
**Institute of Psychology
C.N.R. - Rome**

Phenotypic Plasticity in Evolving Neural Networks

Stefano Nolfi⁺ Orazio Miglino* Domenico Parisi⁺

⁺Institute of Psychology, National Research Council, Rome, Italy.

^{*}Department of Psychology, University of Palermo, Italy

e-mail: stefano@kant.irmkant.rm.cnr.it

orazio@caio.irmkant.rm.cnr.it

domenico@gracco.irmkant.rm.cnr.it

May 1994

Technical Report PCIA-94-05

Department of Cognitive Processes and Artificial Intelligence

15, Viale Marx

00137 - Rome - Italy

voice: 0039-6-86894596

fax: 0039-6-824737

To appear In: Proceeding of the First Conference From Perception to Action, 5-9 September,
Lausanne

Phenotypic Plasticity in Evolving Neural Networks

Stefano Nolfi* Orazio Miglino+ Domenico Parisi*

*Institute of Psychology, National Research Council, Rome
Viale Marx, 15, 00137, Rome, Italy

+Department of Psychology, University of Palermo
Viale delle Scienze, Palermo, Italy

E-mail: stefano@kant.irmkant.rm.cnr.it, orazio@caio.irmkant.rm.cnr.it
domenico@gracco.irmkant.rm.cnr.it

Abstract

We present a model based on genetic algorithm and neural networks. The neural networks develop on the basis of an inherited genotype but they show phenotypic plasticity, i.e. they develop in ways that are adapted to the specific environment. The genotype-to-phenotype mapping is not abstractly conceived as taking place in a single instant but is a temporal process that takes a substantial portion of an individual's lifetime to complete and is sensitive to the particular environment in which the individual happens to develop. Furthermore, the respective roles of the genotype and of the environment are not decided a priori but are part of what evolves. We show how such a model is able to evolve control systems for autonomous robots that can adapt to different types of environments.

Keywords

PHENOTYPIC PLASTICITY - GENETIC ALGORITHMS - NEURAL NETWORKS

1 A Model of Phenotypic Plasticity for Developing Neural Networks

All organisms are subject to some form of change during their life. This change is under the control of both the genetic information contained in the organism itself (the DNA present in the nuclei of all the organism's cells) and the information contained in the external environment. Two individuals living in the same environment but hosting different genotypes will tend to undergo different changes and the same will be true for two individuals that have the same genotype but live in different environments.

The changes that occur in the individual must in general be adaptive, that is, they must lead to an increase or at least to a non-decrease of the individual's reproductive chances. Insofar as such changes are under genetic control, evolution guarantees that they will be adaptive by selecting genotypes that translate into an adaptive succession of phenotypic forms. But evolution's role is not limited to taking care of the genetic influences on development. The environmental influences, too, must lead to adaptive changes, and evolution must also guarantee that this is generally true.

Phenotypic plasticity [Scheiner, 1993] is a notion sometimes used to describe the role of evolution in constraining and guiding the environmental influences on the individual. What an individual inherits from its parents is a genotype. The genotype is initially encoded in the DNA contained in the egg cell. Through cell division and differentiation an entire multicellular organism is constructed which is called

the phenotype. One says that there is phenotypic plasticity when the phenotype constructed on the basis of the instructions encoded in the genotype is not completely specified in the genotype but there is an influence of the external environment which is visible in the observed phenotype. Given a particular genotype two or more different phenotypes can result from the genotype as a function of the particular environment in which the mapping happens to take place, i.e. in which the individual is developing.

However, the environmental influences on development are not arbitrary but they are constrained (in terms of timing, nature, effects, etc.) by the genotype and by the succession of developing phenotypic forms. To describe the relation between genes and environment the notion of a reaction norm can be used. A reaction norm is a genetically inherited constraint that specifies the range of possible developmental reactions of the individual to the environmental influences. If a trait of the individual can be described quantitatively on a single dimension, the genetically inherited reaction norm specifies the maximum and minimum values of the trait but it is the particular environment in which the individual happens to develop that determines the actual value. The notion of a reaction norm is perhaps too simple for dealing with the more complex cases of gene/environment interaction but it appears to be useful to indicate in general terms the type of interaction between genes and environment that cause the emergence of a particular phenotype.

In this paper we propose a simulation model of phenotypic plasticity. Organisms are modeled by neural networks. The evolution of populations of such networks is modeled using the genetic algorithm [Holland, 1975]. Individual networks reproduce as a function of their fitness. An 'offspring' network inherits some genetically transmitted information from its parent(s) but this information interacts with information from the external environment to determine the phenotype.

Evolving populations of neural networks have already been used to model the interaction between genetically inherited and environmental information. One example is Todd, Miller, & Hedge's [1989] simulation of the evolution of network architecture. Individual networks inherit their network architecture from their parents. They learn some task during their life using the backpropagation procedure and starting from initial connection weights which are randomly assigned to the inherited architecture at birth. The final performance on the task determines the reproductive chances of each individual (fitness). Appropriate (adaptive) architectures for learning the task evolve at the population level because architectures vary from one individual to another and are genetically transmitted with mutations and crossover. However, the external environment has also a role in determining the final shape of the phenotype (and its fitness) because the individual learns on the basis of the input and teaching input which originate in the external environment. Therefore, the changes in the connection weights that take place during an individual's life are caused by the external environment together with the inherited network architecture. In other simulations what evolves and is genetically transmitted is not the network architecture but the initial weights for learning the task or the learning parameters (learning rate, momentum) [Belew, McInerney, & Schraudolph, 1991]. In all cases the phenotypic network and how it changes during life are dependent on both inherited and environmental information.

In the simulation on phenotypic plasticity that will be described in the present paper we are looking for a more intimate interaction between genes and environment in neural networks. To this purpose we use a more complex model of what is genetically inherited and of how it maps into the phenotypic individual. In the work just referred to the genotype maps directly and linearly, i.e. one-to-one, into the phenotype. For example, the genotype explicitly says which units will be contained in the phenotypic network's architecture and which unit will be connected with which other unit [Miller, Todd and Hegde, 1989]. Furthermore, the mapping of the genotype into the phenotype is viewed as an instantaneous event, not as a temporal process which develops during the course of an individual's life. Therefore, the genotype determines the initial phenotype but has no role in the successive changes that occur in the phenotype. The environment, i.e. the input and teaching input arriving to the network from outside and causing the network to learn, is the sole determinant of the changes that occur in the individual during the individual's life.

In the model to be described the mapping from genotype to phenotype is indirect and nonlinear, i.e. many-to-many. A single property of the genotype may have a role in determining many different traits of the phenotypic network and, conversely, each phenotypic trait may depend on many properties of the genotype. Furthermore, the genotype-to-phenotype mapping is not abstractly conceived as taking place in a single instant but is a temporal process that takes a substantial portion of an individual's lifetime to complete. Hence, during the life of an individual there occur changes that are under the control of the inherited genetic material.

We have already described some simulations using such a model of development for neural networks [Nolfi and Parisi, in press]. However, in that work the environment had no role in the developmental changes that occurred in the individual under genetic control. In the present model the genetic material is plastic, i.e. it is sensitive to the particular environment in which the individual happens to develop. Furthermore, the respective roles of the genetic and the environmental information are not fixed and are not decided a priori by the researcher, as in the literature referred to above. It is not the case, for example, that the network architecture or the initial weights or the learning parameters are genetically inherited while to the environment is assigned the role of changing the connection weights of the neural network using the backpropagation procedure. In the simulations to be presented the boundary between the influence of the genes and the influence of the environment on the developing phenotype is variable and is itself an adaptive trait which is subject to evolution.

An organism is modeled as a neural network living in a physical environment [Parisi, Cecconi, and Nolfi, 1990]. The organism's network receives sensory information about the environment in its input units and it generates motor actions in its output units. The organism belongs to a population of non-identical individuals that reproduce as a function of their fitness, i.e. of how well they behave in the environment (given a particular criterion of fitness). At any given instant each individual contains both a genotype and a phenotype. The genotype is a set of instructions for constructing a network. The phenotype is the neural network of the individual at the current stage of construction. Therefore, for purely genetic reasons an organism changes during its life because its phenotype (neural network) changes. The genotype, of course, remains identical and it is transmitted to the individual's offspring if the individual happens to reproduce. Reproduction is agamic, i.e. a single reproducing individual generates a certain number of copies of its genotype with some random mutations.

The genotype encodes developmental instructions that control the growth of the branching axon of each of a set of neurons. When the axonal branch of neuron X reaches the body of another neuron Y, a connection between the two neurons is established. However, the genetic control of this process of growth is not absolute. The environment has also a role. The respective roles of the genotype and of the environment in controlling neural development are as follows.

During the life of the organism a particular neuron may happen to be stimulated from outside. The algebraic sum of excitations and inhibitions arriving to the neuron is the neuron's net input. This net input is mapped into the neuron's activation state using a logistic function. Given a sequence of N successive activation states of the same neuron one can measure the variability of this distribution of states. A neuron will grow its branching axon only if this measure of activation variability is higher than a certain threshold. This threshold is specified in the genotype but is different for different neurons and it can be changed by the evolutionary process because it is genetically inherited and is subject to mutations. (Notice that this growth threshold of a neuron has nothing to do with the neuron's activation threshold, i.e. the neuron's bias. The growth threshold of a neuron controls when the neuron will grow its branching axon. The activation threshold controls the activation state of the neuron. Both are specified in the genotype in our model. See the next Section.)

Therefore, if and at what point in development a particular neuron will grow its branching axon and possibly establish connections with other neurons depends on both the genotype and the environment. The environment must send sufficiently variable stimulation to the network's input units so that the particular growth threshold of a neuron is reached and the neuron can grow its axon. However, the

genotype specifies what this threshold is and, therefore, it also controls if and when the neuron will grow its axon. (The reaching of the threshold is of course also controlled by the already developed neural structure which connects the input units to the critical neuron. Hence, what happens in a particular developmental stage depends also on previous development.) But the respective roles of genotype and environment are not fixed because the threshold is a genetically inherited trait and therefore it also can evolve. Concretely, the threshold can go to zero, which implies that the particular neuron will grow its axon at birth independently from the environment. The growth process is completely controlled by the genes. Or the threshold can go to its maximum value which also implies that the environment has no influence on the growth process. The neuron will never grow its axon whatever the environment. Intermediate values of the threshold assign variable roles to the environment and to the genetically inherited threshold.

In Section 2 we will describe the model in more detail and we will present some simulations that apply the model to a population of evolving neural networks. In Section 3 we discuss the results of the simulations.

2 The Model and Some Simulations

Our goal is to develop a neural network that controls an organism trying to reach a target area in its environment. The environment is a 60x35 cm arena surrounded by walls made of naturally colored wood. The organism should find as quickly as possible a target area of circular shape with a diameter of 2 cm which is randomly positioned in the arena. The target area may or may not be illuminated by a 1-watt lightbulb positioned right above the area.

We assume that the organism's body is Khepera, a small mobile robot [Mondada, Franzi & Jenne, 1993]. Khepera has a circular shape with a diameter of 55 mm, a height of 30 mm, and a weight of 70g. It is supported by two wheels and two small teflon balls. The wheels are controlled by two DC motors with an incremental encoder (10 pulses per mm of advancement of the robot) and can move in both directions. The robot is provided with 8 infrared/light proximity sensors. They can detect obstacles (at a distance of about 200 mm) and/or the light originating in the target area (at a distance of about 2000 mm). Six sensors are positioned on the front of the robot, the remaining two on the back. Although the physical sensors are only 8, they can be conceived as two separate sets of 8 obstacle detectors and 8 ambient light sensors. (Cf. Figure 1.)




Figure 1. The picture on the left side shows Khepera, the miniature mobile robot. The picture on the right side shows how the 8 sensors are distributed on the robot's body. The large empty circles represent the ambient light sensors and the small full circles represent the obstacle detectors.

A simulator of the robot and of the robot's environment was built [Nolfi, Floreano, Miglino, and Mondada, in press] by sampling the real environment. We let the physical robot turn 360 degrees and we recorded the sensory activations at different distances with respect to a wall and the lightbulb. The activation level of each of the 8 obstacle detectors was recorded for 180 different orientations and 20 different distances. The activation level of each of the 8 ambient light detectors was recorded for 180 different orientations and 26 different distances. In addition, we recorded how and how much Khepera moved and turned for each of the 20x20 possible states of the two motors. This information was used by the simulator to set the activation level of the neural network input units and to compute the displacements of the robot in the simulated environment. The physical shape of Khepera, the environment structure, and the actual position of the robot were accurately reproduced in the simulator with floating point precision.




Figure 2. Developmental instructions specified in an organism's genotype.

Each organism includes both a genotype and a phenotypic neural network. The genotype contains developmental instructions for the construction of the phenotypic network. The genotype has a fixed length and is divided up into blocks, each block corresponding to a single neuron (Figure 2).

Each block contains instructions that determine the following properties of the corresponding neuron:

(a) The "expression threshold gene" determines at what condition a neuron is allowed to grow its axon and, therefore, to possibly establish connections with other neurons. For a neuron to grow its axon the variability in a sequence of 10 successive activation states of the neuron must exceed the threshold value specified by this gene. If the gene codifies a 0 value the corresponding neuron is allowed to grow immediately.

(b) The two "physical position genes" specify the Cartesian coordinates of the neuron in the bidimensional nervous system of the organism.

(c) The "branching angle gene" and the "segment length gene" determine the angle of branching of the neuron's axon and the length of the branching segments. While angle of branching and segment length can vary from one neuron to another in the same network, all the branchings of the axon of a given neuron have the same branching angle and the same length. The growing axon of all neurons branches a fixed number of times (five) with exclusively binary branchings. In order to obtain feedforward networks without recurrent connections, the branching angles are so restricted that all units send their growing axons in the direction from input units (bottom layer) to output units (top layer).

(d) The "synaptic weight gene" determines the synaptic coefficients of all the connections that will be established by the neuron. All connections departing from the same neuron have the same weight.

(e) The "bias gene" represents the activation bias (threshold) of the neuron itself.

(f) The "neuron type gene" specifies, in the case of a sensory neuron, to which of Khepera's 8 sensors the corresponding neuron should be clamped and if it should detect obstacles or the ambient light, and in the case of a motor neuron, if the corresponding motor neuron determines the speed level of the left or the right motor. If more than a single output neuron codifying the same motor information exist, the actual motor output is taken to be the average of the activation levels of the different neurons.

There are three types of neurons, sensory neurons, internal neurons, and motor neurons. Genotypes are 32 blocks in length, i.e. each organism can have a maximum number of 32 neurons in its nervous system. The first 10 blocks in the genotype correspond to sensory neurons, the last 5 blocks to motor neurons, and the 17 intermediate blocks to internal neurons. Internal neurons can be arranged in a maximum of 7 layers. The standard logistic function is used to compute the activation value of the neurons.

The results of the execution of the genotypic instructions of an evolved organism are shown in Figure 3. Neurons of different types and in different parts of the organism's nervous system are created with their axonal branchings (left box). When the growing axonal branch of a particular neuron reaches another neuron a connection between the two neurons is formed. The resulting neural network is shown in the central box. Nonconnecting branchings have been removed. The right box shows the functional part of the same network, i.e. the network after isolated (nonfunctional) neurons and groups of interconnected neurons have also been removed.




Figure 3. Development of the neural network of an evolved organism. The growing and branching process of axons is shown on the left side of the figure. The resulting neural network after removal of nonconnecting branches is shown in the center. The functional network after elimination of isolated (nonfunctional) neurons and groups of interconnected neurons is shown on the right side of the figure.

To evolve organisms which are able to reach the target area efficiently we applied a genetic algorithm to the organisms' genotypes. We begin with 100 randomly generated genotypes each yielding a network with a different architecture and a different assignment of connection weights. This is Generation 0 (G0). G0 networks are allowed to "live" for 10 epochs, with an epoch consisting of 500 actions. At the beginning of each epoch the organism and of the target area are randomly positioned in the arena. At the end of their life the organisms have the possibility to reproduce. However, only the 20 individuals which have accumulated the most fitness in the course of their life reproduce (agamically) by generating 5 copies of their genotype. These $20 \times 5 = 100$ new organisms constitute the next generation (G1). Random mutations are introduced in the copying process resulting in possible changes in the architecture and/or connection weights. The process is repeated for 500 generations.

During the life of an organism the organism's fitness is increased when the organism reaches the target area and the increase is inversely proportional to the number of actions needed to reach the target area. More precisely, the first time an organism reaches the target area during an epoch its fitness level is increased by $(500 - N)$ units where N represents the number of actions spent before reaching the

target area since the beginning of the epoch.

To test if the model is able to evolve organisms that adapt during their life to changing environmental conditions we exposed the organisms of alternate generations to different environments. The organisms of even generations were exposed to environments in which the target area was illuminated by the lightbulb ("light" environment) while the organisms of odd generations were exposed to environments without light ("dark" environment).

If we look at the fitness of the best individuals of each generation we see that the organisms increase their ability to reach the target area very quickly during the first 50-100 generations and then the fitness continues to grow much more slowly until the end. (Cf Figure 4.)

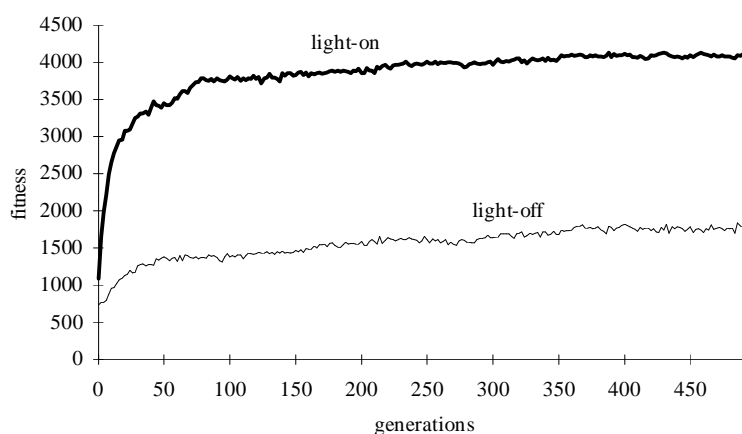


Figure 4. Fitness of the 20 best individuals in each of the 500 generations. The fat curve represents the performance of the even generations (those exposed to the "light" environment), the thin curve represents the performance of the odd generations (those exposed to "dark" environment). Both curves are the average results of 10 simulations starting from initial populations with different randomly generated genotypes.

Figure 4 also shows that the organisms that live in an environment in which the target area is illuminated have a higher fitness than the organisms that live in dark environments. This is due to the fact that organisms that live in lighted environments receive more information from the environment and appear to have an easier task. They explore the arena until they happen to perceive the light source and then they approach the light source and reach the target area. The organisms that are exposed to dark environments are forced to explore the arena and find the target area without the help of sensory information that can inform them about the location of the target area. It is not surprising that organisms developing in "light" environments succeed in reaching the target area more quickly and, therefore, in collecting more fitness than organisms developing in "dark" environments.

To examine the influence of the two different environments on the development of these organisms we let two exact copies (clones) of the same evolved genotype develop and live in the two environments. As Figure 5 and Figure 6 show, the same genotype that develops in different environments produces different phenotypes (i.e. different neural networks), relies on different sensors, and generates different behaviors. This implies that evolution has selected genotypes that exploit their capacity to be influenced by the external environment.




Figure 5. An evolved genotype developing in a "light" environment. The picture on the left side shows the organism's behavior. The environment and the target area are represented as a rectangular box and as a circle, respectively. The picture shows the organism's trajectory from the starting position to the target area. The central picture represents the sensors that have developed in the organism. Large empty circles represent light detectors and small full circles represent wall (obstacle) detectors. Finally, the picture on the right side represents the organism's entire nervous system. The bottom layer circles labeled with progressive numbers represent sensory neurons, the upper layer circles labeled with "L" and "R" represent the left and right motor neurons, and the remaining layers represent internal neurons.




Figure 6. The same genotype of Figure 5 developing in a "dark" environment.

We should now try to determine more precisely how the changing environments affect these organisms. First, we wish to determine if being selected in a particular environment affects an organism's performance in the two environments. In order to test this hypothesis we take the phenotypic (adult) network of the best individual of each generation at the end of its development and we test such frozen network in the two types of environments. We can then compare the performance of an organism selected in environment X and tested in the same type of environment with the performance of an organism selected in environment X and tested in environment Y.

As Figures 7 and 8 show, organisms that have been selected in a type of environment perform better in that type of environment than organisms that have been selected in the alternative environment. When tested in "light" environments organisms which have been selected in "light" environments outperform organisms which have been selected in a "dark" environment. (Cf. the fat and thin curves in Figure 7). Similarly, but less significantly, in "dark" environments organisms which have been selected in "dark" environments outperform organisms which have been selected in a "light" environment (cf. the fat and

thin curves in Figure 8)

We can also determine if the effects of the environment on the development of these organisms are adaptive, i.e. if organisms which happen to develop in an environment with some specific properties acquire phenotypic traits that are useful in that environment and less useful in environments with different properties. In order to test this hypothesis we made two identical copies of the genotype of the best individual of each generation and let each copy develop in one of the two different environments. The two phenotypic networks that were obtained were then frozen and the fitness obtained by each network in the two different environments was computed. As Figure 7 and 8 show, individuals obtain more fitness if the environment in which they develop is of the same type of the environment in which they are tested than if it is different. In "light" environments organisms which have developed in "light" environments outperform organisms which have developed in a "dark" environment (cf. the fat and the dashed curves in Figure 7). In "dark" environments organisms which have developed in "dark" environments outperform organisms that have developed in a "light" environment (cf. the fat and dashed curves in Figure 8).

In other words, if we "artificially" move a *genotype* that has been selected in a environment of type X (cf. the fat curves in Figure 7 and 8) to an environment of type Y (cf. the thin curves in Figure 7 and 8) we observe a certain decrement in performance. However, if we move a *phenotype* that has developed in an environment of type X (cf. the fat curves in Figure 7 and 8) to an environment of type Y (cf. the dashed curves in Figure 7 and 8) we observe a larger decrement in performance. Since in the latter case we deprive the organisms of the possibility to ontogenetically adapt to the environment, we can conclude that development appears to produce phenotypes that are adapted to the environment in which they develop.

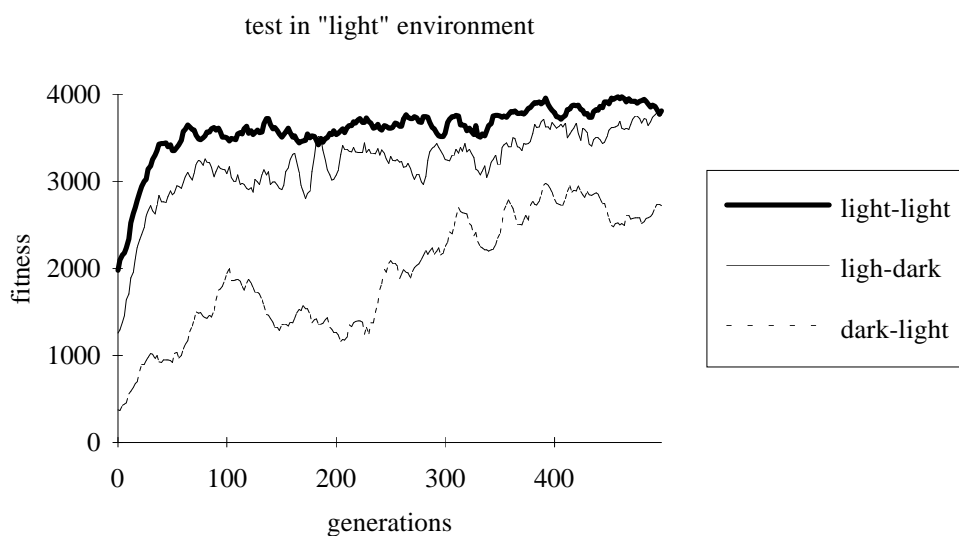


Figure 7. Fitness of the best organisms of each generation tested in a "light" environment for 10 epochs after completion of development. The fat curve represents the performance of organisms that have been selected in a "light" environment (i.e. organisms of even generations) and have developed in the same type of environment (light-light). The thin curve represents the performance of organisms that have been selected in a "dark" environment (i.e. organisms of odd generations) and have developed in a "light" environment (dark-light). The dashed curve represents the performance of organisms that have been selected in a "light" environment (i.e. organisms of even generations) and have developed in a "dark" environment (light-dark). Each curve represents the 7 point moving average of 10 different simulations.




Figure 8. Fitness of the best organisms of each generation tested in a "dark" environment for 10 epochs after completion of development. The fat curve represents the performance of the organisms that have been selected in a "dark" environment and have developed in the same type of environment (dark-dark). The thin curve represents the performance of organisms that have been selected in a "light" environment and have developed in a "dark" environment (light-dark). The dashed curve represents the performance of organisms that have been selected in a "dark" environment and have developed in a "light" environment (dark-light). Each curve represents the 7 point moving average of 10 different simulations.

3 Discussion

To apply genetic algorithms to neural networks appears as the most natural way to study developmental plasticity and evolution. Genetic algorithms are models of evolution and neural networks are models of learning. However, in many simulations based on genetic algorithms and neural networks the evolutionary process and the learning process do not appear to be integrated as in natural organisms. In these models the genotype happens to act only once to determine the network's initial weights [Montana, 1989; Parisi, Cecconi, and Nolfi, 1990; Belew, McInerney, and Schraudolph, 1991] or the network's architecture [Miller, Todd, and Hedge, 1989; Belew, 1993; Kitano, 1990]. The role of the genotype is confined to determining the initial state and then learning operates to modify this initial state. In biological organisms things appear to be different. While the genotype in the germinal cell acts once, the genotype of the somatic cells is continuously active during all of the life of the organism and such activity is controlled by the regulatory genes. Learning and all the processes that occur during an individual's lifetime (for example, cell differentiation, growth of new connections within the nervous system, changes in synaptic conductivity) are driven by the genotype of the somatic cells because all these processes require the synthesis of new molecules. In biological organisms learning (i.e. the effects of the environment on the phenotype) and development (i.e. the effects of the genotype on the phenotype) appear to be very intimately related. In addition, while the main role of individual plasticity is the possibility to adapt to changing environments, the possibility to adapt to changing environments has not been investigated, as far as we know, with this type of models.

We presented a model in which the mapping between genotype and phenotype (i.e. ontogeny) takes place during the individual's lifetime and is influenced both by the genotype and by the external environment. Development (i.e. genetically driven changes) and learning (i.e. environmentally driven changes) operate on the same basic process, the growing and branching process of neuronal axons, with similar effects. Furthermore, the respective roles of the genetic and the environmental information are not fixed and are not decided a priori by the researcher. Evolution is free to transfer the complete control on the development of a particular neural structure to the genes or to assign a more or less important role to the environment in this development. We showed how such model is able to evolve neural networks that adapt during lifetime to different environments.

The main objective of the simulations reported in the paper was to show that simulations with genetic algorithms and neural networks can be used as tools to explore many more aspects of biological phenomena than has been the case so far. Not only plastic neural networks of the type we have described can be used to study epigenesis during development but one is free to select, develop, and test phenotypes in different environments as one can sometimes do with real plants but more rarely with real animals. [But cf. Taylor, 1987.] Plasticity in development allows the evolution of various classes of genotypes: (a) genotypes that are environmentally sensitive, i.e. that develop into different phenotypes as a function of the particular environment vs genotypes that yield more or less the same phenotype in different environments; (b) genotypes that are specialized for specific environments, i.e. that yield phenotypes much more effective in one environment rather than in another environment, vs genotypes that yield phenotypes that although different are equally effective in different environments. Future research should be dedicated to examining in what conditions these different classes of genotypes tend to evolve and how the distribution of the genotypes into the various classes may change across evolutionary time.

The model presented may have some usefulness not only as a tool for investigating how evolution and developmental plasticity interact in organisms but also as a tool for developing control system for autonomous robots. Recently many researchers have suggested the use of evolutionary techniques to evolve autonomous robots [Brooks, 1992; Cliff, Husband, and Harvey, 1993, Nolfi, Floreano, Miglino, and Mondada, 1994]. Unfortunately, the large population size and the number of generations required for the evolutionary emergence of interesting behaviors imply that a large number of robots must be evaluated. The problem can be reduced if the control system of the robot can be evolved in computer simulations and then transferred to the real robot [Miglino, Nafasi, and Taylor, 1994]. On the other hand, since it is extremely difficult to capture all the critical physical characteristics of the robot and of the environment in an abstract model, the behavior of the control system evolved in a computer may differ significantly once the control system is transferred to the real robot. The possibility to evolve plastic control systems can represent a method for solving this problem. The simulated environment and the real environment can be viewed as just two different types of environments. As a consequence, plastic control systems evolved in a simulated environment may be able to adapt to the real environment once they are transferred to the real environment.

References

- Belew, R. K., J. McInerney, and N. N. Schraudolph. 1991. Evolving networks: Using the Genetic Algorithm with connectionist learning. In: *Proceedings Second Artificial Life Conference*, eds. C.G. Langton, J. D. Farmer, S. Rasmussen, C. E. Taylor. New York, Addison-Wesley.
- Belew, R. K. 1993. Interposing an ontogenetic model between Genetic Algorithms and Neural Networks. In: *Advances in Neural Information Processing (NIPS5)*, ed. J. Cowan. San Mateo, CA, Morgan Kaufmann.

- Brooks, R. A. 1992. Artificial life and real robots. In: *Toward a Practice of Autonomous Systems: Proceedings of the First European Conference on Artificial Life*, eds. F. J. Varela, and P. Bourguine. Cambridge, MA, MIT Press/Bradford Books.
- Cliff, D. T., I. Harvey, and P. Husbands. 1993. Explorations in Evolutionary Robotics. *Adaptive Behavior*, 2: 73-110.
- Holland, J. J. 1975. *Adaptation in natural and artificial systems*. Ann Arbor, MI, University of Michigan Press.
- Kitano, H. 1990. Designing neural networks using genetic algorithms with graph generation system. *Complex Systems*, 4:461-476.
- Miglino O., K. Nafasi, and C. Taylor. 1994. Selection for Wandering Behavior in a Small Robot. *Technical Report*. UCLA-CRSP-94-01. Department of Cognitive Science, University of California, Los Angeles.
- Miller, G. F., P. M. Todd, and S. U. Hedge. 1989. Designing neural networks using genetic algorithms. In: *Proceedings Third International Conference on Genetic Algorithms*. San Mateo, CA, Morgan Kaufmann.
- Mondada F., E. Franzi, and P. Ienne. 1993. Mobile Robot miniaturisation: A tool for investigation in control algorithms. In: *Proceedings of the Third International Symposium on Experimental Robotics*. Kyoto, Japan.
- Montana, D.J., and L. Davis. 1989. Training feedforward neural networks using genetic algorithms. In: *Proceedings Eleventh International Conference on Artificial Intelligence*, ed. N. S. Sridharan. San Mateo, CA, Morgan Kaufmann.
- Nolfi, S., D. Floreano, O. Miglino, and F. Mondada. in press. How to evolve autonomous robots: different approaches in evolutionary robotics. *Proceedings Fourth Artificial Life Conference*, Cambridge MA.
- Nolfi, S., and D. Parisi. in press. "Genotypes" for Neural Networks. In: *The Handbook of Brain Theory and Neural Networks*, ed. M. A. Arbib. Cambridge, MA, MIT Press/Bradford Books.
- Parisi, D., F. Cecconi, and S. Nolfi. 1990. Econets: Neural networks that learn in an environment. *Network*, 1:149-168.
- Scheiner, S.M. 1993. Genetics and evolution of phenotypic plasticity. *Annual Review of Ecological Systems*, 24:35-68.
- Taylor, C.E. 1987. Habitat selection within species of *Drosophila*: A review of experimental findings. *Evolutionary Ecology*, 1:389-400.