Washington University in St. Louis

Washington University Open Scholarship

5-24-2012

Neural Adaptation and the Effect of Interelectrode Spacing on Epidural Electrocorticography for Brain-Computer Interfaces

Adam Rouse
Washington University in St. Louis

Follow this and additional works at: https://openscholarship.wustl.edu/etd

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
Eric C. Leuthardt
Robert E. Morley
Lawrence H. Snyder
Kilian Q. Weinberger

Neural Adaptation and the Effect of Interelectrode Spacing on Epidural
Electrocorticography for Brain-Computer Interfaces

by

Adam Gene Rouse

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

May 2012

Saint Louis, Missouri

Abstract

Electrocorticography (ECoG) is increasingly being identified as a safe and reliable recording technique for both Brain-Computer Interface (BCI) applications as well as neurophysiology studies. This thesis describes some of the first real-time closed-loop BCI studies of chronic ECoG in non-human primates. Epidural microECoG electrodes developed in our lab were implanted in three monkeys with the electrode array centered over primary motor cortex (M1). Monkeys were then trained to perform a one-dimensional BCI task. The BCI control scheme was independent of any prior screening for task-related activity. All three monkeys successfully learned to perform the task with multiple control configurations and each time gained significant performance in 10 days or less. Interelectrode distance between control electrodes was also tested for three different distances. 15 and 9 mm spacing resulted in equivalent performance while 3 mm saw a moderate but significant degradation in performance. Finally, post hoc analysis was performed to analyze various decoding parameters. While decoding parameters were generally well matched to the observed signals, several potential decoding improvements were identified. Overall, these results demonstrate the feasibility of epidural ECoG BCIs, highlight the importance of neural adaptation for BCI control, and quantify various metrics of a current ECoG BCI system to drive further studies.

Acknowledgements

Washington University in St. Louis

May 2012

Table of Contents

List of Tables

List of Figures

1 Introduction

Cogito ergo sum. I think, therefore I am. As Descartes so elegantly stated, the ability to think lies at the core of our existence. Someone is unlikely to argue that sensory input or motor output lie at the same fundamental level to justify the reality of being a person. However, the human experience for most of us is often defined as much or more by the interactions with the world around us as it is by our own internal thoughts. Diseases of the nervous system are perhaps the most feared of all human diseases because they disrupt the thoughts and actions we believe are so central to the human experience. Paralysis and other motor disorders are a subset of nervous system diseases that specifically limit a person's ability to turn their thoughts into actions. While many of these patients are able to think and have normal cognitive function, their ability to fully express themselves and their thoughts is severely limited.

Prosthetic devices that reduce a person's reliance on the their peripheral nervous system and tap directly into the neural signals of the central nervous system are known as brain-computer interfaces (BCI) or brain-machine interfaces (BMI). BCIs can be broken down into two main categories. Sensory BCIs are neuroprosthetics that provide sensory information into the brain as a replacement or augmentation of the normal senses (i.e. cochlear implants for deafness). Conversely, motor BCIs take intentions from the brain and translate them to control an external device directly instead of through a user's normal musculoskeletal system. The work presented here focuses exclusively on motor BCIs and their potential for controlling an external device.

There are several different recording modalities that can potentially be used for providing the neural signals used in a motor BCI. The work presented here uses a recording technique known as electrocorticography (ECoG). ECoG measures the voltage at an electrode site below the skull on the surface of the brain. While ECoG can refer to recordings that have electrodes placed either above or below the dura mater covering the brain, all of the recordings described here were done using epidural electrodes. These experiments are one of the first times that chronic electrocorticographic electrodes were implanted into non-human primates. The electrode response and recording stability were measured throughout the experiments to quantify their performance and assess their suitability for long-term ECoG recordings.

1.1 Specific Aims

Three specific aims were selected for examination in this thesis. Aims I and II represent a series of BCI experiments that were designed and conducted while Aim III was a series of analyses that were conducted on the collected dataset. The aims are as follows:

Aim I- Determine the initial time course of adaptive learning in a BCI task where naïve monkeys develop control using chronically implanted epidural ECoG electrodes. Pairs of electrodes will be arbitrarily assigned to control the one dimensional velocity of a computer cursor in a closed-loop BCI task. Through training over the course of days to weeks, the animal will learn to accurately modulate gamma band (75-105 Hz)

brain activity to complete the task. The accuracy and speed with which the task is completed will be measured to assess overall performance as well as the time course of learning.

Aim II – Examine the spatial extent of ECoG signals during a closed-loop BCI task by varying the distances between the two electrodes used for control. This aim will determine the minimal spacing between two independent control electrodes before their signals become too correlated for effective control. To quantify the effect of distance on ECoG recordings for BCI, the task will be completed using interelectrode spacings of 15, 9, and 3 mm. Additionally, the signals of adjacent electrodes to the control electrodes will be analyzed to examine the extent of cortical activation during the BCI task.

Aim III – Analyze both causal and non-causal ECoG signals during closed-loop BCI tasks for potential improvements in predictive power by changing various BCI decoding parameters. At any point, gamma band activity from two channels will be used for control. However, the raw, broadband signal from all 28 channels on the electrode grid will be recorded. The resulting dataset using this basic task provides the opportunity to make predictions about many key issues related to ECoG BCIs. *Post hoc* analysis will be performed by changing various signal processing parameters to identify possible improvements for control and drive future experiments.

1.2 Motivation

The amount of time and money devoted to studying motor BCIs has grown considerably over the past few years. A number of different approaches have been used to make significant advancements in the field. The studies described here have several novel characteristics. First, these are some of the first chronic ECoG studies performed in non-human primates. Second, all of the experiments involved a closed-loop BCI task where the subject was provided with real-time feedback which required adaptation of the neural signals to control a cursor on a video monitor.

The mapping from recorded neural code to intended motor movement is less obvious for electrocorticographic signals as compared to the signal from single neurons. While the ability to predict movement kinematics based on recorded signal is possible with both, single neurons appear to have the advantage of more obvious decoding algorithms with more straightforward mapping. Field potentials appear to be a more gross measure of motor intentions. Much of the information that is obtained is at the level of movement versus rest as compared to movement in a given direction. While signals have been decoded and shown to be correlated with movement in different directions, the predictive accuracy is less than what is observed for a population of individual neurons.

When it comes to predicting the usability of neural signals for a BCI, the level of correlation between these signals and movement kinematics only tell half of the story. In a closed-loop BCI task, the user receives feedback of the external control resulting from the

recorded brain signals. It is possible that the ability of the user to adjust the neural control

signals based on feedback will prove to be just as important as a BCI designer's initial

decoding of the neural signals. In fact, the experiments described here demonstrate that

with virtually no initial decoding and only user adaptation, BCI control is still possible.

These experiments are not designed to prove that that traditional decoding is unnecessary,

but rather to argue that understanding both sides, decoding and adaptation, will be key for

building a better BCI.

To date, many ECoG studies have depended on using human patients with

intractable epilepsy. In neuroscience research, the ability to gain access and conduct

experiments with human subjects is a valuable tool. However, these types of studies have

limitations as clinical considerations must always take precedent over the secondary research

goals. Electrode grids are usually placed where the clinicians believe is best for patient care.

Additionally, once the clinical monitoring goals have been met, the electrodes are removed;

therefore, long time course studies (greater than one to two weeks) are not possible. By

studying ECoG signals in the laboratory setting with non-human primates, experiments are

conducted in a more controlled environment for much longer periods of time.

Our experiments were designed with several key features to take advantage of the

controlled laboratory setting to improve the quality of the data. First, all three monkeys

were implanted with the exact same surgical procedures with the goal of placing electrodes at

the exact same anatomical location over primary motor cortex. Second, by being able to

perform months of stable recordings, the experiment could be designed to examine the

subject's adaptation to a predefined decoding scheme. Finally, the exact same subject could

be switched between several different decoding schemes over the course of the study.

Additionally, great care was taken in these experiments to try to provide real, physical

units as well as measures of signal to noise ratios. Percentage or fractional changes from

baseline are often the only measurements included in some published studies.   While these

results may meet the scientific standard of showing a difference between two experimental

conditions, they are limited in their ability to drive future hardware and device design. For

practical BCI devices to move out of the lab and into the clinical setting, hardware

manufacturers must know the technical specifications required when designing new devices.

Another limitation of many currently published ECoG studies for BCI development

is the fact that they are often mapping or screening studies. Recordings are made while

subjects are performing a task and the recordings are then analyzed off-line to try to predict

some parameter associated with the task such as movement kinematics, language

components, or sensory input. Many of these studies do identify useful BCI control

parameters and a few studies do go on to use the developed decoding schemes in closed-

loop BCI tasks where the subject is provided with direct feedback of how his/her neural

signals are controlling the task.

While these screening experiments are effective initial studies to explore potential

control signals, they are limited by the fact that potential neural signals can only be selected

based on their correlation to either observable behavioral responses or experimentally

6

presented cues. These signals are often graded against one another based solely on their ability to predict these cues or responses. This type of exploration ignores possible differences in the adaptability of signals once they are used in a closed-loop system. For instance, while one signal may have a higher correlation to a certain arm kinematic parameter during screening, another signal may prove to be easier for a subject to adapt to and volitionally control during BCI tasks to improve performance even though it had a lower predictive value in open-loop screening sessions. Additionally, some signals may not be correlated to any observed behavior, but through closed-loop feedback and training the signal may still be a viable BCI signal. For less invasive techniques like ECoG and EEG this ability to create and enhance novel control signals without direct motor correlates may be especially critical to their viability as a BCI modality.

In many cases, creating BCI decoding schemes from observed screening sessions is not a trivial task with a direct mapping. While single unit based BCIs often use paradigms where the subject behaves in a similar task like center out reaching for a screening and brain control task, many ECoG BCI paradigms rely on more abstract mappings. To date, the best ECoG signals with the largest signal-to-noise ratio are often those created by gross motor movements of different parts of the body and utilize the spatial separation across cortex based on the topographic layout of different body regions (the motor homunculus). A great amount of research effort has been directed at examining all of the different behavior parameters that can be extracted from ECoG recordings for various experimental tasks. These studies provide us with useful insight into the underlying neurophysiology for coding and performing various tasks but once again there is no guarantee that these are the best

signals for BCI control. Just as learning to use a steering wheel or joystick allows a user to manipulate a car or complex machinery without the user having to make movements similar to the end effectors, there is no reason for BCI users to only perform BCI movements or tasks that mimic their original behavior.

As a more pragmatic issue, more non-human primate studies are likely needed to further advance ECoG BCI development. These studies can evaluate long-term stability as well as provide a more controlled environment with multiple subjects all implanted with the same surgical procedure and location and experiments performed under identical conditions. The language limitations with monkeys, however, make it unfeasible to instruct them to translate motor imagery into brain control commands. Experimenters are thus left with only two mechanisms for training monkeys for BCI control. First, a training task that very closely matches the BCI task and provides suitable brain signals for decoding can be used that makes the transition from normal behavior to BCI straightforward and natural.   Second, the subject can be operantly conditioned to improve performance once brain control is initiated. While the first is perhaps ideal, it is also limiting. The work presented here will highlight the fact that the second mechanism is quite feasible and in fact BCI control can be achieved with only this mechanism in complete absence of neural recordings during a training task.

2 Background

Neurological diseases that impair basic motor function, such as ALS, spinal cord injuries, and stroke have a dramatic impact on patients' lives. These motor deficits severely limit patients' abilities to perform everyday tasks and place a large burden on their families and society. Non-invasive assistive devices that use the translation of head, mouth, or eye movements to create a desired output allow many patients to restore some function and independence. However, patients are often still limited in the tasks they are able to perform and often these devices are slow to use and require long periods of training. Additionally, assistive devices typically require some residual motor function that the patients with the most severe motor disabilities may not have. Thus, an ideal solution for restoring function would be to link the normal functioning cortical brain activity directly with an output system. This establishment of a direct brain-computer interface would allow these individuals to circumvent damaged neurological connections to use their brain activity to directly accomplish tasks they were previously unable to complete successfully.

Currently, approximately 1.9% of the U.S. population, or about 5.6 million people, reported some form of paralysis (Figure 2.1) (Christopher & Dana Reeve Foundation, 2009). Initially, BCIs for lost motor function would most likely target the most severely paralyzed patients with spinal cord injuries, amyotrophic later sclerosis (Lou Gehrig's disease), and brainstem strokes. However, more successful and less-invasive motor BCIs could eventually allow for an increased candidate population for those with less severe forms of paralysis. Additionally, while the traditional concept of a BCI is one designed for a user with intact,

normally functioning cerebral cortex, recent post-stroke rehabilitation studies have suggested

that motor imagery (Zimmermann-Schlatter et al., 2008) and EEG-based BCIs to foster and

provide feedback to enhance motor imagery (Ang et al., 2010) can lead to increased recovery

from diseases that affect cortical function.

Causes of Paralysis
N= 5,596,000

Other 526,000 **9%**

Post-Polio Syndrome
272,000 **5%**

Cerebral Palsy
412,000 **7%**

Neurofibromatosis
212,000 **4%**

Stroke
1,608,000 **29%**

Unspecified
Birth Defect
110,000 **2%**

Traumatic Brain
Injury 242,000 **4%**

Multiple Sclerosis
939,000 **17%**

Spinal Cord Injury
1,275,000 **23%**

Figure 2.1. Paralysis Statistics.

Paralysis prevalence statistics from a recent study published by the Christopher & Dana Reeve foundation. The data was collected by an extensive household survey and represents an increase in prevalence compared to previous estimates based on hospital-reported clinical data.

2.1 Motor Cortex

The primary motor cortex (area M1, Brodmann area 4) is a strip of cerebral cortex

that runs just anterior to the central sulcus and makes up the posterior portion of the frontal

lobe. Its connection to motor movements was first identified by Penfield (1937) by

observing that electrical stimulation of this area elicited motor movements during surgery of

patients with focal epilepsy. Interestingly, in 1949, Penfield and Jasper (1949) also examined

the beta rhythm (18-30 Hz) over motor cortex and reported that it showed a reduction with

voluntary motor movements. Brindley and Craggs (1972) further identified that specific

activation occurred with different motor movements could be identified from the cortical

surface recordings of a baboon and suggested its usefulness for motor BCIs. These

electrophysiology findings also co-localized with the observed cytoarchitecture of large Betz

cells that project directly from the area to the motor tracts of the spinal cord and helped

identify primary motor cortex as one of the key cortical areas of motor output (Betz, 1874).

In the 1980s, previously developed techniques for examining single neuron activity

led to greater exploration of the underlying neural code observable in primary motor cortex.

The work by Georgopoulos (1982) and others highlighted the correlation between the

activity of individual neurons in the arm area of motor cortex and the direction of

movement that was observed. Further work has highlighted numerous different movement

variables that are coded for by neurons in motor cortex (Georgopoulos et al., 1986;

Schwartz, 1994; Moran and Schwartz, 1999a; 1999b; Schwartz and Moran, 1999; Reina et al.,

2001; Wang et al., 2007). There is still debate about what motor parameters make up the

core variables coded for by motor cortex as evidence for both extrinsic (movement kinematics) and intrinsic (muscle activations) coordinate encoding are often observed (Kakei et al., 1999).

In addition to the observable correlation of neurons to voluntary motor movements, as early as the 1970s, it has also been shown that monkeys can be operantly conditioned to modulate individual neurons in motor cortex (Fetz & Baker, 1973). Additionally, studies using both field potential recordings like EEG (Pfurtscheller & Neuper, 1997) and ECoG (Leuthardt et al., 2004) as well as imaging techniques such as fMRI (Porro et al., 1996) showed that motor imagery tasks where no overt muscle activation actually occurred showed similar cortical activity (albeit sometimes with diminished amplitude) to actual motor movements. Thus, cortical activity from motor areas represented more than just the resulting motor movements but also information about a subject's voluntary cortical state when performing imagery tasks or directing brain control of an external device.

It was the combination of these scientific findings that formed the vision of a potential brain-computer interface that could restore or augment motor function by translating direct cortical signals from the brain into motor intentions for control of some type of external actuator. By taking recorded signals from the identified motor cortical areas, the signals could be decoded using the underlying knowledge of motor neurophysiology to recreate motor movements. Furthermore, since many motor areas showed activity even without actual movements, BCIs that used only motor imagery should also be feasible. Finally, since the cerebral cortex is well-known for its plasticity, it was likely that users would

improve their BCI performance with practice and learning even if initial movement decoding was far from perfect.

2.2 BCIs for Motor Control

Several different electrophysiology recording modalities have been proposed and studied as the potential output control signal for a brain-computer interface. All of these methods record extracellular microvolt-level potentials from a single neuron or collection of neurons. However, these different recording modalities require trade-offs between invasiveness and robustness of signal. At the one extreme, electroencephalography (EEG) based BCIs rely on signals from electrodes placed non-invasively on the scalp. While one and two-dimensional control has been demonstrated with EEG in humans (Kübler et al., 2005; Wolpaw & McFarland, 2004) the accuracy rates in these tasks remain lower than other methods of BCI and require longer training periods to achieve a given level of control.

At the other extreme, single neuron based BCIs rely on fine tip intracortical electrodes. By placing ~20 μm electrode tips into the brain parenchyma, the action potentials of individual neurons can be recorded. These signals have been used successfully in a number of BCI set-ups in non-human primates. First, these systems were used for control of a cursor on a computer screen with multiple degrees of freedom (Wessberg et al., 2000; Serruya et al., 2002; Taylor et al., 2002). Additionally, single unit BCIs have also been designed and implemented for control of external devices through a robotic arm (Carmena et al., 2003; Velliste et al., 2008) . In addition to studies in non-human primates, human trials

13

have also been performed where single-unit activity was used for BCI control in patients with motor disabilities (Kennedy and Bakay, 1998; Kim et al., 2008). However, these types of recordings require complicated, highly invasive surgeries. This type of procedure has the risk of neurological and vascular damage as well as possible CNS infection (Bjornsson et al., 2006). Additionally, the quality of these recordings tends to decay over time as electrodes become encapsulated by the immunologically reactive tissue (Williams et al., 2007).

2.3 ECoG for BCI applications

Electrocorticography has shown growing promise in the field as an intermediate solution between the two extremes of recording single unit activity and EEG. Although the signals are not as direct of representation of the underlying neural code that can be observed at the single neuron level, the surgical procedure to implant ECoG electrodes is less invasive with a lower chance of cortical tissue damage or infection. Compared to EEG, ECoG provides better spatial and spectral resolution albeit at the trade-off of being an invasive technique as compared to the non-invasive EEG electrodes on the scalp.

Various components of ECoG signals have been shown to be well correlated to various parameters when subjects perform various motor tasks. Event-related potential (ERP) changes of the raw waveform in the time-domain of the ECoG signal have been used to identify the onset and timing of various motor actions on individual trials (Levine et al., 1999). Additionally, it has been demonstrated that event-related power changes recorded using ECoG can be used for mapping somatotopic areas of sensorimotor cortex associated

with visually cued movements of different body parts. A Fourier transform or Fourier-like

algorithm is performed over time windows during both a given task as well as at rest. Using

this analysis to estimate the power at different frequencies of the signal for the given epochs

of data, it is possible to identify different frequency components that increase or decrease in

power during the task compared to rest. Historically, two bands that have been specifically

identified were the alpha (8-13 Hz) and beta (15-25 Hz) bands which tended to show a

decrease in power with the onset of motor movement or imagery (Crone, Miglioretti,

Gordon, Sieracki, et al., 1998). Additionally, regions of the gamma band (30-100 Hz) have

been identified that show an increase in power during movement (Crone, Miglioretti,

Gordon, & Lesser, 1998).

These characteristic spectral features have been used for real time, closed-loop BCI

experiments with motor imagery tasks as the training paradigm (Leuthardt et al., 2004). By

first having patients perform various real and imagined motor movements, recorded ECoG

signals were analyzed to identify power spectrum features well correlated to the motor

behavior. These features could then be assigned to control the cursor kinematics during a

brain control task and the subject could be instructed on the necessary motor movement or

imagery necessary to complete the task (i.e. "Imagine opening and closing your hand to

move the cursor to the right"). Successful BCI control with ECoG recordings has been

demonstrated with these types of experimental paradigms in several instances (Blakely et al.,

2009; Felton et al., 2007; Schalk et al., 2008).

2.4 Improving ECoG BCI

ECoG recordings in non-human primates have only recently been conducted. Acute ECoG recordings have proved successful for open-loop mapping of sensorimotor cortex during reaching movements (Heldman, 2007). Although these recordings showed limited directional tuning for specific targets, the signals were well correlated to gross motor movements. Additionally, a series of successful closed-loop BCI tasks were completed using these acute ECoG recordings. In these experiments, microwire electrodes were acutely placed each day above the dura over the arm area of primary motor cortex. Recordings from two electrodes separated by about 1 cm were used to control a two-dimensional cursor to complete center-out and circle drawing tasks. Over the course of several weeks, monkeys were able to successfully complete approximately 40 center-out movements in 6 minutes or approximately 30 circle drawings in 7 minutes. This control was achieved by using the amplitude the high gamma frequency range of the signal between 65-100 Hz.

To achieve two-dimensional control, one electrode was used as the control for the horizontal velocity of the cursor and a separate electrode was used for the vertical velocity of the cursor. The circle drawing task provides a way to analyze how well a subject is able to independently control these two degrees of freedom. In order to draw a perfect circle using standard x and y Cartesian coordinates for control, it is necessary for the velocity control signals to be 90 degrees out of phase from each other such that the cursor can be directed to move in all directions over the course of each individual trial. At the beginning of the experiment the two recording sites tended to be correlated so that both were higher or both

lower power at the same time. In order for the monkey to improve his performance in the circle drawing task, it is necessary for the animal to gain independent control over the two signals being used for control. For a perfectly drawn circle, the overall correlation between the two signals will be zero.

This decorrelation could theoretically be done either indiscriminately across all frequencies or only within the frequency band being used for control. To examine what actually occurred during the experiment, the power spectrum was calculated for the two recorded signals in 300 ms non-overlapping time bins. The correlation between the powers at each given frequency for the two different channels was then calculated for all points in time. Figure 2.2 shows the resulting correlations for the five days of recordings and shows that the correlation between the recording sites dropped between 65-100 Hz. This data shows that through biofeedback, motor cortex is quite adaptable to learning and improving brain-computer interface control.

Figure 2.2. Power Spectrum Correlation

The correlation of the power spectrum of the *x* and *y* control electrodes during a circle drawing task. The signals become more decorrelated over the 5 days of control. The decorrelation is most dramatic within the band used for control but also occurs across most frequencies.

Recent studies have also demonstrated the long-term stability and robustness of ECoG recordings in monkeys over primary motor cortex. (Chao et al., 2010) These results showed that the predictive value of hand position from ECoG signals did not significantly decay over a period of five months. Additionally, the recordings were stable as cross-day predictions (generate coefficients one-day, validate on a different day) did not differ significantly from the accuracies using same-day coefficient prediction. This long-term

stability is a key advantage that ECoG-based BCIs offer compared to BCIs based on single-unit activity.

2.5 Neurophysiology of Field Potentials

Several recent studies have found that field potentials are tightly coupled with the underlying neural population. These studies tend to show that the power at the higher LFP frequencies: 60-150 Hz (Ray et al., 2008), 40-90 Hz (Rasch et al., 2008) provide the best estimates of underlying spike activity. Some recent studies have also argued that many of these high frequency changes are really the result of total broadband power changes which are the best predictor of underlying spike activity (Heldman et al., 2006; Schwartz et al., 2006; Manning et al., 2009; Miller et al., 2009). Studies examining the high gamma signals of ECoG recordings 60-200 Hz (Crone, Sinai et al. 2006) show that these signals as compared to lower frequencies are more specific in their timing and localization and are better aligned with our putative understanding of observable cortical activation. The similarity of high gamma amplitude changes to both electrical cortical stimulation mapping and fMRI imaging has lead to investigation of using high gamma ECoG activity for functional mapping using the amplitude between 70-100 Hz (Brunner et al. 2009). For a review, see Jerbi et al. (2009).

In addition to comparisons to underlying firing rate, LFP power changes have also been shown to display many of the same tuning properties that can be observed with single units. Studies of LFP recordings in motor cortex found that the frequency bands of 60-200 Hz (Heldman et al., 2004; Heldman et al., 2006) and ≤4 Hz and 63-200 Hz (Rickert et al.,

2005) had the most predictive power for direction of reaching and that the higher frequency

changes were more narrow in time and concentrated around only movement onset.

Additionally, a study by Liu and Newsome (2006) in area MT found that the LFP signal

from 40-150 Hz was correlated to the direction and speed tuning to visual stimuli found in

the multi-unit activity recorded from the same sites. These studies were careful to point out

that the high gamma LFP amplitude does not appear directly correlated to the recorded

spiking activity on a trial by trial basis but rather only shows similar tuning properties.

Additionally, (Pesaran et al., 2002) noted that gamma band LFPs in lateral intraparietal

cortex (LIP) reflect the area's columnar organization (Blatt et al., 1990); however, motor

cortex does not have as neatly ordered columns. These previous experiments seem to point

towards the high frequency component of the LFP being highly related to the observable

task-related spiking activity of nearby neurons even if it is not a direct surrogate of

instantaneous firing rate of simultaneously recorded multi-unit activity. Conversely, the low

frequency field potential recordings seem to be much more spatially broad and spread across

time suggesting that these signals are a more global process less related to event related

spiking activity.

3 Experimental Methods

In the following experiments, three male, 6-10 kg monkeys (one *Macaca mulatta* - J

and two *Macaca fascicularis* – M, N) were chronically implanted with epidural ECoG recording

grids. Prior to implantation, each monkey was trained to perform a radial choice task. This

task was learned through operant conditioning with liquid reward while sitting in a primate

chair. The animal sat in front of a standard 17" LCD monitor (Dell Inc.) approximately 20

inches in front of the monkey and controlled the cursor on the screen using a joystick

(APEM 9000 Series Joystick, APEM Components Inc.) before switching to brain control

following implantation.

3.1   EECoG Recordings

Once the monkey had been trained to complete the tasks using the joystick, the

monkey was implanted epidurally with a custom built ECoG grid through a 22 mm hole

made in the skull at a standard location that approximately aligns with the proximal arm area

of primary motor cortex (area M1).   The electrodes were 300 μm diameter platinum-iridium

(90-10 Pt-Ir) wires arranged in a 28 electrode hexagonal pattern with 3 mm interelectrode

distance. The electrodes were connected to standard Omnetics connectors (18 Position

Nano-Miniature Connector, Omnetics Connector Corp.) on a printed circuit board. A

protective chamber and cap made of titanium and stainless steel was cemented around and

above the hole containing the electrodes and connectors for protection when the animal was

Figure 3.1. ECoG electrode design.
300 μm Pt/Ir wires with 3 mm interelectrode distance were used for 28 signal channels and 6 selectable reference channels. The electrodes were embedded in PMMA. PMMA was also used to embed the entire chamber when placed on the skull. Connections were made with two 16 channel Omnetics connectors on a printed circuit board. The stainless steel outside chamber served as the ground. (Drawings courtesy of JJ Wheeler)

in its home cage. The whole ensemble was cemented in poly-methyl methacrylate (PMMA, i.e. dental cement). A drawing of the custom built ECoG grid is shown in Figure 3.1.

Each recording day, the animal's head was restrained, the protective cap was removed and the electrodes connected to a low-impedance head stage and digitized with the equivalent of 17.5-bit analog to digital resolution (digitization with oversampled 16-bit Sigma-Delta A/D conversion). Before digitization, the total amplification of the analog signal was 50x and was band limited between 3-500 Hz. The signal was sampled and processed at 6 kHz using a multi-channel neurophysiology recording system (Tucker-Davis Technologies) before being down sampled to 2 kHz for storage and later analysis. The raw

signal for any given electrode channel was processed by the digital signal processors to generate an amplitude estimate for the signal between 75-105 Hz. These control signals were then sent to a host computer responsible for controlling and displaying the task for the subject. The task scene was updated and displayed based on the incoming control signals. Figure 3.2 gives a schematic of the closed-loop recording system.

Following the experiments, monkey N was sacrificed and electrodes were localized to their anatomical locations. The acrylic headcap was removed principally intact by drilling
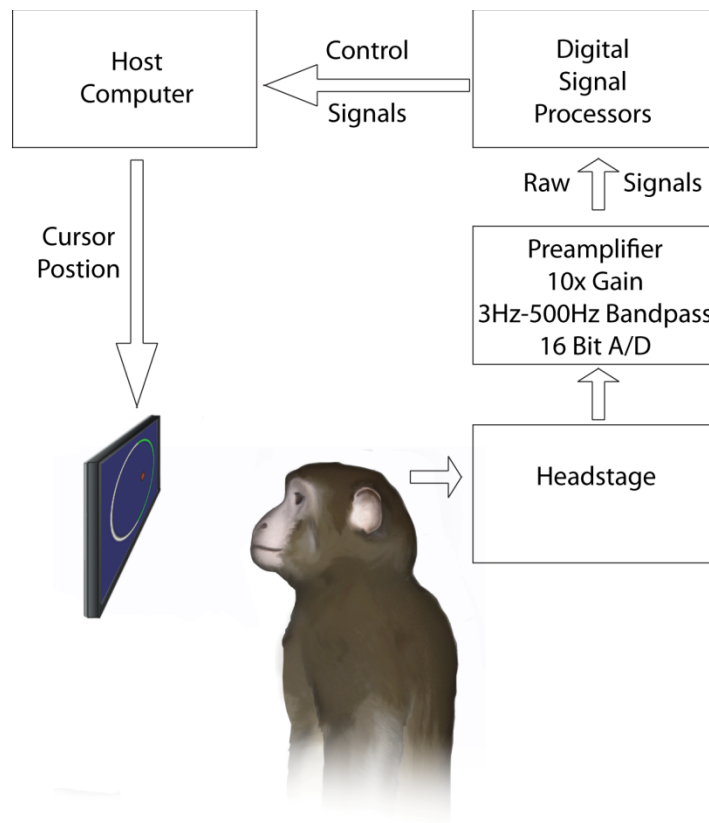
Figure 3.2. Experimental Set-up.
Flow diagram of closed-loop signal processing from recorded ECoG signals to display the cursor on the video monitor. (Drawing courtesy of JJ Wheeler)

out the connecting skull screws. Photographs of the headcap containing the ECoG grid as well as the exposed skull and brain surface were taken. Colocalization was completed by merging the two photographic images based on the observable landmarks of screw locations and chamber outline. The electrode locations overlaid on the exposed cortex is shown in Figure 3.3. The approximate electrode locations are shown for each monkey but are overlaid over the photographs obtained from only monkey N.

3.2 Radial Choice Task

In these experiments a radial choice task was implemented (Figure 3.4). The subject interacts with the task by controlling the velocity of a spherical cursor. A large circular ring serves as the target for the task. At the start of each trial, the cursor is returned to the center of the ring automatically. One hundred milliseconds after the cursor has been moved to the center of the screen, an arc of the ring is highlighted to signify the desired target. The subject must then move the cursor towards the portion of the ring that is highlighted. Once the cursor has moved through the correct target, the monkey is rewarded with a liquid reward. For a one-dimensional two target task, each target is a 180° arc. Therefore, for this task, completely random control of the cursor would result in the correct target being chosen 50% of the time. For a two-dimensional task, the number of targets can be increased with each target being a smaller portion of the circle. For example, a four target task has 90° target arcs with a chance level of 25%.

Figure 3.3. Electrode Placement.
ECoG electrode placement and control channels projected on the cortical surface of monkey N, post-mortem. The control electrodes are labeled for monkeys: a) M, b) J, and c) N. Monkey M and J are approximations based on projecting the grid and control electrodes on to monkey N's brain if the recording chamber was in the exact same location. The central sulcus (CS) and superior precental dimple (SPD) are labeled. The primary motor cortex (M1) is a strip that runs dorsal-ventral (up-down in the middle 3rd of the image) immediately rostral to the central sulcus until approximately the superior precentral dimple. Caudal to the central sulcus is primary sensory cortex (S1) (left portion of images a and c). On the rostral side of the precentral dimple is the dorsal premotor cortex (PMd) (right portion of images a and c). The closed-loop BCI experiments were run twice for each monkey, once controlling the horizontal (*x*) velocity and once controlling the vertical (*y*) velocity of the cursor. Control was always push-pull with two electrodes (+ and -) controlling the cursor in opposite directions.
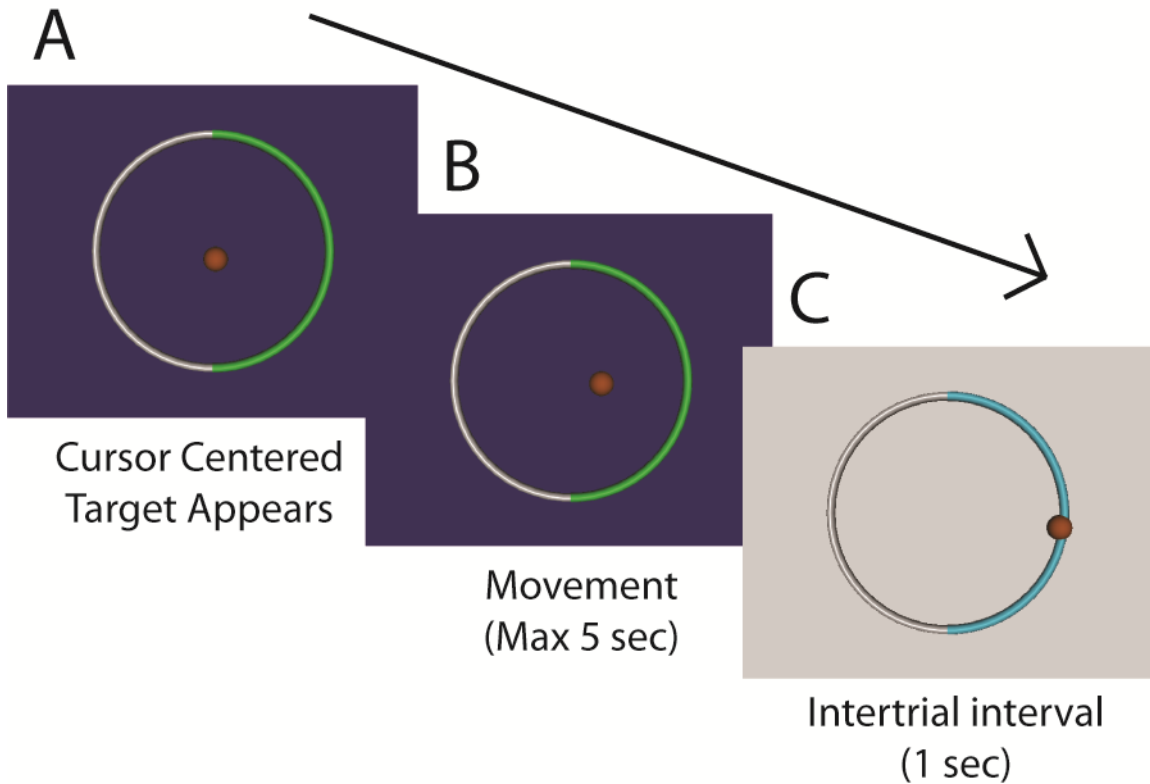
Figure 3.4. Radial Choice Task.
The one-dimensional, two-target radial choice task. a) At the start of the trial, the cursor is moved to the center and one of two targets appears. b) The monkey then has five seconds to select the correct target. c) Once a target is selected or the maximum movement time 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.

3.3 Closed-loop Tasks

Once the monkeys were trained and implanted, all of the actual experiments were closed-loop brain-computer interface tasks where the recorded brain signals were directly used to control the cursor displayed to the monkey. Each monkey attempted to move a cursor in one-dimension with recorded brain signals to complete the task which had the same visual appearance as the joystick task. A pair of recording electrodes was assigned to control either the horizontal or vertical velocity of the cursor in a push-pull control scheme.

Electrodes for were arbitrarily chosen with the only constraint being that the quality of their signal and total RMS amplitude was within the normal range of the rest of the channels on the grid.

For the one dimensional task, a pair of electrodes were assigned to control the cursor in a push-pull decoding scheme. The amplitude estimate of the signal between 75-105 Hz was mapped to control the velocity of the cursor. For one electrode, if the amplitude is greater than the mean value; the cursor moves in the positive direction (right or up). Conversely, if its amplitude is less than the mean value it moves in the opposite direction. The second electrode used for control also controls the velocity but in the opposite direction. Thus, if its amplitude is greater than the mean, the cursor moves in the negative direction (left or down) while a smaller amplitude moved the cursor in the positive direction.

In order to successfully complete a trial, modulation of the amplitude between 75-105 Hz could occur on either electrode. For either target, the subject needed to either increase the amplitude on one electrode or decrease the amplitude on the other electrode to move towards the correct target. It is the differential between these two electrodes' control signal that determined the actual velocity of the cursor and whether the correct target was selected. Using a push-pull scheme provides three main potential benefits to these experiments. First, it provides an added level of noise reduction as any transient increase of external noise will likely affect the recordings on both electrodes similarly and not cause the cursor to dramatically move in one direction or the other. Additionally, a push-pull control system allows us to explore the spatial separation that is necessary to get differential cortical

activity that can be used for control (i.e. independent modulation). Electrodes that are too close and have similar recorded signals will not allow for movement of the cursor in one direction or the other. Finally, it is possible that increasing and decreasing amplitude are processes that do not have similar time courses or range of modulation. By using a push-pull scheme, the cursor can be moved in either direction by an increase on one of the channels or a decrease on the other. All of these issues associated with push-pull control will be further explored in the analysis and discussion.

3.4 Control Signals

Our closed-loop BCI experiments relied on the translation of the recorded signal from two ECoG electrodes to a control signal that controlled the movement of a cursor displayed to the monkey. The sampled signals were digitally processed using our digital signal processing hardware and host computer to display the BCI task scene. The equations of this translation are as follows:

$$b(t) = f_{75-105Hz}[a(t)] \qquad (3.1)$$

$$c(t) = f_{<3Hz}(|b(t)|) \qquad (3.2)$$

$$d(t) = \frac{c(t) - \overline{c(t)}}{\sqrt{\frac{1}{n}\sum[c(t) - \overline{c(t)}]^2}} \qquad (3.3)$$

$$\dot{x}(t) = g \cdot [d_+(t) - d_-(t)] \qquad (3.4)$$

The control signal that was used to control the velocity of the cursor was an amplitude estimate of the signal between 75-105 Hz. This estimate was computed using a band-pass, rectify, low-pass filter algorithm. Starting with Eq. (3.1), the raw signal ($a(t)$) was first band pass filtered. Each given control channel was passed through a series of two 16^{th} order digital Butterworth filters, a 75 Hz high-pass filter and a 105 Hz low-pass filter, to give the resulting band passed signal ($b(t)$). Next, the 75-105 Hz signal ($b(t)$) was full wave rectified by taking the absolute value of the filtered signal. This rectified signal was then low-passed at 3 Hz (1^{st} order Butterworth) to generate an amplitude estimate ($c(t)$).

There was some variation in the overall amplitude and amount of modulation that occurred for various channels on a recording array. To minimize the effect of these differences in signal amplitude, each channels amplitude estimate was normalized as shown in Eq. (3.3). The first step in the normalization process was the subtraction of a running average. This running average was computed on the DSP with a time constant of 100 seconds. It should be noted that this running average was computed using the continuous stream of signals and included the signal recorded during previous trials. Thus, this mean value is not the mean observed only during rest nor inter trial intervals, but rather also included the modulated signal from previous trials. A 100 second time constant was chosen since for a maximum movement trial of five seconds plus one second inter-trial interval at least 16 trials should contribute significantly to the running average. This minimizes the possible effect of an uneven sampling of previous trials where the desired target was significantly biased towards one target. An alternative approach would be to choose a mean value for the day based on an initial calibration or data from the previous day. Our approach

has the benefit of creating a continuous BCI that can run autonomously and does not require any initial calibration or data analysis. Additionally, our system can potentially adapt to possible electrical noise and signal quality changes throughout the course of the day.

In addition to subtracting the mean, each channel's control signal was also normalized by the observed variance of the signal. This variance was calculated by taking the root mean square (RMS) of the amplitude of the signal minus its running average. Once again, this variance estimate was smoothed with a time constant of 100 seconds. The final control signal was then generated by subtracting the mean and dividing by the RMS.

In Eq. (3.3), $c(t)$ represents the amplitude estimate between 75-105 Hz. $\overline{c(t)}$ represents the 100 second running average of x and the denominator represents the RMS of $[c(t) - \overline{c(t)}]$ for the last 100 seconds. The generated control signal $d(t)$ from the two control channels (positive channel, $d_+(t)$ and negative channel $d_-(t)$) were then combined in a push-pull manner to control the horizontal ($\dot{x}$) or vertical ($\dot{y}$) velocity of the cursor. The gain term, g, in Eq. (3.4) controls the speed the cursor moves for any given amplitude change in the control signal.

Figure 3.5 shows a step by step example of the raw wave being processed to generate an amplitude estimate. The channel in this example was assigned to move the cursor to the left when there was an increase in amplitude and to the right when there was a decrease. In this example, the trial in panel A represents when a left target was presented at time zero to the monkey. After observing the desired target, the subject increased the amplitude between

75-105 Hz. This amplitude increase relative to the running average caused the cursor to move to the left to select the correct target. Likewise, in panel B, when a right target appeared, the monkey decreased the amplitude to correctly move to the right.

The mean control signal across five days of recordings for two different monkeys is shown in Figure 3.6. For monkey M, the positive channel (Figure 3.6a) control signal showed little variation between left and right trials. In contrast, the negative channel (Figure 3.6b) shows a large increase in amplitude for left trials as this amplitude was used by the monkey to move the cursor to the left or right and select the proper target. A second pair of plots shows the positive and negative control channels for monkey N. In this case, the positive channel (Figure 3.6c) shows an increase for trials with right targets while the negative channel (Figure 3.6d) shows an amplitude increase for left targets.

A

B

Figure 3.5. Single Trial Control Signals.
Example control signals from two individual trials from the same
control channel. The raw signal (blue) is band-pass filtered between
75-105 Hz (green), rectified (red), and then low pass filtered at 3 Hz
(black). For the first trial, an increase in amplitude (a) moved the
cursor towards the correct target. For the second example trial, the
opposite target appeared and a decrease (b) moved the cursor towards
the correct target.

Figure 3.6. Mean Control Signals
Example mean control signals from two different monkeys. The first monkey (M) held the positive channel (a) relatively constant while the negative channel (b) increased amplitude for left targets to move the cursor in the correct direction. The other monkey (N) used a push-pull scheme with an increase on the positive channel (c) for right targets and an increase on the negative channel (d) for left targets. The error represents the standard error of the mean.

3.5 Experimental Timeline

Each series of experiments had two phases. The first phase represented a training phase where the subject first learned to complete the task using a new pair of electrodes that were arbitrarily assigned to control the cursor. The two electrodes were always 15 mm apart during the training phase of the experiment (Figure 3.7a). During this phase, a bias was initially added to the velocity of the cursor to aid in initial training. Once the bias had been removed and the monkey had reached a plateau in performance, a second phase, the distance phase, was started. During this phase, control was switched each week to two electrodes that were 15, 9, or 3mm apart (Figure 3.7b). The experiments were always done in the same order going from farthest separation to closest. The same three pairs of electrodes were then used again and the experiments repeated for three more weeks. Thus, the distance phase of the experiments was conducted

Figure 3.7. Electrode Arrangement.
a) ECoG electrode arrangement showing the 28 recording channels with an example two channels (red) selected for push-pull control. The green electrode represents the reference electrode. b) During the distance phase of the experiment, three different spacings of 15, 9, and 3 mm were examined along a line of electrodes.

over a course of six weeks total. A timeline is shown in Figure 3.8. Adjacent electrodes along a single line were used for the distance experiments to reduce the amount of training and relearning that would be required of the monkey with the change in control electrodes from week to week. All three monkeys performed the one-dimensional radial choice series of experiments twice. During the first series, the brain signals were used to control the cursor along the horizontal axis. Next, the assigned control was moved to a new area on the array and the cursor moved along the vertical axis. All data will be presented and labeled for each instance of monkey-dimension combination. For example, M-X symbolizes the recordings from monkey M controlling the *x* dimension of the cursor.

Figure 3.8. Timeline.

The experimental timeline for each monkey and dimension. An initial training period is followed by 6 weeks of distance experiments where control electrodes are separated by 15, 9, and 3mm.

3.6 Data Analysis

One challenge with performing brain-control tasks with monkeys is that it is often difficult to discern when the subject is actively attempting to complete the task. Each day of the conducted experiments, the monkey performed the task for approximately 1.5 hours and usually completed between 1000 and 1500 trials. This amount of time was chosen based on previous training sessions such that the monkey should be well satiated with the proper

35

amount of reward by the end of the recording session. In nearly all cases, each monkey's

performance decreased towards the end of the day. Figure 3.9 illustrates this decrease in

performance by examining percentage correct as a function of actual trial number during the

day. A 100 trial sliding window was used for this plot.

Unfortunately, there tended to be considerable variation between monkeys in terms

of how soon after starting they stopped working and also how abruptly the transition was

between appearing to give maximal effort of correctly selecting targets and then performing

randomly at chance. Since we wanted the analysis to incorporate data where the monkey

was actively attempting the task, we chose to select the consecutive block of trials from each

day where the highest percentage of correctly selected targets occurred. To aid in choosing

the proper length of trials to use, we compared the observed data to a purely stochastic

model where each trial was an independent Bernoulli trial with an equally likely probability to

be correct. For example, if a given monkey averaged 1200 trials per day and 90% chance of

being correct on each trial, we calculated what the expected best percentage would be using

sliding trial windows of various sizes from 50 to 1000 trials. In Figure 3.10a, the observed

best percentage for various trial windows for each monkey-dimension combination is

plotted. As a comparison, the expected relationship based on a Bernoulli trial model is also

plotted. For this plot, the model is fitted to the data such that the percentages are equal

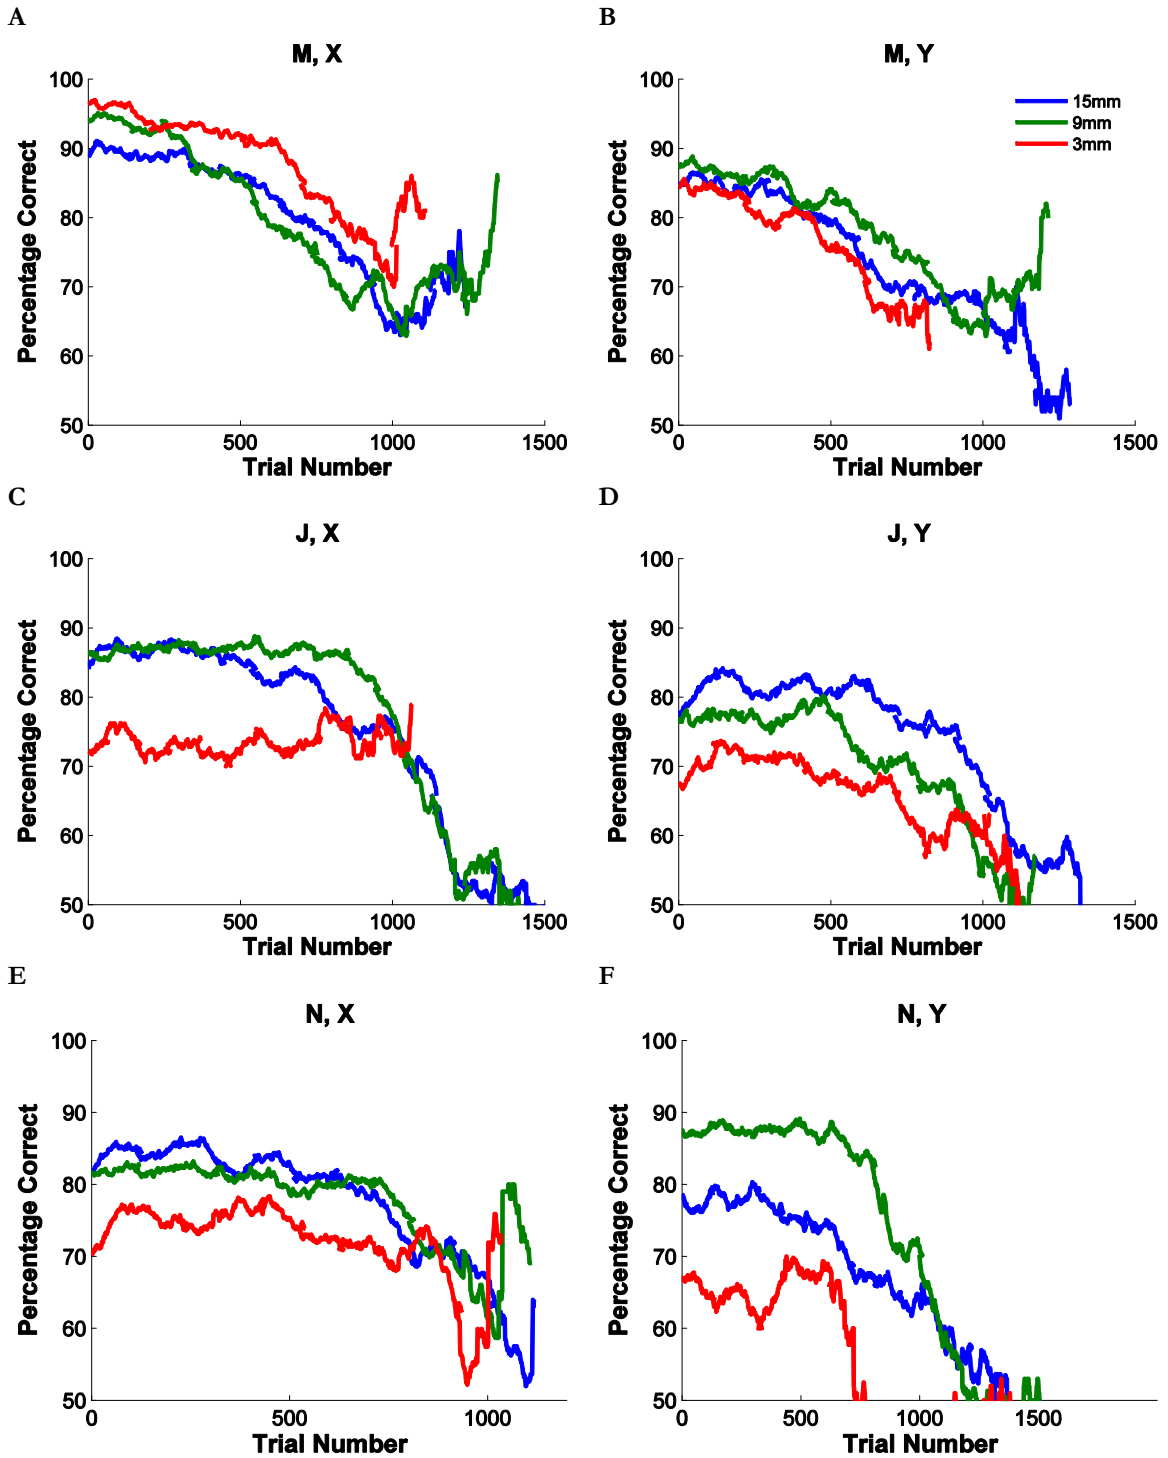when using the 400 trial window.

Figure 3.9. Performance through Day.
The percentage correct for the absolute trial number throughout the day. The six plots are for the different monkey dimension combinations. Percentages were calculated using a 100 trial sliding window. In all cases, the monkeys' performance decreased during the day as he became satiated and less motivated to complete the task.

37

The 400 trial window was chosen based on comparing the mean square error of the modeled data to the actual data for different trial windows. The mean square error was compared between modeled and observed data for all trial windows less than the trial window of interest. Figure 3.10b shows that the mean square error rises dramatically as the larger trial windows are included. This suggests that the monkeys were indeed not performing at the same level throughout the entire day but rather had stretches of better and worse performance. Since the performance level almost always dropped as a function of actual trial number through the day, we are confident that this effect was primarily caused by the monkey becoming satiated and being less motivated to perform the task. To provide for a consistent analysis of the data and only incorporate data where the monkeys were motivated, we chose to use the 400 best consecutive trials from each day since that appears to be a point in all the datasets where the monkey still gave maximum effort. Thus, except for some of the simulations presented in chapter 6, all data presented uses the 400 best consecutive trials from each day. Analysis was also done but not shown using the 100 best trials as well as all trials from each day and resulted in similar trends.

3.7 BCI Performance

Performance of the subjects during the closed-loop radial choice task can be characterized by the percentage correct and the speed each trial is completed. Speed and percentage can also be combined into a single metric of bit rate based on the percentage of targets correctly hit within a given amount of time. The per trial bit rate of data transfer for a discrete target task can be determined with the following equation (Pierce, 1980):

$$B = log_2N + Plog_2P + (1 - P)log_2\left(\frac{1-P}{N-1}\right) \text{ (3.5)}$$

Where N represents the number of targets and P is the percent correct expressed as a

fraction. The bits per trial, B, can then be multiplied by the trial rate (trials/second) to

produce the bit rate for a given amount of time. This relationship allows for the

combination of number of targets, percent correct, and speed into a single metric. The bit

rate is a standard metric used in information theory and represents the maximal rate that

information can be transmitted error free if an efficient error-checking algorithm is

employed. In BCI design and any other human interface, there are practical limitations to

whether this theoretical limit can be achieved in a way that is intuitive and not cumbersome

for the user. Even with this limitation, bit rate still provides a useful metric and is a standard

of the BCI field (McFarland et al., 2003).

A

B

Figure 3.10. Model for Different Time Windows.
a) Comparison between actual data (solid) and expected model (dashed) performance. Each trial window represents the best percentage observed for that given number of consecutive trials. For this plot, the model was matched to observed data at a trial window of 400 trials. b) The mean square error (MSE, in arbitrary units) between the observed data and the fitted model when matched at varying trial windows. MSE is calculated from all data points less than the point where the data was matched to the fit. (Contributions for b) from TM Pearce)

3.8 Amplitude Estimation

The band-pass, envelope detection technique provided a fast signal processing method to estimate the signal amplitude for real-time control and is also an estimation method that can easily be replicated on either an analog circuit or a small implantable chip. However, the reported amplitude can change depending on the response characteristics of the filters. Additionally, because no filter offers perfect frequency cutoffs, the amplitude estimated from a signal filtered between 75-105 Hz is not necessarily the sum of two amplitudes calculated between 75-80 Hz and 80-105 Hz. Also, depending on the shape of the waveform, the relationship of different amplitude estimates can be ambiguous. For example, while the peak amplitude, mean amplitude, and root mean square amplitude have a defined relationship that is the same for any single pure sine wave, these relationships are not fixed across different waveforms. To avoid these ambiguities, many fields that use signal processing use the root-mean-square (RMS) as an unambiguous measure of a signal's amplitude. RMS has several properties that make it useful for reporting signal amplitude. First, it is well defined for any signal even when the signal has a complex waveform and is non-repeating. Second, Parseval's theorem states that the sum of a signal's square is equal to the sum of the square of its Fourier transform. Thus, the RMS can be calculated in either the time- or frequency-domain. Since the RMS can be estimated in the frequency-domain, it is straightforward to calculate the band-limited RMS for any frequency range of interest.

In order to provide a better quantification of the signals and a more precise estimate of different frequency components, a power spectral estimation technique was used for

much of the signal analysis of the recorded data. The power spectral density can be estimated by computing the discrete Fourier transform to create a discrete spectrum. The square of the amplitude spectrum yields the power spectral density with units of V^2/Hz. Since we are performing the spectral estimate over a short time range with truncated end points, a simple periodogram results in spectral bias and has a high variance for each individual frequency point. These effects can be reduced by windowing the time-domain signal, smoothing across adjacent frequencies, and/or averaging multiple estimates. The Thomson multitaper method (pmtm in Matlab) is one such method and is the method we have chosen to use for our power spectral density estimation. It uses a series of Slepian sequences as orthogonal data tapers and is well suited for short time window signals with broadband signals across a wide spectrum of frequencies. For our analysis, a 1025 point spectrum was calculated from 0 Hz up to the Nyquist frequency of 1017 Hz yielding a spectral step resolution of approximately 1 Hz. Seven window tapers were used for the multitaper calculation. An example power spectrum is shown in Figure 3.11.

The amplitude of the signal during movement for a given frequency band was also estimated using the following equation:

$$RMS_{band} = \sqrt{f_{step} \cdot \Sigma_{f_{band}} PSD(f)} \quad (3.6)$$

Here, PSD is the power spectral density in units of $\mu V^2/Hz$ as a function of frequency. The sum is calculated over the desired frequency range. f_{step} is the frequency step of the discretely sampled power spectral density and was approximately 1Hz. The square root of the power was then computed to return an amplitude estimate in units of microvolts. Data analysis of the signals was performed by calculating the power spectrum of a given channel's signal during the entire movement portion of the task.

Figure 3.11. Example Power Spectrum.
**An example power spectrum from monkey M. For this electrode, an increase in amplitude between 75-105 Hz caused the cursor to go to the left. The two power spectrums represent trials where the right target (blue) and the left target (green) were displayed. Thus, this electrode was being properly modulated by changing the power of the signal between 75-105 Hz to correctly complete the task.**

The estimated RMS amplitude within the control band based on the movement power spectrum provides a single value to quantify a given channel's amplitude on a given trial. This single value provides a simple metric that can then be used in a variety of analyses. Once the amplitude of the signal was computed for each trial, a series of analyses were used to quantify the level, location, and frequency of modulation that was achieved by the subject to accomplish the task. Figure 3.12 shows the trial-by-trial histograms for a given channel. These plots represent the same examples used for the mean control signals plotted in Figure 3.6 but now give a better indication of the trial by trial variability. Monkey M used a large separation on one control channel (Figure 3.12b) while monkey N showed more of a push-pull decoding with modulation in the correct direction on both channels (Figure 3.12c and Figure 3.12d) but smaller mean differences.

One way to further analyze the discriminability of the signals is to use a variety of thresholds for a binary classifier to trace out the separability of the signals for the two target conditions. A curve can be created for the percentage of trials where the amplitude was above a given threshold on the two different target conditions. In Figure 3.13, these classification curves are plotted. The vertical axis represents trials where the right (or up) target appears and plots the percentage of these trials where the amplitude was above the given threshold. Similarly, the horizontal axis represents the percentage of trials above threshold but for trials where a left (or down) target appeared. The plotted line traces out a series of potential thresholds. These plots are similar to Receiver Operator Characteristic

A   B
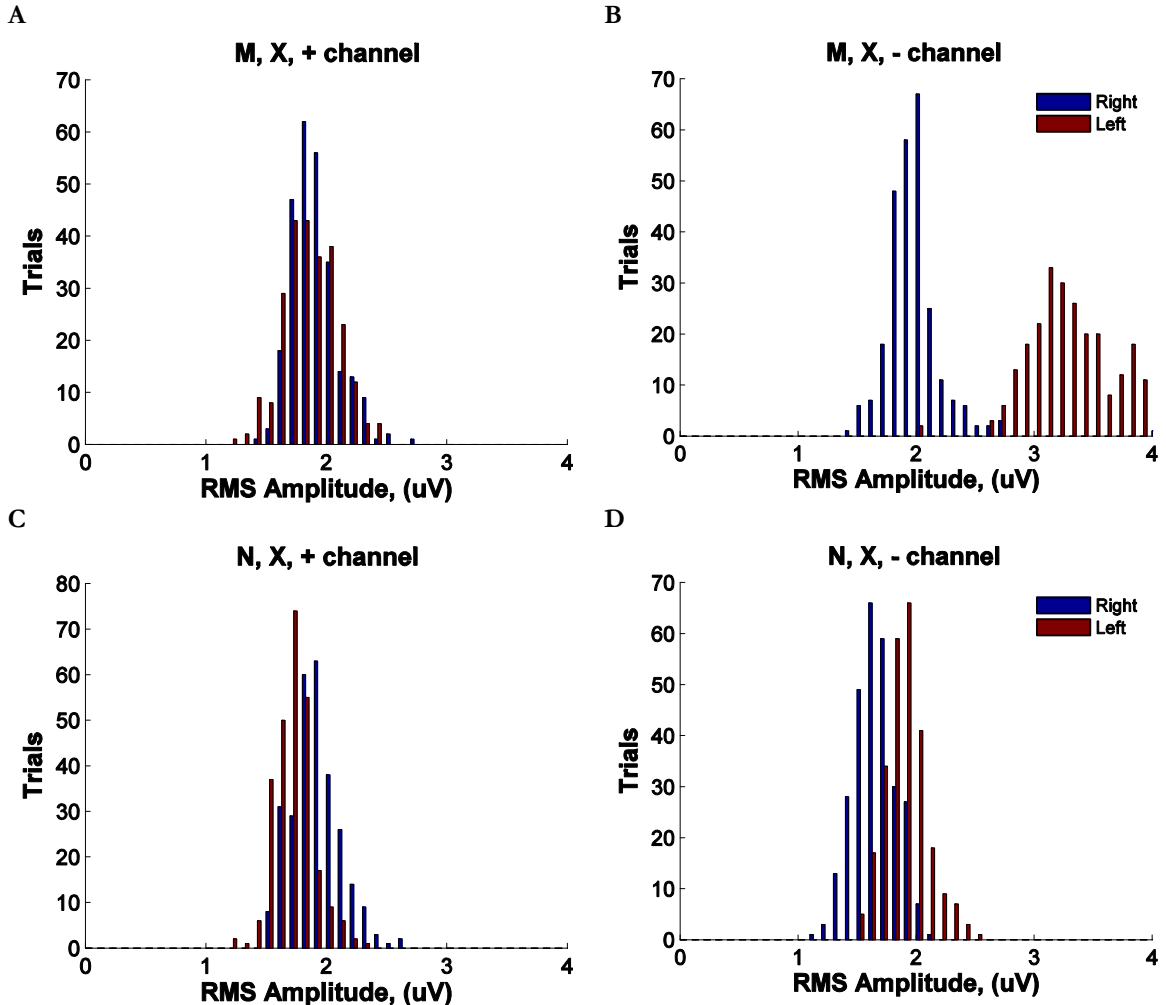
Figure 3.12. Trial Amplitude Histograms.
Example histograms showing separation in amplitudes for right versus left targets. As seen in Figure
3.6, the first monkey (M) held the positive channel (a) relatively constant while creating a wide spread
between right and left targets on the negative channel (b). The other monkey (N) used a push-pull
scheme with smaller differences that occurred on both the positive channel (c) and the negative
channel (d).

(ROC) curves except for the fact that there is no inherent hit versus miss condition but

rather we are trying to classify between two identical target conditions. For the positive

channel, where an increase in amplitude causes the cursor to move to the right (or up), the

right target condition should have a larger number of trials higher than any given threshold

45

for the left target condition. Thus, it is desirable for the positive channels to have values plotted in the upper-left portion of the plot. Conversely, the negative channel should have higher amplitude signals more often for left (or down) targets and thus should have values plotted in the lower-right half of the plot. A line along the 45° line represents equal percentages of right and left trials at any given threshold and thus chance performance in using the available signal to correctly predict left or right.

Figure 3.13 plots these classifier curves for each individual channel for the three different monkeys for each of the two different repetitions of the task they performed. This shows that there were four clear examples (M-X, J-X, J-Y, and N-Y) where one electrode is clearly being used for control while the other electrode stayed relatively constant for both conditions or even modulated slightly in the opposite direction than would be optimal for control. Conversely, the two other cases (M-Y and N-X) show that while neither electrode by itself had as large and distinct separation between the two conditions, both electrodes were being modulated in the correct direction.

Figure 3.13. Signal Discriminability.
Percentage of positive and negative targets above various thresholds for both the positive (blue) and negative (green) electrodes. For correct control, each positive channel trace should be in the upper-left half and the negative channel trace should be in the lower-right half for each monkey-dimension combination. a) monkey M, horizontal dimension b) M, vertical c) J, horizontal d) J, vertical e) N, horizontal f) N, vertical

To obtain a single measure for the amount of amplitude modulation that occurred at a given electrode for the two different target conditions, a d-prime (*d'*) statistic was used. *d'* or the sensitivity index is the difference of the mean of the signals for the two targets divided by their pooled standard deviation:

$$d' = \frac{\mu_R - \mu_L}{\sqrt{\frac{(n_R-1)\sigma_R^2 + (n_L-1)\sigma_L^2}{n_R + n_L - 2}}} \quad (3.7)$$

μ_R and μ_L represent the mean RMS for right and left targets. This difference is then normalized by dividing by the pooled standard deviation calculated from the standard deviation (σ_R, σ_L) of the right and left target distributions. n_R and n_L are the number of trials to right and left desired targets, respectively. *d'* is a classic measure of signal detection theory and should give a direct measure of how well the amplitude of the signal can be used to predict the correct target. If the distribution of the channel's amplitude is normal and of equal variance for both target conditions, *d'* values of 1, 2, and 3 would correspond to correctly classifying the two target conditions at 69, 84, and 93 percent correct, respectively. Figure 3.14 is a plot of the *d'* for each pair of control electrodes for the six different monkey-dimension combinations. Once again it highlights the tendency for M-Y and N-X to be truer push-pull while the other four relied predominantly on only one channel.

Figure 3.14. *d'* Statistic.
***d'* statistic for all 12 channels for the six different monkey-dimension combinations. Two of the combinations (M,Y and N,X) show true push-pull control while the other 4 combinations used primarily only one channel's modulation for control.**

3.9 BCI Control versus Rest

The control signals used from each electrode were normalized by comparing the current signal to the running average of the previous 100 seconds and dividing by the RMS of this control signal for the 100 second period. This constant updating is different than what is often done in other studies by comparing the movement signal to some rest or baseline period. Although we did not use it for closed-loop control, we did at the start of each day take a baseline recording for five minutes before the monkey started performing

49

the brain control task. During this time, no task was displayed on the screen as the monkey was sitting in the same primate chair and experiment room as during the actual task. A power spectrum was performed on these baseline recordings using one second non-overlapping sliding windows. For comparison, in this analysis only, a movement power spectrum was performed on the first second of data after the target appeared. Generally, the one second ended before the target was selected but if the trial ended before one second some of the reward period was included so that every trial resulted in exactly one second of data so it was perfectly matched to the baseline data. For each control electrode, three distributions of the RMS of the signal between 75-105 Hz were analyzed. The trials where the target was presented such that an increase in amplitude on the electrode caused the cursor to move towards the displayed target ("positive target"). The trials where the presented target could be reached by decreasing the amplitude on the electrode ("negative target"). And finally the baseline distribution made up from one second non-overlapping windows during the five minute baseline recordings. The mean RMS for both the positive targets (Figure 3.15a) and negative targets (Figure 3.15b) were compared to the mean RMS during the baseline recording period for each day of the different monkey-dimension combinations. Interestingly, while the positive target mean RMS was higher than baseline as was necessary to complete the task, the mean RMS during negative trials was almost always exactly equal to the mean baseline RMS. This suggests that while it would be beneficial for the subject to modulate the RMS below baseline this is much less likely to occur on an electrode than an increase above baseline. These results lead us to believe that modulating the amplitude up is the active process that the monkeys are actively accomplishing while

50

lowering the amplitude to go to the other target is more similar to simply resting or not actively attempting to modulate the brain signal.

It is also interesting to compare the spread of the distributions for the three conditions. Specifically, we were interested in the trial to trial variability that occurred while the monkey was performing the closed-loop brain control task and when the monkey was resting with no task to perform. Figure 3.16 plots the standard deviation of the trial RMS values for each of the three conditions as a function of the mean RMS values. During the task, there appears to be a consistent relationship between the mean RMS value for the day and the standard deviation of the RMS distribution. This relationship appears to hold for both sets of targets both when the amplitude should be modulated up or modulated down. An ANCOVA statistical test revealed that the there was no significant difference in the slope ($p=0.94$) or the intercept ($p=0.77$) of the two target conditions. When the data is combined together, the following relationship is observed:

$$RMS_{SD} = 0.102 \cdot RMS_{Mean} + 0.023 \text{ (3.8)}$$

This suggests that no matter what the amplitude of the signal the standard deviation of the distribution is always about one tenth of the mean amplitude the signal during the task. Surprisingly, during the baseline, the standard deviation of the amplitude between 75-105 Hz is considerably larger than during movement of the actual task with a significantly different slope than during the task ($p < 1x10^{-4}$). Here the relationship is:

$$RMS_{SD} = 0.194 \cdot RMS_{Mean} - 0.024 \quad (3.9)$$

This represents a near doubling of the amount of variability of in the RMS signal during

baseline for any given mean RMS.

   While these results were surprising, there are several potential explanations. First,

during baseline recordings there was no task or given behavior expected from the monkey.

The monkey was free to look, move, or think however he desired. Since this baseline state

was not as well controlled as during the actual task there could be a wider range of cognitive

states that the monkey moved between that resulted in a wider range of amplitude measures.

Another possibility could be caused by an increase in overall noise and variability caused by

motion or other factors because of increased restlessness of the monkey when no task was
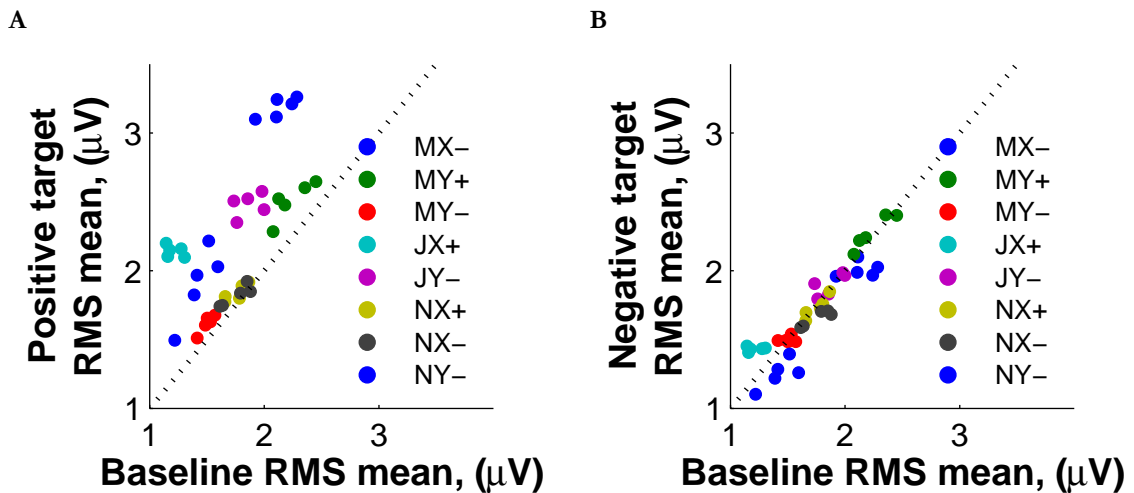
Figure 3.15. Mean RMS vs. Baseline.
The mean RMS values for each day for a) positive direction targets and b) negative target directions
compared to baseline RMS at the beginning of the day. Monkeys increased the amplitude for
positive targets but did not actively decrease the amplitude for negative targets.

Figure 3.16. RMS Variability.
The standard deviation of the trial RMS for each day in the three different
conditions (baseline, positive target, negative target) is plotted. The same
variability was observed for both movement conditions and was less than observed
during the baseline recordings.

occurring. Either way, it appears that when the monkeys were actively engaged in the task

there was decreased variability in the observed RMS amplitudes compared to baseline.

3.10 Spatial Analysis

Although only two channels were being used for control at any given time, all 28

channels were always recorded. Thus, it is possible to analyze the amount of signal

modulation occurring on all channels. Figure 3.17 illustrates the spatial distribution of *d'* for

all 28 channels. In order to successfully complete the task, the *d'* on the positive channel

must be positive and/or the negative channel must have a negative *d'*. It is clear that

53

adjacent electrodes tended to exhibit similar modulation to the electrode actually used for control. Additionally, the magnitude of modulation gradually falls away the farther the distance from the control electrode.

In order to get a more quantitative measure of the spatial scope of the brain signals, the d' values for all 28 recorded channels were fit with a two-dimensional spatial Gaussian. Two model parameters were generated for the fitted function using least-squares regression. The model function was a standard Gaussian:

$$\widehat{d'} = Ae^{\frac{-((x-x_0)^2+(y-y_0)^2)}{2\sigma^2}} \quad (3.10)$$

A and σ were the two fitted parameters. A represented the peak amplitude of d' in the fitted function and σ represents the width parameter of the bell curve in units of mm. The x and y variables represent the spatial location of each individual electrode while x_0 and y_0 represent the central peak of the Gaussian. Regressions were computed in two different manners. When just examining the spatial extent of the cortical modulation, x_0 and y_0 were fixed to be located directly under the control electrodes. For signals where push-pull control was evident, a summation of two Gaussians with one positive and one negative amplitude was used for fitting positive and negative d' values. Figure 3.18 shows two example plots of the Gaussian fits for the spatial distribution of d' from Figure 3.17. In Figure 3.18a, two Gaussians are used to fit the push-pull with positive and negative d's occurring under the two control electrode. In contrast, Figure 3.18b is an example where only a single Gaussian fit the data centered at the single electrode being used for a large majority of the control.

54
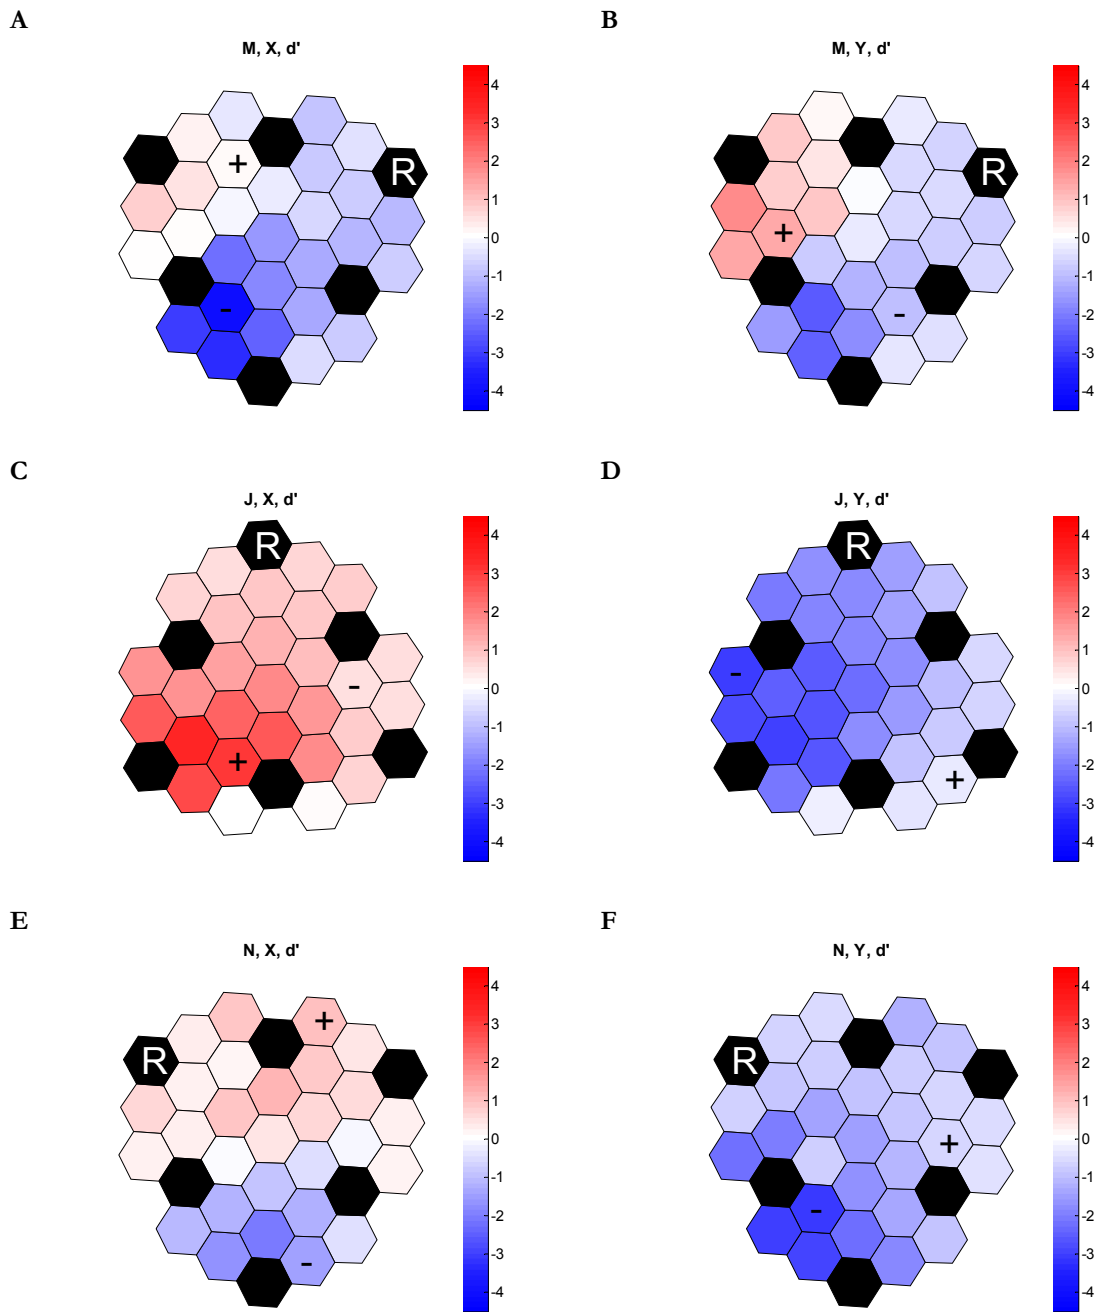
Figure 3.17. *d'* **Spatial Distribution**

Spatial distribution of d' across the 28 recorded channels for the six monkey-dimension combinations. + and − indicate the two channels used for push-pull control for that week's recordings. R represents the reference electrode while the other black locations represent the unused reference electrodes.

A

M, Y

B

J, X
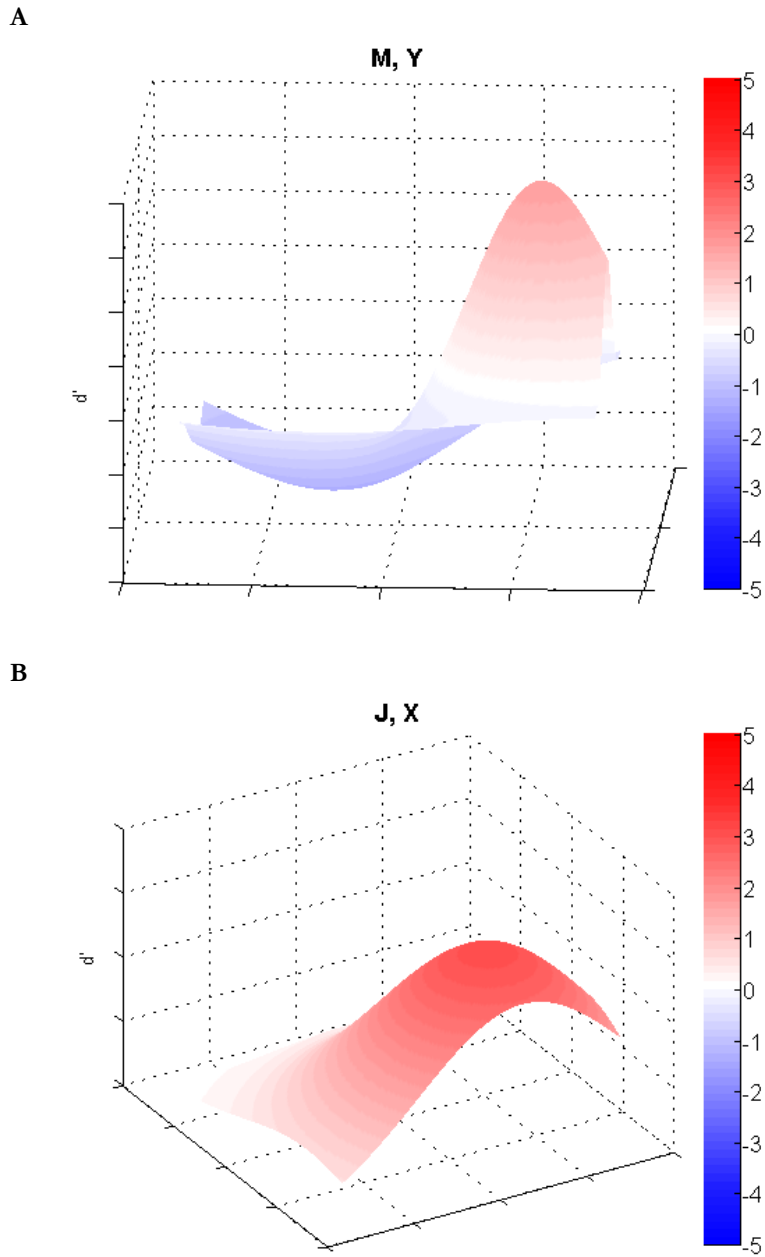
Figure 3.18. Example Gaussian Fits.
Example two-dimensional Gaussian fits across the d' values on 28
recorded channels. In a) the sum of two Gaussians is used for the push-
pull control of M, Y while in b) a single Gaussian centered under a single
control electrode is used for J, X.

A summary of the fit parameters of amplitude and bell curve width is presented in Table 3.1. Interestingly, it seems that monkey M tended to have narrower distributions with a σ value around 3mm for the two larger peaks. In contrast, for monkeys J and N, the σ value seemed to be in the range of 6-9mm. It is unknown how much of this difference represents recording differences such as electrode, impedance, implantation location, and tissue reaction differences as compared to possible cortical architecture and area differences. Nevertheless, these differences highlight the variability that can be observed with ECoG recordings.

Table 3.1. Spatial fit of d' using a two-dimensional Gaussian function with amplitude A and width parameter σ.

	+ channel		- channel		(all p<0.001)
	A (d')	σ (mm)	A (d')	σ (mm)	R^2
M,X	0.31	2.35	-4.14	3.04	0.71
M, Y	2.21	2.68	-1.42	6.37	0.56
J, X	3.06	6.00			0.86
J, Y			-2.99	8.86	0.90
N, X	1.13	8.62	-1.93	5.51	0.80
N, Y			-2.48	6.06	0.61

3.11 Frequency Analysis

In addition to looking at the amplitude between 75-105 Hz, the amplitude modulation across all frequencies between 0-300 Hz was examined. Once again, the power spectrum was calculated for each trial using the multitaper method as previously described section 3.8. Instead of calculating the RMS amplitude across a wider frequency band, the

RMS value for each trial was calculated for each 1 Hz step. Equation 3.7 was then used to calculate the d' as a function of frequency for each 1 Hz bin.

Figure 3.19 plots *d'* as a function of frequency for all 12 of the different control electrodes for the six monkey-dimension combinations. The frequency range of 75-105 Hz that was actually used for control is highlighted in red. Interestingly, there appears to be a fairly conserved frequency pattern where the modulation between 75-105 Hz was also accompanied by an amplitude modulation in the opposite direction at lower frequencies centered around 25 Hz. Additionally, lower frequency modulation <12 Hz tended to show amplitude modulation in the same direction as the control band.

In order to characterize the power spectrum modulation, a heuristic algorithm was developed to identify the various frequency components based on the zero crossings that occurred in the plot of d' as a function of frequency. For the range between 0-70 Hz, two breakpoints were determined that separated three different frequency components. The two breakpoints were chosen such that they maximized the total area under the curve of the d' values with the middle frequency component being of opposite sign to the low-frequency and high-frequency component. These frequency components will be referred to as the mu, beta and gamma bands. If there was only one zero crossing, then it is assumed that there was only a beta and gamma component. Likewise, if there were no zero crossings, then the mu and gamma bands were identified by choosing the minimum d' value as a breakpoint between the two bands. Additionally, the peak d' occurring within the identified gamma band was used to identify the frequency with maximum separation between signals for the

58

two targets. This peak was identified from a low-pass filtered signal (5 Hz boxcar) to minimize the effect of outliers at single frequency points.

The cutoffs and gamma peak frequency for the channels used primarily for control is shown in Table 3.2. Not surprisingly, the peak gamma frequency occurred within the 75-105 Hz range in seven out of eight cases. Perhaps more surprisingly, all of these seven peaks occurred within the range of 80-85 Hz. Additionally, the cutoffs between a low frequency mu band, a middle frequency beta band, and higher frequency gamma band appears to be surprisingly conserved. It appears that all mu bands were less than 12 Hz while the gamma frequency included all frequencies greater than 40 Hz. This conservation of frequency ranges is interesting considering the fact that only the 75-105 Hz frequency range was used for control and there was no incentive for the monkeys to modulate these signals within the other frequency ranges.

Figure 3.19. *d'* Across Frequencies.
D' plotted as a function of frequency. The red area represents the 75-105 Hz band used for control. a)
M,Y,+ b) M,X,- c) M,Y,+ d) M,Y,- e) J,X,+ f) J,X,-

J, Y, + channel

J, Y, − channel

N, X, + channel

N, X, − channel

N, Y, + channel

N, Y, − channel

Figure 3.19 cont. *d'* Across Frequencies.
D' plotted as a function of frequency. The red area represents the 75-105 Hz band used for control. g) J,Y,+ h) J,Y,- i) N,X,+ j) N,X,- k) N,Y,+ l) N,Y,-

61

Table 3.2. Cutoff frequencies between the mu, beta, and gamma bands where the d' zero crossings occur. Also, the peak d' frequency in the gamma frequency band is given.

	+ channel			- channel		
	Mu-Beta cutoff	Beta-Gamma cutoff	Peak Gamma	Mu-Beta cutoff	Beta-Gamma cutoff	Peak Gamma
M,X				13 Hz	36 Hz	84 Hz
M, Y	23 Hz	30 Hz	81 Hz	12	40	85
J, X	21	26	86			
J, Y				16	28	82
N, X	14	40	85	22	24	125
N, Y				16	28	82

4 Learning and Adaptation

When first presented with the brain control task the subject had no insight into how to move the cursor in the correct direction. Two electrodes that were located at opposite sides of the electrode array and were 15 mm apart were chosen to control the velocity of the cursor in a push-pull decoding scheme. The electrodes were screened through baseline recording to make sure that they were within the normal range of RMS and impedance values to avoid using outlier or defective channels, but no other pre-task screening was performed. Only through visual feedback could the monkey learn and adapt to successfully perform the task.

On the first day of closed-loop brain control, the cursor moved randomly based on the modulating amplitude of the ECoG recordings. In addition to these random movements, a biased velocity component was added to the cursor's position. This computer controlled bias moved the cursor toward the correct target and allowed the monkey to achieve a performance level greater than 50% correct. This bias aided in keeping the monkey motivated to continue attempting to complete the task and also provided him with a suitable reward level. At all times, the brain signals were moving the cursor in addition to the bias and the subject could always improve the number of correct targets hit and the speed at which the trial was completed by correctly modulating the brain signals. The bias level was manually adjusted by the experimenter and was adjusted to maintain actual performance levels at approximately 75% correct trials or above.

It should be noted that monkey M (the first animal tested in this project) is a special case in that the decoding scheme was originally simpler during the initial training period. Because of a difference in the overall RMS amplitude of the positive and negative control channels, it was possible for the monkey to modulate both of the control channels to larger or smaller amplitudes but the absolute changes were larger for the electrode with the larger overall signal. Since we desired to use more of a push-pull decoding scheme, the decision was made to normalize the two control signals by the overall range of each individual channel during the closed-loop experiments. This was accomplished by dividing the control signal by the RMS of the control signal from the previous 100 seconds (see Eq. (3.3)). This modification caused monkey M to focus more on modulating individual channels and was subsequently used in the rest of the experiment for all monkeys and dimensions after the first eight days of the x dimension for monkey M.

4.1 Initial Learning Time Course

In order to successfully complete the task, the subject must correctly modulate the amplitude of the signal on one or both of the designated electrodes to move the cursor towards the desired target. In order to monitor the amount of learning going on, we can examine how the d' for the control channels changes as the subject learns to accomplish the task. The d' for each electrode represents the separation of the amplitude of the signal between 75-105 Hz between the two different targets for a series of trials. All of the reported d's use the amplitude of the positive target (right or up targets) minus the negative target (left or down targets) divided by the pooled standard deviation of the two targets. A d'

of zero represents no difference in the mean amplitude when the two different targets appeared while a positive value indicates that the amplitude was larger for trials where the positive target appears while a negative value indicates that the amplitude was larger for trials where the negative target appeared. Therefore, to successfully complete the task, either the positive channel's *d'* must be positive or the negative channel's *d'* must be negative. Additionally, meeting both of these conditions is needed for a true push-pull control scheme.

Figure 4.1 shows the calculated *d'* from a sliding window of 100 trials in 10 trial steps across the series of days that a monkey learned to use a new decoding scheme. The six plots are for each of the three monkeys when they first learned control for both the *x* and *y* dimension. For all of the monkey-dimension combinations there is a clear divergence for at least one of the channels away from a *d'* of zero. In most cases, as the monkey became satiated for the day, the *d'* for both channels falls back to zero by the end of the day. If we examine the peak *d'* for each day, it becomes apparent that there were various rates at which the monkey learned to perform the task. Three of the combinations, J-Y, N-X, and N-Y, appear to increase the performance over a series of days with a new, larger magnitude *d'* occurring each day until reaching a plateau. Additionally, the special case of M-X appears to follow a similar time course of gradual improvement over a few days following the grayed out portion where the control signals were not normalized.

A

B

Figure 4.1. Initial Learning Time Course.
The *d'* values for the positive and negative control channels separated by 15 mm for the initial series of recording days when switched to a new dimension of control. a) Monkey M, horizontal control (Note, control signals were only normalized after the first eight recording days), b) Monkey M, vertical control

66

C

D

Figure 4.1. cont. Initial Learning Time Course.
The *d'* values for the positive and negative control channels separated by 15 mm for the initial series
of recording days when switched to a new dimension of control. c) Monkey J, horizontal control,
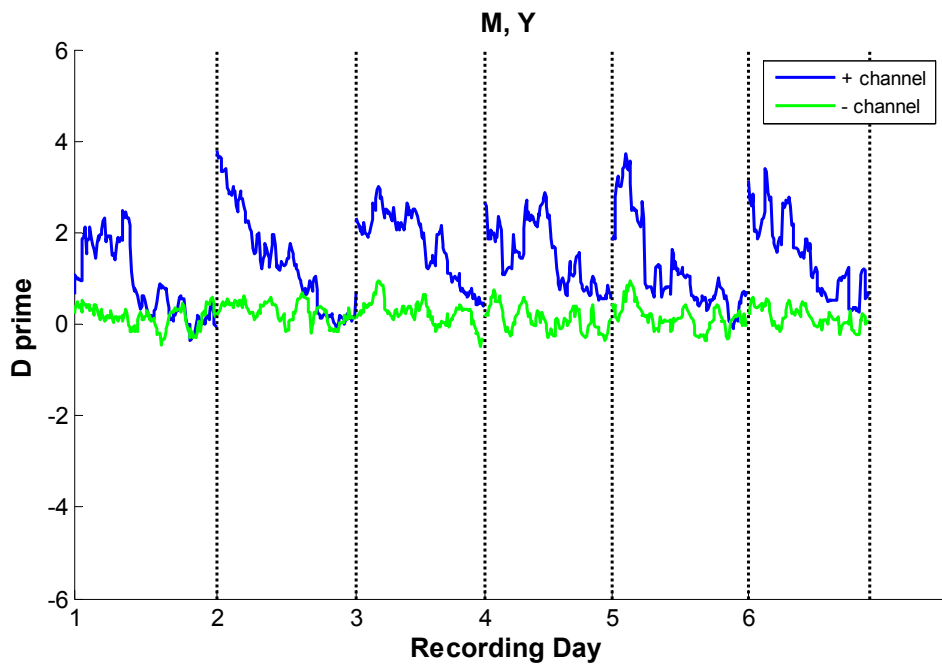**d) Monkey J, vertical control**

E

F

**Figure 4.1. cont. Initial Learning Time Course.**
The *d'* values for the positive and negative control channels separated by 15 mm for the initial series
of recording days when switched to a new dimension of control. e) Monkey N, horizontal control,
f) Monkey N, vertical control

In contrast, two cases appeared to more quickly learn to accomplish the task. Both M-Y and J-X showed significant adaptation within the first day of performing the task. Figure 4.2 enlarges Figure 4.1 to show only the first day from these two monkeys. In both cases, there was a dramatic rise in the d' for the positive channel within the first 300 trials. Both subjects showed peak d' separations greater than two standard deviations on the first day of training.

To help quantify the learning rates for all three monkeys, it is useful to look at the difference between the positive and negative channels' d' values or $\Delta d'$. A larger separation between the two channels in the push-pull decoding scheme should directly lead to increased performance in terms of percentage correct. For example, if the two channels are independent normal distributions with equal and opposite modulations, $\Delta d'$ values of 1, 2, and 3 would correspond to 64, 76, and 86 percent correct, respectively. Using the $\Delta d'$ values calculated from the d' values shown in Figure 4.1, the first time point where a new, larger $\Delta d'$ was achieved by a given monkey was calculated. Figure 4.3 traces out these time points for rising values of $\Delta d'$. All subjects showed a clear rise in performance while learning the task until reaching a plateau at $\Delta d'$ values of at least 2.5. As seen previously, monkey-dimensions M-Y and J-X showed a dramatic rise where each monkey was performing at a level above a $\Delta d'$ of 3.0 by only the second day. In contrast, the other four showed a more gradual rise with the slowest rise taking approximately 10 days to reach a $\Delta d'$ of 2.5.

A

M, Y

B

J, X

Figure 4.2. One Day Learning Time Course.
The *d'* values for the first day for a) Monkey M, vertical control and b) Monkey J, horizontal control to highlight the neural adaptation on the first day of control.

70
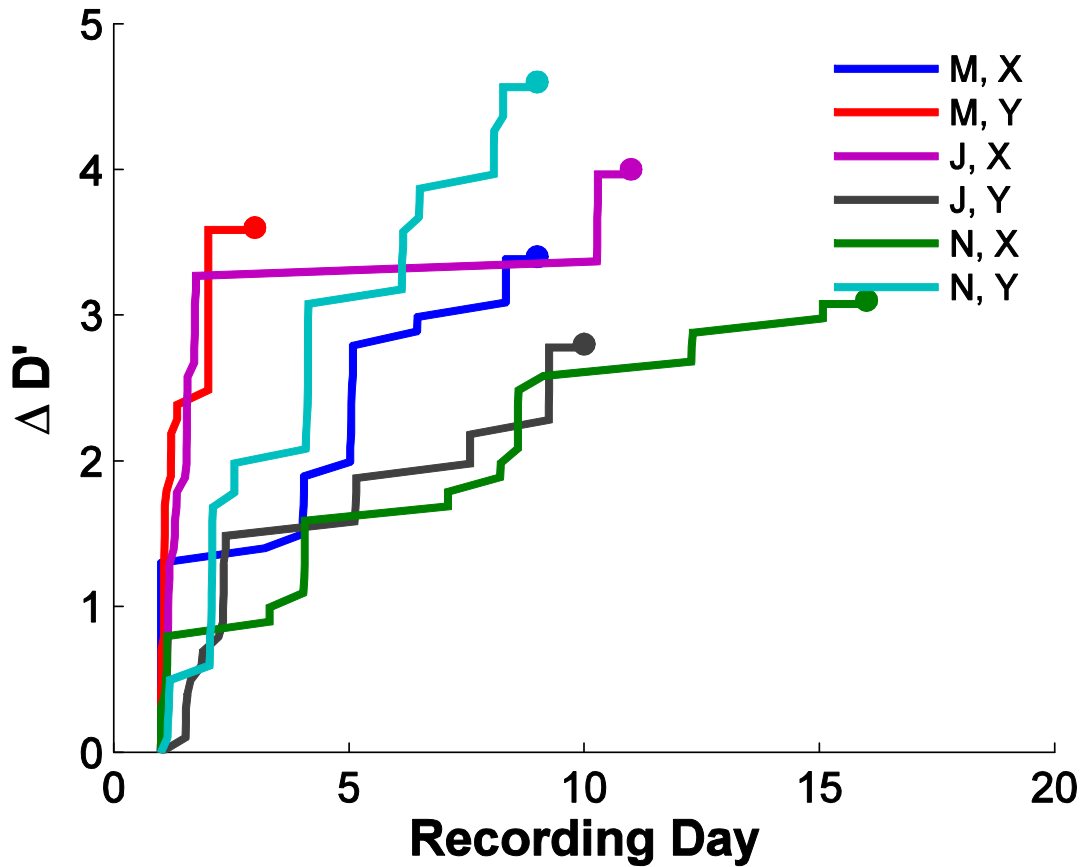
Figure 4.3. **Δ***d'* **Time Course**
The increase in the difference between positive and negative control channel *d'* (Δ*d'*) across recording days as the subject adapts to the task. Each plotted point represents the first time point that a monkey for a given dimension reached a given Δ*d'* value.

Based on these learning rates, two clear facts emerge. First, in all six monkey-dimension cases with arbitrary and predefined electrode configurations, the monkey learned how to successfully control the cursor. No screening, signal processing adaptation, or other form of experimenter intervention was necessary for gaining successful control of a one-dimensional closed-loop BCI. Additionally, learning to accomplish this task occurred for all monkeys in a relatively short time period over a range from one to two days up to a couple weeks.

4.2 Frequency Adaptation

In addition to the modulation within the control band between 75-105 Hz on the designated electrodes, it is also interesting to examine the observed modulation occurring outside the frequency range being used for control as well as the signal changes occurring on the other 26 channels not being used for control. These signals have no direct effect on the movement of the cursor but can provide insight into how the task is being performed.

Figure 4.4 shows a plot of d' as a function of frequency for each of the individual channels that was actively modulated for the five monkey-dimension combinations that showed initial push-only control. In all cases, the beta (25-30Hz) band showed opposite modulation to the gamma (>40 Hz) and mu (<12 Hz) bands. This trend emerged early during the training period and did not deviate significantly throughout the period of learning. Additionally the peak gamma band modulation frequency for each day is also highlighted with a horizontal black bar. Once again, a stable peak frequency of modulation is observed across the series of days for all five combinations. Figure 4.5 is a similar d' versus frequency plot but is for the two control channels for monkey N in the x dimension. In contrast to the other series of data, there is more variability in the frequency cutoffs and peak frequencies from day to day.
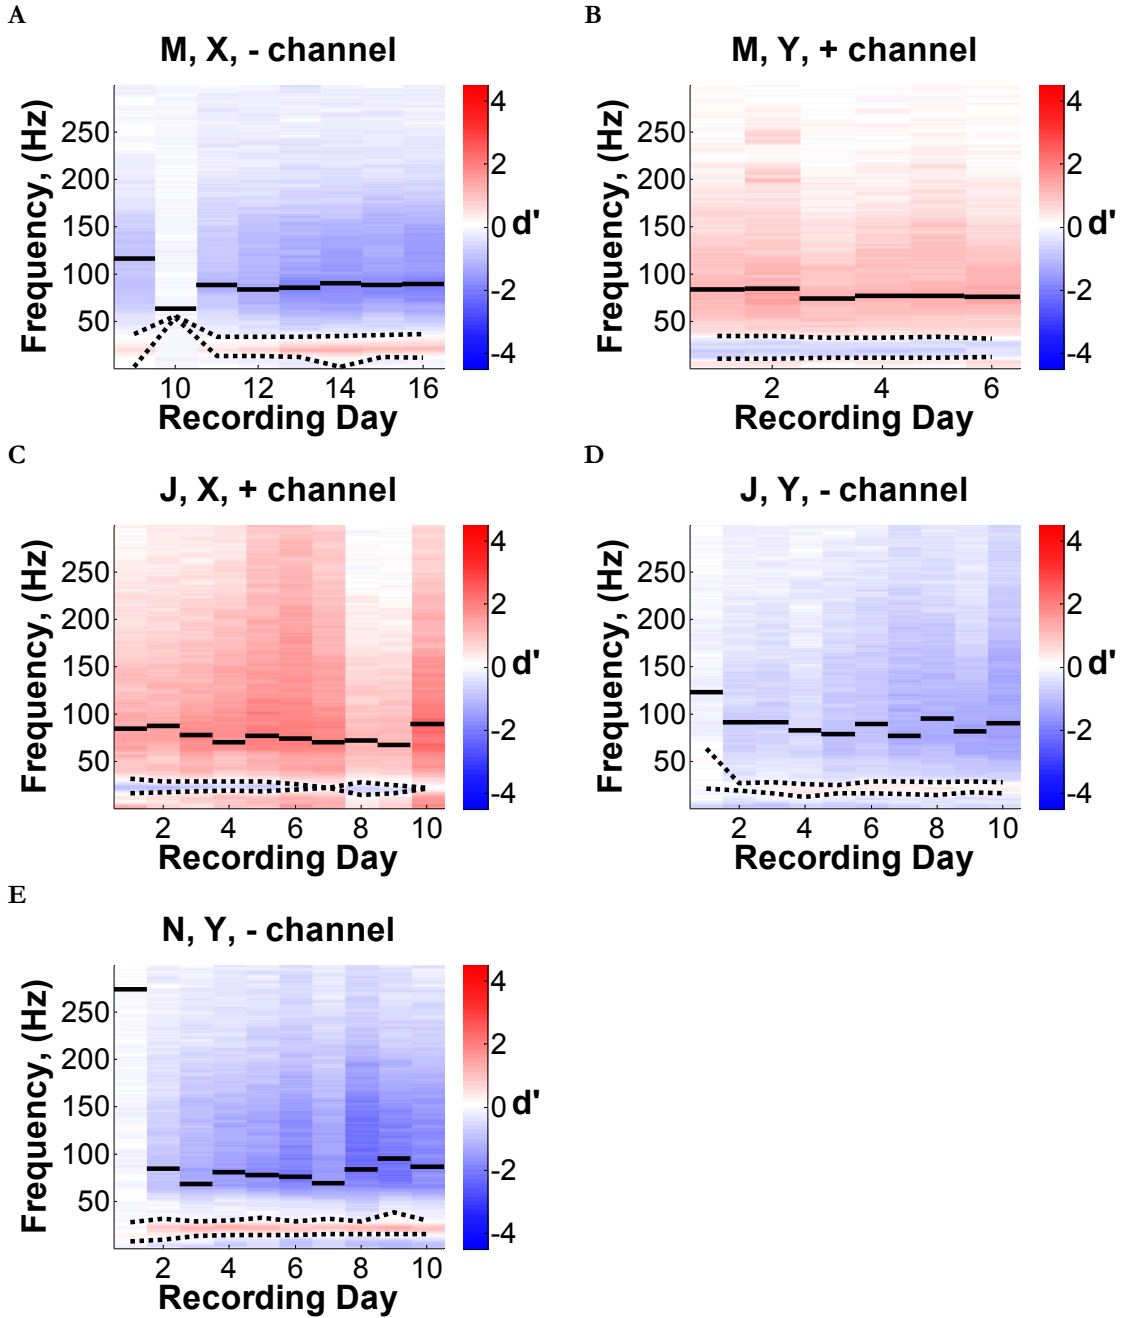
Figure 4.4. Initial Frequency Time Course.
The adaptation in *d'* across frequency plotted across the initial days of control. The horizontal bars represent the peak *d'* within the gamma band. The dashed lines represent the separation between the mu, beta, and gamma bands. Plots are for the following monkey, dimension, and channel:
a) M, X, -   b) M, Y, +   c) J, X, +   d) J, Y, -   e) N, Y, -

A

N, X, + channel

B

N, X, - channel

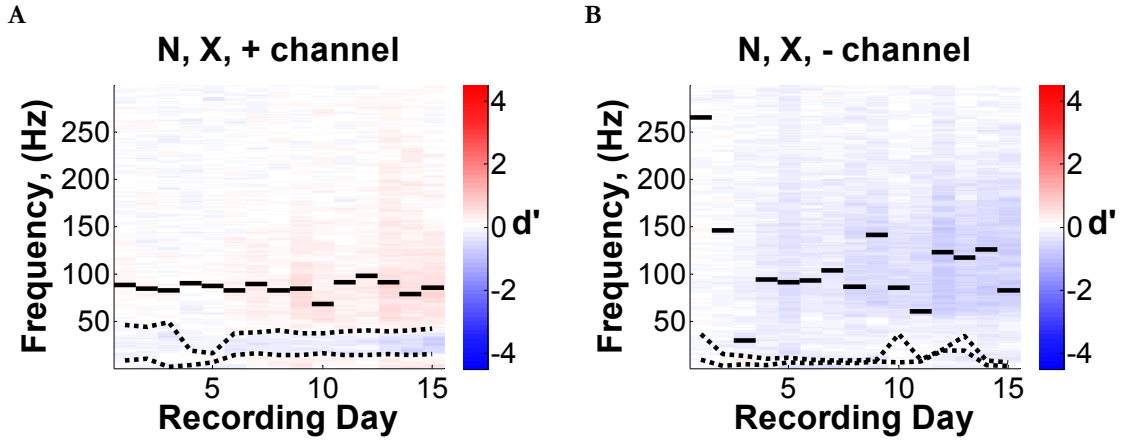**Figure 4.5. Initial Frequency Time Course, Part 2.**
The adaptation in *d'* across frequency plotted across the initial days of control for the one true push-pull control scheme (monkey N in the *x*-dimension). Plots are a) Positive channel and b) Negative channel. The horizontal bars represent the peak d' within the gamma band. The dashed lines represent the separation between the mu, beta, and gamma bands.

5 Interelectrode Distance

A key issue for further development of an ECoG BCI is the number of features that will be available for independent control of multiple degrees of freedom. The feature space when examining ECoG recordings is typically comprised of both a spatial and spectral dimension. Theoretically, either of these feature dimensions or a combination of both dimensions could be used for control of multiple degrees of freedom. From the same recording site, one frequency range could be assigned to one dimension while other frequencies could control another dimension. Conversely, the same frequency from multiple different recording sites could be used for independent control of two degrees of freedom. The experiments in this chapter were designed to determine the independence of gamma band activity from two separate channels along a spatial dimension (when using the same frequency range on both channels).

There is still limited knowledge about the minimum spacing needed between ECoG electrodes for independent control for BCI applications. Historically, ECoG has been used in the clinical setting for seizure detection and localization. The design of the electrode spacing was to maximize coverage of a given region of the brain to yield the best chance for observing seizure activity. Electrodes are often designed with interelectrode spacings on the order of approximately 1cm and diameters around 3mm. It is possible and perhaps likely that electrodes with a finer spatial resolution could yield more information per unit of cortical area for purposes such as a motor BCI.

To assess the spatial resolution of ECoG recordings, our lab built an in-house custom ECoG recording set-up with microwire electrodes with 300 um diameters and 3 mm interelectrode distance. Before the start of the closed-loop BCI experiments, baseline recordings were examined to quantify the spatial properties of the electrodes. To analyze these baseline recordings, the power spectral density was calculated using the multitaper method as previously described in section 3.8. Spectra were calculated for one second non-overlapping sliding windows of the five minutes of baseline recordings from the first day before the monkey performed any BCI control. The correlation of the power spectral values between pairs of channels across time at each given frequency was calculated. The mean correlation of the power values between 75-105 Hz was then determined for all of the pairs of electrodes that were a certain distance apart ranging from 3 mm to 16 mm.

Figure 5.1 shows the signal correlation between electrodes as a function of distance. This analysis uses all 28 channels and thus looks at channels that were eventually used for control as well as those that were not. There is a clear trend showing that electrodes separated by shorter distances are clearly more correlated within the frequency range of 75-105 Hz. There appears to be a steeper change between 3 to 9 mm while the change in correlation between 9 and 15 mm is not nearly as great.

Figure 5.1. Distance Correlation.
Correlation of amplitude between 75-105Hz as a function of distance for baseline recordings for all 28 channels before any closed-loop BCI tasks were completed.

Electrode distances were examined in a closed-loop task for the three different distances of 15mm, 9mm, and 3mm. These different distances allowed the experiments to be conducted by changing the control electrodes to the next adjacent electrode each week for the new experimental distance. Choosing this set-up allowed the experiments to be conducted in successive weeks since each new distance did not require a new cortical area to be trained but rather allowed the subject to learn to use the adjacent area more quickly since it was likely already at least partially activated. The arrangement of electrodes for the three different distances is shown in Figure 3.7b.

Each monkey performed the closed-loop 1D task in two different dimensions. First, the experiment was conducted in the horizontal dimension with push-pull control electrodes located at distances of 15, 9, and 3mm. The monkey performed the task at each distance for one week and then moved to the next distance. This was repeated for the same three distances on the same electrodes for a total of six weeks of horizontal control of the cursor. Next, a different pair of control electrodes separated by 15 mm was assigned to control the vertical velocity of the cursor. After initial training, the subject again spent one week at 15, 9, and 3 mm. The experimental distances were once again repeated for a total of 6 weeks on the vertical condition. Thus, the entire distance experiments were conducted over a total of 12 weeks for each monkey.

5.1 Performance

As a measure of performance, the two main metrics used were percentage correct and bit rate. Figure 5.2 shows the percentage correct for all six monkey-dimension combinations. Two out of the six series of experiments exhibited trends that appeared to be significantly different from the other four series. The first combination of monkey M with the *x*-dimension of control appears to have a significant learning trend where percentage correct appears to improve with each week of performing the task. Additionally, the monkey was able to perform at a high level on all three control distances which leads to the possibility of ceiling effects that make it more difficult to discriminate differences between

78

different control distances.   The sixth combination (N, y-dimension) shows a significant

drop in performance for the second set of the three distances compared to the first.   This

corresponded to a large observed drop in the recorded signals RMS amplitude and poorer

quality recordings overall.   This was likely due to physical changes in the implanted

electrodes that led to the signal degradation.   Therefore, the N-Y combination was not

included in further analysis.

For the other four combinations of monkeys and dimensions, there is no dramatic

trend in performance either up or down across the entire six weeks.   There does appear to

be some weeks where the subject did improve performance within the week as he

presumably became more accustomed to the new control electrodes.   The expected effect

of decreased performance with closer electrode distances is observable with the percentage

correct dropping in weeks where the 3 mm spacing between positive and negative electrodes
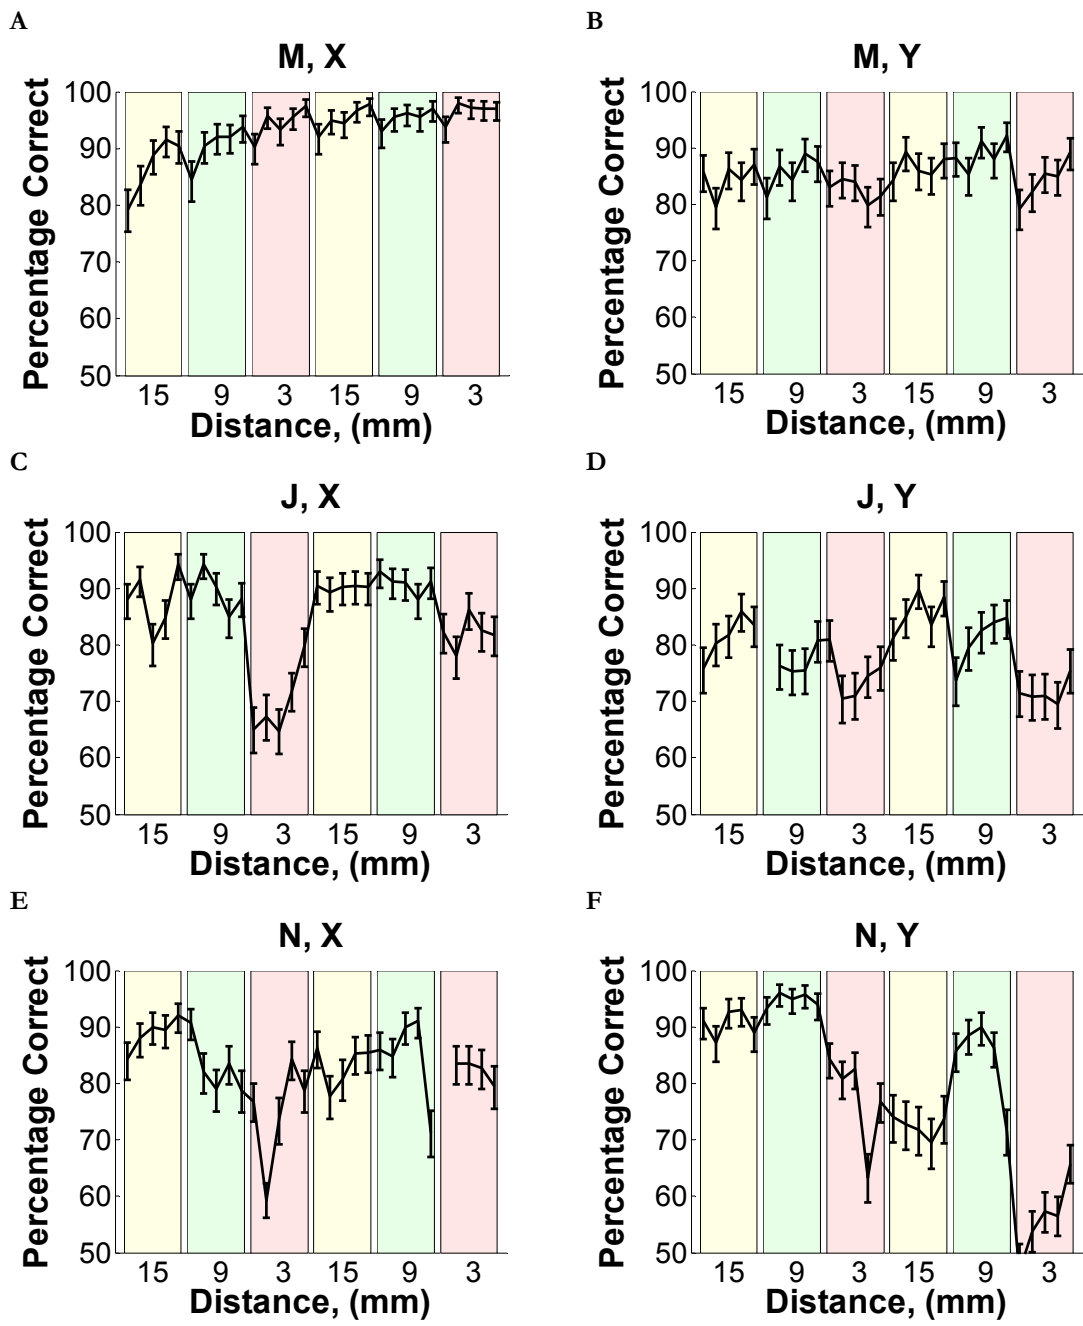
Figure 5.2. Individual Daily Percentages.
Percentage correct for each day throughout the 6 week distance experiment. Percentages represent the best 400 trial block during the recording day. The monkey, dimension combinations are: a) M, X b) M, Y c) J, X d) J, Y e) N, X f) N, Y

Figure 5.3. Individual Percentages by Distance.
Combined percentage correct using multiple comparison test for each individual monkey dimension combination across the three different distances. Percentages represent the best 400 trial block during the recording day. Error bars represent the 95% confidence intervals.

To better test the differences between the three distances, an ANOVA was performed for each monkey-dimension using the ten days of percentages for each of the three control distances. Since the data values used were percentages, an arcsine transformation was used to eliminate the dependence of variance as a function of the mean (Sokal & Rohlf, 1995). Figure 5.3 shows the 95% confidence intervals for each of the three control distances for the five experimental combinations. Two out of the five combinations (both cases of monkey J) showed a significant decrease in percentage correct at 3 mm compared to the 15 and 9 mm control distances. In addition, two other combinations (M-Y and N-X) trended toward the 3 mm control being the poorest with the mean percentage

correct being the worst at the closest distance. Interestingly, the 3 mm spacing had the best

percentage for monkey M when he was controlling along the x-dimension although the

differences were not statistically significant. Based on the individual days plotted in Figure

5.2, it appears at least part of the explanation for this difference is that the overall

performance of the subject improves throughout the different distance phases of the study

and had not reached steady state during the training phase.

Finally, the data was combined in a two-way ANOVA (5 monkey-dimension

combinations **x** three distances) to find the overall means and significant differences

between the 3 different distances. In this combined analysis shown in Figure 5.4, the 15 and

9 mm distances had nearly identical performance around 87%. Moving to 3mm resulted in

degradation in performance to around 82%. While there is still a clear level of control at

3mm that is well above chance, it does appear that the ability to get two independent control

signals from separate spatial locations has been reduced at the 3 mm spacing.

Figure 5.4. Combined Percentages by Distance.
Combined percentage correct using multiple comparison test for each individual monkey dimension combination across the three different distances. Percentages represent the best 400 trial block during the recording day. Error bars represent the 95% confidence intervals.

5.2 Spatial Adaptation

The resulting data can also be examined for adaptive changes that the subject makes to each change in the electrodes being used for control. In order to more successfully complete the task, it is likely that the monkey will have to increase the modulation as well as potentially decrease the correlation between the positive and negative control channels. Thus, to examine this effect, we can examine the changes in electrode activity from a week

prior to closed-loop control to the week of closed-loop control. Also, the week after closed-loop control can be examined to track the de-adaptation or wash-out of the signal once a channel is no longer used for control.

Each day, recordings were taken from all 28 channels on the recording array. It is especially interesting to examine electrodes used for control at some point in the experiment even in the weeks where they were not actually used for control. For example, in Figure 5.5a, the negative y channel used during the 3mm spacing weeks for monkey M is plotted. The first week represents a 9 mm spacing week where an adjacent electrode to the one plotted was used for control. During this week, we see very little separation in the 75-105Hz amplitude between trials where the subject was attempting to move up compared to down trials. However, during the following week when the control was switched to use the electrode that is plotted, the signal separation between the two target conditions grows. In the last week plotted control was switched back to the farthest separation of 15mm and once again there appears to be a narrowing of the signal separation although not back to nearly complete lack of modulation seen in the first week before this electrode was ever used for control.

A similar effect is observed for the 15mm, negative y channel of monkey J in Figure 5.5b. In this case, the 15mm had been previously used for training as well as the first iteration of the series of control distance studies. The same overall trend is observed where an increase in signal separation is observed during the week when the channel is actually being used for control. To test whether this separation in means is not accompanied by an

84

equivalent increase in trial-by-trial variability, the d' metric can be examined. In both example cases (Figure 5.6), there is a decrease in the d' (the correct direction for negative channels) during the middle weeks when the channel was actually the control electrode.

Since a push-pull decoding scheme was being employed, in order for the observed adaptation on one channel to increase performance it must not be accompanied by an equivalent decrease in modulation on the other channel. The $\Delta d'$ metric measures the total modulation that is occurring by comparing the difference in modulation on a pair of electrodes. In Figure 5.7, the $\Delta d'$ increased for both the 3mm electrode pair (monkey M) and 15mm electrode pair (monkey J) on the week where these electrode pairs were used for control. To test the amount of significant adaptation that occurred for all five monkey-dimension combinations that were used in the study, a standard student's t-test was

A

M, Channel used for -Y, 3mm control

B

J, Channel used for -Y, 15mm control

Figure 5.5. Signal Amplitude Adaptation Examples.
Signal modulation for the two different targets for two example channels
for the week before, during, and after it was used for control. The
increased separation between the two signals during middle week
highlights the adaptation that occurred while the channel was actually
used for control. Error bars represent the standard deviation of the signal
RMS. a) The $-y$ channel at 3mm spacing for monkey **M.** b) The $-y$
channel at 15mm spacing for monkey **J.**

86

A

B

Figure 5.6. *d'* **Adaptation Examples.**
d' metric for two example channels for the week before, during, and after
it was used for control. The decreased *d'* (for negative channels) during
middle week highlights the adaptation that occurred to increase the
signal separation between the two target conditions. a) The −*y* channel at
3mm spacing for monkey M. b) The −*y* channel at 15mm spacing for
monkey J.

A

B

Figure 5.7. Δ*d'* Adaptation Examples.
Δ*d'* metric for two example channel pairs for the week before, during, and after they was used for control. The increased Δ*d'* (positive channel *d'* − negative channel *d'*) during middle week highlights the adaptation that occurred to increase the signal separation between the two target conditions. a) The *y* channel pair at 3mm spacing for monkey M. b) The *y* channel pair at 15mm spacing for monkey J.

performed between the $\Delta d'$ values for the five days of the week a pair of electrodes were used for control in comparison to the five days immediately before an electrode pair was used. Figure 5.8 shows the mean $\Delta d'$ values as well as 95% confidence intervals for groups of three weeks. The blue represents the week before a pair of electrodes was used for control, the green represents the week of actual control, and the red represents the week after control for the electrodes. The a-c panels group the control channels by the distance between electrodes. In all cases at least two out of five pairs of control electrodes showed a significant increase in the d' separation for the week when control was assigned to those electrodes compared to the week before. (Pie chart with black = $p<0.05$ and gray = $p<0.10$) This adaptation represents the subject correctly learning through feedback to increase the modulation on one or both electrodes compared to the previous week when another pair of electrodes was being used for control. In the week following control at least one of the subject-dimension combinations showed a significant decrease in $\Delta d'$ for each distance. This shows that, as control is shifted from a given electrode and feedback is no longer given, in some cases the cortical area is no longer reinforced and the level of modulation diminishes.
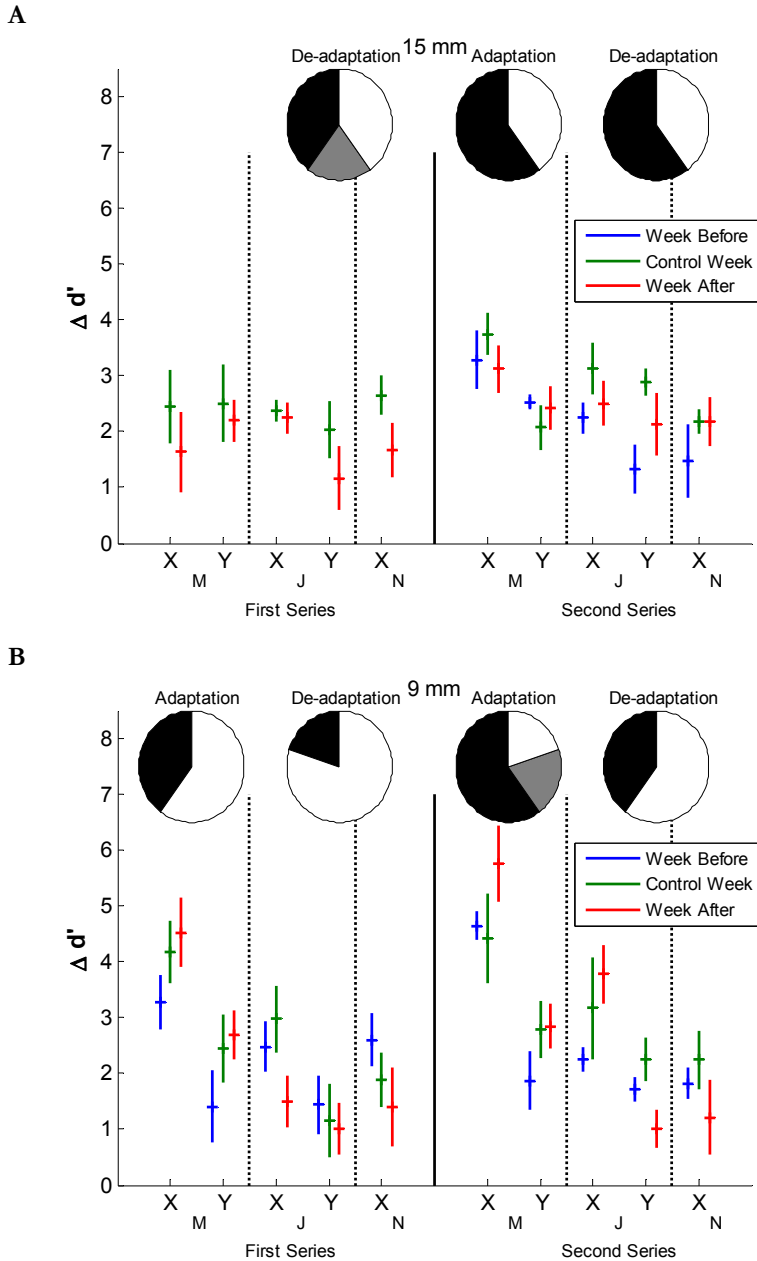
Figure 5.8. Δd' Adaptation Summary.
The mean Δd' (with 95% confidence intervals) for a given pair of control electrodes for each monkey-dimension combination. The three Δd' values represent the week before control (blue), week of control (green), and week after control (red). The pie charts represent what fraction of electrodes showed significant changes (p<0.05 black, p<0.1 grey) illustrating adaptation and de-adaptation between weeks. a) 15mm electrode pairs. b) 9mm electrode pairs.

Figure 5.9. cont. Δ*d'* Adaptation Summary.
The mean Δ*d'* (with 95% confidence intervals) for a given pair of control electrodes for each monkey-dimension combination. The three Δ*d'* values represent the week before control (blue), week of control (green), and week after control (red). The pie charts represent what fraction of electrodes showed significant changes (p<0.05 black, p<0.1 grey) illustrating adaptation and de-adaptation between weeks.   a) 15mm electrode pairs. b) 9mm electrode pairs.

6 Decoding Effects

The closed-loop experiments described here were all performed with fixed decoding parameters that were not changed throughout the experiments. Except for initial bias to facilitate learning and the changing of control channels to study the effect of interelectrode distance, all other parameters were held constant. Two channels were always band pass filtered between 75-105 Hz with a 32^{nd} order filter, rectified, and low pass filtered at 3 Hz to estimate the signal amplitude. These two signals were then normalized to the running average and RMS of the previous 100 seconds and then combined as a differential signal between the two channels for push-pull control of the velocity of the cursor. The gain of the control signal to cursor velocity transformation was also held fixed throughout the experiments. The equations from Chapter 3 are reprinted here:

$$b(t) = f_{75-105Hz}[a(t)] \qquad (6.1)$$

$$c(t) = f_{<3Hz}(|b(t)|) \qquad (6.2)$$

$$d(t) = \frac{c(t) - \overline{c(t)}}{\sqrt{\frac{1}{n}\sum[c(t) - \overline{c(t)}]^2}} \qquad (6.3)$$

$$\dot{x}(t) = g \cdot [d_+(t) - d_-(t)] \qquad (6.4)$$

The filters ($f_{75-105Hz}$ and $f_{<3Hz}$) and gain parameter (g) all represent experimenter-controlled parameters that could potentially be changed to improve performance. For these closed-loop BCI experiments, the parameters were held fixed to study the two primary aims of interest of subject adaptation and interelectrode distance effects. In reality, these various decoding parameters could be tuned to enhance future BCI control. The parameters could be adjusted generally in future closed-loop BCI studies or adaptive algorithms could be devised to find parameters that match the behavior of individual subjects. To analyze the effects of various decoding parameters, the BCI task environment was recreated by processing the recorded signals in software using the same processing algorithm as originally used on the DSP hardware.

6.1 Filter Parameters Simulation

The first simulation was run to examine the effectiveness of the experimenter-selected gain, g from Eq. (6.4), for the control signal to cursor velocity transformation. If the gain is too high, the cursor will move too quickly and touch a target after only a short time period of the neural signal has been observed and the number of targets correctly selected will be. Conversely, if the gain is too low, the average time to target selection will be unnecessarily long with little increase in performance and a decreased bit rate. During the actual experiment, the gain was originally selected by the experimenters and kept fixed at a level that appeared to give a reasonable movement time of approximately 2-3 seconds and a high percentage correct. The gain was chosen relatively conservatively to make sure there was a reasonable length of recorded signal for each trial and with the fact that with post-

analysis it is possible to increase the gain for analysis but impossible to determine what would have happened had the gain been reduced.

The gain simulation test was run by iteratively changing the gain (g) to change the threshold where the cursor first touched either the positive or negative target. For each of the different gains, the number of times the chosen target was the desired target and the simulated time to reach each target could be determined. From these results, the percentage correct, mean movement time, and bit rate were calculated. In Figure 6.1, the bit rate as a function of gain is shown. For all cases, it appears that the bit rate shows improvement as the gain decreases to incorporate a longer time period for each trial. There appears to be a plateau with each trace's peak bit rate occurring between 1-1.25x of the original gain that was used in the actual experiments. Because of this plateau, it appears likely that decreasing the gain any farther would not have improved performance as measured by bit rate. It also appears that increasing the speed of the cursor by increasing the gain would also have been accompanied with a decreased bit rate. The *post hoc* analysis is limited, however, by the fact we are unable to determine how much the monkeys might have been able to adapt to changes in gain. Therefore, it cannot be ruled out that they could have been able to improve reaction time and focus for faster cursor speeds if the gain would have been increased to encourage them to do so.
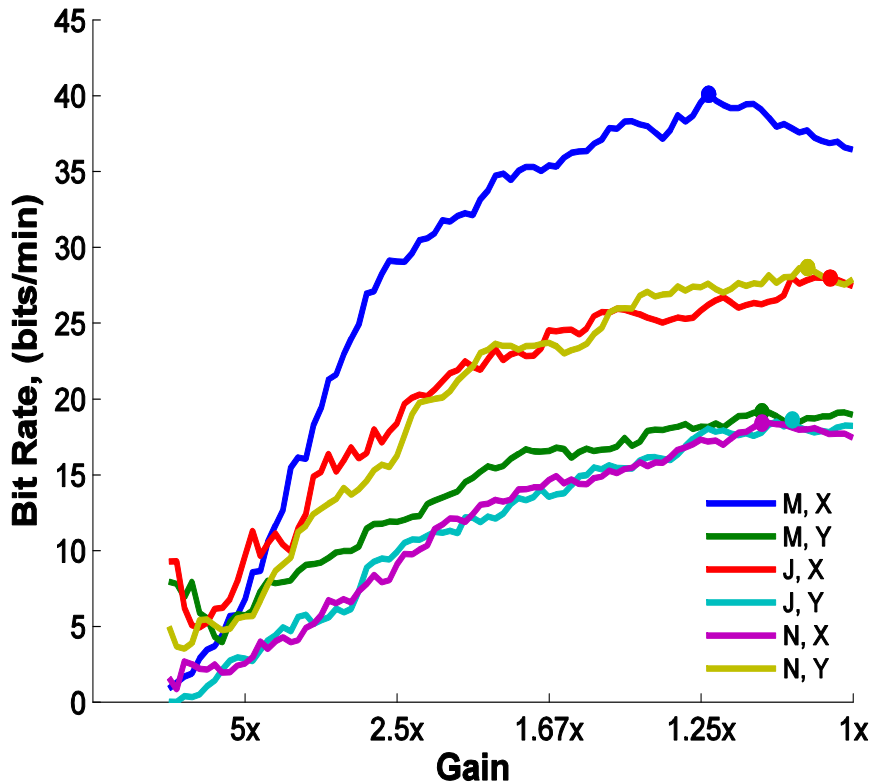
Figure 6.1. Gain Effect
The bit rate was calculated for a range of different gain levels by estimating the time and target that would be chosen by different control signal to cursor velocity gain values. The peak bit rate for each monkey is shown in the filled circles. For all cases, the peak bit rate was between 1-1.25x of the original gain used during the closed-loop experiments.

Next, two simulations were run to test the effect of various filtering parameters on performance. During the experiment a low pass filter of 3 Hz was used for smoothing the amplitude estimate ($f_{<3Hz}$ in Eq. (6.2)). This low pass filtering represents a temporal averaging of the amplitude estimate. Normal human reaching tends to show movement trajectories with frequency components up to approximately 5 Hz. Choosing the proper low pass cutoff requires an optimization between two extreme conditions. If the signal has too high of a cutoff, there is no useful information in these higher frequency ranges and the

95

cursor movements will have unnecessary noisy jitter which make precise movement and target selection difficult. Conversely, if the signal is filtered too much with a lower cutoff, the cursor response will be unnecessarily sluggish with slow response times.

A simulation was run by varying the low pass cutoffs between 1-5 Hz and measuring the predicted bit rate. For these analyses the gain was again allowed to vary since the exact gain for optimal performance can vary based on the other decoding parameters. Figure 6.2 shows the bit rate for the various low pass cutoffs used in the simulation. Generally, some slight improvements in bit rate were observed for higher cutoffs. Interestingly, it appears that there was some difference in the effect of filter cutoff for the different individual monkeys. For combinations M-Y and N-X, there appears to be little effect as a function of low-pass cutoff. In the other cases, it appears that the 1 Hz and also likely the 2 Hz cases would have resulted in slower responses and worse bit rates. In all cases, it appears that there was little useful amplitude modulation occurring at any frequencies above 3 Hz. The data is summarized in Figure 6.2 by plotting the peak bit rate (by selecting the gain level that provided the maximum bit rate) as a function of filter cutoff. The limited improvement about 3Hz may represent an inherent limitation of using amplitude modulation of ECoG signals for BCI control. However, there are several other possible reasons why this may be the case. First, the task is relatively simple and requires no change of direction within a given trial. The subject should always be trying to move in one given direction. Thus, the only truly desired modulation occurs at the start of each trial when the subject ramps up or down the signal amplitude when the target is first observed. Also, once again, there is the inherent bias of using 3 Hz cutoffs for the actual closed-loop experiment.

96

Another filter property that we chose to examine was the effect of changing the

bandpass filter orders ($f_{75\text{-}105Hz}$ in Eq. (6.1)). Butterworth filters are designed to be maximally

flat while having the steepest possible transition from passband to stopband for the given

filter order. Increasing filter order increases the steepness of this transition and how

selectively only the desired frequencies are passed while filtering out all other frequencies.

Ideally, we would use a perfect filter that would allow only the signal occurring between 75-

105Hz to be measured and used for control. In reality, because of the inherent time-

frequency uncertainty trade-off when performing signal processing, more precise frequency

resolution requires more temporal information and thus a delay in the translation from input

signal to amplitude estimate. Thus, higher order filters give more precise estimates of the

amplitude of the signal between 75-105 Hz but the estimate is more delayed which adds time

lags during a closed-loop BCI task.

For our actual experiment, we used a high order bandpass filter with a combination

of 16th order low pass and also 16th order high pass filter for a total order of 32. In reality,

each 16th order filter was made by cascading two 8th order filters together. This represents a

filter roll-off of 320 dB/decade. This bandpass filter had a group delay at 90Hz of

approximately 47 ms. Figure 6.4 shows the Bode magnitude plot as well as measured time

delays at 90 Hz for the 32nd order filter used for control as well as the three other lower

A

M, X

B

M, Y

C

J, X

D

J, Y

E

N, X

F

N, Y
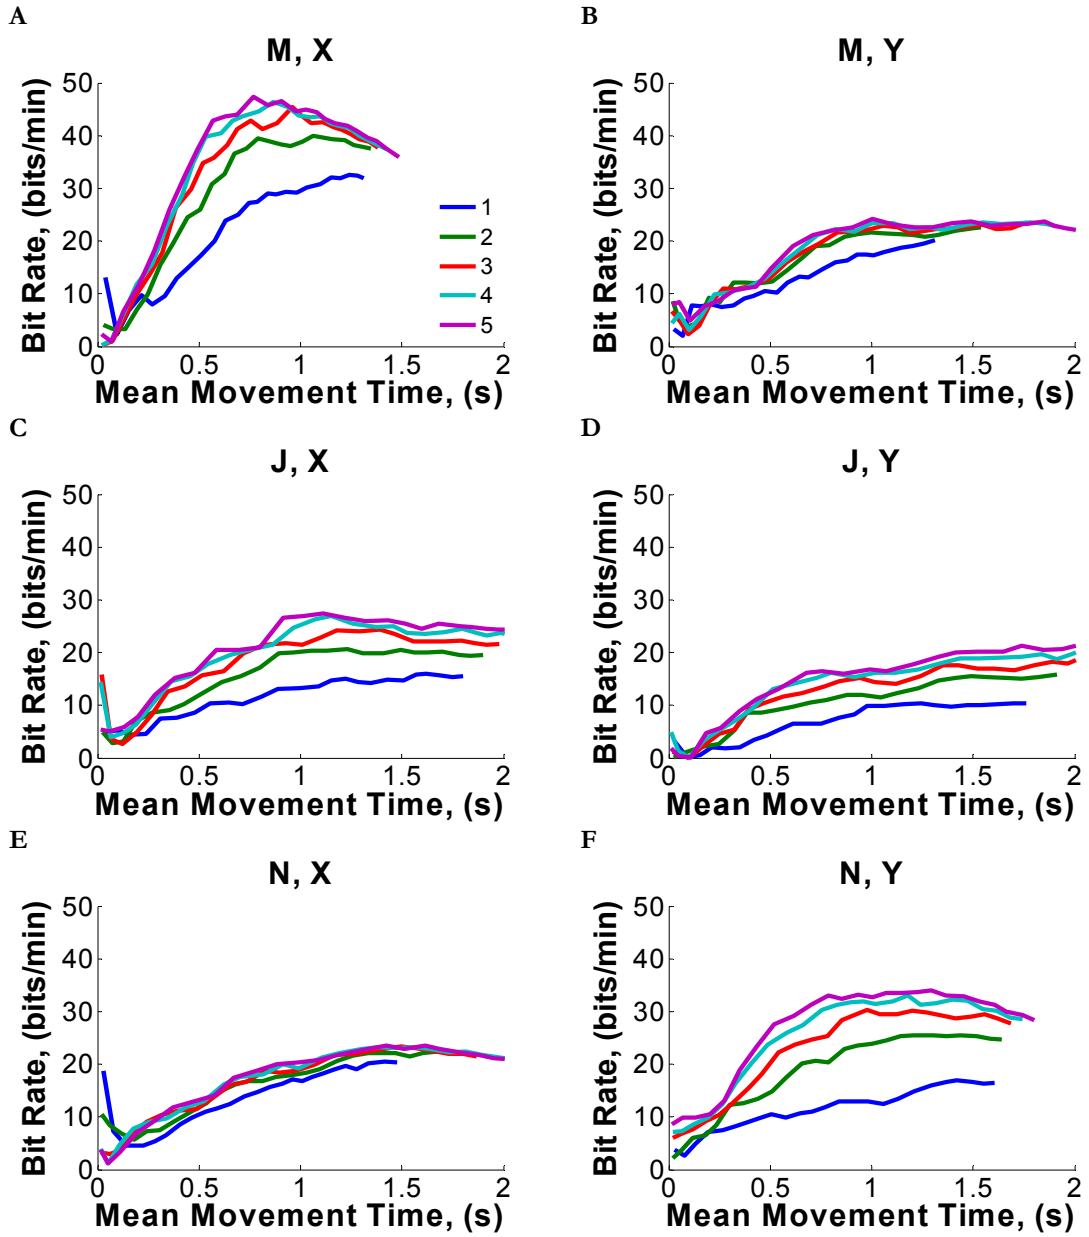
Figure 6.2. Low-pass Cutoff Simulation.
In the band-pass, rectify, low-pass filter envelope detection simulation, the low-pass filter cutoff was varied between 1-5 Hz. For each filter simulation the bit rate was calculated for varying gain levels from the resulting mean movement time and percentage correct.
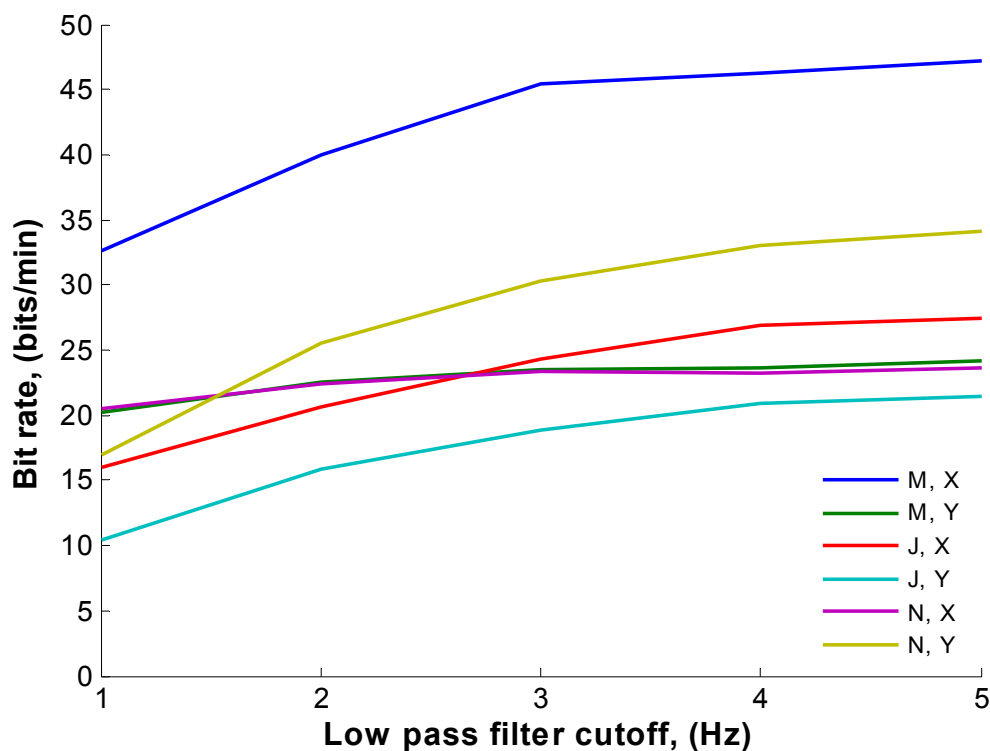
Figure 6.3. Peak Bit Rate for Low Pass Simulation.
The peak bit rate for the low pass cutoff simulation shown in Figure 6.2. In all cases, increasing the low-pass cutoff improved bit rate in all cases with diminishing improvement above the 3 Hz cutoff that was actually used for control.

order filters that were used in the simulation. For our analysis, we wanted to examine if

performance could be sped up with reduced delay by reducing the steepness of the filters

with lower orders. This comes at a cost of less precise frequency discrimination. Figure 6.5

and Figure 6.6 show the results of this simulation. There is once again some interesting

variation across monkeys. In general it appears that the 32^{nd} order filter was unnecessary

compared to the 16^{th} and 8^{th} order filters and in some cases resulted in slightly poorer

performance. Additionally, the 4^{th} order filters showed the most variation with it improving

performance for J,Y but also resulting in worse bit rates in a couple of cases.

In addition to changing the order of the band pass filter ($f_{75\text{-}105Hz}$ in Eq. (6.1)), a simulation was also run to examine the effect of using different frequency ranges for the filter. The simulation was run by varying both the center frequency as well as the bandwidth of the band pass filter. For the actual control, the passband was centered around 90 Hz with a bandwidth of 30 Hz to give the 75-105 Hz frequency band that was desired. Figure 6.7 shows a surface plot of the bit rate for the results of this

A

B

Figure 6.4. Bandpass Filter Response.
a) The frequency responses used for the four bandpass filters used in the simulation. The 32nd Order filter was the filter used during the actual experiments. b) The time delays observed for a 90 Hz signal for each of the of the bandpass filters.

simulation. These results show that there was very little difference in bit rate based on how wide the frequency band that was used for control as evidenced by the relatively constant levels when moving at any position along the horizontally axis.

100
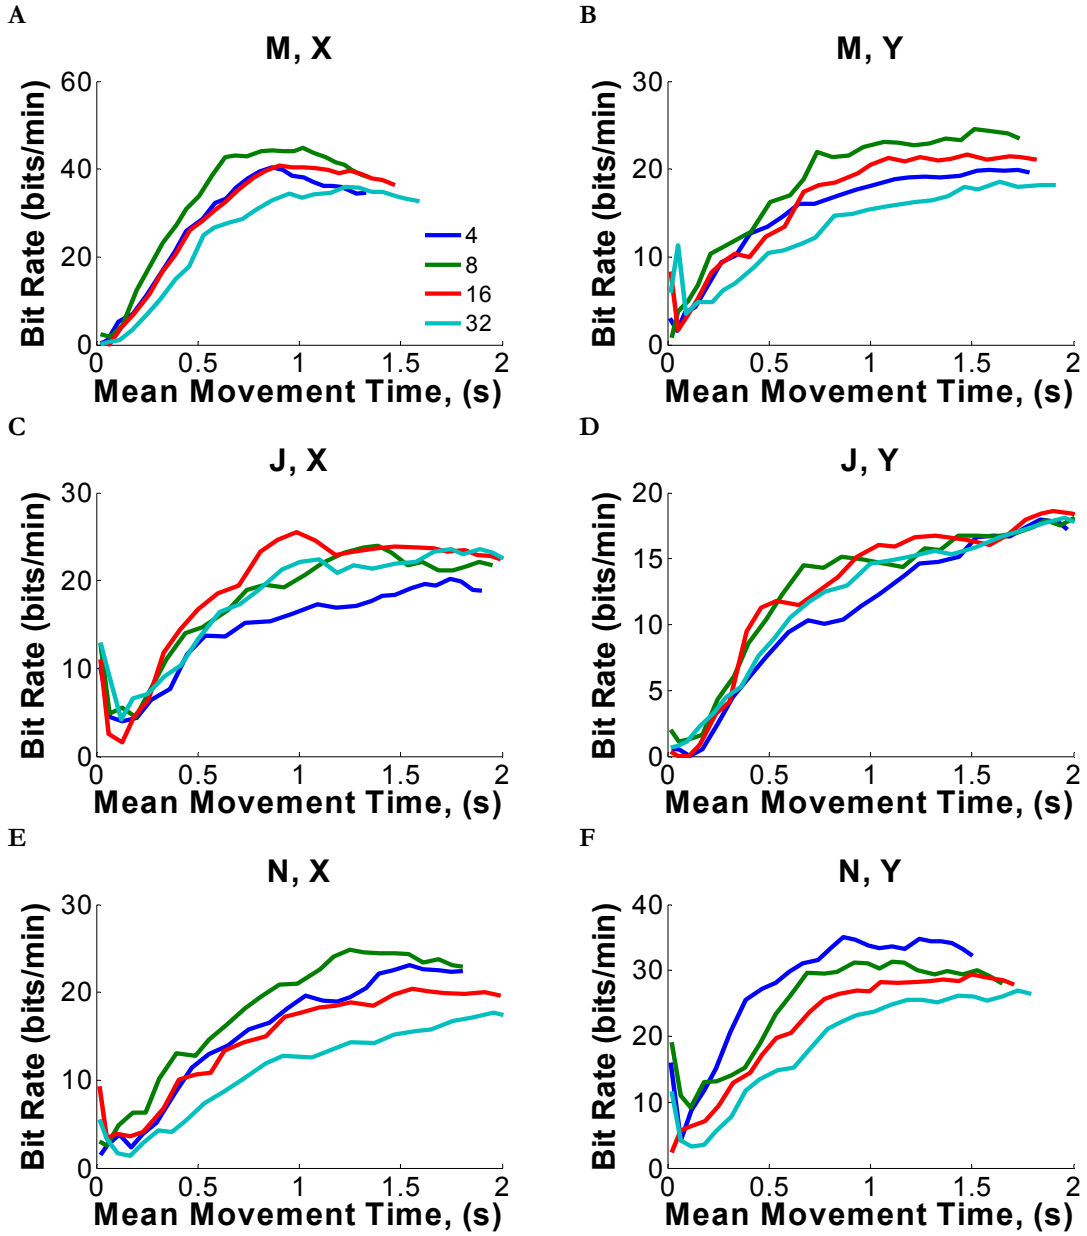
Figure 6.5. Bandpass Order Simulation.
Four different bandpass filters of varying order were used in the band-pass, rectify, low-pass filter envelope detection simulation.

Figure 6.6. Peak Bit Rate for Bandpass Order Simulation.
The peak bit rate for the bandpass order simulation shown in Figure 6.5. In general, the 32nd order filter used in the experiments appears to result in a decreased bit rate in all but one case compared to the 8th and 16th order filters.

When looking at different central frequencies, it appears that the high gamma range close to 90 Hz provided the best bit rate that was observed. Interestingly, there was a broad frequency range for all of the monkeys greater than 60 Hz all the way up to in many cases 200 Hz that gave a very similar high bit rate. Figure 6.8 further illustrates the general findings by plotting one of the two parameters separate while holding the other constant. Figure 6.8a looks at the effect of bandwidth when the center frequency was 90 Hz. There is very little difference depending on the bandwidth used with a possible slight improvement with a wider band. Figure 6.8b shows the simulations using a 30 Hz bandwidth but varying

Figure 6.7. Frequency Band Simulation.
Bit rate plotted as a function of frequency band used for control. The horizontal axis is the width of the pass band while the center frequency varies along the vertical axis. The bandpass filter of 75-105 Hz actually used for control is located at a bandwidth of 30 Hz and center frequency of 90 Hz. White circles indicate the point where the maximum bit rate occurred.

A

B

Figure 6.8. Frequency Band Simulation Summary.
Bit rates for the same simulation as shown in Figure 6.7 but with one
parameter fixed. a) Fixed center frequency of 90 Hz. b) Fixed
bandwidth of 30 Hz.

104

the center frequency. All of the peaks appear to be close to the 90 Hz actually used for control.

These results suggest that our decoding of the amplitude of the frequency range between 75-105 Hz was well matched to the neural activity the monkeys were modulating to complete the closed-loop task. This leads us to believe that either the 75-105 Hz high gamma range is a natural phenomenon that is easy for the monkeys to co-opt for BCI control and/or the monkeys are capable at adapting to the frequency range that was dictated. While it is our belief that the relative quickness that the subjects zeroed in their modulation to the high gamma range as shown in Chapter 4 suggests at least some underlying phenomenon, using only the 75-105 Hz band for closed-loop control makes it difficult to separate the two possibilities. Only with more closed-loop studies actually using different frequency ranges can the true suitability and flexibility of using one frequency range compared to another be fully examined.

This study does show that the modulation observed was happening over a relatively broad frequency range. The ability of our decoding filters to precisely identify the desired frequency and reject activity in neighboring frequencies did not appear to be critical to achieving a high level of BCI performance. This finding should allow one to build BCI algorithms with faster responses which will reduce closed-loop time delays. Additionally, the results suggest that the most important frequencies to reject for high gamma control are those less than 40 Hz in the mu and beta bands. This is especially true when one considers the dominant shape of the power spectrum of the signal is larger power at low frequencies

105

that falls off as frequency increases. This leads us to believe that asymmetric bandpass

filters that maintain the steeper roll-off on the low frequency side while reducing the roll-off

on the high frequency side could be used to increase BCI decoding performance.

6.2 Channel Information

The completed BCI experiments used only two channels for control and always used

a fixed push-pull decoding scheme. While this fixed weighting of two channels provided

good control and was a good starting point for training, it is possible that better classification

could be achieved by adjusting the weights of the two channels as well as by incorporating

the recordings from all 28 channels that were recorded during the experiments. If we return

to Eq. (6.4), the velocity is a combination of the current amplitude estimate of the two

control channels. The two channels were combined with a fixed weighting of +1 and -1.

On each trial, the cursor position is the integration of the calculated cursor velocities from

the start of the trial. Thus, the final cursor position can be calculated by integrating the

velocity of the cursor or equivalently each channel's amplitude estimate:

$$\dot{x}(t) = g \cdot [d_+(t) - d_-(t)] \quad (6.4)$$

$$x_{\text{Final}} = g \cdot [D_+ - D_-] \quad (6.5)$$

where D represents the integrated amplitude estimate, d, for each channel.

To test the effect using more channel information for decoding, a linear least-squared regression was performed. The RMS amplitudes between 75-105 Hz for each channel on each trial were calculated for each movement period as previously described. The RMS amplitudes were normalized by performing the z-transform on each channel. The desired target, T, during the 1D task, was assigned a value +1 and -1 to represent the two different targets. The predicted target value, $\hat{T}$, was then calculated by a linear combination of the normalized RMS amplitude estimate, z, for each channel i using the weights, β. The three different models used were:

$$\hat{T} = z_+ + z_- \; : \; \beta_+ = +1, \beta_- = -1 \quad (6.6)$$

$$\hat{T} = \beta_+ z_+ + \beta_- z_- \quad (6.7)$$

$$\hat{T} = \sum_{i=1}^{28} \beta_i z_i \quad (6.8)$$

In the first model, Eq. (6.6), the normalized RMS amplitudes are combined with fixed, equal, and opposite weights to recreate the push-pull decoding used during the closed-loop tasks. The fixed weights are replaced with regression terms for the second model, Eq. (6.7), such that one channel can have a larger effect on predicting the target than the other. Only the amplitude estimates from the two channels actually used for closed-loop control are used in

this model. Finally, the third model, Eq. (6.8), uses the 75-105 Hz amplitude estimates for all 28 of the recorded channels to try to predict the desired target.

For each of the six monkey-dimension combinations, the same week of test data used in Chapter 3 was used. 5-fold cross-validation was used. Each of the five days of the week was left out with the other four days used as the training data. The predicted values, $\hat{T}$, from the three different models were then tested against different thresholds to generate a receiver operating characteristic (ROC) curve for each validation day. The ROC curves for each of the test data sets from the six monkey-dimension combinations are shown in Figure 6.9. In these plots, all positive targets (right or up) should be larger than the threshold before any negative targets (left or down). Thus, perfect classification would be a curve to the far upper-left going through the point (0,1) while chance would be a line running along the diagonal. These results can be grouped into two different effects. For four out of the six combinations (M-X, J-X, J-Y, and N-Y), adjusting the weights away from pure push-pull resulted in improved classification. These were the four cases where modulation appeared to occur mostly on only one of the two channels and was not true push-pull control. Interestingly, there does not appear to be much of an increase in classification performance by incorporating all 28 channels. In the other two cases (M-Y and N-X) where the adapted strategy included modulating both channels, the two channel regression did not perform significantly better than the fixed weights. This is likely because these two monkeys adjusted their signals to more closely align with the fixed push-pull weights that were being used. In these two cases incorporating information from more channels did significantly improve

classification. A summary of these results is shown by taking the mean area under the curve of the ROC analysis and is shown in Figure 6.10.

There appears to be some interesting differences in how subjects learn ECoG BCI control based on whether a monkey modulated both channels in a push-pull fashion or only relied primarily on modulation of one channel. For monkeys that relied on modulating only one channel, they likely chose this strategy because this single channel had a strong signal that allowed them to perform the task reliably. However, if the monkey could not find a single channel that could reliably be used, more effort was put into adjusting the modulation on the two channels to more closely match the decoding scheme. This hypothesis plays out in our results as the single channel modulation monkeys have better classification using only the two channels for regression. But when the amount of available modulation is more limited on these two channels, as evidenced by worse classification from only two channels that improves with all 28 channels, the monkeys work harder to adapt a more push-pull strategy with equal and opposite modulation. If true push-pull control is desired for subjects, it may be possible to make training protocols that are more likely to achieve this type of control by making sure that the monkey does not have a single channel that becomes too dominant. Instead, the weights could be adjusted to force the subject to use and strengthen weakly modulated channels by weakening the weights and resulting control effect of the strongly modulated channels.
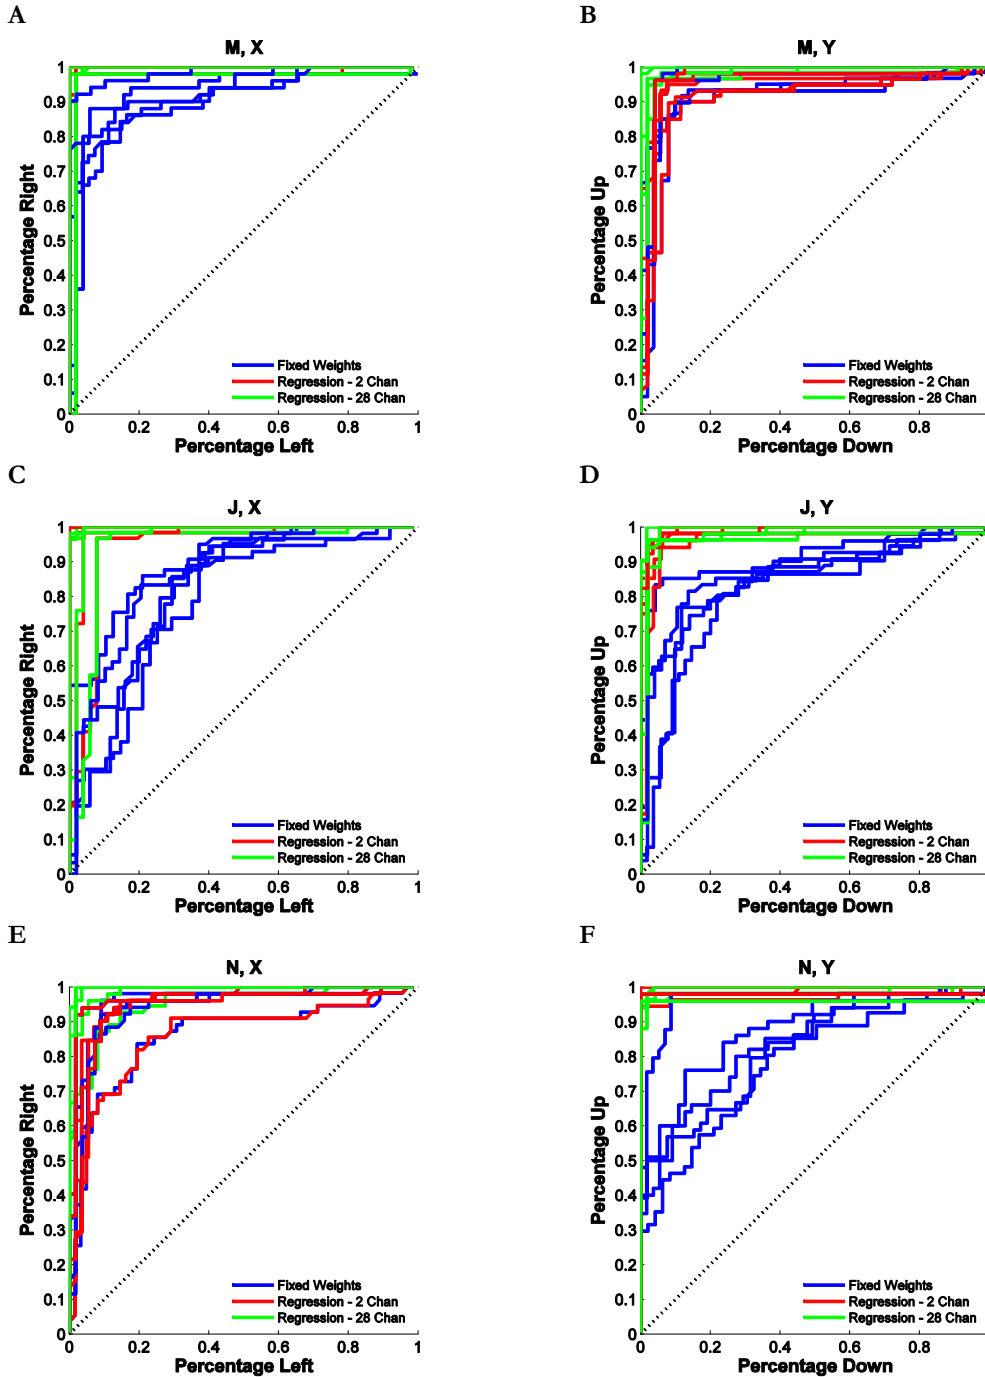
Figure 6.9. ROC Curves for Different Weights.
The observed classifications of the correct target using three different weighting functions. The fixed weights (blue) forced a push-pull classification that was similar to the actual decoding scheme. The 2 channel (red) and 28 channel (green) regressions generated weights based on a regression of the observed features and the known desired targets.

Figure 6.10. Mean Area under Curve for Different Weights.
The mean area under the curve was used for the three different classifications shown in Figure 6.9.
Error bars represent the 95% confidence intervals. Means and intervals were calculated using the
arcsine transformation.

7 Conclusion

These experiments were designed to examine the feasibility of an epidural ECoG BCI in a chronic setting as well as quantify the observed brain control responses in a basic one-dimensional BCI task for the future advancement and design of ECoG BCIs. Perhaps the most clear and promising finding is the fact that all three monkeys were able to learn two different 1D control configurations that were based on no previous screening tasks. While there was some variation in the level of performance achieved, every configuration that was used in these experiments provided control that was significantly above chance. Additionally, the learning rate to adapt to a novel configuration occurred on the order of days and not months.

7.1 BCI Training

It is our belief that the learning was greatly enhanced by the combination brain+bias training scheme that we devised. This scheme kept the monkey focused on the task by always providing a reasonable reward level but also always challenging him by allowing for improved performance by better signal modulation at any point in time. It is likely that this type of training scheme is most critical for non-human primates where task instruction is limited, but we also feel that it will be a useful tool even in human BCI applications. As with nearly all types of learning, we believe that BCI training will be the most successful when attention to the task is at its highest, and the subject is neither too frustrated by low performance nor unmotivated when the task is trivial.

One question that these experiments did not explore and could be perhaps explored in the future is whether the brain+bias training could be modified to further push the subjects to improve signal modulation. In the present experiments the bias was reduced to zero so that the subjects were performing the task under total brain control. However, there was nothing to stop us from attempting to use a negative bias that moves the cursor in the opposite direction as the correct target. The subject would then be required to increase modulation more to overcome the negative bias to select the correct target. This type of resistance training could result in a more reliable and larger modulation than we observed when training was limited to only reducing positive bias until reaching purely brain control.

While the subjects were all able to modulate cortical signals to move the cursor, our push-pull decoding scheme resulted in mixed results. Early on during the experiments we determined that it was necessary to normalize the two control signals to make sure that overall signal amplitudes and electrode property differences did not enable the subject to use global amplitude changes to create a differential signal that would move the cursor. While this did result in what we believe to be more selective cortical areas of modulation, it did not always require the true push-pull modulation that we desired. Instead of amplitude on one electrode being increased for one direction and increased on the other electrode for the opposite direction, we often saw one cortical area performing all of the modulation on one electrode while the other electrode was relatively unmodulated.

When the monkeys chose to adopt this single channel modulation strategy, we believe it resulted in asymmetrical performance to the two targets. Performance was almost

always faster and more reliable to the target where the single electrode being modulated was selected by increasing the signal amplitude. This type of asymmetry is not ideal and could be detrimental to performing BCI control depending on the specific task. As a proposed strategy for addressing this issue, we turn again to the idea of resistance training for BCIs. When one electrode signal modulation is observed to be stronger than the other electrode, the relative weighting between the two push-pull electrodes could be adjusted to increase the relative contribution of the more weakly modulated electrode. This could strengthen the modulation occurring on the weaker channel and create a more balanced push-pull control.

The question of what frequency range is best for BCI control using ECoG is far from being determined. In fact, the underlying electrophysiological processes that result in observed changes within different frequency ranges are still not well understood. Our results here show that the 75-105Hz range appears to be a viable choice. Based on the simulations performed in Chapter 6, the high gamma frequency range around 75-105 Hz appeared to be the optimal frequency range for the signals recorded during the closed-loop experiments. Additionally, the relatively broadband nature of the signals suggests that band pass filters with sharp transitions and slow responses are unnecessary. This should allow for improved filter design with slower time lags during closed-loop tasks. For the closed-loop kinematic cursor tasks that were used in these experiments, the difference in time lags may not be a large factor as delays from cortical activity to BCI cursor movement may be on the order of or even faster than motor cortex activity to natural limb movement times. However, time delays become more critical when the BCI system is used for control of a robotic arm or functional stimulation of the subject's own arm (Taylor et al., 2003).

114

Our closed-loop experiments also used only a single amplitude estimate of one frequency range for BCI control. It is certainly possible that control could be improved by increasing the number of spectral features used from a given channel. Specifically, the observed relationship that the lower frequency beta band tended to decrease in amplitude with high gamma increases suggest that decoding incorporating both frequencies could potentially improve closed-loop accuracy. Although not presented here, our lab is further analyzing our closed-loop data for better frequency decoding techniques and devising ways to incorporate more frequencies into future closed-loop experiments.

7.2 Interelectrode Distance

A key component to making epidural ECoG a viable BCI modality will be determining the number of independent degrees of freedom that can simultaneously be used for control. The data presented here required control in only one dimension. However, since push-pull control from two electrodes was used we can gain insight into what spatial limits there are for BCI control with ECoG. In our experiments, both 15mm and 9mm electrode separations worked very well for control with the two control electrodes having minimal correlation between the two sites. When the control electrodes were switched to only a 3mm separation, the performance diminished. It is important to remember that the BCI performance did not diminish to chance but still had a significant level of control with all cases still selecting over 70% of targets correctly. Additionally, subjects were only given one week at each new configuration before moving on to a different control configuration. It is possible that performance could have improved even further with additional training.

This is one of the first studies to examine the spatial resolution that is obtainable with closed-loop BCI control. It is important to appreciate the difference between spatial resolution during recording and spatial resolution during BCI control. In normal mapping experiments, the recorded signals are correlated to some task that is being performed. The analysis is designed to solve what is commonly referred to as the inverse problem. A grouping of sensors is used to infer the underlying process. In our case, the underlying neural activity is inferred based on the recorded field potential at a series of points. The inverse problem is often ill-posed without a single, unique solution but rather relies on assumptions about the underlying physical system. For field potential recordings, the simplest assumption to make is that each recording site represents the cortical activity directly under the electrode and is minimally influenced by surrounding areas. For improved models, spatial filtering based on the Laplacian operator is often used to improve the precision of the predictions of the spatial activity.

In the BCI experiments described here, the transformation is occurring in the opposite direction. The electrode coding scheme was determined by the experimental set-up and it is the subject's responsibility, based on feedback, to create a cortical activity pattern that generates the proper signal at the recording sites. Once again, there are many potential solutions since the voltage at the recording site is a spatial summation of the underlying cortical activity. The best cortex for the subject to modulate is close to the recording site but there is a whole region that affects the signal that diminishes the farther away the cortical area is from the recording site.

To illustrate this point, a simple model was designed to examine the difference between the natural cortical activity to electrode transformation compared to the new mapping of the control electrode configuration onto the underlying cortical region. To start, the measured potential from a dipole moment was modeled using the following far-field approximation:

$$\Phi(r) = \frac{1}{4\pi\epsilon}\frac{\boldsymbol{p}\cdot\hat{\boldsymbol{r}}}{r^2} \quad (7.1)$$

Here, $\boldsymbol{p}$ represents the dipole moment created by the neuronal activity and ϵ is the permittivity of the tissue. The distance from the dipole source to the recording location is given by r. $\hat{\boldsymbol{r}}$ is a unit vector pointing in the direction from source to recording location.

A simple one-dimensional model of field potential recordings at an electrode above a layer of neuronal activity was created. To calculate the total potential at any given electrode location (x_e), we must integrate all of the dipoles lying along the cortical sheet of activity at each neuron location (x_n).

$$\Phi(r_e) = \int \frac{1}{4\pi\epsilon}\frac{\boldsymbol{p}_n\cdot\hat{\boldsymbol{r}}}{(r_n-r_e)^2}dr_n \quad (7.2)$$

If the dipole orientation is assumed to be constant throughout the cortical layer of activity, the recording site potential is proportional to the strength of each neuronal dipole divided by the squared distance between the electrode and cortical site.

$$\Phi(r_e) \propto \int \frac{p_n(r_n)}{(r_n - r_e)^2} dr_n \qquad (7.3)$$

When recording naturally occurring cortical activity, the measured field potential is simply the summation of all the underlying activity. If trying to selectively identify two different cortical columns, the measureable potential difference ($\Delta\Phi$) between two separately activated cortical site locations (r_{n1} and r_{n2}) at any given recording site (r_e) is:

$$\Delta\Phi(r_e) \propto \int \frac{p_{n1}(r_{n1})}{(r_{n1} - r_e)^2} dr_{n1} - \int \frac{p_{n2}(r_{n2})}{(r_{n2} - r_e)^2} dr_{n2} \quad (7.4)$$

Figure 7.1a. A traditional cortical column is modeled as a 0.5 mm wide area of cortex that is selectively activated. The two columns are separated by 15, 9, and 3 mm spatial separation. The resulting measurable difference, $\Delta\Phi$, is plotted along the brain surface with the six electrode locations showing the potential spatial sampling with our grid with 3mm interelectrode distance.

For the closed-loop BCI task, we are not concerned about mapping cortex but rather the subject is identifying which cortical areas control the velocity of the cursor. Returning to Eq. (3.4), the velocity is determined by the amplitude difference at 75-105 Hz difference between the positive and negative electrode.

$$\dot{x}(t) = g \cdot [d_+(t) - d_-(t)] \quad (7.5)$$

If we want to get the effect on the cursor velocity from any given location of cortical activity $\dot{x}(r_n)$, we must calculate the differential effect from any given dipole on the two recording sites.

$$\dot{x}(r_n) \propto \int \frac{p_n(r_n)}{(r_n - r_{e1})^2} dr_n - \int \frac{p_n(r_n)}{(r_n - r_{e2})^2} dr_n \quad (7.6)$$

Combining the terms, the relative effect of any cortical site on the cursor's velocity can be represented as the difference of the reciprocal of the squares of the distance times the strength of the dipole:

$$\dot{x}(r_n) \propto \left(\frac{1}{(r_n - r_{e1})^2} - \frac{1}{(r_n - r_{e2})^2} \right) \int p_n(r_n) \, dr_n \quad (7.7)$$

Simplifying, the relative effect of any cortical site on cursor velocity is the BCI control strength (Γ_{BCI}) of any cortical column times the dipole strength (p_n).

119

$$\dot{x}(r_n) \propto \Gamma_{BCI} \cdot p_n \quad , where \ \Gamma_{BCI} = \left(\frac{1}{(r_n - r_{e1})^2} - \frac{1}{(r_n - r_{e2})^2} \right) (7.8)$$

Figure 7.1b shows a plot of this BCI control strength (Γ_{BCI}) for a layer of cortex when the pair of assigned push-pull control electrodes are separated by 15, 9, and 3 mm.

While the differential signal that can be observed in part A shrinks when the cortical columns are moved closer together, the optimal spatial recording location is determined by where the peak field potential change is located. Conversely in part B, with the reverse case of BCI induced cortical activation, the ability to successful control the cursor is a summation of the total cortical area being modulated. As the electrodes move closer together, not only does the peak ability of a given cortical location to selectively influence only one electrode decrease, but the total cortical area that can be modulated for control also shrinks.

In our experiments to determine the optimal spacing, the effectiveness of BCI control appeared to fall off once the control electrode separation was 3 mm. We hypothesize that the reorganization of the cortex for BCI control occurs in a broader, less organized fashion than the typical columnar organization of cortex. While the columnar organization of neurons that typically appears is on the order 0.3-0.5 mm, the cortical control columns associated with epidural ECoG BCI appear to be on the order of ten times larger. When performing the task, the monkey is only left with a global reward signal for whether the correct neural modulations are occurring. Therefore, even though a signal may be best affected by the nearest cortical region, it is in the monkey's best interest to modulate

any cortical area that has a positive effect on the cursor. Additionally, by using a larger

cortical area, the parallel adaptation of a larger population of neurons can potentially create

larger signals and foster faster learning.

Choosing the optimal spacing, size, and number of channels for any ECoG grid for

BCI applications will depend on a number of factors. In our view, the proper spacing will

be somewhere on the order of the 3mm that was used in these experiments. A recording

site 3mm from all other sites can still provide useful information. However, in the current

experiments the 3mm control spacing did appear to be approaching a limit of diminishing

independent information compared to the larger spacings.

A   **B**
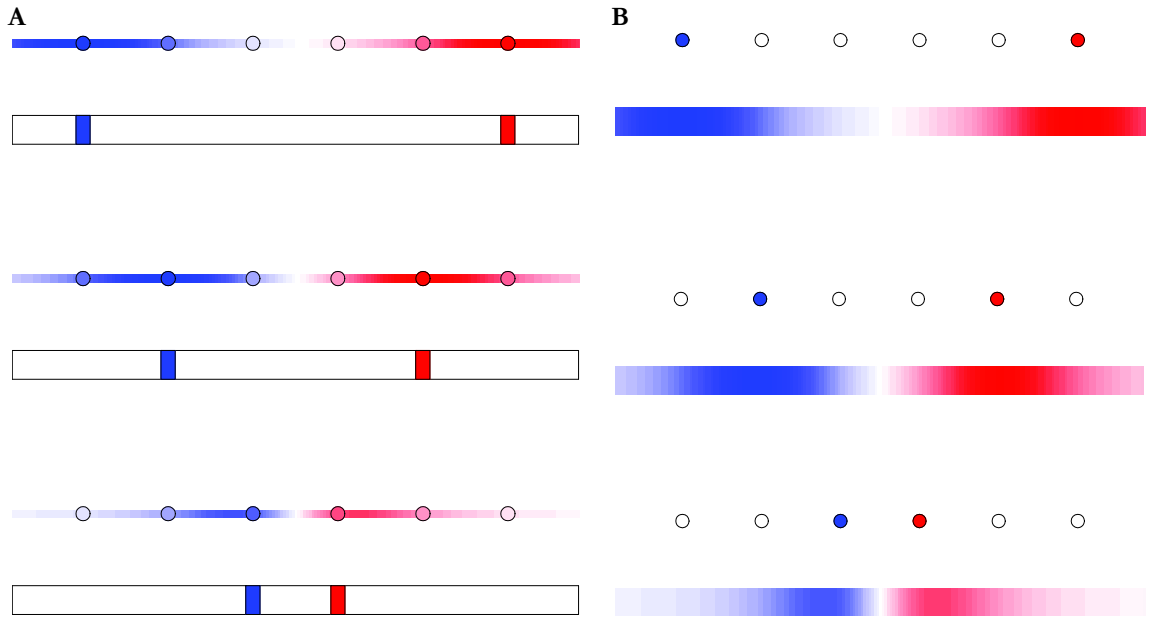
Figure 7.1. Cortical Columns vs. Control Columns.
A model was developed to show the difference in the a) cortical activity to ECoG recording
electrodes and the b) BCI ECoG electrode to cortical activity transformation. The sampled field
potential resulting from activated cortical columns is shown in part a. The relative effects of cortical
activation on the BCI electrodes is shown in part b. All electrodes are 3 mm apart and the electrode
plane is 3 mm from the underlying cortical sheet of dipoles. The field potential is modeled to fall off
at 1/r² as expected for dipoles.

7.3 Future Work

All of the experiments here were done with a fixed and assigned BCI decoding scheme that was completely independent of the results that we observed. We feel that it highlights and better quantifies the amount of modulation and adaptation that the brain can perform to accomplish a BCI task with epidural ECoG. It is important to stress that the results are not meant to represent the optimal BCI performance that can be obtained with our current recording setup. Chapter 6 highlighted several ways to potentially improve performance. Based on the simulations in the chapter, it appears that our amplitude estimation technique between 75-105 Hz was well suited for the signals we observed. It is possible that better performance could be achieved by adjusting the bandpass filters to fit the observed signals as well as improve the temporal response of the system. The cortical modulation that we observed was well matched to the decoding parameters that were chosen. It is difficult to assess how much this effect is a result of the subjects being able to adapt to the task presented to them and how much we truly did pick some of the best parameters for control. This illustrates a key point when doing BCI experiments. There is no substitute for gaining the information that can be obtained by true closed-loop BCI experiments. While simulations as well as mapping and decoding studies are useful, it is impossible to fully appreciate which features are most important and best suited to being adapted for BCI control without completing closed-loop testing.

In addition to the frequency and filtering properties examined in Chapter 6, it was shown that decoding could also be improved by better incorporating the observed

differences across channels instead of the forced, fixed weights used here. Co-adaptive algorithms have been used successfully in the past in single-unit BCI devices (Taylor et al., 2002). Since the described experiments have been completed, our lab has moved on to more co-adaptive (both BCI decoding and subject adaptation) studies using information from all 28 channels to improve performance. Immediately after the 1D experiments were completed, the two sets of fixed channel assignments were combined to perform a 2D task. After a plateau was reached with the fixed 2D task, a series of experiments with adaptive decoding were performed. Each day, partial least squared regression was performed to modify the channel weights from all 28 recorded channels that would have provided the best performance from the previous day. Figure 7.2 shows the performance change that occurred once the co-adaptive strategy was used. Both monkeys clearly improved by using the adaptive weights from all 28 channels. The actual weights used for monkey M are shown in Figure 7.3. It appears that the weights were relatively stable and appeared to incorporate the amplitude from mainly four channels to control the cursor.

It is our belief that this kind of co-adaptive strategy may provide the best way to develop ECoG BCI control. An initial decoding scheme can be devised that is known to typically be successful for BCI control. Through feedback, the subject is trained to modulate cortical activity in a way that is well suited for the given BCI tasks. Once the subject has some moderate level of control, the BCI system can then be adjusted to better fit the observed signals for the given user. A co-adaptive approach can then be pursued where the BCI decoding scheme adjusts to the user while the user is also adapting the updated weights. A challenge for developing future systems will be to determine how quickly BCI decoding

should be changed such that the system and user are converging toward a common solution and not adapting to each other in such a way that both are changing too quickly for the other to catch up.

Figure 7.2. Fixed vs. Adaptive Weight Performance.
The percent correct for monkeys M and J when fixed weight 2-dimensional control was switched to adaptive weights using all 28 channels. Performance had reached a plateau but was improved by using new weights calculated each day. (Courtesy of JJ Williams)

Figure 7.3. Channels Weights.
The weights as they were changed each day to adapt to the observe signals. Approximately 4 channels appear to be used for control. (Courtesy of JJ Williams)

7.4 Final Thoughts

The twentieth century saw tremendous advancements in science and technology. Perhaps the two crowning achievements were the advancements made in the biological sciences as well as computing. Advances in biology led to new medical innovation that prevented and cured countless diseases to improve human health. Even more abruptly, the last half of the century was an age where computing power has been continuously growing and becoming more ubiquitous in people's lives. While there are many ways that medicine can still be improved, neurological diseases represent a final frontier where many of the diseases are poorly understood with very few direct treatment options. It is the combination of computing power available to researchers and the increased fundamental knowledge in biology that make these current times an exciting period to be in the field of neuroscience. Concepts like brain-computer interfaces would likely seem very foreign to people only a short time ago. This author is excited to be a part of these incredible times. Much work is left to be done, but it is this researcher's hope that the experiments described here will play at least a small role in improving future work using electrocorticographic recordings and designing practical brain-computer interfaces.

References

Ang, K. K., Guan, C., Sui, K., Chua, G., Ang, T., Kuah, C., et al. (2010). Clinical study of neurorehabilitation in stroke using EEG-based motor imagery brain-computer interface with robotic feedback. *Conf Proc IEEE Eng Med Biol Soc*, *1*, 5549-5552.

Betz, W. (1874). Anatomischer Nachweis zweier Gehirncentra. *Centralblatt für die medizinischen Wissenschaften*, *12*, 578-580, 595-599.

Bjornsson, C. S., Oh, S. J., Al-Kofahi, Y. A., Lim, Y. J., Smith, K. L., Turner, J. N., et al. (2006). Effects of insertion conditions on tissue strain and vascular damage during neuroprosthetic device insertion. *J Neural Eng*, *3*(3), 196-207.

Blakely, T., Miller, Kai J, Zanos, S. P., Rao, R. P. N., & Ojemann, Jeffrey G. (2009). Robust, long-term control of an electrocorticographic brain-computer interface with fixed parameters. *Neurosurg Focus*, *27*(1), E13.

Blatt, G. J., Andersen, R. a, & Stoner, G. R. (1990). Visual receptive field organization and cortico-cortical connections of the lateral intraparietal area (area LIP) in the macaque. *J Comp Neurol*, *299*(4), 421-45.

Brindley, G. S., & Craggs, M. D. (1972). The electrical activity in the motor cortex that accompanies voluntary movement. *J Physiol*, *223*(1), 28P-29P.

Carmena, J M, Lebedev, M. A., Crist, R. E., O'Doherty, J. E., Santucci, D. M., Dimitrov, D. F., et al. (2003). Learning to control a brain-machine interface for reaching and grasping by primates. *PLoS Biol*, *1*(2), E42.

Chao, Z. C., Nagasaka, Y., & Fujii, N. (2010). Long-term asynchronous decoding of arm motion using electrocorticographic signals in monkeys. *Front Neuroeng*, *3*(March), 3.

Christopher & Dana Reeve Foundation. (2009). One Degree of Separation, Paralysis and Spinal Cord Injury in the United States. *Christopher & Dana Reeve Foundation*.

Crone, N. E., Miglioretti, D. L., Gordon, B., & Lesser, R. P. (1998). Functional mapping of human sensorimotor cortex with electrocorticographic spectral analysis. II. Event-related synchronization in the gamma band. *Brain*, *121*, 2301-2315.

Crone, N. E., Miglioretti, D. L., Gordon, B., Sieracki, J. M., Wilson, M. T., Uematsu, S., et al. (1998). Functional mapping of human sensorimotor cortex with electrocorticographic

spectral analysis. I. Alpha and beta event-related desynchronization. *Brain, 121*, 2271-2299.

Felton, E. A., Wilson, J. A., Williams, Justin C, & Garell, P. C. (2007). Electrocorticographically controlled brain-computer interfaces using motor and sensory imagery in patients with temporary subdural electrode implants. Report of four cases. *J Neurosurg, 106*(3), 495-500.

Fetz, E. E., & Baker, M. A. (1973). Operantly Conditioned Patterns of Precentral Unit Activity and Correlated Responses in Adjacent Cells and Contralateral Muscles. *J Neurophysiol, 36*(2), 179-204.

Georgopoulos, A. P., Kalaska, J. F., Caminiti, R., & Massey, J. T. (1982). On the relations between the direction of two-dimensional arm movements and cell discharge in primate motor cortex. *J Neurosci, 2*(11), 1527-37.

Georgopoulos, A. P., Schwartz, A B, & Kettner, R. E. (1986). Neuronal Population Coding of Movement Direction. *Science, 233*(4771), 1416-1419.

Heldman, D. A. (2007). Epidural Electrocorticography and Intra-Cortical Local Field Potentials in Motor Cortex during Volitional Arm Movements and their Applications to Neural Prosthetic Control. *Doctoral Dissertation, Washington University in St. Louis.*

Heldman, D. A., Hokanson, J. A., & Moran, D W. (2004). Local field potential spectral tuning in primary motor cortex. *Soc Neurosci Abstr* (p. 421.22).

Heldman, D. A., Wang, W., Chan, S. S., & Moran, D W. (2006). Local field potential spectral tuning in motor cortex during reaching. *IEEE Trans Neural Syst Rehabil Eng, 14*(2), 180-183.

Jasper, H., & Penfield, W. (1949). Electrocorticograms in man: Effect of voluntary movement upon the electrical activity of the preeentral gyrus. *Arch Psychiatr Nervenkr, 183*, 163-174.

Jerbi, K., Ossandón, T., Hamamé, C. M., Senova, S., Dalal, S. S., Jung, J., et al. (2009). Task-related gamma-band dynamics from an intracerebral perspective: review and implications for surface EEG and MEG. *Hum Brain Mapp, 30*(6), 1758-71.

Kakei, S., Hoffman, D. S., & Strick, P. L. (1999). Muscle and movement representations in the primary motor cortex. *Science, 285*(5436), 2136-9.

Kennedy, P. R., & Bakay, R. A. (1998). Restoration of neural output from a paralyzed patient by a direct brain connection. *Neuroreport, 9*(8), 1707-1711.

Kim, S.-P., Simeral, J. D., Hochberg, L. R., Donoghue, J. P., & Black, M. J. (2008). Neural control of computer cursor velocity by decoding motor cortical spiking activity in humans with tetraplegia. *J Neural Eng*, *5*(4), 455-76.

Kübler, A., Nijboer, F., Mellinger, J., Vaughan, T. M., Pawelzik, H., Schalk, G., et al. (2005). Patients with ALS can use sensorimotor rhythms to operate a brain-computer interface. *Neurology*, *64*(10), 1775-7.

Leuthardt, E. C., Schalk, G., Wolpaw, J. R., Ojemann, J G, & Moran, D W. (2004). A brain-computer interface using electrocorticographic signals in humans. *J Neural Eng*, *1*(2), 63-71.

Levine, S. P., Huggins, J. E., BeMent, S. L., Kushwaha, R. K., Schuh, L. A., Passaro, E. A., et al. (1999). Identification of Electrocorticogram Patterns as the Basis for a Direct Brain Interface. *J Clin Neurophysiol*, *16*(5), 439.

Liu, J., & Newsome, W. T. (2006). Local field potential in cortical area MT: stimulus tuning and behavioral correlations. *J Neurosci*, *26*(30), 7779-7790.

Manning, J. R., Jacobs, J., Fried, I., & Kahana, M. J. (2009). Broadband shifts in local field potential power spectra are correlated with single-neuron spiking in humans. *J Neurosci*, *29*(43), 13613-20.

McFarland, D. J., Sarnacki, W., & Wolpaw, J. R. (2003). Brain–computer interface (BCI) operation: optimizing information transfer rates. *Biol Psychol*, *63*(3), 237-251.

Miller, K J, Sorensen, L. B., Ojemann, J G, & DenNijs, M. (2009). Power-law scaling in the brain surface electric potential. *PLoS Comp Biol*, *5*(12), e1000609.

Moran, D W, & Schwartz, A. (1999). Motor cortical representation of speed and direction during reaching. *J Neurophysiol*, *82*(5), 2676-2692.

Penfield, W., & Boldrey, E. (1937). Somatic motor and sensory representation in the cerebral cortex of Man as studied by electrical stimulation. *Brain*, *60*(4), 389-443.

Pesaran, B., Pezaris, J. S., Sahani, M., Mitra, P. P., & Andersen, R. A. (2002). Temporal structure in neuronal activity during working memory in macaque parietal cortex. *Nat Neurosci*, *5*(8), 805-811.

Pfurtscheller, G., & Neuper, C. (1997). Motor imagery activates primary sensorimotor area in humans. *Neurosci Lett*, *239*(2-3), 65-8.

Pierce, J. R. (1980). *An Introduction to Information Theory* (pp. 145-165). Dover.

Porro, C. a, Francescato, M. P., Cettolo, V., Diamond, M. E., Baraldi, P., Zuiani, C., et al. (1996). Primary motor and sensory cortex activation during motor performance and motor imagery: a functional magnetic resonance imaging study. *J Neurosci*, *16*(23), 7688-98.

Rasch, M. J., Gretton, A., Murayama, Y., Maass, W., & Logothetis, N. K. (2008). Inferring spike trains from local field potentials. *J Neurophysiol*, *99*(3), 1461-1476.

Ray, S., Crone, N. E., Niebur, E., Franaszczuk, P. J., & Hsiao, S. S. (2008). Neural correlates of high-gamma oscillations (60-200 Hz) in macaque local field potentials and their potential implications in electrocorticography. *J Neurosci*, *28*(45), 11526-11536.

Reina, G. A., Moran, Daniel W, Schwartz, Andrew B, Chase, S. M., Kass, R. E., Ganguly, K., et al. (2001). On the Relationship Between Joint Angular Velocity and Motor Cortical Discharge During Reaching On the Relationship Between Joint Angular Velocity and Motor Cortical Discharge During Reaching. *J Neurophysiol*, 2576-2589.

Rickert, J., Oliveira, S. C., Vaadia, E., Aertsen, A., Rotter, S., & Mehring, C. (2005). Encoding of movement direction in different frequency ranges of motor cortical local field potentials. *J Neurosci*, *25*(39), 8815-8824.

Schalk, G., Miller, K J, Anderson, N. R., Wilson, J. a, Smyth, M. D., Ojemann, J G, et al. (2008). Two-dimensional movement control using electrocorticographic signals in humans. *J Neural Eng*, *5*(1), 75-84.

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

Schwartz, A., & Moran, D W. (1999). Motor cortical activity during drawing movements: population representation during lemniscate tracing. *J Neurophysiol*, *82*(5), 2705-2718.

Schwartz, A., Cui, X. T., Weber, D. J., & Moran, D W. (2006). Brain-controlled interfaces: movement restoration with neural prosthetics. *Neuron*, *52*(1), 205-220.

Serruya, M., Hatsopoulos, N., Paninski, L., Fellows, M. R., & Donoghue, J. P. (2002). Instant neural control of a movement signal. *Nature*, *416*(March), 141-142.

Sokal, R. R., & Rohlf, F. J. (1995). *Biometry: the principles and practice of statistics in biological research. WH Freeman* (Vol. 3, pp. 419-422).

Taylor, D. M. D., Tillery, S. I. H., & Schwartz, A. (2003). Information Conveyed Through Brain-Control: Cursor Versus Robot. *IEEE Trans Neural Syst Rehabil Eng*. IEEE Engineering in Medicine and Biology Society.

Taylor, D. M., Tillery, S. I., & Schwartz, Andrew B. (2002). Direct cortical control of 3D neuroprosthetic devices. *Science*, *296*(5574), 1829-1832.

Velliste, M., Perel, S., Spalding, M. C., Whitford, A. S., & Schwartz, A. (2008). Cortical control of a prosthetic arm for self-feeding. *Nature*, *453*(7198), 1098-1101.

Wang, W., Chan, S. S., Heldman, D. A., & Moran, D W. (2007). Motor cortical representation of position and velocity during reaching. *J Neurophysiol*, *97*(6), 4258-4270.

Wessberg, J., Stambaugh, C. R., Kralik, J. D., Beck, P. D., Laubach, M., Chapin, J. K., et al. (2000). Real-time prediction of hand trajectory by ensembles of cortical neurons in primates. *Nature*, *408*(6810), 361-5.

Williams, J C, Hippensteel, J. A., Dilgen, J., Shain, W., & Kipke, D. R. (2007). Complex impedance spectroscopy for monitoring tissue responses to inserted neural implants. *J Neural Eng*, *4*(4), 410-423.

Wolpaw, J. R., & McFarland, D. J. (2004). Control of a two-dimensional movement signal by a noninvasive brain-computer interface in humans. *Proc Natl Acad Sci U S A*, *101*(51), 17849-17854.

Zimmermann-Schlatter, A., Schuster, C., Puhan, M. a, Siekierka, E., & Steurer, J. (2008). Efficacy of motor imagery in post-stroke rehabilitation: a systematic review. *J Neuroeng Rehabil*, *5*, 8.