Neurocomputing 73 (2010) 2928–2934

Contents lists available at ScienceDirect

Neurocomputing

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

A Pong playing agent modelled with massively overlapping cell assemblies

Kailash Nadh *, Christian R. Huyck

School of Engineering and Information Sciences, Middlesex University, London, United Kingdom

ARTICLE INFO

Article history:
Received 22 March 2010
Received in revised form
1 July 2010
Accepted 30 July 2010
Communicated by G. Palm
Available online 26 August 2010

Keywords:
Neural networks
Fatiguing LIF neurons
Cell assemblies
Associative memory

ABSTRACT

Cell assemblies (CAs) are central to many higher order cognitive processes such as perception, recognition and recollection. These processes stem from the fundamental cognitive tasks of memorisation and association, which CA models are able to perform with a viable degree of biological realism. This paper describes a virtual agent that uses CAs that emerge from fatiguing leaky integrate and fire neurons via learning from dynamic interaction. Learning is continuous and the topology is biologically motivated. The agent is able to visually perceive, learn and play a simplified game of Pong. It can learn from a user playing the game, or playing on its own. The agent's memories are encoded in the form of overlapping CAs that enable it to generalise its associations to account for previously unseen game moves. The trained agent hits the Pong ball correctly over 90% of the time. This work furthers the understanding of associative memory and CAs implemented in neural systems.

1. Introduction

Associative memory is a fundamental cognitive process. Concepts in memory, and different types of associations between them are acquired via learning. Associations can vary semantically, for instance hierarchical associations, sequential associations, containment associations, spatial associations, and higher level semantic associations. These concepts and associations that form associative memory are critical to cognitive processing. Many connectionist accounts for associative memory exist, but cell assemblies (CAs) [1] provide a biologically and psychologically realistic basis for associative memory.

This paper describes a virtual agent modelled in simulated CAs capable of playing a simplified version of the popular arcade game, Pong. The agent is able to learn to play the game by observing a human play, or on its own. The agent learns to associate input from the environment with actions, thus learning game moves.

The agent is made entirely from fatiguing Leaky Integrate and Fire (fLIF) neurons that have a reasonable resemblance to biological neurons. Unlike many simulations, but like human neurons, learning remains on at all times.

The agent learns by encoding shared, overlapping associative memories. This allows generalisation behaviour to emerge, which further assists the agent in game play by enabling it to carry out actions in novel situations.

The paper is organised as follows: Section 2 overviews CAs; Section 3 discusses the fLIF neural network architecture where CAs emerge; Section 4 details the simulation and findings; and Section 5 discusses the impact of the simulation and highlights the findings.

2. Background

Human associative memory is a remarkable process. Throughout life, concepts continue to be acquired, learnt and associated. Any given concept is associated with many other concepts, and retrieval of an associated concept can be based on a combination of a range of base concepts with a range of contexts. Many associative memory models exist, e.g. [2–4], but Hebb's CA theory [1] provides an account that is supported by biological and psychological evidence. CAs exhibit dynamics that provide a unified explanation for long term memory and various short term memories as opposed to higher level box models of memory. There is extensive evidence that CAs are the basis of human associative memory and many other cognitive phenomena [5–10] and they have been used in computational models of associative memory [11–14].

2.1. Cell assemblies

CAs are reverberating circuits of neurons that form the neural basis of concepts. Hebb's CA theory postulates that objects, ideas, stimuli and even abstract concepts are represented in the brain by the simultaneous activation of large groups of neurons with high mutual synaptic strengths [1,11]. CAs are learnt by a Hebbian

* Corresponding author.
E-mail addresses: k.nadh@mdx.ac.uk (K. Nadh), c.huyck@mdx.ac.uk (C.R. Huyck).

0925-2312/$ - see front matter © 2010 Elsevier B.V. All rights reserved.
doi:10.1016/j.neucom.2010.07.013

learning rule, whereby modifications in the synaptic transmission efficacy are driven by correlations in the firing activity of pre-synaptic and post-synaptic neurons [15]. When external stimuli are presented to a network, synaptic strength between neurons is adjusted so as to gain more strength if they undergo repeated and persistent activation or firing, gradually assembling themselves into a CA. A CA thus formed is bound to the object that generated the stimuli, a neural representation of that object; the behaviour of the CA causes psychological behaviour.

Such formation of CAs accounts for long term memory. When a stimulus similar to previously experienced stimuli occurs, it may excite a sufficient number of neurons of an existing CA to cause the spreading of activation within the CA, activating it fully due to recurrent activity and high mutual synaptic strength. The CA can remain active even after the stimulus is removed and this reverberating behaviour accounts for many kinds of short term memories. Ambiguous stimuli can cause contention between similar CAs, where lateral inhibition between those CAs causes competition, with one eventually winning. Thus, the CA hypothesis provides a functional account for a biologically viable associative memory mechanism.

There is extensive evidence of CAs in mammals based on a range of recording techniques and experimental paradigms (see [5,7] for reviews). CAs can also account for various psychological phenomena such as sensing [16] and determining action [17], besides different types of associative memory [18].

2.2. CAs and associative memory

Even though CAs account for memory formation, their precise neural dynamics are far from perfectly understood. In auto-associative memories, an initial state is allowed to settle into a stored memory, allowing subsequent noisy input to retrieve a stored pattern. The Hopfield Model, which is a network of units that are well connected with bidirectional weighted connections, that are used to store a set of binary patterns (typically using a Hebbian calculation), illustrates this property [19]. When an initial set of neurons is ignited, in a discrete version of the system, activation spreads through the system based on the weighted connections. In most cases, the system will settle into a stable state with no neurons switching between on and off. If the input pattern is close to a stored pattern, it will settle into that pattern's state, thus functioning as a content-addressable memory or an auto-associative memory.

Neurons may also belong to more than one CA. Hopfield patterns that share on-bits are models of CAs that share neurons. As mentioned in Section 2.1, neurons in a network may belong to different CAs, and if they are repeatedly co-activated by different versions of the same stimulus, they tend to become associated [1]. This is based on the notion that events that occur together repeatedly should somehow belong together. Every time these events occur in conjunction, they drive certain subgroups of neurons to fire in correlation, resulting in the association of the respective events [11]. A more complete review of CA based associative memory models is [20].

Repeated co-activation of neurons can lead to the formation of CAs. Similarly, repeated co-activation of multiple CAs result in the formation of multiple and sequential associations, and even new CAs. When an external stimulus activates a CA, it may excite neurons shared with a different CA that is not directly stimulated, activating it. This forms the rudimentary, neural level explanation of associative memory.

In prior work, associative memory has been explored with orthogonal and overlapping CAs. Orthogonal CAs were used to encode spatial cognitive maps, many-to-many, and context sensitive associations [18]. With orthogonal CAs, a neuron belongs to at most one CA, but with overlapping CAs, a neuron may belong to several CAs. Learnt overlapping CAs can form hierarchical categories from instances of individuals [21]. The simulations mentioned in this paragraph use a similar neural and topological architecture to the one described below.

3. The fLIF CA architecture

A computational model based on fLIF neurons, using a Hebbian learning mechanism can self-organise to form CAs. Similar to many existing models, the basic architecture of such a mechanism, explained below, is a simplification of the mammalian neural architecture.

3.1. The fLIF neuron

The fLIF neuron model is an extension of the Leaky Integrate and Fire (LIF) model [22,23]. fLIF neurons share many attributes with their biological counterparts. Like the biological neuron, the fLIF neuron integrates coincident pre-synaptic potentials until a critical threshold is reached. On exceeding the threshold, the neuron produces an action potential, or *fires*. This potential further propagates via the neuron's axonal terminal to incident post-synaptic neurons, while the firing neuron loses its activation. The neuron leaks potential if the firing threshold is not attained for prolonged periods. This leaking behaviour of fLIF neurons is similar to that of the biological neuron. This can be represented as follows:

The activation A of a neuron i at time t is

$$A_i(t) = \frac{A_i(t-1)}{\delta} + \sum_{j \in V_i} w_{ji} \tag{1}$$

The current total activation A is the sum of incoming activation and remnant activation from the previous time step $t-1$ divided by decay factor $\delta > 1$. The incoming activation is the sum of total neurons that fired at $t-1$ of all neurons $j \in V_i, V_i$ being all pre-synaptic neurons of i that fired at $t-1$, weighted by the connection from neuron j to i.

When the accumulated activation A exceeds the threshold θ, the neuron fires, losing its potential and thereby resetting A to zero. Firing is a binary event, and activation of w_{ji} is sent to all neurons i to which the firing neuron j has a connection.

Fatiguing causes the threshold to be dynamic, $\theta_{t+1} = \theta_t + F_t$. F_t is positive (F_+) if the neuron fires at t and negative (F_-) otherwise. An increase in the threshold causes the total amount of activation required for neuron firing to increase. Hence, successive firing reduces the ability of the neuron to fire. Similarly, the threshold decreases with each step the neuron does not fire, but is never less than the original threshold.

3.2. Learning

Learning in the fLIF network is dictated by a correlatory Hebbian learning rule [24], whereby synaptic connection weights are modified based on the following equation:

$$\Delta^+ w_{ij} = (1 - w_{ij}) * \lambda \tag{2}$$

$$\Delta^- w_{ij} = w_{ij} * -\lambda \tag{3}$$

w_{ij} is the synaptic weight from neuron i to j and λ is the learning rate. During each step, weights change based on the state of pre-synaptic and post-synaptic neurons. If both neurons fire, the weights increase as per the Hebbian rule (Eq. (2)). If only the