

Evolution of Plastic Control Networks

Dario Floreano

*Institute of Robotic Systems,
Swiss Federal Institute of Technology (EPFL), Lausanne, Switzerland*

Joseba Urzelai

ELCA Informatique SA, Lausanne, Switzerland

Abstract. Evolutionary Robotics is a powerful method to generate efficient controllers with minimal human intervention, but its applicability to real-world problems remains a challenge because the method takes long time and it requires software simulations that do not necessarily transfer smoothly to physical robots. In this paper we describe a method that overcomes these limitations by evolving robots for the ability to adapt on-line in few seconds. Experiments show that this method require less generations and smaller populations to evolve, that evolved robots adapt in a few seconds to unpredictable change -including transfers from simulations to physical robots- and display non-trivial behaviors. Robots evolved with this method can be dispatched to other planets and to our homes where they will autonomously and quickly adapt to the specific properties of their environments if and when necessary.

Keywords: Evolutionary Robotics, Neural Networks, Evolution and Learning

1. Evolutionary Robotics in the Physical World

Evolutionary Robotics (Cliff et al., 1993; Nolfi and Floreano, 2000) implies the use of a population, but in most cases only one physical robot is used to test the performance of each individual in the population one after the other. Since most experiments resort to populations of a few hundred individuals evolved for a few hundred generations, a single evolutionary run with physical robots may take weeks or months. A typical solution to this problem is to use simulations and then transfer the best evolved controller to the physical robot. However, evolved controllers do not always transfer smoothly to physical robots. Researchers have suggested solutions that under some conditions ensure a smooth transfer, such as sampling sensors and motors and storing values with added noise in lookup tables (Miglino et al., 1996) or including in the simulation only those aspects of the interactions between the robot and the environment that are important for the expected behavior while adding noise to them (*minimalistic simulations*) (Jakobi, 1997). However, these two approaches are not always suitable. On the one hand, it is not always possible to sample all robot sensors in all possible environments that the robot will encounter. On the other hand, minimalistic simulations require the the researcher knows in advance



© 2001 Kluwer Academic Publishers. Printed in the Netherlands.

what are the aspects that will matter when the controller is transferred to the physical robot, which implies a lot of trials and errors.

Another limitation is that evolved controllers display behaviors that are fine-tuned to the properties of the environment used during evolution. Environmental change is likely to cause a failure of the evolved control system. Furthermore, since evolution exploits relationships between the robot and its environment that are not always visible to an external observer, it is hard to predict under what type of environmental change the system will fail (it may be lighting conditions, environment layout, electro-mechanical features, etc., or any combination of them). Evolving systems in a variety of different environments (Thompson et al., 1999) is not a long-term solution because the evolved solutions will be robust only for those aspects of the environment that have been varied, but not necessarily for others. Furthermore, testing robots in a variety of different environments may require the use of simulations, which brings us back to the transfer problem mentioned above. Another solution consists of incrementally evolving the robot as the environment changes over generational time (Harvey et al., 1994; Floreano and Mondada, 1998), but this takes long time. Even the issue of smooth transfer from simulations to physical robots is an instance of this more general problem, namely the ability for an evolved system to self-adapt online to unpredictable sources of change without requiring incremental evolution and human supervision. One can think of the transfer from simulated to physical robots as a case where the control system is suddenly faced with slightly different conditions in sensory, motor, and mechanical response.

The ability to cope quickly and reliably with unpredictable change is therefore a top priority for Evolutionary Robotics and its applicability for developing robots that are expected to operate on the surface of other planets and in our homes because *a)* it is almost guaranteed that the operating conditions will be different than those included during evolutionary training; *b)* it is not possible to let the robot incrementally evolve on the surface of the planet or at the customer's home; *c)* even if the initial operating conditions match the evolutionary conditions, it is almost inevitable that some unpredictable change will take place during the life of the robot.

2. Evolution of Plastic Neural Controllers

We have recently suggested a method for evolving neural controllers of robots that can cope on-line with a large variety of unpredictable change (Floreano and Urzelai, 2000; Urzelai and Floreano, 2001). The

core of the method consists in evolving the rules of online adaptation, instead of the connection strengths of the neural controller. Since the evolved adaptative rules are based on local Hebbian learning, evolved controllers do not require external supervision or reinforcement signals and the method is applicable to any neural architecture.

The genetic string encodes the architecture and a set of Hebbian rules, but *not* the synaptic strengths. Every time a genotype is decoded into a neural controller, its synaptic values are *randomly initialized* (always, from the first to the last generation) and are let free to adapt for ever using the genetically-specified Hebbian rules while the robot operates in the environment. We call this approach evolution of *Plastic Controllers* to emphasize the fact that they can continuously change online. The random initialization of synaptic strengths does not allow the genetic string to encode a strategy that fits a particular environment (which may not function properly if the environment changes after evolution), but rather forces evolution to generate individuals capable of developing on the fly a suitable strategy to cope with the features of the environment where they are positioned. Since each robot is evaluated for a short amount of time, there is selection pressure to discover combinations of adaptive rules that allow the controller to develop quickly and reliably the required behavioral abilities.

We have compared a conventional approach based on encoding the strengths of the synapses of the neural controller with our approach where we encode a set of Hebbian rules, the learning rate, and the sign of each neuron (that means that all synapses afferent to a given neuron will use the same genetically-encoded rule and learning rate for that neuron). The genetic code can select for each neuron one of four Hebbian rules: 1) *Plain Hebb* rule increments the synaptic strength proportionally to the correlated activity of the pre- and postsynaptic neurons; 2) *Postsynaptic rule* changes the synaptic strength only if the postsynaptic neuron is active, incrementing it if the presynaptic neuron is active, otherwise decrementing it; 3) *Presynaptic rule* is similar to postsynaptic, but changes the synaptic strength only if the pre-synaptic neuron is active; 4) *Covariance rule* increments and decrements synaptic strength depending whether the correlated activity of the two neurons is above or below a threshold. Details of the genetic encoding and Hebbian rules are to be found in (Floreano and Urzelai, 2000).

We have used a robot (figure 1, left) equipped with a vision module, proximity sensors, and light sensors. This robot can gain fitness points only when it sits on the grey area on the left *when the light is on*. At the beginning of each robot “life”, the light is off but it can be switched on if the robot goes to the black area on the right of the arena.




Figure 1. Left: A Khepera robot equipped with a vision module can gain fitness points only when it is sitting on the light (grey zone on the left) *when the light is on*. Initially the light is off, but the robot can switch it on by going over the black area on the right. No fitness points are given for the light switching behavior. **Right:** Trajectory of an evolved robot that adapts its connection strengths on the fly using the genetically-specified Hebbian rules. The figures on top show the fitness of this trajectory and the average fitness of this individual when tested ten times with different synaptic random initializations and starting positions.

Therefore, in order to receive fitness points this robot must evolve the ability to find the light switching area, go there, and, once the light is switched on, rapidly move towards the fitness area and remain there for the rest of its life.¹ Our adaptive approach can generate very quickly neural controllers that solve this problem in a very reliable and efficient manner, whereas the traditional approach takes almost twice as many generations and the result is a much less efficient.

The key difference is that the conventional approach comes up with a so-called “minimalistic solution” whereby it performs looping trajectories around the arena without any sensitivity to the visual pattern and to the light. The evolved turning angle of these robots is such that the trajectories has a high chance of taking the robot at some point over the light switch and then over the light bulb. The trajectories are finely tuned to the geometry of the environment, the reflection properties of the walls, and to the overall light intensity. Instead, the evolved plastic robot develops on the fly a set of synaptic connections for each of the sub-behaviors required by this environment. Initially, it develops the ability to avoid walls, then to locate and navigate towards the visual signal corresponding to the light switch, then to turn away and become attracted by light, and finally, once under the light bulb, to remain still (figure 1, right). Each of these abilities are developed on the fly through a rapid self-reconfiguration of the synaptic strengths on the basis of the available sensory information and actions taken by the robot.

¹ Note that the robot cannot see the patches on the floor and therefore must look at the patterns on the walls and at the light.

2.1. POPULATION SIZE AND TRANSFER FROM SIMULATIONS

As we mentioned above, the amount of time required by the evolutionary process is directly proportional to the size of the population. Therefore, genetic encodings that require smaller population sizes are very important for evolution of physical robots. We have performed an experiment in order to study the effect of the population size on the different types of genetic encoding methods that we have used in our experiments. We have evolved populations of different sizes (200, 100, 75, 50, 35, and 25 individuals) using both evolution of adaptive rules (Node Encoding of adaptive synapses) and evolution of synaptic strengths (Synapses Encoding of genetically-determined synapses). The results shown in figure 2.1 indicate that evolution of adaptive synapses generate performant solutions even in the case of relatively small populations. Performance is not significantly affected for populations composed of at least 50 individuals. In populations composed of less than 50 individuals, some of the replications cannot find a solution, as indicated by the standard deviation (thin lines) of best fitness values. In the case of Synapse Encoding and genetically-determined synapses instead, there is a significant loss in performance if the population is composed of less than 200 individuals, since no solution is found in some of those replications. In addition, it should be noticed that best fitness curves (thick lines) report higher values in the case of Node Encoding of adaptive synapses for every population size.

In another set of experiments, we have evolved the control systems in simulations and then transferred the evolved controllers on the physical robot. The graph on the left of figure 3 shows that evolved plastic controllers report a relatively higher fitness when transferred to the physical robot than evolved controllers with genetically-determined and fixed synapses. The bars show the average of ten tests for each of the best individual of ten evolutionary runs (that is, the average of one hundred tests). It is important to notice that the fitness drop for evolved adaptive controllers is uniquely due to the fact that the neural controllers take a few more seconds to adapt to the physical robot but always manage to accomplish the entire task, whereas the fitness drop of the genetically-determined and fixed controller is due to the fact that most of the individuals remain stuck against a wall.

An analysis of evolved plastic controllers (3, right) shows that the two motor neurons that set the speeds of the wheels display a marked preference Hebbian rules that can both increment and decrement synaptic strength. In particular, at least one of these neurons always uses the postsynaptic rule. In other words, the strengths of its afferent synapses are changed only when the neuron becomes active

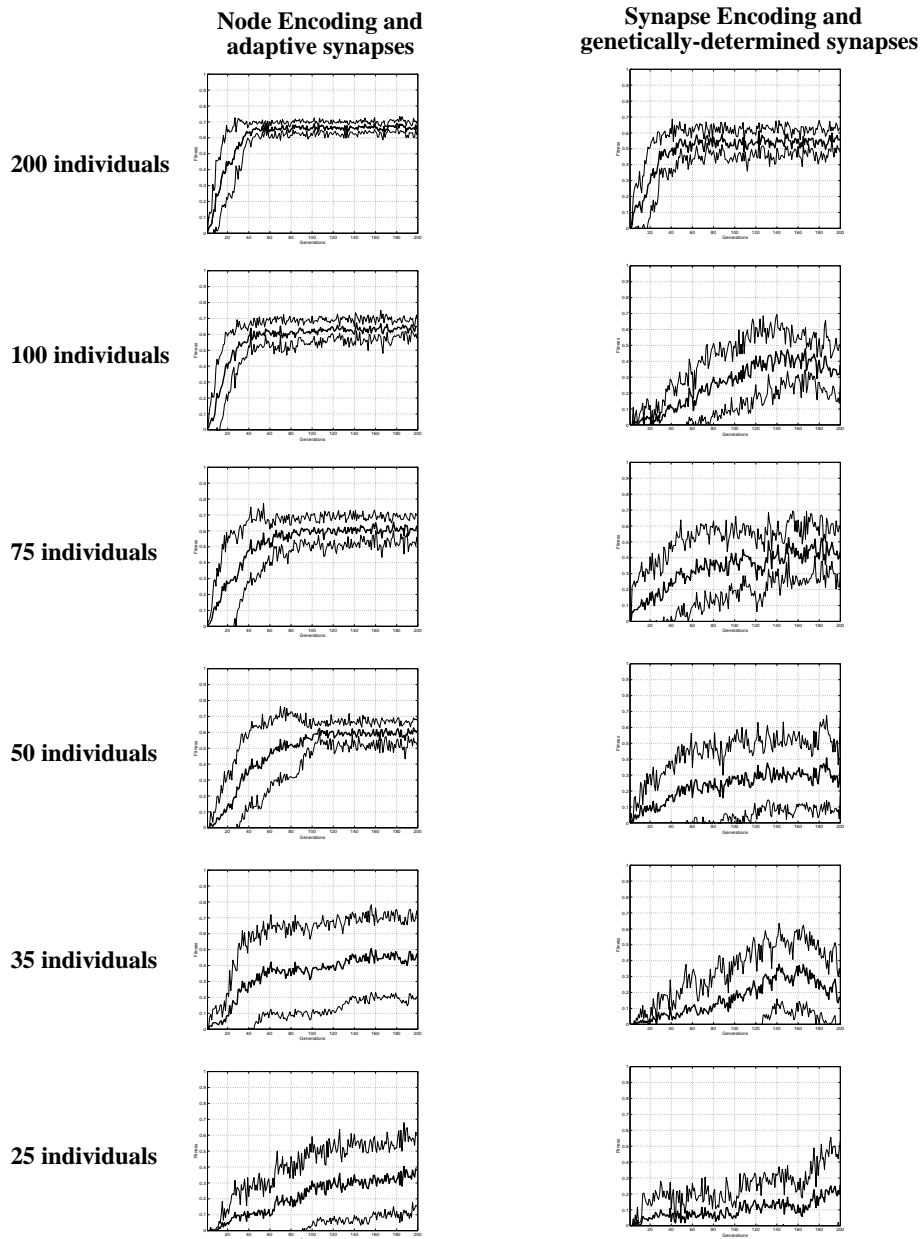


Figure 2. Evolvability of adaptive synapses encoded at node level and of genetically-determined synaptic strengths with respect to the size of the genetic population. Thick lines represent best fitness values and thin lines show the standard deviation. Each data point is an average over 10 replications with different random initializations.




Figure 3. Left: Performance of best controller evolved in simulation and transferred to the physical robot for the two conditions: plastic networks and genetically-determined and fixed networks. Each bar is the average of 10 tests of the best individuals of 10 evolutionary runs (100 data points). *Right:* Frequency of learning rules used for motor neurons during evolution. One of the two motor neurons is always controlled by the postsynaptic rule (frequency = 0.5), and the other motor neuron obeys either the presynaptic or the covariance rule. After the initial 100 generations, the plain Hebb rule is not used anymore.

according to the activation values of the sensory neurons. In other words, this is an indication that variations in the sensory signals that affect motor neurons are used by the learning rules to selectively modify synaptic strengths in order to develop and maintain a stable overall behavior.

3. Conclusions

We have described a method to evolve robot neural controllers for their ability to use synaptic plasticity in order to develop behavioral abilities on the fly. This method can produce controllers that display non-trivial abilities, do not rely on minimalistic solutions, and require less generations and smaller populations when compared to the conventional evolution of connection strengths. We have also shown that evolved plastic controllers can transfer smoothly from simulations to physical robots by adapting online to the new sensory-motor conditions.

In a set of experiments reported elsewhere, we have also shown that such evolved plastic controllers adapt on-line to new conditions, such as *a*) change of wall color, which drastically affect the response of the infrared sensors (Urzelai and Floreano, 2000b); *b*) change of environmental layout (moving around the arena the lightswitch and the lightbulb)

(Urzelai and Floreano, 2000a); *c* transferring the controller evolved for the Khepera robot on another robot with different size, morphology, sensory layout, and motor response (Urzelai and Floreano, 2000a); *d* on-line change of behavior of another robot operating within the same physical space and competing for resources (Floreano et al., 2001).

It would be interesting to compare the method presented in this paper to other methodologies that exploit evolution for generating dynamical controllers that display adaptive behavior without using synaptic plasticity (Yamauchi and Beer, 1995) and for discovering the rules by which the logic functions of FPGA controllers change according to environmental situations (Keymeulen et al., 1999).

We believe that this methodology not only solves the problem of material and time resources often required by conventional evolutionary approaches (Mataric and Cliff, 1996), but also exploits evolution to generate controllers that remain adaptive in the face of unpredictable change. This feature could be useful to develop controllers for real-world applications. For example, the control system of a robot for extraplanetary exploration could be evolved in a laboratory environment that recreates the conditions of that planet. The evolved robot could then be sent to the surface of the planet where, once landed, it would trigger the process of synaptic adaptation and develop the required functionality while remaining adaptive to environmental, mechanical, and electronic sources of unpredictable change. A similar procedure could be applied to develop control systems of personal and service robots that would adapt to the features of the individual environments where they are put in operation.

Acknowledgements

DF thanks the Jet Propulsion Laboratory for the invitation to the workshop on Biomorphic Robotics and the referees for useful comments that helped to improve this article. DF acknowledges support from the Swiss National Science Foundation, grant no. 620-58049. JU has been supported by grant BF197.136-AK from the Basque government.

References

- Cliff, D., I. Harvey, and P. Husbands: 1993, 'Explorations in Evolutionary Robotics'. *Adaptive Behavior* **2**, 73–110.
- Floreano, D. and F. Mondada: 1998, 'Evolutionary Neurocontrollers for Autonomous Mobile Robots'. *Neural Networks* **11**, 1461–1478.

- Floreano, D., S. Nolfi, and F. Mondada: 2001, 'Co-Evolution and Ontogenetic Change in Competing Robots'. In: M. Patel, V. Honavar, and K. Balakrishnan (eds.): *Advances in the Evolutionary Synthesis of Intelligent Agents*. Cambridge, MA: MIT Press.
- Floreano, D. and J. Urzelai: 2000, 'Evolutionary Robots with Online Self-Organization and Behavioral Fitness'. *Neural Networks* **13**, 431–443.
- Harvey, I., P. Husbands, and D. Cliff: 1994, 'Seeing The Light: Artificial Evolution, Real Vision'. In: D. Cliff, P. Husbands, J. Meyer, and S. W. Wilson (eds.): *From Animals to Animats III: Proceedings of the Third International Conference on Simulation of Adaptive Behavior*. Cambridge, MA: MIT Press-Bradford Books.
- Jakobi, N.: 1997, 'Half-baked, Ad-hoc and Noisy: Minimal Simulations for Evolutionary Robotics'. In: P. Husbands and I. Harvey (eds.): *Proceedings of the 4th European Conference on Artificial Life*. Cambridge, MA, MIT Press.
- Keymeulen, D., M. Iwata, Y. Kuniyoshi, and T. Higuchi: 1999, 'On-line evolution for self-adapting robotic navigation system using evolvable hardware'. *Artificial Life* **4**, 359–393.
- Matarić, M. and D. Cliff: 1996, 'Challenges in Evolving Controllers for Physical Robots'. *Robotics and Autonomous Systems* **19**(1), 67–83.
- Miglino, O., H. H. Lund, and S. Nolfi: 1996, 'Evolving Mobile Robots in Simulated and Real Environments'. *Artificial Life* **2**, 417–434.
- Nolfi, S. and D. Floreano: 2000, *Evolutionary Robotics: Biology, Intelligence, and Technology of Self-Organizing Machines*. Cambridge, MA: MIT Press.
- Thompson, A., P. Layzell, and R. S. Zebulum: 1999, 'Explorations in Design Space: Unconventional electronics design through artificial evolution'. *IEEE Trans. Evol. Comp.* **3**(3), 167–196.
- Urzelai, J. and D. Floreano: 2000a, 'Evolutionary Robotics: Coping with Environmental Change'. In: J. e. a. Koza (ed.): *Second International Conference on Genetic and Evolutionary Computation*. San Mateo, CA, Morgan Kaufmann.
- Urzelai, J. and D. Floreano: 2000b, 'Evolutionary Robots with Fast Adaptive Behavior in New Environments'. In: T. C. Fogarty, J. Miller, A. Thompson, and P. Thomson (eds.): *Third International Conference on Evolvable Systems: From Biology to Hardware (ICES2000)*. Berlin, Springer-Verlag.
- Urzelai, J. and D. Floreano: 2001, 'Evolution of Adaptive Synapses: Robots with Fast Adaptive Behavior in New Environments'. *Evolutionary Computation* **9**, in press.
- Yamauchi, B. and R. D. Beer: 1995, 'Sequential Behavior and learning in evolved dynamical neural networks'. *Adaptive Behavior* **2**, 219–246.

