

Towards a robust BCI: Error potentials and online learning

Anna Buttfeld, Pierre W. Ferrez and José del R. Millán

Abstract

Recent advances in the field of Brain-Computer Interfaces (BCIs) have shown that BCIs have the potential to provide a powerful new channel of communication, completely independent of muscular and nervous systems. However, while there have been successful laboratory demonstrations, there are still issues that need to be addressed before BCIs can be used by non-experts outside the laboratory. At IDIAP we have been investigating several areas that we believe will allow us to improve the robustness, flexibility and reliability of BCIs. One area is recognition of cognitive error states, that is, identifying errors through the brain's reaction to mistakes. The production of these error potentials (ErrP) in reaction to an error made by the user is well established. We have extended this work by identifying a similar but distinct ErrP that is generated in response to an error made by the interface, (a misinterpretation of a command that the user has given). This ErrP can be satisfactorily identified in single trials and can be demonstrated to improve the theoretical performance of a BCI. A second area of research is online adaptation of the classifier. BCI signals change over time, both between sessions and within a single session, due to a number of factors. This means that a classifier trained on data from a previous session will probably not be optimal for a new session. In this paper we present preliminary results from our investigations into supervised online learning that can be applied in the initial training phase. We also discuss the future direction of this research, including the combination of these two currently separate issues to create a potentially very powerful BCI.

Index Terms

Brain-computer interface, Cognitive error state recognition, Online learning, Adaptive classifiers

I. INTRODUCTION

The goal of a brain-computer interface (BCI) is to create a direct channel of communication between a user's brain and a computer, completely bypassing the traditional muscle-dependent communication channels. This is still a very young field of research and it encompasses many different approaches to a variety of problems.

The BCI system that we have developed at the IDIAP Research Institute has shown good results in distinguishing between up to three mental states, including imagination of left and right hand movements, 3D visualisation and

This work is supported by the European IST Programme FET Project FP6-003758 and by the Swiss National Science Foundation NCCR "IM2". This paper only reflects the authors' views and funding agencies are not liable for any use that may be made of the information contained herein.

Authors are with the IDIAP Research Institute, CH-1920 Martigny, Switzerland (e-mail: {anna.buttfeld, pierre.ferrez, jose.millan}@idiap.ch).

language tasks [1]. We use a Gaussian mixture classifier to distinguish between the given tasks, which allows us to define an “unknown” output when the probability of none of the defined classes is above a defined confidence threshold, effectively creating an “idle” state without having to model it directly. While this technology has been demonstrated in the laboratory, it is not yet ready to be taken out of the laboratory and used in real-world situations. Our work at the moment is focussing on different ways of improving the robustness of BCIs with the goal of making BCIs a more practical and reliable technology.

One avenue of research that we are investigating is the use of high-level cognitive and emotional information in BCIs. In particular we are looking at error potentials (ErrP), the neural correlates to error awareness. The appearance of error potentials in response to errors made by the user are well established. Our research has identified a new type of error potential that is generated in response to errors made by the BCI rather than the user. In addition to identifying errors and stopping the BCI from executing incorrect commands, this new type of error potential may provide us with performance feedback that could allow us to improve the performance of the classifier while it is being used.

Online learning is another issue that we are currently investigating. EEG signals naturally change over time, both between different sessions and within a single session. Online learning can be used to adapt the classifier throughout its use and keep it tuned to drifts in the signals it is receiving in each session. Our hypothesis is that online learning used during initial training will reduce training time by facilitating mutual adaptation between the user and the BCI, allowing the user to refine his or her mental strategy through rapid feedback. During ongoing use we believe we can improve the performance of the classifier by constantly tuning it with alternative learning paradigms that do not require explicit knowledge of the target class. However, dynamic adaptation of the BCI is risky, since we may confuse the user by rapid and unpredictable changes, or degrade the performance of the classifier through inappropriate adaptation.

In this paper we will be describing the current state of the IDIAP BCI, including the hardware and statistical classifier. We will also present our current research in the two areas outlined above, error potentials and online learning.

II. THE IDIAP BCI

We have recently shown that after a few days of training, subjects are able to control a miniature robot in an indoor environment with several rooms and corridors using mental commands derived from an EEG-based BCI [1]. Key aspects that make it possible are the use of an asynchronous BCI and the combination of the user’s high-level commands with advanced robotics that implement those commands efficiently. We are working to improve this initial demonstrator, in collaboration with several European institutions, along four lines. The first is the development of a more powerful adaptive shared autonomy framework for the cooperation of the human user and the robot in achieving the target. The second line is the use of a technique recently developed by Grave et al. [2] that estimates the local field potentials (LFP) in the whole human brain from scalp EEG. Recent results show significant improvements in the classification of bimanual motor tasks using estimated LFP with respect to scalp EEG [3]. The third and fourth

research avenues seek to improve the robustness of a BCI and will be discussed in this paper.

A. Hardware and signal acquisition

EEG potentials are acquired with a portable BioSemi system using a cap with either 32 or 64 integrated electrodes arranged in the modified 10/20 International System [4]. The EEG recordings are monopolar and taken at 512Hz. The common average reference (CAR) procedure is used to suppress the background brain activity, where at each time step the average potential over all the channels is subtracted from each channel. This re-referencing procedure removes the background activity, leaving activity from local sources beneath the electrodes.

B. Statistical classifier

This is a short summary of the classifier we use in the IDIAP BCI. For more details, see [1]. We use a Gaussian classifier to separate the signal into the different classes of mental task. Each class is represented by a number of Gaussian prototypes, typically less than four. That is, we assume that the class-conditional probability function of class C_k is a superposition of N_k Gaussian prototypes. We also assume that all classes have equal prior probability. All classes have the same number of prototypes N_p , and for each class each prototype has equal weight $1/N_p$. Thus, the activity a_k^i of the i^{th} prototype of class C_k for a given sample \mathbf{x} is the value of the Gaussian with centre $\boldsymbol{\mu}_k^i$ and covariance matrix $\boldsymbol{\Sigma}_k^i$. From this we calculate the posterior probability y_k of the class C_k . The posterior probability y_k of the class C_k is now the sum of the activities of all the prototypes of class k divided by the sum of the activities of all the prototypes of all the classes. The input vector \mathbf{x} can be composed of either temporal or frequency features from a selection of electrodes, depending on the experiment.

The classifier output for input vector \mathbf{x} is now the class with the highest probability, provided that the probability is above a given threshold, otherwise the result is “unknown”. This rejection criteria gives the BCI the flexibility to not make a decision at any point without explicitly modelling an idle state.

Usually each prototype of each class would have an individual covariance matrix $\boldsymbol{\Sigma}_k^i$, but to reduce the number of parameters the model has a single diagonal covariance matrix common to all the prototypes of the same class.

During offline training of the classifier, the prototype centres are initialised by a clustering algorithm, generally self-organising maps [5]. This initial estimate is then improved by stochastic gradient descent to minimise the mean square error $E = \frac{1}{2} \sum_k (y_k - t_k)^2$, where \mathbf{t} is the target vector in the form 1-of-C; that is, if the second of three classes was the desired output, the target vector is (0,1,0). The covariance matrices are computed individually then averaged over the prototypes of each class to give $\boldsymbol{\Sigma}_k$.

III. ERROR POTENTIALS

EEG-based BCIs are prone to errors in the recognition of the user’s intent. However, EEG signals provide us with a tool to help us overcome this problem. In addition to the command signals generated intentionally by the user, EEG signals carry high level information about the cognitive states of the user. These states include attention,

fatigue and motivation as well as the one we are investigating — error recognition, where the user “reacts” to the occurrence of an error.

Brain signals associated with the recognition of an error made by the user are well established [6], [7], [8]. “Response ErrP” occur when the user makes an error and recognizes it immediately [6], [7]. This situation can be provoked by asking a user to perform a task that requires fast reactions so that the user makes some mistakes. A few studies have addressed the recognition of this kind of ErrP in single trials as a potential tool to improve the performance of a BCI [9], [10]. A second type of error potentials has also been identified. Known as “Feedback ErrP”, they occur when a user makes an error but is unaware of it until he or she is informed by feedback [8]. The brain response in this case is similar to that of response ErrP, but is a reaction to the feedback rather than the incorrect action. In both cases the timing of the brain reaction is tied to the point where the user realises his or her mistake, be that at the time of action or on receipt of negative feedback.

An important aspect of the described ErrP (response ErrP and feedback ErrP) is that they always follow an error made by the user. First the user makes a selection, and then the ErrP arise either after the occurrence of an error or after a feedback indicating the error. In the context of a BCI, the central question is whether ErrP are also elicited when the interface makes an error in recognising the user’s intent. Schalk *et al.* [11] have previously observed that such an ErrP appears in a BCI when the cursor reaches wrong targets, an operation that requires a number of consecutive mental commands. Here we are interested in the recognition of such a kind of ErrP in single trials — i.e., after the user delivers each mental command.

A. Experimental Setup

In a previous study, we presented experimental results showing the presence of ErrP in response to an error made by the interface rather than the user [12]. In the three subjects studied we were able to identify these “Interaction ErrP” on a single-trial basis, which is crucial for their integration into a real-time BCI. In these subjects were able to identify to presence of the ErrP in the error trials with an average accuracy of 79.9% and the absence of this ErrP in correct trials with an average accuracy of 82.4%. The results of this study are presented and discussed further here.

The experiment simulated a real interaction with a robot, where the user gives repetitive commands to bring the robot to the left or right side of a room. Feedback is delivered by two progress bars showing the number of times each command has been delivered. In order to separate the complex usage of the BCI from the reaction to the error, the task was presented in a simplified manual form. Instead of using a BCI to deliver the commands, the user presses a key to issue the command. By removing the BCI classifier from this part of the experiment and making the user task trivial, we ensure that the user performance is perfect and the only mistakes are the deliberately introduced mistakes made by the interface. So, the user repeatedly presses a key to issue his or her command (without making mistakes, due to the triviality of the task), but 20% of the time the interface makes a mistake and provides the wrong feedback. Three healthy male subjects between the ages of 24 and 42 participated in this experiment. The EEG signals were recorded with a 32 electrode cap, but only channels Cz and Fz were

used for classification since this combination produced the best results. This selection of electrodes makes sense because the distribution of ErrP is known to be fronto-central along the midline. The signals were filtered with a 1-10 Hz bandpass filter, since error potentials are known to be relatively slow cortical potentials. The classifier then considers the 0.5 second window starting between 150ms and 650ms after the feedback. The input to the classifier is a 128 element vector of the temporal features (0.5 seconds at 128Hz for Cz and Fz).

B. Experimental Results

Examining the average difference error-minus-correct (Figure 1) of our experimental results reveals what seems to be a new kind of ErrP, different from “Response” and “Feedback” ErrP, which for convenience we call “Interaction ErrP”. A first sharp negative peak occurs 270 ms after the feedback and is followed by a positive peak (between 350 and 450 ms after the feedback) and by a later negative peak (~ 550 after the feedback). Since our protocol is quite similar to an oddball paradigm the question arises of whether the potentials we describe are simply oddball N200 and P300. To clarify this issue we ran a second experiment with an error rate of 50%, where error trials are no less frequent than correct trials. The potentials observed in this case have the same timing as those observed in the case of 20% (Figure 2). This means that while we cannot exclude the possibility that N200 and P300 contribute when the error rate is 20%, the oddball N200 and P300 are not sufficient to explain the reported potentials.




Fig. 1. Average EEG for the difference error-minus-correct for channel Cz for the three subjects plus the grand average of them for 20% error. Feedback is delivered at time 0 seconds. The negative and positive peaks show up about 270 ms and between 350 and 450 ms after the feedback, respectively. A later negative peak shows up about 550 ms after the feedback.

For ErrP to be useful in a real-time BCI, we need to be able to identify error potentials in single trials, not just in grand averages. Table I gives the classification rates of error and correct single trials in a 10-fold cross-validation study.

These classification rates can be demonstrated to improve the theoretical performance of a BCI in terms of bit rate. The bit rate is a measure of performance that represents the amount of information communicated per unit




Fig. 2. Grand average EEG of two subjects showing the difference error-minus-correct for channel Cz when the error rate is 50%. Feedback is delivered at time 0 seconds.

TABLE I

CLASSIFICATION RATES OF EACH SUBJECT ON ERROR AND CORRECT SINGLE TRIALS.

	Subject 1	Subject 2	Subject 3	Average
Error trials [%]	87.3±11.3	74.4±12.4	78.1±14.8	79.9±6.6
Correct trials [%]	82.8±7.2	75.3±10.0	89.2±4.9	82.4±7.0

time. It is generally expressed in bits per trial. Traditional bit rate measures can be adapted to model the integration of error potentials into a BCI [12]. The most general method of integrating the error potentials is to stop the execution of a command when an error potential is detected. Alternatively, in the specific case of a 2-class BCI we can replace a command with the other command when we detect an error potential.

Using these bit rate equations we can calculate the theoretical performance of a BCI, assuming that the accuracy of the BCI is 80% and the subject-specific ErrP recognition rates are as reported in Table I. The performance of the BCI integrating ErrP is compared with that of standard 2-class and 3-class BCIs, which have performances of 0.28 and 0.66 bits per trial respectively. Table II shows that we can achieve a significant increase in the bit rate by stopping the commands identified as erroneous, with an average gain over the standard BCI of 72% in the case of a 2-class BCI and 28% in the case of a 3-class BCI. Surprisingly, the average gain from replacing the erroneous commands in a 2-class BCI is much lower, 14%, and the performance of subject #2 actually decreases under this paradigm.

IV. ONLINE LEARNING

For a BCI to be an effective tool it must have the ability to adapt dynamically throughout its use. The EEG signals received by the BCI will always change over time, both within a single session and between sessions, due

TABLE II
PERFORMANCES OF THE BCI INTEGRATING ERRP FOR THE 3 SUBJECTS AND THE AVERAGE OF THEM.

		Subjects			
		#1	#2	#3	Average
$N_c = 3$	BpT	0.91	0.73	0.92	0.85
stop	Increase	37.0%	9.5%	38.0%	27.6%
$N_c = 2$	BpT	0.53	0.40	0.52	0.48
stop	Increase	90.9%	41.9%	85.5%	71.7%
$N_c = 2$	BpT	0.36	0.19	0.44	0.32
replace	Increase	28.9%	-31.5%	58.9%	14.2%

to a number of factors. These factors include change in strategy by the user, user fatigue, and small differences in electrode position. So while at the beginning of each session the classifier will be initialised with respect to previous sessions, it must adapt itself to the particular signals it is receiving in the current session. However, this adaptation must be undertaken with great care — rapid and unpredictable changes in the classifier will confuse the user, and incorrect adaptation will degrade the performance of the classifier.

Work so far has focused on online learning during the training phase, where the user is told which command to give and so the real target is always known, making it a supervised learning task. The goal is to reduce the time a user takes to learn to use the system efficiently by facilitating mutual adaptation between the user and the BCI. This means that the user can refine his or her mental strategies while receiving rapid feedback from the BCI. Once the user has been trained and can produce the BCI commands reliably, he or she is ready to move to unstructured, unsupervised tasks such as controlling a robot. In this situation the techniques described below could be used in short recalibration sessions.

Previous preliminary work on this task [13] evaluated the advantages of using continued online learning with the basic gradient descent algorithm on the Gaussian classifier. Since then we have been investigating extensions of the basic gradient descent algorithm such as stochastic meta descent [14], which accelerates training by adapting individual learning rates for each parameter of the classifier.

A. Stochastic Meta Descent

Stochastic Meta Descent (SMD) [14] is an extension of gradient descent that uses adaptive learning rates to accelerate learning. The SMD algorithm is applied to each parameter in the classifier separately (the centre and covariance of each Gaussian prototype), and each parameter maintains and adapts an individual learning rate. This is in contrast to basic gradient descent, which uses a single learning rate for all parameters. The details of our implementation of SMD can be found in [15]. This approach has the potential for faster adaptation, but requires more computation. It also has more parameters that need to be tuned — basic gradient descent has two parameters, learning rates for the centres and covariances of the prototypes, while SMD has four parameters, initialisation values for the learning rates as well as an additional hyper-parameter.

B. Experimental results

Initial offline experiments have been performed on data recorded while giving no feedback to the subjects and no online learning was used during recording. We tested these algorithms on a three-class problem (imagination of left and right hand movements, and vocabulary search). Each class was performed for one second before switching to a different class in random order. Data is from three subjects, each with four sessions of almost four minutes collected with a break of ten minutes between sessions. Samples are passed to the classifier 16 times per second and the output from eight samples is averaged to give a decision every 0.5 seconds. The features vector used for classification was constructed by estimating the power spectral density over the previous second in the frequency range 8-30Hz at 2Hz resolution for the 8 centro-parietal channels (C3, Cz, C4, CP1, CP2, P3, Pz, and P4), giving a 96 element feature vector.

In this experiment we measured how well the classifier tracked the changing signals by applying online learning through all the sessions and measuring the classification performance. For each subject, the classifier was first trained offline on the data from the first session. The resulting classifier was then applied to sessions two through four, with the final adapted classifier from the end of the previous session used as the initial classifier for the next session. In these sessions we are continually adapting the classifier — each new sample comes in, is classified with the existing classifier, and the classifier is updated based on this sample. We compared basic gradient descent and the SMD algorithm against the static classifier with no adaptation, where the classifier trained on the first session is used to classify the following sessions with no further modification. Preliminary tests of basic gradient descent over a range of learning parameter values showed that the optimal parameters vary between subjects and sessions. We selected learning parameter values (the same for all sessions of all subjects) for gradient descent. This same value also serves as the initialisation value for SMD. In these experiments both gradient descent and SMD outperform the static classifier. Figure 3 shows the performance over time of the different classifiers for the three sessions of the second. It can be seen that the classification rates of the adaptive algorithms are statistically significantly better than the static classifier, with SMD better than gradient descent, especially towards the end of each session. Similar trends are observed with the other subjects.

In addition to applying the online learning algorithms throughout the session, we wanted to see whether there is a performance gain in applying the online learning algorithms for the first half of the session only, then applying the resultant classifier to the second half of the session with no further learning. This is similar to a recalibration scenario, where we want to use supervised learning for only part of the session. Results from this experiment show that there is a small improvement when using the classifier trained on the first half of the data over the classifier with no further training, but there is no statistically significant difference between the performance of basic gradient descent and SMD.

Table III shows the classification results for the three subjects averaged over their three sessions, and the overall average. The results are given as the percentage improvement of the online learning algorithms over the static classifier. There are three parts: the improvement in the first half of the data, the improvement in the second half

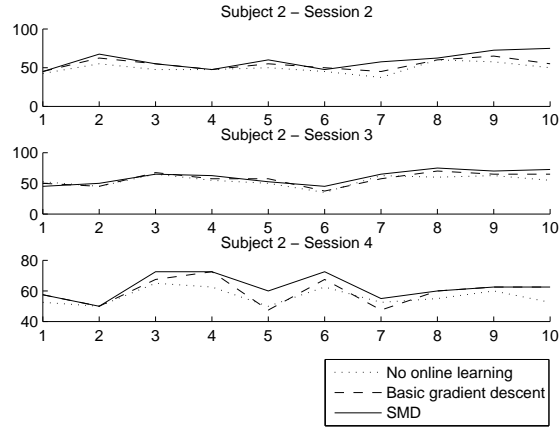


Fig. 3. Percentage of trials correctly classified in 20 second bins for the static classifier, basic gradient descent and SMD. Both of the adaptive algorithms have significantly better performance than the static classifier, with SMD performing better than basic gradient descent, especially towards the end of each session.

of the data when the online learning is continued, and the improvement on the second half of the data when the classifier is only trained on the first half of the data. Although there is variation between the subjects, the figures hold to the general trend as discussed above.

TABLE III

AVERAGE PERCENTAGE IMPROVEMENT OF ONLINE LEARNING, SMD AND BASIC GRADIENT DESCENT (GD) OVER STATIC CLASSIFIER

		First half	Second half	
			Adapting	Not adapting
Subject 1	GD	3.5	7.9	3.9
	SMD	4.9	9.3	5.9
Subject 2	GD	6.1	7.8	3.1
	SMD	9.3	18.6	5.2
Subject 3	GD	18.3	18.7	7.1
	SMD	28.0	33.0	2.5
Average	GD	9.3	11.5	4.7
	SMD	14.0	20.3	4.5

V. CONCLUSION AND FUTURE WORK

The two issues discussed in this paper, error potentials and online learning, have shown promising results in their early stages of investigation. Work needs to be done to integrate these principles into the BCI system. Operating individually these research areas can improve the performance of the BCI. Combined, they open up a new area of

research — continued reinforcement learning during general usage of the BCI using error potentials as feedback on the performance of the BCI. This means that for every decision the BCI makes we do not know what the correct decision should have been, but we are informed by the presence of the error potential when a wrong decision has been made, and we can use this information to adapt the classifier.

Another area of research is the use of estimates of the local field potentials (LFP) to improve the classifier [2] [3]. This technique allows us to generate a three dimensional map of the activity of from the EEG signals. Through this technique we can also gain a better understanding of the nature of the brain activity driving the BCI.

REFERENCES

- [1] J. del R. Millán, F. Renkens, J. Mouriño, and W. Gerstner, “Non-invasive brain-actuated control of a mobile robot by human EEG,” *IEEE Trans. Biomedical Engineering*, vol. 51, pp. 1026–1033, 2004.
- [2] R. Grave de Peralta, M. Murray, C. Michel, R. Martuzzi, and S. Gonzalez Andino, “Electrical neuroimaging based on biophysical constraints,” *NeuroImage*, vol. 21, pp. 527–539, 2004.
- [3] R. Grave de Peralta, S. Gonzalez Andino, L. Perez, P. Ferrez, and J. del R. Millán, “Non-invasive estimation of local field potentials for neuroprosthesis control,” *Cognitive Processing*, vol. 6, pp. 59–64, 2005.
- [4] F. Sharbrough, G. Chatrian, R. Lesser, H. Lders, M. Nuwer, and W. Picton, “Am. electroencephalogr. society guidelines for standard electrode position nomenclature,” *Clin. Neurophysiol.*, vol. 8, pp. 200–202, 1991.
- [5] T. Kohonen, *Self-Organising Maps, 2nd ed.* Berlin: Springer-Verlag, 1997.
- [6] W. Gehring, M. Coles, D. Meyer, and E. Donchin, “The error-related negativity: An event-related brain potential accompanying errors,” *Psychophysiology*, vol. 27, 1990.
- [7] M. Falkenstein, J. Hoormann, S. Christ, and J. Hohnsbein, “ERP components on reaction errors and their functional significance: A tutorial,” *Biological Psychology*, vol. 51, pp. 87–107, 2000.
- [8] C. Holroyd and M. Coles, “The neural basis of human error processing: Reinforcement learning, dopamine, and the error-related negativity,” *Psychological Review*, vol. 109, pp. 679–709, 2002.
- [9] B. Blankertz, G. Dornhege, C. Schäfer, R. Krepki, J. Kohlmorgen, K.-R. Müller, V. Kunzmann, F. Losch, and G. Curio, “Boosting bit rates and error detection for the classification of fast-paced motor commands based on single-trial EEG analysis,” *IEEE Trans. Neural Systems and Rehabilitation Engineering*, vol. 11, pp. 127–131, 2003.
- [10] L. Parra, C. Spence, A. Gerson, and P. Sajda, “Response error correction—A demonstration of improved human-machine performance using real-time EEG monitoring,” *IEEE Trans. Neural Systems and Rehabilitation Engineering*, vol. 11, pp. 173–177, 2003.
- [11] G. Schalk, J. Wolpaw, D. McFarland, and G. Pfurtscheller, “EEG-based communication: Presence of an error potential,” *Clinical Neurophysiology*, vol. 111, pp. 2138–2144, 2000.
- [12] P. Ferrez and J. del R. Millán, “You are wrong!—Automatic detection of interaction errors from brain waves,” in *Proc. 19th Int. Joint Conf. Artificial Intelligence*, 2005.
- [13] J. del R. Millán, “On the need for on-line learning in brain-computer interfaces,” in *Proc. Int. Joint Conf. Neural Networks*, 2004.
- [14] N. Schraudolph, “Local gain adaptation in stochastic gradient descent,” in *Proc. 9th Int. Conf. Artificial Neural Networks*, 1999.
- [15] A. Buttfeld and J. del R. Millán, “Online classifier adaptation in brain computer interfaces,” IDIAP, IDIAP Research Report 06-16, 2006. [Online]. Available: <http://www.idiap.ch/publications.php>