You Are Wrong!—Automatic Detection of Interaction Errors from Brain Waves

Pierre W. Ferrez, José del R. Millán
IDIAP Research Institute
Rue du Simplon 4, 1920 Martigny, Switzerland
{pierre.ferrez, jose.millan}@idiap.ch

Abstract

Brain-computer interfaces, as any other interaction modality based on physiological signals and body channels (e.g., muscular activity, speech and gestures), are prone to errors in the recognition of subject's intent. In this paper we exploit a unique feature of the "brain channel", namely that it carries information about cognitive states that are crucial for a purposeful interaction. One of these states is the awareness of erroneous responses. Different physiological studies have shown the presence of error-related potentials (ErrP) in the EEG recorded right after people get aware they have made an error. However, for human-computer interaction, the central question is whether ErrP are also elicited when the error is made by the interface during the recognition of the subject's intent and no longer by errors of the subject himself. In this paper we report experimental results with three volunteer subjects during a simple human-robot interaction (i.e., bringing the robot to either the left or right side of a room) that seem to reveal a new kind of ErrP, which is satisfactorily recognized in single trials. These recognition rates significantly improve the performance of the brain interface.

1 Introduction

A brain-computer interface (BCI) is a system that translates the intention of a subject, as represented by his brain waves (typically, electroencephalogram signals), into a control signal without using activity of any muscles or peripheral nerves [Millán, 2002; Wolpaw *et al.*, 2002]. Electroencephalogram (EEG) is recorded non-invasively by means of electrodes placed on the scalp. Then some features are extracted from the EEG and sent to a classifier, whose response is translated into some action thus making possible mental control of different devices. For instance, recent studies have shown that after a few days of training, subjects are able to control a miniature robot in an indoor environment with several rooms, corridors and doorways only using the signals derived from an EEG-based BCI [Millán *et al.*, 2003]. This same system could be used to drive a wheelchair so that a subject suffering from severe motor disabilities, but with intact cognitive capabilities, could greatly gain in autonomy. Other applications, such a virtual keyboard, could more generally provide an alternative way of communication with the outside world [Birbaumer *et al.*, 1999; Wolpaw *et al.*, 2003; Millán *et al.*, 2004].

Nevertheless, BCIs—as any other interaction modality based on physiological signals and body channels (e.g., muscular activity, speech and gestures)—are prone to errors in the recognition of subject's intent, and those errors can be frequent. Indeed, even well-trained subjects rarely reach 100% of success. A possible way to reduce errors consists in a verification procedure whereby each output consists of two opposite trials, and success is required on both to validate the outcome [Wolpaw *et al.*, 1998]. Even if this method greatly reduces the errors, it requires much more mental effort from the subject and reduces the communication rate.

In contrast to other interaction modalities, a unique feature of the "brain channel" is that it conveys both information from which we can derive mental control commands to operate a brain-actuated device as well as information about cognitive states that are crucial for a purposeful interaction, all this on the millisecond range. One of these states is the awareness of erroneous responses, which a number of groups have recently started to explore as a way to improve the performance of BCIs [Schalk *et al.*, 2000; Blankertz *et al.*, 2003; Parra *et al.*, 2003]. Since the late 1980s, different physiological studies have shown the presence of error-related potentials (ErrP) in the EEG recorded right after people get aware they have made an error [Carter *et al.*, 1998; Falkenstein *et al.*, 2000; Holroyd and Coles, 2002]. Nevertheless, most of these studies show the presence of ErrP in typical choice reaction tasks [Carter *et al.*, 1998; Falkenstein *et al.*, 2000; Blankertz *et al.*, 2003; Parra *et al.*, 2003]. In this kind of tasks, the subject is asked to respond as quickly as possible to a stimulus and ErrP (sometimes referred to as "response ErrP") arise following errors due to the subject's incorrect motor action (e.g., the subject pressed a key with the left hand when he should have responded with the right hand). More recently, other studies have also shown the presence of ErrP in typical reinforcement learning tasks where the subject is asked to make a choice and ErrP (sometimes referred to as "feedback ErrP") arise following the presentation of a stimulus that indicate incorrect performance [Holroyd and Coles, 2002]. In

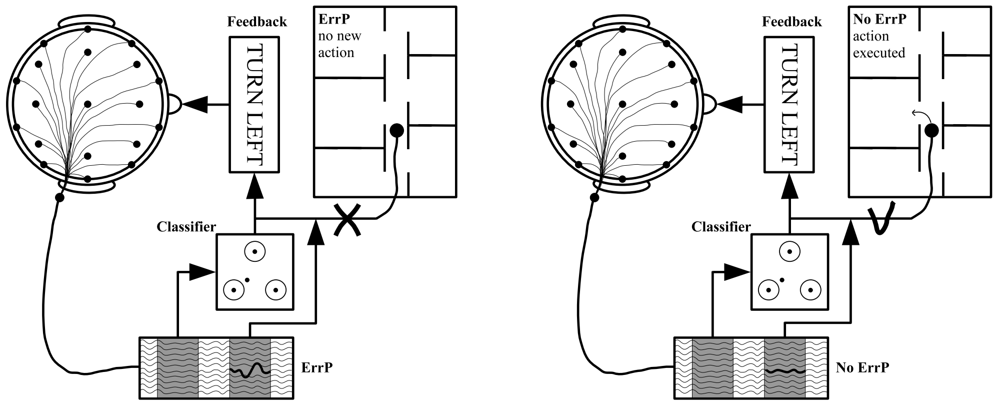

Figure 1: Exploiting error-related potentials (ErrP) in a brain-controlled mobile robot. The subject receives visual feedback indicating the output of the classifier before the actual execution of the associated command (e.g., "TURN LEFT"). If the feedback generates an ErrP (left), this command is simply ignored and the robot will stay executing the previous command. Otherwise, the command is sent to the robot (right).

both cases, response ErrP and feedback ErrP, it has been observed a negative potential in the EEG with a fronto-central scalp distribution (see Figure 3) whenever subjects make errors. This negative potential is most probably generated in a brain area called anterior cingulate cortex, which is crucial for regulating emotional responses.

An important aspect of the two described ErrP is that they always follow an error made by the subject himself. First the subject make a selection, and then ErrP arise either simply after the occurrence of an error (choice reaction task) or after a feedback indicating the error (reinforcement learning task). However, in the context of a BCI, or human-computer interaction in general, the central question is:

"Are ErrP also elicited when the error is made by the interface during the recognition of the subject's intent?"

In order to consider the full implications of this question, let's imagine that the subject's intent is to make a robot reach a target to the left. What would happen if the interface fails to recognize the intended command and the robot starts turning in the wrong direction? Are ErrP still present even though the subject did not make any error but only perceive that the interface is performing wrongly?

The objective of this paper is to investigate how ErrP could be used to improve the performance of a BCI. Thus, we will first explore whether or not ErrP also follow a feedback indicating incorrect responses of the interface and no longer errors of the subject himself. If ErrP are also elicited in this case, then we could integrate them in a BCI in the following way as shown in Figure 1: after translating the subject's intention into a control command, the BCI provides a feedback of that command, which will be actually executed only if no ErrP follows the feedback. This should greatly increase the reliability of the BCI as we will see later. Of course, this new interaction protocol depends on the ability to detect ErrP no longer in averages of a large number of trials [Schalk *et al.*, 2000], but in each single trial using a short window following the feedback that shows the response of the classifier embedded in the BCI.

In this paper we report experimental results with three volunteer subjects during a simple human-robot interaction (i.e., bringing the robot to either the left or right side of a room) that seem to reveal a new kind of ErrP, which is satisfactorily recognized in single trials. These recognition rates significantly improve the performance of the interface.

2 Experimental Setup

To test the presence of ErrP after a feedback indicating errors made by the interface in the recognition of the subject's intent, we have simulated a real interaction with a robot, where the subject wishes to bring the robot to one side of a room (left or right) by delivering repetitive commands until the robot reaches the target. This virtual interaction is implemented by means of two horizontal progress bars made of 10 steps each. One of the bars goes from the center of the screen to the left side (left bar), and the other bar progresses to the right side (right bar). Figure 2 shows the left and right horizontal progress bars used as feedback.

To isolate the issue of the recognition of ErrP out of the more difficult and general problem of a whole BCI where erroneous feedback can be due to non-optimal performance of both the interface (i.e., the classifier embedded into the interface) and the user himself, in the following experiments the subject delivers commands manually and not mentally. That is, he simply presses a left or right key with the left or right hand. In this way, any error feedback is only due to a wrong recognition of the interface of which is the subject's intention.

Three volunteer healthy subjects participated in these experiments. The subject sends a command after a stimulus delivered by the system (the word "GO" appears in the screen). The system filled the bars with an error rate of 20%; i.e., at each step, there was a 20% probability that the robot made a movement in the wrong direction. Subjects performed 10 series of 5 progress bars, the delay between 2 consecutive steps (2 consecutive "GO" from the system) was of 3-4 seconds (random delay to prevent habituation). Duration of each interaction experiment (i.e., filling a progress bar) was about 40 seconds, with breaks of 5-10 minutes between two series but no break between interaction experiments of the same series.

EEG potentials were acquired with a portable system (Biosemi ActiveTwo) by means of a cap with 32 integrated electrodes covering the whole scalp and located according to

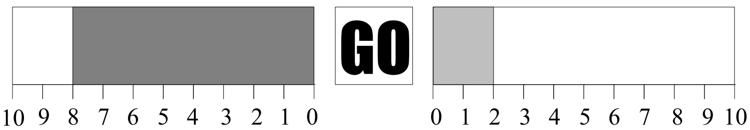

Figure 2: Left and right horizontal progress bars. The goal of an interaction experiment is to fill one of the bars, which simulates a real interaction with a robot that needs to reach one side of a room (left or right). The system fills the bars with an error rate of 20%; i.e., at each step, there was a 20% probability that the robot made a movement in the wrong direction.

the standard 10/20 International system. The sampling rate was 512 Hz and signals were measured at full DC. Raw EEG potentials were first spatially filtered by subtracting from each electrode the average potential (over the 32 channels) at each time step. The aim of this re-referencing procedure is to suppress the average brain activity, which can be seen as underlying background activity, so as to keep the information coming from local sources below each electrodes. Then, we applied a 1-10 Hz bandpass filter as ErrP are known to be a relatively slow cortical potential. Finally, EEG signals were subsampled from 512 Hz to 128 Hz (i.e., we took one point out of 4) before classification, which was entirely based on temporal features. Indeed the actual input vector for the statistical classifier described below is a $1/2$-second window starting 150 ms after the feedback and ending 650 ms after the feedback for channels Cz and Fz. The choice of these channels follows the fact that ErrP are characterized by a fronto-central distribution along the midline (see Figure 3). Thus the dimensionality of the input vector is 128; i.e., concatenation of two windows of 64 points (EEG potentials) each.

No artifact rejection algorithm (for removing or filtering out eye or muscular movements) was applied and all trials were kept for analysis. It is worth noting, however, that after a visual a-posteriori check of the trials we found no evidence of muscular artifacts that could have contaminated one condition differently from the other.

3 Statistical Classifier

The different classes are recognized by a Gaussian classifier trained to classify single trials as "correct" or "error". The output of this statistical classifier is an estimation of the posterior class probability distribution for a single trial; i.e., the probability that a given single trial belongs to class "correct" or class "error". The present classifier is quite similar to that proposed by Millán *et al.* [2004] in the framework of EEG-based BCIs, the main difference being the update rules.

In this statistical classifier, every Gaussian unit represents a prototype of one of the classes to be recognized. We use several prototypes per mental task. We assume that the class-conditional probability density function of class C_k is a superposition of N_k Gaussians (or prototypes) and that classes have equal prior probabilities. In our case, all the classes have the same number of prototypes N_p. In addition, we assume that all prototypes have an equal weight of $1/N_p$. Then, dropping constant terms, the activity a_k^i of the i^{th} prototype of the class C_k for the input vector, or sample, x derived from a trial as described above is

$$a_k^i(x) = |\Sigma_k|^{-1/2} \exp\left(-1/2\left(x-\mu_k^i\right)^T \Sigma_k^{-1}\left(x-\mu_k^i\right)\right) \quad (1)$$

where μ_k^i is the center of the i^{th} prototype of the class C_k, Σ_k is the covariance matrix of the class C_k, and $|\Sigma_k|$ is the determinant of that matrix. Usually, each prototype has its own covariance matrix Σ_k^i. In order to reduce the number of parameters, we restrict our model to a diagonal covariance matrix Σ_k that is common to all the prototypes of the class C_k. Now, the posterior probability y_k of the class C_k is

$$y_k(x) = p(x|C_k) = \frac{a_k(x)}{A(x)} = \frac{\sum_{i=1}^{N_p} a_k^i(x)}{\sum_{k=1}^{2}\sum_{i=1}^{N_p} a_k^i(x)} \quad (2)$$

where a_k is the activity of class C_k and A is the total activity of the network. The response of the classifier for the input vector x is simply the class with the highest probability.

To initialize the center of the prototypes, μ_k^i, of the class C_k we run a clustering algorithm—typically, self-organizing maps [Kohonen, 1997]. We then initialize the diagonal covariance matrix Σ_k of the class C_k by setting

$$\Sigma_k = \frac{1}{|S_k|}\sum_{x\in S_k}\left(x-\mu_k^{i*}\right)\left(x-\mu_k^{i*}\right)^T \quad (3)$$

where S_k is the set of the training samples belonging to the class C_k, $|S_k|$ is the cardinality of this set, and i^* is the nearest prototype of this class to the sample x.

During learning we improve these initial estimations iteratively by stochastic gradient descent so as to minimize the mean square error $E = 1/2\sum_k (y_k - t_k)^2$, where t_k is the k^{th} component of the target vector in the form 1-of-c; e.g., the target vector for class "error" is coded as $(0, 1)$. Taking the gradient of the error function yields

$$\Delta\mu_k^i(x) = \alpha\frac{\partial E}{\partial \mu_k^i}(x) = \alpha\frac{a_k^i(x)}{A(x)}\frac{\left(x-\mu_k^i\right)}{\Sigma_k}\,e_k(x) \quad (4)$$

and

$$\Delta\Sigma_k^i(x) = \beta\frac{\partial E}{\partial \Sigma_k^i}(x) = \beta\frac{a_k^i(x)}{A(x)}\frac{\left(x-\mu_k^i\right)^2}{\left(\Sigma_k^i\right)^3}\,e_k(x) \quad (5)$$

where α and β are the learning rates and $e_k(x) = (t_k(x)-y_k(x))-\sum_j y_j(x)(t_j(x)-y_j(x))$. After updating μ_k^i and Σ_k^i for each training sample, the covariance matrices of all the prototypes of the same class are averaged to obtain the common class covariance matrix Σ_k. This simple operation leads to better performance than if separate covariance matrices are kept for each individual prototype. It is also worth

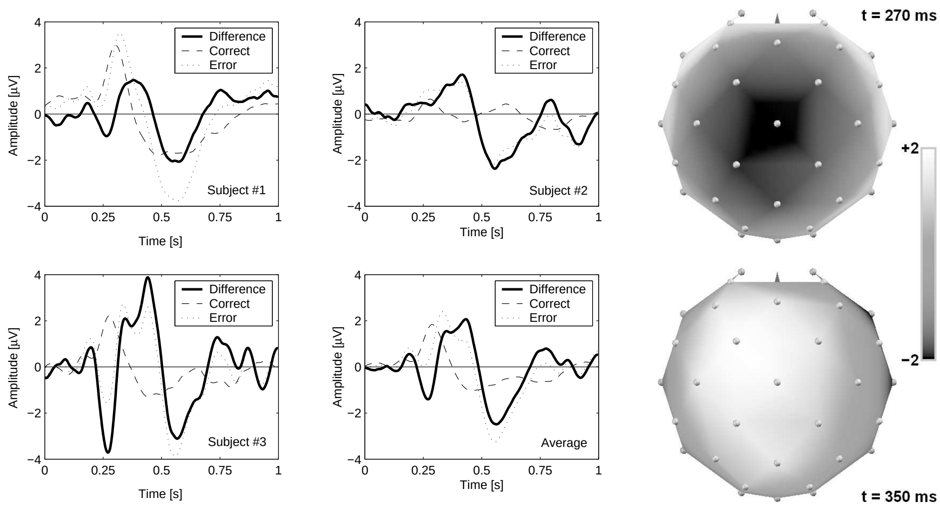

Figure 3: *Left.* Average EEG for error, correct and difference error-minus-correct at channel Cz for the three subjects plus the grand average of them. Feedback is delivered at time 0 seconds. The negative (Ne) and positive (Pe) peaks show up about 270 ms and between 350 and 450 ms after the feedback, respectively. *Right.* Scalp potentials topographies, for the grand average EEG of the three subjects, at the occurrence of the Ne and the Pe. Small filled circles indicate positions of the electrodes (frontal on top), Cz being in the middle of the scalp.

noting that, due to the small number of training samples for the "error" class, it is sometimes preferible to have a single covariance matrix common to all classes, which is obtained by averaging all the individual Σ_k.

4 Experimental Results

With this protocol, it is first necessary to investigate wether or not ErrP are present no more in reaction to an error made by the subject himself, but in reaction to an erroneous response made the interface as indicated by the feedback visualizing the recognized subject's intention. Figure 3 shows the different curves (error, correct and difference error-minus-correct) for channel Cz for the three subjects plus the grand average of the three subjects. A first sharp negative peak (Ne) can be clearly seen 270 ms after the feedback (except for subject 2). A later positive peak (Pe) appears between 350 and 450 ms after the feedback. Finally an additional negative peak occurs ~550 ms after the feedback. Figure 3 also shows the scalp potentials topographies, for the grand average EEG of the three subjects, at the occurrence of the maximum of the Ne and the Pe: a first central negativity appears after 270 ms, followed by a fronto-central positivity at 350 ms.

These experiments seem to reveal a new kind of error-related potentials that, for convenience, we call "interaction ErrP". The general shape of this ErrP is quite similar to the shape of the response ErrP in a choice reaction task, whereas the timing is similar to the feedback ErrP of reinforcement learning tasks[1]. As in the case of response ErrP, the interaction ErrP exhibits a first sharp negative peak followed by a broader positive peak. This is quite different from the shape of feedback ErrP that are only characterized by a small negative deflection. On the other hand, the time course of the interaction ErrP bears some similarities to that of the feedback ErrP: in both cases the first distinctive feature (negative peak and negative deflection, respectively) appears ~250 ms after feedback. The time course of response ErrP is definitely different as the peaks show up much faster.

4.1 Single Trial Classification

To explore the feasibility of detecting single-trial erroneous responses, we have done a 10-fold cross-validation study where the testing set consists of one of the recorded sessions. In this way, testing is always done on a different recording session to those used for training the model. This yields a better estimation of the actual generalization capabilities because brain activity naturally changes over time, as Figure 4 illustrates. Top panels show the average EEG of two consecutive series (5 interaction experiments each) at channel Cz for subject 3, while the bottom panels show samples of single trials of those series. Their variability is rather evident.

Table 1: Percentages of correctly recognized error trials and correct trials for the 3 subjects and the average of them.

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

Table 1 reports the recognition rates (mean and standard deviations) for the three subjects plus the average of them. The different hyper-parameters—i.e., the learning rates of the centers and diagonal covariance matrices, number of prototypes, and common/single covariance matrices for each class—were chosen by model selection in the training sets. Regarding the learning rates, usual values were 10^{-4} to 10^{-6} for the centers and 10^{-6} to 10^{-8} for the variances, while the usual number of prototypes was rather small (from 2 to 4). These results show that single-trial recognition of erroneous responses is almost 80% on average, while the recognition rate of correct responses is slightly better (82.4%). Quite im-

[1]Due to space limitations, experiments showing response and feedback ErrP cannot be reported.

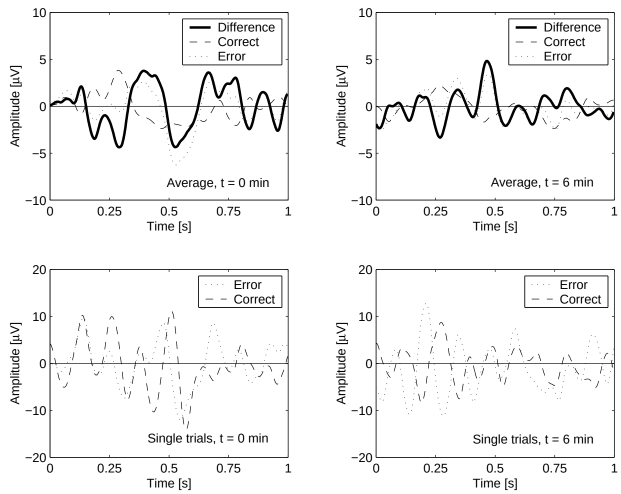

Figure 4: *Top.* Average EEG (error, correct and difference error-minus-correct) of two consecutive experimental series at channel Cz for subject 3. *Bottom.* Samples of single trials (error and correct) of those two series.

portantly, even for the subject with the worse detection rates, they are around 75%. But the key point for the exploitation of the automatic recognition of interaction errors is that they translate into an actual improvement of the performance of the BCI, which we can measure in terms of the bit rate.

4.2 Bit Rate Improvement

A traditional measure of the performance of a system is the bit rate, the amount of information communicated per unit time. The bit rate is usually expressed in bits per trial (bits per selection). If a single trial has N_c possible outcomes, if the probability p that this outcome is correct (accuracy of the BCI), and if finally each of the other outcomes has the same probability of selection (i.e., $(1-p)/(N_c-1)$), then the information transfer rate in bits per trial (BpT) is

$$BpT = \log_2(N_c) + p\log_2(p) + (1-p)\log_2\left(\frac{1-p}{N_c-1}\right) \quad (6)$$

Let's consider now how the performance of the BCI changes after introducing ErrP and that the system detects a proportion e of erroneous trials and a proportion c of correct trials. In the general case, after detecting an erroneous trial the outcome of the interface is simply stopped and not sent to the brain-actuated device. The new accuracy p' of the BCI becomes $p' = p \cdot c/p_t$, where $p_t = p \cdot c + (1-p)\cdot(1-e)$. Now the new information transfer rate in bits per trial is

$$BpT = p_t\left(\log_2(N_c) + p'\log_2(p') + (1-p')\log_2\left(\frac{1-p'}{N_c-1}\right)\right) \quad (7)$$

In the case of a 2-class BCI ($N_c = 2$), after detecting an erroneous trial, it could be possible to replace the "wrong" outcome by the opposite one, what yields an accuracy $p'' = p \cdot c + (1-p)\cdot e$. The information transfer rate in this case is calculated by replacing p by p'' in equation 6, because now there is no stopped outcome.

Table 2 reports the performances of a BCI that integrates ErrP for the 3 subjects and the average of them, where we have assumed an accuracy of 80% in the recognition of the subject's intent. These figures are to be compared to the performance of a standard BCI (i.e., without integrating ErrP). We have also reported the performances in the case of $N_c = 3$, as the mind-controlled robot described by Millán *et al.* [2003]. In the case of standard 2-class and 3-class BCI, their performances are 0.28 and 0.66 bits per trial, respectively. Results indicate that there is a significant improvement in performance in the case of stopping outcomes, which is above 70% on average and higher than 90% for one of the subjects. Surprisingly, replacing leads to smaller improvements and, in the case of subject 2, even to a significant degradation.

Table 2: Performances of the BCI integrating ErrP for the 3 subjects and the average of them.

	$N_c = 3$ Stop		$N_c = 2$ Stop		$N_c = 2$ Replace	
Subjects	BpT	Gain	BpT	Gain	BpT	Gain
#1	0.91	37%	0.53	91%	0.36	29%
#2	0.73	10%	0.40	42%	0.19	-32%
#3	0.92	38%	0.52	86%	0.44	59%
Average	*0.85*	*28%*	*0.48*	*72%*	*0.32*	*14%*

5 Discussion

In this paper we have reported first results on the detection of the neural correlate of error awareness for improving the performance and reliability of brain-computer interfaces. In particular, we have found what seems to be a new kind of error-related potential elicited in reaction to an erroneous recognition of the subject's intention. More importantly, we have shown the feasibility of detecting single-trial erroneous responses of the interface that leads to significant improvements of the information transfer rate of a BCI.

Given the promising results obtained in a simulated human-robot interaction, we are currently working in the actual integration of ErrP detection into our BCI system. In parallel, we are exploring how to increase the recognition rate of single-trial erroneous and correct responses. A basic issue here is to find what kind of feedback elicits the strongest "interaction ErrP". The feedback can be of very different nature—visual, auditory, somatosensory, or even a mix of these different types. More importantly, we will need to focus on alternative methods to exploit at best the current "interaction ErrP". In this respect, Grave *et al.* [2004] have recently developed a technique that estimates the so-called local field potentials (i.e., the synchronous activity of a small neuronal population) in the whole human brain from scalp EEG. Furthermore, recent results show significant improvements in the classification of bimanual motor tasks using estimated local field potentials (LFP) with respect to scalp EEG [Grave *et al.*, 2005]. Consequently, we plan to utilize the ELECTRA method to best discriminate erroneous and correct responses of the interface. As a matter of fact, a key issue for the success in the above-mentioned study was the selection of those relevant voxels inside the brain whose estimated LFP were more discriminant. It turns out that the sources of the ErrP seem to be very well localized into the anterior cingulate cortex and thus we may well expect a significant improvement in recognition rates by focusing on the LFP estimated in this specific brain area.

More generally, the work described in this paper suggests that it could be possible to recognize in real time high-level cognitive and emotional states from EEG (as opposed, and in addition, to motor commands) such as fatigue, stress, frustration, confusion, or attention that are crucial for an effective and purposeful interaction. Indeed, the rapid recognition of these states will lead to truly adaptive interfaces that customize dynamically in response to changes of the cognitive and emotional/affective states of the user.

Acknowledgments

This work is supported by the Swiss National Science Foundation NCCR "IM2" and by the European IST Programme FET Project FP6-003758. 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.

References

[Birbaumer *et al.*, 1999] N. Birbaumer, N. Ghanayim, T. Hinterberger, I. Iversen, B. Kotchoubey, A. Kubler, J. Perelmouter, E. Taub, and H. Flor. A spelling device for the paralysed. *Nature*, 398:297–298, 1999.

[Blankertz *et al.*, 2003] 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 Transactions on Neural Systems and Rehabilitation Engineering*, 11:127–131, 2003.

[Carter *et al.*, 1998] C.S. Carter, T.S. Braver, D.M. Barch, M.M Botvinick, D. Noll, and J.D. Cohen. Anterior cingulate cortex, error detection, and the online monitoring of performance. *Science*, 280:747–749, 1998.

[Falkenstein *et al.*, 2000] M. Falkenstein, J. Hoormann, S. Christ, and J. Hohnsbein. ERP components on reaction errors and their functional significance: A tutorial. *Biological Psychology*, 51:87–107, 2000.

[Grave *et al.*, 2004] R. Grave de Peralta, M.M. Murray, C.M. Michel, R. Martuzzi, and S.L. Gonzalez Andino. Electrical neuroimaging based on biophysical constraints. *NeuroImage*, 21:527–539, 2004.

[Grave *et al.*, 2005] R. Grave de Peralta, S.L. Gonzalez Andino, L. Perez, P.W. Ferrez, and J. del R. Millán. Non-invasive estimation of local field potentials for neuroprosthesis control. *Cognitive Processing*, 6, 2005.

[Holroyd and Coles, 2002] C.B. Holroyd and M.G.H. Coles. The neural basis of human error processing: Reinforcement learning, dopamine, and the error-related negativity. *Psychological Review*, 109:679–709, 2002.

[Kohonen, 1997] T. Kohonen. *Self-Organizing Maps, 2nd ed.* Springer-Verlag, Berlin, 1997.

[Millán *et al.*, 2003] J. del R. Millán, F. Renkens, J. Mouriño, and W. Gerstner. Non-invasive brain-actuated control of a mobile robot. In *Proceedings of the 18th International Joint Conference on Artificial Intelligence*, pages 1121–1126, Acapulco, Mexico, 2003. Morgan Kaufmann.

[Millán *et al.*, 2004] J. del R. Millán, F. Renkens, J. Mouriño, and W. Gerstner. Brain-actuated interaction. *Artificial Intelligence*, 159:241–259, 2004.

[Millán, 2002] J. del R. Millán. Brain-computer interfaces. In M.A. Arbib, editor, *Handbook of Brain Theory and Neural Networks*, pages 178–181. The MIT Press, Cambridge, Massachusetts, 2002.

[Parra *et al.*, 2003] L.C. Parra, C.D. Spence, A.D. Gerson, and P. Sajda. Response error correction—A demonstration of improved human-machine performance using real-time EEG monitoring. *IEEE Transactions on Neural Systems and Rehabilitation Engineering*, 11:173–177, 2003.

[Schalk *et al.*, 2000] G. Schalk, J.R. Wolpaw, D.J. McFarland, and G. Pfurtscheller. EEG-based communication: Presence of an error potential. *Clinical Neurophysiology*, 111:2138–2144, 2000.

[Wolpaw *et al.*, 1998] J.R. Wolpaw, H. Ramoser, D.J. McFarland, and G. Pfurtscheller. EEG-based communication: Improved accuracy by response verification. *IEEE Transactions on Rehabilitation Engineering*, 6:326–333, 1998.

[Wolpaw *et al.*, 2002] J.R. Wolpaw, N. Birbaumer, D.J. McFarland, G. Pfurtscheller, and T.M. Vaughan. Brain-computer interfaces for communication and control. *Clinical Neurophysiology*, 113:767–791, 2002.

[Wolpaw *et al.*, 2003] J.R. Wolpaw, D.J. McFarland, T.M. Vaughan, and G. Schalk. The Wadsworth Center brain-computer interface (BCI) research and development program. *IEEE Transactions on Neural Systems and Rehabilitation Engineering*, 11:204–207, 2003.