Neural Networks 74 (2016) 35–51

Contents lists available at ScienceDirect

Neural Networks

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

Behavioral plasticity through the modulation of switch neurons

Vassilis Vassiliades, Chris Christodoulou *

Department of Computer Science, University of Cyprus, 1678 Nicosia, Cyprus

ARTICLE INFO

Article history:
Received 10 April 2015
Received in revised form 13 October 2015
Accepted 3 November 2015
Available online 12 November 2015

Keywords:
Switch neuron
Behavioral plasticity
Neuromodulation
Gating
Adaptive behavior
Reinforcement learning

ABSTRACT

A central question in artificial intelligence is how to design agents capable of switching between different behaviors in response to environmental changes. Taking inspiration from neuroscience, we address this problem by utilizing artificial neural networks (NNs) as agent controllers, and mechanisms such as neuromodulation and synaptic gating. The novel aspect of this work is the introduction of a type of artificial neuron we call "switch neuron". A switch neuron regulates the flow of information in NNs by selectively gating all but one of its incoming synaptic connections, effectively allowing only one signal to propagate forward. The allowed connection is determined by the switch neuron's level of modulatory activation which is affected by modulatory signals, such as signals that encode some information about the reward received by the agent. An important aspect of the switch neuron is that it can be used in appropriate "switch modules" in order to modulate other switch neurons. As we show, the introduction of the switch modules enables the creation of sequences of gating events. This is achieved through the design of a modulatory pathway capable of exploring in a principled manner all permutations of the connections arriving on the switch neurons. We test the model by presenting appropriate architectures in nonstationary binary association problems and T-maze tasks. The results show that for all tasks, the switch neuron architectures generate optimal adaptive behaviors, providing evidence that the switch neuron model could be a valuable tool in simulations where behavioral plasticity is required.

1. Introduction

Adaptive organisms have the remarkable ability of adjusting their behavior in response to changes in their environment. Such *behavioral plasticity* is believed to be linked with modifications in the neural circuitry that produces the behavior. These modifications are likely to be caused by mechanisms that go beyond the classical neurotransmission (of excitation or inhibition), such as *neural plasticity* (Binder, Hirokawa, & Windhorst, 2009; Churchland & Sejnowski, 1992) and *neuromodulation* (Katz, 1999). Neural plasticity refers to the capacity of neural circuits for functional or organizational modifications due to previous activity or damage (Binder et al., 2009; Churchland & Sejnowski, 1992). For example, synaptic plasticity, i.e., the strengthening or weakening of synapses, is a major process that underlies learning and memory (Martin, Grimwood, & Morris, 2000) and has been validated through neural recordings (see for example Bi & Poo, 1998; Bliss & Lømo, 1973; Kandel & Tauc, 1965; Markram, Lübke, Frotscher, & Sakmann, 1997). Neuromodulation refers to the process where a small number of neurons can influence (modulate) the intrinsic properties of multiple synapses or neurons, through the diffusion of certain neurotransmitters known as neuromodulators (Binder et al., 2009; Katz, 1999; Marder & Thirumalai, 2002). Neuromodulation and neural plasticity can be complementary. For example, the neuromodulator dopamine is believed to play a major role in operant conditioning (Skinner, 1938; Thorndike, 1911) as it was found to encode a reward prediction error (RPE; see for example Schultz, 1998) analogous to temporal difference (TD) error in reinforcement learning (RL, Sutton & Barto, 1998).

While neuromodulation can be used to *gate* plasticity and synaptic transmission, a growing number of studies provide evidence that supports the existence of other types of *synaptic gating* mechanisms, capable of regulating information flow between various sets of neurons (see Gisiger & Boukadoum, 2011, and references therein). Such mechanisms should not be thought of as simply interrupting information flow, as they can also act as permissive gates (Katz, 2003). For example, certain neurons from an area of the brain called the nucleus accumbens (NAcc) were found to implement this type of gating. These NAcc neurons are *bistable*, meaning that they exhibit oscillations between two discrete states: an "up" state (where the membrane potential is depolarized) during which the neuron generates action potentials (spikes), and a

* Corresponding author.
E-mail addresses: v.vassiliades@cs.ucy.ac.cy (V. Vassiliades), cchrist@cs.ucy.ac.cy (C. Christodoulou).

http://dx.doi.org/10.1016/j.neunet.2015.11.001
0893-6080/

resting, "down" state (where the membrane potential is hyperpolarized) during which the production of action potentials is absent (Grace, 2000; O'Donnell & Grace, 1995). They were found to be part of a gating mechanism that controls whether information from the prefrontal cortex (PFC) is allowed to pass through to the ventral pallidum and further, to the thalamus. More specifically, input from PFC neurons arrives at NAcc neurons, but only those that are in their up state allow the input to propagate forward. What modulates the state of the NAcc neurons is an extra input from the hippocampus (HPC). That is, only the NAcc neurons that are stimulated by HPC neurons enter their depolarized state and subsequently, fire upon receiving input from the PFC (Grace, 2000). For this reason, these neurons are said to implement a type of AND gate, since they fire only if they receive input from both the PFC and the HPC (Gisiger & Boukadoum, 2011). Various other logic gates, such as NOT, Switch, XOR, and Flip-Flop, can be implemented by neural circuits, as demonstrated by Vogels and Abbott (2005).

Apart from bistable neurons, which were found to be abundant in the cortex, other gating mechanisms have been observed in theoretical or experimental data, that feature inhibitory neurons or even oscillations (for examples see Anderson & Van Essen, 1987; Barbas & Zikopoulos, 2007; Burchell, Faulkner, & Whittington, 1998; Floresco & Grace, 2003; Olshausen, Anderson, & Van Essen, 1993). In all observations and models, certain "gatekeeper" circuits influence synaptic transmission (Gisiger & Boukadoum, 2011). As mentioned above, it has been observed that for NAcc neurons the gatekeepers originate in the HPC. For cortex neurons, experimental evidence suggests that the gatekeepers could originate in the cortex, thalamus and basal ganglia (Gisiger & Boukadoum, 2011). Gisiger and Boukadoum (2011) present a theoretical model where a copy of the gating signal produced by one gatekeeper circuit, can be fed as an input to another gatekeeper circuit. A key observation in that model is the existence of two types of neural pathways: the first implements normal information processing, whereas the second is formed by the gating mechanisms. They hypothesize that such interacting gating circuits could create sequences of gating events that are responsible for the production of structured behavior.

Gruber, Hussain, and O'Donnell (2009) used multichannel recordings to investigate rats during spatial exploration of an operant chamber, and during reward-seeking afterwards in the same chamber. They observed that during spatial exploration, the activities of neurons in the NAcc core, i.e., the inner part of the NAcc which is located within the basal ganglia (Gerfen & Wilson, 1996, Chap. 2, p. 372), synchronized with the activity of HPC neurons; however, during reward-seeking, they instead synchronized with the activity of PFC neurons. This suggested that the NAcc core can dynamically select its inputs according to environmental requirements, as it is able to switch its synchronization in a task-dependent manner (Gruber et al., 2009). It has to be noted that the basal ganglia is the structure believed to be associated with action selection (Redgrave, Prescott, & Gurney, 1999) and RL (chap. 11, 1995; Montague, Dayan, & Sejnowski, 1996; Schultz, Dayan, & Montague, 1997). For example, Redgrave et al. (1999) proposed that the action selection problem of vertebrates is solved by a central "switching" mechanism that resides in the basal ganglia. It has also been recently suggested that the release of a peptide called "substance P" in the striatum (i.e., the primary input nucleus to the basal ganglia), allows for rapid switching between actions in action sequences (Buxton, Bracci, Overton, & Gurney, 2015).

Gating mechanisms can be seen as implementing a type of "on–off" switch by allowing or interrupting communication between brain regions. Inspired by these gating phenomena in the brain, which seem to play a significant role in various processes such as working memory (Williams & Goldman-Rakic, 1995), attention (Usher, Cohen, Servan-Schreiber, Rajkowski, & Aston-Jones, 1999), and decision making (Cisek, Puskas, & El-Murr, 2009), we ask whether we could design an abstract computational model that can be used for the purpose of adaptive behavior. For this reason, we adopt the artificial intelligence agenda of rational decision making (Russell & Norvig, 2003), where an agent tries to maximize its reward intake (Sutton & Barto, 1998). The agent is controlled by artificial neural networks (NNs), as they are very well suited for simulating adaptive behavior, due to the possibility of implementing memory (through recurrent connections), and learning (through plasticity rules). In this paper, we do not focus on learning behavior per se, but rather on behavior exploration. More specifically, a central hypothesis of this work is that once some general neural circuits are established for certain behaviors through possibly neural plasticity mechanisms (or other methods), neuromodulation alone can be used to switch these behaviors by selectively gating various pathways accordingly.

We implement such a gating mechanism by introducing a novel type of an artificial neuron we call *"switch neuron"* that can be used in NNs. Instead of implementing an on–off switch for certain connections, this unit selects which one of its incoming connections is allowed to propagate its signal forward, by opening its gate while closing the gate of all others. The role of the switch neuron is to endow an agent with different behaviors and the ability to flexibly switch them as needed. The switching activity is controlled by modulatory signals that encode some information about the reward received by the agent. In order to create sequences of gating events and structured exploration, we additionally introduce a way for switch neurons to modulate other switch neurons. This is done by placing them in appropriate switch modules. We assess our model by designing appropriate switch neuron architectures for nonstationary association tasks (Section 3.1) and discrete T-maze problems (Section 3.2). We show in all tasks that these architectures perform optimal deterministic exploration when the goal changes, therefore, illustrating that our approach advances the field of NNs by creating more adaptive networks.

Note that the switch neurons of this paper should not be viewed in a strict biological sense, but rather in a functional sense. They are inspired by biological phenomena, but they are artificially constructed to perform certain computations. Thus, throughout this study, we use the word "neuron" for the switch neuron, but note that this is a purely artificial unit. In other words, despite the strong biological inspiration for the design of the switch neuron model, our paper does not contribute to any advances in biological areas. If mechanisms similar to the switch neuron model exist in the brain, they could either be in the form of individual cells, population of cells, or groups of interconnected neurons.

The remainder of the paper is organized as follows. Section 2 describes our approach by introducing the switch neuron and switch module. These are integrated in NN architectures designed specifically for the experiments reported in Section 3 along with the results. Section 4 discusses our results and directions for future work, and the conclusions are given in Section 5.

2. Approach

2.1. Artificial neurons

The usual formulation of an artificial neuron involves the integration of incoming signals $\mathbf{y} := (y_1, y_2, \ldots, y_N)$ and parameters $\mathbf{w} := (w_1, w_2, \ldots, w_N)$ through an accumulation or integration function $G(.)$ resulting in the neuron's activity $a(t) := G(\mathbf{y}, \mathbf{w})$ at time t. This activity is then fed through an activation function $F(.)$ resulting in the neuron's output $y(t) := F(a(t))$.