Neural Networks 67 (2015) 1–13

Contents lists available at ScienceDirect

Neural Networks

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

A spiking neural network based on the basal ganglia functional anatomy

Javier Baladron, Fred H. Hamker *

Department of Computer Science, Chemnitz University of Technology, Strasse der Nationen 62, 09107 Chemnitz, Germany

ARTICLE INFO

Article history:
Received 5 May 2014
Received in revised form 29 January 2015
Accepted 3 March 2015
Available online 24 March 2015

Keywords:
Spiking neurons
Basal ganglia
Reinforcement learning
Cognitive modeling
Action selection

ABSTRACT

We introduce a spiking neural network of the basal ganglia capable of learning stimulus–action associations. We model learning in the three major basal ganglia pathways, direct, indirect and hyperdirect, by spike time dependent learning and considering the amount of dopamine available (reward). Moreover, we allow to learn a cortico-thalamic pathway that bypasses the basal ganglia. As a result the system develops new functionalities for the different basal ganglia pathways: The direct pathway selects actions by disinhibiting the thalamus, the hyperdirect one suppresses alternatives and the indirect pathway learns to inhibit common mistakes. Numerical experiments show that the system is capable of learning sets of either deterministic or stochastic rules.

1. Introduction

The basal ganglia are a set of nuclei located in the forebrain. Several experiments have associated this brain area to action selection and reinforcement learning (Grillner, Hellgren, Menard, Saitoh, & Wikstrom, 2005; Packard & Knowlton, 2002; Wickens, Reynolds, & Hyland, 2003). The reinforcement signal, i.e. a reward prediction error (Schultz, 2010), is transferred to basal ganglia in form of the neurotransmitter dopamine originating in the substantia nigra pars compacta and the ventral tegmental area.

The basal ganglia are composed of several cortico thalamic loops that start in the cortex and via different pathways converge in the internal globus pallidus, an output structure which projects through the thalamus back to the cortex. All loops include either the striatum or the subthalamic nucleus, areas that are considered as input stages of the basal ganglia. In the basal ganglia each loop is composed of typically three different pathways, a direct pathway, an indirect pathway and a hyperdirect pathway (Schroll & Hamker, 2013).

To allow the simulation of behavioral experiments, several models of the complete basal ganglia have been proposed that do also include synaptic plasticity. Most of them are based on mean rate neurons (see for example: Frank, 2005; Gurney, Prescott, & Redgrave, 2001; Schroll, Vitay, & Hamker, 2012, 2014). Only recently, some models using spiking neurons have appeared (Chersi, Mirolli, Pezzulo, & Baldassarre, 2013; Stewart, Bekolay, & Eliasmith, 2012). The main difference between both relates to the different ways of modifying the synaptic weights. While rate-based models typically adjust their weights based on a three factor rule, pre- and post-synaptic firing rate plus dopamine, spiking models can also consider the exact timing of pre- and post-synaptic spikes by spike timing dependent plasticity (STDP) learning rules (Markram, Gerstner, & Sjöström, 2011; Morrison, Diesmann, & Gerstner, 2008).

Recently, Schroll et al. (2014) presented a set of learning rules for rate-based models that allow to determine the function of each pathway while minimizing hard-wired connections. However, none of the published spiking neural models of the basal ganglia allows for synaptic plasticity in all three pathways. Thus, in the present work we propose a new spiking network, inspired by the rate model presented by Schroll et al. (2014), that allows for learning by means of STDP in all three pathways. As a result of learning novel interpretations in the function associated to each of the cortico-thalamic pathways that include the basal ganglia emerge.

2. Basal ganglia anatomy

The main input structure of the basal ganglia is the striatum. It is composed mainly of medium spiny neurons (MSNs) and of a small amount of interneurons. The input from the cortex and thalamus to the striatum is mediated by glutamatergic synapses (Leh, Ptito, Chakravarty, & Strafella, 2007; Smith, Raju, Pare, & Sidibe, 2004; Wiesendanger, Clarke, Kraftsik, & Tardif, 2004). MSNs

* Corresponding author.
E-mail address: fred.hamker@informatik.tu-chemnitz.de (F.H. Hamker).

http://dx.doi.org/10.1016/j.neunet.2015.03.002
0893-6080/

are quiet at rest and require a strong correlated input to activate (Nisenbaum & Wilson, 1995). Once active they inhibit through GABAergic connections the neurons in the globus pallidus.

The striatum is also the destination of many projections from the dopaminergic cells of the substantia nigra pars compacta. The dopamine signal produced by these connections provides the basal ganglia with information about the performance of the task by means of a reward prediction error (Schultz, 2007). When more reward than expected is obtained (for example, in the form of juice in animal experiments) the level of dopamine is enhanced and when less reward than expected is received the dopamine level is reduced. This dopamine signal modulates learning in the basal ganglia connections. In neurons expressing the type-1 receptor (D1) a rise in the level of dopamine produces long term potentiation, while in neurons expressing the type-2 receptor (D2), it produces long term depression (Shen, Flajolet, Greengard, & Surmeier, 2008). A reduced level of dopamine reverses this effect, producing long term depression in D1 cells and long term potentiation in D2 cells. This difference suggests that both cell types have a different function. D1 expressing cells directly project to the internal globus pallidus (GPi), the main output nucleus of the basal ganglia, while D2 expressing cells first project to neurons in the external globus pallidus (GPe), which then project to the GPi. All these connections are inhibitory.

The pathway comprising striatal D1 cells and its direct connection to GPi is usually called the direct pathway and the one including D2 cells and GPe is usually called the indirect pathway. Both pathways converge in GPi which projects to the thalamus via GABAergic connections. The direct pathway, through the projections from striatum D1 cells, reduces the tonic activity of GPi and thus reduces the level of inhibition from GPi to the thalamus. The indirect pathway, through the inhibitory connections between striatum D2 cells and GPe, is able to remove the continuous inhibition that the tonic firing of GPe provides to GPi. The absence of this inhibition increases the level of activity in GPi. Thus, standard theories associate the direct pathway with a GO-function, i.e., initiating the correct action, and the indirect pathway to a NO-GO-function, i.e., inhibiting the incorrect actions (Braak & Del Tredici, 2008; O'Reilly & Frank, 2006; Schroll & Hamker, 2013).

The Subthalamic Nucleus (STN) is another input structure of the basal ganglia which receives connections from the cortex, the thalamus and the GPe and projects to GPi and GPe (Temel, Blokland, Steinbusch, & Visser-Vandewalle, 2005). The projections from the STN to GPi are excitatory, so this pathway, usually called hyperdirect, is supposed to have a different function than the inhibitory connection from the striatum. Anatomical evidence suggests a center–surround structure where this pathway inhibits competing motor programs while the direct pathway excites the correct one (Nambu, Tokuno, & Takada, 2002). For a recent review of the computational function of the different BG pathways see Schroll and Hamker (2013).

3. Previous spiking models of the basal ganglia

Recently, spiking models of the complete basal ganglia have been proposed which are able to learn stimulus–response associations. Chersi et al. (2013) developed a model to simulate a behavioral task in which a monkey must learn a stimulus–action association. The monkey sits in front of a table with three lights and three buttons. At the beginning of a trial a light is flashed and then the animal must discover which button he has to press to turn the light on again. Only one button is correct for each light. An interesting characteristic of this model is that to reach a successful response in one trial a set of consecutive actions must be performed. The simulated monkey must first look at the flashed light, then reach the button and finally press it.

In their network each nucleus of the basal ganglia, the motor cortex and the prefrontal cortex, are represented by a layer composed of a set of populations of leaky integrate and fire neurons, one for each possible action.

The projections arriving to the striatum, but not to STN, are plastic and adapt according to a three factor rule which depends on both the level of dopamine, the timing between spikes and a trace to solve the temporal credit assignment problem. However, the learning rule is identical for all connections, independent of the pathway or dopamine receptor. All remaining connections of the direct, indirect and hyperdirect pathway and their weights are determined by an optimization routine which assures the fulfillment of a set of biological and functional restrictions. In the resulting network the direct pathway is in charge of selecting the proper action by disinhibiting the corresponding population of the thalamus, while the indirect only maintains the activity of GPi within working limits, avoiding undesired behaviors like oscillations. The hyperdirect pathway basically switches off the BG selection mechanism allowing direct connections between the prefrontal cortex and the motor areas to determine the action to be executed.

Another spiking model has been recently proposed by Stewart et al. (2012) based on the firing rate model of Gurney et al. (2001). The model of Gurney et al. (2001) is basically an action selection model that, by its implemented BG connections, ensures that a single action is determined, the one with output zero, given utility values Q for each action as input. The utility value is the estimate of reward given a particular state and a chosen action. The model of Stewart et al. (2012), uses learning between cortex and striatum but not between cortex and STN to map a sensory state to utility values represented by populations of spiking neurons while the other connections are set to replicate the same exact computations performed by the original rate model. The learning rule does not depend on the timing between the presynaptic and postsynaptic spikes but only on the amount of activity of both cells and on an error signal. The activity of a neuron is estimated using the amount of neurotransmitters released in the synapse. The error signal is computed by subtracting the vector of utility values (one value for each possible action in a particular state) from the received reward. Thus, the error signal is not uniform across striatal populations. The network is capable of learning a probabilistic bandit task, in which an agent must choose between two or more actions. Reward is non deterministic, depending on a probability distribution.

Humphries, Stewart, and Gurney (2006) proposed a different spiking model which was used to study the appearance of slow oscillations in the STN and the GPe. The dopamine signal in this model is not used to modify the synapses, which are all fixed, but only as a modulator of synaptic efficiency. This model is able to reproduce several biological datasets but its lack of plasticity makes it incomparable with this work where one of the main objectives is to test the learning capabilities of the basal ganglia.

The learning process in the previous spiking models does not differentiate between cells with a different type of dopamine receptor. However, recent experiments have shown that the effect of dopamine in plasticity varies according to the type of receptor expressed by the cell (Calabresi, Picconi, Tozzi, & Filippo, 2007; Shen et al., 2008), a factor that none of the spiking models have taken into account.

Also, none of the previous models can explain the effects of pallidotomy, a common treatment for Parkinson's disease in which the GPi is lesioned, reducing the influence of the BG in the thalamus. After the surgery, patients can perform everyday movements (Lozano et al., 1995) but are impaired at learning (Sage et al., 2003). This suggests that there exists another mechanism for action selection that can guide behavior independent of the basal ganglia. We propose that this can be performed by direct cortico thalamic connections which are trained by the BG. However, in order to change