Non-Invasive Brain-Actuated Control of a Mobile Robot

Jose del R. Millan[1,2], Frederic Renkens[2], Josep Mourino[3], Wulfram Gerstner[2]
[1] Dalle Molle Institute for Perceptual Artificial Intelligence (IDIAP)
CH-1920 Martigny, Switzerland. jose.millan@idiap.ch
[2]Laboratory of Computational Neuroscience, Swiss Federal Institute of Technology
CH-1015 Lausanne EPFL, Switzerland
[3]Centre de Recerca en Enginyeria Biomedica, Universitat Politecnica de Catalunya
E-08028 Barcelona, Spain

Abstract

Recent experiments have indicated the possibility to use the brain electrical activity to directly control the movement of robotics or prosthetic devices. In this paper we report results with a portable non-invasive brain-computer interface that makes possible the continuous control of a mobile robot in a house-like environment. The interface uses 8 surface electrodes to measure electroencephalogram (EEG) signals from which a statistical classifier recognizes 3 different mental states. Until now, brain-actuated control of robots has relied on invasive approaches—requiring surgical implantation of electrodes—since EEG-based systems have been considered too slow for controlling rapid and complex sequences of movements. Here we show that, after a few days of training, two human subjects 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.

1 Introduction

There is a growing interest in the use of physiological signals for communication and operation of devices for physically-disabled as well as able-bodied people. Over the last years evidence has accumulated to show the possibility to analyze brainwaves on-line in order to determine the subjects' mental state that is then mapped into actions such as selecting a letter from a virtual keyboard or moving a robotics device [Birbaumer *et al,* 1999; Kennedy *et al,* 2000; Millan, 2002; Millan *et al,* 2002; Pfurtscheller and Neuper, 2001; Roberts and Penny, 2000; Serruya *et al,* 2002; Taylor *et al.,* 2002; Wolpaw and McFarland, 1994; Wolpaw *et al,* 2002]. This alternative communication and control channel, which does not require the user to perform any physical action, is called a *brain-computer interface (BCI).*

A BCI may monitor a variety of brainwave phenomena. Most BCIs use electroencephalogram (EEG) signals; i.e., the brain electrical activity recorded from electrodes placed onto 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 brain electrical activity, but it does not provide detailed information on the activity of single neurons (or small clusters of neurons) that could be recorded from microelectrodes surgically implanted in the cortex.

Some groups exploit evoked potentials—the automatic responses of the brain to external stimuli—recorded from either scalp or intracranial electrodes (for a review, see [Wolpaw *et al.,* 2002]). Evoked potentials are, in principle, easy to pick up but constrain the subject to synchronize themselves to the external machinery. A more natural and suitable alternative for controlling devices is to analyze components associated with spontaneous 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 [Birbaumer *et al.,* 1999]. Other groups look at local variations of EEG rhythms. The most used of such rhythms are related to the imagination of body movements and are recorded from the central region of the scalp overlying the sensorimotor cortex [Pfurtscheller and Neuper, 2001; Wolpaw and McFarland, 1994]. But, in addition to motor-related rhythms, other cognitive mental tasks are being explored [Millan *et al,* 2002; Roberts and Penny, 2000] as a number of neurocognitive studies have found that different mental tasks—such as imagination of movements, arithmetic operations, or language—activate local cortical areas at different extents. In this case, rather than looking for predefined EEG phenomena as when using slow cortical potentials or movement rhythms, the approach aims at discovering mental-specific EEG patterns embedded in the continuous EEG signals. Yet, another kind of spontaneous signals is the direct activity of neurons in the motor cortex measured with implanted electrodes [Kennedy *et al,* 2000; Serruya *et al.,* 2002; Taylor *et al,* 2002; Wessberg *et al,* 2000].

Recent experiments have shown the near possibility to use the brain electrical activity to directly control the movement of robotics or prosthetic devices. In these experiments, several monkeys have been implanted with microelectrodes recording the activity of single neurons (their spiking rate) in the motor and premotor areas of the cortex. Then, the monkey's hand trajectory was predicted (and replicated online with a robot arm) from the activity of the neural populations [Wessberg *et al,* 2000]. Also, after appropriate training, monkeys were able to move a computer cursor to desired targets using only their brain activity [Serruya *et al,* 2002; Taylor *et al,* 2002]. Until now, brain-actuated control of robots has being only tried with this kind of invasive approaches—requiring surgical implantation of electrodes—since EEG-based systems have been considered too slow for controlling rapid and complex sequences of movements. In this paper we show that two human subjects could, within a few days, learn to master a portable EEG-based brain-computer interface that recognized three mental states. Subjects 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.

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/second [Wolpaw *et al,* 2002]. 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 [Birbaumer *et al,* 1999; Pfurtscheller and Neuper, 2001; Wolpaw and McFarland, 1994]. In contrast, our approach uses an asynchronous protocol that analyzes the ongoing EEG to determine the subject's mental state, which they can voluntarily change at any moment. The rapid responses of the BCI, together with its performance (see Section 3), give a theoretical channel capacity in between 1 and 1.5 bits/second.

Two volunteer healthy subjects "A" and "B" wore a commercial EEG cap with integrated electrodes (white spots in Figure 1). EEG potentials were recorded at the 8 standard fronto-centro-parictal 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 [Perrin *et al,* 1989, 1990]. This spatial filtering yields new potentials that should represent better the cortical activity due only to local sources below the electrodes. Then, we used the Welch periodogram algorithm to estimate the power spectrum of each channel over the last second. We averaged 3 segments of 0.5 second with 50% overlap, what 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).

During an initial training period of a few days, the two subjects learned to control 3 mental tasks of their choice. 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[1]. After a short evaluation, the experimental subjects "A" and "B" chose to work with the combination of 3 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).** Neither subject had previous experience with BCIs or mental training.

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 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. After each training session the statistical classifier was optimized offline. After this initial training, subjects learned to control mentally the mobile robot for 2 days. 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 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 classify EEG samples as state **#1, #2, #3** or "unknown". In this statistical classifier, every 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 position, and receptive field, of the prototypes in the high-dimensional input space described above to differentiate the desired classes.

Although Gaussian classifiers are well known, 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

[1] Relax is done with eyes closed, whereas the other tasks arc performed with eyes opened. But the recognition of the task "relax" is *not* based on the detection of eye movements.

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

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

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

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 diagonal and common to all the prototypes of the class. In this way, we reduce the number of parameters and pool data to 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 of 0.85; 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) 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 [Hauser *et al,* 2002], and the actual value was selected based on the performance of the two subjects during the initial period of training.

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 [Kohonen, 1997]) to compute the position of the desired number of prototypes. Then, the covariance matrix is

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

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

We then improve these initial estimations iteratively by stochastic gradient descent so as to minimize the mean square error. For every sample x in the training set, the update rule for all the prototypes of all the classes is

$$\Delta\mu_k^i(x) = \alpha\left(t_k(x) - y_k(x)\right) \Sigma_k^{-1}(x-\mu_k^i)\frac{a_k^i(x)}{A(x)}\sum_{j\neq k}^{N_c} y_j(x) \tag{4}$$

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

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. 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 a threshold value of 0.85. At the end of training, errors and "unknown" responses are below 5% and 30%, respectively. The theoretical channel capacity of the interface is hence above 1 bit/second (operation mode 1). 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 the theoretical channel capacity is about 1 bit/second.

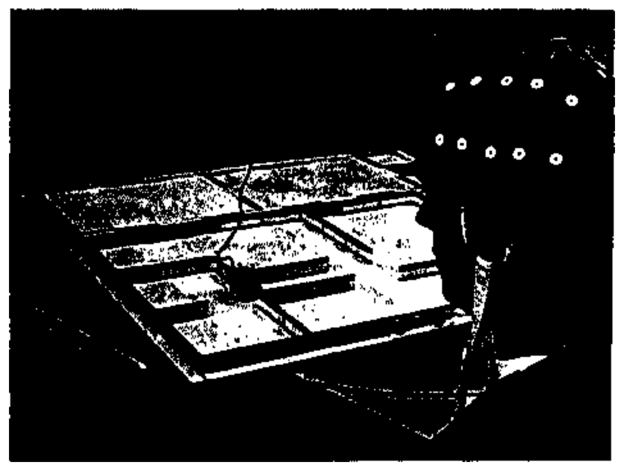

Figure 1. One of the experimental subjects while driving mentally the robot through the different rooms of the environment during the first experiment.

4 Robot Setup and Control

The task was to drive the robot through different rooms in a house-like environment (Figure 1). 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 wheelchair in an office building.

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 subjects can issue high-level commands at any moment. This is possible because the operation of the BC1 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 [Arkin, 1998] 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. Figure 2 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 actual transitions between the behaviors "forward" and "follow left/right wall" are not exactly as indicated in the figure, otherwise the robot would stay following walls forever. 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 experimental subjects felt that our control schema was simple and intuitive to use.

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, what makes difficult the recognition of the different perceptual states from the raw readings. To overcome this limitation, the robot uses a multilayer perceptron that maps the 8 raw infrared sensory readings into the current perceptual state.

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 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).

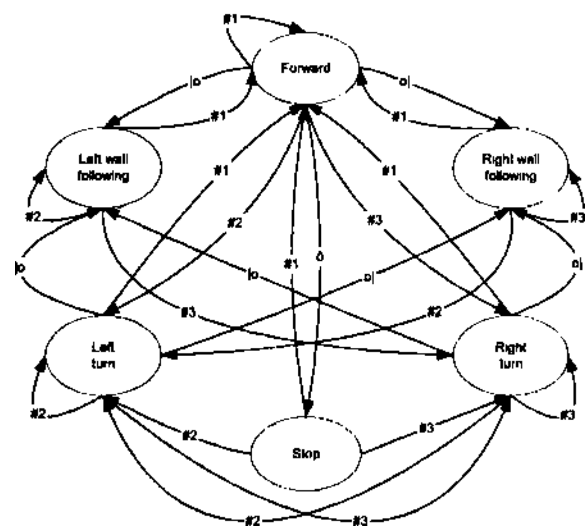

Figure 2. Finite state automaton used for the control of the robot. Transitions between the 6 behaviors were determined by 3 mental states (#1, #2, #3), 6 perceptual states (|o: left wall, o|: right wall, 6: wall or obstacle in front), and some memory variables. The memory variables and some of the perceptual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

5 Experimental Results

After 5 and 3 days of initial training with the interface operating in mode 1, subjects "A" and "B", respectively, achieved a satisfactory level of performance (correct recognition was above 60% while errors were below 5%). At this moment, subjects started to learn to control mentally the robot with the interface operating in mode II. During this second period of training subjects 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.

Figure 3 shows a trajectory generated by subject "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 subject was driving the robot for about 10 minutes continuously. Although the subject brought the robot to the desired room each time, 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 subject'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 subject needed to maneuver mentally the robot to bring it back. In other situations, the robot's sensors perceived a wall or corner too close, thus making the robot stop automatically to avoid collisions. In these situations, the subject needed to turn (by mental control) the robot away from the phantom obstacle and then resume the trajectory.

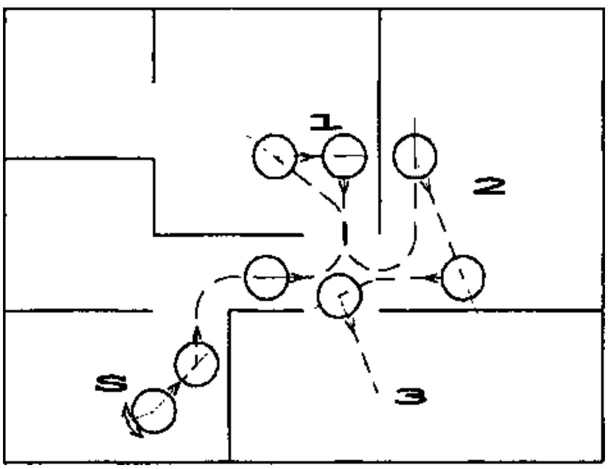

Figure 3. Trajectory followed by the robot under the mental control of subject "A" during one of the trials of the first experiment. The robot started in the bottom left room and then visited 3 other rooms, top center, top right and bottom right, sequentially. The figure docs not show the details of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

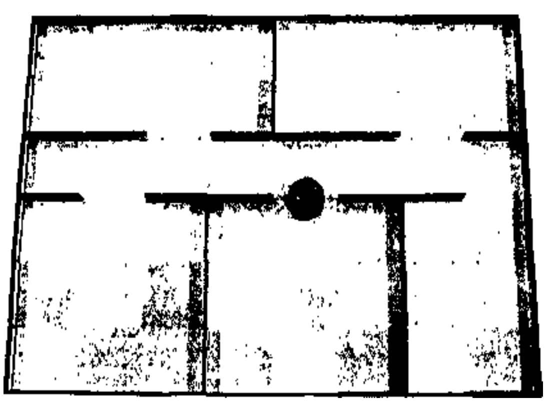

Figure 4. Environment used for the second set of experiments.

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, subjects "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 (Figure 4).

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 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 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 necessary to generate the desired trajectory for three different trials for the two subjects. For each trial, the table indicates the time required for mental control and manual control. Surprisingly, we can see that mental control was only marginally worse than manual control. On average, brain-actuated control of the robot is only 35% longer than manual control for both subjects.

Table 1. Time in seconds for three different trials where subjects "A" and "B" controlled the robot first mentally and then manually.

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

Discussion

In this paper we have reported first results of a brain-actuated mobile robot by means of a portable non-invasive BCI. Although the quality and resolution of the brain signals measured with our EEG system are not comparable to those recorded by implanted electrodes, they are sufficient to 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.

The work described in this paper suggests that it could be possible for human subjects to mentally operate a wheelchair. But 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 subject is seated on a mobile platform. This will require longer training times for the subject. 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 [Michel *et al,* 2001] suggests that a sufficient number of mental states can be recognized to control robotics and prosthetic devices and in a more natural and flexible way. 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.

Acknowledgments

JRM is supported by the Swiss National Science Foundation through the National Centre of Competence in Research on "Interactive Multimodal Information Management (1M2)."

References

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

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

[Hauser *et al,* 2002] A. Hauser, P.-E. Sottas, and J. del R Millan. Temporal processing of brain activity for the recognition of EEG patterns. In *Proc International Conference on Artificial Neural Networks,* pages 1125-1130, Madrid, Spain, August 2002. Springer-Verlag.

[Kennedy *et al,* 2000] P. R. Kennedy, R. Bakay, M. M. Moore, K. Adams, and J. Goldwaithe. Direct control of a computer from the human central nervous system. *IEEE Trans, on Rehabilitation Engineering,* 8:198-202, 2000.

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

[Michel *et al,* 2001] C. M. Michel, G. Thut, S. Morand, A. Khateb, A. Pegna, R. Grave de Peralta, S. Gonzalez, M. Seeck, and T. Landis. Electric source imaging of human brain functions. *Brain Research Reviews,* 36:108-118,2001.

[Millan, 2002] J. del R. Millan. Brain-computer interfaces. In M. A. Arbib (ed.), *Handbook of Brain Theory and Neural Networks,* pp. 178-181. MIT Press, Cambridge, Massachusetts, 2002.

[Millan *et al,* 2002] J. del R. Millan, J. Mourino, M. Franze, 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, on Neural Networks,* 13:678-686,2002.

[Perrin *et al,* 1989] F. Perrin, J. Pernier, O. Bertrand, and J. Echallier. Spherical spline for potential and current density mapping. *Electroencephalography and Clinical Neurophysiology,* 72:184-187, 1989.

[Perrin *et al,* 1990] F. Perrin, J. Pernier, O. Bertrand, and J. Echallier. Corrigendum EEG 02274. *Electroencephalography and Clinical Neurophysiology,* 76:565, 1990.

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

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

[Serruya *et al,* 2002] M. D. Serruya, N. G. Hatsopoulos, L. Paninski, M. R. Fellows, and J. Donoghue. Instant neural control of a movement signal. *Nature,* 416:141-142,2002.

[Taylor *et al,* 2002] D. M. Taylor, S. I. Helms Tillery, and A. B. Schwartz. Direct cortical control of 3D neuroprosthetic devices. *Science,* 296:1829-1832, 2002.

[Wessberg *et al,* 2000] 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,* 408:361-365, 2000.

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

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