Available online at www.sciencedirect.com
ScienceDirect
Cognitive Systems Research 12 (2011) 93–101
Cognitive Systems RESEARCH
www.elsevier.com/locate/cogsys

Conflict resolution and learning probability matching in a neural cell-assembly architecture

Action editor: Richard P. Cooper

Roman V. Belavkin *, Christian R. Huyck

School of Engineering and Information Sciences, Middlesex University, London NW4 4BT, UK

Available online 26 August 2010

Abstract

Donald Hebb proposed a hypothesis that specialised groups of neurons, called *cell-assemblies* (CAs), form the basis for neural encoding of symbols in the human mind. It is not clear, however, how CAs can be re-used and combined to form new representations as in classical symbolic systems. We demonstrate that Hebbian learning of synaptic weights alone is not adequate for all tasks, and that additional meta-control processes should be involved. We describe an earlier proposed architecture (Belavkin & Huyck, 2008) implementing an adaptive conflict resolution process between CAs, and then evaluate it by modelling the probability matching phenomenon in a classic two-choice task. The model and its results are discussed in view of mathematical theory of learning and existing cognitive architectures.

Keywords: Artificial intelligence; Cognitive science; Neuroscience; Decision making; Intelligent agents; Learning; Bayesian modelling; Computational neuroscience; Human experimentation

1. Introduction

There exists a variety of artificial systems and algorithms for learning and adaptation. Most of them can be classified as sub-symbolic (e.g. Bayesian and connectionist networks) or symbolic systems (e.g. rule-based systems). Known natural learning systems use neural networks, and therefore can be classified as using sub-symbolic computations. A distinguishing feature of the human mind, however, is the ability to use rich symbolic representations and language.

From an information-theoretic point of view, symbols are elements of some finite set that are used to encode discrete categories of sub-symbolic information. They enable communication of information about the environment or a complex problem in a compact form. One obvious benefit is that with language, one can learn not only from one's own experience, but also from experiences of others. The benefits of reading a guidebook before going abroad are obvious.

The duality between sub-symbolic and symbolic approaches has been studied in cognitive science. There exist sub-symbolic (i.e. connectionist), symbolic (e.g. SOAR, Newell, 1990) and hybrid architectures (e.g. ACT-R, Anderson & Lebiere, 1998) for cognitive modelling. These different approaches, however, have not yet explained where the symbols are in the human mind, or how the brain implements symbolic information processing (though see Jilk, Lebiere, O'Reilly, & Anderson, 2008).

It was proposed by Hebb (1949) that symbols are represented in the brain not by individual neurons, but by correlated activities of groups of cells, called *cell-assemblies* (CAs). The cell-assemblies robot project (CABOT) set out to test and demonstrate this idea in an engineering task by building an artificial agent, situated in a virtual environment, capable of complex symbolic processing, and implemented entirely using CAs of simulated neurons. Some of the objectives have already been achieved and reported elsewhere (e.g. Huyck & Belavkin, 2006, 2007, 2008). The architecture and some of these works will be discussed in the next section.

* Corresponding author.
E-mail addresses: R.Belavkin@mdx.ac.uk (R.V. Belavkin), C.Huyck@mdx.ac.uk (C.R. Huyck).

1389-0417/$ - see front matter © 2010 Elsevier B.V. All rights reserved.
doi:10.1016/j.cogsys.2010.08.003

The work described in this paper is concerned with a particular aspect of the project—a stochastic conflict resolution and meta-control mechanism that modulates Hebbian learning to allow for re-use and combination of CAs into new representations, such as learning logical implications (i.e. procedural knowledge). As will be discussed in this paper, this cannot be achieved by using a Hebbian learning mechanism alone. A unique contribution of this work is evaluation of the meta-control mechanism in a cognitive model of the probability matching phenomenon in a two-choice experiment (Friedman et al., 1964). The results suggest that a proposed mechanism is a plausible model. Some neurophysiological studies and hypotheses about the brain circuitry will be discussed supporting the biological plausibility of the architecture.

In the next section, we describe briefly the neural model that is used in our architecture, how simulated neurons form cell-assemblies and how we use them to test the CA hypothesis of symbolic processing. Then we discuss the problem of learning connections between existing CAs. This process is important for learning new symbolic knowledge by re-using and combining existing symbolic representations. In particular, we focus on the problem of learning the correct set of rules from the set of all possible rules connecting existing antecedents and consequents. Here we draw the parallel with the ACT-R conflict resolution mechanism. Using a mathematical theory of stochastic learning, we argue that utility (or reinforcement) and stochastic noise are essential components of this process, and that they are not included in the Hebbian principle for adaptation of synaptic weights. The neural architecture implementing the utility-based stochastic learning of the connections between CAs is explained in Section 4, and its performance is demonstrated in an experiment. Section 5 presents the same architecture simulating the probability matching phenomenon as observed by Friedman et al. (1964), and a comparison with the hybrid model based on the ACT-R architecture is drawn. We then summarise contributions of this work and discuss its potential future development.

2. Cell-assemblies as the basis of symbols

In this section, we outline some of the basic features of the CABOT architecture as well as the CA hypothesis.

2.1. Neural information processing in CABOT

It is widely accepted that human cognition is the result of the activity of approximately 10^{11} neurons in the central nervous system that interact with each other as well as with the outside world via the peripheral nervous system. Biological neurons are complex systems, and they have been modelled with various levels of details (McCulloch & Pitts, 1943, 1952). In our system, we use fatiguing, leaky, integrate and fire (fLIF) neurons.

The ‘integrate and fire’ component is based on the classical idea that the neuron ‘fires’ (or spikes) if its action potential, A, exceeds a certain threshold value θ:

$$y = \begin{cases} 1 & \text{if } A \geqslant \theta \\ 0 & \text{otherwise} \end{cases}$$

The action potential, A, is a function of the integral (inner product) $\langle x, w\rangle = \sum_{i=1}^{k} x_i w_i$ of the stimulus (pre-synaptic) vector $x \in \mathbb{R}^k$ and the synaptic weight vector $w \in \mathbb{R}^k$ of the neuron. Here, $\mathbb{R}^k$ is a k-dimensional Euclidean space, where k is the number of synapses to the neuron. We use binary signals, and therefore x is a k-dimensional binary vector.

The ‘leaky’ property refers to a more complex (non-linear) dependency of the action potential on the pre- and post-synaptic activity:

$$A_{t+1} = \frac{A_t}{d_t} + \langle x_t, w_t\rangle, \quad d_t = \begin{cases} \infty & \text{if fired } (y_t = 1) \\ d \geqslant 1 & \text{otherwise} \end{cases} \tag{1}$$

Thus, the action potential is accumulated over several time moments if the neuron does not fire. Parameter $d \geqslant 1$ allows for some of this activation to ‘leak’ away. This is the LIF model (Maas & Bishop, 2001).

The ‘fatigue’ property refers to a dynamic threshold that is defined as follows:

$$\theta_{t+1} = \theta_t + F_t, \quad F_t = \begin{cases} F_+ \geqslant 0 & \text{if fired } (y_t = 1) \\ F_- < 0 & \text{otherwise} \end{cases} \tag{2}$$

where values F_+ and F_- represent the *fatigue* and fatigue *recovery* rates. Thus, if a neuron fires at time t, its threshold increases, and it is less likely to fire at time $t+1$.

The fatiguing and leaky properties of the neural model allow for a non-trivial dynamics of the system. Repetitive stimulation of excitatory synapses increases the probability of a neuron to fire, even if the weights have small (positive) values. On the other hand, if the neuron fires repetitively, its threshold increases reducing the chance of it firing again. Thus, frequencies of pre- and post-synaptic activities are important factors in our system.

The weights of a neuron adapt according to the following compensatory rule Huyck, 2007:

$$\Delta w_{ij} = \begin{cases} \alpha(1 - w_{ij})e^{W_B - W_i} & \text{if } x_t = 1,\ y_t = 1 \\ -\alpha w_{ij} e^{W_i - W_B} & \text{if } x_t = 1,\ y_t = 0 \end{cases}$$

where $\alpha \in [0, 1]$ is the learning rate parameter, W_B is a constant representing the average total synaptic strength of the pre-synaptic neuron, and W_i is the current total synaptic strength (see Huyck (2007), for details). Note that absolute values of the weights w_{ij} here are in the interval $[0, 1]$, and the rule ensures that the new weight depends on the correlation between the pre-synaptic, x_t, and the post-synaptic, y_t, activities, which is an implementation of the Hebbian principle.

The above described properties are known characteristics of biological neurons, and our model is a compromise