Neural Networks 24 (2011) 1120–1127

Contents lists available at SciVerse ScienceDirect

Neural Networks

journal homepage: www.elsevier.com/locate/neunet

On the use of interaction error potentials for adaptive brain computer interfaces

A. Llera [a,b,*], M.A.J. van Gerven [a,b,c], V. Gómez [a,b], O. Jensen [a,b], H.J. Kappen [a,b]

[a] Radboud University Nijmegen, The Netherlands

[b] Donders Institute for Brain, Cognition and Behaviour, Nijmegen, The Netherlands

[c] Institute for Computing and Information Sciences (ICIS), Nijmegen, The Netherlands

ARTICLE INFO

Article history:
Received 19 July 2010
Received in revised form 28 January 2011
Accepted 22 May 2011

Keywords:
Brain computer interfaces
Error potential
Adaptive classification

ABSTRACT

We propose an adaptive classification method for the Brain Computer Interfaces (BCI) which uses Interaction Error Potentials (IErrPs) as a reinforcement signal and adapts the classifier parameters when an error is detected. We analyze the quality of the proposed approach in relation to the misclassification of the IErrPs. In addition we compare static versus adaptive classification performance using artificial and MEG data. We show that the proposed adaptive framework significantly improves the static classification methods.

1. Introduction

The interest in Brain Computer Interfaces has quickly grown in the last few years. The possibility to provide disabled people with new communication channels, such as BCI spellers, new mobility channels such as BCI driven wheel chairs or BCI controlled mechanical prostheses makes this a very attractive research field. However, the applicability of current BCI systems is still limited because of a number of problems. One of these problems is the presence of non-stationarities in the data (Shenoy, Krauledat, Blankertz, Rao, & Müller, 2006). This causes patterns associated with each task during the training of the BCI to be different during testing, leading to a poor performance.

Several approaches have been proposed to overcome this problem by the introduction of adaptive classification methods (Sykacek, Roberts, & Stokes, 2004; Shenoy et al., 2006; Pfurtscheller, Neuper, Schölgl, & Lugger, 1998). In Shenoy et al. (2006), it is shown how the probability distributions associated with class features change between training and test sessions, and assuming that the labels of new incoming trials are known, it is shown that proper updates in the classifier parameters would improve the performance of the original static classifier. Note that in the BCI setting we normally do not know the user intention, so the labels of the trials are unknown. We propose the use of neural feedback to detect incorrect performance of the device, and to be able to recover the labels in the case of a binary classification task. The on-line detection of the wrong performance of a BCI has been addressed before by means of Interaction Error Potentials (IErrP) (Ferrez & Millán, 2008, 2005; Seno, Matteucci, & Mainardi, 2010).

Error-related potentials are potentials detected in the recorded electroencephalogram (EEG) of a subject just after an error occurs. The error is the difference between the expected and the actual result of an action. Error-related potentials have been studied in many different scenarios since the late 1980s (Falkenstein, Hohnsbein, Hoormann, & Blanke, 1990; Miltner, Braun, & Coles, 1997; van Schie, Mars, Coles, & Bekkering, 2004; Ferrez, 2007). It is well known that the presence of an error is usually followed by what are called event-related negativity and positivity which are present in the alpha band in the fronto-central channels. More recently, a study using Magnetoencephalography (MEG) (Mazaheri, Nieuwenhuis, van Dijk, & Jensen, 2009) has shown that an erroneous reaction to stimuli is followed by an increase in the frontal theta and a decrease in the posterior alpha and central beta powers.

Based on the nature of the feedback, the error-related potentials can be categorized as response error potentials (Falkenstein et al., 1990; Mazaheri et al., 2009), feedback error potentials (Miltner et al., 1997), observation error potentials (van Schie et al., 2004) and the most interesting for us, interaction error potentials (IErrP) which are present when a device delivers an erroneous feedback (Ferrez, 2007).

Since the IErrP are present in the recorded EEG of a subject controlling a device just after the device returns an unexpected feedback (the BCI makes a classification error) (Ferrez & Millán, 2008), its detection can help to construct a more robust BCI, either by correcting the BCI output directly (Ferrez, 2007), or more interestingly, by adapting the BCI classifier to prevent similar mistakes in the future. This idea is illustrated in Fig. 1.

* Corresponding author at: Radboud University Nijmegen, The Netherlands.
E-mail address: a.llera@donders.ru.nl (A. Llera).

0893-6080/$ – see front matter © 2011 Elsevier Ltd. All rights reserved.
doi:10.1016/j.neunet.2011.05.006

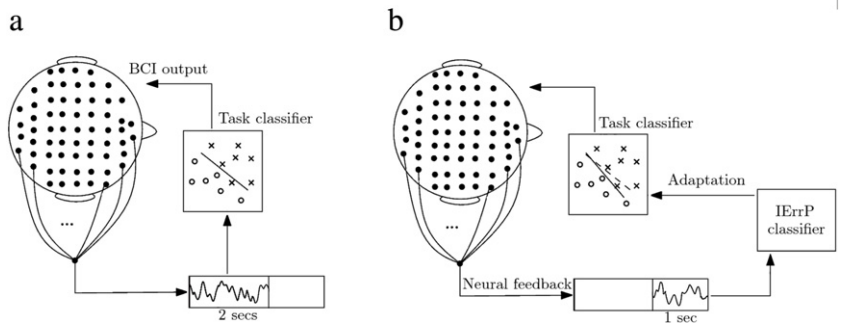

Fig. 1. (a) Classical scenario: brain activity is measured during a period of time, then the *task classifier* decides the class label given the measured activity and the device produces an output. (b) Proposed scenario: after task classification, an *IErrP classifier* uses feedback from the user (subsequent data from zero to one seconds after (a)) to compare previous user intention with the device output. If an IErrP is detected, the parameters of the task classifier are updated.

Although the possibility of BCI adaptation using error feedback from the user has been previously proposed (Chavarriaga, Ferrez, & Millán, 2007), the impact of using an IErrP classifier to improve the original task classifier has not been studied before in a realistic BCI setting. In this article, we explore this idea in detail. In Section 2 we introduce a framework based on reinforcement learning for adaptive BCI using an IErrP classifier as a control signal. We analyze the effect of IErrP misclassification in terms of false positives and false negatives, and we measure how the performance of the task classifier is affected. In Section 3, we first perform the single trial classification of IErrP and then we apply the proposed adaptive method to MEG data collected during a binary forced-choice covert attention task.

2. Adaptive BCI classifier

In this section we introduce the proposed method to design a binary adaptive BCI. We consider a binary task with an adaptive task classifier that learns from the output of a static IErrP classifier in order to minimize erroneous feedback.

2.1. Adaptive learning rule

Consider the (unobserved) subject's intention, left or right, that we denote as target class $t \in \{0, 1\}$ respectively. The generated brain activity is measured and a vector of feature values $\mathbf{x} := (x_1, \ldots, x_n)$ is extracted which is relevant to discriminate between both classes. We use the logistic regression model (Bishop, 2007) which takes the form:

$$p(t = 1|\mathbf{x}, \mathbf{w}) = \sigma(\mathbf{x}, \mathbf{w}) = \frac{1}{1 + e^{-\sum_{i=0}^{n} w_i x_i}}, \tag{1}$$

where $\mathbf{w} \in \mathbb{R}^{n+1}$ is the vector of weights, and $x_0 = 1$ accounts for the bias term.
The error in the prediction is quantified as the log-likelihood of the data:

$$G(\mathbf{x}, \mathbf{w}, t) = -(t \ln \sigma(\mathbf{x}, \mathbf{w}) + (1 - t) \ln(1 - \sigma(\mathbf{x}, \mathbf{w}))). \tag{2}$$

The output of the task classifier is defined as

$$\tilde{t} = \chi\left(p(t = 1|\mathbf{x}, \mathbf{w}) > \frac{1}{2}\right), \tag{3}$$

where χ returns 1 if its argument is true and 0 otherwise. An adaptive learning rule for the parameters $\mathbf{w}$ updates $\mathbf{w}$ in the direction of the gradient of (2):

$$\Delta w_i = \eta \frac{\partial G(\mathbf{x}, \mathbf{w}, t)}{\partial w_i} = \eta(t - \sigma(\mathbf{x}, \mathbf{w}))x_i, \tag{4}$$

where η denotes the learning rate.

In a realistic BCI setting however, the intention of the user t is unknown. We define $E \in \{0, 1\}$ as the user's true absence or presence of surprise following the output of the device. Thus $E = 0$ corresponds to $\tilde{t} = t$ and $E = 1$ to $\tilde{t} \neq t$. After the output of the task classifier ($\tilde{t}$) is delivered, subsequent brain activity (neural feedback) is measured and a feature vector $\mathbf{y} := (y_1, \ldots, y_m)$ is extracted and used by the IErrP classifier to provide an estimation of E, which we denote by $\tilde{E} \in \{0, 1\}$. Updates occur only when a surprise is detected ($\tilde{E} = 1$), in which case the observed output $\tilde{t}$ is presumably incorrect, so $t = 1 - \tilde{t}$ and the learning rule (4) for the task classifier becomes

$$\Delta w_i = \eta\tilde{E}(1 - \tilde{t} - \sigma(\mathbf{x}, \mathbf{w}))x_i, \tag{5}$$

where $1 - \tilde{t}$ is the opposite label from the output of the task classifier.

The performance of this model clearly depends on the flexibility of the model to adapt to changes at the correct time scale (Heskes & Kappen, 1991, 1992), but also on the asymptotic behavior of the task classifier in relation to the misclassification of IErrP. In Sections 2.2 and 2.3 we study this relation.

2.2. Effect of IErrP misclassification

The performance of a BCI system based on the previous framework clearly depends on the accuracy of the IErrP classifier. Previous researchers have reported classification rates of IErrP of around 80%, as well as the stability on IErrP detection across sessions (Ferrez & Millán, 2008). The misclassification of IErrPs can occur in two ways (see Fig. 2):

False positives. Correctly classified trials ($\tilde{t} = t$) are considered to be erroneous, causing an update of the task classifier parameters with the wrong class label. We characterize the rate of false positives with α_1.

False negatives. Erroneously classified trials ($\tilde{t} \neq t$) are considered as correct. As a consequence, the task classifier parameters will not be updated when it is desirable. We characterize the rate of false negatives with α_2.

Note that the effect of false positives results in learning from incorrectly labeled data, whereas false negatives result in discarding potentially useful learning samples.

2.3. Simulations

In order to better understand the asymptotic behavior of the task classifier in relation to the accuracy of the IErrP classifier (α_1 and α_2), we consider an artificial binary class classification problem in a one-dimensional feature space. For each class $t \in \{0, 1\}$, the feature is distributed according to a Normal distribution

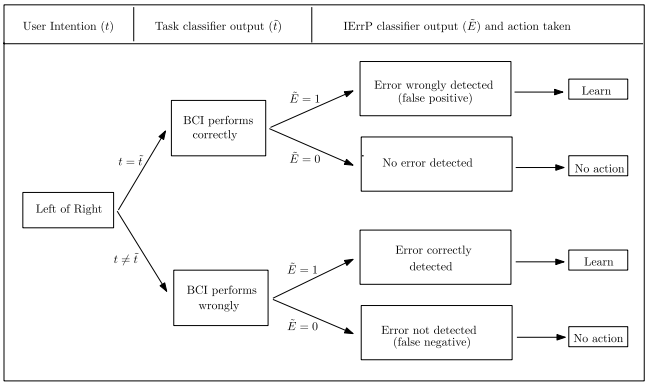

Fig. 2. The first column represents the target (t) intended by the user; the second column compares the user intention with the output of the task classifier ($\tilde{t}$); the third column shows the possible outputs of the IErrP classifier ($\tilde{E}$) as well as the effect produced in the task classifier due to the proposed method. *Learn* means that the parameters of the task classifier are updated using Eq. (5).

parametrized as $p(\mathbf{x}|t) = \mathcal{N}(\mu_t, \sigma^2)$. We choose $\sigma = 1$. The distance between μ_0 and μ_1 determines the overlap between the distributions and different levels of overlap result in classification problems with different optimal classification rates. We consider three different distances $\rho = |\mu_0 - \mu_1| = \{1, 2, 4\}$. For $\rho = 1$, the classification task is difficult and the optimal Bayes classifier accuracy is at most 0.70 whereas for $\rho = 4$ the task is easy and classification rate reaches 0.99 (see Fig. 3).

We start with a random weight vector $\mathbf{w}_0$. For each trial i, $i \in \{1, \ldots, M\}$, we will assume that $p(t^i = 0) = p(t^i = 1) = \frac{1}{2}$ and we draw a sample from $p(\mathbf{x}|t^i)$. This sample represents a feature extracted from the measured brain activity generated by the user while having intention t^i. The output of the BCI device $\tilde{t}^i$ is obtained using Eq. (3) and is then compared with t^i:

- If $t^i = \tilde{t}^i$, we draw $\tilde{E} = 1$ with probability α_1 and apply (5). In this case learning occurs with the wrong label. The classifier is not adapted with probability $1 - \alpha_1$.
- If $t^i \neq \tilde{t}^i$, we draw $\tilde{E} = 1$ with probability $1 - \alpha_2$ and we apply (5). In this case the task classifier is correctly adapted. With probability α_2 the classifier is not updated.

We observe that for any condition (with $0 \leq \alpha_1 < 0.5$ and $0 \leq \alpha_2 < 0.5$), the adaptive classifier improves the initial random boundary and converges after a short initial transient which we neglect for performance evaluation. In the experiments we set $M = 10^5$ and we consider the last 20% of the trials for performance evaluation, but the results do not critically depend on this choice. Fig. 3 shows the classification error of the task classifier as a function of α_1 and α_2 for different values of ρ. For $\alpha_1 = \alpha_2 = 0$, the method converges to the Bayes classification error for all ρ.

For difficult problems such as $\rho = 1$, the baseline classification rates are poor. In these cases, the surface remains constant for large values of α_1 and α_2 suggesting that the effects of misclassification in the IErrP classifier are not severe. On the contrary, for $\rho = 4$ optimal classification is very accurate. The adaptive task classifier is then more sensitive in this scenario and small increases in α_1 decrease significantly the performance of the classifier. We can see also that the performance of the method does not depend so much on α_2. Therefore, as expected, learning from wrong samples is more costly than discarding good samples. This is an important fact, since the relation between false positives and false negatives can be incorporated in the design of the IErrP classifier, in such a way that the effect of IErrP misclassification is minimized.

Note that the result assumes that the misclassification rate of the IErrP classifier is independent of whether the user intention is left or right. More exactly, defining α_i^L and α_i^R for $i \in \{1, 2\}$ as the rate of false positives/negatives associated with left and right user intention respectively, we have that

$$\alpha_i = \frac{\alpha_i^L + \alpha_i^R}{2} \tag{6}$$

and we assumed that for $i \in \{1, 2\}$ $\alpha_i^L \approx \alpha_i^R$, i.e. false positives/negatives are balanced for both right and left. Violation of this assumption introduces additional error, which we have studied in Appendix. We concluded that this effect is in general small.

3. Results on MEG data

In this section we show the applicability of the proposed method using MEG data. In Section 3.1 we describe the BCI experiment used to gather the data that will be used in the rest of the section. In Section 3.2 we construct an IErrP classifier (trials with an unexpected/expected device output), and in Section 3.3 we perform a comparison between a static classifier and the proposed adaptive method.

3.1. Data acquisition

Eight healthy subjects were instructed to direct a forced-task binary BCI device using the covert attention paradigm. Covert spatial attention is a well known paradigm for BCI control, based on the lateralization of the power in the alpha-band in posterior channels (van Gerven & Jensen, 2009; van Gerven, Bahramisharif, Heskes, & Jensen, 2009a; van Gerven, Hesse, Jensen, & Heskes, 2009b). The design has been carefully chosen to avoid any lateralization due to movement or stimulus presentation.
The description of the experiment is as follows (see Fig. 4 for more details).

Two squares and a fixation cross appear on the screen. After 300 ms, the fixation cross turns into an arrow (pointing left or right). The subject is instructed to direct his/her attention to the direction indicated by the arrow while maintaining fixation at the center of the screen. After 2000 ms, the arrow disappears

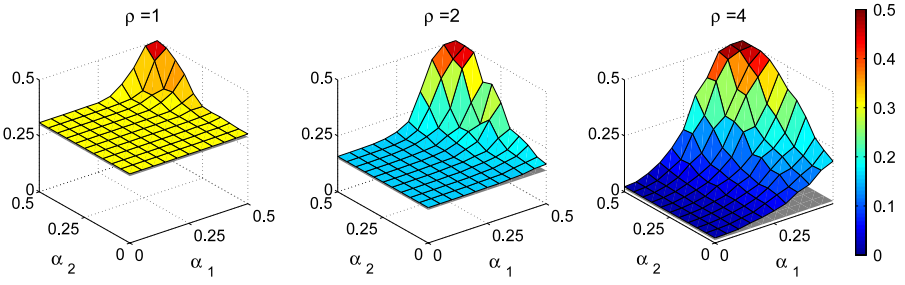

Fig. 3. Classification *error* for the adaptive classification method as a function of the false positive α_1 and false negative α_2 of the IErrP classifier. The results are shown for three different levels of overlap between the two distributions. A constant grey surface on the bottom indicates the optimal Bayes classifier which we use as a baseline. We set $\eta = 10^{-5}$ for this case, although results were equivalent for learning rates in a wide range. For better visualization the figure needs to be printed in color.

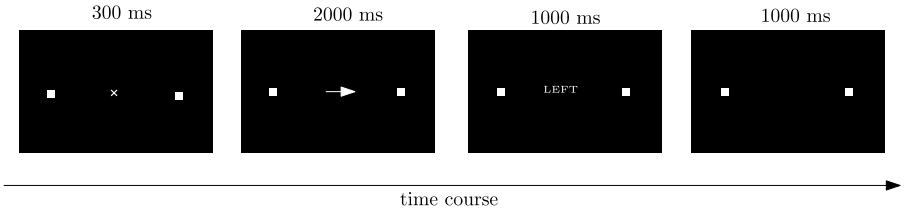

Fig. 4. The experimental protocol is a sequence of four steps: after a fixation cross an arrow appears and the subject has to attend accordingly. Two seconds later, the feedback is displayed as a text. At this point, a mismatch between the feedback and the direction instructed by the arrow is expected to elicit an IErrP.

and is replaced with a text indicating the decision of the device (right or left). The output is chosen at random in such a way that 80% of the trials are equal to the instructed direction and in the remaining 20% are equal to the opposite direction. The text remains visible for 1000 ms. Finally, the text disappears and the two squares persist for 1000 ms before a new trial starts. Although we instructed the subjects to direct the device in the indicated direction, they were not informed of the fact that the BCI output was not under their control but computed by the random protocol. For each subject we collected 504 trials divided in 6 sessions with one minute of rest between each two sessions (84 trials per session). Ongoing brain activity was recorded (sampling rate: 1200 Hz; low-pass filter: 300 Hz) using a whole head MEG system with 275 axial gradiometers (CTF Systems, Canada) inside a magnetically shielded room. Head localization was done before the experiment using coils that were placed at the nasion, right and left ear canal. The magnetic fields produced by these coils were used to measure the position of the subject's head with respect to the MEG sensor array. In addition we also recorded vertical and horizontal electrooculogram (EOG) in order to remove trials contaminated with muscular activity.

3.2. IErrP classification

The recorded data was segmented in trials starting 1500 ms after the onset of arrow presentation and finishing 1500 ms after the output of the device (left/right) was returned. Data segments contaminated with artifacts, eye movements, eye blinks or muscle activity were detected and removed using a semi-automatic routine. Then the planar gradiometer representation of the data was computed. The calculated planar field gradient approximates the signals measured by physical planar gradiometers (Bastiaansen, 2000). Time frequency representations (TFRs) of power were calculated using a multitaper approach applied to short sliding time windows (Percival & Walden, 1993). The data in each time window were multiplied with a Hanning taper with the length of the time window for the frequencies $[3, \ldots, 30]$ Hz. We applied an adaptive time window of length $\Delta T = 3/f$. The power values were calculated for the horizontal and vertical components of the estimated planar gradient and summed.

As features for classification we used the normalized power in all channels with a time window from 150 to 1000 ms after device output. The choice of frequency bands and time window of interest is based on previous work (Mazaheri et al., 2009). Note that we do not restrict the feature space to fronto-central areas as is usually the case for this problem (Ferrez & Millán, 2008).

In order to classify the surprise, E, based on the IErrP, we used a linear support vector machine (Vapnik, 1995), where the regularization parameter was determined based on initial empirical testing. The performance is calculated as the percentage of correctly classified trials. In order to reduce the dependence of the results on the particular split of train and test set, we repeated the ten-fold cross-validation procedure ten times, each time with a different partition of the data. In this way, the performance is estimated by averaging the performances of this repeated ten-fold cross-validation for each subject. This result is shown in the second column of Table 1.

Table 1
Classification results of IErrP.

Subject	Performance (%)	Confusion matrix	α_1	α_2
1	83.84	$\begin{pmatrix}0.841 & 0.159\\ 0.164 & 0.835\end{pmatrix}$	$\alpha_1^R \approx 0.085$ $\alpha_1^L \approx 0.073$	$\alpha_2^R \approx 0.082$ $\alpha_2^L \approx 0.082$
2	74.33	$\begin{pmatrix}0.757 & 0.243\\ 0.270 & 0.730\end{pmatrix}$	$\alpha_1^R \approx 0.110$ $\alpha_1^L \approx 0.133$	$\alpha_2^R \approx 0.138$ $\alpha_2^L \approx 0.132$
3	77.83	$\begin{pmatrix}0.781 & 0.219\\ 0.224 & 0.776\end{pmatrix}$	$\alpha_1^R \approx 0.091$ $\alpha_1^L \approx 0.128$	$\alpha_2^R \approx 0.100$ $\alpha_2^L \approx 0.124$
4	78.21	$\begin{pmatrix}0.788 & 0.211\\ 0.223 & 0.776\end{pmatrix}$	$\alpha_1^R \approx 0.102$ $\alpha_1^L \approx 0.109$	$\alpha_2^R \approx 0.114$ $\alpha_2^L \approx 0.109$
5	73.95	$\begin{pmatrix}0.756 & 0.244\\ 0.276 & 0.724\end{pmatrix}$	$\alpha_1^R \approx 0.126$ $\alpha_1^L \approx 0.118$	$\alpha_2^R \approx 0.121$ $\alpha_2^L \approx 0.155$
6	86.25	$\begin{pmatrix}0.880 & 0.120\\ 0.155 & 0.845\end{pmatrix}$	$\alpha_1^R \approx 0.067$ $\alpha_1^L \approx 0.052$	$\alpha_2^R \approx 0.082$ $\alpha_2^L \approx 0.072$
7	82.63	$\begin{pmatrix}0.835 & 0.165\\ 0.182 & 0.818\end{pmatrix}$	$\alpha_1^R \approx 0.082$ $\alpha_1^L \approx 0.082$	$\alpha_2^R \approx 0.092$ $\alpha_2^L \approx 0.090$
8	88.27	$\begin{pmatrix}0.907 & 0.092\\ 0.142 & 0.857\end{pmatrix}$	$\alpha_1^R \approx 0.050$ $\alpha_1^L \approx 0.042$	$\alpha_2^R \approx 0.069$ $\alpha_2^L \approx 0.073$
Mean	80.66	$\begin{pmatrix}0.819 & 0.181\\ 0.205 & 0.795\end{pmatrix}$	$\alpha_1^* \approx 0.147$	$\alpha_2^* \approx 0.089$

The mean classification rate is 80.66%, which as mentioned before is similar to the accuracy reached by other researchers. Column 3 of Table 1 shows the confusion matrices, where the false positive rate (α_1) and the false negative rate (α_2) are given in the upper right and lower left entries respectively. Column 4 shows the values of α_1^R and α_1^L and column 5 that of α_2^R and α_2^L. For $i \in \{1, 2\}$, the values of α_i^* in columns 4 and 5 are the mean across subjects of the relative difference between α_i^R and α_i^L. More precisely

$$\alpha_i^* = \frac{2}{N}\sum_{n=1}^{N}\frac{|\alpha_{(i,n)}^R - \alpha_{(i,n)}^L|}{\alpha_{(i,n)}^R + \alpha_{(i,n)}^L}, \tag{7}$$

where N is the number of subjects and n in the subindex denotes the subject number. We can conclude that the differences between α_i^L and α_i^R for $i \in \{1, 2\}$ are, for all subjects, relatively small and therefore will not significantly affect the expected performance of the adaptive method.

Fig. 5 shows the absolute values of the classifier parameters associated with each channel for subject 1. They correspond to the classical pattern of brain activity associated with the presence of error-related potentials.

3.3. Task classifier: static vs. adaptive

In this section we assess the expected performance increase of the proposed adaptive BCI in comparison to the static BCI. Ideally, such a comparison should be done in an on-line setting, where

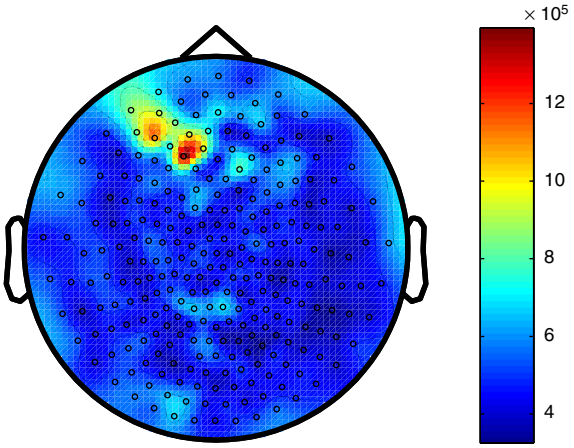

Fig. 5. Absolute values of the classifier parameters for subject 1. For better visualization the figure needs to be printed in color.

the actual output of the IErrP classifier is used to adapt the task classifier. However since the output of the actual IerrP classifier was not available during the reported BCI experiment, we have done this comparison in an off-line setting, reproducing the on-line scenario. Note that in the case of 100% IErrP detection, the off-line setting is equivalent to the on-line one.

For task discrimination we use the MEG data of the experiment reported in Section 3.1. The recorded data was segmented in trials starting at the onset of arrow presentation and finishing 500 ms after the output of the device (left/right) was returned. The data was then preprocessed as in Section 3.2, in this case for the frequencies $[8, \ldots, 14]$ Hz.

As features for the task classifier we used the normalized power in the posterior channels with a time window of 1500 ms starting 500 ms after arrow presentation. The choice of frequency bands and channels of interest is based on previous work (van Gerven & Jensen, 2009).

The static task classifier is trained in the first session and tested in the five remaining sessions for each subject independently. As a classifier we used the regularized logistic regression model as introduced in Section 2. In this static scenario, alternative classifiers such as SVMs gave comparable results.

The adaptive task classifier is initialized with the parameters of the static classifier. That is, it is trained in session 1, and is adapted during sessions 2–6. When the accuracy in IErrP detection is not 100%, we consistently simulate the IErrP classifier output by randomly generating false positives/negatives with the subject-dependent probabilities given by the error rates shown in columns 4 and 5 of Table 1. Since different realizations of the same experiment will return slightly different results, the presented results are the average over 100 different realizations of the simulated IErrP output to account for fluctuations in the adaptive learning. These results indicate the performance improvement that can be expected in the on-line setting.

For each subject, we chose the learning rate η in such a way that it maximized the total increase in performance. More information on how to set the learning rate is provided at the end of this section.

Fig. 6(A) shows for every subject the performance of the static classifier, the simulated adaptive method with realistic IErrP misclassification rates and the adaptive method with 100% error detection. Although there is a big variability in performance between subjects, we can see that the adaptive classifier improves the static classifier for all subjects in both scenarios.

In Fig. 6(B) we consider the effect of modifying the balance between α_1 and α_2. Given a fixed IErrP misclassification rate c, since $c = \frac{\alpha_1+\alpha_2}{2}$, we vary consistently the range of α_1 and α_2 between 0 and $2c$. In practice, α_1 and α_2 cannot be varied independently of the overall performance, so we only consider this result as an illustration to study the behavior of the method. To quantify the relative decrease in classification error for fixed α_1 and α_2 we define $\text{error}_a(\alpha_1, \alpha_2)$ as the error in the performance (number of incorrectly classified trials averaged over 100 realizations) of the adaptive classifier for a given α_1 and α_2. Then, the relative decrease in erroneous classification is defined as:

$$2\frac{\text{error}_a(2c, 0) - \text{error}_a(\alpha_1, \alpha_2)}{\text{error}_a(2c, 0) + \text{error}_a(\alpha_1, \alpha_2)}.$$

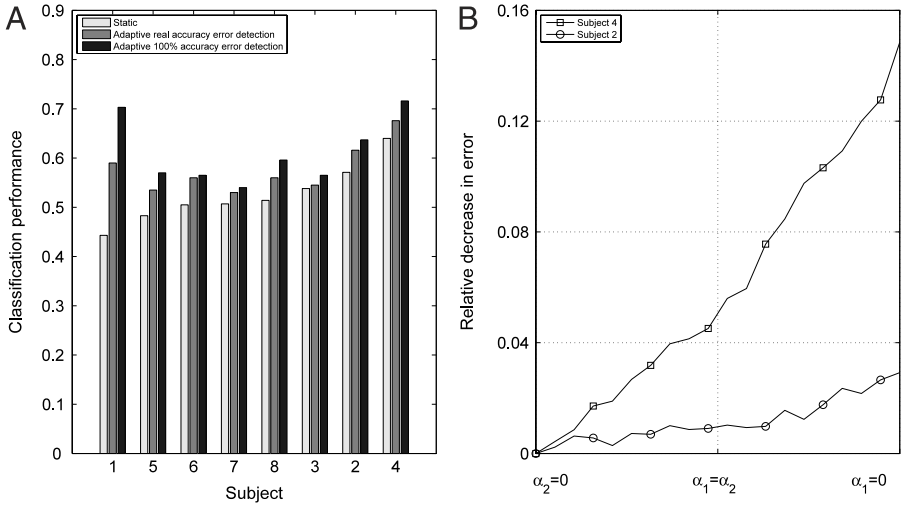

Fig. 6. Panel A depicts performance per subject of the static classifier, the adaptive classifier with 100% accuracy in error detection, and the simulated adaptive classifier with α_1^L, α_1^R, α_2^L and α_2^R rates as given for each subject in Table 1. Subjects are sorted in ascending order of performance of the static classifier. Results on the simulated case are averages between 100 different realizations. Panel B shows the influence on adaptive classification error when varying the false positives and false negatives rates for subjects 2 and 4 with a fixed misclassification rate of IErrP of 20%. The x-axis represents the dependence of the misclassification rate on α_1 and α_2, while the y-axis shows the relative decrease in classification error.

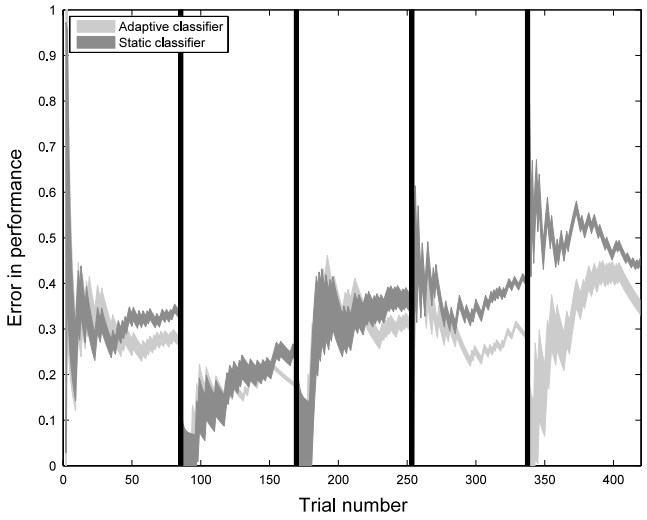

Fig. 7. Cumulative mean error and standard deviation for each session is shown as a function of the trial number for the adaptive and the static classifier for subject 4. Dark and light gray lines correspond to static and adaptive classifiers respectively. The thicknesses of the lines denote standard deviations. Black vertical lines show pauses between sessions.

Fig. 6(B) shows the results of this analysis considering $c = 0.2$ for subjects 2 and 4 since they are the subjects with better static classification rates. Note that the effect of the false positives is more damaging than that of the false negatives, and furthermore this effect becomes more pronounced as static task classification performance increases. These results agree with the simulations results reported in Section 2.3.

To illustrate the behavior of the adaptive and static method over time, we define error(k, j) equal to one if trial k in test session j was incorrectly classified and zero otherwise. Then we compute the cumulative mean error μ_e and standard deviations σ_e in trial i and test session j according to:

$$\mu_e(i,j) = \frac{\sum_{k=1}^{i} \text{error}(k,j)}{i},$$

$$\sigma_e(i,j) = \frac{\sum_{k=1}^{i} (\mu_e(k,j) - \mu_e(i,j))^2}{i-1}.$$

In Fig. 7 we show the cumulative mean performance error and standard deviation of the static and adaptive classifier (with 100% accurate error detection) across sessions for subject 4 versus the trial number. Black vertical lines show the pauses between sessions. As we can see, after a transient with high variances and similar performance of both classifiers, the adaptive classifier shows significantly better performance at the end of each session. Qualitatively similar results were obtained for all subjects.

Finally, we assessed the importance of the size of the learning rate on the performance of the adaptive classifier. Clearly, if the learning rate is too small, adaptation is too slow and when the learning rate is too large, the performance will decrease. This is confirmed in Fig. 8, which shows the performance of the adaptive classifier (with 100% accurate error detection) as a function of the learning rate for each subject separately. We can see that, in general, performance increases with the learning rate up to a subject-dependent threshold where it starts to decrease (except for the data of subject 3, which is very noisy). We observe that setting the learning rate between $[10^{-2}, 10^{-1}]$ is a safe choice for all subjects, but the optimal choice is subject dependent.

4. Discussion

In this article, we have demonstrated for the first time an improvement in BCI performance using neural feedback. We have presented our framework for a forced choice binary task using a

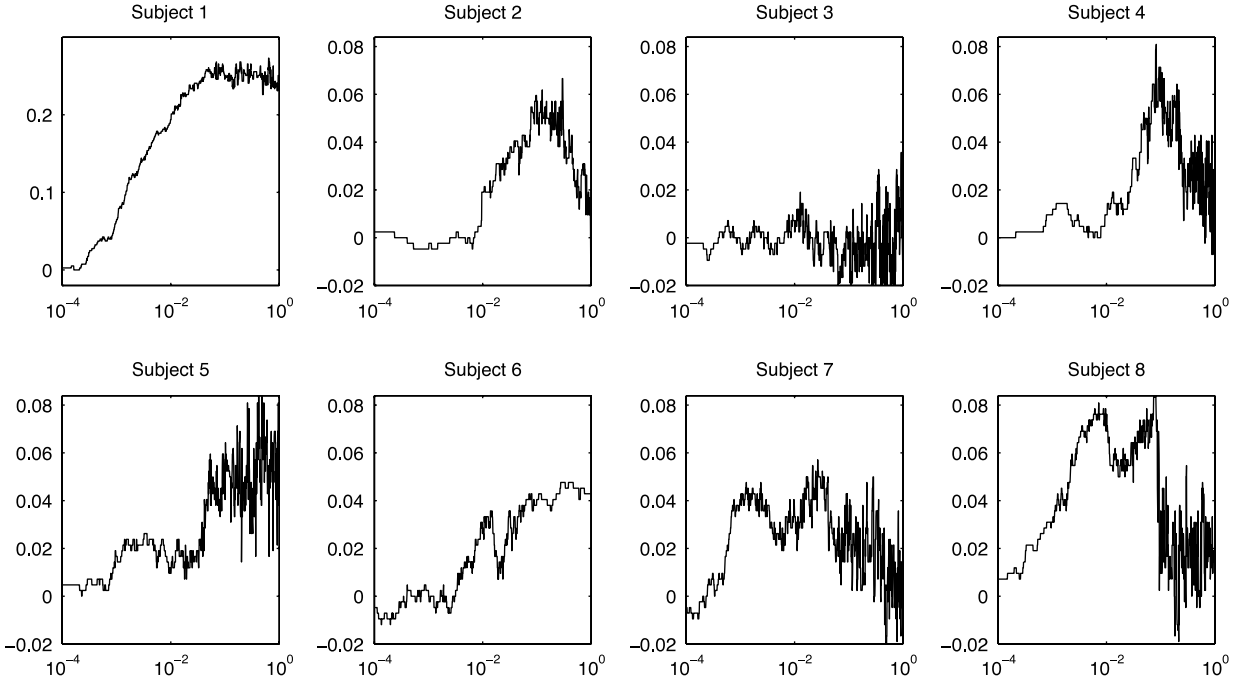

Fig. 8. Increase in the accuracy of the adaptive task classifier as a function of the learning rate with 100% of accuracy in error detection for all the subjects.

linear classifier that is adapted on-line using the IErrPs as neural feedback.

We have shown that the negative effect due to the misclassification of IErrPs is mainly due to the false positives and is practically insensitive to false negatives. One can therefore further optimize the adaptive BCI method by biasing the IErrP classifier such as to reduce the false positive rate at the cost of an increased false negative rate.

The impact of this work clearly depends on whether we can accurately detect the IErrP signal or not. We have shown that relatively high IErrP classification performance can be obtained for all subjects. Furthermore, other researchers have reported high stability in IErrP detection across sessions (Ferrez & Millán, 2008).

Ongoing and future work is oriented to the on-line validation of the method, the construction of a classifier of IErrP biased to decrease $\alpha 1$ and increase $\alpha 2$, and that uses features that generalize across subjects.

It is clear that the optimal choice of learning rate is important to obtain good results for all subjects. There exist methods that can automatically adapt the learning rate to an optimal value that make a trade-off between accuracy for stationary data (low learning rate) and adaptivity to change (large learning rates) (Heskes & Kappen, 1991, 1992). We believe that such methods should be integrated in an on-line adaptive BCI method.

The performance of the task classifier presented in this paper is quite poor for several subjects. We believe that this poor performance is due to a suboptimal choice of the features that we have used, and improvement in both the static and adaptive task classifier could be obtained using optimized features. However, we believe that the reported increase in performance of the adaptive classifier relative to the static classifier will still hold with optimized features. We have tested this by training the static classifier on all data for each subject independently, which results in significantly higher classification rates (60%–80%). Using the adaptive BCI procedure initialized with this static classifier on this same data resulted in an improved classification rate for all subjects.

An open question of considerable importance is how to generalize the proposed adaptive BCI method to non-binary tasks, such as controlling a keyboard. Such learning tasks are more complex, since the error signal will indicate that an error has occurred, but will not provide information on what the correct output should have been. This type of learning paradigm is called reinforcement learning (Sutton & Barto, 1998; Rescorla, 1967; Dayan & Abbott, 2001). An important future research direction is to integrate these reinforcement learning methods in on-line adaptive BCI.

Acknowledgments

The authors gratefully acknowledge the support of the BrainGain Smart Mix Programme of the Netherlands Ministry of Economic Affairs and the Netherlands Ministry of Education, Culture and Science.

Appendix. Effect of unbalanced IErrP classification

We show here that unbalanced classification of the IErrP only degrades the performance of the proposed model for very extreme cases. Following the notations of Section 2.3, we consider $\rho = 2$ (distance between the mean of the distributions), $\alpha_1 = \alpha_2 = 0.15$ (false positive/negative rates) and we vary, for $i \in \{1, 2\}$, the values of α_i^L and α_i^R consistently with (6).

Fig. A.9 shows the error in performance in the z-axis while the x- and y-axes represent the effect of left and right user intention in α_1 and α_2 respectively. The center in this axis represents balanced IErrP classification, i.e $\alpha_i^L = \alpha_i^R$ for $i \in \{1, 2\}$, while *Left* or *Right* indicate a total bias on α_1 (horizontal axis) or α_2 (vertical axis) towards that direction.

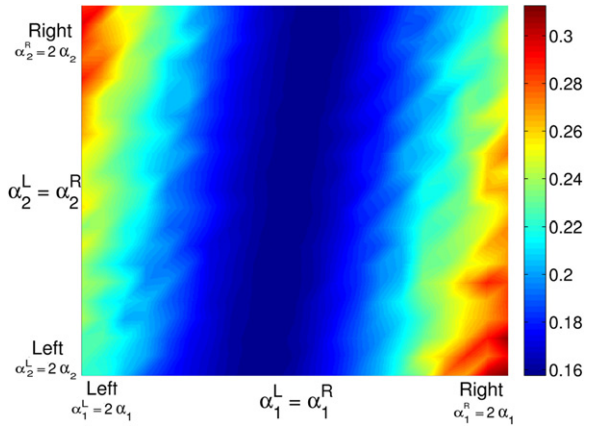

Fig. A.9. Classification error with $\rho = 2$ and $\alpha_1 = \alpha_2 = 0.15$ as a function of α_1^L, α_1^R, α_2^L and α_2^R. For better visualization the figure needs to be printed in color.

First, we see that the best performance is reached in the balanced case, and that the effect caused for unbalanced α_2 is less severe than that for unbalanced α_1. Further, the worst case occurs when false positives and negatives are unbalanced oppositely (see lower right and upper left corners). However, only in these extreme cases is this effect significant. We experienced that in all cases the method converges to a boundary close to the one obtained in the balanced case and further we obtained qualitatively the same results for different values of ρ.

References

Pfurtscheller, G., Neuper, C., Schölgl, A., & Lugger, K. (1998). Separability of EEG signals recorded during right and left motor imagery using adaptive autoregressive parameters. *IEEE Transactions on Rehabilitation and Engineering*, *6*, 316–325.

Bastiaansen, K. (2000). Tangential derivative mapping of axial meg applied to event-related desynchronization. *Clinical Neurophysiology*, *111*, 1300–1305.

Bishop, C. M. (2007). *Pattern recognition and machine learning (information science and statistics)* (1st ed.). Springer.

Chavarriaga, R., Ferrez, P. W., & Millán, J. D. R. (2007). To err is human: learning from error potentials in brain–computer interfaces. In R. Wang, F. Gu, & E. Shen (Eds.), *International conference on cognitive neurodynamics* (pp. 777–782).

Dayan, P., & Abbott, L. F. (2001). *Theoretical neuroscience: computational and mathematical modeling of neural systems* (1st ed.). The MIT Press.

Falkenstein, M., Hohnsbein, J., Hoormann, J., & Blanke, L. (1990). Effects of error in choice reaction tasks on the ERP under focused and divided attention. In C. Brunia, A. Gaillard, & A. Kok (Eds.), *Psychophysiological brain research* (pp. 192–195). Tilburg: Tilburg Univerity Press.

Ferrez, P. W. (2007). *Error-related EEG potentials in brain–computer interfaces. Ph.D. thesis*. Thèse Ecole polytechnique fédérale de Lausanne EPFL, no. 3928.

Ferrez, P. W., & Millán, J. D. R. (2008). Error-related eeg potentials generated during simulated brain–computer interaction. *IEEE Transactions on Biomedical Engineerings*, *55*, 923–929.

Ferrez, P. W., & Millán, J. D. R. (2005). You are wrong!—automatic detection of interaction errors from brain waves. In *Proceedings of the 19th international joint conference on artificial intelligence*. Edinburgh, UK.

Heskes, T., & Kappen, B. (1991). Learning processes in neural networks. *Physical Review A*, *44*, 2718–2726.

Heskes, T., & Kappen, B. (1992). Learning-parameter adjustment in neural networks. *Physical Review A*, *45*, 8885–8893.

Mazaheri, A., Nieuwenhuis, I. L. C., van Dijk, H., & Jensen, O. (2009). Prestimulus alpha and mu activity predicts failure to inhibit motor responses. *Human Brain Mapping*, *30*, 1791–1800.

Miltner, W. H. R., Braun, C. H., & Coles, M. G. H. (1997). Event-related brain potentials following incorrect feedback in a time-estimation task: evidence for a generic neural system for error detection. *Journal of Cognitive Neuroscience*, *9*, 788–798.

Percival, D., & Walden, A. (1993). *Spectral analysis for physical applications: multitaper and conventional univariate techniques*. Cambridge, UK: Cambridge University Press.

Rescorla, R. (1967). Pavlovian conditioning and its proper control procedures. *Psychological Review*, *74*, 71–80.
Seno, B. D., Matteucci, M., & Mainardi, L. T. (2010). Online detection of P300 and error potentials in a BCI speller. *Computational Intelligence and Neuroscience*,.
Sutton, R. S., & Barto, A. G. (1998). *Reinforcement learning: an introduction*. MIT Press.
Sykacek, P., Roberts, S., & Stokes, M. (2004). Adaptive BCI based on variational Kalman filtering: an empirical evaluation. *IEEE Transactions on Biomedical Engineering*, *51*.
Shenoy, P., Krauledat, M., Blankertz, B., Rao, R. P., & Müller, K.-R. (2006). Towards adaptive classification for BCI. *Journal of Neural Engineering*, *3*.
van Gerven, M., Bahramisharif, A., Heskes, T., & Jensen, O. (2009a). 2009 special issue: selecting features for BCI control based on a covert spatial attention paradigm. *Neural Networks*, *22*, 1271–1277.
van Gerven, M., Hesse, C., Jensen, O., & Heskes, T. (2009b). Interpreting single trial data using groupwise regularisation. *NeuroImage*, *46*, 665–676.
van Gerven, M., & Jensen, O. (2009). Attention modulations of posterior alpha as a control signal for two-dimensional brain–computer interfaces. *Journal of Neuroscience Methods*, *179*, 78–84.
van Schie, H. T., Mars, R. B., Coles, M. G. H., & Bekkering, H. (2004). Modulation of activity in medial frontal and motor cortices during error observation. *Nature Neuroscience*, *7*, 549–554.
Vapnik, V. (1995). *The nature of statistical learning theory*. New York: Springer.