Published in final edited form as:
J Neural Eng. 2011 June ; 8(3): 036012. doi:10.1088/1741-2560/8/3/036012.

Goal Selection vs. Process Control while Learning to Use a Brain-Computer Interface

Audrey S Royer[1], **Minn L Rose**[2], and **Bin He**[1,2,*]

[1]Graduate Program in Neuroscience, University of Minnesota, Minneapolis, MN 55455

[2]Department of Biomedical Engineering, University of Minnesota, Minneapolis, MN 55455

Abstract

A brain-computer interface (BCI) can be used to accomplish a task without requiring motor output. Two major control strategies used by BCIs during task completion are process control and goal selection. In process control, the user exerts continuous control and independently executes the given task. In goal selection, the user communicates their goal to the BCI and then receives assistance executing the task. A previous study has shown that goal selection is more accurate and faster in use. An unanswered question is, which control strategy is easier to learn? This study directly compares goal selection and process control while learning to use a sensorimotor rhythm based BCI. Twenty young healthy human subjects were randomly assigned either to a goal selection or a process control based paradigm for 8 sessions. At the end of the study, the best user from each paradigm completed 2 additional sessions using all paradigms randomly mixed. The results of this study were that goal selection required a shorter training period for increased speed, accuracy, and information transfer over process control. These results held for the best subjects as well as in the general subject population. The demonstrated characteristics of goal selection make it a promising option to increase the utility of BCIs intended for both disabled and able bodied users.

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

*Correspondence: Bin He, Ph.D. University of Minnesota Department of Biomedical Engineering 7-105 NHH, 312 Church Street, SE Minneapolis, MN 55455 binhe@umn.edu.

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). 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.

NIH-PA Author Manuscript

NIH-PA Author Manuscript

2. Methods

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) 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 1A, 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 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) 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 1). 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

NIH-PA Author Manuscript

NIH-PA Author Manuscript

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).

The details of how each paradigm progressed is shown in figure 1B. In PCA, the subjects had 6 s to hit a target with the cursor (figure 1B, 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 1B, 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 centred between the two targets (figure 1B, 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 1B, 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'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).

2.3 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).

2.4 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 comparable amounts (Yuan *et al* 2010). 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, 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, 2010). 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 2A). The frequencies were limited to the 3Hz bins centred 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 neighbouring 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 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). 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.

2.6 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:

$$\log 2\mathrm{N}+\mathrm{P}\log 2\mathrm{P}+(1-\mathrm{P})\ \log 2\left[(1-\mathrm{P})/(\mathrm{N}-1)\right] \tag{1}$$

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 3 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 4, 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 5 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 centred on a multiple of three from 0 to 30 Hz. Window length was 160ms with 50% overlap.

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 are presented below.

3.1 Control signal evolution

Figure 2B and C show the evolution of the control signal across the eight sessions, both in terms of electrodes used (figure 2B) and frequencies used (figure 2C). The colour 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 2C. 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.

NIH-PA Author Manuscript

NIH-PA Author Manuscript

3.2 Eight session performance metrics

For both GSP and PCP, accuracy increased over the 8 sessions (figure 3A). 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 levelled 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 3B). 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 3D). 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 3C). 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 levelled 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.

The effort of hit decreased for both paradigms over the 8 sessions (figure 3E). 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.

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 4 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 4A). GSP was 13% more accurate than PCP (figure 4B). GSP had 41% more hits per run than PCP (figure 4C). GSP was 31% faster to a hit than PCP (figure 4D). GSP required 44% less effort for a hit than PCP (figure 4E). Across all 5 measures, the median increase in performance was 41% from PCP to GSP.

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 levelled 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). 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).

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 2, 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). 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). Therefore, valid comparisons between the two studies include comparing the naive subjects (Royer and He 2009) to sessions 1 and 2, the trained subjects (Royer and He 2009) to sessions 7 and 8, and the trained subjects (Royer and He 2009) 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) 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) 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) 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) 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). 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).

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 criteria 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 3, 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* 2010). 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 5. 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 3 and 4, 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 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 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.

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). 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.

Supplementary Material

Refer to Web version on PubMed Central for supplementary material.

Acknowledgments

The authors are grateful to Alex Doud, Han Yuan, Dan Rinker, and Cristina Rios for assistance in data collection and many useful discussions. This work was supported in part by NSF grant CBET-0933067, NIH Grants T32 EB008389, RO1 EB007920, and RO1 EB006433.

References

Bai O, Lin P, Huang D, Fei DY, Floeter MK. Towards a user-friendly brain-computer interface: Initial tests in ALS and PLS patients. Clin Neurophysiol. 2010; 121(8):1293–1303. [PubMed: 20347612]

Bell CJ, Shenoy P, Chalodhorn R, Rao RPN. Control of a humanoid robot by a noninvasive brain-computer interface in humans. J. Neural Eng. 2008; 5:214–220. [PubMed: 18483450]

Birbaumer N. Breaking the silence: brain–computer interfaces (BCI) for communication and motor control. Psychophysiology. 2006; 43:517–32. [PubMed: 17076808]

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

Farwell LA, Donchin E. Talking off the top of your head: toward a mental prothesis utilizing event-related brain potentials. Electroenceph. Clin. Neurophysiol. 1988; 70:510–23. [PubMed: 2461285]

Fink A, Grabner RH, Neuper C, Neubauer AC. EEG alpha band dissociation with increasing task demands. Brain Res Cogn Brain Res. 2005; 24:252–9. [PubMed: 15993763]

Hochberg LR, Serruya MD, Friehs GM, Mukand J, Saleh M, Caplan AH, Branner A, Chen D, Penn RD, Donoghue JP. Neuronal ensemble control of prosthetic devices by a human with tetraplegia. Nature. 2006; 442:164–71. [PubMed: 16838014]

Iversen I, Ghanayim N, Kubler A, Neumann N, Birbaumer N, Kaiser J. Conditional associative learning examined in a paralyzed patient with amyotrophic lateral sclerosis using brain–computer interface technology. Behav Brain Funct. 2008; 4:53. [PubMed: 19025641]

Kamousi B, Amini AN, He B. Classification of motor imagery by means of cortical current density estimation and Von Neumann entropy for brain-computer interface applications. Journal of Neural Engineering. 2007; 4:17–25. [PubMed: 17409476]

Karim AA, Hinterberger T, Richter J, Mellinger J, Neumann N, Flor H, Kübler A, Birbaumer N. Neural internet: Web surfing with brain potentials for the completely paralyzed. Neurorehabil Neural Repair. 2006; 20:508–515. [PubMed: 17082507]

Keil A, Mussweiler T, Epstude K. Alpha-band activity reflects reduction of mental effort in a comparison task: a source space analysis. Brain Res. 2006; 1121:117–27. [PubMed: 17010944]

Kennedy PR, Bakay RA, Moore MM, Adams K, Goldwaithe J. Direct control of a computer from the human central nervous system. IEEE Trans. Rehabil. Eng. 2000; 8:198–202. [PubMed: 10896186]

Kotchetkov IS, Hwang BY, Appelboom G, Kellner CP, Connolly ES Jr. Brain-computer interfaces: military, neurosurgical, and ethical perspective. Neurosurgical FOCUS. 2010; 28:E25. [PubMed: 20568942]

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

Kunst CB. Complex genetics of amyotrophic lateral sclerosis. Am. J. Hum. Genet. 2004; 75:933–47. [PubMed: 15478096]

McFarland DJ, Sarnacki WA, Wolpaw JR. Electroencephalographic (EEG) control of three-dimensional movement. J Neural Eng. 2010; 7:036007. [PubMed: 20460690]

Neuper C, Grabner RH, Fink A, Neubauer AC. Long-term stability and consistency of EEG event-related (de-)synchronization across different cognitive tasks. Clin Neurophysiol. 2005; 116:1681–94. [PubMed: 15922658]

Neuper C, Scherer R, Wriessnegger S, Pfurtscheller G. Motor imagery and action observation: modulation of sensorimotor brain rhythms during mental control of a brain-computer interface. Clin Neurophysiol. 2009; 120(2):239–47. [PubMed: 19121977]

Neuper C, Schlögl A, Pfurtscheller G. Enhancement of left-right sensorimotor EEG differences during feedback-regulated motor imagery. J. Clin. Neurophysiol. 1999; 16:373–82. [PubMed: 10478710]

Pfurtscheller G, Lopes da Silva FH. Event-related EEG/MEG synchronization and desynchronization: basic principles. Clinical Neurophysiology. 1999; 110:1842–57. [PubMed: 10576479]

Pfurtscheller G, Brunner C, Schloegl A, Lopes da Silva FH. Mu rhythm (de)synchronization and EEG single-trial classification of different motor imagery tasks. NeuroImage. 2006; 33:153–9. [PubMed: 16443377]

Pineda JA, Silverman DS, Vankov A, Hestenes J. Learning to control brain rhythms: making a brain-computer interface possible. IEEE Trans Neural Syst Rehabil Eng. 2003; 11:181–4. [PubMed: 12899268]

Ray WJ, Cole HW. EEG alpha activity reflects attentional demands, and beta activity reflects emotional and cognitive processes. Science. 1985; 228:750–2. [PubMed: 3992243]

Royer AS, He B. Goal selection versus process control in a brain-computer interface based on sensorimotor rhythms. J. Neural Eng. 2009; 6:016005. [PubMed: 19155552]

Santhanam G, Ryu SI, Yu BM, Afshar A, Shenoy KV. A high-performance brain-computer interface. Nature. 2006; 442:195–8. [PubMed: 16838020]

Schalk G, McFarland DJ, Hinterberger T, Birbaumer N, Wolpaw JR. BCI2000: a general-purpose brain-computer interface (BCI) system. IEEE Trans. Biomed. Eng. 2004; 51:1034–43. [PubMed: 15188875]

Schalk, G.; Mellinger, J. A Practical Guide to Brain-Computer Interfacing with BCI2000. Springer; London: 2010.

Sellers EW, Donchin E. A P300-based brain-computer interface: initial tests by ALS patients. Clin Neurophysiol. 2006; 117:538–48. [PubMed: 16461003]

Sellers EW, Vaughan TM, Wolpaw JR. A brain-computer interface for long-term independent home use. Amyotroph Lateral Scler. 2010; 11:449–55. [PubMed: 20583947]

Smith ME, McEvoy LK, Gevins A. The impact of moderate sleep loss on neurophysiologic signals during working-memory task performance. Sleep. 2002; 25:784–94. [PubMed: 12405615]

Sykacek P, Roberts S, Stokes M, Curran E, Gibbs M, Pickup L. Probabilistic methods in BCI research. IEEE Trans Neural Syst Rehabil Eng. 2003; 11:192–5. [PubMed: 12899272]

Vallabhaneni, A.; Wang, T.; He, B. Brain computer interface. In: He, B., editor. Neural Engineering. Kluwer Academic/Plenum; New York: 2005. p. 85-122.

Velliste M, Perel S, Spalding MC, Whitford AS, Schwartz AB. Cortical control of a prosthetic arm for self-feeding. Nature. 2008; 453:1098–1101. [PubMed: 18509337]

Wolpaw JR. Brain–computer interfaces as new brain output pathways. J. Physiol. 2007; 579:613–9. [PubMed: 17255164]

Wolpaw JR, McFarland DJ. Control of a two-dimensional movement signal by a noninvasive brain–computer interface in humans. Proc. Natl. Acad. Sci. U. S. A. 2004; 101:17849–54. [PubMed: 15585584]

Wolpaw JR, Birbaumer N, McFarland DJ, Pfurtscheller G, Vaughan TM. Brain–computer interfaces for communication and control. Clinical Neurophysiology. 2002; 113:767–91. [PubMed: 12048038]

Yuan H, Doud AJ, Gururajan A, He B. 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. 2008; 16:425–31.

Yuan H, Perdoni C, He B. Relationship between Speed and EEG Activity during Imagined and Executed Hand Movements. J. Neural Eng. 2010; 7:26001. [PubMed: 20168002]

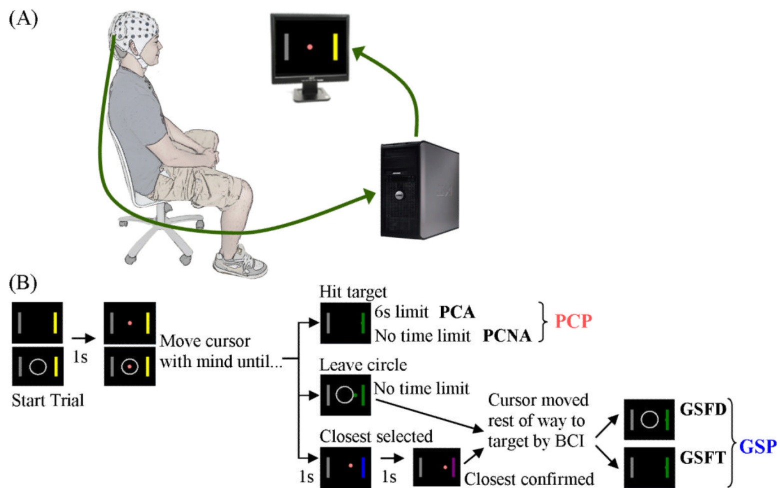

Figure 1.
Experimental setup and paradigms. (A) Healthy human subjects sat motionless in a chair facing a computer monitor and imagined moving their right or left hand, arm, or shoulder. Their EEG signal was sent to a computer, which translated the raw EEG into a control signal that moved a computer 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.

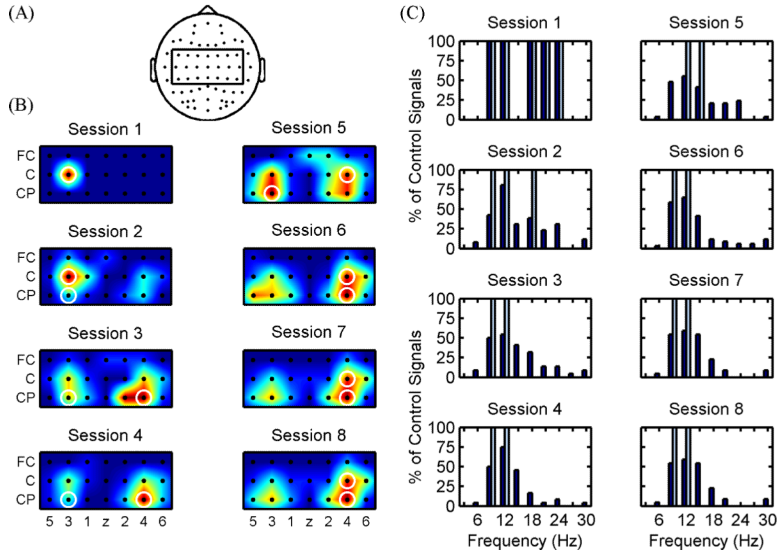

Figure 2.

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.

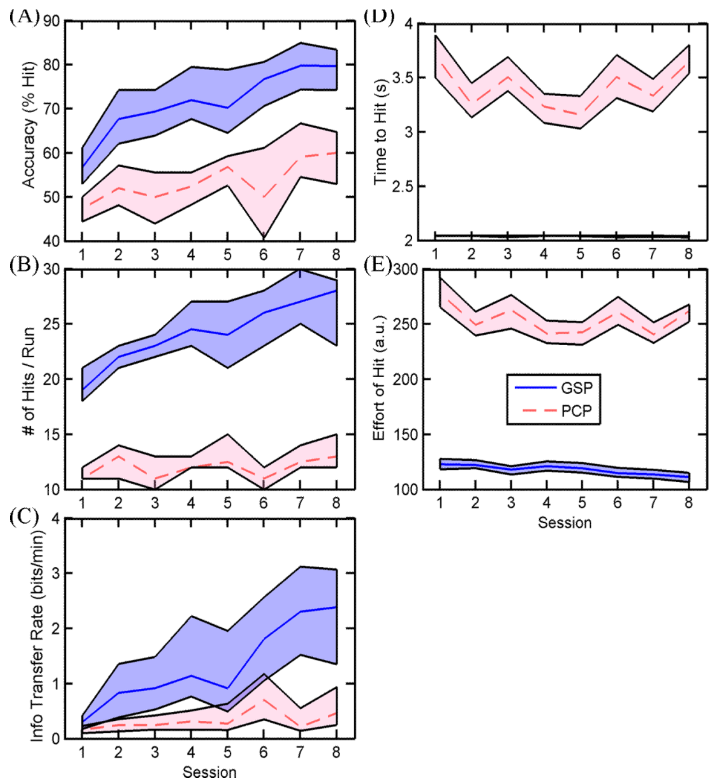

Figure 3.
Goal selection was significantly better and showed significantly more learning than process control over the eight sessions. Plotted lines are the medians of the grouped data for each session. Blue (darker, solid line) represents GSP. Pink (lighter, broken line) represents PCP. The legend in (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.

NIH-PA Author Manuscript

NIH-PA Author Manuscript

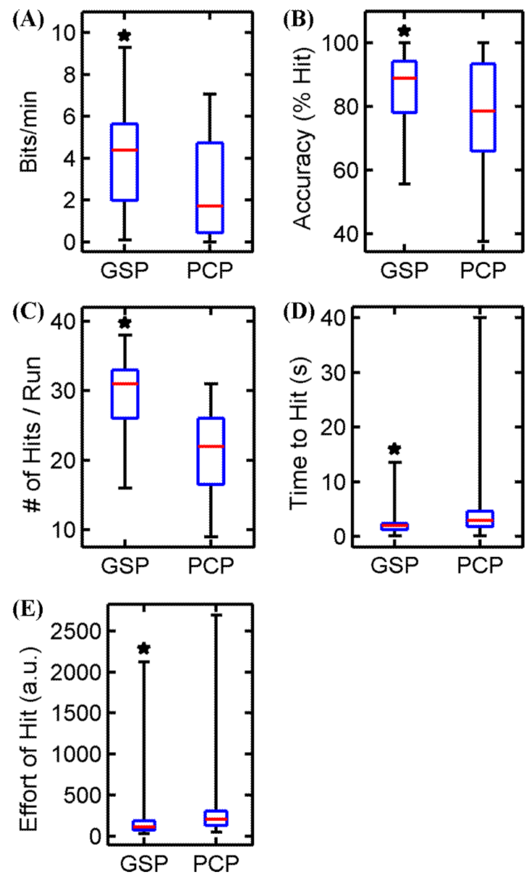

Figure 4.

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.

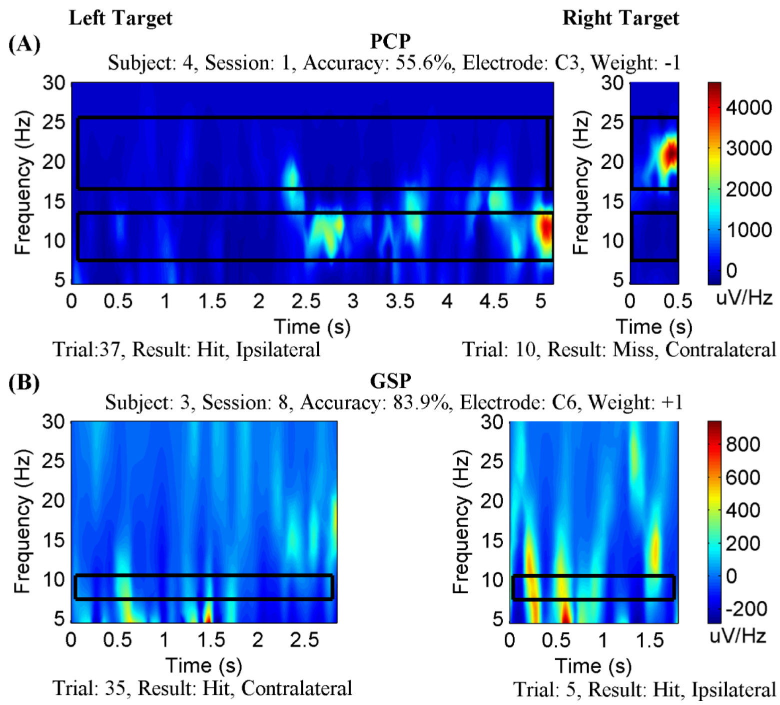

Figure 5.

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 centred across the figure apply to both right and left plots. Words below a figure apply to the individual plot. Colour scale bars on the right apply to the entire row.