Goal Selection as a Control Strategy in a Brain-Computer Interface

A DISSERTATION
SUBMITTED TO THE FACULTY OF THE GRADUATE SCHOOL
OF THE UNIVERSITY OF MINNESOTA
BY

Audrey Nicole Smith Royer

IN PARTIAL FULFILLMENT OF THE REQUIREMENTS
FOR THE DEGREE OF
DOCTOR OF PHILOSOPHY

Bin He, Adviser

September, 2011

Acknowledgements

I would like to acknowledge the other people who contributed to the works contained in this thesis. My advisor Bin He supported me throughout the process. In section 2.1, Andrew McCullough and Cristina Rios assisted with data collection. In section 2.2, Minn Rose contributed intellectually and helped gather data; I am also grateful to Alex Doud, Han Yuan, Dan Rinker, and Cristina Rios for assistance in data collection and many useful discussions. In chapter 3, Alex Doud and Minn Rose contributed intellectually; I am also grateful to Dan Rinker for data collection assistance, Han Yuan for useful discussions, Keith Jamison for assistance with Matlab, and the Blender community. In chapter 4, I am grateful to Alex Doud, Minn Rose, and Keith Jamison for assistance in data collection and many useful discussions. I would like to thank all my subjects for their time and cooperation.

Several groups of people contributed to my overall intellectual development. I would like to thank my thesis committee for their guidance along this path: James Ashe, Geoff Ghose, Sheng He, and David Redish. I would like to thank everyone from the Biomedical Functional Imaging and Neuroengineering Laboratory, both past and present. The many people in lab over time provided intellectually stimulating discussions and technical assistance.

Dedication

This dissertation is dedicated to my loving and patient husband, David Royer

Abstract

A brain-computer interface (BCI) translates signals recorded directly from the brain into commands that control an external device, such as a computer cursor, wheelchair, or neuroprosthetic. BCIs promise to help the nearly 6 million people who live with paralysis by allowing them to interact with the world in ways they are no longer able. BCIs can also be used by able bodied individuals to extend their capabilities. BCIs differ widely in how they implement the translation from raw brain signal to device command. Two competing control strategies, goal selection and process control, differ in how much the BCI assists the user. In process control, the user controls every step of the process and receives minimal to no assistance from the system. Other terms for process control include low-level control or continuous control. In goal selection, the user only needs to determine the goal and the system executes the process to achieve that goal. Other terms for goal selection include high-level control or shared control. This thesis presents the first studies directly comparing goal selection and process control. We found in these studies that the goal selection based paradigms were easier to learn, had a decreased training period, and provided improved speed, accuracy, and information transfer in both the simple and more complex applications studied. This thesis also extends our understanding of the neurophysiology while using a sensorimotor rhythm based BCI. When individual trial data were analyzed and not averaged as is typically done in the literature, we found that duration of sensorimotor rhythm modulation was more correlated to successful use than amplitude of modulation. Additionally, we found that correct modulation that led to either a single hit or overall high accuracy was the same between the two control strategies. This shows that the improved performance in these studies while using the goal selection based paradigms was more attributable to the difference in device command instead of the difference in raw brain signal. By understanding neurophysiology and applying that knowledge to BCI design, we can make a better BCI.

Table of Contents

List of Tables

List of Figures

Chapter 1 Introduction

1.1 Brain-Computer Interfaces

A brain-computer interface (BCI) strives to make a connection directly from a person's brain to a computer without relying on any motor output (Wolpaw et al 2002, Vallabhaneni et al 2005). BCIs promise to help the nearly 6 million people who live with paralysis (www.christopherreeve.org) by allowing them to interact with the world in ways they are no longer able. Those individuals have lost normal motor control through diseases and conditions such as amyotrophic lateral sclerosis (Lou Gehrig's disease), brainstem stroke, spinal cord injury, muscular dystrophies, or cerebral palsy (Kunst 2004). Although people living with these conditions suffer major muscular loss, their cognitive abilities are left intact, thus they still want to communicate and manipulate their environment (Kunst 2004). For these patients, a BCI could allow them to use a computer, a neuroprosthetic, or control a mobile robot (Kennedy et al 2000, Karim et al 2006, Hochberg et al 2006, Bell et al 2008).

BCIs can also be used by able bodied individuals to extend their capabilities. The military hopes to utilize BCIs to enhance a soldier's abilities in the theatre of war (Kotchetkov et al 2010). In the civilian realm, current video games such as the Emotiv allow gamers to journey through a mythical world controlled by their mind (www.emotiv.com).

Although BCIs can be used by a variety of people to do a variety of things, all BCIs share the same common architecture. As presented in figure 1.1, BCIs record the raw signal from the brain, process the signal, and translate the signal to a device command. The user typically receives visual feedback by watching the controlled device. BCIs vary in how they record the brain signal, what features are extracted from the brain signal, how those features are translated to a device command, and what device is being controlled.

We have already mentioned some of the devices BCIs control in the above paragraphs. Below, we briefly present recording methodologies used by BCIs and features of interest, going into detail into the method and features used in this thesis in sections 1.2 and 1.3. Section 1.4 presents two competing ways the brain signal is translated to a device command.

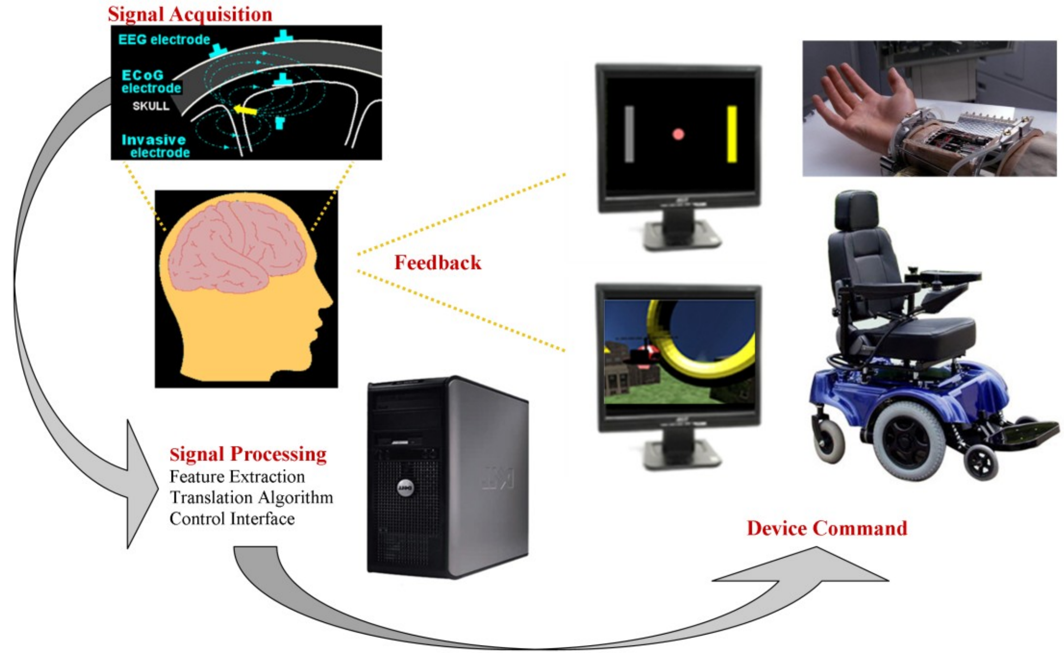

Figure 1.1 Brain-computer interface overview

BCIs use recording methodologies that cover the spectrum from non-invasive to invasive, including electroencephalography (EEG), electrocorticography (ECoG), local field potentials (LFPs), and single neuron action potential recordings (single units), as shown in figure 1.2. The non-invasive methodologies have the advantage of ease of acquisition, but suffer from decreased fidelity of the signal and lack of localization. As the methodology becomes more invasive, the signal quality and localization improves, but at the cost of increased difficulty in signal acquisition (Schwartz et al 2006).

The most popular non-invasive methodology is electroencephalography (EEG), but fMRI and MEG have also been used. EEG records the activity from at least 100 million gyral neurons that cover a spatial extent of 3 cm. Since it is non-invasive, it is easy to record. However, much information, such as amplitude, localization, and the higher frequencies, is lost as the signal travels through the skull. Electrocorticography (ECoG) removes the issues caused by the skull by placing electrodes on the surface of the brain. This methodology greatly improves spatial resolution to 0.5 cm and retains the higher frequencies (30 - 200 Hz), while being minimally invasive. The most invasive recording strategies insert electrodes into the brain to record either the local field potentials of a small population of neurons or the single unit action potentials of a neuron. These recordings retain the fidelity of the signal, but are technically challenging due to biocompatibility issues (Schwartz et al 2006).

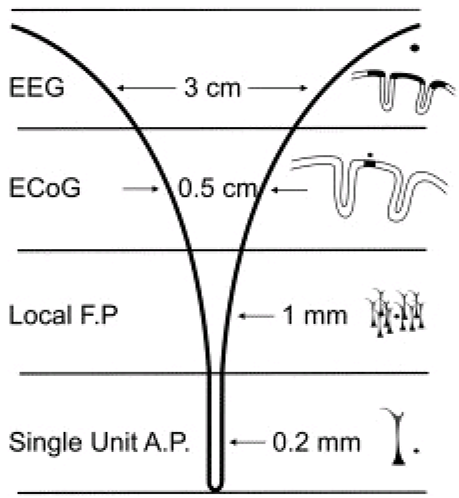

Figure 1.2 Different recording methodologies used in BCIs and their spatial domains adapted from Schwartz et al 2006.

BCIs use a variety of electrophysiological activities as the feature of interest. (1) Slow cortical potentials are changes in the cortical polarization of an EEG signal volitionally controlled by the subject lasting from 300ms to several seconds. As the name suggests, slow cortical potentials are slow and require weeks or months of training in order to produce. (2) The P300 is a positive peak in the EEG about 300ms after an infrequent or particularly significant auditory, visual, or somatosensory stimuli. The stimulus of interest is interspersed with frequent or routine stimuli. (3) Visual evoked potentials are a response to a visual stimulus such as flashing lights, which are strongest in the occipital area. A continuous, steady-state response named steady-state visual evoked potentials can be elicited in the visual pathways if the stimulus is presented at frequencies greater than 5 to 6Hz. (4) Some labs measure the specific distribution of EEG frequency patterns over the scalp in response to different mental tasks, such as mental arithmetic vs. singing a song. (5) Labs performing invasive recordings measure the local field potential or the single unit action potentials while the subject performs, imagines, or observes different tasks. (6) Sensorimotor activity creates changes in the brain rhythms measurable by EEG and will be covered in detail in section 1.3. (7) The previous 6 categories can be combined to form a hybrid BCI (Wolpaw et al 2002, Schwartz et al 2006).

1.2 Electroencephalography

The ease of obtaining EEG recordings has made it a popular tool in the world of BCI, especially for those BCIs using human subjects. What exactly is the EEG recording? At a single cell level, the source of the EEG signal is the post synaptic potential of pyramidal cell assemblies in cortex (Eiccles 1951, Dale and Sereno 1993, Kandel et al. 2000, Baillet et al. 2001). The binding of neurotransmitters at a synapse triggers ionic flow in the apical dendrites of pyramidal cells that resembles a current dipole, as depicted on the left in figure 1.3. That current dipole generates an extracellular electrical field. However, the electrical field is very small, and attenuates at a rate of one over the square of the distance. An EEG

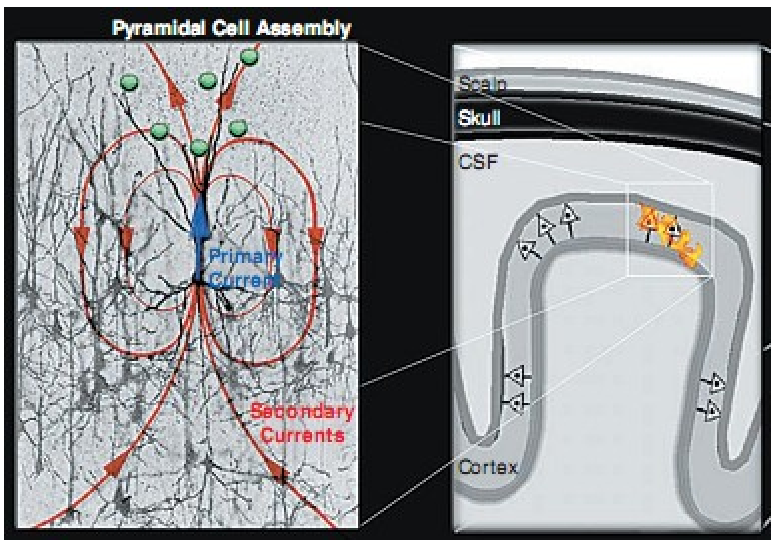

Figure 1.3 Source of EEG signal: Pyramidal cell assemblies
adapted from Baillet et al 2001

electrode lies 2 to 3cm away from the surface of the cortex, with the scalp, skull, and cerebral spinal fluid between the electrode and the electrically active neurons. Conveniently for researchers, the pyramidal cells are aligned with their dendrites perpendicular to the surface of the cortex, as illustrated on the right in figure 1.3. This alignment allows the electrical field generated by the post synaptic potential to sum when multiple neurons are excited simultaneously. Note however that the cortex folds, with gyral pyramidal cells perpendicular to the skull, and those neurons lying in the sulci oriented parallel with the skull. The opposing orientations of the neurons in the sulci lead to the addition of their fields effectively cancelling each other. In order to generate enough of a signal to be measured by EEG, large scale synchrony of pyramidal cells on the order of 100 million gyral neurons is required. That covers an area of nearly $6cm^2$, or a diameter of roughly 3cm (Baillet et al 2001, Schwartz et al 2006).

As described above, EEG measures the summation of brain activity. This summation is only limited by the propagation of the electrical activity through the skull. Frequencies above 70Hz do not survive this propagation. The skull both attenuates and smears the signal. Thus, a large amplitude signal distant from the recording electrode may be represented in the recording as equally as the lower amplitude signal from the cortex immediately below the electrode. This limits the spatial resolution of EEG (Schwartz et al 2006).

1.3 Motor Control and Sensorimotor Rhythms

A common source of the signal for a BCI is the motor and sensory cortices (Wolpaw et al 2002). Therefore, understanding how the brain processes sensory input and motor output sheds light on BCI operation and use.

The process of planning and executing movement in a healthy individual is a multistep process. The key components of a goal directed movement can be broken down into four steps. (1) Through somatosensory, proprioceptive, and visual inputs to the posterior parietal cortex, an individual is aware of the surrounding environment and his position in space. (2) The individual decides what action is desired through activity in the highly interconnected parietal and anterior frontal lobes. (3) Axons from the prefrontal and parietal cortex converge in the premotor cortex and supplementary motor areas to construct a plan that indicates how the actions will be executed. This plan is typically a general plan. The exact muscle activation sequence will be determined in the spinal cord. (4) The primary motor cortex issues a command to begin the action according to the plan (Carlson 1998). The primary motor cortex may initiate this process, but other structures, illustrated in figure 1.4, such as the basal ganglia, cerebellum, thalamus, brainstem nuclei, spinal interneurons, and spinal motor neurons are vital to produce coordinated motion (Wolpaw 2007).

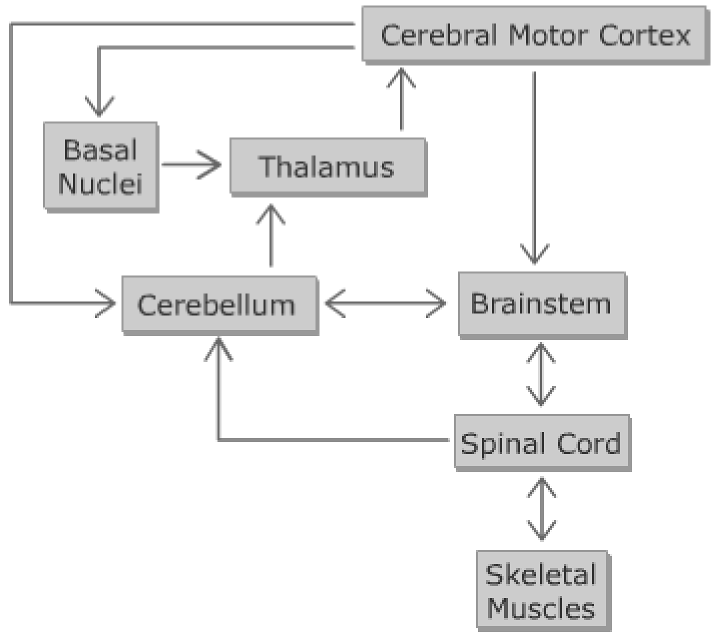

Figure 1.4 Overview of major components of the motor nervous system
adapted from www.getbodysmart.com

During the execution of a motor task, such as an arm motion, a complicated process unfolds that involves internal models and both closed loop feedback and open loop feed-forward control, as illustrated in figure 1.5. As you can see, the process of converting a

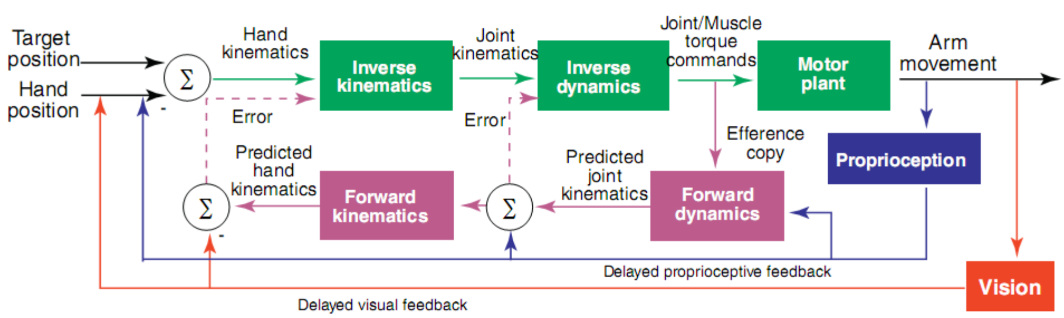

Figure 1.5 Motor control of the arm uses internal models with both feedback and feed-forward control
adapted from Green and Kalaska 2011

desired goal to the necessary spinal motor neuron firing pattern that produces the appropriate muscle contractions is a complicated matter. Humans spend many years learning and fine tuning this process to allow us to perform everyday activities such as drinking from a cup or walking.

One brain region used in both motor control and sensory perception is the thalamus. The thalamus is the gatekeeper to the cortex. Four of the five sensory modalities, motor control, and even cognitive abilities such as attention, memory, and emotion each have their own specific nuclei in the thalamus that connects them to the cortex. Coordinated activity within a corticothalamic loop produces rhythms in the 8-12Hz range measurable with EEG (Lopes da Silva 1991, Pfurtscheller and Lopes da Silva 1999). This circuitry has been specifically implicated in a wide variety of rhythmic activities, from the occipital alpha rhythm, initiation of epileptic seizures, and volitional control of a BCI (Lopes da Silva 1991, Pfurtscheller and Lopes da Silva 1999, Breakspear et al 2006).

These rhythms, when measured over sensory or motor cortex, are called sensorimotor rhythms (SMRs). One type of BCI utilizes scalp-recorded EEG to monitor sensorimotor rhythms (Wolpaw et al 2002, Vallabhaneni et al 2005). Sensorimotor rhythms are produced by both the primary sensory and motor cortices. SMR based BCI's utilize two distinct states: event related synchronization (ERS) and event related desynchronization (ERD). When an awake person is not processing sensory data or producing motor output, the primary sensory and motor cortices are in an idling state, which creates a rhythmic EEG pattern known as ERS. The mu rhythm from 8 to 12Hz and the beta rhythm from 13Hz to 30Hz have been particularly useful as BCI control signals. The sensorimotor rhythms decrease in amplitude, an effect that is known as ERD, when processing sensory data or planning or executing movement. ERD occurs during both actual and imagined movement (Pfurtscheller and Lopes da Silva 1999). Primary motor and primary sensory cortex are somatotopically organized into what is known as the homunculus, as illustrated in figure 1.6. Consequently, actual and imagined movement of specific body parts produces ERD at the corresponding location on the brain (Ehrsson et al 2003, Yuan et al 2008).

One specific type of BCI uses motor imagination to generate sensorimotor rhythms. Imagination of hand motion is particularly common (Wolpaw et al 2002, Wolpaw and McFarland 2004, Yuan et al 2008). In the international 10-20 EEG electrode montage, C3 and C4 lie over the central sulcus between primary motor and sensory cortex at the hand area of the homunculus. C3 is over the left hemisphere of the brain, C4 the right. As

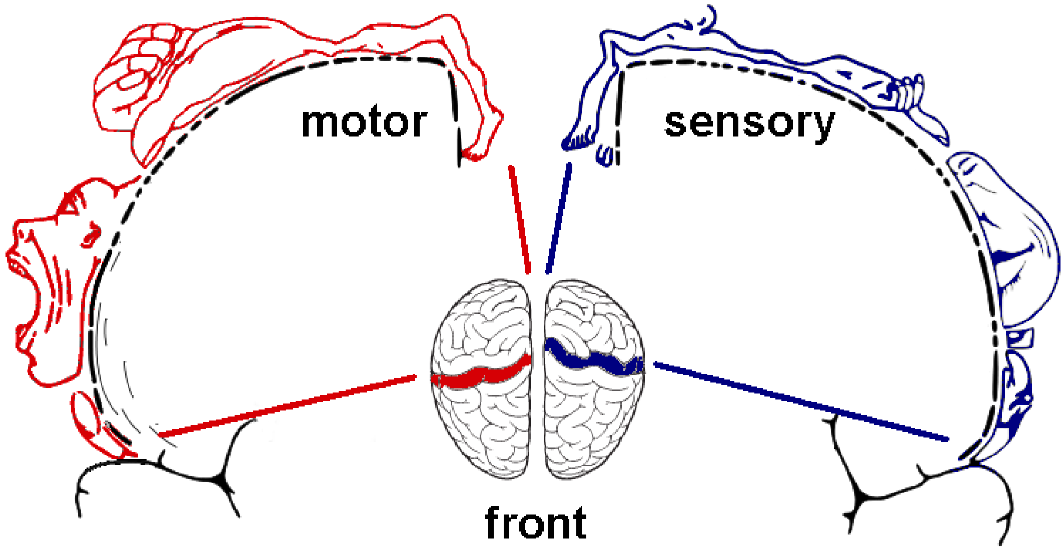

Figure 1.6 Organization of primary motor and sensory cortex
adapted from:
http://www.bci2000.org/wiki/index.php/User_Tutorial:Introduction_to_the_Mu_Rhythm

illustrated in figure 1.7, when a subject imagines moving the right or left hand, the sensorimotor rhythms respond with contralateral event related desynchronization (ERD) and ipsilateral event related synchronization (ERS) (Pfurtscheller and Lopes da Silva 1999, Yuan et al 2008). As the speed of the imagination increases, the ERD and ERS increase in amplitude by equal amounts on both sides of the brain (Yuan et al 2010b) .

The classical method of computing the time course of ERD/ERS involves bandpass filtering all event-related trials, squaring the amplitude measurements to produce power measurements, and then averaging the power across all trials (Pfurtscheller and Aranibar 1979). Since those classic instructions clearly list averaging trials as one of the steps, averaging trials is the norm for EEG analysis (Neuper et al 1999, Wolpaw and McFarland 2004, Pfurtscheller et al 2006, Yuan et al 2008, Neuper et al 2009, McFarland et al 2010, Yuan et al 2010a and b), as one researcher recently bemoaned (Wu et al 2011). Even those studies that addressed duration of ERS or ERD during motor imagination averaged trials (Pfurtscheller et al 2008, Nam et al 2011). However, recent findings have shown that averaging can lead to misleading, and even invalid, conclusions (Golowasch et al 2002, Freyer et al 2009).

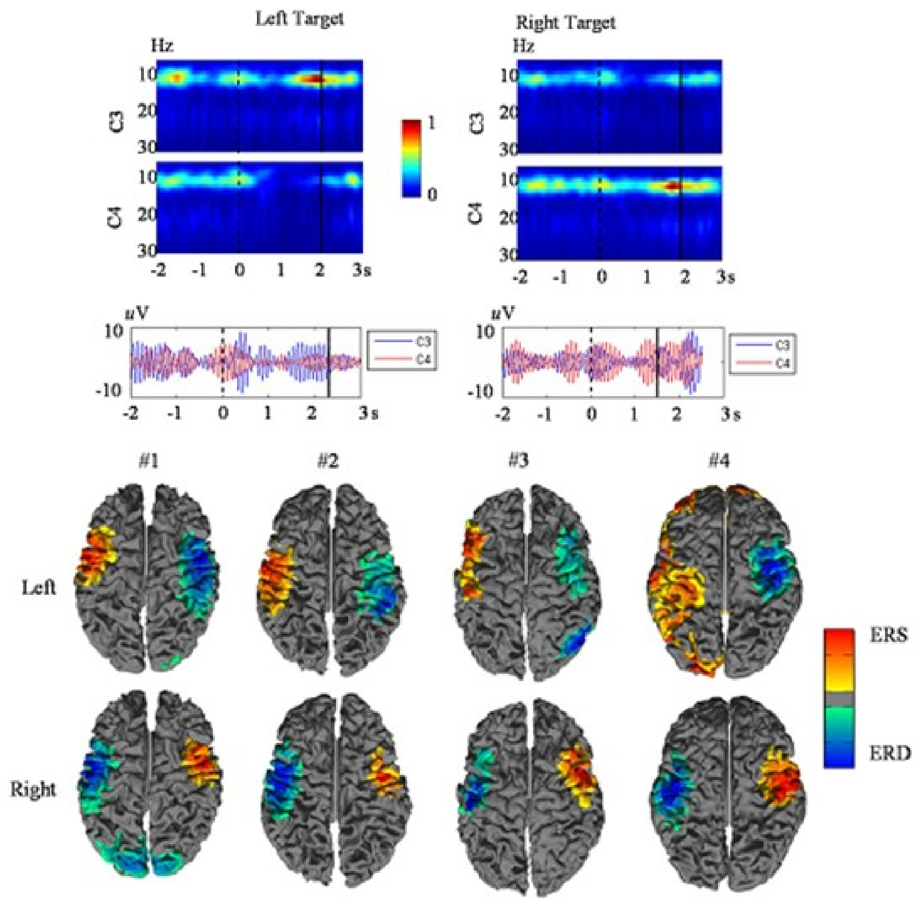

Figure 1.7 Sensorimotor rhythms displaying event related (de)syncronization
adapted from Yuan et al 2008

Characterizing the EEG signal within single trials seems an intuitive way to analyze data from a BCI given that BCIs process the EEG signal and provide feedback to the user in time steps as short as 40 to 80 ms. The subject's experience with the BCI is determined by the moment to moment operation of their brain, and not the average of multiple trials. However, averaging is the norm for EEG analysis (Wu et al 2011).

The prevalence of averaging BCI data in offline analysis, despite the non-relevance of averaging to online BCI performance, may be a contributing factor to the number of subjects that find it difficult to adequately operate a BCI. Work addressing this issue has used questionnaires and fMRI to probe the vividness of motor imagery of different subjects (Guillot et al 2008, Halder et al 2011). Analysis of the EEG signal has been used to predict acceptable BCI performance (Blankertz et al 2010). However, no study has performed analysis of the EEG signal at the same timescale as the BCI in order to address what distinguishes acceptable from unacceptable performance.

The ERD/ERS literature presently consists of mainly ERD/ERS data from either trained or naive subjects (Neuper et al 1999, Wolpaw and McFarland 2004, Pfurtscheller et al 2006, Yuan et al 2008, Neuper et al 2009, McFarland et al 2010, Yuan et al 2010b). An unanswered question is how do the sensorimotor rhythms of ERD and ERS change as a subject progresses from naive to trained?

While learning to operate a motor imagery based BCI, subjects learn to volitionally control their sensorimotor rhythms. The circuitry that produces the sensorimotor rhythms is the same circuitry as that used in perception, motor control, and even more esoteric items such as emotion and epilepsy. By understanding how humans purposely modulate this shared circuitry in the particular application of BCI, perhaps that can shed light on how people can better control their motor behavior, decision making, emotions, and even health issues such as epilepsy.

1.4 Goal Selection and Process Control

A brain-computer interface (BCI) translates signals recorded directly from the brain into commands that control an external device, such as a computer cursor, wheelchair, or neuroprosthetic (Wolpaw et al 2002, Vallabhaneni et al 2005). BCIs differ widely in how they implement the translation from raw brain signal to device command (Wolpaw et al 2002, Wolpaw 2007). One way they differ is in the overall control strategy.

Consider an individual with a right leg neuroprosthetic. When a healthy individual goes for a walk, the person's primary motor cortex issues the command to walk, and then the actual movement is coordinated by central pattern generators in the spinal cord (Grillner and Zangger 1979, Duysensa and A Van de Crommertb 1998). If an individual with a leg neuroprosthetic were to follow a similar procedure, their motor cortex would issue the command to walk and the BCI would recognize that command. Like a healthy individual's spinal cord, the BCI system would then control the prosthetic to generate a coordinated walking motion. That is an example of the BCI using a control strategy called goal selection (Wolpaw 2007). While using goal selection, the user merely needs to convey the goal to the BCI and then the user receives assistance from the BCI to execute that goal.

A second control strategy, called process control (Wolpaw 2007), would be more appropriate for a different scenario, such as when learning a new dance. Imagine that the individual with the neuroprosthetic is doing the Hokey Pokey for the first time. That person is listening carefully to the song's instructions: "You put your right foot in, You put your right foot out; You put your right foot in, And you shake it all about." In that instance, the individual would want to control the actual movement of the neuroprosthetic without additional BCI generated movement commands. At that time, the BCI would be using a control strategy called process control. While using process control, the user controls the entire process with no assistance provided.

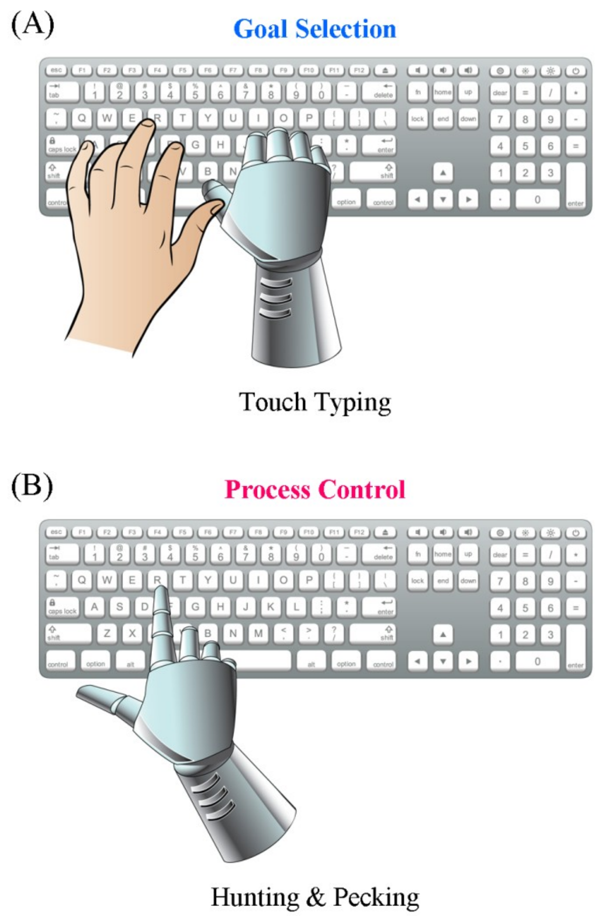

Figure 1.8 An illustration of goal selection vs. process control while typing
adapted from Royer et al, in press

Let us now consider an individual with a right hand neuroprosthetic (fig. 1.8). While typing, a person can either touch type, leaving their hands at home row and using all 10 fingers on the keys, or "hunt and peck", using just one or two digits for the whole keyboard. While touch typing, a healthy typist thinks "N", and through experience gained via training, their fingers automatically move to the N key and push it with no additional conscious thought. If the individual with a right hand neuroprosthetic were to follow a similar procedure (fig. 1.8A), their motor cortex would issue the command "N", and the BCI would recognize that command. Like a healthy individual's trained fingers, the BCI system would control the prosthetic to extend the correct finger (right index finger) to the correct key (N) and push it. That is another example of the BCI using goal selection.

Process control would be more appropriate while "hunting and pecking". When "hunting and pecking", the typist consciously thinks about each of the different component parts of the action to type a single character. First, the typist thinks "R"; then they locate the R key on the keyboard; next they move their finger to the key, and finally they push R. The typist is consciously controlling each part of the action. If the individual with a right hand neuroprosthetic were to "hunt and peck" (fig. 1.8B), the BCI using process control would receive all action commands from the individual's cortex. The BCI would not provide any additional assistance or action commands. The individual would consciously control the BCI the entire time of the keystroke.

In process control, the user controls every step of the process and receives minimal to no assistance from the system. In goal selection, the user only needs to determine the goal and the system executes the process to achieve that goal. In goal selection, the system performs the work that was asked of the user in process control. Since in goal selection, less work is asked of the user, goal selection is intrinsically easier than process control.

The above two examples are somewhat futuristic, since neuroprosthetics such as the ones described do not yet exist. Current BCIs also utilize the two control strategies of process control and goal selection. Significant advancements have been made by invasive BCIs using both control strategies. Information transfer rates of up to 6.5 bits per second have been achieved using goal selection (Santhanam et al 2006). Embodied control of a prosthetic arm was achieved using process control (Velliste et al 2008). As well as invasive BCIs, non-invasive BCIs have met success using both control strategies. The non-invasive P300 systems are intrinsically goal selection based (Farwell and Donchin 1988, Donchin et al 2000). This methodology has enabled an ALS patient who could no longer use conventional assistive devices to communicate and resume professional and social activities (Sellers et al 2010). Process control was used by a non-invasive system to move a computer cursor in both 2- and 3- dimensions (Wolpaw and McFarland 2004, McFarland et al 2010). Two studies implemented goal selection in a sensorimotor rhythm based BCI (McFarland et al 2008, Friedrich et al 2008). They met with modest success.

Goal selection could also be called shared control, since the user and the BCI system share the control process. Shared control has been previously used in BCI research to control robots and wheelchairs (Vanacker et al 2007, Bell et al 2008, Galan et al 2008). In one study, subjects drove a real wheelchair around obstacles (Galan et al 2008). The intelligent wheelchair used environmental sensors and shared control to ensure obstacle avoidance and safe driving. This study demonstrated that shared control can be used by a BCI in the real world to improve performance and safety.

During shared control, there may be time periods where the BCI is controlling the system's operation and the user's input is being ignored. An opposing control strategy to shared control is continuous control, where the user has control of the BCIs operation at all times. Just as goal selection can be called shared control, process control can be called continuous control.

The terms goal selection and process control may not be the best terms to use. One problem with the goal selection vs. process control terminology is the very definition of a goal. Many goals can be accomplished by breaking the goal down into the tasks necessary to complete that goal. Those tasks can then in turn be viewed as goals. This process of breaking the goal down into its necessary tasks can then be repeated. This then complicates the definition of goal selection versus process control since using process control could also be viewed as selecting a series of smaller goals.

Hierarchical control terms from control theory may be more descriptive and less ambiguous. Goal selection, where the user selects the goal and the BCI then executes that goal, could be referred to as high-level control. Process control, where the user controls the entire process without assistance, could be referred to as low-level control.

We have now defined three sets of terms for the two control strategies studied in this thesis: goal selection = shared control = high-level control and process control = continuous control = low-level control. The terms "goal selection" and "process control" will be used throughout this thesis to remain consistent with the previously published works that compose this thesis. It may be beneficial moving forward to use alternate terms other than goal selection and process control. The hierarchical control terms may be preferred since a hierarchy intuitively has multiple levels, and a middle ground.

Process control, when used either by an invasive or non-invasive BCI, produces motion that would be classified as ataxic by neuromuscular control specialists (Wolpaw 2007). As shown in figure 1.9, the time required to reach a target in a simple 2D center-out cursor movement task is slower and much less consistent when using a BCI than when using a joystick (Wolpaw and McFarland 2004, Hochberg et al 2006, McFarland et al 2010). A substantial number of BCI trials did not even reach the target in the 7 seconds allowed. Such inconsistent performance is typical of process control BCIs.

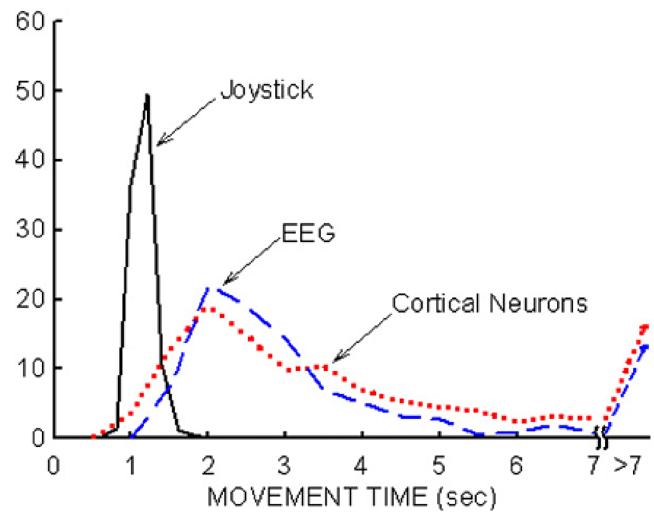

Figure 1.9 Distributions of target-acquisition times on a 2D center-out cursor movement task adapted from McFarland et al 2010

Many BCIs perform a task that otherwise would be performed through motor output. A possible reason for the ataxia highlighted by figure 1.9 may be that BCIs gather their control signal from the cortex, bypassing all the complicated, trained interactions that produce normal motion (Wolpaw 2007). Process control BCIs call for the primary motor cortex to control all the fine motor details normally handled by other parts of the motor nervous system, such as the basal ganglia, cerebellum, or spinal cord. The user must ensure that they are properly encoding position, velocity, and/or acceleration to hit the desired target while using a process control BCI.

In goal selection, the BCI uses the signal it obtains primarily from the cortex to determine the overall end goal of the user, here the selection of the desired target. This is more typical of the role cortex plays in normal motor control. The BCI execution unit then determines the necessary position, velocity, and/or acceleration parameters of the cursor to hit the desired target. Goal selection more closely resembles natural motor control with the BCI system assisting the user akin to how the distributed motor network assists the motor cortex (Wolpaw 2007). Since goal selection is easier and more natural, it follows that it would be more accurate, faster in use, and easier to learn. This system would have a higher information transfer rate, with a decreased training period.

In our daily lives, able bodied individuals receive much assistance from the systems we interact with. Anti-lock braking systems stop cars faster and safer than the driver can do by pumping the brakes himself. Spell-check and grammar-check have improved the quality of the written word. Word completion software makes texting faster. Point-and-shoot cameras dominate the camera market. Consumers have come to expect their electronics to offer them assistance and make the use of the electronic device easier.

Those individuals who have lost normal motor control require a BCI that is effective, easy to use, and easy to learn. As discussed in Bai et al (2010), the minimization of fatigue during BCI use will be important as BCIs move from laboratory to clinical settings. The patient populations that many BCIs are designed to serve, such as those with ALS, have reduced physical and mental endurance (Sykacek et al 2003, Birbaumer 2006). This diminished endurance has decreased the accuracy of a BCI system with 90% accuracy in healthy subjects to levels just over chance in the patient population (Sellers and Donchin 2006, Iversen et al 2008). The proposed advantages of goal selection may help to reduce the strain of using a BCI. Therefore, the reduction of fatigue due to using goal selection as a control strategy may aid the usefulness and adoption of BCIs by individuals who truly need them to restore lost functionality.

1.5 The Organization of this Thesis

This thesis presents research analyzing goal selection as a control strategy in a BCI. Although several goal selection based BCIs exist, we present the first studies directly comparing goal selection and process control. In the analysis, we look at both behavioral performance metrics as well as analyze the underlying differences in the EEG signal. This thesis is structured as a collection of papers that were published (or submitted) on this topic. Each paper is presented in its entirety, as originally published (or submitted). The only changes made were to insert the appropriate section number when referring to other parts of this thesis and to update the figure numbers for sequential numbering throughout a chapter. Two chapters have additional analysis in the last section that was not in the published (or submitted) work.

In chapter 2, we present 2 studies that directly compare goal selection and process control in a relatively simple left/right cursor task. Both studies use the same BCI paradigms. Section 2.1 presents a paper that tests the hypothesis in both naive and trained users that a goal selection BCI should be more accurate and faster to use, which together lead to a higher information transfer rate. Section 2.2 presents a paper that tracks naive subjects through the learning process. In addition to the speed and accuracy measures presented in the previous study, this much larger subject pool allowed the comparison of the ease of learning and mental effort required between goal selection and process control.

In chapter 3, goal selection is applied to a much more complicated, real world scenario. In this paper, subjects are asked to fly a virtual helicopter to any point in three-dimensional space. The task was made simpler by using intelligent control strategies, such as goal selection. This chapter presents the first accomplishment by a non-invasive BCI of a task previously performed only by invasive BCIs. This was made possible by using intelligent control strategies. The last section presents additional, unpublished analysis on the influence goal selection had on the study.

In chapter 4, we address the underlying neural signal while using both goal selection and process control. This chapter analyzes the EEG data from section 2.2, but is the first study to analyze the EEG from a BCI at the same timescale as the BCI, and without averaging trials. It also presents the longest running EEG analysis tracking subjects from naive to trained. The submitted paper presents the results of analyzing the electrodes and frequencies that were used for control. The last section presents additional, unsubmitted analysis on the electrodes and frequencies that were not used for control.

Chapter 2 Goal Selection vs. Process Control in a Simple Task

This chapter presents 2 studies that directly compare goal selection and process control in a relatively simple left/right cursor task. Both studies use the same BCI paradigms. Section 2.1 presents a paper that tests the hypothesis in both naive and trained users that a goal selection BCI should be more accurate and faster to use, which together lead to a higher information transfer rate. Section 2.2 presents a paper that tracks naive subjects through the learning process. In addition to the speed and accuracy measures presented in the previous study, this much larger subject pool allowed the comparison of the ease of learning and mental effort required between goal selection and process control.

2.1 Trained and Naive Subjects

The material in section 2.1 has been reprinted with permission from Royer and He 2009.

2.1.1 Introduction

There are several diseases and conditions, which lead to loss of muscular control. Those include amyotrophic lateral sclerosis (ALS), brainstem stroke, spinal cord injury, muscular dystrophies, and cerebral palsy. Although people living with these conditions suffer major muscular loss, their cognitive abilities are left intact, thus they still want to communicate and manipulate their environment (Kunst 2004). To do this, these patients require an output pathway from their brain that does not rely on muscular output, or a brain-computer interface (BCI) (Wolpaw et al 2002, Vallabhaneni et al 2005).

A brain-computer interface can be either invasive or non-invasive. One of the most promising types of non-invasive BCI utilizes scalp-recorded electroencephalograms (EEG) to monitor sensorimotor rhythms (SMRs) (Wolpaw et al 2002, Vallabhaneni et al 2005). Sensorimotor rhythms are produced by the primary sensory and motor cortices. SMR based BCI's utilize two distinct states: event related synchronization (ERS) and event related desynchronization (ERD). When an awake person is not processing sensory data or producing motor output, the primary sensory and motor cortices are in an idling state, which creates a rhythmic EEG pattern known as ERS. The mu rhythm from 8 to 12Hz and the beta rhythm from 13Hz to 26Hz have been particularly useful as BCI control signals. The sensorimotor rhythms decrease in amplitude, an effect that is known as ERD, when processing sensory data or planning or executing movement. ERD occurs during both actual and imagined movement (Pfurtscheller and Lopes da Silva 1999).

The process of planning and executing movement in a healthy individual is a multistep process that involves many parts of the central nervous system. The primary motor cortex initiates this process, but other structures such as the basal ganglia, cerebellum, thalamus, brainstem nuclei, spinal interneurons, and spinal motor neurons are vital to the proper functioning of the process. Since BCIs gather their control signal from the cortex, a BCI bypasses all the complicated, trained interactions that produce normal motion (Wolpaw 2007).

The best BCI systems, both invasive and non-invasive, produce motion that would be classified as ataxic by neuromuscular control specialists (Wolpaw 2007). A possible reason for this ataxia may be that most BCIs in existence today call for the primary motor cortex to control all the fine motor details normally handled by other parts of the brain. These BCIs use the signal obtained primarily from the cortex to determine the position, velocity, and/or acceleration of the controlled device, here a cursor. The user must ensure that they are properly encoding position, velocity, and/or acceleration to hit the desired target. This is known as process control. Process control is not the only control strategy used in BCIs. An alternative control strategy is goal selection. In goal selection, the BCI uses the signal it obtains primarily from the cortex to determine the overall end goal of the user, here the selection of the desired target. The BCI execution unit then determines the necessary position, velocity, and/or acceleration parameters of the cursor to hit the desired target. The user must only encode the desired action, not the details necessary to achieve that action (Wolpaw 2007).

To date, the majority of BCIs employ process control. Some of the few exceptions which employ goal selection include the non-invasive P300 based BCIs (e.g. Farwell and Donchin 1988) and a few invasive studies (e.g. Musallam et al 2004). Lately, goal selection has been applied to SMR based BCIs. Two studies (McFarland et al 2008, Friedrich et al 2008), implemented goal selection in a sensorimotor rhythm based BCI. They met with modest success.

Since goal selection more closely resembles the normal process for motor execution, it follows that a BCI based on goal selection would be an easier and more natural system than one based on process control. An easier and more natural system would be more accurate, faster in use, and easier to learn. This system would have a higher information transfer rate, with a decreased training period. Although several goal selection based BCIs exist, we present the first study directly comparing goal selection and process control. This study will test the first two ideas: that a goal selection BCI should be more accurate and faster to use, which together lead to a higher information transfer rate.

2.1.2 Methods

2.1.2.1 Data acquisition and cursor control

The human study was conducted according to a human protocol approved by the Institutional Review Board (IRB) of the University of Minnesota. Eight healthy, young volunteers, 1 female and 7 male, participated in the 1-dimensional BCI study. They were seated facing a computer monitor while wearing a 64 channel EEG cap set up according to the 10-20 international system. The scalp-recorded EEG signal passed to a Neuroscan amplifier, and was sampled at 1000Hz. The BCI2000 system (Schalk et al 2004) was used to conduct the online experiments with visual feedback. The control signal was the difference between the auto-regressive (AR) spectral amplitudes from 7.5 to 13.5Hz of electrodes C4 and C3 (Wolpaw and McFarland 2004). The magnitude of cursor movement was determined by the normalized AR amplitude difference. At the end of each trial, the control signal was normalized so that it had a zero mean and unit variance across a multiple trial buffer. The parameters used in the normalization, namely the normalizer offset and normalizer gain, were recorded at the end of each session for use at the beginning of the subject's next session. This is the adaptation built into the BCI2000 system, version 2.0 (Schalk et al 2004). Since the BCI used sensorimotor rhythm as the control signal, the subjects were encouraged to use motor imagination, such as imagining squeezing their right hand to move the cursor right and imagining squeezing their left hand to move the cursor left (Wang & He 2004, Qin et al 2005, Pfurtscheller et al 2006, Kamousi et al 2007). Imaginations were not dictated to the subjects. They were free to imagine whatever worked best for them. Several props, such as squeeze balls and dumb bells, were available during breaks to aid in the imagination.

2.1.2.2 Study design

The 8 subjects fell into one of two groups. The first group received BCI training, and consisted of three subjects. The second group was naïve to BCI usage, and included five subjects. The trained subjects used a BCI approximately once per week for six to eight weeks. The naïve group, who had never used a BCI prior to this study, completed 2 sessions on different days. The trained subjects completed either one or two sessions. During each session, subjects completed 3 runs of five different paradigms. Each run, regardless of paradigm, was four minutes long, and consisted of as many trials as the subject could complete in four minutes. Although the exact details of a trial varied based on paradigm, each trial presented the targets to the user, and the user attempted to select the yellow target. Between each trial, the subject had three seconds of rest. Each paradigm utilized the same control signal as the other paradigms, so acquired skill transferred easily from one paradigm to the next. The order of the paradigms was reversed between the first and second session.

2.12.3 Experimental paradigms

Each session consisted of five similar, yet distinct, paradigms. Figure 1 outlines the flow of each paradigm. Each started with the targets appearing on the screen at time 0. The subject was instructed to hit the yellow target. One second later, the cursor appeared. For all the paradigms except goal selection (GS), the cursor then moved under cortical control.

In GS (figure 1A), the subject did not have the visual feedback of the cursor movement. Instead, the cursor moved invisibly according to the same rules as the other paradigms, but a fixation point identical to the cursor stayed in the middle of the screen. At t = 2s, whichever target the invisible cursor was closest to became selected and turned blue. After selection, the invisible cursor returned to the middle of the screen and began moving again under cortical control. At t=3s, one of two things happened. The first option was that the invisible cursor was closest to the selected target so that the target was reconfirmed, turned purple, and the visible cursor automatically went to that target. The other option was that the invisible cursor was closer to the other target, which became selected and turned blue. In this case, a third round of selection began, and final target selection was determined by being selected two out of three times.

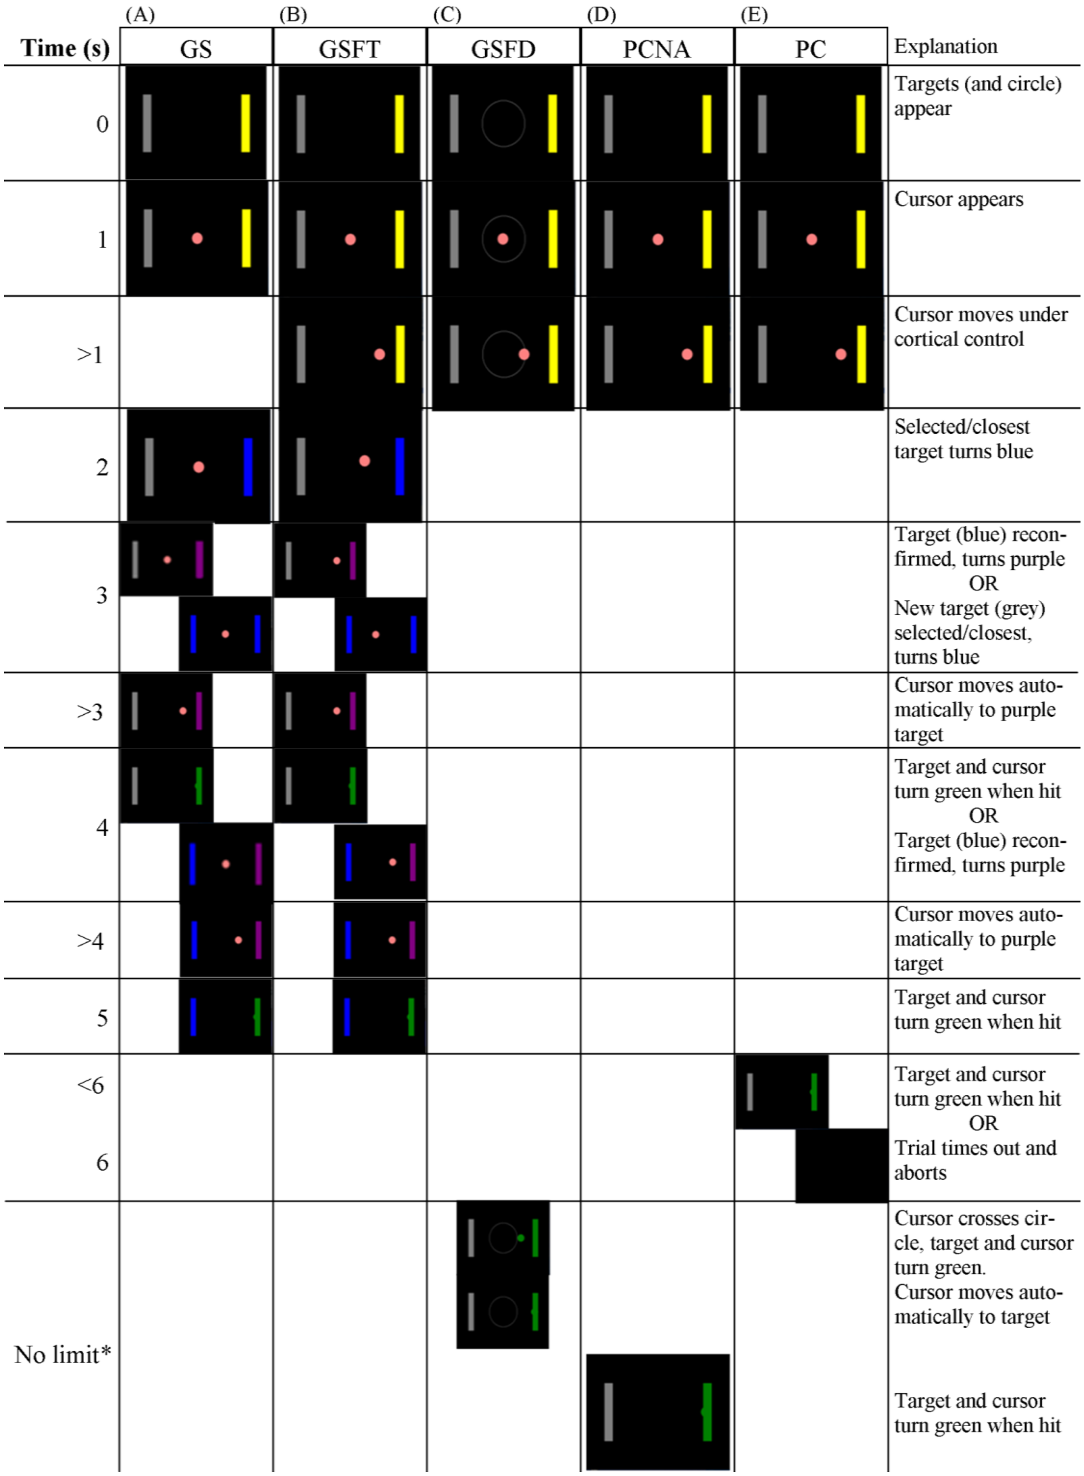

Figure 2.1 Experimental paradigms illustrated for a hit

Experimental paradigms illustrated for a hit. The yellow target is the intended target. Times and explanations apply to paradigms with an image in that row. If the end result of a trial is a miss, the target and cursor will turn red instead of green. (A) Goal selection (GS) and (B) goal selection with feedback limited by time (GSFT). Once a target is selected (blue), it must be reconfirmed (purple) before the cursor will automatically move to that target. Starting at second 3 there are two possible paths. Either the target that was originally selected is reconfirmed (left path), or a new target is selected (right path). If a new target is selected, the cursor will move to the target that was selected two out of three times. Target selection (blue) at seconds 2 and 3, and target reconfirmation (purple) at seconds 3 and 4 may occur slightly earlier if the cortical control signal is especially strong. The times given are the typical, and maximal, times. (C) Goal selection with feedback limited by distance (GSFD) and (D) process control with no aborts (PCNA). No limit* had a limit of 60s. Only one subject ever reached this limit to experience an abort. (E) Process Control (PC). If the subject did not hit a target in under 6s (left picture), the trial timed out and aborted at 6s (right picture).

The GS paradigm involved minimal feedback. Both Hochberg et al (2006) and Hinterberger et al (2005) posited that feedback improves BCI performance. Data to support the role of online feedback in improving performance comes from numerous sources (e.g. Neuper et al 1999, Brunner et al 2006). Without sufficient feedback, goal selection may not live up to its potential. In order to increase the amount of feedback in GS, two other variations of goal selection paradigms were developed. Goal selection with feedback limited by time (GSFT) was almost identical to GS, except that it displayed the movement of the cursor. One can see the similarities in figure 1B. One difference between GS and GSFT was that the cursor did not return to the centre after selection in GSFT.

The other variation of a goal selection paradigm with feedback is illustrated in figure 1C. Goal selection with feedback limited by distance (GSFD) was more similar to the remaining 2 paradigms than it was to GS or GSFT. In GSFD, there was a grey circle in the centre of the screen. This circle was visible whenever the targets were visible. The radius of the circle was set at 20% of the screen. GSFD showed the movement of the cursor. When the cursor crossed the circle, it automatically moved to the closest target. In order to allow the subjects to feel like there was no time limit on the trial, and to eliminate trials that timed out and aborted, the maximum trial time was set to 60s. All trials of all subjects were completed in the 60s.

The next experimental paradigm was process control with no aborts (PCNA), illustrated in figure 1D. PCNA was a typical cursor task used in BCI studies (Wolpaw and McFarland 2004, Shenoy et al 2006, Krusienski et al 2007, Yuan et al 2007, Yuan et al 2008, Blankertz et al 2008), where the user controlled the movement of the cursor to hit a target. Similar to GSFD, PCNA was intended to not allow the subjects to abort a trial through timing out. Therefore, the maximum trial time was also set to 60s. However, one subject did have two trials that extended to 60s and thus aborted. Those were exceptional trials, and all the other subjects were able to complete all trials within 60s.

The last experimental paradigm was process control (PC), illustrated in figure 1E. As in PCNA, the user controlled the movement of the cursor until the cursor hit a target. The only difference between PC and PCNA was that a PC trial only allowed the subject 6s to hit a target. If no target had been hit after 6s, the trial aborted.

The five experimental paradigms, GS, GSFT, GSFD, PCNA, and PC, can be divided into two groups. The first group consists of the paradigms based on goal selection: GS, GSFT, and GSFD. The second group consists of the paradigms based on process control: PCNA and PC. To limit confusion, the paradigms will henceforth be referred to by their acronym, whereas the spelled out words goal selection and process control will be reserved for the groups of paradigms and the concepts the words represent.

Table 2.1 Paradigm program parameters

GS and GSFT	*Value (s)*	*PC, PCNA, and GSFD*	*Value (s)*
PreSelectionDuration	1	PreFeedbackDuration	1
GoalSelectionDuration	1	FeedbackDuration	2
ReactionTime (GS)	0.25	MaxFeedbackDuration (PC)	6
ReactionTime (GSFT)	0	MaxFeedbackDuration (PCNA & GSFD)	60
GoalReconfirmDuration	1		
MovementDuration	0.5	PostFeedbackDuration	1
ResultDuration	0.5		
BufferLength	30	Buffer Length	30
SampleBlockSize	0.04	SampleBlockSize	0.04
ITI Duration	3	ITI Duration	3
MinRunLength	240	MinRunLength	240

In order to allow a valid comparison of subject performance across all five experimental paradigms, the inner workings and programming of each paradigm were made consistent. Key paradigm program parameters are summarized in table 1. The parameter FeedbackDuration in PC, PCNA, and GSFD was used to set the cursor movement speed. Note that it is the sum of the GoalSelectionDuration and

GoalReconfirmDuration for GS and GSFT. Cursor speed was carefully set to be the same across all paradigms. PostFeedbackDuration in PC, PCNA and GSFD was split into MovementDuration and ResultDuration for GS and GSFT. The BufferLength is the amount of time that was used to normalize the control signal. The SampleBlockSize indicates how often the cursor position was updated.

The naïve subjects started with PC followed by PCNA, GSFD, GSFT, and GS. The order was reversed for their second session. The trained subjects used the same order as the naïve subjects and reversed the order each session, resulting in alternating sessions having the same order.

2.1.2.4 Data analysis

Several aspects of subject performance were analyzed in order to compare the different paradigms. The features analyzed included the average number of hits per run, time distribution to a hit, overall accuracy, and information transfer rate in bits per minute. The paradigms were compared both within each subject and across all the data pooled from each group.

The time required for a subject to hit the target was determined from the time during which the cursor was under cortical control. For all paradigms, the time started when the cursor appeared. In GS and GSFT, time ended when a target was reconfirmed and turned purple. In GSFD, time ended when the cursor crossed the circle and turned green. In all three of those cases, the additional time required to actually hit the target was a user settable programmed parameter that was constant for each trial. Therefore, that time was not held against the user. In PCNA and PC, time ended when a target was hit and both the target and cursor turned green. As discussed above, the cursor moved under the same control signal and at the same speed for all five paradigms. Any remaining differences in time to task completion were due to the facets of the paradigms we wished to compare. Because of this, the time required for a subject to hit the target as determined from the time during which the cursor was under cortical control was a valid comparison and measure of the different paradigms.

For the purpose of data analysis, accuracy was defined as the number of hits in a run divided by the total number of trials in a run. For PC runs, aborted trials were counted in the total number of trials. Accuracy could also be viewed as the percentage of trials that ended in a hit. This form of calculating accuracy effectively normalized the data for the different number of trials in each run, which allowed a fair comparison of the different paradigms.

A useful way to compare different BCIs is via their information transfer rate, either in bits/trial or bits/min. As given by Wolpaw et al. (2002), bits/trial can be calculated from the following equation:

$$\log_2 N + P \log_2 P + (1-P) \log_2[(1-P)/(N-1)] \qquad (1)$$

In the above equation, N is the number of targets, and P is the probability of a hit, or the accuracy. The number of trials per minute can then be multiplied by (1) to obtain the information transfer rate in bits/min. One useful feature of this measure is that it incorporates both speed and accuracy into one number. The information transfer rate in bits per minute was calculated for each trial. The average of those trials was then used for plotting.

Significance testing utilized one of two statistical tests. A pair wise t-test was used for measures that had an underlying normal distribution, such as the number of hits per run and the information transfer rate in bits per minute. The standard deviation was pooled between paradigms, and no p-value correction was applied. The chi-squared pair wise proportion test was used for significance testing the accuracies between paradigms since that measure had an underlying binomial distribution. No p-value correction was applied to the chi-squared test. Asterisks on the plots indicate pair-wise significance to the previous paradigms. Tables providing more information on the significant differences follow any figure with asterisks. Error bars on the plots indicate standard error.

The box plots show the distribution of the data. The lower whisker extends from the minimum value to the 25th percentile. The lower, darker box extends from the 25th percentile to the median. The upper, lighter box extends from the median to the 75th percentile. The upper whisker extends from the 75th percentile to the maximum value.

2.1.3 Results

2.1.3.1 Number of hits in an average run

The average run for each subject and each paradigm is shown in figure 2. Significant differences in the number of hits are shown in table 2. Across the board, the goal selection paradigms have more hits in an average run than the process control paradigms. Figure 2A shows the data for the trained subjects. It is interesting to note that the trained subjects still experienced a good number of aborted trials in PC even though the number of hits stayed relatively consistent with or without aborts. Both GSFT and GSFD had significantly more hits than PCNA for all three subjects. Additionally, GS had significantly more hits than PCNA for two out of the three subjects. Significance could not be calculated for subject

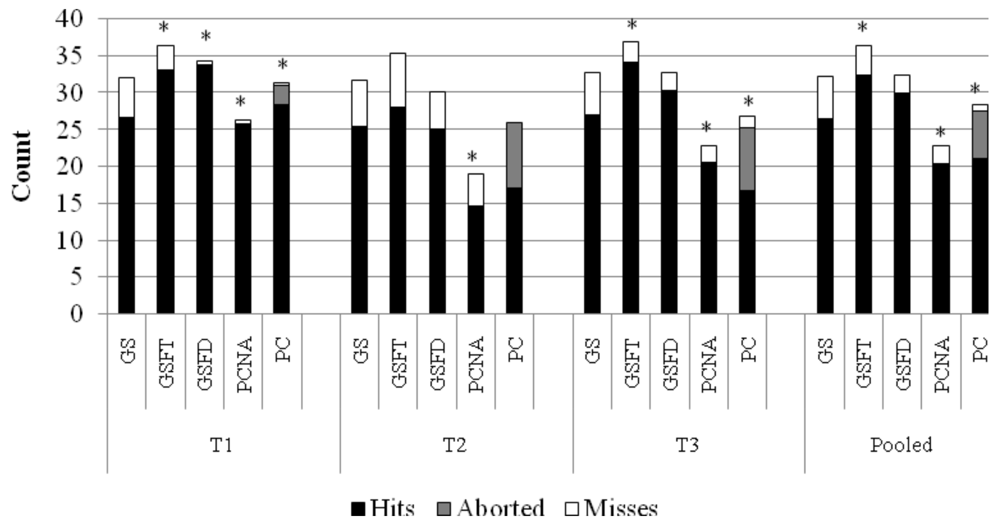

(A)

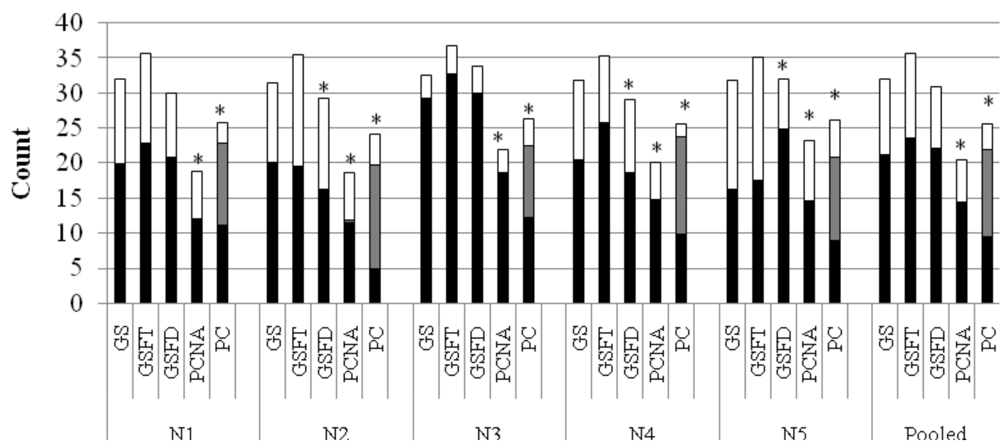

(B)

Figure 2.2 Average run breakdown for both trained and naive subjects
Overall, the goal selection paradigms have significantly more hits than the process control paradigms. Asterisks indicate pair-wise significance to the previous paradigms for the number of hits. (A) trained (B) naive.

T2's process control paradigm. Both of the remaining two subjects showed that PC had significantly fewer hits than GSFT and GSFD. In the pooled data, GSFT and GSFD had more hits than GS, with a significant difference between GS and GSFT. Importantly, all of the goal selection paradigms had significantly more hits than either of the process control paradigms. The increase in hits from process control to goal selection paradigms ranged from 25% to 59%.

Table 2.2 Significant differences for number of hits

	GS	*GSFT*	*GSFD*	*PCNA*	*PC*
T1	GSFT** GSFD**	GS** PCNA*** PC*	GS*** PCNA*** PC**	GSFT*** GSFD***	GSFT* GSFD**
T2	PCNA***	PCNA***	PCNA***	GS*** GSFT*** GSFD***	NA
T3	GSFT* PCNA* PC*	GS* PCNA*** PC***	PCNA*** PC***	GS* GSFT*** GSFD***	GS** GSFT*** GSFD***
Pooled	GSFT** PCNA** PC*	GS** PCNA*** PC***	PCNA*** PC***	GS** GSFT*** GSFD***	GS* GSFT*** GSFD***
N1	PCNA*** PC***	PCNA*** PC***	PCNA*** PC***	GS*** GSFT*** GSFD***	GS*** GSFT*** GSFD***
N2	GSFD* PCNA*** PC***	PCNA*** PC***	GS* PCNA** PC***	GS*** GSFT*** GSFD** PC***	GS*** GSFT*** GSFD*** PCNA***
N3	PCNA*** PC***	PCNA*** PC***	PCNA*** PC***	GS*** GSFT*** GSFD*** PC*	GS*** GSFT*** GSFD*** PCNA*
N4	PC**	GSFD* PCNA** PC***	GSFT* PC*	GSFT**	GS** GSFT*** GSFD*
N5	GSFD*** PC**	GSFD** PC***	GS*** GSFT** PCNA*** PC***	GSFD*** PC**	GS** GSFT*** GSFD*** PCNA**
Pooled	PCNA*** PC***	PCNA*** PC***	PCNA*** PC***	GS*** GSFT*** GSFD*** PC***	GS*** GSFT*** GSFD*** PCNA***

*p<0.05, **p<0.01, ***p<0.001

The data for the naïve subjects is displayed in figure 2B. As in the trained data, the goal selection paradigms resulted in more hits than the process control paradigms. The naïve subjects experienced many aborts during PC, with individual subjects aborting 39% to 61% of the trials. The pooled data shows 49% of the PC trials timed out. Not surprisingly then, PC always had significantly fewer hits than any form of goal selection. PCNA had significantly fewer hits than any form of goal selection for three of the five subjects, whereas PCNA had significantly fewer hits than GSFT or GSFD for the other two subjects. In the pooled data, all of the goal selection paradigms have significantly more hits than the process control paradigms. The increase in hits from PCNA to the goal selection paradigms ranged from 48% to 65%. Astoundingly, the increase in hits from PC to the goal selection paradigms ranged from 124% to 151%. The goal selection paradigms had more than twice the number of hits than PC.

2.1.3.2 Time to a hit

The distribution of time the cursor was under cortical control leading to a hit is displayed in figure 3. The median times are also summarized in table 3. The trained and the naïve subjects showed similar results. For all subjects, the goal selection paradigms were faster than the process control paradigms. Another commonality is that a GSFT hit occurred at approximately the same time, with very little spread both within and between subjects. This implies that a hit occurred when a target was first selected and then reconfirmed. A hit hardly ever resulted from a best two- out-of-three condition. GS, which had the same selection and reconfirmation timing as GSFT, had a larger spread. This means that more hits resulted from a best two-out-of three condition. Although GS and GSFT were inherently time constrained, even the non-time constrained GSFD was typically faster than both of the process control paradigms, as shown by a smaller 25th to 75th percentile distribution in the pooled data. Looking at the pooled data from each group provides further insight. The pooled trained data (figure 3A) showed that the median of all three goal selection paradigms was approximately twice as fast as the median hit time for PCNA and 30% to 60% faster than the median time for PC. The data from the naïve subjects (figure 3E) is qualitatively very similar to the data from the trained subjects. It is not surprising that the non-time constrained paradigms of GSFD, PCNA, and PC had slightly longer medians with a larger spread in times in the untrained subjects.

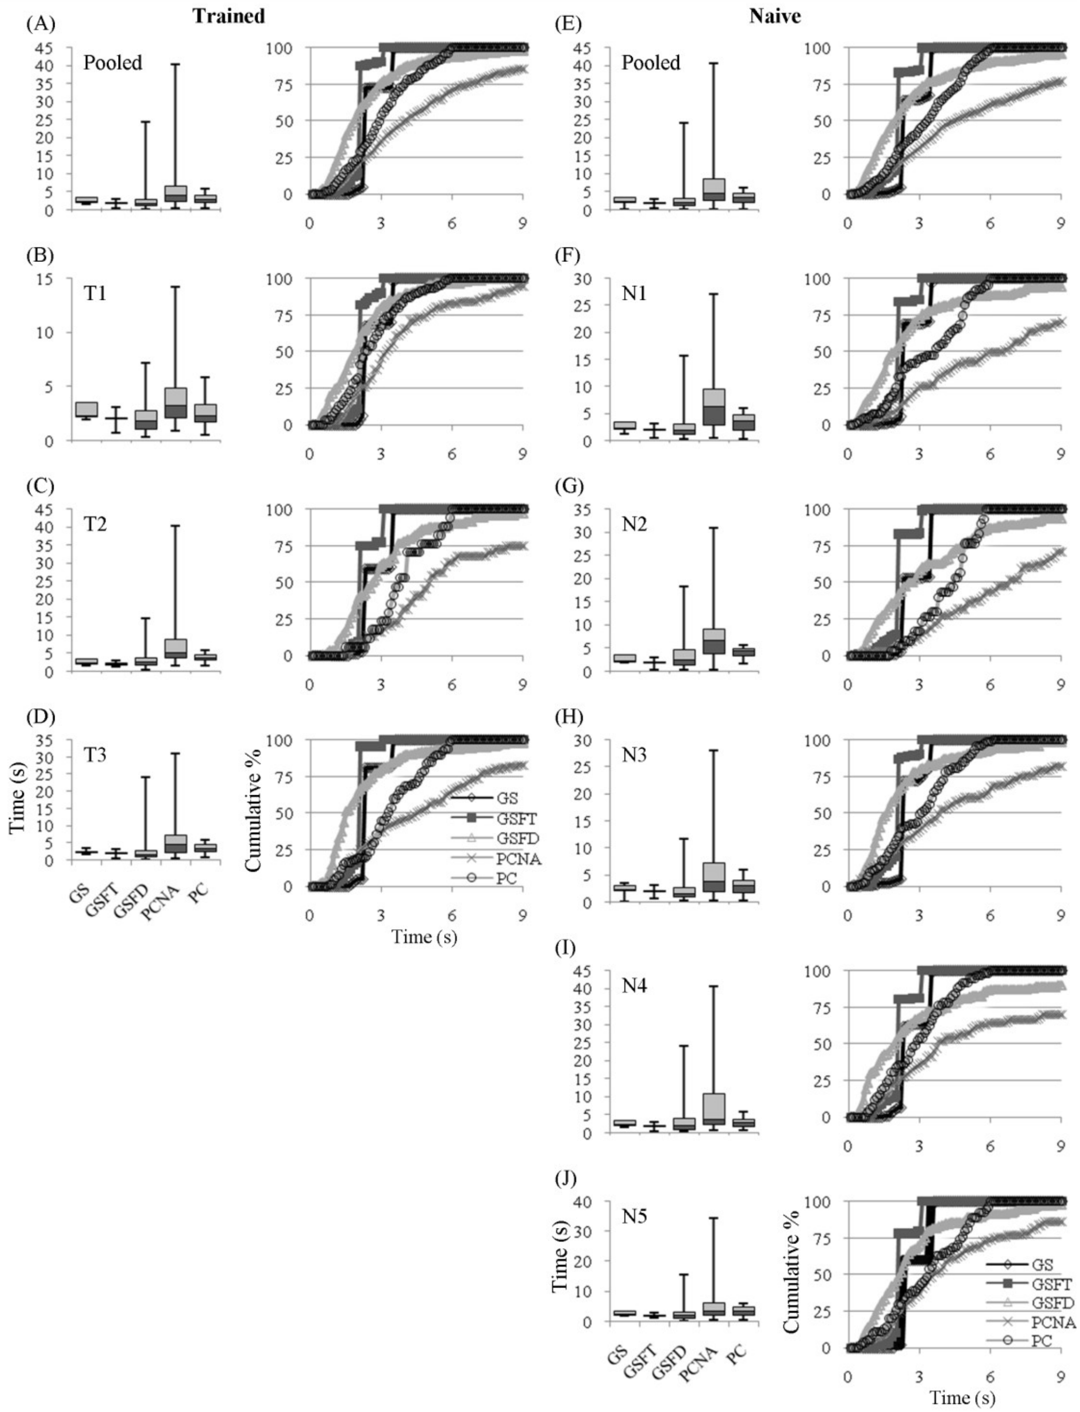

Figure 2.3 Time required for a hit

Time required for a hit across the five paradigms for trained ((A) through (D)) and naive ((E) through (J)) subjects. The goal selection paradigms are faster than the process control paradigms. Axes labels and graph legends on the bottom row of each column ((D) and (J)) apply to all figures in that column. The plots to the right of each box plot are the same data, shown as a cumulative distribution.

This led to all three of the goal selection paradigms being at least twice as fast as PCNA, and 48% to 72% faster than PC in the pooled data.

Table 2.3 Median time to a hit (s)

	GS	*GSFT*	*GSFD*	*PCNA*	*PC*
T1	2.25	2.05	1.83	3.19	2.28
T2	2.25	2.05	2.58	5.01	3.74
T3	2.25	2.05	1.48	4.46	3.28
Pooled	2.25	2.05	1.82	4.05	2.92
N1	2.25	2.05	1.82	6.23	3.65
N2	2.26	2.05	2.48	6.75	4.45
N3	2.25	2.05	1.59	3.81	3.06
N4	2.25	2.05	1.89	3.79	2.82
N5	2.25	2.05	2.12	3.58	3.33
Pooled	2.25	2.05	1.94	4.46	3.33

2.1.3.3 Accuracy

The accuracy of each of the subjects using each paradigm is given in figure 4, with the significant differences shown in table 4. For the trained subjects (figure 4A), GSFD had the highest accuracy. For subjects T1 and T3, GSFD had significantly higher accuracy than both GS and PC. T1 preferred feedback limited by distance over feedback limited by time, but T3's performance did not distinguish between the different forms of feedback. Both were significantly more accurate than GS and PC. In the pooled data, GSFT and GSFD were more accurate than GS or PC. GSFD was significantly more accurate than both GS and PC, whereas GSFT was significantly more accurate than PC.

In the naïve subjects (figure 4B), there was considerable variation from individual to individual. No one paradigm was consistently the most accurate. However, goal selection of some form was the most favored. Four out of the five subjects had the highest accuracy with a form of goal selection. Although one paradigm did not stand out as the best, one did establish itself as the worst. For all subjects, PC was significantly less accurate than all other paradigms.

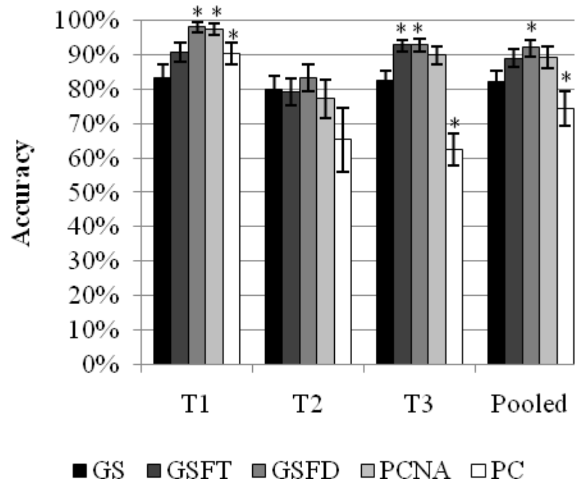

(A)

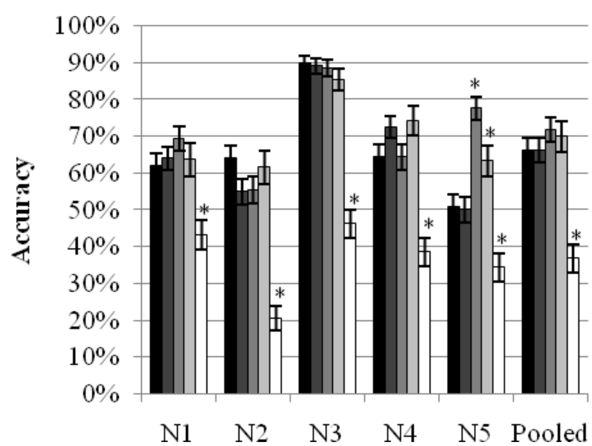

(B)

Figure 2.4 Accuracy across the five paradigms
Accuracy of the trained (A) and naïve (B) subjects across the five paradigms. GSFT or GSFD commonly had the highest accuracy, with PC having the worst accuracy.

Table 2.4 Significant differences for accuracy

	GS	*GSFT*	*GSFD*	*PCNA*	*PC*
T1	GSFD*** PCNA**	GSFD*	GS*** GSFT* PC*	GS**	GSFD*
T2					
T3	GSFT** GSFD** PC***	GS** PC***	GS** PC***	PC***	GS*** GSFT*** GSFD*** PCNA***
Pooled	GSFD*	PC**	GS* PC**	PC*	GSFT** GSFD** PCNA*
N1	PC***	PC***	PC***	PC**	GS*** GSFT*** GSFD*** PCNA**
N2	PC***	PC***	PC***	PC***	GS*** GSFT*** GSFD*** PCNA***
N3	PC***	PC***	PC***	PC***	GS*** GSFT*** GSFD*** PCNA***
N4	PC***	PC***	PC***	PC***	GS*** GSFT*** GSFD*** PCNA***
N5	GSFD*** PCNA* PC**	GSFD*** PCNA* PC**	GS*** GSFT*** PCNA** PC***	GS* GSFT* GSFD** PC***	GS** GSFT** GSFD*** PCNA***
Pooled	PC***	PC***	PC***	PC***	GS*** GSFT*** GSFD*** PCNA***

*p<0.05, **p<0.01, ***p<0.001

2.1.3.4 Information transfer rate

The information transfer rate as measured in bits per minute is a convenient way to compare different BCI systems. This one number incorporates both accuracy and speed. Figure 5 shows the information transfer rate for this experiment, and table 5 shows the

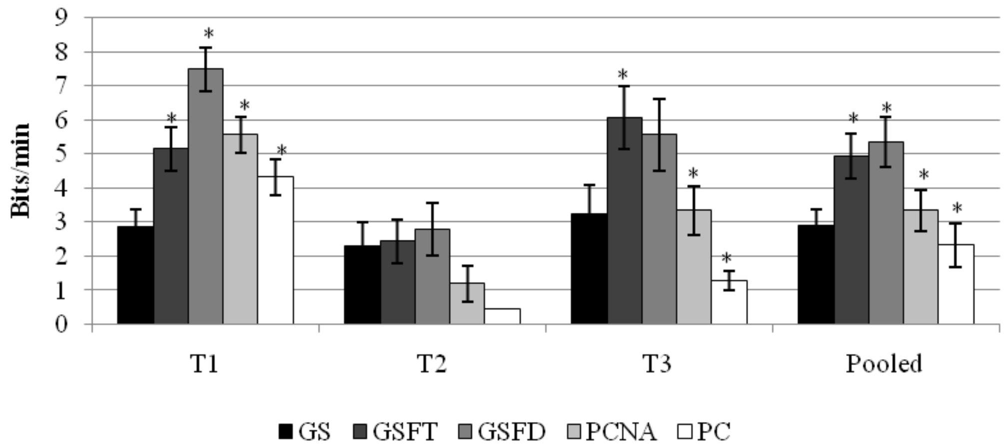

(A)

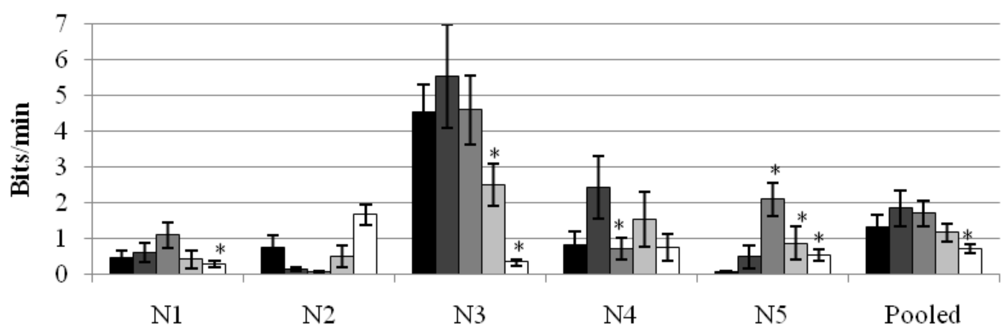

(B)

Figure 2.5 Information transfer in bits/min
Information transfer in bits/min for both the trained (A) and naïve (B) subjects across the five paradigms. GSFT or GSFD consistently had the highest information transfer rate.

significant differences. In general, GSFT and/or GSFD transferred more bits per minute than the process control paradigms. Looking at the trained subjects (figure 5A), for T1 and T3 either GSFT or GSFD had a significantly higher bit rate than GS, PCNA, and PC. In the pooled data, GSFT and GSFD had a significantly increased rate of information transfer over GS. GSFT and GSFD had a significantly higher information transfer rate than one or both forms of process control. Numerically, GS had a bit rate 25% higher than PC. GSFT and GSFD increased the information transfer to rates more than twice as high as PC. PCNA did have a higher bit rate than PC, but GSFT and GSFD still had a 47% to 60% higher information transfer rate than PCNA.

The naïve subjects showed much variability in their information transfer rates (figure 5B). For four of the five subjects, GSFT or GSFD had the highest bit rate. For some of these subjects, the difference was significant. N1 transferred significantly more information using GSFD than PC. N3's information transfer rate for GSFT was significantly higher than both process controls, and PC transferred significantly less information than any of the goal selection paradigms. N5's GSFD had a significantly higher information transfer rate than all other paradigms. In the pooled data, PC was significantly worse than both GSFT and GSFD. Numerically, GS transferred 85% more information than PC. GSFT and GSFD more than doubled PC's information transfer rate. GSFT and GSFD had a higher bit rate than PCNA by 47% and 58%.

Table 2.5 Significant differences for information transfer rate

	GS	*GSFT*	*GSFD*	*PCNA*	*PC*
T1	GSFT* GSFD*** PCNA**	GS* GSFD*	GS*** GSFT* PCNA* PC**	GS** GSFD*	GSFD**
T2					
T3	GSFT*	GS* PCNA* PC**	PC**	GSFT*	GSFT** GSFD**
Pooled	GSFT* GSFD**	GS* PC**	GS** PCNA* PC**	GSFD*	GSFT** GSFD**
N1			PC*		GSFD*
N2					
N3	PC**	PCNA* PC***	PC**	GSFT*	GS** GSFT*** GSFD**
N4		GSFD*	GSFT*		
N5	GSFD***	GSFD**	GS*** GSFT** PCNA* PC**	GSFD*	GSFD**
Pooled		PC*	PC*		GSFT* GSFD*

*p<0.05, **p<0.01, ***p<0.001

2.1.4 Discussion

Although several goal selection based BCIs exist, we presented the first study directly comparing goal selection and process control. We found the following to be true in both trained and naïve populations studied: (1) Goal selection had more hits than process control. (2) Goal selection was faster than process control. (3) Goal selection was more accurate than process control for most subjects, and (4) goal selection had a higher information transfer rate than process control.

Goal selection outperformed process control in every measure studied here. This is summarized well by the information transfer rate data presented in figure 5. However, the goal selection paradigms were not optimized, and further increases in performance could be reasonably expected. For instance, in GS and GSFT, the selection and reconfirm times were set somewhat arbitrarily at 1s. As demonstrated by Santhanam et al (2006), optimizing selection and reconfirm times can significantly increase the information transfer rate. In fact, a version of GS or GSFT could be implemented where the time was not a set time, but was instead determined by a statistical confidence threshold. This could lead to an asynchronous, user paced, BCI. Information transfer using GSFD could likewise be optimized by determining the optimal circle radius for each user. This study produced a 50% to 1600% improvement from a process control based paradigm to an un-optimized goal selection based paradigm. It is exciting to think of the magnitude of improvement that could be possible with an optimized goal selection based paradigm.

Two other sensorimotor rhythm based BCIs have adopted goal selection to some extent. One study (McFarland et al 2008) emulated computer mouse control where the subject used 2-dimensional control to move a computer cursor around a screen. When the cursor was above the target of interest, the subject could then select that target, like a mouse click. Adding the requirement that the target not only be hit by the cursor, but also selected, increased accuracy. This study combined process control, i.e. moving the cursor, with goal selection, i.e. selecting the target, to increase accuracy.

Another study (Friedrich et al 2008) used a scanning protocol similar in some respects to a P300 system (Donchin et al 2000). When the target of interest was highlighted, the subject selected the target through modulation of sensorimotor rhythms. Although it is a novel application of sensorimotor rhythms, it shares a common disadvantage of the P300 systems that the scanning protocol is relatively slow. The best mean information transfer rate was approximately 2.5 bits per minute for the trained subjects in Freidrich et al's study. In the work we presented here, the pooled data from the trained subjects had a higher information transfer rate for all three forms of goal selection. In fact, goal selection with feedback had twice the information transfer rate as the Friedrich et al study. We were able to achieve this despite having double the resting period between trials.

One trend apparent throughout all the results presented here is the importance of feedback. Goal selection with a form of feedback, GSFT and GSFD, tended to outperform goal selection alone, GS, in every measure presented. This agrees with previous literature (Neuper et al 1999, Hinterberger et al 2005, Brunner et al 2006, Hochberg et al 2006). An important distinction between GS and GSFT or GSFD is that GS offered discrete feedback, whereas the other two paradigms offered continuous feedback. In the Neuper et al (1999) study, continuous feedback was demonstrated to be more effective than discrete feedback.

In the present study, continuous feedback was only temporarily removed when the subjects performed the GS runs. When the subjects moved on to the other paradigms, continuous feedback was once again available. This is similar to another study that temporarily removed feedback in a similar fashion (McFarland et al 1998). McFarland et al found that there was no significant effect when feedback was removed. Their trained subjects could maintain performance for short periods of time. Although the group analysis found no significant effect, the removal of feedback either increased or decreased individual accuracy. Both of the above points were corroborated by the comments of the current subjects while they were performing this experiment. On the first point, some subjects thought the GS paradigm was easier when it was last because they had already experienced feedback that day. On the second point, some subjects found the movement of the cursor to be distracting, whereas others liked seeing it.

As McFarland et al showed (2008) in their computer mouse emulation, combining goal selection and process control can improve a system. This is consistent with our findings. Our subjects were moving the cursor in a similar manner across all five paradigms in what could be described as a process control strategy. However, in the goal selection paradigms, how the cursor moved under cortical control was only a small part of the overall process necessary to hit the target. GS did not show the motion of the cursor to the user. Instead, target selection was indicated by the colour of the target. This is similar to the scanning protocol study (Friedrich et al 2008). Once a target was reconfirmed, the cursor moved under automated control to the target. Our other two goal selection paradigms, GSFT and GSFD, combined goal selection with process control in a much more obvious fashion, yet kept the final control of cursor movement under the automated movement of the BCI. GSFT simply allowed the visualization of the underlying motion of the cursor in GS. GSFD allowed the user to point the BCI in the right direction. The BCI would then complete what the user began. In none of the goal selection paradigms did the user have to complete the task under cortical control. The final execution of the task was completed by the execution unit of the BCI.

In the present study we utilized two different process control paradigms. In the first paradigm, PCNA, the trial did not end until a subject hit a target. In the second paradigm, PC, the trial timed out and aborted at 6s if a target had not been hit. PC was chosen purposefully to facilitate comparison to the previously published work that allowed the trial to time out (e.g. Wolpaw and McFarland 2004, Yuan et al 2007, Yuan et al 2008). PCNA is more relevant to non-computer applications, such as controlling a wheelchair or a neuroprosthetic, where an abort might entail the wheelchair or the prosthetic returning to its starting position. That style of abort could be quite frustrating to the user. For a wheelchair, a prosthetic, or some similar application, a user will try until they succeed at their task.

Some would argue that, for trained subjects, there is no difference between a protocol that allows aborts and one that does not. Our data showed that PCNA and PC were different in every measure presented. The trained subjects were faster than the untrained subjects, but they still experienced a good number of aborts. Individually, the percentage of aborted trials ranged from 9% to 35%, and the pooled data showed 22% of the trials ended by timing out. This shows that trained subjects could not always hit a target within 6s. When timing out was not an option, more than 29% of the trained subjects' hits occurred after 6s, with individual subjects having 17%, 34%, and 36% of their hits after the 6s point. Because of the number of aborts, PC had lower accuracy than PCNA. When both speed and accuracy are combined in the information transfer rate, allowing aborts reduced the information transfer rate.

Despite the above statements, allowing aborts in process control does have some redeeming value. If the aborted trials were thrown out during data analysis, as if the trial had never occurred and the subject had an extended rest, the accuracy and the information transfer rate would be approximately equal to or greater than the values for PCNA. This makes process control with aborts, PC, the preferred paradigm when the consequences of an abort are acceptable, such as when interacting with a computer. In conclusion, process control with aborts, PC, and process control with no aborts, PCNA, are not the same paradigm. The proper strategy should be chosen based on the final application.

The present study compared five different BCI paradigms, three based on goal selection and two based on process control. Since goal selection uses the cortical signal in a manner similar to normal motor control, it follows that a BCI based on goal selection would be an easier and more natural system than one based on process control. It was hypothesized that the goal selection paradigms would be more accurate, faster in use, and easier to learn. All of those things would lead to a higher information transfer rate with decreased training time. This study tested only the first two points and showed that the goal selection paradigms were more accurate and faster in use, which together led to a higher information transfer rate. The third point, that goal selection should be easier to learn, was not tested. The design of the study did not and could not address learning. That question will be addressed in the next section.

2.2 Tracking Naive Subjects through the Learning Process

The material in section 2.2 has been reprinted with permission from Royer et al 2011.

2.2.1 Introduction

A brain-computer interface (BCI) strives to make a connection directly from a person's brain to a computer without relying on any motor output (Wolpaw et al 2002, Vallabhaneni et al 2005). BCIs promise to help the nearly 6 million people who live with paralysis (www.christopherreeve.org) by allowing them to interact with the world in ways they are no longer able. Those individuals have lost normal motor control through diseases and conditions such as amyotrophic lateral sclerosis (Lou Gehrig's disease), brainstem stroke, spinal cord injury, muscular dystrophies, or cerebral palsy (Kunst 2004). For these patients, a BCI could allow them to use a computer, a neuroprosthetic, or control a mobile robot (Kennedy et al 2000, Karim et al 2006, Hochberg et al 2006, Bell et al 2008). BCIs can also be used by able bodied individuals to extend their capabilities (Kotchetkov et al 2010).

In our daily lives, able bodied individuals receive much assistance from the systems we interact with. Anti-lock braking systems stop cars faster and safer than the driver can do by pumping the brakes himself. Spell-check and grammar-check have improved the quality of the written word. Point-and-shoot cameras dominate the camera market. However, photographers have a choice in what type of camera to use. These cameras differ in how much is required of the photographer, and how much the camera does for the user. A casual photographer, like myself, may choose the point-and-shoot model, where all that is asked of me is to frame the image and push the button. The camera then chooses the ISO speed, adjusts the lens, focuses, sets the aperture, sets the shutter speed, sets the white balance, and captures the image. On the other hand, professional photographers prefer to have more control. They frame the image, choose the ISO speed, choose the lens, focus, set the aperture, set the shutter speed, set the white balance, and then push the button. The camera merely captures the image.

The professional photographer used a control strategy named process control, whereas the casual photographer used a control strategy called goal selection. In process control, the user controls every step of the process and receives minimal to no assistance from the system. In goal selection, the user only needs to determine the goal and the system executes the process to achieve that goal. In goal selection, the system performs the work that was asked of the user in process control. Since in goal selection, less work is asked of the user, goal selection is intrinsically easier than process control.

BCIs also utilize the two control strategies of process control and goal selection. Significant advancements have been made by invasive BCIs using both control strategies. Information transfer rates of up to 6.5 bits per second have been achieved using goal selection (Santhanam et al 2006). Embodied control of a prosthetic arm was achieved using process control (Velliste et al 2008). As well as invasive BCIs, non-invasive BCIs have met success using both control strategies. The non-invasive P300 systems are intrinsically goal selection based (Farwell and Donchin 1988, Donchin et al 2000). This methodology has enabled an ALS patient who could no longer use conventional assistive devices to communicate and resume professional and social activities (Sellers et al 2010). Process control was used by a non-invasive system to move a computer cursor (Wolpaw and McFarland 2004).

Although advancements have been made using both control strategies of goal selection and process control, goal selection requires less of the user than process control, making goal selection intrinsically easier. In addition, many BCIs perform a task that otherwise would be performed through motor output. Since the majority of BCIs record their input signal from cortex alone, many other locations in the normal motor pathway are ignored, such as the cerebellum and spinal motor neurons. Goal selection more closely resembles natural motor control with the BCI system assisting the user akin to how the distributed motor network assists the motor cortex (Wolpaw 2007). Since goal selection is easier and more natural, it follows that it would be more accurate, faster in use, and easier to learn. A previous study from our lab was the first to directly compare goal selection and process control (Royer and He 2009, section 2.1). That study tested the first two points of accuracy and speed with the finding that goal selection was superior to process control in both trained and naive subjects. However, the study design had limited subjects and was unable to test if goal selection was easier to learn than process control.

Those individuals who have lost normal motor control require a BCI that is both effective and easy to learn. The previous study showed in a small sample of people the effectiveness of goal selection over process control. However, the ease of learning has not been directly compared between goal selection and process control. This study hypothesises that goal selection is more accurate, faster to use, easier to learn, and requires less mental effort than process control. This will test the results of the previous study in a larger sample size while being the first study to address the issues of ease of learning and required mental effort of goal selection vs. process control.

2.2.2 Methods

2.2.2.1 Data collection

This study was conducted according to a human protocol approved by the Institutional Review Board of the University of Minnesota. Twenty young, healthy human subjects participated in a one-dimensional BCI study using similar methods as in Royer and He (2009, section 2.1) which are described below. The subjects ranged in age from 18 to 28. Seven were male and 13 were female. Eighteen were right handed and two were left handed. Subjects were recruited from the university community. All subjects were included; none were rejected or omitted from analysis. All subjects were naive to BCI usage prior to the study.

Subjects used motor imagination to modulate the sensorimotor rhythms of their primary sensory and motor cortex. Subjects were instructed to imagine moving their right hand, arm, or shoulder to move the cursor to the right, and to imagine moving their left hand, arm, or shoulder to move the cursor to the left. Subjects were encouraged to imagine movements familiar and comfortable to them, such as hitting a ball with a tennis racquet if they played tennis or dribbling a basketball if they played basketball. Other motor imaginations that were suggested included squeezing a tennis ball, punching, and lifting weights. Each subject was free to use whatever motor imagination worked best for them. By merely imaging moving their right or left hands, the subjects created event related (de)synchronization (ERD or ERS) of their neurons that was measured via scalp recorded electroencephalography (EEG) as a decrease (ERD) or increase (ERS) in spectral amplitude in the mu and beta frequency bands (Pfurtscheller and Lopes da Silva 1999). As illustrated in figure 6A, subjects wore a 64-channel EEG cap connected to a Neuroscan amplifier. The particular EEG cap used was the Compumedics NeuroMedical Supplies Quik-Cap, with setup taking approximately 20 minutes per subject. The signal from all 64 channels was fed into the general purpose system BCI2000 (Schalk et al 2004).

2.2.2.2 Experimental paradigms

The subjects were split into 4 groups of 5 subjects. Each group was assigned one of the paradigms from Royer and He (2009, section 2.1) that are described below. Each subject completed 8 sessions of their assigned paradigm. Sessions occurred approximately once per week and consisted of 10 four minute runs. Between runs, subjects rested for a user-determined period of time. Each run had as many trials as the subject could complete in 4 minutes with right and left block randomized cues presented. Subjects had 3 s of rest after each trial.

In the four paradigms, the underlying signal processing, operation of the paradigms, and movement of the cursor were identical. The paradigms differed only in control strategy. Two of the paradigms were based on process control, and two were based on goal selection. The two process control based paradigms were process control with aborts (PCA) and process control with no aborts (PCNA). The two goal selection based paradigms were goal selection with feedback limited by distance (GSFD) and goal selection with feedback limited by time (GSFT). For purposes of analysis and presentation, the two paradigms based on process control (PCA and PCNA) were grouped into the process control paradigms (PCP). Similarly, the two paradigms based on goal selection (GSFD and GSFT) were grouped into the goal selection paradigms (GSP). For all paradigms, the subject was instructed to move the computer cursor to the yellow target located on either the right or left side of the screen (figure 6). The targets were shown for 1 s before the cursor appeared, then at time 0, the cursor appeared and moved under cortical control. In PCP, the subjects had to move the cursor all the way to the target themselves in order to get a hit. In GSP, once the BCI determined the subject's goal through either time or distance, the BCI moved the cursor the rest of the way to the target to get a hit. The subject received the assistance of the BCI and did not have to do all the work themselves. In both PCP and GSP, one paradigm was time constrained (PCA and GSFT) and the other paradigm had no time limit (PCNA and GSFD).

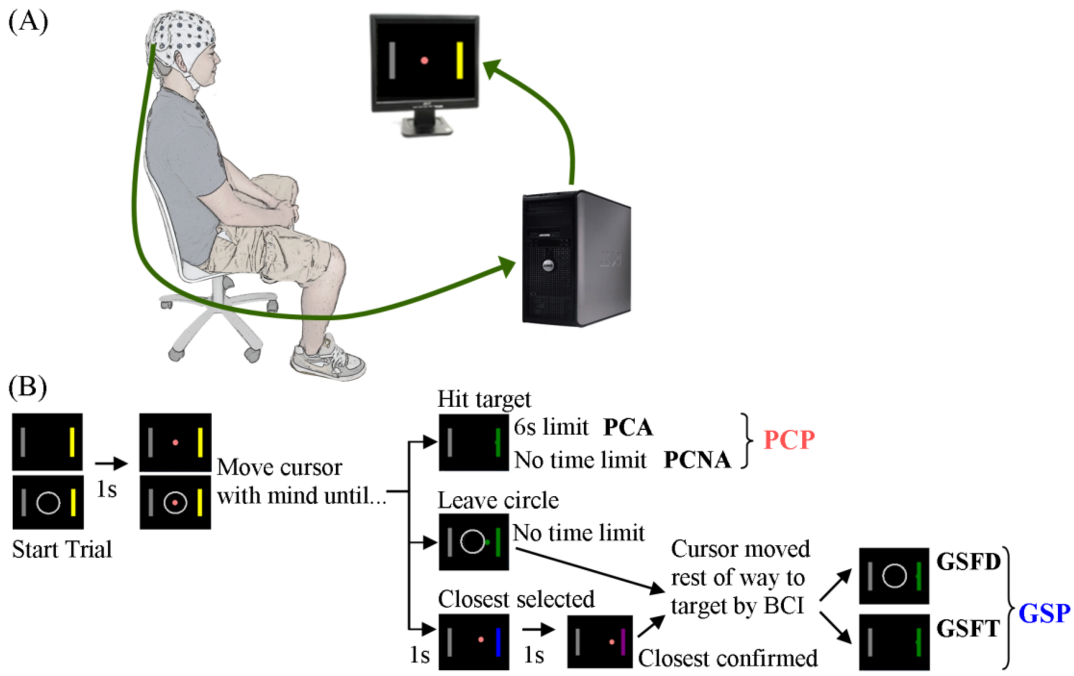

Figure 2.6 Experimental setup and paradigms
(A) Healthy human subjects sat motionless in a chair facing a computer monitor and imagined moving their right or left hand, arm, or shoulder. Their EEG signal was sent to a computer, which translated the raw EEG into a control signal that moved a computer cursor right or left on the screen. (B) Experimental paradigms. At the start of each trial, two targets appeared on the screen. The subject was instructed to hit the yellow one (here: right). For all paradigms, a correctly hit target turned green and an incorrectly hit target, or a miss, turned red. In the bottom row describing the GSFT paradigm, the closest target after one second of mind control was selected and turned blue. If the cursor was still closest to that target after an additional second, the target would be confirmed and turn purple. However, if the cursor were closer to the other non-selected (non-blue) target, both targets would turn blue and the final destination determined by a third second of cursor control, with best two out of three winning. For analysis purposes, process control with aborts (PCA) and process control with no aborts (PCNA) were grouped into the process control paradigms (PCP). Similarly, goal selection with feedback limited by distance (GSFD) and goal selection with feedback limited by time (GSFT) were grouped into the goal selection paradigms (GSP). Supplementary videos demonstrating PCA (1.08 MB .avi), GSFD (1.97 MB .avi), and GSFT (1.37 MB .avi) are available online.

The details of how each paradigm progressed is shown in figure 6B. In PCA, the subjects had 6 s to hit a target with the cursor (figure 6B, top). If no target was hit within 6 s, the trial timed out and aborted. In this paradigm, the words "time out" and "abort" are used interchangeably.

In PCNA, subjects also had to hit a target with the cursor (figure 6B, top). The only difference between PCA and PCNA was that the subjects had no time limit in PCNA. In order to move on to the next trial, the cursor had to hit one of the targets.

In GSFD, there was a grey circle with a radius of 20% of the screen centered between the two targets (figure 6B, middle). Once a subject moved the cursor outside of the circle, it automatically went to the closest target. GSFD had no time limit and the cursor had to exit the circle before progressing to the next trial.

In GSFT, subjects did not have to move the cursor any particular distance, but instead a hit target was determined by time. After 1 s of cortical control, the closest target to the cursor was selected (figure 6B, bottom). This was indicated to the subject by the selected target turning blue. After another 1s of cortical control, the closest target to the cursor was again selected. If it was the same target as in the previous 1 s (the blue target), the target turned purple and the cursor travelled automatically to it. If the closest target at the end of the 2nd 1 s interval was the opposite, non-blue target, the new target turned blue and a third 1 s cortical control period selected the final target by whichever target was selected twice in the three 1 s intervals.

Supplementary videos illustrating GSFD, GSFT, and PCA are available on the Journal of Neural Engineering’s website. These videos are not videos of subjects performing the paradigms, but were created using the BCI2000 signal generator controlled by the mouse. The objective of these videos is to demonstrate the operation of the paradigm, and not to represent the capabilities of the paradigm. However, all videos show trial times that are within typical subject performance.

In order to allow for a valid comparison, the four paradigms were designed to be as similar as possible, with consistent inner workings and programming. All paradigms had a consistent cursor speed with the position of the cursor updated every 40 ms. The time before (1 s) and after (1 s) cortical control of the cursor was the same for all paradigms, as was the time between trials (3 s).

2.2.2.3 Control of the cursor

The movement of the cursor was determined by a value that we called "the control signal". The method of calculation of the control signal was as follows. Once the EEG signal was fed into BCI2000, the AR spectral amplitudes were calculated for 3Hz bins centred on a multiple of three from 0 to 30 Hz. Then, in the classifier, the spectral amplitudes from the set frequency bins of the set electrodes were given a weight and added together. The frequency bins and electrodes were selected as described in the next section, 2.4 Control signal selection. The signal from the classifier was then passed through a normalizer which linearly transformed the signal into the control signal, as described in the next paragraph. Positive values of the control signal moved the cursor to the right, and negative values of the control signal moved the cursor to the left. The magnitude of the cursor movement was determined by the amplitude of the control signal. .

The normalizer linearly transformed the signal by multiplying the classifier signal by a gain and adding an overall offset. Adaptation was built into this process. After each trial, the gain and offset of the normalizer were adjusted to create a control signal with zero mean and unit variance. This was used to reduce the effect of session to session, and even within session, recording differences. It also helped normalize the cursor movement speed between subjects. The zero mean and unit variance were determined by a buffer that was updated with the control signal data at the end of each trial. This buffer was a set length that was chosen to be long enough for multiple trials of both left and right, keeping in mind the timed versus untimed nature of each paradigm. The buffer was 60 s for PCNA and 30 s for PCA, GSFT, and GSFD. This allowed the buffer to contain 12.7, 7.3, 15.0, and 17.3 median length trials for each paradigm, respectively. Only the gain and the offset of the normalizer changed each trial. The specific electrodes, frequency bins, and weights were fixed and only changed manually. The exact combination of electrodes, frequency bins, and weights defined what we are calling a user's control signal. This is the adaptation that is built into BCI2000 version 2.0 (Schalk et al 2004).

2.2.2.4 Control signal selection

For the first session, all subjects used the same control signal of the negatively weighted auto-regressive (AR) spectral amplitudes from 7.5 to 13.5 Hz and 16.5 to 25.5 Hz of electrode C3 in the 10-20 international system. Relating to the above description, the control signal for the first session was electrode C3, using 9 Hz, 12 Hz, 18 Hz, 21 Hz, and 24 Hz, all with a -1 weight. This was chosen based on previous research that showed that naive BCI subjects could more easily produce similar levels of 8-12 Hz activity than they could differential activity (Pineda et al 2003). By limiting the control signal to only one side of the head, subjects had more flexibility in EEG signals that could adequately control the system. Since we were not rejecting any subjects, it was important that every subject had the best possibility of succeeding. We did not want our subjects losing motivation due to frustration. Another reason that this control signal was chosen was that recent research showed that increased speed of motor imagery produced a greater EEG signal on both C3 and C4 of equal amounts (Yuan et al 2010b). If the signal of those two electrodes were subtracted, all speed information was lost. In contrast to previous studies that set the control signal as the difference between electrodes on opposite sides of the head (Royer and He 2009 (section 2.1), Wolpaw and McFarland 2004), in this study, the subjects had finer control of the magnitude of the cursor movement because we did not subtract the signal across both hemispheres.

An example of the desired outcome of the initial chosen control signal is, if a subject was imaging a right handed motion, that would cause an event related desynchronization visible in the chosen frequencies of C3 (Pfurtscheller and Lopes da Silva 1999, Wolpaw and McFarland 2004, Pfurtscheller et al 2006, Kamousi et al 2007, Yuan et al 2008, 2010a). Since the spectral amplitudes were negatively weighted, this reduction in spectral amplitude was translated to an increase in the control signal, which moved the cursor to the right. The greater the change in spectral amplitude, the greater the distance of cursor movement.

The data from the first session was used to customize each subject's control signal for the second session according to the guidelines in the BCI2000 Offline Analysis online tutorial (www.bci2000.org/wiki/index.php/User_Tutorial:Performing_an_Offline_Analysis_of_EEG_Data, Schalk and Mellinger 2010). In brief, electrodes and frequencies were selected that had the highest r^2 for the conditions right target versus left target. Since the subjects were encouraged to use motor imagination to generate SMRs, the electrodes were limited to FCz-6, Cz-6, and CPz-6 (box in figure 7A). The frequencies were limited to the 3Hz bins centered on 6 to 30 Hz. Control signals for all sessions were also generally limited to a single side of the head with a single positive or negative weighting for the reasons described above. The second session's data was then used to update the control signal for the third session. This continued until session 7, when the control signal was locked and remained the same for sessions 7 and 8. An additional constraint was that the control signal in session 7 could not be new to the subject, but had to be one the subject had used previously. This was done to minimize the likelihood that the subject would be locked for the final two sessions into a control signal that did not work for them. In general, we did not see major changes in a subject's r^2 values from session to session. Rather, the typical case was that the control signal was customized for session 2, and then tweaked with minor changes that might have added or subtracted neighboring electrodes or frequencies. In a few of the early sessions, multiple control signals were used in a single session. Because of this, we tracked the control signal used for each individual run.

In the early stages, when subjects were trying multiple imagination strategies in an effort to find what worked, we encouraged subjects to use a particular imagination strategy for an entire run. We often recorded which mental strategy was being used. We would then customize the control signal for the most successful strategy. The next session, we would inform the user of the strategy that they had successfully used the previous time. The same control signal customization procedure was followed for each subject, regardless of assigned paradigm. Therefore, all subjects are included in the results shown in figure 7.

2.2.2.5 All-stars

Since subjects were not excluded from the study based on ability, or inability, to use a BCI, a group named "the all-stars" was formed to serve the purpose of a skill level control. At the end of the 8 sessions, the best subject from each group, or the subject with the highest average information transfer rate for session 7 and 8, was designated as an "all-star." The all-stars completed two additional sessions intended to better allow for a clear comparison across paradigms, as well as to the previous study (Royer and He 2009, section 2.1). Each session consisted of 3 runs of each paradigm in block-wise random order. The results are presented as the grouped data. The conclusions were the same for each individual as presented for the group.

2.2.2.6 Data analysis

As in the previous study, subject performance was measured via four factors: accuracy, number of hits per run, time to hit, and information transfer rate. Accuracy was determined as the number of hits divided by the number of trials. Time to hit was the time that the cursor was under cortical control. Information transfer rate was calculated first as bits/trial (Wolpaw et al 2002) according to the following equation where N is the number of targets and P is the probability of a hit, or the accuracy:

$$\log2N + P \log2P + (1-P) \log2[(1-P)/(N-1)] \quad (1)$$

Information transfer rate in bits/min was obtained by multiplying the results of equation (1) by the number of trials per minute. Accuracy, number of hits per run, and information transfer rate were calculated for each run for each subject. A fifth measure, effort of hit, was also used. Since it is widely recognized (Ray and Cole 1985, Pfurtscheller and Lopes da Silva 1999, Fink et al 2005, Neuper et al 2005, Keil et al 2006) that increased effort is reflected in a greater alpha spectral power reduction, effort of hit was calculated as the integral of the squared control signal during the time the cursor was under cortical control before a hit. Given that the control signal consisted of the spectral amplitude, squaring the control signal is the equivalent of the spectral power. As the integral increased, that indicated that more modulation of the spectral power was necessary. More modulation indicates more effort. Effort of hit and time to hit were calculated for each hit for each subject.

All measures were tested for normality using a 2-sided Lilliliefors test. All measures were found to be non-normal. Therefore, we used medians and a 2-sided sign test to test for statistical significance. Alpha = 0.05 for all statistical analysis. No p-value correction was applied.

Measures are presented in figure 8 as the median of the grouped data for each session. The shaded area in the figure indicates the 95% confidence interval of the median. The measures are significantly different from each other if the confidence intervals do not overlap. In order to look at learning over time, the percent change from session 1 was calculated for each paradigm and each measure. Statements such as, "GSP showed significantly more improvement than PCP" refer to the percent change from session 1 being significantly different between GSP and PCP.

In figure 9, the box plots show the distribution of the all-star data. The lower whisker extends from the minimum value to the 25th percentile. The box extends from the 25th percentile to the 75th percentile with the median drawn across the box. The upper whisker extends from the 75th percentile to the maximum value. Asterisks above the upper whisker indicate significantly different medians between GSP and PCP.

The time frequency plots in figure 10 present the AR spectral amplitudes during single trials with the baseline subtracted. Baseline was the median AR spectral amplitude for all 1 s intervals after the targets were displayed but before the cursor appeared. Calculation of the spectral amplitudes was performed in the same manner as done real time by the BCI with a 16th order AR model calculating 3Hz bins centered on a multiple of three from 0 to 30 Hz. Window length was 160ms with 50% overlap.

2.2.3 Results

During the course of the study, we looked at three main categories of data: how the subject specific control signals evolved over time, how the subjects performed on their single assigned paradigm over the eight sessions, and how the all-stars performed using all paradigms in the same session. Multiple measures for each of those categories is presented below.

2.2.3.1 Control signal evolution

Figure 7B and C show the evolution of the control signal across the eight sessions, both in terms of electrodes used (figure 7B) and frequencies used (figure 7C). The color in figure 2B indicates the percent of control signals that used that channel. The circled electrodes indicate the electrodes that were used in the most runs. The number of circled electrodes indicates the average number of electrodes that were used in that session's control signals. The frequencies used are indicated by figure 7C. The dark bars represent the percent of control signals that used each frequency. The light bars extending to 100% represent the frequencies that were used in the most runs. The number of light bars indicates the average number of frequencies that were used in that session's control signals. Subjects typically used two electrodes and two frequency bins. Over the course of the 8 sessions, the electrodes shifted from the left side of the head to the right side of the head. Not surprisingly, C3, C4, CP3 and CP4 were the most commonly used electrodes. The most commonly used frequencies were 9, 12, and 15 Hz. The final control signal for all but one subject involved at least one of the 9, 12, or 15 Hz bins.

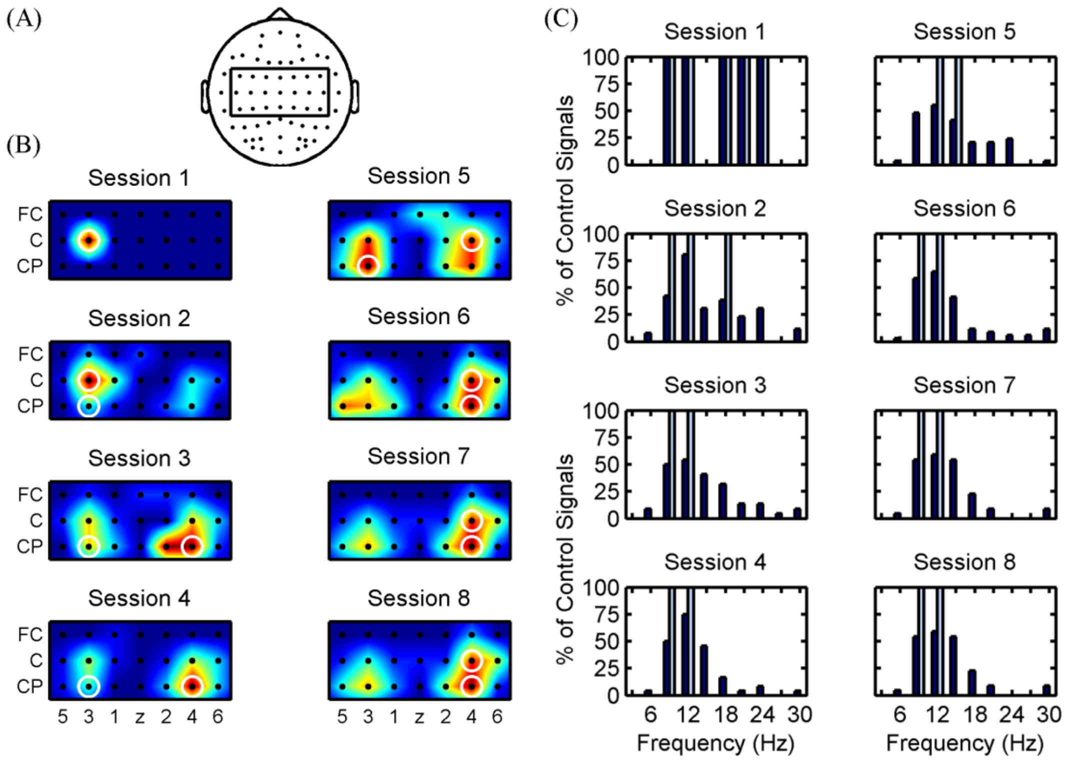

Figure 2.7 The evolution of the control signal across sessions for all subjects
For session 1, all subjects used the same control signal. Then, each subject's control signal was customized based on their previous session's data. For session 7, the control signal was restricted to a control signal that the subject had previously used. The control signal used in session 7 had to be used in session 8. (A) Possible electrodes were limited to the outlined box containing FCz-6, Cz-6, and CPz-6 . (B) Electrodes used in each session. Each box represents the same outlined box from (A). Colour indicates the percent of control signals that used that channel (red = most, blue = none). Circles indicate the electrodes that were used in the most number of runs. The number of circles indicates the average number of electrodes that were used in that session's control signals. (C) Centre frequencies of the 3Hz wide frequency bins used in each session. Dark bars represent the percent of control signals that used each frequency bin. Light bars extending to 100% represent the frequencies that were used in the most runs. The number of light bars indicates the average number of frequencies that were used in that session's control signals.

2.2.3.2 Eight session performance metrics

For both GSP and PCP, accuracy increased over the 8 sessions (figure 8A). GSP was significantly more accurate than PCP in all sessions: 34% more on average, and 33% more in session 8. GSP also showed significantly more improvement in accuracy than PCP. This improvement occurred earlier and was more sustained for GSP than PCP. By session 2, GSP was already significantly more accurate than it was in session 1. PCP did not significantly improve on its session 1 accuracy until session 4. Both GSP and PCP continued to significantly improve on their session 1 accuracy until both leveled off in sessions 7 and 8. By session 8, GSP showed 52% more improvement in accuracy than PCP.

The number of hits per run increased for both paradigms over the 8 sessions (figure 8B). GSP had significantly more hits than PCP in all sessions: 102% more on average, and 115% more in session 8. GSP also showed a significantly greater increase in number of hits than PCP. By session 2, GSP exhibited a significant increase in the number of hits per run. GSP continued to demonstrate a significant steady increase in number of hits per run that did not level off. PCP's number of hits fluctuated up and down across sessions and exhibited no sustained significant change across sessions. By session 8, GSP showed 161% more learning than PCP in terms of number of hits per run.

Even though both GSP and PCP consisted of one timed paradigm (GSFT and PCA) and one untimed paradigm (GSFD and PCNA), GSP had significantly less time to a hit than PCP in all sessions: 40% better on average, and 44% better in session 8 (figure 8D). GSP was also much more consistent in time to a hit, resulting in a very narrow 95% confidence interval for GSP. GSP had a slow but sustained decrease in time to hit that was significant in sessions 6 through 8, whereas PCP lost all significant gains and showed no significant change in time to hit by session 8.

The information transfer rate is one metric that combines the speed and accuracy presented by the previous figures into one measure. As expected, the information transfer rate increased for both paradigms over the 8 sessions (figure 8C). GSP transferred significantly more information than PCP: 324% more on average, and 411% more in session 8. GSP showed a significantly greater increase in information transfer rate than PCP. GSP showed consistent improvement. By session 2, GSP transferred significantly more information than it had in session 1. Another significant improvement in the information transfer rate occurred between sessions 5 and 7, then GSP leveled off. PCP was slow to improve its information transfer rate, showing the first significant gain in session 6. However, sessions 7 and 8 were quite volatile and PCP lost almost all significant improvement. By session 8, GSP showed 282% more improvement in information transfer rate than PCP.

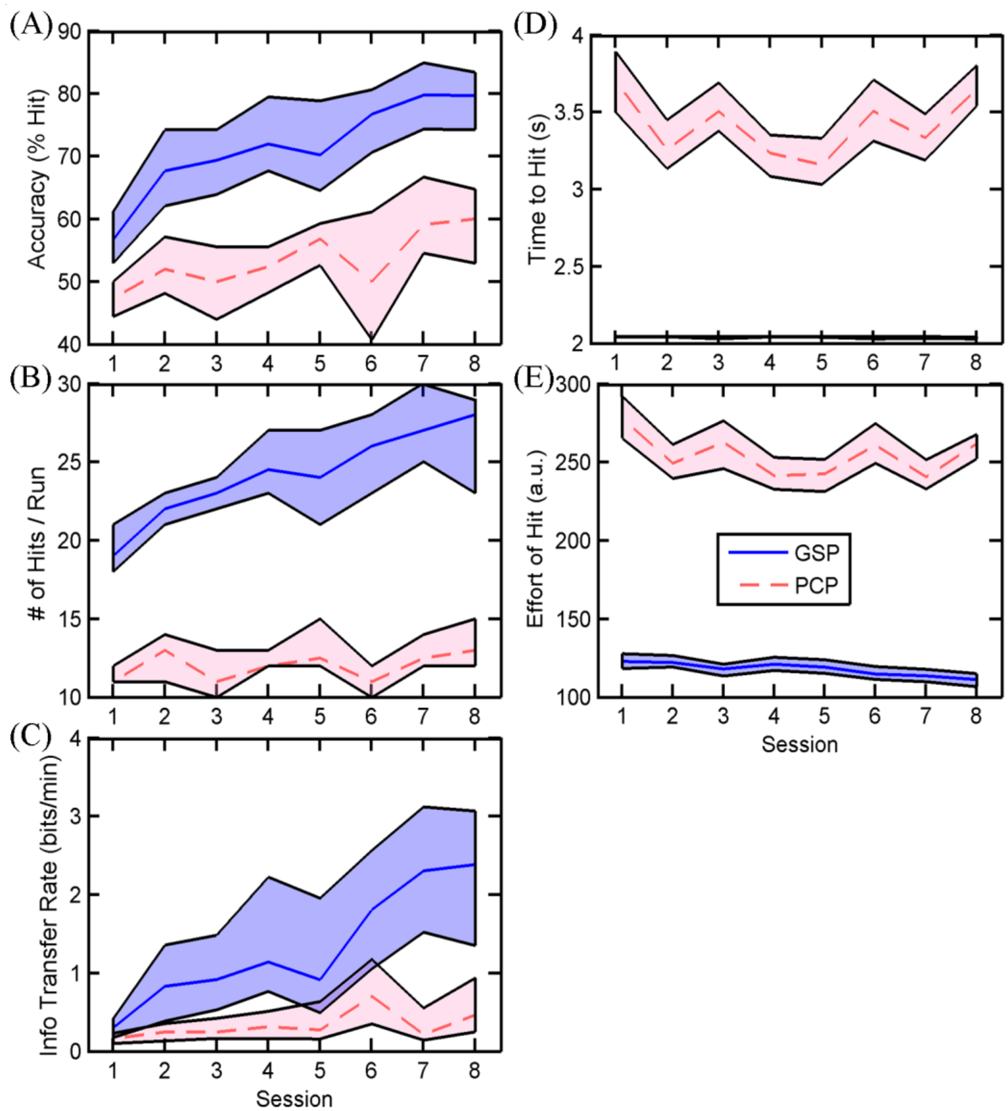

Figure 2.8 Eight session performance metrics
Goal selection was significantly better and showed significantly more learning than process control over the eight sessions. Plotted lines are the medians of the grouped data for each session. Blue (darker, solid line) represents GSP. Pink (lighter, broken line) represents PCP. The legend in (E) applies to (A-E). Shaded areas indicate the 95% confidence interval of the median. Significance is indicated by non-overlapping areas. a.u. = arbitrary units.

The effort of hit decreased for both paradigms over the 8 sessions (figure 8E). GSP required significantly less effort than PCP in all sessions: 52% less on average, and 57% less in session 8. GSP showed a significantly greater decrease in effort of hit than PCP. PCP was quite variable and lost almost all of the significant reductions in effort obtained, whereas GSP showed steady significant improvement eventually outpacing PCP. By session 8, GSP showed 63% more reduction in effort of hit than PCP.

For all 5 performance measures, GSP was significantly better with a median increase in performance of 54% from PCP to GSP across all sessions and measures. GSP also showed significantly more improvement than PCP for all 5 measures, showing on average twice the learning of PCP.

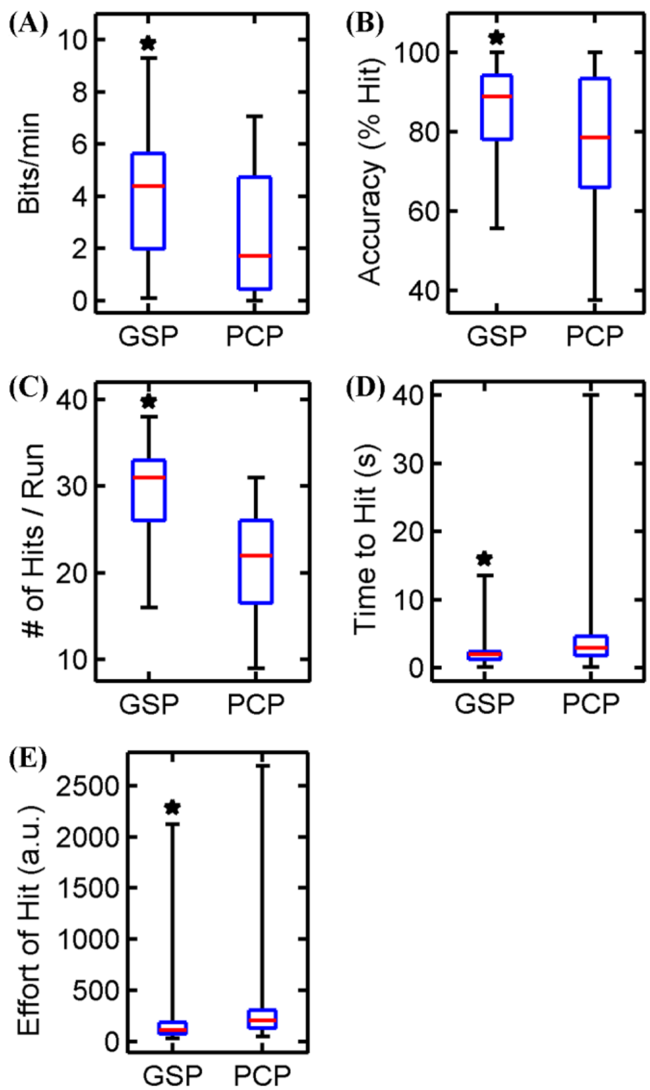

Figure 2.9 All-star performance metrics
Goal selection was significantly better than process control when performed in the same session. Asterisks on the data from the all-stars indicate significantly different medians between GSP and PCP. a.u. = arbitrary units.

2.2.3.3 All-star performance metrics

In order to more directly compare the influence of control strategy on performance, the best subject from each group performed two additional sessions of three runs of each of the four different paradigms. Figure 9 shows the 5 performance measures for these sessions. For all 5 measures, GSP was significantly better than PCP. GSP transferred 155% more information than PCP (figure 9A). GSP was 13% more accurate than PCP (figure 9B). GSP had 41% more hits per run than PCP (figure 9C). GSP was 31% faster to a hit than PCP (figure 9D). GSP required 44% less effort for a hit than PCP (figure 9E). Across all 5 measures, the median increase in performance was 41% from PCP to GSP.

2.2.4 Discussion

This study focused on the effect of control strategy, goal selection or process control, on a subject's ability to learn to use a BCI. The measures studied were accuracy, number of hits per run, time to a hit, information transfer rate, and effort of hit. From the very first session, goal selection outperformed process control. Goal selection was more accurate and faster to use, which led to a higher information transfer rate. This was achieved with less effort than process control required. As the sessions progressed, goal selection showed significantly more improvement than process control across all measures. This indicates that the goal selection subjects demonstrated more learning than their process control counterparts. These conclusions held even when the paradigms were not grouped into GSP and PCP but analyzed separately.

Did learning actually occur in this study? For GSP, all five performance measures were significantly better by session 8 than they were in the first several sessions. GSP certainly demonstrated learning. For PCP, four of the five measures were significantly better by session 8 than they were in the first session. However, time to hit did not show sustained significant improvement. The results do show significant improvement did occur, but were not maintained. This could be due to the fact that the time to hit measure only included hits. In the first session, there were not as many hits as in the later sessions. As subjects progressed, targets that previously would have resulted in an abort or a miss now resulted in a hit that required protracted lengths of time. That argument combined with the fact that the other four measures showed significant improvement leads to the conclusion that PCP did demonstrate learning.

The fact that learning occurred does not necessarily indicate that subjects were fully trained. In terms of accuracy and information transfer rate, GSP leveled off for the last two sessions. By session 8 GSP was still improving in terms of the number of hits per run and the effort of hit, and the time to hit was still inconsistent. Given those results, GSP subjects could be considered trained, but still refining their skills.

Were the PCP subjects fully trained? PCP was not nearly as consistent as GSP in performance. In all measures, PCP was quite volatile, improving significantly in one session and then losing those gains in the subsequent sessions. For all measures but accuracy, that pattern of performance continued for all 8 sessions. PCP accuracy did somewhat level off in sessions 7 and 8. Given those results, PCP subjects were not trained, but were still learning.

Although GSP transferred over five times the information as PCP did in session 8, the goal selection paradigms were not optimized. The 1 s time interval in GSFT was somewhat arbitrarily chosen, as was the radius of the circle in GSFD. Those times and distances could be optimized for each user. The very design of the goal selection paradigms could be radically changed and improved upon. For example, the goal of right or left could have been decided by analyzing the motion of the cursor and choosing a goal when a certain confidence threshold had been crossed. What was used in this study were only two possible ways to determine a goal. Much better ways exist, including methods that do not rely on the motion of the cursor. We chose to keep cursor motion a part of the paradigms in order to most effectively compare the goal selection paradigms to the process control paradigms. Some might argue that, since cursor motion was part of the paradigms, we were not actually using goal selection. However, in our goal selection paradigms, the final execution of the task was performed by the BCI system and not the user. The goal selection subjects did not have to do all the work themselves, whereas the process control subjects did. Although the goal selection paradigms could have been improved upon, these methods were chosen to facilitate comparison with previous studies (Royer and He 2009, section 2.1). Similarly, other improvements could have been made to all paradigms to improve performance, such as changing the methods of control signal selection, classification, or adaptation. Those improvements were not implemented in order to allow a fair comparison between this and previous studies (Royer and He 2009, section 2.1).

How did changes of the control signal affect the study? The control signal changed throughout the study in two primary ways. First, the control signal was customized between sessions to those electrodes and frequencies that the subject could best manipulate. Although no formal blinding procedures were followed, we generally did not know which paradigm a subject was assigned to while performing the customization. We followed the same control signal customization procedure on all individuals, regardless of paradigm. Hence, the customization did not influence the overall conclusions of the relative merits of goal selection vs. process control. Second, the control signal experienced adaptation within sessions. The result of the control signal adaptation was a signal with zero mean and unit variance. This was used to reduce the effect of session to session, and even within session, recording differences. It also helped normalize the cursor movement speed between subjects. As mentioned in the methods, the programming parameters governing the adaptation were adjusted to account for the longer trial lengths of the untimed vs. time based paradigms. Therefore, neither the adaptation of the control signal within a session, nor customization of the control signal between sessions should have influenced the overall purpose and conclusions of this study.

A common question that arises when discussing motor imagery based systems is the handedness of the subjects. Here, both GSP and PCP had nine right handed subjects and one left handed subject. Because we customized the control signal to each individual, we argue that handedness did not affect the study. Each subject was able to use electrodes from whichever hemisphere they could best control, regardless of handedness. The predominance of right handed subjects was the reason why the chosen initial control signal was from the left hemisphere. However, as can be seen in figure 7, by session 8 the majority of control signals were from the right hemisphere (55%), with an almost equal minority staying on the left (45%).

A previous study compared goal selection to process control in two populations of users: naive and trained (Royer and He 2009, section 2.1). Each subject used both goal selection based and process control based paradigms each session. This study followed a naive population as they learned to use one particular paradigm across 8 sessions. Afterwards, the best subject from each group, or the all-stars, performed two additional sessions using all paradigms, like in Royer and He (2009, section 2.1). Therefore, valid comparisons between the two studies include comparing the naive subjects (Royer and He 2009, section 2.1) to sessions 1 and 2, the trained subjects (Royer and He 2009, section 2.1) to sessions 7 and 8, and the trained subjects (Royer and He 2009, section 2.1) to the all-stars. Adding to the validity of the comparison is that the trained subjects had 6 to 8 weeks experience with approximately one session per week.

In general, the two studies have similar results. The current study confirmed the findings of Royer and He (2009, section 2.1) that goal selection was more accurate and faster to use. That combined to create a higher information transfer rate (ITR). The current study confirmed those results in a larger sample size. However, there were a few interesting differences. The subjects in Royer and He (2009, section 2.1) may have been intrinsically better since both the naive and trained subjects displayed better accuracy and information transfer then shown here in the 8 sessions. This is further supported by the fact that the all-star data is comparable to the trained data. Another factor influencing these results is the difference in overall study design between the two studies. In the current study, all subjects began with the understanding that they had committed themselves to 8 sessions worth of experiments. Even if they became frustrated at their lack of progress, they had committed to complete the study. The subjects in Royer and He (2009, section 2.1) did not have to commit to a certain number of experiments. This led to a natural selection effect in subject ability. Those subjects that were not very good would become frustrated with the experiment, and remove themselves from the subject pool. Hence, the previous trained data is naturally composed of subjects that would have made the current all-star group.

Other discrepancies in the data relate to the transfer of learning between control strategies in Royer and He (2009, section 2.1) and not this study. Although the naive number of hits per run is comparable to this study for both GSP and PCP, the trained number of hits per run is only comparable for GSP. The PCP subjects never improved to the number of hits per run seen in Royer and He (2009, section 2.1). Since the PCP subjects in the current study did not benefit from the learning achieved with a goal selection based paradigm, the PCP subjects could not perform at the same level. This fact made an important impact in the ITR of GSP versus the ITR of PCP between the two studies. In the current study, GSP had an ITR four to five times that of PCP, whereas in both the naive and trained subjects GSP only had an ITR approximately twice that of PCP. This is more similar to the current all-star data where the subjects were performing goal selection based paradigms in succession with process control based paradigms. These results support the fact that goal selection is easier to learn than process control. They also demonstrate that learning did transfer between the goal selection based paradigms and the process control based paradigms in Royer and He (2009, section 2.1).

Neuper et al (2009) conducted a multiple session study similar to the current study in many ways. They used motor imagery of left and right hand movements to control a one-dimensional BCI. Their criterion for classification accuracy is most similar to the GSFT paradigm in this study, but they only had one selection period that lasted 4 s. At the end of the 4 s, the trial was classified as right or left if it had been classified that for at least 3 s out of the 4 s. Their average feedback classification result was 68-70%. This is nearly identical to the accuracy of the current study for comparable sessions (GSP sessions 2-4). Neuper et al (2009) also had twenty subjects and customized the control signal to each subject. In both their study and the current study, the chosen frequencies had a distribution that was biased towards 10 to 12 Hz, with some subjects using higher frequencies up to 30 Hz. In both their study and the current study, the control signal was not static but was updated throughout the study.

A surprising result of the Neuper et al (2009) study was that there was no improvement of right/left classification accuracy across the sessions. They hypothesized that the reason could be that their three feedback sessions scheduled sometimes weeks apart were not numerous or frequent enough to show learning. The current study supports their hypothesis in three ways. First, our results did not show significant improvement over session 2's data until session 6 for GSP, and longer for PCP. This shows that our subjects needed at least five sessions of feedback, two more than in Neuper et al (2009), in order to show significant improvement. The significant improvement that we often saw from session 1 to session 2 could possibly be attributed to customizing the control signal for each user. Neuper et al (2009) had a customized control signal in the first feedback session from data they had gathered in a separate screening session. Second, each GSP session had about three times, and each PCP session had about twice, the total number of trials as Neuper et al (2009). Third, our sessions were regularly scheduled. The frequency and longer length of the current study both in number of sessions and in total number of trials allowed our subjects to demonstrate significant improvement in all measures.

Neuper et al (2009) was not the only BCI study that failed to show learning. Kubler et al (2010) performed an exhaustive literature search investigating how much learning is involved in BCI control of non-invasive and ECoG systems in human studies. The vast majority of 137 studies consisted only of one to four BCI sessions. Their conclusion was that most BCI studies do not involve learning. As shown in figure 8, we were able to show significant learning.

The ERD/ERS literature presently consists of mainly ERD/ERS data from either trained or naive subjects, but does not address the progression when learning is involved (Neuper et al 1999, Wolpaw and McFarland 2004, Pfurtscheller et al 2006, Yuan et al 2008, Neuper et al 2009, McFarland et al 2010, Yuan et al 2010b). Not only did the current study show significant learning, to our knowledge it is the longest running sensorimotor rhythm study that presented data tracking the progress of subjects from naive to trained. The current study is unique in its combination of duration and large subject pool. Twenty subjects completed 160 sessions of 1,600 runs consisting of 46,036 trials. Sample time frequency plots of trials featuring both right and left targets of GSP and PCP are shown in figure 10. Those four trials represent less than 1/10,000 of the data and yet demonstrate the important point that there are many factors in this study. The trials are dissimilar in many ways: patterns of ERD/ERS/rest amplitude and duration, trial lengths, targets, sessions, control strategies, end results, overall subject skill levels, control signals from different sides of the head, and different weights of the control signal. Additionally, these plots only feature the electrodes used for control. What was the EEG signal on the non-control electrodes? Future work will tease apart these factors to determine the important aspects of EEG signal changes while learning to use a sensorimotor rhythm based BCI, and the influence that control strategy has on those changes. We hypothesize that the underlying EEG signal will be more conducive to goal selection, and that the EEG signal controlling a typical goal selection trial will change more over time compared to process control. The full analysis of the evolution and importance of different features of the EEG signal in a study of this duration and subject pool will be a useful addition to the ERD/ERS literature.

As shown in figures 8 and 9, goal selection requires less mental effort than process control. An additional advantage of goal selection is that the user can "take a mental break" while the BCI system is completing the execution. The combination of requiring less effort and naturally introducing breaks leads to less overall mental fatigue from using a goal selection based BCI when compared to a process control based BCI. This has been seen in our personal experience with subjects. As discussed in Bai et al (2010), the minimization of fatigue during BCI use will be important as BCIs move from laboratory to clinical settings. The patient populations that many BCIs are designed to serve, such as those with

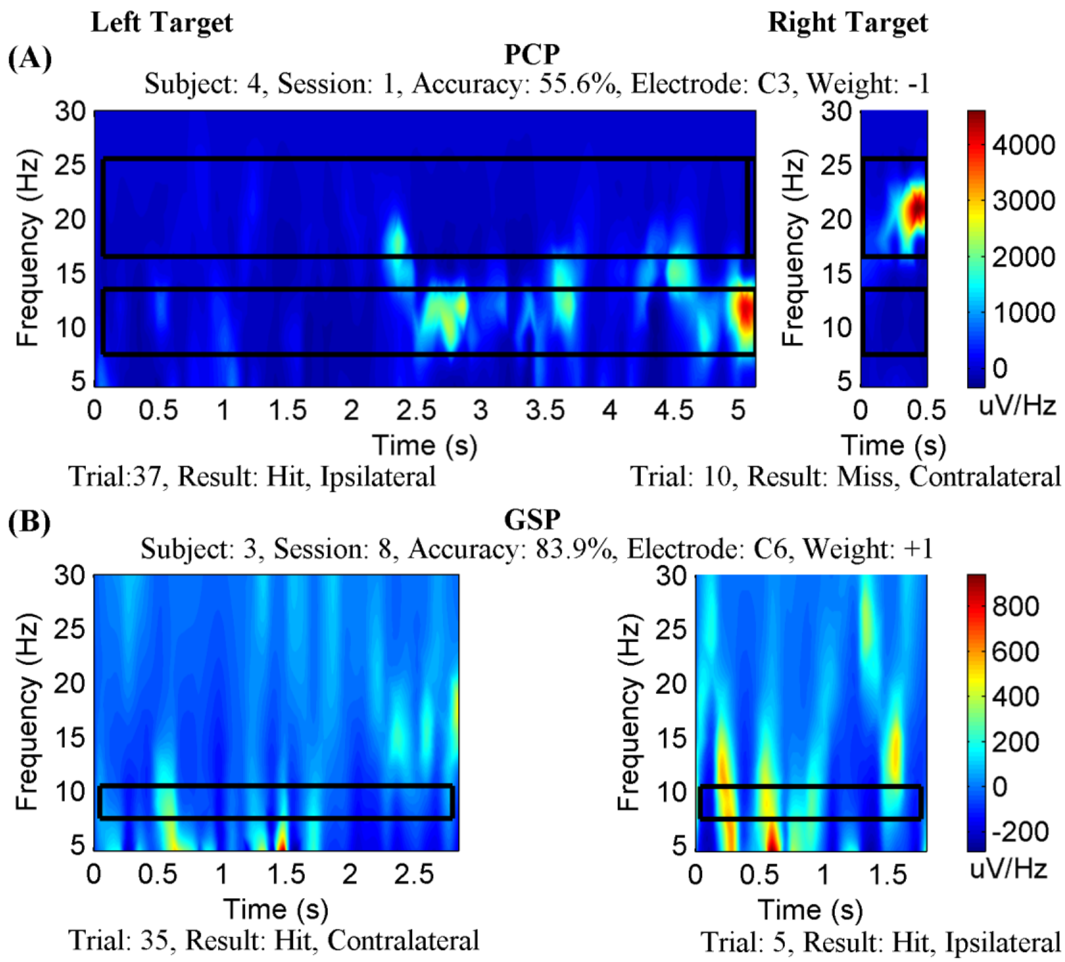

Figure 2.10 Time frequency plots of representative left and right trials
Time frequency plots of representative left and right trials of PCP (A) and GSP (B) illustrate the numerous factors influencing the EEG signal and subject performance. The plotted colour corresponds to the AR spectral amplitude minus baseline. Black squares on the plots indicate the frequencies used for control. Words centered across the figure apply to both right and left plots. Words below a figure apply to the individual plot. Color scale bars on the right apply to the entire row.

ALS, have reduced physical and mental endurance (Sykacek et al 2003, Birbaumer 2006). This diminished endurance has decreased the accuracy of a BCI system with 90% accuracy in healthy subjects to levels just over chance in the patient population (Sellers and Donchin 2006, Iversen et al 2008). Therefore, reduction of fatigue due to using goal selection as a control strategy may aid the usefulness and adoption of BCIs by individuals who truly need them to restore lost functionality.

Goal selection demonstrated many advantages over process control in this study. However, goal selection does have limitations. In order for the BCI system to assist the user, the system needs to be pre-programmed to provide the correct assistance. That requires that the situation be an anticipated, known event. The major advantage of process control is that it provides unlimited possibilities for action, making it indispensible when encountering a novel situation or event. Ideally, a BCI would be able to assist the user as often as possible using goal selection, while still allowing the freedom that process control provides. Additionally, an ideal BCI would learn from the new encounter to possibly provide assistance in the future. This main distinction between goal selection and process control implies that, however much benefit a BCI derives from implementing goal selection, process control will continue to be employed as BCI use increases in society.

2.2.5 Conclusion

This study confirmed the hypothesis that goal selection is more accurate, faster to use, easier to learn, and requires less mental effort than process control. This study validated previous findings concerning speed and accuracy in a larger sample size (Royer and He 2009, section 2.1). Median improvement from process control to goal selection across all sessions was 71% for accuracy, number of hits per run, time to hit, and information transfer rate. This study was also the first to show that goal selection is easier to learn and requires less mental effort than process control. Goal selection showed on average twice the learning and required 54% less effort than process control. If we wish to use BCIs to help individuals that can no longer rely on their own natural motor output, it will be important to make using the BCI as effective and as simple as possible. Applying goal selection in the BCI's control strategy will make the system easier to learn, decrease the training period, and provide improved speed, accuracy, and information transfer. These improvements will also help make BCIs more appealing to able-bodied users.

Chapter 3 Goal Selection in a More Complicated Application

In chapter 3, goal selection is applied to a much more complicated, real world scenario. In this paper, subjects are asked to fly a virtual helicopter to any point in three-dimensional space. The task was made simpler by using intelligent control strategies, such as goal selection. This chapter presents the first accomplishment by a non-invasive BCI of a task previously performed only by invasive BCIs. This was made possible by using intelligent control strategies. Section 3.5 presents additional, unpublished analysis on the influence goal selection had on the study.

3.1 Introduction

In the 1982 movie Firefox, a fighter jet's weapons were controlled by thought. Although the movie was science fiction, brain-computer interfaces (BCIs) are approaching such capabilities. Current invasive systems allow paralyzed humans to operate a computer (Hochberg et al 2006, Karim et al 2006, Kennedy et al 2000) and monkeys to select keys, move a computer cursor, and control a prosthetic arm to feed themselves (Taylor et al 2002, Velliste et al 2008, Musallem et al 2004, Santhanam et al 2006). Non-invasive systems, such as those based on scalp-recorded electroencephalography (EEG) (Wolpaw et al 2002, Vallabhaneni et al 2005) have been able to emulate the performance of invasive systems up to establishing control in two dimensions. This includes controlling a computer cursor (Wolpaw and McFarland 2004), a humanoid robot (Bell et al 2008), or a wheelchair in a real environment (Vanacker et al 2007, Galan et al 2008). Limited exploration of 3D space has been demonstrated by a non-invasive BCI (McFarland et al 2010). However, in that study, the targets were confined to the corners of a virtual box, the cursor was confined to the inside of the box, and did not exhibit the capacity for continuous control.

Full and continuous exploration of 3D space is an unaccomplished goal for non-invasive systems. Here we report our investigation which demonstrates that human subjects can fly a virtual helicopter to any point in 3D space using an EEG-based BCI. We achieved this result through a conventional four class system typically used for 2-dimensional (2D) control (Wolpaw et al 2002, Wolpaw and McFarland 2004, Yuan et al

2008), but used intelligent control strategies (Bell et al 2008, Vanacker et al 2007, Galan et al 2008, Royer and He 2009 (section 2.1), Kim et al 2006) to navigate in 3D space quickly and fluently. Non-invasive BCI control of continuous movement to any point in 3D space represents a major advancement in the field. The potential of this methodology extends beyond the scope of this study. It represents a new way of thinking about 3D control that expands the user population, reduces training barriers, and optimizes control signal economy. The material in sections 3.1-3.4 has been reprinted with permission from Royer et al 2010.

3.2 Methods

3.2.1 Participants

Four young, healthy, human subjects participated in the study. This study was conducted according to a human protocol approved by the Institutional Review Board of the University of Minnesota. None of the four subjects were particularly skilled in video games, or in virtual reality navigation. None identified themselves as gamers. All four had been trained in BCI usage in previous studies. Subjects began by using a 1D left/right cursor control BCI similar to that used in for 8 to 11 sessions (Royer and He 2009, section 2.1). Sessions lasted approximately two hours including capping and hair washing time. Time of actual BCI usage during a session was typically around 45 minutes. After mastering left/right, the subjects attempted 1D control of up/down. Once subjects were comfortable with up/down control, the subjects performed 2D cursor control before attempting to control the helicopter. Eye movements were monitored during training on the cursor task to ensure that subjects were not using eye movements to control the cursor. In total, subjects 1-4 completed 33, 31, 24, and 21 sessions, respectively, before the data presented here.

3.2.2 Data Acquisition and System Design

Subjects wore an EEG cap that recorded sensorimotor rhythms from motor imagination (Yuan et al 2010b, Wang et al 2004). EEG recording methods and processing in BCI2000 were the same as in previous studies (Yuan et al 2008, Royer and He 2009 (section 2.1)). In brief, a 64 channel EEG cap was plugged into a Neuroscan amplifier. The signal was then fed to a computer running BCI2000. We used BCI2000's standard 2-dimensional cursor task (Schalk et al 2004) to generate two control signals: left/right and up/down. Each subject's control signal was individualized during their training to those seen in Table I. Electrodes were limited to those recording sensorimotor cortex. The autoregressive spectral amplitude was calculated for each of the electrodes and frequency bins indicated. The left/right control signal was the subtraction of the electrodes and frequencies of the left hemisphere from the electrodes and frequencies of the right hemisphere. The up/down control signal was the inverted addition of left and right. Using this scheme, subjects imagined moving their right hand to go right, their left hand to go left, both hands for up, and rest for down to create a four class system. Each control signal was normalized to zero mean and unit variance. The control signals were continuously fed through a UDP port (every 20-50 ms) to a virtual world modeled in Blender. The subjects sat motionless in a comfortable chair in front of a flat screen computer monitor. They only saw the virtual world of Blender.

Table 3.1 Subject specific electrodes and frequencies used for control

Subject	Right	Left
1	C4: 12, 15, and 18 Hz CP2: 9 and 15 Hz	C3: 12 Hz
2	C6: 15 and 18 Hz	C3: 12 Hz
3	C4: 9, 12, and 15 Hz CP4: 9, 12, and 15 Hz	C3: 15 Hz C1: 15 Hz
4	C4: 9, 12, and 15 Hz CP4: 12 and 15 Hz	CP5: 18 Hz

Electrode names are given according to the 10-20 international system. For each cell above, each frequency bin designated by the center frequency listed was weighted equally with all the other frequency bins listed (bin width = 3 Hz). Left/right control in BCI2000 was the subtraction of the left column from the right column. Up/down control in BCI2000 was the addition of the left column to the right column.

3.2.3 Helicopter Control and Virtual World

Although the subjects were using two control signals, they were able to control the helicopter in 3D space. The left/right control signal from BCI2000 was used to control the helicopter's rotation (Fig. 1a). If the subject imagined left, the helicopter would turn to the left. If the subject imagined right, the helicopter would turn to the right. The up/down control signal from BCI2000 was used to control the helicopter's vertical position in space

(Fig. 1b). The helicopter had a constant forward speed of 0.5 Blender units per second (bu/s).

The subjects' control signals were continuously controlling the helicopter, but in a different manner for left/right versus up/down. The left/right control signal was linearly tied to the rotation of the helicopter, with gains that could be independently adjusted for right and left for each subject. An additional restraint on the rotation was that a turn was limited to a maximum of 20 degrees per screen update, which occurred at approximately 5 Hz. The up/down control signal was cubically tied to a force that acted upon the helicopter. The magnitude of the force could be independently adjusted for up and down for each

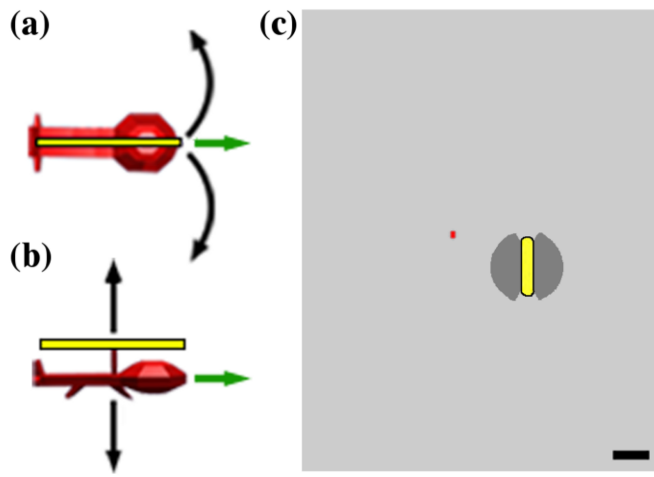

Figure 3.1 Control strategy and world size
(a) and (b) Two control signals adjust the helicopter's position in 3D space. The helicopter has a constant forward speed (green arrows). (a) Top view. The left/right control signal adjusts rotation (black arrows). (b) Side view. The up/down control signal adjusts elevation (black arrows). (c) Size comparison of ring (yellow), helicopter (red), cone of guidance (dark grey), and allowable space for the helicopter (light grey). The black scale bar in the lower right hand corner has a length of 2 Blender units (bu).

subject. Three of the four subjects preferred to dampen the effects of the force by reducing the helicopter's vertical linear velocity to 97% of its value each screen update. The up/down control signal was further limited by capping it to a range of -1.5 to 2, a reasonable range for a zero mean, unit variance signal. Subjects' left, right, up, and down gain coefficients were often the same from session to session, but were adjusted according to subject preference if the subject requested.

The virtual world was large and the helicopter was not confined to the ring space. The rings were confined to a space of 69 bu^3. The helicopter was allowed within a space of

4,285 bu^3, which is more than 600% of the volume occupied by the rings. However, if the helicopter did reach the edge of this large space (Fig. 1c), it was reset to the starting position. To give the subjects reference, the virtual world was based on the Northrop Mall area of the University of Minnesota (Fig. 2). The buildings were added to the virtual world mainly to give the subjects reference. They were not designed to be significant obstacles. Given that the helicopter could easily fly above all the buildings, they occupied minimal flying space. In fact, the buildings only occupied five percent of the total allowed volume.

3.2.4 Experimental Paradigm

Each session, subjects completed 7-13 five-minute runs. The exact number depended on the subject's availability. At the start of a run, the helicopter hovered stationary slightly above ground at the starting position for 3 seconds and then began to move forward. At that time, the subject gained control of the helicopter's motion. A run consisted of as many ring attempts as could be completed in a five-minute period. During each session, approximately half the runs ± 2 used the cone of guidance, a form of shared control. By sharing control between the subject and the BCI system (Bell et al 2008, Vanacker et al 2007, Galan et al 2008, Royer and He 2009 (section 2.1)), we were able to leverage the expertise of both brain and computer to create a system more powerful than either individually. The cone of guidance directed the helicopter through the ring during the final approach. It was implemented as a training aid and to reduce frustration. The cone of guidance was not visible to the subject. The area of assistance resembled a cone originating from the center of the ring with a radius of 2 bu. The upper and lower boundaries were defined such that the helicopter had to approach no more than 60 degrees from horizontally even with the ring. The side boundaries required that the helicopter approach the ring with a clear path through it. Each session alternated whether the first half or the second half of the session used the cone of guidance.

The subjects had to find the ring in a large space and then fly through it (Fig 1c). This can be compared to finding a needle in a haystack, and then threading it. Both tasks are difficult on their own, and combining the two could be very frustrating. The cone of guidance can be thought of as a needle threader. Effectively, the cone of guidance changed the ring into a balloon-shaped target shown in Fig. 1c. This target size was consistent with many other invasive and non-invasive BCI studies (Santhanam et al 2006, Wolpaw et al

2002, Vallabhaneni et al 2005, Wolpaw and McFarland 2004, McFarland et al 2010, Yuan et al 2008, Royer and He 2009 (section 2.1)). Each session, half the runs ± 2 used the cone of guidance. The exact number was chosen to minimize subject frustration. In some runs, the cone of guidance also prevented the user from hitting a building. Some users requested that they be allowed to hit the buildings since their strategy to maximize the number of rings flown through involved resetting upon colliding with a building when quite distant from the ring. As expected, the cone of guidance improved performance, but it was not necessary for successful performance. All subjects were able to fly through the ring without the assistance of the cone of guidance.

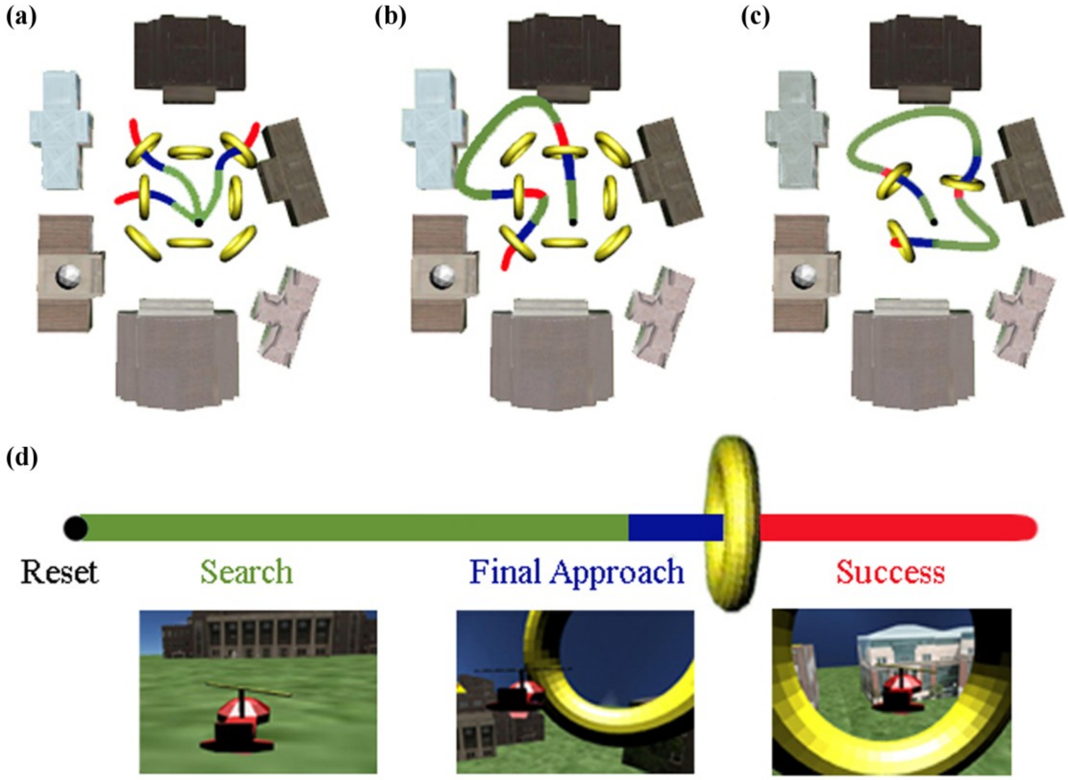

Figure 3.2 Experimental paradigms and virtual world
(a) - (d) The colors correspond to the different phases of a trial: black dot for reset, green for search, blue for final approach, and red for success. Reset is the time the helicopter was motionless. Success was the time after fly through, but before the ring changed position. (a) Trial-based discrete paradigm. After successful completion, the helicopter reset to the starting position. (b) Continuous discrete paradigm. The helicopter did not reset to the starting position after flying through a ring. (c) Continuous random paradigm. The rings could be anywhere within the original ring space at any vertical orientation. (d) Examples of what the subject saw during each phase of a trial.

The subjects controlled the helicopter under three different paradigms: trial-based discrete, continuous discrete and continuous random (Fig. 2a-c). The subjects began with trial-based discrete (Fig. 2a). From the starting position, the subjects attempted to find and then fly through a ring that was located at one of eight discrete positions. Only one ring was present at a time. A map similar to Fig. 2a-c in the lower portion of the subject's screen indicated the position of the ring. The subjects had 57 s to fly through the ring. Otherwise, the helicopter reset to the starting position. If a subject hit a building, the helicopter also reset to the starting position.

After completing 3 sessions of trial-based discrete (4 for subject 2), subjects then completed 3 sessions (4 for subject 2) of continuous discrete (Fig. 2b). The major difference between the two paradigms was that, in continuous discrete, the helicopter did not reset after passing through a ring. Instead, the helicopter kept flying as another of the 8 possible rings appeared. Subjects then completed 2 sessions (3 for subject 1) of continuous random (Fig. 2c). In this paradigm, the ring could appear anywhere within the box defined by the original 8 rings. The ring could be at any rotation about its vertical axis, but it had to be at least 3 bu from the previous ring. In both of the continuous paradigms, there was no time limit on how long a subject could attempt to pass through a given ring. In all three paradigms, subjects were instructed to pass through as many rings as possible in the five-minute run. Any discrepancy between subjects in number of sessions was due to subject availability.

Fig. 2d shows example images of what the subject saw during a trial. Additional features were on the subjects' monitor during a run. An "instrument panel" was located on the bottom of the screen which contained the map mentioned above, a joystick whose position corresponded to the position of the cursor on the 2D BCI2000 screen, the number of rings flown through during the run, and the time remaining in the trial. Small white numbers indicating programming variables such as BCI2000 source time and target code were located in the upper left hand corner of the screen for the experimenter's reference. These variables, 4 to 6 depending on paradigm, were designed to be as unobtrusive as possible to the subject. Personal monitoring of the EEG during the helicopter experiments showed that processing the complex environment was not affecting the control signal.

3.2.5 Experimental Controls

In order to gauge the subject's performance, we implemented two experimental controls. Each session started with the subject playing the game using the keyboard for approximately 5 minutes. The keyboard controlled the helicopter in the same manner as the BCI, with rotation controlled by the left/right arrows, and vertical forces controlled by the up/down arrows. This allowed us to analyze the degree of control subjects were able to achieve using the BCI as compared to when they had standard control using the keyboard. The subjects completed 21, 13, and 9 runs of trial-based discrete, continuous discrete, and continuous random, respectively. That resulted in 451, 267, and 236 rings being flown through for the three paradigms, respectively. For the second experimental control, the helicopter flew for two sessions of ten runs of each paradigm with environmental noise as the input to BCI2000. This was achieved by running Neuroscan and BCI2000 without an EEG cap attached to the amplifier. This generated two random control signals with zero mean and unit variance. These sessions served as our chance performance. Chance performance was able to fly through a few rings during its 20 runs of each paradigm: 28, 18, and 28 for trial-based discrete, continuous discrete, and continuous random, respectively.

3.2.6 Data Analysis

Average path to a ring plots (Fig. 3) present averages of all paths to a ring that concluded with the helicopter passing through the ring (946 total trials of 1370). Trials that ended in a reset (297) or a time out (127) were excluded from Fig. 3. All paths were resampled to 1001 points for the purposes of averaging and other data analysis. The standard error of the mean was calculated for the x, y, and z coordinate for each of the 1001 plotted points based on the 946 trials averaged. The width of the shaded area in Fig. 3 at each of the 1001 plotted points is the standard error of the mean averaged over the 3 dimensions, x, y, and z, for that particular point. The following measures were not normal according to a 2-sided Lillilifors test: the percent of flight time spent at each normalized distance to the ring, information transfer rate, rings/min, time to ring, and path length. The statistical analysis performed on those measures was a 2-sided sign test. Alpha = 0.05 for all tests. The percent of flight time spent at each normalized distance to the ring is plotted as a cumulative percent. The metric was calculated for distance intervals representing 1% of normalized distance where the starting distance to the ring was defined as a distance of 1 for all trials. Statistics were calculated for each 1% distance bin.

Information transfer rate and rings/min were calculated for each run of each subject, for a total of 132, 112, and 87 values for trial-based discrete, continuous discrete, and continuous random, respectively. Time to ring and path length were calculated for each ring successfully flown through, for a total of 946, 566, and 608 for the three paradigms. Since chance successfully flew through so few rings in comparison to the keyboard and BCI data, a statistical comparison against chance is not meaningful. Thus, the measures calculated for each ring are not presented for chance. Each ring successfully flown through counted as 3 bits of information transferred for the discrete paradigms. Time to ring and path length were normalized to the keyboard path length for trial-based discrete. Otherwise, they were normalized to the starting distance to the ring.

The subject's goal was to fly through as many rings as possible. Some subjects used resets as part of this strategy. If these subjects wandered off and found themselves very far away from the ring, they would purposely hit a building or the bounding box to reset. This way, they would not need to spend the time flying back to the ring and could go through more rings in the same run. Because this strategy was not penalized, presenting the number of rings obtained as a fraction of total trials is unfair. A better metric, and the one presented in Table I, is the number of rings flown through per minute. Another fair metric is the percent of total flight time spent at each normalized distance to the ring, as presented in Fig. 5.

3.3 Results

All four subjects successfully navigated in 3D space. A supplementary video of each subject is available at http://ieeexplore.ieee.org. The videos feature all three paradigms. Fig. 3 shows the trial-based discrete group average path (Fig. 3a-b) to each of the 8 rings. The paths closely resemble the paths the subjects took when they were using the keyboard (Fig. 3c-d). The cross-correlation at 0 delay between the average BCI and keyboard paths was greater than 0.99 for all 8 rings. For individual subjects, the average cross-correlation for all 8 paths was 0.9946, 0.9940, 0.9915, and 0.9664 for subjects 1 through 4, respectively.

Subjects achieved movement control by successfully controlling their EEG signal in a goal-directed manner. Fig. 4 plots the position of the helicopter and the BCI control signals during subject 1's trial-based discrete, no cone of guidance trials shown in supplementary video 1. The combination of video 1 and Fig. 4 show that the subject appropriately modulated their EEG to create a control signal that flew the helicopter in a very goal directed manner through the ring. In five of the seven trials, the subject progressed swiftly through the ring. In the third trial, the subject started flying to the right, and then realized that the target was actually off to the left. When they had circled to the left, the subject realized the helicopter was too high, so the subject descended. In the fifth trial, the subject missed on their first approach to the ring, and so had to circle around for two more

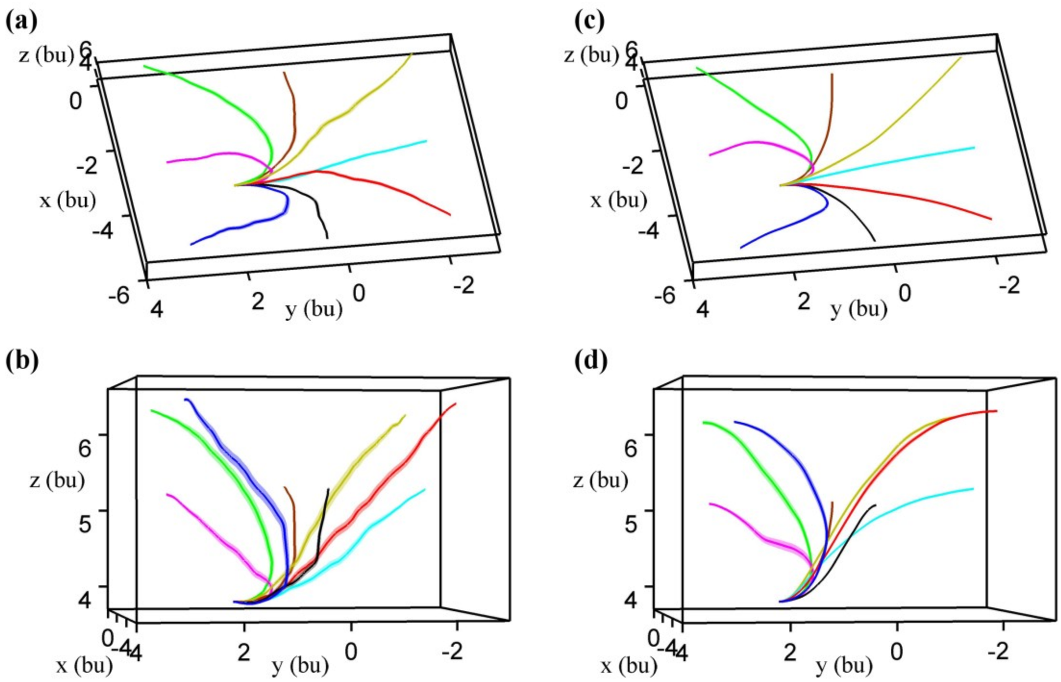

Figure 3.3 Average paths in trial-based discrete correlate strongly between BCI and keyboard control

(a) - (d) The path to an individual ring has the same color throughout. The shaded area represents the standard error of the mean around average trajectories (solid line) to each ring. bu indicates Blender units. (a) and (b) The group average BCI paths when viewed from above (a) and the side (b) with a slight rotation. Rotate Fig. 2a-c ~90 degrees clockwise to achieve the orientation of (a) and (c). n = 946 total trials, with 106-131 per ring. (c) and (d) The group average keyboard paths viewed from the same perspectives as in (a) and (b). n = 451 total trials, with 53-62 per ring.

attempts before successfully flying through the ring. The other three videos show similar purposeful control from the other three subjects. Fig. 4 shows that the control of the helicopter as seen in video 1 is a direct result of modulation of the EEG creating the control signal since Blender translated the control signal to helicopter position with such high fidelity. In Fig. 4, the cross-correlation between the helicopter's heading and the subject's left/right control signal is 0.997 at a delay of 0. The cross-correlation between the vertical position and up/down control signal is 0.817 at a delay of 0 and reaches a max of 0.837 at a delay of approximately 2.1s. Between the videos and Fig. 4, we show that the subjects were purposely modulating their EEG to control the flight of the helicopter.

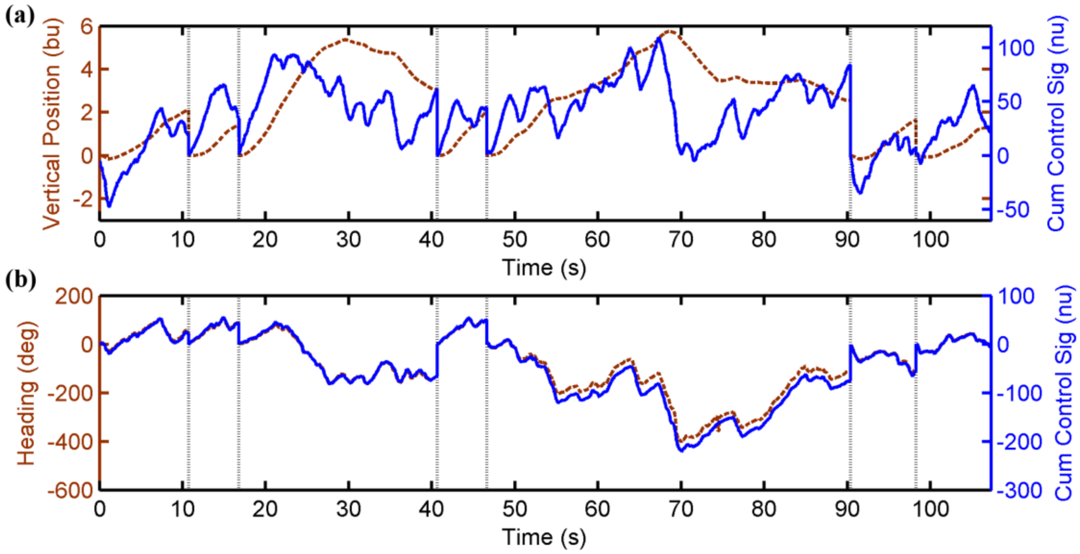

Figure 3.4 Control signal and helicopter position/heading correlate strongly
(a) and (b) Vertical lines indicate different trials shown in Supplemental Movie S1 of subject 1 performing trial-based discrete. (a) The cumulative up/down control signal (right axis, solid blue line) plotted with the vertical position (left axis, dotted brown line) as a function of time. (b) The cumulative left/right control signal (right axis, solid blue line) plotted with the helicopter's heading (left axis, dotted brown line) as a function of time. Upward slope indicates rightward rotation. Downward slope indicates leftward rotation. bu = Blender units. deg = degrees. nu = normalized units.

We assessed the quality of the movement control through several measures. We first analyzed the percent of time the helicopter spent closer to the ring than it was upon ring presentation. Fig. 5a presents the grouped data. In trial-based discrete, the helicopter spent 93 % of flight time closer to the ring than it started when controlled by the keyboard. It spent 67% of flight time closer to the ring than it started when controlled by the BCI, and 44% when controlled by chance, or the random controller. The same data for continuous discrete are 76%, 63%, and 38% when controlled by keyboard, BCI, and chance, respectively. For continuous random the data are 81%, 59%, and 34%. Since the helicopter had a constant forward speed, subjects were forced to move away from rings behind them while turning around. This is reflected in the fact that the keyboard data do not achieve 100%. Individual subject data is shown in Fig. 5b. For all three paradigms, subjects performed significantly better than chance. For some distance intervals, subject performance was statistically the same whether using the BCI or keyboard. This was true of subject 1 for 34% of the distance intervals between 0 and 2 in continuous discrete (all intervals were greater than or equal to 1.24) and 14.5% of continuous random (all intervals were greater than or equal to 1.26). This was true of subject 2 for 26.5% of continuous discrete (all intervals were greater than or equal to 1.48) and 16% of continuous random (1.27-1.58). For the remaining distance intervals, subject performance was not statistically the same whether using the BCI or keyboard. Also of interest, in continuous discrete, subjects 1 and 2's BCI performance approached subjects 3 and 4's keyboard performance, as can be seen by the proximity of the blue and red dotted lines to the green and orange solid lines.

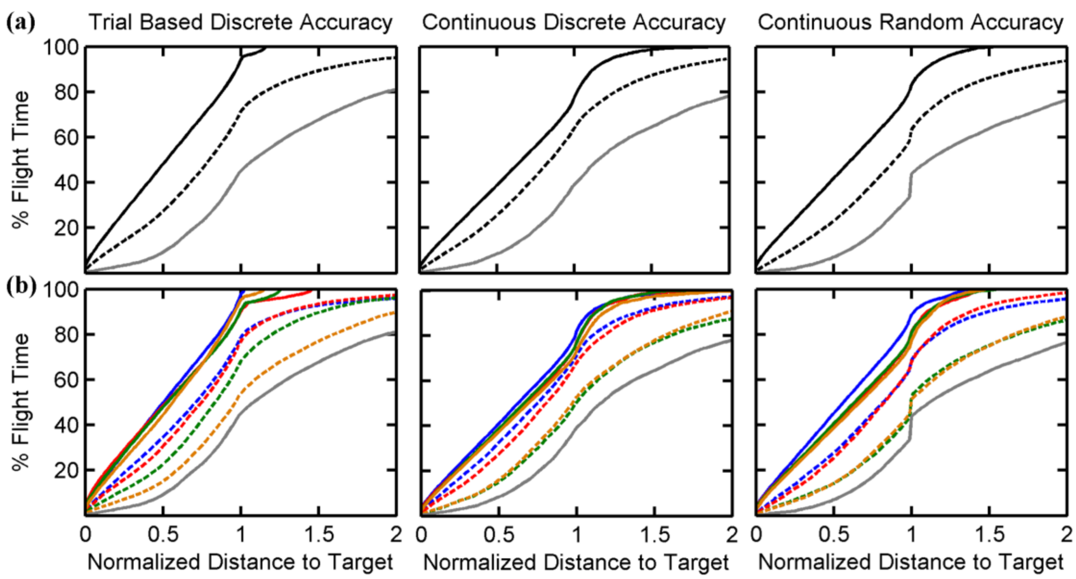

Figure 3.5 Subjects effectively and accurately controlled the helicopter in all paradigms
(a) and (b) Percent of total flight time spent at or closer to each normalized distance to the target. 1 represents starting position. 0 represents ring position. Solid black or colored lines indicate keyboard performance. Dotted lines indicate BCI performance. Solid grey lines indicate chance. Each column presents the data for the labeled paradigm. (a) Grouped data. (b) Individual subject data. Subject 1's data are in blue, subject 2 in red, subject 3 in green, and subject 4 in orange.

Additional quality measures assessed include information transfer rate (ITR), number of rings/min, time to ring, and path length. Fig. 6 shows all subjects' data for these measures. The ITR and rings/min using the BCI were over 5 and 6.5 times that of chance, respectively. All subjects' ITRs and rings/min were significantly better than chance. The range of ITRs and rings/min achieved with the BCI overlapped the range achieved with the keyboard. In fact, in trial-based discrete, Subjects 2 and 3 achieved an ITR using BCI which was not significantly different than that with keyboard control. In the grouped data, the time to ring was only 34% longer with the BCI than with the keyboard, with a path length 38% longer for trial-based discrete. The continuous paradigms' time to ring were 70% longer using the BCI than using the keyboard, with a path length roughly 90% longer.

For both time to ring and path length, subjects performed slightly better in continuous discrete than continuous random, when compared to keyboard performance.

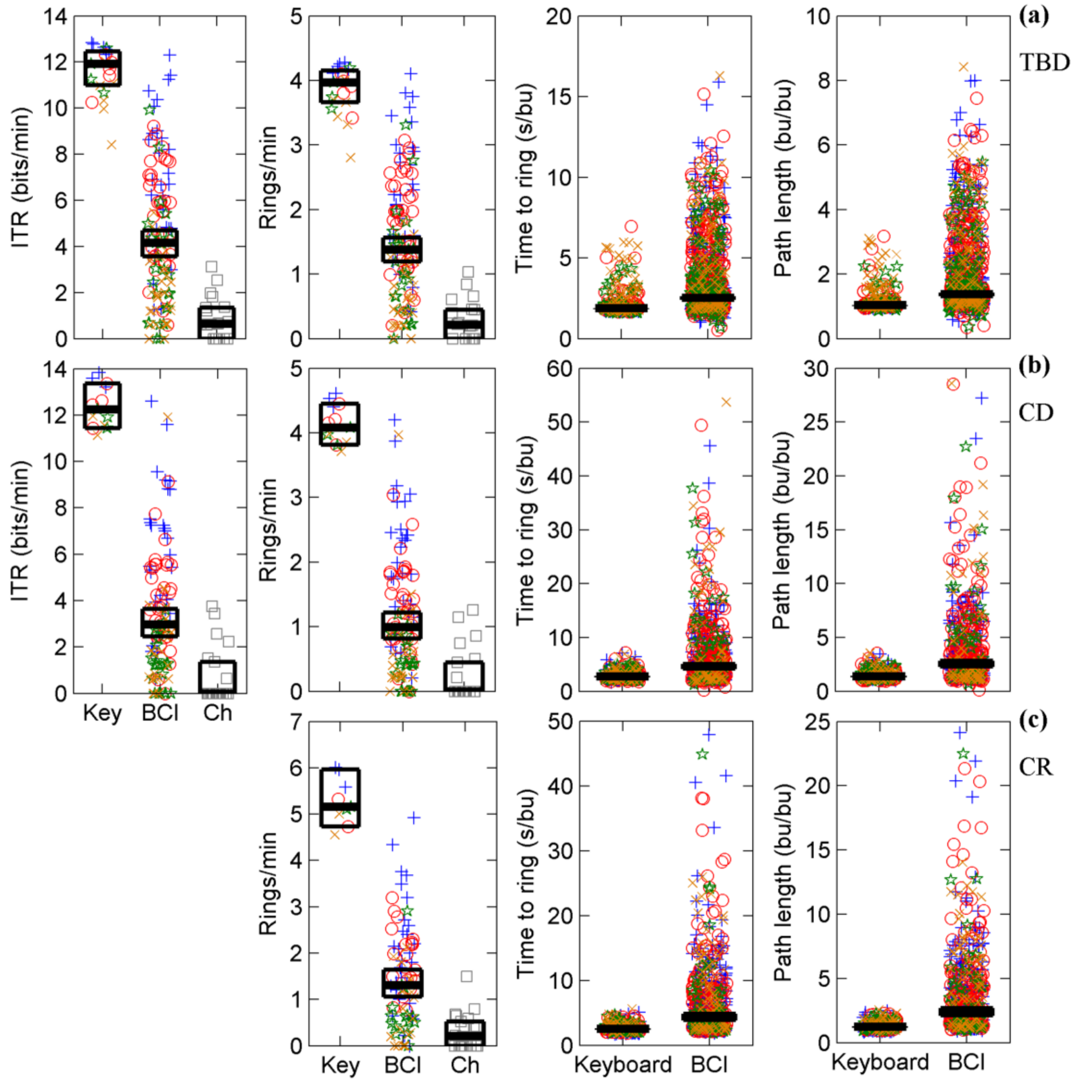

Figure 3.6 Movement quality measures
(a) - (c) Scatter plots of the data by subject and paradigm. Each row shows the paradigm labeled on the right. Subject 1 is represented by a blue +, subject 2 by a red O, subject 3 by a green star, and subject 4 by an orange x. Chance performance is represented by grey squares. Group medians are indicated by the thick black line. The black box indicates the 95% confidence interval of the median. Median and box are present on all plots but may not be distinguishable from each other due to proximity. ITR = Information transfer rate, Key = Keyboard, Ch = Chance, bu = Blender units. (a) Data for trial-based discrete, TBD. (b) Data for continuous discrete, CD. (c) Data for continuous random, CR.

The above results show that continuous control is harder than trial-based control, even when controlling for distance travelled. Continuous discrete had a median keyboard path length over 41% longer than trial-based discrete, whereas continuous random had a median keyboard path length only 4% longer. (Median keyboard path lengths were 4.95bu, 7.00bu, and 5.16bu for trial-based discrete, continuous discrete, and continuous random, respectively.) However, both paradigms had similar BCI-to-keyboard ratios that were higher than those for trial-based discrete.

3.4 Discussion

In this study, the challenge of navigating 3D space was surmounted through novel means. Subject expertise in 2D control was leveraged by identifying the crucial components of 3D flight and directing control capabilities to their accomplishment. In the field of BCI, this reductionist method is revolutionary in its efficient approach to 3D navigation. This implementation opens the doors of 3D BCI navigation by reframing it as a 2D control problem. Through these means we have developed a system that allows subjects to fluently navigate 3D space. The system requires minimal training beyond an established 2D BCI of any variety. Since the system does not require 3D expertise, subjects who can master 2D but struggle with 3D can now navigate in 3D space. It allows subjects to reach targets that require both positional and rotational control accuracy while providing control in both trial-based and continuous control scenarios. The potential of this methodology extends beyond the scope of this study. It represents a new way of thinking about 3D control that expands the user population, reduces training barriers, and optimizes control signal economy.

Some may find the control strategy in this study limited in its lack of a stop command. However, an airplane in flight is a common real world example of a system with similar limitations. If a plane arrives at its destination and needs to wait before it can land, it begins a holding pattern, effectively circling around the intended destination. Our subjects exhibited similar behavior and circled around if they missed the ring on their first attempt. In that way, the control exhibited by the subjects emulated the control available to an airplane pilot once in the air. The goal of this study was to navigate within 3D space and not address the takeoff or landing. Such improvements could be a goal for future work. The current study only utilized left and right hand motor imagery. Other types of motor imagery could be utilized to provide additional control (Wolpaw et al 2002).

Subjects in this study mastered 3D space in a manner that has not been previously accomplished by a non-invasive BCI. Not only did the subjects demonstrate positional and rotational control, but also they did not have the assistance of being confined to the same space as their targets. In this study, the 3D space through which subjects navigated was much larger than the space occupied by the targets. The subjects were not able to navigate to a wall and then follow the wall to their desired location. That added a great amount of complexity to the task when compared to previous studies (McFarland et al 2010). Furthermore, subjects demonstrated continuous control, going directly from one ring to another. Finally, subjects successfully flew through rings that were randomly placed within the target space.

In a recent study, three dimensional hand movements were reconstructed from non-invasive EEG recordings (Bradberry et al 2010). Their work could serve as the foundation for yet another 3D BCI. However, this research is still in preliminary stages. In this study, actual hand motions were reconstructed after data collection in an offline system. The algorithm has not yet run in real time, which is necessary in order to provide feedback to a BCI user. It will also be interesting to see if motor imagery can be decoded with the same accuracy.

Subjects in the current study were able to master 3D space more easily because of the assistance of the intelligent control strategies used by the BCI. In this study, the subject and the BCI shared control when using the cone of guidance. As expected, the cone of guidance improved performance. However, it was not necessary for successful performance. For example, the cross-correlation in the grouped data between the average BCI and keyboard path went from 0.990 without using the cone of guidance to 0.996 when using the cone of guidance.

An interesting implication of the constant forward velocity of the helicopter is that timing became more important for successful task completion. In order to successfully fly through a ring, the subjects had to not only navigate to the ring, but also approach the ring from the correct angle. If the subject was approaching from the side, they had to wait for the right moment to turn through the ring. Similarly, subjects had to be at the right altitude before arriving at the ring. The subjects were well acquainted with this constraint. Many of the delays in flying through the ring were caused by either misjudging the time required to maneuver, or executing the wrong maneuver at the worst possible time. In those cases, the subjects circled around and tried again. The timing factor can be seen in Fig. 4a. In that figure, peak correlation between control signal and helicopter position occurred at a delay of approximately 2.1s. This was interesting since the subject had commented on how they felt as if they always had to be planning about 2s ahead. When the subjects were using the cone of guidance, the subjects still had to navigate to the ring. However, the timing constraints were not quite as challenging since the actions at pivotal moments were executed for the subjects. As seen in the improvement in cross-correlation between BCI and keyboard path given above, this was the primary cause of performance improvement using the cone of guidance.

Another interesting timing constraint was seen in the transition from trial-based control to continuous control. As seen in the BCI to keyboard ratios, the subjects found continuous control harder than trial-based, even when distances were similar. One possible reason for this, as mentioned in interactions with the subjects, was that continuous control did not provide any mental breaks. An additional advantage of the cone of guidance was that, while the helicopter was under the control of the BCI system, the user could mentally relax and not attempt to control their EEG. Our personal experience with the subjects made it clear that they took advantage of this opportunity to take a small mental break until the next ring was presented.

By using the cone of guidance, we were able to leverage the expertise of both brain and computer to create a system more powerful than either individually. Shared control is widely used in systems that we interact with daily. These systems are multidisciplinary solutions to complex goals. Anti-lock braking systems in vehicles perform a function the human driver performed in older car models. The car pumps the brakes faster and more efficiently than the driver, improving safety. Point and shoot cameras dominate the camera market, allowing any person to take quality pictures without detailed knowledge of f-stops and shutter speeds. Even real helicopters have a variety of shared control functions such as landing assistance and obstacle avoidance. Likewise, brain-computer interfaces are intrinsically multi-disciplinary. Neuroscience, engineering, and computer science combine to create a complex system. The future of BCIs lies in leveraging the potential of all disciplines involved.

Not only is shared control used in everyday systems, shared control has been previously used in BCI research (Bell et al 2008, Vanacker et al 2007, Galan et al 2008, Royer and He 2009 (section 2.1)). In one study, subjects drove a real wheelchair around obstacles (Galan et al 2008). The intelligent wheelchair used environmental sensors and shared control to ensure obstacle avoidance and safe driving. Although limited to 2D due to the constraints of a wheelchair, this study demonstrated that shared control can be used by a BCI in the real world to improve performance and safety.

In conclusion, we achieved movement to any point in 3D space using scalp-recorded EEG. Prior to this study, such navigation was only possible through invasive means. The building blocks of this system, rotational control (Vanacker et al 2007, Galan et al 2008, Scherer et al 2007, Ron-Angevin and Diaz-Estrella 2009), virtual environments (Karim et al 2006, Kennedy et al 2000, Wolpaw et al 2002, Vallabhaneni et al 2005, Wolpaw and McFarland 2004, Vanacker et al 2007, Galan et al 2008, McFarland et al 2010, Yuan et al 2008, Royer and He 2009 (section 2.1)), and continuous control (Kim et al 2006) are well-established in the field of BCI. Through the synthesis of these elements, we were able to create a system capable of quickly and fluently navigating 3D space. The system's efficient use of control signals allows a subject trained in 2D control to be directly translated to our 3D control system with little additional training. This work provides a platform for the development of 3D non-invasive BCIs that are open to a wide subject population. The three-dimensional world we live in demands such functionality from BCI systems. Here we demonstrate the potential of non-invasive systems to meet that demand.

3.5 Further Analysis on the Impact of the Cone of Guidance

3.5.1 Introduction

In this subsection, we further describe the influence of the cone of guidance on the study. This was not included in the published work for fear it would hurt the acceptance of the overall accomplishment of navigating to any point in 3-dimensional space. In this section, similar figures are presented as in section 3.3, but these break out the runs and trials into 3 categories: those that used the cone of guidance, those that did not use the cone of guidance, and those that did not use the cone of guidance but had the effects of the cone added in post analysis. These results show that the subjects had general control of the helicopter sufficient to navigate to any point in 3D space, as argued by the previous sections. However, the cone of guidance helped with fine control, influencing the metrics of the weaker subjects more than the stronger subjects.

3.5.2 Methods

The methods to calculate the average path, accuracy, rings per minute, time to ring, and path length were identical to those described in section 3.2, with the added addition of separating out the runs and trials into those that did use the cone of guidance and those that did not. As described in section 3.2, the information transfer rate for trial-based discrete and continuous discrete was purely the rings per minute multiplied by 3. Therefore, information transfer rate contains no new information and will not be presented here.

Figures 3.9 through 3.14 present the data broken down into the three categories of those runs and trials that: used the cone of guidance, did not use the cone of guidance, and did not use the cone of guidance but had the effects of the cone added in post analysis. These three categories are referred to as cone (C), no cone (NC), and post cone (PC), respectively.

The post cone category was implemented in as simple of a manner as possible. All sessions (except 1 trial-based discrete session from subject 4) included in the data a variable that indicated whether or not the helicopter was in the cone of guidance. Figure 3.2 visually indicates this portion of the helicopter's path as the blue colored final approach. We have this variable for both when the subject was using the cone of guidance and for when the subject was not. All trials were split into the time before entering the cone of guidance and the time after entering the cone of guidance. If a trial never entered the cone, the entire trial occurred before the split. The time and distance were calculated for both before and after entering the cone in the same manner as calculated for the entire trial and described in section 3.2.

The post cone data were calculated from the no cone runs. A post cone trial was exactly the same as the no cone trial up until the helicopter first entered the cone of guidance. At that point, a tail was added to the trial so that the trial would count as if the helicopter successfully flew through the ring. In the trial-based discrete paradigm, this was only done if enough time remained to successfully fly through the ring within the time limit of the trial. The tail had a time and distance equal to the median time and distance of the trials in the same paradigm that used the cone of guidance and ended with the helicopter successfully flying through the ring. This tail was applied even if the original trial ended with the helicopter flying through the ring.

The new post cone trial had the elapsed time of the time before the cone plus the time of the tail; the distance travelled was the distance before the cone plus the distance of the tail. If the helicopter never entered the cone of guidance, or did not have time to successfully fly through the ring, the trial was unchanged and the original times, distances, and result were used. The measures plotted in figures 3.10 through 3.14 were then calculated for post cone as described in section 3.2.

3.5.3 Results

3.5.3.1 Subjects had general control of the helicopter sufficient to navigate in 3D space

As previously stated, the cone of guidance was not necessary for successful task completion. All subjects were able to successfully fly through multiple rings in all paradigms without using the cone of guidance. Subjects had general control to effectively

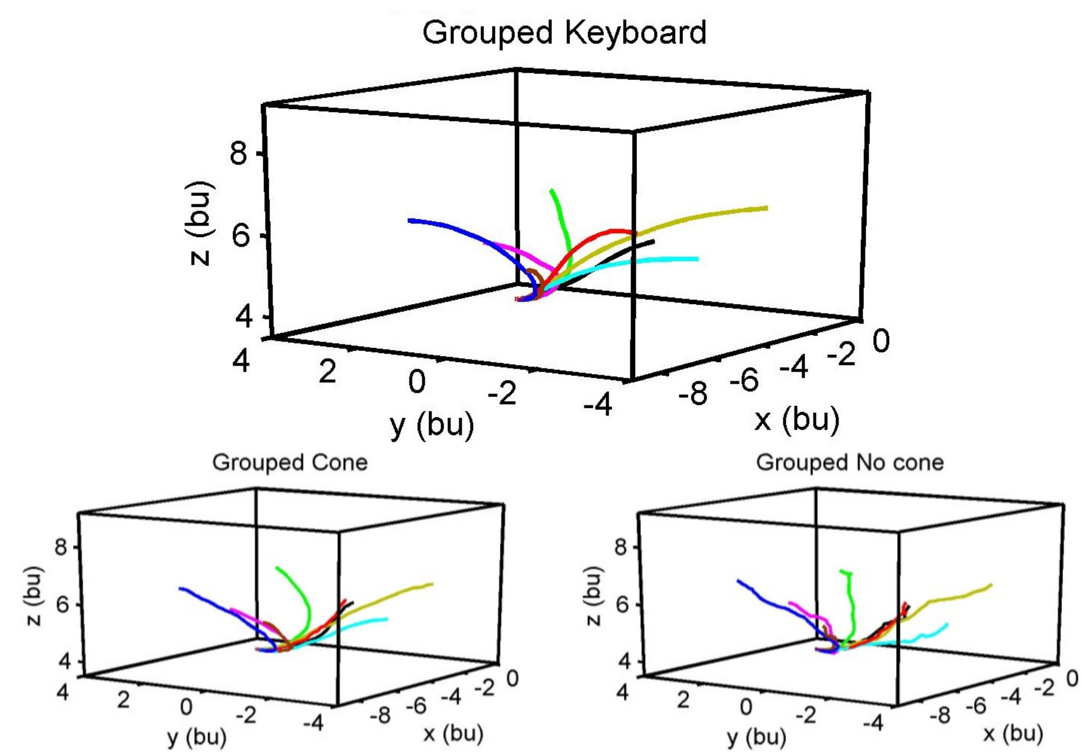

Figure 3.7 Average paths in trial-based discrete while using the keyboard, the cone of guidance, and no cone.

fly the helicopter to the ring in a large space. To demonstrate this point, similar to figure 3.3, figure 3.7 shows the trial-based discrete group average path for when subjects were using the keyboard, cone of guidance, or no cone. The paths closely resemble each other.

The average cross-correlation for all 8 paths at 0 delay between the average BCI and keyboard paths is shown in table 3.2. Little improvement in correlation was seen when using the cone of guidance for three of the four subjects, with the overall group change negligible.

Table 3.2 Cross correlation to the keyboard path

Subject	*Cone*	*No Cone*	*% Improvement*
S1	0.995	0.982	1.3%
S2	0.994	0.981	1.3%
S3	0.990	0.961	3.1%
S4	0.966	0.780	23.8%
Grouped	0.996	0.990	0.6%

Figure 3.8 presents another measure that shows that subjects had general control, or the percent of time the helicopter spent closer to the ring than it was upon ring presentation. This is the same measure presented in figure 3.5. Figure 3.8 plots the grouped accuracy for trial based discrete, continuous discrete, and continuous random, as well as the percent improvement from no cone to cone for the four subjects individually, grouped, and chance. Both discrete paradigms show almost no difference in accuracy between BCI cone and no cone, only showing about 3% improvement. The cone of guidance made a slightly larger impact in continuous random, with a grouped improvement of 11%. When the data from each subject are considered individually, it can be seen that the cone of guidance made very little difference in overall subject accuracy. Subject 4's accuracy was the most improved using the cone of guidance, as was subject 3's continuous random accuracy. However, of those four bars, only subject 4's trial-based discrete accuracy was significantly different between using the cone and not using it.

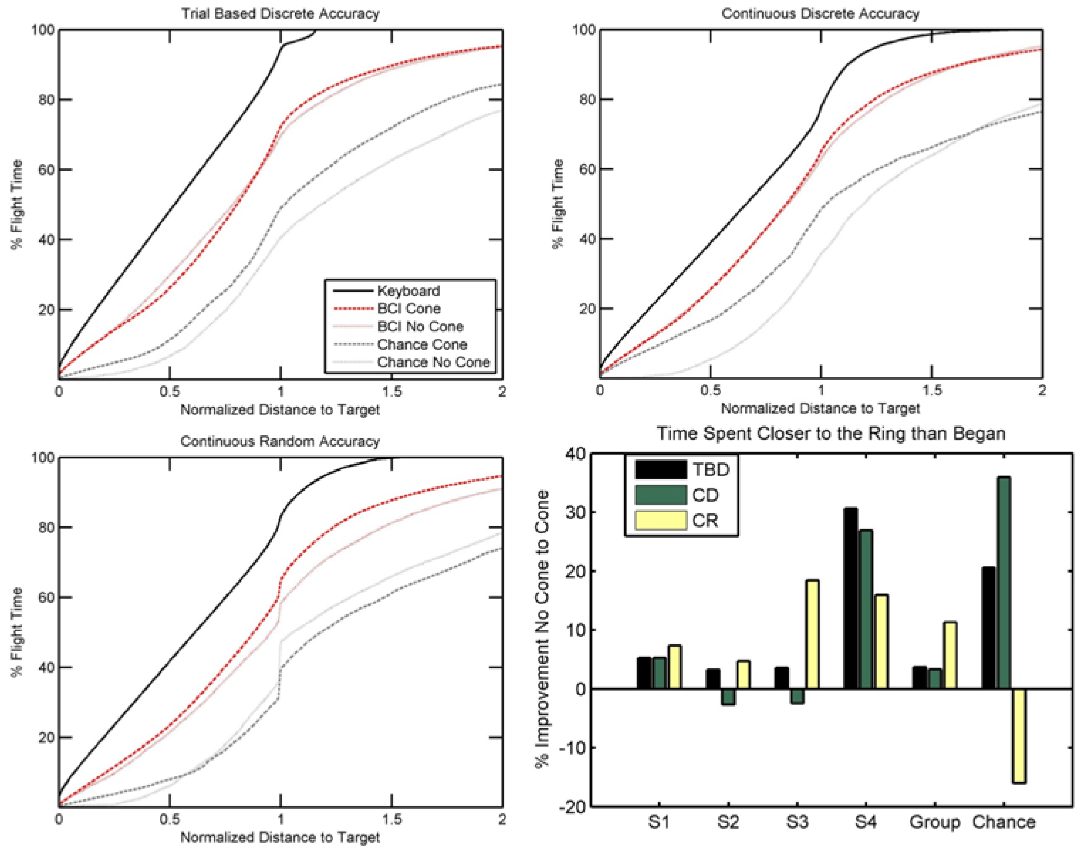

Figure 3.8 Accuracy while using the keyboard, cone of guidance, and no cone

3.5.3.2 The cone of guidance improved fine control for weaker subjects

As discussed in the previous sections, flying through the ring once arriving at it was a difficult task that could be compared to threading a needle after finding it in a haystack. Flying through the ring required precise timing and exact positional and rotational control. The cone of guidance could be viewed as a needle threader, or a tool that reduced the amount of precision and control required to complete the task. The cone of guidance had the most impact on fine control, or the ability of the subjects to successfully fly through the ring once arriving at it. What if the cone of guidance had been used during the no cone runs? How would that have changed successful completion of the trial?

Figure 3.5.3 addresses those questions by applying the effects of the cone of guidance to the no cone trials in post analysis, called post cone. The figure describes how post cone changed the end result of the no cone trials, split out by paradigm. In a trial-based discrete trial, the end result could be successfully flying through the ring, resetting, or a timeout. Post cone could change a reset or a timeout to a hit. The left panel of figure 3.9 has 3 bars for each of the subjects, plus the grouped data. The first bar contains the no cone end result of ring, reset, or timeout. The second bar subdivides the reset and timeout trials into those that entered the cone of guidance with enough time to fly through the ring and those trials that did not. The third bar presents the post cone data, where those trials identified in the middle bar as entering the cone of guidance with enough time to fly through the ring had an end result of ring. The middle and right panel present the same data for continuous discrete and continuous random. Since the continuous paradigms had no time limit on individual trials, there were no timeouts. The percent of no cone resets and timeouts converted to post cone rings varied by subject. Overall, considering both resets and timeouts in the grouped data, 51%, 47%, and 40% of all non-ring trials were converted to a successful ring for trial-based discrete, continuous discrete, and continuous random, respectively. In those converted trials, the subjects had successfully flown to the ring, but were not able to fly through on their own. The cone of guidance would have provided the assistance they needed to complete the task.

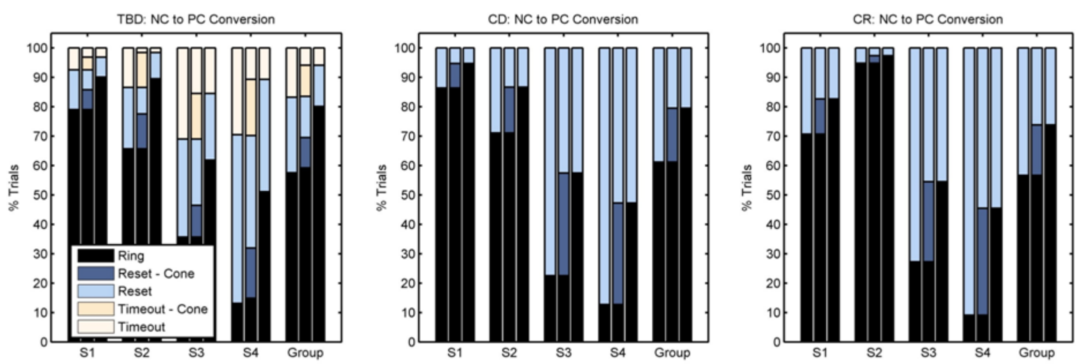

Figure 3.9 The end result of no cone trials if the cone of guidance had been used

Further effects of the cone of guidance can be seen in the number of rings per minute, time to ring, and path length measures. These measures are presented as in figure 3.6, but are further broken down into cone, no cone, and post cone.

Figure 3.10 presents the number of rings per minute for the three paradigms as well as the percent improvement from no cone to cone and no cone to post cone for all four subjects individually and grouped. The asterisks in the percent improvement graph indicate a significant difference between the two measures represented by the bar. The cone of guidance significantly improved the number of rings per minute for all subjects and for all paradigms in the grouped data. The average percent improvement across all paradigms for the grouped data in the number of rings per minute using the cone of guidance was 111%, with some subjects having an average improvement as high as 571%. Using the cone of guidance, subjects were able to achieve the same number of rings per minute as they were with the keyboard, as shown by the overlap in ranges.

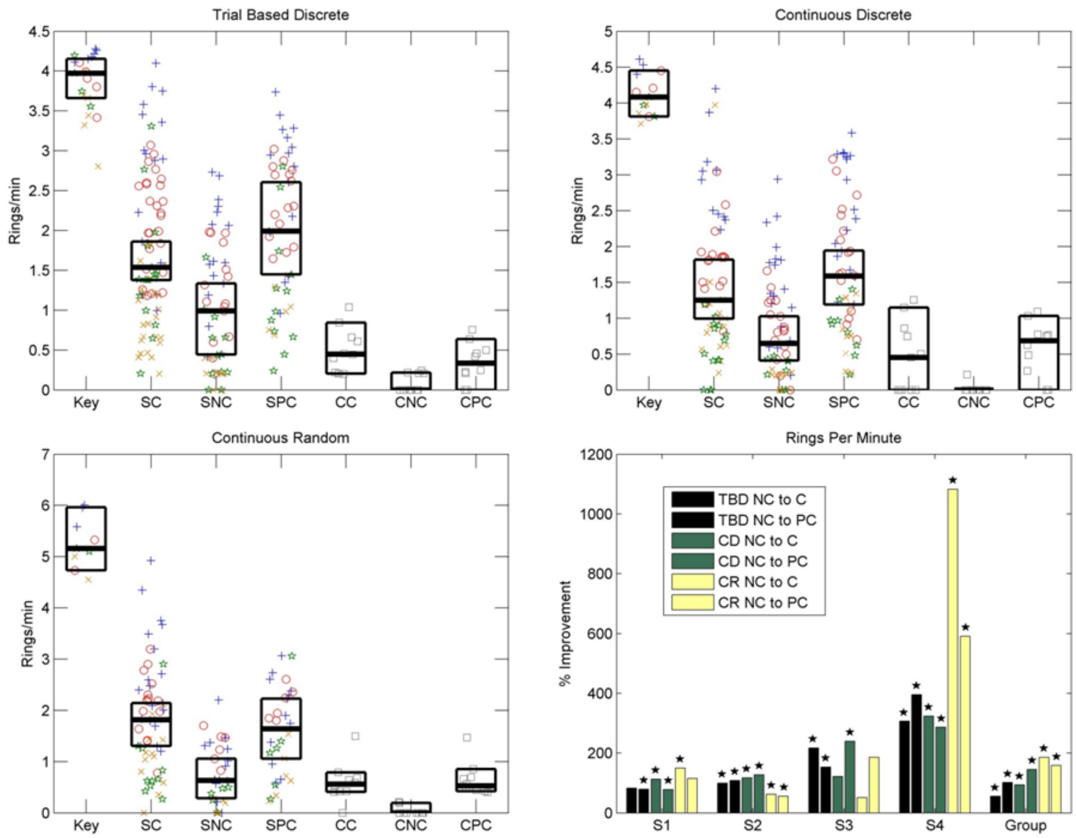

Figure 3.10 Rings per minute broken down by cone, no cone, and post cone for the subjects (S), and chance (C)

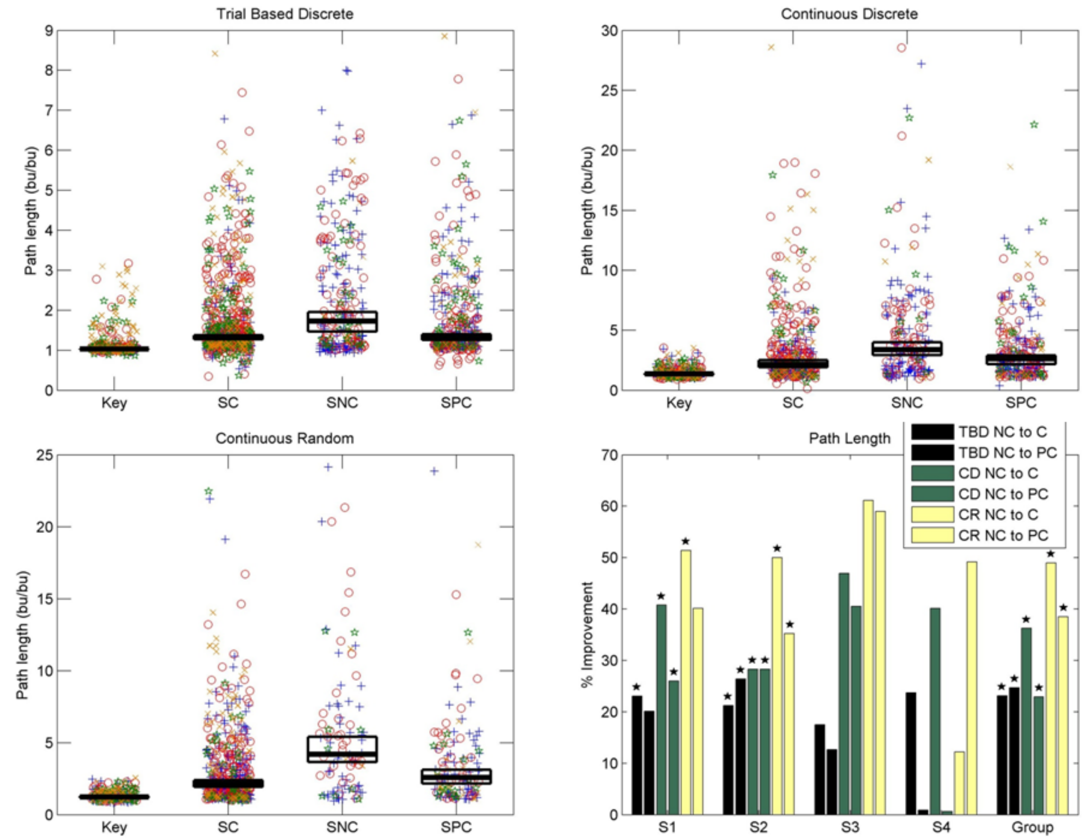

Figure 3.11 Path length to a ring when using the keyboard, cone, no cone, or post cone

Figure 3.11 presents the path length to a ring in a similar manner. In the grouped data, the cone of guidance significantly improved the path length for all paradigms. The average path to a ring across all paradigms using the cone of guidance was 36% shorter than without the cone. Figure 3.12 compares the BCI data to the keyboard data more explicitly. Asterisks indicate that the BCI path length was significantly longer than the keyboard path length. On average across all paradigms in the grouped data, when using the cone of guidance subjects had a path length only 55% longer than their keyboard path length. That number grew to 154% longer when not using the cone of guidance. In terms of acceptable BCI performance to a user, subjects would be willing to excuse 55% much easier than they would 154%, or more than 2.5 times the distance.

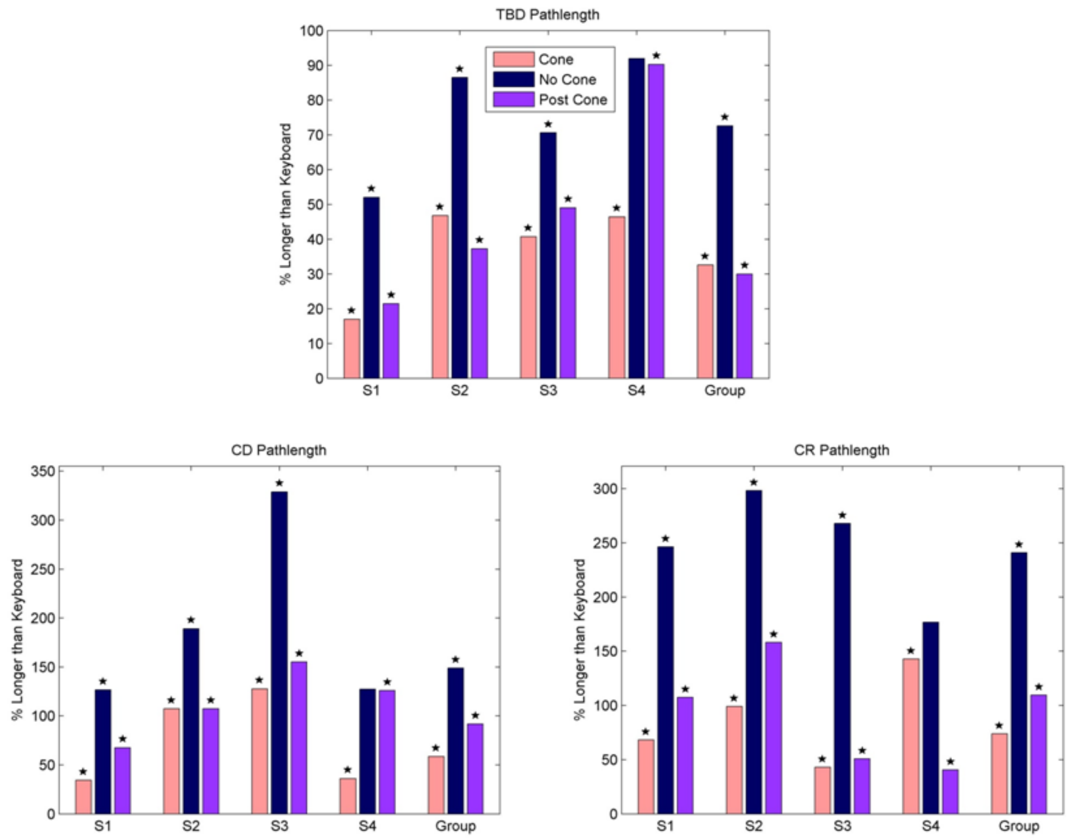

Figure 3.12 Comparison of path length when using the keyboard versus cone, no cone, and post cone

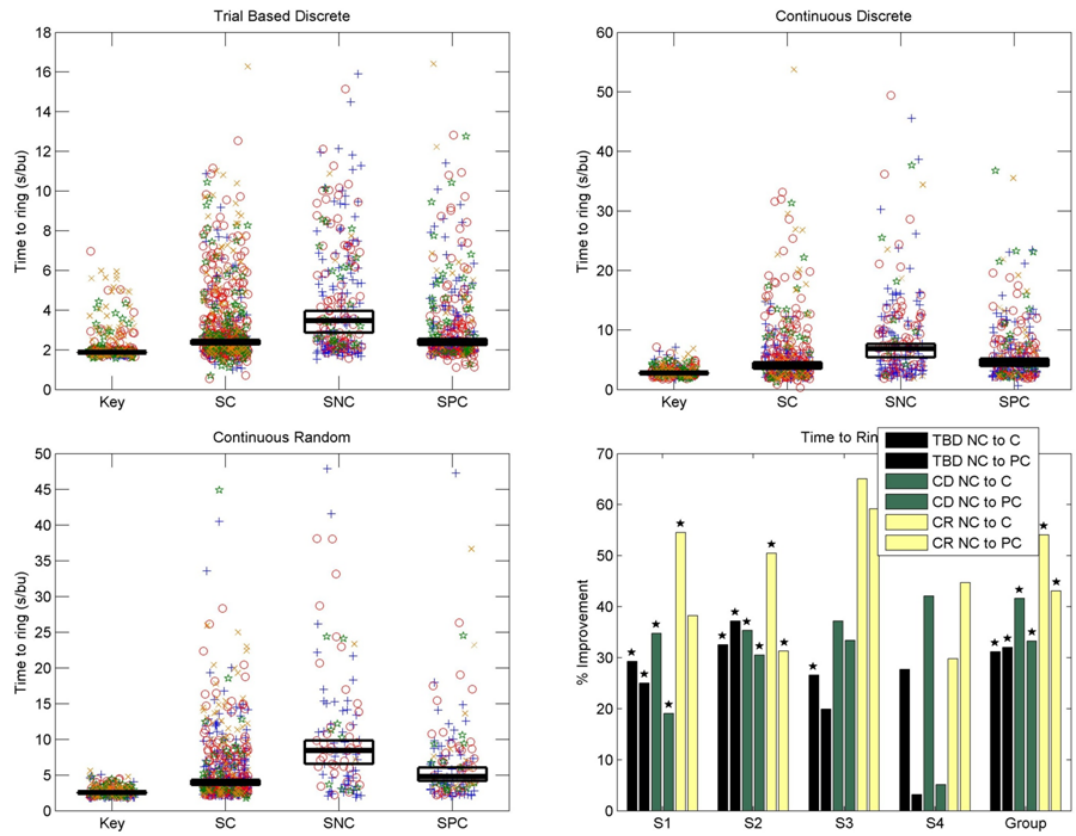

Figure 3.13 Time to ring when using the keyboard, cone, no cone, and post cone

Figure 3.13 presents the data for time to ring. In the grouped data, the cone of guidance significantly improved the time to a ring for all paradigms. The median time to a ring across all paradigms was 42% shorter using the cone of guidance than without the cone. Figure 3.14 compares the BCI data to the keyboard data more explicitly, similar to figure 3.12. On average across all paradigms in the grouped data, using the cone of guidance led to a time to a ring only 42% longer than when using the keyboard. That value was as low as only 13% longer for subject 1 performing trial-based discrete. When not using the cone of guidance, times to ring were on average 155% longer than when using the keyboard, with some subjects taking over 4 times as long to fly through a ring when using the BCI than the keyboard. Once again, times to ring using the cone of guidance were much more acceptable to the user than when not using the cone of guidance.

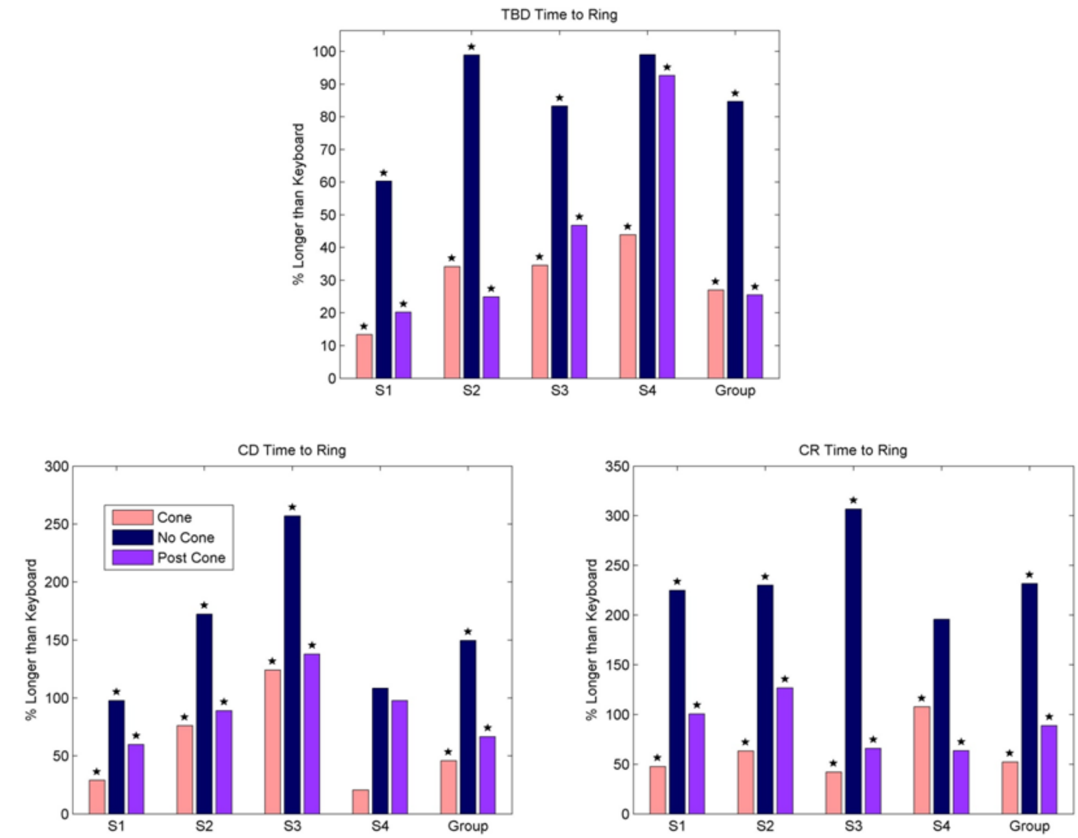

Figure 3.14 Comparison of time to ring when using the keyboard versus cone, no cone, and post cone

3.5.4 Conclusion

All subjects were able to successfully fly through the ring without the assistance of the cone of guidance. However, the cone of guidance did improve performance as quantified above. Subjects had general control of the helicopter and were able to fly to the ring, as shown by the minimal change between using the cone of guidance or not in the average path and time spent closer to the ring than the helicopter began. Subjects had varying degrees of fine control, and that is where the assistance of the cone of guidance significantly influenced the metrics. In the grouped data, the cone of guidance significantly improved the number of rings per minute, path length, and time to a ring. The cone of guidance doubled the number of rings per minute, reduced the time to ring by 27-65%, and reduced the time to a ring by a third. These effects were seen even if the cone of guidance was applied in post analysis. The impact of the cone of guidance varied by subject, and the metrics of the weaker subjects improved more than for the better subjects. Using the cone of guidance allowed this study to have 4 subjects instead of 2 since, using both the cone and no cone data, subjects 3 and 4 did not perform statistically the same as chance for any paradigm. That was not the case if only no cone data was used. Fine control may come with training, reducing the need for the cone of guidance. We cannot say anything more definitive since this was not a long term study. However, the assistance provided by the cone of guidance allowed the system to be useful to a wider population.

Chapter 4 Analyzing the Underlying Neural Signal

In chapter 4, we address the underlying neural signal while using both goal selection and process control. This chapter analyzes the EEG data from section 2.2, but is the first study to analyze the EEG from a BCI at the same timescale as the BCI, and without averaging trials. It also presents the longest running EEG analysis tracking subjects from naive to trained. Sections 4.1-4.4 were submitted for publication. Those sections present the results of analyzing the electrodes and frequencies that were used for control. Section 4.5 presents additional, unsubmitted analysis on the electrodes and frequencies that were not used for control.

4.1 Introduction

A brain-computer interface (BCI) translates signals recorded directly from the brain into commands that control an external device, such as a computer cursor, wheelchair, or neuroprosthetic (Wolpaw et al., 2002; Vallabhaneni et al., 2005). BCIs differ widely in how they implement the translation from raw brain signal to device command (Wolpaw et al., 2002; Wolpaw, 2007). One way they differ is in the overall control strategy.

Consider an individual with a right leg neuroprosthetic. When a healthy individual goes for a walk, the person's primary motor cortex issues the command to walk, and then the actual movement is coordinated by central pattern generators in the spinal cord (Grillner and Zangger, 1979; Duysensa and A Van de Crommertb, 1998). If an individual with a leg neuroprosthetic were to follow a similar procedure , their motor cortex would issue the command to walk and the BCI would recognize that command. Like a healthy individual's spinal cord, the BCI system would then control the prosthetic to generate a coordinated walking motion. That is an example of the BCI using a control strategy called goal selection (Wolpaw, 2007). While using goal selection, the user merely needs to convey the goal to the BCI and then the user receives assistance from the BCI to execute that goal.

A second control strategy, called process control (Wolpaw, 2007), would be more appropriate for a different scenario, such as when learning a new dance. Imagine that the individual with the neuroprosthetic is doing the Hokey Pokey for the first time. That person is listening carefully to the song's instructions: "You put your right foot in, You put your right foot out; You put your right foot in, And you shake it all about." In that instance, the individual would want to control the actual movement of the neuroprosthetic without additional BCI generated movement commands. At that time, the BCI would be using a control strategy called process control. While using process control, the user controls the entire process with no assistance provided.

As the story above demonstrates, goal selection has the potential to enhance the usability of BCIs. However, current BCIs predominantly use process control (Wolpaw, 2007). Two recent studies directly compared goal selection to process control in a non-invasive, sensorimotor rhythm based BCI (Royer and He, 2009; Royer et al., 2011). Royer and He (2009) found that goal selection was more accurate and faster to use in both trained and naive subjects. Royer et al., (2011) tracked naive subjects through the learning process and confirmed in a larger sample size that goal selection was more accurate and faster to use, while showing that goal selection was easier to learn and required less mental effort than process control.

Although those studies presented behavioral and performance data directly comparing goal selection to process control, the differences in the underlying neural signal while using goal selection vs. process control have not been addressed. What were the neural mechanisms underlying the improved performance of goal selection over process control?

Both studies (Royer and He, 2009; Royer et al., 2011) used motor imagination to generate sensorimotor rhythms. When a subject imagines moving the right or left hand, the sensorimotor rhythms respond with contralateral event related desynchronization (ERD) and ipsilateral event related synchronization (ERS) (Pfurtscheller and Lopes da Silva, 1999; Yuan et al., 2008). The classical method of computing the time course of ERD/ERS involves bandpass filtering all event-related trials, squaring the amplitude measurements to produce power measurements, and then averaging the power across all trials (Pfurtscheller and Aranibar, 1979).

Recent work has shown that averaging trials and samples can lead to misleading, or even invalid, conclusions (Golowasch et al., 2002; Freyer et al., 2009). Work in the pyloric network of the crustacean stomatogastric ganglion has shown that similar network performance is achieved by very different network and cellular properties across different animals. If measurements of those properties are averaged, the resulting network behaves in a manner completely inconsistent from the networks from which the measurements were taken (Golowasch et al., 2002; Prinz et al., 2004; Grashow et al., 2010). In a more pertinent example, recent research has shown that the human alpha rhythm does not "wax and wane" in amplitude but shifts erratically from a low-power to a high-power state in a multi-stable manner (Freyer et al., 2009).

As discussed in Vidaurre et al. (2011), consistently across different motor imagery studies 15-30% of subjects cannot modulate their sensorimotor rhythms well enough to operate a BCI at an accuracy level that would allow operation of an assistive device (Kuebler et al., 2004; Vidaurre et al., 2011). That accuracy level is commonly accepted to be a minimum of 70% (Kuebler et al., 2004; Vidaurre et al., 2011). The majority of naive subjects struggle to consistently achieve 70% accuracy across paradigms (Guger et al., 2003; Royer and He, 2009). Even with training, a non-negligible portion of healthy subjects do not operate a sensorimotor rhythm based BCI at accuracy levels of 70% (Neuper et al., 2009; Royer et al., 2011). Perhaps the accepted practice of averaging trials while performing offline analysis has contributed to this problem.

As a BCI researcher recently bemoaned, variability of the EEG amplitude has not been considered by the BCI community (Wu et al., 2011). That study looked at trial by trial variability, but looked at the trial-average amplitude and still averaged trials to produce the single-trial EEG time course. A BCI processes the EEG in real-time, providing near instantaneous feedback to the user. We argue that it is in fact the single-trial EEG time course that is important to good BCI performance. We therefore analyzed within single trials, without averaging, in an effort to determine what distinguishes acceptable from unacceptable BCI performance. We hypothesize that the underlying EEG signal will also be more conducive to a goal selection control strategy.

4.2 Methods

4.2.1 Subjects and data acquisition

This study involved 20 healthy human subjects recruited from the community. Seven subjects were male and 13 were female. Ages ranged from 18 to 28. All subjects gave written consent according to a protocol approved by the Institutional Review Board of the University of Minnesota. As previously described (Royer et al., 2011), all subjects sat in a comfortable chair in front of a computer monitor while wearing a 64 channel EEG cap

(Compumedics NeuroMedical Supplies Quik-Cap) set up in accordance with the 10-20 layout. Data from all electrodes were fed to a Neuroscan amplifier sampling at 1000Hz and filtered from DC to 200Hz. The signal was then passed into the general purpose system BCI2000 (Schalk et al., 2004).

Subjects were asked to use motor imagination of the right hand, arm, or shoulder vs. the left hand, arm, or shoulder to control a one-dimensional cursor task. The 20 subjects were randomly split into 4 groups of 5 subjects. Each group performed one of the four paradigms depicted in figure 1 and further described in Royer and He (2009) and Royer et al. (2011, includes supplementary movies). Two of the paradigms were based on process control and two on goal selection. The paradigms were designed to be as similar as possible. The underlying signal processing that controlled the movement of the cursor was the same across all four paradigms, as was trial timing. What differed was what the subjects had to do in order to have a successful trial.

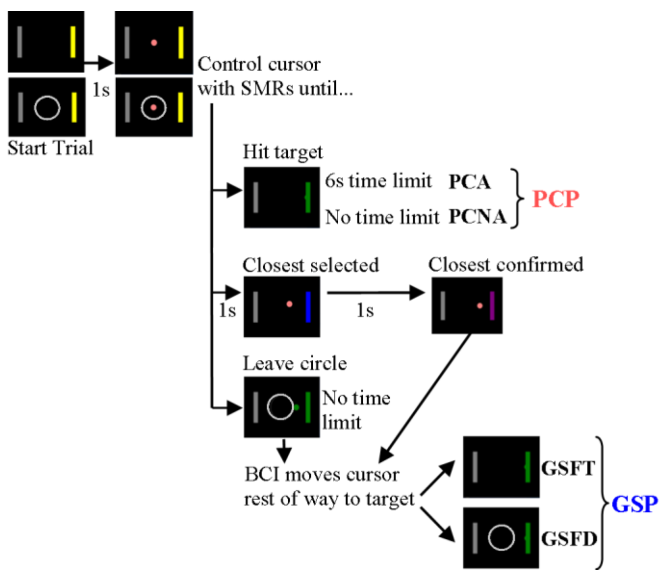

Figure 4.1 Experimental paradigms
Subjects were instructed to use motor imagery to move the cursor toward the yellow target (here, right). The subject had to hit the target with the cursor in the process control paradigms (PCP): process control with aborts (PCA) and process control with no aborts (PCNA). The subject received assistance from the BCI system in the goal selection paradigms (GSP): goal selection with feedback limited by time (GSFT) and goal selection with feedback limited by distance (GSFD). SMRs = sensorimotor rhythms.

All four paradigms began with two targets appearing on the lateral edges of the screen (figure 1). After 1s, a red cursor appeared in the middle. Subjects were instructed to use right vs. left motor imagination to move the cursor to the yellow target. Subjects were given guidance on suggested imaginations, but were free to use whatever worked for them. For all paradigms, if the cursor hit the yellow target, the target turned green, and if the cursor hit the grey target, the target turned red.

In the two paradigms based on process control, the subjects had to move the cursor all the way to the target themselves. In process control with aborts (PCA), the subjects had a 6s time limit. In process control with no aborts (PCNA), the trial did not end until the cursor hit a target. The data from both those paradigms were grouped into the process control paradigms (PCP).

In the two paradigms based on goal selection, the BCI moved the cursor the rest of the way to the target once either a time or distance criteria was met. In goal selection with feedback limited by distance (GSFD), there was a grey circle with a radius of 20% of the screen centered between the two targets. Once a subject moved the cursor outside the circle, the BCI automatically moved the cursor the rest of the way to the closest target. There was no time limit on how long the cursor could remain within the circle.

In goal selection with feedback limited by time (GSFT), the cursor did not have to move any predetermined distance. Instead, the BCI automatically moved the cursor to the closest target after a predetermined time (2 or 3 s). In order to provide more feedback to the subject, after 1s, the closest target turned blue to indicate that it had been selected. If, after an additional 1s, the cursor was still closest to that target, the target turned purple and the BCI automatically moved the cursor the rest of the way to the target. If the cursor had moved so it was closer to the other target (not blue), that target would then turn blue and a third 1s time period would begin. The target that had the cursor nearest for 2 of the 3 1s intervals would then be the final destination for the cursor.

The two paradigms based on goal selection, GSFD and GSFT, were grouped into the goal selection paradigms (GSP). Both GSP and PCP consisted of a timed and an untimed paradigm, with the main difference between GSP and PCP being that the user received assistance from the BCI in GSP.

All subjects were naive to BCI usage prior to the study. During the study, subjects completed 8 sessions of 10 runs of their assigned paradigm. Each run was four minutes long and contained as many trials as would fit within the four minutes. Subjects had 3s of rest between trials, and a user determined time of rest between runs. Sessions occurred about once per week.

4.2.2 Control of the cursor

The EEG signal from all 64 channels was fed into BCI2000, which used a 16th order autoregressive model to calculate the spectral amplitudes of 3 Hz bins centered on a multiple of three from 0 to 30 Hz. The window size was 160ms with a 50% overlap, producing a new set of spectral amplitudes every 80ms. As described in Royer et al. (2011), only the spectral amplitudes from a predetermined set of electrodes and frequencies were used to control the motion of the cursor. The electrodes and frequencies were determined by using the BCI2000 Offline Analysis tools to select those that had the highest r2 to the condition left vs. right target. Those electrodes and frequencies were limited to those over sensorimotor cortex (FCz-6, Cz-6, and CPz-6) and 6-30Hz. The exact combination of electrodes and frequencies used for each subject was called their control signal.

The control signal during session 1 for all subjects was C3: 9, 12, 18, 21 and 24 Hz. After session 1, the control signal was individualized to each subject based on their previous session's data. On average, subjects used two electrodes and two frequency bins in their control signal. Control signals typically did not change much for a single subject after session 2, with only minor tweaks of neighboring electrodes or frequency bins. For all subjects, the control signal used in session 7 had to be one that had been used previously, and that control signal was also used for session 8. By the end of the 8 sessions, the most commonly used electrodes were C3, C4, CP3, and CP4. All subjects but 1 used at least one of the 9, 12, or 15Hz bins in their final control signal. For further information on control signal selection and customization please see Royer et al. (2011).

The electrodes that were chosen to form the control signal were typically limited to one side of the head. This was done for several reasons. First, previous research has shown that subjects can more easily produce similar activity across both hemispheres of the brain than they can differential activity in the 8-12Hz band (Pineda et al., 2003). Second, more recent research has shown that the speed of motor imagination scales the amplitude of the recorded EEG signal equally on both sides of the head (Yuan et al., 2010b). Previous studies that used the difference of amplitude across the hemispheres as the control signal would not be able to see this change of signal (Wolpaw and McFarland, 2004; Royer and He, 2009). By limiting the electrodes that were used to control the cursor to a single side of the head, we increased the chances that the subjects could generate a useful signal and allowed for increased control through the scaling of the amplitude of the signal.

The EEG signal controlled the movement of the cursor as illustrated in figure 2 and further described in Royer et al. (2011). For each 80ms time step, the spectral amplitudes of the preselected frequency bins (the light colored box in fig 2A) were averaged (fig 2B left axis) and then normalized (fig 2B right axis) to a signal with zero mean and unit variance. The normalization process used a rolling buffer set to accommodate multiple trials of both left and right. Positive values of the normalized signal moved the cursor to

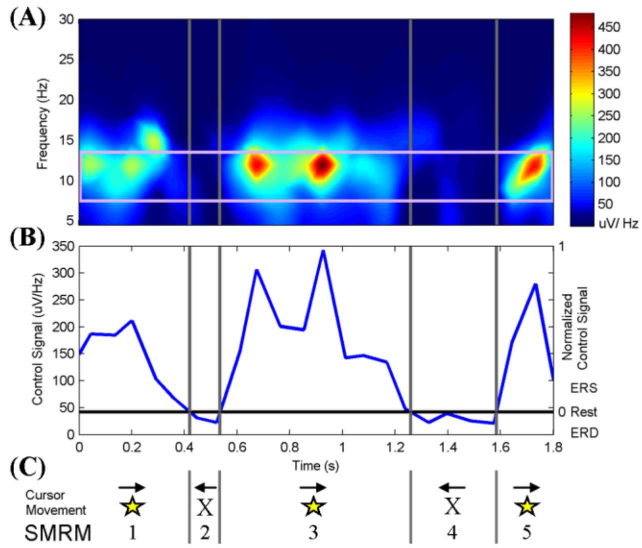

Figure 4.2 The decomposition of a single trial to cursor movement
The decomposition of a single trial from the time-frequency domain to individual sensorimotor rhythm modulations (SMRMs) and the movement of the cursor. The vertical grey lines across all parts of the figure break the trial into 5 distinct SMRMs numbered at the bottom of the figure. The time axes labels in (B) apply to the entire figure. (A) The signal from CP4 during a typical trial in the time-frequency domain. The frequency bins of 9 and 12 Hz (light colored box) controlled the cursor movement. (B) The average spectral amplitude (left axes) of the boxed area above was normalized (right axes) to create the control signal. The horizontal black line indicates the resting, baseline value. Event related synchronization (ERS) and event related desynchronization (ERD) were both present. (C) Positive values of the control signal moved the cursor to the right, or in the correct direction as indicated by the star. Negative values of the control signal moved the cursor to the left, or in the incorrect direction as indicated by the X. This trial resulted in a hit.

the right, whereas negative values moved the cursor to the left (fig 2C). Motion towards the yellow target, the right one in the particular trial illustrated in figure 2, is indicated by the stars in figure 2C. Motion towards the undesired target is indicated by the X in figure 2C.

4.2.3 Data analysis

This study looked at the EEG signal controlling the motion of the cursor throughout the course of the eight sessions. Behavioral results including accuracy, information transfer rate, number of hits per run, time to a hit, and effort of a hit were previously reported (Royer et al. 2011).

Both real time control of the cursor and offline analysis of the EEG signal used the same signal processing techniques described above. As seen in figure 2, the spectral amplitude of the sensorimotor rhythm forming the control signal fluctuated during the course of a trial, resulting in a series of jittery cursor movements back and forth. Each trial was thus broken into a series of sensorimotor rhythm modulations (SMRMs, or modulations). A single SMRM could be either correct if it moved the cursor toward the desired target, or incorrect if it moved the cursor away from the desired target. Since correct and incorrect modulations alternated, there were approximately equal numbers of both.

The amplitude and duration of individual modulations as well as the entire trial were quantified. During the course of the 8 sessions, there were a total of 46,036 trials: 18,890 PCP trials, with 10,398 hits and 8,492 misses; and 27,146 GSP trials, with 18,927 hits and 8,219 misses. Those trials broke down into a total of 432,722 SMRMs: 163,421 GSP SMRMs, 82,916 during left target trials and 80,505 during right target trials; and 269,301 PCP SMRMs, 132,365 during left target trials and 136,936 during right target trials.

We characterized the EEG signal, more specifically the control signal, during the trials by quantifying different aspects of both the total trial and its individual modulations. We quantified each trial by looking at two main categories, duration measures and amplitude measures. The duration measures were further categorized into correct modulation, incorrect modulation, and the difference between correct and incorrect modulation. The exact duration measures included, for each trial, the total time of modulation, the average time of a single modulation, and the longest time of a single modulation. The amplitude measures were calculated separately for the right target, left target, and difference between the two targets. For each trial, we calculated the average spectral amplitude of the control signal during the entire trial, the average amplitude of a correct modulation, and the average amplitude of an incorrect modulation.

Although the data are presented as "right target" and "left target," that was done for ease of understanding and was not literally true for every control signal. As previous research has shown (Pfurtscheller and Lopes da Silva, 1999; Yuan et al., 2008), motor imagination normally produces ipsilateral ERS and contralateral ERD. Since we were only using electrodes from one side of the head to control the movement of the cursor, the control electrodes would theoretically record ipsilateral ERS for motor imagery of the target ipsilateral to the electrodes and would record contralateral ERD for motor imagery of the target contralateral to the electrodes. By the very definition of ERS and ERD, the synchronized activity of ipsilateral ERS would have a higher amplitude from baseline and the desynchronized activity of contralateral ERD would have a lower amplitude from baseline. Thus, ipsilateral motor imagery would produce a higher amplitude signal on the electrodes used for control than contralateral motor imagery.

Assuming that the subject produced the normal ipsilateral ERS and contralateral ERD, if the electrodes in a subject's control signal were on the right side of the head, those electrodes would record a higher amplitude for right motor imagination (ipsilateral ERS) than left motor imagination (contralateral ERD). However, if the electrodes in a subject's control signal were on the left side of the head, those electrodes would record a smaller amplitude for right motor (contralateral ERD) imagination than left motor imagination (ipsilateral ERS). To account for this reversal in sign, the control signal for each subject was assigned a positive or negative weight to ensure that right motor imagery produced an overall higher control signal value than left motor imagery and moved the cursor to the right. A normal subject, one who generated ipsilateral ERS and contralateral ERD, with a control signal from the right hemisphere would have a positive weight, and a normal subject with a control signal from the left hemisphere would have a negative weight.

Previous research has also shown a large amount of subject specific variation in ERD/ERS spatial patterns as a result of motor imagery (Blankertz et al., 2008). This study was no exception, with 26% of all control signals having a non-normal weight for the hemisphere of the control signal (-1 for right and +1 for left). Without this non-normal population, the targets could have been presented merely as ipsilateral and contralateral to the control signal electrodes, similar to what has been done previously (Yuan et al., 2011). However, the 26% of non-normal control signals muddied the conclusions from that analysis. The overall trends and conclusions were the same, just less distinct.

The BCI performed mathematical calculation of the control signal as described above according to the electrodes, frequency bins, and weights provided by the experimenter. The BCI moved the cursor accordingly, without regard for the above theory regarding motor imagination and contralateral/ipsilateral ERD/ERS. The electrodes, frequency bins, and weights provided by the experimenter were chosen on a subject specific basis so that each subject's right motor imagination moved the cursor to the right and left motor imagination moved the cursor to the left, regardless of whether or not their imaginations were producing the normal, theoretical output. Within the signal processing of the BCI, the only constraint was that higher amplitude on the electrodes used for control would move the cursor toward one target and lower amplitude would move the cursor toward the other target. The amplitudes could be from ERS, ERD, or rest. The source of the amplitude did not influence the movement of the cursor. For the sake of fairness to the subjects, clarity of analysis, and ease of understanding, what is labeled as the "right" target was the high amplitude target for each control signal. Similarly, what is labeled as the "left" target was the low amplitude target for each control signal. This was literally true for the majority of trials.

Figures 3 and 4 quantify the presence of ERD and ERS within the trials and the portion of subjects that used ERD, ERS, or rest as the correct modulation for each target. For this analysis, the baseline for each control signal was calculated from the average spectral amplitude of all trials, both right and left targets, in the one second after the targets were presented but before the cursor appeared. The area plot for GSP in figure 4 does not reach 100% at session 2 because one individual did not consistently use ERD or ERS as the correct modulation for either target, but had median amplitudes for correct modulations both to the right and left targets that were not significantly different from baseline. The amplitude measures in figure 5 are also given as the difference from baseline.

Measures were tested for normality using a 2-sided Lilliliefors test. All measures were found to be non-normal. Therefore, figures 3B, 5, 6 and 7 plot the median of the grouped data and the 95% confidence interval as indicated by the shaded region. Figure 3C presents the duration of all modulations across all sessions as both a histogram and as a cumulative percent. Figures 5, 7, and 11 break the trials down into "hit trials" and "miss trials." A hit trial was a trial that ended with the cursor hitting the desired target. A miss trial was a trial that did not end with the cursor hitting the desired target. Correlations between the different measures in figures 8 and 9 were performed using the median value for each session. For both GSP and PCP, there were 79 sessions of data. Figure 10 used the Clopper-Pearson method to calculate the 95% confidence intervals as indicated by the shaded region. Figure 11 presents the grand average time frequency response of all hits following procedures previously described (Yuan et al., 2008).

4.3 Results

4.3.1 Trials were a collection of sub-second sensorimotor rhythm modulations

We analyzed the EEG signal from the electrodes and frequencies each subject used to control a cursor during a left/right BCI task. Each trial was broken down into individual sensorimotor rhythm modulations (SMRMs, or modulations) as shown in figure 2. As displayed in figure 2, the single trial contained both ERD and ERS. Neither modulation

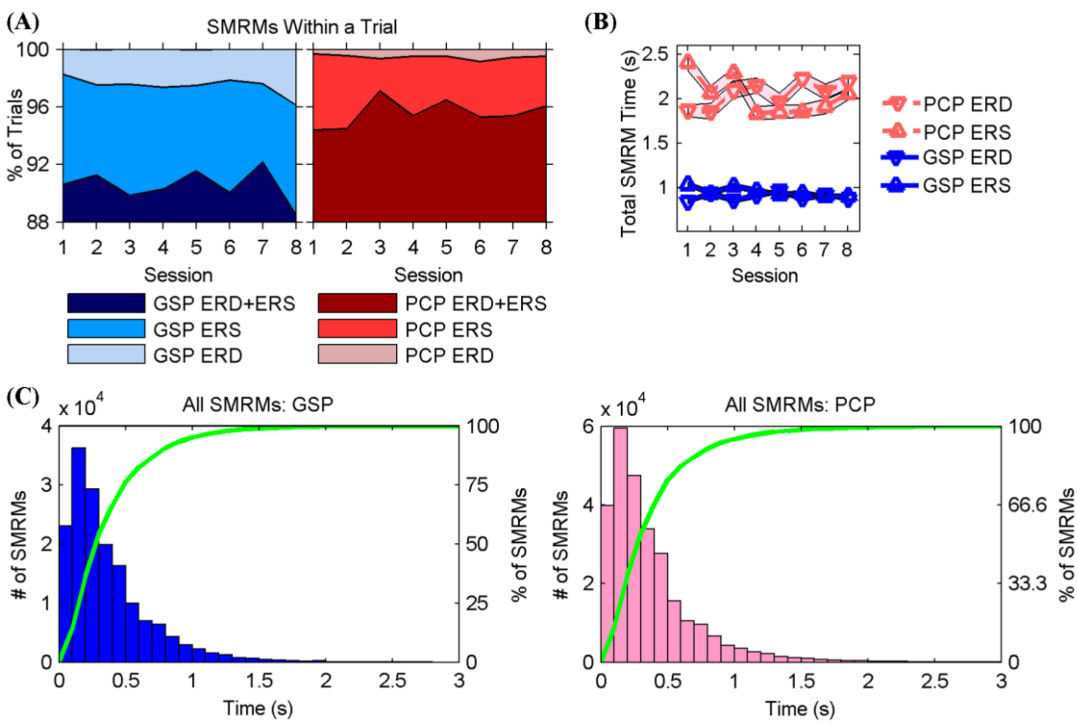

Figure 4.3 Sensorimotor rhythm modulation (SMRM) within individual trials
(A) 93% of all trials contained both ERD and ERS, with GSP having more trials with only ERD or ERS than PCP. (B) Trials contained approximately equal amounts of ERD and ERS for both GSP and PCP. Shaded area represents the 95% confidence interval of the median. (C) Time Distribution of all SMRMs. 95% of all SMRMs were under 1s for both GSP and PCP. The number of SMRMs between 0.1s time steps is shown by the bars and the left axis. The cumulative percent of SMRMs at each time step is shown by the green line and the right axis.

was sustained throughout the trial. Figure 3 quantifies the prevalence of this behavior within all the individual trials. As figure 3A shows, 93% of all trials contained both ERD and ERS. Only 7% of all trials contained only ERD or ERS. The median time of ERD or ERS within a single trial was approximately equal across all trials (fig 3B).

On average, the time periods of ERD and ERS were spread across 6 modulations for GSP and 13 modulations for PCP. How long were these modulations? Figure 3C presents the duration of all modulations as both a histogram and as a cumulative percent. The majority of modulations were quite short, with 22% lasting between 100ms and 200ms. 95% of all modulations lasted less than 1s.

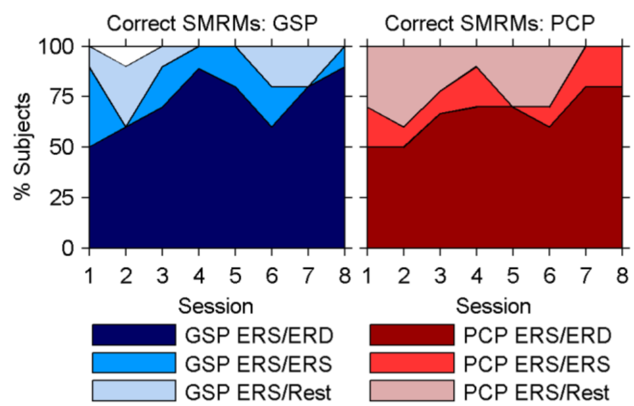

Figure 4.4 Correct sensorimotor rhythm modulations for the left and right targets
A correct sensorimotor rhythm modulation (SMRM) moved the cursor closer to the desired target. The majority of subjects used ERS as the correct SMRM for the right target, and ERD for the left target. Fewer GSP subjects than PCP subjects used rest to move the cursor.

As mentioned in the methods, one amplitude of modulation moved the cursor toward the right target, and a lower amplitude of modulation moved the cursor toward the left target. Those amplitudes could either be the same level as rest, higher than baseline (ERS), or lower than baseline (ERD). The exact combination of amplitudes varied by subject and session. Figure 4 shows an area plot of the percent of subjects that used each combination of rest, ERS, or ERD. From the very first session, 50% of subjects used ERS as the correct modulation moving the cursor closer to the right target, and ERD as the correct modulation moving the cursor closer to the left target. The percentage of subjects using ERS and ERD as their correct modulations grew over time so that, by session 8, 85% of all subjects used ERS and ERD. Across the 8 sessions, 10% more GSP subjects used ERS and ERD than did PCP subjects. Interestingly, 15% of subjects used two amplitudes higher than baseline (ERS and ERS) as the correct modulation for the two targets. That percentage was relatively constant across the 8 sessions. By the final session no subject used rest as their primary means to move the cursor towards a target. However, throughout the 8 sessions, 80% more PCP subjects used rest to control the cursor than did GSP subjects.

4.3.2 Amplitudes of correct and incorrect modulation

Figure 5 presents the average amplitudes for the entire trial, a correct modulation, and an incorrect modulation when looking at all the data, and when breaking the trials out into hit and miss trials. The trial average amplitude for all trials (figure 5A, left) behaved

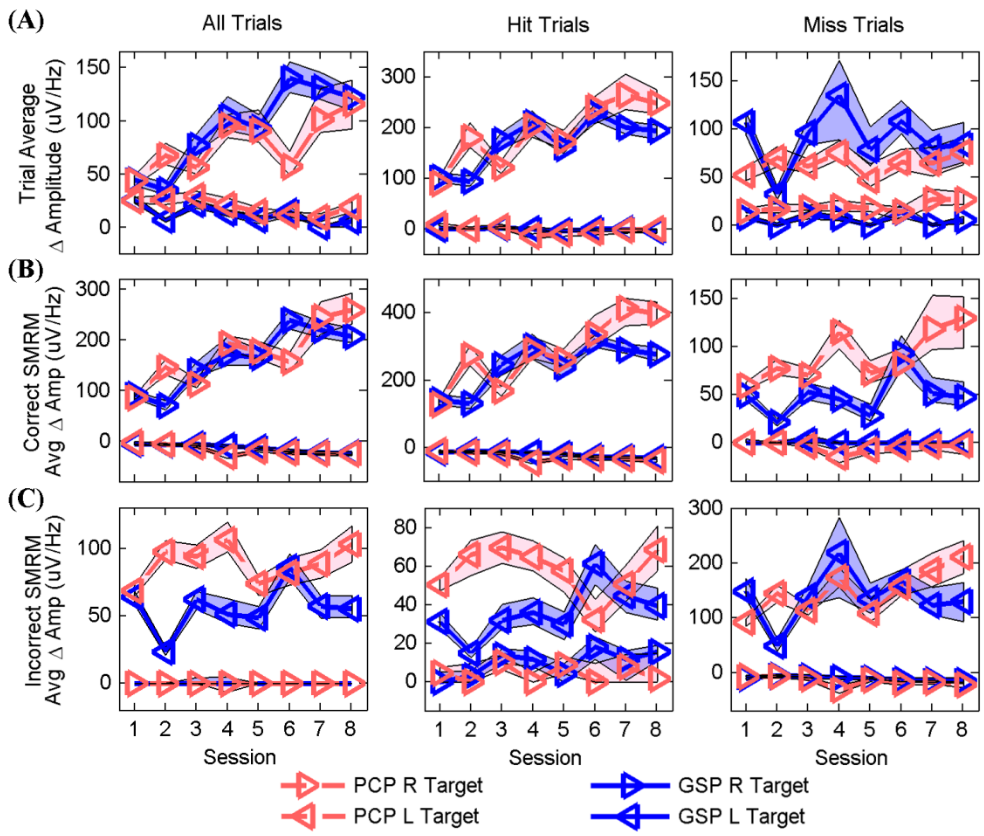

Figure 4.5 Sensorimotor rhythm modulation spectral amplitude
Sensorimotor rhythm modulation (SMRM) spectral amplitude within individual trials for each target. The left column presents the data for all trials, the middle column presents the data for only trials that ended in a hit, and the right column presents the data from trials that did not end in a hit. The label on the left of each row applies to the whole row. (A) Average trial amplitude during a hit was the same for GSP and PCP. In PCP, a miss had an average amplitude closer to a hit than did GSP. (B) The average amplitude for a correct SMRM was the same for GSP and PCP except for within miss trials, where PCP had better correct amplitudes than GSP but the trial still resulted in a miss. (C) GSP minimized the amplitude of individual incorrect SMRMs overall and in hit trials, but equalled PCP in misses. (A-C) Plotted lines are the medians of the grouped data for each session. Shaded areas indicate the 95% confidence interval of the median. Significance is indicated by non-overlapping areas. R = right. L = left.

similarly for both GSP and PCP across the 8 sessions. The amplitude for the right target gradually increased whereas the amplitude for the left target gradually decreased. GSP decreased the left target amplitude significantly more than did PCP, ending 67% lower. That resulted in the difference in average amplitude between the right and left targets in the last session 20% larger for GSP than PCP. When the data was broken out into hit trials and miss trials, an interesting trend emerged. The average trial amplitude for a hit trial (figure 5A, middle), was the same for both GSP and PCP. In a miss trial (figure 5A, right), PCP had an average trial amplitude closer to a hit than did GSP.

Figure 5B and 5C breaks the trials down into their individual correct and incorrect modulations. Across all trials and during all hit trials (figure 5B, left and middle), an individual correct modulation was the same for GSP and PCP. However, in the miss trials (figure 5B, right), PCP had higher amplitudes of correct modulation than GSP but the trial still resulted in a miss. The average amplitude of an incorrect modulation in a miss trial was the same for GSP and PCP (figure 5C, right), but GSP minimized the amplitude of an incorrect modulation in the hit trials and in all trials compared to PCP (figure 5C, left and middle). Overall, correct modulation, either an individual modulation or the trial average for a hit, had the same amplitude for GSP and PCP. However, PCP had higher amplitudes of incorrect modulation than GSP and PCP had higher amplitudes of correct modulation that still resulted in a miss.

4.3.3 Durations of correct and incorrect modulation

Figure 6 presents for all data the duration of correct and incorrect modulations for the entire trial as well as for individual modulations. An important factor behind the interpretation of figure 6 is that the median time of a PCP trial (4.4s) was considerably longer than for GSP (2s). The median values held constant across all 8 sessions. Therefore, it is not surprising that PCP had much longer total times of modulation than did GSP (figure 6A, left). However, the total time of correct modulation for PCP did not change significantly over time, whereas GSP significantly increased the time of correct modulation within a trial. When the shorter trial times of GSP are accounted for (figure 6A, right), GSP trials had significantly more percentage of time of correct modulation than PCP and significantly less incorrect modulation.

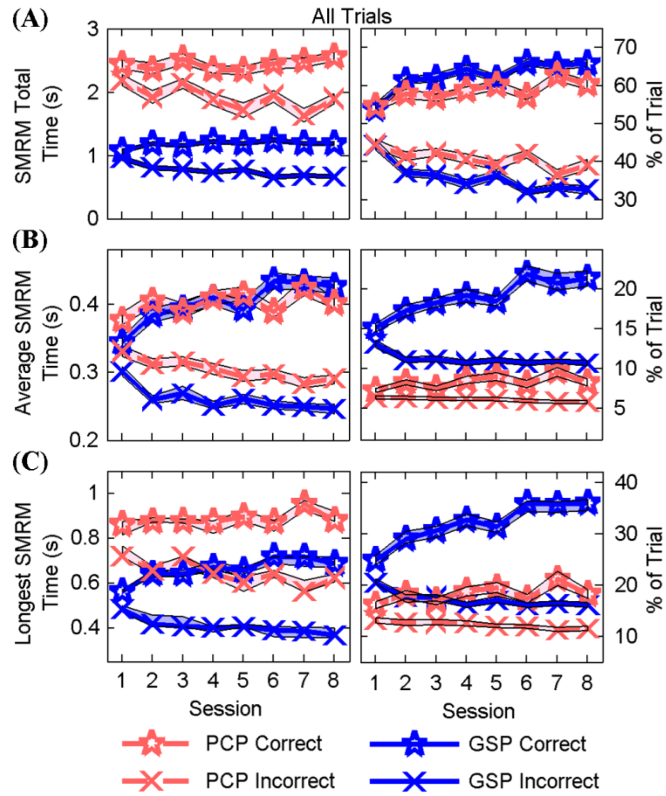

Figure 4.6 Sensorimotor rhythm modulation duration
Sensorimotor rhythm modulation (SMRM) duration within individual trials. The left column plots the time in seconds whereas the right column plots the same data as a percentage of the trial. (A) GSP trials had significantly more percentage of time of correct modulation than PCP and significantly less incorrect modulation. (B) The average time of an individual correct SMRM was the same for GSP and PCP, although PCP had longer incorrect SMRMs. (C) The longest duration of an individual correct SMRM in a trial was significantly shorter for GSP than PCP, but was a significantly longer percentage of the trial.

A single average correct modulation had the same duration between GSP and PCP, but PCP had longer incorrect modulations (figure 6B, left). Despite being the same actual duration, when looked at as the percentage of the trial, a single correct average modulation was 129% longer for GSP than PCP (figure 6B, right). The longest correct and incorrect modulation in a trial was significantly longer for PCP than GSP (figure 6C, left). The longest correct PCP modulation did not change significantly over time, whereas GSP was able to significantly lengthen the longest correct modulation in a trial. Although the longest duration of an individual correct modulation was significantly shorter for GSP than PCP, it was a significantly longer percentage of the trial (figure 6C, right). Overall, the shorter length of GSP trials allowed correct modulations to occupy a longer percentage of the trial than in PCP. Comparing the longest correct modulation to an average correct modulation, PCP had a longest modulation 121% longer than its average modulation across all sessions. GSP had a longest modulation only 65% longer than its average modulation.

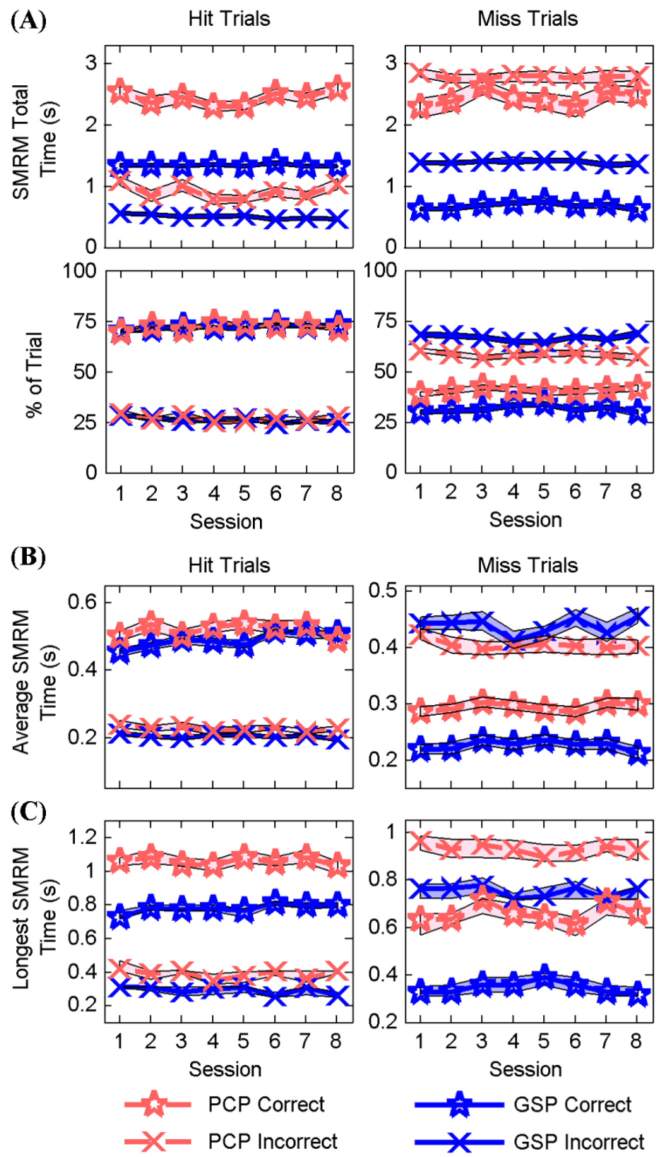

Figure 4.7 Sensorimotor rhythm modulation duration within hit vs. miss trials
(A) A hit trial for both GSP and PCP had the same percentage of correct and incorrect SMRM total time. A miss trial for PCP had the same total time of correct SMRM as a hit. (B) Although the average length of an SMRM in a hit trial was similar to PCP, GSP significantly lengthened a correct SMRM and significantly shortened an incorrect SMRM over the 8 sessions, whereas PCP remained constant. In miss trials, both GSP and PCP remained constant, but PCP had longer durations of correct SMRM that still resulted in a miss. (C) The longest SMRMs in a PCP hit trial were significantly longer than for GSP. In a miss trial, the longest correct SMRM for PCP was almost as long as the longest correct SMRM for a GSP hit trial.

Figure 7 presents the same duration measures applied to only hit or miss trials. When viewed as the percentage of the trial, the total time of correct and incorrect modulation is statistically the same in hit trials for both GSP and PCP (figure 7A, bottom left). Total time of correct modulation for PCP was similar in both hit and miss trials, with the difference between a hit and a miss being the total duration of incorrect modulation. On the other hand, GSP hits and misses exchanged durations of total correct vs. incorrect modulation (figure 7A, top row). The duration of an average correct or incorrect modulation in a hit trial was similar between GSP and PCP (figure 7B, left). However, GSP significantly lengthened a correct modulation and significantly shortened an incorrect modulation over the 8 sessions, whereas PCP remained constant. In miss trials, the duration of an average correct or incorrect modulation remained constant for both GSP and PCP, with PCP having longer durations of correct modulations that still resulted in a miss (figure 7B, right). PCP had significantly longer durations of longest correct and incorrect modulation in both hits and misses than GSP (figure 7C). PCP durations were so long that the longest correct modulation in a PCP miss trial was almost as long as the longest correct modulation in a GSP hit trial. In general, hit trials for both GSP and PCP have similar durations of modulation for total time as percent of trial and a single average modulation. PCP miss trials had much longer correct modulations than GSP that still resulted in a miss.

4.3.4 Duration of modulation was more important than amplitude of modulation

Each modulation was quantified in terms of duration, amplitude, and correct vs. incorrect as presented in figures 5 through 7. What characteristics of the within trial modulations were important to successful BCI performance? To answer that question, the trial modulation measures were correlated to the performance data (previously reported in Royer et al., 2011). Figure 8 shows that the duration measures had higher significant correlations to performance than did the amplitude measures. In particular, the duration measures were highly correlated to information transfer rate, accuracy, and number of hits per run. Those metrics addressed the successful use of the BCI. The duration measures were less correlated to the metrics that addressed the ease of use of the BCI: time to a hit, effort of a hit, and total trial time. The amplitude measures were less correlated than the duration measures to both successful use and ease of use of the BCI.

Additionally, total trial time was significantly, negatively correlated to many of the amplitude and duration measures. For both GSP and PCP, as trial time increased, a correct modulation had a shorter average duration (GSP r=-0.43, PCP r=-0.73, p<0.001 for both) and a decreased average amplitude difference between the right and left targets (GSP r=-0.28, p<0.05, PCP r=-0.54, p<0.001). Similarly, longer trials had less of a difference between right and left targets in overall trial average amplitude (GSP r=-0.47, PCP r=-0.48, p<0.001 for both). Those are just examples of some of the more highly correlated measures.

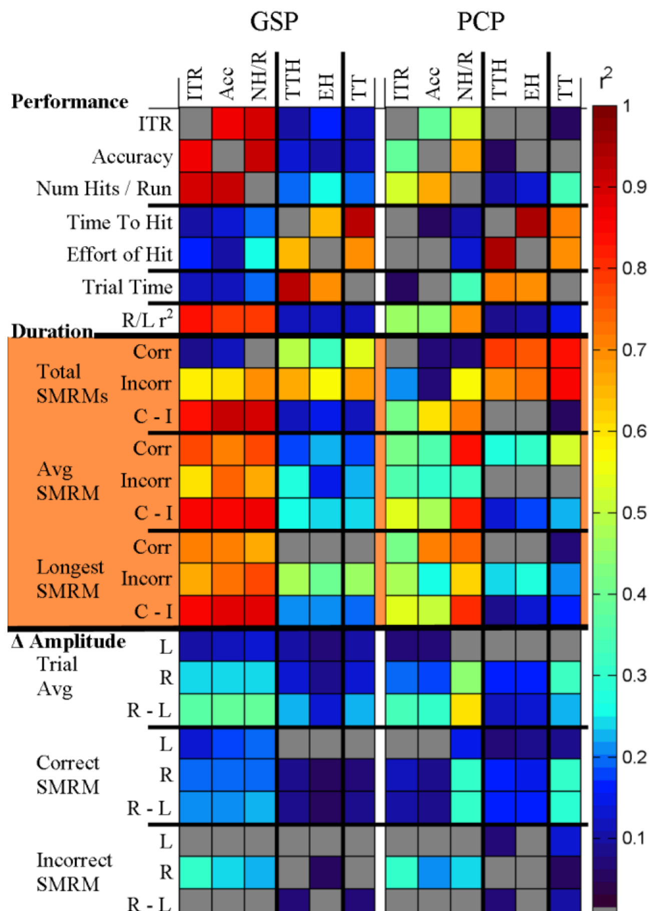

Figure 4.8 SMRM duration (orange box) was more correlated to performance than amplitude The performance metrics fell into one of two categories, successful use (information transfer rate (ITR), accuracy (Acc), and number of hits/run (NH/R)) or ease of use (time to hit (TTH), effort of hit (EH), and trial time (TT)). Grey boxes indicate that the correlation between the metrics was not significant at the $p < 0.05$ level. R = right target, L = left target, Corr = correct SMRM, Incorr = incorrect SMRM, C - I = correct SMRM minus incorrect SMRM, Avg = average. n=79 GSP sessions and 79 PCP sessions.

4.3.5 Correct modulation longer than incorrect modulation led to successful use

An accuracy level of 70% is commonly accepted as the minimum accuracy necessary to allow BCI operation of an assistive device (Kuebler et al., 2004; Vidaurre et al., 2011). Did particular duration lengths correspond to this acceptable level of performance? Figure

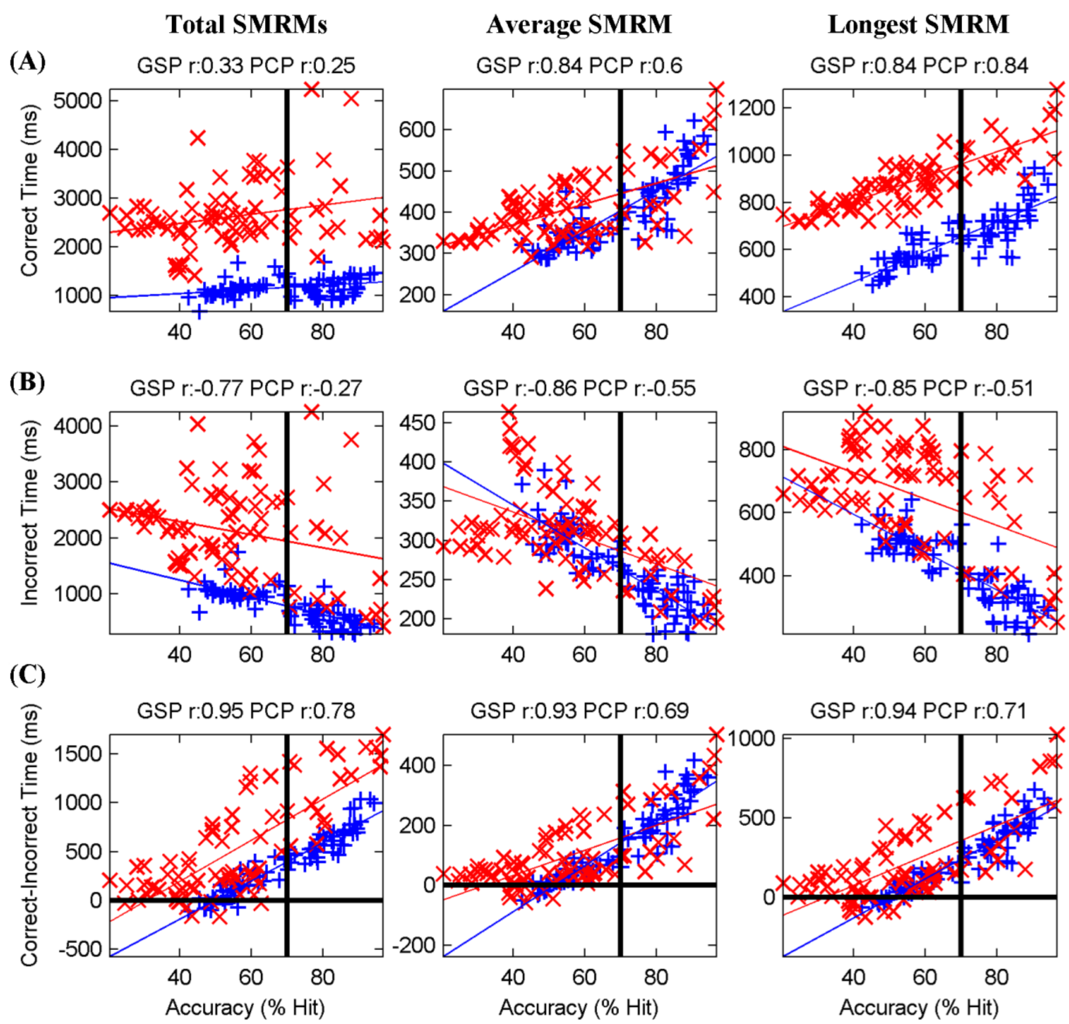

Figure 4.9 Accuracy correlations with modulation duration
Correct SMRM durations longer than incorrect SMRM durations led to acceptable BCI performance regardless of actual SMRM durations for both GSP and PCP. Each point is one subject's median accuracy vs. median SMRM duration for one session. The black vertical line emphasizes 70% accuracy, the minimum level of acceptable performance. The black horizontal line in C emphasizes correct longer than incorrect durations. Note in A and B the overlapping range of durations on each side of 70%, whereas in C ,all points greater than 70% accuracy fall in the upper quadrant. The headings in A apply to A-C. The legend in the middle column applies to all columns. All correlations were significant, 15 at p < 10-6 and 3 at p < 0.05.

9 plots the overall accuracy for each session against the median correct modulation (fig 9A), incorrect modulation (fig 9B), and the difference between time of correct and

incorrect modulation (fig 9C) for total time of modulation, average time of a single modulation, and the longest time of a single modulation. All of the duration measures were significantly correlated to accuracy. Total time of modulation, the average time of a single modulation, and the longest time of a single modulation exhibited the same trends. In general, a more accurate session had longer correct durations, shorter incorrect durations, and a larger difference between time of correct and incorrect modulation. However, the exact duration of correct or incorrect modulation did not predict acceptable performance. Note that in figure 9A-B the range of durations below 70% and above 70% overlap extensively. What did predict acceptable performance was the difference in time of correct modulation versus incorrect modulation. In figure 9C, all sessions with an accuracy greater than 70% had longer times of correct modulation than incorrect modulation, as can be seen from all points right of 70% falling in the upper quadrant of the graph. In brief, successful sessions had a wide variety of actual times of correct and incorrect modulation, but a successful session always had longer correct modulation than incorrect modulation.

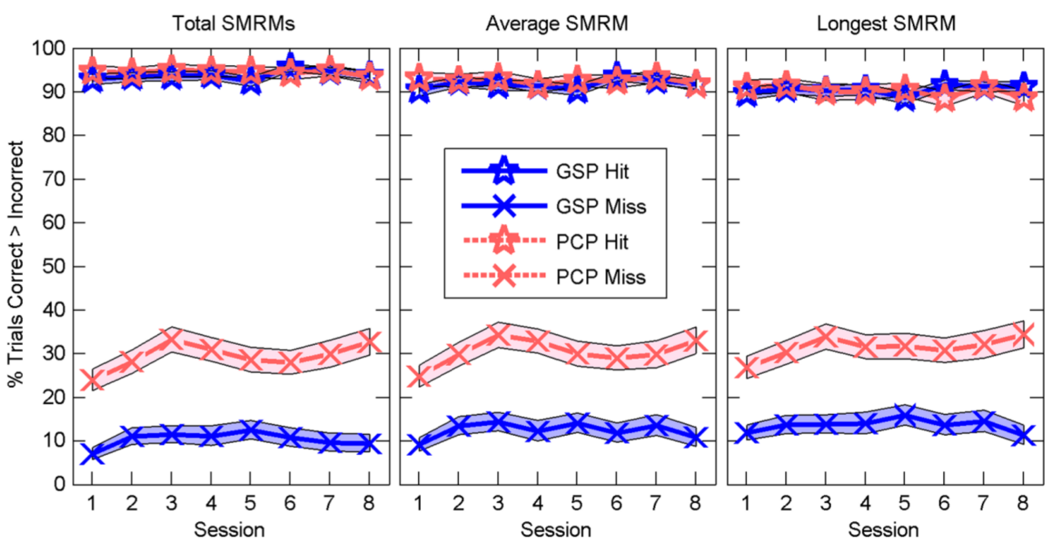

Figure 4.10 Correct modulation longer than incorrect modulation by trial
Correct modulations had a longer duration than incorrect modulations in more than 90% of all hit trials for both GSP and PCP. That held for total SMRMs as well as the average SMRM and the longest SMRM. Shaded area represents the 95% confidence interval.

Figure 9 presented median data from sessions. Figure 10 shows that the same conclusions can be drawn from the individual trials. In more than 90% of all trials that resulted in a hit, time of correct modulation was longer than time of incorrect modulation. That held for total duration of modulation, average duration of a single modulation, and longest duration of a single modulation. Correct modulation outlasted incorrect modulation for only about 20% of the trials that resulted in a miss. This shows that duration of correct modulation longer than incorrect modulation was very close to a necessary condition for a hit, even though it was not a sufficient condition.

4.3.6 Averaging trials removes information important to successful use

Figure 11 presents the data from this study using the classic methodology of averaging trials. As expected, the plots show continuous contralateral desynchronization and ipsilateral synchronization and look nearly identical to previous work (Yuan et al., 2008). The plots imply a near steady brain state that only waxes and wanes in amplitude for the length of the trial, in contrast to figures 2 through 10 above. The plots retain the amplitude information while removing the duration information.

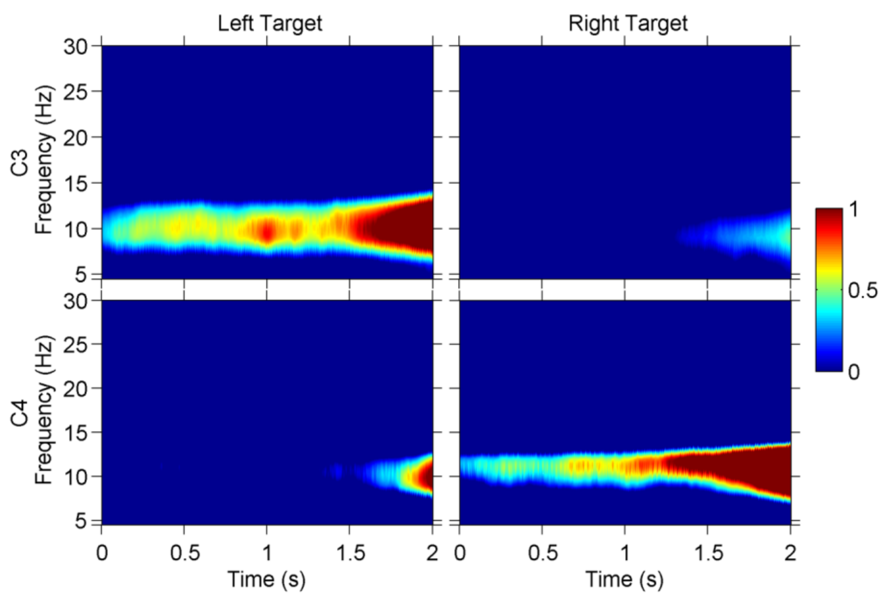

Figure 4.11 Grand average time frequency response of all data
Grand average time frequency response of all data, plotted using standard procedures of normalizing feedback time to 2s and averaging the power of all hit trials, yielded the standard expected result of contralateral desynchronization and ipsilateral synchronization that contains none of the duration information presented in the previous figures.

4.4 Discussion

In this study we analyzed individual trial data without performing the averaging that dominates the literature. We found trials to be a collection of sub-second sensorimotor rhythm modulations that alternated between correct and incorrect modulation. When both amplitude and duration of the modulations within a trial were analyzed, the duration measures were more correlated to successful BCI use than the amplitude measures. Successful use of a BCI exhibited durations of correct modulation longer than incorrect modulation, irrespective of the actual lengths of the modulations. That was the case for both success within a session leading to high accuracy, and success within a trial leading to a hit. When the trials were averaged according to the classical methodology of analyzing sensorimotor rhythms, we obtained the expected results of contralateral desynchronization and ipsilateral synchronization. Averaging the trials retained none of the duration information that was more correlated to success than was the amplitude information retained.

Amplitude of modulation was not entirely inconsequential. Brief periods of high-amplitude correct modulation accounted for the 8% of hit trials where the duration of incorrect modulation was longer than the duration of correct modulation. Similarly, brief periods of high-amplitude incorrect modulation led to the 12% of GSP miss trials that had longer durations of correct vs. incorrect modulation.

This study also analyzed the effect of control strategy. Subjects received assistance while using GSP and did not while using PCP. We found that correct modulation, either an individual correct sensorimotor rhythm modulation or the total modulation that led to a hit, was similar in amplitude and duration between PCP and GSP. Incorrect modulation in PCP had longer durations and higher amplitudes than in GSP. PCP also had longer durations and higher amplitudes of correct modulation in a miss trial than GSP.

The shorter trials of GSP proved to be a major advantage. Shorter trial times were significantly correlated to numerous measures of the EEG signal, from increased duration of correct modulation to greater amplitude differences between the right and left target. Shorter trials also allowed modulations of shorter durations to last for a longer percentage of the trial, allowing a single modulation to have a larger effect.

In this study, subjects performed a simple one-dimensional cursor task with a particular implementation of goal selection. Many other ways exist to implement goal selection into a cursor task, and they may be better than the ones tested here. These forms of goal selection were chosen to allow for ease of comparison to the process control paradigms. The key aspect of goal selection is that the user receives assistance executing their goal. Exactly how the user receives assistance needs to be customized to each individual task. Unfortunately, a direct comparison of process control to goal selection has not been presented for a more complicated task.

Surprisingly, the findings of this study only partially support our hypothesis that the underlying EEG signal would be more conducive to goal selection. The data did support the hypothesis in two ways. First off, GSP did have more favorable incorrect modulations with lower amplitude and shorter durations than PCP. Secondly, a PCP miss trial was much more similar to a hit trial than was a GSP miss trial, indicating that PCP worked harder to produce a more correct signal, but those efforts still resulted in a miss. However, the findings that correct modulation, either a single individual correct modulation or the total modulation that led to a hit, was similar in amplitude and duration between GSP and PCP is inconsistent with our hypothesis. This implies that the brain activity that properly operates a sensorimotor rhythm based BCI is the same regardless of control strategy. GSP was more able to take advantage of that brain activity and translate it into successful use of the BCI.

This study had a large amount of statistical power. We analyzed over 106 hours of EEG data. 20 subjects performed 8 sessions apiece of 10 runs of 15 to 45 trials. That resulted in a total of 46,036 trials. With each trial consisting of multiple sensorimotor rhythm modulations, we had a total of 432,722 modulations to quantify. With an n that high, the 95% confidence intervals were very narrow. Also with that many trials and modulations, it was impressive when more than 90% shared the characteristics discussed above.

This study is unique in the time span of the presented ERD/ERS data. Current literature mainly presents data from either naive or trained subjects (Neuper et al., 1999; Wolpaw and McFarland, 2004; Pfurtscheller et al., 2006; Yuan et al., 2008; Neuper et al., 2009; McFarland et al., 2010; Yuan et al., 2010a). This is the longest running sensorimotor rhythm study that presents data tracking subjects from naive through the learning process. This study is also unique in its combination of duration and large subject pool.

The duration of modulation was more correlated to successful use of the BCI than was amplitude of modulation even though the control signals were purposely chosen to allow increased control of the cursor through amplitude modulations. As presented in the methods, amplitude of the recorded EEG signal scales equally on both sides of the brain

(Yuan et al., 2010b). Since we limited the electrodes that controlled the movement of the cursor to a single side of the head, we retained the scaling of amplitude that other studies lost by subtracting amplitudes across the two hemispheres (Wolpaw and McFarland, 2004; Royer and He, 2009). Compared to previous studies, increasing the amplitude of modulation in this study had more effect on the movement of the cursor. However, the enhanced amplitude effects were still not as important to success than was duration.

Characterizing the EEG signal within single trials seems an intuitive way to analyze data from a BCI given that BCIs process the EEG signal and provide feedback to the user in time steps as short as 40 to 80 ms. The subject's experience with the BCI is determined by the moment to moment operation of their brain, and not the average of multiple trials. However, averaging is the norm for EEG analysis (Wu et al., 2011). Even studies that investigated durational effects of motor imagery averaged trials in their analysis (Pfurtscheller et al., 2008; Nam et al., 2011).

The prevalence of averaging BCI data in offline analysis, despite the non-relevance of averaging to online BCI performance, may be a contributing factor to the number of subjects that find it difficult to adequately operate a BCI. Work addressing this issue has used questionnaires and fMRI to probe the vividness of motor imagery of different subjects (Guillot et al., 2008; Halder et al., 2011). Analysis of the EEG signal has been used to predict acceptable BCI performance (Blankertz et al., 2010). However, to our knowledge, this is the first study that has performed analysis of the EEG signal at the same timescale as the BCI in order to address what distinguishes acceptable from unacceptable performance.

The thalamus is the gatekeeper to the cortex. Four of the five sensory modalities, motor control, and even cognitive abilities such as attention, memory, and emotion each have their own specific nuclei in the thalamus that connects them to the cortex. Coordinated activity within a corticothalamic loop produces rhythms in the 8-12Hz range measurable with electroencephalography (EEG) (Lopes da Silva, 1991; Pfurtscheller and Lopes da Silva, 1999). This circuitry has been specifically implicated in a wide variety of rhythmic activities, from the occipital alpha rhythm, initiation of epileptic seizures, and volitional control of a brain-computer interface (BCI) (Lopes da Silva, 1991; Pfurtscheller and Lopes da Silva, 1999; Breakspear et al., 2006).

Freyer et al. (2009) analyzed the human alpha rhythm at multiple time scales, including those at which a BCI operates. Contrary to the popular notion that resting alpha rhythm "waxes and wanes", they found that the rhythm burst erratically between a low-power and a high-power state in a multistable manner. This came about when the synaptic gains of the corticothalamic loop placed the system near a dynamical instability called a subcritical Hopf bifurcation (Freyer et al., 2011). When the corticothalamic loop was in that state, fluctuating noisy inputs into the thalamic neurons caused the jump from one power state to another. This noisy input may come from a wide variety of sources ranging from thermal effects to synaptic inputs from regions of the brain not directly involved in the corticothalamic loop. However, when the synaptic gains of the corticothalamic loop were higher, the system was no longer at the subcritical Hopf bifurcation, and consistent alpha activity was produced that merely waxed and waned with the addition of noise (Freyer et al., 2011).

Motor imagination activates the same circuitry as motor behavior, but at a reduced level (Yuan et al., 2010b). With training on a motor imagery based BCI, the activation can surpass that of motor behavior (Miller et al., 2010). These findings, combined with the work of Freyer et al. (2009, 2011), lead to some interesting theories. Perhaps a subject new to motor imagination activates the corticothalamic loop at reduced levels of gain that would place the system at the subcritical Hopf bifurcation. Then, when a subject imagines motor behavior, the two power modes compete, switching back and forth like in ocular dominance (Haynes et al., 2005). Our results of sub-second sensorimotor rhythm modulation are consistent with that theory. Learning to control a BCI may involve learning how to control the switch of those two states by increasing the gain of the system, increasing activation, and causing a more consistent alpha power output. Over the course of the 8 sessions of this study, GSP subjects exhibited more learning than PCP subjects (Royer et al., 2011). Interestingly enough, GSP subjects showed a significant negative correlation between number of modulations and accuracy ($r=-0.42$, $p<0.001$). As subjects improved their performance, they switched between power states less frequently, producing a more consistent modulation of their sensorimotor rhythms.

The same model used in Freyer et al. (2011) has been used to predict human epileptic seizures in the same frequency band, using the same mechanism (Breakspear et al., 2006). Neurofeedback therapy has long been considered as an alternative aid to epileptic patients

(Sterman et al., 1974). Despite the long duration of use, the exact technique is still being refined (Legarda et al., 2011; Weber et al., 2011). Currently, sensorimotor rhythm neurofeedback training sessions are part of the refining process (Weber et al., 2011). Increasing our knowledge of how subjects learn to use a sensorimotor rhythm based BCI as presented here may translate into improved therapies.

This study extends our understanding of the neurophysiology while using a sensorimotor rhythm based BCI. It shows that the improved performance while using a different control strategy in the translation process between raw brain signal and device command was more attributable to the difference in device command instead of the difference in raw brain signal. Our motor system itself uses different control strategies at different times, as illustrated in the opening accounts of walking versus dancing the Hokey Pokey for the first time. If researchers wish to circumvent and replace the normal motor output pathways with a BCI, it may be beneficial to learn from physiology and adapt the control strategy based on what is most appropriate for the current task. By understanding neurophysiology and applying the lessons learned to BCI design, we can make a better, more effective, and more useful BCI.

While learning to operate a motor imagery based BCI, subjects learn to volitionally control their sensorimotor rhythms. The circuitry that produces the sensorimotor rhythms is the same circuitry as that used in perception, motor control, and even more esoteric items such as emotion and epilepsy. The multistability of modulation presented by this study is also a feature of perception, decision making, and behavior (Schoener and Kelso, 1988; Haynes et al., 2005; Deco and Rolls, 2006). By understanding how humans purposely modulate this shared circuitry in the particular application of BCI, perhaps that can shed light on how people can better control their motor behavior, decision making, emotions, and even health issues such as epilepsy.

4.5 Electrodes and Frequencies Not Used for Control

4.5.1 Introduction

In this subsection, we describe the analysis of the electrodes and frequencies not used for control of the brain-computer interface in the study described in the previous two subsections as well as originally described in section 2.2.

4.5.2 Methods

The methods for both data collection and data analysis were as previously described in sections 2.2.2 and 4.2. Figure 4.12 was calculated and is displayed in the same manner as figure 2.2.2. Figure 4.13 includes all data, even that from the electrodes and frequencies used for control. Figures 4.14 and 4.15 were plotted in similar manners. Both plotted the median of the grouped data for each session excluding electrode and frequency combinations for each subject that contributed to their control signal. The rectangle represents the electrodes boxed in figure 4.12A. All data is presented with the subject's control signal on the right side of the rectangle, marked ipsi for ipsilateral to the control signal and cont for contralateral to the control signal. If the subject's control signal was on the left side of the head, their data was reflected across the z electrode line. Electrode color was scaled between a dark grey and white to indicate what proportion of control signals used each electrode and frequency combination, with dark grey indicating one and white indicating all control signals. In figure 4.14, only significant values were included in the grouped data. In figure 4.15, if a subject had a control signal on the left side of the head, the sign of the amplitude was flipped along with the plotted electrode locations.

4.5.3 Results

In order to better address the electrodes and frequencies not used for control, figure 4.12 serves as a reminder of the electrodes and frequencies that were used for control. Electrodes used in the control signals were limited to those over sensorimotor cortex, or the ones within the box drawn in figure 4.12A: FCz-6, Cz-6, and CPz-6. The color surrounding an electrode in figure 4.12B indicates the percentage of control signals that used each electrode in GSP (left column) and PCP (right column) with red indicating the highest number and blue indicating none. The circled electrodes were used in the most number of runs. The number of circles indicates the average number of electrodes used in that session's control signals. Figure 4.12C presents the frequencies used in each session for GSP (left column) and PCP (right column). The dark blue bars show the percentage of control signals that used each frequency bin. The light blue bars show the frequency bins that were used in the most number of runs, with the number of bars indicating the average number of frequency bins used in that session's control signals. For both GSP and PCP, subjects typically used two electrodes and two frequency bins. Electrodes C3, C4, CP3,

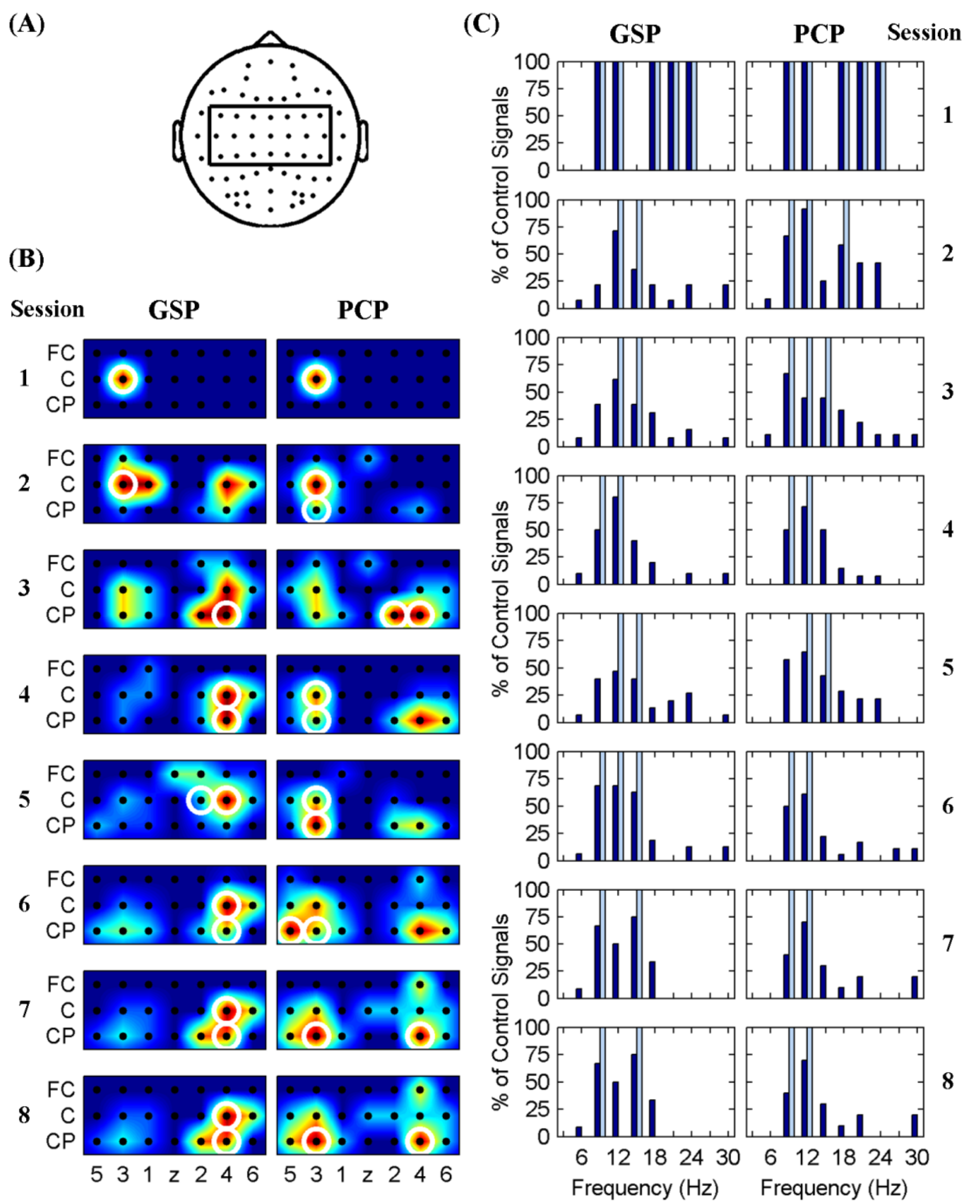

Figure 4.12 Electrodes and frequencies used in subjects' control signals

and CP4 were the most commonly used, with all subjects but one using at least one of the 9, 12, or 15Hz frequency bins in their final control signal.

As discussed in chapter 1, the statistical measure r^2 calculates the fraction of the total signal variance that is accounted for by the task right target vs. left target. Figure 4.13 presents the frequencies that had the highest significant r^2 in sensorimotor cortex (the boxed region in fig 4.12A) for all sessions grouped (fig 4.13A) as well as broken out by session (fig 4.13B). As can be seen in part A, the highest r^2 were from the 12 and 15 Hz bins for both

GSP and PCP. Frequencies above 30 Hz only had a smattering of sessions where those bins had the highest significant r^2. Those sessions are composed entirely of one-off sessions from various subjects. No two sessions came from the same subject. For 90% of all sessions, the highest r^2 in sensorimotor cortex was in a frequency bin less than or equal to 30Hz (fig 4.13B). Since the gamma band did not account for the left/right signal variance as well as the lower frequency bands, further analysis will focus on frequencies 30Hz and below.

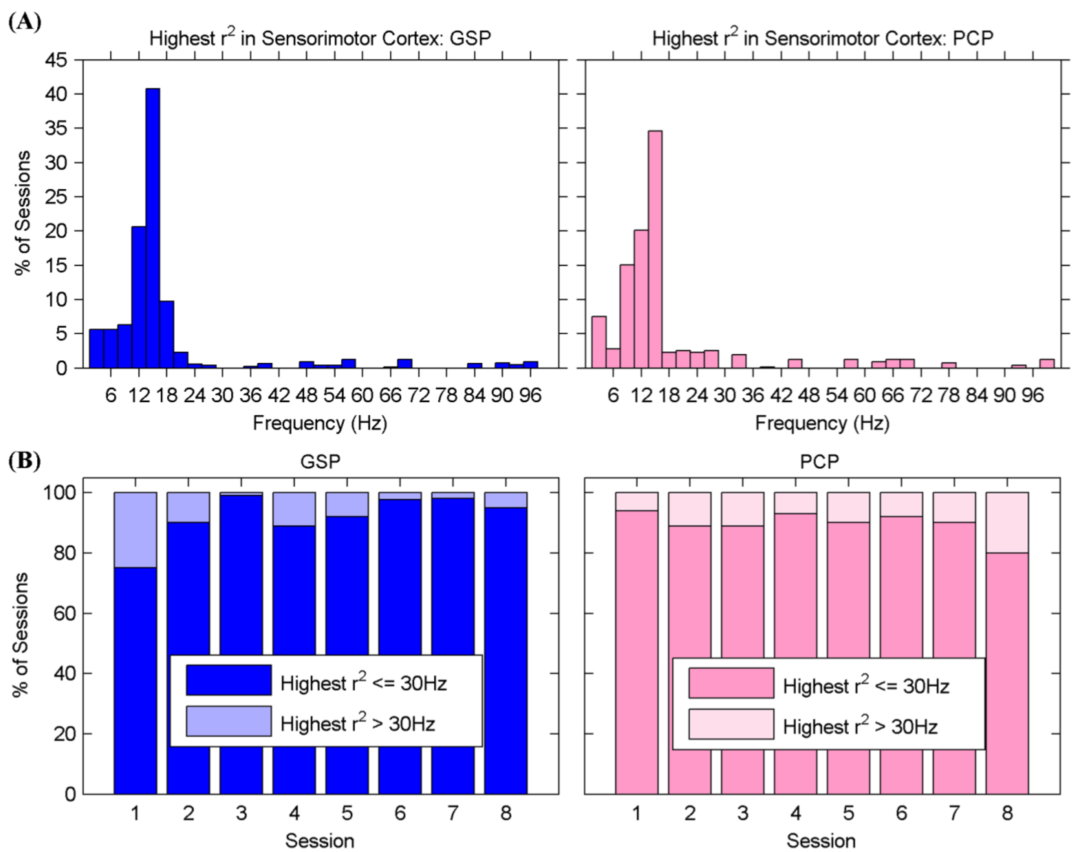

Figure 4.13 The highest significant r^2 in sensorimotor cortex

How closely did the signal on the non-control electrodes relate to the task at hand? To answer that question, r^2 was calculated for the non-control electrodes and frequencies and tested for significance. Figure 4.14 plots the median of the grouped data for sensorimotor cortex across all 8 sessions. In order to compare the data from the non-control electrodes to the control electrodes, the color of the electrode indicates to what extent it was included in the control signals. Electrode color was scaled between a dark grey and white to indicate what proportion of control signals used each electrode and frequency combination, with dark grey indicating one and white indicating all control signals. In order to compare ipsilateral and contralateral effects, all data is presented with the subject's control signal on the right side of the rectangle, marked ipsi for ipsilateral to the control signal and cont for contralateral to the control signal. If the subject's control signal was on the left side of the head, their data was reflected across the z electrode line. Median r^2 increased over time even for non-control electrodes and frequencies, with the effect focused in the 9 to 18Hz

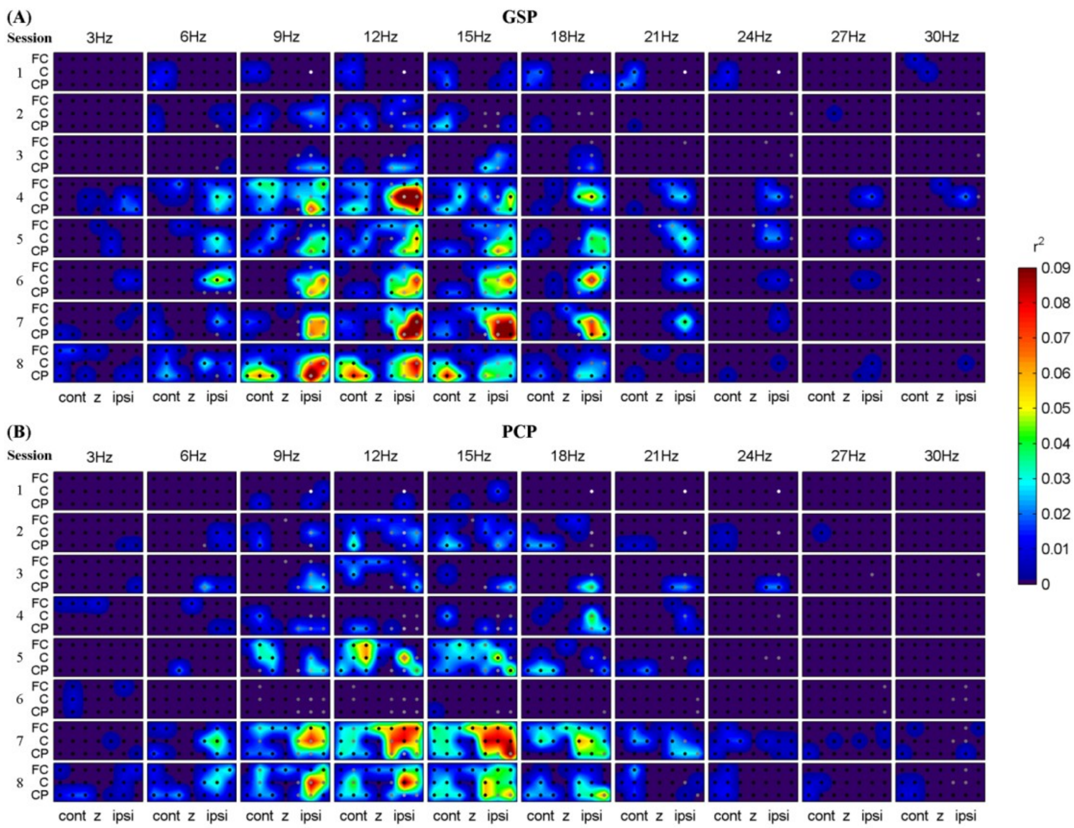

Figure 4.14 Median significant r^2 for non-control electrodes and frequencies
Purple indicates no significant r^2.

bins. GSP achieved higher r^2's than PCP, with the correlated electrodes mainly ipsilateral to the control signal whereas PCP had a more bilateral response.

How did the amplitudes of the non-control electrodes change over time? Figure 4.15 plots the data for the median right target spectral amplitude minus the left target spectral amplitude in a similar manner as used in figure 4.14. An additional modification was made that if a subject had a control signal on the left side of the head, the sign of the amplitude was flipped along with the plotted electrode locations. Spectral amplitude trends are not as clear as r^2 trends, as is expected from the lower correlation to performance shown in figure 4.8, but still focused in frequencies at 18Hz and lower. GSP and PCP showed bilateral results, with GSP better able to establish the expected pattern of ipsilateral increase and contralateral decrease. For GSP, amplitudes increased and then declined. PCP amplitudes increased steadily.

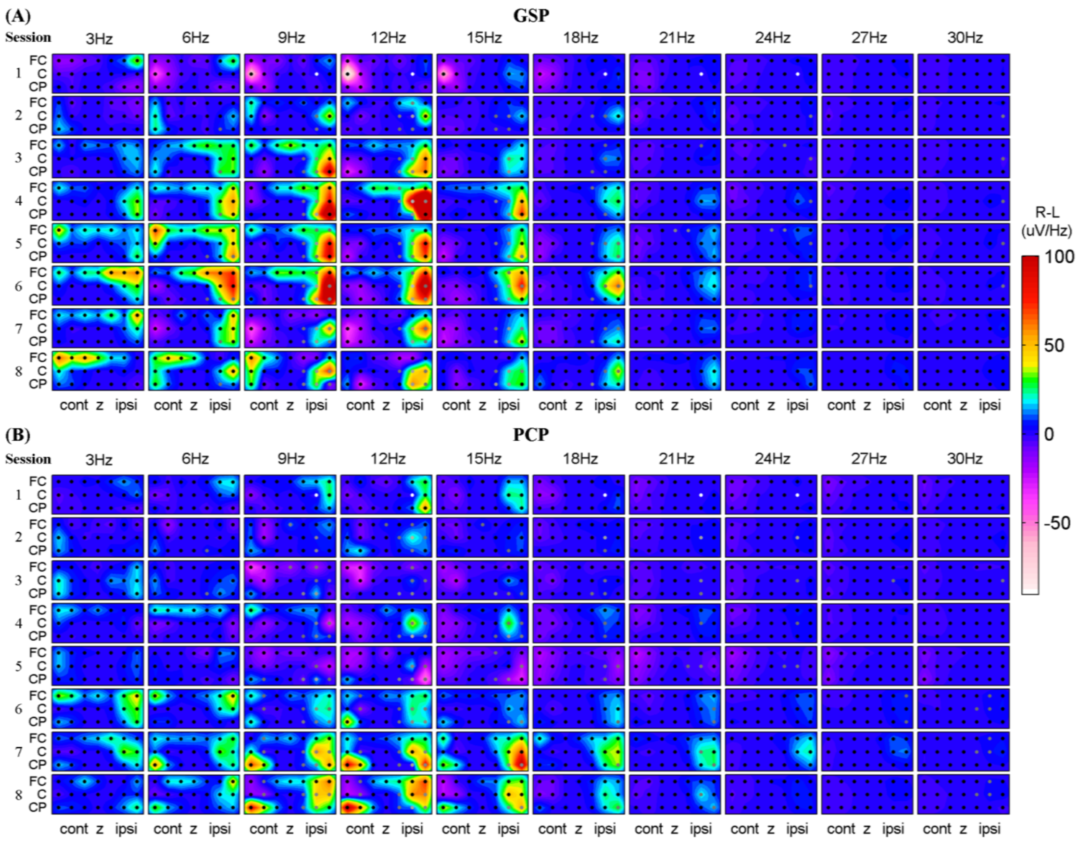

Figure 4.15 Median right - left spectral amplitude for non-control electrodes and frequencies

4.5.4 Discussion

As expected, the skull had the effect of smearing the signal. That made the EEG measures similar for neighboring electrodes. This smearing effect could have been reduced by using the surface Laplacian method (Pfurtscheller and Lopes da Silva 1999). In brief, the surface Laplacian can be approximated by subtracting the average value of the neighboring electrodes from the electrode of interest (Pfurtscheller and Lopes da Silva 1999). This methodology was purposely not chosen for control of the BCI during the study. In an effort to increase the chances that subjects could easily control the BCI, we wished to take advantage of the skull smearing effect to extend the area from which a useful control signal could be recorded. This helped compensate for small shifts in the EEG cap from session to session, or small variances in which electrode provided the best control.

Not surprisingly, the frequencies from 9 to 18 Hz were the most reactive for the presented measures. These were the frequencies that were most used in the control signal. These frequencies lie squarely in the mu and beta bands, where most motor imagination research is focused (Wolpaw et al 2002, Wolpaw and McFarland 2004, Yuan et al 2008, Yuan et al 2010a, Blankertz et al 2010).

Recently the BCI field has shown interest in the frequencies below 9 Hz (personal communication). Our results show activity in the 1.5-7.5 Hz range that equaled or exceeded that found in the more traditional beta band, especially the high beta. These results support the belief that the lower frequencies may be useful for BCI control.

Other recent research has shown promising results in the gamma band, or in frequencies above 30Hz (Grosse-Wentrup et al 2011). As the figures 4.13, 4.14, and 4.15 demonstrated, our results showed decreasing information in frequencies above 18Hz. The recording methodology of this study may have influenced this finding. Sensorimotor rhythms follow a 1/f noise shape. As the frequency increases, the power contained in the EEG signal rapidly falls off (Wolpaw et al 2002, Blankertz et al 2008). A minority of the sessions occurred in a noise shielded room. However, in the remainder of the sessions, the subjects sat in an area that was completely unshielded, increasing the noise in the recording. This noise was not sufficient to influence control of the BCI (personal experience). However, it may have been enough to reduce the usefulness of the gamma band. The studies that utilize the gamma band take extra noise precautions beyond just a shielded room, and this study did not even have that benefit consistently (Grosse-Wentrup et al 2011).

GSP showed more differential activity across the two hemispheres than did PCP in the data presented here. Previous research has shown that it is more difficult to produce differential activity than similar activity in the 8-12Hz range (Pineada et al 2003). As shown by section 2.2, GSP users were more fully trained than PCP users, and may have been able to more fully develop the more difficult differential activity. PCP users were still early enough in the training process that they may not had yet developed differential activity.

In conclusion, frequencies below 30Hz in this study contained the most information pertinent to successful completion of the task. The skull smeared the EEG signal, making the EEG measures similar for neighboring electrodes. The frequencies from 9 to 18 Hz were the most reactive for all measures. Additionally, GSP showed more differential brain activity across the two hemispheres than did PCP.

Chapter 5 Discussion

A brain-computer interface (BCI) strives to make a connection directly from a person's brain to a computer without relying on any motor output (Wolpaw et al 2002, Vallabhaneni et al 2005). BCIs promise to help the nearly 6 million people who live with paralysis (www.christopherreeve.org) by allowing them to interact with the world in ways they are no longer able. BCIs can also be used by able bodied individuals to extend their capabilities. The military hopes to utilize BCIs to enhance a soldier's abilities in the theatre of war (Kotchetkov et al 2010). In the civilian realm, current video games such as the Emotiv allow gamers to journey through a mythical world controlled by their mind (www.emotiv.com).

A BCI translates signals recorded directly from the brain into commands that control an external device, such as a computer cursor, wheelchair, or neuroprosthetic (Wolpaw et al 2002, Vallabhaneni et al 2005). BCIs differ widely in how they implement the translation from raw brain signal to device command (Wolpaw et al 2002, Wolpaw 2007). One way they differ is in the overall control strategy.

The two control strategies this thesis focused on are process control and goal selection. In process control, the user controls every step of the process and receives minimal to no assistance from the system. In goal selection, the user only needs to determine the goal and the system executes the process to achieve that goal. In goal selection, the system performs the work that was asked of the user in process control. Since in goal selection, less work is asked of the user, goal selection is intrinsically easier than process control.

Goal selection and process control may be referred to by other terminology: shared control vs. continuous control and high-level control vs. low-level control. Both sets of alternate terms have their own advantages, and may be preferred to goal selection vs. process control. The terms "goal selection" and "process control" were used throughout this thesis to remain consistent with the previously published works that compose this thesis.

Many BCIs perform a task that otherwise would be performed through motor output. A possible reason for the ataxic motion commonly produced by process control BCIs may be that BCIs gather their control signal from the cortex, bypassing all the complicated, trained interactions that produce normal motion (Wolpaw 2007). Process control BCIs call for the primary motor cortex to control all the fine motor details normally handled by other parts of the motor nervous system, such as the basal ganglia, cerebellum, or spinal cord. The user must ensure that they are properly encoding position, velocity, and/or acceleration while using a process control BCI.

In goal selection, the BCI uses the signal it obtains primarily from the cortex to determine the overall end goal of the user. This is more typical of the role cortex plays in normal motor control. The BCI execution unit then determines the necessary position, velocity, and/or acceleration parameters. Goal selection more closely resembles natural motor control with the BCI system assisting the user akin to how the distributed motor network assists the motor cortex (Wolpaw 2007). Since goal selection is easier and more natural, it follows that it would be more accurate, faster in use, and easier to learn. This system would have a higher information transfer rate, with a decreased training period.

This thesis presented a collection of papers on research analyzing goal selection as a control strategy in a BCI. Although several goal selection based BCIs exist, we presented the first studies directly comparing goal selection and process control. In the analysis, we looked at both behavioral performance metrics as well as analyzed the underlying differences in the EEG signal.

Chapter 2 presented 2 studies that directly compared goal selection and process control in a relatively simple left/right cursor task. This was a purposely simple task so we could more easily isolate the effect of changing the control strategy. Other possible implementations of goal selection into a cursor task exist, many of which may be better than what is described here. The paradigms described here were chosen for ease of comparison to the process control paradigms. Using these paradigms, we tested the hypothesis that goal selection would be more accurate, faster in use, and easier to learn, with a higher information transfer rate and a decreased training period. Section 2.1 focused on the first part of the hypothesis concerning speed and accuracy, whereas section 2.2 addressed learning.

In Section 2.1, we found the following to be true in both the trained and naïve populations studied: (1) The goal selection paradigms had more hits than the process control paradigms. (2) The goal selection paradigms were faster than the process control paradigms. (3) The goal selection paradigms were more accurate than the process control paradigms for most subjects and (4) the goal selection paradigms had a higher information transfer rate than the process control paradigms. In this study, the goal selection paradigms significantly outperformed the process control paradigms in every measure studied here. The primary goal of this study was to directly compare goal selection to process control, not produce the optimal goal selection based cursor task. However, even the un-optimized goal selection based paradigms produced 50% to 1600% improvement over the process control paradigms as tested here. Additionally, the goal selection based paradigms here outperformed all previous goal selection based sensorimotor rhythm BCIs. It is exciting to think of the performance that could be achieved with an optimized goal selection based paradigm. That is left to future work.

Section 2.2 completed testing the hypothesis by presenting a large scale, long term learning study. This study found that the goal selection paradigms were more accurate, faster to use, easier to learn, and required less mental effort than the process control paradigms. This study reproduced the findings of the previous section concerning speed and accuracy in a larger sample size. Median improvement from the process control paradigms to the goal selection paradigms across all sessions was 71% for accuracy, number of hits per run, time to hit, and information transfer rate. This study was also the first to show that the goal selection paradigms were easier to learn and required less mental effort than the process control paradigms. The goal selection paradigms showed on average twice the learning and required 54% less effort than the process control paradigms. In this study, across all five measures presented, the goal selection paradigms were significantly better and showed significantly more learning than the process control paradigms. Even in the best subjects, the goal selection paradigms were significantly better than the process control paradigms, showing a median improvement of 41% with these paradigms.

The study presented in section 2.2 is unique in its combination of duration and large subject pool. It is the longest running sensorimotor rhythm study that presented data tracking subjects from naïve to trained. This study was long enough and the subjects used the BCI frequently enough to actually demonstrate significant learning. That is rare for BCI experiments, even those that consider learning effects (Kubler et al 2010). Despite this study's length, it was not quite long enough to fully train the process control paradigms' subjects. The goal selection paradigms' subjects could be considered trained, but still refining their skills. However, the process control paradigms' subjects were not fully trained, but were still learning. This study also demonstrated that learning did transfer between the goal selection based paradigms and the process control based paradigms in section 2.1.

The task performed in chapter 2 was a simple left/right cursor task. Chapter 3 presented data from a more complicated, real world task. Subjects were asked to continuously fly a virtual helicopter to any point in three-dimensional space. This chapter presented the first accomplishment by a non-invasive BCI of a task previously performed only by invasive BCIs. This work provides a platform for the development of 3D non-invasive BCIs that expands the user population, reduces training barriers, and optimizes control signal economy. The three-dimensional world we live in demands complex functionality from BCI systems. Here we demonstrated the potential of non-invasive systems to meet that demand.

The accomplishments of chapter 3 were made possible by using intelligent control strategies, including goal selection. In approximately half the runs, subjects shared control with the BCI system while using the cone of guidance, which directed the helicopter through the ring during the final approach. This effectively transformed the target ring into a target "balloon". This is analogous to the GSFD (goal selection with feedback limited by distance) paradigm from chapter 2, where close is good enough, the BCI took care of the final details.

All subjects were able to successfully fly through the ring without the assistance of the cone of guidance. However, the cone of guidance did improve performance. Subjects had general control of the helicopter and were able to fly to the ring, as shown by the minimal change between using the cone of guidance or not in the average path and time spent closer to the ring than the helicopter began. Subjects had varying degrees of fine control, and that is where the assistance of the cone of guidance significantly influenced the metrics. In the grouped data, the cone of guidance significantly improved the number of rings per minute, path length, and time to a ring. The cone of guidance doubled the number of rings per minute, reduced the time to ring by 27-65%, and reduced the path length by a third. These effects were seen even if the cone of guidance was applied in post analysis. Using the cone of guidance allowed subjects to have comparable control with the BCI as they did when using the keyboard. The impact of the cone of guidance varied by subject, and the metrics of the weaker subjects improved more than the metrics of the better subjects. Using the cone of guidance allowed this study to have 4 subjects instead of 2 since, using both the cone and no cone data, subjects 3 and 4 did not perform statistically the same as chance for any paradigm. That was not the case if only no cone data was used. Fine control may come with training, reducing the need for the cone of guidance. We cannot say anything more definitive since this was not a long term study. However, the assistance provided by the cone of guidance allowed the system to be useful to a wider population.

As seen in the BCI to keyboard ratios, the subjects found continuous control harder than trial-based, even when distances were similar. One possible reason for this, as mentioned in interactions with the subjects, was that continuous control did not provide any mental breaks. An additional advantage of the cone of guidance was that, while the helicopter was under the control of the BCI system, the user could mentally relax and not attempt to control their EEG. Our personal experience with the subjects made it clear that they took advantage of this opportunity to take a small mental break until the next ring was presented.

The combination of requiring less effort and naturally introducing breaks leads to less overall mental fatigue from using a goal selection based BCI when compared to a process control based BCI. This has been seen in our personal experience with subjects across all studies presented in this thesis. As discussed in Bai et al (2010), the minimization of fatigue during BCI use will be important as BCIs move from laboratory to clinical settings. The patient populations that many BCIs are designed to serve, such as those with ALS, have reduced physical and mental endurance (Sykacek et al 2003, Birbaumer 2006). This diminished endurance has decreased the accuracy of a BCI system with 90% accuracy in healthy subjects to levels just over chance in the patient population (Sellers and Donchin 2006, Iversen et al 2008). Therefore, reduction of fatigue due to using goal selection as a control strategy may aid the usefulness and adoption of BCIs by individuals who truly need them to restore lost functionality.

If we wish to use BCIs to help individuals that can no longer rely on their own natural motor output, it will be important to make using the BCI as effective and as simple as possible. Applying goal selection in the BCI's control strategy may make the system easier to learn, decrease the training period, and provide improved speed, accuracy, and information transfer. These improvements will also help make BCIs more appealing to able-bodied users.

In chapter 4, we addressed the underlying neural signal while using goal selection. This chapter analyzed the EEG data from section 2.2, but was the first study to analyze the EEG from a BCI at the same timescale as the BCI, and without averaging trials. It also presented the longest running EEG analysis tracking subjects from naive to trained. Sections 4.1- 4.4 analyzed the effect of goal selection vs. process control on the underlying neural signal from the electrodes and frequencies that were used for control. Section 4.5 presented the effect of goal selection vs. process control on the electrodes and frequencies that were not used for control. Taken together, chapter 4 extends our understanding of the neurophysiology while using a sensorimotor rhythm based BCI.

When we analyzed individual trial data without performing the averaging that dominates the literature, we found trials to be a collection of sub-second sensorimotor rhythm modulations that alternated between correct and incorrect modulation. When both amplitude and duration of the modulations within a trial were analyzed, the duration measures were more correlated to successful BCI use than the amplitude measures. Successful use of a BCI exhibited durations of correct modulation longer than incorrect modulation, irrespective of the actual lengths of the modulations. That was the case for both success within a session leading to high accuracy, and success within a trial leading to a hit. When the trials were averaged according to the classical methodology of analyzing sensorimotor rhythms, we obtained the expected results of contralateral desynchronization and ipsilateral synchronization. Averaging the trials retained none of the duration information that was more correlated to success than was the amplitude information retained. The prevalence of averaging BCI data in offline analysis, despite the non-relevance of averaging to online BCI performance, may be a contributing factor to the number of subjects that find it difficult to adequately operate a BCI (Vidaurre et al 2011).

This study had a large amount of statistical power. We analyzed over 106 hours of EEG data. 20 subjects performed 8 sessions apiece of 10 runs of 15 to 45 trials. That resulted in a total of 46,036 trials. With each trial consisting of multiple sensorimotor rhythm modulations, we had a total of 432,722 modulations to quantify. With an n that high, the 95% confidence intervals were very narrow. Also with that many trials and modulations, it was impressive when more than 90% shared the characteristics discussed above.

When looking at the effects of goal selection vs. process control on the EEG signal, we found that correct modulation, either an individual correct sensorimotor rhythm modulation or the total modulation that led to a hit, was similar in amplitude and duration between the process control paradigms and the goal selection paradigms. The process control paradigms had longer durations and higher amplitudes of incorrect modulation, while having longer durations and higher amplitudes of correct modulation in a miss trial than the goal selection paradigms. Shorter trial times were significantly correlated to numerous measures of the EEG signal, from increased duration of correct modulation to greater amplitude differences between the right and left target. Hence, the shorter trials of the goal selection paradigms proved to be a major advantage over the process control paradigms.

Section 4.5 showed that frequencies below 30Hz in this study contained the most information pertinent to successful completion of the task. The skull smeared the EEG signal, making the EEG measures similar for neighboring electrodes. The frequencies from 9 to 18 Hz were the most reactive for all measures. Additionally, the goal selection paradigms showed more differential brain activity across the two hemispheres than did the process control paradigms, which is consistent with goal selection paradigms' subjects being more fully trained as shown by section 2.2.

Chapter 4 showed that the improved performance while using a different control strategy in the translation process between raw brain signal and device command was more attributable to the difference in device command instead of the difference in raw brain signal. Our motor system itself uses different control strategies at different times, as illustrated in section 1.4 in the accounts of walking versus dancing the Hokey Pokey for the first time. If researchers wish to circumvent and replace the normal motor output pathways with a BCI, it may be beneficial to learn from physiology and adapt the control strategy based on what is most appropriate for the current task. By understanding neurophysiology and applying the lessons learned to BCI design, we can make a better, more effective, and more useful BCI.

Goal selection demonstrated many advantages over process control in this thesis. However, goal selection does have limitations. In order for the BCI system to assist the user, the system needs to be pre-programmed to provide the correct assistance. That requires two things: 1) that the situation be an anticipated, known event, and 2) that the programmed assistance actually benefit the user. The assistance provided by a goal selection based BCI is only as effective as it is programmed to be. Situations could exist where the programmed assistance may not help the user.

The major advantage of process control is that it provides unlimited possibilities for action, making it indispensible when encountering a novel situation or event. Ideally, a BCI would be able to assist the user as often as possible using goal selection, while still allowing the freedom that process control provides. Additionally, an ideal BCI would learn from the new encounter to possibly provide assistance in the future. This main distinction between goal selection and process control implies that, however much benefit a BCI derives from implementing goal selection, process control will continue to be employed as BCI use increases in society.

An interesting example of many of the concepts discussed in the previous paragraph can be found in Rebsamen et al (2010). They described a brain controlled wheelchair to navigate in familiar environments intended to be used by the disabled population. Because of the intended users, Rebsamen et al prioritized safety and efficiency. The wheelchair used a P300 based BCI to select one of the predetermined paths through the home that had been previously programmed into the BCI. Once the path was selected, the wheelchair smoothly and safely travelled to the intended destination. While the wheelchair was in motion, it could be stopped by a faster, motor imagery based BCI. Subjects could stop within 4.9s with no false activations. Here, the predetermined paths represent the goal selection portion of the BCI, whereas the stop command represents the process control portion of the BCI.

As demonstrated throughout this thesis, goal selection leverages the expertise of both brain and computer to create a system more powerful than either individually. Brain-computer interfaces are intrinsically multi-disciplinary. Neuroscience, engineering, and computer science combine to create a complex system. The future of BCIs lies in leveraging the potential of all disciplines involved. BCI systems with intelligent control strategies like goal selection may prove to be the foundation for complex BCIs capable of doing more than we ever imagined.

This thesis presents the first studies directly comparing the control strategies of process control and goal selection in a BCI. In these studies, we found that the goal selection paradigms were easier to learn, had a decreased training period, and provided improved speed, accuracy, and information transfer in both simple and more complex applications. This thesis also extends our understanding of the neurophysiology while using a sensorimotor rhythm based BCI. When individual trial data were analyzed and not averaged as is typically done in the literature, we discovered that duration of sensorimotor rhythm modulation was more correlated to successful use than amplitude of modulation. Additionally, we found that correct modulation that led to either a single hit or overall high accuracy was the same between the two control strategies. This shows that the improved performance while using the goal selection paradigms was more attributable to the difference in device command instead of the difference in raw brain signal. By understanding neurophysiology and applying the lessons learned to BCI design, we can make a better, more effective, and more useful BCI.

References

Bai O, Lin P, Huang D, Fei DY and Floeter MK 2010 Towards a user-friendly brain-computer interface: Initial tests in ALS and PLS patients *Clin Neurophysiol*. **121**(8) 1293-1303

Baillet S, Mosher JC, Leahy RM 2001 Electromagnetic brain mapping *IEEE Sign. Proc. Mag*. **18** 14-30

Bell CJ, Shenoy P, Chalodhorn R and Rao RPN 2008 Control of a humanoid robot by a noninvasive brain-computer interface in humans *J. Neural Eng*. **5** 214-220

Birbaumer N 2006 Breaking the silence: brain–computer interfaces (BCI) for communication and motor control *Psychophysiology* **43** 517–32

Blankertz B, Dornhege G, Krauledat M, Müller KR and Curio G 2007 The non-invasive Berlin Brain-Computer Interface: fast acquisition of effective performance in untrained subjects *Neuroimage* **37** 539-50

Blankertz B, Losch F, Krauledat M, Dornhege G, Curio G and Mueller KR 2008 The Berlin brain–computer interface: accurate performance from first-session in BCI naive subjects *IEEE Trans. Biomed. Eng.* **55** 2452-62

Blankertz B, Sannelli C, Halder S, Hammer EM, Kübler A, Müller KR, Curio G, Dickhaus T 2010 Neurophysiological predictor of SMR-based BCI performance *Neuroimage* **51** 1303-1309

Bradberry TJ, Gentili RJ, Contreras-Vidal JL 2010 Reconstructing three-dimensional hand movements from noninvasive electroencephalographic signals *J Neurosci* **30** 3432-3437

Breakspear M, Roberts JA, Terry JR, Rodrigues S, Mahant N, Robinson PA 2006 A unifying explanation of primary generalized seizures through nonlinear brain modeling and bifurcation analysis *Cereb Cortex* **16** 1296–1313

Brunner C, Scherer R, Graimann B, Supp G and Pfurtscheller G 2006 Online control of a brain-computer interface using phase synchronization *IEEE Trans. Biomed. Eng.* **53** 2501-6

Carlson, N Physiology of Behavior, 6th ed. Needham Heights, MA: Allyn and Bacon, 1998, pp.234-236, 242-245

Dale AM, Sereno MI 1993 Improved localization of cortical activity by combining EEG and MEG with MRI cortical surface reconstruction: A linear approach *J. Cogn. Neurosci* **5** 162-176

Deco G, Rolls ET 2006 Decision-making and Weber's law: a neurophysiological model *Eur J Neurosci* **24** 901–916

Donchin E, Spencer KV and Wijesinghe R 2000 The mental prosthesis: Assessing the speed of a P300-based brain–computer interface *IEEE Trans. Rehab. Eng*. **8** 174–9

Duysensa J and A Van de Crommertb HWA 1998 Neural control of locomotion: Part 1: The central pattern generator from cats to humans *Gait & Posture* **7** 131-41

Ehrsson HH, Geyer S, Naito E 2003 Imagery of voluntary movement of fingers, toes, and tongue activates corresponding body-part-specific motor representations *J Neurophysiol* **90** 3304-16

Eiccles JC 1951 Interpretation of action potentials evoked in the cerebral cortex *Electroencephalogr. Clin. Neurophysiol* **3** 449-464

Farwell A and Donchin E 1988 Talking off the top of your head: toward a mental prothesis utilizing event-related brain potentials *Electroenceph. Clin. Neurophysiol*. **70** 510–23

Fink A, Grabner RH, Neuper C and Neubauer AC 2005 EEG alpha band dissociation with increasing task demands *Brain Res Cogn Brain Res*. **24** 252-9

Freyer F, Aquino K, Robinson PA, Ritter P and Breakspear M 2009 Bistability and non-gaussian fluctuations in spontaneous cortical activity *J. Neurosci*. **29** 8512-24

Freyer F, Roberts JA, Becker R, Robinson PA, Ritter P, Breakspear M 2011 Biophysical mechanisms of multistability in resting-state cortical rhythms *J Neurosci* **31** 6353-6361

Friedrich EVC, McFarland DJ, Neuper C, Vaughan TM, Brunner P and Wolpaw JR 2008 A scanning protocol for a sensorimotor rhythm-based brain–computer interface *Biological Psychology* doi:10.1016/j.biopsycho.2008.08.004

Friston K 2005 A theory of cortical responses *Philos Trans R Soc Lond B Biol Sci* **360** 815–836

Galán F, Nuttin M, Lew E, Ferrez PW, Vanacker G, Philips J, Millán JDR 2008 A brain-actuated wheelchair: asynchronous and non-invasive Brain-computer interfaces for continuous control of robots *Clin. Neurophysiol*. **119** 2159-2169

Golowasch J, Goldman MS, Abbott LF, Marder E 2002 Failure of averaging in the construction of a conductance-based neuron model *J. Neurophysiol*. **87** 1129-31

Grashow R, Brookings T, Marder E 2010 Compensation for variable intrinsic neuronal excitability by circuit-synaptic interactions *J Neurosci* **30** 9145-9156

Green AM, Kalaska JF 2011 Learning to move machines with the mind *Trends Neurosci* **34** 61-75

Grillner S and Zangger P 1979 On the central generation of locomotion in the low spinal cat *Exp. Brain Res* **34** 241-61

Grosse-Wentrup M, Schölkopf B, Hill J 2011 Causal influence of gamma oscillations on the sensorimotor rhythm *Neuroimage* **56** 837-42

Guger C, Edlinger G, Harkam W, Niedermayer I, Pfurtscheller G 2003 How many people are able to operate an EEG-based brain-computer interface (BCI)? *IEEE Trans. Neural Syst. Rehabil. Eng.* **11** 145-7

Guillot A, Collet C, Nguyen VA, Malouin F, Richards C, Doyon J 2008 Functional neuroanatomical networks associated with expertise in motor imagery *Neuroimage* **41** 1471-1483

Halder S, Agorastos D, Veit R, Hammer EM, Lee S, Varkuti B, Bogdan M, Rosenstiel W, Birbaumer N, Kübler A 2011 Neural mechanisms of brain-computer interface control *Neuroimage* **55** 1779-1790

Haynes JD, Deichmann R, Rees G 2005 Eye-specific effects of binocular rivalry in the human lateral geniculate nucleus *Nature* **438** 496–499

Hinterberger T, Veit R, Wilhelm B, Weiskopf N, Vatine J and Birbaumer N 2005 Neuronal mechanisms underlying control of a brain–computer interface *European Journal of Neuroscience* **21** 3169–81

Hochberg LR, Serruya MD, Friehs GM, Mukand J, Saleh M, Caplan AH, Branner A, Chen D, Penn RD and Donoghue JP 2006 Neuronal ensemble control of prosthetic devices by a human with tetraplegia *Nature* **442** 164-71

Iversen I, Ghanayim N, Kubler A, Neumann N, Birbaumer N and Kaiser J 2008 Conditional associative learning examined in a paralyzed patient with amyotrophic lateral sclerosis using brain–computer interface technology *Behav Brain Funct* **4** 53

Kamousi B, Amini AN and He B 2007 Classification of motor imagery by means of cortical current density estimation and Von Neumann entropy for brain-computer interface applications *Journal of Neural Engineering* **4** 17-25

Kandel ER, Schwartz JH, Jessell TM 2000 Principles of neural science, 4th Ed. McGraw Hill

Karim AA, Hinterberger T, Richter J, Mellinger J, Neumann N, Flor H, Kübler A and Birbaumer N 2006 Neural internet: Web surfing with brain potentials for the completely paralyzed *Neurorehabil Neural Repair* **20** 508-515

Keil A, Mussweiler T and Epstude K 2006 Alpha-band activity reflects reduction of mental effort in a comparison task: a source space analysis *Brain Res*. **1121** 117-27

Kennedy PR, Bakay RA, Moore MM, Adams K and Goldwaithe J 2000 Direct control of a computer from the human central nervous system *IEEE Trans. Rehabil. Eng.* **8** 198-202

Kim HK, Biggs SJ, Schloerb DW, Carmena JM, Lebedev MA, Nicolelis MA, Srinivasan MA 2006 Continuous shared control for stabilizing reaching and grasping with brain-machine interfaces *IEEE Trans. Biomed. Eng* **53** 1164-1173

Kotchetkov IS, Hwang BY, Appelboom G, Kellner CP and Connolly ES Jr 2010 Brain-computer interfaces: military, neurosurgical, and ethical perspective *Neurosurgical FOCUS* **28** E25

Krusienski DJ, Schalk G, McFarland DJ and Wolpaw JR 2007 A mu-rhythm matched filter for continuous control of a brain-computer interface *IEEE Trans. Biomed. Eng.* **54** 273-80

Kubler A, Donatella M, Harry G, Benaya D and Neuper C 2010 How much learning is involved in BCI-control? http://www.bcimeeting.org/2010/poster_abstracts.shtml#A7

Kübler A, Neumann N, Wilhelm B, Hinterberger T, Birbaumer N 2004 Predictability of brain–computer communication *Int. J. Psychophysiol* **18** 121–129

Kunst CB 2004 Complex genetics of amyotrophic lateral sclerosis *Am. J. Hum. Genet.* **75** 933–47

Legarda SB, McMahon D, Othmer S, Othmer S 2011 Clinical Neurofeedback: Case Studies, Proposed Mechanism, and Implications for Pediatric Neurology Practice *J Child Neurol* May 16. [Epub ahead of print]

Leuthardt EC, Schalk G, Moran D, Ojemann JG 2006 The emerging world of motor neuroprosthetics: a neurosurgical perspective *Neurosurgery* **59** 1-14

Lopes da Silva F 1991 Neural mechanisms underlying brain waves: from neural membranes to networks *Electroencephalogr Clin Neurophysiol* **79** 81-93

McCormick DA1992 Neurotransmitter actions in the thalamus and cerebral-cortex and their role in neuromodulation of thalamocortical activity *Prog Neurobiol* **39** 337–388

McFarland DJ, Dean J Krusienski DJ, Sarnacki WA and Wolpaw JR 2008 Emulation of computer mouse control with a noninvasive brain–computer interface *J. Neural Eng.* **5** 101–10

McFarland DJ, McCane LM and Wolpaw JR 1998 EEG-based communication and control: short-term role of feedback *IEEE Trans. Rehabil. Eng.* **6** 7-11

McFarland J, Sarnacki W A and Wolpaw JR 2010 Electroencephalographic (EEG) control of three-dimensional movement J Neural Eng. 7 036007

Miller KJ, Schalk G, Fetz EE, den Nijs M, Ojemann JG, Rao RP 2010 Cortical activity during motor execution, motor imagery, and imagery-based online feedback *Proc Natl Acad Sci U S A* **107** 4430-5 Erratum in: *Proc Natl Acad Sci U S A* (2010) 10715:7113

Musallam S, Corneil BD, Greger B, Scherberger H and Andersen RA 2004 Cognitive control signals for neural prosthetics *Science* **305** 258–62

Nam CS, Jeon Y, Kim YJ, Lee I, and Park K 2011 Movement imagery-related lateralization of event-related (de)synchronization (ERD/ERS): motor-imagery duration effects *Clin. Neurophysiol*. **122** 567-77

Neuper C, Grabner RH, Fink A and Neubauer AC 2005 Long-term stability and consistency of EEG event-related (de-)synchronization across different cognitive tasks *Clin Neurophysiol*. **116** 1681-94

Neuper C, Scherer R, Wriessnegger S and Pfurtscheller G 2009 Motor imagery and action observation: modulation of sensorimotor brain rhythms during mental control of a brain-computer interface *Clin Neurophysiol*. **120**(2) 239-47

Neuper C, Schlögl A and Pfurtscheller G 1999 Enhancement of left-right sensorimotor EEG differences during feedback-regulated motor imagery *J. Clin. Neurophysiol.* **16** 373-82

Pfurtscheller G and Aranibar A 1979 Evaluation of event-related desynchronization (ERD) preceding and following voluntary self-paced movement *Electroencephalogr. Clin. Neurophysiol*. **46** 138-46

Pfurtscheller G and Lopes da Silva FH 1999 Event-related EEG/MEG synchronization and desynchronization: basic principles *Clinical Neurophysiology* **110** 1842-57

Pfurtscheller G, Brunner C, Schloegl A and Lopes da Silva FH 2006 Mu rhythm (de)synchronization and EEG single-trial classification of different motor imagery tasks *NeuroImage* **33** 153–9

Pfurtscheller G, Scherer R, Müller-Putz GR and Lopes da Silva FH 2008 Short-lived brain state after cued motor imagery in naive subjects *Eur. J. Neurosci.* **28** 1419-26

Pineda JA, Silverman DS, Vankov A and Hestenes J 2003 Learning to control brain rhythms: making a brain-computer interface possible *IEEE Trans Neural Syst Rehabil Eng* **11** 181–4

Prinz AA, Bucher D, Marder E 2004 Similar network activity from disparate circuit parameters *Nat Neurosci* **7** 1345-1352

Qin L and He B 2005 A wavelet-based time-frequency analysis approach for classification of motor imagery for brain-computer interface applications *Journal of Neural Engineering* **2** 65-72

Ray WJ and Cole HW 1985 EEG alpha activity reflects attentional demands, and beta activity reflects emotional and cognitive processes *Science* **228** 750-2

Rebsamen B, Guan C, Zhang H, Wang C, Teo C, Ang MH Jr, Burdet E 2010 A brain controlled wheelchair to navigate in familiar environments *IEEE Trans Neural Syst Rehabil Eng* **18** 590-8

Ron-Angevin R and Díaz-Estrella A 2009 Brain-computer interface: changes in performance using virtual reality techniques *Neurosci Lett* **449** 123-127

Royer AS and He B 2009 Goal selection versus process control in a brain-computer interface based on sensorimotor rhythms *J. Neural Eng.* **6** 016005

Royer AS, Doud AJ, Rose ML, He B 2010 EEG control of a virtual helicopter in 3-dimensional space using intelligent control strategies *IEEE Trans Neural Syst Rehabil Eng*. **18** 581-9

Royer AS, Rose ML and He B 2011 Goal Selection vs. Process Control while Learning to Use a Brain-Computer Interface *J Neural Eng.* **8** 036012

Royer AS, Rose ML, He B 2011b Goal Selection vs. Process Control in Non-invasive Brain-Computer Interface *Conf Proc IEEE Eng Med Biol Soc*. (in press)

Santhanam G, Ryu SI, Yu BM, Afshar A and Shenoy KV 2006 A high-performance brain-computer interface *Nature* **442** 195-8

Schalk G and Mellinger J 2010 *A Practical Guide to Brain-Computer Interfacing with BCI2000* (London: Springer)

Schalk G, McFarland DJ, Hinterberger T, Birbaumer N and Wolpaw JR 2004 BCI2000: a general-purpose brain-computer interface (BCI) system *IEEE Trans. Biomed. Eng*. **51** 1034–43

Scherer R, Schloegl A, Lee F, Bischof H, Jansa J, Pfurtscheller G 2007 The self-paced graz brain-computer interface: methods and applications *Comput. Intell. Neurosci* **2007** 79826

Schöner G, Kelso JA 1988 Dynamic pattern generation in behavioral and neural systems *Science* **239** 1513–1520

Schwartz AB, Cui XT, Weber DJ, Moran DW 2006 Brain-controlled interfaces: movement restoration with neural prosthetics *Neuron* **52** 205-20

Sellers EW and Donchin E 2006 A P300-based brain–computer interface: initial tests by ALS patients. *Clin Neurophysiol* **117** 538–48

Sellers EW, Vaughan TM, Wolpaw JR 2010 A brain-computer interface for long-term independent home use *Amyotroph Lateral Scler* **11** 449-55

Shenoy P , Krauledat M, Blankertz B, Rao RPN and Mueller KR 2006 Towards adaptive classification for BCI J. *Neural Eng.* **3** R13–R23

Stephan KE, Kasper L, Harrison LM, Daunizeau J, den Ouden HE, Breakspear M, Friston KJ 2008 Nonlinear dynamic causal models for fMRI *Neuroimage* **42** 649–662

Sterman MB, Macdonald LR, Stone RK 1974 Biofeedback training of the sensorimotor electroencephalogram rhythm in man: effects on epilepsy *Epilepsia* **15** 395-416

Sykacek P, Roberts S, Stokes M, Curran E, Gibbs M and Pickup L 2003 Probabilistic methods in BCI research *IEEE Trans Neural Syst Rehabil Eng* **11** 192–5

Taylor DM, Helms Tillery SI, Schwartz AB 2002 Direct cortical control of 3D neuroprosthetic devices *Science* **296** 1829-1832

Vallabhaneni A, Wang T and He B 2005 Brain computer interface *Neural Engineering* ed B He (New York: Kluwer Academic/Plenum) pp 85-122

Vanacker G, Millán JDR, Lew E, Ferrez PW, Moles FG, Philips J, Van Brussel H, Nuttin M 2007 Context-based filtering for assisted brain-actuated wheelchair driving, *Comput. Intell. Neurosci* **2007** 25130

Velliste M, Perel S, Spalding MC , Whitford AS and Schwartz AB 2008 Cortical control of a prosthetic arm for self-feeding *Nature* **453** 1098-1101

Vidaurre C, Sannelli C, Müller KR, Blankertz B 2011 Machine-Learning-Based Coadaptive Calibration for Brain-Computer Interfaces *Neural Comput* **23** 791-816

Wang T, Deng J, He B 2004 Classifying EEG-based Motor Imagery Tasks by means of Time-frequency Synthesized Spatial Patterns *Clinical Neurophysiology* **115** 2744-2753

Weber E, Köberl A, Frank S, Doppelmayr M 2011 Predicting successful learning of SMR neurofeedback in healthy participants: methodological considerations *Appl Psychophysiol Biofeedback* **36** 37-45

Wolpaw JR 2007 Brain–computer interfaces as new brain output pathways *J. Physiol*. **579** 613-9

Wolpaw JR and McFarland DJ 2004 Control of a two-dimensional movement signal by a noninvasive brain–computer interface in humans *Proc. Natl. Acad. Sci. U. S. A*. **101** 17849–54

Wolpaw JR, Birbaumer N, McFarland DJ, Pfurtscheller G and Vaughan TM 2002 Brain–computer interfaces for communication and control *Clinical Neurophysiology* **113** 767–91

Wu W, Chen Z, Gao S and Brown EN 2011 A hierarchical Bayesian approach for learning sparse spatio-temporal decompositions of multichannel EEG *Neuroimage* **56** 1929-45

Yuan H, Doud AJ, Gururajan A and He B 2007 Cortical imaging of sensorimotor rhythm during on-line control of brain-computer interface *Proc. of NFSI&ICFBI Conf.* 327-9

Yuan H, Doud AJ, Gururajan A and He B 2008 Cortical imaging of event-related (de)synchronization during online control of brain-computer interface using minimum-norm estimates in the frequency domain *IEEE Trans. Neural Sys. & Rehab. Eng.* **16** 425-31

Yuan H, Liu T, Szarkowski R, Rios C, Ashe J, He B 2010a Negative covariation between task-related responses in alpha/beta-band activity and BOLD in human sensorimotor cortex: an EEG and fMRI study of motor imagery and movements *Neuroimage* **49** 2596-2606

Yuan H, Perdoni C and He B 2010b Relationship between Speed and EEG Activity during Imagined and Executed Hand Movements *J. Neural Eng.* **7** 26001

Yuan H, Perdoni C, Yang L and He B 2011 Differential Electrophysiological Coupling for Positive and Negative BOLD Responses during Unilateral Hand Movements *J. Neurosci.* **31** 9585-93

Copyright Permissions

The following were used with permission:

Royer AS and He B 2009 Goal selection versus process control in a brain-computer interface based on sensorimotor rhythms *J Neural Eng*. **6** 016005

Royer AS, Rose ML and He B 2011 Goal Selection vs. Process Control while Learning to Use a Brain-Computer Interface *J Neural Eng.* **8** 036012

Royer AS, Doud AJ, Rose ML, He B 2010 EEG control of a virtual helicopter in 3-dimensional space using intelligent control strategies *IEEE Trans Neural Syst Rehabil Eng*. **18** 581-9