Error-Related EEG Potentials Generated During Simulated Brain–Computer Interaction

Pierre W. Ferrez* and José del R. Millán*

Abstract—Brain–computer interfaces (BCIs) are prone to errors in the recognition of subject's intent. An elegant approach to improve the accuracy of BCIs consists in a verification procedure directly based on the presence of error-related potentials (ErrP) in the electroencephalogram (EEG) recorded right after the occurrence of an error. Several studies show the presence of ErrP in typical choice reaction tasks. However, in the context of a BCI, the central question is: "Are ErrP also elicited when the error is made by the interface during the recognition of the subject's intent?" We have thus explored whether ErrP also follow a feedback indicating incorrect responses of the simulated BCI interface. Five healthy volunteer subjects participated in a new human–robot interaction experiment, which seem to confirm the previously reported presence of a new kind of ErrP. However, in order to exploit these ErrP, we need to detect them in each single trial using a short window following the feedback associated to the response of the BCI. We have achieved an average recognition rate of correct and erroneous single trials of 83.5% and 79.2%, respectively, using a classifier built with data recorded up to three months earlier.

Index Terms—Anterior cingulate cortex (ACC), brain–computer interface (BCI), electroencephalogram (EEG), error-related potentials (ErrP), inverse models, presupplementary motor area (pre-SMA), single-trial classification.

I. INTRODUCTION

BRAIN–COMPUTER INTERFACEs (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 [1]. 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 [2]–[5]. Since the late 1980s, different physiological studies have shown the presence of error-related potentials (ErrP) immediately following the occurrence of a perceived error [6]–[8]. Apart from Schalk *et al.* who investigated ErrP in real BCI feedback, most of these studies show the presence of ErrP in typical choice reaction tasks [3], [4], [6], [7]. 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/she should have responded with the right hand). The main components here are a negative potential showing up 80 ms after the incorrect response followed by a larger positive peak showing up between 200 and 500 ms after the incorrect response. 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 indicates incorrect performance [8]. The main component here is a negative deflection observed 250 ms after presentation of the feedback indicating incorrect performance. Finally, other studies reported the presence of ErrP (that we will refer to as "observation ErrP") following observation of errors made by an operator during choice reaction tasks [9], where the operator needs to respond to stimuli. As in the feedback ErrP, the main component here is a negative potential showing up 250 ms after the incorrect response of the operator performing the task. ErrP are most probably generated in a brain area called anterior cingulate cortex (ACC), which is crucial for regulating emotional responses [8]. An important aspect of the first two described ErrP is that they always follow an error made by the subject himself. First, the subject makes 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 us imagine that the subject's intent is to make a robot reach a target to the left. What would happen if the interface failed to recognize the intended command and the robot started turning in the wrong direction? Are ErrP still present even though the subject did not make any error but only perceived that the interface was performing wrongly?

The objective of this study 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 simulated BCI interface and no
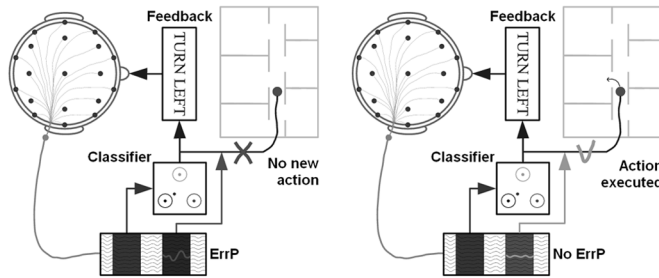
Fig. 1. Exploiting ErrP in a brain-controlled mobile robot. The subject receives visual feedback indicating the output of the classifier before the actual execution of the associated command (e.g., "turn left"). If the feedback generates an ErrP (left), this command is simply ignored and the robot will stay executing the previous command. Otherwise, the command is sent to the robot (right).
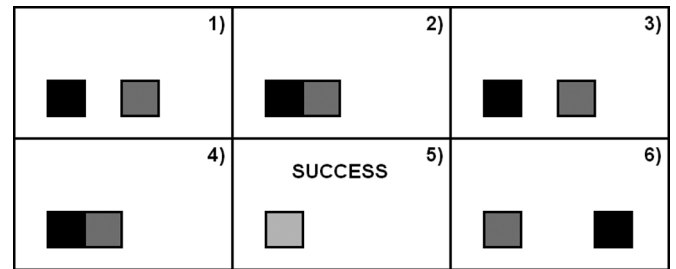

Fig. 2. Illustration of the protocol. (1) Target appears two steps on the left-hand side of the cursor (green). (2) Subject pressed the left key and the cursor moves one step to the left. (3) Subject still pressed the left key, but the system moves the cursor in the wrong direction. (4) Correct move to the left, compensating the error. (5) Cursor reaches the target. (6) New target appears three steps on the right-hand side of the cursor.

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 Fig. 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 [2], 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 new experimental results with five volunteer subjects during a simple human–robot interaction that confirm the previously reported existence of a new kind of ErrP [5], which is satisfactorily recognized in single trials. The new protocol does not exhibit the two main characteristics of an oddball paradigm, namely, there is no habituation of the subject to one of the stimuli and, in case of an error rate of 50%, there is no longer a frequent and infrequent stimulus. Furthermore, using a well-known inverse model called sLORETA [10] that noninvasively estimates the intracranial activity from scalp electroencephalogram (EEG), we show that the main focus of activity at the occurrence of ErrP seems to be located in the presupplementary motor area (pre-SMA) and in the ACC, as expected.

II. 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 human–robot interaction task where the subject has to bring the robot to targets 2 or 3 steps either to the left or to the right. This virtual interaction is implemented by means of a green square cursor that can appear on any of 20 positions along a horizontal line. The goal with this protocol is to bring the cursor to a target that randomly appears either on the left (blue square) or on the right (red square) of the cursor. The target is no further away than three positions from the cursor (symbolizing the current position of the robot). Contrarily to the protocol used by Ferrez *et al.* [5] and Buttfield *et al.* [12], this new protocol is more realistic and engaging and prevents the subject from habituation to one of the stimuli since the cursor reaches the target within a small number of steps. Fig. 2 illustrates the protocol with the target (black) initially positioned two steps away on the left-hand side of the cursor (gray). An error occurred at step 3) so that the cursor reaches

the target in five steps. 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 because of the nonoptimal 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/she simply presses a left or right key with the left or right hand. In this way, any error feedback is only because of a wrong recognition of the interface of which is the subject's intention. Five volunteer healthy subjects participated in these experiments. After the presentation of the target, the subject pressed the corresponding key until the cursor reached the target. The system moved the cursor 1 s after the key press (to avoid movement-related artifacts) with an error rate of 20%; i.e., at each step, there was a 20% probability that the cursor moved in the opposite direction. After the cursor reached the target, the word "success" appeared above the target, and a new target was randomly selected by the system. If the cursor did not reach the target after ten steps, the word "failed" appeared under the target and a new target was selected. Subjects performed ten sessions of 3 min on two different days, corresponding to ∼75 single trials per session. The delay between the two days of measurements was about three months.

EEG potentials were acquired with a portable system (Biosemi ActiveTwo) by means of a cap with 64 integrated electrodes covering the whole scalp uniformly. The sampling rate was 512 Hz and signals were measured at full direct current (dc). Raw EEG potentials were first spatially filtered by subtracting from each electrode the average potential (over the 64 channels) at each time step. The aim of this rereferencing 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 electrode. Then, we applied a 1–10-Hz bandpass filter as ErrP are known to be a relatively slow cortical potential [11]. Finally, EEG signals were subsampled from 512 to 64 Hz (i.e., we took one point out of eight) before classification, which was entirely based on temporal features. Indeed, the actual input vector for the statistical classifier described below is a half-second window starting 150 ms after the feedback and ending 650 ms after the feedback for channels "FCz" and "Cz." The choice of these channels follows the fact that ErrP are characterized by a fronto–central distribution along the midline. Thus, the
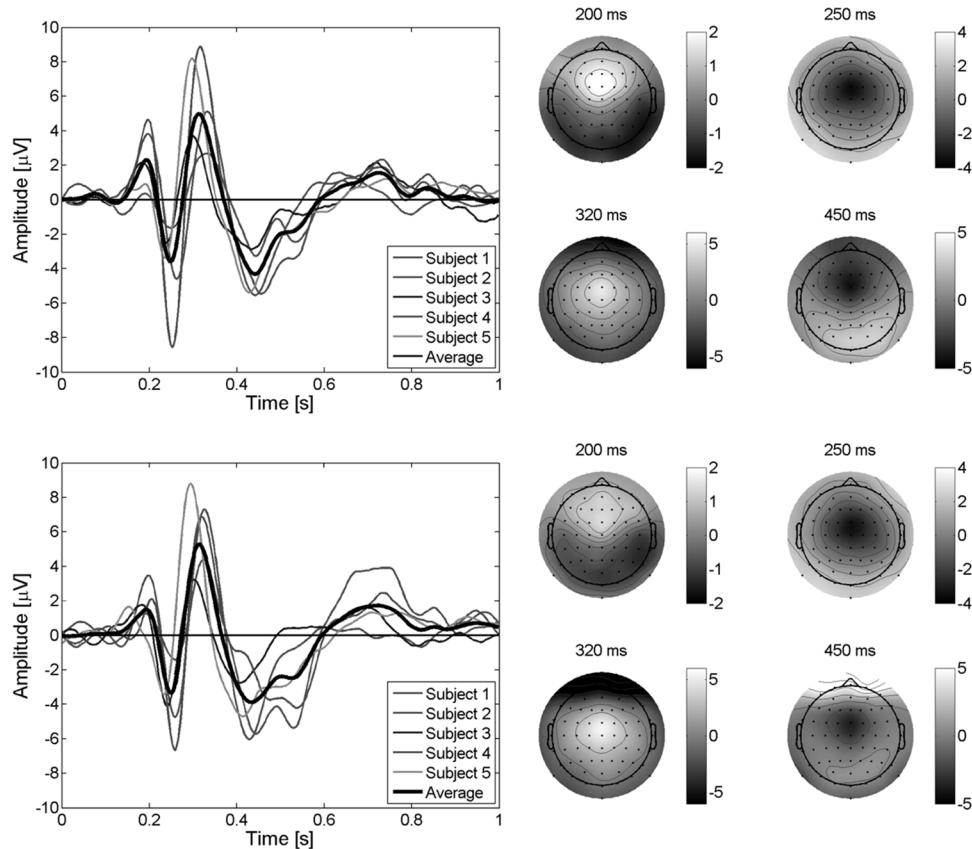
Fig. 3. (Left) Average EEG for the difference error–minus–correct at channel "FCz" for the five subjects plus the grand average of them for the first day (top) and for the second day (bottom) with an error rate of 20%. Feedback is delivered at time 0 s. A first positive peak shows up after 200 ms after the feedback. Negative and positive peaks show up about 250 and 320 ms after the feedback, respectively. A second broader negative peak occurs about 450 ms after the feedback. (Right) Scalp potential topographies, for the grand average EEG of the five subjects, at the occurrence of the peaks. Small filled circles indicate positions of the electrodes (frontal on top), "Cz" being in the middle.

dimensionality of the input vector is 64; i.e., concatenation of two windows of 32 points (EEG potentials) each. The two different classes are recognized by a Gaussian classifier trained to classify single trials as "correct" or "error" as in [5]. The output of the 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." In this statistical classifier, every Gaussian unit represents a prototype of one of the classes to be recognized, and we use several prototypes per class. During learning, the centers of the classes of the Gaussian units are pulled towards the trials of the class they represent and pushed away from the trials of the other class. 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 (see also Section III-E).

As mentioned, this protocol is more realistic and engaging and avoids some drawbacks. Furthermore, this study is a new one and is more complete than the one reported in [5] and [12] even if we performed the same kind of investigations to validate our new protocol [5]–[12]. Indeed, we have a new graphical interface, more subjects took part in the experiments, and we now record EEG with 64 electrodes. This gives us access to

critical electrode "FCz" and allow us to successfully use inverse models to estimate intracranial activity.

III. EXPERIMENTAL RESULTS

A. Error-Related Potentials

With this protocol, it is first necessary to check whether or not ErrP are evoked also by errors of the interface following a correct subject response. Fig. 3 shows the difference error–minus–correct for channel "FCz" for the five subjects plus the grand average of the five subjects for the two days of recordings. A first positive peak shows up 200 ms after the feedback. Negative and positive peaks can be seen 250 and 320 ms after the feedback, respectively. Finally, a second broader negative peak occurs about 450 ms after the feedback. Fig. 3 also shows the scalp potentials topographies, for the grand average EEG of the five subjects, at the occurrence of the four previously described peaks: a first fronto–central positivity appears after 200 ms, followed by a fronto–central negativity at 250 ms, a fronto–central positivity at 320 ms, and a fronto–central negativity at 450 ms. All five subjects show very similar ErrP time courses whose amplitudes slightly differ from one subject to the other. It is important to note here that, even if the delay between the two days of recordings was about three months, potentials as well as scalp topographies
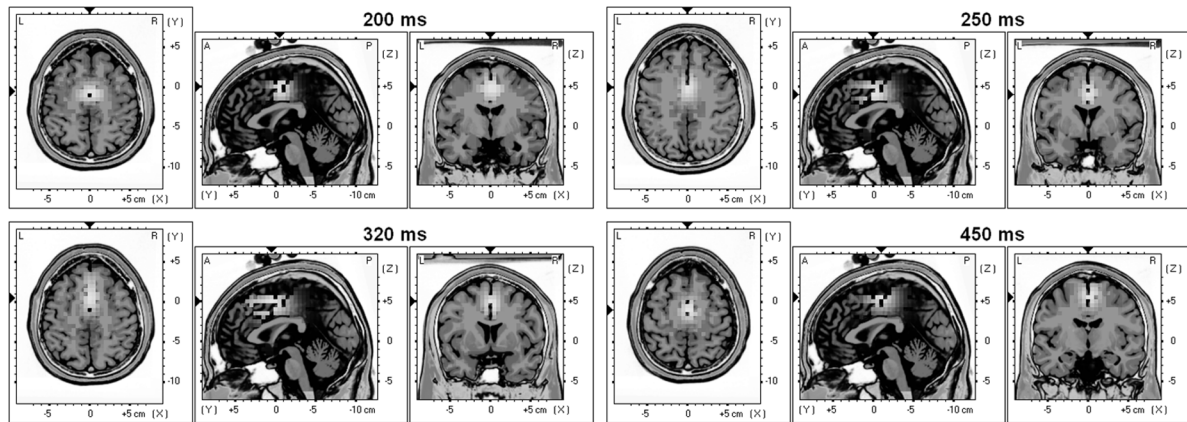
Fig. 4.   Talairach slices of localized activity for the grand average of the five subjects for the first day of recording at the occurrence of the four peaks described in Section III-A. Supplementary motor cortex and ACC are systematically activated. Furthermore, for the second positive peak (320 ms) and the second negative peak (450 ms) parietal areas are also activated. This parietal activation could reflect the fact that the subject is aware of the error.

are very similar for both days of recordings. These experiments seem to confirm the existence of a new kind of ErrP 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 and to observation ErrP. As in the case of response ErrP, interaction ErrP exhibit a first sharp negative peak followed by a broader positive peak. However, interaction ErrP are also characterized by a second negative peak that does not appear in response ErrP. Interaction ErrP are clearly different from both feedback and observation ErrP. Indeed, feedback and observation ErrP are only characterized by a small negative deflection 250 ms after the feedback and no other components are reported. We report three other clear components and, furthermore in Fig. 3, it appears that the 250-ms component is not the main ErrP component. Some subjects only show a small 250-ms negative component whereas the 320-ms positive peak and the 450-ms negative peak have large amplitudes.

B. Estimation of Intracranial Activity

Estimating the neuronal sources that generate a given potential map at the scalp surface (EEG) requires the solution of the so-called inverse problem. This inverse problem is always initially undetermined, i.e., there is no unique solution since a given potential map at the surface can be generated by many different intracranial activity maps. The inverse problem requires supplementary *a priori* constraints in order to be univocally solved. The ultimate goal is to unmix the signals measured at the scalp and to attribute to each brain area its own estimated temporal activity. The sLORETA inverse model [10] is a standardized low-resolution brain electromagnetic tomography. This software was used as a localization tool to estimate the focus of intracranial activity at the occurrence of the four ErrP peaks described in Section III-A. Fig. 4 shows Talairach slices of localized activity for the grand average of the five subjects for the first day of recording at the occurrence of the four peaks. As expected, the areas involved in error processing, namely, the pre-SMA (Brodmann area 6) and the rostral cingulate zone (RCZ; Brodmann areas 24 and 32) are systematically activated [8], [13]. Indeed, for the first positive peak (200 ms

after the feedback), the focus of activity is located at $X = 0$ mm, $Y = -5$ mm, and $Z = 50$ mm. The best match is Brodmann area 24 (ventral anterior cingulate cortex). For the first negative peak, (250 ms after the feedback), the focus is at $X = 0$ mm, $Y = 0$ mm, and $Z = 40$ mm and the best match is again Brodmann area 24. For the second positive peak (320 ms after the feedback), the focus is at $X = 0$ mm, $Y = 5$ mm, $Z = 50$ mm and the best match is in this case Brodmann area 32 (dorsal anterior cingulate cortex). Finally, for the second negative peak (450 ms after the feedback), the focus is on Brodmann area 6 (pre-SMA) at $X = 0$ mm, $Y = -10$ mm, and $Z = 55$ mm. For the second positive peak (320 ms) and the second negative peak (450 ms), parietal areas starts to be also activated. These associative areas (somatosensory association cortex; Brodmann areas 5 and 7) could be related to the fact that the subject becomes aware of the error. It has been proposed that the positive peak generated in a reaction task was associated with conscious error recognition [14]. In our case, activation of parietal areas about 300 ms after the feedback agrees with this hypothesis.

C. Single-Trial Classification

To explore the feasibility of detecting single-trial erroneous responses, we have done a tenfold 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. Table I reports the recognition rates (mean and standard deviations) for the five subjects plus the average of them for both days of recordings. This table also shows the recognition rates when classifying data of the second day using a classifier built with all data of the first day. The different hyperparameters—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 of the cross validations show that single-trial recognition of erroneous and correct responses is above 82% for both days. More importantly, classification using a classifier built with data recorded up to three months earlier show similar results (79.3% and 83.5%

TABLE I
PERCENTAGES (MEAN AND STANDARD DEVIATIONS) OF CORRECTLY RECOGNIZED ERROR TRIALS (E) AND CORRECT TRIALS (C) FOR THE FIVE SUBJECTS AND THE AVERAGE OF THEM FOR BOTH DAYS PERFORMING A TENFOLD CROSS VALIDATION AND FOR THE SECOND DAY, USING A CLASSIFIER BUILT WITH DATA OF THE FIRST DAY

|	Subject 1	Subject 2	Subject 3	Subject 4	Subject 5	**Average**
	10-fold cross-validation Day I					
E	80.5 ± 11.6	76.9 ± 13.7	74.3 ± 17.5	89.5 ± 9.6	89.0 ± 9.4	**82.0 ± 6.9**
C	83.4 ± 4.3	69.3 ± 8.6	78.8 ± 8.0	93.1 ± 4.6	87.3 ± 5.2	**82.4 ± 9.0**
	10-fold cross-validation Day II					
E	82.4 ± 10.4	81.2 ± 13.4	83.8 ± 13.1	85.0 ± 12.6	86.7 ± 13.8	**83.8 ± 2.2**
C	85.4 ± 6.7	72.6 ± 5.4	85.3 ± 3.5	86.4 ± 6.4	88.6 ± 4.9	**83.7 ± 6.3**
	Day II classified with Day I					
E	83.4 ± 10.7	69.5 ± 18.8	73.5 ± 9.4	86.0 ± 8.7	84.0 ± 11.8	**79.3 ± 7.3**
C	85.8 ± 4.8	76.2 ± 6.9	80.2 ± 7.3	85.5 ± 6.6	89.7 ± 4.5	**83.5 ± 5.3**
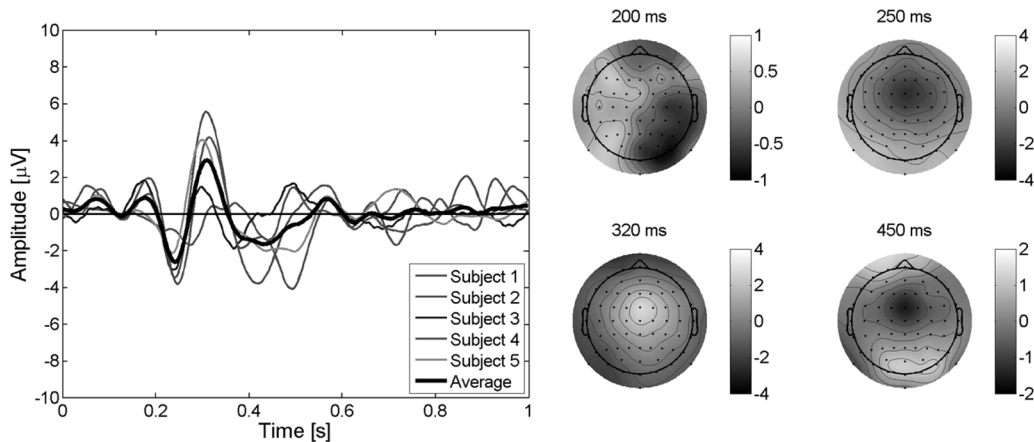
Fig. 5. (Left) Average EEG for the difference error–minus–correct at channel "FCz" for the five subjects plus the grand average of them with an error rate of 50%. Feedback is delivered at time 0 s. The same potentials as in Fig. 3 can be clearly seen, only the amplitudes are slightly smaller. (Right) Scalp potentials topographies, for the grand average EEG of the five subjects, at the occurrence of the peaks.

for error and correct trials, respectively). This stability is in accordance with the stability of the potentials described previously. Beside the crucial importance to integrate ErrP in the BCI in a way that the subject still feels comfortable, for example, by reducing as much as possible the rejection of actually correct commands, a 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. Indeed, as previously shown [5], the integration of error potentials detection theoretically leads to an increase of more than 70% of performance expressed in terms of bit rate.

D. ErrP and Oddball N200 and P300

Because our protocol is quite similar to an oddball paradigm, the question arises whether the potentials we describe are simply oddball N200 and P300. An oddball paradigm is characterized by an infrequent or especially significant stimulus interspersed with frequent stimuli. The subject is used to a certain stimulus and the occurrence of an infrequent stimulus generates a negative deflection (N200) about 200 ms after the stimulus followed by a positive peak (P300) about 300 ms after the stimulus. Our protocol may look similar to an oddball paradigm in the sense that the cursor usually moves in the correct direction, and a step in the wrong direction is the infrequent stimulus. To check this issue, we have run a series of experiments where the interface executed the subject's command with an error rate of 50% and, so, error trials are no longer less frequent than correct trials.

Analysis of the ErrP for the five subjects at channel "FCz" using error rates of 20% and 50% show no difference between them except that the amplitude of the potentials are smaller in the case of an error rate of 50%, but the time course remains the same as shown in Fig. 5. This is in agreement with all previous findings on ErrP that show that the amplitude is directly proportional to the error rate. We can conclude then that, while we cannot exclude the possibility that N200 and P300 contributes to the potentials in the case of an error rate of 20%, the oddballs N200 and P300 are not sufficient to explain the reported potentials.

E. Ocular Artifacts

In the reported experiments, subjects look at the cursor, awaiting its movement after they pressed the key corresponding to the target. After the feedback, the subject gets aware of the correct or erroneous response and he/she will shift gaze to the new position of the cursor, so that there is a gaze shift in every single trial. Nevertheless, it is possible that the subject looks at the target rather than the cursor. Consequently, there will be a gaze shift only after erroneous trials or, in any case, a larger gaze shift in erroneous trials. The statistical classifier could, therefore, pick those gaze shifts since several prototypes per class were used. To demonstrate that there is no systematical influence of gaze shifts on the presented ErrP as well as on classification results, we have calculated the different averages of the single trials with respect to the side of the target: left
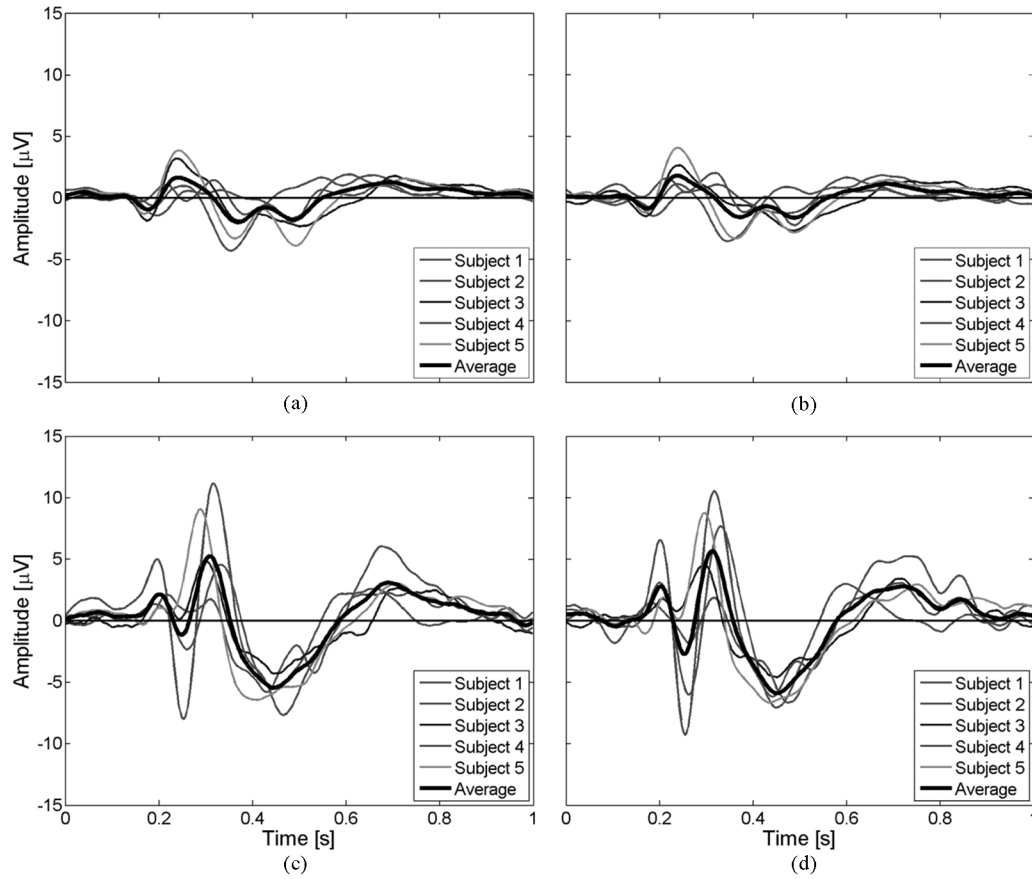
Fig. 6. Averages of the single trials at channel "FCz" with respect to the side where the target appeared for the five subjects and the average of them. There are four cases: (a) correct trials when the target appeared on the left-hand side, (b) correct trials with the right-hand side, (c) erroneous trials with the left-hand side, and (d) erroneous trials with the right-hand side. The left and right correct averages as well as the left and right erroneous averages are very similar whereas the left correct and erroneous as well as the right erroneous and correct are very different. This probably excludes any artifacts due to gaze shifts.

correct, right correct, left error, right error. Fig. 6 shows these four averages at channel FCz. Fig. 6(a) shows the average of correct single trials when the target appeared on the left for the five subjects and the average of them. Fig. 6(b) shows the average of correct single trials with respect to the right-hand side. Fig. 6(c) and (d) shows the average of erroneous trials with respect to the left- and right-hand sides, respectively. The left and right correct averages as well as the left and right erroneous averages are very similar whereas the left correct and erroneous as well as the right correct and erroneous are very different; so it appears that there is no systematical influence of gaze shifts on the reported potentials. Eye blinks are another potential source of artifacts. Indeed, it is conceivable that subjects may blink more frequently after one of the two conditions, and so the classifier could partly rely on eye blinks to discriminate error and correct trials. However, the scalp topographies of Fig. 3 show that the three ErrP components do not have a frontal focus, which would be expected in blink related potentials. Furthermore, we checked the different prototypes used for classification, and all of them have the shape of the ErrP shown in Fig. 3. It seems clear that the frontal foci present in some scalp topographies are generated by different phenomena than those responsible for ErrP and because we only use channels "FCz"

and "Cz" for classification, eye blinks do not systematically influence the reported results.

IV. CONCLUSION

In this paper, we have reported results on the detection of the neural correlate of error awareness for improving the performance and reliability of BCI. In particular, we have confirmed the existence of a new kind of ErrP elicited in reaction to an erroneous recognition of the subject's intention. An important difference between response ErrP, feedback ErrP, and observation ErrP on one side and the reported interaction ErrP on the other side is that the former involve a stimulus from the system for every single trial whereas the latter involve a choice of a long-term goal made by the subject himself/herself (where he/she wants to bring the cursor). More importantly, we have shown the feasibility of detecting single-trial erroneous responses of the interface and we have shown the stability of these potentials over time.

However, the introduction of an automatic response rejection strongly interferes with the BCI. The user needs to process additional information which induces higher workload and may considerably slow down the interaction. These issues will be investigated when running online BCI experiments integrating automatic error detection. 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.

In this paper, we have also shown that, as expected, typical cortical areas involved in error processing such as pre-SMA and ACC are systematically activated at the occurrence of the different peaks. The software used for the estimation of the intracranial activity (sLORETA) is only a localization tool. However, Babiloni et al. [15], [16] have recently developed the so-called cortical current density (CCD) inverse model that estimates the activity of the cortical mantle. Since ErrP seem to be generated by cortical areas, we plan to use this method to best discriminate erroneous and correct responses of the interface. In this respect, it is a normal practice in machine learning to apply feature selection techniques to reduce the dimensionality of the input vector to the classifier, especially, when the input space has a huge dimensionality and many of the features are not relevant for discriminating among the classes. This is particularly the case for the analysis of EEG signals and we have previously shown the advantages of such a feature selection for BCI [17]. In this case, we will apply them for selecting the most relevant solution points from the few thousands provided by the CCD model. In fact, the very preliminary results using the CCD inverse model confirm the reported localization in the pre-SMA and in the ACC and thus we may well expect a significant improvement in recognition rates by focusing on the dipoles estimated in those specific brain areas.

More generally, the work described here 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 alarm, fatigue, 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.

REFERENCES

[1] J. Wolpaw, H. Ramoser, D. McFarland, and G. Pfurtscheller, "EEG-based communication: Improved accuracy by response verification," *IEEE Trans. Rehabil. Eng.*, vol. 6, no. 3, pp. 326–333, Sep. 1998.

[2] G. Schalk, J. Wolpaw, D. McFarland, and G. Pfurtscheller, "EEG-based communication: Presence of an error potential," *Clin. Neurophysiol.*, vol. 111, pp. 2138–2144, 2000.

[3] 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 Syst. Rehabil. Eng.*, vol. 11, no. 2, pp. 127–131, Jun. 2003.

[4] 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 Syst. Rehabil. Eng.*, vol. 11, no. 2, pp. 173–177, Jun. 2003.

[5] P. Ferrez and J. del R. Millán, "You are wrong!—Automatic detection of interaction errors from brain waves," in *Proc. 19th Int. Joint Conf. Artif. Intell.*, 2005, pp. 1413–1418.

[6] C. Carter, T. Braver, D. Barch, M. B. , D. Noll, and J. Cohen, "Anterior cingulate cortex, error detection, and the online monitoring of performance," *Science*, vol. 280, pp. 747–749, 1998.

[7] M. Falkenstein, J. Hoormann, S. Christ, and J. Hohnsbein, "ERP components on reaction errors and their functional significance: A tutorial," *Biol. Psychol.*, 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," *Psychol. Rev.*, vol. 109, pp. 679–709, 2002.

[9] H. van Schie, R. Mars, M. Coles, and H. Bekkering, "Modulation of activity in medial frontal and motor cortices during error observation," *Nature Neurosci.*, vol. 7, pp. 549–554, 2004.

[10] R. Pascual-Marqui, "Standardized low resolution brain electromagnetic tomography (sLORETA): Technical details," *Methods Find. Exp. Clin. Pharmacol.*, vol. 24D, pp. 5–12, 2002.

[11] W. Gehring, B. Goss, M. Coles, D. Meyer, and E. Donchin, "A neural system for error detection and compensation," *Psychol. Sci.*, vol. 4, pp. 385–390, 1993.

[12] A. Buttfield, P. Ferrez, and J. del R. Millán, "Towards a robust BCI: Error potentials and online learning," *IEEE Trans. Neural Syst. Rehabil. Eng.*, vol. 14, no. 2, pp. 164–168, Jun. 2006.

[13] K. Fiehler, M. Ullsperger, and Y. von Cramon, "Neural correlates of error detection and error correction: Is there a common neuroanatomical substrate?," *Eur. J. Neurosci.*, vol. 19, pp. 3081–3087, 2004.

[14] S. Nieuwenhuis, K. Ridderinkhof, J. Blom, G. Band, and A. Kok, "Error-related brain potentials are differently related to awareness of response errors: Evidence from an antisaccade task," *Psychophysiology*, vol. 38, pp. 752–760, 2001.

[15] F. Babiloni, C. Babiloni, L. Locche, F. Cincotti, P. Rossini, and F. Carducci, "High-resolution EEG: Source estimates of Laplacian-transformed somatosensory-evoked potentials using a realistic subject head model constructed from magnetic resonance images," *Med. Biol. Eng. Comput.*, vol. 38, pp. 512–519, 2000.

[16] F. Babiloni, F. Cincotti, C. Babiloni, F. Carducci, D. Mattia, L. Astolfi, A. Basilico, P. Rossini, L. Ding, Y. Ni, J. Cheng, K. Christine, J. Sweeney, and B. He, "Estimation of the cortical functional connectivity with the multimodal integration of high-resolution EEG and fMRI data directed by transfer function," *NeuroImage*, vol. 24, pp. 118–131, 2005.

[17] J. del R. Millán, M. Franzé, J. Mouriño, F. Cincotti, and F. Babiloni, "Relevant EEG features for the classification of spontaneous motor-related tasks," *Biol. Cybern.*, vol. 86, pp. 89–95, 2002.

Pierre W. Ferrez received the engineering degree in materials science from the Swiss Federal Institute of Technology, Lausanne, Switzerland, in 2001. Currently, he is working towards the Ph.D. degree at the IDIAP Research Institute, Martigny, Switzerland.

He joined the brain–computer interface research team of Prof. J. del R. Millán at the IDIAP Research Institute, in 2004 as a Research Assistant. He has been a Research Assistant at the Institute for Research in Ophthalmology (IRO), Sion, Switzerland. His research interests include biological signals acquisition and processing and noninvasive brain–computer interfaces, with a focus on the integration of high-level cognitive states such as error-related potentials in EEG-based brain–computer interfaces.

José del R. Millán received the Ph.D. degree in computer science from the Universita Politècnica de Catalunya, Barcelona, Spain, in 1992.

Currently, he is a Senior Researcher at the IDIAP Research Institute, Martigny, Switzerland, where he explores the use of brain signals for multimodal interaction and, in particular, the development of noninvasive brain-controlled robots and neuroprostheses. He is also an Adjunct Professor at the Swiss Federal Institute of Technology (EPFL), Lausanne, Switzerland. He was an Assistant Professor for three years at the Universita Politècnica de Catalunya, a Research Scientist at the Joint Research Centre of the European Commission, Ispra, Italy, a Visiting Professor at the EPFL, and a Visiting Scholar at Stanford University, Stanford, CA.

Dr. del R. Millán was named "Research Leader 2004" by the *Scientific American* for his work on brain-controlled robots. He was finalist of the European Descartes Prize 2001 for his research on brain–computer interfaces.