A Prospectus on the Obstacles Inhibiting the Implementation of Advanced Artificial Neural Systems - Part 1

Marko A. Rodriguez[1], Michael I. Ham[1], Vadas Gintautas[1], and Benjamin S. Kunsberg[2]

[1]Center for Nonlinear Studies - Los Alamos National Laboratory [2]Computational Neuroscience - New Mexico Consortium

Introduction

- The functional capabilities of advanced neural systems such as the mammalian brain are immense. These systems have evolved over millions of years to perform specialized tasks both quickly and efficiently.
- The functional capabilities of such neural systems are very different from those of the modern day computer. However, given the current state of knowledge in both biological neuroscience and theoretical computing, it is hypothesized that the computer, being Turing complete, could one day be used to effectively model an advanced neural system.
- In order to implement such systems, there exist a set of core requirements and impediments that must be achieved and overcome.
- This poster will focus on three advances required in the biological sciences and three advances required in the computational sciences.

1. Biological Obstacles

1.1 Measurement Resolution

A detailed understanding of neural interaction depends on precisely measuring the activity patterns of many individual neurons simultaneously and non-invasively. The current techniques that afford high temporal and spatial resolution measure a relatively small number of cells and are often invasive (e.g. micro electrodes). Current noninvasive techniques only provide aggregate data from many neurons with low temporal and spatial resolution (e.g. MRI, FMRI, MEG, EEG). A significant advance would combine the ability to image large portions of the brain with detailed temporal firing patterns of individual neurons.

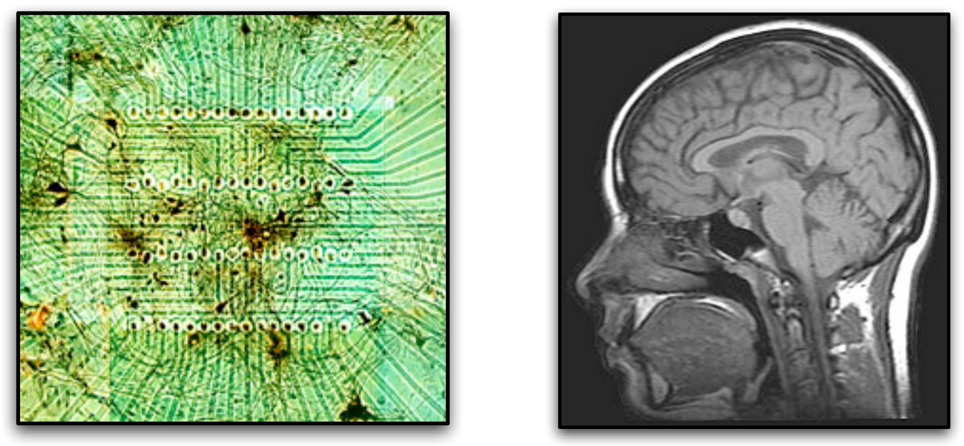

Figure 1: *Neurons on a microelectrode array (left – photo courtesy of the Center for Network Neuroscience, University of North Texas) and an FMRI image of the brain (right – photo courtesy of NASA).*

1.2 The Relationship Between Input and Memory

Neuronal interactions have been studied extensively, both in vitro and in vivo. Yet the mechanisms that neurons in large groups use to interpret and represent data in higher order brain regions remain unknown. The hypothesized existence of "grandmother cells" [?], neurons which respond to many different representations of one's grandmother (or Fluffy the dog), suggests that the brain represents objects in a transformation and modality invariant manner. For example, the smell of Fluffy's wet fur can trigger the thought of Fluffy as can the sound of his bark. While a great deal is known about neurons that are closely linked to sensory data (for example, simple/complex cells), little is known about how these contribute to the creation and triggering of more abstract representations. Additionally, the role that memories play in data analysis is an important question that must be addressed in order to build accurate models of the brain. Memories are very important for allowing animals to learn from their environment, but no experimental results have demonstrated the interaction between memories and incoming sensory data.

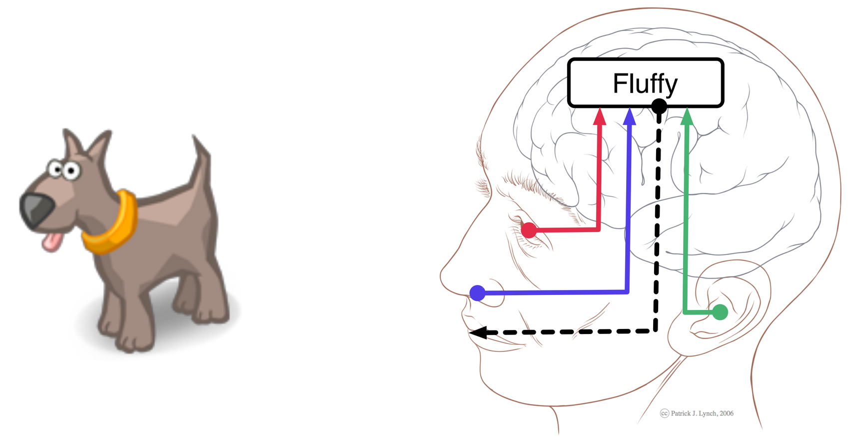

Figure 2: *The sight, smell, and sound (input) of a dog all coalesce into an abstract representation of a dog. Output corresponding to this abstract dog can exist as speech, writing, or in the imagination. The merger of categorization from the various sense modalities is poorly understood.*

1.3 Purpose of Idling Activity

Electrophysiological activity is always present in neural systems. The role of such activity is hypothesized to range from development and maintenance [?] to anticipatory states [?] that help animals make rapid decisions. It is likely that these spontaneous dynamic interactions perform many important tasks, but until this is better understood, it cannot be easily included in neural models. Therefore, a better understanding of this activity could lead to more biologically accurate models.

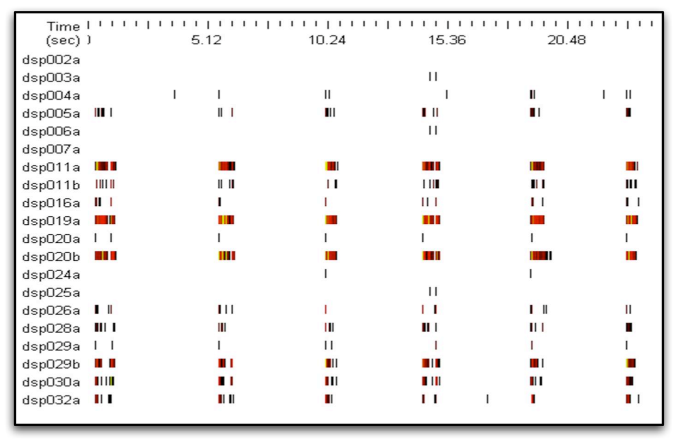

Figure 3: *Spontaneous activity in a cultured neural network. Labels of identified neurons form the vertical axis, while time is represented by the horizontal axis. Marks denote the timestamps of neuronal firing events; brighter colors indicate faster spiking rate.*

2. Computation Obstacles

2.1 Symbolic Functionality in a Sub-Symbolic System

There are two general approaches to artificial intelligence: sub-symbolic (connectionism) and symbolic (knowledge representation and reasoning). The sub-symbolic level is concerned with the implementation of biologically-plausible models of individual neurons and their connectivity into artificial neural networks [?]. The assumption is that by modeling the fundamental components of a neural system, high-level functionality will emerge as these components are connected to form a larger system. The symbolic approach deals with higher-order structures of cognition such as objects, their relations to one another, and cognitively plausible algorithms for reasoning over such structures [?]. While the sub-symbolic approach yields a biologically-plausible implementation of the processing capabilities of relatively simple neural systems, the symbolic approach implements the functionality of advanced neural systems.

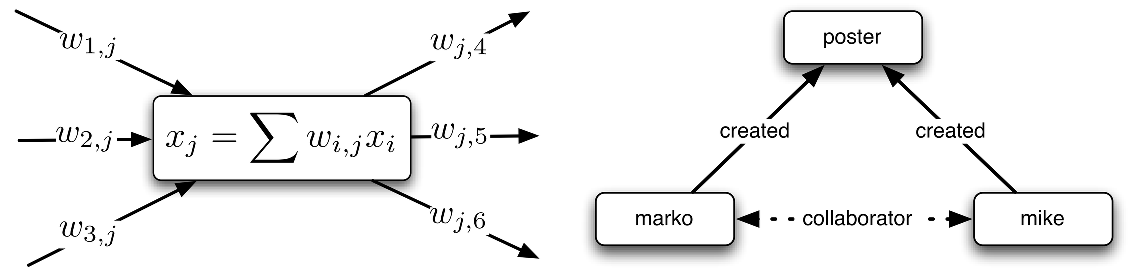

Figure 4: *The representational schemas utilized in sub-symbolic (left) and symbolic (right) artificial intelligence. The sub-symbolic approach models the simple processes of neurons within a network, as they perform a larger, more complex computation. The symbolic approach represents labeled relationships between "things", where the labels denote the space of possible logical inferences.*

Classification and recognition is typical of sub-symbolic systems. State of the art synthetic visual systems have been able to implement the V1, V2, and V4 areas of the visual cortex both structurally and functionally [?]. Such systems can discriminate objects in an image – for example, they can classify animal and non-animal images. On the other hand, symbolic systems, while not based on a biologically-plausible substrate, utilize reasoning algorithms to perform more abstract inferences such as "Fluffy is a dog $\wedge$ a dog is an animal" $\longrightarrow$ "Fluffy is an animal." Advances in cognitively realistic reasoning [?] within biologically-plausible architectures is a necessary requirement for advanced artificial neural systems.

2.2 From Intelligent Design to Neurogenesis

Most artificial neural systems are designed with neurons and connections between them as the basic building blocks. Due to this level of granularity, very few systems have been designed that utilize feedback and spike timing during computation. This is because it is very difficult for a system designer to manage recurrent systems and ensure synchronized timing. However, feedback and spike timing appear to significantly contribute to the computations carried out by advanced neural systems. Such complexity calls for a new design philosophy that is predicated on the principle of growth and experience. Thus neurogenesis, embryonic development, and situated and embodied cognition must take center stage to hard-wired connectivity and supervised learning algorithms.

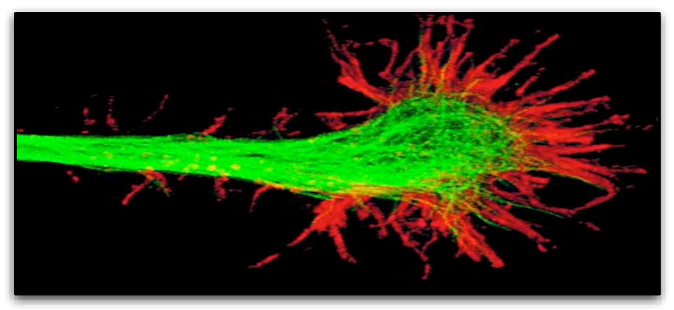

Figure 5: *An image of a neural growth cone (red) guiding an axon (green) in 3D space as it searches for synaptic connections (photo courtesy of Paul Letourneau, University of Minnesota).*

2.3 Distributed Representation and Processing

The Los Alamos National Laboratory is currently building a synthetic visual cortex on the Roadrunner petaflop-scale supercomputer. The implementation of an advanced neural system (including those yielding functionality beyond the human) is restricted by the amount of computational resources that can be allocated. More computational resources could be allocated if the neuroscience community takes inspiration from standards-oriented disciplines such as astronomy and the World Wide Web.

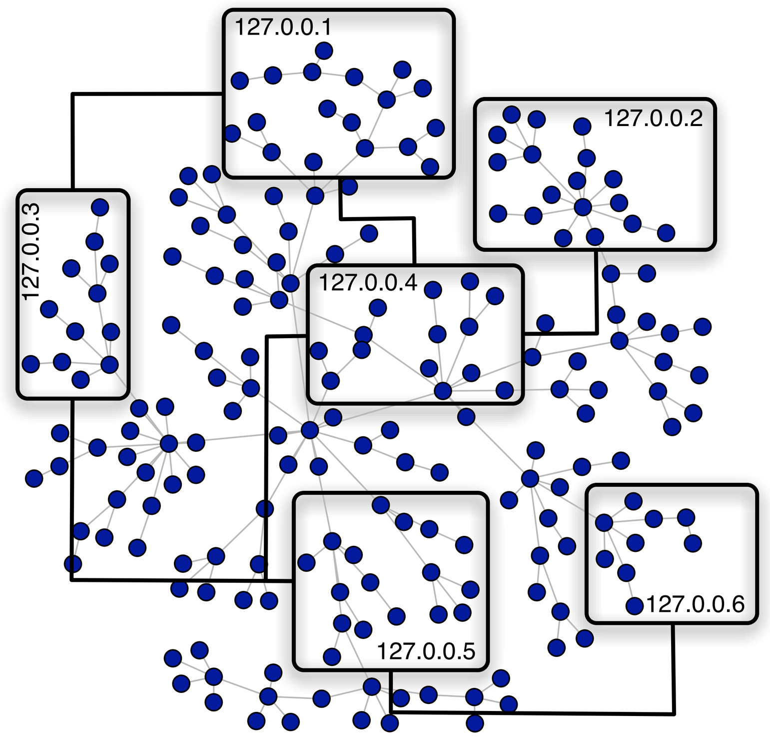

Figure 6: *Implementing an advanced neural system requires massive computational resources. It may benefit the neuroscience community to derive standards for the representation and distribution of a neural model.*

It is possible to provide a distributed representation of an advanced neural system across computers worldwide. The "web of data" is an emerging data representation paradigm that is being developed by the World Wide Web community. Instead of only representing documents and images within the URI address space of the Web, every minutia of data can be represented and thus, the Web will serve as a massive global database. The underlying data structure of the web of data is a multi-relational network (that is, a directed labeled graph). With this flexible data model, it is possible to create a distributed representation of an advanced neural system. Moreover, within the same URI address space as other data, such neural systems could contribute novel, non-mammalian, neural-based information processing to the world's digital data.

LA-UR-09-00043 – Research conducted through the Synthetic Cognition through Petascale Models of the Primate Visual Cortex project, LDRD-2009006DR, funded by the Los Alamos National Laboratory.

Decade of Mind IV, Albuquerque, New Mexico January 13-15, 2009