BCI Competition 2003: Data Set IIa – Spatial Patterns of Self-Controlled Brain Rhythm Modulations

Gilles Blanchard and Benjamin Blankertz

Abstract— **A brain-computer interface (BCI) is a system that should in its ultimate form translate a subject's intent into a technical control signal without resorting to the classical neuromuscular communication channels. By using that signal to, e.g., control a wheelchair or a neuroprosthesis, a BCI could become a valuable tool for paralyzed patients. One approach to implement a BCI is to let users learn to self-control the amplitude of some of their brain rhythms as extracted from multi-channel EEG. Here we present a method that estimates subject-specific spatial filters which allow for a robust extraction of the rhythm modulations. The effectiveness of the method was proved by achieving the minimum prediction error on data set IIa in the BCI Competition 2003, which consisted of data from three subjects recorded in 10 sessions.**

Index Terms— **brain-computer interface, single-trial classification, self-regulation of brain rhythms, common spatial patterns, feedback control**

I. Introduction

The goal of brain-computer interface (BCI) research is to provide humans with a new communication channel that allows to translate brain states via a computer into application specific actions. At the first international meeting for BCI technology it was agreed to reserve the term *BCI* for a system that does *not depend on the brain's normal output pathways of peripheral nerves and muscles*, [1]. Some researcher use this restriction to form the notion of an *independent BCI*. Such a system may become a valuable tool for paralyzed patients who may ultimately use it to control a wheelchair, a neuroprosthesis or a computer application.

One of the pioneering labs of BCI research is the Wadworth Center, NYS Department of Health, headed by Jonathan R. Wolpaw. They established a BCI system in which users learn to control the amplitude of their μ or β brain rhythms over sensorimotor cortices, [2]. A feedback signal is calculated from multi-channel EEG and used to control the movement of a cursor on a computer screen which is observed by the subject. This closed feedback loop allows subjects to adapt their strategies in order to improve the recognition rates of the BCI system.

A comprehensive data set of this type recorded from three subjects was given in the *BCI Competition 2003*, cf. [3], [4]. Here we present our method to classify this data set, which attained the minimal classification error on the competition's test set. The labels of the test set were unknown during analysis and could thus not be used for tuning the model. The main difference between our approach and the original one is the choice of spatial filters. While in the original approach spatial filters are selected manually from a small repertoire of very general filters (common average reference, small or large laplacians) we determine specific spatial filters in a data-driven manner by a common spatial pattern (CSP) analysis. While this technique is not new from an algorithmic view point, the way it is applied here is tailored to the requirements and novel in (1) using CSPs for different frequency bands at the same time and (2) extracting CSPs only with respect to one class.

GB and BB are with Fraunhofer FIRST (IDA), Berlin, Germany, E-mail: gilles.blanchard@first.fraunhofer.de. GB is supported by the CNRS (France). The work of BB was partly supported by a grant of the BMBF, FKZ 01IBB02A (Germany).

II. The data set

A. Description of the data

Since this data set is described in detail in [4], see this issue, we give only a summary here. EEG was recorded from 64 scalp electrodes at a sampling rate of 160 Hz during an experiment of the following trial design. After a resting period of 1 s during which the screen stays blank, a target appears on the right side in one of four possible positions. Another second later a cursor appears in the middle of the left side and travels at a constant speed to the right. The vertical movement of the cursor is determined by a linear combination of the subject's amplitude power in a μ and/or β frequency band for 1 to 3 spatially filtered channels. When the cursor reaches the right edge, the height of the cursor defines the result target of the trial. The screen is cleared and the next trial begins. The aim of the subject is to steer the cursor such that the result target coincides with the target indicated at the beginning of the trial. The choice of spatial filters, feedback electrodes, and frequency bands is made subject-specific according to preliminary experiments and stays fixed then. In contrast, certain parameters like slope and intercept of the linear equation that controls the cursor movement, are updated online after each trial. Spatial filters are chosen from a fixed repertoire like common average reference, small or large laplacians, cf. [5].

B. The classification goal and its challenges

The data available to the competitors consisted of the full recordings for ten 30-minute sessions from three subjects (*A*, *B*, *C*). Each session contained 192 trials. The training set consisted of all trials of the first 6 sessions which were labeled with the target position code (*top, up, down* or *bottom*) and the result position code (corresponding to the position actually reached by the cursor at the end of the 2 s feedback period), giving 1152 labeled trials in all for each subject (except for subject *C*, for which one of the sessions was corrupt). The goal for the competitors was to classify the 768 unlabeled trials for sessions 7 to 10, for each of the three subjects; the class to predict was the unknown target position code. The classifier had to be causal in the sense that it had to make a prediction using only information coming from the present trial and earlier trials. In the end, the classifier we used was actually static (i.e., only depended on the present trial and labeled training trials), so that this requirement was of course satisfied.

The difficulty of this data set came from the fact that the experiment included a feedback part (not known to the competitors in extensive detail) mediated by the movement of the cursor on the screen. While the original feedback procedure tried to "learn" from the subjects (since some parameters of the feedback procedure were updated after each trial), it could be expected that conversely, the subjects were probably trying to "learn" how to use the feedback procedure to have a better control over the cursor's movement. Somehow the classification task implied to learn in an indirect way this reciprocal feedback loop. The paradox lies in the fact that the ultimately envisioned goal was to improve the feedback procedure; but of course with a new feedback procedure the subjects' reactions would also probably be different.

C. *Our approach to the classification task*

From a high level perspective the task is a classification problem with 4 classes (target positions). But regarding the feedback design on a more basic level, control is accomplished via a one-dimensional control variable that is defined by two opposing brain states, one that makes the cursor go up and the other makes the cursor go down. Accordingly the basic task is to extract the subject's intent with respect to those brain states. So we came to the conclusion that it would be more reliable to first take into account only the *top* and *bottom* classes. Indeed, we hoped that for these two classes the subjects would try to make the cursor always go up (or down) whatever the actual movement of the cursor on the screen, in other words that the feedback loop would be less influential. By contrast, we expected that to reach the two intermediate target positions, the feedback would be more influential since the subject would constantly try to "correct" the position and movement of the cursor (note that no information about the actual position of the cursor on the screen during the trial was made available to the competitors).

Induced by the original feedback method the two controlling brain states are defined by the power of the μ/β rhythms over sensorimotor cortices (strong power makes the cursor rise in subjects *A* and *C* while it makes the cursor go down in subject *B*). The next question to attack was how to extract features that robustly reflect the subject's intent. By robust we mean that that the features should as little as possible be influenced by task-unrelated signal components, such as physiological or measurement artifacts or other ongoing brain activity. Since

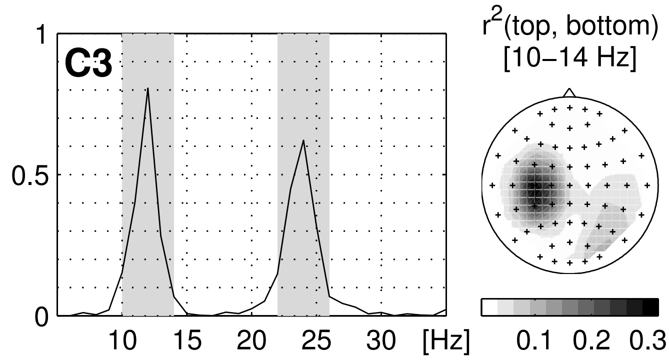

Fig. 1. Left: The r^2 values across the spectrum for a single particular channel (C3, common average reference) for discrimination of the classes *top, bottom* (subject *C*). Intervals marked in gray indicate the frequency intervals that have been chosen for this subject. Right: Scalp topography of the r^2-values for subject *C* in the frequency band 22–26 Hz.

the brain rhythms which are used to control the feedback have (subject-) specific spatial distributions, see Fig. 1, our choice was to use a Common Spatial Pattern (CSP) analysis to determine spatial filters for each subject as described below.

III. Feature extraction and classification

A. *Preliminary analysis*

The goal of the preliminary analysis was to find good discriminative power bands for each subject. For each channel separately, we plotted the power spectra for the *top* and *bottom* classes, then the individual channels r^2-values (i.e., the proportion of the variance of the spectral power values accounted for by the label information [6]) across the spectrum corresponding to the discrimination of these two classes, see Fig. 1. This allowed us to select the more discriminative power bands (corresponding to peaks of r^2-values). These were slightly different for each subject; curiously enough, while we kept 2 discriminative power bands for each subject corresponding to the μ and β rhythms, the frequency bands corresponded to noticeably higher frequencies that what had been suggested by the providers of the data (for each subject the typical frequency bands we kept were roughly 10–15 Hz and 23–28 Hz while the providers reported to have used 8–13 Hz or 18–24 Hz [4]). No trials were rejected due to artifacts.

B. *Feature extraction by CSPs*

In the end, for each subject we kept two band-passed signals (using a linear causal filter) for each channel. The next step was to find appropriate spatial filters to extract features that robustly reflect the subject's intention concerning the cursor movement. To this end we used a Common Spatial Patterns analysis. CSP is a technique known from statistical pattern recognition [7] and was suggested by [8] for spatial analyses of EEG signals, and more specifically by [9] to find spatial structures of event-related (de-)synchronization (ERD/ERS) in a BCI context. The CSP analysis for a two-class problem consists in finding linear subspaces, i.e., linear combinations of channels, for which the variance of the signal for one class

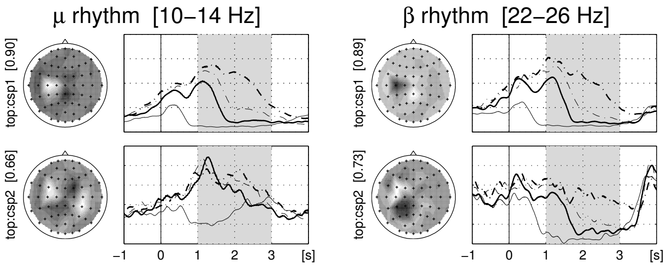

Fig. 2. Spatial repartition of CSP coefficients on the scalp: shown are the 2 most significant CSPs of the μ (left side) and the β (right side) rhythm for class *top* and subject *C*. On the left of each scalp is given the associated eigenvalue. On the right of each scalp map the respective ERD curves are shown, split into the four target classes (*bottom* (–), *down* (**–**), *up* (- -), *top* (**- -**)). The x–axis is time. The grayed interval [1 s, 3 s] is the period during which the cursor is visible on the screen. This interval was fed into the CSP algorithm.

is maximized while the variance of the other is minimized (we then talk of CSPs associated to class x when the variance of class x is maximized). More precisely, let Σ_i be the covariance matrix of the trial-concatenated matrix of dimension [channels $\times$ concatenated time-points] belonging to the respective label $i \in \{1, 2\}$. The CSP analysis consists in calculating a matrix R and diagonal matrix D with elements in $[0, 1]$ such that

$$R\Sigma_1 R^T = D \qquad \text{and} \qquad R\Sigma_2 R^T = 1 - D.$$

which can easily be done by whitening and spectral theory. The projection that is given by the ith row of matrix R has a relative variance of d_i (ith element of D) for trials of class 1 and relative variance $1 - d_i$ for trials of class 2. Typically one would retain some projections corresponding to the highest eigen-values d_i, i.e., CSPs for class 1, and some corresponding to the lowest eigenvalues, i.e., CSPs for class 2. For an extension of the CSP algorithm to multi-class problems, see [10], [11].

We performed this analysis for the segment of recorded signals corresponding to the time period [1 s, 3 s] where the cursor is visible on the screen. (The point t= 0 s for each trial corresponds to the time at which the target appears on the screen.) Here again, only the classes *top, bottom* were considered to extract CSPs. It turned out that, for all of the subjects, the CSPs associated to only one of the two classes {*top, bottom*} were significant (the significance of a CSP is measured by its associated eigenvalue), but not the same class for everyone: for subjects *A* and *C*, the *top* class produced the most significant eigenvalues; for subject *B*, it was the *bottom* class, see Sec. II-C. For each subject, we decided to keep the first two CSPs for the significant class in each frequency band, thus keeping in total 4 CSPs for each subject, see Fig. 2.

C. Final classifier

For investigation purposes, we also drew the ERD/ERS curves of the CSP channels, see Fig. 2. They were calculated by squaring the band-pass filtered CSP channels, smoothing them in time, and averaging over training trials for the different target classes. This is a standard procedure as described, e.g., in [12]. The observation that the ERD/ERS curves had specific temporal evolutions for each target motivated to take the following features for classification: the channels in the time period [1 s,4 s] were projected onto the 4 CSPs. Fourier power coefficients were calculated for these four signals (for the frequency bands corresponding to the bands from which the CSPs were respectively extracted), in successive windows of size 1 s with an overlap of 0.5 s. The feature vector of one trail was the concatenation of the power values of all 5 subwindows and all four channels.

The final classifier was a regularized linear discriminant trained over the full set of training examples (i.e., using all four classes this time), with the regularization parameter determined by cross-validation. Note that this method classifies trials as a whole and is not suitable for a continuous feedback. This issue is addressed in the next section.

D. Continuous feedback classifier

One legitimate question is to know whether the procedure and features we used could be translated into a "feedback" algorithm which would allow to continuously control the cursor movement on the screen and to use the final position of the cursor for classification.

We tried to build such an algorithm based on the same principles we used for the classifier described above. The difference with the previous classifier was that the new classifier should be able to discriminate the intention of moving up from the intention of moving down from short time windows in order to give continuous feedback. For this reason we divided the period [0.75 s, 3 s] into 8 subintervals I_k of width 0.5 s with an overlap of 0.25 s. For each of these subintervals the same features were used as before, i.e., Fourier power coefficients for CSP signals in the relevant frequency bands.

To do this we just collected all the subintervals of the classes *top* and *bottom* as separate training examples (each labeled with the corresponding target class) and trained a Fisher linear discriminant f on these.

Then we used the output of this classifier (a real number) to represent the movement of the cursor on the screen. To avoid extreme values this output was thresholded (by threshold values -3 and 3 which were determined empirically in order to minimize the cross-validation error on the training set); then the position of the "virtual cursor" on the screen at the end of a trial is just the sum of these thresholded classifier values: $p_{\text{end}} = \sum_k t(f(I_k))$ where t is the threshold function $t(x) = 3$ for $x > 3$, $t(x) = -3$ for $x < -3$ and $t(x) = x$ elsewhere. Finally, to test this procedure we had to translate back this final position into the initial classification task; for this we chose three cut values $c_1 < c_2 < c_3$ such that, if $p_{\text{end}} < c_1$, then the class *bottom* is predicted; if $c_1 \leq p_{\text{end}} < c_2$ the class *down* is predicted, etc.

We emphasize that for this method

- To build the classifier f, only the information of the training classes *top* and *bottom* is used; the other training classes are only used to determine the best cut values c_i for each subject.
- The temporal information available in the global classifier is lost, since all subintervals are treated equally.

- The cut values c_i are different for each subject, but could be normalized to arbitrary values if we used instead an affine rescaling of the classifier output for each subject.

These handicaps notwithstanding, the overall test misclassification error only raised by 4 %. This shows that some information was lost in the process, but that the continuous feedback procedure is nevertheless practically viable – although ultimately such a feedback algorithm would have to be tested with a new set of experiments on the subjects.

E. What we did expect

Of course we hope that our spatial filters are capable of capturing the relevant dynamics of the subject's brain state more robustly and allows a better discrimination of the two controlling brain states. On the other hand this does not imply that the classification accuracy on the given data set will be better compared to the original feedback. The strategy of the subject was directly coupled to the original feedback. It had to include counterbalancing possible shortcomings of the original feedback. But this counterbalancing is counterproductive in the off-line analysis of those feedback sessions with different algorithms that have other (or less) shortcomings. This means that an algorithm which is better in extracting the subject's intent from the EEG compared to the original one possibly performs not better or even worse on the test set.

Another question is how the process of learning strategies to produce suitable brain states would be affected when using our approach. The main issue here is that the spatial filters we are using are more specific. The consequences could be twofold: It could either be that the learning process of the subject is impaired because s/he is confined to the specific patterns as extracted from the initial training period. On the other hand it could be helpful that the spatial filters are tailored to the subject's brain signals, releasing the subject from the need of adapting to a predefined feedback control. To handle the anticipated possible problem, the CSPs should be recalculated quite regularly from the most actual data during the initial sessions. Even using the conventional method in the very beginning might be preferable. The performance in real environment of our algorithm will stay an open question until it is implemented in a BCI system and several feedback experiments have been conducted and evaluated.

F. Results on the test dataset

Table I gives the test classification results of our classification methods. The trial-wise classifier, cf. Sec. III-C, achieved the best results among the contestants of the BCI Competition 2003 for this data set, see [13], [4]. While the results averaged over all three subjects gives a performance slightly worse than the original online prediction algorithm, it is interesting to note that this is actually mainly due to Subject *B*, where our method compares poorly to the original one. By contrast, for Subject *C* we observe a clear improvement over the original method. (For Subject *A*, results are almost the same, actually slightly better for our method). This shows, if anything, that the two methods exhibit important differences: the next step would be to understand why a method is better on a given subject than another, possibly to combine their respective capabilities for increased performance. One important issue here seems to be the non-stationarity of the EEG. In consequence, some parameters of our method should probably be updated online, as was the case for the original prediction method. For example, in the analysis of the results of the continuous classifier we noticed that the cut values c_i estimated on the training set were not optimal on the test set. An additional online update for these values would probably lead to a significant increase in performance. The failure of our method in subject *B* might also be caused by a shift of the spatial distribution of the μ rhythm which could distort our CSP channels.

TABLE I

CLASSIFICATION RESULTS FOR THREE SUBJECTS FROM CLOSED-LOOP BCI EXPERIMENTS. THE METHODS WERE TUNED AND TRAINED ON A TRAINING DATA SET AND EVALUATED ON A FIXED TEST DATA SET.

Subject	Original online feedback	CSP-based classifier (trial-wise classifier)	CSP-based classifier (continuous classifier)
A	26.6 %	25.5 %	29.6 %
B	22.8 %	34.0 %	35.0 %
C	31.0 %	25.3 %	29.0 %
mean	26.9 %	28.2 %	31.2 %

IV. CONCLUSION

We presented our approach to the classification analysis of data set IIa from the BCI Competition 2003 ([3], [4]). Since multi-channel EEG was available and brain states are discriminated by rhythmic features, a common spatial pattern analysis seemed promising. In order to make such an approach work for this data set some thoughts were necessary. Although being primarily a four class problem the discrimination of two opposed brain states (cursor up vs. cursor down) is a sufficient basic ingredient. In contrast to the original CSP application in BCI context (left vs. right hand imagery) here the role of the classes was asymmetric (i.e. maximizing the variance for one class while minimizing variance for the other class resulted in significant CSPs, but not *vice-versa* – and the order was subject-dependent). To catch the intent of the user more robustly CSPs we calculated separately for the μ and for the β rhythm. For competition purpose a method was chosen that used the CSP filtered signals to classify trials was a whole. Additionally we presented an approach that allows to calculate a continuous feedback signal, while the classification of a trial depends only on the position of the cursor when it arrives at the right side of the screen. While the first approach reaches somewhat better results in terms of classification error, the latter approach seems to be more appropriate for implementation in a BCI feedback system. A line for improvement is to incorporate strategies for a continuous adaptation of the feedback algorithm to account for the non-stationary characteristics of the EEG. While this is straight forward for the thresholds between the classes, the development of an adaptive version of the CSP algorithm is subject of present research in our team.

ACKNOWLEDGMENT

The authors would like to thank Guido Dornhege for helpful discussions and two anonymous reviewers for their suggestions which helped to improve the quality of this paper.

REFERENCES

[1] J. R. Wolpaw, N. Birbaumer, W. J. Heetderks, D. J. McFarland, P. H. Peckham, G. Schalk, E. Donchin, L. A. Quatrano, C. J. Robinson, and T. M. Vaughan, "Brain-computer interface technology: A review of the first international meeting," *IEEE Trans. Rehab. Eng.*, vol. 8, no. 2, pp. 164–173, 2000.

[2] J. R. Wolpaw and D. J. McFarland, "Multichannel EEG-based brain-computer communication," *Electroencephalogr. Clin. Neurophysiol.*, vol. 90, pp. 444–449, 1994.

[3] B. Blankertz, "BCI competition 2003 (web page)." [Online]. Available: http://ida.first.fhg.de/projects/bci/competition/

[4] B. Blankertz, K.-R. Müller, G. Curio, T. M. Vaughan, G. Schalk, J. R. Wolpaw, A. Schlögl, C. Neuper, G. Pfurtscheller, T. Hinterberger, M. Schröder, and N. Birbaumer, "The BCI competition 2003: Progress and perspectives in detection and discrimination of EEG single trials," *IEEE Trans. Biomed. Eng.*, 2004, to appear.

[5] D. J. McFarland, L. M. McCane, S. V. David, and J. R. Wolpaw, "Spatial filter selection for EEG-based communication," *Electroencephalogr. Clin. Neurophysiol.*, vol. 103, pp. 386–394, 1997.

[6] H. Sheikh, D. J. McFarland, W. A. Sarnacki, and J. R. Wolpaw, "Electroencephalographic(EEG)-based communication: EEG control versus system performance in humans," *Neurosci. Lett.*, vol. 345, no. 2, pp. 89–92, 2003.

[7] K. Fukunaga, *Introduction to statistical pattern recognition*, 2nd ed. Boston: Academic Press, 1990.

[8] Z. J. Koles and A. C. K. Soong, "EEG source localization: implementing the spatio-temporal decomposition approach," *Electroencephalogr. Clin. Neurophysiol.*, vol. 107, pp. 343–352, 1998.

[9] H. Ramoser, J. Müller-Gerking, and G. Pfurtscheller, "Optimal spatial filtering of single trial EEG during imagined hand movement," *IEEE Trans. Rehab. Eng.*, vol. 8, no. 4, pp. 441–446, 2000.

[10] G. Dornhege, B. Blankertz, G. Curio, and K.-R. Müller, "Boosting bit rates in non-invasive EEG single-trial classifications by feature combination and multi-class paradigms," *IEEE Trans. Biomed. Eng.*, 2004, in revision.

[11] ——, "Increase information transfer rates in BCI by CSP extension to multi-class," in *Advances in Neural Inf. Proc. Systems (NIPS 03)*, vol. 16, 2004, to appear.

[12] G. Pfurtscheller and F. H. L. da Silva, "Event-related EEG/MEG synchronization and desynchronization: basic principles," *Clin. Neurophysiol.*, vol. 110, no. 11, pp. 1842–1857, Nov 1999.

[13] B. Blankertz, "BCI competition 2003 results (web page)." [Online]. Available: http://ida.first.fhg.de/projects/bci/competition/results/