

- [24] J. F. Hetke, J. C. Williams, D. S. Pellinen, R. J. Vetter, and D. R. Kipke, "3-D silicon probe array with hybrid polymer interconnect for chronic cortical recording," presented at the *1st Int. IEEE EMBS Conf. Neural Engineering*, Capri, Italy, 2003.
- [25] J. U. Meyer, T. Stieglitz, O. Scholz, W. Haberer, and H. Buetel, "High density interconnects and flexible hybrid assemblies for active biomedical implants," *IEEE Trans. Adv. Packaging*, vol. 24, pp. 366–375, Aug. 2001.
- [26] J. Chen, K. D. Wise, J. F. Hetke, and S. C. Bledsoe Jr., "A multichannel neural probe for selective chemical delivery at the cellular level," *IEEE Trans. Biomed. Eng.*, vol. 44, pp. 760–769, Aug. 1997.

Dynamical Dimension of a Hybrid Neurobotic System

Michael Kositsky, Amir Karniel, Simon Alford, Karen M. Fleming, and Ferdinando A. Mussa-Ivaldi

Abstract—The goal of this work is to understand how neural tissue can be programmed to execute predetermined functions. We developed a research tool that includes the brainstem of a lamprey and a two-wheeled robot interconnected in a closed loop. We report here the development of a framework for studying the dynamics of the neural tissue based on the interaction of this tissue with the robot.

Index Terms—Brain-machine interface, dynamical dimension, neuro-controllers.

I. INTRODUCTION

Recent studies on primates have shown that information measured in the motor cortex can be used to guide a robotic arm [1]–[3]. A critical challenge for brain machine interfaces is to train the neural system connected to a device in the execution of a variety of tasks. While some studies are focusing on learning mediated by visual feedback mechanisms, the focus of our work is on the control of plasticity at a specific synaptic site, through the interaction of pre- and post-synaptic activities associated with feedback and command signals. To address this issue, we have developed a research tool that includes a portion of living brain tissue of a lamprey and a two-wheeled robot interconnected in a closed loop [4].

The Lamprey is an eel-like fish, whose nervous system has been extensively studied, particularly for what concerns its ability to generate and modulate locomotor behavior [5]. We have selected a portion of the neural circuitry that in normal circumstances receives vestibular and other sensory signals and issues motor commands to stabilize the orientation of the body during swimming [6], [7]. This system has been shown to be adaptive, as unilateral lesions of the vestibular capsules are followed by a slow reconfiguration of neuronal activities until the correct postural control is recovered [6].

In our neurobotic system, vestibular signals are replaced by light-intensity signals. Two electrodes apply stimulations to the axons of the

octavomotor nuclei. The stimulations rates are proportional to the light intensity measured by sensors on the right and left sides of the mobile robot. The mobile robot has two wheels that receive velocity commands proportional to the population spike rates recorded by two electrodes in the spinal cord. Therefore, the natural stabilizing behavior—in which the lamprey tracks the vertical axis—correspond, in the hybrid system to tracking a source of light. Indeed, in most cases, the robot orients itself toward the source of light.

The nervous system between the four electrodes assumes the function of a controller, with two inputs and two outputs that determine the behavior of the robot in a closed feedback loop. In this communication, we focus on the identification of dynamical properties of neural tissue embedded in the neurobotic system. In particular, we are interested to establish the dynamical dimension of the neural element, as this is a critical parameter to assess what can and cannot be learned.

Section II describes the neurobotic system, i.e., the neural tissue, the robot, and the interface. Section III briefly describes previous research with the neurobotic system, which laid the foundation for the work being reported. Section IV describes the experimental protocol. Section V describes the computational analysis method used to establish the dynamic dimension of the system. Then the results are presented in Section VI and discussed in Section VII.

II. NEUROBOTIC SYSTEM

The neurobotic system includes three elements: neural tissue, a robot, and an interface.

A. Neural Tissue

The neural component of the hybrid system is a portion of the brainstem of the Sea Lamprey (*Petromyzon marinus*) in its larval state. After anesthetizing the larvae with tricane methanesulphonate (MS222, 100–200 mg/L), the whole brain was dissected and maintained in continuously superfused, oxygenated, and refrigerated (9 °C–11 °C) Ringer's solution (NaCl, 100.0 mM; KCl, 2.1 mM; CaCl₂, 2.6 mM; MgCl₂, 1.8 mM; glucose, 4.0 mM; NaHCO₃, 25.0 mM), see details in [8].

We recorded extracellularly the activity of neurons in a region of the reticular formation, a relay that connects different sensory systems (visual, vestibular, tactile) and central commands to the motor centers of the spinal cord. We placed two recording electrodes among the axons of the right and left posterior rhombencephalic reticular nuclei (PRRN). We also placed two unipolar tungsten stimulation electrodes among the axons of the intermediate and posterior octavomotor nuclei (nOMI and nOMP). These nuclei receive inputs from the vestibular capsules and their axons form synapses with the rhombencephalic neurons on both sides of the midline. The recorded signals were acquired at 10 kHz by a data acquisition board (National Instruments PCI-MIO-16E4) on a Pentium III 2GHz computer (Dell Computer Corporation).

While the axons of the nOMI remain ipsilateral, those of the nOMP cross the midline. Accordingly, the activity of one vestibular capsule affects both the ipsi- and contralateral reticulospinal (RS) nuclei. We placed each stimulating electrode near the region in which the axons of the nOMI and nOMP cross [see Fig. 1(a)]. This placement of the electrodes induced predominantly excitatory responses in the downstream neurons. The recording electrodes were placed on either side of the midline, near the visually identifiable neurons of the PRRN. To verify the placement of the stimulating electrodes we delivered brief single stimulus pulses (200 μs) and observed the response in both the ipsi- and contralateral PRRN neurons. Once it was determined that the stimulation electrodes were properly placed, the recording electrodes were moved caudally in order to pick up population spikes.

Manuscript received July 22, 2002; revised April 30, 2003. This work was supported by The Office of Naval Research under Grant N000140110524.

M. Kositsky, K. M. Fleming, and F. A. Mussa-Ivaldi are with the Department of Physiology, Northwestern University Medical School, Chicago, IL 60611 USA (e-mail: kositsky@northwestern.edu; kfleming@northwestern.edu; sandro@northwestern.edu).

A. Karniel is with the Department of Electrical Engineering, Technion-Israel Institute of Technology, Haifa 32000, Israel (e-mail: karniel@ee.technion.ac.il).

S. Alford is with the Department of Biological Sciences, University of Illinois at Chicago, Chicago, IL 60607 USA (e-mail: sta@uic.edu).

Digital Object Identifier 10.1109/TNSRE.2003.814444

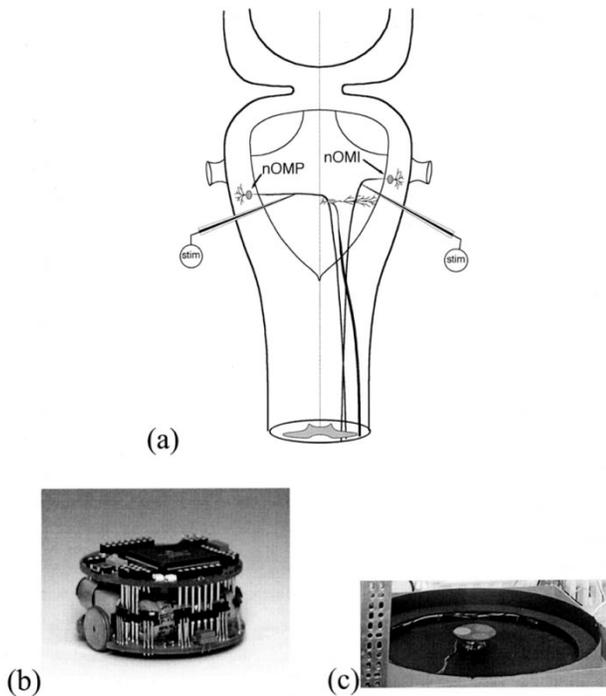


Fig. 1. (a) Circuitry of the preparation and placement of stimulating electrodes. (b) Two-wheeled robot (Khepera module). (c) Robot in the workspace.

B. Robot

The robot system is the base Khepera module [Fig. 1(b)]. Placed along the circumference of the robot are eight sensors providing light-intensity information. The sensors are located on opposite sides of the robot's midline at 10° , 45° , 85° , and 165° from the front direction. Two wheels provide a means of locomotion for the small robot. Our computer system communicates with the robot through the serial port.

The small size of the robot (5.3 cm in diameter) allowed us to use a small workspace [Fig. 1(c)]. A circular wall was constructed with a 2-ft diameter and then painted black to reduce the amount of reflected light. Eight lights are mounted at the edge of the workspace at 45° intervals. The lights are computer controlled using the digital outputs of the acquisition card. These lights generate the stimulus that elicits a phototactic response.

In some of the experiments, we have used a simulated version of the robot. This proved to be quite useful in generating a large volume of data, as the simulated robot does not require continuous supervision by the experimenter.

C. Interface

The interface acts as an interpreter between the neural signals and the robot control system. It is responsible for the transformation of the robot's light sensor information into vestibular inputs as well as for processing in real time the neural activity of the reticulo-spinal nuclei and translating it into motor commands for the robot.

The rates of pulses emitted by the right and left stimulating electrodes are proportional to a combined intensity measured by optical sensors respectively on the right and left sides of the mobile robot.

The spiking activities of the PRRN as recorded near the axons are analyzed through a five-step process. The signal picked up by the recording electrodes contains a combination of spikes, stimulus artifacts, excitatory and inhibitory postsynaptic potentials (PSP), and noise. To suppress the slow PSP components, this signal is first put

through a high-pass filter (cut off at 200 Hz). The output of this filter contains high-frequency noise, stimulus artifacts, and the spikes generated by multiple neurons in the vicinity of the electrode. Stimulus artifacts are canceled by zeroing the recorded signals over temporal windows of 5 ms following the delivery of each 200- μ s stimulus pulse. The remaining signal is rectified, and a threshold is applied to separate the spikes from the background noise under the assumption that the spike amplitude is much larger than the noise amplitude. The resulting train of spikes is put through a low-pass filter (5 Hz), which is effectively equivalent to frequency to voltage conversion.

III. PREVIOUS WORK

The key achievement of earlier research with the neurobotic system was the establishment of stable phototaxis, reproducible across preparations. In these experiments, the two wheels of the robot receive velocity commands proportional to the estimated population spike rates (computed as described in Section II-C) recorded by two electrodes in the brainstem. In this arrangement, the natural stabilizing behavior—in which the lamprey tracks the vertical axis—translates into the robot tracking a source of light [9]. At each trial, the robot was initially placed at the center of the workspace, one of the light sources on the perimeter of the workspace was turned on, and the closed-loop neurobotic system was run for 5–10 s. The neural preparation between the four electrodes assumed the function of a controller with two inputs and two outputs that determine the behavior of the robot in a closed feedback loop. Fig. 2(a) shows sample trajectories produced with a preparation exhibiting positive phototaxis, i.e. moving the robot toward the light source. Fig. 2(b) shows sample trajectories produced with a preparation exhibiting negative phototaxis, i.e. moving the robot away from the light source. Whether a preparation exhibits positive or negative phototaxis depends mostly on fine details of electrode placement. A preparation exhibiting negative phototaxis can easily be changed into one exhibiting positive phototaxis by swapping the right and left commands to the wheels.

IV. EXPERIMENTAL PROTOCOL

In the dynamical dimension experiments reported here, a simulated robot was used. To simplify the procedure and the analysis, movements of the simulated robot are restricted to a single degree of freedom. Namely, the robot is restricted to only revolving around its center. The tissue-to-robot interface in this experiment differs from the one used to obtain phototactic behaviors (see Section III). The two robot wheels always rotate in opposite directions with equal speeds. The magnitude of the speeds is proportional to the difference between the spike rates recorded by the right and left electrodes. Thus, the robot never moves forward or backward, but only rotates around its center with an angular velocity proportional to the difference between the right and left signals. The collected trajectories are represented as a time course of the robot orientation. We are only interested in the difference between the robot orientation θ and the direction toward the light source θ_l , which we will call $\Delta\theta$ (see Fig. 3).

The experiment consists of a series of episodes. At the beginning of each episode, the initial position of the robot is chosen at random. Then, the entire closed-loop system, which includes the robot and the neural tissue, is run for a certain time fixed throughout the entire experiment. Each episode produces a trajectory, $\Delta\theta(t)$. Between each two consecutive episodes, the neural signal is not recorded and the robot is set into a random initial position for the next episode.

V. IDENTIFYING THE DYNAMICAL DIMENSION OF THE SYSTEM

The dynamical dimension of the neural tissue is the information that we are seeking to extract. If one were to consider the tissue as a memory element, then this would introduce a dependence of the output at any

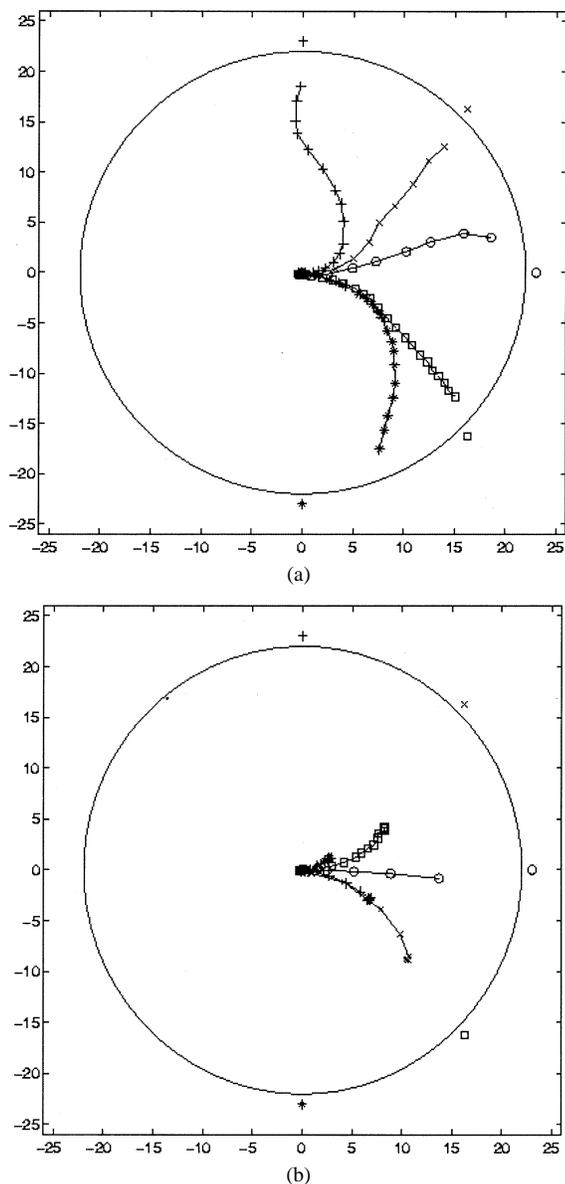


Fig. 2. (a) Positive phototaxis trajectories. (b) Negative phototaxis trajectories. Each trajectory originates at the center of the workspace. Different symbols around the workspace mark location of the light sources. For each light source, the corresponding trajectory is marked by the same symbol. Notice that each negative phototaxis trajectory is heading in a direction opposite to the corresponding positive phototaxis trajectory.

given time upon inputs or outputs at earlier times. To assess the dimension of neural tissue, we have developed a method that allows us to represent the phase space of the neurorobotic system by constructing a surrogate state vector, composed of position samples $\Delta\theta$ at different multiples of a fixed time-delay Δt . A well-known embedding theorem states that the dimension of a dynamical system can be deduced from the minimal number of components in the surrogate state vector that lead to trajectories without self-intersections.

In order to determine intersections in the trajectories due to an insufficient dimension of the surrogate state vector, we used a method known as false nearest neighbors elimination [10]. To assess whether a given dimension d is the adequate dynamical dimension for the trajectory in hand, we perform the following analysis. For each set of d consecutive points along the trajectory, a surrogate state vector is formed where each point determines one component of the vector. For each sample of the surrogate state vector, we find the nearest neighbor in the d -dimensional state space. Then, each surrogate state is analyzed

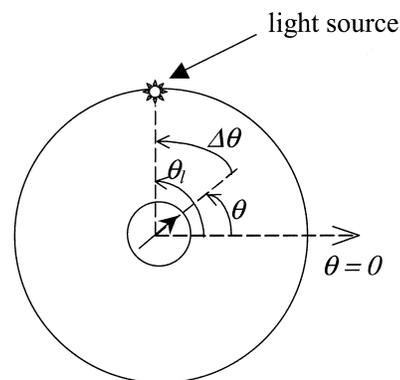


Fig. 3. Experimental setup. The large circle depicts the workspace and the smaller circle in the middle depicts the robot. The robot's orientation is shown by the arrow. The difference between the robot orientation θ and the direction toward the light source θ_i is denoted by $\Delta\theta$.

in $(d + 1)$ -dimensional space by adding the next point along the trajectory as another component. If the nearest neighbor remains close to the state in question in the $(d + 1)$ -dimensional space, it is considered a true nearest neighbor. Otherwise, it is a false nearest neighbor, which happens to be close to the state in question due to projecting from a higher dimensional adequate state space.

These transitions from dimension d to $d + 1$ are repeated for d equal to 1, 2, etc. For each magnitude of d , the percentage of false nearest neighbors is evaluated. The minimal d , which provides the 0% of false nearest neighbors is the dimension of the embedding space.

VI. RESULTS

The results were collected using four preparations. Observed trajectories and the deduced dynamical dimension are similar for different preparations. Fig. 4 shows a representative collection of trajectories produced by the closed-loop neurorobotic system using one of the preparations. The time delay between consecutive samples Δt is 200 ms.

In phototactic behavior experiments (see Section III), achieving a clear phototactic behavior requires tuning of some interface parameters, like threshold magnitudes used for spike detection and balance between the motors of the robot. In the experiment reported here though, we did not perform any such tuning or adjustments. Thus, the stable or meta-stable values of $\Delta\theta$ observed in the trajectories are not close to zero.

We applied the false nearest neighbors analysis described above to the time series of the robot orientation relative to the light source, $\Delta\theta$. The results for one of the trajectories are shown in Fig. 5. In this case, the dynamical dimension is equal to six.

We should stress that more trajectories need to be analyzed and collected before a firm conclusion on dimensionality can be reached. The observed data are very noisy, different trajectories and different preparations produce slightly different results. Nevertheless, these observations allow us to draw a very important conclusion: this preparation appears to have a rather extended "memory capacity," as its behavior depends on its previous states, not only on the current value of the input signal.

For different trajectories, the reconstructed dynamical dimension takes values between 5 and 10.

VII. DISCUSSION AND CONCLUSION

We report preliminary results allowing us to assess the mathematical dimension of the state space for a closed-loop neurorobotic system. These observations might be related to a temporal memory span of the preparation. Earlier work in our group has attempted to identify specific mechanisms underlying this memory function. Regardless of the

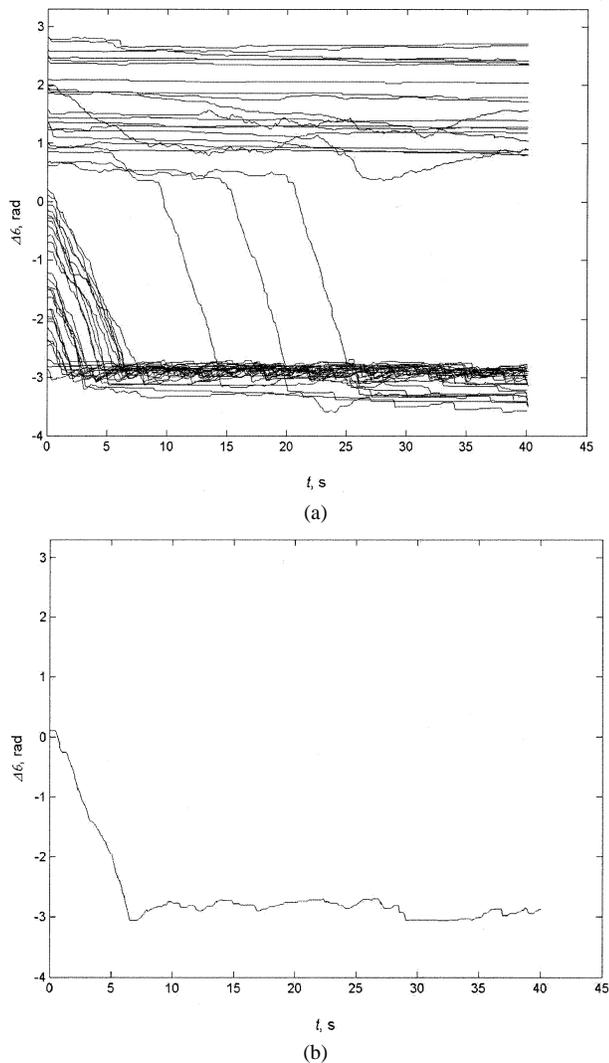


Fig. 4. (a) Sample trajectories produced during the experiment. (b) One of the trajectories. Notice that in this experiment the trajectories are represented in terms of $\Delta\theta$ versus t , which is different from trajectories in phototactic behavior experiments, like those shown in Fig. 2.

particular mechanisms though, this information is critical to determine what functions the neural components may be trained to execute. This memory function can be at the basis of the sustained oscillatory activities that generate swimming motion of the lamprey's body.

The dynamical dimension of the entire neurorobotic system depends on the dynamical properties of each of the components—the preparation, the robot, and the interface. The robot used in these experiments (both physical and simulated) is inertialess, hence its dynamical dimension is equal to 1. However, the interface between the neural preparation and the robot probably introduces an extra dimension to the entire system. To address this question, a series of computational experiments should be conducted with models of known dimensions in place of the neural preparation. Results obtained using simulated preparations of different dimensions will be matched against the results observed in the experiments reported here.

The results reported in this communication are in good agreement with findings, both modeling and experimental. In an earlier analysis [11], ten two-input/two-output neural network models of the same neural structure were trained with the observed phototactic trajectories of the neurorobotic system, which included a physical robot. The generalization error of each model, that is the ability of each model to reproduce observed robot trajectories, was analyzed. It was found

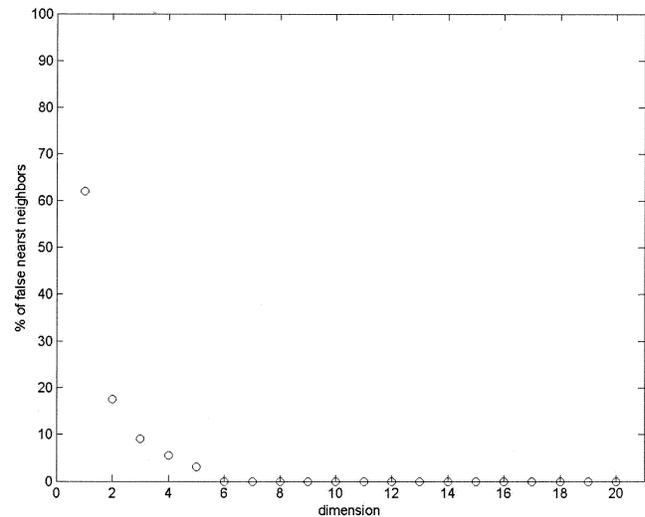


Fig. 5. False nearest-neighbors analysis reveals a high-order system.

that dynamic models, in which the output depends upon the input but also upon earlier outputs, were significantly superior to static models, in which the output depends only upon the current input. This finding can be attributed to two distinct causes: 1) recurrent neural pathways that make a neuron's current activity dependent upon past activities, or 2) intrinsic dynamic properties of neural behavior, not involving any recurrent connectivity. We have tested these alternatives by observing the behavior of the neurorobotic system, following a transection of the spinal cord, caudal to the recording electrodes. This reduced preparation displayed the same dynamical properties as the intact preparation, thus, ruling out a critical role for recurrent connections in generating these properties. An alternative candidate mechanism responsible for the dynamical properties of the preparation was investigated using neuropharmacological techniques [12]. This mechanism is based on the functioning of metabotropic glutamate receptors (mGluR) as autoreceptors: when repeated action potentials arrive on the presynaptic terminal, glutamate is released [13]. The mGluRs on the same terminal are then activated leading to an enhancement in synaptic transmission. This mechanism is a form of short-term plasticity, mediated by release of calcium ions that leads to history-dependent pattern of activity. In other words, this specific type of receptor seem to be at the origin of a positive feedback mechanism that enhances the efficacy of a synapse as a function of past use.

We are now investigating the possibility of using reinforcement learning schemes based on combinations of electrical stimuli and pharmacological agents for programming desired behaviors of this neurorobotic system.

REFERENCES

- [1] M. Serruya, N. Hatsopoulos, L. Paninski, M. Fellows, and J. Donoghue, "Brain-machine interface: Instant neural control of a movement signal," *Nature*, vol. 416, pp. 141–142, 2002.
- [2] D. Taylor, S. Tillery, and A. Schwartz, "Direct cortical control of 3D neuroprosthetic devices," *Science*, vol. 296, pp. 1829–1832, 2002.
- [3] J. Wessberg, C. Stambaugh, J. Kralik, P. Beck, M. Laubach, J. Chapin, J. Kim, S. Biggs, M. Srinivasan, and M. Nicolelis, "Real-time prediction of hand trajectory by ensembles of cortical neurons in primates," *Nature*, vol. 408, pp. 361–365, 2000.
- [4] B. Reger, K. Fleming, V. Sanguineti, S. Alford, and F. Mussa-Ivaldi, "Connecting brains to robots: an artificial body for studying the computational properties of neural tissues," *Artif. Life*, vol. 6, pp. 307–324, 2000.
- [5] S. Grillner, L. Cangiano, G.-Y. Hu, R. Thompson, R. Hill, and P. Wallen, "The intrinsic function of a motor system—from ion channels to networks and behavior," *Brain Res.*, vol. 886, pp. 224–236, 2000.

- [6] T. Deliagina, "Vestibular compensation in lampreys: impairment and recovery of equilibrium control during locomotion," *J. Exp. Biol.*, vol. 200, pp. 1459–1471, 1997.
- [7] C. Rovainen, "Electrophysiology of vestibulospinal and vestibuloreticulospinal systems in lampreys," *J. Neurophysiol.*, vol. 42, pp. 745–766, 1979.
- [8] S. Alford, I. Zompa, and R. Dubuc, "Long-term potentiation of glutamatergic pathways in the lamprey brainstem," *J. Neurosci.*, vol. 15, pp. 7528–7538, 1995.
- [9] K. M. Fleming, B. D. Reger, V. Sanguineti, S. Alford, and F. A. Mussa-Ivaldi, "Connecting brains to robots: an artificial animal for the study of learning in vertebrate nervous system," in *Proceedings of the VI International Conference on Simulation of Adaptive Systems*. Cambridge, MA: MIT Press, 2000.
- [10] H. Abarbanel, *Analysis of Observed Chaotic Data*. New York: Springer-Verlag, 1996.
- [11] A. Karniel, K. Fleming, V. Sanguineti, S. Alford, and F. Mussa-Ivaldi, "Dynamic properties of the Lamprey's neuronal circuits as it drives a two-wheeled robot," in *Proceedings of the SAB'2002 Workshop on Motor Control in Humans and Robots: On the Interplay of Real Brains and Artificial Devices*, J. M. Carmona and G. Maistros, Eds. Edinburgh, Scotland, 2002, pp. 29–36.
- [12] M. Takahashi and S. Alford, "The requirement of presynaptic metabotropic glutamate receptors for the maintenance of locomotion," *J. Neurosci.*, vol. 22, pp. 3692–3699, 2002.
- [13] N. Schwartz and S. Alford, "Physiological activation of presynaptic metabotropic glutamate receptors increases intracellular calcium and glutamate release," *J. Neurophysiol.*, vol. 84, pp. 415–427, 2000.

Asynchronous BCI and Local Neural Classifiers: An Overview of the Adaptive Brain Interface Project

José del R. Millán and Josep Mouriño

Abstract—In this communication, we give an overview of our work on an asynchronous brain–computer interface (where the subject makes self-paced decisions on when to switch from one mental task to the next) that responds every 0.5 s. A local neural classifier tries to recognize three different mental tasks; it may also respond "unknown" for uncertain samples as the classifier has incorporated statistical rejection criteria. We report our experience with 15 subjects. We also briefly describe two brain-actuated applications we have developed: a virtual keyboard and a mobile robot (emulating a motorized wheelchair).

Index Terms—Asynchronous protocol, brain-actuated applications, brain–computer interfaces (BCIs), electroencephalogram (EEG), local-neural classifier.

I. INTRODUCTION

Over the last seven years, our brain–computer interface (BCI) laboratory, in cooperation with the Institute Santa Lucia, Rome, Italy, and

Manuscript received July 19, 2002; revised April 2, 2003. This work was supported in part by the European ESPRIT Program under LTR Project 28193-ABI. The work of J. d. R. Millán was supported by the Swiss National Science Foundation through the National Centre of Competence in Research (NCCR) on "Interactive Multimodal Information Management (IM2)."

J. d. R. Millán was with the Joint Research Centre, European Commission, I-21020 Ispra (VA), Italy. He is now with the Dalle Molle Institute for Perceptual Artificial Intelligence, CH-1920 Martigny, Switzerland (e-mail: jose.millan@idiap.ch).

J. Mouriño was with the Joint Research Centre, European Commission, I-21020 Ispra (VA), Italy. He is now with the Centre de Recerca en Enginyeria Biomèdica, Universitat Politècnica de Catalunya, E-08028 Barcelona, Spain (e-mail: jmourino@eic.icnet.es).

Digital Object Identifier 10.1109/TNSRE.2003.814435

the Computational Engineering Laboratory of the Helsinki University of Technology, Finland, has developed a portable BCI, called *Adaptive Brain Interface (ABI)*, based on the on-line analysis of spontaneous electroencephalogram (EEG) signals measured with eight scalp electrodes and able to recognize three mental tasks. Our approach relies on an *asynchronous* protocol where the subject decides voluntarily when to switch between mental tasks and uses a simple *local neural classifier* to recognize, every 0.5 s, the mental task on which the subject is concentrating [1]. ABI is being used to operate two brain-actuated devices: a virtual keyboard and a mobile robot (emulating a motorized wheelchair) [2]–[4].

Like some of the other BCIs reported in the literature, our BCI is based on the analysis of EEG signals associated with spontaneous mental activity. In particular, we look at local variations of EEG over several cortical areas related to different cognitive mental tasks such as imagination of movements, arithmetic operations, or language. The approach aims at discovering EEG patterns embedded in the continuous EEG signal and associated with different mental states [1], [5], [6]. It applies machine-learning techniques to train the classifier and follows a *mutual learning* process where the user and the brain interface are coupled and adapt to each other [1], [6], [7]. This accelerates the training process. In the presence of feedback, our subjects achieve good performance in just a few hours of training. Analysis of learned EEG patterns confirms that for a subject to operate satisfactorily his/her personal BCI, the personal BCI must fit the individual features of the subject [1], [8].

Most BCIs are based on synchronous protocols where the subject must follow a fixed repetitive scheme to switch from one mental task to the next [7], [9], [10]. In these synchronous BCI systems, the EEG phenomena to be recognized are time-locked to a cue, and a trial typically lasts from 4 to 10 s or longer. In contrast, ABI and a few other systems rely on asynchronous protocols in which the subject makes voluntary, self-paced decisions on when to stop performing a mental task and when to start the next one [1], [11]. This makes the system very flexible and natural to operate and yields rapid response times (e.g., 0.5 s in our case).

Typically, EEG-based BCIs make binary decisions as they seek to recognize two different mental states and reach accuracy levels that, in general, are around 90%. ABI achieves error rates below 5% for three mental tasks, while correct recognition is 70% (or higher). In the remaining cases (around 20%–25%), the classifier doesn't respond, since it considers the EEG samples as uncertain. The incorporation of rejection criteria to avoid making risky decisions is an important concern in BCI. From a practical point of view, a low classification error is a critical performance criterion for a BCI; otherwise users can become frustrated and stop utilizing the interface. The system of Roberts and Penny [6] applies Bayesian techniques for rejection purposes.

The classification rates of our system, together with the number of recognizable tasks (3) and the 0.5-s response times, yields a theoretical maximum transmission rate of approximately 2.0 b/s for our system. However, as will be discussed in the following, this bit rate was rarely achieved in practice for long periods.

The use of statistical rejection criteria also helps to deal with an important aspect of a BCI, namely "idle" states where the user is not involved in any particular mental task. In an asynchronous protocol, idle states appear during the operation of a brain-actuated device, while the subject does not want the BCI to carry out any action. Although the neural classifier is not explicitly trained to recognize those idle states, the BCI can process them adequately by giving no response.