in press In Bentley & Kumar (eds).

# Artificial Life Models of Neural Development

Angelo Cangelosi[1] Stefano Nolfi[2] Domenico Parisi[2]

[1]University of Plymouth
Institute of Neuroscience and School of Computing
Drake Circus, PL4 8AA, Plymouth UK

[2]Institute of Cognitive Science and Technologies, CNR
15 Viale Marx, 00137 Rome, Italy

## 1. Introduction

Artificial neural networks are computational models of nervous systems. Natural organisms, however, do not possess only nervous systems but also genetic information stored in the nucleus of their cells (genotype). The nervous system is part of the phenotype which is derived from the genotype through a process called development. The information specified in the genotype determines aspects of the nervous system which are expressed as innate behavioral tendencies and predispositions to learn. When neural networks are viewed in the broader biological context of Artificial Life, they tend to be accompanied by genotypes and to become members of evolving populations of networks in which genotypes are inherited from parents to offspring (Parisi, 1997; 2001).

Artificial neural networks can be evolved by using evolutionary algorithms (Holland, 1975; Schwefel, 1995; Koza, 1992). An initial population of different artificial genotypes, each encoding the free parameters of an individual neural network (e.g., the connection strengths and/or the architecture of the network and/or the learning rules), are created randomly. Each individual network is evaluated in order to determine its performance in some task (fitness). The fittest networks are allowed to reproduce (sexually or nonsexually) by generating copies of their genotypes with the addition of changes introduced by some genetic operator (e.g., mutations, crossover, duplication). This process is repeated for a number of generations until a network that satisfies the performance criterion (fitness function) set by the experimenter is obtained (for a review of methodological issue see Yao, 1993).

### 1.1 Evolution and Development

A cornerstone of biology is the distinction between inherited genetic code (genotype) and the corresponding organism (phenotype). What is inherited from the parents is the genotype. The phenotype is the complete individual that is formed according to the instructions specified in the genotype.

In simulations with evolving neural networks, the genotype might encode all the free parameters of the corresponding artificial neural network or only the initial value of the parameters and/or other parameters that affects learning. In the former case the network is entirely innate and there is no learning. In the latter networks change both philogenetically across a succession of generations and ontogenetically during the life of the individual, i.e., during the period of time in which they are evaluated.

Evolution is critically dependent on the distinction between genotype and phenotype, and on their relation, i.e., the genotype-to-phenotype mapping. The fitness of an individual, that affects selective reproduction, is based on the phenotype but what is inherited is the genotype, not the phenotype. Furthermore, while the genotype of an individual is one single entity, the organism can be considered as a succession of different phenotypes taking form during the genotype-to-phenotype mapping process, each derived from the previous one under genetic and environmental influences.

When the genotype-to-phenotype mapping process takes place gradually during an individual's lifetime we can talk of development. In this case, each successive phenotype, corresponding to a given stage of development, has a distinct fitness. The total fitness of a developing individual is a complex function of these developmental phases. Evolution must ensure that all these successive forms are viable and, at the same time, that they make a well-formed sequence where each form leads to the next one until a more or less stable (adult) form is reached. This puts various constraints on evolution but it also offers new means for exploring novelty. Small changes in the developmental rates of different components of the phenotype, for example, can have huge effects on the resulting phenotype. Indeed it has been hypothesized that in natural evolution changes affecting regulatory genes that control rate of development have played a more important role than other forms of change such as point mutations (Gould, 1977).

Although the role of the genotype-to-phenotype mapping and of development has been ignored in most of the experiments involving artificial evolution, there is now an increasing awareness of its importance. Wagner & Altenberg (1996) write: "In evolutionary computer science it was found that the Darwinian process of mutation, recombination and selection is not universally effective in improving complex systems like computer programs or chip designs. For adaptation to occur, these systems must possess *evolvability*, i.e. the ability of random variations to sometimes produce improvement. It was found that evolvability critically depends on the way genetic variation maps onto phenotypic variation, an issue known as the representation problem." (p. 967).

## 1.2 Artificial life Approaches to Modeling Neural Development

In the next sections, different approaches to modeling neural development in artificial life simulations will be presented. They range from simple direct genotype-phenotype encoding to more complex methods such as axonal growth, cellular encoding, and regulatory models. Furthermore, we discuss some models of the interaction between evolution and learning. These models address a different type of plasticity in neural network development, that is, the effects of ontogenetic learning in the overall evolutionary process. (For a review of neural network models of development, see Parisi, 1996, and Parisi and Nolfi, 2001.)

## 2. Genetic Encoding: Direct Genotype-Phenotype Mapping

To evolve neural networks one decision that has to be taken is how to encode the network in the genotype in a manner suitable for the application of genetic operators. In most cases, all phenotypical characteristics are coded in an uniform manner so that the description of an individual at the level of the genotype assumes the form of a string of identical elements (such as binary or floating point numbers). The transformation of the genotype into the phenotypical network is called genotype-to-phenotype mapping.

In direct encoding schemes there is a one-to-one correspondence between genes and the phenotypical characters that are subjected to the evolutionary process (e.g. Miller et al., 1989). Aside from being biological implausible, simple one-to-one mappings have several drawbacks. One problem, for example, is scalability. Since the length of the genotype is proportional to the complexity of the corresponding phenotype, the space to be searched by the evolutionary process increases exponentially with the size of the network (Kitano, 1990).

Another problem of direct encoding schemes is the impossibility to encode repeated structures (such as network composed of several sub-networks with similar local connectivity) in a compact way. In one-to-one mappings, in fact, elements that are repeated at the level of the phenotype must be repeated at the level of the genotype as well. This does not only affect the length of the genotype and the corresponding search space, but also the evolvability of individuals. A full genetic specification of a phenotype with repeated structures, in fact, implies that adaptive changes affecting repeated structures should be independently rediscovered through changes introduced by the genetic operators.

## 3. Growing methods

The genotype-to-phenotype process in nature is not only an abstract mapping of information from genotype to phenotype but it is also a process of physical growth (growth in size and in physical structure). By taking inspiration from biology, therefore, one can decide to encode growing instructions in the genotype.

The phenotype is progressively built by executing the inherited growing instructions.

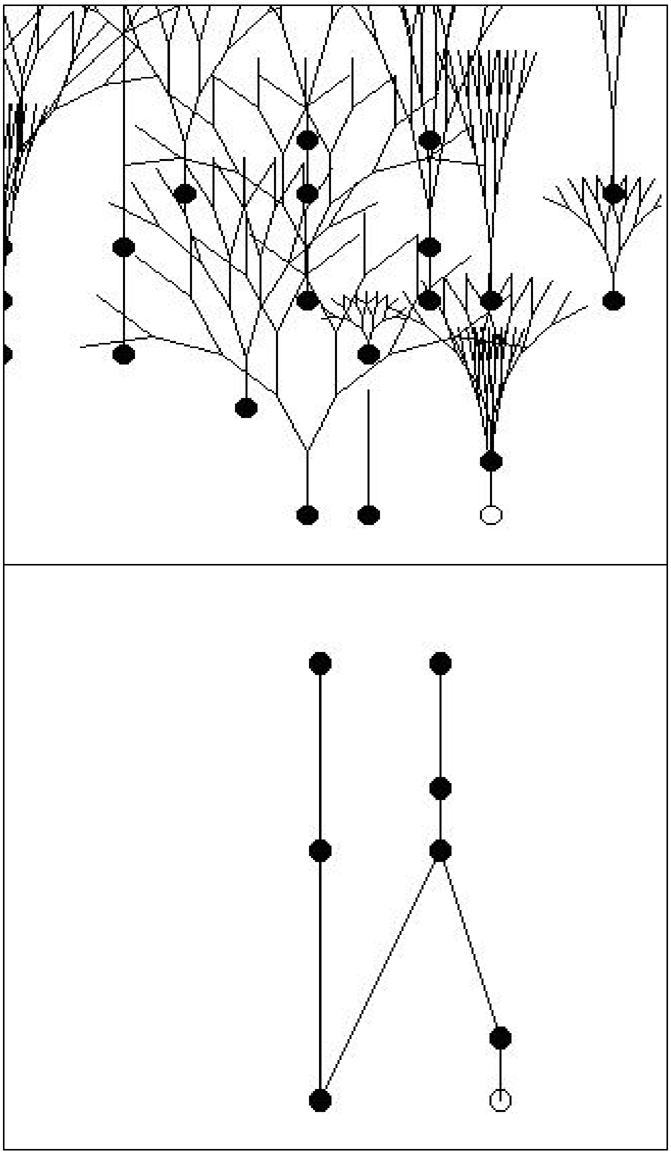

Figure 1. Development of an evolved neural network. Top: The growing and branching process of the axons. Bottom: the resulting neural network after

removal of nonconnecting branches and the elimination of isolated neurons and groups of interconnected neurons.

Nolfi, Miglino and Parisi (1994) used a growing encoding scheme (see also Nolfi and Parisi, 1995) to evolve the architecture and the connection strengths of neural networks that controlled a small mobile robot (for a similar method see Husband et al., 1994). These controllers are composed of a collection of artificial neurons distributed over a 2-dimensional space with growing and branching axons (Figure 1, top). Inherited genetic material specifies instructions that control the axonal growth and the branching process of neurons. During the growth process, when a growing axonal branch of a particular neuron reaches another neuron a connection between the two neurons is established. On the bottom of Figure 1 you can see the network resulting from the growth process displayed in the top of the Figure after the elimination of nonconnecting branches and isolated and non-functional neurons. However, axons grow and branch only if the activation variability of the corresponding neurons is larger than a genetically-specified threshold. This simple mechanism is based on the idea that sensory information coming from the environment has a critical role in the maturation of the connectivity of the biological nervous system and, more specifically, that the maturation process is sensitive to the activity of single neurons (see Purves, 1994). Since the actual sequence of sensory states experienced by the network influences the process of neural growth, in this model the developmental process is influenced not only by genetic factors but also by environmental factors.

This type of genotype-to-phenotype mapping allows the evolutionary process to select neural network topologies that are better suited to the task chosen. Moreover, by being sensitive to environmental conditions, the developmental process might display a form of plasticity. Indeed, as shown by the authors, if some aspects of the task are allowed to vary during the evolutionary process, evolved genotypes display an ability to develop into different final phenotypical structures that are adapted to the current conditions.

## 4. Cellular Encodings

In natural organisms the development of the nervous system begins with a folding in of the ectodermic tissue which forms the neural crest. This structure gives origin to the mature nervous system in a succession of three phases: the genesis and proliferation of different classes of neurons by cellular duplication and differentiation, the migration of neurons toward their final destination, and the growth of neurites (axons, dendrites). The growth process described in the previous section characterizes very roughly only the last of these three phases. A number of attempts have been made to include other aspects of neural development in artificial evolutionary experiments.

Cangelosi et al. (1994), for example, extended the model described in the previous section by adding a cell division and migration stage to the already existing stage of axonal growth. The genotype, in this case, is a collection of rules governing the process of cell division (a single cell is replaced by two "daughter" cells) and migration (the new cells can move in 2D space). The genotype-to-phenotype process starts with a single cell which, by undergoing a number of duplication and migration processes, produces a collection of neurons arranged in a 2D space. At the end of this stage the neurons grow their axons and establish connections until a viable neural controller is formed (for a related approach, see Dellaert and Beer, 1994).

Gruau (1994) proposed a genetic encoding scheme for neural networks based on a cellular duplication and differentiation process. The genotype-to-phenotype mapping starts with a single cell that undergoes a number of duplication and transformation processes ending up in a complete neural network. In this scheme the genotype is a collection of rules governing the process of cell divisions (a single cell is replaced by two "daughter" cells) and transformations (new connections can be added and the strengths of the connections departing from a cell can be modified). In this model, therefore, connection links are established during the cellular duplication process.

In Gruau's model the instructions contained in the genotype are represented as a binary-tree structure as in genetic programming (Koza, 1992). During the genotype-to-phenotype mapping process, the genotype tree is scanned starting from the top node of the tree and then following each ramification. The top node represents the initial cell that, by undergoing a set of duplication processes, produces the final neural network. Each node of the genotype tree encodes the operations that should be applied to the corresponding cell and the two sub-trees of a node specify the operations that should be applied to the two daughter cells. The neural network is progressively built by following the tree and applying the corresponding duplication instructions. Terminal nodes of the tree (i.e., nodes that do not have sub-trees) represents terminal cells that will not undergo further duplications. Gruau also considered the case of genotypes formed by many trees where the terminal nodes of a tree may point to other trees. This mechanism allows the genotype-to-phenotype process to produce repeated phenotypical structures (e.g., repeated neural sub-networks) by re-using the same genetic information. Trees that are pointed to more than once, in fact, will be executed more times. This encoding method has two advantages: (a) compact genotypes can produce complex phenotypical networks, and (b) evolution may exploit phenotypes where repeated sub-structures are encoded in a single part of the genotype. Since the identification of sub-structures that are read more than once is an emergent result of the evolutionary process, Gruau defines this method Automatic Definition of Neural Subnetworks (ADNS) (Gruau, 1994).

## 5. Heterochrony in Neural Development

The existence of a variable and plastic ontogenetic development is strictly related to the evolution of regulatory genotypes, i.e., genotypes whose main role is to control the functioning of simple ontogenetic events. Even though some genes directly encode structural molecules, most genetic products consist of regulatory elements such as enzymes. These regulatory genes act as ON-OFF switches on the complex chain of biochemical events that constitute the three main phenomena of cellular development: mitoses, cell differentiation, and migration. A regulatory ontogenetic development consists of a variety of interactions between the growing organism and its environment.

In such a regulatory development, the timing of the events, i.e., their temporal activation/inhibition, and their rate, i.e., the frequency of occurrence of the phenomena, both have a strong impact. The temporal co-occurrence of two or more events can prove essential for allowing the activation of a biological phenomenon. Even the spatial relation between sub-structures of the developing organism is a key factor. The spatial interaction between cells can induce the phenomena of cell differentiation or cell migration. These classes of interactions, especially the temporal relations occurring during the organism's development, constitute the phenomenon known as heterochronic change. Heterochrony (McKinney and McNamara, 1991) is the study of the effect of changes in timing and rate of the ontogenetic development in an evolutionary context. In particular, heterochonic classifications are based on the comparison of ontogenies that differ in terms of (1) onset of growth, (2) offset of growth, and (3) rate of growth of an organ or other biological traits. These three kinds of change correspond respectively to the following couples of heterochonic phenomena: Predisplacement and Postdisplacement for an anticipated and postponed growth onset, Hypermorphosis and Progenesis for a late and early offset, and Acceleration and Neoteny for a faster and slower rate of growth (see also Gould, 1977).

Cangelosi and Elman (1995; Cangelosi 1999) have developed a model of development that simultaneously simulates many biologically-inspired phenomena for the development of neural networks in artificial organisms. They use a regulatory genotype in which most of the genes produce elements whose role is to control the activation, inhibition, and delay of the developmental events. The phenomena occurring during neural network development (cell duplication, differentiation, migration, axonal growth and synaptogenesis) are directly inspired by their real biological functioning (Purves & Lichtheim, 1985).

At the beginning of neural development the organism's neural system consists of a single egg cell with its own genome and a set of elements present in the intercellular environment. Some of these elements act as “receptors” for extracellular signaling. Others are “structural” elements for the activation and execution of developmental events. Others are pure “regulatory” elements for the modulation of gene expression, and do not play any direct role in development. The structural elements can regulate gene expression while the receptors cannot.

The physical environment in which the egg cell will grow consists of a 2D grid of 7*20 cells. The grid has a polarized orientation in the y dimension. The upper pole corresponds to the organism’s muscle tissue side, and the lower pole to the sensory tissue side. The initial intracellular elements are considered to be inherited from the parent organism. Their distribution, i.e., the initial amount of each element, will function as the zygote's pattern formation mechanism. During development, the amount of these elements, together with the other environmental conditions, will determine the activation, inhibition or delay of developmental events. Moreover, these elements act also as regulators of gene expression.

Five developmental events cyclically occur during the neural network’s growth: Cell duplication, Cell differentiation, Cell migration, Axonal growth, and Synaptogenesis. For example, the cell duplication process consists in the replacement of the mother cell with two new daughter cells. The physical displacement of the new cells, and their differentiation (i.e., the splitting of the mother cell's elements), is determined by the environment available around the mother cell and by the amount of the two elements responsible for mitosis. In this case, the choice of the two elements is inspired by the role played by cyclin and the kinase enzyme, two mayor regulatory proteins for mitosis (Marx, 1989).

The feasibility of regulatory development for adaptation to environmental changes is the hypothesis tested using this model. A two-stage, two-task simulation setting was used. In the first evolutionary stage, the organisms are selected according to their performance in a foraging task. In the second stage, dangerous elements are introduced in the environment, together with food. This task requires that organisms adapt their food approaching strategy to a new behavioral pattern for approaching only foods and avoiding dangers. To do this, organisms need to restructure their neural network, for example by adding or readapting some sensory and hidden neurons to the new processing needs. The way to re-adapt the neural network is by modifying its architecture.

Analysis of the distribution of neurogenetic changes that allow organisms to successfully re-adapt for the food and danger task shows that all five developmental events, except migration, are involved in re-adapting ontogenesis for coping with new behavioral requests. The events related to axon growth and synaptogenesis are the ones most frequently used as a re-adaptation strategy. Even small adjustments of the connectivity pattern can prove very functional for the evolution of good networks.

In the simulation different examples of heterochronic changes were observed. For example, a case of adaptive local Progenesis and contemporaneous cell-cell induction effect due to spatial interaction was observed. Few mutations in an offspring were enough to cause significant changes in its neural development and to allow the organism to adapt to the new environment. Figure 2 shows the morphogenetic tree of the ancestor and descendant organisms where such changes happened. The morphogenetic tree is a graphic representation of a cell duplication tree using the two dimensions of time and space. It facilitates the understanding of the developmental events, and their temporal and spatial interactions (Arthur, 1984). In the ancestor organism, two sensory cells for food location input originate from a common founder cell. In the descendant individual, this cell stops duplicating early, leaving two free spots in the sensory area of the developing grid. This is a case of local Progenesis, because the offset of the mitotic sequence is anticipated. At the same time, there is a change in the cell displacement of other cell duplication branches. In the upper side of the developmental grid, two new cells, coming from a different mitosis branch, occupy the space left free in the sensory area. What happens in the descendant is that in later stages of development a newly formed cell changes position moving to the lower row. This new displacement will induce a dividing cell to place one of its daughters in the lower input area. Because of these spatial interactions, the progeny of this cell ends up in the two spots left free by the Progenesis.

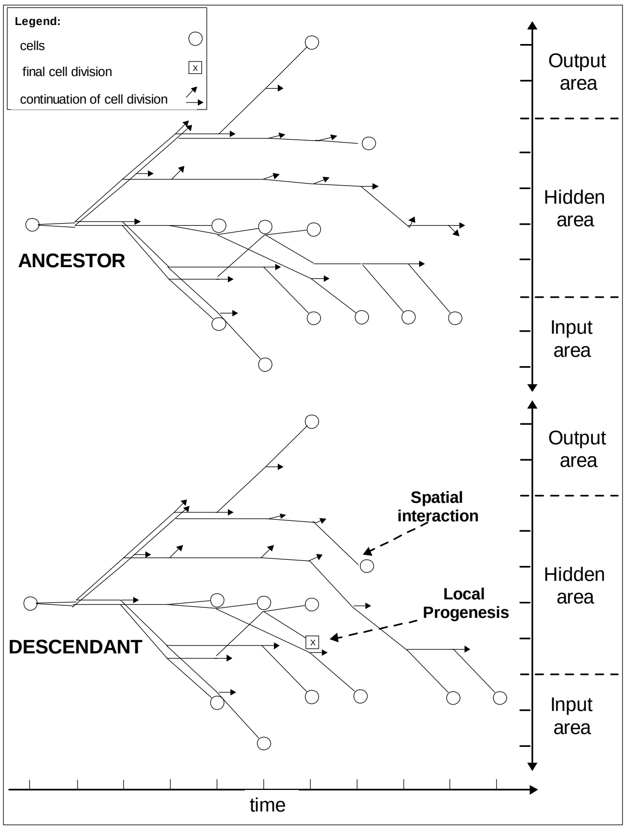

Figure 2: Morphogenetic trees that show heterochronic changes in the development of a neural network (see text for explanation). From Cangelosi 1999.

## 6. Evolution and Learning

Evolution and learning are two forms of adaptation that operate on different time scales. Evolution is capable of capturing relatively slow environmental changes that might encompass several generations. Learning, instead, allows an individual to adapt to environmental changes that are unpredictable at the generational level. Moreover, while evolution operates on the genotype, learning affects the phenotype and phenotypic changes cannot directly modify the genotype. Recently, the study of artificial neural networks that are subjected to both an evolutionary and a lifetime learning process has received an increasing attention. These studies (see also Nolfi and Floreano, 1999) have been conducted with two different purposes: (a) looking at the advantages, in terms of performance, of combining two different adaptation techniques; (b) understanding the role of the interaction between learning and evolution in natural organisms. The general picture that emerges from this body of research is that, within an evolutionary perspective, learning has several different adaptive functions:

- It can help and guide evolution by channelling evolutionary search toward promising directions. For example, learning might significantly speed up the evolutionary search.
- It can supplement evolution by allowing individuals to adapt to environmental changes that, by occurring during the lifetime of the individual or within few generations, cannot be tracked by evolution.
- It can allow evolution to find more effective solutions and increase the ability to scale up to problems that involve large search space.

However, learning also has costs and, in particular, it might increase the unreliability of evolved individuals (Mayley, 1997). Since an individual's abilities are determined by the individual's learning experiences, learning individuals might fail to acquire the required abilities in unfavorable conditions.

## 6.1 How learning might help and 'guide' evolution

A simple and clear demonstration of how learning might influence evolution even if the characteristics that are learned are not communicated to the genotype was provided by Hinton and Nowlan (1987). The authors considered a simple case in which (a) the genotype of the evolving individuals consists of 20 genes that encode the architecture of the corresponding neural networks, and (b) only one architecture, i.e., only a single combination of gene values, confers added reproductive fitness. Individuals have a genotype with 20 genes that can assume two alternative values (0 or 1). The only combination of genes that provide a fitness value above 0 consists of all ones. In this extreme case, the probability of finding the good combination of genes would be very small given that the fitness surface looks like a flat area with a single spike in correspondence of the good combination. The fitness surface is a metaphor often used to visualize the search space on an evolutionary algorithm. Any point on the search space corresponds to one of the possible combinations of genetic traits and the height of each point on the fitness surface corresponds to the fitness of the individual with the corresponding genetic traits. In the fitness surface of Hinton and Nowlan's model, artificial evolution does not perform better than random search. Finding the right combination is like looking for a needle in a haystack..

The addition of learning simplifies evolutionary search significantly. One simple way to introduce learning is to assume that, in learning individuals, genes can have three alternative values [0, 1, and ?] where question marks indicate modifiable genes whose value is randomly selected within [0, 1] at each time step during an individual's lifetime. By comparing learning and non-learning individuals one can see that performance increases throughout generations much faster in the former than in the latter. The addition of learning, in fact, produces an enlargement and a smoothing of the fitness surface area around the good combination that, in this case, can be discovered much more easily by the genetic algorithm. This is due to the fact that not only the right combination of alleles but also combinations which in part have the right alleles and in part have unspecified (learnable) alleles might report an average fitness greater than 0 (fitness monotonically increases with the number of fixed right values because the time needed to find the right combination is inversely proportional, on the average, to the number of learnable alleles). As claimed by the authors, “it is like searching for a needle in a haystack when someone tells you when you are getting close” (Hinton and Nowlan, 1987, p. 496). (For a variation of this model that has been used to study the interaction between evolution, learning, and culture, see Hutchins and Hazlehurst (1991).)

The Hinton-Nowlan model is an extremely simplified case that can be analyzed easily but that makes several unrealistic assumptions: (1) there is no distinction between genotype and phenotype, (2) learning is modeled as a random process that does not have any directionality, and (3) there is no distinction between the learning task (i.e., the learning function that individuals try to maximize during their lifetime) and the evolutionary task (i.e., the selection criterion that identify the individuals that are allowed to reproduce). Further research conducted by Nolfi, Elman and Parisi (1994) showed how, when these limitations are released, learning and evolution display other forms of interactions that are also mutually beneficial.

Nolfi, Elman and Parisi (1994) studied the case of artificial neural networks that 'live' in a grid world containing food elements. Networks evolve (to become fitter at one task) at the population level and learn (a different task) at the individual level. In particular, individuals are selected on the basis of the number of food elements that they are able to collect (evolutionary task) and try to predict the sensory consequences of their motor actions during their lifetime (learning task).

The genotype of the evolving individuals encode the initial weights of a feedforward neural network that, each time step, receives sensory information from the environment (the angle and the distance of the nearest food element and the last planned motor action), determines a given motor action selected within four options (move forward, turn left, turn right or stay still) and predicts the next sensory state (i.e., the state of the sensors after the planned action will be executed). Sensory information is used both as input and as teaching input for the output units encoding the predicted state of the sensors - the new sensory state is compared with the predicted state and the difference (error) is used to modify the connection weights through back-propagation. As in the case of the Hinton-

Nowlan model, modifications due to learning are not transferred back into the genotype.

The experimental results show that (a) after a few generations, by learning to predict individuals increase their performance during life not only with respect to their ability to predict but also with respect to their ability to find food, and (b) the ability to find food increases evolutionarily faster and achieves better results at the end of evolution in the case of learning populations than in the case of control populations in which individuals are not allowed to learn during lifetime. Further analyses demonstrate that result (a) can be explained by considering that evolution tends to select individuals that are located in regions of the search space where the learning and evolutionary tasks are dynamically correlated (i.e., where changes due to learning that produce an increase in performance with respect to the learning task also produce positive effects, on the average, with respect to the evolutionary task), and that result (b) can be explained by considering that, since learning tends to channel evolution toward solutions in which the learning task and the evolutionary task are dynamically correlated, learning allows individuals to recover from deleterious mutations (Nolfi, 1999).

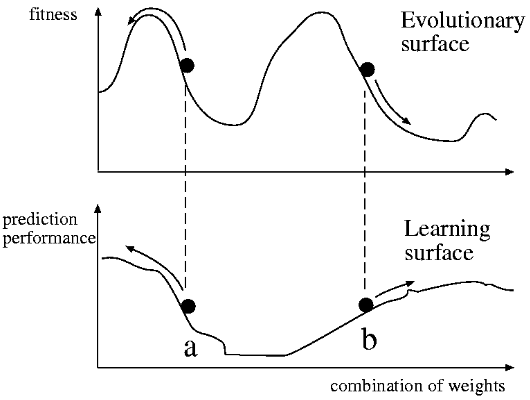

Figure 3. Fitness surface for the evolutionary task (finding food) and performance surface for the learning task (sensory prediction) for all possible weight matrices. Movement due to learning is represented as arrows.

Consider for example two individuals, *a* and *b*, which are located in two distant locations in weight space but have the same fitness at birth, i.e., the two locations correspond to the same height on the fitness surface (see Figure 3). However, individual *a* is located in a region in which the fitness surface and the learning surface are dynamically correlated, i.e., a region in which movements that result in an increase in height with respect to the learning surface also cause, on average, an increase with respect to the fitness surface. Individual *b*, on the other hand, is located in a region in which the two surfaces are not dynamically correlated. If individual *b* moves in weight space it will go up in the learning surface but not necessarily in the fitness surface. Because of learning, the two individuals will move during their lifetime in a direction that improves their learning performance, i.e., in a direction in which their height on the learning surface tends to increase. This implies that individual *a*, which is located in a dynamically correlated region, will end up with a higher fitness than individual *b* and, therefore, will have a better chance to be selected. The final result is that evolution will have a tendency to progressively select individuals which are located in dynamically correlated regions. In other words, learning forces evolution to select individuals which improve their performance with respect to both the learning and the evolutionary task.˙

## 6.2 Adapting to changing conditions on the fly

As we claimed above, learning might complement evolution by providing a mean to master changes that occur too fast to be tracked by the evolutionary process. However, as we will see in this section, the combination of learning and evolution deeply alter both processes so that, in individuals that evolve and learn, adaptive characteristics emerge as the result of the interaction between evolutionary and lifetime adaptation and cannot be traced back to only one of the two processes.

Nolfi and Parisi (1997), evolved neural controllers for a small mobile robot that was asked to explore an arena of 60 x 20 cm surrounded by walls. The robot was provided with 8 infrared sensors that could detect walls up to a distance of about 4 cm and two motors that controlled the two corresponding wheels. The colors of the walls switched from black to white and vice versa each generation. Given that the activity of the infrared sensors is highly affected by the color of the reflecting surface (white walls reflect more that black walls), to maximize their exploration behavior evolved robots should modify their behavior on the fly. In the environment with dark walls, in fact, robots should move very carefully when sensors are activated given that walls are detected only when they are very close. In the environment with white walls, on the contrary, robots should begin to avoid walls only when the sensors are strongly activated in order to explore also the area close to the walls.

Individuals learn during their lifetime by means of a self-generated teaching signals. The genotype of the evolving individuals encodes the connection strengths of two neural modules: a teaching module that each time step receives the state of the sensors as input and produce a teaching signal as output, and an action module that receives the state of the sensors as input and produce motor actions as output. The self-generated teaching signal is used to modify the connection strengths of the action module (for a similar architecture, see Ackley and Littman, 1991). This implies that not only the initial behavior produced by the evolving individuals but also what individuals learn is the result of the evolutionary process and is not determined by the experimenter.

The results of the simulation show that evolved robots display an ability to discriminate the two types of environments and to modify their behavior accordingly, thus maximizing their exploration capability. An analysis of the obtained results reveals that this ability results from a complex interaction between the evolutionary and the learning process. For example, evolved individuals display an inherited ability to behave so as to enhance the perceived differences between the two environments. This in turns allows the learning process to progressively modify the behavior of the robots in such a way that they adapt to the different environmental conditions.

More generally, this and other researches show that evolution, in the case of individuals that are able to change during life as a result of learning, do not tend to develop directly an ability to solve a problem but rather tend to develop a predisposition to acquire such ability through learning.

Other experiments conducted by co-evolving two competing populations of predator and prey robots (Nolfi and Floreano, 1998) emphasized how lifetime learning might allow evolving individuals to achieve generality, i.e., the ability to produce effective behavior in a variety of different circumstances. Predators consist of small mobile robots provided with infrared sensors and a linear camera with a view angle of 36° with which they could detect the prey. Prey consist of mobile robots of the same size provided only with infrared sensors but that have a maximum available speed set to twice that of the predators. Each individual is tested against different competitors for 10 trials. Predators are scored with 1 point for each trial in which they were able to catch the prey while prey are scored with 1 point for each trial they were able to escape predators.

What is interesting about this experimental situation is that, since both populations change across generations, predators and prey are facing ever-changing and potentially progressively more complex challenges. Interestingly, the authors observe that, in this situation, evolution alone displays severe limitations and progressively more effective solutions can be developed only by allowing evolving individuals to adapt on the fly through a form of lifetime learning. Indeed, any possible fixed strategy is able to master only a limited number of different types of competitors and therefore only by combining evolution and learning was it possible to synthesize individuals able to deal with competitors adopting qualitatively different strategies. Indeed, only by evolving learning individuals it was possible to observe the emergence of predators able to detect the current strategy adopted by the prey and to modify their behavior accordingly.

### 6.3 Evolving the Learning Rules

Floreano and Urzelai (2000) conducted a set of experiments in which the genotype of the evolving individuals encoded the learning properties of the neurons of the corresponding neural network (see also Belew, McInerney, and Schraudolph, 1991). These properties included one of four possible hebbian learning rules, the learning rate, and the sign of all the incoming synapses of the corresponding neuron. When the genotype is decoded into a neural controller, the connection strengths are set to small random values. As reported by the authors, after some generations, the genetically specified configuration of learning rules tend to produce changes in the synaptic strengths that allow individuals to acquire the required competencies through lifetime learning. By comparing the results obtained with this method with a control experiment in which the strength of the synapses were directly encoded into the genotype, the authors observed that evolved controllers able to adapt during lifetime can solve certain tasks faster and better than standard non-adaptive controllers. Moreover, they demonstrated that their method scales up well to large neural architectures.

The authors applied their method in order to evolve neural controllers for mobile robots. Interestingly, the analysis of the synaptic activity of the evolved controllers showed that several synapses did not reach a stable state but keep changing all the time. In particular, synapses continue to change even when the behavior of the robot became rather stable.

Similar advantages has been reported by Husband et al. (1999) who evolved a type of neural network in which neurons, that were distributed over a 2D surface, emitted 'gases' that diffused through the network and modulated the transfer function of the neurons in a concentration-dependent fashion, thus providing a form of plasticity. Finally, in the experiments performed by Di Paolo (2000) it has been shown how learning can play the role of a homeostatic process whereby evolved neural networks adapt in order to remain stable in the presence of external perturbations.

## 7. Discussion

All changes that occur during the life of an individual organism, especially those that concern the organism's nervous system and the resulting behavior, are due both to the influence of the information contained in the organism's inherited genotype and to the influence of the individual's specific experience in the specific environment. When the first type of influences (genetic) are prevalent, one talks about maturation while learning is the term used when the second type of influence (environmental) are predominant. When both are equally important the term most frequently used is development.

Simulations using neural networks of the "classical" type (Rumelhart and McClelland, 1986) tend to concentrate on learning and to ignore the organisms' genotypes and the evolutionary process at the population level which results in those genotypes. This makes it difficult to study development using "classical" neural networks. Neural networks viewed in an Artificial Life perspective (Parisi, in press), on the contrary, are used in simulations in which what is simulated is not only the nervous system of organisms but also their body, environment, and genotype. Furthermore, the object of any particular simulation is not a single individual but a population of different individuals which reproduce selectively and evolve across of a succession of generations. In this framework it becomes possible to study neural development, i.e., the changes that occur in an individual organism's nervous system (neural network) during the individual's lifetime and that are due to both the individual's inherited (evolved) genotype and the individual's experience in the specific environment.

In the past decade a number of simulations using neural networks have been conducted which have attempted to model how genetic information can be mapped in the organism's phenotype, how genetic information can determine changes in an individual's neural network and therefore in the individual's behavior, and how genetic information and information from the environment can interact in determining these lifetime changes. In many circumstances this interaction results in better, more rapid, more flexible adaptation of the organisms to the environment in which they happen to live. Therefore, these simulations shed some light on the question why learning has emerged as an evolutionary adaptation.

Of course, like all simulations using neural networks these simulations greatly simplify everything, from the genetic encoding to the genotype-to-phenotype mapping, from the genetically-based changes that occur in an individual's neural network during the individual's life to the role of learning in modulating and directing these changes. These simplifications are all the more regrettable given the steadly increasing corpus of new and detailed knowledge that research in "real" genetics and developmental biology is accumulating in these years. However, simulations using neural networks in an Artificial Life perspective can progressively incorporate new, altough always highly selected, knowledge generated by the "real" sciences and in any case they can play a useful role for testing in detailed and controllable ways ideas on general and specific mechanisms underlying neural and behavioral development in organisms.

## References

Ackley D.H., and Littman M.L. (1991). Interaction between learning and evolution. In C.G. Langton et. al (eds.), *Proceedings of the Second Conference on Artificial Life*. Reading, MA: Addison-Wesley.

Arthur W. (1984). *Mechanisms of Morphological Evolution*. New York: Wiley.

Belew, R.K., McInerney, J., and Schraudolph, N.N. (1991). Evolving networks: Using the genetic algorithm with connectionist learning. In C.G. Langton, J.D. Farmer, S. Rasmussen, and C.E. Taylor (eds.), *Artificial Life II*. Reading, MA: Addison-Wesley.

Cangelosi A. (1999). Heterochrony and adaptation in developing neural networks. In W. Banzhaf et al. (Eds), *Proceedings of GECCO99 Genetic and Evolutionary Computation Conference*. San Francisco, CA: Morgan Kaufmann.

Cangelosi A. and Elman J.L. (1995). Gene Regulation and Biological Development in Neural Networks: an Exploratory Model. *Technical Report*. CRL-UCSD, University of California at San Diego.

Cangelosi A., Nolfi S., and D. Parisi, (1994). Cell division and migration in a 'genotype' for neural networks. *Network - Computation in Neural Systems*, 5:497-515.

Dellaert F. and Beer R.D. (1994). Toward an evolvable model of development for autonomous agent synthesis. In R. Brooks and P. Maes, (eds.), *Proceedings of the Forth Conference on Artificial Life*. Cambridge, MA: MIT Press.

Di Paolo E.A. (2000). Homeostatic adaptation to inversion in the visual field and other sensorimotor disruptions. In J-A Meyer, A. Berthoz, D. Floreano, H.L. Roitblat, and S.W. Wilson (eds.), *From Animals to Animats 6. Proceedings of the VI International Conference on Simulation of Adaptive Behavior*. Cambridge, MA: MIT Press.

Floreano, D. and Urzelai, J. (2000) Evolutionary Robots with on-line self-organization and behavioral fitness. *Neural Networks*, 13:431-443.

Gould, S.J. (1977). *Ontogeny and Phylogeny*. Cambridge, MA: Harward University Press.

Gruau F. (1994). Automatic definition of modular neural networks, *Adaptive Behavior*, 3:151-183.

Hinton G.E. and Nowlan S.J. (1987). How learning guides evolution. *Complex Systems*, 1:495-502.

Holland J.J. (1975). *Adaptation in natural and artificial systems*. Ann Arbor, MI: University of Michigan Press.

Husband, P., Smith T., Jakobi N. and O'Schea M. (1999). Better living through chemistry: Evolving GasNets for robot control. *Connection Science*, 3-4:185-210.

Husbands P., Harvey I., Cliff D. and Miller G. (1994). The use of genetic algorithms for the development of sensorimotor control systems. In P.

Gaussier and J-D. Nicoud (eds.), *From Perception to Action*. Los Alamitos CA: IEEE Press.

Hutchins E., and Hazlehurst B. (1991). Learning in the cultural process. In (C Langton, C. Taylor, J.D. Farmer, and S. Rasmussen (eds.), *Artificial Life II*. Reading, MA: Addison-Wesley.

Kitano H. (1990). Designing neural networks using genetic algorithms with graph generation system. *Complex Systems*, 4:461-476.

Koza J.R. (1992). *Genetic Programming: On the Programming of Computers by Means of Natural Selection*. Cambridge, MA: MIT Press.

Marx J.L. (1989). The cell cycle coming under control. *Science*, 245:252-255.

Mayley G. (1997). Landscapes, learning costs, and genetic assimilation. *Evolutionary Computation*, 4:213-234.

Miller, G.F., P.M. Todd, and S.U. Hedge (1989). Designing neural networks using genetic algorithms. In L. Nadel, and D. Stein (eds.), *Proceedings Third International Conference on Genetic Algorithms*. San Mateo, CA: Morgan Kaufmann.

McKinney M.L. and McNamara K.J. (1991). *Heterochrony: the Evolution of Ontogeny*. New York: Plenum Press.

Nolfi S. (1999). How learning and evolution interact: The case of a learning task which differs from the evolutionary task. *Adaptive Behavior*, 2:231-236.

Nolfi S., Elman J.L. and Parisi D. (1994). Learning and evolution in neural networks. *Adaptive Behavior,* 1:5-28.

Nolfi S. and Floreano D. (1998). Co-evolving predator and prey robots: Do ‘arm races’ arise in artificial evolution? *Artificial Life*, 4:311-335.

Nolfi S. and Floreano D. (1999). Learning and evolution. *Autonomous Robots,* 1: 89-113.

Nolfi S., Miglino O. and Parisi D. (1994). Phenotypic Plasticity in Evolving Neural Networks. In D. P. Gaussier and J-D. Nicoud (eds.), *Proceedings of the Intl. Conf. From Perception to Action*. Los Alamitos, CA: IEEE Press.

Nolfi S. and Parisi D. (1995). Genotypes for neural networks. In M. A.Arbib (ed.), *Handbook of brain theory and neural networks*. Cambridge, MA: MIT Press.

Nolfi S. and Parisi D. (1997). Learning to adapt to changing environments in evolving neural networks. *Adaptive Behavior,* 1:75-98.

Parisi D. (1996). Computational models of developmental mechanisms. In R. Gelman and T.K. Au (eds.), *Perceptual and Cognitive Development*. San Diego, Academic Press.

Parisi D. (1997). Artificial life and higher level cognition. *Brain and Cognition*, 34:160-184.

Parisi D. (in press). Neural networks and Artificial Life. In D. Amit and G. Parisi (eds.) *Frontiers of Life*. San Diego, Academic Press.

Parisi, D. and Nolfi, S. (2001). Development in neural networks. In M. Patel, V. Honovar, and K. Balakrishnan (eds.), *Advances in the Evolutionary Synthesis of Intelligent Agents*. Cambridge, MA: MIT Press.

Purves D. (1994). *Neural activity and the growth of the brain*. Cambridge: Cambridge University Press.

Purves D. and Lichtheim J.W. (1985). *Principles of Neural Development*. Sunderland, MA: Sinauer Ass.

Rumelhart, D.E. and McClelland, J.L. (1986). *Parallel Distributed Processing. Explorations in the Microstructure of Cognition. Volume I. Foundations*. Cambridge, MA: MIT Press.

Schwefel H.P. (1995). *Evolution and Optimum Seeking*. New York: Wiley Press.

Wagner G.P., and Altenberg L. (1996). Complex adaptations and the evolution of evolvability. *Evolution*, 50:967-976.

Yao, X. (1993). A review of evolutionary artificial neural networks. *International Journal of Intelligent Systems*, 4:203-222.