

# The EvoNeuro Approach to Neuro-Evolution

Stéphane Doncieux, Jean-Baptiste Mouret, Tony Pinville, Paul Tonelli and Benoît Girard

ISIR, Université Pierre et Marie Curie-Paris 6, CNRS UMR 7222

4 place Jussieu, F-75252, Paris Cedex 05, France Email: {doncieux,mouret,pinville,tonelli,girard}@isir.upmc.fr

**Abstract**—Computational neuroscience uses networks of artificial neurons to model cognitive functions of animals. Neuro-evolution uses the same abstract models of artificial neurons with the goal to control the behavior of artificial agents or real robots and make them exhibit adapted behaviors. Despite these similarities, neural networks obtained in these two fields are surprisingly different. The approach developed in the EvoNeuro project aims at reducing this gap to build via Neuro-evolution cognitive functions studied in a Computational Neuroscience context. This article summarizes the first results of this project on this topic, results focused on the inclusion of a feature that is absent from Neuro-evolution and ubiquitous in Computational Neuroscience: the structuring in maps of neurons with regular connection schemes instead of isolated neurons and connections.

## I. INTRODUCTION

Computational neuroscience [1], [2] and neuro-evolution [3]–[6] both produce “artificial nervous systems” that aim at reproducing on a machine some of the abilities found in animals. Models produced in computational neuroscience, however, display almost no similarities with evolved neural networks. In particular, neuro-evolution usually involve up to a few dozens of neurons with finely tuned topology and synaptic weights; by contrast, computational neuroscience often manipulate hundreds of neurons in a very organized fashion.

An analysis of published models (especially [7]–[13]) reveals that the basic building blocks make at least one fundamental difference: evolutionary methods mostly use individual neurons, ideally organized in modular and repetitive fashion [5], [6], [14]–[16], but many neuroscience models rely on the concepts of *maps* (a 1D or 2D grid of *identical* neurons) connected by regular connection schemes (either one to one connections or one to all with a regular assignation of weights). This allows such neural networks to scale up to larger maps (e.g. to handle higher-dimensional inputs) while maintaining the same overall structure. This description of neural networks as connected maps can be seen as the result of a developmental process in which a network of maps is developed to form a neural network.

Besides this analysis, computational neuroscience can also provide an efficient approach to benchmark neuro-evolution methods in the context of cognitive functions. On one hand, neuro-evolution ended up with substantially good results to control non-linear systems [4], [16], [17], but despite Beer’s preliminary proposals [18], [19], no clear benchmark has been widely accepted to evaluate their potential for generating more cognitive functions. On the other hand, experimental

neuroscience isolated several modules of the nervous systems, and precisely described their inputs and outputs such that computational neuroscience can model them. On the road to automatically design artificial nervous systems, the minimum benchmark for any neuro-evolution method should be to reproduce the functions modeled by neuro-scientists and to equal the efficiency of hand-designed neural networks.

The EvoNeuro project builds on this inspiring parallel and the present paper showcases a few salient results from this project<sup>1</sup> [20]–[23]. After an introduction to the map-based encoding, we summarize the studies made around the properties of the encoding on the three following aspects: scaling, versatility and learning abilities.

## II. MAP-BASED ENCODING

Many computational neuroscience models can be described as a graph of neural maps in which each map and each connection is described by a set of parameters (figure 1). In our model, each edge is associated with three parameters:

- connection type (1 to 1, 1 to all with uniform synaptic weights, 1 to all with Gaussian weights);
- synaptic weight (all connections between maps have the same strength) or parameters of the Gaussian (if the connection is of Gaussian type).

Similarly, three parameters describe each node:

- isolated neuron or map of neurons (a Boolean value);
- inhibitory or excitatory (the whole map will be inhibitory or excitatory);
- parameters of the neuron (time constant, threshold).

In this work, all maps have the same size. Such a graph of neural maps is developed into a full neural network by analyzing each node and each edge of the graph to create the corresponding neurons, maps and connections.

Mutation can modify such a graph structurally (add/remove a connection or a node) and parametrically (change of a label).

## III. SCALING: WINNER-TAKES-ALL INSPIRED BY THE BASAL GANGLIA

The basal ganglia are a set of interconnected subcortical nuclei [24], that are thought to be involved in action selection [25], [26], i.e. the problem, for an agent, of choosing which action to perform within a repertoire, given internal and external sensory information, in order to achieve its goals. The BG performs the two main aspects of action selection:

<sup>1</sup>Parts of the current text and all the figures have been previously published [20]–[23].




Fig. 1. Overview of the development process. From left to right: (1) the genotype is a labeled graph with evolvable labels; (2) the labels are interpreted to a neuroscience-inspired description of the neural network; (3) for a given size of maps, this neural network can be fully developed into a neural network (for instance to evaluate its fitness).




Fig. 2. Proportion of runs that satisfy the constraints (fitness > 0) after 200,000 evaluations and average fitness of those runs (10 runs of each variants have been launched). All differences are statistically significant (Student T-test,  $p < 0.01$ ) except between map-based (6) and map-based (15). CBG corresponds to a computational neuroscience model [9] whose parameters are optimized by an evolutionary algorithm.

the central process of selection of one action only among conflicting ones, similar to a winner-takes-all (WTA), and the learning process necessary to bias the selection process towards the most profitable option. We focus here on the WTA-like process for which various models have been proposed [8], [9], [27].

As a first test of the map-based encoding, we wanted to know (1) if this new encoding can reproduce an important function of the brain (the WTA), (2) if the new encoding outperforms a classic direct encoding for this task and, (3) how the new encoding can scale up to many inputs/outputs.

The objective function (fitness) counts the number of correct outputs vectors for 1000 random inputs: for each input vector, if the maximum salience corresponds to input  $i$ , then output  $i$  must be the lowest output. Furthermore, the fitness rewards the contrast between outputs, so that the selected action is as desinhibited as possible and the others actions are as inhibited as possible.

Results (Fig. 2) show that evolving neural network with the proposed map-based encoding leads to efficient networks which are at least as efficient as the optimized reference model [9]. The new encoding also outperforms the classic direct encoding when 6 or 15 input channels are used. Since

the neural network description is independent of the size of the map, results obtained with 6 or 15 channels are statistically not different. This result demonstrates that the map-based approach can potentially generate very large, functional and organized neural networks. More details about these results are to be found in [20].

#### IV. VERSATILITY: WORKING MEMORY CIRCUIT

Defined as “the ability to transiently hold and manipulate goal-related information to guide forthcoming actions” [28], working memory is a basic cognitive ability that has motivated lots of modeling research in computational neuroscience [28]–[30].

The synthesis of a neural network exhibiting this functionality was studied to first check whether the proposed encoding can generate it or not. The chosen setup is the AX-CPT task [29], [31], [32]. In this setup, the network has four inputs: A, B, X and Y and two outputs: target and non-target. The target output should be above the non-target one each time an A is followed by a X. The rest of the time, the non-target output should be above the target one. To solve the task, the network should then memorize each A and trigger the target output if and only if an X follows.

In a second step, we considered the case in which the signal to keep in memory, i.e. A in the original setup, is changed. To mimic a learning mechanism, we tried to evaluate whether the generated networks can be adapted to a new situation by means of a parameter optimization only, i.e. without change in the structure of the network. The goal of this experiment was to evaluate the versatility of the networks generated by the proposed encoding. It should be highlighted here that we did not reward versatility during the topology evolution experiment, the goal was to evaluate whether the proposed encoding had a tendency towards versatility or not.

The questions that were investigated in this work are then the following: (1) does the map based encoding facilitate the synthesis of a working memory functionality? (2) are the networks generated with the map based encoding more versatile?

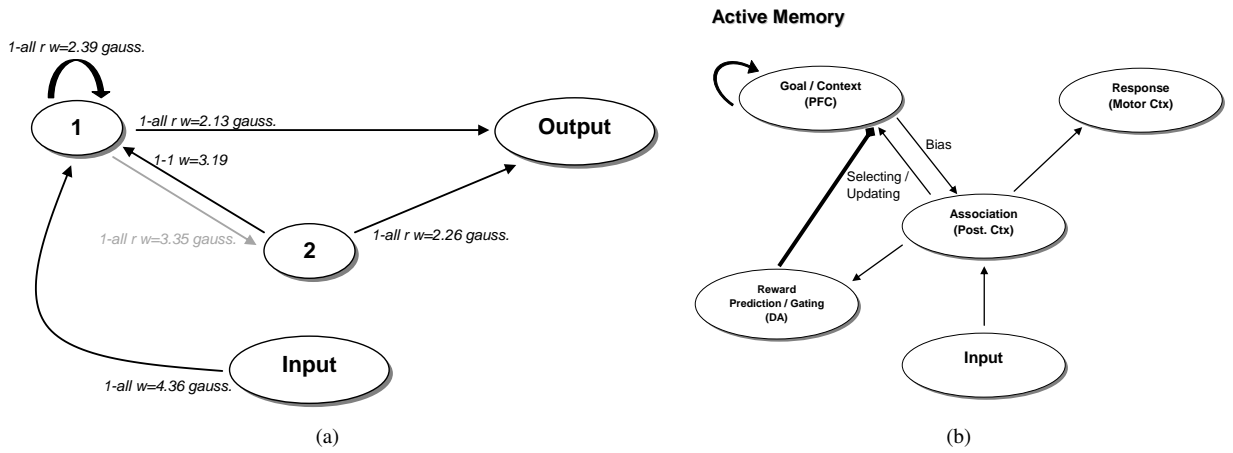


Fig. 3. (a) Example of a neural network obtained with map-based encoding for the AX-CPT working memory task. In this case, each map is composed of 4 neurons. 1-1 represents one to one connections between maps, 1-all, one to all connections, gauss., weights following a gaussian distribution; (b) Minimalist canonical model of cognitive control proposed by Braver [32]




Fig. 4. Minimalist neural network obtained by direct encoding for the AX-CPT working memory task. In this case, each circle represents one single neuron.

For the first question we compared the synthesis capability of the map-based encoding to that of a simpler direct encoding. The map based encoding reveals to be able to generate working memory networks with a similar success rate (5 out of 10 runs versus 4 out of 10 runs), although it requires more generations (1016 generations on average versus 656 generations).

The networks generated with the map based encoding look more generic and even are, sometimes, similar to computational neuroscience models (fig. 3), while those generated with a direct encoding are specific to the sequence to recognize (fig. 4). To validate this impression and answer the second question, the topologies of those solutions were kept and their parameters optimized to solve a task in which the target sequence was BY instead of AX. None of the topologies generated with the direct encoding succeeded in solving this new task, while 4 out of 5 networks generated with the map based encoding lead to solutions.

These results show that the map-based encoding generates networks that exhibit the desired functionality while remaining much more adaptable than a simple direct encoding. The regularity included in the encoding results in a versatility “for free”, i.e. without the need to explicitly reward it during the

evolutionary process. For further details, see [21].

## V. LEARNING ABILITIES: THE SKINNER BOX

Synaptic plasticity underlies most models of learning, memory and development in animals [33], [34]. It has been described at many levels of detail, but studies on the evolution of plastic artificial neural networks for intelligent agents are mainly focused on Hebbian-like adaptation rules, according to which the strength of connection is modified with regard to pre- and post-synaptic activity [35]–[40]. Here we rely on *modulated Hebbian plasticity* [38]–[40], i.e. for each synapse, Hebbian plasticity is modulated by the output of special modulatory neurons, whose activation depends on the particular circuit of the evolved neural network.

Nonetheless, evolving plastic ANNs currently requires long fitness evaluations because one must ensure that each possible learning scenario (e.g. different positions of reward or different sizes of robots) can be learned. The number of scenarios tends to grow exponentially with the number of alternatives, therefore testing most of them when evaluating the fitness arguably prevents the evolution of plastic ANNs for anything else than toy problems. Besides this computational issue, one of the goals of designing plastic ANNs is to make agents able to react to *unknown situations* which will obviously not be known during the evolutionary process. Put differently, a lot of computation time is employed to encourage the evolutionary process to find a *general learning system* and not adaptation rules that are specialized for the situations tested during the fitness evaluation.

We hypothesized that using a map-based encoding will drastically reduce the number of tests required to evolve neural networks with *synaptic general learning abilities* (sGLA): if a learning system is found to work for a few sample situations (ideally one situation), the regularity implied by the map system should have replicated this system to make a generic structure that will be able to learn any similar task.

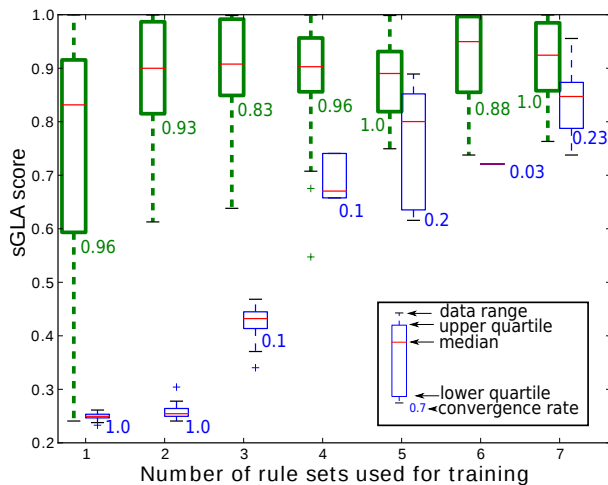


Fig. 5. Success rate (30 experiments) on learning association sets different from the ones used in evolution with the map-based encoding (green bold boxes) and with a direct encoding (blue thin boxes). Only successful runs ( $Fit(x) = 1$ ) are taken into account and numbers show the proportion of runs that reached this perfect fitness.

We tested this hypothesis in a simulation of the classic Skinner box used in operant conditioning [41]: an agent is placed in a cage with four stimuli (lights), four actions (levers), positive rewards (food) and punishments (electric shocks). The goal of the agent is to learn the right associations between each stimulus and each action (for each light, the agent must press the matching lever; since we only allow one to one matches, there are 256 possible sets of 4 associations). The goal of the experiment is to evolve an agent that is able to learn any set of them during the evaluation of the fitness function. We call the number of tests required to obtain such neural networks with GLA the “sTLA-level” (synaptic Transitive Learning Abilities level) and the proportion of successfully learned association set the “sGLA score”.

We launched 7 sets of experiments, each one using a different number of association sets during the fitness evaluation (from 1 to 7). After 2000 generations of 200 individuals, we tested the ability of the individuals with a perfect fitness to learn each of the 256 possible association sets.

Experimental results validate our hypothesis (figure 5): when the map-based encoding was used, individuals that perfectly learned 1 or 2 association sets during the fitness evaluation were able to learn about 90% of the other, unknown sets. By contrast, using a direct encoding led to networks with significantly lower sTLA scores, even when a large number of sets were used.

## VI. CONCLUSIONS

Taking inspiration from some computational neuroscience models, we proposed a new encoding based on the evolution of a graph of neural maps. Using such maps instead of individual neurons makes this encoding a very simple model of a developmental encoding. Successive works to evaluate this

encoding [20]–[22] demonstrated that it allows the evolution of networks that are (1) scalable [20], (2) versatile [21] and, if Hebbian plasticity is added to the neural networks, (3) with good general learning abilities [23]. These results show how fruitful the interactions between neuro-evolution and computational neuroscience can be.

## VII. ACKNOWLEDGEMENT

This project was funded by the EvoNeuro ANR-09-EMER-005-01 project.

## REFERENCES

- [1] P. Dayan, L. Abbott, and L. Abbott, *Theoretical neuroscience: Computational and mathematical modeling of neural systems*. MIT Press, 2001.
- [2] T. P. Trappenberg, *Fundamentals of Computational Neuroscience (2. ed.)*. Oxford University Press, 2009.
- [3] J.-A. Meyer, S. Doncieux, D. Filliat, and A. Guillot, *Biologically Inspired Robot Behavior Engineering*. Springer-Verlag, 2002, ch. Evolutionary Approaches to Neural Control of Rolling, Walking, Swimming and Flying Animals or Robots, pp. 1–43.
- [4] K. O. Stanley and R. Miikkulainen, “Evolving neural networks through augmenting topologies,” *Evolutionary Computation*, vol. 10, no. 2, 2002.
- [5] J.-B. Mouret and S. Doncieux, “MENNAG: a modular, regular and hierarchical encoding for neural-networks based on attribute grammars,” *Evolutionary Intelligence*, vol. 1, pp. 187–207, 2008.
- [6] —, “Evolving modular neural-networks through exaptation,” in *Proc. of IEEE-CEC*, 2009.
- [7] K. Gurney, T. J. Prescott, and P. Redgrave, “A computational model of action selection in the basal ganglia. I. A new functional anatomy,” *Biological cybernetics*, vol. 84, no. 6, pp. 401–410, 2001.
- [8] —, “A computational model of action selection in the basal ganglia. II. Analysis and simulation of behaviour,” *Biological cybernetics*, vol. 84, no. 6, pp. 411–423, 2001.
- [9] B. Girard, N. Tabareau, Q. C. Pham, A. Berthoz, and J.-J. Slotine, “Where neuroscience and dynamic system theory meet autonomous robotics: a contracting basal ganglia model for action selection,” *Neural Networks*, vol. 21, no. 4, pp. 628–641, 2008.
- [10] T. Trappenberg, M. Dorris, D. Munoz, and R. Klein, “A model of saccade initiation based on the competitive integration of exogenous and endogenous signals in the superior colliculus,” *Journal of Cognitive Neuroscience*, vol. 13, no. 2, pp. 256–271, 2001.
- [11] J. Vitay, N. Rougier, and F. Alexandre, “A distributed model of spatial visual attention,” in *Biomimetic Neural Learning for Intelligent Robots*. Springer, 2005, pp. 54–72.
- [12] N. Rougier and J. Vitay, “Emergence of attention within a neural population,” *Neural Networks*, vol. 19, no. 5, pp. 573–581, 2006.
- [13] N. Rougier, “Dynamic neural field with local inhibition,” *Biological cybernetics*, vol. 94, no. 3, pp. 169–179, 2006.
- [14] H. Lipson, “Principles of modularity, regularity, and hierarchy for scalable systems,” *Journal of Biological Physics and Chemistry*, vol. 7, no. 4, pp. 125–128, 2007.
- [15] G. S. Hornby and J. B. Pollack, “Creating high-level components with a generative representation for body-brain evolution,” *Artificial Life*, vol. 8, no. 3, pp. 223–246, 2002.
- [16] S. Doncieux and J.-A. Meyer, “Evolving modular neural networks to solve challenging control problems,” in *Proc. of EIS 2004*, 2004.
- [17] F. Gruau, “Automatic definition of modular neural networks,” *Adaptive Behaviour*, vol. 3, no. 2, pp. 151–183, 1995.
- [18] R. Beer, “Toward the evolution of dynamical neural networks for minimally cognitive behavior,” in *From Animals to Animals 4*, 1996, pp. 421–429.
- [19] A. Slocum, D. Downey, and R. Beer, “Further experiments in the evolution of minimally cognitive behavior: From perceiving affordances to selective attention,” in *From Animals to Animals 6*, 2000, pp. 430–439.
- [20] J.-B. Mouret, S. Doncieux, and B. Girard, “Importing the computational neuroscience toolbox into neuro-evolution—application to basal ganglia,” in *Proceedings of GECCO’10*, 2010, pp. 587–594.

- [21] T. Pinville and S. Doncieux, "Automatic synthesis of working memory neural networks with neuroevolution methods," in *Proceedings of Neurocomp 2010*, 2010, pp. 89–94.
- [22] P. Tonelli and J.-B. Mouret, "Using a map-based encoding to evolve plastic neural networks," in *Proceedings of IEEE Symposium Series on Computational Intelligence*, 2011.
- [23] —, "On the relationships between synaptic plasticity and generative systems," in *Proceedings of GECCO'11*, 2011, pp. 1531–1538.
- [24] P. Redgrave, "Basal ganglia," *Scholarpedia*, vol. 2, no. 6, p. 1825, 2007.
- [25] P. Redgrave, T. Prescott, and K. N. Gurney, "The basal ganglia: A vertebrate solution to the selection problem?" *Neuroscience*, vol. 89, pp. 1009–1023, 1999.
- [26] T. Prescott, "Action selection," *Scholarpedia*, vol. 3, no. 2, p. 2705, 2008, doi:10.4249/scholarpedia.2705.
- [27] A. Gillies and G. Arbruthnott, "Computational models of the basal ganglia," *Movement Disorders*, vol. 15, no. 5, pp. 762–770, 2000.
- [28] D. Durstewitz, J. K. Seamans, and T. J. Sejnowski, "Neurocomputational models of working memory." *Nat. neurosci.*, vol. 3 Suppl, pp. 1184–1191, November 2000.
- [29] T. S. Braver, J. D. Cohen, and D. Servan-Schreiber, "A computational model of prefrontal cortex function," *Nips*, p. 141148, 1995.
- [30] R. C. O'Reilly and M. J. Frank, "Making working memory work: a computational model of learning in the prefrontal cortex and basal ganglia." *Neural comput.*, vol. 18, no. 2, pp. 283–328, February 2006.
- [31] T. S. Braver, D. M. Barch, and J. D. Cohen, "Cognition and control in schizophrenia: a computational model of dopamine and prefrontal function." *Biol. psychiat.*, vol. 46, no. 3, pp. 312–28, August 1999.
- [32] T. S. Braver and D. M. Barch, "A theory of cognitive control, aging cognition, and neuromodulation." *Neurosci. biobehav. r.*, vol. 26, no. 7, pp. 809–17, November 2002.
- [33] L. F. Abbott and S. B. Nelson, "Synaptic plasticity: taming the beast." *Nature neuroscience*, vol. 3, pp. 1178–83, 2000.
- [34] T. P. Trappenberg, *Fundamentals of Computational Neuroscience (2. ed.)*. Oxford University Press, 2009.
- [35] J. Blynel and D. Floreano, "Levels of dynamics and adaptive behavior in evolutionary neural controllers," in *Proc. of SAB*, 2002.
- [36] E. Tuci and M. Quinn, "Behavioral Plasticity in Autonomous Agents: a Comparison Between Two Types of Controller," *Proc. of Applications of evolutionary computing*, 2003.
- [37] J. Urzelai and D. Floreano, "Evolution of adaptive synapses: Robots with fast adaptive behavior in new environments," *Evolutionary Computation*, vol. 9, no. 4, pp. 495–524, 2001.
- [38] Y. Niv, D. Joel, I. Meilijson, and E. Ruppin, "Evolution of Reinforcement Learning in Uncertain Environments: A Simple Explanation for Complex Foraging Behaviors," *Adaptive Behavior*, vol. 10, no. 1, pp. 5–24, 2002.
- [39] A. Soltoggio, J. Bullinaria, C. Mattiussi, P. Dürri, and D. Floreano, "Evolutionary advantages of neuromodulated plasticity in dynamic, reward-based scenarios," *Artificial Life*, vol. 11, p. 569, 2008.
- [40] S. Risi, S. D. Vanderbleek, C. E. Hughes, and K. O. Stanley, "How novelty search escapes the deceptive trap of learning to learn," *Proc. of GECCO*, 2009.
- [41] B. F. Skinner, "the Experimental Analysis of Operant Behavior," *Annals of the New York Academy of Sciences*, vol. 291, no. 1 The Roots of, pp. 374–385, Apr. 1977.