Institute of Psychology
C.N.R. - Rome

The influence of learning on evolution

Domenico Parisi Stefano Nolfi
Institute of Psychology,
National Research Council,
15, Viale Marx, 00137, Rome, Italy
e-mail:domenico@gracco.irmkant.rm.cnr.it
stefano@kant.irmkant.rm.cnr.it

June 1994

Technical Report PCIA-94-09

Department of Cognitive Processes and Artificial Intelligence
15, Viale Marx
00137 - Rome - Italy
voice: 0039-6-86894596
fax: 0039-6-824737

To appear In: In Belew, R,K, Mitchell, M. (eds.) Plastic Individuals in Evolving Populations. SFI Series, Addison-Wesley

The influence of learning on evolution

Domenico Parisi Stefano Nolfi
Institute of Psychology,
National Research Council,
15, Viale Marx, 00137, Rome, Italy
e-mail:domenico@gracco.irmkant.rm.cnr.it
stefano@kant.irmkant.rm.cnr.it

1. Evolution and learning

Evolution and learning are two different ways in which the behavior, and other traits, of organisms can change. Evolution is change at the population level. Organisms reproduce selectively and subject to mechanisms (mutation, sexual recombination) which maintain interindividual variability. This causes changes in the population from one generation to the next. Learning, on the other hand, is change at the individual level. By interacting during its life with a specific environment an organism can change its behavior by incorporating, through its experience, aspects of the environment in its internal structure.

Evolutionary change is generationally cumulative. The changes which occur in a particular generation are superimposed upon changes that have occurred in previous generations. Learning is individually (but not generationally) cumulative. The changes that occur in an individual at a particular time of the individual's life are influenced by changes at preceding times but no changes due to learning are inherited by the individual's offspring.

2. Evolution's influence on learning

Although evolution and learning are two distinct kinds of change which occur in two distinct types of entities (populations and individual organisms), they may influence each other. The influence of evolution on learning is not surprising. Evolutionary change leaves its trace in the genotype. Hence, each individual inherits a genome which is the cumulative result at the level of the individual of the past evolutionary changes which have occurred at the level of the population. Since an individual's genome partially specifies the resulting phenotypic individual and it constrains how the individual will behave and what it will learn, the way is open for an influence of evolution on learning.

Simulations applying genetic algorithms to populations of neural networks have shown this to occur. In these simulations a strictly Darwinian framework is adopted: the changes that occur in an individual network as a consequence of learning are not inherited by the network's offspring. (For an example of a Lamarckian inheritance of learned changes, cf. Ackley and Litman, 1994). However, evolution progressively selects for networks that incorporate a predisposition to learn some specific task. Although the networks of later generations do not perform better in the task at birth, i.e., prior to learning, in comparison to networks of earlier generations they do learn more or more quickly. Therefore, what is genetically inherited is not a congenital ability to perform the task but only a predisposition to learn the task if given the appropriate experience.

A predisposition to learn some particular task can be incorporated in a neural network in a variety of different ways, for example as the network's initial matrix of connection weights or the network's architecture. Evolution may select initial weight matrices or network architectures that cause better learning. This has been shown to happen both in the case where the learning task and the evolutionary task are the same (for weight matrices, cf. Belew, McInerney, and Schraudolph, 1991; for network architectures, cf. Miller, Todd, and Hedge, 1989) and in the case where they are different (for weight matrices, cf. Parisi, Nolfi, and Cecconi, 1991; Nolfi, Elman, and Parisi, in press). The learning task and the evolutionary task are the same if the terminal performance in the learning task is identical with the fitness in terms of which individuals are selected for reproduction. The two tasks are different if the individual learns task A but its fitness is independently evaluated based on its performance on task B.

A predisposition to learn created by evolution can also be expressed in other ways, for example as the inheritance of an appropriate learning rate or momentum for learning some particular task (cf. Belew, McInerney, and Schraudolph, 1991) or, in more ecological conditions, as an inherited tendency to behave in such a way that the individual is exposed to the appropriate learning experiences. For example, assuming that the teaching input for some supervised learning task is only physically available in some specific locations in the environment, networks can be shown to evolve and, therefore, to inherit genetically a tendency to move in the environment in such a way that they tend to approach and stay close to those locations (Denaro and Parisi, 1994).

3. Learning's influence on evolution

The influence of evolution on learning can be easily understood within a Darwinian framework because this framework provides a mechanism through which the influence can be realized: the inherited genotype is simultaneously the result of evolution and a partial cause of the phenotypic individual and its learning tendencies. A Darwinian framework makes the opposite influence, of learning on evolution, much more difficult to explain. Learning is realized as changes in the phenotypic individual but these changes are not inherited. Therefore, learning should have no direct effect on the course of evolution.

This orthodoxy has been challenged by some simple simulations done by Hinton and Nowlan (1987) (H&N) (reprinted as chapter X). H&N have shown that if evolution must search for a very specific result (genome) among many useless alternative results (genomes), learning can aid this search even if it takes the simple form of random changes during an individual's life. Genomes which are close to the desired genome but which evolution alone would have a hard time finding can be discovered if inherited genomes are subject to random changes during life and if they reproduce as a function of how quickly the random changes transform them into the desired genome. In this way, evolution can converge on the desired genome more quickly than if learning is absent, although it remains true that learned changes are not inherited. If evolution is unaided by learning, its chances of success are restricted to the case that the single desired genome suddenly emerges because of the chance factors operating at reproduction (random mutations and sexual recombination). Learning solves half of the problem and therefore it makes the task of the chance factors much easier.

The H&N simulations have been influential in giving a new plausibility to the idea that, in addition to the more biologically plausible causal influence of evolution on learning, there may be a more indirect causal influence of learning on evolution. However, the H&N simulations have various limitations and it remains to be seen if learning continues to influence evolution when these limitations are removed.

A first limitation of the H&N simulations is that learning is equated to random changes. Changes due to learning appear to be nonrandom but to have a direction, a goal. There is a task to be learned. A second limitation is that in H&N simulations there is no independent physical environment in which individual organisms live and, therefore, evolve and learn. Evolution and learning are changes which result from the interactions between organisms and environments. Evolution is based on fitness. But fitness is not a property of genotypes but of genotype/environment pairs (Odling-Smee, 1988).

The third limitation is perhaps the most serious one because it could be interpreted as restricting the significance of H&N results. In their simulations the evolutionary task and the learning task coincide (see above). What determines the reproductive chances of a particular individual (its fitness) is its performance on the learning task. Individuals are more likely to reproduce if they happen to find a solution to their learning task early in their life. Therefore, one might object that it is this equating of fitness with learning performance that explains why learning can influence evolution in H&N simulations.

In fact, the identity of evolutionary task and learning task in H&N and other simulations can represent a limitation of these simulations. In real organisms it is not clear that a good learning performance automatically translates into high fitness. Learning tasks, insofar as they are evolved, are likely to converge with fitness but, especially in more advanced organisms, the relationship between the two can be very complex and indirect.

Simulations in which the evolutionary task and the learning task are distinct tasks have also been done (e.g., Nolfi, Elman, and Parisi, in press). In these simulations an attempt is also made to overcome the other two limitations of the H&N simulations. Learning is not random changes but is directed learning (using backpropagation) and the organisms (networks) evolve and learn in a physical (simulated) environment.

Hence, these simulations may represent a further and more robust test of an influence of learning on evolution.

4. How learning to predict can have an effect on evolution

In these simulations, organisms reproduce as a function of their performance in the task of capturing food present in the environment. This is the evolutionary task. The organisms are modeled by neural networks living in an environment that contains randomly distributed pieces of food. Each network has input units encoding the position of the nearest food element relative to the organism. The network's output units encode various possibile motor actions that allow the organism to move and turn. The network architecture is fixed. The problem for evolution is to find weight matrices for this fixed network architecture that cause an organism to respond to sensory information about food location with motor actions allowing the organism to capture food efficiently.

Furthermore, during its life each organism learns to predict how the position of food changes with its motor actions. This is the learning task. The learning task is a conceptually distinct task with respect to the evolutionary task although it may be causally related to fitness. The individual's reproductive chances are determined by its performance in the evolutionary task (capturing food), not in the learning task (predicting food position). In principle, an individual that learns quickly how to predict could have poor fitness and not reproduce, or an individual can be a slow learner but it can have high fitness and reproduce.

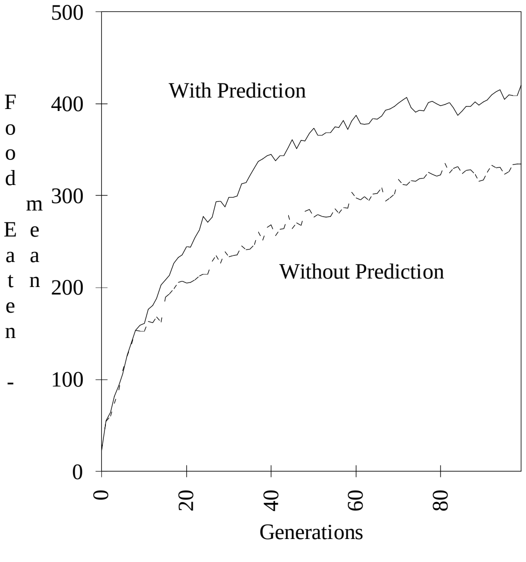

Figure 1. Average number of food elements eaten by successive generations of organisms evolved with and without learning during life. Each of the two curves represents the average performance of ten different simulations with different random assignment of weights.

We have observed that in these conditions, learning has a positive influence on evolution (for a demonstration of the same effect on other tasks c.f. Cecconi and Parisi, 1991). Populations in which both evolution and learning occur exhibit a better growth of fitness across generations than populations that only evolve but do not learn (cf. Figure 1).

In these simulations, reproduction is nonsexual (single parent) and the addition of random variability at reproduction is restricted to mutations. If we assume that the portion of weight space which is explored by learning overlaps significantly with the portion explored by mutations, we can advance the following explanation for the positive influence of learning on evolution. (The explanation is similar to the one given by H&N for their results but is extended to cover the more complex situations studied in our simulations).

In a population which only evolves, the decision regarding which individuals should reproduce can only be based on the fitnesses of genomes. If what is genetically inherited is a matrix of connection weights, each individual genome can be represented as a particular location in weight space which corresponds to a given height on the fitness surface. This height represents the fitness value of the particular genome. (Phenotypes, not genomes, have fitness. However, for simplicity we are identifying genomes and phenotypes.) To determine this value the individual is tested during its life.

The situation is different in a population that both evolves and learns. Learning causes changes in weights. This means that the individual moves in weight space and, therefore, its total fitness is a function of the various fitnesses (heights on the fitness surface) associated to the various locations traversed during learning. Since learning determines this path and the fitness of an individual is a function of the path traversed during learning, this appears to be the mechanism that allows learning to influence evolution.

But why has learning a *beneficial* effect on evolution? Consider the decision whether a particular individual should reproduce or not. In a population without learning this decision can only be based on the fitness of a single location in weight space, i.e., the location occupied by the genome of the individual. By testing the individual the fitness associated with this specific location can be known. However, the region surrounding the individual's location remains unknown. In other words, the selective mechanism ignores what fitnesses are associated with the locations that are near the given location.

Consider now two individuals, *a* and *b*, which have two rather different genomes and therefore are located in two distant locations in weight space (cf. Figure 2). Assume, however, that the genome of *a* and the genome of *b* have the same fitness (case 1); alternatively, assume that the genome of *a* has a higher fitness than the genome of *b* (case 2). In a population without learning *a* and *b* would have the same reproductive chances (case 1) or *a* would reproduce rather than *b* (case 2).

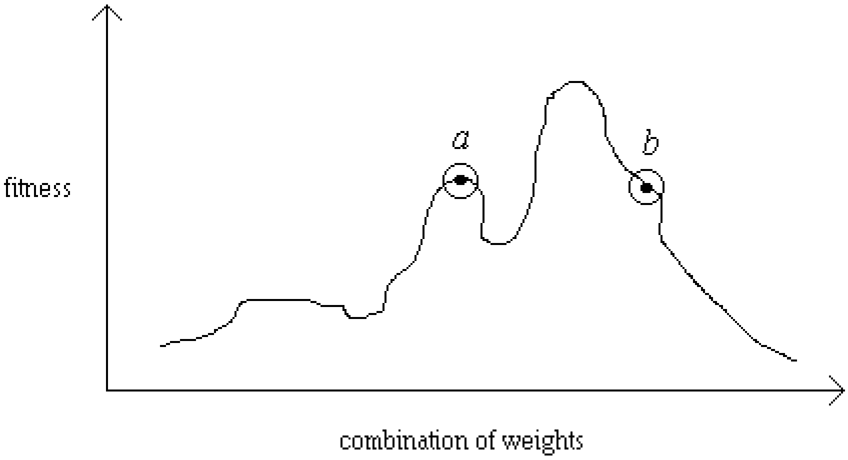

Figure 2. Fitness of all possible weights matrices. Point *b* has a better surrounding region than point *a* even if the fitnesses of the two points are identical. For practical reasons the N dimensions of the weight space are represented as a single dimension.

Because of mutations the offspring of a reproducing individual won't be located in precisely the same location as their parent, but rather somewhere in the region surrounding it. Imagine now that the region surrounding individual *b* is better than the region surrounding individual *a*; i.e., it contains more locations with a higher fitness than the region surrounding *a*. This implies that the offspring of *b* will have a higher fitness than the offspring of *a*. Therefore, instead of leaving the decision to chance (case 1) or favouring individual *a* (case 2), it would be more effective for evolution to select individual *b* for reproduction.

Learning is a mechanism which allows evolution to find out about the regions in weight space surrounding the location of candidates for reproduction. Learning involves movement in weight space. This implies that the total fitness of an individual which learns will be a function of both (i) the fitness of the starting point (inherited genome), and (ii) the fitnesses of the locations traversed during learning. If we assume that at least some of these locations are situated in the surrounding region explored by mutations, we can see that in populations that both evolve and learn, reproductive decisions can be based on knowledge of surrounding regions. In other words, evolution in based on the fitnesses of the currently living individuals but it would be more effective if it could be based also on the fitnesses of the offspring of these individuals. Learning allows evolution "to look into the future". Therefore, evolution with learning can be more effective than evolution alone.

The argument so far implies that even random movement in weight space during life can benefit evolution. And, in fact, this is what H&N simulations demonstrate. As a matter of fact, a population of networks which are selected for the ability to capture food and at the same time "learn" during life on the basis of randomly generated teaching input, has a better evolutionary fitness curve than a population with no learning at all (Parisi, Nolfi, and Cecconi, 1991). (H&N simulations are a pure example of learning as random movement in weight space. Backpropagation learning on the basis of a randomly generated teaching input, as described in Parisi, Nolfi, and Cecconi (1991), may involve movement which has some directionality.)

However, the effect of learning on evolution is likely to vary as a function of the learning task and how the learning task is related to the evolutionary task. Learning has a tendency to push evolution towards particular kinds of solutions and different types of learning (i.e., different learning tasks) can have different effects on the evolutionary process.

In order to understand how learning affects the course of evolution, we should consider two different surfaces, the fitness surface that represents the performance of each individual with respect to the evolutionary task and the learning surface that represents the performance of each individual with respect to the learning task. The fitness surface determines which individuals will reproduce while the learning surface determines how individuals move in weight space during their lifetime. But, as we have seen, the way in which individuals move in weight space during their lifetime affects their fitness. The learning task and the corresponding learning surface determine the nature of the movement and, as a consequence, have an influence on which individuals are selected for reproduction.

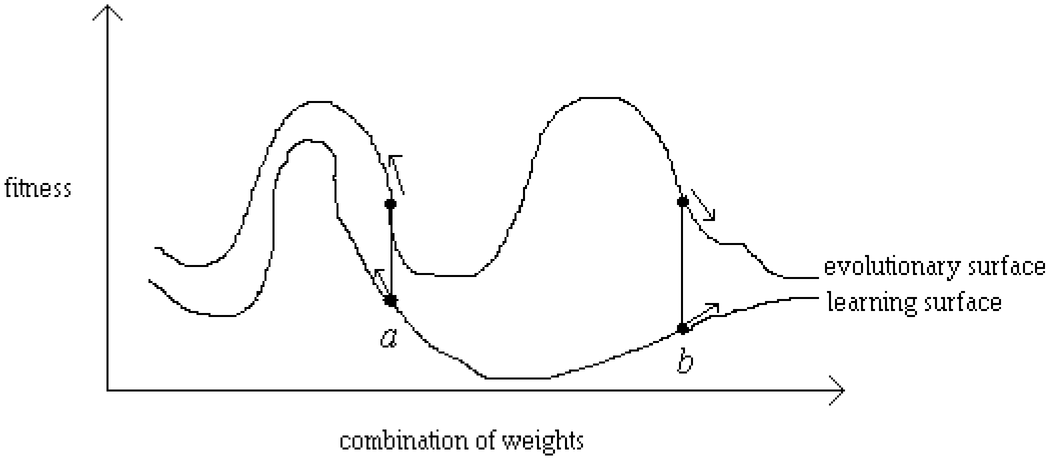

Figure 3. Fitness surface for the evolutionary task and performance surface for the learning task for all possible weights matrices. Movements due to learning are represented as arrows. Point *a* is in a region in which the two surfaces are dinamically correlated. Even if *a* and *b* have the same fitness on the evolutionary surface at birth, *a* has more probability to be selected than *b* since it is more likely to increase its fitness during life than *b*.

Consider two individuals, *a* and *b*, which are located in two distant locations in weight space but have the same fitness; i.e., the two locations correspond to the same height on the fitness surface (cf. 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 increases in height with respect to the learning surface tend to cause increases also 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 more chances to be selected. The final result is that evolution will have a tendency to progressively select individuals which are located in dynamically correlated regions. Learning represents an evolutionary pressure to select individuals which improve their performance with respect to both the learning and the evolutionary task.

We have observed this influence of learning on evolution in our simulations with networks that are selected for the ability to collect food elements and at the same time learn during their life how to predict changes in sensory input which are caused by their movements. As we show in Figure 4, while the individuals of generation 0 do not increase their ability to capture food during life, the individuals of successive generations eat more food elements in successive epochs of their life. Now, all individuals are subject to prediction learning during their life and all individuals, including those of generation 0, do learn how to predict. (The error on the prediction task has more or less the same decrease across epochs for all individuals.) However, learning has a beneficial effect on the fitness of an individual only in later generations. Thus, although learning to predict does not by itself increase the food collecting ability (generation 0), it is able to guide evolution so that it progressively selects individuals for which learning does have a beneficial effect on the evolutionary task, i.e. individuals located in regions of dynamic correlation between the evolutionary surface and the learning surface.

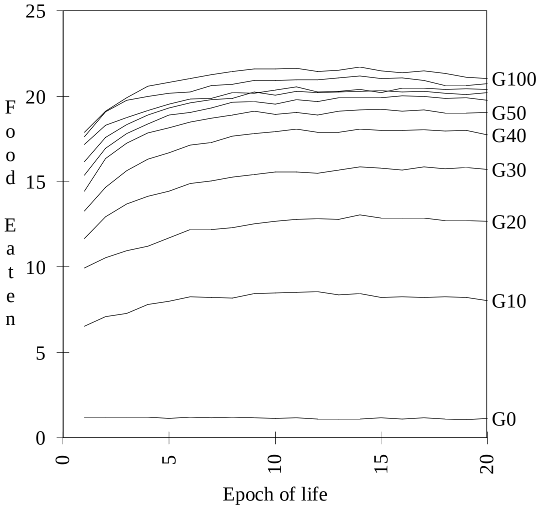

Figure 4. Average number of food elements eaten by organisms during their lifetime in successive generations. Each curve represents the average performance of ten different

simulations with different random initial assignment of weights. For practical reasons, only generations 0, 10, 20, 30, 40, 50, and 100 are shown.

Although the direction in which learning guides evolution depends on the particular learning task, some effect may be expected whatever the learning task. For example, if one trains networks to do the XOR task and the networks belong to a population selected for the ability to capture food, one obtains the same result (cf. Parisi, Nolfi, and Cecconi, 1991). After a certain number of generations, individuals that are learning the XOR task show an improvement in their food collecting ability during life which parallels their improved performance on the XOR task; i.e., learning forces evolution to find individuals that are located in dynamically correlated regions. However, the strength of the effect is likely to be a function of the relation between learning and evolutionary task. If the two tasks are statically correlated, i.e., a weight matrix which has a high (low) value on one surface tends to have a high value (low) on the other surface, the effects will be stronger. In the extreme case in which the learning task and the evolutionary task are the same, their performance surfaces will coincide and, therefore, by definition, an increase in performance with respect to one surface will correspond to an increase in performance with respect to the other surface.

References

Ackley, D. E., and M. L. Littman. "A Case for Lamarkian Evolution." In *Artificial Life III*, edited by C. G. Langton. Reading, MA: Addison-Wesley, 1994.

Belew, R. K., J. McInerney, and N. N. Schraudolph. "Evolving Networks: Using the Genetic Algorithm with Connectionist Learning." In *Artificial Life II*, edited by C. G. Langton, J. D. Farmer, S. Rasmussen, and C. E. Taylor. Reading, MA: Addison-Wesley, 1991.

Cecconi, F., and D. Parisi. "Evolving Organisms that can Reach for Objects." In *From Animals to Animats*, edited by J. A. Meyer, and S. W. Wilson. Cambridge, MA: MIT Press, 1994.

Denaro, D., and D. Parisi. "Cultural Transmission in a Population of Neural Networks." Technical Report, Institute of Psychology, National Research Council, Rome, 1994.

Hinton, G. E., and S. J. Nowlan. "How Learning Guides Evolution." *Complex Systems*, **1**, (1987): 495-502.

Miller, G. F., P. M. Todd, and S. U. Hedge. "Designing Neural Networks Using Genetic Algorithms." In *Proceedings of the Third International Conference on Genetic Algorithms*. San Mateo, CA: Morgan Kaufmann, 1989.

Nolfi, S., J. Elman, and D. Parisi. "Learning and evolution in neural networks.". *Adaptive Behavior*, in press.

Odling-Smee, R. "Niche-constructing phenotypes." In *The role of behavior in evolution*, edited by H. C. Plotkin, Cambridge, MA: MIT Press, 1988.

Parisi, D., S. Nolfi, and F. Cecconi. "Learning, Behavior, and Evolution." In *Toward a practice of autonomous systems*., edited by F. Varela and P. Bourgine. Cambridge, MA: MIT Press, 1991.