BioSystems 136 (2015) 80–89

Contents lists available at ScienceDirect

BioSystems

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

On learning time delays between the spikes from different input neurons in a biophysical model of a pyramidal neuron

Achilleas Koutsou[a], Guido Bugmann[b], Chris Christodoulou[a,*]

[a] Department of Computer Science, University of Cyprus, P.O. Box 20537, 1678 Nicosia, Cyprus
[b] School of Computing, Electronics and Mathematics, Plymouth University, Drake Circus, PL4 8AA Plymouth, United Kingdom

ARTICLE INFO

Article history:
Received 2 March 2015
Received in revised form 21 August 2015
Accepted 22 August 2015
Available online 01 September 2015

Keywords:
Dendritic propagation delays
Coincidence detection
Membrane noise
Synaptic scaling

ABSTRACT

Biological systems are able to recognise temporal sequences of stimuli or compute in the temporal domain. In this paper we are exploring whether a biophysical model of a pyramidal neuron can detect and learn systematic time delays between the spikes from different input neurons. In particular, we investigate whether it is possible to reinforce pairs of synapses separated by a dendritic propagation time delay corresponding to the arrival time difference of two spikes from two different input neurons. We examine two subthreshold learning approaches where the first relies on the backpropagation of EPSPs (excitatory postsynaptic potentials) and the second on the backpropagation of a somatic action potential, whose production is supported by a *learning-enabling* background current. The first approach does not provide a learning signal that sufficiently differentiates between synapses at different locations, while in the second approach, somatic spikes do not provide a reliable signal distinguishing arrival time differences of the order of the dendritic propagation time. It appears that the firing of pyramidal neurons shows little sensitivity to heterosynaptic spike arrival time differences of several milliseconds. This neuron is therefore unlikely to be able to learn to detect such differences.

1. Introduction

The detection of sequences of sensory inputs with specific short time delays (e.g., velocity sensitive motion detection or decoding of the firing of Geniculate lagged cells, see Saul, 2008) is a function of biological systems. Sequence detectors are usually modelled as coincidence detectors that exploit appropriate delays of asynchronous individual input to cause a coincidence after the arrival of the last input of the sequence (see for example Branco et al., 2010). Given the adaptability of neural systems, the question arises as to whether learning mechanisms exist that develop appropriate coincidence detectors and then stabilize them during use.

The widely used Spike Timing Dependent Plasticity (STDP) (Markram et al., 27; Bi and Poo, 1998; Zhang et al., 1998; Froemke and Dan, 2002; Dan and Poo, 2004) learning rule normally requires the postsynaptic neuron to fire a spike and will reinforce all synapses with inputs arriving shortly before that spike. Synapses on distant dendrites whose earlier inputs also contribute to the spike undergo a much weaker reinforcement than proximal dendrites and end up disappearing when resource limitations are considered in the model, as proposed by Letzkus et al. (2006). Branco et al. (2010) have shown that, on the contrary, synapses at various distances from the soma stay strong and contribute to sequence-specific neuronal responses. They did that by activating a succession of synapses by optical uncaging and noted that if the uncaging sequence moves from distal to proximal synapses, the soma showed a higher increase in potential than if the sequence moved away from the soma. Given the results by Branco et al. (2010), it should be possible to reinforce synapses at any distance.

In this paper, we are interested in reinforcing pairs of synapses that are separated by a propagation time delay corresponding to the arrival time difference of spikes from two different input neurons. We initially examined whether a detector based on dendritic propagation delays in a biophysical model of a pyramidal neuron (Letzkus et al., 2006) can be developed in a bottom-up, unsupervised fashion, i.e., without the soma firing a prior spike to trigger learning on pre-synaptic inputs, following a hypothesis formulated by Bugmann and Christodoulou (2001). A bottom-up approach is in the spirit of experiments conducted by Marom and Shahaf (2002) showing learning without supervisory spiking by the target. The examined mechanism is based on non-linear summation of synaptic EPSPs (Excitatory Postsynaptic Potentials) and their effects, as described for example in Denham and Denham (2001), followed by

* Corresponding author.
E-mail addresses: achilleas.k@cs.ucy.ac.cy (A. Koutsou), g.bugmann@plymouth.ac.uk (G. Bugmann), cchrist@cs.ucy.ac.cy (C. Christodoulou).

http://dx.doi.org/10.1016/j.biosystems.2015.08.005
0303-2647/

the backpropagation of the summed EPSP to the dendrites, triggering a learning mechanism at the originating synapse. Simulating this initial approach revealed that the learning mechanism appears to be insufficiently sensitive to differences in time delays. This lead to the development of a second approach using a backpropagating action potential (AP).

In the second approach, a background input current is added to the neuron (at the somatic compartment), to allow the coincidence of pairs of small EPSPs to generate a spike that can then activate learning mechanisms when backpropagating. That background current can be seen as a "learning-enable" signal that is activated when the organism decides that there is a need to learn the current input situation. These processes are designed to allow learning of weights in conditions where they are initially too weak to induce output spikes.

A key element of both approaches is the assumption that inputs from each presynaptic neuron initially target several pre-existing synapses at various positions on the dendrite. These synapses have a probabilistic behaviour and will activate at most one at a time, thus probing various dendritic propagation times (Bugmann and Christodoulou, 2001). The learning rule should then select synaptic pairs separated by the appropriate distance and reinforce them.

This approach differs from the supervised approach used by van Leeuwen (2004) who assumes synaptic relocation along the dendritic tree, or the model by Hüning et al. (1998) that assumes delay modification. The principle of selection of existing synapses is also used by Gerstner et al. (1998), where the time differences between pre- and postsynaptic spikes determine weight changes, or the work by Eurich et al. (2000) who use a Hebbian learning rule depending on correlations between pre- and postsynaptic activity within a certain time window. Senn et al. (2002) also proposed the use of stochastic synapses, for adapting synaptic delays. Note that the problem treated here is different from that of detecting temporal patterns in a single input spike train, like in Hunzinger et al. (2012), or global oscillations in multiple spike trains like in Kerr et al. (2013). In the context of dendritic delays selection, in this paper we examine the capability of a pyramidal neuron to provide a learning signal selective enough to certain input time differences.

2. Methods

2.1. Overview

Fig. 1 shows a simplified sketch of our model's architecture. Four synapses attach to a neuron's dendrite at increasing distances from the soma. The synapse that is closest to the soma, synapse B, originates from presynaptic neuron B. The rest of the synapses, A1, A2 and A3, originate from presynaptic neuron A.

Presynaptic neurons A and B fire the same spike trains with a fixed time delay, Δt. In other words, whenever neuron A fires at a time t, B fires at $t + \Delta t$. In this paper, we aim at reinforcing synapse A2. We consider two scenarios (Fig. 1):

(i) An EPSP from A2 arrives at B after Δt time, thus coinciding with the time the EPSP at B is created. The coinciding EPSPs are amplified, creating an increase in postsynaptic potential at B, which travels back to the A synapses.
(ii) The EPSPs from A2 and B coincide at their arrival at the soma and trigger a somatic spike, creating a back-propagating action potential which travels back to the synapses.

In both cases, the back-propagating potential is expected to cause weight changes in the active synapses (i.e., the synapses that have recently been active).

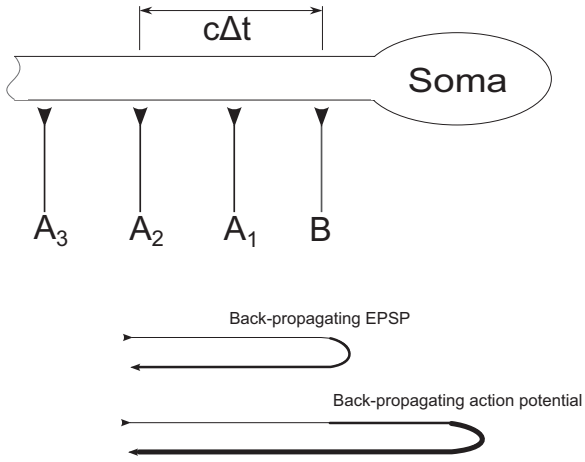

Fig. 1. Schematic of our model's architecture consisting of a simple neuron with 4 synapses. B is a proximal synapse, while synapses A1–3 are at increasing distances from the soma. All A-synapses originate from the same pre-synaptic neuron (neuron A) and B originates from a different one (neuron B). See the text for an explanation of the two back-propagation diagrams.

The purpose of both scenarios is to make the post-synaptic neuron sensitive to the firing delay between pre-synaptic neurons A and B, by reinforcing only synapse B and the corresponding A-synapse whose distance from B is such that the EPSP from A coincides with the EPSP from B, at location B. In other words, if c is the propagation speed and Δt is the firing delay between pre-synaptic neurons A and B, the learning mechanism should reinforce an A-synapse that is at a distance $c\Delta t$ from synapse B. In all our scenarios, the A-synapse that is located at the ideal distance from B will be labelled A2. Our methods require that synapses are stochastic with a low probability of release (Pun et al., 1986; Redman, 1990; Senn et al., 2002), since synapses A1, A2 and A3 all originate from the same pre-synaptic neuron, but should receive individual reinforcement. By setting the release probability sufficiently low, we can consider that the probability of having two or more A-synapses active at the same time is negligible.

The main difference between the two approaches is the lack of somatic spiking in the first approach. The first scenario relies on the amplification and backpropagation of a potential, caused by the coinciding EPSPs at the dendritic location of synapse B. Plasticity, in this scenario, would occur as a result of the changes caused along the dendrite by the backpropagating amplified EPSP, in the absence of somatic spiking. The second scenario follows a more traditional approach to learning, where the coinciding EPSPs trigger a somatic action potential that is able to cause synaptic changes based on a STDP-type learning rule.

2.2. Model

For our simulated experiments, we used the NEURON simulation environment (Hines and Carnevale, 1997) using a reconstructed layer 5 pyramidal neuron model, originally built by Stuart and Spruston (1998). This model was modified by Letzkus et al. (2006)[1] to account for active properties, by including voltage-gated ion channels at the following densities (in pS μm^{-2}):

- Soma: $g_{Na} = 3000$, $g_{Kv} = 30$, $g_{Ka} = 0.06$, $g_{Kca} = 2.5$, $g_{Km} = 2.2$, $g_{CaT} = 0.0003$.
- Dendrites: $g_{Na} = 40$, $g_{Kv} = 30$, $g_{Ka} = 0.03$, $g_{Kca} = 2.5$, $g_{Km} = 0.05$, $g_{CaT} = 0.0003$.

[1] The biophysical pyramidal neuron model is available at https://senselab.med.yale.edu/ModelDB/ShowModel.cshtml?model=108459.