompass the model's discriminative capability.

In light of this, the state decoding models constructed during these experiments were evaluated using Receiver Operator Characteristic (ROC) analysis in addition to total accuracy levels. ROC analysis is a non-parametric characterization of a binary classifier's discriminatory power by sweeping out a series of threshold levels for the classifier's output metric and classifying all samples above the threshold to the "positive" class (i.e. task samples in these experiments) and all samples below threshold to the "negative" class (i.e. rest samples). At each threshold level, the proportion of true positive instances that are correctly classified as positives (true positive rate, TPR) is plotted against the proportion of true negative instances that are incorrectly classified with a positive label (false positive rate, FPR). Sweeping out a series of thresholds that spans the entire range of the classifying signal of interest produces a curve ranging from (0,0) (all samples classified as negatives) to (1,1) (all samples classified as positives). The shape of the curve between these extremes indicates the degree of overlap of the two class distributions and the expected performance of the classifier. States with completely overlapping distributions will follow the diagonal line from (0,0) to (1,1) and

show chance performance in predicting the correct class. States with completely separated distributions will trace a vertical line from (0,0) to (0,1) and a horizontal line from (0,1) to (1,1) and will show perfect performance in using the signal of interest to predict a sample's state. Between these curves, the performance of a classifying signal can be described by the Area Under the Curve (AUC) metric, the area between the ROC trace and the x-axis. Perfect classifiers have an AUC of 1.0 while chance performance classifiers have an AUC of 0.5. Classifiers with an AUC less than 0.5 (curve below the 45° diagonal) do still hold predictive value, but it indicates that scores for true positive instances are likely to be less than those for negative instances. In this case, one should actually label samples opposite what the classifier would suggest.

Although we chose to examine the ECoG spectral features used in our cursor movement decoding model as the basis for our task/rest model as well, there are other obvious features that an outside observer could potentially use for trying to determine the subject's attention to the BCI task. For example, even though a desired target on the computer screen of a subject using a BCI to move a computer cursor may not be readily apparent to a bystander as it is in our experimental setup, observing the movement of the cursor itself might give some indication of whether the BCI user is actively using the BCI cursor or not. Due to this consideration, we compared the performance of our ECoG-based classifier with that of a simple classifier using the cursor's speed as the raw classification metric. Because the speed distribution of the cursor varies with different periods of the center-out task (i.e. Hold A period, Movement period, etc.), we also compared these periods individually against rest periods using both neural and speed features.

7.2.6 Temporal Filtering of Raw Classifier Output

Because the linear summation of noisy features in our model still produces a relatively noisy raw decision value output and the frequency of shifting between task and rest blocks is much lower than the frequency content of the classifier output, we explored whether low-pass filtering the raw decision value could improve classifier performance. The raw decision value, $y_{raw}(t)$, was passed through a 1st order low-pass

Butterworth filter. The low-pass frequency was swept logarithmically from 10^{-3} Hz to 10^1 Hz, and the AUC of the filtered decision value was evaluated at each frequency for comparison with the unfiltered decision value.

7.2.7 Temporal Analysis of State Transitions

To determine the possible time course of each monkey transitioning from actively participating in the BCI task and rest and vice versa, we examined the temporal characteristics of the raw decision value output during these transition periods. For the rest-to-task transition, the raw decision value traces were overlaid from 10 seconds before the beginning of a task block to 10 seconds after. The response time (RT) of this transition was defined as the average time of the decision value to go from its initial value at the beginning of the task block at t = 0 to classification as a task sample (the first zero crossing). For the task-to-rest transition period, the raw decision value traces were overlaid from 20 seconds prior to the start of a rest block through the duration of the rest block (60 seconds). ROC analyses were then performed for individual time segments of the rest block (e.g. 20-30 seconds after the start of each rest block) versus all task samples to inspect the discriminability of the decision value as a function of time into the rest block. Each of these temporal analyses were repeated for decision values low-pass filtered at selected frequencies.

7.3 Results

7.3.1 Intra-task and Task/Rest Feature Modulation

Figure 7.1 displays the spatial distribution of absolute directional discriminability indices, d'_{target}, and task/rest discriminability indices, $d'_{task,rest}$, within the five spectral bands used for BCI control. As seen in Figure 7.1A, all of the frequency bands used for control demonstrated some degree of modulation in average spectral amplitude between task and rest trials. However, features in the alpha (8-15 Hz) and beta (15-24 Hz) frequency ranges on a given channel tended to show greater discriminability between task and rest trials than their high gamma band (70-150 Hz) counterparts. Conversely, Figure 7.1B demonstrates much greater target discriminability within task trials in the
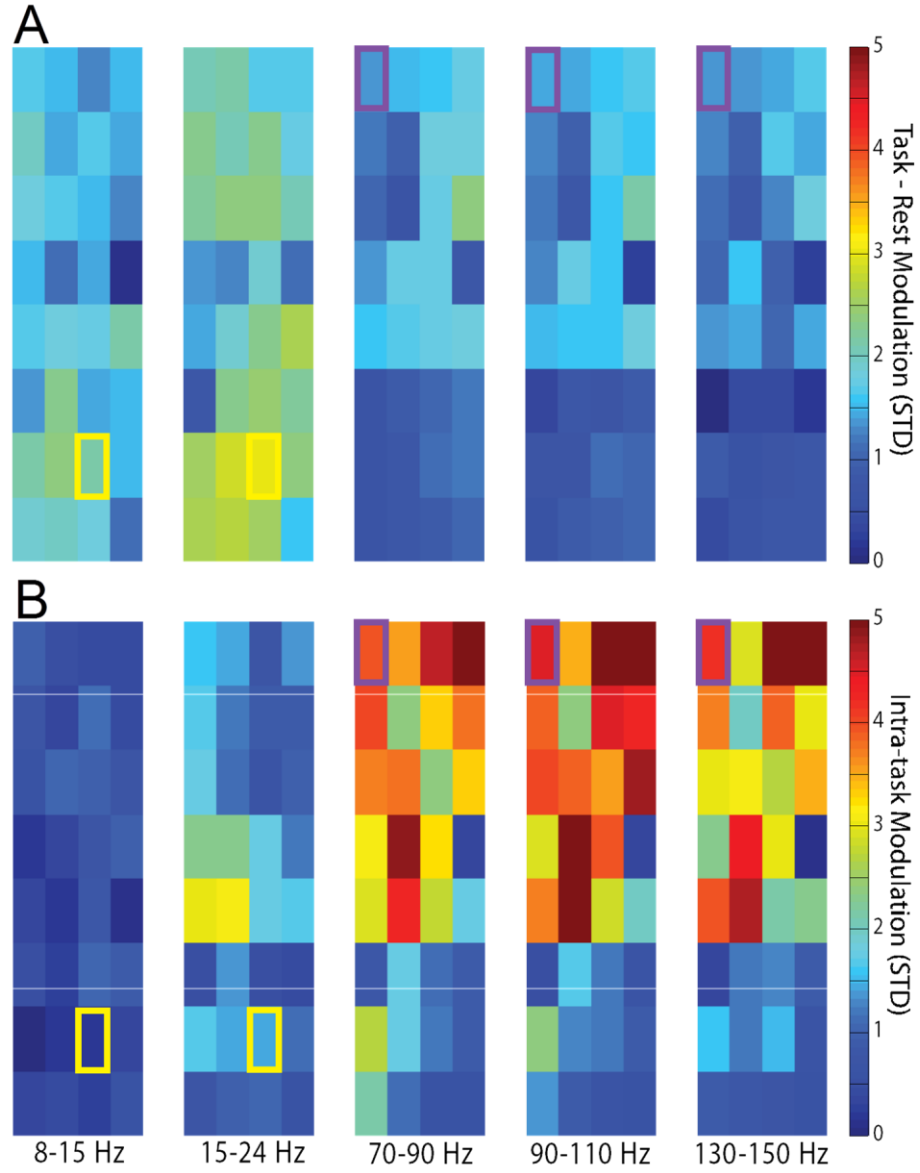
Figure 7.1. Spatial task vs. rest and intra-task modulation. Panel (A) shows spatial channel heat maps (each 4x8 colored array) of absolute normalized modulation between task and rest trials within each of the five frequency bands used for control. Panel (B) shows channel heat maps of absolute task-related modulation. As seen from the differences between frequency bands, high gamma control bands (70-150 Hz) exhibit strong task-related modulation while lower alpha (8-15 Hz) and beta (15-24 Hz) bands show weaker task modulation. Conversely, the lower alpha and beta bands show stronger separation between rest and task states. Purple boxes highlight one channel that shows strong task-related modulation in the gamma frequency control bands but less task-rest separation, while yellow boxes highlight a channel that shows stronger task-rest separation in the alpha and beta frequencies but less task-related modulation (see Figure 7.2). (Data selected from monkey K, right hemisphere)

high gamma frequencies compared to the lower frequency features. Figure 7.2 further illustrates on a trial-by-trial basis these spectral differences in discriminability for the two exemplar channels outlined in the previous figure. From the strongly task/rest modulated low-frequency features on one channel depicted in Figure 7.2A, a dramatic increase in spectral amplitude can be observed during rest blocks. However, the task trials grouped by target demonstrate negligible modulation with target direction. For the strongly task modulated channel (Fig 7.2B), these trends are reversed in high gamma features as power drops during rest blocks while clear modulation across target direction is evident.
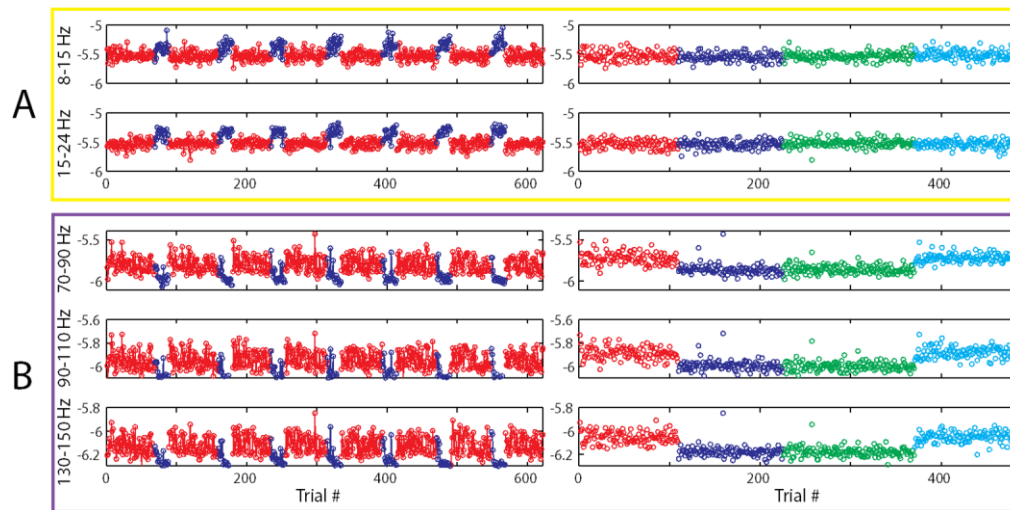
Figure 7.2. Examples of task modulated and rest/task modulated spectral features. *Rest/Task Modulated Channel (Yellow):* **Box A illustrates individual trial band powers for the channel features outlined in the yellow boxes in Figure 8.1 (8-15 Hz, 15-24 Hz). The left panel illustrates trial-by-trial band powers over the course of a day's experiment with task trials in red and rest trials in blue. The right panel shows task trials grouped by the target presented (0, 90, 180, and 270 degrees as red, blue, green, and cyan, respectively). As seen on this channel, there is a significant change in band power between task and rest channels, manifested as an increase in power in the lower frequency bands during rest trials. However, very little modulation within these frequency bands is shown between targets in the right panel.** *Task Modulated Channel (Purple):* **Box B illustrates strongly task modulated channel features outlined in the purple boxes in Figure 8.1. This channel exhibits some power differences between task and rest in the gamma frequency range, but less separation than that seen in the lower frequencies in the first example channel. Unlike the lower frequency bands of the top channel, the higher gamma frequencies on this channel do show modulation between targets. (All vertical axes are in log(V²)).**

7.3.2 Trial-by-Trial State Decoding

An example of trial-by-trial performance of a task/rest state decoding model for one monkey hemisphere is portrayed in Figure 7.3. This model was trained and evaluated on separate data sets from one day after the monkey had demonstrated proficiency at the brain-controlled center-out task, and the same model was again evaluated on trial data from a training day two weeks later. Total accuracy for day 1 in Figure 4 was 0.98 (true positive rate (TPR) = 1.0, false positive rate (FPR) = 0.08). Testing the same model on a new data set recorded two weeks later revealed no appreciable loss in accuracy (total accuracy = 0.99, TPR = 1.0, FPR = 0.05). This generalizability of a model trained on a given day to future dates was consistent across subjects.
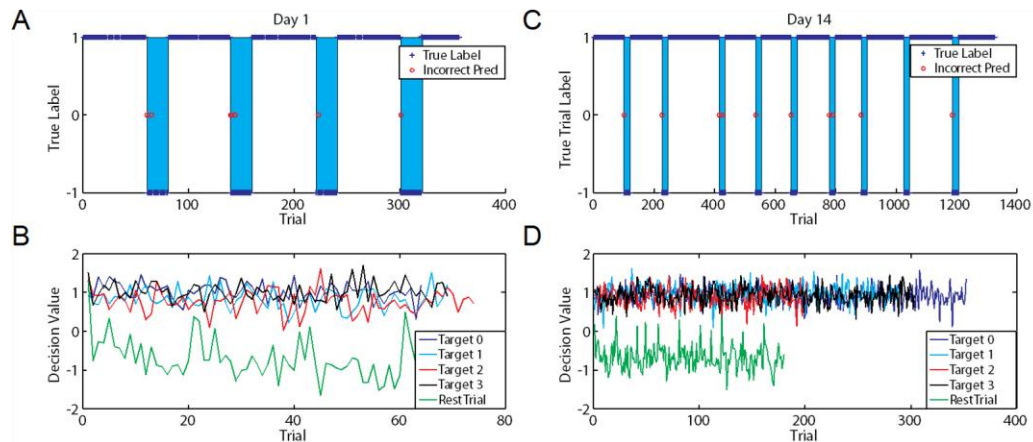
Figure 7.3. Linear model trial-by-trial decoding on separate weeks. An L1-regularized least squares model was trained to predict rest/task trials given only ECoG trial-averaged control band amplitudes (monkey K, right hemisphere). A) Classification performance of the model on a separate test set from the same day as the original training data. Blue crosses indicate the actual trial type (+1 = task, -1 = rest), and rest trial blocks are shaded in light blue. Red circles indicate an incorrect model prediction. As seen in this panel, the model often misses trials at the beginning of a rest trial block. This could be consistent with the monkey anticipating another task trial to appear. B) Raw classifier output from model used in (A) for trials to each target. This panel demonstrates a distinct separation between rest trials and all target trials regardless of the desired target. (C-D) Prediction results as in (A-B) using the original model on a data set taken two weeks later. As seen here, the model retains its predictive capabilities and separation between rest/active trials, and again demonstrates incorrect predictions at the beginning of rest blocks due to possible residual attention to the task.

Figure 7.3B and D also depicts the raw decision values of the classifier model for task trials grouped by the presented target. This figure shows a distinct separation

between rest trials and trials to all targets. This observation is evident despite the fact that the model was trained with the objective of separating rest trials from task trials as a whole rather than on a target-by-target basis. It should also be noted that the rest trials occur in blocks of 20, so the spikes in raw decision value observed in the bottom panels of the figure coincide with the beginning of rest blocks (*i.e.* near rest trial numbers 0, 20, 40, etc.).

7.3.3 Real-time State Decoding

Figure 7.4 summarizes the ROC curves of real-time state decoding models based on ECoG features from each of the three monkey-hemisphere combinations used for cursor control in these experiments. These example models were trained on data well after each monkey had achieved accurate control of the BCI cursor. In addition, the figure also shows the ROC curves using the absolute cursor speed as the decision criterion. In each case, the ECoG-based model offers much better discriminability for detecting active usage of the BCI than observing the BCI cursor speed alone. Interestingly, in the case of monkey P, the cursor speed ROC curve lies below the chance level diagonal, indicating that the cursor actually moves faster during rest trials.
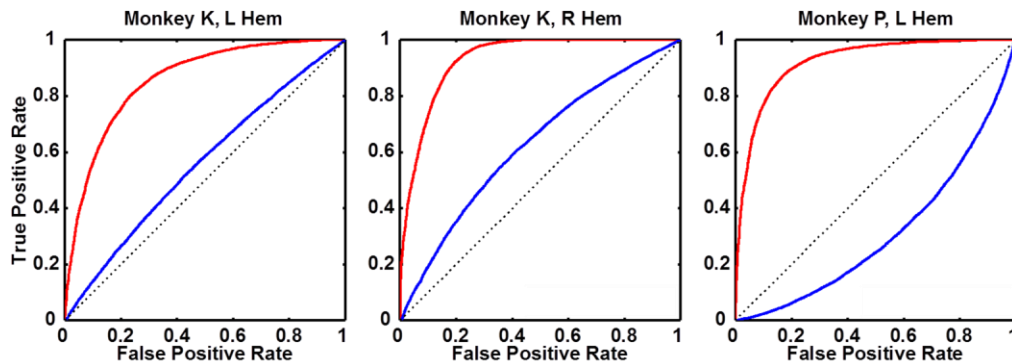
Figure 7.4. Real-time decoding. Each panel depicts Receiver-Operator Characteristic (ROC) curves for real-time discrimination between task and rest states for the three monkey hemispheres in real-time using either the output of the task/rest classifier (red) or the output cursor speed from the brain-derived velocity control signal (blue). While simply observing the output speed of the brain-controlled cursor on the screen does offer some insight into whether the subject is actively engaged in a BCI task, the task/rest state classifier offers much better separation between states in each case.

We further dissected the relationship between cursor speed and state discriminability by examining the cursor speed and EGoG-based decision value during different periods of the behavioral task. Figure 7.5A-B depict the relationship between the cursor speed vs. ECoG-based decision value during the "Hold A" and "Movement" periods, respectively, for Monkey K. These two periods represent polar opposites in terms of probable speed profiles and target uncertainty. During the Hold A period, the monkey has no information about at which location the target will be displayed. Thus, it is to his advantage to hold the cursor as still as possible in the center during this period until target information becomes available. Conversely, during the Movement period, the monkey has complete knowledge of the target location but a limited amount of time to reach it, encouraging him to move as fast as possible without sacrificing accuracy or the ability to stop for the "Hold B" period. Hence, one might expect a difference in discriminability during these two periods using only cursor speed if the periods were analyzed separately. After regressing real-time decision values from the ECoG model against coincident cursor speeds as seen in Figure 7.5A-B, the movement period had a slope significantly different than zero at a significance level of $\alpha = 0.05$ (*i.e.* the decision value indicated a sample was more "task-like" with increasing cursor speed). During the Hold A period, however, the ECoG decision value did not trend
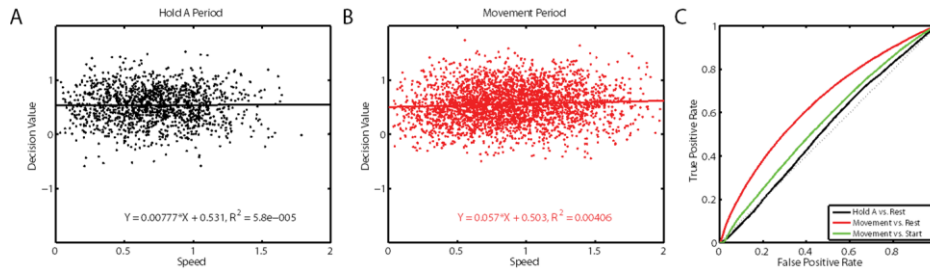
Figure 7.5. BCI cursor movement characteristics. Decision value output from the rest/task classifier versus the corresponding ECoG-derived cursor speed during A) Hold A and B) Movement periods over the course of one center-out task day (monkey K, right hemisphere). Equations give the least-squares model fit of the decision value (Y) and speed (X) data, and the proportion of variance in decision value explained by cursor speed (R^2). The slope of the model for Movement periods was significantly different than 0 at an $\alpha=0.05$ significance level, while the slope for Hold A periods was not (p=0.27). However, both models accounted for less than 1% of the variance in the rest/task decision value. C) ROC discriminability analysis between individual periods using cursor speeds. Although the Movement periods show slightly better discriminability versus rest than Hold A periods, both underperform the ECoG model ROC curve in Figure 7.4's middle panel.

significantly with cursor speed. In both cases, cursor speed accounted for less than 1% of the variance in the ECoG decision value. Figure 7.5C illustrates the ROC curves when cursor speeds during that period are compared with rest period speeds. As seen by the AUC of each curve in this panel, Movement period speed samples offered more discriminative capacity against rest samples than Hold A samples. More importantly, this increased discriminability is still considerably less than that of the ECoG-based model (see Figure 7.4, middle panel).

7.3.4 Decoding Improvement with Filtering

While the real-time decoding model does offer significant predictive power of which trial type the monkey is currently engaged in or observing, the accuracy results demonstrated by the model using trial-averaged features suggests that temporal smoothing of the real-time model output could offer improved performance. Figure 7.6 illustrates the improvements and detriments in classifier performance after low-pass filtering the raw ECoG-based task/rest decision value at various frequencies. As the band amplitude estimates for each feature were low-pass filtered at 3Hz, low-pass
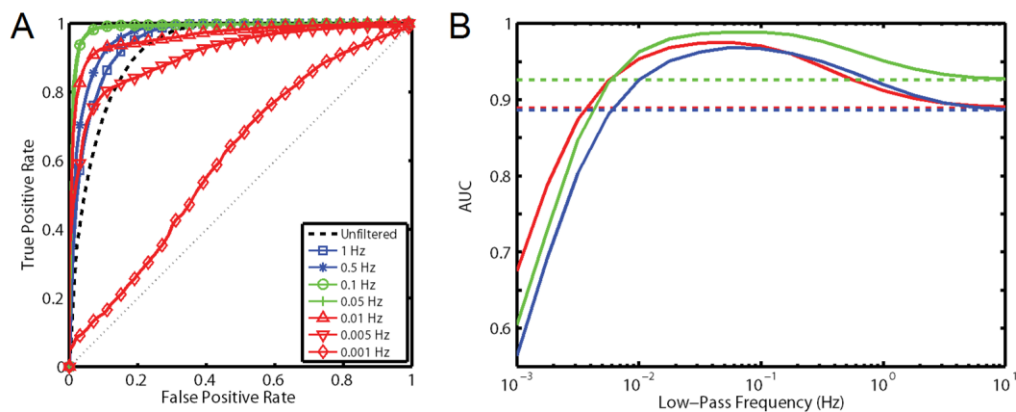
Figure 7.6. Low-pass filtering effect on classifier performance. A) ROC curves for low-pass filter cutoff frequencies from 0.001 to 5 Hz for monkey K's right hemisphere, and (B) the corresponding AUC metric as a function of cutoff frequency for all three monkey hemispheres (red – monkey K, left hemisphere; green – monkey K, right hemisphere; blue – monkey P, left hemisphere). The horizontal colored dashed lines in (B) represent the unfiltered AUC performance for the correspondingly colored solid curve. Filtering the classifier output with a cutoff frequency between 10^{-2} and 10 Hz shows improvement in classifier performance, while a maximum in AUC near 0.1 Hz to 0.01 Hz appears to be optimal for the typical task/rest durations used in these experiments.

filtering the ensemble raw decision value near or above this cutoff frequency offered negligible if any improvement in the AUC metric. Low-pass filtering at lower frequencies, however, did offer monotonically increasing discriminative power over the unfiltered decision values for each monkey hemisphere until reaching a plateau between 0.01 and 0.1 Hz (Figure 7.6B). Below approximately 0.01 Hz, filtered AUCs declined beneath that of the unfiltered decision value.

7.3.5 Temporal Characteristics of State Transitions

The temporal characteristics of transitioning from rest blocks to task blocks and from task to rest blocks are shown in Figure 7.7. The response time of the unfiltered decision value transitioning from a rest block to a task block (the average time after the beginning of a task block for the decision value to increase above the 0 threshold) was approximately 400 ms (Figure 7.7A). Referring to the results in Figure 7.7B, a low-pass cutoff frequency of 0.05 Hz was chosen that approximately resided on the plateau of the AUC vs. cutoff frequency for each monkey hemisphere. This cutoff frequency was used to filter the traces in Figure 7.7A. Low-pass filtering the decision value at decreasing cutoff frequencies decreased the number of false positives (Figure 7.7C). However, this improvement in classifier performance comes at the cost of an increased response time. For example, the response time at a near-peak AUC frequency (0.025 Hz) was nearly 5 seconds while that of a decision value filtered at 0.25 Hz was less than 1 second. The other transition to consider within these experiments is the transition from task blocks to rest blocks. Figure 7.7B, portrays the unfiltered time course of transitioning from task to rest blocks. As seen from the decision value traces over time, the time constant of decaying to a "resting" state appears to be approximately 20 seconds during these experiments. For this transition, low-pass filtering of the decision value does not significantly impact the time constant of decay as it did for transitioning from rest to task. This was further validated by an ROC analysis shown in Figure 7.8 in which samples taken from segments of time later and later in the rest blocks results in greater discriminability from task samples. For samples taken after 30 seconds into the rest block, the ROC curves demonstrate nearly perfect discriminability with AUCs near

unity.
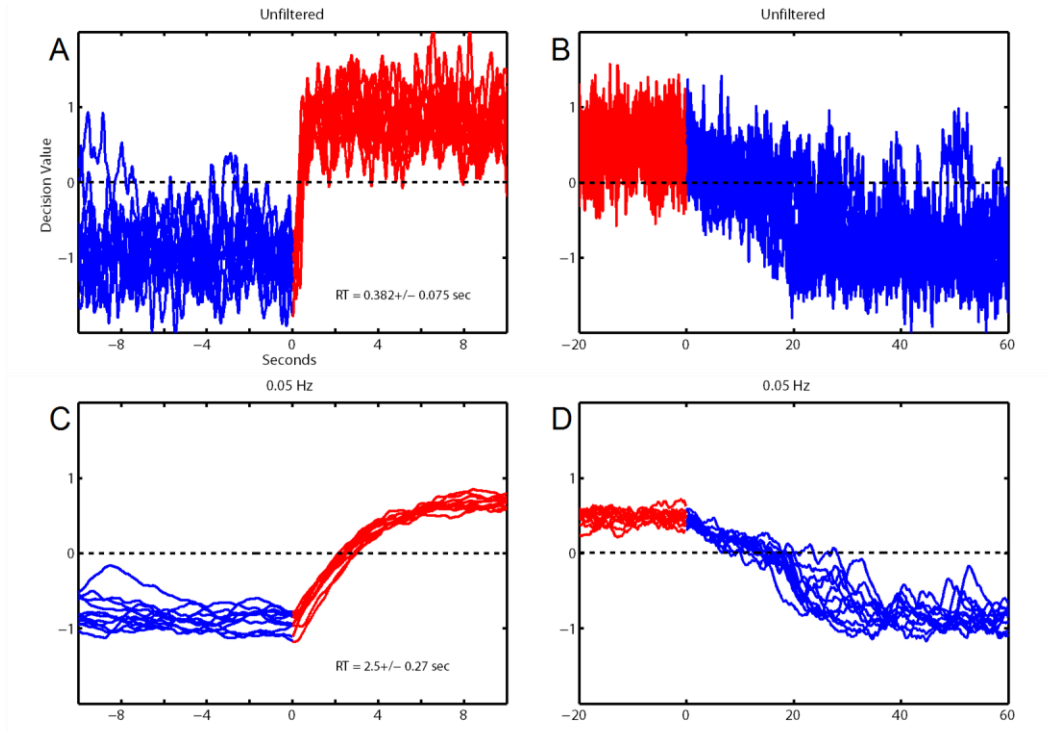
Figure 7.7. Classifier temporal characteristics. Panel (A) shows the unfiltered classifier output transition from the last 10 seconds of a rest block (blue traces) to the beginning of a task block (red traces). The time from the end of a rest block to the classifier decision value crossing 0 was averaged over 11 rest blocks on this example day to compute an average Response Time (RT). Panel (B) shows the unfiltered classifier output transition from the last 20 seconds of a task block (red traces) through the entirety of a rest block (blue traces, approximately 1 minute block duration). Panels (C-D) show the same transition traces from (A-B) when low-pass filtered at 0.05 Hz. As seen in (C), low-pass filtering significantly decreases the number of false positives during a rest block, but also noticeably increases the response time in transitioning from rest to task. In (D), however, the natural decay time constant of the task-to-rest transition remains largely unchanged by the filtering.

7.4 Discussion

7.4.1 Intra-task and Task/Rest Feature Modulation

Examining the spectral differences between task and rest trials as well as modulation within the BCI task itself during these experiments revealed several potential key considerations for the design of a neural prosthetic device. First, the observation that significant differences in mean power between task and rest trials exist raises the question of how to choose the correct baseline or threshold value for a

particular decoding model. For example, one employing a linear BCI algorithm for cursor movement might choose to subtract a baseline value from each feature and omit a constant offset term from the decoding model. This omission is valid if the features are truly mean centered around zero after subtracting the designated baseline. However, the results from these experiments suggest that the monkeys did not modulate around the original baseline for directional control of the cursor, but shifted to a different spectral energy state for the BCI task: higher power in the gamma bands and lower power in the alpha and beta bands. This finding could suggest that it is easier for a subject to modulate around this new "task baseline" than modulating in an unnatural direction from a "resting baseline" (i.e. gamma power below baseline or alpha and beta powers above baseline). This argument is supported by a prior ECoG study in our lab using high gamma activity for BCI control in which high gamma power was unlikely to decrease below resting baseline levels during the task [19]. Thus, centering features around a "resting baseline" without including a constant bias term in the decoding model could result in difficulty for the subject to move in certain directions.
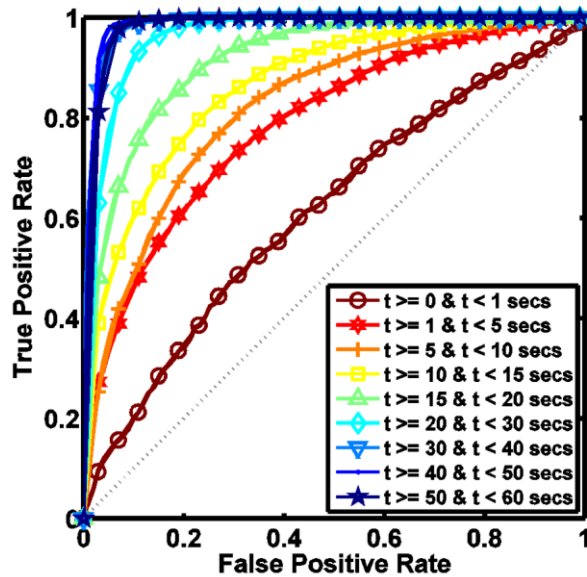
Figure 7.8. AUC analysis of task-to-rest transition. ROC analysis of the rest blocks in Figure 7.7. Each trace depicts the discriminability of classifier output samples taken from a given segment of each rest block versus all task samples (i.e. the first brown line with circles takes the first second of all rest blocks versus all task samples, while the last line compares the last ten seconds versus all task samples).

Similar to the consideration of a constant offset term is whether to calibrate any baseline and normalization terms prior to operation of an asynchronous BCI device, or

to calculate these terms online using methods such as a running average for each feature over a limited past history. While the latter could lessen the technical burden on the user of recalibrating the system due to drift in signal properties, it could also significantly lengthen the time needed to switch between active and resting states as past periods of rest would have to decay from any buffer before the user could fully transition to a suitable "task baseline".

Finally, the spectral differences in modulation patterns seen in these experiments suggest possible specialized decoding roles for various spectral bands. For example, although the high gamma frequency bands demonstrated significant differences in amplitude between task and rest states, they clearly showed greater discriminatory capacity for directional decoding related to the task. Conversely, the lower frequency bands showed only weak directional modulation with a more appreciable modulation between task and rest states. Given these properties, one could envision a sequential BCI processing chain similar to the EEG "hybrid BCIs" proposed by Pfurtscheller et al [50], [51] in which low spectral bands might serve as a binary "on/off" state decoder, after which higher frequencies would provide continuous valued movement information.

7.4.2 State Decoding Results

The results in Section 7.3.2 demonstrate excellent performance of the classifier in classifying task and rest periods when performed on a trial-by-trial basis, reaching 98-99% total accuracy levels. Of the trials that were incorrectly classified, the majority of classification errors occurred on rest trials at the beginning of a rest block. Since labels for task and rest trials were defined by the experimental block design and not by any external indicators of attention to the task, these may not necessarily represent ground truth. These "errors" could represent lingering attention to the task as the monkey remains prepared for another task trial, and accuracy results reported here could be conservative estimates.

Of course, trial-by-trial decoding is, by the definition laid out earlier, a synchronous design. Thus, perhaps more important are the real-time decoding results

reported here that compare favorably with past EEG experiments of the same nature in which AUCs or a derived measure (TPR at a fixed FPR) were reported [48], [49], [73]. These measures show further improvement when taken in light of the observed early rest block false positives seen in trial-by-trial analyses and the decay profile of the decision value during the task to rest transition. If the first 30 seconds of each rest block are removed from the analyses such that only rest data from after this decay period is included, the filtered AUC jumps to near unity (>0.99) for each monkey hemisphere. In addition, this performance was achieved within the framework of detecting distinct states during a two-dimensional task with a continuous valued variable of velocity as the control output. The possible overlap of task/rest state distributions in this scenario could likely be much greater than that between states in a binary classifier of movement/no movement, and theoretically makes it a more difficult classification problem.

Comparing the performance of our ECoG-based classifier with that of a simple cursor speed-based model also raises several important points. If the brain-controlled cursor were to simply stop moving during rest blocks due to lack of spectral modulation by the subject, an additional asynchronous decoding system would be unnecessary. However, the results in Figures 7.4 and 7.5 magnify the necessity for this decoding component. First, as shown in Figure 7.5, decoding of the state decision value is largely independent from that of the cursor speed, even though both are extracted from the same spectral source. This separation of states was not strictly imposed upon the monkeys in order to perform the task as in other studies [47], but it may be a phenomenon that emerges naturally as a consequence of the adaptive cursor decoding model we employed (with no *a priori* movement screening) along with the baseline and modulation characteristics discussed earlier. Secondly, the independence of state and cursor speed as well as the discriminability results during specific movement periods could prove advantageous in real-life applications. For example, a common action during everyday computer use is to hold a cursor over an icon for a period of time in order to view an associated help window. A similar brain control scheme should allow the user to hold the brain-controlled cursor in place without entering an idle state, and

112

then move on to the next action without pause after the user has finished reading. As seen in Figure 7.5A, the state decoder is equally as likely to classify low speeds as active during Hold A periods as it is for higher speeds, indicating the decoder can maintain an active classification even during periods of little movement. This so-called "active hold" may not be as easily implemented in systems that do not differentiate between a baseline or "no control" state and an intentional state to not move. Conversely, the speed distributions during rest blocks overlapped significantly with those during movement and hold periods, resulting in poor discriminability if using cursor information alone. Thus, a prosthetic device utilizing a movement decoding scheme similar to the one in these experiments would likely experience numerous undesired movements during periods of inattention if a state detection scheme was not implemented.

7.4.3 Filtering and Temporal Characteristics of Real-time Signal

We demonstrated a simple option of low-pass filtering the raw decision values of our state classifier in order to improve performance. Of course, more complicated statistical methods may indeed offer further improved performance with fewer temporal effects. However, low-pass filtering is a first-pass option that exhibited significant improvements, thus evaluation of other models was deemed outside the scope of this study.

Based on the results in Section 3.4, we suggested a low-pass filtering cutoff frequency around 0.05 Hz. Several caveats accompany this selection, however. First, selection of the cutoff frequency must balance the design criteria of classification accuracy and device response time. The response time for transitioning from task to rest and rest to task was likely contributed to by three major sources: 1) the visual and cognitive response time of registering and reacting to the appearance of the task-related objects, 2) the filtering delays associated with transforming raw ECoG signals into each spectral feature, and 3) delays associated with filtering of the raw decision value output. In the context of these experiments, the first two sources were held constant and represent a minimum response time from rest to task of approximately 400 ms (Figure

8A). Filtering of the decision value at a low-pass cutoff of 0.05 Hz (RC time constant of approximately 3.2 s) resulted in decreased false positive rates but increased the response time nearly six-fold. However, the transition from task to rest occurred over a much longer time scale (~20 s) than the transition from rest to task, and filtering therefore did not impact this period nearly as greatly. Thus, it appears that a desired rest-to-task response time (along with a desired accuracy level) should garner the most design consideration when choosing an output filter. A response time of 2.4 s (low-pass frequency of 0.05 Hz) seems reasonable when compared to the startup time of many modern day electronics (television, computer, cell phone, etc.), but a BCI's specific application will likely dictate this design criterion.

Secondly, the filtering results presented here may be partly biased by the temporal design of the task and rest blocks. With rest blocks lasting approximately one minute, the maximal portions of the curves in Figure 7B contain frequencies with time constants less than that of the task/rest switching frequency. If the task switching frequency in these experiments were shortened considerably (5-20 seconds per block), it is unlikely that the AUC vs. low-pass frequency curves would remain unchanged as data from non-adjacent blocks would begin to average together when filtered at lower frequencies. It is also unclear whether the decay profile observed during the task-to-rest transition would shorten in this situation or whether it would disappear completely. If that profile is truly related to the monkey's residual attention to the task or a desire to work for additional rewards, subjects may find it easier or more efficient to simply remain in a "ready" state during short rest blocks as opposed to placing a BCI device into standby between tasks.

One might question whether the differences between states observed in these experiments might be related to visual perception of the task since the cursor and targets were removed from the computer screen during rest blocks. Several observations discount the hypothesis that visual processing was the driving force behind these differences rather than conscious intent to control the cursor. The decay profile of the task-to-rest transition was much slower than the time needed to visually process the computer scene, and remnants of an active state persisted long after the screen had

114

been blanked. It was also not uncommon to observe extended periods of rest classifications during task trials near the end of a training session as the monkey's task motivation and task performance declined. During these task periods "incorrectly" classified as rest, the monkey would decline to start the trial by moving the cursor to the center target and complete the Hold A period. Accordingly, a rest state was observed despite visual presentation of the task, thereby lessening the role basic visual processing is likely to contribute to the state signal.

The observation that rest state classifications appears late in a training day during periods of declining performance could also provide additional motivation for employing this state classifier during early BCI training. Since it is difficult to determine when subjects, especially non-human primates, are providing maximal effort in the early stages of BCI training in which *a priori* movement correlations are not used to train a decoder, including periods of inattention in the training data could have deleterious effects on future BCI performance. Using a state decoder of the nature presented here to identify these periods could dramatically improve performance and decrease a subject's training time.

7.5. Conclusions

In this study, we examined the spectral differences in micro-ECoG signals between BCI control and rest periods when monkeys were allowed to develop their own neural strategies for completing a two-dimensional BCI task. After adapting our cursor velocity decoding algorithm to each monkey's modulation patterns on a daily basis, it appeared that the monkeys would transition to a different spectral energy state than baseline (higher gamma powers and lower alpha and beta powers) during the BCI task, and modulate around this new task baseline to complete the task. Examining specific spectral modulation patterns revealed possible specialized roles as high gamma frequencies showed greater modulation with target direction within the task, while lower frequency components showed greater amplitude differences between task and rest periods. Upon post-hoc training of a classifier to discriminate between active brain-control and rest periods, accurate decoding performance was demonstrated in each case,

while low-pass smoothing of the output signal further improved this performance. To date, this study is the first to our knowledge to employ this type of state detector along with a 2-D regression (continuous valued) decoding model for movement, and the first to apply this framework to closed-loop micro-ECoG control. Within this experimental framework, we observed likely evidence of residual readiness to perform the BCI task early in rest periods as well as inattention to the task late in training sessions. Overall, the results presented here could provide a useful framework for controlling the idle state of an ECoG-based BCI device as well as a method for improving BCI training data.

8 BCI Learning and Adaptation

8.1 Introduction

For most of the experiments discussed in this thesis, an adaptive algorithm was used in which the decoding model for cursor movement under brain control attempted to infer the neural strategy of the monkey using only past cursor data and ECoG signals. This modeling was done without the use of *a priori* correlations between movements and neural activity, yet each monkey developed convincing control of the virtual cursor over time in a variety of tasks as demonstrated in Chapters 5 and 6.

Previous work in our lab, however, examined whether non-human primates could adapt to using a simple, arbitrarily chosen decoding model [19], [65]. Through biofeedback, subjects identified ECoG modulation patterns that were in line with a given decoding scheme and allowed the monkeys to successfully perform a BCI task similar to those described here.

In the analyses presented in this chapter, we examine learning trends, evidence of adaptation – both by the subject and the decoding algorithm – as well as the robustness of decoding models and features over time. Examining these factors in both of the previously mentioned experimental contexts may provide insights into avenues for optimizing BCI algorithms.

8.2. BCI Trends using an Arbitrary Decoding Algorithm

In this section, we examine evidence of subject adaptation to an arbitrary/fixed decoding algorithm. In addition, we examine further changes in adaptation and performance when the decoding algorithm later adapts to the subject on a daily basis. Finally, we examine the theoretical robustness of the ECoG feature set in terms of lower limits on the number of channels needed for 2-D control as well as durability of decoding capabilities with channel loss.

8.2.1 Methods

In the following set of experiments, two monkeys (J and M) first learned to control a 1-D BCI cursor using ECoG signals from two arbitrarily chosen electrodes along a row of electrodes in a circular 28-channel recording array (e.g. +X, -X channels, see Figure 8.1). Signals were processed and normalized mostly as done throughout this thesis to extract spectral modulation in the 75-105 Hz frequency band on each control channel. One notable difference between normalization methods was that these experiments utilized a 100 second running average of the control signal as the baseline rather than the fixed 5 minute session baseline used throughout the rest of this thesis (see Chapter 7 for further rationale). Modulation above baseline on the positive channel resulted in cursor velocity in the positive direction while modulation below baseline resulted in movement in the negative direction. Conversely, modulation above baseline
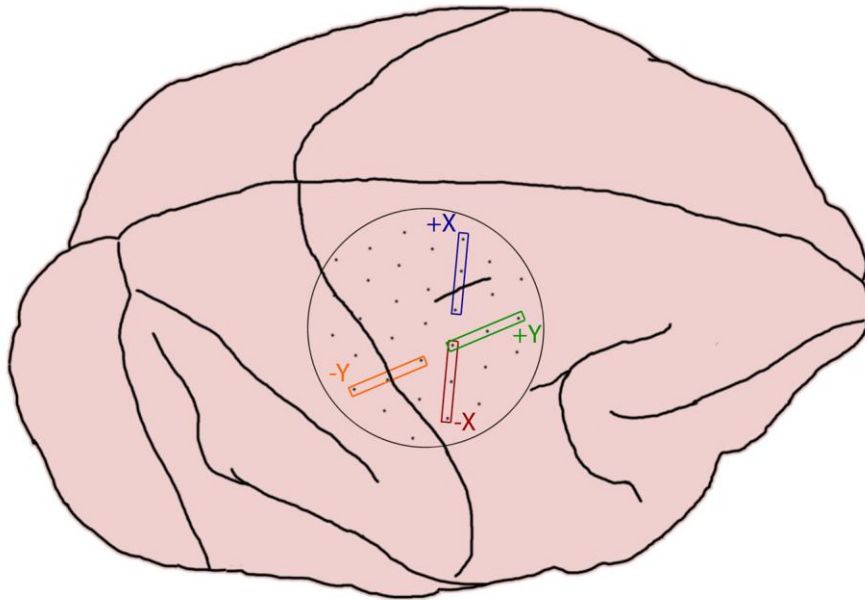
Figure 8.1 Array placement and control scheme for arbitrary control experiments. Light gray circle and dots outline the approximate array and channel locations for a 28-channel, 300 μm electrode diameter ECoG array used in monkeys J and M. Colored boxes indicate an example of the arbitrary control scheme used in which two channel's were assigned per dimension – one to "push" the cursor in the positive direction when modulated above baseline, and one to "pull" the cursor in the negative dimension when modulated above baseline. These control started at opposite ends of a row and were moved closer together on subsequent weeks.

on the negative channel resulted in negative direction velocity or positive velocity for modulation below baseline. (Note: For the purposes of the original thesis for which the 1-D experiments discussed here were intended [19], the control channels for a given dimension initially started on opposite ends of the chamber and were moved closer together within a row (colored boxed in Figure 8.1) during subsequent weeks to examine distance effects on BCI control.)

The monkey first had to learn to use this control scheme to perform a 1-D, two target radial-choice task as described in section 3.7.1. Subjects were aided with a training bias as in our other experiments until reaching proficiency well above chance. After finishing experiments for the first dimension (i.e. horizontal axis), the monkey proceeded to progress through the same experiments with a different pair of channels for a 1-D radial-choice task on the vertical axis. Following completion of both 1-D experiments, the four control channels at the periphery of the rows outlined in Figure 8.1 were combined to control the 2-D velocity of a cursor in a four target radial-choice task using their same directional assignments. This configuration was used until task performance plateaued.

After reaching a stable performance level using arbitrary control channels, the decoding algorithm was switched to an adaptive algorithm similar to those described in earlier chapters in this thesis:

$$y(t) = w^T X(t) + b_0 \qquad (8.1)$$

Here, $y(t)_{2x1}$ is the two dimensional control signal (translated to velocity in these experiments), $X(t)_{28x1}$, is the set of normalized spectral ECoG features, w_{28x2} is the decoding weight vector, and $b_{0_{2x1}}$ is a constant bias term for the control signal. This decoding model was trained using an ℓ_1-regularized least-squares regression model at the end of every training day as in previous chapters (see section 5.2.3). This first attempt at day-to-day adaptive decoding used only the 75-105 Hz frequency band per channel. Further adaptation iterations with these monkeys included expanding the number of targets in the 2-D radial-choice task to 8 while still using one frequency band, and finally performing the 8 target task using the five frequency bands used for

control throughout the rest of this thesis. For most of the following analyses, however, we will focus on adaptation during the fixed channel decoding algorithm as well as the transition to day-to-day adaptation of the decoding algorithm using the 75-105 Hz frequency band.

8.2.2 Subject-to-Model Adaptation

Although only two equally weighted channels were used for control on a given week during the 1-D experiments, a decoding model similar to that described by Eqn. (8.1) was formulated on a daily basis using all 28 channels to predict whether a positive or negative direction target was presented to the monkey. In addition to using all 28 channels, the spectral content of each channel on a given trial was subdivided into 5 Hz bins from 0-200 Hz. This model was used to infer the neural strategy of the monkey and its adaptation toward the arbitrary control scheme used for these experiments.

Figure 8.2 depicts the predictive weights of two regularized least-squares models trained separately on ECoG data from two days of a single week of BCI training.
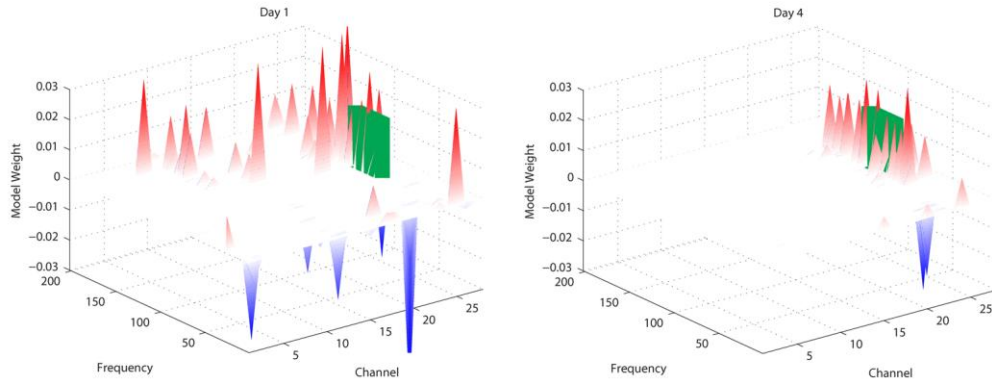
Figure 8.2. Subject Adaptation and convergence of predictive model weights during BCI training. A regularized least-squares model was used to predict whether a left or right target was displayed to the monkey (M) based on ECoG spectral features on 28 channels from 0-200 Hz. Red peaks indicate that a spectral increase in power above baseline at a channel/frequency feature was predictive of a left target having been displayed, while blue peaks indicate an increase in power was predictive of a right target having been displayed to the monkey. Conversely, a decrease below baseline for red peaked features would predict a right target, etc. Green bars indicate the actual BCI control algorithm assigned for this task. Note how the monkey's initial neural strategy on Day 1 (left panel) appears relatively broad and random, while by Day 4 (right panel), the monkey has localized his neural activity to match the imposed control algorithm.

During this period, an increase in power above baseline between 75-105 Hz on an arbitrarily assigned control channel (green bars, Channel 24) was coded to move the computer cursor to the left, while a decrease below baseline in this frequency band would move the cursor to the right. (A second channel could be used in the opposite fashion to move the computer cursor to the right. However, the monkey did not modulate this channel as much as the other and is thus not displayed in the figure for simplicity). The least-squares model was trained to predict which target the monkey was attempting to move to provided only an ECoG spectra (28 channels, 0-200 Hz) for a given trial, and the best predictive weights for each day are shown above. As seen in the figure, the monkey appears to employ a relatively broadband and non-specific strategy to accomplish the brain-control task early in the training week. After a few days of training and biofeedback, however, it is evident that the monkey has modified his neural output pattern to more closely fit the fixed control algorithm on one of the assigned control channels.

In addition to much of the informative spectral content residing in the high gamma band around the 75-105 Hz band of interest, it is also interesting to note that an associated band of predictive weights with the opposite sign emerged in the lower beta frequency range on the same channel. This spectral feature emerged even though it was not causal to control of the cursor, suggesting a possible neurophysiological linkage to the associated changes in gamma power.
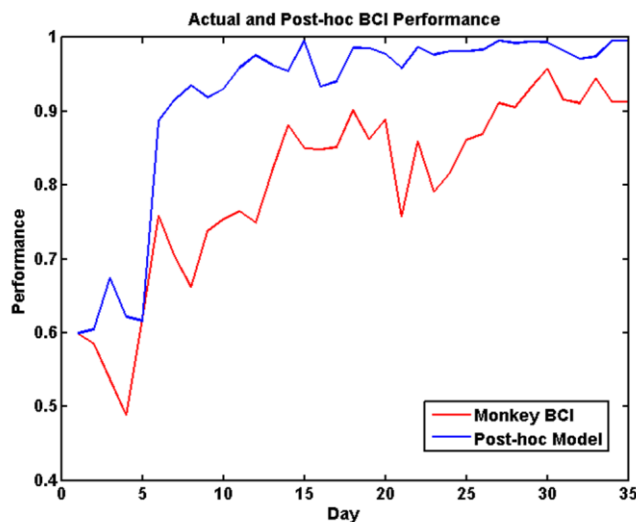
A timeline of the

Figure 8.3. Actual vs. post-hoc modeling performance, 1-D task. Red lines indicate the actual performance of the monkey using arbitrarily assigned control features on a given day, while blue lines indicate the cross-validation performance of the post-hoc regression model utilizing all channels in predicting the displayed target.

adaptation suggested in the previous figure is depicted in Figure 8.3. This figure shows two performance curves: one for the actual 1-D task performance of the monkey using the arbitrary control scheme, and one for the cross validation accuracy of post-hoc models trained for each day employing all channels as features. A starting point near random chance is shown for both curves as the monkey just begins training on the BCI task. The performance of the post-hoc model soon jumps up as the monkey develops a neural strategy consistent within a given day, suggesting that additional information is present that could aid in BCI performance. However, the gap that exists between the two performance curves gradually closes as the monkey adapts his neural strategy to the fixed algorithm. Taken with Figure 8.2, these data suggest that although additional features may hold predictive value early on in training, biofeedback may be used to adapt to a number of arbitrarily assigned control algorithms and reduce the need for pre-screening in BCI control feature selection.

8.2.3 Model-to-Subject Adaptation

Figure 8.4 illustrates 2-D task performance for each monkey during the transition from a fixed, arbitrary decoding algorithm using only four channels to an
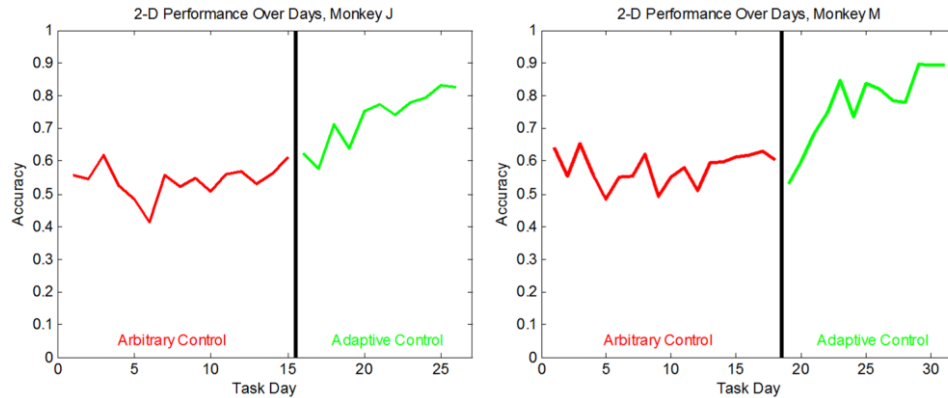
Figure 8.4. Performance for arbitrarily assigned vs. adaptively learned control schemes. The performance of two monkeys over the best 200 trials on a 2-D, four target choice task is shown over the course of experiments involving two different BCI control algorithms. Red lines indicate performance on days where the same four arbitrary electrodes were assigned to control the horizontal and vertical velocity of a cursor. Green lines indicate performance on days where the control was distributed amongst all electrode channels based on a least-squares model trained from the previous day's data. The black vertical line indicates a two week break between experiments.

adaptive algorithm using all 28 channels. As seen from the figure, performance on the four target radial-choice task using the fixed decoding channels appears to level off for each monkey near 60%. After switching to an adaptive decoding model, each monkey's performance sharply increases within days of the switch until leveling off at a higher plateau near 80-90%.

While the previous figure demonstrates a stark contrast in the absolute target selection accuracy, Figure 8.5 expands on this comparison by demonstrating trends in
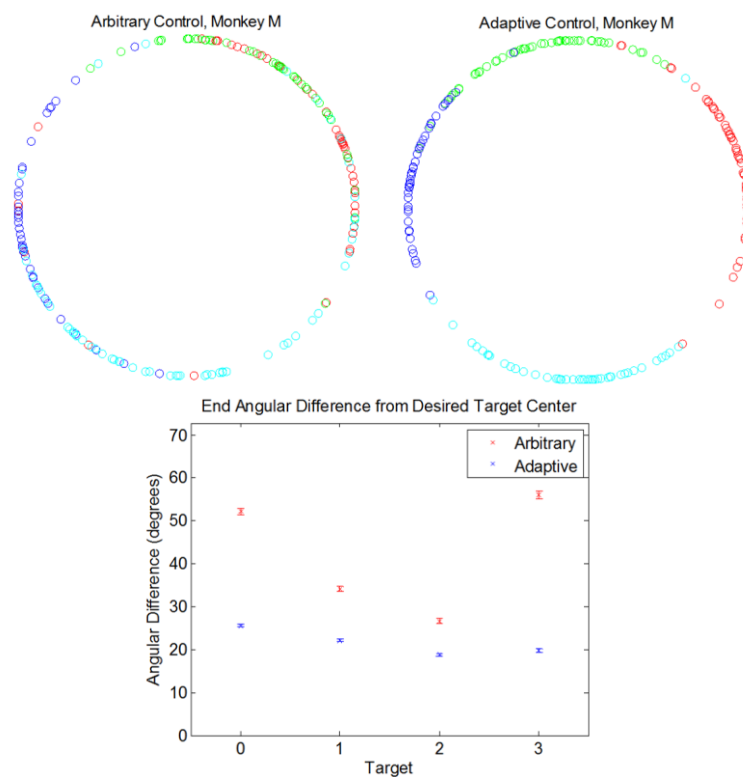
Figure 8.5 End cursor position for 2-D, BCI radial-choice task. The cursor position at the end of a trial when one of four targets has been selected is shown. Two days near the end of training for BCI tasks using either arbitrary (left panel) or adaptive (right panel) control schemes are represented. The four colors of markers indicate which target was displayed to the monkey (the target was ¼ of a circle being highlighted), and the targets were centered on the horizontal and vertical axes: red targets were at the rightmost quadrant of the circle, green – top quadrant, blue – left quadrant, cyan – bottom quadrant. As seen in the left panel for arbitrary control, the monkey appears to be biased toward target selection along the diagonal in quadrants I and III with relatively sparse end positions in quadrants II and IV. With adaptive control, the monkey appears to have more precise control over the cursor with uniform target hits among the four quadrants and more concentrated target hits near the centers of the targets.

the precise location of the BCI cursor when it crosses the target ring. Data for this figure were taken from task days at the end of each decoding algorithm strategy in Figure 8.4 for monkey M. As seen from the figure, target ring crossings during adaptive control trials were likely to lie much closer to the center of the intended target than during fixed control trials. Fixed control crossings also tended to lie closer to the lower left-to-upper right diagonal, suggesting a degree of correlation between the horizontal and vertical control signals that is not apparent with the adaptive control scheme.

Overall, these results suggest that while biofeedback may aid a subject considerably in adapting to a fixed algorithm, additional task-related information is likely present that may be exploited by adapting the decoding algorithm to the subject.

8.2.4 Task-related Feature Modulation

The directional or target-related discriminability index, d'_{target}, introduced in Chapter 7 was calculated for each task day throughout the fixed and adaptive control experiments for monkeys J and M. To recap the utility of this metric, it quantifies the amount of amplitude modulation a feature exhibits with target direction and then scales it by the pooled standard deviation in spectral amplitude within trials grouped by target as follows:

$$d'_{target} = \frac{A_{pp}}{\sqrt{\Sigma_{j=0}^{N-1} p_j \sigma_j^2}} \qquad (8.2)$$

Again, A_{pp} is the peak to peak difference in mean amplitude modulation across all N targets, σ_j^2 is the standard deviation in amplitude within trials to target j, and p_j is the proportion of trials to target j. In other terms, this discriminability index could also be thought of as a signal-to-noise ratio for directional modulation. Since d'_{target} was not the primary focus of Chapter 7, Figure 8.6 provides a pictorial description of its calculation.
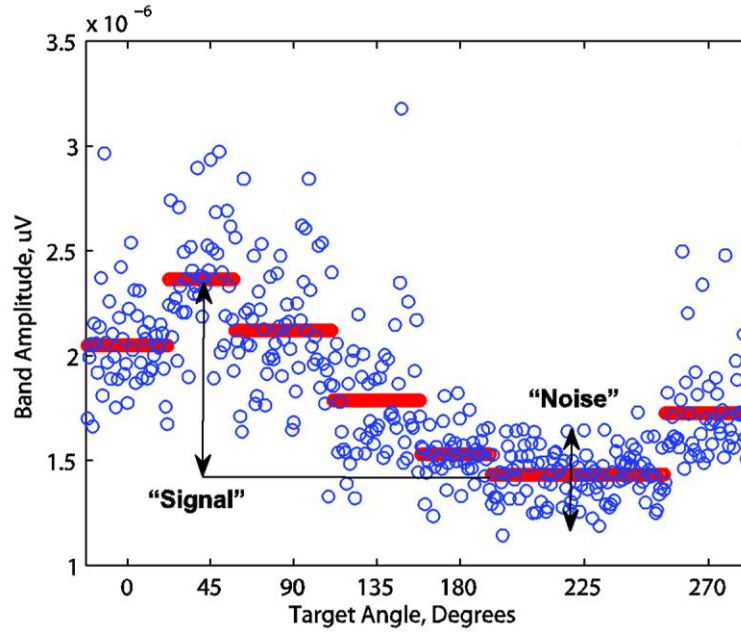
Figure 8.6 Example tuning curve for d'_{target} calculation. Blue dots indicate band amplitudes for individual trials to a given target. Red bars indicate mean band amplitudes over all trials to a given target. The "signal" indicates the peak-to-peak difference between maximum and minimum mean target amplitudes. The "noise" represents the intra-target variability in band amplitude.

Figure 8.7 depicts changes in d'_{target} on each channel (75-105 Hz amplitude) over time for monkey J as he progresses through the fixed and adaptive control experiments. From the figure, several "flavors" of channels materialize with respect to their directional modulation. Some channels that were originally a control channel or near a control channel in the arbitrary control scheme start with a high d'_{target} early on and maintain this throughout the rest of the experiments. In the channel with the exploded view in the top left panel (a channel originally a control channel for the +X direction), comparison of its d'_{target} with the progression of control experiments reveals that the channel starts with a high d'_{target} during horizontal 1-D experiments. The channel's modulation then quiets down during vertical axis 1-D experiments, and ramps up again after starting 2-D experiments. Other channels initially do not show a strong d'_{target} during arbitrary control experiments. Interestingly, upon entering the adaptive control phase, several of these channels develop a stronger d'_{target}. Presumably, this shift occurred after biofeedback revealed that more neural features could affect the

125

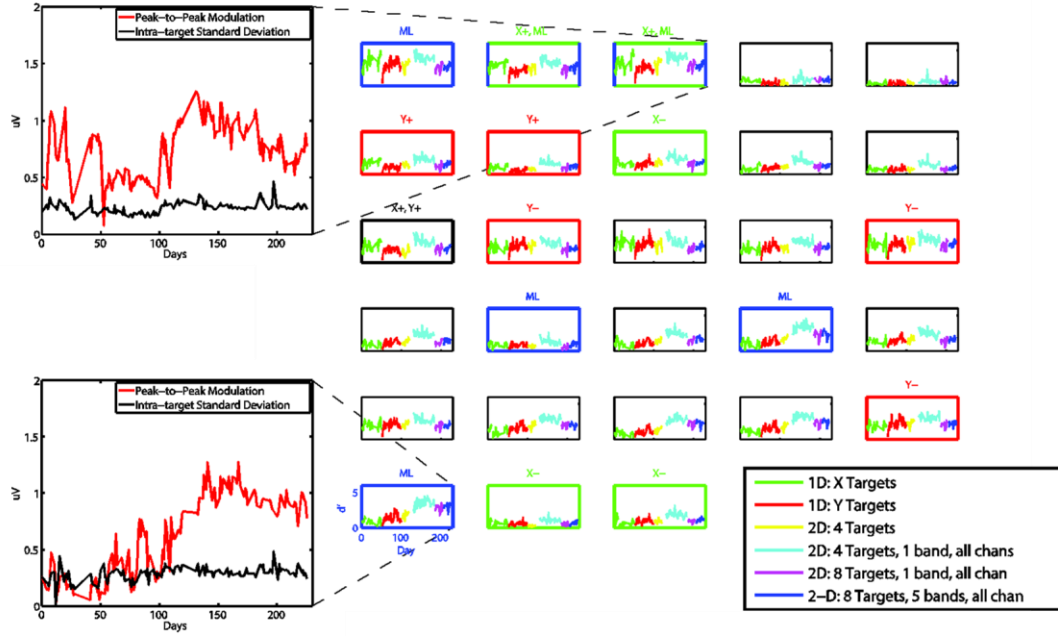**Figure 8.7 Target discriminability, modulation, and noise over days for monkey J. <u>Top Right:</u> Target discriminability index, d'_{target}, for 75-105 Hz over time for all 28 channels. Colored line segment plots indicate the BCI experiment the data was taken from for a particular day. Red and green box borders indicate a channel was arbitrarily assigned control to either the X or Y direction of movement on a particular week. A channel's assigned direction is given above the axes (i.e. X+ indicates a band amplitude above baseline results in cursor velocity in the positive x direction). The one thick black box channel was used both as a X+ and a Y+ channel at some point (see Figure 8.1). Blue axes indicate a given channel held considerable weight during experiments using day-to-day adaptive decoding algorithms. Some axes share more than one color (channels 2 and 3), indicating that they were both used as a fixed control channel at some point as well as a significant adaptive control feature. Other features show only blue axes, indicating they were not originally assigned control, but were later assigned significant contribution during adaptive control experiments. <u>Exploded channel views:</u> Band amplitude modualtion and noise over time. These figures show the peak-to-peak modulation of a channel between 75-105 Hz and the pooled intra-target variation in modulation over time for a channel showing a fairly strong and consistent d'_{target} over time (left panel) and for a channel showing a weak d'_{target} early on that increases significantly over time (bottom right panel).**

movement of the cursor, and the monkey adapted his neural strategy accordingly by recruiting more features.

Along with changes in d'_{target}, Figure 8.7 also breaks down this metric into its constituent peak-to-peak modulation and intra-target variation components as a function of time. Theoretically, an increase in d'_{target} could result from two changes in these components: either the monkey could learn to increase the modulation with target direction for a given feature, or he could make a feature's modulation more consistent (lower variation) within each target. As seen from the exploded views

126

depicting these two components, the peak-to-peak modulation follows the d'_{target} trend closely while the intra-target "noise" or variation remains flat. This result implies the monkey is adapting and becoming more proficient at the task by increasing absolute modulation differences rather while the variance in spectral amplitudes per target remains relatively constant with training.

8.2.5 Model Stability and Feature Robustness

Two important considerations for a BCI system are the need to retrain or recalibrate the decoding model and the robustness of the system to maintaining decoding performance even with the loss of some channels' signals. To investigate the first consideration, we examined the evolution of the decoding model over time after each monkey began using the adaptive control scheme. Figure 8.8 depicts the time course of decoding model weights for monkeys J and M immediately following the transition from arbitrary to adaptive decoding weights. The timelines show that each monkey's weights settle relatively quickly (within 5-10 days) to a stable decoding model with roughly 3-6 prominent features per dimension. These prominent features aligned well with channels previously assigned as fixed control channels. This might suggest that the monkeys had already adapted to the arbitrary control scheme, and adaptation by the decoding algorithm merely fine tuned this model. In addition, the relative stationarity of this model might require minimal recalibration on a day-to-day basis for real-life applications.

By taking the absolute magnitude of each channel's preferred direction in the decoding model, Figure 8.9 shows a general picture of the number of heavily weighted features for each day's decoding model during the same period as the previous figure. The number of these features was quantified by counting the number of channels with a preferred direction magnitude at least greater than 40% of the maximum weight magnitude. Consistent with the number of prominent features identified by visual inspection in the previous figure, the number of feature weights from day-to-day above the 40% threshold was generally between 3-8 features.
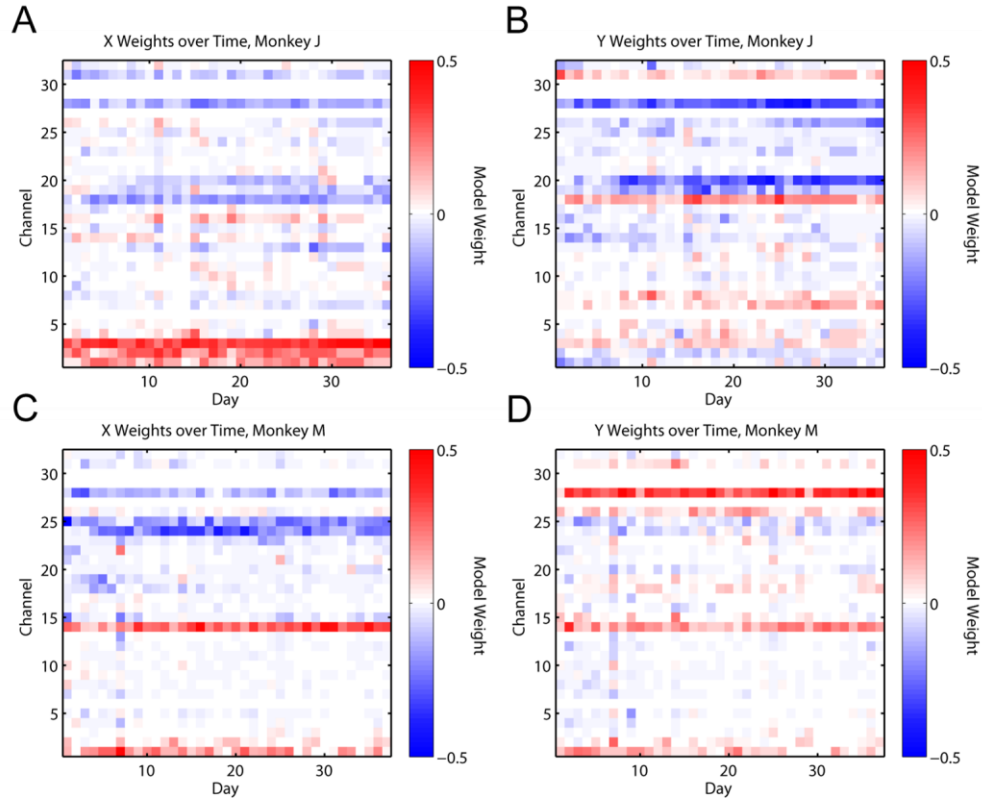
Figure 8.8 Adaptive model weights over time, monkeys J and M. Decoding model weights immediately following the transition to an adaptive decoding model are plotted as a function of training days. Horizontal and vertical dimension weights are shown for monkey J in (A) and (B), while monkey M's weights are shown in (C) and (D).
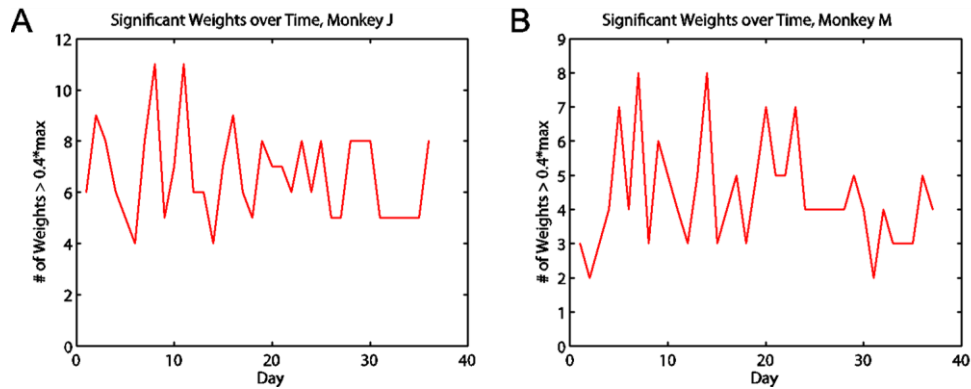
Figure 8.9 Significant weights over time, monkeys J and M. The absolute magnitude of the weight vectors in Figure 8.8 were sorted, and those greater than a threshold of 40% of the max weight were counted as significant. Weight magnitudes exceeding threshold are plotted for (A) monkey J and (B) monkey M.

Based on the number and relative magnitude of prominent features seen in the previous two figures, it is plausible that a BCI system might need fewer channels than the total number that were recorded from to perform near maximal performance. To support this claim, a set of post-hoc analyses was performed to examine the robustness of the ECoG feature ensemble in predicting the displayed target. For these analyses, four-target radial-choice task data from monkeys J and M was used from later experiments in which 5 frequency bands per channel were used for control as in experiments with monkeys K and P. The first analysis involved gradually adding more channels to the control ensemble to see how performance improved with additional model features. Channels were first examined individually to see which had the best solo predictive performance. This channel was retained, and the remaining channels were iteratively paired with the first to see which pair held the best predictive value. This process was repeated until all channels were included in the model. The results of this analysis in Figure 8.10A for monkey J suggest that 4-8 channels are needed for maximum accuracy in this particular 2-D, 4 target task.

The second robustness analysis assumed that the best 8-channel combination from the previous analysis was used for decoding. Performance bounds were evaluated as channels were dropped from the decoding model to simulate the effect of electrode failure on task performance. An example of the results of this simulation for monkey J are continued in Figure 8.10B. From these data, it can be seen that there is a strong probability of maintaining classification accuracy even after losing half of the signals on this limited subset of channels.
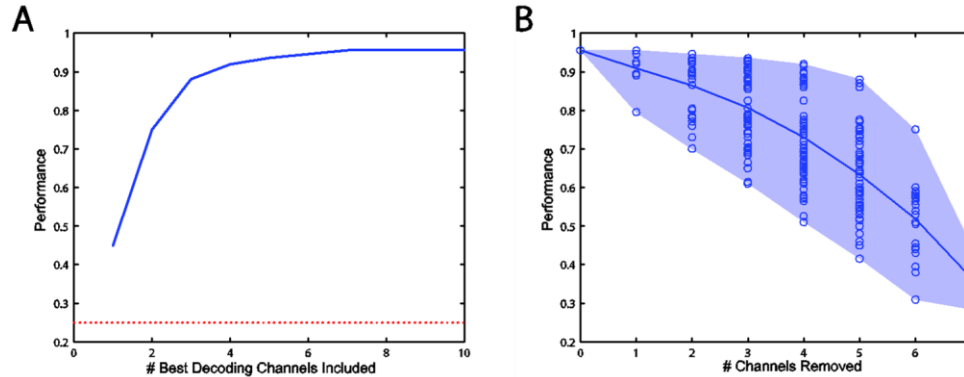
Figure 8.10 Decoding robustness analysis, monkey J. (A) *Post-hoc analysis of decoding performance (blue line) with increasing number of channels used in decoding algorithm.* **A regularized least squares decoding model was trained and validated on each channel individually to determine the channel most predictive of displayed target position in a 2-D, 4 target BCI selection task. This channel was retained and iteratively combined with the remaining 27 channels to determine the most predictive combination of two channels. These two channels were retained, and the procedure was repeated for the best three channels, four channels, etc.   The trend in validation performance shown suggests that decoding performance plateaus for this 2-D task after only eight channels, while inclusion of the best four channels in the model still results in over 95% of the maximum accuracy achieved (0.92 versus 0.955). The red line indicates random chance performance for this task (0.25). (B)** *Simulation of decoding performance with channel loss.* **The best combination of eight channels from the previous analysis was used to simulate how decoding performance might suffer with random channel loss. All combinations of channels (n-choose k channels, n = 8, k = 1:8) were trained and validated as before. The solid blue line indicates mean validation performance for a given group size, and the shaded region indicates the range of performances. Blue circles indicate individual performances of channel combinations. The displayed data suggest that even after failure of half of the channels on this small subset of channels, reasonable performance would be expected to persist and may decrease only slightly.**

8.3 BCI Trends using an Adaptive Decoding Algorithm

In this series of analyses, we examine trends in BCI learning and adaptation when the algorithm adapts to the subject from the very first day. In contrast to experiments in the last section where adaptation by the decoding algorithm occurred late in training, the subjects in the current experiments were not initially forced to adapt to an arbitrary algorithm. Thus, we can examine how the neural strategy of each subject for BCI control may naturally evolve when no prior model is forced upon them. In addition, we can compare how the evolution of feature modulation and decoding

models for active BCI tasks over time compares with the development of differences between the BCI task/rest states discussed in the last chapter.

8.3.1 Methods

For the purposes of examining learning and adaptation trends, early radial-choice task data from monkeys K (both hemispheres) and P were primarily used for analysis (see Chapter 5 for details of task description and multi-band decoding algorithm). Radial-choice data was primarily analyzed as it presents the best opportunity for comparison with data from monkeys J and M as well as easy quantification of classification performance. From this task data, trial-averaged data was used to calculate the d'_{target} and $d'_{task/rest}$ metrics on a day-to-day basis.

Center-out task data (both velocity and acceleration control) was used as needed for analysis of weight matrices and feature modulation at later time points. However, it is important to note how differences in the task and control methods influence these analyses. First, as decoding matrices were not updated from day-to-day during acceleration control experiments, adaptation would not be visible by viewing these matrices alone. Also, because the center-out task requires modulation of the control signal to both accelerate and decelerate the cursor in order to satisfy hold times, whereas the radial-choice task allows one to maintain a high control signal magnitude in a given direction without the need for deceleration, artificial differences in modulation (or lack thereof) across targets may emerge when a given feature is averaged over the course of a trial. In general, target and task/rest related modulation was most prominent in radial-choice task data, but the other mapping tasks still showed notable evidence of modulation if it was present.

8.3.2 Task-related Feature Modulation

Figure 8.11 illustrates the d'_{target} metric as a function of time for all three monkey hemispheres. From the figure, d'_{target} appears to increase dramatically after the first 5

Figure 8.11 Time course of d'_{target} development. The target discriminability index, d'_{target}, is plotted over the initial BCI training period (radial-choice and center-out tasks under velocity control) for **(A)** monkey K's left hemisphere, **(B)** monkey K's right hemisphere, and **(C)** monkey P's left hemisphere. For panel **(B)**, Day 0 represents the first day the monkey K's right hemisphere was used for causal control of the cursor. Frequency labels on the vertical axis indicate a single group of 32 spectral control features with the designated frequency band.

days for monkey K's first hemisphere (left) and monkey P. Figure 8.11B shows the d'_{target} time course for monkey K's right hemisphere. This array was recorded from during the last two weeks of experiments in which the left hemisphere was used for control before beginning its own set of BCI control experiments. In this hemisphere's figure panel, Day 0 marks the point from after which its signals controlled cursor movement. As seen from the training sessions leading up to this point, many features showed significant directional modulation before they were actually used for control. These estimates of d'_{target} are likely conservative due to the difficulties associated with averaging profiles for center-out trials as mentioned in the previous section. However, once that hemisphere was given control of the cursor, target discriminability remained

fairly constant with those levels for a few days and then increased dramatically on the third day of control. It is also interesting to note that following this transition in hemispheric control, monkey K's left hemisphere continued to modulate with the task even though it had no effect on the cursor.

As proposed in Chapter 7, directional modulation appears to have evolved the most in the high gamma frequency control bands (70-150 Hz). Monkeys K's and P's left hemispheres showed some increase in modulation in the beta frequency range (15-24 Hz), but not as consistently as the gamma range. Alpha frequency features (8-15 Hz) showed very little modulation for each monkey hemisphere.

8.3.3 Task vs. Rest Feature Modulation

Figure 8.12 depicts the time course for the development of feature differences between active task and rest states as measured by the state discriminability index, $d'_{task/rest}$, introduced in Chapter 7. The trends depicted for each monkey indicate that differences between BCI task and rest states may emerge more quickly than the target discriminability indices in the previous figure. Strong differences appear almost immediately from the start of a hemisphere's control for both of monkey K's hemispheres. Task/rest modulation did not seem to emerge as quickly for monkey P during initial training on the radial-choice task. However, as noted in Chapter 3, it was suspected that this monkey often took long breaks within active blocks of the radial-choice task, which would reduce the ability to discriminate between the two states. In contrast to the target modulation spectral trends in Figure 8.11, task/rest modulated features mostly resided in the lower frequency alpha and beta bands with less modulation in the higher gamma bands.

8.3.4 Decoding Model Evolution and Robustness

Because the decoding scheme for monkeys K and P adapted on a daily basis from the beginning of their experiments, it is difficult to firmly decipher who/what was adapting to whom. However, we can examine how the total sum of adaptation represented in the decoding weight matrix evolved over time. Figure 8.13 shows an
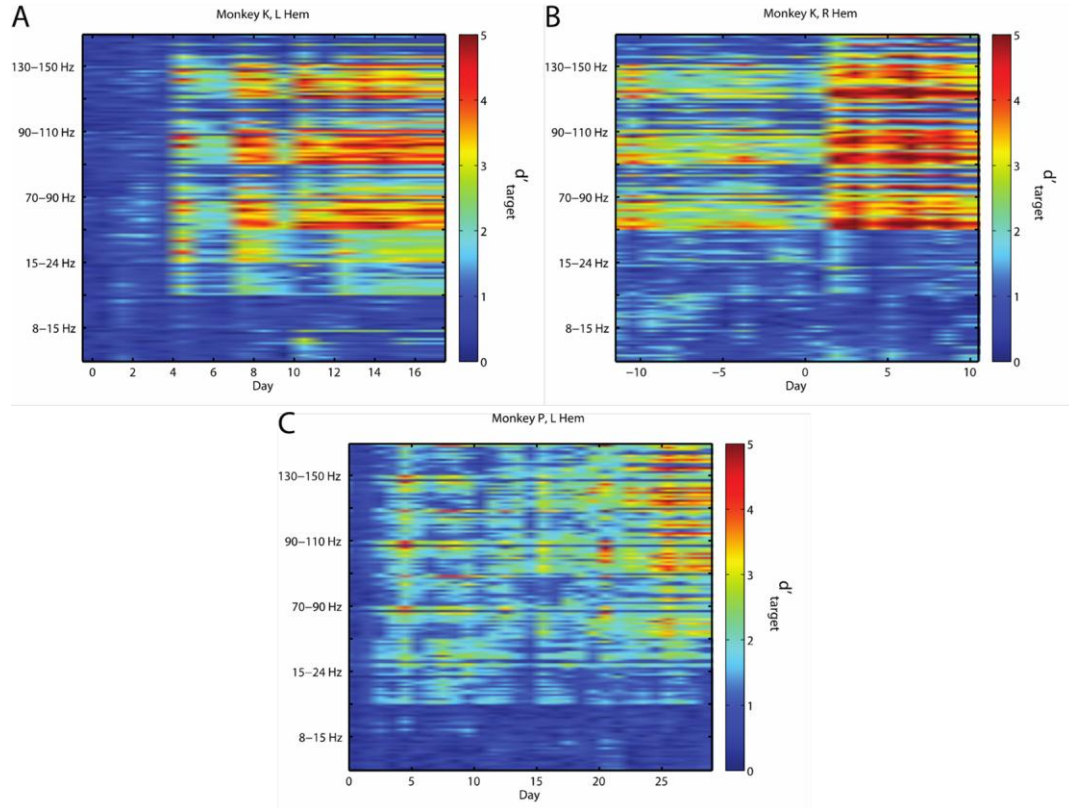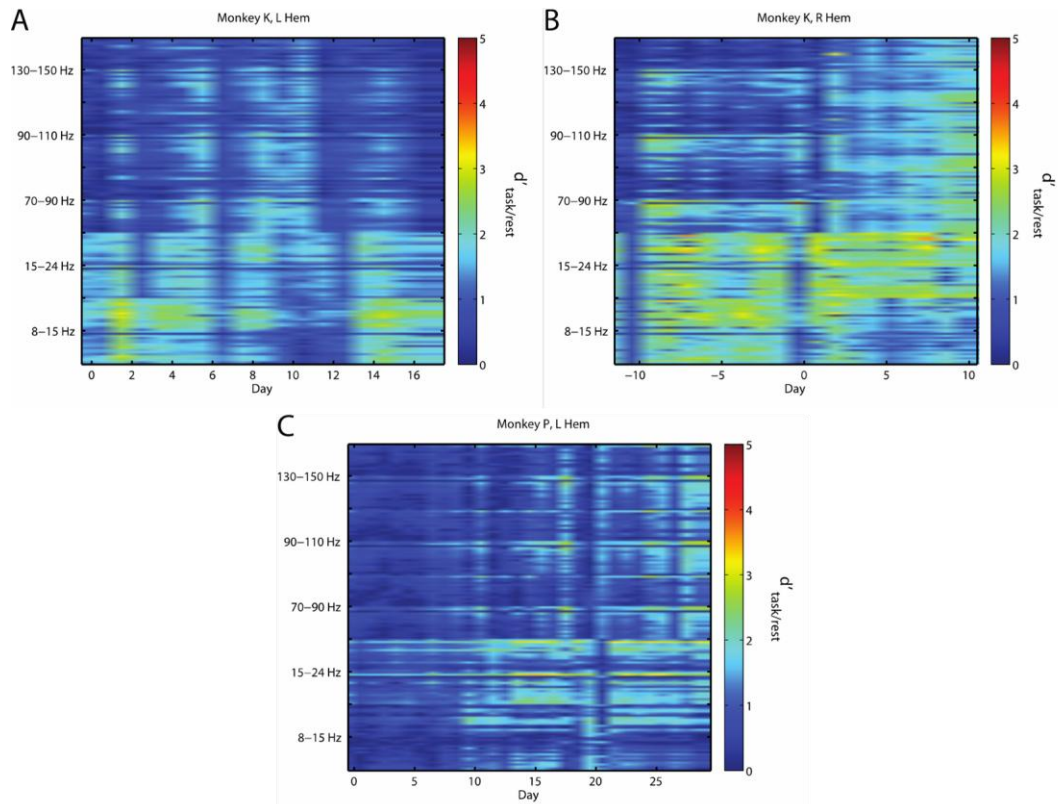
Figure 8.12 Time course of *d'*_{task/rest} development. The target discriminability index, *d'*_{target}, is plotted over the initial BCI training period (radial-choice and center-out tasks under velocity control) for (A) monkey K's left hemisphere, (B) monkey K's right hemisphere, and (C) monkey P's left hemisphere. For panel (B), Day 0 represents the first day the monkey K's right hemisphere was used for causal control of the cursor.

example of the evolution of the X and Y weights for monkey K's right hemisphere. As seen from the figure, when the algorithm adapts to the subject from Day 1, the distribution of model weights appears to be more randomly distributed across features than was the case for monkeys J and M after training on arbitrarily assigned channels for control.

For monkeys J and M, one could discern distinct feature bands which held a relatively constant weight over the course of many weeks. This was likely due to the earlier training to adapt to a limited subset of channels for control. No such consistent features readily emerge for the weight matrices when adapted from the very beginning, however. This could partly be due to the regularization process we employed in training

134

each day's model. If multiple features contained similar or redundant information, the sparseness inducing ℓ_1-regularization process might very well pick out one feature to retain one day while zeroing other similar features, and then single out another feature the next day depending on which feature had better signal quality on a given day.
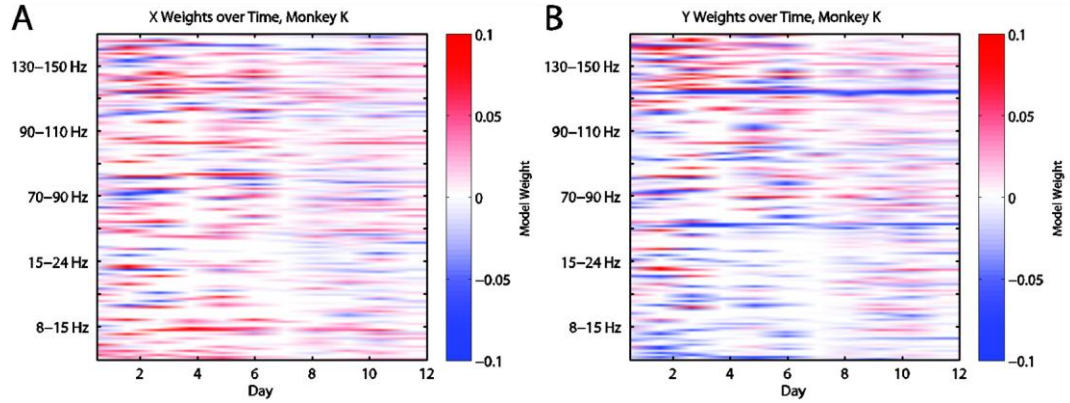
Figure 8.13 Adaptive model weights over time, monkey K. Example of decoding model weights over time when BCI training is initiated with day-to-day adaptation by the algorithm. Horizontal and vertical weight components are displayed in (A) and (B), respectively. Control features are grouped into their respective spectral bands.

Although the overall structure of the day-to-day adaptive weight matrix looks less consistent than those following arbitrary model control, it is possible that models on consecutive days actually produce very similar control signals due to redundancy of information across channels. To examine this possibility, the model for each day was used to predict the actual control signal on all other days. A normalized dot product between the actual control signal and the predicted control signal was calculated to measure the directional agreement between the two signals. Figure 8.14 shows an example of this temporal relationship between the date of training for the models in Figure 8.13 and their predictive utility for future dates. According to the figure, between model days 5-7, the end product of the decoding model (and presumably the neural strategy of the monkey) appears to have begun to stabilize as it able to predict control signal intent relatively well several weeks into the future.
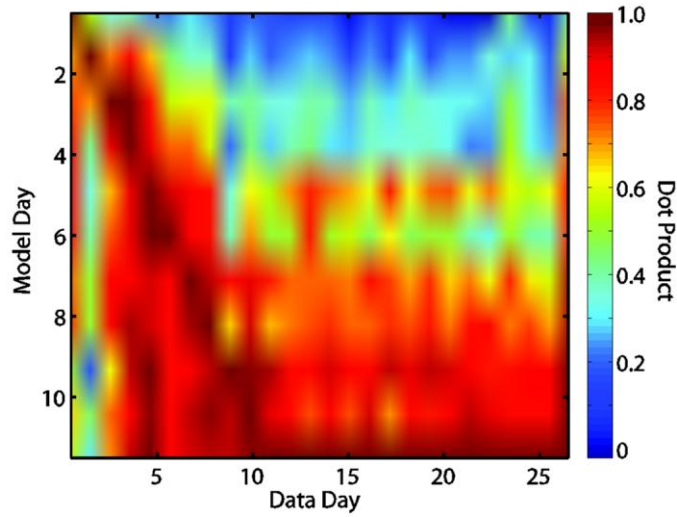
Figure 8.14 Model predictive stability over time. Colors indicate the normalized dot product between a model's prediction of control vector data on a given test day's data set and the actual control vector data produced from the model used on that day. Model weights were fixed after Day 11.

The previous result speaks to the longevity of the decoding model once it has stabilized assuming all control channels remain intact. However, we still desired to examine how robust the entire feature set would be with possible channel failures assuming the decoding model would be retrained. To this end, a similar analysis to that originally shown in Figure 8.10 was performed with results for monkey K shown in Figure 8.15. Comparing the two figures demonstrates a similar trend in the minimum number of channels needed per dimension (2-4) for maximum performance as well as redundancy in the feature information such that even an array of smaller size could theoretically be very robust to channel loss.
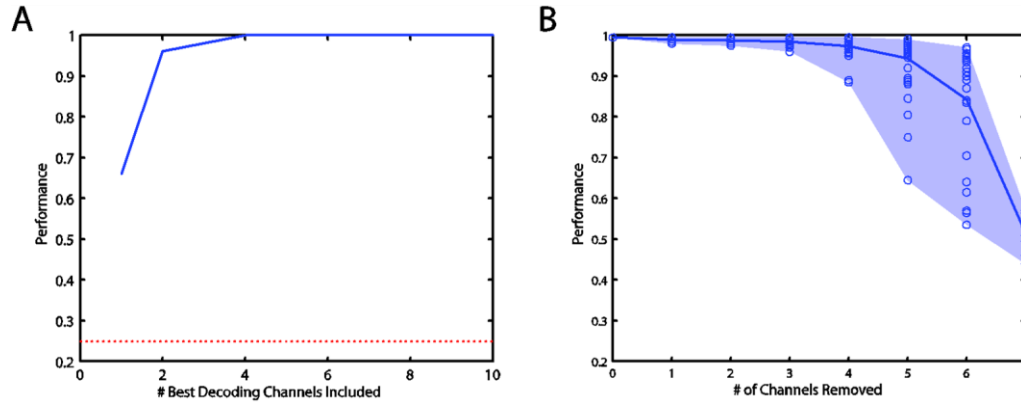
Figure 8.15 Decoding robustness analysis, monkey K. (A) Post-hoc analysis of decoding performance (blue line) with increasing number of channels used in decoding algorithm. (B) Simulation of decoding performance with channel loss. See Figure 8.10 for full description of analyses.

8.4 Conclusions

Directly comparing the two sets of experiments presented in this chapter is difficult as many differences in their experimental setups exist. First, although the scale and electrode spacing of the two ECoG array designs were similar, other design considerations such as signal referencing and anatomical coverage varied to a greater extent. The training and control progression for monkeys J and M was also much different than that for monkeys K and P. The first two monkeys progressed from a small number of fixed control channels on 1-D tasks, then to 2-D tasks, then to adaptive multi-channel control, and finally to multi-channel/multi-band control. The latter monkeys used the last control scheme from the very start and thus did not experience any biases the other 1-D training or control schemes may have introduced. In addition, individual differences within a small group of subjects could certainly account for some of the performance differences observed.

The two sets of results do offer anecdotal evidence, however, for differences in a BCI decoding model that may arise as a result of the training progression that is used. With initiation of BCI training using a few arbitrary channels, it appears that

biofeedback allows the subject to hone in on those select features and modulate them very well toward the task. However, it is also apparent that adaptation on the part of the subject is not completely sufficient, and he or she may still benefit from adaptation by the decoding algorithm.

When BCI training is initiated with adaptation from the beginning, it appears that changes in directional modulation patterns might take a more global scope over the course of 5-7 days while differences between active and rest states appear closer to the onset of training. Following this adaptation, however, the ECoG array presents a robust feature set from which control can be achieved with a small number of channels. Whether or not this ensemble neural feature set could support multiple independent degrees of freedom in this manner is unclear from these experiments, but this is definitely an avenue of research that needs to be addressed in the future.

9 Conclusions

9.1 Summary of Results

The series of experiments and analyses in this thesis were intended to address the feasibility of a force-based BCI system using micro-ECoG signals for possible use in applications such as FNS. The general approach to this goal was to investigate it from three different research avenues: open-loop neurophysiological recordings, closed-loop BCI tasks, and post-hoc analysis of learning and adaptation.

Our first approach to this primary goal was to examine ECoG signals over various cortical points along the motor planning chain to identify possible flexibility in task-related movement parameters. We utilized a joystick center-out task in which we remapped the physical position of the joystick to different perceived movement characteristics of a cursor on a screen in order to further dissociate physical movements from the end-effector movements directly responsible for completing the task. Our original hypothesis suspected that cortical areas further up the motor plan might hold higher level representations of the desired cursor movement. The results we observed, however, suggested that most areas anterior of sensory/parietal areas correlated with the physical speed of the monkey's hand/joystick, regardless of which mapping was utilized. More posterior areas showed some task-related activity possibly aligned with the beginning of a trial, but correlations with cursor movement profiles were not readily apparent in this study. Given these results, we hypothesize that lower-level planning areas normally closer to actual motor output might actually be well suited to controlling a range of BCI signal applications. As humans, we are quite adept at utilizing one set of motor commands to effect movement of an external tool or device whose movement characteristics do not directly or linearly correlate with the original control signal. Yet, we are still confident in the resulting movement further down the tool or object's chain of cause and effect as if the tool is merely an extension of our body.

For our second approach, we presented monkeys with several closed-loop BCI tasks in which they had to utilize neural signals in a manner more akin to dynamics-based control than has been presented before. The monkeys initially developed their own neural strategy to complete kinematics- (velocity) based control in radial-choice and center-out tasks by adapting the decoding algorithm to their inferred strategy on a daily basis. Following these experiments, the "velocity" decoding model was fixed and recast as both acceleration and force-based decoding models. Use of an acceleration-based control signal in a center-out BCI task revealed that the monkeys, although with more difficulty in certain aspects, were able to produce cursor movement profiles resembling natural reaching profiles. Under those control conditions, the monkeys exhibited active biphasic (positive and negative peaked) control signal profiles in order to accelerate and decelerate the cursor appropriately. Depending on the particular implementation of the force control signal and the number of independent features required, the ability to modulate a feature in this manner could be important to the success of such a system. Examination of the areas presumably most causal for the output control signal revealed that similar areas were primarily responsible for producing both the velocity and acceleration control signals. This finding suggests that the particular movement parameter to be decoded for a BCI sysem may not be as important as the modulatory characteristics of the features that will control it. For example, it is likely to be more effective to adapt a feature strongly modulated by velocity features to instead control a force command rather than use a feature significantly but weakly modulated by force.

Next, the monkeys were expected to once again adapt their neural strategy within a static decoding model in order to produce forces toward interacting in a simple virtual environment. During these experiens, monkeys demonstrated proficiency at the novel mass-grab task as well as adaptation to various environmental variables (target mass, gravity, force gain). The discriminability metric of the adaptation to these variables also provided valuable insight into possible system specifications unique to our ECoG-derived control signal (signal-to-noise ratio for graded force, etc.). More importantly, this demonstration is one of the first to showcase force-mediated

interaction with variable objects within a simulated environment as opposed to simply "pointing" or selecting objects. While other BCI systems have demonstrated interactions with external objects such as using a robotic arm to grab food for feeding [38], subjects in this study were responsible for calculating the necessary forces for desired movements rather than an inverse-model controller. We do not necessarily want to postulate that a force-based control signal is necessarily the best method for controlling such interactions, but the results in this thesis at least open the door for further exploration of this approach.

Following the closed-loop experiments, we used post-hoc analysis of two different BCI setups - one in which subjects were initially forced to adapt to an arbitrary control scheme and one in which the algorithm adapted to the user - to investigate training times and neural adaptation to a BCI system. In both cases, it is clear that biofeedback is likely to be crucial for subjects to successfully learn any new BCI system. In addition, adaptation of the decoding model to the subject can produce striking improvements in performance. At this point, however, it is unclear whether one approach to initial BCI training is strictly "better" for the long term success of a system. Constraining control to specific channels may be difficult to master initially, but it may conserve potentially useful neural features that could be applied toward additional degrees of freedom later on. Conversely, allowing the subject to utilize all features from the very start of training may significantly decrease training time. However, if all available neural features are then used redundantly toward the current operating task, it could theoretically be more difficult to use some features toward a new independent control dimension later if the first control dimension must first be "unlearned". Perhaps a combination of the two training paradigms as mentioned in Chapter 8 (subject adaptation to algorithm followed by algorithm adaptation to subject) might hold an optimal compromise. The studies in this thesis were not designed to address the neural changes associated with later expansion to more dimensions, but it is an important consideration that should be investigated in future experiments.

Finally, the task vs. rest decoding in Chapter 7 provided interesting insight into how the brain might possibly transition into a state most suitable for BCI control. First,

we suggested that different spectral bands might hold specialized utility in a hybrid BCI system, i.e. high gamma frequencies for fine decoding of movement intent, and lower frequencies for gross determination of active or resting states. A hybrid system of this nature could use lower spectral features to turn on or off a device, and then use higher gamma features for the actual operation of the device. From those results, it also appears that effective BCI control might require a transition to a different spectral energy state (higher gamma powers, lower alpha and beta amplitudes) around which to opreate rather than modulating around baseline. An appropriate analogy might be how we carry and use our arms in everyday life. During much of the day, our arms typically hang down by our sides in their "baseline" state. While this state is more energy efficient, our options for movement around this state are limited: we do not often move our arms posteriorly or handle any objects below waist level in this state. On the other hand, when we want to utilize our hands, we often bring our arms and hands up to a more outstretched position at chest level. From this position, we can move our hands up, down, side to side, etc., with relative ease to accomplish a number of tasks. Again, it may require more energy to maintain our limbs in this active position, but it also makes performing a variety of tasks easier as well.

9.2 Future Directions and Final Thoughts

9.2.1 Force Control Algorithms

In this first attempt at dynamic control of an object, we utilized a relatively simple scheme of linear transformation from neural signals to a given control dimension, both positive and negative valued. In context of the neuromuscular system, however, perhaps a non-linear transformation to a force signal might be more appropriate. For example, muscles in the body typically act in a paired, complimentary fashion... generally a muscle to flex a joint and a muscle to extend it. With this reasoning, perhaps neural control features should be processed in a similar manner. Muscles can only contract in one direction or else relax. Conversely, in the scheme used in this thesis, a muscle linked to the control signal would have to exert a "negative"

force. Thus, if features were thresholded in a way such that they only produced a positive control signal for their designated muscle once activiated beyond a certain point, perhaps total acceleration and deceleration of an object in such a BCI setting would feel more natural to the user such as stepping on the gas and brake in a car to accelerate and decelerate.

9.2.2 Force and Acceleration Models

In this study, we relied on adaptation by the subject to produce a force-like control signal. This rationale was based on a lack of a consistent model with which to infer the subject's desired force or acceleration. Future experiments in this area might benefit from experiments that could bring such models to light. One example might be a circle drawing task in which the position, velocity, and acceleration profiles should be very similar but are synced out of phase with each other.

In addition, the current experiments were based in extrinsic coordinates. For an FNS system designed to reanimate a natural arm, however, the coordinate system is more likely to be based on the intrinsic coordinates of the arm. Thus, a simple next step for the current system would be to apply a similar control model to apply joint torques to a simulated two-dimensional arm model.

9.2.3 Furthering Adaptive Decoding

The adaptive control scheme used in this thesis appeared by and large to be very successful. Subjects appeared to gain significant insight into controlling the cursor within days of beginning training. However, the adaptation we employed was on a relatively slow and possibly noisy time scale. For example, if the monkey were to become frustrated and change neural strategies mid-session, this could hinder performance improvement provided by the decoder. Current work by other members in our lab, however, is investigating faster time courses for adaptation by the decoding algorithm on the order of minutes. This could very well be a staple for future BCI devices. It is also interesting from a neuroscience point of view to examine how fast this adaptation by the subject can occur while learning a new BCI skill.

9.2.4 Future BCI Control Signal Applications

Finally, as mentioned in Chapter 6, the experiments in this thesis were designed with the overall goal of one day applying them to restore natural motor function through techniques such as FNS. In reality, though, limiting the results of these results to only a few applications would be shortsighted. As demonstrated in this thesis, the brain is quite capable of applying signals from a relatively limited area of cortex to a variety of output functions. It may be well in the future, but brain-assisted technologies may very well revolutionize how we drive our vehicle to work or operate complicated machinery while leaving our hands available for other tasks and vice versa. It may seem farfetched now, but it is exciting to think of the seemingly limitless possibilities along those lines. In the future, BCI devices may not only be thought of as medical devices intended for those with debilitating diseases, but devices made to change the way we proceed with our everyday life, similar to the way many of us wonder how we ever lived without being able to check our e-mail and surf the internet on our smart phones. The author is excited to be a part of the research that may one day lead to those exciting technologies.

References

[1]   J. R. Wolpaw and D. J. McFarland, "Control of a two-dimensional movement signal by a noninvasive brain-computer interface in humans," *Proc. Natl. Acad. Sci. U.S.A*, vol. 101, no. 51, pp. 17849–17854, Dec. 2004.

[2]   A. Kübler, F. Nijboer, J. Mellinger, T. M. Vaughan, H. Pawelzik, G. Schalk, D. J. McFarland, N. Birbaumer, and J. R. Wolpaw, "Patients with ALS can use sensorimotor rhythms to operate a brain-computer interface," *Neurology*, vol. 64, no. 10, pp. 1775–1777, May 2005.

[3]   M. D. Serruya, N. G. Hatsopoulos, L. Paninski, M. R. Fellows, and J. P. Donoghue, "Brain-machine interface: Instant neural control of a movement signal," *Nature*, vol. 416, no. 6877, pp. 141–142, 2002.

[4]   D. M. Taylor, S. I. H. Tillery, and A. B. Schwartz, "Direct cortical control of 3D neuroprosthetic devices," *Science*, vol. 296, no. 5574, pp. 1829–1832, Jun. 2002.

[5]   C. S. Bjornsson, S. J. Oh, Y. A. Al-Kofahi, Y. J. Lim, K. L. Smith, J. N. Turner, S. De, B. Roysam, W. Shain, and S. J. Kim, "Effects of insertion conditions on tissue strain and vascular damage during neuroprosthetic device insertion," *Journal of neural engineering*, vol. 3, p. 196, 2006.

[6]   J. C. Williams, "Performance of chronic neural implants: measurement, modeling and intervention strategies," Arizona State University, 2001.

[7]   W. J. Freeman, M. D. Holmes, B. C. Burke, and S. Vanhatalo, "Spatial spectra of scalp EEG and EMG from awake humans," *Clinical Neurophysiology*, vol. 114, no. 6, pp. 1053–1068, Jun. 2003.

[8]   G. Pfurtscheller, B. Graimann, J. E. Huggins, S. P. Levine, and L. A. Schuh, "Spatiotemporal patterns of beta desynchronization and gamma synchronization in corticographic data during self-paced movement," *Clinical neurophysiology*, vol. 114, no. 7, pp. 1226–1236, 2003.

[9]   E. C. Leuthardt, G. Schalk, J. R. Wolpaw, J. G. Ojemann, and D. W. Moran, "A brain–computer interface using electrocorticographic signals in humans," *Journal of Neural Engineering*, vol. 1, pp. 63–71, 2004.

[10]   J. A. Wilson, E. A. Felton, P. C. Garell, G. Schalk, and J. C. Williams, "ECoG factors underlying multimodal control of a brain-computer interface," *IEEE Trans Neural Syst Rehabil Eng*, vol. 14, no. 2, pp. 246–250, Jun. 2006.

[11] G. Schalk, K. J. Miller, N. R. Anderson, J. A. Wilson, M. D. Smyth, J. G. Ojemann, D. W. Moran, J. R. Wolpaw, and E. C. Leuthardt, "Two-dimensional movement control using electrocorticographic signals in humans," *J Neural Eng*, vol. 5, no. 1, pp. 75–84, 2008.

[12] Z. C. Chao, Y. Nagasaka, and N. Fujii, "Long-term asynchronous decoding of arm motion using electrocorticographic signals in monkeys," *Front Neuroeng*, vol. 3, p. 3, 2010.

[13] S. Kellis, K. Miller, K. Thomson, R. Brown, P. House, and B. Greger, "Decoding spoken words using local field potentials recorded from the cortical surface," *Journal of neural engineering*, vol. 7, p. 056007, 2010.

[14] W. Wang, A. D. Degenhart, J. L. Collinger, R. Vinjamuri, G. P. Sudre, P. D. Adelson, D. L. Holder, E. C. Leuthardt, D. W. Moran, M. L. Boninger, and others, "Human motor cortical activity recorded with Micro-ECoG electrodes, during individual finger movements," in *Engineering in Medicine and Biology Society, 2009. EMBC 2009. Annual International Conference of the IEEE*, 2009, pp. 586–589.

[15] E. C. Leuthardt, Z. Freudenberg, D. Bundy, and J. Roland, "Microscale recording from human motor cortex: implications for minimally invasive electrocorticographic brain-computer interfaces," *Neurosurgical focus*, vol. 27, no. 1, p. E10, 2009.

[16] S. Thongpang, T. J. Richner, S. K. Brodnick, A. Schendel, J. Kim, J. A. Wilson, J. Hippensteel, L. Krugner-Higby, D. Moran, A. S. Ahmed, D. Neimann, K. Sillay, and J. C. Williams, "A Micro-Electrocorticography Platform and Deployment Strategies for Chronic BCI Applications," *Clinical EEG and Neuroscience*, vol. 42, no. 4, pp. 259–265, Oct. 2011.

[17] C. Henle, M. Raab, J. G. Cordeiro, S. Doostkam, A. Schulze-Bonhage, T. Stieglitz, and J. Rickert, "First long term in vivo study on subdurally implanted micro-ECoG electrodes, manufactured with a novel laser technology," *Biomed Microdevices*, vol. 13, no. 1, pp. 59–68, Feb. 2011.

[18] A. G. Rouse, S. R. Stanslaski, P. Cong, R. M. Jensen, P. Afshar, D. Ullestad, R. Gupta, G. F. Molnar, D. W. Moran, and T. J. Denison, "A chronic generalized bi-directional brain–machine interface," *J. Neural Eng.*, vol. 8, no. 3, p. 036018, Jun. 2011.

[19] A. G. Rouse, "Neural Adaptation and the Effect of Interelectrode Spacing on Epidural Electrocorticography for Brain Computer Interfaces," Washington University, 2011.

[20]   J. Simeral, S. Kim, M. Black, J. Donoghue, and L. Hochberg, "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*, vol. 8, p. 025027, 2011.

[21]   E. V. Evarts, "Relation of pyramidal tract activity to force exerted during voluntary movement," *J. Neurophysiol*, vol. 31, no. 1, pp. 14–27, Jan. 1968.

[22]   W. T. Thach, "Correlation of neural discharge with pattern and force of muscular activity, joint position, and direction of intended next movement in motor cortex and cerebellum," *J. Neurophysiol*, vol. 41, no. 3, pp. 654–676, May 1978.

[23]   D. R. Humphrey, E. M. Schmidt, and W. D. Thompson, "Predicting measures of motor performance from multiple cortical spike trains," *Science*, vol. 170, no. 3959, pp. 758–762, Nov. 1970.

[24]   M. C. Hepp-Reymond, U. R. Wyss, and R. Anner, "Neuronal coding of static force in the primate motor cortex," *J. Physiol. (Paris)*, vol. 74, no. 3, pp. 287–291, 1978.

[25]   A. P. Georgopoulos, J. Ashe, N. Smyrnis, and M. Taira, "The motor cortex and the coding of force," *Science*, vol. 256, no. 5064, pp. 1692–1695, Jun. 1992.

[26]   A. P. Georgopoulos, R. E. Kettner, and A. B. Schwartz, "Primate motor cortex and free arm movements to visual targets in three-dimensional space. II. Coding of the direction of movement by a neuronal population," *J Neurosci*, vol. 8, no. 8, p. 2928–37, 1988.

[27]   R. E. Kettner, A. B. Schwartz, and A. P. Georgopoulos, "Primate motor cortex and free arm movements to visual targets in three-dimensional space. III. Positional gradients and population coding of movement direction from various movement origins," *J Neurosci*, vol. 8, no. 8, p. 2938–47, 1988.

[28]   A. B. Schwartz, R. E. Kettner, and A. P. Georgopoulos, "Primate motor cortex and free arm movements to visual targets in three-dimensional space. I. Relations between single cell discharge and direction of movement," *J Neurosci*, vol. 8, no. 8, p. 2913–27, 1988.

[29]   Q. G. Fu, D. Flament, J. D. Coltz, and T. J. Ebner, "Temporal encoding of movement kinematics in the discharge of primate primary motor and premotor neurons," *Journal of Neurophysiology*, vol. 73, no. 2, pp. 836 –854, Feb. 1995.

[30]   W. Wang, S. S. Chan, D. A. Heldman, and D. W. Moran, "Motor Cortical Representation of Position and Velocity During Reaching," *Journal of Neurophysiology*, vol. 97, no. 6, pp. 4258 –4270, Jun. 2007.

[31]   J. Ashe and A. P. Georgopoulos, "Movement parameters and neural activity in motor cortex and area 5," *Cereb Cortex*, vol. 4, no. 6, pp. 590–600, 1994.

[32]   D. W. Moran and A. B. Schwartz, "Motor cortical representation of speed and direction during reaching," *J Neurophysiol*, vol. 82, no. 5, p. 2676–92, 1999.

[33]   A. B. Schwartz, D. W. Moran, and G. A. Reina, "Differential representation of perception and action in the frontal cortex," *Science*, vol. 303, no. 5656, pp. 380–383, Jan. 2004.

[34]   A. Marathe and D. Taylor, "Decoding position, velocity, or goal: Does it matter for brain–machine interfaces?," *Journal of Neural Engineering*, vol. 8, p. 025016, 2011.

[35]   J. M. Carmena, M. A. Lebedev, R. E. Crist, J. E. O'Doherty, D. M. Santucci, D. F. Dimitrov, P. G. Patil, C. S. Henriquez, and M. A. L. Nicolelis, "Learning to Control a Brain–Machine Interface for Reaching and Grasping by Primates," *Plos Biol*, vol. 1, no. 2, p. e2, 2003.

[36]   A. B. Schwartz, X. T. Cui, D. J. Weber, and D. W. Moran, "Brain-controlled interfaces: movement restoration with neural prosthetics," *Neuron*, vol. 52, no. 1, pp. 205–220, 2006.

[37]   S. I. Helms Tillery, D. M. Taylor, and A. B. Schwartz, "Training in Cortical Control of Neuroprosthetic Devices Improves Signal Extraction from Small Neuronal Ensembles," *Reviews in the Neurosciences*, vol. 14, no. 1–2, pp. 107–119, 2003.

[38]   M. Velliste, S. Perel, M. C. Spalding, A. S. Whitford, and A. B. Schwartz, "Cortical control of a prosthetic arm for self-feeding," *Nature*, vol. 453, no. 7198, p. 1098–101, 2008.

[39]   J. B. Zimmermann, K. Seki, and A. Jackson, "Reanimating the arm and hand with intraspinal microstimulation," *Journal of Neural Engineering*, vol. 8, p. 054001, 2011.

[40]   N. Nannini and K. Horch, "Muscle recruitment with intrafascicular electrodes," *Biomedical Engineering, IEEE Transactions on*, vol. 38, no. 8, pp. 769–776, 1991.

[41]   P. H. Peckham, K. L. Kilgore, M. W. Keith, A. M. Bryden, N. Bhadra, and F. W. Montague, "An advanced neuroprosthesis for restoration of hand and upper arm control using an implantable controller," *J Hand Surg Am*, vol. 27, no. 2, pp. 265–276, Mar. 2002.

[42]   C. T. Moritz, S. I. Perlmutter, and E. E. Fetz, "Direct control of paralysed muscles by cortical neurons," *Nature*, vol. 456, no. 7222, pp. 639–642, Oct. 2008.

[43]   D. Blana, R. F. Kirsch, and E. K. Chadwick, "Combined feedforward and feedback control of a redundant, nonlinear, dynamic musculoskeletal system," *Medical & Biological Engineering & Computing*, vol. 47, no. 5, pp. 533–542, Apr. 2009.

[44]   E. Chadwick, D. Blana, J. Simeral, J. Lambrecht, S. Kim, A. Cornwell, D. Taylor, L. Hochberg, J. Donoghue, and R. Kirsch, "Continuous neuronal ensemble control of simulated arm reaching by a human with tetraplegia," *Journal of neural engineering*, vol. 8, p. 034003, 2011.

[45]   L. R. Hochberg, M. D. Serruya, G. M. Friehs, J. A. Mukand, M. Saleh, A. H. Caplan, A. Branner, D. Chen, R. D. Penn, and J. P. Donoghue, "Neuronal ensemble control of prosthetic devices by a human with tetraplegia," *Nature*, vol. 442, no. 7099, p. 164–71, 2006.

[46]   R. Wahnoun, J. He, and S. I. Helms Tillery, "Selection and parameterization of cortical neurons for neuroprosthetic control," *Journal of neural engineering*, vol. 3, p. 162, 2006.

[47]   K. A. Ludwig, R. M. Miriani, N. B. Langhals, T. C. Marzullo, and D. R. Kipke, "Use of a Bayesian maximum-likelihood classifier to generate training data for brain–machine interfaces," *Journal of neural engineering*, vol. 8, p. 046009, 2011.

[48]   S. Mason and G. Birch, "A brain-controlled switch for asynchronous control applications," *Biomedical Engineering, IEEE Transactions on*, vol. 47, no. 10, pp. 1297–1307, 2000.

[49]   A. Bashashati, R. K. Ward, and G. E. Birch, "Towards Development of a 3-State Self-Paced Brain-Computer Interface," *Computational Intelligence and Neuroscience*, vol. 2007, pp. 1–8, 2007.

[50]   G. Pfurtscheller, B. Z. Allison, C. Brunner, G. Bauernfeind, T. Solis-Escalante, R. Scherer, T. O. Zander, G. Mueller-Putz, C. Neuper, and N. Birbaumer, "The hybrid BCI," *Front Neurosci*, vol. 4, p. 30, 2010.

[51]   G. Pfurtscheller, T. Solis-Escalante, R. Ortner, P. Linortner, and G. R. Müller-Putz, "Self-paced operation of an SSVEP-Based orthosis with and without an imagery-based 'brain switch:' a feasibility study towards a hybrid BCI," *IEEE Trans Neural Syst Rehabil Eng*, vol. 18, no. 4, pp. 409–414, Aug. 2010.

[52]   N. Achtman, A. Afshar, G. Santhanam, B. M. Yu, S. I. Ryu, and K. V. Shenoy, "Free-paced high-performance brain–computer interfaces," *Journal of neural engineering*, vol. 4, p. 336, 2007.

[53]   M. S. Fifer, M. Mollazadeh, S. Acharya, N. Thakor, and N. Crone, "Asynchronous decoding of grasp aperture from human ECoG during a reach-to-grasp task," in

Engineering in Medicine and Biology Society, EMBC, 2011 Annual International Conference of the IEEE, 2011, pp. 4584–4587.

[54]   A. Kubler, F. Nijboer, J. Mellinger, T. M. Vaughan, H. Pawelzik, G. Schalk, D. J. McFarland, N. Birbaumer, and J. R. Wolpaw, "Patients with ALS can use sensorimotor rhythms to operate a brain-computer interface," *Neurology*, vol. 64, no. 10, p. 1775–7, 2005.

[55]   P. Horki, T. Solis-Escalante, C. Neuper, and G. Müller-Putz, "Combined motor imagery and SSVEP based BCI control of a 2 DoF artificial upper limb," *Med Biol Eng Comput*, vol. 49, no. 5, pp. 567–577, May 2011.

[56]   E. C. Leuthardt, G. Schalk, J. R. Wolpaw, J. G. Ojemann, and D. W. Moran, "A brain-computer interface using electrocorticographic signals in humans," *J Neural Eng*, vol. 1, no. 2, pp. 63–71, 2004.

[57]   D. P. Doane and L. E. Seward, "Measuring Skewness: A Forgotten Statistic?," *Journal of Statistics Education*, vol. 19, no. 2, 2011.

[58]   A. P. Georgopoulos, A. B. Schwartz, and R. E. Kettner, "Neuronal population coding of movement direction," *Science*, vol. 233, no. 4771, p. 1416–9, 1986.

[59]   D. W. Moran and A. B. Schwartz, "Motor cortical representation of speed and direction during reaching," *J Neurophysiol*, vol. 82, no. 5, p. 2676–92, 1999.

[60]   D. A. Heldman, W. Wang, S. S. Chan, and D. W. Moran, "Local Field Potential Spectral Tuning in Motor Cortex During Reaching," *IEEE Trans. Neural Syst. Rehabil. Eng.*, vol. 14, no. 2, pp. 180–183, Jun. 2006.

[61]   D. M. Taylor, S. I. Tillery, and A. B. Schwartz, "Direct cortical control of 3D neuroprosthetic devices," *Science*, vol. 296, no. 5574, p. 1829–32, 2002.

[62]   D. W. Moran and A. B. Schwartz, "Motor Cortical Representation of Speed and Direction During Reaching," *Journal of Neurophysiology*, vol. 82, no. 5, pp. 2676 – 2692, Nov. 1999.

[63]   D. Moran, "Evolution of brain-computer interface: action potentials, local field potentials and electrocorticograms," *Current Opinion in Neurobiology*, vol. 20, no. 6, pp. 741–745, Dec. 2010.

[64]   D. A. Heldman, W. Wang, S. S. Chan, and D. W. Moran, "Local field potential spectral tuning in motor cortex during reaching," *IEEE Trans Neural Syst Rehabil Eng*, vol. 14, no. 2, p. 180–3, 2006.

[65]   A. G. Rouse and D. W. Moran, "Neural adaptation of epidural electrocorticographic (EECoG) signals during closed-loop brain computer

interface (BCI) tasks," *Conf Proc IEEE Eng Med Biol Soc*, vol. 2009, p. 5514–7, 2009.

[66]   L. H. Snyder, A. P. Batista, and R. A. Andersen, "Coding of intention in the posterior parietal cortex," *Nature*, vol. 386, no. 6621, pp. 167–170, Mar. 1997.

[67]   J. Wolpaw and E. W. Wolpaw, *Brain-Computer Interfaces: Principles and Practice*. Oxford University Press, 2012.

[68]   J. Liu, S. Ji, and J. Ye, "SLEP: Sparse Learning with Efficient Projections." Arizona State University, 08-Oct-2010.

[69]   P. Shenoy, K. J. Miller, J. G. Ojemann, and R. P. N. Rao, "Generalized Features for Electrocorticographic BCIs," *IEEE Transactions on Biomedical Engineering*, vol. 55, no. 1, pp. 273–280, Jan. 2008.

[70]   A. Levin and J. Wyman, "The Viscous Elastic Properties of Muscle," *Royal Society of London Proceedings Series B*, vol. 101, pp. 218–243, Apr. 1927.

[71]   A. Desplantez, C. Cornu, and F. Goubel, "Viscous properties of human muscle during contraction," *J Biomech*, vol. 32, no. 6, pp. 555–562, Jun. 1999.

[72]   S. Haufe, M. S. Treder, M. F. Gugler, M. Sagebaum, G. Curio, and B. Blankertz, "EEG potentials predict upcoming emergency brakings during simulated driving," *Journal of Neural Engineering*, vol. 8, no. 5, p. 056001, Oct. 2011.

[73]   G. Townsend, B. Graimann, and G. Pfurtscheller, "Continuous EEG classification during motor imagery--simulation of an asynchronous BCI," *IEEE Trans Neural Syst Rehabil Eng*, vol. 12, no. 2, pp. 258–265, Jun. 2004.