To appear in the journal *Artificial Intelligence*

Brain-Actuated Interaction

José del R. Millán[a,b,*], Frédéric Renkens[b], Josep Mouriño[c], Wulfram Gerstner[b]

[a] *IDIAP Research Institute, Rue du Simplon 4, 1920 Martigny, Switzerland*

[b] *Lab of Computational Neuroscience, School of Computer and Communication Science & Brain and Mind Institute, Swiss Federal Institute of Technology,1015 Lausanne EPFL, Switzerland*

[c] *Centre de Recerca en Enginyeria Biomèdica, Universitat Politècnica de Catalunya, 08028 Barcelona, Spain*

Abstract

Over the last years evidence has accumulated that shows the possibility to analyze human brain activity on-line and translate brain states into actions such as selecting a letter from a virtual keyboard or moving a robotics device. These initial results have been obtained with either invasive approaches (requiring surgical implantation of electrodes) or synchronous protocols (where brain signals are time-locked to external cues). In this paper we describe a portable noninvasive brain-computer interface that allows the continuous control of a mobile robot in a house-like environment and also the operation of a virtual keyboard. The interface works asynchronously (the person makes self-paced decisions on when to switch from one mental task to the next) and uses 8 surface electrodes to measure electroencephalogram signals from which a statistical classifier recognizes 3 different mental states. Here we report results with five volunteers during their brain-actuated interaction experiments with the mobile robot and the virtual keyboard. Two of the participants successfully moved the robot between several rooms, while the other three participants managed to write messages with the virtual keyboard. One of the latter volunteers is a physically impaired person suffering from spinal muscular atrophy.

Keywords: brain-computer interface, electroencephalogram signals, asynchronous protocol, statistical classifier, robotics, communication devices.

[*] Corresponding author.

E-mail addresses: jose.millan@idiap.ch (J. del R. Millán)

1. Introduction

There is a growing interest in the use of brain signals for communication and operation of devices, in particular for physically disabled people. Brain states can be detected and translated into actions such as selecting a letter from a virtual keyboard or moving a robot arm [4],[5],[7],[9],[12],[15], [19],[28],[29],[30],[32],[39],[40]. Such devices, which do not require the user to perform any physical action, are called *brain-computer interfaces (BCI)* or *brain-machine interfaces* (for reviews see [17],[23],[37],[38]). It is worth noting that, although BCI prototypes have only been developed recently, the basic ideas were already put forward in the 1970s. Initial successful experiments were based on the analysis of brain electrical activity—the visual evoked potential—generated in response to changes in gaze direction [34] (see also [16],[31]).

A BCI may monitor brain activity via a variety of methods, which can be coarsely classified as invasive and non-invasive. Most non-invasive BCI systems use electroencephalogram (EEG) signals; i.e., the electrical brain activity recorded from electrodes placed on the scalp. The main source of the EEG is the synchronous activity of thousands of cortical neurons. Measuring the EEG is a simple noninvasive way to monitor electrical brain activity, but it does not provide detailed information on the activity of single neurons (or small brain areas). Moreover, it is characterized by small signal amplitudes (a few μVolts) and noisy measurements (especially if recording outside shield rooms).

Besides electrical activity, neural activity also produces other types of signals, such as magnetic and metabolic, that could be used in a BCI. Magnetic fields can be recorded with magnetoencephalography (MEG), while brain metabolic activity—reflected in changes in blood flow—can be observed with positron emission tomography (PET), functional magnetic resonance imaging (fMRI), and optical imaging. Unfortunately, such alternative techniques require sophisticated devices that can be operated only in special facilities. Moreover, techniques for measuring blood flow have long latencies and thus are less appropriate for interaction.

In invasive BCI systems the activity of single neurons (their spiking rate) is recorded from microelectrodes implanted in the brain. In a series of experiments with rats and monkeys, researchers have monitored different areas of the cortex related to execution and planning of movements—motor, premotor and posterior parietal cortex. From a real-time analysis of the activity of the neuronal population, it has been possible to determine the animal's movement intention [7],[15], predict the monkey's hand trajectory [32],[36], and to drive a computer cursor to desired targets [30],[32]. In human patients, first steps towards invasive approaches have been made [12]. One of the patients was eventually able to drive a cursor and write messages. His performance, however, was similar to that achieved with noninvasive BCI systems [17].

Given the risks generated by permanent surgically implanted devices in the brain, and the associated ethical concerns, we concentrate only on non-invasive approaches, in particular electrical brain signals as measured by EEG. For certain stimuli, such as flashed images and lights, the EEG exhibits a strong characteristic signal, the so-called evoked potential, which reflects the immediate responses of the brain to those external stimuli. Evoked potentials are, in principle, easy to pick up with scalp electrodes and have been used in the context of BCIs [9],[16],[31]. The necessity of external stimulation does, however, restrict the applicability of evoked potentials to a limited range of tasks. In our view, a more natural and suitable alternative for interaction is to analyze components associated with spontaneous “intentional” mental activity. Thus, some researchers measure slow cortical potentials—whose negative amplitudes are related to the overall preparatory excitation level of a given cortical network—over the top of the scalp [4]. Other groups look at local variations of EEG rhythms. A particularly significant EEG rhythm can be recorded from the central region of the scalp overlying the sensorimotor cortex during the imagination of body movements [5],[27],[28],[39],[40]. But, in addition to motor-related rhythms, other cognitive mental tasks are being explored [19],[29] as a number of neurocognitive studies have found that different mental tasks—such as mental rotation of geometric figures [41], arithmetic operations [8], or language [26]—activate local cortical areas to a different extent.

Current EEG-based BCIs are limited by a low channel capacity and are considered too slow for controlling rapid and complex sequences of movements. So far control tasks based on human EEG have been limited to simple exercises such as moving a computer cursor to the corners of the screen [39] or opening a hand orthosis [28]. In this paper we show that human volunteers could, within a few days, learn to master a portable EEG-based brain-computer interface that recognized three mental states. Two participants successfully moved a robot between several rooms by mental control only. Furthermore, mental control was only marginally worse than manual control on the same task. As a second demonstration of brain-actuated interaction, we describe a communication tool that enables people to select letters from a virtual keyboard and write messages. We report results with three volunteers who have used this virtual keyboard, one advanced and two beginners. The advanced user has irregularly worked with the virtual keyboard for several months, while the beginners did their training in a single day. One of the latter participants is a physically impaired person suffering from spinal muscular atrophy.

2. Brain Interface Protocol

EEG-based brain-computer interfaces are limited by a low channel capacity. Most of the current systems have a channel capacity below 0.5 bits/s [38]. One of the main reasons for such a low bandwidth is that they are based on synchronous protocols where EEG is time-locked to externally paced cues repeated every 4-10 s and the response of the BCI is the average decision over this period [4],[28],[29],[39],[40]. Such synchronous protocols facilitate EEG analysis since the starting time of mental states are precisely known and differences with respect to background EEG activity can be amplified. Unfortunately, they are slow and BCI systems that use them normally recognize only 2 mental states, independently of the number of electrodes from which EEG is measured. In contrast, our approach and a few other systems [5] use an asynchronous protocol that analyzes the ongoing EEG to determine the person's mental state, which they can voluntarily change at any moment without waiting for external cues. The rapid responses of this asynchronous BCI—0.5 s in our case—, together with its performance (see Section 3), give a theoretical channel capacity between 1 and 1.5 bits/s.

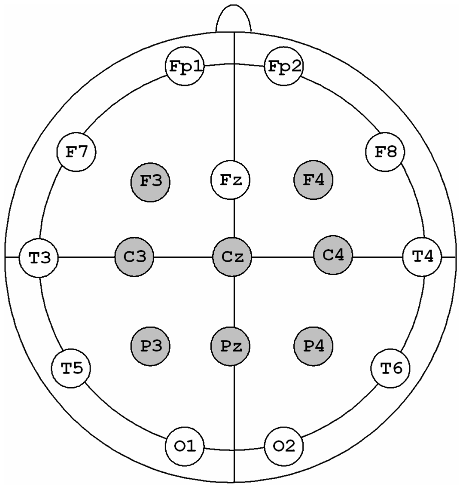

Fig. 1. Electrode configuration. Electrodes are placed according to the standard International 10-20 system. The EEG system acquires brain signals from the 8 fronto-central-parietal electrodes indicated in grey. All the signals are recorded with respect to a linked-ear reference (average potentials measured in both ear lobes).

EEG potentials were recorded at the 8 standard fronto-centro-parietal locations F3, F4, C3, Cz, C4, P3, Pz, and P4 (see Fig. 1). Participants, all of them volunteers, wore a commercial EEG cap with integrated electrodes (small circles in Fig. 3). The sampling rate was 128 Hz. Raw EGG potentials are too noisy and variable to be analyzed directly. Thus the first step is to preprocess them to increase their signal-to-noise ratio and extract relevant features that better describe the mental states to be recognized. The raw EEG potentials were first transformed by means of a surface Laplacian (SL). Specifically, we first interpolated using spherical splines of order 2 and then took the second spatial derivative which is sensitive to localized sources of electrical activity [24],[25]. The second derivative is evaluated only at the 8 locations of the electrodes. Normally, the SL is estimated with a high number of electrodes. But Babiloni et al. [2] have shown that, for the operation of a BCI, SL waveforms with either a low or a high number of electrodes give statistically similar classification results. Mouriño et al. [22] compare different ways to compute the SL based on a few electrodes. The superiority of SL-transformed over raw potentials for the operation of BCI has been demonstrated in different studies [3],[14],[20].

The features we extract from the 8 SL-transformed electrode signals are based on a temporal Fourier transform. To estimate the power spectrum of each channel over the last second we used the Welch periodogram algorithm. Specifically, we averaged the FFT of 3 segments of 0.5 second with 50% overlap, which yields a frequency resolution of 2 Hz. The values in the frequency band 8-30 Hz were normalized according to the total energy in that band. Thus an EEG sample has 96 features (8 channels times 12 components each). EEG samples were computed every 62.5 ms (i.e., 16 times per second). It is worth noting that, for our experimental protocol, periodogram features lead to better or similar performances than more elaborated features such as parameters of autoregressive models and wavelets [33].

Fig. 2 shows two samples of 1s of raw EEG signals measured at locations C3 and C4, and their corresponding periodograms. Both samples were taken at an interval of 3 s while the subject "B" was continuously performing the mental task "left" (see below). This figure illustrates two points. Firstly, despite the variability of the raw EEG signals, their periodograms are rather stable. Secondly, despite the fact that channels C3 and C4 are located over the right and left motor areas of the brain, respectively, their temporal activity is very similar for imagined movements of one single hand (left, in this case). This indicates that the spatial resolution of EEG signals in this type of task is rather low. Hence, the discriminative power of EEG is not proportional to the number of electrodes.

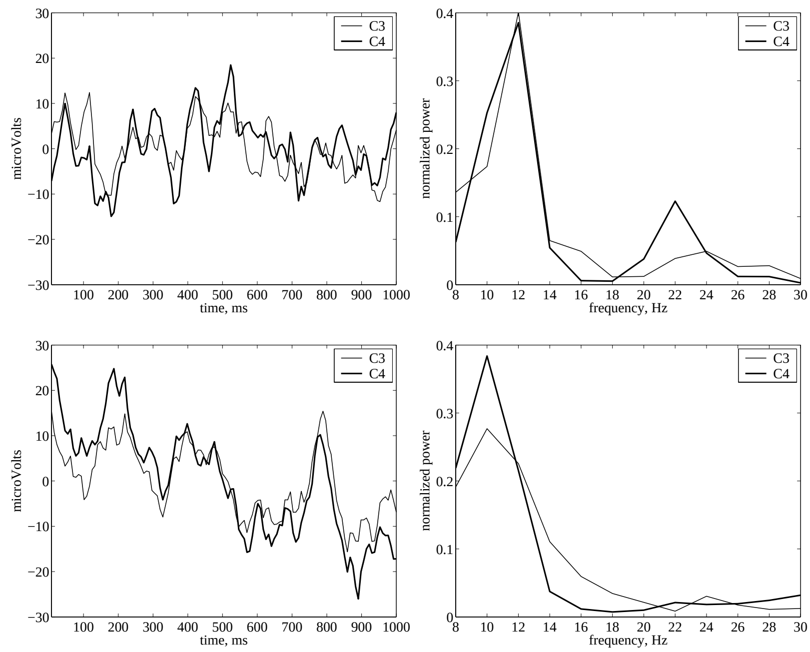

Fig. 2. Examples of two typical raw EEG signals measured at locations C3 and C4, and the periodograms obtained after preprocessing. The ensemble of the periodograms (power in bins of 2 Hz) for all the channels is the input to the statistical classifier. *Left:* raw EEG signals. *Right:* periodograms.

In this paper we report results obtained with 5 volunteers. Four of them are able-bodied ("A", "B", "C", and "D"), and the fifth suffers from spinal muscular atrophy (volunteer "E"). This is an abnormal bone and muscle development impairment, and currently he only has a very limited control of his right arm. Participants "A" and "B" controlled the mobile robot, while participants "C", "D", and "E" communicated through the virtual keyboard. Volunteers "A" and "C" were directly involved in the research reported in this paper. The other three volunteers did not have any link to the project: "B" was a student at the Swiss Federal Institute of Technology in Lausanne (Switzerland), "D" was a visitor during an exhibition in Nice (France), and "E" lives in London (UK). No participant received any kind of reward (monetary or other) for their participation, but all were highly motivated.

In order to interact with the brain-actuated mobile robot or virtual keyboard, participants underwent an initial training period, ranging from a couple of hours (participants "D" and "E") to a few days (participants "A", "B", and "C"), where they learned to control 3 mental tasks of their choice. Participants tried the following mental tasks: "relax", imagination of "left" and "right" hand (or arm) movements, "cube rotation", "subtraction", and "word association". The tasks consisted of getting relaxed, imagining repetitive self-paced movements of the limb, visualizing a spinning cube, performing successive elementary subtractions by a fixed number (e.g., 64–3=61, 61–3=58, etc.), and generating words that begin with the same letter. Three of our volunteers ("C", "D", "E") did "relax" with eyes closed[1], while volunteers "A" and "B" (who controlled the mobile robot) did "relax" with eyes open. All other tasks were performed with eyes open. After a short evaluation, participant "B" chose to work with the combination of 3 tasks relax-left-right, while the other four participants selected the tasks relax-left-cube. In the sequel, we will refer to these mental tasks as #1, #2 and #3 (i.e., relax is #1, left is #2, and right or cube is #3). Neither participant had previous experience with a BCI or mental training.

Each day, volunteers participated in four consecutive training sessions (except participant "E" who did seven training sessions) of about 5 min, separated by breaks of 5-10 min. During each training session participants switched randomly every 10-15 s between the three tasks. Participants received feedback online through three colored buttons on a computer screen. Each button is associated to one of the mental tasks to be recognized. A button flashed when an EEG sample is classified as belonging to the corresponding mental task[2]. After each training session the statistical classifier was optimized offline. After this initial training, participants learned to control mentally either the mobile robot (volunteers "A" and "B") or the virtual keyboard (volunteers "C", "D", and "E"). During this training period, the user and the BCI engaged in a mutual learning process where they were coupled and adapted to each other.

3. Statistical Classifier

The mental tasks (or classes) are recognized by a Gaussian classifier trained to discriminate EEG samples as state #1, #2, #3 or "unknown". This family of classifiers has been shown to perform better than support vector machines and temporal-processing neural networks [11], as well as committees of multilayer perceptrons, learning vector quantization, and incremental radial basis networks [19].

---

[1] Even for volunteers who shut their eyes, the recognition of the task "relax" is *not* based on the detection of eye lid movements.

[2] Since "relax" can be performed with eyes closed, participants received auditory feedback indicating the recognition of this mental task.

In this statistical classifier, every Gaussian unit represents a prototype of one of the classes to be recognized. Its output gives an estimation of the posterior class probability distribution for an EEG sample. The challenge is to find the appropriate centers and widths of the Gaussians in the high-dimensional input space described above to differentiate the classes.

Although mixtures of Gaussians are well known classifiers [6], our implementation differs from classical ones in a few respects. We assume that the class-conditional 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, namely 4. In addition, we assume that all four prototypes have an equal weight of 1/4. Then, dropping constant terms, the posterior probability y_k of class C_k for sample x is

$$y_k(x) = \frac{\sum_{i=1}^{N_k} a_k^i(x)}{\sum_{j=1}^{N_c} \sum_{i=1}^{N_j} a_j^i(x)} \tag{1}$$

where N_c is the number of classes and a_k^i is the activation level of the ith prototype of the class C_k

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

where μ_k^i corresponds to the center of the ith prototype of class C_k, Σ_k is the covariance matrix of class C_k, and $|\Sigma_k|$ is the determinant of that matrix. In our case, Σ_k is taken to be diagonal and common to all the prototypes of the class. In this way, we reduce the number of parameters and increase the accuracy of their estimation.

The response of the network for sample x is the class C_k with the highest posterior probability provided that is greater than a given probability threshold; otherwise the response is “unknown.” This rejection criterion keeps the number of errors (false positives) low, because recovering from erroneous actions (e.g., robot turning in the wrong direction or deleting a letter) has a high cost. The choice of this probability threshold was guided by a previous ROC study where a different set of volunteers only carried out the initial training described before [11], and the actual value was selected based on the performance of the present volunteers during the initial period of training. In particular, the probability threshold was 0.85 for participants “A” and “B”, 0.90 for participant “C”, and 0.75 for participants “D” and “E”.

To initialize the center of the prototypes and the covariance matrix of the class C_k we run a clustering algorithm (typically, a self-organizing map [13]) to compute the position of the desired number of prototypes. Then, the initial value of the covariance matrix is

$$\Sigma_k = \tfrac{1}{S_k} \sum_{n=1}^{S_k} (x^n - \mu_k^{i(n)})(x^n - \mu_k^{i(n)})^T \tag{3}$$

where S_k denotes the number of training samples belonging to the class C_k and *i(n)* is the nearest prototype of this class to the sample x^n.

Since the centers μ_k^i can be interpreted as prototype units, we move them, during training, towards the EEG samples of the mental task they represent. Moreover, to avoid interference the prototype units are pushed away from EEG samples of other tasks. Specifically, for every sample x in the training set, the update rule for all the prototypes of all the classes is

$$\Delta\mu_k^i(x) = \alpha\big(t_k(x) - y_k(x)\big)\ \Sigma_k^{-1}(x - \mu_k^i)\frac{a_k^i(x)}{A(x)} \tag{4}$$

where α is the learning rate, t_k is the kth component of the target vector in the form *1-of-c* and A is the total activity of the network—i.e., the denominator in (1).

Finally, after every iteration over the training set, we estimate again the new value of Σ_k using expression (3). It is possible to estimate the covariance matrices in more elaborated ways, including through gradient descent in order to minimize their contribution to the error function, but we did not do so in this paper.

The brain-computer interface responds every 0.5 s. Firstly, it computes the class-conditioned probability for each class—i.e., the mixture of Gaussians in the numerator of Eq. (1). Secondly, it averages the class-conditioned probabilities over 8 consecutive samples[3]. Thirdly, it estimates the posterior probability based on the average class-conditioned probability of each class using Bayes' formula; cf Eq. (1). Finally, it compares the posterior probability with the confidence threshold. At the end of training, errors and "unknown" responses are below 5% and 30%, respectively. These are online performances obtained on a new session using the classifier trained with data of previous sessions. The theoretical channel capacity of the interface is hence above 1bit/second (operation mode I). In addition, the interface could also operate in another mode (operation mode II) where classification errors are further reduced by requiring that two (or three) consecutive periods of 0.5 s give the same classification response. In this mode II errors and "unknown" responses are below 2% and 40%, respectively, and the theoretical channel capacity is about 1 bit/second (two consecutive responses) or 0.6 bits/second (three consecutive responses).

4. Brain-Actuated Control of a Mobile Robot

The task was to drive the robot through different rooms in a house-like environment (see Fig. 3). The robot was a small Khepera (5.7 cm diameter) that closely mimics a motorized wheelchair. The robot moved at a constant speed of one third of its diameter per second, similar to the speed of a

[3] The goal of the preprocessing procedure (surface Laplacian followed by Welch periodogram) is to improve the signal-to-noise ratio of EEG measurements. Nevertheless, since the Welch is over short time windows, the EEG samples fed to the classifier are still highly variable, hence this averaging step.

wheelchair in an office building.

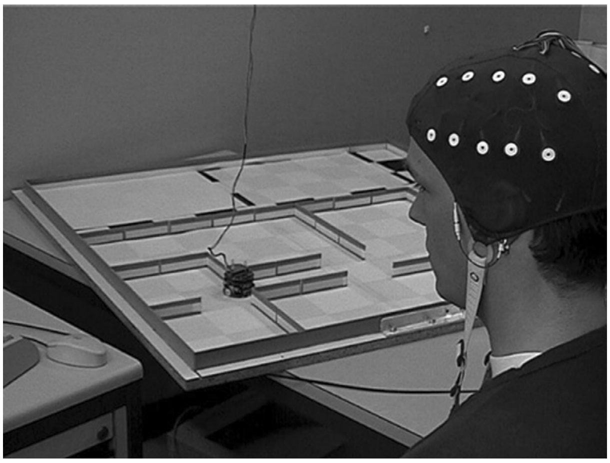

Fig. 3. One of the particpants while driving mentally the robot through the different rooms of the environment during the first experiment.

4.1. Robot Setup and Control

To make the robot move along a desired trajectory it is necessary to determine the speed of the motors controlling the wheels at each time step. Obviously, this is impossible by means of just three mental commands. A key idea is that the user's mental states are associated to high-level commands (e.g., "turn right at the next occasion") that the robot executes autonomously using the readings of its on-board sensors. Another critical aspect for the continuous control of the robot is that users can issue high-level commands at any moment. This is possible because the operation of the BCI is asynchronous and does not require waiting for external cues, unlike synchronous approaches. The robot will continue executing a high-level command until the next is received.

The robot relies on a behavior-based controller [1] to implement the high-level commands that guarantees obstacle avoidance and smooth turns. In this kind of controller, on-board sensors are read constantly and determine the next action to take. The mapping from the user's mental states (or commands) to the robot's behaviors is not simply one-to-one, but, in order to achieve a more flexible control of the robot, the mental states are just one of the inputs for a finite state automaton with 6 states (or behaviors). The transitions between behaviors are determined by the 3 mental states (#1, #2, #3), 6 perceptual states of the environment (as described by the robot's sensory readings: left wall, right wall, wall or obstacle in front, left obstacle, right obstacle, and free space) and a few internal memory variables. Fig. 4 shows a simplified version of the finite state automaton. The memory variables were required to implement correctly the different behaviors. Thus, if the robot is performing the behavior "forward" and perceives a wall to the left, it switches automatically to the behavior "follow left wall". The transition to the behavior "forward" is necessary, for example, in the case the robot is approaching an open door and the user wants the robot not to enter into the room. On the other hand, the robot "stops" whenever it perceived an obstacle in front to avoid collisions (not all the transitions to the behavior "stop" appear in the figure for the sake of simplicity). Briefly, the interpretation of a mental state depends on the perceptual state of the robot. Thus, in an open space the mental state #2 means "left turn" while the same mental state is interpreted as "follow left wall" if a wall is detected on the left-hand side. Similarly, mental state #3 means "right turn" or "follow right wall"; mental state #1 always implied "move forward". Altogether participants felt that our control schema was simple and intuitive to use because each mental state was associated with the same category of movement (for example, category "left") and the perceptual state determined the subcategory (for example, the category "left" has the subcategories "turn left" or "follow left wall").

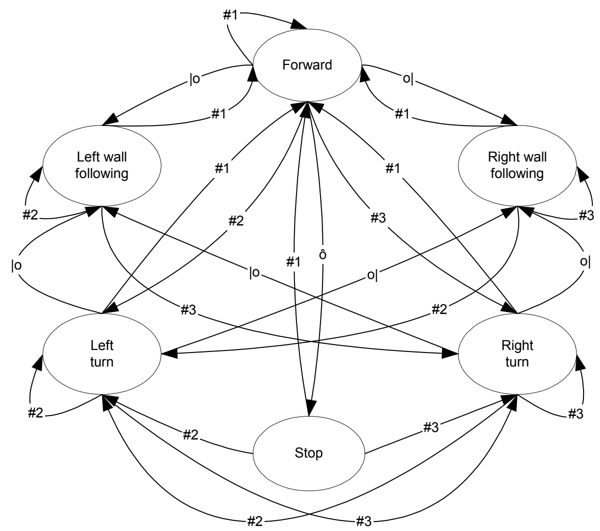

Fig. 4. Finite state automaton used for the control of the robot. Transitions between the 6 behaviors were determined by 3 mental states (#1, #2, #3), 6 perceptual states (|o: left wall, o|: right wall, ô: wall or obstacle in front), and some memory variables. The memory variables and some of the perceptual states are not shown for the sake of simplicity.

The Khepera robot is a two-wheeled vehicle. It has 8 infrared sensors around its diameter to detect obstacles. The sensors have a limited perception range, which makes the recognition of the different perceptual states from the raw readings difficult. To overcome this limitation, the robot uses a multilayer perceptron that maps the 8 raw infrared sensory readings into the current perceptual state. This neural classifier is trained offline by gradient descent [6] to estimate the probability distribution of a sensory sample to belong to each of the 6 perceptual states. Training samples (500 for each of the perceptual states) were generated by randomly placing and orienting the robot in a known environment.

A final element is the use of an appropriate feedback indicating the current mental state recognized by the embedded classifier. This is done by means of three lights on top of the robot, with the same colors as the buttons used during the training phase. The front light is green and is on when the robot receives the mental state #1. The left light is blue and is associated to the mental state #2, whereas the right light is red and is associated to the mental state #3. Thus, if the robot is following the left wall and is approaching an open door, a blue feedback indicates that the robot will turn left to continue following the left wall (and, so, it will enter into the room). On the contrary, a green feedback indicates that robot will move forward along the corridor when facing the doorway and will not enter into the room. This simple feedback allows users to correct rapidly the robot's trajectory in case of errors in the recognition of the mental states or errors in the execution of the desired behavior (due to the limitations of the robot's sensors).

4.2. Experimental Results with the Mobile Robot

After 5 and 3 days of initial training with the interface operating in mode I, participants "A" and "B" achieved a satisfactory level of performance (correct recognition was above 65% while errors were below 7%—the remaining were "unknown" responses). Volunteers participated in 4 training sessions per day. Fig. 5 shows the performance curves of both participants for the same number of training sessions. Note that the very first training session was used to gather the initial EEG samples to train the statistical classifiers and so users did not receive any feedback at this time. In the case of participant "A", there was a break of several weeks between the first 2 days of training and the remaining three. Fig. 5 shows the performance of participant "A" in the first day (training sessions 2 to 4), the third day (training sessions 5 to 8), and the fifth day (training sessions 9 to 12). This break explains the significant decrease in performance from session 4 to session 5.

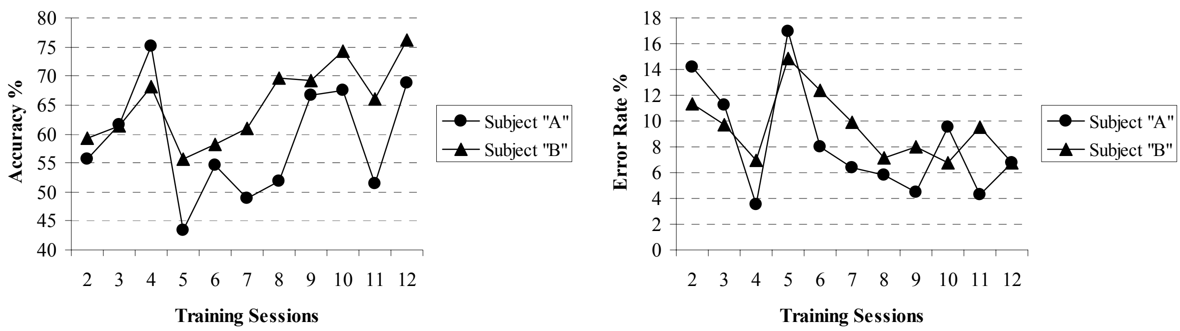

Fig. 5. Performance curves for participants "A" and "B" over the different training sessions performed along three days (four training sessions per day). The left panel shows the accuracy (percentage of true positives) and the right panel shows the error rates (percentage of false positives). The remaining percentage for each training session corresponds to "unknown" responses. During the first training session users did not receive feedback.

Apart from this observation, both participants exhibit common trends in their performance curves. First, a clear improvement can be observed during the first day, with an excellent final performance. Second, the performance degrades at the beginning of the second day but recovers at the end. This shows the difficulty in generalizing from one day to the next due to both the natural variability of brain signals, the refitting of the EEG cap, and the fact that the user's ability to operate the BCI is not yet stable [18]. Third, the performance in the third day is more stable, although it decreases in one of the sessions to recover later. A similar trend has been observed with other volunteers in a previous study [19]. Furthermore, for these latter volunteers a geometrical study [21] has shown that the prototypes learned during training evolved so as to increase the distance between the mental classes and to reduce the interclass variance, while at the same time it decreases the distance between the prototypes of a same class between consecutive sessions. The first two facts are correlated with the improvement in performance at the end of training, while the latter fact supports that the user and her personal BCI are approaching stability.

Fig. 6 shows the performance during the second and last training session of the two participants for different numbers of prototypes per mental task. The probability threshold was 0.85 in all the cases and each classifier was trained with EEG samples of the previous training sessions. A McNemar statistical test shows that the performance of the classifiers with 4 prototypes is significantly better or equally good than classifiers with any other number of prototypes. In particular, using more than 4 prototypes usually leads to overfitting at the beginning of training, probably due to more variability of the EEG samples which decreases by the end of training.

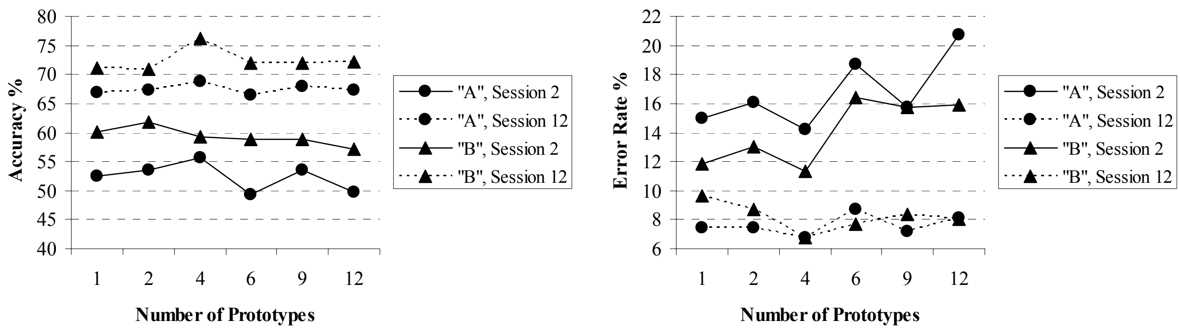

Fig. 6. Performance of the participants “A” and “B” using different number of prototypes per mental task for the second and last sessions of the initial training phase. The left panel shows the accuracy, while the right panel shows the error rates. The remaining percentage corresponds to “unknown” responses.

Table 1. Performance of the two participants during their second training session according to the number of responses averaged over 0.5 seconds. Performance is measured in terms of accuracy (acc) and errors (err). The remaining percentage corresponds to “unknown” responses. The probability threshold was 0.85 in all the cases. Each classifier was trained with EEG samples of only the first training session.

		Number of Averaged Responses				
Participant		1	2	4	8	16
“A”	*acc*	62.6%	57.0%	58.8%	55.7%	55.0%
	err	27.1%	17.2%	23.2%	14.2%	15.1%
“B”	*acc*	62.2%	62.2%	60.4%	59.3%	57.8%
	err	31.7%	21.9%	13.4%	11.3%	12.6%

Finally, Table 1 shows how performance varies according to the number of consecutive responses averaged over 1/2 second for participants "A" and "B" during their second training session. This means that, if we average the responses to n EEG samples, each is obtained every $1/(2*n)$ seconds from the power spectrum over the last second. As before, the probability threshold was 0.85 in all the cases and each classifier was trained with EEG samples of only the first training session. It can be observed that averaging less than eight responses leads to significantly higher errors, while averaging more than eight responses does not yield improvements. The reason for this latter case is that the temporal shifts between consecutive EEG samples are so short that the samples are very similar.

After this initial training, participants learned to control mentally the mobile robot for 2 days with the brain interface operating in mode II. During this second period of training participants had to drive the robot mentally from a starting position to a first target room; once the robot arrived, a second target room was selected and so on. The starting position and the target rooms were drawn at random. The results reported here were obtained at the end of the second day of work with the robot.

Fig. 7 shows a trajectory generated by participant "A" after two days of training. The robot had to visit 3 different rooms, drawn randomly, starting from location "S". Although the figure does not show the details of the trajectory inside the rooms, the robot made a short exploration in each of them. During the experiment, the participant was driving the robot for about 10 minutes continuously. He brought the robot to the desired room each time, but there were a few occasions where the robot did not follow the optimal trajectory. This was mainly because the brain interface took a longer time than usual to recognize the participant's mental state. For instance, in one case the robot missed a turn because the brain interface did not recognize the appropriate mental state until the robot had passed the doorway of the desired room, and so the participant needed to maneuver mentally the robot to bring it back to the doorway. In other situations, the robot's sensors perceived a wall or corner as too close, thus making the robot stop automatically to avoid collisions. In these situations, the participant needed to turn (by mental control) the robot away from the phantom obstacle and then resume the trajectory.

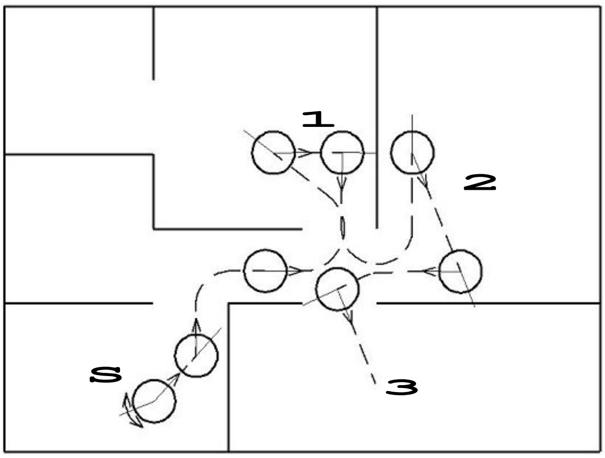

Fig. 7. Trajectory followed by the robot under the mental control of participant "A" during one of the trials of the first experiment. The robot started in the bottom left room and then visited 3 other rooms, top center, top right and bottom right, sequentially. The figure does not show the details of the trajectory inside the rooms.

Qualitatively, the trajectory is rather good as the robot visited the 4 rooms in the desired order and it was never necessary to make significant corrections to the robot's active behaviors. But in order to evaluate *quantitatively* the performance of the brain-actuated robot, participants "A" and "B" also carried out a second set of experiments in a slightly different arrangement of the rooms that were now located along the two sides of a corridor.

In a given trial, the robot must travel from a starting room to a target room as well as also visiting an intermediate room. The rooms and their order were selected at random. First, the participant made the robot visit the desired sequence of rooms by mental control. In a later session, the participant drove the robot along the same sequence of rooms by manual control. In this case, the participant used the same controller as described above but, instead of sending mental commands to the robot, he simply pressed one of three keys. This procedure allowed us to compare mental and manual control for a system that is identical in all other aspects. In addition, the manual trajectory should be quite close to the optimal path that can be generated with the current controller. It is worth mentioning that the reason why the participant controls the robot mentally first and only afterwards manually is to avoid any learning process that could facilitate mental control.

Table 2 gives the time in seconds necessary to generate the desired trajectory for three different trials for the two participants. For each trial, the table indicates the time required for mental control and manual control. Remarkably, we can see that mental control was comparable to manual control. On average, brain-actuated control of the robot is only 35% longer than manual control for both participants. This relationship would be worse for mental control if the speed of the robot were faster because the likelihood of failing to deliver the right command at the right moment would increase. As mentioned previously, the speed of the robot chosen for the experiments was similar to that of a wheelchair in an indoor environment.

Table 2. Time in seconds for three different trials where participants "A" and "B" controlled the robot first mentally and then manually.

Participant	Trial	Mental	Manual
"A"	1	149	124
	2	183	135
	3	191	129
	Average	*174*	*129*
"B"	1	219	156
	2	189	155
	3	175	117
	Average	*194*	*143*

5. Brain-Actuated Communication with a Virtual Keyboard

The brain interface can also be used to select letters from a virtual keyboard on a computer screen and write a message (see Fig. 8). Initially, the whole keyboard (26 English letters plus the space to separate words, for a total of 27 symbols organized in a matrix of 3 rows by 9 columns) is divided in three blocks, each associated to one of the mental tasks. The association between blocks and mental tasks is indicated by the same colors as during the training phase. Each block contains an equal number of symbols, namely 9 at this first level (3 rows by 3 columns). Then, once the statistical classifier recognizes the block on which the user is concentrating, this block is split in 3 smaller blocks, each having 3 symbols this time (1 row). As one of this second-level blocks is selected (the classifier recognizes the corresponding mental task), it is again split in 3 parts. At this third and final level, each block contains 1 single symbol. Finally, to select the desired symbol, the user concentrates in its associated mental task as indicated by the color of the symbol. This symbol goes to the message and the whole process starts over again. Thus, the process of writing a single letter requires three decision steps.

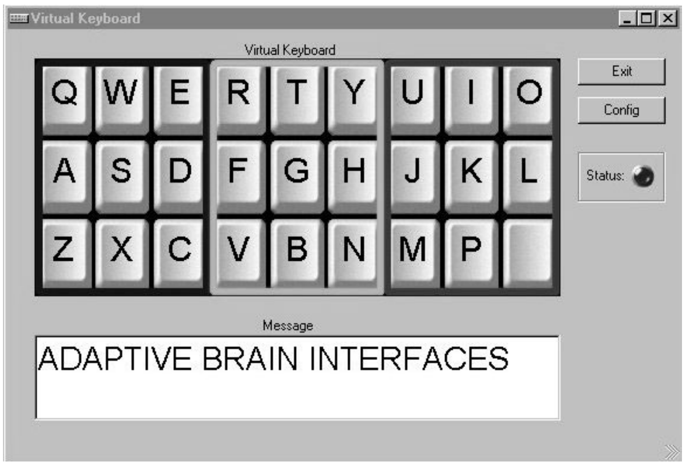

Fig. 8. Virtual keyboard. Initially, the whole keyboard is divided in three blocks, each associated to one of the mental tasks. Then, as the statistical classifier recognizes a mental task, the corresponding block is split in smaller blocks until a letter a selected. This letter goes to the message (at the bottom of the screen) and the whole process starts over again. The association between blocks and mental tasks is indicated by the same colors as during the training phase. The left block is blue and is associated to the mental task #2, the central block is green and is associated to the mental task #1, and the right block is red and is associated to the mental task #3.

The actual selection of a block incorporates some additional reliability measures (in addition to the statistical rejection criteria) for the purpose of increasing the likelihood of correct functioning, especially during the different public demonstrations we have given (including live TV programs). In particular, a part of the keyboard is only selected when the corresponding mental task is recognized three times in a row (mode II). Also, in the case of an eventual wrong selection, the user can undo it by concentrating immediately on one of the mental tasks of his/her choice. Thus the system waits a short time after every selection (3.5 seconds) before going down to next level. The mental task used to undo selection is that for which the user exhibits the best performance.

For the advanced participant "C", it took 22.0 seconds on average to select a letter. This time includes recovering from eventual errors. Despite all these additional checks before the actual selection of a block, this time compares favorably with some other noninvasive brain-actuated virtual keyboards requiring 1 minute [28] or 2 minutes [4] per letter, and it is similar to an invasive approach allowing a patient to write 3 letters per minute [12]. In some trials where those additional checks were relaxed (requiring that two consecutive responses were the same and eliminating the waiting period, but adding an additional symbol for undoing the last selection), participant "C" could write letters at an average speed of 7.0 s. It goes without saying that the incorporation of statistical language models [35], or other techniques for word prediction such as T9 in cellular phones, will eventually speed up writing. Participant "C" was using the virtual keyboard for almost one year in an irregular way—there were periods of intense working, 3 or 4 days per week and several hours per day, and also breaks of several weeks—to write short texts (in between one and five words). Only exceptionally he made typing errors. Interestingly, once this participant reached a good performance at the beginning of his training, it was not necessary to modify the statistical classifier very often and he continued to use the same classifier for over six months in order to operate the virtual keyboard. The average performance of participant "C", when tested in operation mode I, was 70.5% correct recognition, 2.8% error rate, and the remaining 26.7% "unknown" responses.

Table 3. Performances of two participants over consecutive training sessions carried out in a single day. Performance is measured in terms of accuracy (acc) and error rates (err). The remaining percentage corresponds to "unknown" responses. During the first training session users did not receive feedback.

		Sessions					
Participant		2	3	4	5	6	7
"D"	*acc*	76.1%	55.4%	70.6%			
	err	4.9%	7.3%	7.2%			
"E"	*acc*	55.6%	62.0%	65.0%	59.9%	65.4%	65.9%
	err	13.1%	7.4%	9.2%	8.8%	11.8%	4.8%

In addition to this advanced user, two other volunteers learned to operate the virtual keyboard in a single day. Participant "D" is a female who did her training in an exhibition area. After an initial training of 4 sessions with the system operating in mode I, she was able to write a short text without errors. Participant "E" was trained at his home in London. He did 7 training sessions with the system operating in mode I, and afterwards he also succeeded in controlling the virtual keyboard. Unfortunately, we did not measure the writing speed of these participants, which was significantly lower than that of volunteer "C". Table 3 shows the performance of participants "D" and "E" over the different training sessions. Note that the first training session was used to gather the initial EEG samples to train the statistical classifiers and so users did not receive any feedback at this time.

6. Discussion

In this paper we have reported first results with five volunteers operating two brain-actuated applications, namely a mobile robot and a virtual keyboard, by means of a portable noninvasive brain-computer interface. Key elements of this BCI are its asynchronous protocol for the analysis of online EEG signals and the use of machine learning techniques for fitting the individual EEG features of each user. Also, the combination of advanced robotics with the previous two elements has made possible to control robots in indoor environments.

One may wonder how significant are our findings and how likely it would be for an untrained person to solve the two tasks described in this paper by mere chance. Both tasks are three-class problems with rejection thresholds. Hence random classification would yield at most 33.3% of correct responses, while, at the end of training with interface in mode I, all participants reached a performance above 65% of correct classifications during the last session of about 3000 samples. In order to assess statistical significance, we have split the session into 6 blocks with 500 samples each. The accuracy was 70.5±9.0%, which is clearly above the threshold of randomness (33.3%). It follows that it is rather unlikely to operate mentally the two applications purely by chance. For instance, in order to select a given letter, the interface has to make three selections. Hence chance level is $(1/3)^3$ per letter and writing a 5-letter word is $(1/3)^{15}$. Similarly, chance level is also close to zero in the case of controlling the robot since the user needs to deliver a minimum of 13 mental commands during a typical trajectory. Hence the chance of generating a correct trajectory is $(1/3)^{13}$.

Scaling up the proposed brain-actuated robot control system to more complex tasks could go in two directions. Firstly, we could keep the same small number of mental states and increase the complexity of the automaton, which chooses the robot behavior by a combination of the recognized mental state and the robot perceptual state. Such an approach would lead to a BCI that is less intuitive to use. Secondly, we could keep the complexity of the automaton and increase the number of mental states that the classifier can recognize. This is limited by the spatial resolution of EEG technology. Moreover, a high number of mental states would make it difficult to provide rapid visual feedback of the recognized mental state, as we have done in these experiments.

Despite these initial encouraging results, our BCI should be still improved to achieve faster writing times and future operation of wheelchairs. In particular, porting the current results to the wheelchair is not straightforward for, at least, two reasons. First, the performance of the BCI will suffer once the user is seated on a mobile platform. This will require longer training times for the user. Second, the current finite state automaton only allows for simple control actions, and so the resulting wheelchair could be too constrained for practical use in cluttered environments. In this respect, recent progress in EEG analysis [10] suggests that a sufficient number of mental states can be recognized to control robotics devices and in a more natural and flexible way, as well as to speed up communication with the virtual keyboard. In this approach we will transform scalp potentials—recorded with a sufficiently high number of electrodes (32, 64 or more)—to brain maps to get detailed information on the activity of small cortical areas. The Gaussian classifier embedded in the BCI would work upon selected parts of these brain maps instead of using EEG features. Prior information (derived for example from fMRI studies) could be included in the process.

Another key concern is to keep the BCI constantly tuned to its owner. This requirement arises because, as users gain experience, they develop new capabilities and change their EEG patterns. In addition, brain activity changes naturally over time. In particular, this is the case from one session (with which data the classifier is trained) to the next (where the classifier is applied). The challenge is to adapt on-line the classifier while the user operates a brain-actuated device. Initial results show the benefit of this approach for users at the very beginning of their training [18], an important issue to keep their motivation and to help them reach stable performances quickly.

7. Acknowledgements

Part of this work was done while JdRM and JM were at the Joint Research Centre of the European Commission in Ispra, Italy. This work has been supported in part by the ESPRIT Programme of the European Commission (LTR project number 28193). JdRM is currently supported by the Swiss National Science Foundation through the National Centre of Competence in Research on "Interactive Multimodal Information Management (IM2)."

8. References

[1] R.C. Arkin, Behavior-Based Robotics, MIT Press, Cambridge, Massachusetts, 1998.

[2] F. Babiloni, F. Cincotti, L. Bianchi, G. Pirri, J.d.R. Millán, J. Mouriño, S. Salinari, M.G. Marciani, Recognition of imagined hand movements with low resolution surface Laplacian and linear classifiers, Medical Engineering & Physics 23 (2001) 323–328.

[3] F. Babiloni, F. Cincotti, L. Lazzarini, J.d.R. Millán, J. Mouriño, M. Varsta, J. Heikkonen, L. Bianchi, M.G. Marciani, Linear classification of low-resolution EEG patterns produced by imagined hand movements, IEEE Trans. on Rehabilitation Engineering 8 (2000) 186–188.

[4] N. Birbaumer, N. Ghanayim, T. Hinterberger, I. Iversen, B. Kotchoubey, A. Kübler, J. Perelmouter, E. Taub, H. Flor, A spelling device for the paralysed, Nature 398 (1999) 297–298.

[5] G.E. Birch, Z. Bozorgzadeh, S.G. Mason, Initial on-line evaluation of the LF-ASD brain-computer interface with able-bodied and spinal-cord subjects using imagined voluntary motor potentials, IEEE Trans. on Neural Systems and Rehabilitation Engineering 10 (2002) 219–224.

[6] C.M. Bishop, Neural Networks for Pattern Recognition, Oxford University Press, Oxford, 1995.

[7] J.K. Chapin, K.A. Moxon, R.S. Markowitz, M.A.L. Nicolelis, Real-time control of a robot arm using

simultaneously recorded neurons in the motor cortex, Nature Neuroscience 2 (1999) 664–670.

[8] F. Chochon, L. Cohen, P.F. van de Moortele, S. Dehaene, Differential contributions of the left and right inferior parietal lobules to number processing, J. Cognitive Neuroscience 11 (1999) 617–630.

[9] L.A. Farwell, E. Donchin, Talking off the top of your head: Toward a mental prosthesis utilizing event related brain potentials, Electroencephalography and Clinical Neurophysiology 70 (1988) 510–523.

[10] R. Grave de Peralta, S.L. Gonzalez Andino, S. Morand, C.M. Michel, T. Landis, Imaging the electrical activity of the brain: ELECTRA, Human Brain Mapping 9 (2000) 1–12.

[11] A. Hauser, P.-E. Sottas, J.d.R. Millán, Temporal processing of brain activity for the recognition of EEG patterns, in: Proceedings 12th Int. Conf. on Artificial Neural Networks, 2002, pp. 1125–1130.

[12] P.R. Kennedy, R. Bakay, M.M. Moore, K. Adams, J. Goldwaithe, Direct control of a computer from the human central nervous system, IEEE Trans. on Rehabilitation Engineering 8 (2000) 198–202.

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

[14] D.J. McFarland, L.M. McCane, S.V. David, J.R. Wolpaw, Spatial filter selection for EEG-based communication, Electroencephalography and Clinical Neurophysiology 103 (1997) 386–394.

[15] D. Meeker, S. Cao, J.W. Burdick, R.A. Andersen, Rapid plasticity in the parietal reach region demonstrated with a brain-computer interface, in: Proceedings Society for Neuroscience Abstracts, 2002.

[16] M. Middendorf, G. McMillan, G. Calhoun, K.S. Jones, Brain-computer interfaces based on the steady-state visual-evoked response, IEEE Trans. on Rehabilitation Engineering 8 (2000) 211–214.

[17] J.d.R. Millán, Brain-computer interfaces, in: M.A. Arbib (Ed.), Handbook of Brain Theory and Neural Networks, MIT Press, Cambridge, Massachusetts, 2002.

[18] J.d.R. Millán, On the need for on-line learning in brain-computer interfaces, in: Proceedings Int. Joint Conf. on Neural Networks, 2004.

[19] J.d.R. Millán, J. Mouriño, M. Franzé, F. Cincotti, M. Varsta, J. Heikkonen, F. Babiloni, A local neural classifier for the recognition of EEG patterns associated to mental tasks, IEEE Trans. on Neural Networks 13 (2002) 678–686.

[20] J. Mouriño, EEG-based analysis for the design of adaptive brain interfaces, Ph.D. thesis, Centre de Recerca en Enginyeria Biomèdica, Universitat Politècnica de Catalunya, Barcelona, Spain, 2003.

[21] J. Mouriño, S. Chiappa, R. Jané, J.d.R. Millán, Evolution of the mental states operating a brain-computer interface, in: Proceedings 2nd European Medical and Biological Engineering Conf., 2002, pp. 400–401.

[22] J. Mouriño, J.d.R. Millán, F. Cincotti, S. Chiappa, R. Jané, F. Babiloni, Spatial filtering in the training process of a brain computer interface, in: Proceedings 23rd Annual Int. Conf. of the IEEE Engineering in Medicine and Biology Society, 2001, pp. 639–642.

[23] M.A.L. Nicolelis, Actions from thoughts, Nature 409 (2001) 403–407.

[24] F. Perrin, J. Pernier, O. Bertrand, J. Echallier, Spherical spline for potential and current density mapping, Electroencephalography and Clinical Neurophysiology 72 (1989) 184–187.

[25] F. Perrin, J. Pernier, O. Bertrand, J. Echallier, Corrigendum EEG 02274, Electroencephalography and Clinical Neurophysiology, 76 (1990) 565.

[26] S.E. Petersen, P.T. Fox, M.I. Posner, M. Mintun, M.E. Raichle, Positron emission tomographic studies of the cortical anatomy of single-word processing, Nature 331 (1988) 585–589.

[27] G. Pfurtscheller, C. Neuper, Motor imagery activates primary sensorimotor area in humans, Neuroscience

Letters 239 (1997) 65–68.

[28] G. Pfurtscheller, C. Neuper, Motor imagery and direct brain-computer communication, Proceedings of the IEEE 89 (2001) 1123–1134.

[29] S.J. Roberts, W.D. Penny, Real-time brain-computer interfacing: A preliminary study using Bayesian learning, Medical & Biological Engineering & Computing 38 (2000) 56–61.

[30] M.D. Serruya, N.G. Hatsopoulos, L. Paninski, M.R. Fellows, J. Donoghue, Instant neural control of a movement signal, Nature 416 (2002) 141–142.

[31] E.E. Sutter, The brain response interface: Communication through visually-induced electrical brain response, J. Microcomputing Applications 15 (1992) 31–45.

[32] D.M. Taylor, S.I. Helms Tillery, A.B. Schwartz, Direct cortical control of 3D neuroprosthetic devices, Science 296 (2002) 1829–1832.

[33] M. Varsta, J. Heikkonen, J.d.R. Millán, J. Mouriño, Evaluating the performance of three feature sets for brain-computer interfaces with an early stopping MLP committee, in: Proceedings 15th Int. Conf. on Pattern Recognition, 2000, pp. 911–915.

[34] J.J. Vidal, Real-time detection of brain events in EEG, Proceedings of the IEEE 65 (1977) 633–664.

[35] D.J. Ward, D.J.C. MacKay, Fast hands-free writing by gaze direction, Nature 418 (2002) 838.

[36] J. Wessberg, C.R. Stambaugh, J.D. Kralik, P.D. Beck, M. Laubach, J.K. Chapin, J. Kim, S.J. Biggs, M.A. Srinivassan, M.A.L. Nicolelis, Real-time prediction of hand trajectory by ensembles of cortical neurons in primates, Nature 408 (2000) 361–365.

[37] I. Wickelgren, Tapping the mind, Science 299 (2003) 496–499.

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

[39] J.R. Wolpaw, D.J. McFarland, Multichannel EEG-based brain-computer communication, Electroencephalography and Clinical Neurophysiology 90 (1994) 444–449.

[40] J.R. Wolpaw, D.J. McFarland, T.M. Vaughan, Brain-computer interface research at the Wadsworth center, IEEE Trans. on Rehabilitation Engineering 8 (2000) 222–226.

[41] A. Yoshino, M. Inoue, A. Suzuki, A topographic electrophysiologic study of mental rotation, Cognitive Brain Research 9 (2000) 121–124.