

Non-Invasive Brain-Actuated Control of a Mobile Robot by Human EEG

José del R. Millán^{a,b}, Frédéric Renkens^b, Josep Mouriño^c, and Wulfram Gerstner^b

^a*IDIAP Research Institute, 1920 Martigny, Switzerland*

^b*Laboratory of Computational Neuroscience, Swiss Federal Institute of Technology, 1015 Lausanne
EPFL, Switzerland*

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

***Abstract*—Brain activity recorded non-invasively is sufficient to control a mobile robot if advanced robotics is used in combination with asynchronous EEG analysis and machine learning techniques. Until now brain-actuated control has mainly relied on implanted electrodes, since EEG-based systems have been considered too slow for controlling rapid and complex sequences of movements. We show that two human subjects successfully moved a robot between several rooms by mental control only, using an EEG-based brain-machine interface that recognized three mental states. Mental control was comparable to manual control on the same task with a performance ratio of 0.74.**

***Index Terms*—Brain-machine interface, EEG, asynchronous protocol, robotics.**

I. INTRODUCTION

The idea of moving robots or prosthetic devices not by manual control, but by mere “thinking” (i.e., the brain activity of human subjects) has fascinated researchers over the last couple of years [1]-[9]. Initial demonstrations of the feasibility of such an approach have relied on intracranial electrodes implanted in the motor cortex of monkeys [1]-[5]. For humans, non-invasive methods based on electroencephalogram (EEG) signals are preferable, but they suffer from a reduced spatial resolution and increased noise due to measurements on the scalp. So far control tasks based on human EEG have been limited to simple exercises such as moving a computer cursor or opening a hand orthosis [6]-[9]. Here we show that the signals derived from an EEG-based brain-machine interface (BMI) are sufficient to continuously control a miniature mobile robot in an indoor environment with several rooms, corridor, and doorways (Fig. 1). Moreover, experimental results obtained with two volunteer healthy subjects show that brain-actuated control of the robot is nearly as efficient as manual control. The subjects achieved these results after a few days of training with a portable non-invasive BMI that uses 8 scalp electrodes.

Human EEG signals represent the global activity of millions of neurons. In standard clinical protocols, EEG signals are synchronized to an external cue and averaged over tens of trials in order to increase the signal-to-noise ratio and resolve spatial and temporal activation patterns. For the control of mechanical devices via an EEG-based BMI, averaging over several trials is not possible. Single-trial analysis (also called “online” analysis) is, however, typically limited by a low channel capacity below 0.5 bits/second [8], and so EEG-based BMIs have been considered too slow for controlling rapid and complex sequences of movements. Nevertheless, previous studies have succeeded in recognizing a few mental states that have been used for communication [8]-[10]. One of the main reasons for such a low bit rate is the use of synchronous protocols where EEG is time-locked to externally paced cues repeated every 4-10 s. In this paper, we use an asynchronous protocol and analyze the ongoing EEG activity to determine the subjects’ mental state which can change at any moment. This approach nearly doubles the usual bit rate of EEG-based brain-machine interfaces.

II. METHODS

How is it possible to control a robot that has to make accurate turns at precise moments in time using signals that arrive at a rate of about one bit per second? There are three key features of our approach. First, the user's mental states are associated with high-level commands (e.g., "turn right at the next occasion") and the robot executes these commands autonomously using the readings of its on-board sensors. Second, the subject can issue high-level commands at any moment. This is possible because the operation of the BMI is asynchronous and, unlike synchronous approaches, does not require waiting for external cues. The robot will continue executing a high-level command until the next is received. Third, the robot relies on a behavior-based controller [11] to implement the high-level commands that guarantee obstacle avoidance and smooth turns.

In our controller, both the user's mental state and the robot's perceptual state can be considered as inputs to a finite state automaton with 6 states (or behaviors). These behaviors are "forward movement", "left turn", "follow left wall", "right turn", "follow right wall", and "stop". Fig. 2 shows the essentials of this finite state automaton. The transitions between behaviors are determined by the 3 mental states (#1, #2, #3) of the user, supplemented by 4 perceptual states of the environment determined from the robot's sensory readings (left wall, right wall, wall or obstacle in front, and free space). In addition, the controller uses two other perceptual states (left obstacle and right obstacle) and a few internal memory variables for obstacle avoidance and stable implementation of the different behaviors. The robot's interpretation of a particular mental state depends on the perceptual state of the robot. Thus, in an open space, mental state #2 means "left turn"; on the other hand, if a wall is detected on the left-hand side, mental state #2 is interpreted as "follow left wall". Similarly, depending on the perceptual state of the robot, mental state #3 can mean "right turn" or "follow right wall". However, mental state #1 always means "move forward". Moreover, the robot "stops" whenever it perceived an obstacle in front to avoid collisions. Altogether experimental subjects felt that our control schema was simple and intuitive to use.

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 command #1. The left light is blue and is associated to the mental command #2, whereas the right light is red and is associated to the mental command #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 rapidly correct 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).

A. EEG Signals

Two volunteer healthy subjects "A" and "B" wore a commercial EEG cap with integrated scalp electrodes. EEG potentials, referenced to the average of the left and right ear lobes, were recorded at the 8 standard fronto-centro-parietal locations F3, F4, C3, Cz, C4, P3, Pz, and P4. The sampling rate was 128 Hz. The raw EEG potentials were first transformed by means of a surface Laplacian (SL) computed globally by means of a spherical spline of order 2 [12]-[14]. Every 62.5 ms the power spectrum in the band 8-30 Hz was estimated over the last second of data. To do so, we used Welch's periodogram algorithm on segments of 0.5 s and averaged the estimations for 3 segments with 50% overlap. This yields a frequency resolution of 2 Hz. Each 96-dimensional vector (8 channels times 12 frequency components in the band 8-30 Hz) was then normalized. The resulting EEG sample was analyzed by a statistical classifier. No artifact rejection or correction was employed.

B. Statistical Classifier

The different mental tasks (or states) are recognized by a classifier trained to classify EEG samples as class #1, #2, #3 or "unknown" [15]. In our statistical classifier, we have for each mental task a mixture of several Gaussians units. We think of each unit as a prototype of one of the mental tasks (or classes) to be recognized. The challenge is to find the appropriate position of the Gaussian prototype as well as an

appropriate variance.

We assume that the class-conditional probability density function of class C_k for sample x is a superposition of several Gaussians

$$p_k(x) = p(x | C_k) = \sum_{i=1}^{N_k} a_k^i(x) p_k^i \quad (1)$$

where N_k denotes the number of prototypes (Gaussians) of the class C_k , a_k^i and p_k^i are the activation level and the amplitude of the i th prototype of the class C_k , respectively. The amplitudes are constrained so that $\sum_i p_k^i = 1$. In our case, we set $p_k^i = 1/N_k$. Also, in the experiments reported below, all the classes have the same number of prototypes, namely $N_k = 4$. This choice is discussed in Sect. III.A, see also Table 4. The activation level is given by

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

where d is the dimensionality of the input space, μ_k^i corresponds to the center of the i th prototype of class C_k , Σ_k^i is the covariance matrix of this prototype, and $|\Sigma_k^i|$ is the determinant of the covariance matrix. In order to reduce the number of parameters we restrict our model to a diagonal covariance matrix Σ_k that is common to all the prototypes of the same class. Imposing diagonality equals an assumption of independence between the features. Even though we do not believe this assumption holds for our experiments in a strict sense, this has demonstrated to be a valid simplification of the model given the a posteriori good performance of the system. Before classification we average the class-conditioned probability over $N_{av}=8$ consecutive samples

$$\bar{p}_t(x | C_k) = \frac{1}{N_{av}} \sum_{i=1}^{N_{av}} p(x_{t-i+1} | C_k) \quad (3)$$

Finally the posterior probability of class C_k at time t is

$$y_k(t) = \frac{\bar{p}_t(x | C_k) P(C_k)}{\sum_j \bar{p}_t(x | C_j) P(C_j)} \quad (4)$$

where j ranges over the number of classes, three in our case, and $P(C_k)$ denotes the prior probability

of class C_k . In the following, we assume equal prior probabilities.

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 of 0.85; otherwise the response is classified as “unknown” so as to avoid making risky decisions for uncertain samples. This rejection criterion keeps the number of errors (false positives) low, which is desired since recovering from erroneous actions (e.g., robot turning in the wrong direction) has a high cost. The choice of this probability threshold was guided by a previous ROC study where different subjects only carried out the initial training described before [16], and the actual value was selected based on a non-exhaustive evaluation of the performance of the subjects during the first training session.

To initialize the center of the prototypes and the diagonal covariance matrix of the class C_k we run a clustering algorithm—typically, self-organizing maps [17]—to compute the position μ_k^i of the four prototypes per class. We then initialize the diagonal covariance matrix by setting

$$(\Sigma_k)_{mm} = \frac{1}{|S_k|} \sum_{n \in S_k} (x^n - \mu_k^{i^*(n)})_m^2 \quad (5)$$

where S_k denotes the set of indexes of samples belonging to the class C_k , $|S_k|$ is the cardinality of this set, $i^*(n)$ is the nearest prototype of this class to the sample x^n , and $\mu_k^{i^*(n)}$ is its center. The index m denotes the element of a vector, and mm the diagonal element of a matrix.

During learning we improve these initial estimations iteratively by stochastic gradient descent so as to minimize the mean square error $E = \frac{1}{2} \sum_j (y_j - t_j)^2$, where t_j is the j th component of the target vector in the form *l-of-c*; e.g., the target vector for class #3 is coded as (0,0,1). We compute y_j from equations (3) and (4) with $N_{av}=1$; i.e., each sample is used separately. Taking the gradient of the error yields

$$\frac{\partial E}{\partial \mu_k^i}(x) = \sum_j \frac{(y_j - t_j)}{p(x)} \left[\delta_{ki} - \frac{p_j(x)}{p(x)} \right] a_k^i(x) \Sigma_j^{-1} (x - \mu_k^i) \quad (6)$$

In order to simplify the algorithm we neglect the second term in the square brackets so that the final update rule is

$$\Delta\mu_k^i(x) = \alpha(t_k(x) - y_k(x))\Sigma_k^{-1}(x - \mu_k^i)\frac{a_k^i(x)}{p(x)} \quad (7)$$

where α is the learning rate. The interpretation of this rule is that, during training, the centers of the Gaussians are pulled towards the EEG samples of the mental task they represent and are pushed away from EEG samples of other tasks.

Finally, after every iteration over the training set, we estimate again the new value of Σ_k using equation (5).

The brain-machine interface responds every 0.5 s. Firstly, during each frame of 62.5 ms it computes the class-conditioned probability for each class—i.e., the mixture of Gaussians, equations (1) and (2). Secondly, it averages the class-conditioned probabilities over 8 consecutive samples, equation (3). Thirdly, it estimates the posterior probability based on the average class-conditioned probability of each class using Bayes' formula, equation (4). Finally, it compares the posterior probability with a threshold value of 0.85.

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 1 bit/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 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 so the theoretical channel capacity has a lower bound of approximately 0.85 bit/second. The channel capacity is estimated using the equation:

$$I = (1 - p_r) \left[\log_2(N) + (1 - p_e) \log_2(1 - p_e) + p_e \log_2\left(\frac{p_e}{N-1}\right) \right] \quad (8)$$

where N is the number of mental classes, p_r is the probability of “unknown” response, and p_e is the probability of error. Eqn. (8) is finally divided by the response interval of the brain interface (0.5 s for mode I and 1 s for mode II) to get the maximum bit rate that could be transmitted theoretically.

The actual bit rate in a control experiment is, however, lower for two reasons. The first one is that the operation of the brain-actuated robot does not require frequent switches between mental tasks. The second reason is that, although subjects can rapidly switch between tasks (see an example in Fig. 3A), they cannot maintain the maximum speed for a long time due to loss of attention.

C. Brain-Machine Interface Protocol

During an initial training period of 5 or 3 days, respectively, the two subjects learned to control 3 mental tasks of their choice with the interface operating in mode I. Neither subject had previous experience with meditation or specific mental training. The subjects 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 concatenating related words. All tasks (including relax) were performed with eyes opened. After a short evaluation, the experimental subjects “A” and “B” chose to work with the tasks relax-left-cube and relax-left-right, respectively. In the sequel, we will refer to these mental tasks as #1, #2 and #3 (i.e., relax is #1, left is #2, and cube or right is #3).

Each day, subjects participated in four consecutive training sessions of about 5 min, separated by breaks of 5-10 min. During each training session subjects switched randomly every 10-15 s between the three tasks. Subjects received feedback through three coloured buttons on a computer screen. The green button flashed if the mental state #1 was recognized, the blue button was associated to state #2, and the red button to state #3. After each training session the statistical classifier was optimized off-line. After this initial training, subjects learned to control mentally the mobile robot for 2 days with the interface operating in mode II. The results reported here were obtained at the end of the second day of work with the robot. During this training period, the user and the BMI engaged in a mutual learning process where they were coupled and adapted to each other.

A feature of the statistical classifier embedded in our brain-machine interface is the use of a

probability rejection criterion, which helps also to deal with “idle” states. In an asynchronous protocol, idle states appear during the operation of a brain-actuated device while the subject does not want the interface to carry out any new action. Although the statistical classifier is not explicitly trained to recognize those idle states, it can process them adequately by responding “unknown”. It is worth noting, however, that our subjects reported that the task of steering the robot between rooms was so engaging that they preferred to emit continuously mental commands rather than to go through idle states. Actually, one of our subjects reported that when he tried to stay in an idle state, he had a tendency to anticipate the next behavior the robot should execute and, instinctively, concentrated on the corresponding mental state—thus delivering a wrong mental command.

D. Mobile Robot

The mobile robot was a small Khepera (Fig. 1) that closely mimics a motorized wheelchair. The robot moved at a maximum speed of one third of its diameter per second, similar to the speed of a wheelchair in an office building. 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 environmental situations difficult if the raw readings were used directly. To overcome this limitation, we implemented a multilayer perceptron that maps the 8 raw infrared sensory readings into 6 classes of environmental states, or robot’s perceptual states; i.e., wall to the left, wall to the right, obstacle to the left, obstacle to the right, wall or obstacle in front, and free space. The mapping was optimized on an independent set of experiments where the robot was put at various locations in the environment.

III. RESULTS AND DISCUSSION

The task was to drive the robot through different rooms in a house-like environment (Fig. 1). During training the subject had to drive the robot mentally from a starting position to a first target room; once the robot arrived, a second target room was drawn at random and so on. At the end of the second day of training, the trajectories were qualitatively rather good and the robot never failed to visit the target room.

Fig. 3A shows the responses of the brain-machine interface over a continuous period of 16 s where subject “B” switched quickly between mental tasks to make the robot navigate between two rooms. The line (◆) gives the probability that the current EEG sample should be classified as #1, the line (■) that for #2 and the line (▲) that for #3. If none of the three probabilities is above the decision threshold of 0.85 (dotted horizontal line), the response is “unknown”. Operating in mode II, the robot executed a new behavior only after two consecutive identical responses (e.g., #1-#1; #2-#2; or #3-#3). Initially the robot moved forward (#1) to exit a room. Then, at step 9, it turned right (#3) to pass the doorway into the corridor. Note that at decision steps 12 and 13 the responses were “unknown”, since all lines stayed below the decision threshold and so the robot continued executing its current behavior (#3). At step 15, the robot moved forward (#1) and one second afterwards it switched to left wall following (#2) until it entered into a second room where it moved forward (#1) before it finally turned right (#3).

It is remarkable that Fig. 3A shows rapid transitions between mental tasks. There are two reasons for this. First, Fig. 3A corresponds to a difficult segment of the robot trajectory where many commands were required. It is therefore not representative of the responses of the brain interface over the whole experiment because normally the subject does not need to switch between mental tasks so quickly. But because the subject wanted to steer the robot so as to move in a relatively complex way, he needed to deliver a rapid and accurate sequence of mental commands—and the brain interface enabled him to do so. In other segments of the trajectory, precise commands were not required and the situation looks different: the experimental results reported in Sect. III.A indicate that the classifier outputs have intermediate values (that correspond to “unknown” responses) a significant number of times. Fig. 3B shows a plot of responses with intermediate values and slow transitions between mental commands during an “easy” segment of the robot trajectory. In this case, the robot was following a left wall for a few seconds and, after it has crossed a doorway, it moved forward. Note that even if the classifier responses were not always the correct one (there are several “unknown” responses and even an error at step 7), the robot still performed the correct behavior because of the control strategy implemented by our finite state automaton. A second reason for the relatively clear responses is that averaging class-conditioned probabilities before

using Bayes' rule (see equations (3) and (4)) helps to stabilize the responses of the classifier.

Although the subject brought the robot to each of the desired rooms, there were a few occasions where the robot did not follow the optimal trajectory. We may, therefore, wonder how efficient the mental control of the robot really is. In order to evaluate quantitatively the performance of the brain-actuated robot, subjects "A" and "B" also carried out a second set of experiments. In a given trial, the robot must travel from a starting room to a target room as well as also visit an intermediate room. The rooms and their order were selected at random. First, the subject made the robot visit the desired sequence of rooms by mental control. In a later session, the subject drove the robot along the same sequence of rooms by manual control. In this case, the subject used the same controller described above but, instead of sending mental commands (#1, #2, #3) 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 noting that the reason why the subject controls the robot mentally first and only afterwards manually is to avoid any learning process that could facilitate mental control. Table 1 gives the time in seconds for three different trials for the two subjects. For each trial, the table indicates the time required for mental control, manual control and also the relation between the two. Mental control was significantly worse than manual control, but still the ratio of operating times was as high as 0.74 on average. Thus, mental control is worse than manual control, but by less than a factor of 1.5.

Although users were emitting mental commands continuously, the theoretical minimum number of control commands to achieve the task is 13 for manual control during a typical task. However, in order to reach the target as fast as possible, subjects do emit more control commands than the minimum (almost twice). Also, in the case of mental control the number of control commands is significantly larger due to the less accurate control of the robot. On average, subjects switched between mental commands every 5.0 seconds.

A. Performance of the Statistical Classifiers

Tables 2, 3, and 4 give some additional details about the performance of the statistical classifiers for the two subjects. Table 2 shows the learning curves of the 2 subjects during the first day of training (4 training sessions in total with the brain interface operating in mode I). The classifiers are trained offline with data of a given training session and tested online on the *next* session. 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. A clear improvement in performance can be observed, in terms of errors and “unknown” responses. We note that already in session four (i.e., after 3 iterations of training) the system has excellent performance in mode II, corresponding to a theoretical bit rate of 1.02 for subject “A” and 0.91 for subject “B”.

Table 3 shows the average distance between the different prototypes of a given mental task and between prototypes of different tasks for the two subjects. The prototypes are those learned at the end of the first training period, before using the mobile robot. Interclass distances are always significantly larger than their corresponding intraclass distances, which is a clear indication that the learned prototypes are modeling relatively compact and distant subspaces, which in turn translates into good discrimination capabilities.

Table 4 shows the performance, measured in terms of errors and “unknown” responses, during the second training session of the two subjects 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 first training session. 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.

B. Analysis of EEG Control Signals

An important issue in any BMI system is to rule out the possibility that subjects may use electrooculogram (EOG) and electromyogram (EMG) activity as the control signals. Most EOG activity

occurs in the delta frequency range (0-4 Hz) [18] (cited in [19]) and so EOG activity should be nearly absent from the band 8-30 Hz we use for analysis. It is still possible that EOG and facial EMG activity is present in this band, but, if so, these artifacts should be more prominent in anterior electrodes than in posterior ones. In fact, we have found that this is not the case for any of the two experimental subjects. In particular, we have calculated the proportion of energy between the frontal and posterior locations

$$\frac{(F3 + F4) - (P3 + P4)}{F3 + F4 + P3 + P4} \quad (9)$$

This expression gives a real value between -1 (all the energy lies in the posterior sites) and 1 (all the energy is located in the anterior electrodes). This value is never larger than 0.03 and in most cases even significantly below zero for all the mental tasks chosen by the subjects. Table 5 gives the proportions for our two experimental subjects for two different bands, namely the whole 8-30 Hz we use for analysis and the higher band 22-30 Hz. Fig. 4 shows the power maps for the three mental tasks chosen by subjects “A” and “B” generated from the average prototypes learned during training¹. This is possible because prototypes give a direct interpretation of the EEG features that represent a given mental task and differentiate it from others. This figure illustrates for two different bands, 8-12 Hz and 22-30 Hz, that the energy is mainly located in the centro-parietal sites. Also, most of the energy lies in the lower frequencies, an indication of the absence of EMG artifacts that, if present, should appear in high frequencies and have larger amplitude than normal EEG signals. The figure shows clear differences between the learned prototypes of the different mental tasks, key for correct operation of the BMI, and between the prototypes of the same task for the two experimental subjects, a definitive evidence for the need of a personal BMI that fits the individual features of the subject.

As additional evidence that subjects are not using EMG activity, which is broad-band, if we apply machine-learning techniques for the selection of those relevant features that best differentiate the mental

¹ The main purpose of these maps is to show that no muscular activity was used as control signals. Its exact neurophysiological interpretation is not clear.

tasks, we find that the classifier performance improves with only a small proportion of features, which are not grouped in a cluster [20]. All these elements together support the fact that only EEG signals account for the control achieved.

There is, however, the possibility that, while imagining movements, subjects actually performed very light movements and the resulting EMG activity might influence the EEG signals. Although subjects did not make any visible movement, we cannot exclude this possibility, but we think this is rather unlikely.

IV. CONCLUSION

Although the quality and resolution of the brain signals measured with our EEG system are not comparable to those recorded by implanted electrodes [1]-[5], they are sufficient to mentally operate robots in indoor environments. This is possible because of the combination of advanced robotics, an asynchronous protocol for the analysis of online EEG signal, and machine learning techniques. In particular, this paper introduces the novel idea of controlling robots by mapping asynchronously high-level mental commands into a finite state automaton. This automaton is a key feature for the efficient control of the mobile robot. Such an efficient control could not be achieved with a naive, direct mapping of mental commands into robot motor actions. For instance, one could map the mental commands #2 and #3 into 90 degrees turns (to the left and right, respectively), but then the robot would fail to navigate in an indoor environment (even if walls and doorways are at right angles) because the robot accumulates orientation errors due to drifts and inaccurate motor actions. Alternatively, one could map those mental commands #2 and #3 into 10 degrees turns, but then the robot must stop exactly in front of a doorway and rotate 9 times to go through. On the contrary, thanks to the finite state automaton, the user needs only to send the mental command #2 or #3 when the robot is approaching a doorway (to the left or right, respectively) and the controller will bring the robot into the room automatically.

Our results open the possibility for physically disabled people to use a portable EEG-based brain-machine interface for controlling wheelchairs and prosthetic limbs. However, we will need to scale up the number of recognizable mental states to provide a more flexible and natural control of these robotics

devices. This could be done by estimating local field potentials of small cortical areas from the scalp potentials recorded with a sufficiently high number of electrodes (32, 64 or more) [21]. The Gaussian classifier embedded in the BMI would work upon the local field potentials of selected cortical areas instead of using EEG features. This is an area of current research. In a more general sense, our results indicate that online EEG analysis could, potentially, be more powerful than previously thought.

ACKNOWLEDGEMENTS

We would like to thank the two referees for their insightful comments. The initial phase of this work was partially supported by the European ESPRIT Programme, LTR Project 28193. J. del R. Millán is currently supported by the Swiss National Science Foundation through the National Centre of Competence in Research on “Interactive Multimodal Information Management (IM2).”

REFERENCES

- [1] J. K. Chapin, K. A. Moxon, R. S. Markowitz, and M. A. L. Nicolelis, “Real-time control of a robot arm using simultaneously recorded neurons in the motor cortex,” *Nature Neurosci.*, vol. 2, pp. 664–670, 1999.
- [2] J. Wessberg, C. R. Stambaugh, J. D. Kralik, P. D. Beck, M. Laubach, J. K. Chapin, J. Kim, S. J. Biggs, M. A. Srinivassan, and M. A. L. Nicolelis, “Real-time prediction of hand trajectory by ensembles of cortical neurons in primates,” *Nature*, vol. 408, pp. 361–365, 2000.
- [3] M. D. Serruya, N. G. Hatsopoulos, L. Paninski, M. R. Fellows, and J. Donoghue, “Instant neural control of a movement signal,” *Nature*, vol. 416, pp. 141–142, 2002.
- [4] D. M. Taylor, S. I. Helms Tillery, and A. B. Schwartz, “Direct cortical control of 3D neuroprosthetic devices,” *Science*, vol. 296, pp. 1829–1832, 2002.
- [5] M. A. L. Nicolelis, “Brain-machine interfaces to restore motor function and probe neural circuits,” *Nature Rev. Neurosci.*, vol. 4, pp. 417–422, 2003.
- [6] J. R. Wolpaw, D. J. McFarland, and T. M. Vaughan, “Brain-computer interface research at the Wadsworth center,” *IEEE Trans. Rehab. Eng.*, vol. 8, pp. 222–226, 2000.
- [7] G. Pfurtscheller and C. Neuper, “Motor imagery and direct brain-computer communication,” *Proc. IEEE*, vol. 89, pp. 1123–1134, 2001.
- [8] J. R. Wolpaw, N. Birbaumer, D. J. McFarland, G. Pfurtscheller, and T. M. Vaughan, “Brain-computer interfaces for communication and control,” *Clin. Neurophysiol.*, vol. 113, pp. 767–791, 2002.
- [9] J. del R Millán, “Brain-computer interfaces,” in *Handbook of Brain Theory and Neural Networks*, 2nd ed, M. A. Arbib, Ed. Cambridge, MA: MIT Press, 2002.
- [10] N. Birbaumer, N. Ghanayim, T. Hinterberger, I. Iversen, B. Kotchoubey, A. Kübler, J. Perelmouter, E. Taub, and H. Flor, “A spelling device for the paralysed,” *Nature*, vol. 398, pp. 297–298, 1999.
- [11] R. C. Arkin, *Behavior-Based Robotics*. Cambridge, MA: MIT Press, 1998.
- [12] F. Perrin, J. Pernier, O. Bertrand, and J. F. Echallier, “Spherical spline for potential and current density mapping,” *Electroenceph. Clin. Neurophysiol.*, vol. 72, pp. 184–187, 1989.
- [13] F. Perrin, J. Pernier, O. Bertrand, and J.F. Echallier, “Corrigendum EEG 02274,” *Electroenceph. Clin. Neurophysiol.*, vol. 76, pp. 565, 1990.
- [14] F. Babiloni, F. Cincotti, L. Bianchi, G. Pirri, J. del R. Millán, J. Mouriño, S. Salinari, and M. G. Marciani,

- “Recognition of imagined hand movements with low resolution surface Laplacian and linear classifiers,” *Med. Eng. Phys.*, vol. 23, pp. 323–328, 2001.
- [15] J. del R. Millán, J. Mouriño, M. Franzé, F. Cincotti, M. Varsta, J. Heikkonen, and F. Babiloni, “A local neural classifier for the recognition of EEG patterns associated to mental tasks,” *IEEE Trans. Neural Networks*, vol. 13, pp. 678–686, 2002.
- [16] A. Hauser, P.-E. Sottas, and J. del R. Millán, “Temporal processing of brain activity for the recognition of EEG patterns,” in *Proc. 12th Int. Conf. Artificial Neural Networks*, 2002, pp. 1125–1130.
- [17] T. Kohonen, *Self-Organizing Maps*, 2nd ed. Berlin, Germany: Springer-Verlag, 1997.
- [18] P. Berg, “Comments on EOG correction methods,” *J. Psychophysiol.*, vol. 3, pp. 41–44, 1989.
- [19] R. J. Croft and R. J. Barry, “EOG correction: A new perspective,” *Electroenceph. Clin. Neurophysiol.*, vol. 107, pp. 387–394, 1998.
- [20] 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.
- [21] R. Grave de Peralta, S. L. Gonzalez Andino, S. Morand, C. M. Michel, and T. Landis, “Imaging the electrical activity of the brain: ELECTRA,” *Human Brain Mapping*, vol. 9, pp. 1–12, 2000.

TABLES

TABLE I. Time in seconds for three different trials for subjects “A” and “B”.

Subject / mental tasks	Trial	Mental	Manual	Manual / Mental
“A” relax-left-cube	1	149	124	0.83
	2	183	135	0.74
	3	191	129	0.68
	<i>Average</i>	<i>174</i>	<i>129</i>	<i>0.75</i>
“B” relax-left-right	1	219	156	0.71
	2	189	155	0.82
	3	175	117	0.67
	<i>Average</i>	<i>194</i>	<i>143</i>	<i>0.73</i>

TABLE II. Performances over the consecutive training sessions of the first day. For each session, performance is measured in terms of errors (*err*, left column) and “unknown” responses (*rej*, right column).

Training Sessions								
Subject	MODE I				MODE II			
	2		3		4		4	
“A”	14.2%	30.1%	11.2%	27.2%	3.5%	21.4%	1.7%	29.1%
“B”	11.3%	29.4%	9.7%	28.9%	7.0%	24.9%	1.8%	36.5%
	<i>err</i>	<i>rej</i>	<i>err</i>	<i>rej</i>	<i>err</i>	<i>rej</i>	<i>err</i>	<i>rej</i>

TABLE III. Average distances between the different prototypes of a given mental task (d_{kk}) and between prototypes of different tasks (d_{kj}).

Task	Subject "A"		Subject "B"	
	d_{kk}	d_{kj}	d_{kk}	d_{kj}
Relax	0.0325	0.0740	0.0344	0.0651
Left	0.0260	0.0857	0.0386	0.0948
Cube	0.0445	0.0785		
Right			0.0386	0.0949

TABLE IV. Performance of the two experimental subjects during their second training session using different numbers of prototypes per mental task. Performance is measured as in TABLE II.

Subject	Number of Prototypes						
	1	2	4	6	9	12	
"A"	<i>err</i>	15.0%	16.1%	14.2%	18.7%	15.7%	20.7%
	<i>rej</i>	32.5%	30.3%	30.1%	32.0%	30.7%	29.6%
"B"	<i>err</i>	11.8%	13.0%	11.3%	16.4%	15.7%	15.9%
	<i>rej</i>	28.0%	25.3%	29.4%	24.8%	25.5%	27.0%

TABLE V. Proportion of energy between frontal and posterior locations for the two subjects and for the different mental tasks.

Task	Subject "A"		Subject "B"	
	8-30 Hz	22-30 Hz	8-30 Hz	22-30 Hz
Relax	-0.53	-0.13	-0.43	+0.03
Left	-0.15	-0.33	-0.37	-0.69
Cube	+0.03	-0.40		
Right			-0.01	-0.02

FIGURES




Fig. 1.

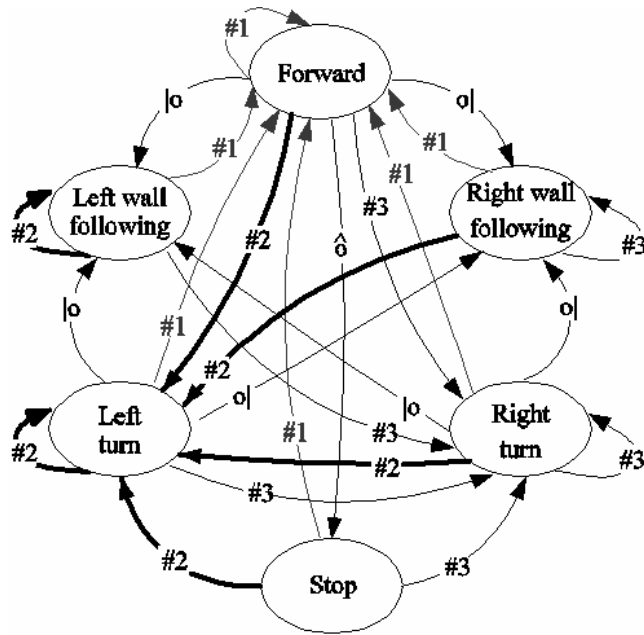
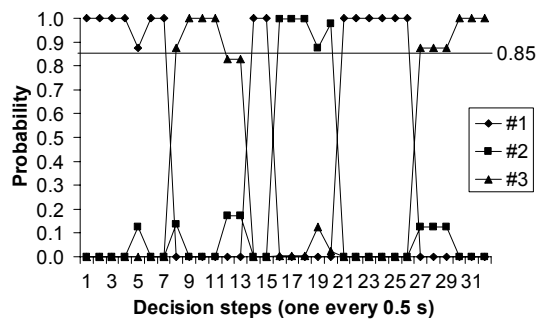



Fig. 2.



A



B

Fig. 3.

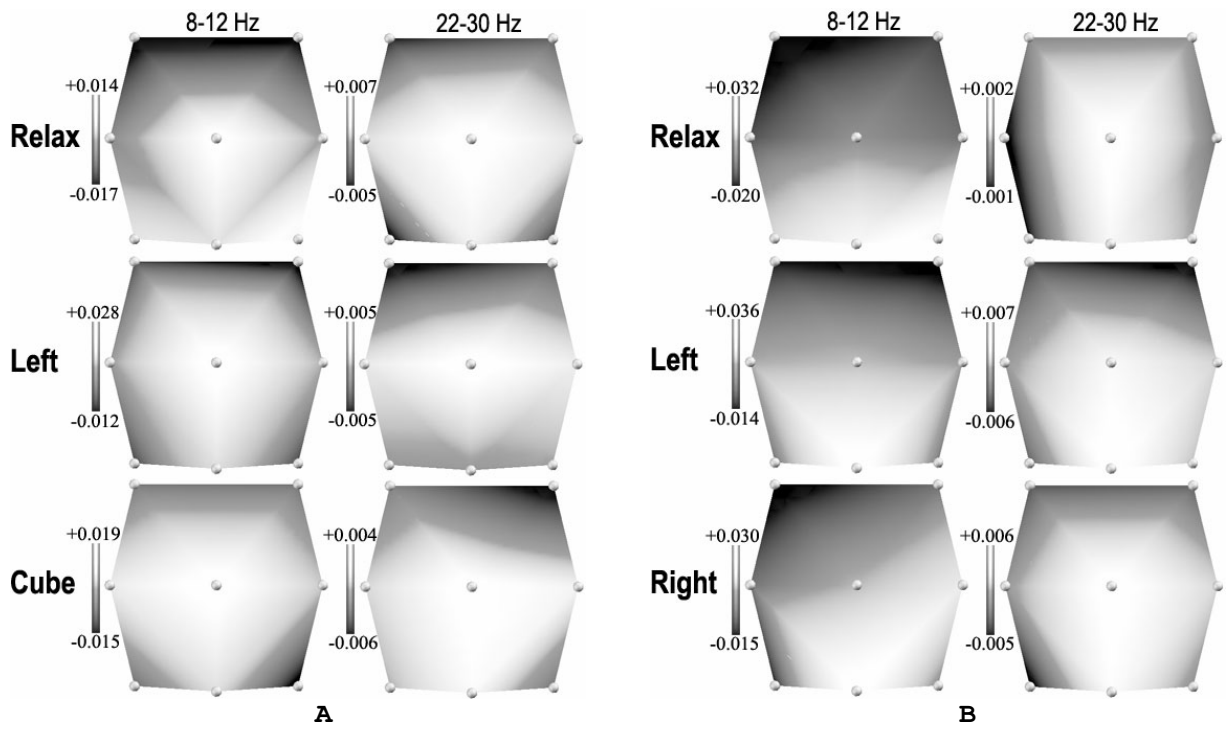


Fig. 4.

LEGENDS

Fig. 1. The mobile robot in its environments. The environment (80.0 cm × 60.0 cm) consists of several rooms along a corridor. The Khepera robot (5.7 cm diameter) is a two-wheeled vehicle. It has 3 lights on top to provide feedback to the user and 8 infrared sensors around its diameter to detect obstacles. The readings of the infrared sensors, which have limited perception ranges, are used by a multilayer perceptron to determine the probability to be in one of 6 perceptual states: open space, obstacle to left, obstacle to right, wall to left, wall to right, wall in front.

Fig. 2. Finite state automaton used for the control of the robot. Transitions between the 6 behaviors (ellipses) are triggered by 3 mental states (#1, #2, #3) and 4 perceptual states (l : left wall, r : right wall, \hat{o} : wall or obstacle in front, and free space). For example, the mental state #2 always causes a transition to “left turn” or, in the presence of a wall to the left, to “wall following”. A similar interpretation holds for the other mental states. For the sake of simplicity, this figure does not represent the obstacle-avoidance routine nor the full set of transitions to the behavior “stop”.

Fig. 3. Responses of the brain-machine interface while subject “B” was mentally controlling the robot during the first trial of the second set of experiments. A. Plot over a period of 16 s (32 decision steps) where the subject delivered a rapid and accurate sequence of mental commands. B. Plot where the responses exhibit intermediate values and slow transitions between mental commands.

Fig. 4. Power maps for the three mental tasks (#1, #2, #3, from top to bottom) chosen by the experimental subjects “A” and “B” (panels A and B, respectively) in two different bands, namely 8-12 Hz (left) and 22-30 Hz (right). Small filled circles indicate positions of the electrodes (frontal on top). The color scales are not the same for the different power maps, but are similar for all the tasks in a given band. Each scale is given next to the corresponding power map. Note that energy in the band 8-12 Hz is about 2 to 15 times larger than in the band 22-30Hz.