Neural Networks 22 (2009) 305–315

Contents lists available at ScienceDirect

Neural Networks

journal homepage: www.elsevier.com/locate/neunet

2009 Special Issue

Exploiting co-adaptation for the design of symbiotic neuroprosthetic assistants

Justin C. Sanchez [a,b,c,*], Babak Mahmoudi [c], Jack DiGiovanna [c], Jose C. Principe [c,d]

[a] Department of Pediatrics, Division of Neurology, University of Florida, Gainesville, FL 32610, United States
[b] Department of Neuroscience, University of Florida, Gainesville, FL 32610, United States
[c] Department of Biomedical Engineering, University of Florida, Gainesville, FL 32610, United States
[d] Department of Electrical and Computer Engineering, University of Florida, Gainesville, FL 32610, United States

ARTICLE INFO

Article history:
Received 31 January 2009
Received in revised form 9 March 2009
Accepted 21 March 2009

Keywords:
Brain–machine interface
Brain–computer symbiosis
Reinforcement learning
Perception–action cycle
Tool use

ABSTRACT

The success of brain–machine interfaces (BMI) is enabled by the remarkable ability of the brain to incorporate the artificial neuroprosthetic 'tool' into its own cognitive space and use it as an extension of the user's body. Unlike other tools, neuroprosthetics create a shared space that seamlessly spans the user's internal goal representation of the world and the external physical environment enabling a much deeper human–tool symbiosis. A key factor in the transformation of 'simple tools' into 'intelligent tools' is the concept of co-adaptation where the tool becomes functionally involved in the extraction and definition of the user's goals. Recent advancements in the neuroscience and engineering of neuroprosthetics are providing a blueprint for how new co-adaptive designs based on reinforcement learning change the nature of a user's ability to accomplish tasks that were not possible using conventional methodologies. By designing adaptive controls and artificial intelligence into the neural interface, tools can become active assistants in goal-directed behavior and further enhance human performance in particular for the disabled population. This paper presents recent advances in computational and neural systems supporting the development of symbiotic neuroprosthetic assistants.

1. Introduction

The evolution of mankind is intrinsically coupled with the invention and the use of new tools to expand the richness of the interaction with other individuals and with the environment. However, tools have primarily served as passive instruments that enhance the brain–body system and do not shape goal-directed behavior as users express their intent. The concept of a "body schema" as classically used in psychology, neurology, and the cognitive sciences involves the development of specific internal mental structures that represents some aspect of the external world (Maravita & Iriki, 2004). To efficiently develop rich and meaningful interactions with the world our brains are dynamically involved in cyclical motor and sensory scenarios that report the outcomes of behavior (Grossberg, 1982). The "body schema" is the principal enabler of tool use to fulfill many everyday life activities (Johnson-Frey, 2003). Tool use is very unique in our development as a species because it has allowed us to extend the natural reaching space and induce further plastic changes in neural system representation (Holmes et al., 2004; Johnson-Frey, 2004; Maravita & Iriki, 2004). As a consequence, new stimuli which may have been out of reach from the body's extremities became accessible, assimilated and increased the integration of new environmental diversity into our internal representation. Driving a car is the modern archetype example where through training the human adapts to speed, anticipation, and size that are far beyond natural body experiences. However, tools have primarily served as passive instruments that do not share the definition of the goal-directed behavior as users express their intent. Indeed, the relationship between user and the tool is inherently lopsided. On one end, users are intelligent and can use dynamic brain organization and specialization while tools are passive devices that enact commands. We submit that the nature of the interface with tools is one of the main limitations impeding the evolution of a more seamless binding between users and tools (and effectively richer environments); and it is perhaps one of the reasons for the chasm between artificial and natural intelligence, because we do not properly share our expectations with artificial systems. Ideally, interfaces with machines should be as active and bidirectional as the interactions with other human beings or animals where the connection between the user and tool is such that both can experience the unique abilities of the counterpart. This egocentric man–machine interface design methodology is the norm today because the interactions are indirectly controlled,

* Corresponding address: Department of Pediatrics, Division of Neurology, University of Florida, P.O. Box 100296, Gainesville, FL 32610, United States. Tel.: +1 352 846 2180; fax: +1 352 846 2180.
E-mail address: jcs77@ufl.edu (J.C. Sanchez).

0893-6080/$ – see front matter © 2009 Elsevier Ltd. All rights reserved.
doi:10.1016/j.neunet.2009.03.015

communicate in a unidirectional manner, lack the ability to operate on the cognitive level of the user, and they are not adaptable.

Cognitive and computational neuroscience has utilized the universal computing power of Turing/Von Neumann machines to implement models of cognition. Perhaps fuelled by Newell's work (Newell, 1990), cognitive architectures have implemented first principles that their designers believe are relevant for intelligent behavior. Two of the better known models in the engineering community are the ACT-R (Anderson, 1993) and SOAR (Laird et al., 1987) and they evolved from pure tools to model and simulate cognitive processes to general architectures that also support direct interaction with the external world through sensory inputs (Bugajska et al., 2002; Jones et al., 1999). When the interaction with an unknown stochastic world becomes center stage, different principles are required. Using ideas from Markov Decision Processes (MDPs), Weng proposed the Incremental Hierarchical Discriminant Regression (IHDR), which is a family of models of different complexity having at the core a self-aware self-effecting architecture (Weng & Hwang, 2006). Another approach closer to the biological reality has used neural dynamics and neural network principles to model brain subsystems as exemplified by the K set hierarchy (Freeman, 1975; Kozma & Freeman, 2009), action networks for the frontal cortical loop (Taylor & Taylor, 1999), working memory (Taylor & Taylor, 2000) the thalamic cortical loop (Hecht-Nielsen, 2007), visomotor transformations (Jeannerod et al., 1995), language (Arbib, 2005), dynamics of perception (Carpenter & Grossberg, 2003), up to consciousness (Edelman, 1990).

An area of engineering that has benefited from all this work has been robotic research because of the importance of autonomous behavior. In the last 10 years, several subfields in robotics have emerged, from behavior based robotics (Brooks, 1999), evolutionary robotics (Nolfi & Floreano, 2000), intentional robotics (Kozma & Fukuda, 2006), developmental robotics (Schmidhuber, 2006) and brain based systems (Krichmar & Edelman, 2005), just to name a few. These systems are being designed more and more using neurobiology knowledge. However, we submit that an important factor that is missing in robotics is the ability to interact with the human brain. The appeal is that these advanced robots already have the computational power, the sensors, and sophisticated architectures for processing and reasoning, but what is missing is a paradigm for co-adaptation with humans. This will be immensely important for neural rehabilitation and will open a new window for symbiotic human machine research. We believe that it is possible to establish a direct communication channel between the user's brain and the machine with the goal of sharing the perception–action cycle of the user. This paper presents a new framework and experimental results that illustrate symbiosis between biological and artificial systems. In Section 2, we will briefly present the state of the art in brain–machine interfaces. Section 3 discusses the architectural prerequisites for co-adaptation and Section 4 develops how to deliver such requirements. Section 5 presents our experimental work on co-adaptive brain–machine interfaces, and Section 6 concludes the paper.

2. Review of brain–machine interface research

Brain–machine interfaces are creating new pathways to interact with the brain and they can be roughly divided into four categories: the sensory BMIs which substitute sensory inputs (like visual (Chelvanayagam et al., 2008; Zrenner, 2002) or auditory (Miller et al., 1995; Nie et al., 2006; Rouger et al., 2007) and are the most common (120,000 people have been implanted worldwide with cochlear implants)); the motor BMIs that substitute parts of the body to convey intent of motion to prosthetic limbs; the cognitive BMIs that repair communication between brain areas such as the hippocampus (Berger et al., 2001) that mediates short term to long term memories; and the clinical BMIs that stimulate specific brain areas to repair normal function, such as deep brain stimulation for Parkinson's disease (Lozano & Mahant, 2004) or to avoid or abort epileptic seizures (Ludvig et al., 2005). We will concentrate this review on motor BMIs which are revolutionizing the way paralyzed users interact with the environment because they offer a direct link between the brain and a tool that interacts with the environment, bypassing the body to express intent (Donoghue, 2002; Nicolelis, 2003; Sanchez & Principe, 2007). Within the motor BMIs there are three basic types: the trajectory BMIs, the goal driven BMIs, and the command and control BCIs. Trajectory BMIs as the name indicates learn how to control a robotic arm to follow a trajectory. They are basically signal translators to actuate prosthetics; they collect firing patterns of dozens to hundreds of neurons in the motor cortex and surrounding areas to decode the user's intent expressed in the neural signal time structure. Since the pioneering work of Chapin in 1999 that showed this was possible in real-time, trajectory BMIs are probably the most popular (Chapin et al., 1999). The goal driven BMIs extract the location in space for the intended movement from a set of predetermined targets using electrodes in the parietal cortices and they can be used for high level coarse command for robots to implement the motion to the desired location in space (Shenoy et al., 2003). The command and control brain–computer interfaces (BCIs) utilize multiple electrodes placed on the scalp (or directly over the cortex) to translate signature of cognition related to imagined movement, expectation, or simply an imagined set of brain states that can control cursors on a screen for action selection (Wolpaw et al., 2000).

Many groups have conducted research in trajectory BMIs and the approach has been strongly signal processing based without much concern to incorporate the design principles of the biologic system in the interface. The implementation path has either taken an unsupervised approach by finding causal relationships in the data (Buzsáki, 2006), a supervised approach using (functional) regression (Kim et al., 2006), or more sophisticated methods of sequential estimation (Brown et al., 2004) to minimize the error between predicted and known behavior. These approaches are primarily data-driven techniques that seek out correlation and structure between the spatio-temporal neural activation and behavior. Once the model is trained, the procedure is to *fix the model parameters* for use in a test set that assumes stationarity in the functional mapping. Some of the best known linear models that have used this architecture in the BMI literature are the Wiener filter (FIR) (Serruya et al., 2002; Wessberg et al., 2000) and Population Vector (Helms Tillery et al., 2003), generative models (Moran & Schwartz, 1999; Taylor et al., 2002; Wu et al., 2002), and nonlinear dynamic neural networks (a time delay neural network or recurrent neural networks (Chapin et al., 1999; Gao et al., 2003; Sanchez et al., 2002)) models that assume behavior can be captured by a static input–output model and that the spike train statistics do not change over time. While these models have been shown to work well in specific scenarios, they carry with them these strong assumptions and will likely not be feasible over the long term.

The success of BMI control is due in part to the remarkable ability of the brain to incorporate the artificial neuroprosthetic 'tool' into its own cognitive space (Velliste et al., 2008) and use it as an extension of the biologic body (Holmes et al., 2004). If we analyze in detail the trajectory BMI paradigm it still follows the egocentric approach of privileging the user versus the computer (hereafter referred to as a *computer agent*) controlling the robot. We can argue that from an engineering perspective this is fine, as long as the combined system solves the task. Unfortunately, there have been difficulties in translating the trajectory paradigm to clinical environments because it requires too much information from the setting, namely the existence of a desired trajectory to train the decoding algorithms. Quadriplegics, the intended clinical group for trajectory BMIs, cannot move so there is no trajectory in real settings and the current solutions are rather poor. Moreover, with continuous neural interface use, the neural representation supporting such behavior will change (Carmena et al., 2003). It has been shown unequivocally in animals and humans that intelligent users can switch to brain control seamlessly (Carmena et al., 2003; Hochberg et al., 2006). However, it has also been shown that the time that it takes to achieve a certain level of "mastery" of the prosthetic device can be extremely slow especially when the details of the dynamics of control are unknown to the user. From a behavioral perspective, even simple issues of scale (i.e. dynamic range of reaching) can create problems for input–output models if the full range of values was not encountered during training (Moody, 1992). Even with the great adaptability of the user's brain, it can take significant time for the performance to recover. To contend with these issues, it has been suggested by a few groups that adaptability of the interface is a critical design principle for engineering the next generation BMIs (del R Millan, 2003; Helms Tillery et al., 2003; Taylor et al., 2003). In these studies, the concept of adaptability typically refers only to very detailed aspects of the signal translation to include automatic selection of features, electrode sites, or training signals (Birbaumer et al., 2000; McFarland et al., 2006). We submit that this concept of adaptability does not go far enough, because it is unable to raise the level of the bidirectional dialogue with the user and still does not provide opportunities to build intelligence into the tool to model the user's goals.

We present here a new computational architecture that is not only adaptable but also intelligent because it can serve as an assistant to the user to facilitate the control and can serve as an equalizer to share the burden of learning the rules of control. An intelligent system is defined here as a system that uses a model of the external world in its interaction with it. A conceptual drawing of this new class of interface is presented in Fig. 1. Here, the user interacts directly via their neural signals with a prosthetic arm. However, the decoder of the neural signals also shares the same goals as the user. It is precisely in the sharing of bidirectional goal-directional behavior, that the decoder can act as an intelligent assistant that co-evolves with the user to solve tasks. This aspect of goal-directed behavior has been overlooked in the present BMI computational paradigms, which assign rudimentary roles to the computer that is controlling the robotics, which constrains the type of tasks, the level of attainable performance and the time required for proper training. This framework is a significant departure from other BMI architectures because it implements the interface at a much higher functional level.

3. Minimal prerequisites for intelligent neuroprosthesis

The design of a new framework to transform BMIs begins with the view that intelligent tools emerge from the process where the user and tool cooperatively seek to maximize goals while interacting with a complex, dynamical environment. Emergence as discussed here and in the cognitive sciences depends on a series of events or elemental procedures that promote specific brain or behavioral syntax, feedback, and repetition over time (Calvin, 1990); hence, the sequential evaluative process is always ongoing, adheres to strict timing and cooperative–competitive processes, and is very different from the notion of static computational methods. With these elemental procedures, intelligent motor control and more importantly goal-directed behavior can be built with closed-loop mechanisms which continuously adapt internal and external antecedents of the world, express intent though behavior in

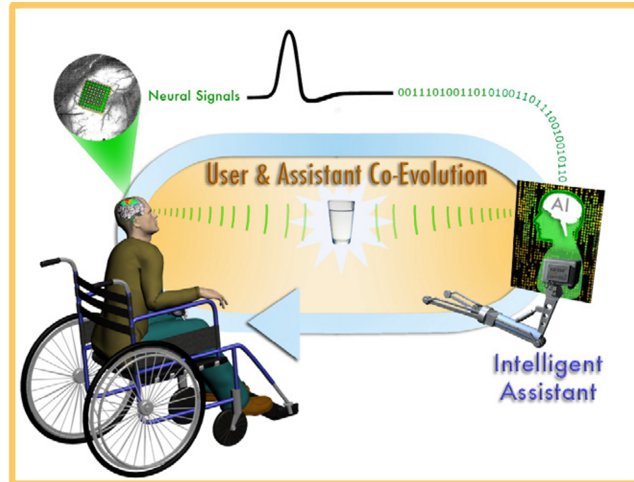

Fig. 1. Conceptual drawing of an intelligent neuroprosthesis system.

Table 1
User–neuroprosthetic prerequisites for co-adaptation.

	User	Neuroprosthetic tool
Representation	Brain states	Environmental states
Valuation	Goal-directed	Goal-directed
Action selection	Neuromodulation	Competition
Outcome measures	Internal reward expectation	Predicted error
Learning	Reinforcement based	Reinforcement based
Co-adaptation	Dynamic brain organization	Optimization of parameters

the environment, and evaluate the consequences of those behaviors to promote learning. Collectively these components contribute to forming a Perception–Action Cycle (PAC) which plays a critical role in organizing behavior in the nervous system (Fuster, 2004). This form of adaptive behavior relies on continuous processing of sensory information that is used to guide a series of goal-directed actions. Most importantly, the entire process is regulated by external, environmental and internal neurofeedback, which is used to guide the adaptation of computation and behavior. The PAC in goal-directed behavior provides several key concepts in the formation of a new framework for BMI. However, unlike the PAC that is central in the animal interaction with the world, the PAC in a co-adaptive BMI will be distributed between the user and the computer agent. Next, we introduce the prerequisites for modifying the PAC to incorporate two intelligent entities.

In order to symbiotically link artificial intelligent tools with neural systems, a new set of protocols must be derived to enable and empower dialogue between two seemingly different entities. A minimal set of six prerequisites given in Table 1 describe the essential computation that is required to enact a symbiotic PAC. These prerequisites are based on concepts considered to be key in value-based decision making (Rangel et al., 2008). Unique to the development of intelligent BMIs is that the user and neuroprosthetic each have their own perspective and contribution to each prerequisite as described below.

Representation: Internal to the user, the spatio-temporal structure of neural activation forms a coding of intent for action in the external world. Understanding the properties of information transmission in the code and determining how this information translates into commands of the motor system as a whole is one of the cornerstones of BMI development. At any given moment, the neural code can be sampled as a brain state defined as the vector of