Fast and Slow Learning in a Neuro-Computational Model of Category Acquisition

Francesc Villagrasa(✉), Javier Baladron, and Fred H. Hamker

Artificial Intelligence, Chemnitz University of Technology,
Straße der Nationen 62, 09111 Chemnitz, Germany
vies@hrz.tu-chemnitz.de

Abstract. We present a neuro-computational model that, based on brain principles, succeeds in performing a category learning task. In particular, the network includes a fast learner (the basal ganglia) that via reinforcement learns to execute the task, and a slow learner (the prefrontal cortex) that can acquire abstract representations from the accumulation of experiences and ultimately pushes the task level performance to higher levels.

Keywords: Categorization · Basal ganglia · Fast-learner · Reinforcement learning · Prefrontal cortex · Slow-learner

1 Introduction

Categorization is the capacity to group items according to specific commonalities, in order to generalize or predict responses to new future stimuli and to build concepts that provide the world with meaning. Humans are especially good in this ability and therefore, we believe that to build a synthetic system capable of categorization we should look at mechanisms grounded in neuroscientific data.

The basal ganglia (BG) are a set of subcortical nuclei shown to be involved in a large number of categorization paradigms [1–5], especially those in which learning occurs via trial and error [6,7]. The BG are also associated with action selection [8], reinforcement learning [9] and it has been proposed to be involved in the training of cortico-cortical connections [10].

The prefrontal cortex (PFC) has also been shown to be involved in category tasks. It has been shown that neurons in this area can represent different categories [11–13]. Furthermore, the PFC plays a well-known role in executive functions [14].

We propose a novel principle of computation in which a fast-learner system (the BG) executes a category learning task via reinforcement learning while training a slower-learner system (the PFC) to acquire category information.

A.E.P. Villa et al. (Eds.): ICANN 2016, Part I, LNCS 9886, pp. 248–255, 2016.
DOI: 10.1007/978-3-319-44778-0_29

2 Methodology

2.1 Network Description

In our network, each nucleus of the BG and the PFC are represented by rate-coded neurons and their function arises from plastic synapses. Synaptic learning rules and neural activity are governed by differential equations which are solved by employing the Euler method with a time step of 1 ms. The network was built using the ANNarchy neural simulator [15] version 3.0.

We used a cortico-basalganglio-thalamic (CBGT) loop model based on preceding work [16] and connected it to the inferior temporal cortex (IT), the PFC and the premotor cortex (PM) (see Fig. 1). The IT represents stimuli information which is read by the Striatum (STR) and the subthalamic nucleus (STN), the two main input nuclei of the BG; the PM represents the different possible motor responses which can be initiated by the BG via the thalamus; and the PFC contains category information which will be acquired during the execution of the task. For further information about our BG model architecture consult [16].

In order to provide a sophisticated mechanism to learn categories, we have mainly introduced three connections linking the PFC with the thalamus and the IT. The IT is connected to the PFC with variable, excitatory synapses so that categories can be learned from the stimuli information. Each of the two thalamic neurons projects to a different PFC cell with fixed excitatory synapses, which allows the BG to transmit its action decision signal to the PFC, selecting one of the two PFC neurons to learn the current input pattern.

Finally, each PFC cell excites its afferent thalamic neuron back in order to bias the motor decision once enough information has been acquired. In particular, each PFC neuron is the only excitatory source of its corresponding thalamic neuron. Thus, when learning in the $IT \mapsto PFC$ connections still has not occurred, PFC provides both thalamic cells with the same input which is then just modulated or controlled by the inhibitory projections from the BG output. However, when category learning has been fully established in the PFC, the thalamic activity is completely biased by this knowledge. Then, a stimulus of a particular category will activate just one PFC cell and in doing so, just one of both thalamic cells will be activated by the PFC, impeding that BG can select the other thalamic neuron and consequently, it is now PFC which mainly rules thalamic activity and motor decision.

2.2 Experiment

We performed a numerical categorization experiment which consists of 400 trials in which learning took place. The model is taught to classify stimuli into one of two possible categories. Each trial starts with a rest period of 100 ms and is followed by 50 ms in which a randomly chosen stimulus is exposed. At the end of this period, the decision of the model is probabilistically evaluated via a

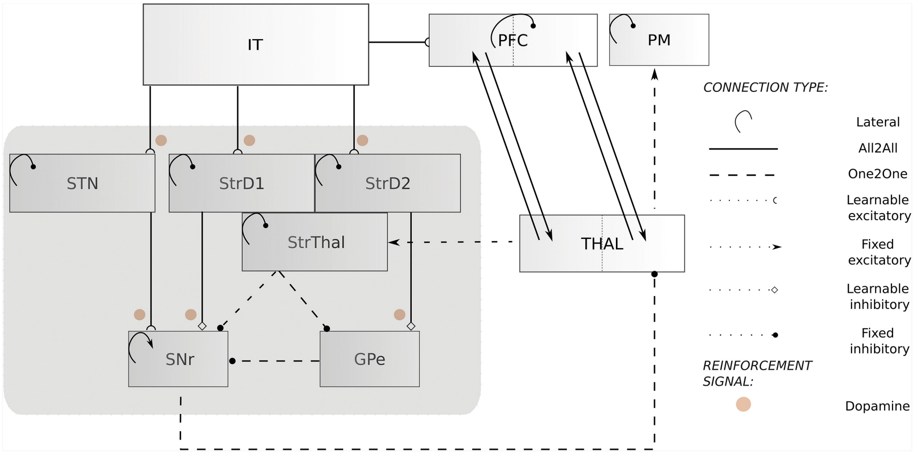

Fig. 1. Connections and connection types of the novel parts of the network. THAL: thalamus. PFC: prefrontal cortex. IT: inferior temporal cortex. PM: premotor cortex. The boxes inside the big shadow are the nuclei of the basal ganglia which correspond to the one introduced in [16]. One2One: each presynaptic population is only connected by its corresponding neuron of the postsynaptic population. All2all: every neuron in the presynaptic population is connected with every neuron in the postsynaptic population. $PFC \leftrightarrow THAL$ connections link half of the THAL neurons with half of the PFC neurons, constructing two loop structures.

soft-max rule depending on the activation of the PM and, if the model produces a correct response, reward is delivered for 500 ms.

Each stimulus or exemplar is defined by a 6×6 numerical matrix whose columns have one element set to one and the rest to zero, allowing for a total of 6^6 different exemplars. Category A represents all the stimuli whose first column has its first element set to one; the rest of exemplars belong to category B.

Each column of this matrix is considered as one of the stimuli's dimensions (color, shape, brightness, texture, size and orientation) and each row in the matrix as a value of each dimension (green, yellow, pink, red, purple and black for the color dimension), therefore category A encompasses all exemplars with green color and category B the rest of stimuli.

3 Results

We ran a set of 100 experiments with randomly initialized weights in the learnable connections and measured, among the trials of all experiments, the percentage of correct decisions (see Fig. 2). At the beginning, the model starts with a performance around 50 % which progressively increases until it reaches around 95 %, meaning that our model can successfully learn this category task.

The fast learning in the BG leads the STR cells to encode small sections of the stimuli (see Fig. 3c) and due to the high diversity of exposed exemplars,

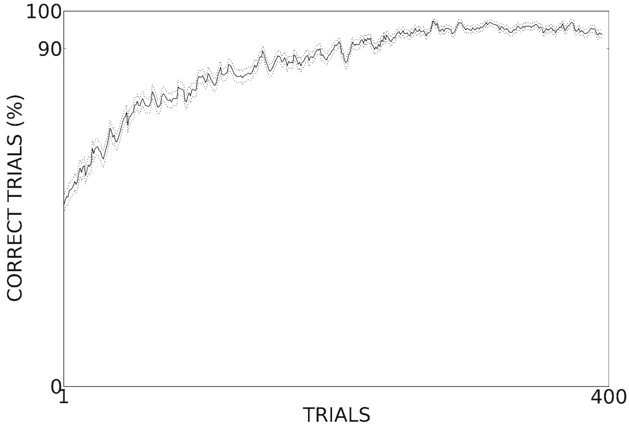

Fig. 2. Model performance in the 100 experiments per each of the 400 trials. The continuous line represents the average of successful trials and the two dotted lines indicate the corresponding standard error. To smooth the line plot, the average and the standard error are calculated from the data of the corresponding trial and the next 4 trials.

the input representation of most STR neurons strongly varies over time (see Fig. 3d). The slow learning in the PFC allows to eventually extract more generic and stable knowledge (see Fig. 3a and d).

In the exposed stimuli of each category, the IT neurons representing "color" are more probable to be active and thus more strongly encoded in PFC than the other IT cells. In the case of category A, the IT "green color" cell is always active and therefore, the PFC representations of category A eventually become selective just to this IT cell. In the case of category B, IT "color" neurons are just slightly more probable to be active than the rest of IT cells and thus, the PFC category B representations encode a broader range of IT neurons than just IT "color" cells. However, due to the random experiences in each experiment, the bias for encoding IT "color" cells can only be clearly observed in the average of the category B representations across all experiments (see Fig. 3b).

To compare both the STR and the PFC representations, we tested, at the end of each experiment, the capacity of both the PFC and the BG to correctly classify new, unseen, stimuli without the influence of the other while suppressing learning in the network. First, the effect of the PFC was tested by running 400 trials after removing the connections between the output nucleus of the BG and the thalamus. Then, the effect of the BG was tested by also running 400 trials after setting back the weights of the projections from the input neurons to the PFC to its initial conditions, thus removing any knowledge stored in them. The results of both are presented in Fig. 4.

Figure 4 shows that PFC can correctly classify around 99 % of the new stimuli while the BG, around the 68 %. This confirms that the slow learning system can acquire a better category representation. However, this is only possible due to the initial training performed by the BG.

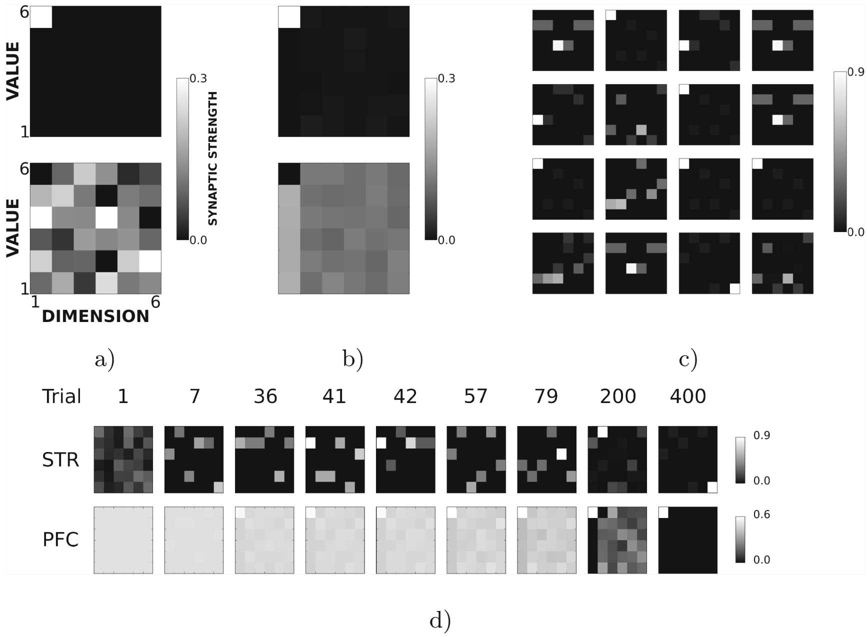

Fig. 3. "Receptive fields" of example PFC and STR cells at the end of the learning period. As illustrated in a), the x-axis, the y-axis and the grey scale of each subplot indicate the dimension of the input, the value of the input and the synaptic strength, respectively. (a) Two example PFC cells; one cell specializes in category A and the other in category B. (b) Mean across all experiments of the synaptic input representations in the two PFC cells. (c) Example of 16 STR neurons, which encode only small parts of stimuli information. (d) Evolution of the synaptic weights in one neuron of STR (first row) and one cell of PFC (second row) over time at nine different moments (trials).

4 Discussion

Our model succeeds in a category classification task and provides insight into the role of fast and slow learning in category acquisition. The high variability in the exposed stimuli impedes the input nuclei of the BG to clearly encode each category due to fast changes in the synaptic weights that force the neurons to extract only features of the most recent stimuli, while forgetting the rest. For this reason, the fast learning of the BG fails to produce a stable and complete category representation required for generalizing. However, it proves to be good enough to teach correct associations to the PFC.

Contrary to the stimuli specific knowledge acquired by the fast leaner, the slow learning in $IT \mapsto PFC$ projections allow them to gather a more general or broader amount of information. And because the fast leaning in BG produces a high enough number of correct associations, the BG train the PFC to more strongly encode those elements common in the stimuli of the corresponding

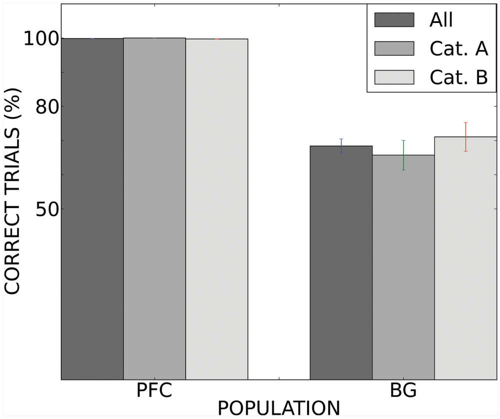

Fig. 4. Performance of the PFC and the BG alone in generalizing new stimuli. For each of both the PFC and the BG, there are three bars representing the percentage of trials correctly classified. Each bar disposes of an error bar indicating its standard error.

category, thus generating stable category representations that generalize across individual exemplars.

However, the slower learner would not be required if the number of stimuli were small enough. Hence, the fast learning could produce stable representations of the stimulus-response (SR) associations necessary to quickly achieve the highest performance [16–18]. Nevertheless, the task is no longer a category learning task but an SR learning task.

Although the fast learning system is not suitable for achieving the greatest scores in our task, it is appropriate for executing the initial trials, as it can reach a high performance in a small amount of time. Using a slow learner alone (without a fast learner) would take too much time to reach a high performance. Therefore, we believe that the brain requires a combination of both fast and slow learning for acquiring category representations while having a high task performance from the beginning.

Assigning fast and slow learning to BG and PFC respectively is in accordance with physiological considerations [19] built in a brain theory that proposes: first, a fast learning for extracting specific stimulus information, and a slow learning for identifying the commonalities among the elements of the same class [19]; and second, the BG to learn fast specific stimulus-motor associations while slowly teaching the PFC to learn categories [19–21]. Likewise, our model's function and architecture is in agreement with brain anatomy and neuroscience studies [22–24] which, for example, understand BG-cortex interaction via a cortico-basal ganglio-thalamic loop.

As in neuroscience, the machine learning community has been investigating the category recognition phenomenon. In particular, the state of the art of this field in machine learning belongs to deep neural networks with supervised learning (SL) which have won the latest contests and outperformed records in this domain [25,26]. However, learning to distinguish between entities just from the correct examples provided by a teacher (i.e. SL) works well if learning is done off-line. Category learning as discussed here relates more closely to reinforcement learning (RL), in which an agent pursues a goal by exploring and exploiting actions in its environment in order to maximize reward sensation [27].

Specifically, our model brings together both RL and categorization in such a way that fast RL learns useful actions while trains the slow learning to acquire category knowledge. Therefore, this model could be proposed as a novel deep networks' classifier, which could efficiently perform SR tasks, in case of a small number of stimuli, and categorization tasks, otherwise. Finally, further work should focus on adapting this model to more complex stimuli, for example, to learn to classify real objects by their shape.

Acknowledgments. This work has been funded by DFG HA2630/4-1 and in part by DFG HA2630/8-1.

References

1. Poldrack, R.A., Prabhakaran, V., Seger, C.A., Gabrieli, J.D.E.: Striatal activation during acquisition of a cognitive skill. Neuropsychology **13**, 564 (1999)
2. Poldrack, R.A., Clark, J., Pare-Blagoev, E.J., Shohamy, D., Moyano, J.C., Myers, C., Gluck, M.A.: Interactive memory systems in the human brain. Nature **414**, 546–550 (2001)
3. Nomura, E., Maddox, W.T., Filoteo, J.V., Ing, A.D., Gitelman, D.R., Parrish, T.B., Mesulam, M.M., Reber, P.J.: Neural correlates of rule-based and information-integration visual category learning. Cere. Cortex **17**, 37–43 (2007)
4. Seger, C.A., Cincotta, C.M.: The roles of the caudate nucleus in human classification learning. J. Neurosci. **11**, 2941–2951 (2005)
5. Zeithamova, D., Maddox, W.T., Schnyer, D.M.: Dissociable prototype learning systems: evidence from brain imaging and behavior. J. Neurosci. **28**, 13194–13201 (2008)
6. Merchant, H., Zainos, A., Hernández, A., Salinas, E., Romo, R.: Functional properties of primate putamen neurons during the categorization of tactile stimuli. J. Neurophysiol. **77**, 1132–1154 (1997)
7. Cincotta, C.M., Seger, C.A.: Dissociation between striatal regions while learning to categorize via feedback and via observation. J. Cognit. Neurosci. **19**, 249–265 (2007)
8. Humphries, M.D., Stewart, R.D., Gurney, K.N.: A physiologically plausible model of action selection and oscillatory activity in the basal ganglia. J. Neurosci. **26**, 12921–12942 (2006)
9. Grillner, S., Hellgren, J., Menard, A., Saitoh, K., Wikström, M.A.: Mechanisms for selection of basic motor programs-roles for the striatum and pallidum. Trends Neurosci. **28**, 364–370 (2005)
10. Hélie, S., Ell, S.W., Ashby, F.G.: Learning robust cortico-cortical associations with the basal ganglia: an integrative review. Cortex **64**, 123–135 (2015)
11. Freedman, D.J., Riesenhuber, M., Poggio, T., Miller, E.K.: A comparison of primate prefrontal and inferior temporal cortices during visual categorization. J. Neurosci. **23**, 5235–5246 (2003)

12. Freedman, D.J., Riesenhuber, M., Poggio, T., Miller, E.K.: Visual categorization and the primate prefrontal cortex: neurophysiology and behavior. J. Neurophysiol. **88**, 929–941 (2002)
13. Freedman, D.J., Riesenhuber, M., Poggio, T., Miller, E.K.: Categorical representation of visual stimuli in the primate prefrontal cortex. Science **291**, 312–316 (2001)
14. Miller, E.K., Cohen, J.D.: An integrative theory of prefrontal cortex function. Ann. Rev. Neurosci. **24**, 167–202 (2001)
15. Vitay, J., Dinkelbach, H., Hamker, F.H.: ANNarchy: a code generation approach to neural simulations on parallel hardware. Front. Neuroinf. **9**, 19 (2015)
16. Schroll, H., Vitay, J., Hamker, F.H.: Dysfunctional and compensatory synaptic plasticity in Parkinson's disease. Eur. J. Neurosci. **39**, 688–702 (2014)
17. Baladron, J., Hamker, F.H.: A spiking neural network based on the basal ganglia functional anatomy. Neural Netw. **67**, 1–13 (2015)
18. Packard, M.G., Knowlton, B.J.: Learning and memory functions of the basal ganglia. Ann. Rev. Neurosci. **25**, 563–593 (2002)
19. Seger, C.A., Miller, E.K.: Category learning in the brain. Ann. Rev. Neurosci. **33**, 203 (2010)
20. Antzoulatos, E.G., Miller, E.K.: Differences between neural activity in prefrontal cortex and striatum during learning of novel abstract categories. Neuron **71**, 243–249 (2011)
21. Miller, E.K., Buschman, T.J.: Rules through recursion: how interactions between the frontal cortex and basal ganglia may build abstract, complex rules from concrete, simple ones. In: Neuroscience of Rule-Guided Behavior, pp. 419–440 (2007)
22. Seger, C.A.: The basal ganglia in human learning. Neuroscientist **12**, 285–290 (2006)
23. Haber, S.N.: The primate basal ganglia: parallel and integrative networks. J. Chem. Neuroanat. **26**, 317–330 (2003)
24. McHaffie, J.G., Stanford, T.R., Stein, B.E., Coizet, V., Redgrave, P.: Subcortical loops through the basal ganglia. Trends Neurosci. **28**, 401–407 (2005)
25. Stallkamp, J., Schlipsing, M., Salmen, F., Igel, C.: Man vs. computer: benchmarking machine learning algorithms for traffic sign recognition. Neural Netw. **32**, 323–332 (2012)
26. Schmidhuber, J.: Deep learning in neural networks: an overview. Neural Netw. **61**, 85–117 (2015)
27. Sutton, R.S., Barto, A.G.: Reinforcement Learning: An Introduction. A Bradford Book, Cambridge (1998)