

Evolutionary Robots with Fast Adaptive Behavior in New Environments

Joseba Urzelai and Dario Floreano

Laboratory of Microprocessors and Interfaces (LAMI)
Swiss Federal Institute of Technology, CH-1015 Lausanne (EPFL)
Joseba.Urzelai@epfl.ch, Dario.Floreano@epfl.ch

Abstract. This paper is concerned with adaptation capabilities of evolved neural controllers. A method consisting of encoding a set of local adaptation rules that synapses obey while the robot freely moves in the environment [6] is compared to a standard fixed-weight network. In the experiments presented here, the performance of the robot is measured in environments that are different in significant ways from those used during evolution. The results show that evolutionary adaptive controllers can adapt to environmental changes that involve new sensory characteristics (including transfers from simulation to reality) and new spatial relationships.

1 Evolution and Adaptation

Evolutionary algorithms are widely used in autonomous robotics in order to solve a large variety of tasks in several kind of environments. However, evolved controllers become well adapted to environmental conditions used during evolution, but often do not perform well when conditions are changed. Under these circumstances, it is necessary to carry on the evolutionary process, but this might take long time.

Combination of evolution and learning has been shown to be a viable solution to this problem by providing richer adaptive dynamics [1] than in the case where parameters are entirely genetically-determined. A review of the work combining evolution and learning for sensory-motor controllers can be found in [5, 9].

Instead of simply combining off-the-shelf evolutionary and learning algorithms, in previous work we presented an approach capable of generating adaptive neural controllers by evolving a set of simple adaptation rules [6]. The method consists of encoding on the genotype a set of modification rules that perform Hebbian synaptic changes [2–4] through the whole individual’s life. The results showed that evolution of adaptive individuals generated viable controllers in much less generations and that these individuals displayed more performant behaviors than genetically-determined individuals.

In this paper, we describe two new sets of experiments conceived to measure the *adaptation* capabilities of this approach in environments that are different from those used during evolution. The results are compared to standard evolution of synaptic weights and to evolution of noisy synaptic weights (control condition).




Fig. 1. A mobile robot Khepera gains fitness by finding as fast as possible the stick under the floor. The walls are covered with white paper and the floor is transparent. The robot has a sensor pointing downwards that can detect the stick when it passes over it. The stick can be positioned at any location under the floor.

The sources of change address two major aspects of behavioral robustness: sensory appearance and spatial relationships of key-features of the environment.

2 Experiment I: Changing Sensory Appearances

A mobile robot Khepera is positioned in the rectangular environment shown in figure 1. The walls are covered with paper and the floor, which is transparent, is placed on four supports. A stick is positioned at a random location under the floor¹. Each individual of the population is tested on the same robot, one at a time, for a maximum of 500 sensory motor cycles, each cycle lasting 100 ms. At the beginning of an individual's life, the robot and the stick are positioned at random positions.

The fitness function selects individuals capable of finding the stick in the shortest time,

$$\Phi = 1 - \frac{t}{500},$$

where t represents the number of sensory motor cycles spent by the robot before finding the stick. Since the robot is not allowed to be on the target at the initial cycle, the fitness will never be 1.0. A robot that cannot manage to find the target will be scored with 0.0 fitness.

A light sensor placed under the robot is used to detect the stick and compute the fitness, but it is not given as input to the controller. Once the robot has found the stick or 500 cycles have passed, the robot and the target are randomly repositioned.

Under these circumstances, the exploration strategies used by the robots will depend much on the sensory appearance of the walls. Therefore, in this

¹ A similar environment has been used in simulation by Nolfi [10] for different experimental purposes.




Fig. 2. The neural controller is a fully-recurrent discrete-time neural network composed of 8 neurons giving a total of $8 \times 8 = 64$ synapses (here represented as small squares of the unfolded network). 6 sensory neurons receive additional input from the infrared sensors positioned around the body of the robot and from the motors (l=left; r=right; f=front; b=back). **IR**=Infrared Proximity sensors; **MF**=Motor feedback. Two motor neurons **M** do not receive sensory input; their activation sets the speed of the wheels ($M_i > 0.5$ forward rotation; $M_i < 0.5$ backward rotation)

experiment the sensory characteristics of the wall surfaces will be changed *after* evolution.

The controller is a fully-recurrent discrete-time neural network. Neurons are updated every 100 ms according to the following equation,

$$y_i \leftarrow \sigma \left(\sum_{j=0}^N w_{ij} y_j \right) + I_i,$$

where y_i is the activation of the i th neuron, w_{ij} is the strength of the synapse between presynaptic neuron j and postsynaptic neuron i , N is the number of neurons in the network, $0 \leq I_i < 1$ is the corresponding external sensory input, and $\sigma(x) = (1 + e^x)^{-1}$ is the sigmoidal activation function. $I_i = 0$ for the motor neurons. The controller has access to two types of sensory information: infrared light (object proximity) and speeds of the wheels (motor feedback). The active infrared sensors positioned around the robot detect proximity of walls (up to 4 cm). Their values are pooled into four pairs and the average reading of each pair is passed to a corresponding neuron. Rotation speeds of the wheels² are normalized and passed to the corresponding feedback neurons. An additional neuron is used as bias in order to excite the network when it does not receive any sensory input. Two output neurons are used to set the rotation speeds of the wheels.

Each synaptic weight w_{ij} can be updated after every sensory-motor cycle (100 ms) using one of the four modification rules specified in the genotype.³ The four rules are called Hebbian because they are a function of the pre-synaptic ac-

² Rotation speeds of the wheels can be different from the values set by the motor neurons. For example, when the robot pushes against a wall, motor neurons may output forward rotations but real rotations are 0.

³ These four rules co-exist within the same network.

tivation, of the post-synaptic activation, and of the current value of the weight itself. The *Plain Hebb rule* strengthens the synapse proportionally to the correlated activity of the two neurons. The *Postsynaptic rule* behaves as the plain Hebb rule, but in addition it weakens the synapse when the postsynaptic node is active but the presynaptic is not. Conversely, in the *Presynaptic rule* weakening occurs when the presynaptic unit is active but the postsynaptic is not. Finally, the *Covariance rule* strengthens the synapse whenever the difference between the activations of the two neurons is less than half their maximum activity, otherwise the synapse is weakened. Synaptic strength is maintained within a range $[0, 1]$ (notice that a synapse cannot change sign) by adding to the modification rules a self-limiting component inversely proportional to the synaptic strength itself [2, 3, for more details].

Two types of genetic (binary) encoding are considered: *Synapse Encoding* and *Node Encoding*. Synapse Encoding is also known as direct encoding [11]. Every synapse is individually coded on 5 bits, the first bit representing its sign and the remaining four bits its properties (either the weight strength or its adaptive rule). Node Encoding instead codes only the properties of the nodes of the network. These properties are then applied to all its incoming synapses (consequently, all incoming synapses to a given node have the *same* properties). Each node is characterized by 5 bits, the first bit representing its sign and the remaining four bits the properties of its incoming synapses. Synapse Encoding allows a detailed definition of the controller, but for a fully connected network of N neurons the genetic length is proportional to N^2 . Instead Node Encoding requires a much shorter genetic length (proportional to N), but it allows only a rough definition of the controller.

Independently of the type of genetic encoding, the following three types of properties can be encoded on the last 4 bits. A) *Genetically determined*: Weight strength. The synaptic strength is genetically determined and cannot be modified during “life”. B) *Adaptive synapses*: Adaptive rule on 2 bits (four rules) and learning rate (0.0, 0.3, 0.6, 0.9) on the remaining 2 bits. The synapses are always randomly initialized when an individual starts its life and then are free to change according to the selected rule. C) *Noisy synapses*: Weight strength on 2 bits and a noise range on the remaining two bits (0.0, ± 0.3 , ± 0.6 , ± 0.9). The synaptic strength is genetically determined at birth, but a random value extracted from the noise range is freshly computed and added after each sensory motor cycle. A limiting mechanism cuts off sums that exceed the synaptic range $[0, 1]$. This latter condition is used as a control condition to check whether the effects of Hebbian adaptation amount to random synaptic variability [6, for more details].

2.1 Results

An initial set of experiments has been carried out in simulations sampling sensor activations separately for white, gray, and black walls and adding 5% uniform noise to these values [8]. Robots are evolved in environments with white walls. For each experimental condition (adaptive, genetically-determined, noisy), 10




Fig. 3. Comparison of adaptive synapses with Node Encoding (*left*) versus genetically-determined synapses with Synapse Encoding (*center*) and noisy synapses with Node Encoding (*right*) in white, gray, and black environments.

different⁴ populations of 100 individuals each have been independently evolved for 200 generations. Each individual is tested three times and the fitness value is averaged. The 20 best individuals reproduce by making 5 copies of their genetic string. Strings are crossed over with probability 0.2 and mutated with probability 0.05 (per bit). In the case of adaptive synapses, synaptic weights of individuals are randomly initialized within the range $[0.0, 0.1]$ at the beginning of each test.

Since we are interested in adaptation capabilities of evolved individuals, after 200 generations evolution is stopped and the best individual of the last generation for each of the 10 populations is tested 10 times in the original environment (white walls), 10 times in an environment where walls have been covered with gray paper, and 10 times in an environment with black walls.

Figure 3 shows average fitnesses corresponding to environments with white, gray, and black walls in the case of individuals with adaptive synapses and Node Encoding (*left*), individuals with genetically-determined synapses and Synapse Encoding⁵ (*center*), and individuals with noisy synapses and Node Encoding (*right*). Although performance decreases in gray and black environments, adaptive individuals are capable of successfully finding the target area in all conditions. Instead, genetically-determined individuals can find the target area only in the environment that has been used during evolution (white walls). When tested in gray and black environments, only a few lucky individuals that encounter the target before a wall have non-zero fitness values. Individuals with noisy synapses score very low fitness values in all conditions but they generalize better than genetically-determined individuals⁶.

⁴ Using different sequences of random number.

⁵ Node Encoding for fixed synapses was not capable of solving the original problem, therefore we report results for Synapse Encoding.

⁶ Notice that adaptive individuals report better fitness also in the evolutionary environment. The performance issue has been addressed in another paper [6].

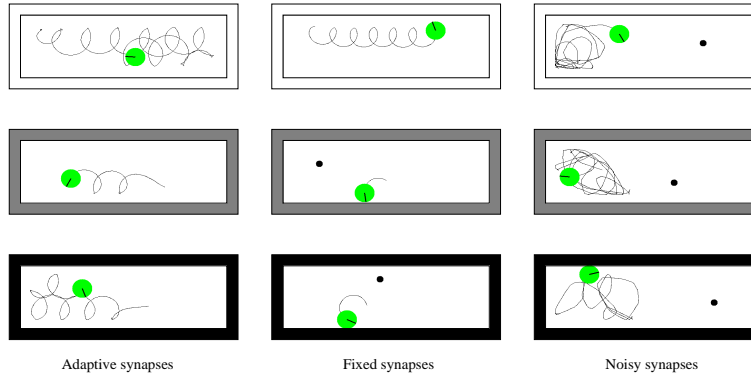


Fig. 4. Behaviors of individuals with adaptive synapses (left), genetically-determined synapses (center), and noisy synapses (right) in environments with white (up), gray (center), and black (bottom) walls. Individuals belong to the last generation.

Figure 4 displays the behaviors of individuals with adaptive synapses (left), genetically-determined synapses (center), and noisy synapses (right) in environments with white (up), gray (center), and black (bottom) walls. The behavior of the adaptive individual is not considerably affected by the color of the walls and it reaches the target area in all conditions. Instead, the genetically-determined individual can reach the target area only when walls are covered with white paper, but gets stuck on gray and black walls. Since darker walls are detected only when the robot gets closer, a behavioral strategy successful for white walls can cause collisions for dark walls. The individual with noisy synapses takes advantage of the random variability to get away from the walls but it scores a low performance because its strategy is based in a local random search. The robot eventually manages to reach the target when its initial position is relatively close to the stick but it fails when it is far away.

2.2 From Simulations to Real Robots

Another way of measuring the adaptive abilities of evolved controllers is to transfer them from simulated to real robots. Since physical robots and environments inevitably have characteristics different from simulations, solutions evolved in simulation typically fail when tested on real robots.

The solutions envisaged so far consist of incorporating special types of noise tailored to sensory-motor properties of the robot [8], or to vary key-features of the environment during simulated evolution [7]. The success of both methods depends upon the ability of the experimenter to spot crucial aspects of variation that must be considered in the simulations. Another solution consists of carrying on artificial evolution in the new conditions [3], but this can take long time.

In another set of experiments, we have tested the best individuals of the last generation for each of the 10 populations evolved in simulation on a real




Fig. 5. Comparison of generalization capabilities in the transfer from simulations to real robots in the case of adaptive synapses with Node Encoding (*left*) versus genetically-determined synapses with Synapse Encoding (*center*) and noisy synapses with Node Encoding (*right*). White bars correspond to the performance in the white simulated environment and striped bars represent the performance on real robots in a white environment.

Khepera robot positioned in an environment where walls are covered with white paper (figure 1). Each individual is tested 3 times and the fitness is averaged over. Figure 5 shows that the performance of adaptive individuals is not affected by the transfer to the physical environment, whereas genetically-determined individuals report a significant fitness loss. Individuals with noisy synapses are not affected by the transfer because their behavior is always random and not effective in both simulated and physical environments. A major reason for failure of genetically-determined individuals is that their spiralling strategy often results in rotation without displacement probably caused by the new sensory and motor responses of the real robot. The same pattern of results holds for tests in gray and black physical environments (data not shown).

3 Experiment II: Changing Spatial Relationships

Whereas the experiments described above were conceived to address mostly variation induced by new sensory responses, in this section we address variation induced by changed spatial relationships. To this end, we resort to an experimental situation where behavioral success is linked to the ability to relate different parts of the environment [6].

A mobile robot Khepera equipped with a vision module is positioned in the rectangular environment shown in figure 6. A light bulb is attached on one side of the environment. This light is normally off, but it can be switched on when the robot passes over a black-painted area on the opposite side of the environment. A black stripe is painted on the wall over the light-switch area. Each individual of the population is tested on the same robot, one at a time, for 500 sensory




Fig. 6. A mobile robot equipped with a vision module gains fitness by staying on the gray area only when the light is on. The light is normally off, but it can be switched on if the robot passes over the black area positioned on the other side of the arena. The robot can detect ambient light and the color of the wall, but not the color of the floor.

motor cycles, each cycle lasting 100 ms. At the beginning of an individual's life, the robot is positioned at a random position and orientation and the light is off.

The fitness function is described as the number of sensory motor cycles spent by the robot on the gray area beneath the light bulb *when the light is on* divided by the total number of cycles available (500). In order to maximize this fitness function, the robot should find the light-switch area, go there in order to switch the light on, and then move towards the light as soon as possible, and stand on the gray area⁷. Since this sequence of actions takes time (several sensory motor cycles), the fitness of a robot is never 1.0. Also, a robot that cannot manage to complete the entire sequence is scored with 0.0 fitness.

A light sensor placed under the robot is used to detect the color of the floor—white, gray, or black—and passed to a host computer in order to switch on the light bulb and compute fitness values. The output of this sensor is *not* given as input to the neural controller. After 500 sensory motor cycles, the light is switched off and the robot is repositioned by applying random speeds to the wheels for 5 seconds.

In a previous article we showed that evolution of adaptive synapses provides a number of advantages with respect to evolution of synaptic weights for this behavioral task. It can generate viable controllers in much less generations and evolved controllers display more performant behaviors [6].

Here we describe a new set of experiments, where the best individuals of the last generation are tested in environments where the light-switching area,

⁷ Notice that the fitness function does not explicitly reward this sequence of actions, but only the final outcome of the overall behavior chosen by the robot. Therefore, we call it a *behavior-based* fitness function.




Fig. 7. Comparison of adaptive synapses with Node Encoding (*left*) versus genetically-determined synapses with Synapse Encoding (*center*) and noisy synapses with Node Encoding (*right*) in the environment used during evolution (white) and in an environment where light-switching area and fitness area are randomly positioned (striped).

the fitness area, and the robot are located at random positions at the beginning of each individual's life. Since in the original experiment the positions of the light-switching area and the fitness area were constant for every individual, this experiment gives us a measure of adaptation capabilities of evolved individuals to new spatial relationships. In order to automate the re-arrangement of the environment, these experiments have been carried out in simulation. The best individuals for each of the 10 populations evolved in the environment of figure 6 are tested in 3 new environments with different random spatial relationships.

The results reported in figure 7 show that individuals with adaptive synapses are much more robust to new configurations of the environment than individuals with genetically-determined synapses. Average performance loss is 25% in the case of adaptive individuals (*left*), but is about 65% in the case of genetically-determined individuals (*center*). Individuals with noisy synapses (*right*) score very low fitness in both cases.

The fact that genetically-determined individuals performed very poorly in new environments indicates that the solutions generated by evolution here are tightly coupled to the geometry and the disposition of the environment. Evolution shapes the individuals in order to take advantage of specific environmental aspects, such as the size of the arena and the position of the light-switching area and of the fitness area. Instead, evolution of adaptive synapses is capable of generating more general solutions that produce performant behavior for a large variety of environmental dispositions. This is shown in figure 8: a genetically-determined individual (*center*) is capable of solving the task in the original environment by performing circular movements and avoiding the walls until it reaches the fitness area. However, these circular movements are not effective to approach the fitness area in the new environmental disposition. Instead, an adaptive individual (*left*) that is capable of solving the task in the original envi-




Fig. 8. Behaviors of individuals with adaptive synapses (left), genetically-determined synapses (center), and noisy synapses (right) in the original environment used during evolution (top) and in an environment where light-switching area and fitness area are positioned at new locations (bottom). The trajectory line is thin when the light is off and becomes thick when the light is turned on. Individuals are the best of the last generation evolved in the environment of top row.

ronment changes the strategy by performing some additional manoeuvres that allow the robot to reach the fitness area in the new environment. Individuals with noisy synapses (right) perform random trajectories in both cases.

4 Conclusions

We have shown through a set of systematic comparisons that evolution of adaptive synapses provides better adaptation capabilities than standard evolution of synaptic weights. Adaptive individuals are capable of successfully performing in environments that are different from the one used during evolution by adapting their strategy to the new constraints of the environment. Instead, genetically-determined individuals often fail in adapting to different environments because their behavior is tightly coupled to the characteristics of the environment used during evolution.

We have studied adaptation to two major sources of environmental change: new sensory appearances and new spatial relationships. In both cases, evolved adaptive controllers can autonomously modify their parameters and behavior online without requiring additional evolutionary training or ad-hoc manipulation of the evolutionary procedure. The control experiments with noisy synapses where changes are induced by random numbers rather than by genetically-determined rules, indicate that evolved adaptive networks modify their parameters in ways that are functionally related to the survival criterion. We are currently analyzing how these evolved controllers work and try to explain how they can adapt to new environmental conditions.

We have also shown that this approach is effective for the transfer of evolved controllers from simulations to real robots. Keeping in mind this idea, one of our current projects aims at studying the applicability of this approach to cross-platform evolution [3]. The experiment will consist in transferring the controller

evolved for the Khepera robot to a bigger Koala robot and in testing its performance in a scaled-up version of the light-switching environment.

Acknowledgements

Joseba Urzelai is supported by grant nr. BF197.136-AK from the Basque government.

References

1. R. K. Belew and M. Mitchell, editors. *Adaptive Individuals in Evolving Populations. Models and Algorithms*. Addison-Wesley, Redwood City, CA, 1996.
2. D. Floreano and F. Mondada. Evolution of plastic neurocontrollers for situated agents. In P. Maes, M. Mataric, J.-A. Meyer, J. Pollack, H. Roitblat, and S. Wilson, editors, *From Animals to Animats IV: Proceedings of the Fourth International Conference on Simulation of Adaptive Behavior*, pages 402–410. MIT Press-Bradford Books, Cambridge, MA, 1996.
3. D. Floreano and F. Mondada. Evolutionary neurocontrollers for autonomous mobile robots. *Neural Networks*, 11:1461–1478, 1998.
4. D. Floreano and S. Nolfi. Adaptive behavior in competing co-evolving species. In P. Husbands and I. Harvey, editors, *Proceedings of the 4th European Conference on Artificial Life*, Cambridge, MA, 1997. MIT Press.
5. D. Floreano and J. Urzelai. Evolution and learning in autonomous robots. In T. Mange and M. Tomassini, editors, *Bio-inspired Computing Systems*. PPUR, Lausanne, Switzerland, 1998.
6. D. Floreano and J. Urzelai. Evolution of Neural Controllers with Adaptive Synapses and Compact Genetic Encoding. In D. Floreano, J.D. Nicoud, and F. Mondada, editors, *Advances In Artificial Life: Proceedings of the 5th European Conference on Artificial Life (ECAL'99)*, pages 183–194. Springer Verlag, Berlin, 1999.
7. N. Jakobi. Half-baked, ad-hoc and noisy: Minimal simulations for evolutionary robotics. In P. Husbands and I. Harvey, editors, *Advances in Artificial Life: Proceedings of the 4th European Conference on Artificial Life*, pages 348–357. MIT Press, 1997.
8. O. Miglino, H. H. Lund, and S. Nolfi. Evolving Mobile Robots in Simulated and Real Environments. *Artificial Life*, 2:417–434, 1996.
9. S. Nolfi and D. Floreano. Learning and evolution. *Autonomous Robots*, 7(1):forthcoming, 1999.
10. S. Nolfi and D. Parisi. Learning to adapt to changing environments in evolving neural networks. *Adaptive Behavior*, 5:75–98, 1997.
11. X. Yao. A review of evolutionary artificial neural networks. *International Journal of Intelligent Systems*, 4:203–222, 1993.