

Cultural Evolution in a Population of Neural Networks

Daniele Denaro
Department of Computer Science University of Ancona
Ancona Italy

Domenico Parisi
Institute of Psychology National Research Council
Rome Italy

Abstract

Genetic algorithms are computational models of the evolution of good solutions to problems based on the selective reproduction of the best variants and the constant addition of random variability to the population of variants. In biological evolution variants are inherited genotypes that are transmitted from parents to offspring. In cultural evolution behavioral variants are transmitted from one individual to another because one individual (the learner) imitates another individual (the teacher). If the two individuals belong to successive generations, reproduction of teachers is selective, and random noise is added to the cultural transmission of behaviors from teachers to learners, good solutions to problems can evolve by cultural rather than biological evolution.

We describe a model of imitation learning (inspired by Hutchins and Hazelhurst, 1995) according to which the learner learns via backpropagation using the output of the teacher in response to some shared input as its teaching input. Cultural transmission of behaviors from one generation to the next via imitation learning leads to the progressive deterioration of performance across generations. However, if only the best individuals of each generation function as teachers and, furthermore, the teaching input provided by the teacher is modified by noise before it is used by the learner, not only culturally transmitted behaviors do not deteriorate but initially nonexistent behavioral capacities can emerge evolutionarily via pure cultural transmission as they can be shown to emerge via genetic transmission.

1. Introduction

Genetic algorithms are computational models of biological evolution (Holland, 1975; Goldberg, 1989; Mitchell and Forest, 1994). They may be applied to populations of neural networks that control the behavior of artificial organisms living

in some particular environment to evolve genetically transmitted behaviors well adapted to the environment. An initial population of "genotypes" is created randomly. Each genotype maps into the phenotypical neural network that controls the behavior of one particular individual. The reproductive chances of each individual in the population are dependent on the individual's behavior in the environment. Each individual will tend to behave differently from all other individuals because of the random generation of the initial genotypes. Although the individuals of the initial population will behave rather inefficiently, some individuals will behave more efficiently than others and these individuals will be more likely to reproduce. Reproduction can be agamic or sexual. Agamic reproduction implies that a single individual generates one or more "offspring", i.e., copies of its genotype. Sexual reproduction requires that two different individuals mate and generate a new combination of parts of the genotype of the "father" and parts of the genotype of the "mother". In both cases random mutations are added to the offspring genotype. The differential reproduction of individuals based on the quality of their behavior and the constant addition of variability to the genetic pool via sexual recombination and genetic mutations cause an evolutionary increase in the average quality of the behavior exhibited by successive generations of organisms. In other words, a genetic algorithm can be used to create some particular behavior in a population of organisms that initially lacks the behavior (Figure 1). (For a more detailed description of the organisms and of the task, see the next Section.)




Figure 1. Evolutionary increase in the ability to find (and eat) food elements in the environment across 200 generations in a population that reproduces selectively with genetic mutations. Reproduction is agamic. The three curves represent the behavior of the best individual (top curve), average individual (middle curve), and worst individual (bottom curve) of each generation.

In the organisms we have just described the behavior which is exhibited is completely determined by the inherited genotype. In other simulations, the inherited

genotype has phenotypical plasticity, that is, the phenotypical neural network that is constructed under the control of the genetic instructions contained in the genotype is sensitive to the particular environment in which the individual organism happens to live and, therefore, the same inherited genotype can result in different phenotypical networks according to the particular environment in which the genotype maps into the phenotypical network. In this case the behavior is partly genetically transmitted and partly learned in response to the particular environment (cf. Nolfi, Miglino, & Parisi, 1994). In still other simulations the phenotypical network is constructed under complete control of the genotype but then it is changed on the basis of the "experience" of the individual organism in the environment using some learning algorithm such as backpropagation (Cf. Belew, McInerney, & Schraudolph, 1992; Nolfi, Elman, & Parisi, 1994.)

In some organisms, however, a second system of transmission of behaviors is observed besides the genetic transmission system from parents to offspring. In organisms that live socially the other individuals are part of the environment for any particular individual. An individual, therefore, can learn some behavior by interacting with the other individuals. Assuming that the other individuals already possess some particular behavior a newborn individual can learn the behavior from them. This type of transmission of behaviors from one individual to another which is based not on genetic inheritance but on learning from others is called cultural transmission. As we will see, using the backpropagation procedure to model learning from others we can study the cultural transmission of behaviors in a population of networks as we can study genetic transmission.

However, biological evolution is not just genetic transmission. As we have observed, the selective reproduction of the best individuals and the addition of variability through sexual recombination and genetic mutations cause the gradual (or, in some cases, sudden) emergence of some behavioral capacity that was initially absent. Can cultural transmission also result in cultural evolution? Can the transmission of behaviors from one generation to the next via learning from others cause the emergence of some behavioral capacity that is initially absent? In this paper we present a model of cultural evolution that provides an affirmative answer to this question.

2. A model of learning from others

Following Hutchins and Hazelhurst (1995), we adopt the following model of learning from others (imitation learning). Imagine two neural networks with the same network architecture. One network is the "teacher". The teacher possesses a set of connection weights that allow it to execute efficiently some task, i.e., it knows how to map a set of inputs into the appropriate outputs. (We will ask how the teacher has come to possess these weights later on.) The second network is the "learner". The learner has random connection weights and, therefore, it is initially unable to execute the task efficiently. The learner learns the task by imitating the teacher. In any given learning trial both the teacher and the learner are exposed to the same input. Both

respond to this input with some output. At the beginning of learning the two outputs will tend to be rather different. The learner learns by backpropagation. It uses the teacher's output as its teaching input. In each trial the learner modifies its connection weights in such a way that the discrepancy (error) between the teacher's output and its own output in response to the same input tends to decrease. After a certain number of trials the learner will respond to some input in the same way as the teacher. (Actually, not exactly in the same way because in backpropagation learning the error rarely or never goes to zero and, in real life, no one can imitate perfectly his or her model. This is an important point to which we will return in a moment.) The learner will have learned from the teacher. The teacher's behavior has been transmitted to the learner (Figure 2).

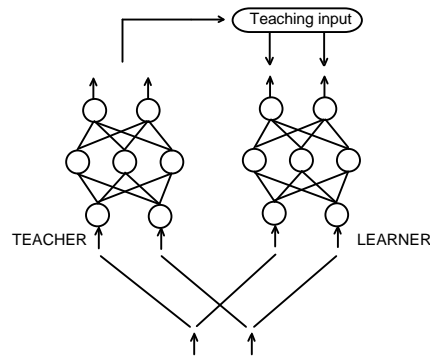


Figure 2. The network on the left is the teacher. The network on the right is the learner. Both teacher and learner are exposed to the same input. The learner learns by using the teacher's output as its teaching input in backpropagation learning.

Imagine a succession of networks with each individual network acting as the teacher for the network that follows it in the sequence. The first network in the sequence learns the Exclusive OR (XOR) task as usual, i.e., by being exposed to the correct teaching input. When the network has learned the task (say, after 10,000 epochs, when the error has reached a very low stable value) its learning is stopped and the network acts as the teacher for the second network. The second network learns the XOR task not by being exposed to the correct teaching input but by using the first network's output as its teaching input. The second network is stopped when its error with respect to the teacher also has reached a stable state. Then the second network acts as the teacher for the third network, and so on.

What is the error of each network at the end of learning if the error is calculated not with respect to the network's teacher but with respect to the correct teaching input? The results are shown in Figure 3.

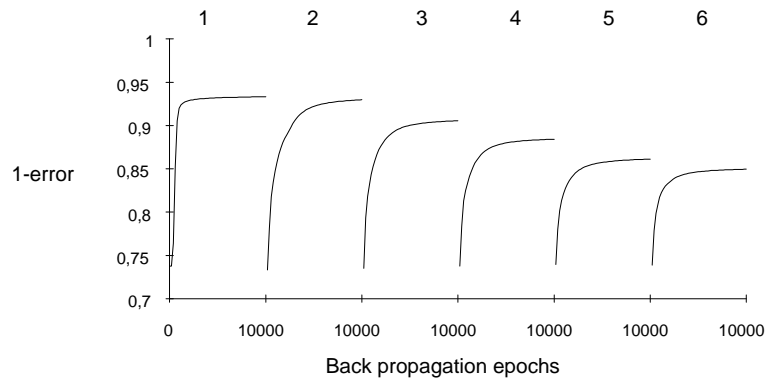


Figure 3. Imitation chain of 6 successive networks. The first network learns the XOR task from the correct teaching input. Each successive network learns the task from the preceding network (teacher) by using the teacher's output as its teaching input. The six curves show the error of each successive network with respect to the correct teaching input. (For presentation purposes learning performance is shown as 1 minus error.)

From figure 3 it is clear that culturally transmitted knowledge dissipates gradually. Since the learning error never goes to zero in backpropagation learning (and presumably in any kind of learning) the behavior of the learner at the end of learning won't be identical to the behavior of the teacher but it will be slightly worse. Therefore, the behavior which is offered by each successive teacher as a model for its "student" to imitate will be slightly worse than the behavior of its teacher and the deterioration of the behavior will accumulate across the successive generations of teachers.

Similar results are obtained if we use not an abstract task such as the XOR task but an ecological task such as finding food in an environment (Parisi, Cecconi, & Nolfi, 1990). The neural network controls the behavior of an organism in an environment that contains food elements. The network's input units encode the location of the nearest food element and the output units encode the movements of the organism. More precisely, the input units encode direction and distance of the currently nearest food element, both mapped in the interval between 0 and 1 and the output units encode angle of turning of the organism and speed of movement forward after turning, also mapped in the interval between 0 and 1. When an organism happens to step on a food element, the food element disappears (it is eaten). An individual's performance is measured as number of food elements found (eaten) during a fixed lifetime. The population is composed by a succession of generations each with a fixed number of 100 individuals. We start from a generation of individuals that are able to eat about 130 food elements during their life on the average. The next generation is composed by 100 new networks with randomly assigned connection weights each of which learns from a different individual of the preceding generation. (Since the individuals of each generation do not inherit anything genetically from the individuals of the preceding generation, these are purely cultural generations.)

During learning both teacher and learner are exposed to the same input (position of nearest food) and they respond by some movement. (Since in ecological networks the input of any given trial is partly determined by the network's motor output in the preceding trial, we imagine that the learner is carried on the teacher's shoulders during the entire period of learning.) The learner learns by backpropagation using the movement encoded in the teacher's output units as its teaching input. At the end of learning the learner steps down from the teacher's shoulders and it is allowed to live an entire life all alone in its own environment.

Figure 4 shows the performance of 6 successive cultural generations of organisms with each generation learning how to find food elements from the individuals of the preceding generation. The average eating ability decreases from about 130 food elements eaten by the first generation to about 10 food elements eaten by the individuals 5 generations later.




Figure 4. Eating performance of 6 successive cultural generations of individuals learning to eat by imitating the individuals of the preceding generation. The performances of the average, best, and worst individual in each generation are shown.

3. Adding selective reproduction and random variability

In the preceding Section we have seen that cultural transmission by imitating others results in the gradual loss of culturally inherited information. Each individual that learns by imitating its teacher won't learn perfectly and, as a consequence, performance deterioration will accumulate generation after generation until any culturally transmitted ability virtually disappears. How can we explain that that is not what appears to be the case in real organisms that have cultural transmission of behaviors? Another limitation of the results reported in the preceding Section is that we start with initial teachers that already possess some behavioral capacity, either the XOR or the eating capacity. Where this behavioral capacity comes from? In real organisms cultural evolution appears to be able not only to maintain some already

existing capacity but also to create some capacity that is initially absent. Can our model capture this creative aspect of cultural evolution?

The model of learning by imitation described in the preceding Section refers only to the mechanism of cultural transmission of behaviors from one individual to another. If we want to study cultural evolution, and not just cultural transmission, we cannot ignore that biological evolution is not just a mechanism for the genetic transmission of behaviors and other phenotypical traits. What we should ask is the following: What is the equivalent in cultural evolution of selective reproduction and the constant addition of variability through sexual recombination and genetic mutations in biological evolution?

The results reported in Figure 4 refer to a succession of cultural generations in which each individual in one generation functions as teacher for one individual of the next generation. In other words, all the individuals that constitute one generation function as teachers for the individuals of the next generation. But let us assume that there is selective reproduction of teachers, i.e., selective cultural reproduction. Only the best individuals in each generation, i.e., the individuals that eat the largest number of food elements, act as teachers for the individuals of the next generation. The culturally nonreproducing individuals have no "students". What would happen in these circumstances? The results for our population of food eaters are shown in Figure 5.

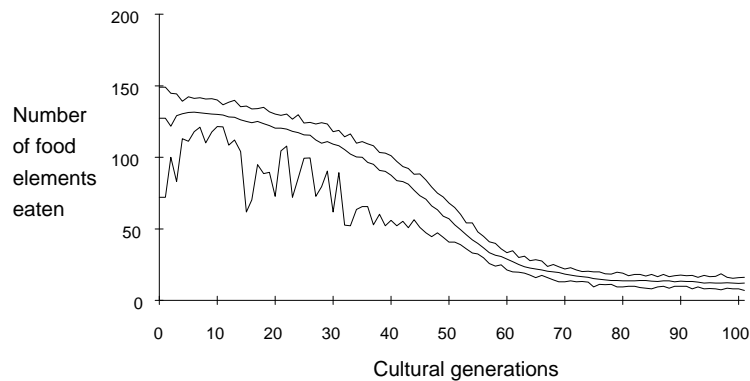


Figure 5. Eating performance of a succession of 100 cultural generations with selective reproduction of teachers. The 20 best individuals in each generation act as teachers for the 100 individuals of the next generation. Each of the 20 teachers is assigned 5 learners, i.e., 5 networks with initial random weights. Average, best, and worst performance in each generation are shown.

It is clear that selective cultural reproduction is able for a while to contrast the cumulative degradation of eating ability in successive generations. Without selective reproduction of teachers eating performance is near zero after 2 or 3 generations (cf. Figure 3). With selective reproduction performance deteriorates much more slowly but after 60 or 70 generations it again approaches zero. Therefore, selective cultural

reproduction may have some role in maintaining culturally transmitted behaviors, but by itself is unable to prevent cultural deterioration.

Let us turn to the second mechanism that appears to be critical in biological evolution: the constant addition of variability to the genetic pool. What could be the equivalent of this mechanism in cultural evolution? We hypothesize that the equivalent of genetic mutations in cultural transmission might be the addition of random noise to the teaching input provided by the teacher to the learner before the teaching input is actually used by the learner to modify its connection weights.

Let us apply this idea to the cultural transmission of eating ability starting from a generation of teachers that already possess the ability to eat. There is selective reproduction of teachers but, in addition, in each learning trial some random noise is added to the teaching input provided by the teacher. More specifically, the activation value of each of the teacher's output units is changed by adding a quantity randomly selected from a normal distribution with a mean of 0 and a range of $-0,5/+0,5$. (The activation value of the teacher's output units can vary from 0 to 1.) The learner uses this modified output as its teaching input. Hence, what the learner is trying to approximate (to imitate) is not its teacher but some random deformation of the teacher. The results are shown in Figure 6.

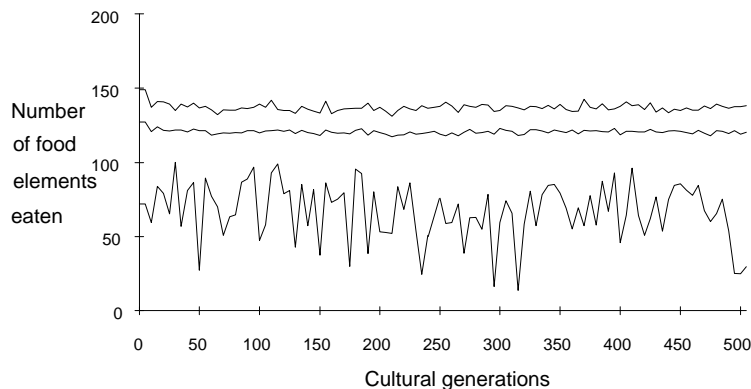


Figure 6. Eating performance of 500 cultural generations starting from a generation of individuals that already know how to eat. Noise is added to the teaching input provided by teachers to learners. The three curves represent the performance of the best, average, and worst individual in each generation.

It seems then that if we adopt selective cultural reproduction and, furthermore, we add noise to the teaching input in imitation learning, it is possible to completely eliminate the progressive deterioration of performance observed in all cases so far in cultural transmission. Without selective reproduction of teachers and without noise in cultural transmission, eating performance quickly deteriorates after a few generations (cf. Figure 4). With selective reproduction of teachers but no noise performance deterioration is much slower but inevitable (cf. Figure 5). With both

selective reproduction and noise added to cultural transmission after 500 generations the average eating ability has the same level as in the first generation (cf. Figure 6).

While selective cultural reproduction and noise added to cultural transmission apparently are sufficient to contrast the cumulative deterioration of performance due to the fact that learners do not perfectly imitate their teachers, we still have to solve the problem of moving from cultural transmission to cultural evolution. Selective reproduction together with noise can maintain some already existing ability but can cultural transmission create a nonexisting ability?

Let us start from an initial population of neural networks with randomly assigned connection weights. The average eating performance of these individuals obviously is extremely limited. At the end of life the 20 individuals with the best eating performance are selected as teachers for the next cultural generation. To each teacher are assigned 5 learners, i.e., 5 networks with initial random weights. Each learner learns by imitating its teacher but some noise is added to the teacher's output before it is used by the learner as its teaching input. At the end of learning (i.e., when the error in the backpropagation procedure reaches a stable low value) the 100 learners become the second generation and their eating performance is measured by allowing them to live in the environment. (Notice that these networks have two "performances": one is learning to respond to the input like the teacher, and the error in the backpropagation procedure is a measure of the goodness of this performance; the other one is finding food in the environment, and the number of food elements eaten is a measure of this second performance.) The process of cultural evolution is continued for 200 generations. The results are shown in Fig. 7.

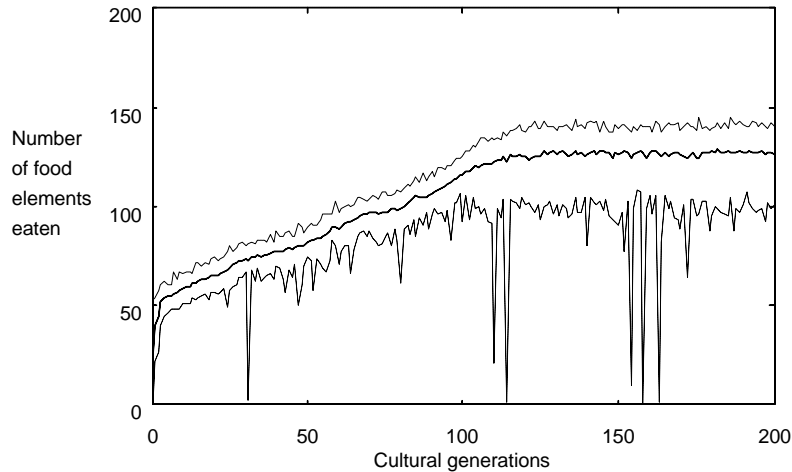


Figure 7. Evolutionary increase in the ability to find food elements in the environment across 200 generations in a population that reproduces culturally by selecting as teachers for the next generation the best individuals in each generation and by adding random noise to cultural transmission. The three curves refer to the performance of the best, average, and worst individual in each generation.

The similarity of Figure 7 with Figure 1 is striking. Our cultural algorithm yields the same results of the genetic algorithm. Cultural transmission of behaviors through learning from teachers replaces the inheritance of genetic information from parents to offspring. Selection of teachers from among the best individuals in each generation replaces selective genetic reproduction. Random noise added to the teaching input provided by teachers replaces genetic mutations. In both cases evolution causes the emergence of a new behavioral ability that was initially absent.

4. Conclusion

Selective cultural reproduction and noise added to cultural transmission are able to maintain an already existing ability by contrasting its progressive deterioration due to imperfect transmission from teachers to learners and they are also able to create an initially nonexisting ability. These two mechanisms of cultural evolution are parallel to the two mechanisms of selective genetic reproduction and genetic mutations in biological evolution. (Sexual recombination can have a cultural analogue in a learner network learning from the output of various teacher networks.) In both cases evolution is based on randomly generating variants of already good solutions and selecting for reproduction the (rare) variants that represent improvements on an already good solution while discarding the (frequent) variants that are less good than the original solution. (Interesting mathematical models of cultural transmission and evolution can be found in Cavalli-Sforza & Feldman, 1981, and Boyd & Richerson, 1985).

It is interesting that evolution will succeed even if the improved variants are very rare. Genetic mutations change randomly some of the inherited genetic information and in rare cases these random changes result in a phenotype that performs better than its parent. However, mutations are one-shot events and it is understandable that in some rare cases they may have lucky phenotypical consequences. On the other hand, the addition of random noise to the teaching input in our model of cultural transmission occurs in all learning trials and one may wonder why the successive random changes in the teaching input do not just neutralize each other without ever resulting in cultural offspring that are better or worse than their teachers.

Figure 8 shows that this is not the case. The figure contains 3 learning curves on the left and 3 curves on the right. The 3 curves on the left refer to a set of 100 networks with initial random weights that learn the XOR task based on the correct teaching input. The 3 curves show the learning performance of the average, best, and worst individual in the set, respectively. The 3 curves on the right refer to another set of 100 individuals that learn using the output of the first set of individuals at the end of learning as their teaching input. Noise is added to this teaching input in each learning trial.

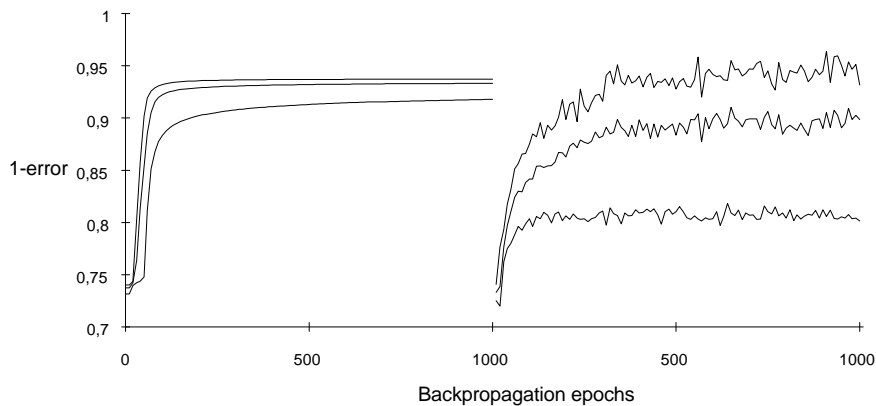


Figure 8. Learning curves for the XOR task for 100 networks learning on the basis of externally provided correct teaching input (left) and for other 100 networks that learn using the output of the first set of networks at the end of learning as their teaching input (right). Noise is added to the teaching input in this second case. The curves show the performance of the average, best, and worst individual in each set.

It is clear that adding noise to the teaching input in cultural transmission increases the variability of the observed performances. Cultural transmission without noise results in a slow and progressive deterioration of performance (cf. Figure 3). Cultural transmission with noise may produce individuals that have a much lower performance level than the worst individual in the preceding generation but at the same time it can produce some rare individuals that have a slightly higher performance level than the best individual of the preceding generation. This appears to be sufficient to trigger cultural evolution. The selective cultural reproduction of the best individuals of each cultural generation will insure that the rare and probably slight improvements produced by noise will be retained in the population of cultural variants while the frequently produced deteriorated variants will be discarded. This process becomes cumulative and it not only is able to neutralize any global tendency to deterioration but can also result in the evolutionary cultural emergence of initially nonexistent capacities.

References

1. Belew, R.K., McInerney, J., & Schraudolph, N.N. Evolving networks: using the genetic algorithm with connectionist learning. In C.G. Langton et al. (eds.) Artificial Life II. Reading, Mass., Addison-Wesley 1992.
2. Boyd, R. & Richerson, P.J. Culture and the Evolutionary Process. Chicago, Chicago Press, 1985.
3. Cavalli-Sforza, L.L. & Feldman, M.W. Cultural Transmission and Evolution. A Quantitative Approach. Princeton, Princeton university Press, 1981.

4. Goldberg, D.E. Genetic Algorithms in Search, Optimization, and Machine Learning. Reading, Mass., Addison-Wesley, 1989.
5. Holland, J.H. Adaptation in Natural and Artificial Systems. Ann Arbor, University of Michigan Press, 1975 (reprinted by MIT Press, Cambridge, Mass., 1992).
6. Hutchins, E. & Hazelhurst, B. How to invent a lexicon: the development of shared symbols in interaction. In N. Gilbert & R. Conte (eds.) Artificial Societies: The Computer Simulation of social life. London, UCL Press, 1995.
7. Mitchell, M. & Forest, S. Genetic Algorithms and Artificial Life. Artificial Life, 1994, 1, 267-290.
8. Nolfi, S., Elman, J.L., & Parisi, D. Learning and evolution in neural networks. Adaptive Behavior, 1994, 3, 5-28.
9. Nolfi, S., Miglino, O. & Parisi, D. Phenotypic plasticity in evolving neural networks. In D.P. Gaussier & J.-D. Nicoud (eds.) From Perception to Action. Los Alamitos, CA, IEEE Computer Society Press, 1994.
10. Parisi, D., Cecconi, F., & Nolfi, S. Econets: networks that learn in an environment. Network, 1990, 1, 149-168.