Published in final edited form as:

J Neural Eng. 2010 February ; 7(1): 16003. doi:10.1088/1741-2560/7/1/016003.

A Comparison of Regression Techniques for a Two-Dimensional Sensorimotor Rhythm-Based Brain Computer Interface

Joan Fruitet¹, Dennis J. McFarland², and Jonathan R. Wolpaw²

¹Ecole Normale Supérieure, 45 rue d'Ulm, 75230 Paris Cedex 05, France

²Laboratory of Neural Injury and Repair, Wadsworth Center, New York State Department of Health and State University of New York, Albany, NY 12201 USA

Abstract

People can learn to control electroencephalogram (EEG) features consisting of sensorimotor-rhythm amplitudes and use this control to move a cursor in one, two, or three dimensions to a target on a video screen. This study evaluated several possible alternative models for translating these EEG features into two-dimensional cursor movement by building an off-line simulation using data collected during on-line performance. In offline comparisons, support-vector regression (SVM) with a radial basis kernel produced somewhat better performance than simple multiple regression, the LASSO, or a linear SVM. These results indicate that proper choice of a translation algorithm is an important factor in optimizing BCI performance, and provide new insight into algorithm choice for multidimensional movement control.

Keywords

Brain Computer interface (BCI); electroencephalography (EEG); regression; classification; support vector regression; kernel

1. Introduction

Many people with severe motor disabilities require alternative methods for communication and control. Numerous studies over the past two decades show that scalp-recorded EEG activity can be the basis for non-muscular communication and control systems, commonly called brain-computer interfaces (BCIs) (Birbaumer et al, 1999; Farwell and Donchin, 1988; Pfurtscheller et al, 1993; Wolpaw et al, 1991). EEG-based communication systems measure specific features of EEG activity and use the results as control signals. Some BCI systems use features that are potentials evoked by stereotyped stimuli (Farwell and Donchin, 1988). Others use EEG components in the frequency domain that are spontaneous in the sense that they are not dependent on specific sensory events (e.g., Wolpaw and McFarland, 2004).

Effective BCI operations depends on use of appropriate methods for recording brain signals, extracting features from these signals, and translating these features into device commands (Wolpaw et al, 2002). The feature translation algorithm is extremely important, and has been the subject of numerous studies and data competitions (e.g. Blankertz et al, 2006; see McFarland et al, 2006 for a review). An effective translation algorithm weights relevant features in proportion to the information they contain about the user's desired outcome. Many kinds of algorithms, linear as well as nonlinear are possible (e.g., Muller et al, 2003).

With the Wadsworth sensorimotor rhythm (SMR)-based BCI system, users learn over a series of training sessions to use SMR amplitudes in mu (8–13 Hz) and/or beta (18–27 Hz) frequency bands recorded over left and/or right sensorimotor cortex to move a cursor in one,

two, or three dimensions (see McFarland et al, 2006 for full system description). At present, the translation algorithm used for online control is a simple linear multiple regression.

This study compares the present translation algorithm to several promising alternatives in terms of performance in two-dimensional SMR-based cursor control. The algorithms were compared through offline analyses of data recorded while subjects controlled the cursor with the standard algorithm. This is the first formal comparison of alternative translation algorithms in regard to two-dimensional control. Multidimensional control presents unique challenges, such as the need for independence among the signals that control the separate dimensions (e.g., Wolpaw and McFarland, 2004).

2. Methods

2.1. Data Collection and Preprocessing

2.1.1. BCI Users—The BCI users were 8 adults (5 men and 3 women, ages 20–39). Six had no disabilities, while two had spinal cord injuries (at C6 and T7) and were confined to wheelchairs. All gave informed consent for the study, which had been reviewed and approved by the New York State Department of Health Institutional Review Board. After an initial evaluation defined the frequencies and scalp locations of each person's spontaneous mu and beta rhythm (i.e., sensorimotor-rhythm(SMR)) activity, he or she learned EEG-based cursor control over several months (2–3 24-min sessions/week). The standard online protocol, which has been described in previous publications (McFarland et al, 1997; Schalk et al, 2004; Wolpaw and McFarland, 2004), is summarized here.

2.1.2. Standard Online Protocol—The user sat in a reclining chair facing a 51-centimeter video screen three m away, and was asked to remain motionless during performance. Scalp electrodes recorded 64 channels of EEG (Sharbrough et al, 1991), each referenced to an electrode on the right ear (amplification 20,000; bandpass 0.1–60 Hz; sampling rate 160 Hz).

A daily session consisted of 8 three-min runs separated by one-min breaks, and each run had 20–30 trials. Each trial consisted of a 1-sec period from target appearance to the beginning of cursor movement, a period of cursor movement of 15-sec maximum, a 1.5-sec post-movement reward period, and a 1-sec inter-trial interval. Users participated in 2–3 sessions/week at a rate of one every 2–3 days. The last seven sessions for each user were used for this offline analysis. The total number of available trials ranged between 1022 and 1579.

To control each dimension of cursor movement (vertical or horizontal), one EEG channel over left sensorimotor cortex (i.e., electrode locations C3 or CP3) and/or one channel over right sensorimotor cortex (i.e., C4 or CP4) were derived from the digitized data according to a Laplacian transform (McFarland et al, 1997). Every 50-msec, the most recent 400-msec segment from each channel was analyzed by a 16th-order (McFarland and Wolpaw, 2008) autoregressive model using the Berg algorithm (Marple, 1987) to determine the amplitude (i.e., square root of power) in a 3-Hz-wide mu or beta frequency band, and the amplitudes of the one or two channels were used in a linear equation that specified a cursor movement in that dimension as described by Wolpaw and McFarland (2004). Thus, 20 times/sec the cursor moved in the vertical and horizontal dimensions simultaneously. Complete EEG and cursor movement data were stored for later offline analyses.

2.1.3. Offline Preprocessing and Feature Extraction—First, because the 64 electrodes are digitized sequentially (the value number n of the electrode number e is not

recorded at the time $t_0 + \frac{n}{\text{Sampling_rate}}$ but $t_0 + \frac{n}{\text{Sampling_rate}} + \frac{e-1}{64 \times \text{Sampling_rate}}$ where t_0

is the time when the recording start) and further processing compares values at different electrodes that are assumed to occur at the same time, the data were temporally aligned using a cubic spline interpolation. A cubic spline is a piecewise cubic polynomial which is twice continuous differentiable. The interpolation of a data set $\{x_i, y_i\}_{i=1}^n$ of n points such as $x_i < x_{i+1}$ with a cubic spline consists of finding a function F that satisfies:

- $$F(x) = \begin{cases} P_1(x), & x \in [x_1, x_2] \\ \dots & \\ P_{n-1}(x), & x \in [x_{n-1}, x_n] \end{cases}$$
 where $P_1 \dots P_{n-1}$ are cubic polynomial.
- $\forall 1 \leq i \leq n, F(x_i) = y_i$
- $\forall 1 \leq i < n, P'_i(x_{i+1}) = P'_{i+1}(x_{i+1}), P''_i(x_{i+1}) = P''_{i+1}(x_{i+1})$
- $P''_1(x_1) = P''_{n-1}(x_n) = 0$

Second, in accord with McFarland et al (1997), a large Laplacian spatial filter was applied to enhance the signal to noise ratio.

Third, a spectral analysis using a 16th-order autoregressive model determined the power in the 17 3-Hz wide spectral bins centered at frequencies from 8 to 24 Hz. The logarithms of these power values served as the features that were translated into cursor movements. As in online cursor control, the offline analysis used a 400-ms window that shifted in 50-ms steps. Thus, cursor movement was computed 20 times/sec. To explore new control possibilities more fully, these features were determined for 10 electrodes located over sensorimotor cortex (FC3, C5, C3, C1, CP3, FC2, C2, C4, C6 and CP4), not simply for the 2 electrodes (C3 and C4 for most users) that had actually been used for online control.

Thus, each 400-msec window of data was described by a vector of features, the dimension of which was the number of electrodes (10) times the number of frequency bins (17), for a total of 170 features. This differs from the protocol that had been used online, in which cursor movement in each dimension was a linear function of only 1–4 features.

2.2. Two approaches to comparing the different translation algorithms

2.2.1. Using the whole trial—The simplest approach to offline analysis is to determine from the full data of each trial the position of the target to which the user had to move the cursor. For this approach, the features used by the translation algorithm are the means of the feature values for all the 400-msec windows of each trial. The number of features used to describe each trial is then equal to the number of electrodes times the number of frequencies used (i.e., 170).

Although this determination of target position might be addressed by a classification algorithm, we used instead a set of regression algorithms for reasons detailed in McFarland and Wolpaw (2005). (Briefly, regression is preferable to classification because it is better suited to controlling continuous cursor movements in real-time and it more readily generalizes to novel target locations.)

The primary result of offline analysis, and the measure used to compare alternative regression algorithms, was the proportion of the variance in target position that was predicted by the algorithm (i.e., r^2).

2.2.2. Using the individual 400-msec windows—The whole-trial approach to offline analysis is not readily transferable to actual online use. During online cursor control, cursor movements are made every 50 ms based on the most recent 400-msec window. Thus, in

order to obtain offline results from methods that would be compatible with online use, we also compared the alternative translation algorithms using the features extracted from the individual 400-msec windows.

Another difference between online and whole-trial offline analysis is the user's ability to use visual feedback (i.e., cursor position) to make corrective movements. For example, as seen in Figure 1, prior cursor movement positioned the cursor so that to reach the target the user needed to move the cursor up, even though the target was lower than the initial starting point. To take account of such effects, the desired output for the regression should not be the absolute coordinates of the target, but rather the coordinates of the target relative to the cursor position.

To make our model realistic, we introduced a delay to take into account the time the user needs to perceive the feedback and correct his trajectory. This is why the relative coordinate of the target was computed as the coordinates of the target minus the coordinate of the cursor a short time (i.e., the delay) before the start of the sliding window.

Unfortunately, it is not possible to directly extract the position of the target from a single window of data since the quality of the EEG signal and the precision of the extracted features are not sufficient to locate the target with only 400ms of signal. On the other hand, a single window should have more information than just the direction of the target. That is, when the cursor is far from the target the user may try to move faster, so that each window may also contain information about the distance between the cursor and the target. This is why a sigmoid function (Figure 1C) is used to take into account the distance of the target while limiting extreme cursor movements.

This individual windows approach consists of deducting from each 400-msec window of the signal the relative direction/position of the target. The performances of the different regressions methods are then evaluate in term of person's r^2 .

2.2.3. Evaluating the performance of the translation algorithms—Estimating the performance of a method is an important issue in machine learning (Kohavi, 1995). The simplest technique, sometimes called Holdout or test sample estimation (Lachenbruch and Mickey, 1968), consists of partitioning the data (i.e. the vectors of features of all the trials when using the whole-trial method and of all the windows when using the single-window method) in two subsets. A training set that is used to train the method and a testing set that is used to evaluate its performances. To obtain unbiased results, the data used for testing must be different then the data used for training, otherwise, we would evaluate the capacity of the method to “remember” the data and not how it will be able to generalize what it learnt to a new sample. This technique reaches its limits when the total amount of data is limited, which is the case here. Indeed, if there are not enough training data, the algorithm will not be able to learn what best predicts the criterion variable and therefore will not generalize well to new examples. On the other hand, if there are not enough testing data, the results may not be reliable.

More complex techniques, like the bootstrap (Efron and Tibshirani, 1993) or the cross validation (Mosier, 1951; Stone, 1974; Kohavi, 1995), were developed to obtain more precise evaluations of the performances. The n -fold cross validation technique consists of splitting the data in n parts (instead of two for the Holdout) each of which will be used alternatively to train or to test the algorithm (Figure 2). The final performance is the mean of the performances on the n different testing sets.

In our case the data is already split in 7 parts (the seven last sessions recorded during seven different days). Using a 7-fold cross validation enable us to obtain results compatible with online experiment: The methods are trained on 6 sessions and are tested on a different session that was recorded during a different day. In this way we can test the robustness of the methods, meaning that the signal characteristics may change from the training and testing data sets due to slight changes of the electrodes position or of the different moods of the user.

Some of the methods have parameters that need to be adjusted (like the C of SVR or the width σ of the RBF kernel). To obtain unbiased results, the parameters are set by doing a 6-fold cross validation on each of the training sets of the 7-fold cross validation (Figure 2). To optimize the results, the features are normalized (for every trial, each feature is normalized according to the mean and standard deviation of its values over the training trials only, in this way it is compatible with online analysis).

In order to evaluate the use of doing Cross Validation a 3-way split analysis was also performed, using the first four sessions for training, session five for selecting the parameters and the last two sessions to evaluate the performances. The 3-way split analysis resulted in a lower performance on the test set of about 15% of r^2 .

2.2.4. Quantifying the differences between methods—During online experiments the performances of the users can be measure by the percentage of hit targets. Although this could be done for the off-line analysis, it would be biased for two reasons: First during online analysis the users correct their trajectories according to the feedback and those same corrections are used during offline analysis even if they are not appropriate. Secondly, during online experiment there is only one target present on the screen for each trial, thought the only way to fail a trial is if the user cannot reach this target within 20 seconds. (The online experiment was designed to teach the users how to move the cursor not to evaluate their performances. That is why only one target was present at the same time.)

Because we are evaluating regression techniques, a natural measurement of the performances is correlation of the output of the regression with the objective (i.e. the position of the target for the global approach or the relative direction of the target for each sliding window).

Unfortunately, it is not possible directly translate r^2 into performance accuracy. The only thing that can be said is that the closer it is to 1, the more the user would have been moving toward the target if he had used this method.

Another issue is that the results vary a lot across users and vary depending which series are used for training and testing, though the benefit from using a particular method for the regression can be hidden by the natural variation of the results. A solution is to measure the performance of each method as the ratio of its results (given as Person's r^2) over the results of a method of reference.

The objective is then to compute the increase (in percent) of the results by using a particular method compared to a method of reference, namely multiple linear regression, as well as a confidence interval of this increase.

First, for each of the n parts of the n -fold cross validation and each user, we compute the logarithm of the ratio of the result (in term of r^2) of the two methods. The mean and standard deviation is then calculated and used to compute the confidence interval of the logarithm of the ratio according to:

$$[\bar{r} - t_{\frac{1-\alpha}{2}}^{m-1} \frac{S}{\sqrt{m}}, \bar{r} + t_{\frac{1-\alpha}{2}}^{m-1} \frac{S}{\sqrt{m}}]$$

where $\bar{r}$ and S are the mean and standard deviation of the logarithm of the ratio of the results, m equal n times the number of user, $1-\alpha$ is the confidence level and t_{γ}^k the quantile of order γ of the student law with k degrees of freedom.

The increase can be computed by the following equation:

$$incr = 100 \times (e^{\bar{r}} - 1)$$

which is also applied to the boundaries of the previous interval to compute the final confidence interval. The use of the logarithm and then exponential is necessary because we are using ratios and not differences – for the mean it is equivalent to compute the geometric mean.

2.3. The different regression techniques

Three different regression techniques are compared. The first technique is “multiple linear regression”. It is closest to the technique that was actually used online, and thus is the one to which the other techniques are compared.

The second technique is the “LASSO regression” (Tibshirani, 1996), which adds a penalty to the multiple linear regression to make it generalize better. More precisely, LASSO regression consists of minimizing the square of the L_2 norm of the difference between the goal and output, and an L_1 penalty. It can be written as:

$$\min_{\beta} \|Y - X\beta\|_2^2 + \lambda \sum_{i=1}^p |\beta_i|$$

where $Y \in \mathbb{R}^n$ is the desired outputs, $X \in \mathbb{R}^{n \times p}$ the features matrix, and $\beta \in \mathbb{R}^p$ a parameter vector.

The third technique is the application of Support Vectors Machines (Vapnik, 1995 and 1998) to regression, also call SVR (Smola, 1996). It is used with two different kernels: a linear kernel and a radial basis function kernel:

$$k(x, y) = e^{-\frac{\|x-y\|^2}{2\sigma^2}}$$

2.4. Feature Selection

2.4.1. Critical issues—In one sense the higher the dimension of the feature space is (i.e., the more electrodes and frequencies used) the more information on each trial or each sliding window is available, and so the more accurate the cursor control can be. Unfortunately, not all features contain useful information and instead may add mostly noise. This is why using more features can appear to increase the information, but actually decreases the information to noise ratio. When too many features are used, every sample (trial or sliding window) of the training data set can be predicted from noise alone. The performances will then be close

to perfect on the training set and catastrophic on the testing one. The solution to this problem is to try to select the features that will add useful information without causing a decrease in global performance.

Two different approaches to features selection can be found in the machine learning literature: the so-called filter and wrapper methods (Blum and Langley, 1997). The filter methods consist of eliminating the irrelevant features (or selecting the useful ones) before the use of the learning algorithm. On the contrary the wrapper methods test the learning algorithm with different subsets of features to try to select the optimal subset. Wrapper approach can provide better feature selection (Kohavi and John, 1995), but generally imply a higher computational cost than filter methods due to the multiple trainings of the learning algorithm with the different feature subsets.

The time needed to train the SVR (especially with the RBF kernel when using large amount of features) being too long to use the wrapper method, we will focus on the filter approach and propose a fast and efficient way to select the features.

The filter methods can be subdivided in two categories. First, the methods that evaluate the use of each feature separately. The usefulness of a feature can be measured by its correlation with the desired output (here the optimal cursor movement) or by its mutual information, to take into account non-linearity (McFarland et al, 2006a). Secondly, the methods that evaluate in one comprehensive analysis the usefulness of all the features. For example, if the features are normalized, the coefficients of a regression that incorporates all the features as independent variables indicate the usefulness of each feature (e.g. the LASSO method is often used this way to select features (Tibshirani, 1996)).

The major drawback of most of the methods of the first category is that they do not take into account the redundancy of the information. For example, if two features contain very similar information, they will certainly be selected together, although only one may be useful, the other could just add more noise. This is particularly true in our case, where features extracted from neighbor electrodes and/or at close frequencies, will contain similar information and highly correlated noise. To deal with this issue it is possible to recursively select the features that add the most information to the group of features already selected. We can then order the features in a way that takes into account the redundancies of the information, and uses only those features that add useful information.

This is a greedy algorithm (i.e. it makes locally optimal choice at each stage with the hope of finding the global optimum) that is not necessarily optimal, since it is not certain the subset of features selected is the one that contains the most information. Nevertheless, determining the best subset would be extremely time-consuming, and this algorithm should provide a good approximation to the optimum solution.

2.4.2. The feature selection algorithm—Let $A \in \mathbb{R}^{m \times I}$ be the matrix representing the training data, where $I = n_{Elec} \times n_{frequencies}$ is the dimension of the features space (or the number of features) and the number of data points (in our case the total number of sliding windows for all the trials). $A_{i,j}$ is then the value of the j^{th} feature for the i^{th} sliding window.

Let $B \in \mathbb{R}^{m \times 2}$ be the matrix of the desired output. The two dimensions of the matrix represent the horizontal and vertical control command for every sliding window.

The algorithm is described in Table 1. The first step is to select the feature that contains the most information about the desired output. In our case, it is evaluated by the correlation of the feature with the output. The second step consists of removing from the other features the information that is contained in the feature that has just been selected. If we consider the

features as vectors in the Euclidian space $\mathbb{R}^m$, this step is exactly the orthogonal projection of the remaining features in the direction of the previously selected feature. The algorithm will continue by selecting the next feature according to the first step and so on. More precisely, the algorithm saves the result of the first evaluation of the usefulness of the features and selects the new features according to the new evaluation (meaning after the multiple projections) and the original evaluation. This guarantees that it selects only features that originally contained unique information.

Once the features have been ordered, the p first ones will be used for the regression (where p is a parameter that has been adjusted by cross validation).

2.4.3. Improving the RBF kernel—The feature selection algorithm does not take into account the fact that the first selected feature is generally more important than the feature number p . In the case of the SVR with the RBF kernel, we can define a norm that takes into account this difference in importance between the features by adding a weight to each component. If x is the j^{th} data point ($x = A_{i,j}$), x_i being the value of the i^{th} feature for the j^{th} data point ($x_i = A_{i,j}$), we define the norm of x by:

$$\|x\|_w^2 = \frac{1}{\sum_{i=1}^n w_i} \sum_{i=1}^n w_i x_i^2 \text{ with } w \in [0, 1]^n \text{ and } \forall i < j, w_i \geq w_j$$

The modified kernel is now:

$$k(x, y) = e^{-\frac{\|x-y\|_w^2}{2\sigma^2}}$$

The vector of weights is defined by:

$$w_i = \frac{\beta}{\beta + i^\alpha - 1}$$

where α and β are parameters that have to be adjusted. To decrease the time necessary for the training of the SVR, it is possible to set the weights w_i to 0 when $i > p$. This leads to use of only the p first features, but is a good approximation when w_p is small.

3. Results

3.1. Comparison of regression techniques when using the whole trial

Figure 3 shows that the linear techniques (multiple regression, LASSO regression and SVR with a linear kernel) give very similar results on the test data set. The test data set results are improved by the SVR technique with a radial basic function kernel, particularly when more data is used for training. The results on the test set for the SVR with the RBF kernel are much lower than the ones on the training set, reflecting most likely an over-fitting issue. This can be solved by a better features selection, which is done by modifying the RBF kernel (see 2.4.3 and results in Figure 6).

However, those results have to be interpreted with caution. First, in order to obtain results from training data sets of different size (in % of the total data) trials from the same session were split between the training and testing data set, and so, the result of Figure 3 may be biased. In the remainder of the paper all the result will be obtained with a training data set of

85% of the total data corresponding to a 7-fold cross validation (this way session will not be split between training and testing data sets).

Secondly, Figure 3 was obtained with the global trial approach. One could make the hypothesis that even if the translation of each sliding window to cursor movement is a linear process, the translation of the global trial into the position of the target could be non-linear, which would explain the better result of the SVR with RBF kernel.

Therefore, the next section will be dedicated to adjust and validate the model used for the sliding-window approach and section 3.3 will focus on a deeper comparison of SVR with a RBF kernel to multiple linear regression.

3.2. Using the multiple linear regression technique to compare the influence of different parameters when working with sliding windows

To do a realistic offline analysis using sliding windows, a model of the interaction between the user and the BCI had to be built. Before using this model to compare feature translation algorithms it has to be validated. This model is mostly based on the assumption that the users take into account the position of the cursor and aimed toward the target rather than just focusing on the target. One way of validating this assumption is to calculate the influence of using the cursor position (which is done in the next section) as well as the benefit of using the sigmoid function.

To be as close as possible to online analysis, the multiple linear regression technique was used, with only two electrodes and no features selection.

3.2.1. Influence of the delay used for the sliding windows method—A very long delay (e.g., Inf in Figure 4) represents the case in which the cursor's prior movement is not taken into account (i.e., the infinite delay means that the position stays at the starting point). For all the users, except two, taking into account the cursor position increases the correlation between the output and the target direction (Figure 4 and Table 2). This increase is maximal for a specific delay, which varies across the users. These results confirm our hypothesis that the users take into account the position of the cursor.

3.2.2. Using the sigmoid function to compute the relative position of the target—The effects on performance from using the sigmoid function are shown in Table 3.

3.2.3. Comparison of interpolation techniques—Figure 5 shows the influence of the interpolation method on the results of the multiple linear regression when working with sliding windows. A cubic spline interpolation improves the results by 15% ([9.7% 20.3%] $p < 0.05$) compared to no interpolation. This result underscores the importance of temporally aligning the signal from the different electrodes. The cubic spline interpolation also enhances the result by 3.6% ([0.9% 6.3%] $p < 0.05$) compared to a linear interpolation that was used during online analysis and by 2.5% ([0.1% 4.9%] $p < 0.05$) compared to a cubic interpolation.

These results may be explained by the fact that cubic splines are twice continuously differentiable, which makes them more adequate to model continuous phenomenon, such as oscillatory rhythms, than piecewise affine functions (used in linear interpolation) or piecewise cubic Hermite polynomials (which are just once continuously differentiable).

3.3. Comparing multiple linear regression and SVR with an RBF kernel, with and without features selection, when using sliding windows

For each user, the seven sessions were used for a 7-fold cross validation. The feature normalization and the adjustment of the parameters (number of features to use for the multiple regression; weights of the features, σ and C for the SVR) were done on the training data sets only, by another level of cross validation (Figure 2).

Early results, like the global trial approach, shown that SVR with a linear kernel was not able to enhance the performances compare to multiple linear regression. We then focus on the comparison of SVR with RBF kernel and multiple linear regression as well as the impact of selecting features.

Figure 6 shows the average increase, for all users, of SVR with RBF compare to using multiple linear regression with 10 electrodes. The relative increase of r^2 resulting from using SVR with the modified RBF kernel over the use of multiple regression when using 10 electrodes is 12.6% and is significant ([7.9%; 17.4%] $p < 0.05$). The relative increase of r^2 reaches 24.2% ([15.8%; 33.7%] $p < 0.05$) when comparing SVR with 10 electrodes to multiple regression with the two electrodes that are used during online analysis. The results for SVR with the standard RBF kernel are not significantly different from MLR with 10 electrodes, which underline the importance of an appropriate feature selection.

The important differences between the results on the training sets and the testing sets mean generalization to new examples is difficult and suggest, especially for the SVR, that the use of additional training data could improve the result.

4. Discussion

4.1. Importance of the user correction

The user's task was simply to move the cursor from the middle of the screen to one of the targets, and so could in theory be done without any feedback. Nevertheless, taking into account the cursor position during offline analysis improves performance by about 13%. This fact suggests that the users (or at least some of them) are using the feedback to control the cursor trajectory in real-time, meaning that they are able to correct the cursor movement and not just aim at the target during the whole trial. Moreover, the delay that represents the user reaction time, about 600 ms, is very similar to the reaction time when subjects perform a conventional motor task. These results are in general agreement with studies of manual reaction times (e.g., Albert et al, 2007) and our previous study of a different mu-rhythm task (Friedrich et al, 2009).

It is possible to look at this BCI as a way to select among eight possible targets and so to compute the bit rate transmission, which is often used as a measure of BCI performance. But this view does not take into account the main purpose of this BCI, namely, allowing the user to control the cursor in real-time. The results suggest that using this BCI is similar to using a standard pointing device such as a computer mouse.

One of the goals of this study was to develop a model of online experiments to enable realistic offline analysis using online recorded data. The benefits of taking into a count the cursor movement support our model. The next step to validate our model will be to confirm with an online analysis the prediction of our model (namely the advantage of SVR with the modified RBF kernel over multiple linear regression).

4.2. Feature Selection

As Figure 6 shows, when the number of electrodes used increase from 2 to 10 the results on the training data set were greatly improved (+83% for MLR) whereas the result on the testing data set just slightly increase (+10% for MLR). The use of the features selection algorithm prevents the over-fitting issue. This result in a decrease on the training set but also an increase on the testing set. It is particularly true for the SVR with RBF kernel: the modified RBF kernel, which use feature selection, decrease the result by 54% on the training data set and improves them by 9% on the testing data set compare to SVR with RBF kernel and no feature selection.

It is also important to note that the modified kernel produced better results than just using a standard RBF kernel on selected features. This justified ordering the features and giving them different weights in contrast to using uniform weighting of the subset of the features.

Another advantage of feature selection is that when fewer features are used the training and testing of the SVR is faster. For example, computing the results in Figure 6 took about twice as much time when using all the features then when selecting the features and using the modified RBF kernel with some weights w_i set to zero.

4.3. SVR with modified RBF kernel versus multiple regression

One important difference is that the training of the SVR, especially when using a nonlinear kernel, is 10 to 100 times slower than the training of the multiple linear regression. Fortunately this is not a real issue for the online analysis. Evaluating the output for each sliding window can be done in real-time, and training the SVR takes about a minute and so can be done between sessions.

It is important to note that this study consisted of offline analysis of previously recorded data. The users were trained on and used an online method (McFarland et al., 2006) that was very close to multiple regression used here. For this reason, the comparison of the multiple regression and the SVR with RBF kernel might be biased. Moreover, to save some computation time, some parameters, like the time delay for the use of the cursor positions, were optimized using the multiple regression only. As a result, the benefit of using a SVR with RBF kernel during an online study may be even more than the +12.6% compare to multiple linear regression reported here.

Another way to increase the accuracy of this BCI was explored in this study, namely increasing the number of features by using the electrodes surrounding the two electrodes used in the online study. This augmentation of the number of features potentially created a problem of over-fitting, in particular for the SVR with RBF kernel, but it was successfully solved by defining a new kernel using a RBF kernel and a features selection algorithm. The use of this modified RBF kernel with 10 electrodes resulted in an average improvement of +24.2% compared to the multiple regression using only two electrodes.

All the results of this study have been obtained with an offline analysis and so will have to be confirmed by online studies. It will also be interesting to investigate the impact of these new techniques on the evolution of the users' performance. It is as yet unknown whether the SVR will decrease the training time necessary to achieve multidimensional cursor control, or whether it will improve the accuracy of the control the user ultimately achieves.

Acknowledgments

This work was supported in part by grants from NIH (HD30146 (NCMRR, NICHD) and EB00856 (NIBIB & NINDS)) and the James S. McDonnell Foundation.

References

- Albert NB, Weigelt M, Hazeltine E, Ivry RB. Target selection during bimanual reaching to direct cues is unaffected by the perceptual similarity of targets. *Journal of Experimental Psychology*. 2007; 33:1107–1116. [PubMed: 17924810]
- Birbaumer N, Ghanayim N, Hinterberger T, Iversen I, Kotchoubey B, Kubler A, Perlmutter J, Taub E, Flor H. A spelling device for the paralyzed. *Nature*. 1999; 398:297–298. [PubMed: 10192330]
- Blankertz B, Müller K-R, Krusienski DJ, Schalk G, Wolpaw JR, Schlogl A, Pfurtscheller G, Millan JR, Schröder M, Birbaumer N. The BCI Competition III: Validating alternative approaches to actual BCI problems. *IEEE Transactions on Neural Systems and Rehabilitation Engineering*. 2006; 14:153–159. [PubMed: 16792282]
- Blum A, Langley P. Selection of relevant features and examples in machine learning. *Artificial Intelligence*. 1997; 97:245–271.
- Efron, B.; Tibshirani, R. An introduction to the bootstrap. Chapman & Hall; 1993.
- Farwell LA, Donchin E. Talking off the top of your head: toward a mental prosthesis utilizing event-related brain potentials. *Electroencephalogr Clin Neurophysiol*. 1988; 70:510–523. [PubMed: 2461285]
- Friedrich EVC, McFarland DJ, Neuper C, Vaughan TM, Brunner P, Wolpaw JR. A scanning protocol for a sensorimotor rhythm-based brain-computer interface. *Biol Psychol*. 2009; 80:169–175. [PubMed: 18786603]
- Kohavi R. A study of cross-validation and bootstrap for accuracy estimation and model selection. *Proceedings of the Fourteenth International Joint Conference on Artificial Intelligence*. 1995; 2:1137–1143.
- Kohavi R, John GH. Wrappers for feature subset selection. *Artificial Intelligence*. 1995; 97:273–324.
- Lachenbruch PA, Mickey MR. Estimation of Error Rates in Discriminant Analysis. *Technometrics*. 1968; 10:1–11.
- McFarland DJ, Lefkowitz T, Wolpaw JR. Design and operation of an EEG-based brain-computer interface (BCI) with digital signal processing technology. *Behav Res Meth Instrum and Comput*. 1997a; 29:337–345.
- McFarland DJ, McCane LM, David SV, Wolpaw JR. Spatial filter selection for EEG-based communication. *Electroencephalography and Clinical Neurophysiology*. 1997b; 103:386–394. [PubMed: 9305287]
- McFarland DJ, Wolpaw JR. Sensorimotor rhythm-based brain-computer interface (BCI): feature selection by regression improves performance. *IEEE Transactions on Neural Systems and Rehabilitation Engineering*. 2005; 14:372–379. [PubMed: 16200760]
- McFarland DJ, Anderson CW, Müller KR, Schlogl A, Krusienski DJ. BCI meeting 2005- Workshop on BCI signal processing: Feature extraction and translation. *IEEE Transactions on Neural Systems and Rehabilitation Engineering*. 2006a; 14:135–138. [PubMed: 16792278]
- McFarland DJ, Krusienski DJ, Wolpaw JR. Brain-computer interface signal processing at the Wadsworth Center: mu and sensorimotor beta rhythms. *Prog Brain Res*. 2006b; 159:411–419. [PubMed: 17071245]
- McFarland DJ, Wolpaw JR. Sensorimotor rhythm-based brain-computer interface (BCI): model order selection for autoregressive spectral analysis. *Journal of Neural Engineering*. 2008; 5:155–162. [PubMed: 18430974]
- Mosier CI. Symposium: The need and means of cross-validation. *Educational and Psychological Measurement*. 1951
- Müller KR, Anderson CW, Birch GE. Linear and non-linear methods for brain-computer interfaces. *IEEE Transactions on Neural Systems and Rehabilitation Engineering*. 2003; 11:165–169. [PubMed: 12899264]
- Pfurtscheller G, Flotzinger D, Kalcher J. Brain-computer interface - a new communication device for handicapped persons. *J Microcomp Applic*. 1993; 16:293–299.
- Schalk G, McFarland DJ, Hinterberger T, Birbaumer N, Wolpaw JR. BCI2000: A general-purpose brain-computer interface (BCI) system. *IEEE Trans Biomed Eng*. 2004; 51:1034–1043. [PubMed: 15188875]

- Smola, AJ. Master's thesis. Technische Universität München; 1996. Regression estimation with support vector learning machines.
- Stone M. Cross-Validatory Choice and Assessment of Statistical Predictions. *Journal of the Royal Statistical Society Series B (Methodological)*. 1974; 36:111–147.
- Tibshirani R. Regression shrinkage and selection via the Lasso. *Journal of the Royal Statistical Society B*. 1996; 58:267–288.
- Vapnik, V. The nature of statistical learning theory. Springer-Verlag; New York: 1995.
- Vapnik, V. Statistical learning theory. John Wiley; New York: 1998.
- Vapnik, V. The support vector method of function estimation. In: Suykens, JAK.; Vandewalle, J., editors. *Nonlinear Modeling: Advanced Black-Box Techniques*. Kluwer Academic Publishers; Boston: 1998.
- Wolpaw JR, McFarland DJ, Neat GW, Forneris CA. An EEG-based brain-computer interface for cursor control. *Electroencephalogr Clin Neurophysiol*. 1991; 78:252–259. [PubMed: 1707798]
- Wolpaw JR, Birbaumer N, McFarland DJ, Pfurtscheller G, Vaughan TM. Brain-computer interfaces for communication and control. *Clin Neurophysiol*. 2002; 113:767–791. [PubMed: 12048038]
- Wolpaw JR, McFarland DJ. Control of a two-dimensional movement signal by a noninvasive brain-computer interface in humans. *Proc Nat Acad Sci*. 2004; 101:17849–17854. [PubMed: 15585584]

Figure 1.

Left: EEG recorded from two electrodes during one trial. At zero time the target appears. One sec later the cursor appears and begins to move. One of the sliding windows is highlighted in red. *Center:* Hypothetical example showing the eight possible target locations, and the target for the current trial (in green). The green (gray) curve represents the cursor movement generated by the offline analysis up to (after) the time of the red window. The red sliding window must be used to move the red dot. *Right:* The sigmoid function. If the target is far from the cursor, the exact relative position of the target does not matter, just the direction is important. This is why, for example, the output for the relative coordinate of (0.5; 0.5) is very close to the one of (1; 1), meaning that in both cases the user should move fast to the right and to the top.

Figure 2.

Holdout: one data set is used for training (gray) and the other for testing (red). Cross Validation: the data are split into n parts. A total of n analyses are performed. Each part is used as the testing set (red) for one analysis while the $n-1$ other parts are used as the training set (gray). To adjust the parameters, a second level of cross validation is made: before each of the n analysis, $n-1$ analyses are made (training sets are in grey and testing sets in blue). To obtain unbiased result, the testing data set of the first level of cross validation (red) is not used during the second level of cross validation (white).

Figure 3. Results on the training (dash lines) and testing (full lines) data sets in term of the square of Pearson's correlation coefficient as a function of the size of the training data set for the four different methods.

Figure 4.

Results on the testing data set (as r^2) as a function of the delay. Each user is represented by a different color. The two users with spinal cord injuries are represented by dash lines. The right end of the curves is with an infinite delay (i.e., the previous cursor movement is not taken into account, so that the desired output of the regression is the absolute coordinate of the target and does not depend on the cursor position). (The multiple linear regression technique was used with the sliding windows approach, two electrodes and features selection.)

Figure 5.

Average results in term of r^2 for eight users (with the multiple linear regression technique, the sliding windows approach, two electrodes and features selection), when no interpolation is used (dark blue) and for four different interpolation methods. The error bars represent the interval of confidence at 95% of the increase of the results compare to no interpolation. Except for the nearest neighbor, all the interpolation techniques significantly increase the results.

Figure 6.

Average results (as r^2) for eight users, for multiple linear regression (M.L.R.) and SVR with RBF kernel, when working on sliding windows. Feature-unique weights for the RBF were calculated by the feature selection algorithm. The error bars represent the interval of confidence at 95% of the increase of the results compare to MLR with 10 electrodes and no features selection. The starts show methods significantly different from 10 electrodes and no feature selection.

Table 1

Feature ordering algorithm. $best_feature(A, B)$ is a function that gives the index of the column of A that contain the most information about B . $A_{i,j}$ represent the j^{th} column of A seen as a vector of $\mathbb{R}^m$, $\langle X|Y \rangle = X^T \cdot Y$ is the Euclidian scalar product, and $\|X\| = \sqrt{\langle X|X \rangle}$ is its associated norm.

Feature Ordering Algorithm	
(1)	<i>for</i> $i = 1 \dots n$
(2)	$ind = best_feature(A, B)$
(3)	$features_order(i) \leftarrow ind$
(4)	<i>for</i> $j = 1 \dots n$
(5)	$A_{-,j} \leftarrow A_{-,j} - \frac{\langle A_{-,j} A_{-,ind} \rangle}{\ A_{-,ind}\ ^2} A_{-,ind}$
	<i>end</i>
	<i>end</i>

Increase of the results on the testing data set (in percent of r^2) by taking into account the cursor position compared to using only the target position. The best delay is the delay for which the increase is maximum. The values in bold are significant ($p < 0.05$).

Table 2

Users	A	B	C	D	E	F	G	H	Mean
Increase (%)	29	22	13	-5	18	6	17	6	12.8
Confidence Interval at 95%	15; 46	7; 39	3; 24	-14; 6	8; 30	-18; 38	8; 26	-7; 21	7.7; 18.1
Best Delay (s)	0.75	0.5	0.75	2	0.3	0.2	0.5	0.2	0.65

Increase of the results on the testing data set (in percent of r^2) by using the sigmoid function to compute the relative position of the target. The values in bold are significant ($p<0.05$). (The multiple linear regression technique was used with the sliding windows approach, two electrodes and features selection.)

Table 3

User	A	B	C	D	E	F	G	H	Mean
Increase (%)	-1.53	8.22	1.5	-3.49	3.12	23.9	5.77	7.7	5.4
Confidence Interval at 95%	-3.8; 0.8	-0.4; 17	-3.3; 7.7	-2.2; 9.1	2.3; 4.6	-1.9; 56	2.7; 9.7	3.6; 11.9	4.1; 6.7