Implementation of Attentional Bistability in a Computational Model of the Dragonfly Visual System

Juan Carlos Farah, Christos Kaplanis, Chris Snowden, Luka Milic
Zafeirios Fountas, Pedro A.M. Mediano ({zfountas,pmediano}@imperial.ac.uk)

Department of Computing, Imperial College London, United Kingdom

Abstract

The dragonfly is notoriously effective at hunting moving prey. Recent research has provided evidence that the selective visual attention of one of the dragonfly's visual neurons plays a key role in the neural mechanisms underlying the dragonfly's hunting ability. In this work, we present a hybrid computational model that includes multi-compartmental spiking neurons of the dragonfly visual system, as well as spike-timing-dependent plasticity (STDP)-based pattern recognition and action selection mechanisms to replicate this predatory behaviour in a simplified simulated environment. We find that under certain conditions two coupled visual neurons are capable of demonstrating bistable switching between input patterns, in agreement with empirical electrophysiological findings that evidence the role of these neurons in target selection. We also demonstrate the feasibility of training an end-to-end dragonfly visual system to map retinal input to motor actions in a biologically plausible way.

Keywords: Spiking neurons; Bistability; Dragonfly; CSTMD1; STDP; Reinforcement learning

Introduction

The Centrifugal Small Target Motion Detector 1 (CSTMD1) is a neuron found in the dragonfly brain, and there are open areas of research concerning its role in dragonflies' ability to effectively track moving targets and perform selective attention (Bart R. H. Geurten & O'Carroll., 2007). Wiederman and O'Carroll (2013) proposed that this phenomenon could be attributed to the CSTMD1 selecting a single prey from a set of competing targets and "encod[ing] the winning stimulus" as if it were the sole target. This would be exhibited as a system with only two stable equilibria, otherwise called a bistable system. Hence, the observation of bistability in our CSTMD1 is an important way to validate our model.

We designed a modular system that models the hunting behaviour of the dragonfly, from its retina to its motor actions. Our primary goal was to observe and analyse the target selection mechanism of the CSTMD1 in the context of the other components of the dragonfly's visual system.

Methods

The Dragonfly Visual System Our system, depicted in Figure 1, shows the high-level architecture of our dragonfly simulator. A visual field containing targets superimposed on a moving natural scene is rendered in the environment module, which is passed as a two-dimensional image to the Elementary Small Target Motion Detectors (ESTMD), neurons in the dragonfly's early visual pipeline. The output of the ESTMD provides the input to the CSTMD1, which produces spike trains for the pattern recognition (PR) module. The PR module deciphers recurring patterns in the activity of the CSTMD1 and outputs spike trains for the action selection (AS) module, which in turn—via reinforcement learning—maps the processed visual input to the motor actions of the dragonfly.

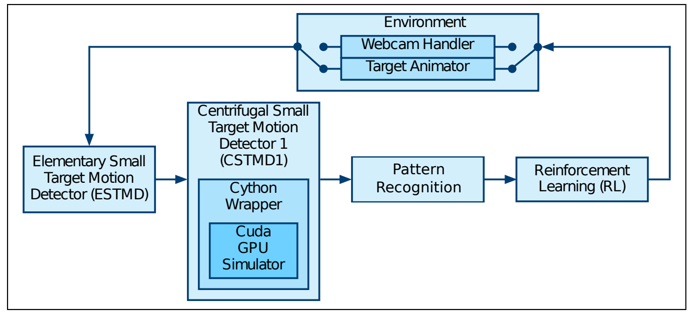

Figure 1: The architecture of the dragonfly model.

ESTMD The ESTMD module aims to emulate the function of the initial layers of visual processing in the dragonfly; for this we replicated an existing model by Halupka, Wiederman, Cazzolato, and O'Carroll (2011) which consists of a series of spatiotemporal transforms. Its primary role is to identify small fast-moving targets in the visual field even against a moving, cluttered background.

CSTMD1 The specific morphological structure of the CSTMD1 is putatively important for its role in target selection, so we decided to create a biologically detailed multi-compartmental model. The CSTMD1 neuron is characterised by a long axon which traverses the protocerebrum (Bart R. H. Geurten & O'Carroll., 2007) and at one end a large "dendritic cloud" that forms a three-dimensional tree-like structure. Morphologies were generated using the TREES Octave toolbox. To model the complex behavior of the dendrites, each neuron is broken down into a series of compartments, each with its own voltage and state variables. For each compartment we use the Hodgkin and Huxley (1952) conductance model to simulate the dynamics of the compartment's voltage with Euler's method.

Pattern Recognition The purpose of the pattern recognition module is to discern recurring spike patterns within the output of the CSTMD1, which encode information on the motion of the target observed in the visual field of the dragonfly. Our implementation follows an unsupervised learning technique using spike-timing-dependent plasticity (STDP) from Masquelier, Guyonneau, and Thorpe (2009). The module consists of a population of leaky integrate-and-fire neurons, each of which becomes selective to one pattern in the CSTMD1 output. Mutual inhibition between the neurons in this module enforces them to become selective to distinct patterns.

Action Selection The role of the action selection module is to map the output of the pattern recognition module to motor actions of the dragonfly that result in the target's capture. The motor actions of the dragonfly are encoded in the activity of four integrate-and-fire motor neurons, which jointly determine the velocity of the dragonfly's attentional focus point in the image. The synapses between the pattern recognition neurons and the motor neurons are trained with a biological form of reinforcement learning called *reward-modulated STDP* (Izhikevich, 2007). In this setup, STDP updates an eligibility trace for each synapse that decays over time; and if the dragonfly receives a reward (i.e. the target is captured), dopamine is administered to the system and eligibility traces are effectively crystallised into weight updates.

Results

We replicated the experimental framework of Wiederman and O'Carroll (2013) to investigate the CSTMD1's bistable behaviour. We presented our visual system, which contained an array of five CSTMD1 neurons (Fig. 2), with laterally ($f_{\{\rightarrow\}}$) or longitudinally ($f_{\{\downarrow\}}$) moving targets and recorded the spiking response at the somas of the CSTMD1s. We then compared these responses with the response in the presence of both stimuli ($f_{\{\rightarrow,\downarrow\}}$). We judged that bistability would most likely be observed in the soma of the CSTMD1 since, due to its morphology, there is a higher chance of spiking events being inhibited by the time they reach the soma from the dendrites. As shown in Fig. 3, we found that the $f_{\{\rightarrow,\downarrow\}}$ response curve closely matched just *one* of the response curves for either $f_{\{\downarrow\}}$ or $f_{\{\rightarrow\}}$ at any given time, and occasionally switched from one to the other.

This demonstration of bistability in our CSTMD1 model mirrors the effect observed by Wiederman and O'Carroll (2013) in CSTMD1 neurons *in vivo* . By replicating this finding in a computational model, we provide evidence that the specific morphology of the CSTMD1 contributes to its function as a low-level mechanism for selective attention, since the bistable activity of the CSTMD1 ensures that the dragonfly is actually only "seeing" one of the targets at a given point in time.

In future work, we will tune the pattern recognition and action selection modules in order to evaluate the hunting performance of our dragonfly model and analyse how its hunting strategy correlates with the target selected by the CSTMD1.

All authors contributed equally to this work. We thank Panagiotis Almpouras, Alex Carver, Erik Grabljevec, Ioannis Kasidakis, Zoe Landgraf and Klaus Stiefel for contributions on earlier phases of this project.

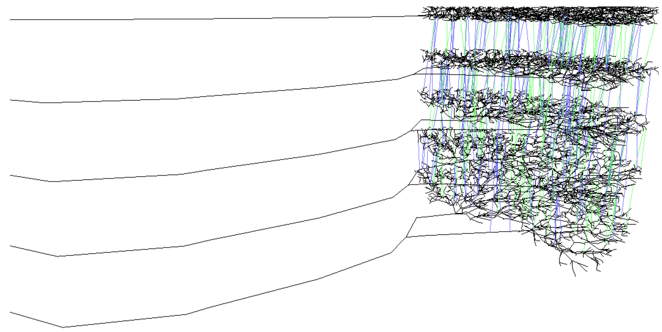

Figure 2: Visualisation of the soma and the dendritic cloud of five stacked CSTMD1 neurons. Green and blue edges represent points of contact between neighbouring neurons. Neurons have been separated vertically for visualisation purposes.

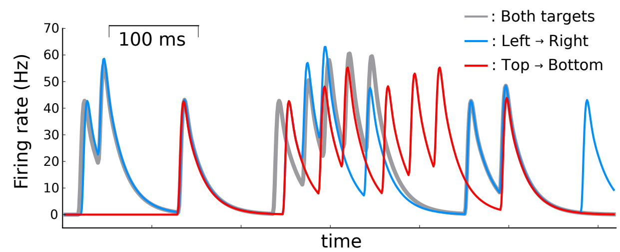

Figure 3: Firing rate bistability at the soma of a CSTMD1 neuron stimulated with longitudinally and laterally moving targets. The bistability can be observed by noticing that the grey line tracks one of the blue line and the red line, sometimes switching between the two.

References

Bart R. H. Geurten, J. D. H. S. D. M. B., Karin Nordstrm, & O'Carroll., D. C. (2007). Neural mechanisms underlying target detection in a dragonfly centrifugal neuron. *Journal of Experimental Biology*, *210*(18), 3277-3284.

Halupka, K., Wiederman, S., Cazzolato, B., & O'Carroll, D. (2011). Discrete implementation of biologically inspired image processing for target detection.

Hodgkin, A. L., & Huxley, A. F. (1952). A quantitative description of membrane current and its application to conduction and excitation in nerve. *The Journal of Physiology. USNLM*.

Izhikevich, E. M. (2007). Solving the distal reward problem through linkage of STDP and dopamine signaling. *Cerebral Cortex*, *17*(10), 2443–2452.

Masquelier, T., Guyonneau, R., & Thorpe, S. J. (2009). Competitive STDP-based spike pattern learning. *Neural Computation*, *21*(5), 1259–1276.

Wiederman, S. D., & O'Carroll, D. C. (2013). Selective attention in an insect visual neuron. *Current Biology*, *23*(2), 156–161.